APPLICATIONS REFERENCE MANUAL

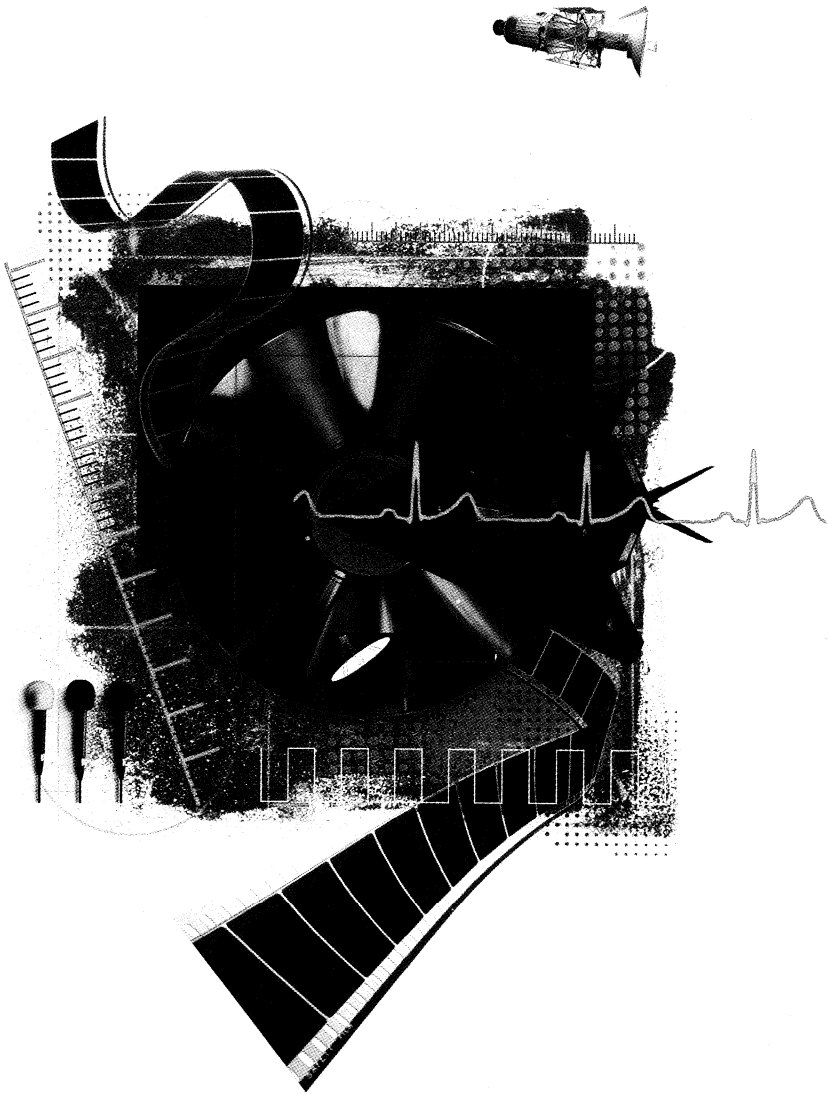

APPLICATION NOTES, TECHNICAL ARTICLES AND OTHER DESIGN TUTORIALS COVERING AUDIO AND VIDEO CIRCUITS, A/D AND D/A CONVERSION, DATA ACQUISITION AND SIGNAL CONDITIONING, DIGITAL SIGNAL PROCESSING AND SIGMA-DELTA CONVERSION

ANALOG DEVICES

1993 APPLICATIONS REFERENCE MANUAL

©Analog Devices, Inc., 1993
All Rights Reserved

General Information 1
Analog I/O Ports 2
Analog-to-Digital Converters 3
Audio Products 4
Communications Products 5
Computational Products 6
Data Acquisition Subsystems 7
Digital-to-Analog Converters 8
Digital Signal Processing Products 9
Instrumentation Amplifiers 10
Isolation Amplifiers 11
Multiplexers and Switches 12
Operational Amplifiers 13
Programmable Delay Generators 14
References 15
Resolver-to-Digital Converters 16
RMS-to-DC Converters 17
Sample/Track-Hold Amplifiers 18
Sensors 19
Sigma-Delta Converters 20
Signal Conditioning Products 21
Voltage-to-Current Converters 22
Voltage-to-Frequency Converters 23
Miscellaneous 24
Appendix 25
Index 26

1993
APPLICATIONS REFERENCE MANUAL

© Analog Devices, Inc., 1993
All Rights Reserved

Information furnished by Analog Devices is believed to be accurate and reliable. However, no responsibility is assumed by Analog Devices for its use, nor for any infringements of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under any patent or patent rights of Analog Devices.

U.S.:
RE29,992, RE30,586, RE31,850, 3,729,660, 3,793,563, 3,803,590, 3,842,412, 3,868,583, 3,890,611, 3,906,486, 3,909,908, 3,932,863, 3,940,760, 3,942,173, 3,942,173, 3,946,324, 3,950,603, 3,961,326, 3,978,473, 3,979,688, 4,016,559, 4,020,486, 4,029,974, 4,034,366, 4,054,829, 4,055,773, 4,056,740, 4,068,254, 4,088,905, 4,092,639, 4,092,698, 4,109,215, 4,118,699, 4,123,698, 4,131,884, 4,136,349, 4,138,671, 4,141,004, 4,142,117, 4,168,528, 4,210,830, 4,213,806, 4,228,367, 4,250,445, 4,260,911, 4,268,759, 4,270,118, 4,272,656, 4,285,051, 4,286,225, 4,309,693, 4,313,083, 4,323,795, 4,333,047, 4,338,591, 4,340,851, 4,349,811, 4,363,024, 4,374,314, 4,374,335, 4,383,222, 4,395,647, 4,399,345, 4,400,689, 4,400,690, 4,404,529, 4,427,973, 4,439,724, 4,444,309, 4,449,067, 4,454,413, 4,460,891, 4,471,321, 4,475,103, 4,475,169, 4,476,538, 4,481,708, 4,484,149, 4,485,372, 4,491,825, 4,503,381, 4,511,413, 4,521,764, 4,538,115, 4,542,349, 4,543,560, 4,543,561, 4,547,766, 4,547,961, 4,556,870, 4,558,242, 4,562,400, 4,565,000, 4,572,975, 4,583,051, 4,586,019, 4,586,155, 4,590,456, 4,596,976, 4,601,760, 4,604,532, 4,608,541, 4,622,512, 4,626,769, 4,633,165, 4,639,683, 4,644,253, 4,646,056, 4,646,238, 4,675,561, 4,677,369, 4,678,936, 4,683,423, 4,684,922, 4,685,200, 4,687,984, 4,694,276, 4,697,151, 4,703,283, 4,707,682, 4,709,167, 4,717,883, 4,722,910, 4,739,281, 4,742,331, 4,751,455, 4,752,900, 4,757,274, 4,761,636, 4,769,564, 4,771,011, 4,774,685, 4,791,318, 4,791,551, 4,800,524, 4,804,960, 4,808,908, 4,811,296, 4,814,767, 4,833,345, 4,839,653, 4,855,585, 4,855,618, 4,855,684, 4,857,862, 4,859,944, 4,862,073, 4,864,454, 4,866,505, 4,878,770, 4,879,505, 4,884,075, 4,885,585, 4,888,589, 4,891,533, 4,891,645, 4,899,152, 4,902,959, 4,904,921, 4,924,227, 4,926,178, 4,928,103, 4,928,934, 4,929,909, 4,933,572, 4,940,980, 4,957,583, 4,962,325, 4,969,823, 4,970,470, 4,973,978, 4,978,871, 4,980,634, 4,983,929, 4,985,739, 4,990,797, 4,990,803, 4,990,916, 5,008,671, 5,010,297, 5,010,337, 5,021,120, 5,026,667, 5,027,085, 5,030,849, 5,036,298, 5,036,322, 5,039,945, 5,041,795, 5,043,295, 5,043,657, 5,043,675, 5,043,732, 5,053,653, 5,055,723, 5,055,843, 5,065,144, 5,065,214, 5,070,331, 5,075,633, 5,075,677, 5,077,494, 5,077,541, 5,084,753, 5,086,370, 5,087,894, 5,087,889, 5,091,701, 5,095,274, 5,097,223, 5,101,126, 5,103,281, 5,111,431, 5,113,362, 5,115,202, 5,119,094, 5,120,990, 5,124,596, 5,124,648, 5,126,586, 5,126,653, 5,132,931, 5,134,401, 5,136,184, 5,141,898, 5,146,181, 5,150,074, 5,159,341, 5,166,637, 5,170,335, 5,175,550, 5,179,293, 5,184,130, 5,192,922, 5,195,827, 5,196,422

France:
111.833, 70 10561, 75 27557, 76 08238, 77 20799, 78 10462, 79 24041, 80 00960, 80 11312, 80 11916, 81 02661, 81 14845, 83 03140, 96 08238, 0393027, 0343177

Japan:
1,092,928, 1,242,936, 1,242,965, 1,306,235, 1,337,318, 1,401,661, 1,412,991, 1,432,164, 1180 463, 1,651,578, 1,694,024

Germany:
2,014034, 25 40 451.7, 26 11 858.1, 31 04 331, 34 04 652, 37 79139.7, 37 82067.2

U.K.:
1,310,591, 1,310,592, 1,537,542, 1,590,136, 1,590,137, 1,599,538, 2,008,876, 2,032,659, 2,040,087, 2,050,740, 2,054,992, 2,075,295, 2,081,040, 2,087,656, 2,103,884, 2,104,288, 2,107,951, 2,115,932, 2,118,386, 2,119,139, 2,119,547, 2,126,445, 2,126,814, 2,135,545, 0393027, 0343177

Canada:
984,015, 1,006,236, 1,025,558, 1,035,464, 1,054,248, 1,140,267, 1,141,034, 1,141,820, 1,142,445, 1,143,306, 1,150,414, 1,153,607, 1,157,571, 1,159,956, 1,177,127, 1,177,966, 1,184,662, 1,184,663, 1,191,715, 1,192,310, 1,192,311, 1,192,312, 1,203,628, 1,205,920, 1,212,730, 1,214,282, 1,219,679, 1,219,966, 1,223,086, 1,232,366, 1,233,913, 1,234,921, 1,295,051

Sweden:
7603320-8

USE IN LIFE SUPPORT APPLICATIONS

Products sold by Analog Devices are not designed for use in life support equipment where malfunction of such products can reasonably be expected to result in personal injury or death. Buyer uses or sells such products for life support application at buyer's own risk and agrees to defend, indemnify, and hold harmless Analog Devices from any and all damages, claims, suits, or expense resulting from such use.

General Information
Contents

Page

Introduction 1–3
Table of Contents 1–5

1

Introduction

This volume contains the new *Applications Reference Manual* from Analog Devices. This manual is a compilation of the "best" from a noteworthy collection of technical articles, application notes, *Analog Dialogue* articles and design tutorials written and published over the past 25 years by members of Analog Devices' own technical engineering community.

It is our hope this manual will serve as a valuable reference tool to engineers involved in analog and/or digital circuit design. This volume not only contains new and innovative solutions to many common (and even some uncommon!) circuit design problems, but also revisits standard solutions to problems such as noise reduction and overvoltage protection, to name a few.

The various application notes and technical articles have been grouped into sections alphabetically by product category (e.g., A/D Converters, Analog I/O, Audio Components, etc.). Within each product category, the articles and application notes have been further grouped into like topics, with the more detailed topics at the end of the section.

Since a reference manual is only as good as its index, you will also find at the back of the manual a very complete and detailed index sorted by product number, subject, application note (AN) number and author.

ACKNOWLEDGMENTS

It would be impossible to list here all who contributed to the production and compilation of this the first edition of our *Applications Reference Manual*. Credit most certainly must be given to the numerous authors and/or contributors whose works are included within this manual, and whose valuable time was given to reviewing the material. However, since a manual directly reflects the efforts of the people that put it all together, much of the credit must go to Marie Etchells and her Communications Services group, without whose tireless efforts this manual could not have been completed.

Steve Guinta, Editor

Table of Contents

Page

Analog I/O Ports – Section 2 2–1

AN-221 Power-Down Circuit Cuts Power to AD7769 and AD7774 2–3
AN-222 AD7569/AD7669 Operation with Offset Signal Grounds for Disk Drive Applications 2–5
AN-224 Input/Output Level Shifting with the AD7769 2–7
AN-284 Implement Infinite Sample-and-Hold Circuits Using Analog Input/Output Ports 2–11
AN-285 A Function Generator and Linearization Circuit Using the AD7569 2–15

Analog-to-Digital Converters – Section 3 3–1

AN-289 8-Bit A/D Converter Mates Transducers with μPs 3–3
AN-290 AD7578 and AD7582 Performance with Reduced V_{DD} Supply (12 V ± 10%) 3–11
AN-291 Asynchronous Clock Interfacing with the AD7878 3–13
AN-223 AD7575 Operation with an Offset Signal Ground for Disk Drive Applications 3–15
AN-292 Interfacing the AD7572 to High Speed DSP Processors 3–17
AN-293 The AD7574 Analog to Microprocessor Interface 3–21
AN-294 AD7672 Converter Delivers 12-Bit 200 kHz Sampling Systems 3–27
AN-295 Using Multiple AD1334s in Many-Channel Synchronous Sampling Applications 3–35
AN-296 Build a Single-Shot Recorder to Catch Fast Transients 3–37
AN-297 Test Video A/D Converters Under Dynamic Conditions 3–49
AN-298 AD671 12-Bit, 2 MHz ADC Digitizes CCD Outputs for Imaging Applications 3–59
AN-299 Simple Circuit Provides Ratiometric Reference Levels for AD782x Family of Half-Flash ADCs 3–67
AN-215A Designer's Guide to Flash-ADC Testing
Part 1: Flash ADCs Provide the Basis for High Speed Conversion 3–71
AN-215B Designer's Guide to Flash-ADC Testing
Part 2: DSP Test Techniques Keep Flash ADCs in Check 3–79
AN-215C Designer's Guide to Flash ADC Testing
Part 3: Measure Flash-ADC Performance for Trouble-Free Operation 3–85
AN-301 Undersampling Techniques Simplify Digital Radio 3–95
AN-302 Exploit Digital Advantages in an SSB Receiver 3–103
AN-228 Evaluation Board for the AD7884 16-Bit A/D Converter 3–109
AN-229 Evaluation Board for the AD7885 16-Bit A/D Converter 3–115
AN-368 Evaluation Board for the AD7701/AD7703 Sigma-Delta A/D Converters 3–121

Audio Products – Section 4 4–1

AN-111 A Balanced Summing Amplifier 4–3
AN-112 A Balanced Input High Level Amplifier 4–5
AN-113 An Unbalanced, Virtual Ground Summing Amplifier 4–7
AN-114 A High Performance Transformer-Coupled Microphone Preamplifier 4–9
AN-115 Balanced Low Noise Microphone Preamplifier Design 4–11
AN-116 AGC Amplifier Design with Adjustable Attack and Release Control 4–13
AN-121 High Performance Stereo Routing Switcher 4–15
AN-122 A Balanced Mute Circuit for Audio Mixing Consoles 4–19
AN-123 A Constant Power "Pan" Control Circuit for Microphone Audio Mixing 4–21
AN-124 Three High Accuracy RIAA/IEC MC and MM Phono Preamplifiers 4–23
AN-125 A Two-Channel Dynamic Filter Noise Reduction System 4–29

Page

AN-127 An Unbalanced Mute Circuit for Audio Mixing Channels 4–31
AN-128 A Two-Channel Noise Gate 4–33
AN-129 A Precision Sum and Difference (Audio Matrix) Circuit 4–35
AN-130 A Two-Band Audio Compressor/Limiter 4–37
AN-131 A Two-Channel VCA Level (Volume) Control Circuit 4–39
AN-133 A High Performance Compandor for Wireless Audio Systems 4–41
AN-134 An Automatic Microphone Mixer 4–45
AN-135 The Morgan Compressor/Limiter 4–47
AN-136 An Ultralow Noise Preamplifier 4–55
AN-211 The Alexander Current-Feedback Audio Power Amplifier 4–57
AN-232 Bootstrapped IC Substrate Lowers Distortion in JFET Op Amps 4–73

Communications Products – Section 5 5–1

AN-238 AD2xx Family for RS-232 Communications 5–3
AN-48 Designing Digital Repeaters with ICs 5–9
AN-50 A Variable-Frequency Clock-Recovery Circuit Using the RPT-82 or RPT-83 5–15
AN-119 Improved T148 (CEPT) PCM Repeater for #22AWG (0.7 mm) Unshielded Twisted-Pair Wire 5–17
AN-303 Understanding and Applying the AD7341/AD7371 Switched Capacitor Filters 5–19

Computational Products – Section 6 6–1

AN-304 One-Chip "Slide Rule" Works with Logs, Antilogs for Real-Time Processing 6–3
AN-306 Synchronous System Measures μΩs 6–11
AN-307 Modem Circuit Techniques Simplify Instrumentation Design Techniques 6–13
AN-308 Commutating Amp Multiplies Precisely 6–27
AN-309 Build Fast VCAs and VCFs with Analog Multipliers 6–29
AN-212 Using The AD834 in DC to 500 MHz Applications: RMS-to-DC Conversion, Voltage Controlled Amplifiers and Video Switches 6–39
AN-216 Video VCAs and Keyers Using the AD834 and AD811 6–49

Data Acquisition Subsystems – Section 7 7–1

AN-312 DAS Packs Memory for Easy μP Interface 7–3

Digital-to-Analog Converters – Section 8 8–1

AN-313 Getting the Most from High Resolution D/A Converters 8–3
AN-342 Analog Signal-Handling for High Speed and Accuracy 8–15
AN-109 Understanding and Preventing Latch-Up in CMOS DACs 8–23
AN-210 Adding Additional Input Channels to the AD7773/AD7775 8–31
AB-6 Single-Supply Operation of the DAC-08 and DAC-20 8–33
AN-225 12-Bit Voltage-Output DACs for Single-Supply 5 V and 12 V Systems 8–35
AN-310 Single-Supply 12-Bit DAC Circuits Using the AD7568 8–39
AN-314 14-Bit DACs Maintain High Performance Over Extended Temperature Range 8–43
AB-3 Digital Nulling of Precision Op Amps 8–49
AN-17 DAC-08 Applications Collection 8–51
AN-19 Differential and Multiplying D/A Converter Applications 8–63

Page

AN-343 Circuit Ideas for IC Converters . . . 8–71
AN-21 4–20 mA Digital to Process Current Transmitter . . . 8–75
AN-23 D/A Converter Generates Hyperbolic Functions . . . 8–77
AN-208 Understanding LOGDACs . . . 8–79
AN-137 A Digitally Programmable Gain and Attenuation Amplifier Design . . . 8–83
AN-142 Voltage Adjustment Applications of the DAC-8800 TrimDAC, An Octal, 8-Bit D/A Converter . . . 8–85
AN-219 Electronic Adjustment Made Easy with the TrimDAC . . . 8–99
AN-316 AD7224 Provides Programmable Voltages Over Varying Ranges . . . 8–103
AN-317 Circuit Applications of the AD7226 Quad CMOS DAC . . . 8–107
AN-318 AD7528 Dual 8-Bit CMOS DAC Application Note . . . 8–113
AN-319 Simple Interface Between D/A Converter and Microcomputer Leads to Programmable Sine-Wave Oscillator . . . 8–121
AN-320A CMOS DACs and Op Amps Combine to Build Programmable Gain Amplifiers – Part I . . . 8–127
AN-320B CMOS DACs and Op Amps Combine to Build Programmable Gain Amplifiers – Part II . . . 8–139
AN-321 3-Phase Sine Wave Generation Using the AD7226 Quad DAC . . . 8–149
AN-209 8th Order Programmable Low-Pass Analog Filter Using Dual 12-Bit DACs . . . 8–153
AN-322 Improve Function Generators with Matched D/A Converters . . . 8–161
AN-323 Microstepping Drive Circuits for Single-Supply Systems . . . 8–167
AN-324 Simple DAC-Based Circuit Implements Constant Linear Velocity (CLV) Motor Speed Control . . . 8–175
AN-325 12-Bit Analog I/O Port Uses AD7549 and 8051 Microcomputer . . . 8–181
AN-326 Interfacing the AD7549 to the MCS-48 and MCS-51 Microcomputer Families . . . 8–185
AN-327 DAC ICs: How Many Bits Is Enough? . . . 8–189
AN-206 Analog Panning Circuit Provides Almost Constant Output Power . . . 8–193
AN-207 Interfacing Two 16-Bit AD1856 (AD1851) Audio DACs with the Philips SAA7220 Digital Filter . . . 8–197
AN-328 Generate Four Channels of Analog Output Using AD7542 12-Bit D/A Converters and Control the Lot with Only Two Wires . . . 8–201
AN-329 Dynamic Performance of CMOS DACs in Modem Applications . . . 8–207
AN-37 Eight-Channel Codec Applications . . . 8–213
AN-38 Four-Channel Shared Codec . . . 8–221
AN-39 Companding D/A Converter . . . 8–227
AN-41 Improved Shared-Channel Codec Design with Analog Devices' Companding DACs . . . 8–239
AN-42 A 1 kHz, 0 dBm0 Standard Signal Generator . . . 8–247
AN-101 Cross-Plot Generator Allows Quick A/D Converter Evaluation . . . 8–249
AN-205 Video Formats and Required Load Terminations . . . 8–253
AN-330 Animation Using the Pixel Read Mask Register of the ADV47x Series of Video RAM-DACs . . . 8–257
AN-331 Changing Your VGA Design from a 171/176 to an ADV471 . . . 8–265
AN-332 Improved PCB Layouts for Video RAM-DACs Can Use Either PLCC or DIP Package Types . . . 8–267
AN-333 Design and Layout of a Video Graphics System for Reduced EMI . . . 8–269
AN-237 Choosing DACs for Direct Digital Synthesis . . . 8–285

Digital Signal Processing Products – Section 9 . . . 9–1

AN-334 Digital Signal Processing Techniques . . . 9–3
AN-344 Digital FIR Filters Without Tears . . . 9–19
AN-218 DSP Multirate Filters . . . 9–27
AN-217 Audio Applications of the ADSP Family (IIR Filters) . . . 9–37

Page

AN-336A Image Compression: Spelling Out the Options . . . 9–41
AN-336B JPEG Compression . . . 9–45
AN-227 Digital Control System Design with the ADSP-2100 Family . . . 9–49
AN-338 Loading an ADSP-2101 Program via the Serial Port . . . 9–61
AN-240 Considerations for Selecting a DSP Processor (ADSP-2101 vs. WE DSP16A) . . . 9–65
AN-233 Considerations for Selecting a DSP Processor (ADSP-2101 vs. TMS320C50) . . . 9–77
AN-235 Considerations for Selecting a DSP Processor (ADSP-21020/ADSP-21010 vs. TMS320C30) . . . 9–93

Instrumentation Amplifiers – Section 10 . . . 10–1

AN-244 A User's Guide to IC Instrumentation Amplifiers . . . 10–3
AN-245 Instrumentation Amplifiers Solve Unusual Design Problems . . . 10–11

Isolation Amplifiers – Section 11 . . . 11–1

AN-246 Isolation Barrier Precisely Passes Transducer Signals . . . 11–3

Multiplexers and Switches – Section 12 . . . 12–1

AN-355 Behind the Switch Symbol: Use CMOS Analog Switches More Effectively When You Consider Them As Circuits . . . 12–3
AN-248 ADG201A/202A and ADG221/222 Performance with Reduced Power Supplies . . . 12–5
AN-32 Single-Supply Operation of JFET Multiplexers . . . 12–9
AN-249 Overvoltage Protection for the ADG5xxA Multiplexer Series . . . 12–11
AN-35 Understanding Crosstalk in Analog Multiplexers . . . 12–13
AN-250 Bandwidth, OFF Isolation and Crosstalk Performance of the ADG5xxA Multiplexer Series . . . 12–21
AN-251 R_{ON} Modulation in CMOS Switches and Multiplexers; What It Is and How to Predict Its Effect on Signal Distortion . . . 12–27

Operational Amplifiers – Section 13 . . . 13–1

AN-356 User's Guide to Applying and Measuring Operational Amplifier Specifications . . . 13–3
AN-357 Operational Integrators . . . 13–11
AN-202 An I.C. Amplifier User's Guide to Decoupling, Grounding, and Making Things Go Right for a Change . . . 13–17
AN-358 Noise and Operational Amplifier Circuits . . . 13–25
AN-102 Very Low Noise Operational Amplifier . . . 13–35
AN-253 Find Op Amp Noise with Spreadsheet . . . 13–39
AN-107 Active Feedback Improves Amplifier Phase Accuracy . . . 13–41
AN-106 A Collection of Amp Applications . . . 13–49
AN-40 A Buffer Amplifier Applications Collection . . . 13–67
AN-25 Applying the OP-06 Op Amp As a High Precision Comparator . . . 13–71
AB-107 Two-Wire, 4–20 mA Current Transmitter . . . 13–75
AB-108 A Micropower Single-Supply Precision Rectifier . . . 13–77
AB-109 High Speed Precision Rectifier . . . 13–79
AB-111 Single-Supply Wien Bridge Oscillator . . . 13–81
AB-112 Single Resistor Controls Wien Bridge Oscillator Frequency . . . 13–83
AB-113 Precision Ramp Generator . . . 13–85
AB-114 Precision Current Regulator . . . 13–87

Page

AB-115 Micropower 1.23 V Bandgap Reference 13–89
AN-254 Add Programmable Gain, Attenuation 13–91
AN-255 Voltage-Controlled Amp Covers 55 dB Range 13–93
AN-359 Settling Time of Operational Amplifiers 13–95
AN-256 Accurately Testing Op Amp Settling Times 13–105
AN-108 JFET-Input Amps Are Unrivaled for Speed and Accuracy 13–107
AN-257 Careful Design Tames High Speed Op Amps 13–115
AN-110 OP-64 Advanced SPICE Macro-Model 13–125
AN-117 OP-42 Advanced SPICE Macro-Model 13–131
AN-120 OP-400 SPICE Macro-Model 13–137
AN-126 OP-260 Advanced SPICE Macro-Model 13–143
AN-132 OP-470 SPICE Macro-Model 13–149
AN-138 SPICE-Compatible Op Amp Macro-Models 13–155
AN-259 AD9617/AD9618 Current Feedback Amplifier Macro-Models 13–171
AN-360 Ask the Applications Engineer—6: Op Amp Issues 13–175
AN-281 Passive and Active Analog Filtering 13–177

Programmable Delay Generators – Section 14 14–1

AN-260 Using Digitally Programmable Delay Generators 14–3
AN-261 Programmable-Delay ICs Control System Timing 14–9

References – Section 15 15–1

AN-262 A Low Noise, Low Drop-Out Regulator for Portable Equipment 15–3

Resolver-to-Digital Converters – Section 16 16–1

AN-263 Resolver-to-Digital Conversion—A Simple and Cost Effective Alternative to Optical Shaft Encoders 16–3
AN-264 Dynamic Characteristics of Tracking Converters 16–7
AN-265 Circuit Applications of the AD2S81A and AD2S80A Resolver-to-Digital Converters 16–11
AN-252 Using the 2S80 Series Resolver-to-Digital Converters with Synchros: Solid-State Scott-T Circuit 16–17
AN-266 Passive Component Selection and Dynamic Modeling for the 2S80 Series Resolver-to-Digital Converters 16–21
AN-267 Vector Control Using a Single Vector Rotation Semiconductor for Induction and Permanent Magnet Motors 16–33
AN-230 Circuit Applications of the AD2S90 Resolver-to-Digital Converter 16–45
AN-234 Digital Resolver Integration 16–49
AN-236 Speed Measurement and Control with the AD2S83, a "Silicon Tachogenerator" 16–53

RMS-to-DC Converters – Section 17 17–1

AN-268 RMS-to-DC Converters Ease Measurement Tasks 17–3
AN-269 Extend RMS-to-DC Converter's Range 17–15

Page

Sample/Track-Hold Amplifiers – Section 18 18–1

AN-270 Applying IC Sample-Hold Amplifiers 18–3
AN-271 Build Precise S/H Amps for Fast 12-Bit ADCs 18–11
AN-27 Polarity Programmable Peak Detector 18–17
AN-204 Applications of the SMP-04 and the SMP-08/SMP-18, Quad and Octal Sample-and-Hold Amplifiers 18–19

Sensors – Section 19 19–1

AN-272 Accuracies of the AD590 19–3
AN-273 Use of the AD590 Temperature Transducer in a Remote Sensing Application 19–7

Sigma-Delta Converters – Section 20 20–1

AN-283 Sigma-Delta ADCs and DACs 20–3

Signal Conditioning Products – Section 21 21–1

AN-274 Taking the Uncertainty Out of Thermocouple Temperature Measurement (With the AD594/AD595) 21–3

Voltage-to-Current Converters – Section 22 22–1

AN-275 Versatile Transmitter Chip Links Strain Gages and RTDs to Current Loop 22–3

Voltage-to-Frequency Converters – Section 23 23–1

AN-361 Ask the Applications Engineer—3: V/F Converters 23–3
AN-276 Analog-to-Digital Conversion Using Voltage-to-Frequency Converters 23–5
AN-277 Applications of the AD537 IC Voltage-to-Frequency Converter 23–13
AN-278 Operation and Applications of the AD654 IC V-to-F Converter 23–31
AN-279 Using the AD650 Voltage-to-Frequency Converter As a Frequency-to-Voltage Converter 23–59

Miscellaneous – Section 24 24–1

AN-345 Grounding for Low and High Frequency Circuits 24–3
AN-214 Ground Rules for High Speed Circuits 24–7
AN-346 Understanding Interference-Type Noise 24–11
AN-347 Shielding and Guarding 24–15
AN-348 Avoiding Passive-Component Pitfalls 24–21
AN-349 Keys to Longer Life for CMOS 24–27
AN-311 How to Reliably Protect CMOS Circuits Against Power Supply Overvoltaging 24–29
AN-350 Take the Guesswork Out of Settling-Time Measurements 24–33
AN-351 Ask the Applications Engineer—2: Offset and Gain Trimming 24–43
AN-352 Ask the Applications Engineer—5: High Speed Comparators 24–45
AN-362 Ask the Applications Engineer—9: Decoupling; Component Selection 24–47
AN-353 Ask the Applications Engineer—10: PC Board Layout 24–49
AN-354 Ask the Applications Engineer—1: ADC Problems; Log Compensation Resistors 24–51
AN-280 Mixed Signal Circuit Techniques 24–53
AN-282 Fundamentals of Sampled Data Systems 24–85

Page

Appendix – Section 25 25–1

Analog-to-Digital and Digital-to-Analog Conversion Factors 25–2
Digital Codes 25–3
Settling Time (Expressed in Time Constants) Required to Achieve Specified Accuracy 25–4
Voltage Gain to dB Conversion 25–5
Worldwide Sales Directory 25–6

Index – Section 26 26–1

Product Index 26–3
Subject Index 26–7
Application Note (AN) Number Index 26–17
Author Index 26–21

Analog I/O Ports
Contents

		Page
AN-221	Power-Down Circuit Cuts Power to AD7769 and AD7774	2–3
AN-222	AD7569/AD7669 Operation with Offset Signal Grounds for Disk Drive Applications	2–5
AN-224	Input/Output Level Shifting with the AD7769	2–7
AN-284	Implement Infinite Sample-and-Hold Circuits Using Analog Input/Output Ports	2–11
AN-285	A Function Generator and Linearization Circuit Using the AD7569	2–15

ANALOG DEVICES

AN-221
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Power-Down Circuit Cuts Power to AD7769 and AD7774

by John Wynne

Both the AD7769 and the AD7774, I/O ports intended for HDD servo applications, operate off of +5 V and +12 V supply voltages. Although neither of these devices offer power-down as a standard feature, implementing such a function external to the devices is simple and inexpensive. Figure 1 shows the suggested circuit. The +5 V and +12 V supplies are switched off by means of MOSFETs, Q1 and Q2, in series with the supplies; however, the sequencing of these supplies going into power-down and coming out of power-down is also important, and this is handled by the cross-coupled NOR gates and their output delays. The diode protection schemes recommended in the AD7769 and AD7774 data sheets to protect against power supply mis-sequencing are now redundant when this power-down circuit is used.

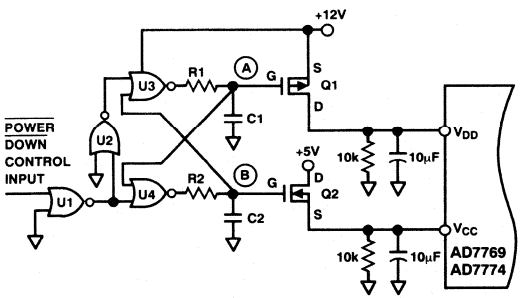

Figure 1. Power-Down Circuit with Correct Power Supply Sequencing

Figure 2 shows the power on/off control input and the resulting gate voltages, A and B. The P-channel Q1 is ON for a negative gate-source voltage while the N-channel Q2 is ON for a positive gate-source voltage.

In the power-down mode ($\overline{\text{POWER DOWN}}$ control input low) the output of U4 is low giving a Q2 gate-source voltage of V_{OL}, typically 50 mV with V_{CC} = 12 V on a CD4001, which is too small to turn on Q2. Similarly, the output of U3 is high giving a Q1 gate-source voltage of V_{OH} – 12 V, typically –50 mV, again too small to turn on Q1.

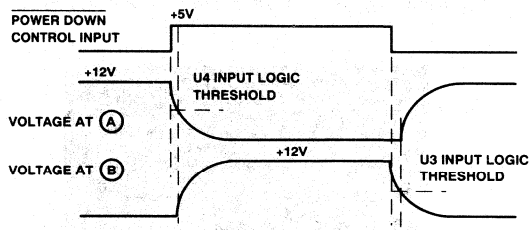

Figure 2. Typical Voltage Waveforms for Figure 1

Coming out of the power-down mode ($\overline{\text{POWER DOWN}}$ control input brought high) the output of U3 goes low and capacitor C1 discharges exponentially to V_{OL} with time constant R1C1 to turn Q1 ON. When the voltage level across C1 drops to the logic input threshold of gate U4, the output of U4 goes high and capacitor C2 charges exponentially to V_{OH}, approximately 12 V, to turn Q2 ON. Thus Q1 turns ON before Q2 turns ON.

Going into the power-down mode ($\overline{\text{POWER DOWN}}$ control input brought low) the RC delays work in reverse, capacitor C2 being discharged to the logic input threshold voltage of U3 before capacitor C1 is allowed to charge. Hence, Q2 turns OFF before Q1 turns OFF.

Power supply decoupling capacitors for the AD7769 and AD7774 should remain on the V_{DD} and V_{CC} pins of the devices. The 10 kΩ resistors from V_{DD} and V_{CC} to ground are not vital to the performance of the circuit but simply act to weakly pull the V_{DD} and V_{CC} pins to ground in order to discharge the decoupling capacitors. The CD4001 quad NOR gate must be powered from a 12 V supply which remains alive during power-down of the I/O port.

Figures 3 and 4 show actual V_{DD} and V_{CC} waveforms for the AD7769 when powering up (Figure 3) and powering down (Figure 4) using the circuit of Figure 1. For these photographs, R1 = R2 = 1 kΩ, C1 = C2 = 150 pF, Q1 = VP0300M and Q2 = VN0300M.

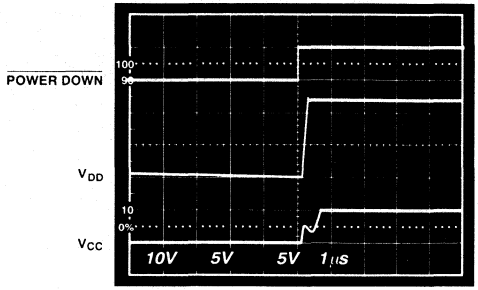

Figure 3. V_{DD} & V_{CC} Waveforms on Power-Up

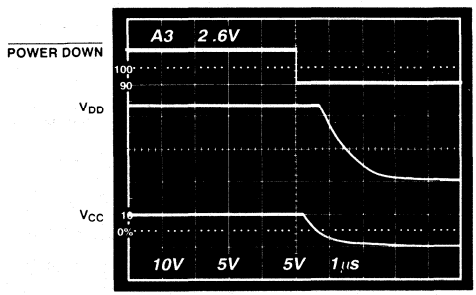

Figure 4. V_{DD} & V_{CC} Waveforms on Power-Down

Tables I and II demonstrate the effectiveness of the circuit in shutting off all power to the AD7769 and AD7774 I/O ports.

Table I. Effectiveness of Power-Down Circuit with AD7769

Control Input	I_{DD}	I_{CC}
High (Power On)	17.5 mA	1.94 mA
Low (Power Off)	3 µA	0

Table II. Effectiveness of Power-Down Circuit with AD7774

Control Input	I_{DD}	I_{CC}
High (Power On)	20.5 mA	3.24 mA
Low (Power Off)	3 µA	0

Additionally, under normal operating conditions, small voltage drops occur across the ON-resistance, $R_{DS}(ON)$, of transistors Q1 and Q2. Table III compares the voltage drops generated across two popular types of N-channel and P-channel MOSFETs measured using the AD7769.

Table III. Comparison of MOSFET Performance for Q1, Q2

Transistor	V_{DS}	Effective $R_{DS}(ON)$
Q1, VP0300M*	28 mV	1.6 Ω
Q2, VN0300M*	2 mV	1.03 Ω
Q1, ZVP2106A**	46.9 mV	2.28 Ω
Q2, ZVN3306A**	7.3 mV	3.76 Ω

*Siliconix Inc.
**Zetex, Inc.

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

AD7569/AD7669 Operation with Offset Signal Grounds for Disk Drive Applications

by Matt Smith

Single rail power supplies are the normal source of power for the electronics in many disk drive circuits. For a typical application such as head positioning, the sensor and control signals are often referenced to a voltage above ground potential. This is done in order to ensure that both the signal conditioning and the control circuitry is operating within its linear region. With any linear circuitry, such as op amps, there must be sufficient headroom between the signal extremes and the power supplies in order to maintain linear operation over the entire signal range. With single supply linear circuitry, there must be sufficient headroom between the minimum signal level and ground for the same reason. To achieve this, the linear circuitry is often operated with an offset or biased signal ground scheme which creates a pseudo signal ground above the system ground, and thereby provides the necessary headroom.

The sensor signals are normally converted to digital format for processing. When using standard ADCs, the offset signal level can result in a reduced dynamic range as several of the bottom codes are unused. Similarly, for the control signals which are generated using a DAC, a deadband can exist which again reduces the full dynamic range of the controller. This application note discusses how the AD7569/AD7669 analog I/O ports may be configured to operate with offset grounds and thereby maintain the full dynamic range of both the ADC and the DAC.

The AD7569/AD7669 is a complete monolithic 8-bit I/O system. A block diagram for the AD7569 is shown in Figure 1, while the AD7669 is shown in Figure 2. The AD7569 contains an internal voltage reference, a sample/hold amplifier, an 8-bit ADC, an 8-bit DAC, a buffer amplifier plus all the digital control circuitry. The AD7669 is similar with the addition of an extra 8-bit DAC. Both parts may be operated from a single 5 V power supply with a choice of 0 to 1.25 V or 0 to 2.5 V signal ranges for both the ADC and the DAC. For further information on both parts consult the data sheet.

To accommodate input signals which are not referenced to 0 V but to some offset voltage, V_{OFFSET}, it is necessary to bias the ground of the ADC, $AGND_{ADC}$, with the same voltage offset. This ensures that no unused codes or deadband will exist at the bottom of the ADC code table. Additionally, in order to ensure that the full dynamic range of the ADC is utilized, the maximum input swing should equal the full-scale input swing of the ADC (either 1.25 V or 2.5 V). Similarly the DAC ground, $AGND_{DAC}$ may also be biased up to provide an offset voltage range for the control circuitry. The AD7569/AD7669 has independent ground connections for the ADC, DAC and the digital circuitry. These are designated $AGND_{ADC}$, $AGND_{DAC}$ and DGND. These ground connections are not internally connected together. In most applications these would be externally connected together, but in this application they may be utilized to provide an offset voltage range for the ADC or the DAC, or both. A novel internal voltage reference structure, which derives the reference with respect to V_{DD} instead of the more usual ground, is used in the part. This allows the flexibility to bias up $AGND_{ADC}$ or $AGND_{DAC}$ independently without severely affecting the accuracy of the converters.

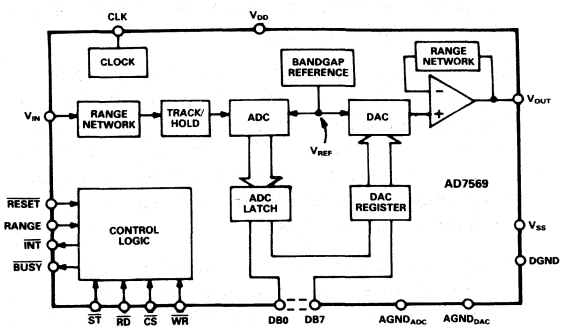

Figure 1. AD7569 Functional Block Diagram

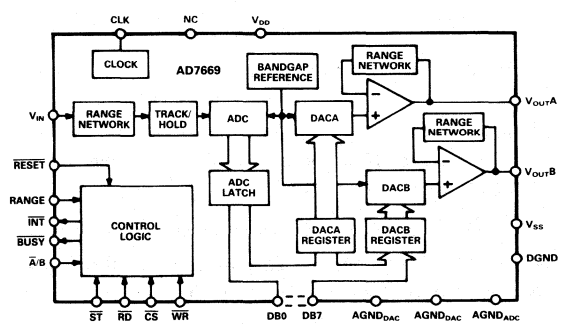

Figure 2. AD7669 Functional Block Diagram

Figure 3 shows a typical disk drive application circuit using the AD7669. The circuit implements a servo control loop using the ADC to monitor the head position while the DAC is used to control the position. Separate ground offset voltages are applied but these could be connected together if required.

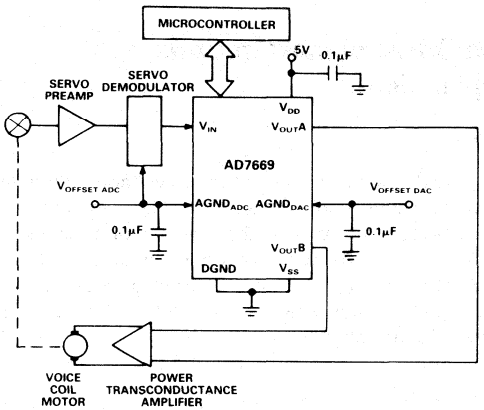

Figure 3. Servo Control Loop

The ADC input voltage range is shifted with respect to the system ground (DGND) by $V_{OFFSET\ ADC}$. This results in an offset transfer function as shown in Figure 4. The first code transition occurs at an input voltage of $V_{OFFSET\ ADC}$ + 0.5 LSB.

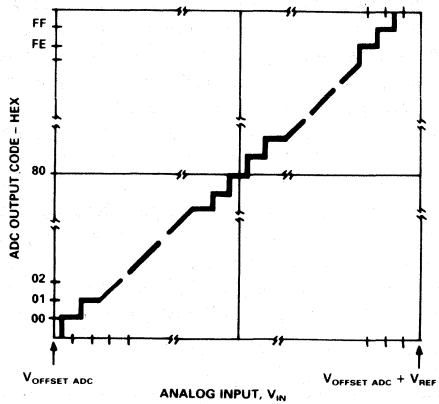

Figure 4. ADC Transfer Function with Offset Ground

Similarly, the $AGND_{DAC}$ may be offset and the transfer function for the DAC is shown in Figure 5. The output voltage may be expressed as

$$V_{OUT} = V_{OFFSET\ DAC} + V_{REF} \cdot D$$

where D is the fractional representation of the digital word. D may be set from 0 to 255/256.

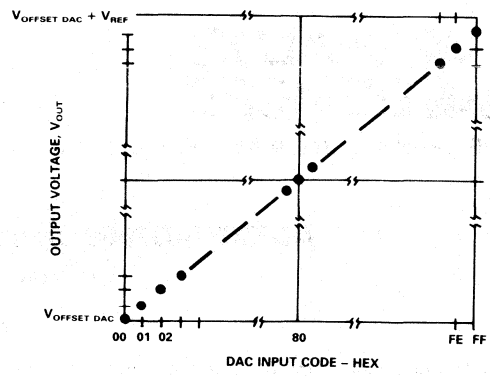

Figure 5. DAC Transfer Function with Offset Ground

For both analog grounds, a low impedance voltag source is required which can sink the currents flowing i the ground lines. With a 0.5 V bias, the $AGND_{ADC}$ curren is approximately 4 mA, and the $AGND_{DAC}$ current is ap proximately 2 mA. In order to maintain a lov dynamic impedance, the pseudo grounds should be de coupled to DGND using 0.1 μF capacitors. The capaci tors will absorb the current transients which occu during conversion.

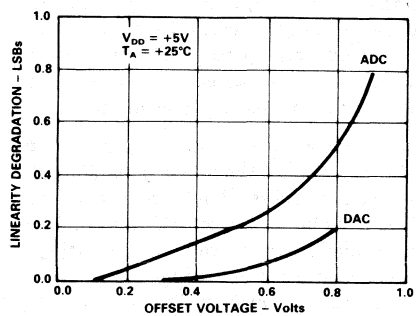

Figure 6. Accuracy vs. Offset Voltage

A plot of linearity error degradation for both the ADC and the DAC is shown in Figure 6. With a ground offse of 0.5 V, the ADC linearity error degradation is 0.2 LSBs while the DAC degrades by less than 0.1 LSB. The DNL error for both the ADC and the DAC remains well below 1 LSB thus ensuring no missed codes for the ADC anc monotonic DAC operation.

Larger ground offsets than those shown in Figure 6 may be applied at the expense of accuracy. If significantly larger ground offsets are desired, then another I/O port manufactured by Analog Devices, the AD7769, should be used. This part allows user programmable offsets over a much wider range. Please consult the data sheet for further details on this part.

Information furnished in this application note is based on results taken from three fabrication lots and is believed to be accurate and reliable. However no responsibility is assumed by Analog Devices for its use.

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Input/Output Level Shifting with the AD7769

by John Wynne

The AD7769 contains a fast, 8-bit sampling ADC with two input channels and two 8-bit DACs with output buffer amplifiers. See Figure 1. A unique feature of the device is the input and output signal conditioning circuitry (U.S. Patent No. 4,990,916) which allows the analog input and output signal voltages to be referred to a level other than analog ground. The input range and offset of the ADC and the output swing and offset of the DACs may be adjusted independently by the application of ground-referenced, positive control voltages; V_{BIAS}(ADC), V_{SWING}(ADC) for the ADC inputs and V_{BIAS}(DAC) and V_{SWING}(DAC) for the DAC outputs.

For example, with V_{BIAS}(ADC) = 6 V and V_{SWING}(ADC) = 2 V, the ADC can convert, with full 8-bit resolution, input signals which swing 2 V above and below 6 V. Note that both input channels have the same range. Similarly for the DACs, with V_{BIAS}(DAC) = 5 V and V_{SWING}(DAC) = 3 V, the output voltage of the DACs will swing from 3 V above to 3 V below 5 V. Again both DACs will have the same range.

However, there may be certain applications where the two input signals to be converted have different bias voltage levels or different signal swing ranges. Similarly, it may be necessary to generate DAC output signals which have different bias voltage levels or signal swing ranges. This application note suggests some simple circuits for these situations. Table I summarizes the level shifting operations possible and relates them to particular figures. Voltages $V_{LEVEL\ IN}$ and $V_{LEVEL\ OUT}$ refer to bias voltage levels which are different to V_{BIAS}(ADC) and V_{BIAS}(DAC), respectively. Similarly $V_{SWING\ IN}$ and $V_{SWING\ OUT}$ refer to voltages swings which are different to V_{SWING}(ADC) and V_{SWING}(DAC), respectively. Note that there are two different circuits when changing the signal swing levels at either the input or output. The choice of circuit depends on whether the swing level is being stepped up or stepped down.

	Condition	Circuit
Input Level Shifting	$V_{LEVEL\ IN} \neq V_{BIAS}$(ADC)	Figure 2
	$V_{SWING\ IN} < V_{SWING}$(ADC)	Figure 3
	$V_{SWING\ IN} > V_{SWING}$(ADC)	Figure 4
Output Level Shifting	$V_{LEVEL\ OUT} \neq V_{BIAS}$(DAC)	Figure 5
	$V_{SWING\ OUT} > V_{SWING}$(DAC)	Figure 6
	$V_{SWING\ OUT} < V_{SWING}$(DAC)	Figure 7

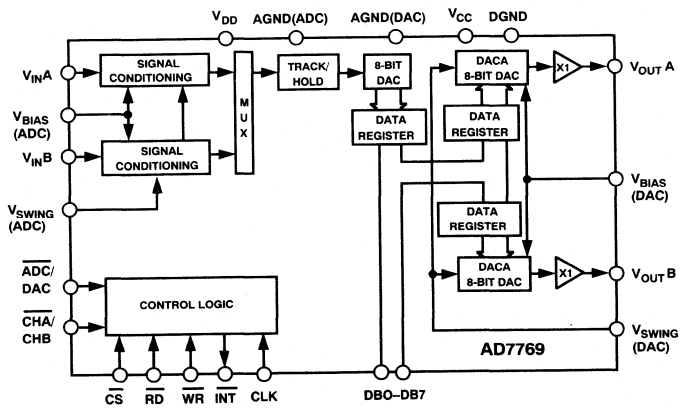

Figure 1. AD7769 Block Diagram

Input Level Shifting Circuits

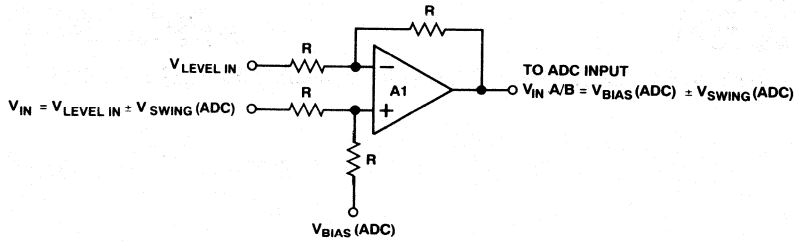

Figure 2. Changing from $V_{LEVEL\ IN} \pm V_{SWING}(ADC)$ to $V_{BIAS}(ADC) \pm V_{SWING}(ADC)$

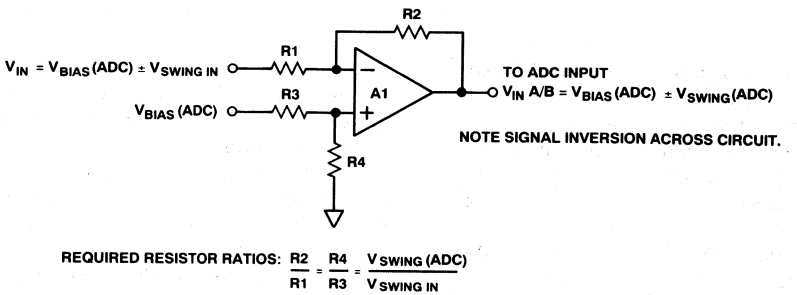

Figure 3. Changing from $V_{BIAS}(ADC) \pm V_{SWING\ IN}$ to $V_{BIAS}(ADC) \pm V_{SWING}(ADC)$ When $V_{SWING\ IN}$ Is Less than $V_{SWING}(ADC)$

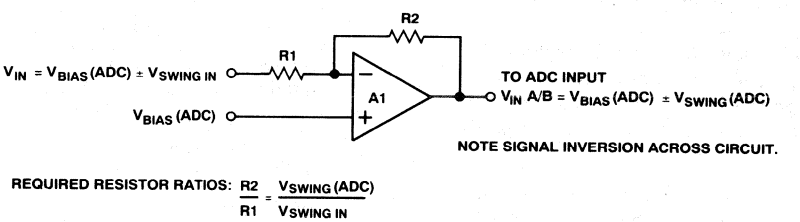

Figure 4. Changing from $V_{BIAS}(ADC) \pm V_{SWING\ IN}$ to $V_{BIAS}(ADC) \pm V_{SWING}(ADC)$ When $V_{SWING\ IN}$ Is Greater than $V_{SWING}(ADC)$

Output Level Shifting Circuits

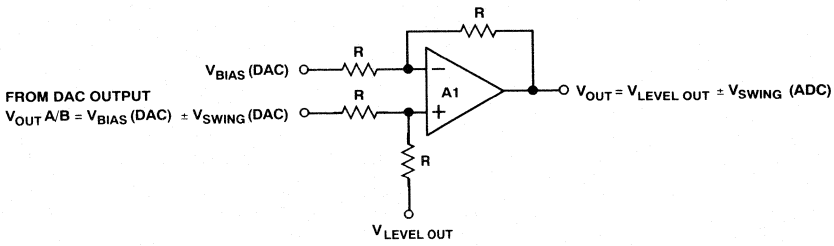

Figure 5. Changing from V_{BIAS}(DAC) ± V_{SWING}(DAC) to $V_{LEVEL\ OUT}$ ± V_{SWING}(DAC)

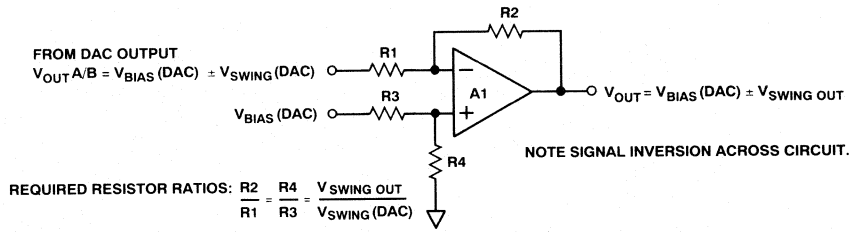

Figure 6. Changing from V_{BIAS}(DAC) ± V_{SWING}(DAC) to V_{BIAS}(DAC) ± $V_{SWING\ OUT}$ When $V_{SWING\ OUT}$ Is Greater than V_{SWING}(DAC)

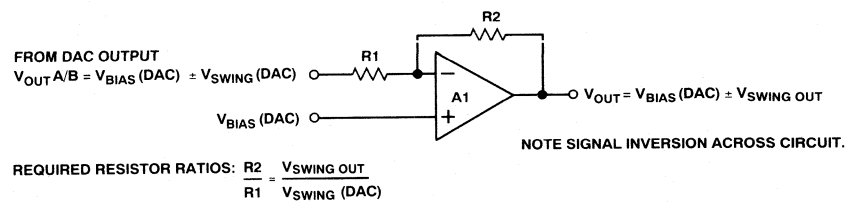

Figure 7. Changing from V_{BIAS}(DAC) ± V_{SWING}(DAC) to V_{BIAS}(DAC) ± $V_{SWING\ OUT}$ When $V_{SWING\ OUT}$ Is Less than V_{SWING}(DAC)

ANALOG DEVICES

AN-284
APPLICATION NOTE

NE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Implement Infinite Sample-and-Hold Circuits Using Analog Input/Output Ports

by Mike Byrne

ample-and-hold circuits are used in a variety of applications where the status of a parameter at a particular oint in time needs to be acquired and held for some me after the measurement instant. The term infinite ample-and-hold is given to sample-and-hold circuits nat are required to hold the acquired signal for a relavely long period of time. This type of circuit is particularly useful in monitoring changes in parameters over eriods of time versus a fixed or "held" reference level. his reference level may be updated occasionally, hence ne need for the sample-and-hold function. It is also seful in peak-detecting circuits where the held output rovides an analog representation of the peak analog nput voltage. This application note discusses the use of hree 8-bit analog input/output ports, the AD7569, D7669 and AD7769, in the implementation of infinite ample-and-hold circuits.

he block diagram for the AD7569 is shown in Figure 1, vhile the block diagram for the AD7669 is shown in igure 2. The AD7769 block diagram is illustrated in igure 3. The AD7569 contains on-chip track/hold, refernce, buffer amplifiers as well as an 8-bit DAC and an -bit ADC. The AD7669 is similar but contains two DACs nd one ADC. The AD7769 has two input channels and two output channels. The AD7569 and AD7669 operate from a single +5 V supply or from ±5 V supplies, while the AD7769 operates from +5 V and +12 V supplies. For more detailed information on these parts, see Table I and consult the respective data sheets.

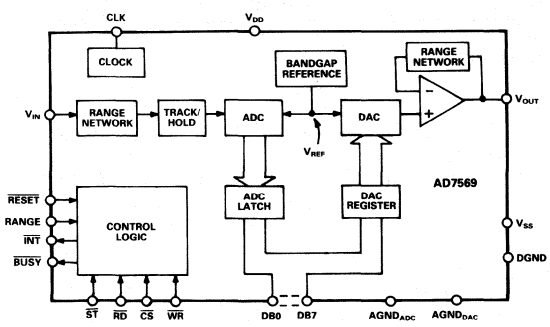

Figure 1. AD7569 Functional Block Diagram

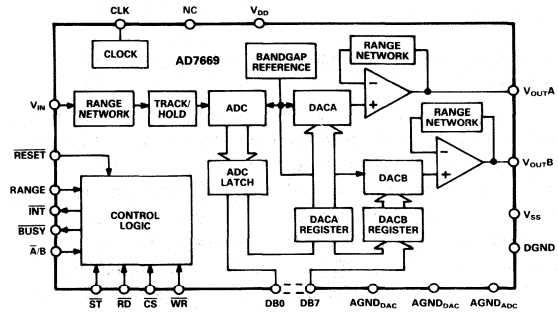

Figure 2. AD7669 Functional Block Diagram

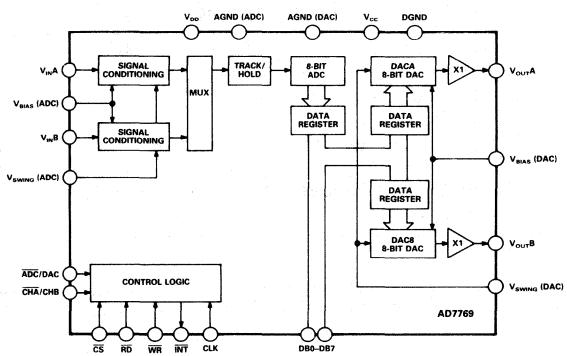

Figure 3. AD7769 Functional Block Diagram

PRINCIPLE OF OPERATION

The basic principle of operation of the infinite sample-and-hold circuits described here is the same regardless of which part is used to implement the function. The input signal to be held is converted to digital form using the on-chip track/hold and ADC. This digital data is restored to analog form using the associated DAC contained on the same chip. This scheme has a number of advantages over traditional schemes which store the value on a hold capacitor. Firstly, because the stored voltage is recreated using a DAC, the circuit does not suffer from any droop problems associated with storing voltages on capacitors for relatively long periods of time. Additionally, the entire sample-and-hold circuit comes in a single dual-in-line package with no external components required to implement the basic sample-and-hold function. Also, when the circuit is not being used for the sample-and-hold function, it provides a DAC and ADC function for the user.

INPUT/OUTPUT CHANNELS

The AD7569, AD7669 and AD7769 allow a number of different input and output voltage ranges, operate from different power supplies and contain different numbers of input/output channels. Table I provides a quick reference as to the differences between the parts to allow selection of the one most suitable for a particular application.

SINGLE-CHANNEL INFINITE SAMPLE-AND-HOLD USING AD7569

The circuit of Figure 4 illustrates the implementation of a single-channel infinite sample-and-hold circuit using the AD7569. No additional components are required to implement the function with the AD7569. The external sample signal drives the $\overline{CS}$, $\overline{RD}$ and $\overline{WR}$ inputs of the AD7569 (see timing diagram in Figure 5). The input signal is held and the ADC conversion is initiated on the falling edge of $\overline{CS}$ and $\overline{RD}$. Data from a previous conversion is output to the data bus. The $\overline{WR}$ input is also low

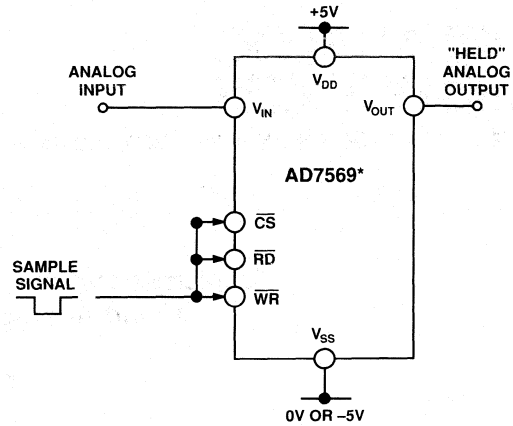

*ADDITIONAL PINS OMITTED FOR CLARITY

Figure 4. Single-Chip, Single-Channel, Infinite Sample-and-Hold Using AD7569

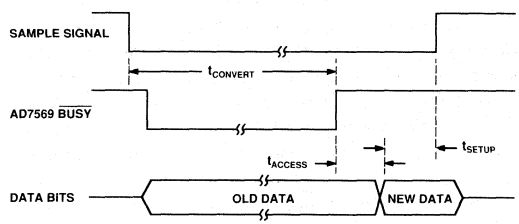

Figure 5. Timing Diagram for Circuit of Figure 4

	AD7569	AD7669	AD7769
Power Supplies	+5 V or ±5 V (for Bipolar Ranges)	+5 V or ±5 V (for Bipolar Ranges)	+5 V and +12 V
No. of Input Channels	1	1	2
No. of Output Channels	1	2	2
Input Voltage Ranges	0 to +1.25 V, 0 to +2.5 V, ±1.25 V, ±2.5 V	0 to +1.25 V , 0 to +2.5 V, ±1.25 V, ±2.5 V	V_{BIAS} (ADC) ± V_{SWING} (ADC)*
Output Voltage Ranges	Same as Input Ranges	Same as Input Ranges	V_{BIAS} (DAC) ± V_{SWING} (DAC)**
Conversion Time	2 μs	2 μs	2.6 μs

*V_{BIAS} (ADC) voltage range is +2 V to +6.8 V; V_{SWING} (ADC) voltage range is +2 V to +3 V.
**V_{BIAS} (DAC) voltage range is +3 V to +6.8 V; V_{SWING} (DAC) voltage range is +2 V to +3 V.

Table I. Selection Table

at this point, but since the $\overline{WR}$ input is rising edge triggered no data is written to the DAC register and nothing happens to the DAC output at this time. When the conversion is complete, data from the conversion is placed on the data bus; the sample signal is then brought high and data from the conversion is latched to the DAC register on the rising edge of $\overline{WR}$. The sample signal must be as long as the ADC conversion time plus the ADC data access time plus the data setup time required by the DAC register. Operating at +25°C means that the AD7569 requires a sample pulse of 2.12 μs. This is the effective acquisition time of the infinite sample-and-hold function, i.e., the time it takes the infinite sample-and-hold to acquire a new sample. This should not be confused with the acquisition time of the track/hold on the AD7569 which starts when $\overline{BUSY}$ goes high and is 200 ns typical.

The output from the DAC provides a low impedance output voltage source which corresponds to the sampled analog input. The output voltage will be held until the DAC register is updated. The full-scale matching between the ADC and DAC ensures a typical error of less than 1% between the analog input voltage and the "held" output voltage. If necessary, the gain and offset differences can be adjusted out during initialization or calibration cycles by adjusting the ADC offset and full-scale errors using an external op amp.

Using an external gate on the $\overline{WR}$ input, the ADC can continue to provide an A/D function while the held analog value is maintained on the DAC output. This gate would only allow the $\overline{WR}$ input to become active when the DAC output is to be updated with a new sampled value.

DUAL-OUTPUT-CHANNEL INFINITE SAMPLE-AND-HOLD USING AD7669

A similar scheme is used to implement an infinite sample-and-hold function using the AD7669 (see Figure 6). This time there are two output channels to which the digital data can be transferred. Selection of the output to be updated is achieved using the $\overline{A}$/B input of the AD7669. In the circuit of Figure 6, each successive conversion is stored to alternate DAC registers, i.e., the first to DAC A register, the second to DAC B register, the third to DAC A, and so on. Updating of the $\overline{A}$/B line takes place at the end of conversion on the rising edge of the sample input. A power-on reset on the 74HC74 (R1 and C1 in Figure 6) ensures that the first conversion result is stored in the DAC A register. An alternative method is to use an external control line to drive the A/B input which allows flexibility in choosing which output the conversion result is to be transferred to.

The timing requirements for the incoming sample pulse are similar to those outlined for the AD7569. The sample pulse width must again be as long as the ADC conversion time plus the ADC data access time plus the DAC register data setup time, resulting in an acquisition of 2.12 μs for the infinite sample-and-hold function.

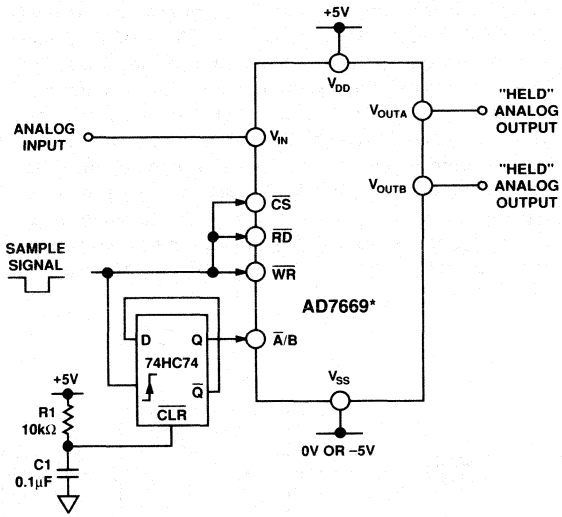

*ADDITIONAL PINS OMITTED FOR CLARITY

Figure 6. Infinite Sample-and-Hold Using AD7669

DUAL-CHANNEL INFINITE SAMPLE-AND-HOLD USING AD7769

The control logic required to turn the AD7769 into an infinite sample-and-hold differs slightly from those outlined for the AD7569 and AD7669. Unlike the previous cases where the same signal is used to drive the $\overline{RD}$ and $\overline{WR}$ signals of the part, separate signals must be generated here because the AD7769 $\overline{WR}$ input controls both the ADC conversion start and DAC register updates. The $\overline{ADC}$/DAC control input determines whether a DAC write or a conversion start takes place when $\overline{WR}$ is active. The $\overline{WR}$ line, therefore, must be pulsed twice in the infinite sample-and-hold application, first to initiate conversion on the ADC and second to transfer data to the DAC register. The $\overline{RD}$ signal to activate the ADC latches must be generated when conversion is complete. The $\overline{INT}$ output of the AD7769, which becomes active when conversion is complete, is used to generate this $\overline{RD}$ signal. Figure 7 shows the circuit used to configure the AD7769 as a dual infinite sample-and-hold.

The timing diagram for the circuit is shown in Figure 8. When the input sample signal is pulsed low, the $\overline{WR}$ input of the AD7769 goes low two gate delays later; and since the $\overline{ADC}$/DAC input is low, conversion is initiated on the rising edge of this signal. $\overline{INT}$ goes low when conversion is complete. One gate delay later, the $\overline{ADC}$/DAC line goes high; and a further gate delay later, the $\overline{WR}$ and $\overline{RD}$ inputs go low. $\overline{RD}$ going low places data on the data bus, and this data is written to the respective DAC register on the rising edge of $\overline{WR}$. $\overline{INT}$, and hence $\overline{RD}$ and $\overline{WR}$, is driven high by $\overline{RD}$ going low. The delay between $\overline{RD}$ and $\overline{INT}$ is dependent on the capacitance on the $\overline{INT}$ pin (C_L in Figure 7). A capacitor in the range 50 pF to 100 pF guarantees a $\overline{WR}$ pulse width which

gives sufficient time for the ADC data access time plus the DAC register data setup time. The $\overline{ADC}$/DAC line returns low one gate delay before the $\overline{WR}$ line goes high when writing data to the DAC register. However, internal delays on the AD7769 ensure that data is written correctly to the DAC register. The resulting acquisition time for the infinite sample-and-hold is a sum of the ADC conversion time plus the ADC access time plus the DAC register setup time (the time from the sample input going high to the AD7769 $\overline{WR}$ and $\overline{RD}$ inputs going high). In the circuit of Figure 7, the width of the AD7769 $\overline{WR}$ and $\overline{RD}$ pulses varies slightly with value of C_L givin slight variations in the acquisition time which is typicall 2.75 µs.

The AD7769 contains two input ADC channels and tw output DAC channels. Selection of either of the DA channels (or either of the ADC channels) is achieve using the $\overline{CHA}$/CHB input. This input is driven from a external control line and should not change for the dura tion of the ADC conversion and the DAC update. As result $V_{OUT}A$ will hold the $V_{IN}A$ signal while $V_{OUT}B$ wi represent the input voltage on $V_{IN}B$.

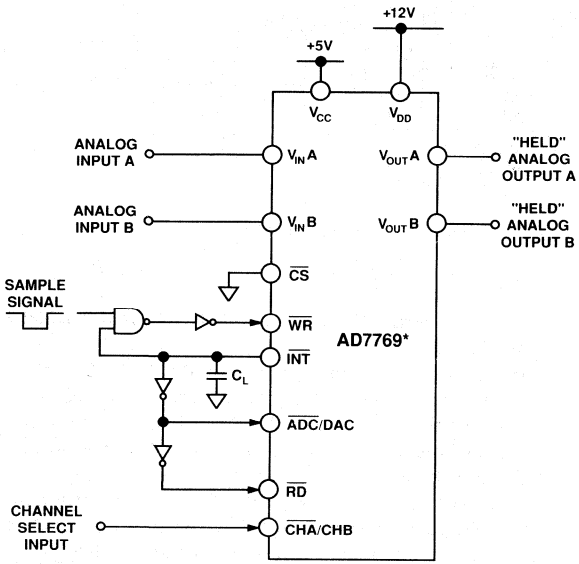

*ADDITIONAL PINS OMITTED FOR CLARITY

Figure 7. Dual Infinite Sample-and-Hold Using the AD7769

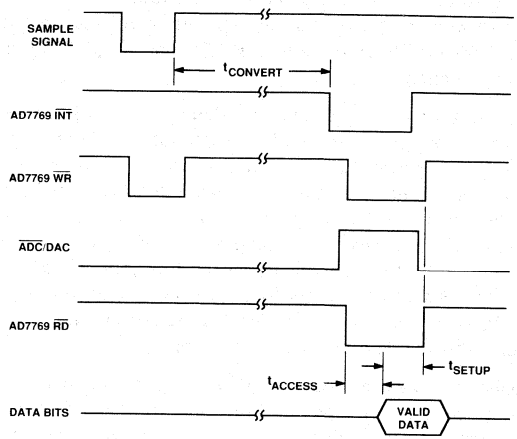

Figure 8. Timing Diagram for Circuit of Figure 7

ANALOG DEVICES

AN-285
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

A Function Generator and Linearization Circuit Using the AD7569

by James M. Bryant, MIERE

INTRODUCTION

A function generator is an electronic circuit which has an input/output relationship defined as a function of the form:

$$V_O = f(V_I) \quad \ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots (1)$$

(It is perfectly possible to have function generators with current inputs and outputs, or even transconductance or transimpedance function generators – with voltage in and current out or current in and voltage out, respectively – but the commonest case is the simple one of voltage in and voltage out.)

Function circuits are used in analog computations of all sorts to generate functions and are also used to linearize the responses of nonlinear transducers and circuits – this last application wil be described in detail towards the end of this application note.

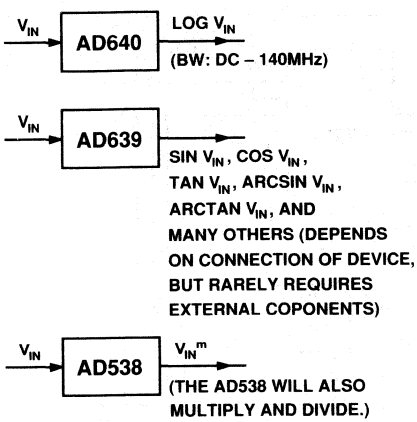

Figure 1. Some Monolithic Function Generators

Such circuits (some of which are shown in Figure 1) include logarithmic circuits such as the AD640, whose output is the logarithm of its input over a 45 dB dynamic range (and with a bandwidth of dc to over 140 MHz), trigonometric ones such as the AD639, which can be configured to have transfer functions corresponding to most trigonometric functions (including sine, cosine, tangent, contangent, versine, haversine, arcsine, arccosine and many more), and root and power circuits, which can be made with such devices as the AD538, and whose output is a power of the input of the form

$$V_O = V_{IN}{}^m. \quad \ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots (2)$$

where m may have any integer or noninteger value between 0.2 and 5. (The AD538 is much more versatile than this, being a log-amp based multiplier/divider/power circuit with an output of the form

$$Y\left(\frac{Z}{X}\right)^m$$

but in the present application note we are interested in function generators and not in more general computation.)

Equation (2) is not dimensionally correct since the output is a voltage, rather than (voltage)^m, and should strictly be expressed as

$$V_O = 1\,V\left(\frac{V_{IN}}{1\,V}\right)^m. \quad \ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots (3)$$

Where the function is more complex, or perhaps discontinuous, there is unlikely to be a ready-made integrated circuit that will perform the necessary operation. In such a case a function generator may be made from an analog-digital converter (ADC), a digital-analog converter (DAC) and a read-only memory (ROM). It is a function generator of this type that this application note describes.

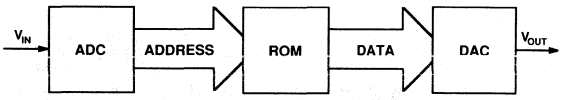

Figure 2. Basic "Convert/Lookup/Convert" Function Generator

The basic design of such a function generator is obvious: the input voltage is applied to the ADC and the resulting digital code is used as the address of a memory location in a ROM (Figure 2). The data word at that location is then applied to a DAC to produce an output voltage. With suitable ROM programming *any* function may be produced by such a circuit – the accuracy will depend on the resolution of the ADC and DAC.

THE AD7569 USED AS A FUNCTION GENERATOR

The AD7569 (which is illustrated in Figure 3) consists of an 8-bit ADC and an 8-bit DAC, with a integral voltage reference. It is available in a 24-pin "skinny-DIP" dual-in-line (DIP) package, a 28-pin leadless chip carrier (LCC) and a 28-pin plastic leaded chip carrier (PLCC).

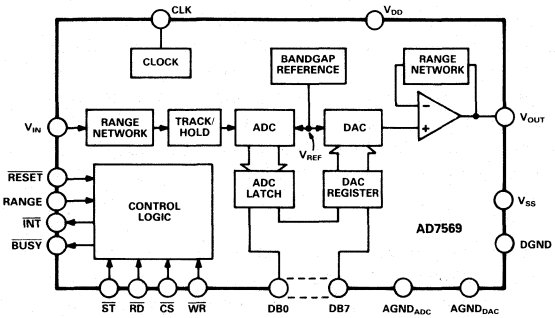

Figure 3. AD7569 Functional Block Diagram

The ADC contains an integral sample-and-hold (SHA) circuit and performs a conversion in 2 μs, and the DAC settles in 1 μs. Input and output ranges of 0–1.25 V, 0–2.5 V, ±1.25 V and ±2.5 V may be programmed by a range pin and the use, or not, of a negative supply. The ADC and DAC ranges are identical – it is not possible to program them separately. The table in Table I shows how these ranges (and the data formats associated with them) are programmed.

Range	V_{SS}	Input/Output Voltage Range	DB0–DB7 Data Format
0	0 V	0 to +1.25 V	Binary
1	0 V	0 to +2.5 V	Binary
0	−5 V	±1.25 V	2s Complement
1	−5 V	±2.5 V	2s Complement

Table I. AD7569 Ranges and Data Formats

The data output from the ADC and the data input to the DAC use a common data port. While it would be perfectly possible to multiplex this data onto the separate address and data buses of a normal ROM, it is obviously simpler, if possible, to use a ROM with a multiplexed address/data bus. The SC22102 EEPROM from Sierra Semiconductor is an 18-pin 256 × 8 device which is ideal for the application.

The input voltage is applied to the V_{IN} pin of the AD7569. The input impedance is quite low and the bias current quite large (and both vary with range – for details see the AD7569 data sheet), so it is well to use a buffer amplifier. It will also be necessary to use a buffer amplifier if the input range is not one of the ranges available from the AD7569.

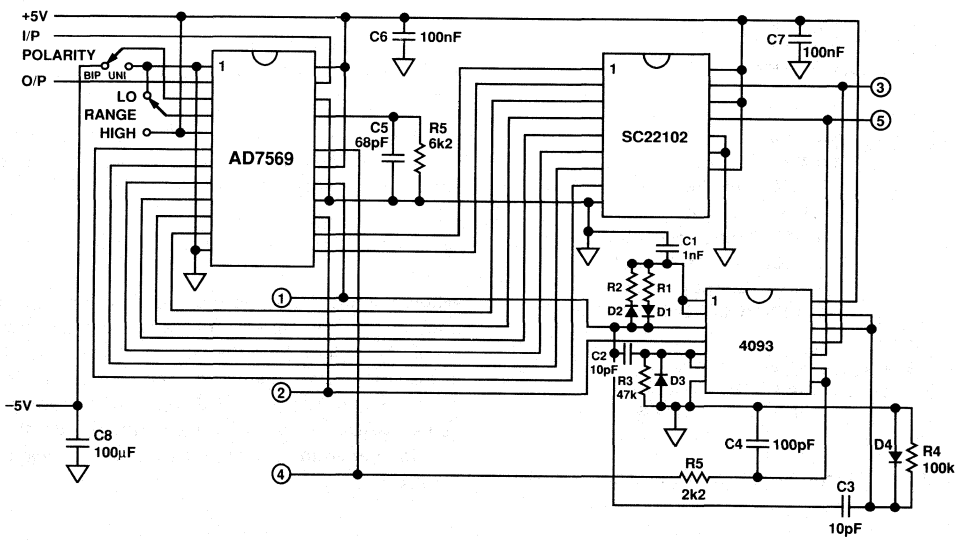

Figure 4. Circuit Diagram of Basic Function Generator

he ADC of the AD7569 performs an A-D conversion and ıe resulting 8-bit word is latched into the SC22102 EPROM as an address. The data byte at that address is ıen read back from the SC22102 into the AD7569 DAC vhere its analog value is output onto the V_{OUT} pin. \gain, the range available from the AD7569 is limited to he range used for the input, so if a different output ange is required, an output buffer with appropriate gain ınd level shifting must be used.

'he basic system uses only the AD7569, the SC22102, ınd a 4093 CMOS quad schmitt NAND gate which per-orms all the logic and timing functions (unused 4093 nputs are tied to +5 V).

Gate 1 of the 4093 acts as a relaxation oscillator. Its ɔutput consists of 3 μs negative pulses separated by ɔositive periods which must be at least 1 μs in duration. R1 sets the duration of the negative part of the cycle T1), and R2 sets the duration of the positive part (T2). The overall sampling frequency is 1/(T1+T2). The mini-mum sampling frequency depends on the bandwidth of :he input signal and should be 256 times the maximum component of the input signal – in practice, since the maximum sampling frequency is 250 kHz, this limits the maximum component of the input spectrum to about 1 kHz.

A-D conversion is triggered by the falling edge of the gate 1 output. Shortly afterwards (about 100 ns) the AD7569 $\overline{BUSY}$ line goes low. $\overline{BUSY}$ is applied to gate 4 of the 4093 after a delay of about 250 ns set by C4 and R5 (the ease with which delays and sharp edges may be produced is the reason for the use of the 4093 schmitt device rather than a simple quad NAND gate). The output of gate 4 is applied to the ALE input of the SC22102 – thus, 250 ns after $\overline{BUSY}$ goes high the result of the A-D conversion is latched into the address latch of the SC22102 (250 ns is much longer than is necessary but allows an ample margin for the rather inexact thresholds of the 4093).

When the output of gate 1 goes high again, two monostables are triggered. One, formed by C2, R3, D1 and gate 2, has a period of about 250 ns and is applied to the $\overline{WR}$ input of the AD7569; the other, formed by C3, R4, D2 and gate 3, has a longer period and is applied to the $\overline{OE}$ of the SC22102. This reads the contents of the memory at the latched address into the DAC. In most cases the output of gate 2 could be used top drive both $\overline{WR}$ and $\overline{OE}$, but a potential race condition is involved and since the fourth gate of the 4093 would otherwise be unused, it is safer to hold the memory output enabled until the DAC is safely latched.

The basic function generator circuit is shown in Figure 5 – it uses the three integrated circuits (AD7569, SC22102 & 4093), four diodes, five resistors, five capacitors (plus three supply decoupling capacitors), and two switches. In most practical applications the switches will not be necessary and the gain range required will be hard-wired. If the circuit will be used only in the unipolar mode, the negative supply and its decoupling capacitor, C8, are also unnecessary.

The waveforms at the main points of the system, indicated by circled numbers in the circuit diagram Figure 4, are shown in Figure 5.

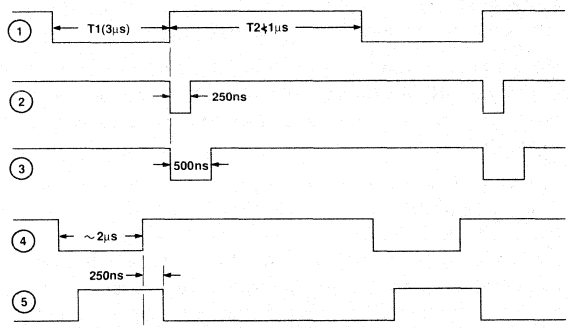

(1) RELAXATION OSCILLATOR WAVEFORM f=1/(T1+T2)
(2) MONOSTABLE 1 OUTPUT – DRIVES $\overline{WR}$ OF THE AD7569 DAC.
(3) MONOSTABLE 2 OUTPUT – DRIVES $\overline{OE}$ OF THE SC22102.
(4) AD7569 $\overline{BUSY}$ OUTPUT.
(5) BUSY DELAYED BY 250ns TO LATCH ADC O/P INTO SC22102 ADDRESS LATCH WHEN CONVERSION IS COMPLETE.

Figure 5. Control Waveforms of the Basic Function Generator

The circuit is so simple that no set-up adjustments are necessary unless the sampling frequency must be set particularly accurately – in this case R1 is set so that T1 is at least 500 ns longer than the BUSY signal of the AD7569 and T2 is then set to give the correct sampling frequency (T2 must be at least 1 μs). The timing of the remaining 4093 circuits is imprecise, but the component values chosen will work over the full range of 4093 thresholds.

PROGRAMMING THE SC22102

The data in the SC22102 EEPROM define the function of the function generator. The word at each of the 256 addresses defines the DAC output when the ADC output (and, of course, input) corresponds to that address.

The function is defined, therefore, by a lookup table programmed into the SC22102. Both the creation of the table, which depends on the function required, and the programming itself, which is done according to quite simple instructions given on the SC22102 data sheet, are beyond the scope of this application note. It is worthwhile mentioning, however, that the author has built a programmer for the SC22102 which plugs into the Centronics printer port of a personal computer and uses very simple software. Since the design and software depends on the computer used (a 12 year old Commodore PET with a heavily modified operating system – the unmodified PET does not have a Centronics printer port!), it is scarcely worthwhile to publish it – but the total design and construction time was less than two hours and a similar exercise could easily be done on most other PCs.

APPLICATIONS

In its simplest application a signal from a low impedance source is applied to the input of the basic circuit and an output is taken from its output. This is appropriate when the input and output ranges are identical and equal to one of the AD7569 ranges, and the signal source is low impedance. The resolution of both input and output being 8 bits, the accuracy will be 0.4% FS.

If the input and output range differ, or do not conform to one of the AD7569 ranges, or both, buffer amplifiers will be necessary. A single buffer can alter both span and offset of a signal.

If a signal has a span from 0 to V volts, or from −V to +V, then no offset is necessary and a simple amplifier, with suitable gain, between the signal source and the input of the ADC, or between the DAC and the output, is all that is required. If the gain required is greater than 1, a simple noninverting amplifier, using an op amp and two resistors, is all that is required. Such an amplifier is illustrated in Figure 6 (a).

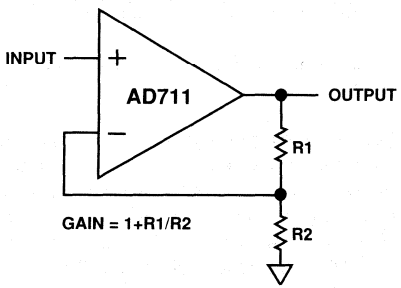

Figure 6a. Simple Noninverting Buffer

It is not possible to design such an amplifier with a gain of less than 1. In cases where the necessary gain is less than unity the input should be attenuated with two resistors, and a unity gain amplifier used to buffer the attenuated signal. Figure 6 (b).

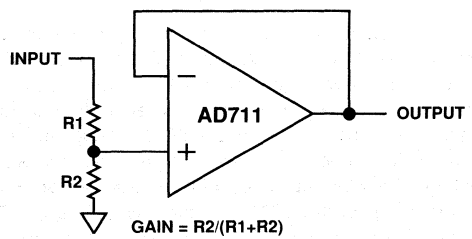

Figure 6b. Buffer for Gain ≤1. (R2 Omitted and R1 = 0 when G = 1)

If the signal span is from V_1 to V_2 ($V_1<V_2$), then both th amplitude and the offset of the signal must be altered. I this case let us assume that the AD7569 is set to it −2.5 V to +2.5 V range. The design and construction c the amplifier is simplest if we use an inverting buffer s that an input of V_1 gives $+2.5\ V_{OUT}$ and V_2 in give $-2.5\ V_{OUT}$. The offset required on this amplifier i $\frac{(V_1+V_2)}{2}$ and the gain is $-\frac{5}{(V_2-V_1)}$.

A suitable amplifier is shown in Figure 7. Its gain i $\frac{-R1}{R2}$, and its offset is $\frac{-V_R \cdot R1}{R3}$. These two formula allow us to choose values for R1, R2 and R3 and th correct polarity of V_R, the reference voltage (an externa voltage reference is required for this circuit because although the AD7569 contains a band gap voltage refer ence, it does not have its output available externally) The advantage of this circuit is that any value of gair may be set, greater or less than, or equal to, unity.

The disadvantages are that (a) the input impedance is only R2, which can only be a few tens of kilohms, and (b) the signal is inverted – so the lookup table in the ROM must also be inverted (i.e., V_1 will correspond to FS positive at the ADC and V_2 to FS negative). Neither is very important and (b) is trivial – provided the programmer remembers the inversion when calculating the EEPROM program!

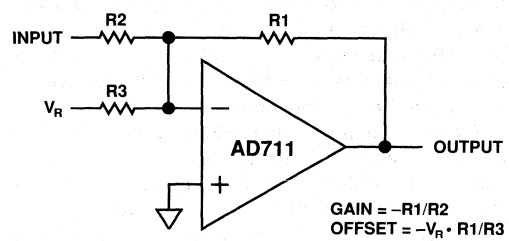

Figure 7. Inverting Amplifier. Gain May Be Set to Any Value <1< Offset is Independent of Gain

The same circuit can be used for the output. If the output range is V_1 and V_2 ($V_1<V_2$), then the gain required is $-\frac{(V_2-V_1)}{5}$, the offset is again $\frac{(V_1+V_2)}{2}$, and the resistors may be calculated as before. Again, when programming, remember that with this inverting amplifier +2.5 V from the DAC will give V_1 and −2.5 V will give V_2.

INEARIZATION

Suppose that the output of a temperature transducer is a voltage V_T defined by the law

$$V_T = V_1 \cdot T + V_2 \cdot T^2 + V_3 \cdot T^3$$

where T is temperature in degrees Celcius and V_1, V_2 and V_3 are constant voltages. If we are using this transducer to measure temperature and display it on a meter, we will find that the meter must have a nonlinear scale.

Many transducers and electronic circuits behave in this way. In some cases the nonlinearity is accepted by the user, but there is a steady demand for circuits which will linearize such responses – that is, in the case above, a circuit which accepts as an input V_T and has an output V_{T_1} such that

$$V_{T1} = V_K \cdot T$$

where the output voltage is proportional to the temperature.

If a circuit or transducer is very nonlinear, the basic function generator described above will be required to linearize it, and the accuracy will be that of the basic function generator (in the present case 8 bits or about 0.4%). Many circuits or transducers, however, are only slightly nonlinear, perhaps deviating from a straight line response by only a few percent. In this case it is evident that a correction to 0.4% is not great gain.

By suitable design, however, a function generator with 0.4% accuracy may be used to linearize such a slightly nonlinear circuit to much higher accuracy. The principle is shown in Figure 8.

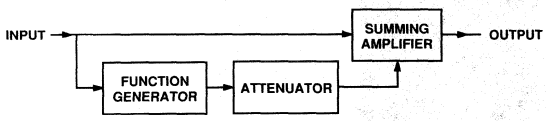

Figure 8. The Basic Principle of Linearization

If the maximum nonlinearity of the signal is ±1%, then a corection is applied to the signal from the DAC of the function generator. If the maximum possible correction is arranged to be ±1% (=2%), then the resolution to which the correction can be applied is 0.4% of 2% (0.008% or almost 14 bits).

Thus an 8-bit function generator, as described in this application note, can linearize a slightly nonlinear transducer to 14-bits or even more, with suitable summing circuitry. Of course if 14-bit performance is required, the accuracy and linearity of the remainder of the system must also be 14-bit and great care must be exercised in its design and the choice of components.

A practical circuit is shown in Figure 9. The use of two cascaded inverting amplifiers results from the difficulty of summing a signal into a unity gain buffer (of course many amplifier arrangements do allow this but the one in Figure 9 is one of the simplest). The design of input and output amplifiers for high accuracy linearization of this type is similar in principle to that described earlier, but considerably more care is necessary if amplifier offset and gain errors, and even thermal effects in resistors, are not to spoil the performance.

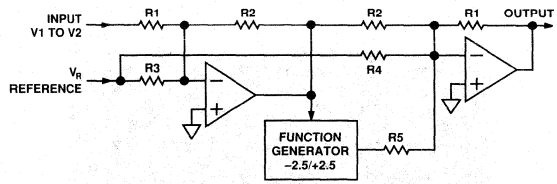

Figure 9. High Accuracy Linearizing Circuit

The resistor values are calculated as follows, assuming that the input and output range is V_1 to V_2 and that we are using end-point linearization (i.e., V_1 and V_2 are unchanged by the linearization procedure but intermediate values are corrected). We shall use the basic function generator in its −2.5 V to +2.5 V range.

We first of all choose a value for R1. This should not be so large that Johnson noise or offsets due to amplifier bias currents are objectionable, nor so low that it unnecessarily loads the signal source. A value in the range of 3–10 kΩ is likely to be suitable.

If $(V_1 + V_2)$ is zero then the voltage reference, R3 and R4 are unnecessary. Otherwise we must next choose a voltage reference, V_R, which must have the OPPOSITE polarity to $(V_1 + V_2)$. Any standard voltage reference circuit, such as the AD586 or AD587, is suitable, but the AD588 is particularly suitable when high stability is required – and it has the additional convenience of both positive and negative outputs.

Once R1 has been selected, R2 may be calculated.

$$R2 = R1 \cdot \frac{5\ V}{(V_2 - V_1)} \quad \ldots\ldots\ldots\ldots\ldots\ldots (4)$$

If necessary we then use V_R (inluding its sign – i.e., if V_R is −5 V then −5 V must be used in the equation, not +5 V), R1 and R2 to calculate R3 and R4.

$$R3 = R2 \cdot \frac{-2 \cdot V_R}{(V_1 + V_2)} \quad \ldots\ldots\ldots\ldots\ldots\ldots (5)$$

$$R4 = R1 \cdot \frac{-2 \cdot V_R}{(V_1 - V_2)} \quad \ldots\ldots\ldots\ldots\ldots\ldots (6)$$

Finally, to get the best possible resolution from the linearization circuit full-scale output from the function generator must correspond to the maximum error. Therefore if the maximum nonlinearity to be corrected is ±A% then

$$R5 = \frac{50 \cdot R1}{A} \quad \ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots (7)$$

When programming the function generator for linearization, one must set each data word so that when the input calls that address the output, added to the signal with the gain of R1/R5, brings the output of the circuit to its desired value. Where the nonlinearity obeys a predictable law the program in the function generator may be obtained by calculation but where it does not simple empiricism rules: the response of the circuit or transducer is measured at 256 points across its curve and the necessary correction at each point programmed into the EEPROM.

CONCLUSION

The function generator described in this application note is simple (using only three chips), easily constructed, and requires little or no setting up. It can perform as a function generator with accuracies of 0.4% FS and, with the addition of one or two functions to much higher accuracies. There are in development combined ADC/DAC circuits with higher resolutions (e.g., the AD7869 14-bit circuit) which will make higher resolution versions of the system as easy to implement – but for many everyday tasks 8-bit performance is more than adequate, and remarkably easy to achieve.

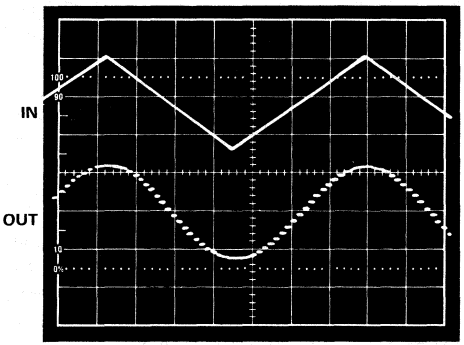

Figure 10a. Sin X

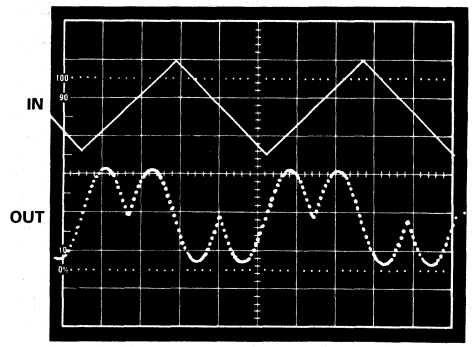

Figure 10b. Sin 2 X

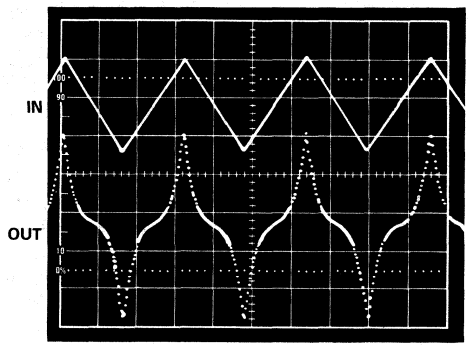

Figure 10c. X^3

Figure 10. Typical Function Generation

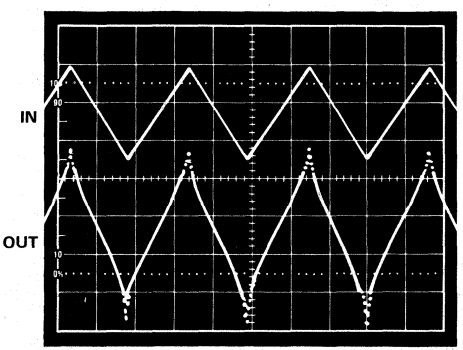

Figure 11a. Arc Sin X with Triwave Input

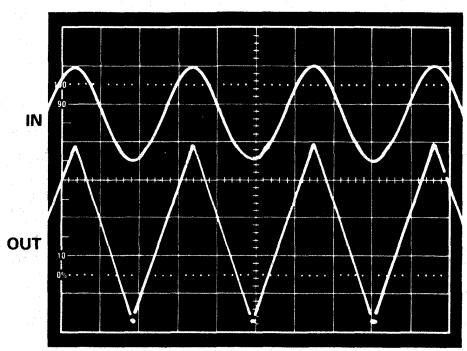

Figure 11b. Arc Sin X with Sine Wave Input Gives Linear Triwave Output

Figure 11. Linearization with a Function Generator

Analog-to-Digital Converters
Contents

		Page
AN-289	8-Bit A/D Converter Mates Transducers with μPs	3–3
AN-290	AD7578 and AD7582 Performance with Reduced V_{DD} Supply (12 V ± 10%)	3–11
AN-291	Asynchronous Clock Interfacing with the AD7878	3–13
AN-223	AD7575 Operation with an Offset Signal Ground for Disk Drive Applications	3–15
AN-292	Interfacing the AD7572 to High Speed DSP Processors	3–17
AN-293	The AD7574 Analog to Microprocessor Interface	3–21
AN-294	AD7672 Converter Delivers 12-Bit 200 kHz Sampling Systems	3–27
AN-295	Using Multiple AD1334s in Many-Channel Synchronous Sampling Applications	3–35
AN-296	Build a Single-Shot Recorder to Catch Fast Transients	3–37
AN-297	Test Video A/D Converters Under Dynamic Conditions	3–49
AN-298	AD671 12-Bit, 2 MHz ADC Digitizes CCD Outputs for Imaging Applications	3–59
AN-299	Simple Circuit Provides Ratiometric Reference Levels for AD782x Family of Half-Flash ADCs	3–67
AN-215A	Designer's Guide to Flash-ADC Testing Part 1: Flash ADCs Provide the Basis for High Speed Conversion	3–71
AN-215B	Designer's Guide to Flash-ADC Testing Part 2: DSP Test Techniques Keep Flash ADCs in Check	3–79
AN-215C	Designer's Guide to Flash ADC Testing Part 3: Measure Flash-ADC Performance for Trouble-Free Operation	3–85
AN-301	Undersampling Techniques Simplify Digital Radio	3–95
AN-302	Exploit Digital Advantages in an SSB Receiver	3–103
AN-228	Evaluation Board for the AD7884 16-Bit A/D Converter	3–109
AN-229	Evaluation Board for the AD7885 16-Bit A/D Converter	3–115
AN-368	Evaluation Board for the AD7701/AD7703 Sigma-Delta A/D Converters	3–121

ANALOG DEVICES

AN-289
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

8-Bit A/D Converter Mates Transducers with μPs

by Doug Mercer and Doug Grant

Through the innovative use of standard design concepts, a monolithic a-d converter simplifies the job of connecting transducers to a micropressor bus. The chip needs only +5 V.

Monolithic 8-bit analog-to-digital converters usually operate slowly, require one or more external components (such as a voltage reference or clock), and frequently draw their power from both positive and negative supplies. Even more important, they can generally handle only high-level input signals between 2 and 5 V. As a result, placing them between transducers and microprocessors is not an easy design task.

Bucking tradition is a successive-approximation analog-to-digital converter that not only performs an 8-bit conversion in 10 μs, but also runs from a +5-V power supply. An impressive roster of internal circuits permits system designers to drop the monolithic converter directly onto a microprocessor bus without adding any external components.

Unlike other successive-approximation devices, the AD670 converter contains a combined instrumentation amplifier-buffer—with a true differential input—which lies between the signal sources and the conversion circuitry (Fig. 1). With this dual-purpose component and the chip's precision digital-to-analog converter (which is based on pnp rather than the more customary npn transistors), the 670 can connect directly to a low-level floating-output signal source, such as a strain gauge or other type of transducers. Together these two circuits give the chip a full-scale sensitivity of 255 mV, or 1 mV/LSB.

These features would mean nothing to the system designer if the analog-to-digital converter were internally too complicated to manufacture. To build this complete analog circuit, innovative design techniques merged with precision bipolar processing, laser wafer trimming of thin-film resistors, and integrated injection logic (I^2L)—all of which are proved technologies for high-volume production. The linear-compatible I^2L is the key to incorporating dense logic functions on a bipolar chip hosting precision linear components.

The conventional successive-approximation a-d converter, whether bipolar or CMOS, must be externally driven by a fast op amp or some other low-impedance component. A bipolar converter relies on

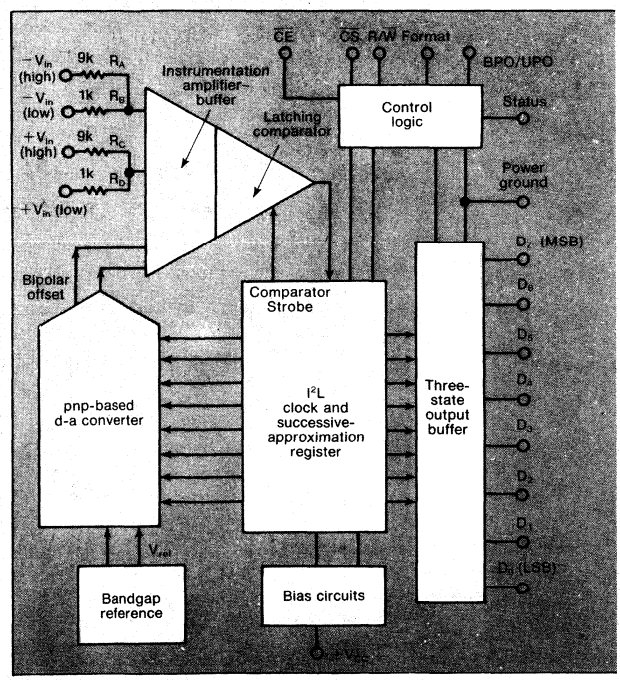

1. The instrumentation amplifier-buffer of the AD670 8-bit a-d converter works with millivolt-level transducer signals. The chip's design permits pin-strapping for a 1- or 10-Mv LSB. Two other pin-strapping or software-controlled options afford unipolar or bipolar operation, as well as two's complement or binary output coding.

Reprinted from ELECTRONIC DESIGN – January 12, 1984

the amplifier to eliminate the effects caused when test-bit currents from the d-a converter are fed back into the signal source. The CMOS version, which must keep its settling time at a minimum, uses the op amp to lower the charge or discharge time constants for the capacitors in the data-sampling comparators. With the 670, in contrast, the combination of input buffer and comparator effectively negates the need for a separate, external amplifier.

The front end of the 670 is resistive and typically needs only 200 nA of input bias current. If the source impedances for the positive and negative voltages are reasonably well balanced, an offset current of 5 nA enables the source impedance to rise as high as 100 kΩ.

The basic input ranges of the converter are +255 mV and from −128 to +127 mV. At each input, an attenuator (R_AR_B and R_CR_D) provides a 10-to-1 attenuation factor, which increases the full-scale ranges to +2.55 V and from −1.28 V to +1.27 V. When operating in this higher voltage range, the nominal input impedance is 10 kΩ. The input resistors supply overrange protection from latch-up for momentary shorts of up to ±30 V on either input.

Within the input section of the instrumentation amplifier is a voltage-to-current converter (Fig. 2). Vertical pnp transistors (Q_1, Q_2) and npn emitter-followers (Q_3, Q_4) achieve the ground-inclusive common-mode range. The differential input voltage is applied across resistor R_1, leading differential current to flow through resistors R_2 and R_3. The signal voltage between the emitters of Q_5 and Q_6 is now twice the input voltage, V_{in}.

The outputs from both the pnp-based d-a con-

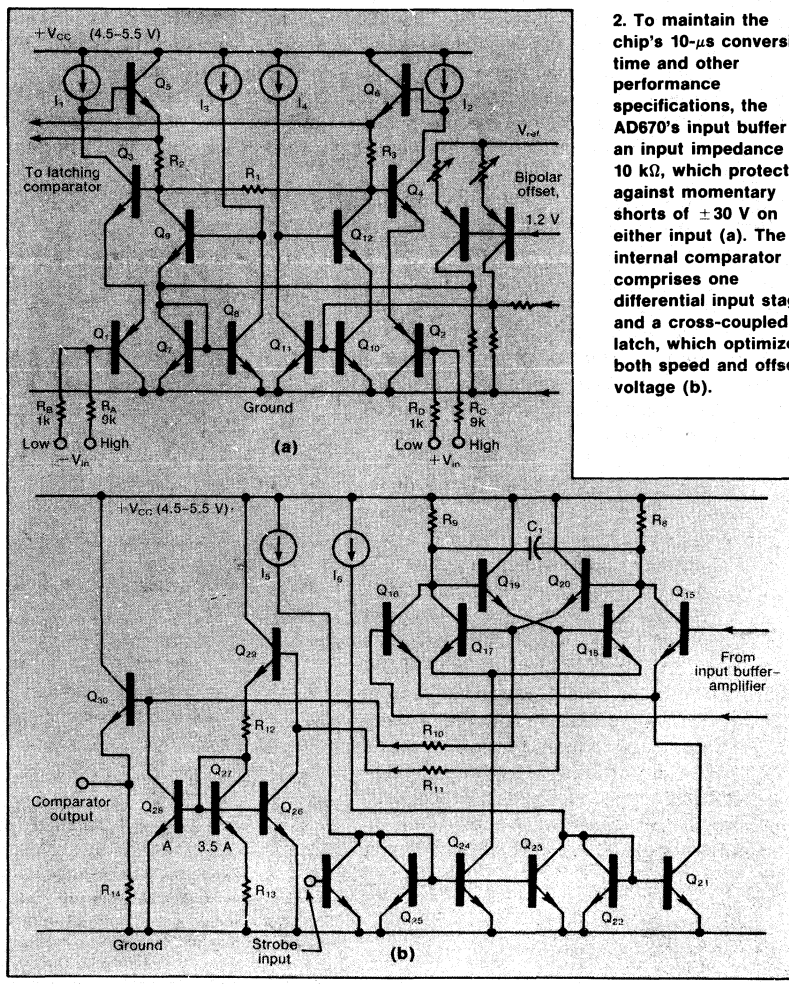

2. To maintain the chip's 10-μs conversion time and other performance specifications, the AD670's input buffer has an input impedance of 10 kΩ, which protects it against momentary shorts of ±30 V on either input (a). The internal comparator comprises one differential input stage and a cross-coupled latch, which optimize both speed and offset voltage (b).

verter and the bipolar current source are then reflected by the two Wilson current mirrors made up of Q_7-Q_9 and Q_{10}-Q_{12}. The output current from the mirrors is summed with the current in R_1, and the level-shifted difference between the amplifier's input and the d-a converter's output appears at the input to the latching comparator.

Since these signal levels are rather small (the full-scale signal at the comparator's input is only 512 mV), clamp diodes are no longer required. These small-signal conditions also increase bandwidth and eliminate the need to slew large voltages on stray capacitance.

A good heart

An a-d converter is only as accurate as its comparator. Therefore, like the input stage, the comparator in the 670 also embodies innovative design concepts (Fig. 2b). For example, the accuracy of the comparator lets the converter resolve a 1-mV LSB with less than 200 μV of code transition uncertainty. Its basic setup is exactly like that of the industry-standard 12-bit 574A a-d converter. The comparator's offset is adjusted by laser-trimming current sources I_1 and I_2, which are located in the input stage of the buffer amplifier.

The comparator consists of a single differential stage (Q_{15}-Q_{16}) and a cross-coupled latch (Q_{17}-Q_{20}), which is activated and deactivated by diverting the current from I_5 through an I^2L gate. The differential output of the latch is converted into a single-ended signal by Q_{26}-Q_{30} through resistors R_{10} and R_{11}. The output of the comparator swings through a temperature-dependent amplitude that is centered at 1.2 V, a value set by the ratio of R_{12} to R_{13}. The dependence on temperature is essential for the comparator to interface properly with the I^2L latches that drive the pnp-based d-a converter.

The d-a converter employed in this design comprises eight identical 100-μA current sources, each built with a lateral pnp transistor. They drive an R-2R ladder, which binarily weights each bit. The converter uses the identical current sources for each bit, thereby avoiding the need to match the base-to-emitter voltage and the α current gain of the lateral transistors. Each switch in the converter contains one common-base stage, yielding the lowest charge transfer and the fastest current settling time of any switch configuration. Consequently, the converter settles to within ½ LSB in less than 200 ns.

A ninth current source generates the bipolar offset current. The output current from the converter and the bipolar source is mirrored and scaled in the input buffer and summed with the differential current from the input voltage. The result is a full-scale voltage of 512 mV at the input of the comparator.

Table 1. Code conversions for the AD670

BPO/UPO input	Format input	Address	Code
0	0	Base	Unipolar straight binary
1	0	Base + 1	Bipolar offset binary
0	1	Base + 2	Unipolar two's complement*
1	1	Base + 3	Bipolar two's complement

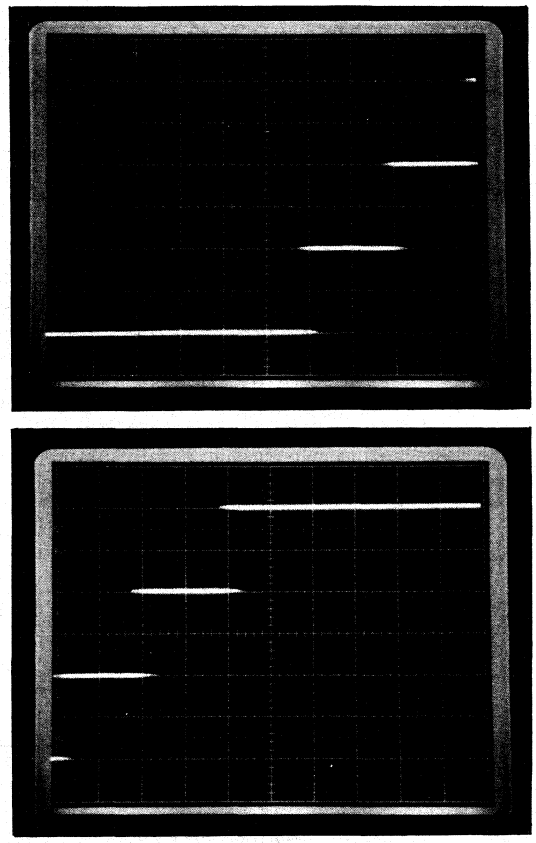

3. The 670's reference reduces noise so that millivolt-level inputs can be digitized. Linearity cross-plots of the zero and full-scale codes (top, bottom) show that noise levels caused by the comparator are less than 1/4 LSB. Noise is discerned by the overlapping between each trace step in the staircase, and the vertical lines mark off ½ LSB.

Another novel circuit in the a-d converter merges the classic functions of a voltage reference with those of its control amplifier. The dynamic loop behavior and dc performance of the all-in-one reference are superior to those achieved if the circuits were separate. Moreover, it substantially saves chip space.

Bandgap references are often thought of as hav-

ing more noise on their outputs than do buried zener diodes. However, a dynamic crossplot can demonstrate nonlinearity and noise at each bit-code transition. The plot is created by summing a small ac signal with an analog reference voltage at the converter's input while simultaneously dithering or sweeping around the output code of interest. A simple 2-bit d-a converter creates a four-step output, corresponding to the states of the two least-significant bits. The output is applied to the Y input of an oscilloscope, and the ac dither signal to the X input.

The reference in the 670 increases neither code-edge uncertainty nor noise significantly, as evident in linearity crossplots of the zero and full-scale codes (Fig. 3).

Getting on the bus

The microprocessor interface of the 670 consists of eight three-state output buffers, a Status output, and five control inputs, which are compatible with TTL, low-power Schottky TTL, and 5-V CMOS logic (see Fig. 1 again). The Chip Select and Chip Enable inputs ($\overline{CS}$ and $\overline{CE}$) are interchangeable and must be low to read data or to start a conversion.

Initiating a conversion entails a minimum pulse of 250 ns on $\overline{CS}$ and $\overline{CE}$, with the R/$\overline{W}$ line low. The normally low Status output will go high about 500 ns after $\overline{CE}$, $\overline{CS}$, and R/$\overline{W}$ go low, disabling the converter and telling it to ignore all inputs that might direct it to read or restart. If the same control lines are held low longer than the 10-μs conversion time, the 670 will convert continuously until one of the lines goes high. The rising edge of the high line must occur only when Status is high and not during the short times between cycles when the 670 is in a reset state prior to the next conversion cycle.

When $\overline{CS}$ or $\overline{CE}$ goes high, the data outputs re main in a high impedance state until a read cycle t the a-d converter is initiated by bringing $\overline{CS}$ and $\overline{CE}$ low, with R/$\overline{W}$ high. The outputs of the three-stat buffers supply data to the bus within 250 ns.

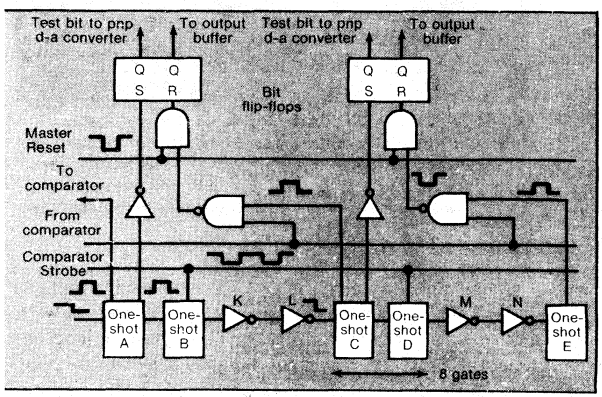

4. The successive-approximation register in the AD670 has no shift register, relying instead on the inherent gate delays of 80 cascaded inverters. The eight taps along the delay line generate both the set or reset bits as well as the strobe for the comparator.

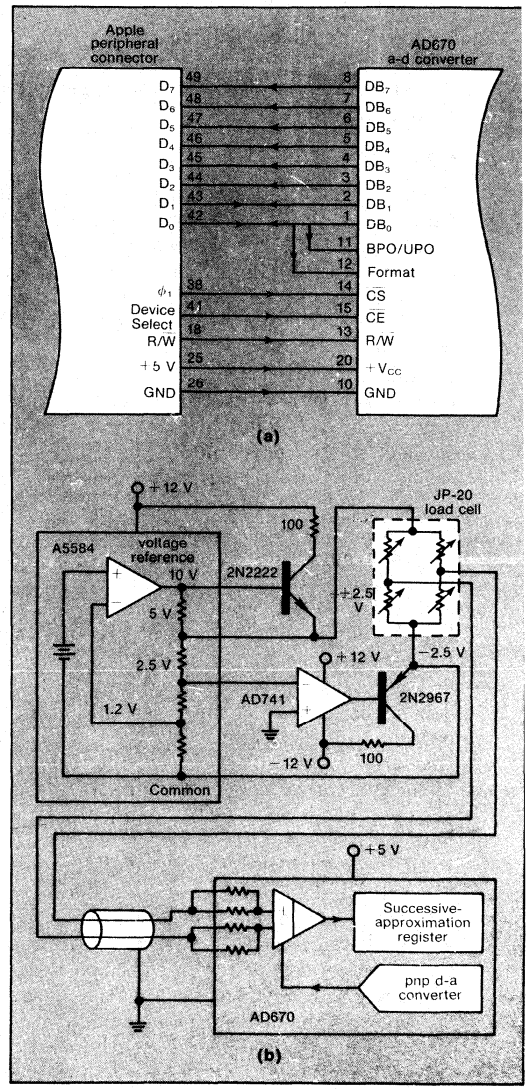

5. Through its microprocessor interface, the a-d converter links directly to the I/O bus of an Apple II (a). Thanks to the 670's differential input, the microcomputer can operate with signals from bridge circuits like load cells (b). The AD584 reference supplies a ±2.5-V excitation for a ±150-mV output when a force of ±20 lb is applied to the bridge. This is compatible with the 670's ±128-mV input range for a resolution of 2.1 oz/LSB.

During a write, or a start-conversion cycle, the Bipolar/Unipolar Offset input (BPO/UPO) and the Format input are latched into the converter. Tied to the data bus, these lines can be changed as each conversion commences. BPO/UPO controls the offset current: A logic 0 sets the converter to a unipolar range, and a logic 1 gives it a bipolar range (Table 1). The Format input selects either straight binary or two's complement as the form of the data output.

The inputs are latched on the first rising edge of $\overline{CS}$, $\overline{CE}$, R/$\overline{W}$, or Status, and the same edge initiates conversion, with the first cycle resetting the successive-approximation register. Conversion data may be read any number of times between conversions, but it is lost when a new conversion starts.

No circuit like an old circuit

Conventional successive-approximation registers are designed around shift registers and are basically static in nature. The speed at which they step from bit to bit is determined by an external clocking signal, which may range from dc to the maximum allowed by logic delays. (These delays can be significant in a linear-compatible I^2L process, since gate delays may be an order of magnitude longer than in, say, TTL.)

The register in the 670 is more dynamic because it has no shift register and needs no external clock to step through the successive-approximation algorithm. Instead, it relies on the inherent gate delay of the linear-compatible I^2L. The shift register is replaced by a delay line comprising 80 cascaded inverters, 10 for each bit.

The various pulses that set or reset bits and strobe the comparator are obtained by combining different taps along the inverter string. Sections of the delay line between the taps may be viewed as one-shots, whose outputs are delayed versions of their inputs. Strung one after another, the one-shots generate a series of properly timed pulses that execute the successive approximation. This technique is similar to the one used over 20 years ago in the first successive-approximation converters.

Looking at the delay line (Fig. 4), one-shot A unlatches the comparator, which samples the difference between the output of the pnp-based d-a converter and the 670's analog input. Once the d-a converter has settled, the rising edge of Comparator Strobe relatches the comparator.

Gates K and L delay the signal long enough to ensure that the comparator's output also has settled before gating it to the reset side of the bit flip-flop. A pulse coincident with the comparator's output sets the next bit, simultaneously switching one bit off and one bit on and thereby minimizing glitches in the output of the internal d-a converter.

Since the timing of the a-d converter is a function of gate delays, the design of its gates has been modified. Unlike the case with most I^2L circuits, the collector current rather than the injector current is controlled, supplying a first-order correction for transistor beta. As a result, gate delay is more consistent, ensuring a conversion time of no more than 10 μs.

Apple bites back

The control signals of the microprocessor interface make the 670 compatible with a variety of popular microcomputers, adding few if any extra parts. For example, when the chip is connected to an Apple II, the Device Select signal on the computer's peripheral connector serves as the converter's Chip Enable line (Fig. 5a). The $\overline{CS}$ line is connected to the computer's ϕ_1 clock, and the R/$\overline{W}$ lines are the same.

The simple, direct connections mean that no preamplifiers are needed for transducer-generated

Table 2. Applesoft driver for a-d converter

```
1 REM  PROGRAM IN APPLESOFT
2 REM  TO INTERFACE AD670
3 REM  8 BIT A/D CONVERTER
4 REM  TO DATA INSTRUMENTS
5 REM  MODEL JP-20 LOAD CELL
6 REM  WHICH HAS +/-150 mV
7 REM  OUTPUT FOR +/-20 POUNDS
8 REM  AND +/-2.5 VOLT INPUT
10  PRINT : PRINT : PRINT : PRINT
20  PRINT "TARE (T) OR WEIGH (W)";
30  INPUT A$
40  IF A$ < > "T" GOTO 100
50  PRINT : PRINT "CLEAR SCALE"
60  FOR I = 1 TO 1000: NEXT I
65  GOSUB 670
70 TARE = W
80  PRINT "TARE IS ";TARE;" POUNDS"
90  PRINT : PRINT : PRINT
95  GOTO 20
100 REM  ACTUAL WEIGHT ROUTINE
105 PRINT "PUT THE OBJECT ON THE SCALE"
110 GOSUB 670
120 NETWT = W - TARE
125  PRINT : PRINT : PRINT : PRINT
130  PRINT "NET WEIGHT IS ";NETWT;" POUNDS"
140  PRINT : PRINT : PRINT
150  GOTO 110
670  REM  THIS ROUTINE INTERFACES
671  REM  THE AD670 TO THE APPLE
672  REM  AND AVERAGES 100 READINGS
673  REM  THEN CONVERTS THE ANSWER
674  REM  TO POUNDS
680  POKE 49360,2
681  REM  THE 2 SETS THE AD670 FOR
682  REM  OFFSET BINARY OPERATION
683  REM  AND BIPOLAR INPUT
684 W = 0
685  FOR I = 1 TO 100
690 X = PEEK (49360)
695 X = X - 128
700 X = (X / 150) * 20
710 W = W + X
720  NEXT I
730 W = INT (W / 10)
740 W = W / 10
750  RETURN
1000 END
```

signals. For instance, the 670 can link directly to semiconductor strain gauges, pressure transducers, and load cells, which typically produce 30 mV full scale per volt of excitation. With this setup, microprocessor-based systems can monitor real-world parameters, for example, in another application, the a-d converter can mate with Data Instruments' JP-20 load cell (Fig. 5b). Excited by ±2.5 V, the transducer delivers ±150 mV for a force of ±20 lb, providing a good match when the 670 is operating in a ±128-mV range. Resolution is about 2.1 oz/LSB over approximately ±17 lb, but it can be scaled down to exactly 2 oz/LSB by trimming the excitation voltage generated by the AD584 reference.

An Applesoft Basic program demonstrates how the converter meshes with both the Apple II and the load cell (Table 2). Coincidentally, the subroutine at line 670 creates the actual interface for the 670 a-d converter. The POKE instruction (line 680) configures the device for a bipolar input with offset binary output coding. It also initiates the conversion cycle, which the 670 completes in 10 μs. The results can be read immediately after POKE by the PEEK instruction. (In Applesoft Basic, "immediately" implies a delay of about 1 ms.) In the interest of reducing the probability of unstable readings or noise, the microcomputer averages 100 readings—doing so without slowing the program's execution.

Let the computer do it

As analog system designers finally come to the realization that microprocessors are their friends not their foes, they can revel in the true symbiotic relationship between converters and the processors.

Because of the 670's high input impedance, external circuits can be employed to scale the analog input. By placing one-half of a dual CMOS 8-bit multiplying d-a converter (AD7528) between V_{in} and the 670's positive input and the other half between V_{ref} and the negative input (Fig. 6), a microprocessor can control the a-d converter's gain and offset values. D-a converter A, acting as an attenuator, controls gain. D-a converter B sets the offset by controlling the voltage from the AD580 precision voltage reference.

Both CMOS converters work in the voltage-

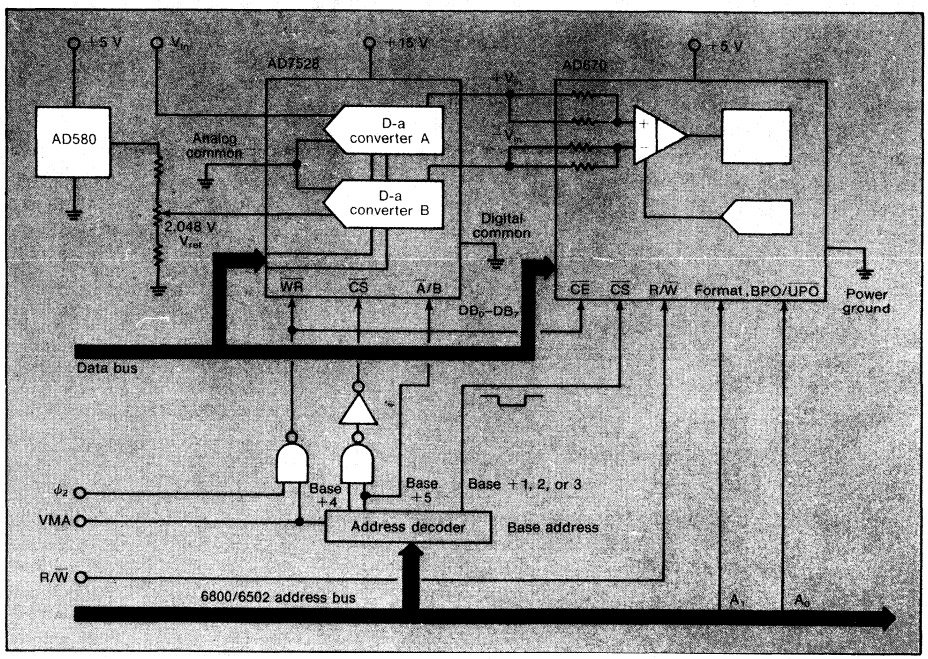

6. A dual 8-bit multiplying converter, the AD7528, controls the input levels of the 670 a-d converter. D-a converter A scales the voltage to the positive input of the 670, and converter B scales the voltage from the AD580 going to the negative input.

switching, or backward, mode, in which the reference inputs (pins 4 and 18) become the voltage outputs and the current outputs (pins 2 and 28) become the inputs. Their voltage swings about 2.5 V; thus they maintain an accuracy of 8 bits—more than enough accuracy for the 670, which has a common-mode range of about 2 V. Moreover, CMOS d-a converters working in this mode present a code-invariant impedance between their outputs and the analog common. Consequently, the offset caused by the input bias current of the 670 is constant, regardless of the code applied to the d-a converters.

The dual d-a converter and the a-d converter interface with a 6800/6502 bus. The address logic decodes a base address, setting A_0 and A_1 to a logic 0. The 670's $\overline{CS}$ input is low (true) for the base address, as well as for the base plus 1, 2, and 3.

The BPO/UPO and Format inputs are tied respectively to A_0 and A_1. Writing to the base address yields a unipolar offset with a straight binary output; to the base address plus 1, bipolar with offset binary; to the base address plus 2, unipolar with two's complement; and to the base address plus 3,

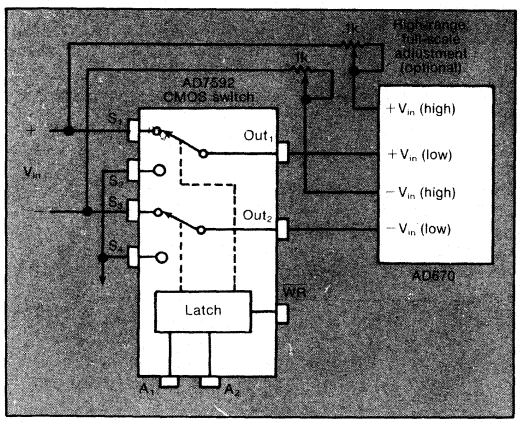

7. The gain ranges of the 670 can be regulated by a CMOS analog latching switch. When a microprocessor detects an overrange at 1-mV LSB, it automatically switches the converter to 10 mV/LSB.

bipolar with two's complement (see Table 1 again). The base address plus 4 loads d-a converter A, and the base address plus 5 loads d-a converter B. By reading any of the first four addresses, the microprocessor can find out the results of the a-d conversion.

D-a converter A scales down the positive input of the 670. A full-scale code (all 1s) goes through untouched with a gain of 255/256 (near unity), whereas a zero code cuts off the input completely. D-a converter B sets V_{ref}, which is connected to the negative input of the 670. The effective value of V_{in} is set by the following equation:

$$V_{in} = \frac{NC}{NA} 256 \text{ mV} + \frac{NB}{NA} V_{ref}$$

where:

NA = code fed into d-a converter A
NB = code fed into d-a converter B
NC = code generated by a-d converter
V_{ref} = 2.048 V

The gain and the offset of the 670 are controlled by software and can vary for different input signals.

Measurement can reach a resolution higher than 8 bits with a two-step conversion technique. Here converter A divides the input by eight and sets its code equal to 32; meanwhile converter B is setting its code to 0, and the 670 is completing a unipolar conversion. Converter B is then loaded with the result of that conversion, applying a fraction of V_{ref}, approximating V_{in}, to the 670's negative input. The fraction is then subtracted from the positive input voltage. If converter A is now set to full scale and the result of the subtraction undergoes a second a-d conversion (bipolar two's complement), the output of the 670 will resolve 1 mV over the input's common-mode range.

The V_{ref} input to d-a converter B could easily be changed to a second signal source rather than remain a fixed reference voltage. In this instance, the altered V_{ref} and V_{in} can be independently scaled by the dual d-a converter, with the difference between their outputs provided by the differential input stage of the 670, whose output is determined by the equation:

$$NC = \frac{(NA/256)\, V_A - (NB/256)V_B}{256 \text{ mV}}$$

where V_A and V_B are the input voltages for d-a converters A and B, respectively, and NC is the output code.

Switching gain-range

An analog CMOS switch, the AD7592, can shift the a-d converter's gain between high and low ranges (Fig. 7). The two positive inputs are tied together to obtain the converter's full-scale range of 0 to 255 mV. In the alternate position, the switch shorts the low V_{in} inputs to ground, creating a 10-to-1 division of V_{in} and changing the converter's range to 0 to 2.55 V.

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

AD7578 and AD7582 Performance with Reduced V_{DD} Supply (12V ± 10%)

by Dan Sheehan

The AD7578 is a single-channel, medium-speed, monolithic 12-bit CMOS A/D converter, which uses the successive approximation technique to provide a conversion time of 100μs. The AD7582 is a four-channel version of the AD7578, with similar specifications. Both devices feature very low offset voltage (typically less than 100μV), low gain error and easy interfacing to most 8- and 16-bit microprocessors using standard control signals ($\overline{CS}$, $\overline{WR}$, $\overline{RD}$).

The devices require three power supply voltages and a reference voltage, which have the following data sheet specifications: $V_{DD} = +15V \pm 5\%$, $V_{CC} = +5V \pm 5\%$, $V_{SS} = -5V \pm 5\%$ and $V_{REF} = +5V (\pm 5\%)$. However, the industry standard +15V (or ±15V) power supply is giving way to a new standard +12V (or ±12V) supply. The majority of modern personal computers (e.g., IBM PC*) have also adopted the 12V standard, with a tolerance of ±10% on this nominal figure. This note discusses the typical performance of the AD7578 (AD7582) with a reduced V_{DD} of 12V ±10% (V_{SS} and V_{CC} remain unchanged), over the operating temperature range of 0 to +70°C. The implied tolerances on V_{DD} throughout this application note are ±5% and ±10% for 15V and 12V respectively. A tolerance of ±5% is implied on V_{REF}. The relevant test circuits and test conditions used to generate the results are given in the data sheets.

ACCURACY

The internal 12-bit DAC and comparator determine the accuracy of the AD7578 (AD7582). The operation of both the DAC and comparator are influenced by reducing V_{DD}.

We will now look at the response of the various parameters with reduced V_{DD}.

Offset Error

This error results from the internal comparator circuit only. The offset voltage of the comparator changes with temperature and V_{DD}. However, the comparator circuit is an autozero type, with an autozero cycle occurring at the start of each conversion. This compensates for any offset voltage change and results in very low system offset voltages, typically less than 100μV.

The offset error of the AD7578 (AD7582) remains virtually unchanged with a V_{DD} of 12V versus 15V over the 0 to +70°C temperature range. The AD7578 (AD7582) meets the data sheet specification for offset error, of ±1/4LSB max, over these conditions.

Full-Scale Error

Full-scale error of the AD7578 (AD7582) remains practically unchanged for 12V versus 15V operation over the 0 to +70°C temperature range. The data sheet specification for full-scale error, of ±1/4LSB max, is met by the ADC over these conditions.

Endpoint Nonlinearity Error

This parameter is not separately specified on the data sheet but is included in the total unadjusted error (T.U.E.) specification. The internal DAC is the principal source of endpoint nonlinearity error for the AD7578 (AD7582), with reducing V_{DD}, because it is scaled and laser trimmed for operation with a V_{DD} of 15V and V_{REF} of 5V. The comparator has a negligible effect on this parameter with reducing V_{DD}. The endpoint nonlinearity error of the DAC, and thus of the AD7578 (AD7582), degrades (negative bow) by typically −1/2LSB with a V_{DD} of 12V (0 to +70°C). However, the negative bow can be avoided by reducing V_{REF} in proportion to V_{DD}, e.g., $V_{DD} = 12V$ and $V_{REF} = 4V$. Note that the previous results given for offset error and full-scale error apply equally for a V_{REF} of 4V or 5V, due to the inherently low errors of the ADC.

Total Unadjusted Error (T.U.E.)

This is a comprehensive specification which includes offset error, full-scale error and endpoint nonlinearity error. The T.U.E. specification given in the data sheet for 15V operation ($V_{REF} = 5V$, $T_{min} - T_{max}$) is ±1LSB max. This specification can be maintained over the 0 to +70°C temperature range, with 12V operation, by reducing V_{REF} to 4V. The T.U.E. specification must be increased to ±2LSBs max (±3/2LSB typ) with 12V operation and V_{REF} of 5V (0 to +70°C).

*IBM is a trademark of International Business Machines Corp.

Differential Nonlinearity

This parameter remains virtually unchanged with a V_{DD} of 12V versus 15V, over 0 to +70°C temperature range (V_{REF} = 4V or 5V).

Note that reducing V_{DD} below 12V, with a corresponding reduction in V_{REF} (less than 4V), does not preserve the accuracy of the AD7578 (AD7582) due to the decreasing LSB size. The smaller LSB size increases the susceptibility of the device to noise. Table I gives a summary of all the previous results.

V_{DD}	15 ±5%[1]	12 ±10%[2]		Units
V_{REF}	5V ±5%	5V ±5%	4V ±5%	
Resolution	12	12	12	Bits
Total Unadjusted Error	±1	±2	±1	LSB max
Differential Nonlinearity	±3/4	±3/4	±3/4	LSB max
	(No Missing Codes)			
Offset Error	±1/4	±1/4	±1/4	LSB max
Full-Scale Error	±1/4	±1/4	±1/4	LSB max
Channel-to-Channel Mismatch (AD7582 Only)	±1/4	±1/4	±1/4	LSB max

NOTES
[1]Data Sheet Specifications.
[2]Typical Specifications. Not Guaranteed.

Table I. Summary of Accuracy Results. V_{CC} = 5V ±5%, V_{SS} = −5V ±5%, T_A = 0°C to +70°C.

TIMING

In general, the various parameters of the AD7578 (AD7582) deteriorate with decreasing V_{DD}. However, any deterioration to the timing remains well within the guardbanded specifications given in the data sheet for the 0 to +70°C (KN grade) temperature range, with the exception of one parameter, $\overline{RD}$ to Valid Data (Bus Access Time). The specification, for the $\overline{RD}$ to Valid Data parameter, must be increased for 12V operation. The increases necessary are: from 200ns to 250ns for +25°C operation and from 240ns to 300ns for 0 to +70°C operation. This timing change is very small and should have no effect on most µP interfaces.

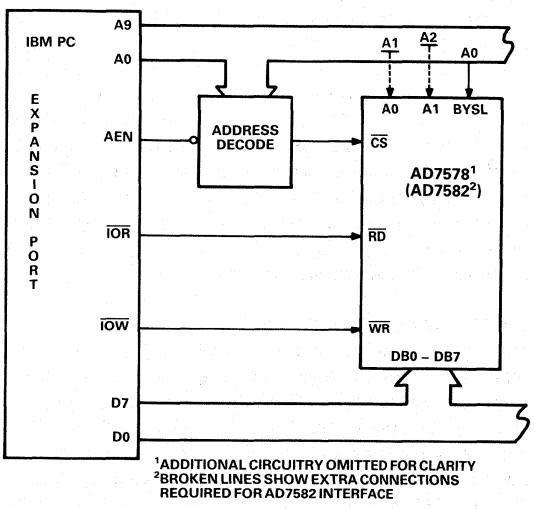

[1]ADDITIONAL CIRCUITRY OMITTED FOR CLARITY
[2]BROKEN LINES SHOW EXTRA CONNECTIONS REQUIRED FOR AD7582 INTERFACE

Figure 1. AD7578 (AD7582) to IBM PC Interface

IBM PC INTERFACE: The AD7578 (AD7582) is ideal for i plementing an Analog input port for the IBM PC. T AD7578 (AD7582) interface given in Figure 1 satisfies timing requirements of the IBM PC, with a V_{DD} of 1 (±10%).

You can assign I/O ports either to the memory addre space (memory mapped I/O port) or to the separate I/O a dress space, (isolated I/O port) depending on the read a write signals you select. In Figure 1 the computer uses $\overline{IOR}$ and $\overline{IOW}$ lines (versus $\overline{MEMR}$ and $\overline{MEMW}$) to treat t AD7578 (AD7582) as an isolated I/O port.

Ports in the I/O address space versus the memory addre space are easier to decode and troubleshoot because t computer uses only the lower 10 address lines (A_0 – A versus all 20 address lines. However, the decoding log for ports in the I/O space must accept the AEN signal mask out DMA cycles that also use the same addre space. You can use any language to execute the Read ar Write cycles that control the A/D converter. In Basic, yc use PEEK and POKE commands if the device is memo mapped, while INP and OUT statements apply if the co verter is operating as an isolated I/O device.

CONVERSION TIME

Conversion time, using internal CLK, increases with re ducing V_{DD}. It does not change for external CLK. The aver age percentage increase in the conversion time of a larg sample of parts, with a V_{DD} of 12V versus 15V, is 7% (10% maximum). The autozero cycle time remains virtually un changed (internal CLK) and does not affect the accuracy of the parts.

The conversion time increases because the LOW anc HIGH trigger levels of the Schmitt trigger circuit, whicl monitors the voltage on the CLK inputs, move slightly apart with reducing V_{DD}. This movement increases the charge time of the CLK capacitors from the LOW to HIGH trigger levels and thus the conversion time also. The remedy for the conversion time increase, with 12V operation, is to reduce the C_{CLK2} capacitor value from 560pF to 520pF. See Figure 6 in data sheet. This change has a minute effect on the autozero cycle time and does not affect the accuracy of the parts.

POWER SUPPLY CURRENT

I_{DD} is proportional to V_{DD}. The value of I_{DD} decreases with reducing V_{DD} voltage, and thus remains well within the specifications for 12V operation. I_{SS} and I_{CC} stay within their specifications with reducing V_{DD} and are altered only by V_{SS} and V_{CC} respectively.

LOGIC INPUTS AND OUTPUTS

The specifications for these parameters remain unchanged for reducing V_{DD} and are influenced by V_{CC} only.

POWER SUPPLY REJECTION

The data sheet specification for this parameter must be increased, from ±1/8LSB typ to 1/4LSB typ, with 12V (±10%) operation. The increase is due to the greater tolerance on V_{DD}.

ANALOG DEVICES

AN-291
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Asynchronous Clock Interfacing with the AD7878

by John Reidy

The AD7878 is a fast, complete, 12-bit A/D converter with a versatile DSP interface consisting of an 8-word, first-in, first-out (FIFO) memory and associated control logic. The FIFO memory allows up to eight samples to be digitized before the microprocessor is required to service the A/D converter. An on-chip status/control register allows the user to program the effective length of the FIFO and contains the FIFO out of range, FIFO empty and FIFO word count information.

Due to the complexity of the AD7878 internal logic, there are timing constraints which must be adhered to when performing read/write operations to the device. One method of abiding by these constraints is to use synchronous clock interfacing as recommended in the current data sheet. However, such interfacing is not suitable for every microprocessor. Another disadvantage is a possible reduction in throughput rate depending on the microprocessor clock frequency. This application note discusses the AD7878 timing constraints and shows an alternative interfacing method that is suitable for any microprocessor and will operate at any sampling frequency up to the maximum of 100 kHz.

The AD7878 is designed so that all internal logic operations are performed on a rising CLK IN edge, e.g., FIFO memory and control register updating and $\overline{\text{ALFL}}$ status output updating occur on a rising CLK IN edge. A read operation, an activity which is controlled external to the device, must be initiated around a falling CLK IN edge. Initiating a read operation (i.e., taking $\overline{\text{CS}}$ and $\overline{\text{DMRD}}$ low) on or near the rising edge of a CLK IN signal may cause the device to clock itself into an idle state with the only method of recovery being to reset the device.

There are two schemes which ensure correct timing operation when interfacing to microprocessors:

1. Synchronous clock operation.
2. Stopping the CLK IN input while reading/writing to the device.

Synchronous clock operation is covered extensively in the data sheet. Briefly, for synchronous clock operation the ADC clock must be the same as or an inverted version of the microprocessor clock out. As mentioned previously, $\overline{\text{CS}}$ and $\overline{\text{DMRD}}$ must not go low near a rising AD7878 CLK IN edge. Also, when writing to the device, the internal latch signal is the overlap between the CLK IN low signal and the $\overline{\text{CS}}$ and $\overline{\text{DMWR}}$ signals (see data sheet, page 7). To satisfy the above restrictions, the relative phase of the microprocessor CLK OUT signal and the memory control signals ($\overline{\text{CS}}$, $\overline{\text{DMRD}}$ and $\overline{\text{DMWR}}$) must be correct for synchronous operation. For example the ADSP-2100 CLK OUT can drive the AD7878 CLK IN directly; whereas, the TMS32010/32020 CLK OUT must be inverted before being applied to the AD7878.

A leading issue with synchronous interfacing is the processor CLK OUT frequency: frequencies greater than 8 MHz make synchronous interfacing impossible, e.g., ADSP-2100A and the TMS320C25 when operated at their maximum speed. Frequencies less than 8 MHz reduce the ADC's maximum throughput rate, examples of these are the TMS32010 and the TMS32020. Some processors such as the DSP56000 do not have a CLK OUT. The above issues can be avoided by using an alternative interfacing option as discussed below.

The second option is to "gate off" the ADC clock when reading/writing to the ADC. The following constraints apply when using this mode of interfacing:

1. $\overline{\text{DMRD}}$ cannot go low within 20 ns before a rising CLK IN edge or 5 ns after a rising CLK IN edge, see Figure 1.
2. When writing to the AD7878, the clock must be stopped in the low condition.
3. The clock must not be stopped during conversion or when $\overline{\text{CONVST}}$ is low.

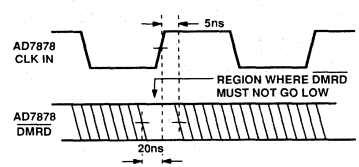

Figure 1. $\overline{\text{DMRD}}$ Asynchronous Timing Constraint

Figure 2 shows a general purpose asynchronous clock interface. The ADC clock is gated off with $\overline{CS}$ during read/write operations. To comply with constraint 1 above, the delay from $\overline{CS}$ to $\overline{DMRD}$ must be 5 ns plus the propagation through gate 1. A delay in series with the AD7878 $\overline{DMRD}$ input (shown in Figure 2) can be included to ensure a sufficient $\overline{CS}$ to $\overline{DMRD}$ setup time. The delay can be realized with a simple RC network or by inserting extra gates acting only as delays in the $\overline{DMRD}$ signal path.

The $\overline{CS}$ to $\overline{RD}$ (t_{CSRD}) delay for any processor may be found by making the following simple calculation:

$$t_{CSRD} = (t_{ADDR} - t_{RD}) - t_{PROP} \quad \ldots\ldots\ldots\ldots\ldots\ldots (1)$$

where: t_{ADDR} is the Address valid-time instant
t_{RD} is the processor $\overline{RD}$ output valid-time instant
t_{PROP} = Address decoder propagation delay

Assuming that the propagation delay of G1 is 5 ns, the $\overline{DMRD}$ delay (t_{DEL}) must be chosen such that $t_{CSRD}+t_{DEL} > 10$ ns. If t_{CSRD} is already greater than 10 ns, then the delay can be omitted. If t_{CSRD} is less than 10 ns, then $t_{DEL} = 10\text{ ns} - t_{CSRD}$.

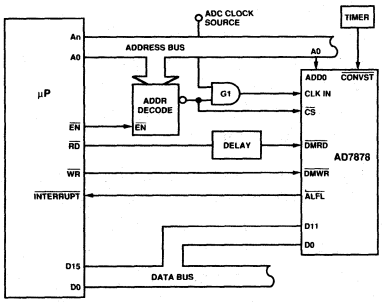

Figure 2. Asynchronous Clock Interfacing with the AD7878

INTERFACE EXAMPLES

Figures 3 to 5 show interfaces for the ADSP-2100A, TMS320C25 and the DSP56000. All three interfaces rely on a single wait state to relax the fast data access times and data setup times required by these DSP machines. The single wait state is hardware controlled for the ADSP-2100A and the TMS320C25. For the DSP56000 the wait state is software controlled by programming the bus control register.

The address valid to $\overline{RD}$ setup time ($t_{ADDR} - t_{RD}$ in Equation 1 above) can be calculated from the microprocessor timing specifications. Next, assuming an addresss decoder delay (t_{PROP}) of 10 ns and a gate propagation delay of 5 ns, it is possible to calculate the $\overline{DMRD}$ delay required for each processor interface. As an example take the ADSP-2100A, where the $t_{ADDR} - t_{RD}$ setup time is 4 ns, the $\overline{DMRD}$ delay required is 16 ns. The delay is realized with two gate propagation delays (2 × 5 ns) and a 6 ns, RC delay in Figure 3. Shown in Table I are the $t_{CSRD} - t_{RD}$ and the calculated t_{DEL} delays for the interfaces shown in this application note.

	ADSP-2100A	TMS320C25	DSP56000
$t_{ADDR} - t_{RD}$ (ns)	4	10	15
t_{DEL} (ns)	16	10	5

Table I. Processor Address Valid to $\overline{RD}$ Setup Times and AD7878 $\overline{DMRD}$ Delay Required

To the processor the interface is similar to synchronous interfacing. The ADC asserts the processor interrupt input when its internal FIFO memory has reached the preprogrammed word count, and then the conversion results are read from the ADC. Data transfers are not allowed during conversion.

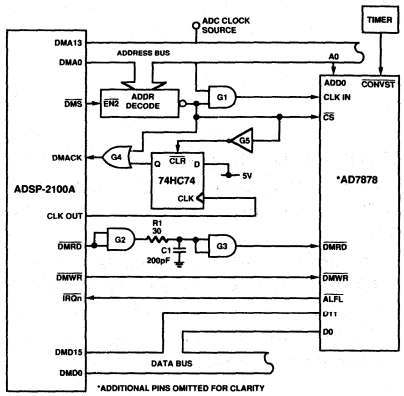

Figure 3. ADSP-2100A–AD7878 Asynchronous Clock Interface

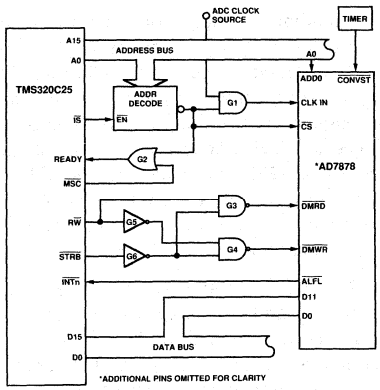

Figure 4. TMS320C25–AD7878 Asynchronous Clock Interface

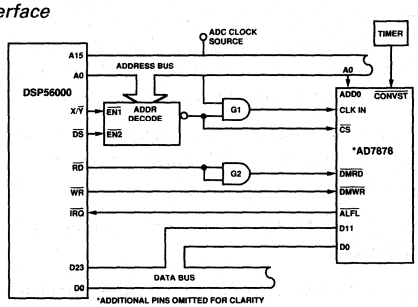

Figure 5. DSP56000–AD7878 Asynchronous Clock Interface

ANALOG DEVICES

AN-223
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

AD7575 Operation with an Offset Signal Ground for Disk Drive Applications

by Matt Smith

The AD7575 is a low cost, high speed, 8-bit sampling ADC with a built-in track/hold amplifier. It uses the successive approximation technique to achieve a fast conversion time of 5 μs. The internal track/hold amplifier allows full scale signals up to 50 kHz to be accurately sampled and digitized. Operating from a single +5 V supply, the device is ideal for single supply applications and has an input signal range from 0 V to 2 V_{REF}. The recommended reference voltage is +1.23 V giving an input range of 0 V to +2.46 V.

In many single supply disk drive applications, the signal conditioning circuitry is operated with a pseudo-ground. The pseudo-ground scheme is used to ensure linear operation of the conditioning circuitry over the complete input signal range. The signal swings around the pseudo-ground allowing adequate headroom between the signal extremities and the V_{DD}/0 V levels as illustrated in Figure 1.

In order to accommodate input signals which are referenced to a pseudo-ground, and yet maintain the full dynamic range, it is possible to bias the ADC signal ground (AGND pin) of the AD7575 to a potential above 0 V by an offset voltage, V_{OFFSET}. This results in an ADC input signal span of V_{OFFSET} to V_{OFFSET} +2.46 V. The pseudo-ground is, therefore, located midway between these signal extremes at V_{OFFSET} + V_{REF}. This corresponds to the midscale code in the ADC transfer function.

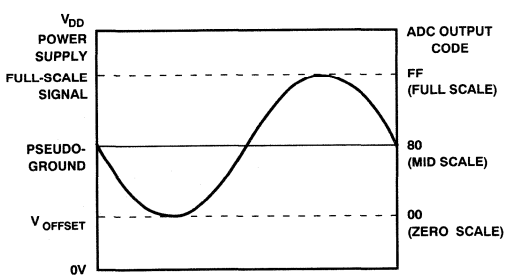

Figure 1. Pseudo-Ground Biasing Arrangement

The recommended biasing scheme for optimum performance from the AD7575 is shown in Figure 2. An AD589 provides the recommended reference voltage of +1.23 V. The AGND pin is biased up above the system ground using a single 5 V operational amplifier such as the TLC271. The amplifier is necessary in order to provide a low impedance signal ground return path. The amplifier must be capable of sinking the analog ground currents which flow in this line. In order to maintain a low dynamic impedance, both the V_{REF} pin and the AGND pin must be decoupled as shown in Figure 2. The optimum decoupling scheme uses 10 μF in parallel with 0.1 μF capacitors between V_{REF} and AGND terminals and a similar arrangement between AGND and DGND. The decoupling is essential in order to remove the ground current transients which occur during the ADC conversion process. Lower values of capacitance will result in degraded performance.

Using the circuit as shown in Figure 1, the performance of the part was evaluated under worst case conditions with V_{DD} = +4.75 V and with AGND bias voltages up to +0.7 V. With a bias voltage of +0.7 V, the input signal range of the ADC is from +0.7 V to +3.16 V or 1.93 V±1.23 V. No degradation in accuracy occurred with this bias voltage. Lower bias voltages may be used required but higher bias voltages may result in performance degradation.

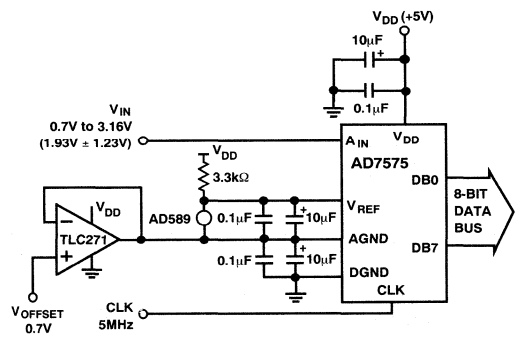

Figure 2. AD7575 AGND Bias Scheme

ANALOG DEVICES

AN-292
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Interfacing the AD7572A to High-Speed DSP Processors

by John Reidy

The AD7572A is a complete 12-bit CMOS analog-to-digital converter, with a conversion time of 3μs (7572AXX03) or 10μs (7572AXX10). An on-chip, buried Zener diode provides a stable reference voltage to give low drift performance over the full temperature range. An on-chip clock circuit is provided, which may be used with a crystal/ceramic resonator or an external clock source such as a divided-down microprocessor clock. The only other external components required for basic operation of the AD7572A are decoupling capacitors for the supply voltages and reference output.

The AD7572A is designed to interface with microprocessors as a memory-mapped device. Interface timing is sufficiently fast to allow the AD7572A to support today's most popular microprocessors. The more advanced processors such as the TMS32020 from Texas Instruments or ADSP-2100 from Analog Devices demand faster data access times than the AD7572A can meet when operating at their maximum clock frequency. This note addresses the problem of interfacing the AD7572A to such high performance processors.

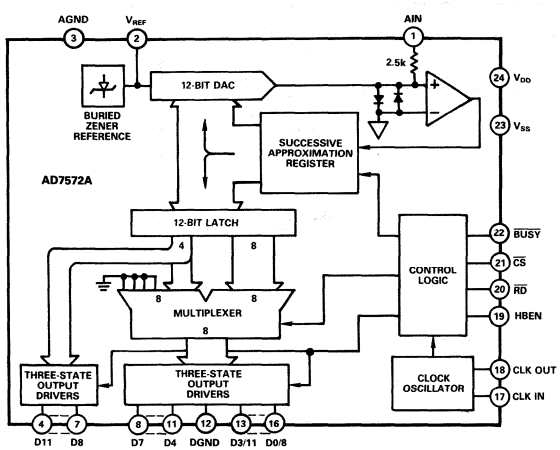

Figure 1. AD7572A Functional Block Diagram

THE USE OF WAIT STATES

A useful feature when interfacing ADCs to high speed microprocessors is a WAIT state input. This feature simplifies the hardware required for slow memory devices by extending the microprocessor bus access time, thereby, accommodating slower data access times. An example is the AD7572A/TMS32020 interface of Figure 2. One WAIT state (i.e., one clock cycle) is added to the microprocessor bus timing during an ADC data read cycle.

The AD7572A is mapped at a port address. The following I/O instruction both starts a conversion and reads the previous conversion result into a data memory location.

IN A, PA (PA = PORT ADDRESS)

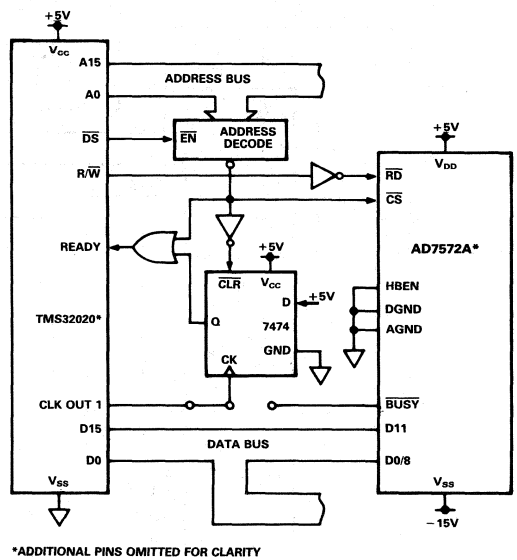

*ADDITIONAL PINS OMITTED FOR CLARITY

Figure 2. AD7572A/TMS32020 Interface Using WAIT States

If the previous conversion result is too old, then more up-to-date data may be read when conversion is complete. A delay at least as long as the ADC conversion time must be allowed between read operations. Note, this interface is not meant specifically for the TMS32020. It works equally as well for the ADSP-2100 or any other processor that has the same WAIT state input configuration.

The interface of Figure 2 is designed to extend the ADC data read cycle by just one WAIT state (i.e., one clock cycle or 200ns for a 20MHz TMS32020 clock input). The circuit can be easily modified to extend the ADC read cycle by the AD7572A conversion time. This is done by driving the 7474 clock input from the AD7572A $\overline{BUSY}$ output rather than the TMS32020 CLK OUT1. In this way once a conversion has started the TMS32020 READY input does not get asserted until the AD7572A conversion is complete. Stopping the microprocessor for the entire ADC conversion is costly in processing time, but it does have the advantage of reducing microprocessor noise.

MICROPROCESSOR NOISE

High resolution ADCs, like the AD7572A, are sensitive to the noisy environments created by high speed microprocessors. Careful attention should be given to a PCB layout and power supply decoupling. A continuously active microprocessor bus can cause LSB errors in the conversion results. These errors are due to feedthrough from the microprocessor to the ADC comparator. The problem of microprocessor noise can be eliminated by using WAIT states for the duration of the ADC conversion time as suggested in the AD7572A/TMS32020 interface. Alternatively three-state buffers can be used to isolate the ADC data bus as shown in the ADSP-2100 interface of Figure 3.

EQUAL SAMPLING INTERVALS

In Digital Signal Processing, and many other applications, it is important that the signal sampling occurs at exactly equal intervals. This is to avoid errors due to sampling uncertainty or jitter. Trying to achieve precise timing with a microprocessor necessitates counting clock cycles and calculating software loop delays. This is especially difficult in interrupt driven systems where there is uncertainty in interrupt servicing delays. A preferred solution is to use an external timer to provide conversion control. Such timers can be software programmable like the 8254 from Intel or a simple counter as shown in the AD7572A/ADSP-2100 interface of Figure 3.

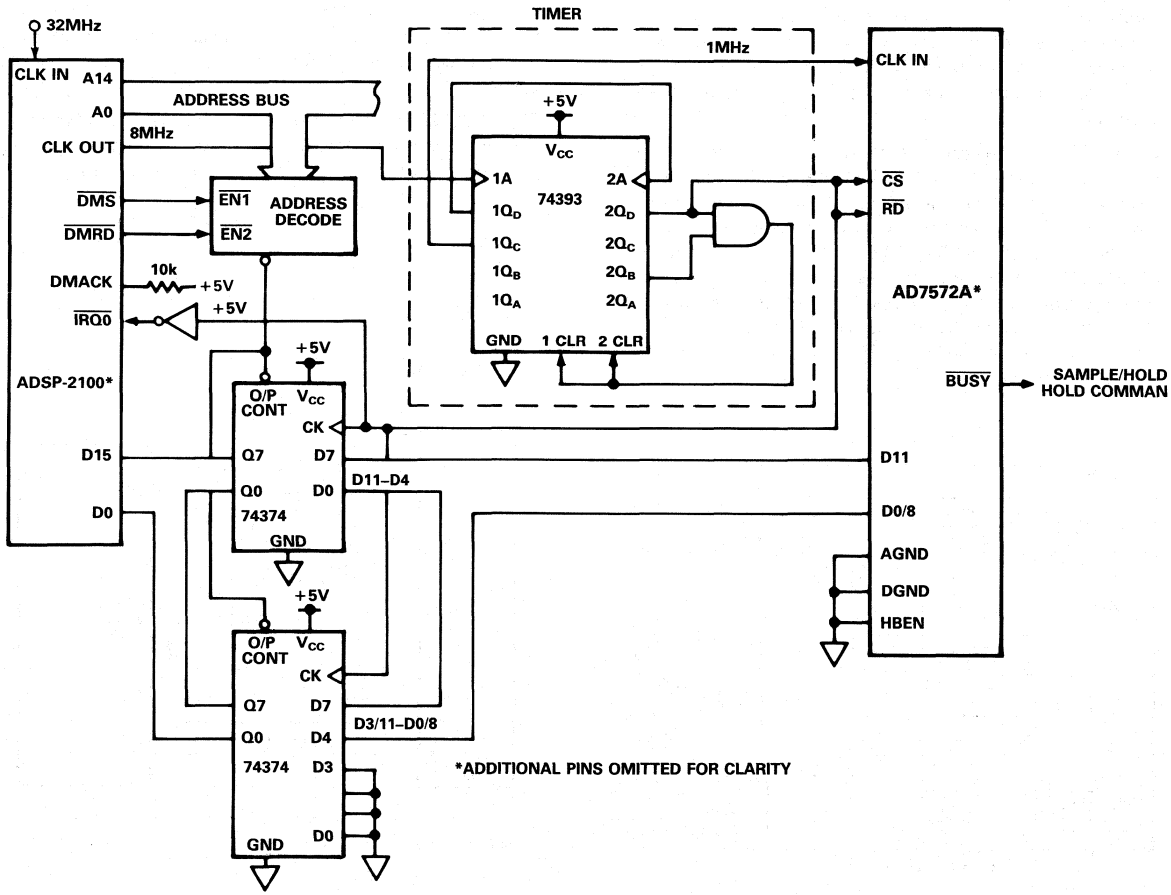

Figure 3. AD7572A/ADSP-2100 Interface

AD7572A-ADSP2100 INTERFACE

The 74393 timer counter of Figure 3 uses the ADSP-2100 CLK OUT to generate an AD7572A CLK IN and provide accurate conversion control. With an ADSP-2100 CLK IN of 32MHz, CLK OUT is 8MHz – the counter is configured so that it starts a conversion every 160 CLK OUT cycles resulting in a sampling rate of 50kHz (see Figure 4 timing diagram). An AD7572A conversion starts when the timer takes $\overline{CS}$ and $\overline{RD}$ inputs low.

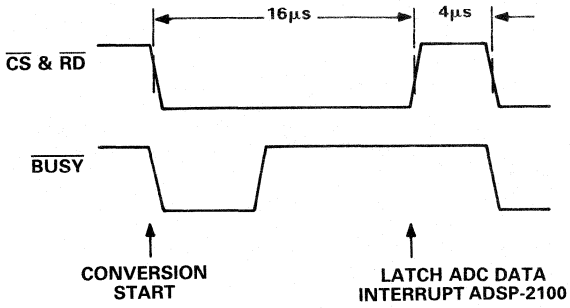

Figure 4. AD7572A/DSP-2100 Timing Diagram

The converter acknowledges by taking $\overline{BUSY}$ low, which asserts the HOLD input of the Sample-and-Hold amplifier (not shown in Figure 3) for the duration of the conversion. $\overline{CS}$ and $\overline{RD}$ remain low for 16μs; the rising edge of $\overline{CS}$ and $\overline{RD}$ latches the conversion result into the 74374 latches and interrupts the ADSP-2100 processor. The interrupt input is set to be edge sensitive during program initialization, so that it does not have to be reset during the interrupt service routine. The following single load instruction in the interrupt service routine reads the ADC data into the MR0 data register.

MR0 = DM (ADC ADDRESS)

The AD7572A data bus of Figure 3 is connected so that the conversion result is placed at the MSB end of the 16-bit MR0 register (i.e., D15-D4). This suits the internal left-justified fractional data format of the ADSP-2100. The same does not apply for all processors; for instance, right-justified data would probably be more suitable for the TMS32020. This allows for word extension in the event of an overflow when adding or multiplying.

Again, the interface of Figure 3 is not specific to the ADSP-2100, it is suitable for practically every processor available today. It has the advantage of allowing the processor to operate at maximum speed while the AD7572A performs conversions at constant sampling intervals.

ANALOG DEVICES

AN-293
APPLICATION NOTE

NE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

The AD7574 Analog to Microprocessor Interface

by Paul Toomey

ITRODUCTION

he AD7574 is a low cost 8-bit ADC designed for easy in-rface to microprocessors as a memory mapped input evice.

uses a successive-approximations conversion tech-ique, runs with an internal or external clock and can com-lete an 8-bit A/D conversion in 15 microseconds.

he analog to digital conversion operations are controlled y two logic inputs labelled $\overline{CS}$ (Chip Select) and $\overline{RD}$ READ). Conversion-in-progress indication is provided by $\overline{BUSY}$ output signal (see Figure 1).

his note is intended to describe the AD7574 in its three ain microprocessor interface modes with references to xternal clock source, input range switching and input ultiplexing applications.

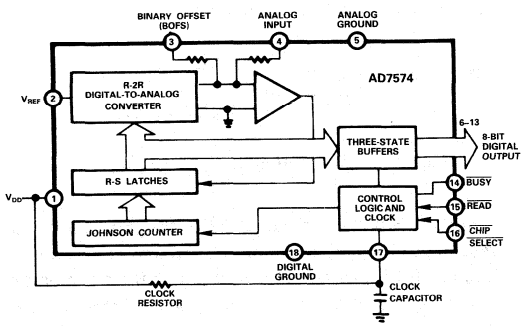

Figure 1. Inside the AD7574

NTERFACE MODES

Since the AD7574 is designed for use in memory mapped applications it can simulate RAM, ROM, or SLOW-ROM memories and can be controlled using standard chip select, READ and WRITE signals common to all memory systems.

ROM MODE

The ROM Mode is the easiest mode in which to use the AD7574. It appears in the processor memory map as one byte of Read Only Memory. One instruction from the CPU both reads conversion data and initiates a new conversion.

Basic Operation: The processor reads the 8 bits of data generated by a previous A/D conversion by executing a READ instruction from the memory address location assigned to the AD7574. When the processor $\overline{RD}$ signal goes LOW the AD7574 three state drivers are activated, placing the conversion data onto the processor data bus.

At the end of the READ instruction the AD7574 $\overline{RD}$ input is returned HIGH, resetting the device and initiating a new A/D conversion sequence (see Figure 2).

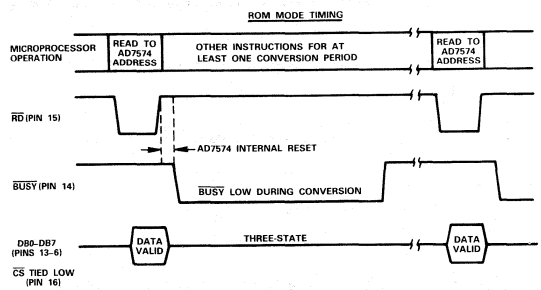

Figure 2.

Typical ROM MODE interface circuits for 8080 and 8085 microprocessors are shown in Figures 3 and 4. Most processors can be configured to operate with the AD7574 in this mode. Note: Any attempt to read data from the AD7574 during a conversion operation will result in incorrect data being read.

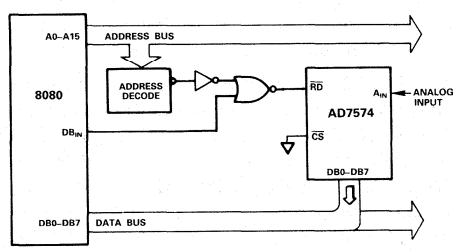

Figure 3. AD7574 to 8080 ROM-MODE Interface

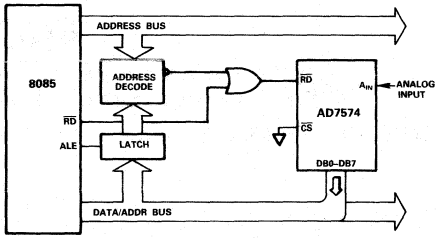

Figure 4. AD7574 to 8085 ROM-MODE Interface

Applications Of The Rom Mode

The advantage of this mode is its simplicity in both software and hardware terms. Reading conversion data is accomplished by just one memory READ instruction. However, it must be remembered that the data read will be the result of the previous conversion operation. This means that the time reference for the data sample will depend on when the previous READ operation finished. In applications where this uncertainty creates problems it can be eliminated by executing two READ operations separated by a software delay as shown in Figure 5.

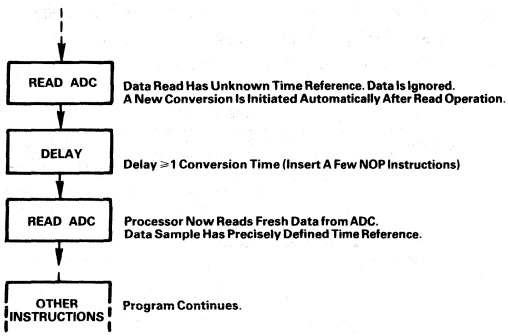

Figure 5. Defining the Data Sample Time Reference (ROM- MODE)

RAM MODE

Basic Operation

In the RAM MODE the AD7574 appears in the processor memory map as one byte of memory. A WRITE command initiates a conversion and a READ command reads the conversion data.

This mode offers complete control over the converter operation. The time reference for the data sample is defined by the WRITE command. The conversion data can be read any time after the conversion has finished.

In the RAM MODE a memory WRITE operation starts a conversion without modifying any of the microprocessor accumulators. (In the ROM mode, using a READ instruction to start a conversion means modifying an accumulator unnecessarily.) So, in situations where the age of a sample is important the RAM MODE makes for simpler, more logical software than the ROM MODE at the expense of requiring slightly more logic to drive $\overline{CS}$ and $\overline{RD}$ (see Figures 8 and 9).

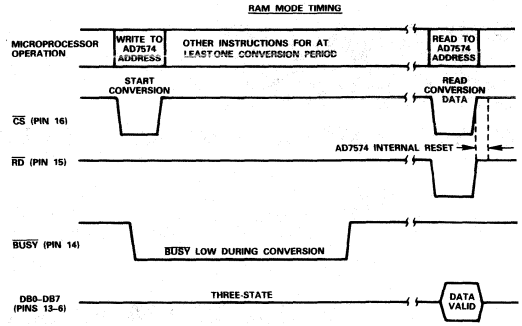

Figure 6.

Ram Mode Timing Considerations

1. $\overline{BUSY}$ must be high before a data read is attempted i.e., delay from conversion start to data read must b at least as great as the AD7574 conversion time. I some situations it is possible to use the AD7574 $\overline{BUSY}$ output to halt the microprocessor for the duration o the conversion period thus simplifying the software re quirements.
2. $\overline{CS}$ must return HIGH within 250ns after $\overline{RD}$ goes HIGH otherwise a new conversion may be initiated and the AD7574 will begin operating in the ROM mode (see Figure 7).

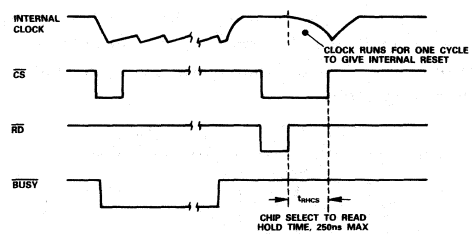

Figure 7. Chip Select to Read Timing

When $\overline{RD}$ returns high the AD7574 begins an internal reset cycle to prepare for the next conversion sequence. If $\overline{CS}$ has not returned high before the end of this reset operation, a conversion start will be detected and a new conversion initiated.

The duration of the internal reset cycle is dependent on the amount of capacitance on the clock pin (pin 17). The maximum value for t_{RHCS} guaranteed is 250ns at 25°C (see data sheet).

3. Conversion data may only be read from the AD7574 once, as after each read operation an internal reset is performed which destroys the data.
4. $\overline{CS}$ LOW has no effect while a conversion is in progress.
5. $\overline{RD}$ has no effect while $\overline{CS}$ is HIGH.

Typical RAM Mode interface circuits for the 8085 and 6800 are shown in Figures 8 and 9 respectively.

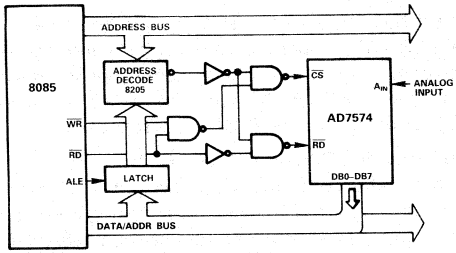

Figure 8. AD7574 to 8085 RAM Mode Interface

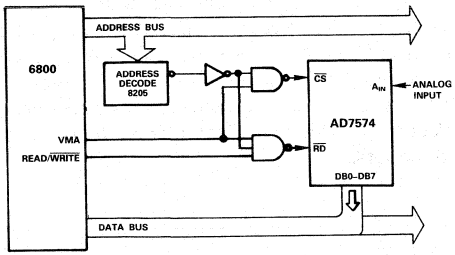

Figure 9. AD7574 to 6800 RAM Mode Interface

SLOW MEMORY MODE

This is by far the most elegant mode of operation for the AD7574. Every read instruction produces fresh data so that there is no doubt about the age of the sample.

Basic Operation

The slow memory mode is intended for use with processors which can be forced into a wait state for at least 15μs such as the 8080, 8085 and SC/MP). It allows the processor to start a conversion, wait until the $\overline{BUSY}$ flag is HIGH and then read data, all during execution of a single memory read instruction (see Figure 10).

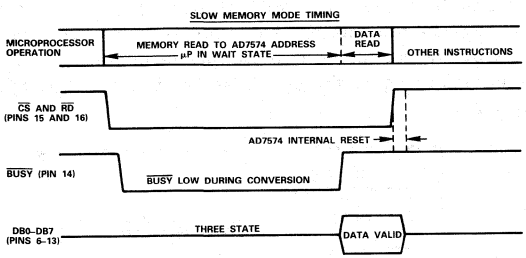

Figure 10.

Wait State

Many processors test the condition of the READY/WAIT input very soon after the start of an instruction cycle. For this reason the timing of the AD7574 and its associated address decode logic must be such that $\overline{BUSY}$ goes LOW early enough in the processor instruction cycle for the READY/WAIT input to be effective in forcing the processor into a WAIT State.

Bus Conflict

In applications where the processors memory READ/WRITE signal is not available early enough in the machine cycle for it to be used to enable or disable the address decode logic, the system software must be such that a WRITE operation to the AD7574 address is never attempted. If this precaution is not taken, Bus Conflicts will occur due to the AD7574 outputting data onto the data bus while the CPU is also driving the data bus.

Typical slow memory mode interface circuits for 8085 and SC/MP microprocessors are shown in Figures 11 and 12.

8085 Interface (Figure 11)

For simplicity, only the upper 8 address bits of the 8085 address bus are decoded to select the AD7574. Invalid address states are eliminated by using ALE to drive an address latch.

The processor SO status signal provides the earliest possible indication that a READ operation is about to occur, SO = 0 for a READ operation. Since, with the processor on a fast clock the READ signal could occur too late to enable the address decode logic, the SO signal is a convenient alternative, elminating any possibility of a bus conflict during WRITE operations.

SC/MP Interface (Figure 12)

Similar to the 8085 application. Address decode is gated with negative READ strobe and $\overline{BUSY}$ drives the SC/MP negative hold input.

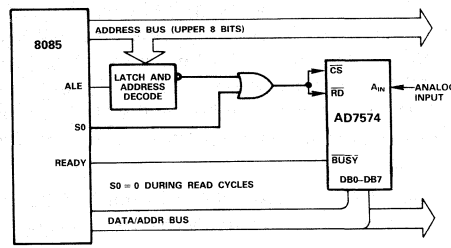

Figure 11. AD7574 to 8085 Slow Memory Mode Interface

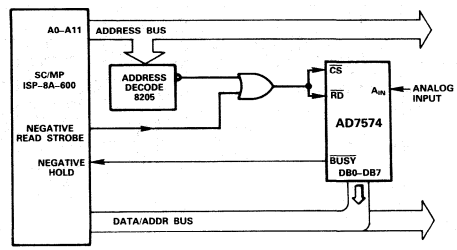

Figure 12. AD7574 to SC/MP Slow Memory Mode Interface

RANGE SWITCHING APPLICATIONS

By means of suitable switching on the AD7574 B_{OFS} and V_{REF} pins, the AD7574 can be made to operate with a range of different attenuation or gain factors. The B_{OFS} or bipolar offset input (pin 3) is used to modify the effective analog input voltage and is normally used to obtain bipolar operation. Figure 13 shows a simple, 3 range, switching arrangement for a 0–20V analog input signal. The full scale input ranges and LSB weights are given in Table I.

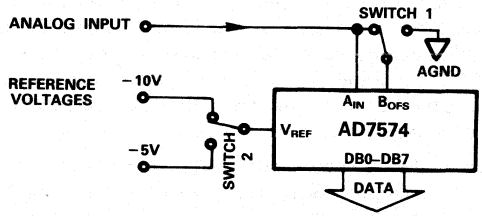

Figure 13. Range Switching the AD7574

BOFS	V_{REF}	Full Scale Input Range	LSB Weight
To A_{IN}	−5V	0V to +5V	19mV
To A_{IN}	−10V	0V to +10V	39mV
To GND	−10V	0V to +20V	78mV
To GND	−5V	0V to +10V	39mV

Table I. AD7574 Input Ranges

The range switching can easily be controlled by the processor using analog switches such as the AD7592 which contains 2 independent SPDT CMOS switches with data latches specifically designed for microprocessor interface.

Figure 14 shows a typical range switching applications circuit with the AD7574 operating in the the slow-memory mode. Since bus conflicts have been eliminated, a WRITE to the AD7574 address selects the analog input voltage range (see Table II) and a READ from the AD7574 address initiates a conversion and reads data when the conversion is complete.

The range selection code is latched by the AD7592 and so need only be considered when a change in analog input range is required.

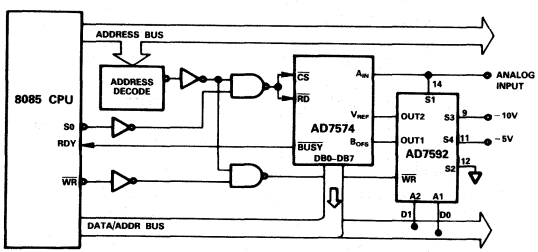

Figure 14. AD7574/8085 Interface with Range Switching

D1	D0	Analog Input Range
0	0	0V to +10V
0	1	0V to +5V
1	0	0V to +20V
1	1	0V to +10V

Table II. Analog Input Range vs. Digital Code to AD7592 for Figure 14

Note: If it is required to use this approach in a syste where the WRITE operation bus conflict cannot be elim nated due to timing difficulties, then by decoding a sep rate address for the AD7574 and the AD7592 the need fc a WRITE operation to the AD7574 address can be avoidec

INPUT MULTIPLEXING APPLICATIONS

This application uses the same principles as the rang switching application. The analog switching devices hav been changed to AD7590s to reduce the package coun The AD7590 consists of four independent SPST CMO analog switches with on chip data latches (see Figure 15).

Using two AD7590 devices and connecting the data bu directly to the AD7590 control inputs, 8 analog input chan nels can be multiplexed. The software design must en sure that only one channel is selected at a time. Error introduced by the analog switch resistance are not signifi cant but can be eliminated by buffering the AD7574 B_{OF} and A_{IN} inputs.

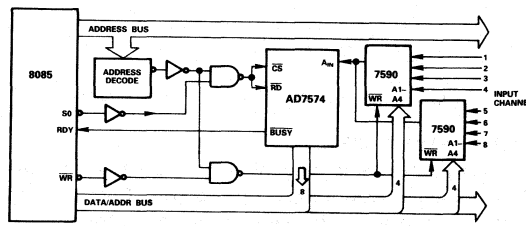

Figure 15. AD7574/8085 Interface with Input Multiplexing of 8 Analog Inputs

A WRITE to the AD7574 address selects an input channel and latches it. A READ from the AD7574 address initiates a conversion and reads data when the conversion is complete.

INTERNAL CLOCK OPERATION

Figure 16 shows a simplified equivalent circuit for the AD7574 internal clock. It operates only during conversion and reset operations, pulses are generated by the action of C charging through R, and discharging through switch 1.

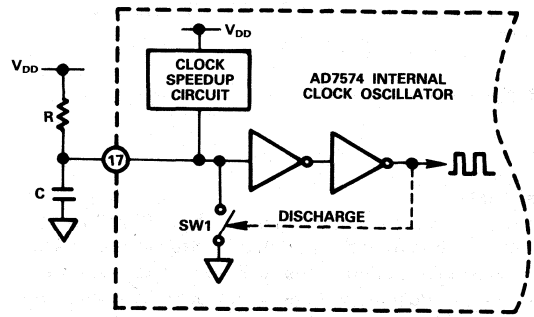

Figure 16. Simplified AD7574 Internal Clock Circuit

A clock speedup circuit shortens the last clock space period in each conversion cycle to reduce overall conversion time. Figure 17 shows a RAM mode timing sequence using the internal clock; other modes have similar timing.

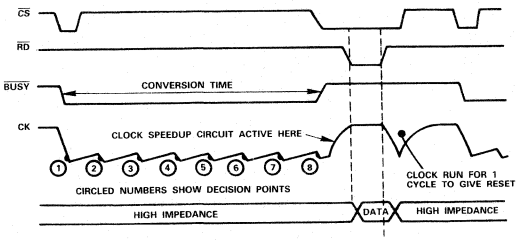

Figure 17. RAM Mode Timing Sequence

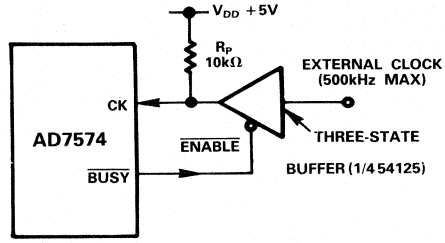

Figure 18. External Clock Connection

'ypical values for internal clock timing components can ıe determined from the graph in the AD7574 data sheet Figure 7a).

:XTERNAL CLOCK INFORMATION

'o obtain dynamic conversion accurcy to rated specifica- ion the clock frequency must not exceed 500kHz. The ıser should understand that normal lot to lot variations in ıMOS transistor characteristics will cause lot to lot differ- ;nces in the internal clock oscillator frequency for a given ;lock R and C.

Additionally, temperature dependence of these MOS ;haracteristics results in thermal drift of internal clock 'requency. For this reason, Analog Devices recommends ısing an external clock in the following situations:

1. Applications having clock frequency within 10% of the 500kHz maximum.
2. Applications where software constraints on time cannot accommodate conversion time differences which may occur due to temperature drift of the internal clock.

Conflicts between the external clock and the internal speedup and reset operations can be avoided by connecting the external clock using the arrangement shown in Figure 18. The three-state buffer is only enabled when the $\overline{\text{BUSY}}$ output is LOW. This ensures that the AD7574 uses the external clock only during the A/D conversion period, i.e., while $\overline{\text{BUSY}}$ is LOW. Internal clock operation is then used with R_P and internal capacitance to give the single clock pulse required for the internal reset.

Since the internal logic of the AD7574 is triggered on falling clock edges, conversion time is reduced if the clock input is at a HIGH level before a conversion starts. This is accomplished by using resistor R_p to pull up the device clock input after the internal reset operation. R_p can be in the range 6kΩ to 100kΩ and simply provides a charge path for the CLK pin capacitance.

ANALOG DEVICES

AN-294
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

AD7672 Converter Delivers 12-Bit 200 kHz Sampling Systems

by John Reidy

The AD7672 is a low power 12-bit analog to digital converter manufactured on Analog Devices' LC²MOS process. Its fast conversion time of 3 μs and microprocessor interfacing capabilities make it popular in DSP applications. Sampling rates of over 200 kHz are easily achieved by combining the AD7672 with a suitable sample-and-hold amplifier (SHA). However, the choice of SHA and other analog conditioning circuitry is critical in achieving 12-bit performance over the full 0 to FS/2 analog bandwidth. Any distortion or noise generated in the analog circuitry will ripple through to the digital domain, degrading system performance.

This application note discusses the noise and distortion levels acceptable for 12-bit performance and details a SHA and antialiasing filter suitable for up to 200 kHz sampling rates.

AD7672 NOISE MODEL

All analog systems are corrupted by noise. An accepted figure of merit for noise performance is signal-to-noise ratio (SNR). The numerical expression is:

$$SNR = 20 \log (V_{SIG}/V_N) \quad \ldots\ldots\ldots\ldots\ldots\ldots (1)$$

where: V_{SIG} = rms value of the signal
V_N = rms value of the noise plus distortion.

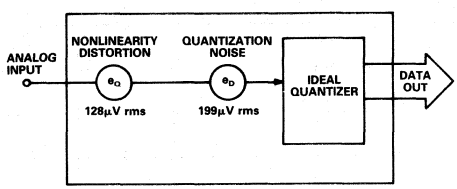

Figure 1. AD7672 Noise Model

A simplified noise model for the AD7672 is shown in Figure 1. The model consists of an ideal quantizer and two noise sources, one for quantization noise and the other for nonlinearity or accuracy errors. Another source of noise for ADCs is code transition noise. The AD7672 code transition noise is less than 0.2 LSBs pk-to-pk and is ignored in this application note.

Quantization noise is random and can be quantified for any N-bit ADC as that noise which will yield a signal to noise ratio of 6.02N+1.76 dB in the 0 to FS/2 bandwidth. For an ideal 12-bit converter the SNR is 74 dB.

Linearity errors for a statically tested ADC, like the AD7672, are quantified by either integral nonlinearity (INL) or differential nonlinearity (DNL). The AD7672 L, C, U grades are guaranteed to have less than 0.5 LSBs of INL and less than 0.9 LSBs of DNL. If these errors are modelled as random noise, then the theoretical SNR will degrade by 3 dB down to 71 dB. Such a model is incorrect for two reasons. Firstly, the error pattern for a successive approximation ADC does not follow a white noise pattern. There is not an equal probability of the error for any given code being between 0 and 0.5 LSBs. The error pattern is much closer to a normal distribution where the probability of a zero error is much higher than that of a 0.5 LSB error. Secondly, linearity errors follow a fixed pattern, i.e., the error associated with any code remains fixed. For the AD7672, the errors tend to repeat themselves within the transfer function of the ADC, and furthermore the errors will repeat when multiple cycles of a sinusoidal input are digitized. This repeatability means that the ADC nonlinearity generates harmonics of the input frequency rather than white noise.

Figure 2 shows a typical integral nonlinearity plot for the AD7672. Examination of the plot shows an error pattern repeating itself every 512 codes or 8 times in the overall transfer function. The maximum INL error in this case is 0.25 LSBs.

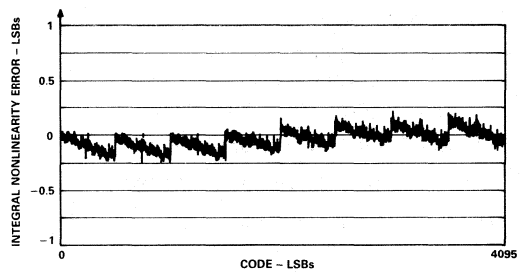

Figure 2. Typical AD7672 Integral Nonlinearity Plot

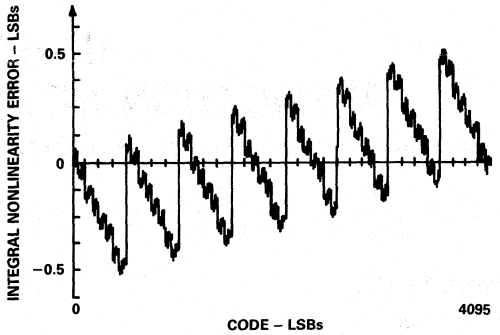

Figure 3a. Simulated AD7672 Nonlinearity Plot

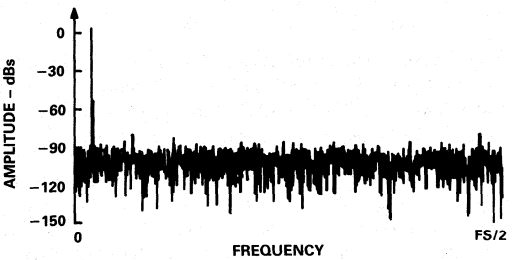

Figure 3b. FFT Plot Using a Simulated ADC

The effect of these INL errors on ac performance can be determined by applying a single tone sinusoidal signal via a SHA to the input of the ADC and performing a fast Fourier transform (FFT) on the conversion results. For a controlled analysis, this task is best done on a computer using a software simulation of the ADC. With a software simulation, all error sources can be regulated so that the SHA and purity of the input signal can be made ideal. Figure 3a shows the INL error profile of a simulated AD7672. The maximum INL error has been increased from the measured value of 0.25 LSBs to 0.5 LSBs reflecting a worst case for the L, C and U grades. Figure 3b shows a corresponding FFT plot; the SNR calculated from the plot is 72.5 dB, a degradation of 1.5 dB from the ideal case of 74 dB. In other words, the degradation in SNR due to INL errors is 1.5 dBs.

The AD7672 is modelled as having two predominant noise sources which add in an rms fashion as shown below:

$$e_{ADC} = (e_Q^2 + e_D^2)^{1/2} \quad \ldots\ldots\ldots\ldots\ldots\ldots (2)$$

where: e_{ADC} = total ADC rms noise
e_Q = quantization noise
e_D = noise due to nonlinearity.

For an ideal 12-bit ADC without errors:

$$SNR = 20\ log\ (A/e_Q) = 74\ dB$$

where: A = ADC full-scale input range.

For the purpose of analysis assume the full-scale input range (A) is 1 V rms, then e_Q can be calculated as follows:

$$e_Q = log^{-1}(-74/20) = 199\ \mu V \quad \ldots\ldots\ldots\ldots\ldots (3)$$

Similarly, for the AD7672 with an SNR of 72.5 dB, e_{ADC} can be calculated to be 237 μV. Finally, e_D can be calculated from Equation 2 to be 128 μV. Note, these figures have to be scaled according to the input range of the ADC, e.g., for the ±5 V input range of the AD7672 the rms full-scale value is 5×0.7071 and the corresponding noise sources are e_D= 453 μV and e_Q = 704 μV.

Any additional noise in the external analog circuitry will add in an rms fashion to the ADC noise. Assuming a full scale range of 1 V, the following expression gives the overall SNR performance:

$$SNR = 20\ log\ 1/(e_{ADC}^2 + e_{EXT}^2)^{1/2}$$

where: e_{ADC} = quantization plus nonlinearity noise
e_{EXT} = external noise.

A corresponding expression for any degradation, D, in SNR performance may be found by subtracting the overall SNR from the ADC SNR;

$$D = 20\ log\ [1/e_{ADC}] - 20\ log[1/(e_{ADC}^2 + e_{EXT}^2)^{1/2}]$$
$$= 10\ log\ [1 + e_{EXT}^2/e_{ADC}^2]$$

A plot showing the degradation in SNR performance from 72.5 dBs due to noise plus distortion in the external analog circuitry is shown in Figure 4. Taking an example from the graph, if the overall SNR is to be kept to 71.5 dBs (an SNR degradation of 1 dB), any external ADC noise plus distortion must be kept to within 78 dBs relative to ADC full scale.

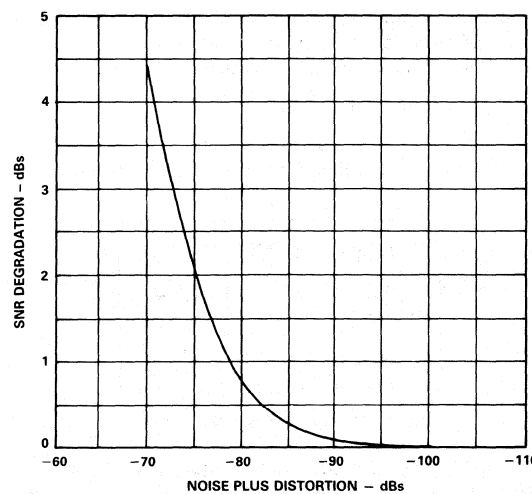

Figure 4. SNR Degradation Due to External Noise

SAMPLE-AND-HOLD AMPLIFIER SELECTION

One of the first considerations when selecting a sample-and-hold amplifier is to ensure that the acquisition time satisfies the throughput rate required. The minimum conversion time of the AD7672 is 3 μs and allowing, say, 2 μs for the SHA acquisition time and other overheads gives a minimum repetition rate of 5 μs or a maximum sampling frequency of 200 kHz. The corresponding analog bandwidth is 100 kHz. However, 100 kHz bandwidth to 12-bit performance is beyond the reach of present low cost monolithic SHAs and requires more expensive hybrid devices.

From Figure 4 the total noise plus distortion must be more than 78 dBs below the fundamental to maintain an overall SNR of 71.5 dBs. In general the noise floor for all SHAs will be good enough, but the limiting factor for monolithic SHAs is third harmonic distortion due to slew rate limiting. Examples of monolithic SHAs are the AD585 and the AD684 from Analog Devices. The AD585 exhibits 12-bit performance up to a frequency of 35 kHz while the AD684 is good enough for a 40 kHz bandwidth. The AD684 and AD7672 are an ideal combination for multiplexed systems. The AD684 has four SHAs in one package each with an acquisition time of 1 μs per channel. When used with the AD7672 the maximum throughput rate per channel is in the region of 70 kHz..

Hybrid SHAs are needed for systems requiring 12-bit performance beyond 40 kHz signal bandwidth. Examples of suitable devices are the AD346 and the HTC0300, available from Analog Devices. An alternative, more cost effective, solution is to use a discrete SHA as discussed in the next section.

DISCRETE SHA DESIGN EXAMPLE

Figure 5 shows a simplified block diagram of a discrete SHA. The integrated circuits used in the design are an AD845 op amp from Analog Devices and an SD5000 quad DMOS switch from Topaz.

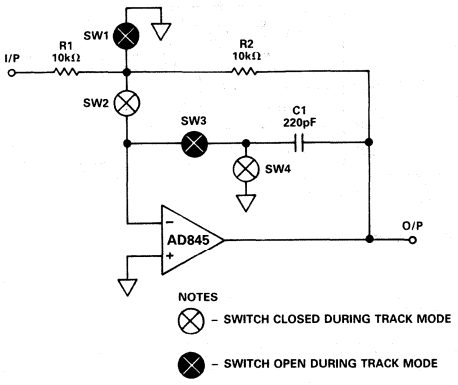

Figure 5. Discrete SHA

In the track mode, switches SW2 and SW4 are closed while SW1 and SW3 are open. The amplifier is in the standard unity gain inverting configuration and charges the capacitor, C1, to the negative value of the input voltage. In the hold mode, the switch settings are reversed, placing C1 across the output and the input inverting terminal of the op amp and holding the output voltage to the charged voltage on the capacitor.

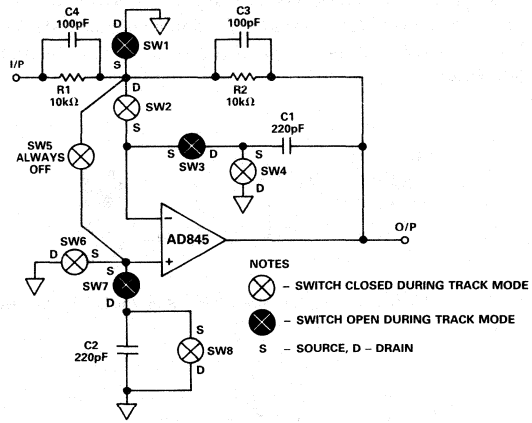

Figure 6. Detailed SHA Diagram with Additional Switches for Charge Injection and Feedthrough Compensation

The droop rate in the hold mode depends on the capacitor leakage, the op amp inverting terminal bias current and the combined leakage of switches: SW2, SW3 and SW4. It is important that a high quality capacitor is used such as polystyrene. A JFET input op amp such as the AD845 ensures that any leakage due to bias current is a minimum. Considerable improvement in droop rate and pedestal error can be obtained at the expense of a few additional switches as shown in Figure 6. Pedestal error is a dc offset between the track and hold mode. The dc offset is caused by switching transients which inject charge into the hold capacitor. One method of minimizing pedestal error is to have equal switching transients at both inputs of the op amp and to rely on a good common mode rejection ratio (CMRR) to cancel the transients. In Figure 6, SW6, SW7 and SW8 have been added to the noninverting input of the op amp to balance the switching effects of SW2, SW3 and SW4. The SD5000 switches contain four switches per package. To make full use of the charge cancelling, corresponding switches on each of the op amp input terminals should come from the same package, i.e., SW2 and SW6 should come from the same package and similarly with SW3 and SW7 and with SW4 and SW8.

A switch SW5 has been added to compensate for any feedthrough that may occur from the analog input while the SHA is in the hold mode. Again this compensation relies on common mode signals cancelling each other out. Finally a 100 pF capacitor C3 has been added to

speed up track to hold settling time. This capacitor introduces a pole in the overall transfer function so a corresponding zero has been included by adding capacitor C4.

One of the key benefits of this design is that all the switch analog terminals are at either true ground or virtual ground. As the switches only conduct near zero voltage level signals they do not need a large drive voltage on the digital control inputs. HCMOS logic or equivalent provides adequate voltage levels as shown in Figure 7. Note, all switch substrates (backgates) must be tied to analog ground.

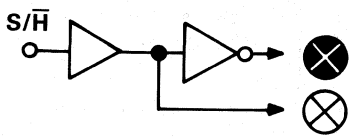

Figure 7. HCMOS Logic Drives SHA's SD5000 Switches

DATA ACQUISITION CIRCUIT

Figure 8 shows a data acquisition circuit designed for a 200 kHz sampling frequency. The sample and hold shown in block diagram form can be either an AD346 AD781 or the discrete version shown in Figure 6. A voltage reference (AD588) provides the appropriate biasing for an analog input range of ±5 V. The data bus outputs are buffered with 74HC374 latches. These provide data bus isolation while improving the data access time. Data access time is under 30 ns allowing interfacing to most microprocessors. Data format can either be a complete parallel load for 16-bit processors or a two byte load for 8-bit processors.

Bus activity on the AD7672 $\overline{CS}$ and $\overline{RD}$ inputs during conversion can feed through to the comparator and cause LSB errors. Ideally, these should be inactive during conversion. One way of achieving this is to force them into an inactive state by gating them with $\overline{BUSY}$ as shown in Figure 8. R2 and C3 are included to provide a delay of approximately 100 ns. This compensates for the data setup time after $\overline{BUSY}$ goes high, ensuring valid data gets loaded into the 74HC374 latches.

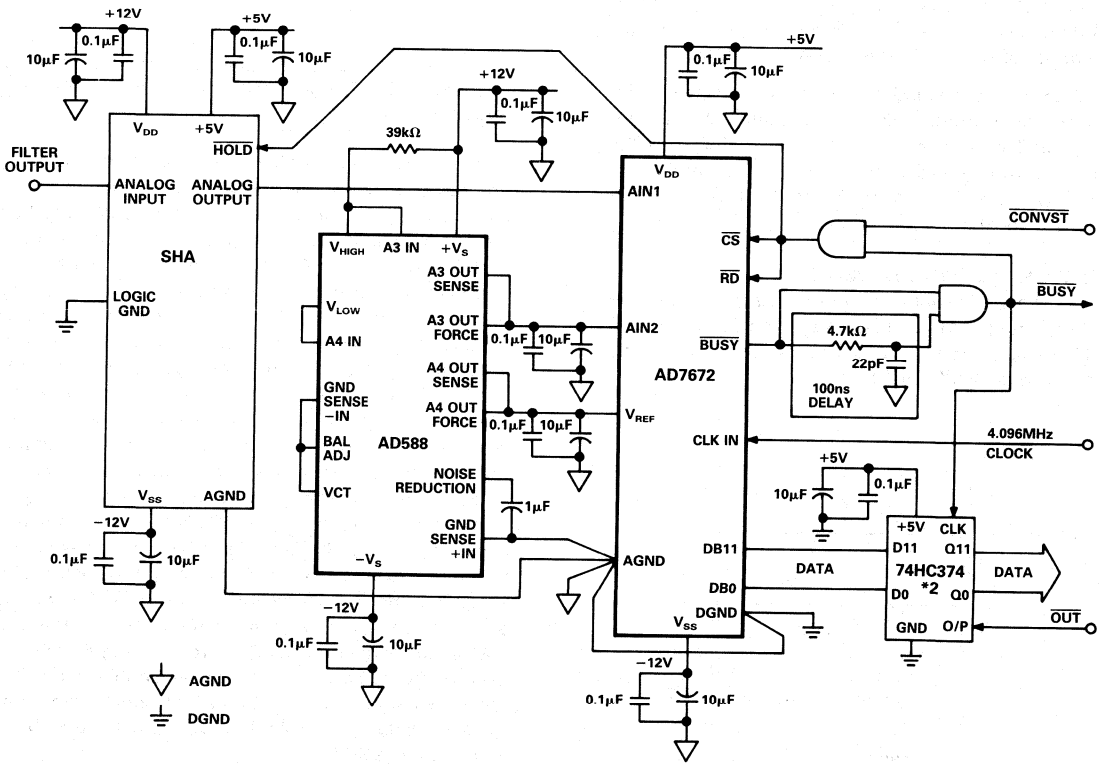

Figure 8. AD7672 Data Acquisition Circuit

'ESTING SYSTEM PERFORMANCE

'o test the performance of Figure 8 a pure tone sinusoi-lal waveform was applied at the input and an FFT was)erformed on 2k samples. The procedure was repeated or multiple frequencies for both the AD346 and the liscrete SHA of Figure 6. The AD346 requires an acqui-ition time of 2 μs, combined with the ADC conversion ime of 3 μs and allowing 500 ns for other overheads, gives a maximum sampling rate of 180 kHz. With the liscrete SHA, the system can be run up to 200 kHz sam-)ling rate. The results are shown in Figures 9 to 12. As :an be seen from the plots, the system SNR perfor-nance is above 71 dB for both SHAs.

NOTE ON THE FAST FOURIER TRANSFORM

Fast Fourier transforms are widely recognised as a very useful tool when evaluating sampling systems. An FFT provides a means of plotting the spectral content of a sample of digital data; examples are shown in Figures 9 and 11. Two important considerations when examining FFT plots are the frequency and amplitude resolutions of the FFT.

Frequency resolution of an FFT is equal to the sampling frequency divided by the number of points in the FFT. In the example of Figure 9, the frequency resolution is 204.8 kHz/2048 = 100 Hz. This figure of 100 Hz is akin to resolution bandwidth for spectrum analyzers. Every doubling of the FFT length will spread the noise floor over twice as many bins and reduce the noise floor by 3 dB. However, any harmonics which are present in the analog domain or are generated by the ADC will remain constant regardless of the FFT length. This includes low level harmonics which have a smaller magnitude than the dynamic range of the ADC. This may seem contradictory if we consider the smallest signal a 12-bit A/D converter can digitize is −72 dB relative to full scale. However, this is only true for a small signal on its own. In the presence of a larger signal, the small signal will act as a dither around the A/D converter's transition points and will appear in the frequency spectrum of an FFT as long as its magnitude is greater than the noise floor. The average noise floor level for a 2k FFT, computed from 12-bit data, can be obtained in the following manner:

From Equation 3

RMS noise for an ideal 12-bit converter = 199 μV.

If we assume this noise is spread evenly across 1024 points between 0 and FS/2

Then the average noise floor = 199 μV/$(1024)^{1/2}$

= 6.21 μV

Average noise = 20 log(6.21μV) = 104 dB.

The FFT noise floor tends to be Gaussian distributed, so the worst case is expected to show a power of 10 or 11 dB higher than the average bin.

So, any harmonic or spurious noise greater than 93 dB will be discernible for a 12-bit converter and a 2048 point FFT.

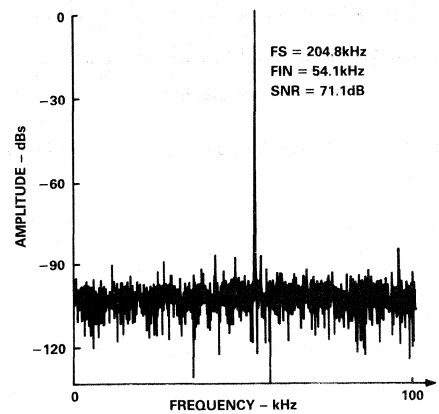

Figure 9. FFT Plot Using the Discrete SHA, Input Frequency Is 54 kHz

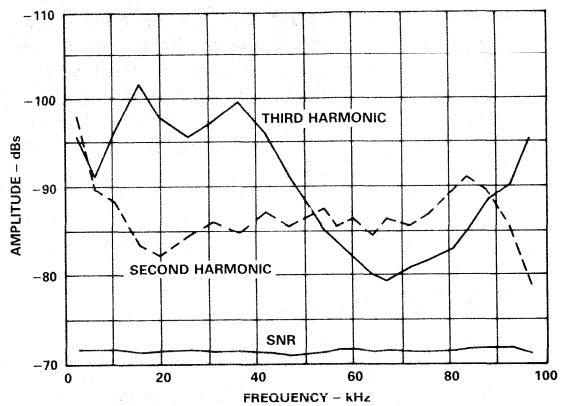

Figure 10. Discrete SHA Harmonic Amplitudes and SNR Performance Versus Frequency

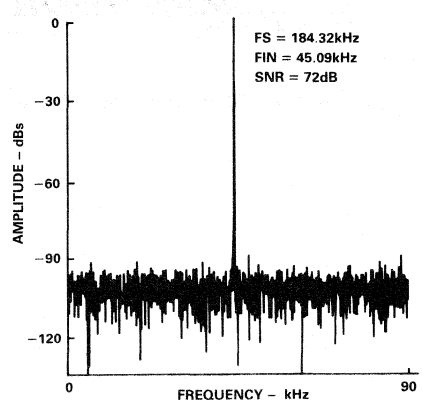

Figure 11. FFT Plot Using the AD346 SHA, Input Frequency Is 45 kHz

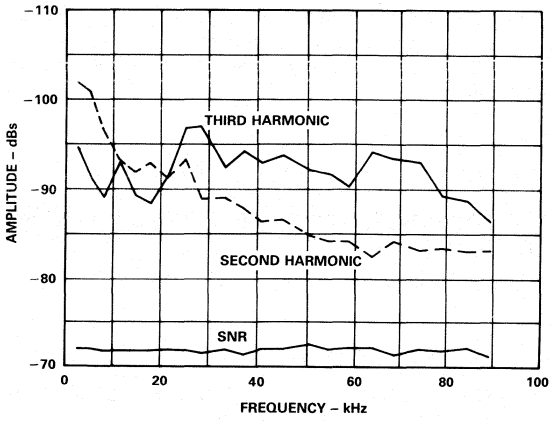

Figure 12. AD346 SHA Harmonic Amplitudes and SNR Performance Versus Frequency

ANTIALIASING FILTER

The purpose of an antialiasing filter is to attenuate any noise or spurious signals outside the analog bandwidth. Two important parameters in any filter design are stopband attenuation and the slope of the roll-off in the transition band. Restrictions on both of these parameters are eased by oversampling. An oversampled system is where the Nyquist limit FS/2 is greater than the analog bandwidth. Figure 13 shows the filter attenuation profile for an oversampled system.

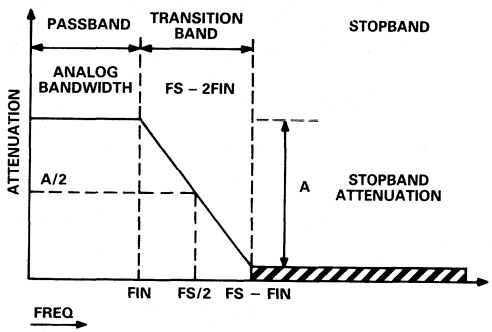

Figure 13. Antialiasing Filter Attenuation Profile

The filter roll-off is determined by drawing a straight line from the upper limit of the analog bandwidth, FIN, to the stopband attenuation frequency, FS−FIN. The difference between the filter cutoff and the stopband frequencies is FS −2FIN. For example, if an antialiasing filter requires one frequency octave to roll off, then the sampling ratio (FS/FIN) must be 3 to 1.

Any residue from the filter output, after attenuation, will add in an rms fashion to the digital domain noise floor, degrading system performance as shown in Figure 4. Oversampling, beyond the Nyquist rate, spreads the noise floor over a wider bandwidth. This is readily explained by considering that the SNR figure for a 12-bit ADC is 74 dBs over the 0 to FS/2 bandwidth regardless of the sampling rate. Increasing the sampling rate increases the FS/2 bandwidth and therefore must reduce the noise floor level to preserve the 74 dBs SNR figure. The decrease in total rms noise for a given bandwidth, FIN, is given by 10 log(FS/2FIN). A plot showing the decrease in noise versus sampling ratio is shown in Figure 14. A lower noise floor means that more external noise can be tolerated for a given system performance. As an example consider the circuit of Figure 8 with a sampling rate of 200 kHz. For an analog bandwidth of 75 kHz, the sampling ratio is 2.66 and the total rms noise reduction is 1.25 dB. Correspondingly, the SNR over the 75 kHz bandwidth improves by 1.25 dB because of the oversampling.

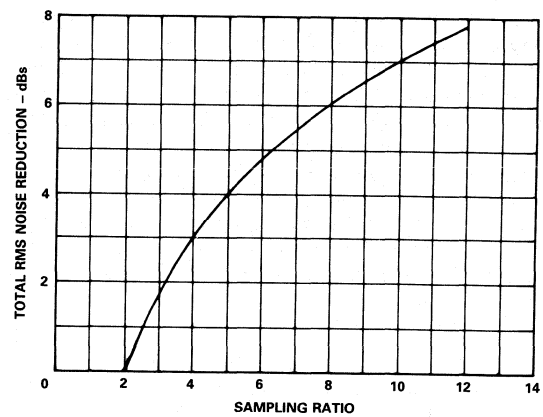

Figure 14. Total Noise Reduction Versus Sampling Ratio

ESIGN EXAMPLE

nlike a SHA design which can be optimized for a given DC, the filter design depends on the signal SNR and quired phase characteristics. Applications requiring near phase may use a Butterworth filter and pay the rice of a higher order, more complicated design, or a gher sampling ratio because of the relatively poor roll-ff in the transition band. Where phase is not important r where phase can be compensated in software then a ıore practical filter design is an elliptic. An elliptic low-ass filter contains both poles and zeros giving the eepest roll-off for any given order allowing the small-st sampling ratio.

igure 15 shows a 9th order elliptic filter design for the ircuit of Figure 8. The cutoff frequency is 75 kHz giving stopband frequency of 125 kHz for a 200 kHz sampling ate. The designed stopband attenuation is 100 dB and a 0.1 dB ripple in the passband. The actual frequency response can be seen in the spectrum analyzer photograph of Figure 16. The stopband attenuation level is over 94 dBs below the passband level at a spectrum analyzer resolution bandwidth of 10 Hz.

The filter design uses passive components to avoid poor noise performance arising from cascaded op amp stages. A figure of merit for inductors is the ratio of the reactance to the series resistance or Q.

$$Q = WL/R$$

All inductors used in the filter construction had a Q of greater than 75 at 75 kHz. All capacitors used were polystyrene.

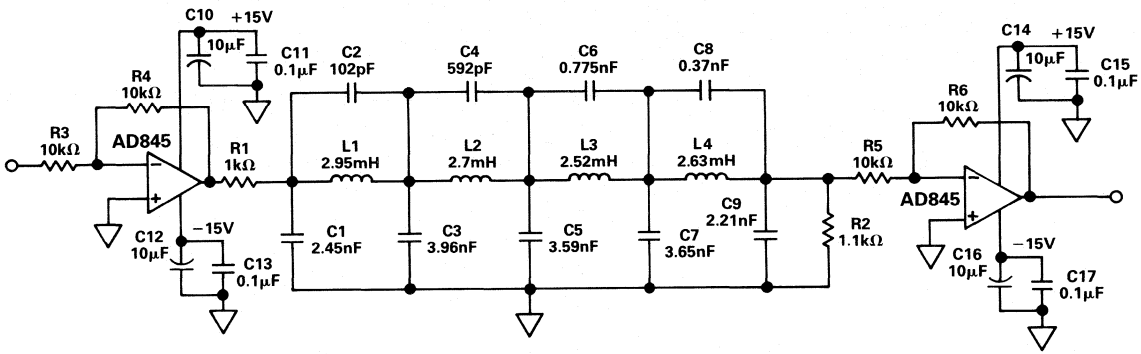

Figure 15. Ninth Order Elliptic Low Pass Filter

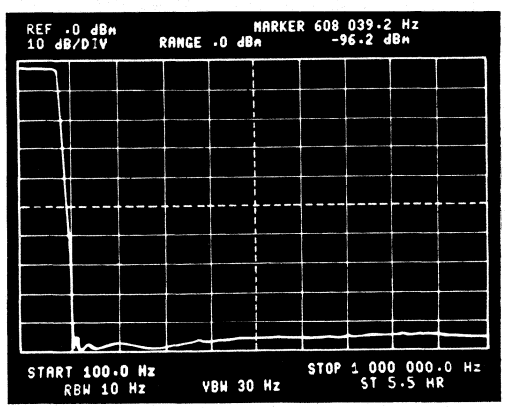

Figure 16. Frequency Response of the Elliptic Filter Shown in Figure 15

ANALOG DEVICES

AN-295
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Using Multiple AD1334s in Many-Channel Synchronous Sampling Applications

by Stephan Goldstein

INTRODUCTION

The development of fast, inexpensive DSP microprocessors has made practical new approaches to signal processing. These applications typically involve the real-time processing of data derived from many simultaneously-sampled channels. One example is phased-array sonar, in which the many digitized sensor outputs may be transformed, filtered, and correlated to perform beamforming or imaging.

A fundamental requirement in dynamic multichannel applications is to maintain phase coherency across all input channels. The tight group-delay specifications of the AD1334's four input SHAs simplify error-budget calculations in simultaneous-sampling applications. The addition of an expandable control function and a data FIFO makes the AD1334 uniquely suited to acquire and buffer the large amounts of data required to support these modern signal processing techniques.

This note details the AD1334's control structure and shows how its features can be used to implement a low-overhead software control scheme for any number of parallel devices operating in synchrony.

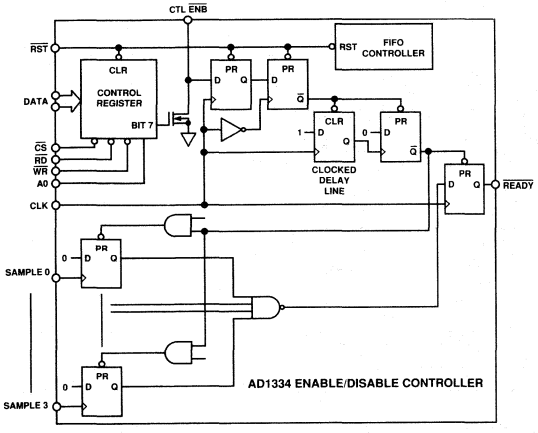

Figure 1. Simplified Schematic of the AD1334 Enable/Disable Controller

Controller Operation

The AD1334 can be enabled and disabled under software or hardware control. The CONTROL $\overline{\text{ENB}}$ pin (Pin 32) is connected to the output of an n-channel open-drain inverter and to the D input of a flip-flop as shown in Figure 1. This pin requires an external pull-up resistor to V_{DD}. The inverter's input is driven by control register Bit 7. A logical "1" will appear at the CONTROL $\overline{\text{ENB}}$ pin when Bit 7 is a "0" and conversions will be disabled. Bit 7 is set to 0 by pulling $\overline{\text{RST}}$ low (hard reset) or by writing a 0 to the control register via the data bus (soft reset). A soft reset does *not* affect data in the AD1334's FIFO. A hard reset destroys all FIFO data and clears all control register bits.

The two types of resets have similar effects. Both resets will:

- Place all four SHAs into Track mode,
- Lock out all four SAMPLE inputs,
- Set $\overline{\text{READY}}$ to 1,
- Halt an in-process A/D conversion, and
- Clear any pending conversions.

Hard and soft resets differ in their immediacy. A hard reset is asynchronous and begins as soon as the $\overline{\text{RST}}$ pin is pulled low. Soft resets are synchronized by the clock so there is a delay of 1–2 clock periods, depending upon the timing of CONTROL $\overline{\text{ENB}}$ relative to the clock before a soft reset takes effect. An A/D conversion begun prior to a hard reset is terminated immediately by that reset. A conversion started before a soft reset terminates when $\overline{\text{READY}}$ goes high.

The AD1334 remains inactive following a hard reset until it is enabled by forcing CONTROL $\overline{\text{ENB}}$ low. This is done by writing a 1 to Bit 7 of the control register or by driving the package pin to a logical 0. *Only open-drain or open-collector gates should be used to drive the CONTROL $\overline{\text{ENB}}$ pin. The AD1334 may be damaged if this pin is driven from a standard totem-pole output.*

An uncertainty of one clock period in the enable and disable timing arises because the state of CONTROL $\overline{\text{ENB}}$ is latched by the rising edge of the 2.5 MHz clock.

The CONTROL $\overline{ENB}$ setup time is 20 ns. Violation of this requirement may delay the effects of the soft reset or enable by one clock period. The uncertainty is eliminated when transitions of CONTROL $\overline{ENB}$ are synchronized by the clock's falling edge.

Timing for enabling and disabling the AD1334's controller appears in Figure 2.

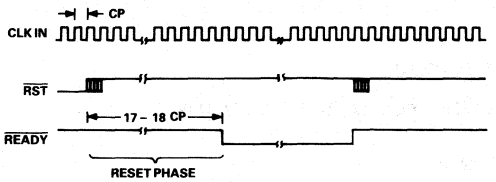

Figure 2a. Timing Diagram for Resetting Controller Using $\overline{RST}$ (CONTROL $\overline{ENB}$ = 0)

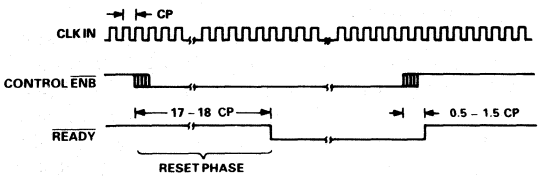

Figure 2b. Timing Diagram for Disabling and Enabling Controller Using CONTROL $\overline{ENB}$ ($\overline{RST}$ = 1)

Controlling Multiple AD1334s

Figure 3 shows the digital signal connections for a simple 16-channel simultaneous sampling system using four AD1334s. Each device's $\overline{SIMULT}$ input must be grounded. The host system's reset signal drives all four $\overline{RST}$ inputs. The 16 SAMPLE inputs, 4 per device, are driven by the system's sample command. These parallel-connected inputs represent a capacitive load in excess of 80 pF, and an appropriate buffer with good drive capability is required to maintain edge fidelity and preserve the system's aperture delay and aperture jitter performance. All parallel data outputs are connected to the system's data bus. Each AD1334 requires unique address decoding for programming and data readback. The four CONTROL $\overline{ENB}$ pins are tied together and share a common pull-up resistor to V_{DD}.

Any one of the AD1334s may be designated as the master device; the remaining devices will be slaves. In the simplest case only the master's $\overline{IRQ}$ need be monitored because all four devices are perfectly synchronized and always contain the same amount of data. (When interrupt-on-overrange operation is desired, the four $\overline{IRQ}$ outputs should be tied together and pulled up to V_{DD} with a resistor. After every interrupt each device must then be polled to determine if the interrupt was caused by valid data, e.g., FIFO half-full, or by an out-of-range result. A specific overrange channel can be identified only by checking each FIFO word stored in the interrupting device.)

The AD1334s must first be initialized using a hard reset. This will reset each unit's FIFO controller, clear the control registers, and leave the devices disabled with all SAMPLE inputs locked out. Next, the desired operating mode(s) should be written to the slave units. All slave devices must be programmed identically (FIFO enabled-/disabled, interrupt on FIFO full/half-full, overrange enabled/disabled). *Bit 7 must be set to 0 when programming the slave units.* Finally, the master device is programmed for the same operating conditions as the slaves, but with its Bit 7 set to 1. This enables all four AD1334s simultaneously, with each device's operation starting from the same state and with all units perfectly synchronized. All sixteen SHAs will enter Hold mode on the first rising edge of the system SAMPLE input after the enable sequence is completed. This scheme can be expanded to provide any number of simultaneously sampled input channels.

System interrupt overhead is minimized when using the FIFOs in each AD1334. A single interrupt can signal the availability of either 64 or 128 data words, depending on the chosen FIFO interrupt condition.

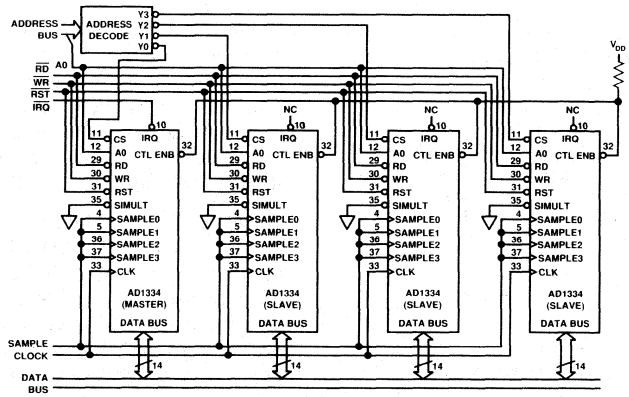

Figure 3. Digital Signal Connections for a Simple 16-Channel Simultaneous Sampling Data Acquisition System Using Four AD1334s

ANALOG DEVICES

AN-296
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Build a Single-Shot Recorder to Catch Fast Transients

by Ken Deevy, Dan Sheehan and Mike Byrne

Capturing fast transients places special requirements on filters, track/hold amplifiers, and A/D converters. By using an A/D converter with a high input bandwidth and oversampling at a 10:1 ratio, you can digitize and then analyze transients without using an expensive analog or digital storage scope.

Don't tie up expensive equipment trying to capture transients that occur infrequently. If you build a low-cost transient recorder or event sampler, you can dedicate it to capturing single-shot events. Typical applications for transient recorders include monitoring power-mains transients, evaluating power supplies, and capturing pressure and vacuum-line transients in medical equipment.

To build a transient recorder or burst-mode event sampler, you need a high-speed A/D converter, a wide-band track/hold amplifier, and an antialiasing filter. The A/D converter must have a sampling rate of at least twice the bandwidth to satisfy the Nyquist criterion. In practice, you should oversample the input signal. At 2× oversampling (a sampling frequency of twice the input bandwidth), you'll need to use a filter with an infinite roll-off rate to avoid aliasing effects. At 3× oversampling, the roll-off requirement drops to 50 dB/octave in an 8-bit system. With an oversampling ratio of 10:1, the filter roll-off need be only about 16 dB/octave. (See **box,** "Oversampling reduces antialiasing requirements.")

High-speed sampling A/D converter chips routinely include track/hold amplifiers on the same chip. The AD7821 is an example of this trend. It combines a 100-kHz track/hold amplifier with a 1M-sample/sec 8-bit A/D converter. Because the A/D conversion rate is 10 times the input bandwidth, you don't have to design a complex antialiasing filter. In fact, if the input signal exhibits only a low-power spectral content at and above 500 kHz, you can eliminate the filter altogether.

The AD7821 uses a half-flash conversion technique to perform an 8-bit conversion in 660 nsec. A requirement of a 350-nsec signal-acquisition period between conversions results in a maximum acquisition rate of 1M samples/sec. You can operate the A/D converter with a single or dual supply for either unipolar or bipolar inputs.

Capture single-shot waveforms

One of the difficulties in capturing single-shot events is the speed at which the transient recorder circuit responds once the input signal has crossed a predetermined trigger point. If the recorder circuit responds too slowly, it can miss fast transients altogether.

Reprinted from EDN – June 20, 1991

To build a transient recorder or burst-mode event sampler you need a high-speed A/D converter, a wideband track/hold amplifier, and an antialiasing filter.

Therefore, to accurately capture fast events, you need a high-speed A/D converter and a wide-bandwidth track/hold amplifier. For example, an 8-bit A/D converter that has a 1-μsec conversion time can capture 1-μsec transients only if it's not preceded by a track/hold amplifier. If you match this A/D converter with a track/hold amplifier that has a 100-kHz bandwidth, the ADC can recover 6-μsec-wide 5V transients.

To simplify fault detection or take corrective measures, you need a transient recorder that can grab pretransient information. You can use this pretransient data to learn timing relationships between the transient and another waveform. Additionally, your recorder should be able to react to both positive and negative transients.

Another important criterion for transient recorders is cost. Although you could use a digital storage oscilloscope (DSO) to capture frequently occurring or very fast transients, dedicating a DSO to capturing random events would be an expensive proposition.

Transient recorder

A block diagram of a transient recorder (**Fig 1**) shows the minimum hardware you'll need to build a high-speed transient recorder with playback. For sim-

Oversampling reduces antialiasing requirements

In the spectrum of a periodically sampled waveform, the spectrum of the (unsampled) input-signal repeats around harmonics of the sampling frequency. Any frequency contained in the input signal is repeated above and below each harmonic of the sampling frequency. Therefore, in the spectrum of the sampled signal, the band between 0 and f_{IN} (the input spectrum), appears—among other places—between $f_S - f_{IN}$ and f_S, where f_S is the sampling frequency. Though you may be under the impression that the input-signal bandwidth is 100 kHz, if the sampling frequency is 1M sample/sec, a signal at 991 kHz in the input spectrum would appear as a 9-kHz "alias" component in the sampled signal's spectrum.

The purpose of an antialiasing filter is to remove or at least attenuate any noise or spurious signals that could be aliased back into the bandwidth of interest. **Fig A** shows the frequency response of an antialiasing filter for a generalized A/D converter. You determine the filter roll-off by drawing a straight line between the highest signal frequency of interest, f_{IN}, and the stopband attenuation frequency, $f_S - f_{IN}$. As the ratio of f_S to f_{IN} increases (that is, as the oversampling ratio increases) the slope of the line decreases.

In an 8-bit system, an ideal ADC's signal-to-noise ratio (S/N) is slightly greater than 256:1 or 48 dB. To avoid having noise limit the system performance, the ratio of the input signal to noise should exceed the approximately 48-dB limit imposed by the ADC. Here, the signal is the peak-to-peak value of the signal within

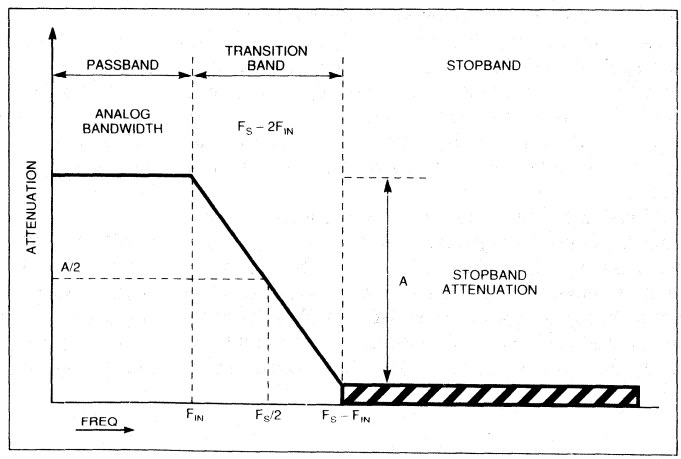

***Fig A—The antialiasing filter** that precedes the transient recorder's ADC can be simple or complex depending on the degree of oversampling. When the sampling frequency is 10× the highest frequency of interest, the filter has 3 octaves to roll in its attenuation. A simple 3-pole filter has 18 dB/octave roll-off. An 8-bit ADC needs slightly more than 48 dB of attenuation. Hence, a 3-pole filter is usually sufficient for 8-bit resolution.*

plicity, the design uses a clock with an even mark/space ratio. The clock's 50% duty cycle limits the acquisition rate to 660k samples/sec rather than the A/D converter's 1-M sample/sec maximum rate. (This simplification reduces the oversampling ratio to 6.6:1.) A memory chip stores the digitized data for later playback on an X-Y plotter or oscilloscope via a dual 12-bit D/A converter and a quad op-amp. One half of the samples are pretransient information; the other half are transient data.

A more detailed schematic (**Fig 2**) shows that two counters, IC_1 and IC_2, control where the circuit stores pretransient and transient data. The counters also

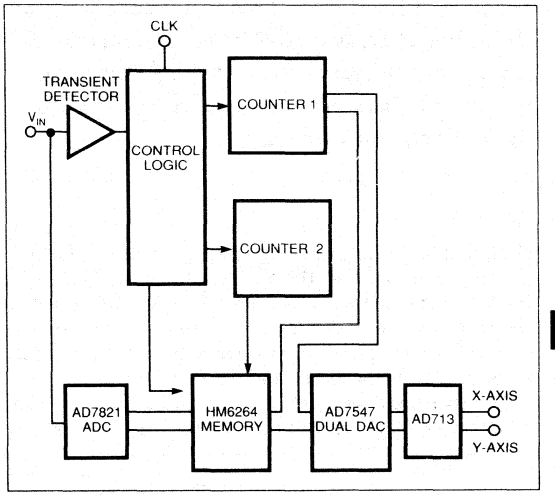

***Fig 1—If you think of a transient recorder** as merely an ADC, look again. This recorder contains several purely digital blocks as well as the DAC that drives the display.*

the band of interest, and the noise is the square root of the sum of the squares of the amplitudes of all of the frequency components outside that band. The attenuation required on signals outside the band of interest depends on the application and the expected magnitude of the out-of-band signals. In most cases, the magnitude of these signals is much lower than that of the desired signal.

Usually, 8-bit systems require 50 dB of attenuation for signals that can be aliased into the band of interest. Even if 50 dB is not the desired number, the following calculations show the kind of reduction in antialiasing filter requirements brought about by oversampling. With 2× oversampling, (that is with $f_S = 2 \cdot f_{IN}$), f_S and f_{IN} are at the same point and the filter has to have infinite rolloff to attenuate signals at $f_S - f_{IN}$. With $f_S = 3 \cdot f_{IN}$ (3× oversampling), the filter's attenuation must drop from 0 dB at f_{IN} to 50 dB at $2 \cdot f_{IN}$. In other words, the slope of the attenuation vs frequency curve must be 50 dB/octave; the filter (if it has a Butterworth characteristic) must have more than eight poles.

With 10× oversampling, there are three octaves for the attenuation to drop from 0 to 50 dB. The required slope is a little more than 16 dB/octave; a 3-pole Butterworth filter will do the job.

The above analysis of the antialiasing filter holds true regardless of the type of ADC that follows the filter. No matter what the conversion technique, oversampling reduces the antialiasing filter requirements. Oversampling also reduces the ADC noise within the signal bandwidth because it spreads the quantization noise over a wider bandwidth. Oversampling has recently gained considerable popularity in connection with sigma-delta ADCs. In the case of these converters, the advantages of oversampling are much greater than with successive-approximation or flash ADCs because noise shaping produces dramatic improvements in noise performance as the oversampling ratio increases.

The relationship between antialiasing-filter performance and oversampling is, however, exactly the same for an oversampled sigma-delta modulator as for a half-flash or a successive-approximation ADC. A sigma-delta ADC and a half-flash ADC with the same oversampling ratio place the same requirements on the antialiasing filter.

The disadvantage of the sigma-delta process for transient recording is the pipelining or averaging technique inherent in sigma-delta converters. Because of the pipelining, a step change requires a significant time (the settling time of the ADC's digital filter) to ripple through to the output. Therefore, there is a delay before a sigma-delta converter produces an output that represents an input change. Between the time the input changes and the sigma-delta converter's output reflects the change, the ADC's output does not accurately represent the converter's input. Such performance is not appropriate for transient recorders of the type discussed here.

To avoid aliasing effects, at 2× oversampling (a sampling frequency of twice the input bandwidth) you need a filter with an infinite roll-off rate.

clock out data to either the oscilloscope or the X-Y plotter. You can use the fast clock input, CLK IN_1, for the clock source when you're using the circuit in the record mode or displaying stored data on an oscilloscope. The design also provides a slower clock input, CLK IN_2, to print data on an X-Y plotter.

The transient recorder operates in two basic modes: record and playback. You select the record mode by placing switch S_1 in the record position. (IC_{18A} and IC_{18B} provide debouncing for this switch.) Having the MODE output of IC_{13D} low makes one input of both IC_{15D} and IC_{14D} low. Hence the clock inputs of IC_{9A} and IC_{9B} (pins 10 and 2, respectively) are disabled, ensuring that the $1\overline{Q}$ and $2\overline{Q}$ outputs of IC_{9A} and IC_{9B} are high. Besides disabling the chip-select inputs of the D/A converter, $\overline{CSA}$ and $\overline{CSB}$, the circuit disables the output-

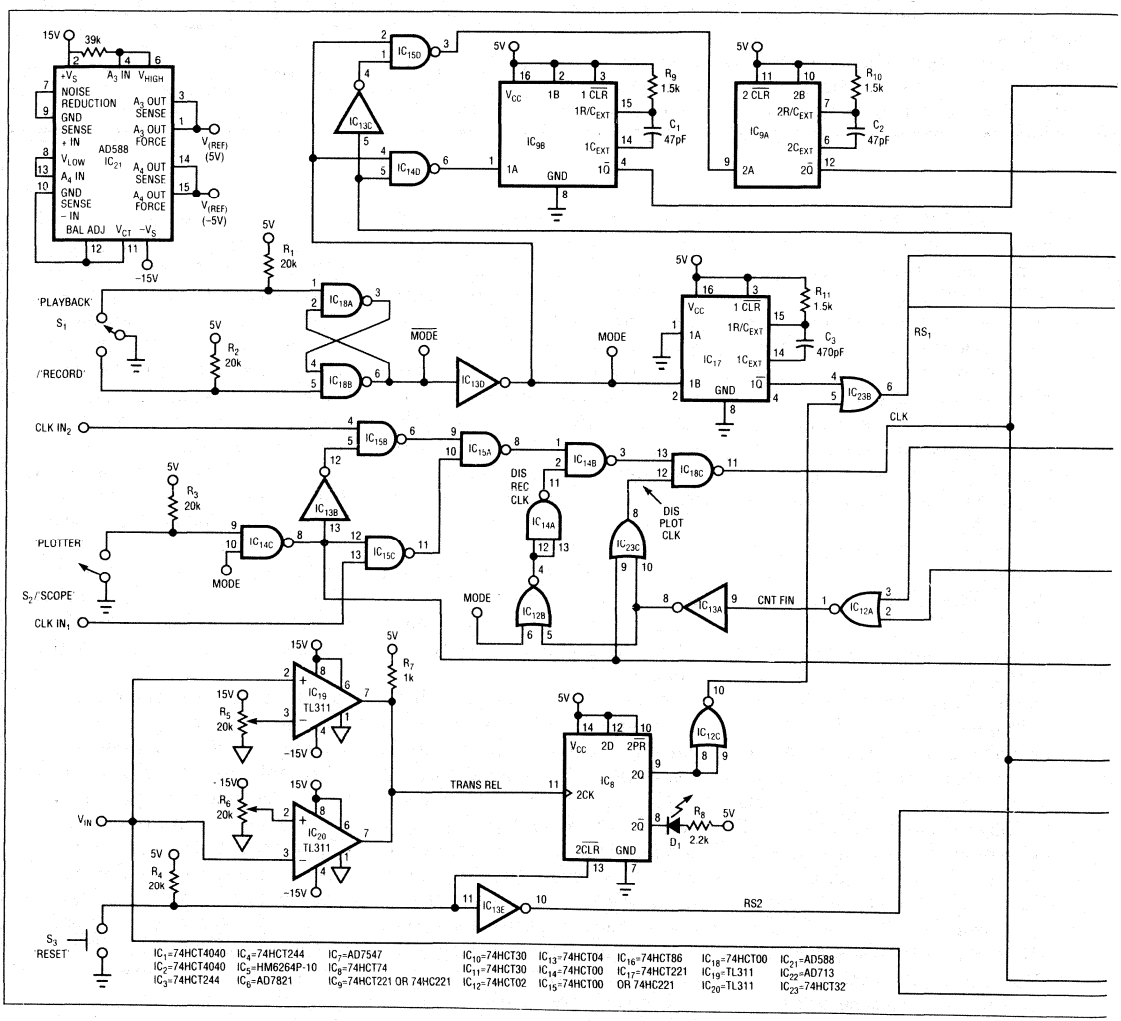

***Fig 2—At the component level,** the basic transient recorder is a circuit of moderate complexity, requiring 23 ICs.*

enable signals of IC_3, IC_4, and IC_5 (the HM6264 memory chip), ensuring that the playback portion of the transient recorder is turned off.

CLK IN_1 serves as the clock source for the counters via IC_{14C}, IC_{-13B}, IC_{15B}, and IC_{15C}. While the MODE signal is low, CLK is the clock input for both counters and provides the $\overline{RD}$ (convert) signal for the A/D converter, IC_6. At the same time, IC_6's $\overline{CS}$ input is active, ensuring that the device is selected. After a reset from S_8 initializes the circuit, counter 2 begins counting. IC_{17} and IC_{23B} hold the reset (CLR) input of counter 1 high from power-up and keep counter 1 in a reset condition until the circuit detects a transient.

You configure the A/D converter by tying its MODE input (pin 7) to GND. (Note that the MODE pin of the AD7821 shown bears no relation to the signal labeled

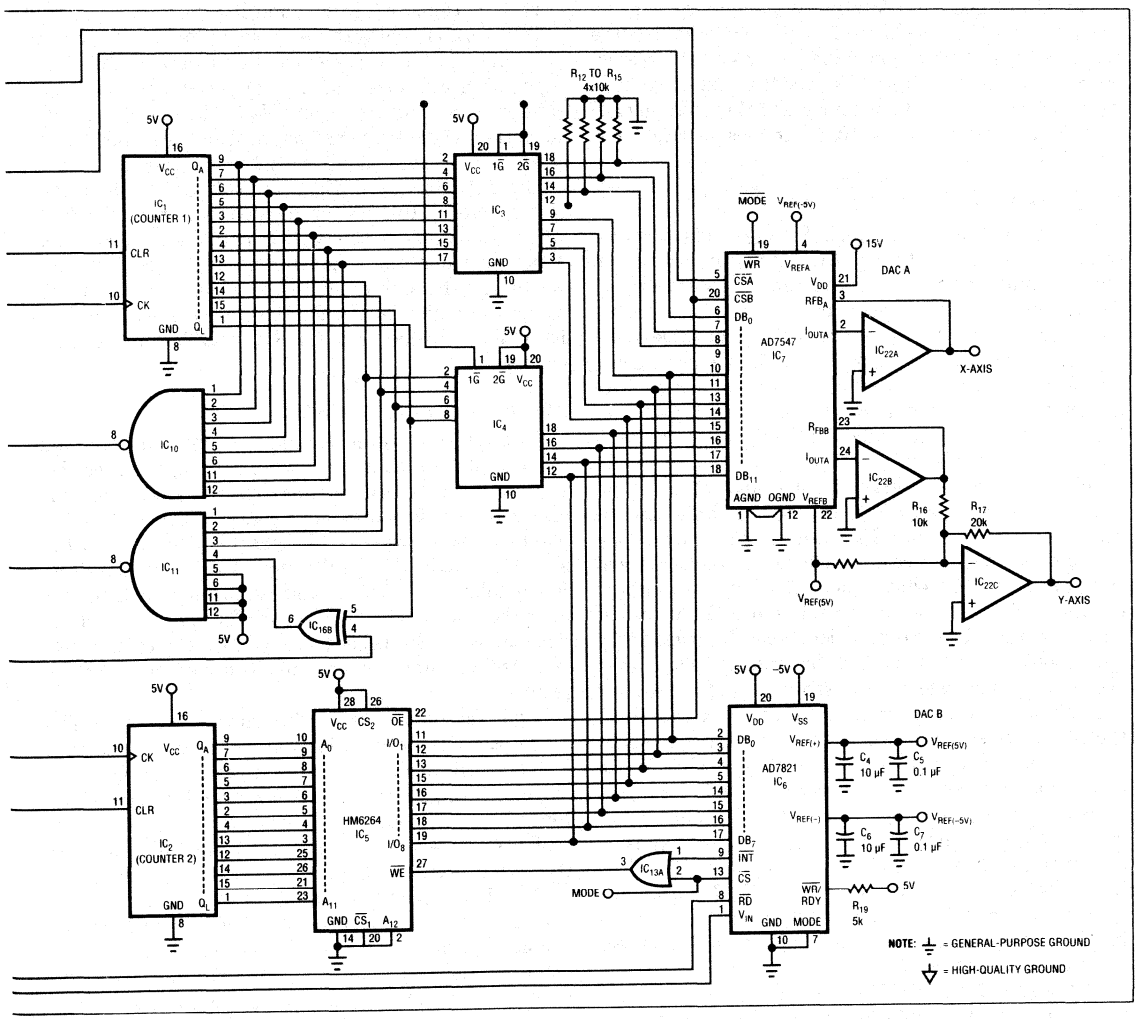

With an oversampling ratio of 10:1, the filter roll-off need be only about 16 dB/ octave.

MODE in the circuit diagram.) When the CLK signal toggles its $\overline{RD}$ input, the A/D converter executes continuous conversions of the input signal, V_{IN}. Counter 2 provides the memory addresses for the A/D conversion results. Data transfers from the digital outputs of IC_6 to IC_5 employ the $\overline{INT}$ output of IC_6 to drive the $\overline{WE}$ input of IC_5.

The circuit automatically loads the first conversion result after reset into location 0 of memory and the second result into location 1. After transferring the result of the 4096th conversion to memory location 4095, the counter resets and stores the next conversion result in location 0. The memory will therefore always contain the most recent 4096 samples of the input waveform.

Detect fast transients

You apply the input signal to V_{IN}. This terminal connects directly to two TL311 comparators and the analog input of the A/D converter. Comparator IC_{19} detects positive transients, and IC_{20} detects negative ones. To set the threshold level for a positive-going signal, adjust resistor R_5; adjust R_6 for negative-going transients. Wire the outputs of both comparators together to ensure that they produce a rising edge to the clock input of IC_8 when either a negative or a positive transient occurs.

Once the circuit detects a rising edge at pin 11 of IC_8, it illuminates an LED, D_1. At the same time, it releases counter 1 from its reset condition by taking RS_1 low. Now the circuit clocks both counter 1 and counter 2 as A/D conversions continue. Counter 2 counts up from the value it held before the transient was detected. The memory locations determined by the output of counter 2 store the transient data while overwriting the oldest 2048 samples of pretransient data already stored in memory. Counter 1 counts off the 2048 clock states that correspond with the samples.

Because the output of IC_{16B} is always high in the record mode, when counter 1 reaches 2047, all inputs to IC_{10} and IC_{11} are high and both IC's outputs go low. As a result, the output of IC_{12A} goes high, causing the output of IC_{14A} to go low via IC_{13A} and IC_{12B}. This DIS REC CLK signal gates off CLK IN_1 from the rest of the circuit at IC_{14B}. The output of IC_{18C} ensures that the CLK signal is held low, stopping both counters and the A/D converter.

At the end of the transient-record cycle, the memory will contain 4096 samples of the input waveform. One half of these samples are transient data, the other half represent pretransient information. Whatever value is in counter 2 will be the last memory location for the transient data. The next memory location will hold the first of the 2048 words of pretransient data. Now when you start the playback mode, the first output from the counter will correspond to the memory location of the first pretransient sample. (To alter the ratio of transient to pretransient samples, simply alter the connections from counter 1 to IC_{10} and IC_{11}.)

To accurately convert the input waveform to stored data, you must pay close attention to the circuit. Use a precision reference, IC_{21}, to generate 5V and −5V references for the V_{REF+} and V_{REF-} inputs of the A/D converter. Make sure that you properly decouple these reference voltages along with the V_{DD} and V_{SS} lines of the A/D converter. Connect IC_6's GND pin to the star ground of the system (the point in the circuit at which you connect the analog and digital grounds). Make sure that the conductor between the A/D converter and the star ground is as wide as circuit-board layout constraints allow. Further, ensure that the $\overline{WR}$/RDY line is pulled high via R_{19} to avoid noise pickup on this pin.

Play back captured signals

Once you've captured that bothersome transient, you can play it back at any convenient time; the recorder will retain the information as long the power remains on, or until you depress the reset button. Select the playback mode with S_1. Playback takes the MODE signal low, activates the $\overline{WR}$ input to IC_7, and deselects IC_6 by taking its $\overline{CS}$ high. You can display the transient on either an analog oscilloscope or an X-Y plotter, depending on the position of S_2. Make sure to select the oscilloscope or the plotter before starting playback.

If you decide to display the transient on an oscilloscope, the clock source for the circuit is the same as in the record mode. If you use a plotter for playback, the clock frequency is much lower and is applied via the CLK IN_2 input. CLK (from either CLK IN_1 or CLK IN_2) passes through gates IC_{15D} and IC_{14D} because the MODE signal is high. IC_{9A} and IC_{9B} generate the $\overline{CSA}$ and $\overline{CSB}$ pulses for IC_7 from this CLK signal.

IC_{9A} drives the $\overline{CSA}$ input of IC_7 as well as the enable signals for IC_3 and IC_4. When you choose the playback mode, counter 1 resets and starts counting from 0 to 4095. The counter's output is the digital input code to DAC A of IC_7. This DAC drives the X axis of either the oscilloscope or the plotter. DAC A produces

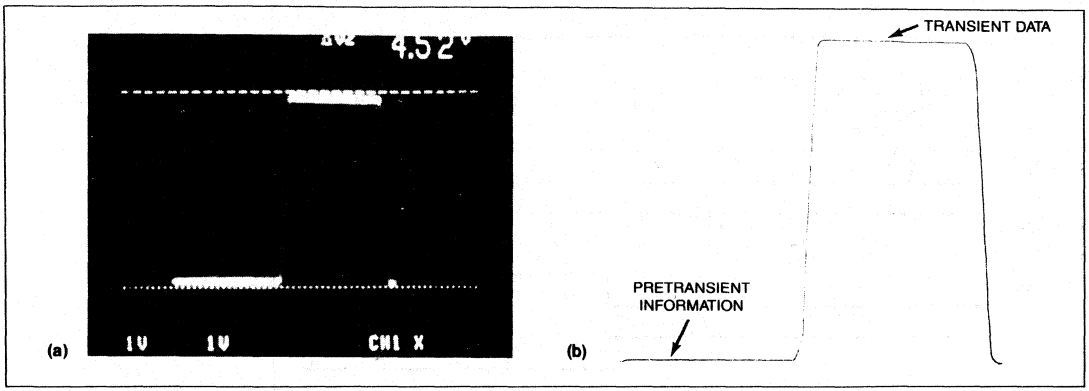

Fig 3—The recorder can display *captured data on a scope,* ***a,*** *or on an X-Y pen plotter,* ***b.*** *The first half of each display shows data acquired before the triggering event.*

a unipolar output range from 0 to 5V, with a resolution of 4096 steps.

The output of IC_{9B} drives the $\overline{CSB}$ input of IC_7 and also sets the logic level on IC_5's output-enable line, $\overline{OE}$. This action latches the data from memory into DAC B, which drives the Y axis of the oscilloscope or plotter. By using dual supplies, you can set DAC B for a bipolar output range to reconstruct both positive and negative transients.

Counter 2 starts its count from the point at which it stopped at the end of the record mode; the first memory output word to IC_7 is the oldest sample in memory. The scan will then proceed through the 2048 samples of pretransient information and the 2048 samples of transient information. The output of each sample from memory to the Y axis, via DAC B, corresponds to the output of a count value from counter 1 to the X axis via DAC A. In this way, the circuit reconstructs the pretransient and transient waveforms.

For oscilloscope display of waveforms, place S_2 in the scope position. Doing so locks out CLK IN_2 from the rest of the circuit but allows CLK IN_1 to operate as clock signal for the circuit. Unlike the plotter display option, where counter 1 runs through once and then stops, CLK runs continuously. CNT FIN does go high when counter 1 reaches a count of 4095, but because the output of IC_{14C} is high, the DIS PLOT CLK signal does not go low. You can see the typical oscilloscope waveform display in **Fig 3(a)**.

You display the stored waveform on an X-Y plotter by placing S_2 in the plotter position. Doing so locks out the CLK IN_1 input from the rest of the circuit and permits CLK IN_2 to generate the clock signal for the circuit. IC_{16B}, IC_{10}, IC_{11}, and IC_{12A} function in a manner similar to the record mode to generate a high CNT FIN signal. But this time, IC_{10} and IC_{11} go low when counter 1 reaches a count of 4095. IC_{13A} goes low, and, because the output of IC_{14C} is already low, the DIS PLOT CLK signal goes low, turning off CLK IN_2 at IC_{18C} and holding the CLK signal high. **Fig 3(b)** shows a captured transient displayed using a plotter as the display method.

Record-mode timing and clock waveforms

The timing diagrams in **Fig 4** show the logic relationships for the record mode. The MODE signal (not shown) is low and the DIS REC CLK signal is high. The RS_2 signal goes high when the recorder receives a reset command via S_3 resetting counter 2. The next falling edge of the CLK signal clocks out an address for IC_5 from counter 2. A conversion is also initiated on this falling CLK edge, and, within 700 nsec, the $\overline{INT}$ signal of IC_6 goes low, activating the $\overline{WE}$ input of IC_5. The rising edge of CLK resets the $\overline{INT}$ line 50 nsec later.

When the circuit detects a transient, the TRANS REC signal goes high, causing the RS_1 line to go low and releasing counter 1 from its reset state. The next falling edge of CLK clocks out the outputs from counter 1. When the count output from counter 1 reaches 2047, the CNT FIN signal goes high and causes the DIS REC CLK signal to go low, shutting off the CLK signal.

To accurately capture fast events, you need a high-speed A/D converter and a wide-bandwidth track/hold amplifier.

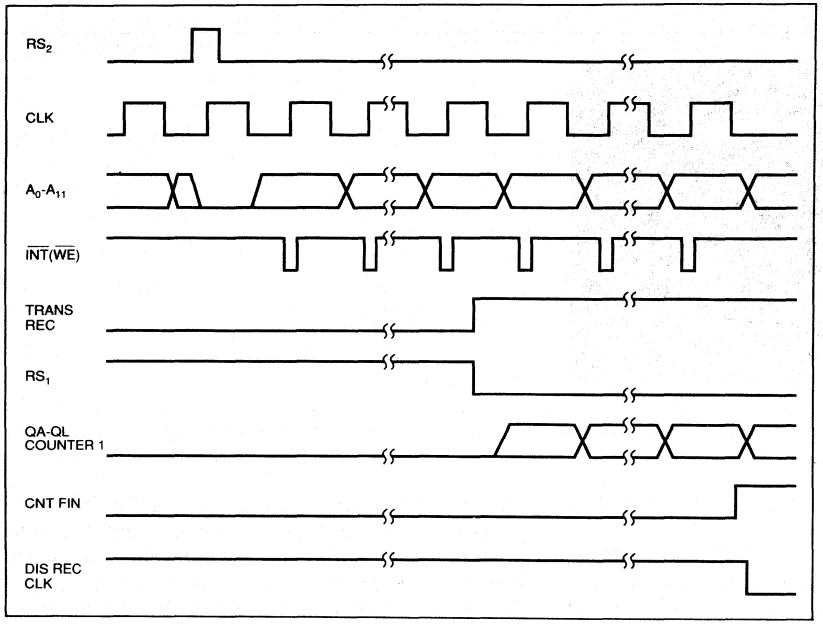

***Fig 4—The timing relationships** in the record mode show that recording starts on the first falling clock edge after the reset line goes high.*

In the record mode in the waveforms shown, the 50/50 mark/space ratio of the CLK signal limits the clock frequency to 660 kHz. You need a CLK-low time of 750 nsec for the A/D converter to perform its conversion correctly and latch the data into IC_5. However, the CLK-high time can be as short as 350 nsec, the time required between conversions by the AD7821. Therefore, if the input to CLK IN_1 has a low time of 750 nsec and a high time of 350 nsec, the circuit can make one conversion every 1100 nsec—equivalent to approximately 900k samples/sec.

Playback to a scope

During playback to an oscilloscope, **(Fig 5(a))**, the MODE signal, the $\overline{WE}$ input of IC_5, and the DIS REC CLK signal are high. When you place S_1 in the playback mode, RS_1 goes high, resetting counter 1. The CLK signal generates a $\overline{CSA}$ signal for IC_7 on its rising edge and a $\overline{CSB}$ signal on its falling edge. The falling edge of the CLK signal clocks data from counter 1, and the rising edge of $\overline{CSA}$ updates the X axis. The falling edge of $\overline{OE}$ outputs stored data from memory, and the rising edge of $\overline{CSB}$ updates the Y axis. The CLK signal runs continuously when the circuit is in the scope-playback mode.

The timing diagram of **Fig 5(b)** shows operation of the circuit for playback on a plotter. Once again, MODE, the $\overline{WE}$ input of IC_5, and the DIS REC CLK signals are high. The circuit generates $\overline{CSA}$ and $\overline{CSB}$ to update the X and Y axes. Compared with scope playback, the difference in the circuit's operation is that when the output count from counter 1 reaches 4095 and the CNT FIN signal goes high, the DIS PLOT CLK signal goes low, forcing the CLK signal into a high state.

Burst-mode event sampling places requirements on an A/D converter similar to those for transient recording. In burst-mode sampling, the recorder looks at the input waveform infrequently, but when it does, it must acquire a large number of samples in a short time. With slower microprocessors or microcontrollers, you'll find that because of instruction- and bus-timing constraints, you can't achieve anything like the A/D converter's maximum throughput.

You can overcome timing limitations in a burst-mode sampler by using a DMA controller to initiate A/D conversions and transfer conversion data to memory. Doing so lets you run the A/D converter at or near its maximum sample rate, permitting high oversampling ratios and the acquisition of short transients.

Building a burst-mode sampler is a simple matter with the popular 8052 microcontroller **(Fig 6)**. Al-

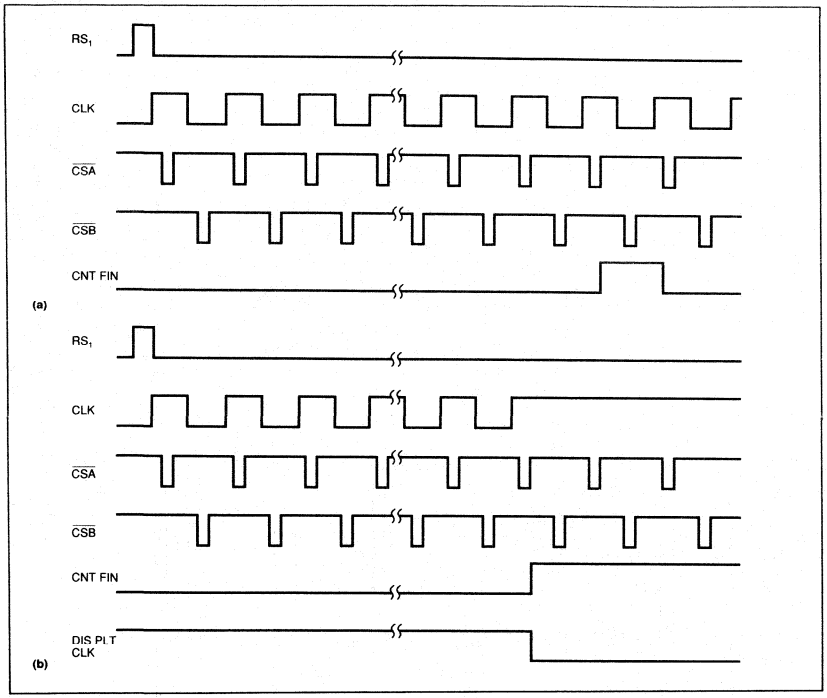

Fig 5—The timing for waveform playback *is almost the same in the scope mode* ***(a)*** *as in the plotter mode* ***(b)****. The difference is that the scope playback runs continuously, whereas the circuit produces a plotter output only once.*

though the 8052 does not support hardware DMA, it does support what is termed "fake DMA." However, expect the response time to DMA requests to be much slower than what is possible with microcontrollers that support genuine DMA.

The HM6264P memory chip, IC_3, stores the control program for IC_1. The first part of the control program is the initialization routine. This routine (**Listing 1**) sets up the sense of the DACK0 line of the 8237, IC_2, to be active high. It also loads the starting data address into IC_2 for the first conversion results. IC_1 initializes the counting register to control the number of conversions before IC_2 returns control to IC_1. The program must also set up IC_1 for "fake DMA."

Once you've run the initialization program, IC_2 is ready to take control when requested to do so. Although IC_2 has four interrupt-request lines, this circuit uses only one, DREQ0. An external command signal drives this interrupt line high, telling IC_2 to take control of the circuit and start the A/D converter sampling the input waveform.

After IC_2 receives the DREQ0 request (**Fig 7**), its HRQ line goes high and feeds IC_{14C}, which takes the $\overline{INT0}$ line of IC_1 low. IC_1 responds to this "fake DMA" request by bringing its P1.6 line low and the output of IC_{14A} high, selecting inputs of IC_7, IC_8, IC_9, and IC_{10}. When the output of IC_{14A} goes high, it shuts off IC_1's address and data lines from the rest of the circuit and deselects the output's address decoder, IC_{13}. The inverted P1.6 line also feeds the HLDA input of IC_2, acknowledging IC_2's request for control. IC_2 then takes control of the address and data bus and the sampling of the input waveform.

To reduce pin count, IC_2 multiplexes the eight higher-order address bits on the data lines. You need an external device to latch these address bits. The address strobe signal, ADSTB, takes AEN high and switches the $\overline{OC}$ line of IC_6 low. ADSTB drives the C input of IC_6 to latch the higher address lines to the outputs of IC_6. The inverted AEN line also drives one input of IC_{16D}. The decoded output, Y_0, of IC_{13} controls the other input of this gate. Therefore, either a high on AEN or a low on the decoder output selects IC_3. You need this control logic because both IC_2 and IC_1 must be able to access IC_3.

The DACK0 line goes high at about the same time that ADSTB latches the address and drives one input of IC_{15A}. IC_{15A} and IC_{15B} ensure that the $\overline{CS}$ line of IC_4 goes low only when an input/output read operation of IC_2 occurs. IC_{15C} provides the correct polarity for the $\overline{RD}$ input and equalizes the delay paths for the $\overline{CS}$ and $\overline{RD}$ lines, ensuring that the circuit obeys the $\overline{CS}$-to-$\overline{RD}$ setup time.

Once IC_4 receives a $\overline{CS}$ signal, it acknowledges receipt of the signal by bringing its RDY line low, placing the controller, IC_2, into a wait state for as long as its READY input is low. When the device completes a conversion, the RDY line goes high, releasing IC_2 from its wait state. Because IC_4's RDY output is an open-drain output, you need to install an external pullup resistor, R_2.

When the circuit releases IC_2 from its wait state,

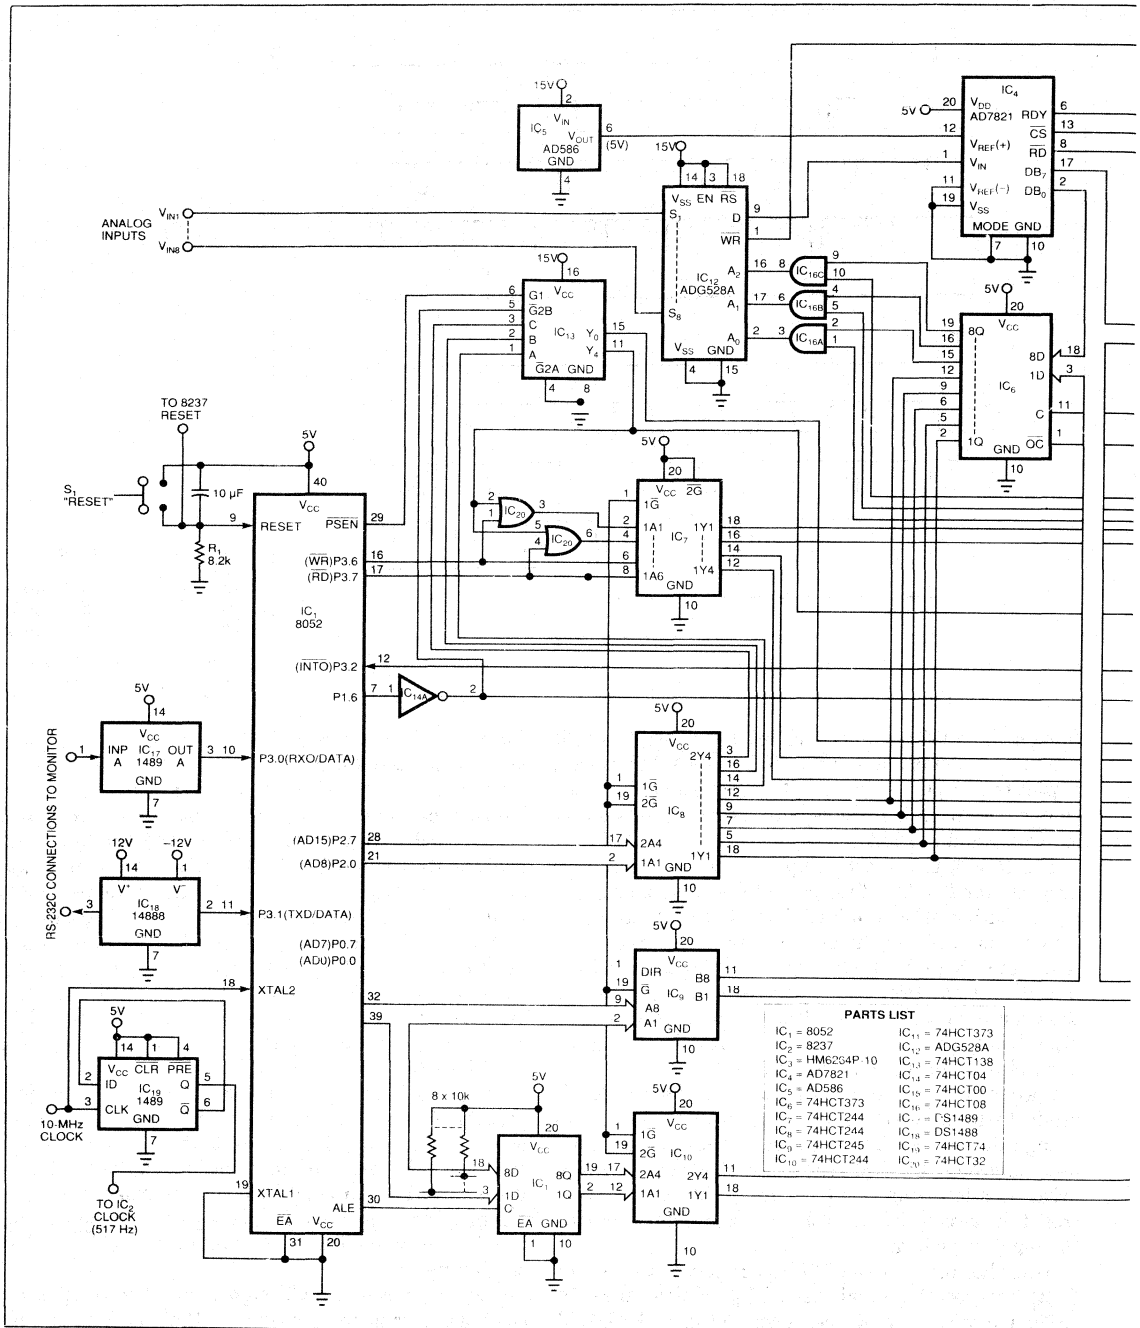

***Fig 6—A burst-mode sampler** using an 8052 microcontroller is roughly equal in complexity to the transient recorder shown earlier.*

data from IC_4 is valid. The address lines of IC_2 determine where data loads into memory. IC_2 performs all of these operations automatically because a memory write accompanies each input/output read. Depending on the value loaded into the counting register, IC_2 will continue to issue read commands to IC_4 until the circuit completes the required number of conversions. IC_2 automatically increments the memory address after every write operation.

The multiplexer, IC_{12}, accommodates eight input

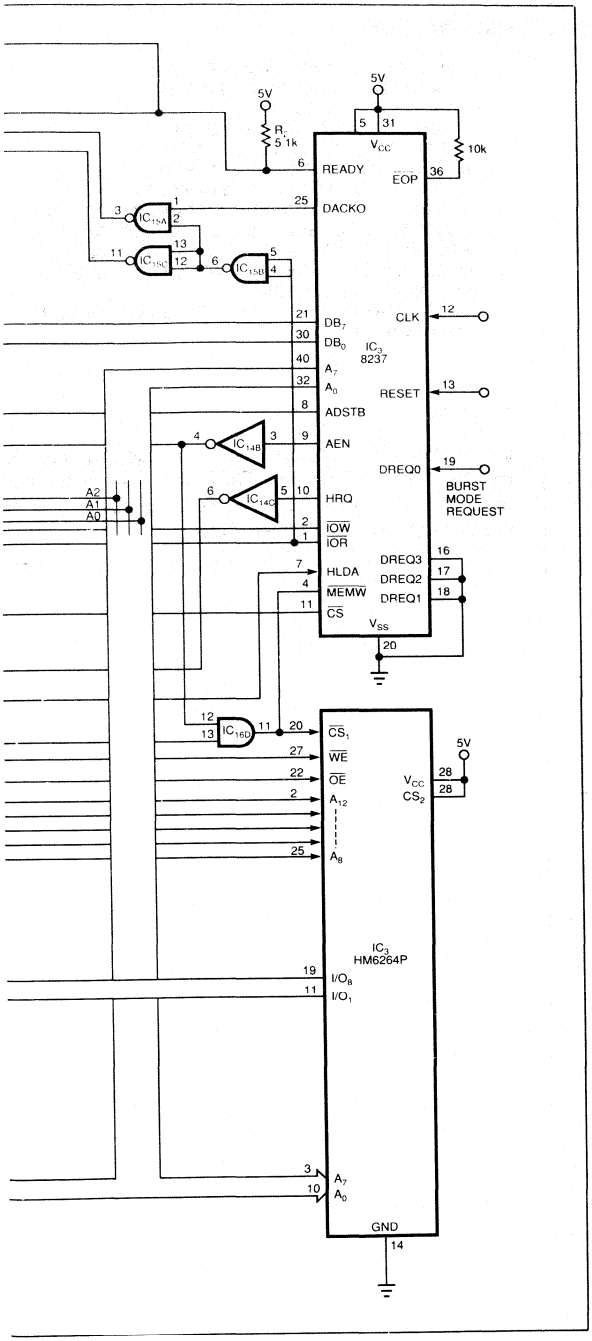

channels. The three upper and three lowest address lines of IC_2, gated through IC_{16A}, IC_{16B}, and IC_{16C}, select the input channel. If the three upper address lines are set to all 1s, IC_4 will convert each channel in sequence, and the conversion results will be stored in consecutive memory locations. For example, if the first conversion takes place on the channel-1 input voltage, V_{IN1}, and the result is stored in location M of IC_3, the next conversion will take place on V_{IN2}, and the result will be stored in location M + 1. If the three uppermost address bits are set to 011, the circuit will sequence through channels 1 to 4 only.

3

Ready or not

The RDY line of IC_4 drives the $\overline{WR}$ input of IC_{12}, loading the address for the next channel to be converted into the multiplexer. When you have only one input channel to convert, you can use an alternate design: Remove IC_{16A}, IC_{16B}, and IC_{16C}, and drive the A0, A1, and A2 inputs of IC_{12} directly from the three uppermost address lines. Using this scheme, the program chooses the input channel.

IC_1 selects the device it talks to using a 1-of-8 address decoder, IC_{13}. The outputs of IC_{13} provide signals for IC_{12}'s $\overline{WR}$ line, IC_3's $\overline{CS}$ input, and IC_2's $\overline{CS}$ input. One of the outputs also gates the P3.7 ($\overline{RD}$) and the P3.6 ($\overline{WR}$) outputs from IC_1 to drive the $\overline{IOW}$ and $\overline{IOR}$ inputs of IC_2. The upper three address lines of IC_1 select the required device. The lower address lines are multiplexed with the data lines in a manner similar to the way IC_2's address and data lines are multiplexed. IC_{11} demultiplexes the lower eight address lines. IC_1's ALE signal latches these address lines. The 3-state buffers, IC_7, IC_8, and IC_{10}, isolate IC_1 outputs from the address bus when IC_2 takes control. You must use these buffers because IC_1 can't place its address and data buses into a high-impedance state when IC_2 takes control of the circuit. IC_9 also acts as a buffer but is bidirectional because IC_1 must read data from and write data to memory.

IC_1 uses a 10-MHz input-clock frequency. A 74HCT74 counter (IC_{19}) divides down this clock to form the clock input to IC_2. The standard 8237 operates from a 3-MHz maximum clock frequency, so you can divide the 10-MHz clock by 4 to provide IC_2's clock. You'll have a resultant acquisition rate of 608k samples/sec. A faster version of the 8237, the 8237-5, operates from a 5-MHz input clock, allowing you to divide the clock frequency by 2 and enabling the circuit to take 812k samples/sec. If you were to use IC_1 on its own to control the sampling of the input waveform, the best acquisition rate you could obtain would be approximately 100k samples/sec.

The entire circuit operates from 15V and 5V sup-

Listing 1—Initialization Routine

```
10      XBY(8008H) = 80H          : SETS DACK SENSE ACTIVE HIGH
20      XBY(800FH) = 0EH          : CLEARS DREQ0 MASK REGISTER
30      XBY(800BH) = 94H          : SETS MODE REGISTER
40      XBY(800CH) = 00H          : CLEARS FIRST/LAST FLIP-FLOP
                                  :  (ONLY REQUIRED IF 8237 IS
                                  :  NOT RESET BETWEEN DMA REQUESTS)
50      XBY(8000H) = 00H          : LOADS LOWER BYTE OF STARTING DATA
                                  :  ADDRESS TO BASE AND CURRENT ADDRESS
60      XBY(8000H) = 08H          : LOADS HIGHER BYTE OF STARTING DATA
                                  :  ADDRESS TO BASE AND CURRENT ADDRESS
70      XBY(8001H) = 00H          : LOADS LOWER BYTE OF COUNTING NUMBER
                                  :  TO COUNT REGISTER
80      XBY(8001H) = 02H          : LOADS HIGHER BYTE OF COUNTING NUMBER
                                  :  TO COUNT REGISTER
90      DBY(38) = DBY(38).OR.02H
100     IE = IE.OR.81H
110     GOTO 10
```

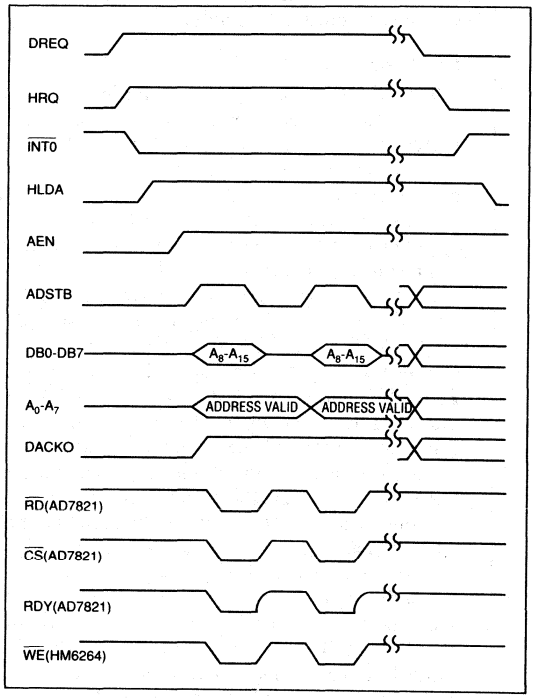

***Fig 7—Although the 8052** does not support true DMA, you can create a "fake DMA" mode, which, though not as fast as real DMA, lets you transfer blocks of data directly to memory.*

plies. If there isn't a 5V supply in your system, you can add a regulator to generate 5V. In addition, plan to use a precision 5V reference (IC_5) for the A/D converter, allowing an input range of 0 to 5V. To obtain accurate conversion results, you must obey the same guidelines regarding decoupling and grounding as apply to the transient-recorder circuit.

You can use the same design (**Fig 6**) with slow- and medium-speed microprocessors that support DMA requests. With these microprocessors, you'll find the DMA response time will be much faster than the response of the 8052's "fake DMA." Because microprocessors that support genuine DMA will 3-state their address and data lines during a DMA transfer, you can eliminate the 3-state driver chips.

ANALOG DEVICES

AN-297
APPLICATION NOTE

NE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Test Video A/D Converters Under Dynamic Conditions

by Walt Kester

To adequately characterize video-speed analog-to-digital converters, collect and evaluate the test results from comprehensive dynamic measurements.

To check out video-speed analog-to-digital converters—those operating at sampling rates exceeding 10 MHz—use dynamic performance testing. Such testing accumulates data that accurately describes these converters' ability to digitize fast-acting analog input signals.

An alternative test approach traditionally used with low- to medium-speed converters, static testing proves deficient for high-speed devices because video A/D converters demonstrate virtually ideal digitizing transfer functions under dc or low-frequency ac stimulation. However, under dynamic stress—ac test signals approaching the Nyquist sampling-rate limit $f_S/2$—video-speed units often exhibit otherwise hidden shortcomings, such as missing codes, nonmonotonic conditions or linearity errors.

On the debit side, though, dynamic testing lacks the standardization of industry-accepted static testing. Consequently, users tend to focus only on those ac specs that highlight a particular video-converter application. For example, differential gain and phase are paramount in digital-video uses, and noise-power-ratio checks dominate evaluation of parts aimed at data-communication tasks.

Moreover, in practice, no single specification or test can completely characterize a video A/D converter's performance under dynamic conditions. Therefore, you must implement a variety of probing analog and digital tests. Equipped with the test results, however, you can then generate a set of converter specs that accurately describe those parameters affecting your application.

DAC eases ADC testing

Connecting a high-performance D/A converter to the video A/D converter under test permits dynamic performance measurements with conventional analog test equipment. To obtain valid test results, however, the DAC's static and dynamic performance must exceed that of the A/D converter by two bits or more.

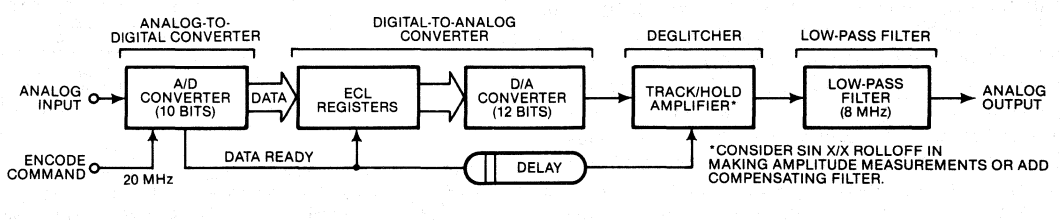

Fig 1—Connecting a high-performance digital-to-analog converter *to the video A/D converter under test in a back-to-back setup allows you to perform dynamic measurements using common analog test equipment. The DAC, however, must function at least four times better than the A/D converter to ensure accurate test data. The track/hold amplifier/deglitcher removes undesirable harmonics or discontinuities from the output signal.*

Reprinted from EDN — August 18, 1982

Use a D/A converter to dynamically test an A/D converter

Fortunately, though, because many DACs achieve such higher performance, back-to-back A/D-to-D/A testing proves successful (**Fig 1**). In this method, the DAC must operate transparently over the input frequency range and with the A/D converter's sampling rate.

For accurate dynamic testing, the DAC's settling time must not surpass the period corresponding to the A/D converter's maximum sampling rate. In addition, a deglitcher connected to the DAC's output removes unwanted harmonics caused by DAC glitches or analog-output discontinuities. Essentially a track/hold amplifier, the deglitcher switches to its Hold mode immediately before the DAC gets updated. During DAC updating in the Hold mode, the glitch settles out. Then, the track/hold enters its Track mode and acquires the DAC's new output value, typically in less than 30 nsec.

Note that the deglitcher inserts a glitch in the DAC's output. However, the glitch occurs at the test system's sampling rate and is therefore filterable. For analysis purposes, a sinc function—a (sin x)/x curve—describes the deglitcher output's frequency response. This response stems from the deglitcher's signal-reconstruction process: Its output is a series of rectangular pulses whose widths equal the sampling frequency's reciprocal.

When this pulse stream passes through the low-pass filter, the attenuation at frequency f with a sampling rate of f_S equals

$$A_f = \frac{\sin(\pi f/f_s)}{\pi f/f_s}.$$

System measurements that depend on frequency response must take into account this theoretical rolloff or must be made using a compensating filter.

This sinc correction filter's output goes through the low-pass filter, which has a cutoff of approximately $f_S/2.2$. The filter's attenuation at frequencies equal to or greater than $f_S/2$ should be at least 10 dB higher than the A/D converter's dynamic range to prevent aliasing errors.

To ensure accuracy, how do you verify the DAC's and deglitcher's dynamic performance independently of the A/D converter? One method involves driving the DAC with a PROM loaded with an ideally quantized recirculating sine wave and examining the deglitcher's output for harmonics using a spectrum analyzer. If you store two sine waves of slightly different frequencies in PROM, you can measure dual-tone intermodulation products. A less desirable method centers on driving the DAC with an A/D converter having known characteristics.

SNR tests deserve priority

Probably the most powerful indicator of an A/D converter's dynamic performance is its signal-to-noise ratio (SNR) when stimulated by a spectrally pure sine wave. Ideally, an A/D converter generates $q/\sqrt{12}$ rms quantizing noise in an $f_S/2$ bandwidth, where q equals the weight of the least significant bit (LSB). This rms noise level is independent of the input sine wave's level and frequency so long as the level lies within the A/D converter's operating range.

Theoretically, the ratio of a full-scale sine wave's rms level to an ideal N-bit A/D converter's rms quantizing noise (measured over an $f_S/2$ bandwidth) equals

$$\text{SNR} = 6.02N + 1.8 \text{ dB}.$$

Practically, though, an A/D converter's noise floor increases with the full-scale sine wave's increasing input frequency; the corresponding SNR thus decreases. Conversely, holding the input frequency constant but reducing the sine-wave amplitude decreases the A/D converter's noise floor.

Fig 2 shows an analog test setup for measuring an A/D converter's SNR. During operation, a spectrally pure full-scale sine wave feeds the A/D converter, which is followed by several signal-processing stages. The output of these stages in turn drives a true-rms voltmeter, first bypassing the bandstop filter for

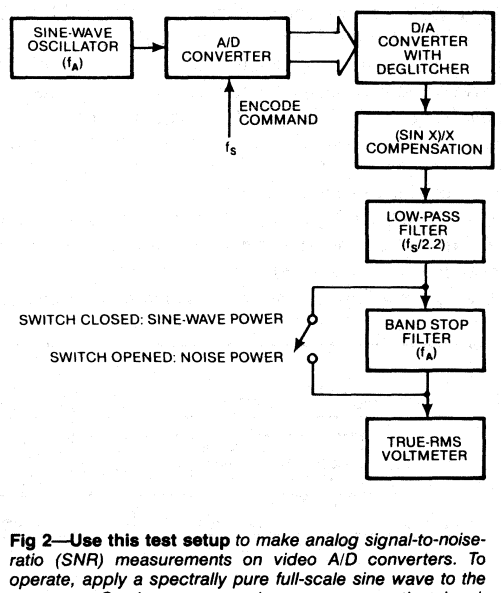

Fig 2—Use this test setup *to make analog signal-to-noise-ratio (SNR) measurements on video A/D converters. To operate, apply a spectrally pure full-scale sine wave to the converter. On the true-rms voltmeter, measure the signal-plus-noise level with the bandstop filter out of the circuit. Then, switch in the bandstop filter to remove the fundamental signal and to measure the rms noise level alone. The two readings, expressed in decibels, make up the signal-plus-noise-to-noise ratio.*

measurement of the signal's level plus noise. Switching in the bandstop filter then removes the fundamental signal to allow the rms voltmeter to measure the noise level alone. The ratio of these two readings expressed in decibels equals the ratio of the signal plus noise to the noise, which approximates the SNR for a signal at least 10 dB greater than the noise level.

Minimizing DAC effects

Fig 3 depicts a useful test method for reducing DAC effects in making back-to-back A/D-to-D/A measurements. The technique employs an analog input sine wave slightly lower in frequency than one-half the sampling frequency, and the registers driving the DAC get updated at an even submultiple of the sampling rate ($f_S/2N$). In turn, the DAC emits a sine wave having a frequency equal to the difference between one-half the sampling rate and the analog input frequency.

In test operation, the A/D converter receives a signal whose frequency approaches the device's Nyquist limit. The DAC, however, updates at a much lower rate ($f_S/2N$), thereby reducing the effects of glitches and other dynamic errors. You can therefore make signal-to-noise measurements over an $f_S/4N$ bandwidth. Use an oscilloscope to inspect the low-frequency beat for missing codes and other nonlinearities; a standard

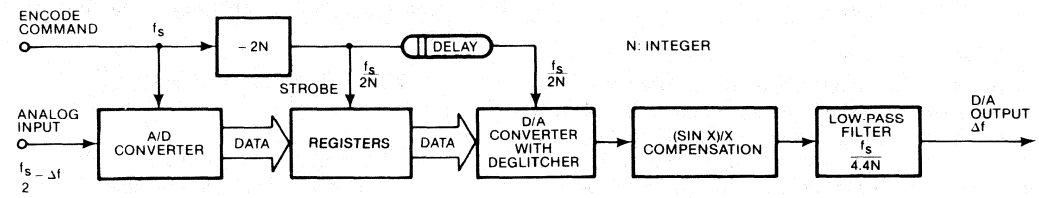

Fig 3—To reduce the DAC's effects *on A/D-converter back-to-back measurements, consider this beat-frequency test setup. In this approach, the A/D converter, operating at its maximum sampling rate, receives a near-Nyquist-frequency signal ($f_S/2$). The DAC updates at a slower rate ($f_S/2N$), reducing glitches and dynamic errors. You can thus make signal-to-noise measurements over an $f_S/4N$ bandwidth. Examine the low-frequency beat on an oscilloscope for missing codes and nonlinearities and on a spectrum analyzer for harmonic content.*

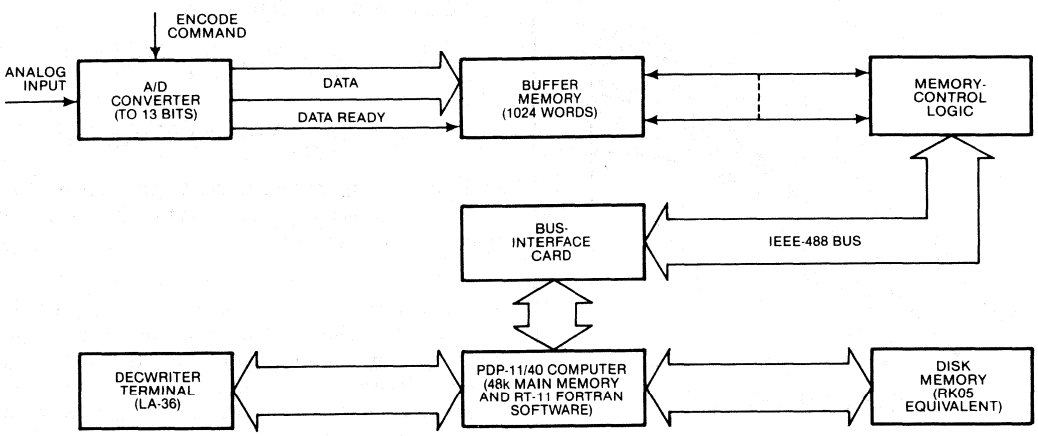

Fig 4—Digitally measuring a video A/D converter's signal-to-noise ratio (SNR) *calls for a computerized test setup that eliminates the DAC and its potential errors. During operation, a spectrally pure sine wave feeds the A/D converter, and contiguous samples store in the high-speed buffer. The computer analyzes the stored data using a discrete Fourier transform to determine the A/D converter's true SNR and harmonic content.*

low-frequency spectrum analyzer lets you measure harmonic content.

Observe that the low-frequency-beat harmonics relate directly to the A/D converter's analog-input-frequency harmonics. In practice, a beat frequency of a few hundred kilohertz performs satisfactorily. Derive both the analog input sine wave and the sampling frequency from frequency synthesizers or crystals to prevent low-frequency-beat-signal smearing.

The beat-frequency test also proves effective in measuring the A/D converter's performance for input signals near the sampling frequency. For a typical test, make the A/D unit's input frequency slightly less than the sampling rate (f_S). Check that a low-frequency beat occurs even when the DAC updates at the sampling rate. Note, however, that updating the DAC at $f_S/2N$ reduces dynamic errors.

Measure SNR digitally

Measuring SNR using the methods previously described proves valid only when the DAC's static and dynamic performance exceeds that of the A/D converter under test by at least four times. To eliminate the DAC's errors, therefore, and to measure the A/D converter's SNR digitally, consider **Fig 4**'s test setup.

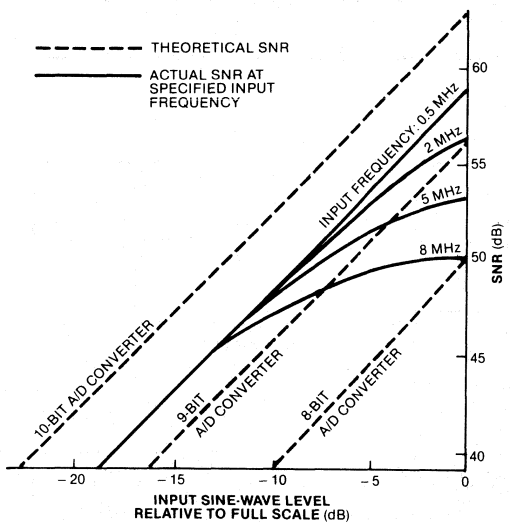

Fig 5—Plotted curves *help analyze a video A/D converter's theoretical and actual SNR test results. They track SNR values vs input sine-wave levels for 8-, 9- and 10-bit converters. With this information, you can determine the number of effective resolution bits for various input signal amplitudes and frequencies.*

To initiate testing, apply a spectrally pure sine wave to the A/D converter. Store contiguous samples in high-speed buffer memory for computer analysis using a discrete Fourier transform (DFT) or a fast Fourier transform (FFT). Note that although the FFT algorithm reduces processing time, the DFT mathematics involves less complexity and thus becomes the preferred method.

To use a DFT, first determine an appropriate time-weighting function for the A/D converter samples in order to reduce frequency side lobes. Exercise care in this determination, though: Selecting an improper weighting function spills main-lobe energy into other frequency bands and makes noise measurements impossible. Accordingly, a $\cos^2$ function without a pedestal (often termed Hanning weighting) is appropriate. The input data multiplied by the weighting function thus becomes

$$WD_n = D_n \left[0.5 - 0.5 \cos \left(\frac{2\pi n}{N} \right) \right],$$

where:

WD_n=nth weighted data sample
D_n=nth input data sample
N=total number of samples.

This weighting function compresses the spillover energy into a few frequency bands centered on the fundamental sine-wave frequency. The result? It eliminates contamination over most of the spectrum.

Next, program a computer to establish the DFT of the sequence of weighted data samples for N/2 frequencies. The program must solve the following two equations for the Kth frequency:

$$A_K = \frac{1}{N} \sum_{n=1}^{N} WD_n \cos \left[\frac{2\pi k\,(n-1)}{N} \right]$$

$$B_K = \frac{1}{N} \sum_{n=1}^{N} WD_n \sin \left[\frac{2\pi k\,(n-1)}{N} \right],$$

where A_K and B_K represent the magnitudes of the Kth spectral line's cosine and sine parts. The line's magnitude therefore equals

$$MAG_K = \sqrt{A_K^2 + B_K^2}.$$

The computer uses this equation to determine SNR. It then computes the residual noise floor in all frequency bands except those occupied by the input signal. During execution, the program squares the magnitude of each frequency band that doesn't contain the input signal, adds all N/2 squared magnitudes and evaluates the following square root:

$$NOISE = \sqrt{\sum_{K=1}^{N/2} MAG_K^2}.$$

Set several frequency bands on either side of the

fundamental sine-wave frequency band to zero to account for main-lobe widening caused by the $\cos^2$ weighting function.

For the final SNR test, measure harmonic distortion using the following expression:

$$\text{HARMONIC DISTORTION} = 20 \log \left[\frac{MAG_f}{MAG_{mf}} \right],$$

where:

MAG_f=magnitude of fundamental

MAG_{mf}=magnitude of mth harmonic.

In making digital SNR measurements, don't lock the fundamental input sine-wave frequency to the sampling frequency or its harmonics. If you do, the sample-to-sample variations in the sampled input signal and the A/D converter's encoded output voltage become predictable and repeatable. Obeying this rule achieves measurement results that indicate true SNR and harmonic content under normal A/D-converter operating conditions, where sample-to-sample variations are unpredictable and nonrepeatable.

Analyzing SNR results

You can display and analyze A/D-converter SNR test results in several ways. For example, you can plot SNR versus sine-wave frequency for a constant-amplitude full-scale sine-wave input. You can also plot SNR versus input-signal amplitude for a fixed sine-wave frequency. What's more, you can combine the two plots (**Fig 5**). Using these curves, you can determine the number of effective bits of resolution for various amplitudes and frequencies and compare them with the theoretical values.

Note that you can measure the SNR for input signals greater than the Nyquist limit ($f_S/2$). Be aware, however, that the fundamental sine wave then appears as an alias component within the $f_S/2$ bandwidth. For example, a 12-MHz sine wave's in-band component sampled at 20 MHz occurs at 8 MHz.

NPR tests demand attention

Another important indicator of an A/D converter's dynamic performance centers on noise-power-ratio (NPR) testing. This ratio testing finds extensive use in measuring the transmission characteristics of frequency-division-multiplexed (FDM) communications links. In a typical FDM system, 4-kHz-wide voice channels get stacked in frequency for transmission over coaxial, microwave or satellite equipment. At the receiving end, the FDM data becomes demultiplexed and returned to 4-kHz individual baseband channels. In an FDM system having more than about 100 channels, you can approximate the FDM signal by Gaussian noise with the appropriate bandwidth. **Fig 6** illustrates how to measure a 4-kHz channel for quietness using a narrow-band notch (bandstop) filter and a special tuned receiver. The receiver measures the noise power within the 4-kHz notch.

To implement testing, you first switch out the notch filter. Then using the receiver, measure the rms noise power of the signal within the notch. Next, switch in the notch filter and evaluate the residual noise within the slot. NPR thus equals the ratio of the two readings expressed in decibels. Check several slot frequencies across the noise bandwidth (eg, low, midband and high) to adequately characterize the FDM system.

For analysis purposes, plot NPR as a function of rms noise level referred to the FDM system's peak range. At low noise-loading levels, undesired noise exists primarily as thermal activity, independent of the input noise level. Over this region of the NPR curve, a 1-dB increase in noise level causes a 1-dB increase in NPR. At increased noise-loading levels, the FDM system amplifiers overload, creating intermodulation products that increase the noise floor. With further input-noise

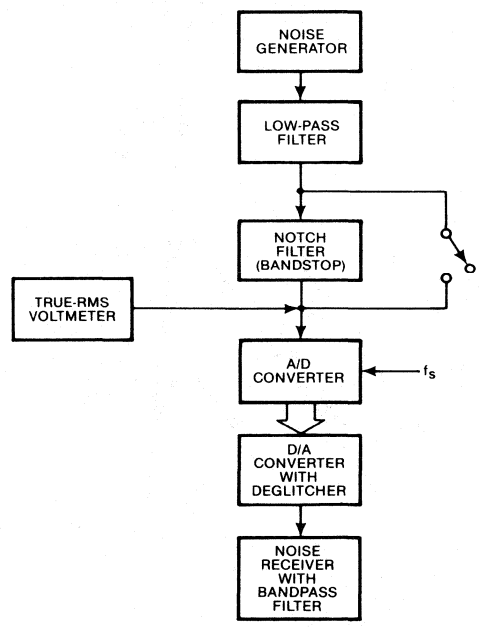

Fig 6—In this noise-power-ratio (NPR) test setup, *the special tuned receiver measures the input signal's rms noise power inside a narrow-band notch with the bandstop filter out of the circuit. Next, it measures the noise power with the bandstop filter included. NPR equals the ratio of the two readings expressed in decibels.*

increases, overload noise effects predominate and dramatically reduce NPR. In practice, FDM systems usually operate at a noise-loading level of a few decibels below the point of maximum NPR.

A digital system containing an A/D converter primarily generates quantizing noise within the notch filter's slot on receiving low values of noise input signals. The NPR curve tracks linearly in this region. As the noise input level increases, though, clipping noise caused by the A/D converter's hard-limiting action starts to prevail. **Fig 7** presents a set of theoretical NPR curves for various A/D converters.

In a practical A/D converter, however, dc or ac nonlinearities cause departures from the theoretical NPR values. Although the NPR's peak value occurs at a low input-noise level (rms noise=¼V_0, where ±V_0=A/D converter's input range), the broadband nature of the noise signal stresses the A/D converter. Consequently, NPR tests provide worthwhile dynamic-performance measurements.

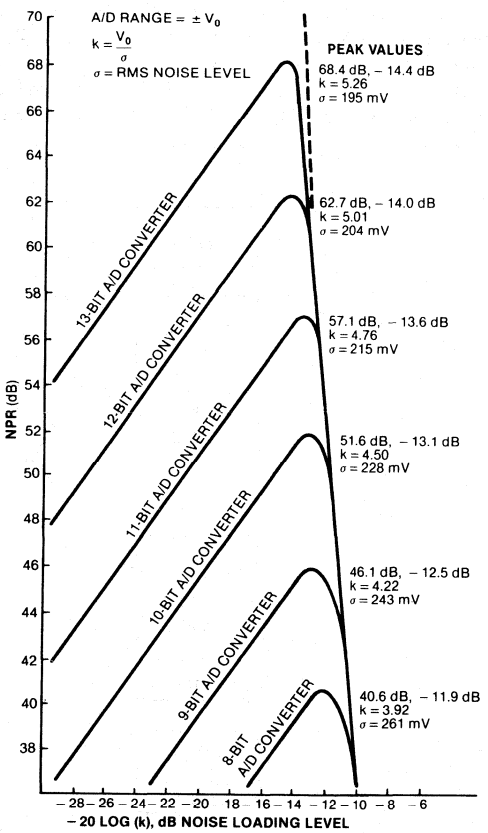

Fig 7—In A/D-converter-based digital systems, *theoretical NPR curves are linear at low noise-input-signal levels. As the noise input levels increase, however, the A/D converters produce clipping noise because of hard-limiting performance. Note that the NPR values peak at corresponding higher levels and at lower rms noise levels as converter bit resolutions increase.*

Histograms disclose nonlinearity effects

Another key parameter in video A/D-converter dynamic testing is differential nonlinearity, which often degrades as the input-signal frequency increases. You can employ histograms to help analyze this problem: To obtain a histogram for an A/D converter, make the sine-wave input frequency noncoherent with the sampling frequency. Then, take many A/D-converter samples and use a computer to calculate the number of times each output code occurs and to plot these results directly. Analyze the results for missing codes and differential nonlinearity.

Because the probability of obtaining digital codes around the zero-crossing point proves less than the probability of obtaining codes around the positive or negative peaks, the resulting curve traces a cusp shape rather than a flat one. Normalize this curve by using the sine wave's probability density function. Or apply a triangular waveform, rather than a sine wave, to the A/D converter.

Obtaining a histogram for a 10-bit A/D converter requires a large buffer memory. To get 100 samples per code, for example, buffer memory must contain 100×1024 or 102,400 locations.

The histogram approach furnishes accurate representations of an A/D converter's statistical differential nonlinearity. However, a video A/D converter's actual differential-nonlinearity characteristic might depend on the input signal's direction of change and rate of change. Thus, positive-slewing signals might produce different differential nonlinearity characteristics than negative-slewing signals. This subtle effect clouds the histogram's results for A/D converters used in transient- or pulse-analysis applications, where the encoded signals are not statistical in nature.

Aperture-time measurements differ

Heading the list of misunderstood and misused video-A/D-converter specifications is aperture time. Originally, aperture-time measurements centered on the classic sample/hold circuit (EDN, April 14, pg 41). Ideally, a sample/hold's switch possesses zero resistance when closed and opens instantly. In practice, though, the sampling switch transits from a low to high resistance over some finite time interval. Accordingly, the original definition holds—an error occurs because the input signal gets averaged over the finite time interval required for opening the switch. The sampled

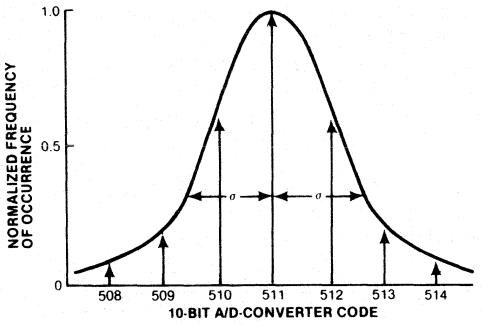

Fig 9—An ideal A/D converter *without aperture jitter shows only one code on a histogram. A practical converter exhibits a distribution of codes. The distribution's calculated sigma (σ) value corresponds to the rms error voltage produced by the rms aperture jitter.*

converters that operate in Continuous mode, in which the encode command is always present. Some applications, however, require that an A/D converter operate in Burst mode, in which the ADC operates on a finite number of encode-command pulses. This mode commonly occurs in radar return-pulse analysis because the A/D converter functions in an idle condition except for the brief time that encompasses each return pulse.

Burst-mode operation usually places more stringent requirements on the A/D converter than continuous operation. In a continuous mode, timing and track/hold circuits reach steady-state conditions. In Burst mode, however, samples taken before reaching steady-state conditions produce errors. A track/hold that uses a transformer to couple the switching pulses to the diode bridge proves particularly error susceptible.

Fig 10 outlines a test setup for evaluating Burst-mode converter operation. Use the variable delay to slide the burst of encode commands across the flat pulse. Analyze the results of each sample pulse with a computer or appropriate gating circuits on the A/D converter's output. Each sample's results should be identical except for the time delay that separates each individual pulse in the burst. Differences in results indicate the presence of analog memory, which might limit the A/D converter's use in certain applications.

A preferred method for handling Burst-mode applications calls for operating the A/D converter continuously and gating the digital outputs into a memory during the desired command intervals. This method relieves unnecessary stress on the A/D converter and accomplishes the same result as actual Burst operation.

Capturing bandwidth clipping

You can use **Fig 11**'s test setup to measure a video A/D converter's frequency response and bandwidth (small or large signal). During checkout, hold the input-signal amplitude constant while varying the input frequency over the desired range. Analyze the A/D converter's output digitally or with a DAC followed by an rms voltmeter. Using a voltmeter, however, limits the filtered DAC's output frequency to $f_S/2.2$.

A quick method for determining an A/D converter's large-signal bandwidth involves adjusting the input

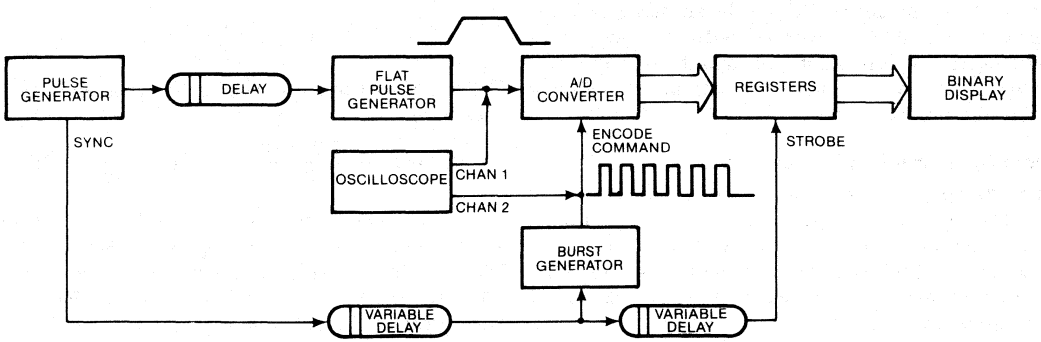

Fig 10—To check out an A/D converter's Burst mode, *in which it operates as directed by a finite number of encode-command pulses, you adjust the variable delays in this test setup to slide the burst of encode commands across the flat input pulse to the A/D converter. Analyze the sample pulses with a computer or appropriate gating circuits connected to the converter's output. Each sample should produce identical test results except for the time delay that separates each pulse in the burst.*

voltage, therefore, does not exactly correspond to the voltage at the instant the switch starts to open, and the time required to open the switch represents the aperture time. The aperture error, according to the original definition, therefore equals

$$E_a = t_a dV/dt, \quad (1)$$

where:

E_a= aperture error

t_a=aperture time

dV/dt=rate at which the input signal changes.

A simple first-order analysis, which neglects nonlinear effects, reveals that no real *aperture* error exists for such a switch. Instead, so long as the switch opens repeatably, an effective fixed *sampling-time delay* exists—the time between encode-command arrival and the opening of the ideal zero-opening-time switch. (Note that you can include logic propagation delays in the encode-command path to modify this time interval.) Unfortunately, most manufacturers still refer to this delay as aperture time.

True aperture errors, on the other hand, result from variable time delays. These errors generally emanate from several sources: In a practical A/D converter, the encode-command signal often gets phase-modulated by an unwanted source—random noise, power-line frequency or digital noise—because of faulty grounding techniques. You can express the resulting error in terms of an rms time jitter, termed aperture jitter. The corresponding rms voltage error caused by aperture jitter qualifies as a valid aperture error.

An A/D converter's aperture-jitter spec sometimes gets interpreted as a measure of the device's ability to encode rapidly changing input signals accurately. However, although important, the aperture-jitter specification does not define the whole encoding process. Therefore, an A/D converter with an impressive aperture-jitter specification still might not accurately digitize a sine wave having a maximum slew-rate calculated from **Eq 1**.

Typical aperture-time test examples

For example, assume that a 20-MHz, 10-bit A/D converter has a $\pm V_0$ bipolar input range ($2V_0$ p-p) and a 10-psec rms aperture-jitter specification. To calculate the maximum aperture error, convert the rms aperture jitter into a maximum value rather than an rms one. If you assume that aperture jitter follows a Gaussian distribution similar to white noise, the rms aperture jitter (t_a) corresponds to the distribution's σ value. The distribution's 2σ point thus becomes the proper place to set the maximum value; the maximum aperture jitter therefore equals $2t_a$.

If you set the corresponding maximum voltage error (ΔV) at a full-scale sine wave's zero crossing to ½ LSB (where ½ LSB=$2V_0/2^{N+1}$ and N=number of resolution bits), you can then calculate the maximum full-scale sine-wave frequency (f_{MAX}) that produces the ½-LSB aperture error:

$$v(t) = V_0 \sin 2\pi ft \quad (2)$$

$$\frac{dv}{dt} = 2\pi f V_0 \cos 2\pi ft \quad (3)$$

$$\left.\frac{dv}{dt}\right|_{MAX} = \frac{\Delta v}{2t_a} = 2\pi V_0 f_{MAX},$$

thus yielding

$$f_{MAX} = \frac{\Delta v}{4\pi V_0 t_a} = \frac{2\pi V_0/2^{N+1}}{4\pi V_0 t_a} = \frac{1}{2\pi t_a 2^{N+1}}.$$

If t_a=10 psec rms and N=10, f_{MAX}=7.8 MHz. These calculations imply that the 20-MHz A/D converter under test can accurately digitize a 7.8-MHz FS sine wave. In practice, however, the A/D converter might suffer from skipped codes, decreased SNR and ac nonlinearities at much lower frequencies. These other limitations to good high-frequency performance, might, in fact, eclipse errors caused by the sample/hold's aperture jitter.

Similarly, you can calculate aperture-jitter effects on full-scale sine-wave SNR using **Eqs 2** and **3** and the following expression:

$$\left.\frac{dv}{dt}\right|_{RMS} = \frac{2\pi f V_0}{\sqrt{2}}. \quad (4)$$

Substituting rms error voltage (ΔV_{RMS}) and rms aperture jitter (t_a) into **Eq 4** yields

$$\frac{\Delta V_{RMS}}{t_a} = \frac{2\pi f V_0}{\sqrt{2}}$$

$$\Delta V_{RMS} = \frac{2\pi f V_0 t_a}{\sqrt{2}}.$$

The rms-signal-to-rms-noise ratio expressed in decibels thus becomes

$$SNR = 20 \log \left[\frac{V_0/\sqrt{2}}{\Delta V_{RMS}}\right] = 20 \log \left[\frac{1}{2\pi f t_a}\right] \text{ dB}. \quad (5)$$

Consider again the 10-bit, 20-MHz A/D converter with a 10-psec rms aperture jitter. For an 8-MHz full-scale input, the SNR arising only from aperture jitter totals 66 dB, as calculated from **Eq 5**. The theoretical SNR due to quantizing noise in a 10-bit A/D converter equals 62 dB. Combining the two SNR decibel values results in a theoretical 60.5-dB SNR; this final value encompasses the ideal quantizing noise and the aperture-jitter noise. From **Fig 5**'s curves, observe that a practical A/D converter having the same aperture-jitter spec achieves only a 50-dB SNR.

From this analysis, you can logically conclude that

the SNR specification more closely defines video-A/D-converter performance than does the aperture-jitter specification. Relying solely on the aperture-jitter specification therefore creates an erroneous deduction about the converter's actual dynamic performance.

Measuring aperture jitter

Fig 8 depicts the test setup for measuring an A/D converter's aperture jitter. The encode-command signal and the analog input signal derive from the same low-jitter pulse generator to minimize the phase jitter between them.

To test, adjust the phase shifter until the A/D converter repetitively samples the sine wave at its midscale point of maximum slew rate. Set the bit switches to the A/D converter's midscale code: 0111...1. Adjust the phase shifter again for a maximum reading on the frequency counter. Examine the midscale code's frequency of occurrence as well as the frequency of occurrence of several codes above and below this value. Using this information, plot a histogram (**Fig 9**).

An ideal A/D converter with no aperture jitter would present only one code on the histogram. In contrast, a practical converter exhibits a distribution of codes. The distribution's calculated σ corresponds to the rms error voltage (ΔV_{RMS}) produced by the rms aperture jitter (t_a). Calculate the aperture jitter (t_a) as follows:

$$t_a = \frac{\Delta V_{RMS}}{dV/dt},$$

where dV/dt equals the sine wave's rate of change at a zero crossing.

If you attenuate the input sine wave sufficiently, the distribution spread around the nominal code results from intrinsic A/D-converter noise. As the input sine wave increases in amplitude, the slew rate (dv/dt) becomes proportionally greater, and the distribution begins to spread because of aperture jitter. Take care when interpreting a histogram for high-slew-rate inputs because high slew rates also affect the converter's ac differential linearity.

The offset trimming pot lets you position the sine wave at different A/D-converter range points. In this manner, you can plot histograms around several nominal codes and attribute variations to range-dependent differential-linearity characteristics. Don't exceed the A/D converter's range during this offset adjustment, however.

The tests presented so far aim at video A/D

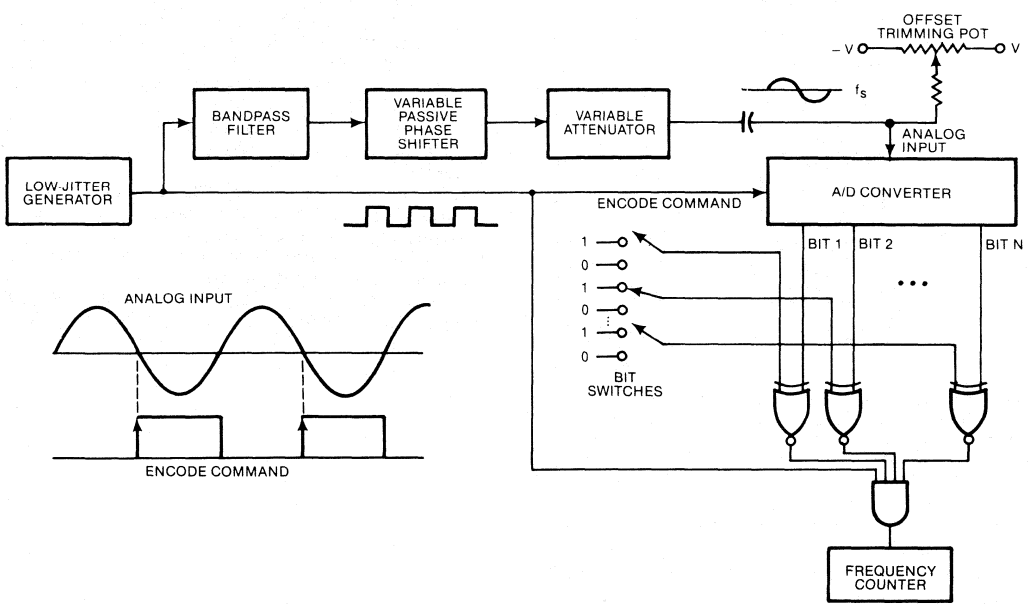

Fig 8—This test setup *measures an A/D converter's aperture jitter. To operate, you adjust the phase shifter until the A/D converter repetitively samples the analog input sine wave at its midscale point of maximum slew rate. Next, you set the bit switches to the converter's midscale code. Then, you readjust the phase shifter for a maximum reading on the frequency counter. The frequency of occurrence of the midscale and adjacent codes indicate aperture jitter when plotted on a histogram.*

sine wave until the DAC indicates A/D-converter output clipping; observe this condition on an oscilloscope connected across the DAC's output, before deglitching and filtering stages. Increase the input frequency and readjust the input level to cause clipping. An increase in input-signal level required to cause clipping corresponds to a decrease in A/D converter gain; a decrease in input level implies an increase in gain. This test method also yields frequency

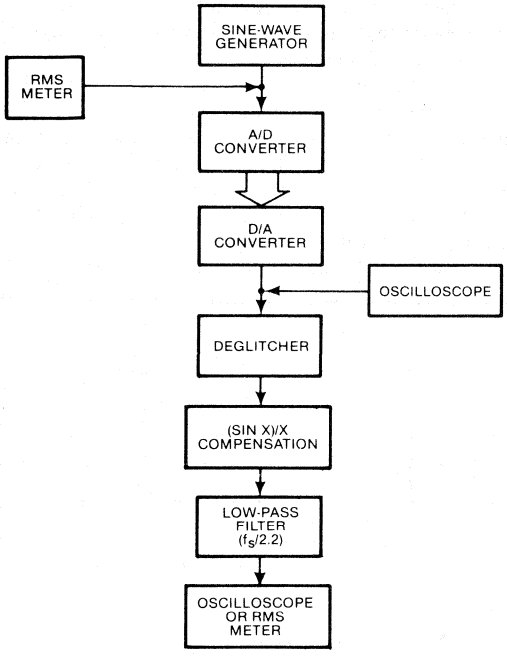

Fig 11—This test setup *measures an A/D converter's frequency response and small- or large-signal bandwidth. During testing, you hold the input signal amplitude constant while varying the input frequency over the desired range. You can analyze the converter's output digitally or with a DAC followed by an rms voltmeter.*

measurements above the Nyquist limit.

Note that a good bandwidth specification does not always ensure that the A/D converter's dynamic performance is satisfactory at higher input frequencies. Thus, an A/D converter having a wide bandwidth might still exhibit inadequate SNR, NPR or harmonic distortion at higher frequencies; carefully check it out.

References

1. Bennett, W R, "Spectra of Quantized Signals," *Bell System Technical Journal 27,* July 1948, pgs 446-472.
2. Blackman, R B, and Tukey, J W, *The Measurement of Power Spectra,* Dover Publications Inc, New York, NY, 1958.
3. Gray, G A, and Zeoli, G W, "Quantization and Saturation Noise Due to Analog-to-Digital Conversion," *IEEE Transactions on Aerospace and Electronic Systems,* January 1971, pgs 222 and 223.
4. Kester, W A, "PCM Signal Codes for Video Applications," *Society of Motion Picture and Television Engineering Journal 88,* November 1979, pgs 770-778.
5. Neil, Martin, and Muto, Art, "Tests Unearth A/D Converter's Real-World Performance," *Electronics,* February 24, 1982, pgs 127-132.
6. Pratt, W J, "High Linearity and Video Speed Come Together in A/D Converters," *Electronics,* October 9, 1980.
7. Ramirez, Robert W, "Digitizer Specifications and Their Applications to Waveforms," *Electronics Test,* September 1981, pgs 74-80.
8. Tant, M J, *The White Noise Book,* Marconi Instruments, Northvale, NJ, July 1974.
9. Wilhelm, Tim, "Test A/D Converters Quickly and Efficiently," *Electronic Design,* October 15, 1981, pgs 193-198.

NE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

AD671 12-Bit, 2 MHz ADC Digitizes CCD Outputs for Imaging Applications

by Elwyn Davies

NTRODUCTION

harge coupled devices (CCDs) are widely available in wo basic configurations – linear arrays and area arrays. CD area arrays are commonly used in video applica- ions. Linear arrays are used in a broad range of appli- ations including facsimile machines, graphics scanners nd pattern recognition equipment.

his application note explains how the AD671 12-bit esolution, 2 MHz analog-to-digital converter (ADC) can e used to digitize the output from CCD linear arrays for se in some popular imaging applications, in particular raphics image scanners, a simplified functional dia- ram of which is shown in Figure 1.

CD OVERVIEW

A brief review of CCD operation and terminology is use- ul at this stage in order to appreciate the following pplication examples.

The CCD contains a large number of photosensitive sub- elements called photosites or pixels which are arranged either in a single row in the case of linear arrays or in a matrix pattern as for area arrays.

The resolution (spatial) of a CCD is expressed as the number of pixels or photosites which it contains.

The dynamic range of a CCD is the ratio of the saturation exposure level divided by the rms noise equivalent exposure level. Alternatively, dynamic range is expressed as the ratio of the saturation exposure level divided by the peak-to-peak noise equivalent exposure level. For the purpose of CCDs, peak-to-peak noise is assumed to be four to six times greater than rms noise. Thus the peak-to-peak noise definition of dynamic range yields a four to six times lower figure than the rms noise definition.

In contrast, the resolution of an analog-to-digital converter (ADC) is expressed as the number of bits of its digital output. This figure is the number of discrete ADC input levels which the converter output can represent. The dynamic range of the ADC is the ratio of the full-scale output code (2^n for an n-bit ADC) to the smallest of these output codes (1). Resolution and dynamic range of an ADC are commonly but incorrectly assumed to mean the same thing. This mistake can lead to confusion with regard to CCD specifications.

The basic topology of a linear array CCD consists of a row of image sensor elements otherwise known as photosites or pixels which are illuminated by light from the object or document.

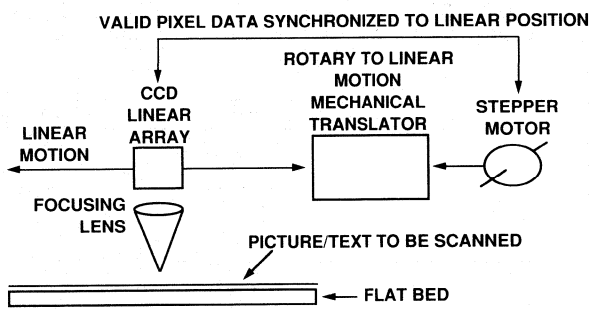

Figure 1. Graphics Image Scanner Functional Diagram

During one exposure period each photosite is loaded with an amount of charge which is proportional to the illumination level present on it. In the case of the CCD191 (a 6000 photosite linear array CCD manufactured by Fairchild Weston) the exposure period is the time between the trailing edges of the ϕx pulse train. (See Figure 4.)

Each of these photosite charge packets is subsequently switched simultaneously via a transfer gate to an analog transport shift register. The charge packets on this shift register are clocked serially to a charge detector and buffer amplifier which convert them into a string of photo dependent output voltage levels.

An extension of this basic scheme consists of two trans fer gate/analog shift register combinations each c which transfers the charge packets from either the od or even photosites, respectively. This system allow faster output of data from the CCD.

While the charge packets from the preceding exposu are being clocked out serially to the charge detecto another exposure period is underway.

CHOICE OF ADC

When choosing an appropriate ADC to use with a parti ular CCD, the resolution of the ADC is usually chosen t be as close to (but usually lower than) the dynam

THE AD671 SUBRANGING ADC

The AD671 is a monolithic 12-bit ADC which converts at rates up to 2 MHz. Figure 2 shows the internal functional diagram of the AD671. It can be configured to accept either unipolar 0 to +5 V or 0 to +10 V or bipolar −5 V to +5 V input ranges.

The AD671 uses a successive subranging conversion technique with digital error correction for errors introduced in the first part of the conversion cycle. An on-chip timing generator provides strobe pulses for each of the four internal flash cycles which combine to form the overall conversion algorithm. A single active high encode pulse applied by the user to the device initiates the entire conversion.

Upon receipt of an encode command, the first 3-bit flash ADC converts the analog input voltage. The 3-bit result is passed to a correction logic register and a 3-bit resolution, 13-bit accurate DAC. The DAC output is subtracted from the analog input creating a residue voltage which is less than one eighth of the full-scale analog input.

This residue is applied to the input of a second 3-bit flash converter and to a high speed differential gain of four amplifier. The result from this second flash conversion is passed to the correction logic. It is also reconverted back to analog form by a second 3-bit resolution, 13-bit accurate DAC. The output of this DAC is subtracted from the voltage applied to the second flash converter by the differential amplifier and multiplied by a factor of four.

The differential output is connected to a two stage 8-bit flash converter which consists of coarse and fine 4-bit flash converters. The coarse 4-bit flash converter is also configured to overlap one bit of the second 3-bit DAC. The result of this coarse 4-bit converter is connected to the correction logic register and also selects one of 16 resistors from which the fine 4-bit flash will establish its span voltage. The result of the fine 4-bit flash converter is directly applied to the output latches.

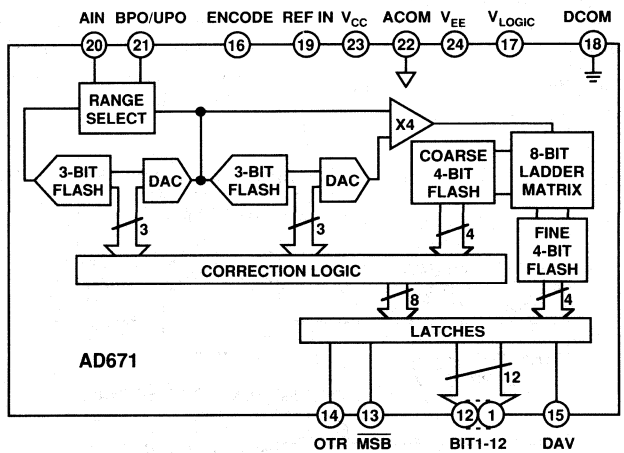

Figure 2. AD671 Internal Functional Diagram

ıge of the CCD defined by the peak-to-peak noise ⁄el.

e choice of conversion speed of the ADC is deter-ned by the maximum clock frequency at which the :D is required to be operated.

many cases, however, the minimum conversion time the particular ADC, which system cost budgets can ford, is a major factor in limiting the actual rate at ıich the CCD is clocked.

RAPHICS IMAGE SCANNERS

gure 1 shows the functional diagram of a graphics ıage scanner. This diagram shows the popular system hich uses a stationary flat bed, on which is placed the cture and/or text to be scanned, and a moving CCD ıear array. Less common alternative systems use a ationery CCD and a movable bed.

ıe axis of the CCD linear array of photosites is orthog-ıal to the axis of motion of the bed and parallel to the ane of the bed.

ght from the scanned picture/text is focused on the CD via a focusing lens. The position of the CCD array ensor assembly relative to the bed is accurately set by ıeans of a stepper motor and worm gear or other pre-sion rotary to linear mechanical transformer. More ac-ırate systems use dc linear motors and position servo ontrol loops instead of stepper motors.

he CCD array is typically mounted on a printed circuit oard which also contains components for generating ıe digital control signals for the CCD and the ADC for igitizing the CCD output signals. This PCB would con-titute the movable sensor assembly. It is essential that areful attention is paid, on this movable PCB, to the ecoupling of power supplies, screening of stray noise om the motor and reduction of digital feedthrough to void corrupting the analog signal from the CCD prior to s digitization by the ADC.

he CCD's digitized output is subsequently routed, typi-ally via flexible ribbon cable, to a stationary digital rocessing board where various digital algorithms can ıe performed to compensate for those photosites re-ponse nonuniformities (PRNU) within the CCD array. or example, if the output of a particular photosite is out ıf range, then the average value of the previous current ınd following pixels' outputs is used instead. This can enhance the quality of the final image. It also allows lower grades of CCD, which typically exhibit greater numbers of PRNUs, to be used, thus lowering system costs.

The control of the exposure timing of the CCD must be synchronized to the movement of the stepping motor. This can be done, for example, by using the high going edge of the ϕx pulse of the CCD191 (see Figure 4) to clock the stepping motor to its next detent position. In this case one must ensure that the length of the ϕx pulse is sufficient for mechanical settling to the new position to have occurred prior to the next exposure period.

Alternatively the CCD timing can be slaved to the stepper motor master control signals. In this case the stepper motor pulses would gate the start of the linear array exposure period.

The monochrome system described above can be extended to full color capability by using three such systems in parallel. Red, green and blue filters, respectively, are applied in front of the individual CCDs. The three color components, which are digitized simultaneously, can be subsequently used to reconstruct the original color image.

Alternatively, color can be provided by a single channel system by applying red, green and blue filters sequentially in front of the same CCD. Each filter is placed in front of the CCD for one complete monochrome scan period. This system costs less to implement but is three times slower than the aforementioned scheme.

DIGITIZING THE PHOTO SIGNALS

The two analog outputs from a two output channel CCD linear array can be digitized either simultaneously, using one ADC per output, or sequentially using a single ADC preceded by a suitable multiplexer to switch between the two channels.

In practice, the multiplexed output method is the most popular because it uses one ADC as opposed to two for the simultaneous conversion technique. It also obviates the need for subsequent digital multiplexing hardware, or software overhead for odd/even pixel data interleaving.

Figures 3a, 3b, and 3c together show the circuit diagram of a system for digitizing the multiplexed outputs of the CCD191 linear CCD array.

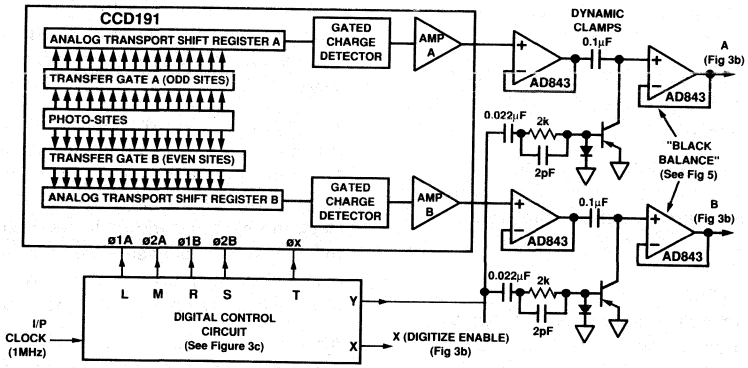

Figure 3a. Linear Array CCD – Dynamic Clamps and Digital Control Circuit

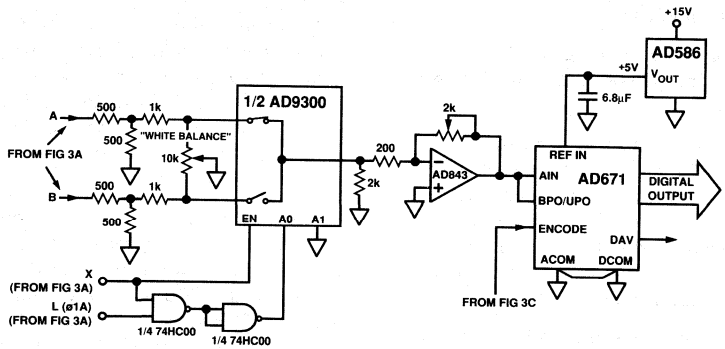

Figure 3b. Linear Array CCD – Multiplexed Output Digitizing Circuit

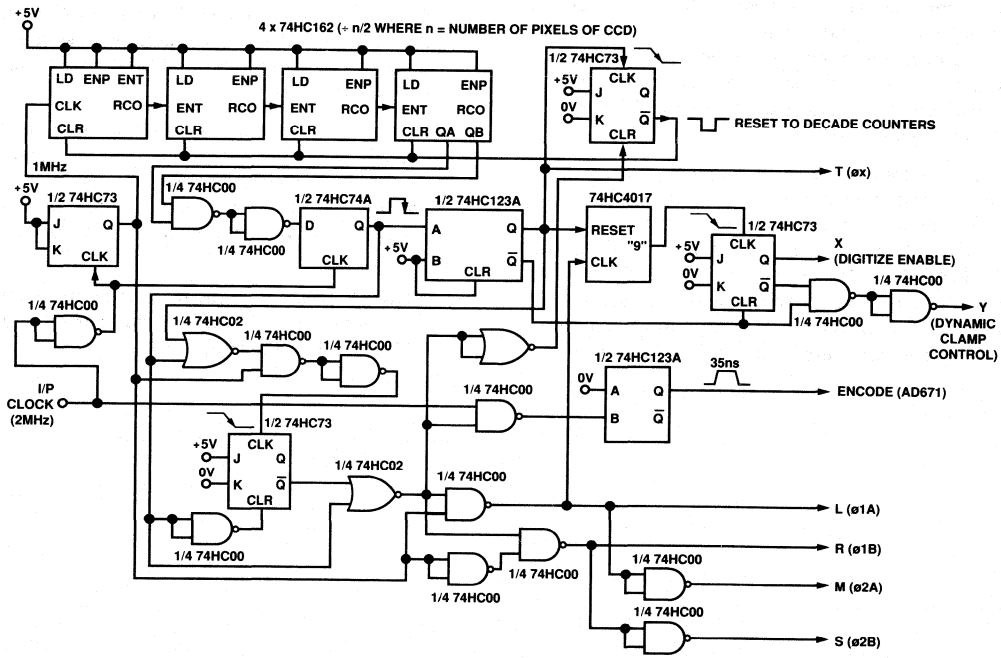

Figure 3c. Linear Array CCD – Digital Control Circuit

sic operation of the CCD is achieved by applying a mber of synchronized clocks and control pulses to the opropriate inputs of the CCD. The phase relationships the various inputs necessary for multiplexed output eration are shown in Figure 4. Also shown is their ase relationship to the odd pixel (Channel A) and even el (Channel B) analog outputs.

D OUTPUT CONSIDERATIONS

ny of the processing steps which are used for the rication of CCDs are similar to those used for NMOS hnology. It is not surprising, therefore, that NMOS cuit elements are the most convenient and cost effec- e to be integrated on the same chip as the CCD. MOS is ideal for the low leakage analog transport shift gisters and high input impedance gated charge detec- s which are found on most CCD ICs. The gated feature these charge detectors allows them to operate like mple-and-hold amplifiers (SHAs). Thus it is not neces- y to add external SHAs though it may be useful to do in nonmultiplexed configurations..

wever, NMOS is not suitable for producing wide ndwidth and low output impedance amplifiers. The tput impedance of the output amplifier of the CCD191, instance, is 1 kΩ. Any bias current or leakage due following circuitry which flowed through this high tput impedance could introduce intolerably high er- rs into the system.

s, therefore, essential to connect a suitable high input pedance, wideband and low output impedance buffer each analog output of the CCD.

The AD843 FET input wideband operational amplifier, which offers 34 MHz unity gain bandwidth and 0.6 nA bias current is ideal for this function.

PRESCAN REFERENCE PHASE

The sequential pixel analog outputs from the CCD are applied to the following subsequent digitization circuitry during the readout phase. The first ten odd and first ten even pixels of this readout phase contain black level reference information which effectively allows for temperature calibration of the CCD and black level matching of the odd and even pixel channels.

During this reference phase the dynamic clamp circuits store any standing dc offset which may be present on the CCD output. This is done by switching on the PNP transistors in the dynamic clamp stages shown in Figure 3 thus allowing the 0.1 μF capacitors to charge up to the standing dc output voltage. When the reference phase is complete, the PNPs are switched off thus allowing the 0.1 μF capacitors to couple any residual ac signal through to the following circuitry and also to strip the dc content, which may be several volts, from it.

The output from the dynamic clamp stage is a 0 to −1 V ac signal. The 74HC4017 decade counter shown in Figure 3c, which is configured to count ten clock pulses before being reset, is used to time out the prescan reference phase. Only during this phase are the dynamic clamp PNP switches turned on.

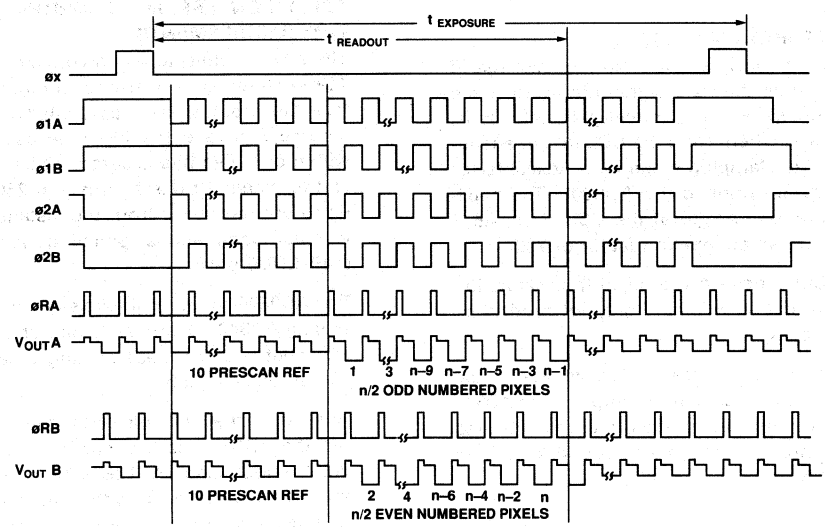

Figure 4. Linear Array CCD (CCD191) – Multiplexed Odd/Even Pixel Output Control Timing

BLACK BALANCE

The black levels from each channel may not be matched due to onboard CCD amplifier offset mismatches and transport shift register leakage mismatches. In order to compensate for these effects a two channel offset correction trim is performed which is known as "black balance." A simple circuit which allows the user to perform "black balance" is shown in Figure 5.

This circuit can be inserted into the signal path at the output of the dynamic clamps as indicated in Figure 3a.

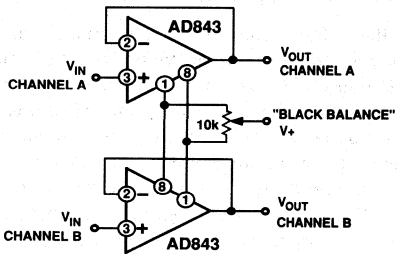

Figure 5. Linear Array CCD – Odd/Even Pixel Channel Black Level Balance Circuit

WHITE BALANCE

Gain mismatch between the CCD's odd and even pixel channel outputs produce mismatches in white level of the final digitized signal. Compensation for this mismatch is known as "white balance" and can be performed by adjusting the ground connected wiper of the 10 k potentiometer as shown in Figure 3b.

CCD OUTPUT SATURATION VOLTAGE

The CCD191 has a typical output saturation voltage of 3 V. This is higher than the ±2 V signal range which the AD9300 video multiplexer can directly handle. It is, therefore, necessary, when using the CCD191, to attenuate the output of the clamping stage by a typical factor of two (see Figure 3b) prior to the AD9300. The maximum output saturation voltage of the CCD191 is 4.5 V which would require an attenuation factor of 2.25.

However, most CCDs have a typical output saturation voltage of only 1 V and would, therefore, not need attenuation of their outputs.

In order to match the output saturation voltage of the particular CCD to the full-scale input range of the AD671 it is necessary to trim the gain of the pre amp (AD843) to the ADC via a potentiometer in its feedback path as shown in Figure 3b.

ODD/EVEN PIXEL MULTIPLEXING

The odd and even pixel output signals can be alterna switched to the input of a singie ADC using the AD9 video multiplexer.

Figure 6 shows the timing diagram of the control sig needed to implement the odd/even pixel multiple and their relationship to the A (odd) and B (even) ch nel CCD video outputs.

The multiplexer must have adequate bandwidth slew rate to accommodate the worst case change signal level from one pixel to its neighbor. This co sponds to a 0 V to full-scale change or vice versa every pixel transition. For an input clock frequency 1 MHz (corresponding to a pixel rate of 2 MHz) this le change would occur every 500 ns. A raw rule of thu guide is that the analog bandwidth of the multiple and amplifiers should be at least five times this cl rate in order to retain adequate gain flatness over 1 MHz pixel bandwidth (1/2 the pixel rate). The 34 M full power bandwidth of the AD9300 easily satisfies t criterion.

The settling time of the AD9300 (to 0.1% of full-sc output) is typically 70 ns (100 ns maximum). Thus fo maximum pixel to pixel signal change of 1 V, it wo seem that at least 70 ns must be allowed for the outp of the multiplexer to have settled before the ADC co version begins.

However, the following section explains why this init 100 ns need not be expended prior to the start of t ADC conversion.

AD671 CONVERSION ALGORITHM EASES SETTLIN TIME REQUIREMENTS

The novel architecture and conversion algorithm of t AD671 allow the converter to commence conversion b fore the input signal has settled to the final desir accuracy. In fact, the input needs to have settled within only five bits accuracy of its final value in 50 and to within 10 bits accuracy in 250 ns. Both of the times are measured from the instant at which the co vert start command is applied to the AD671 (rising ed of ENCODE pulse).

The AD9300 video multiplexer combined with the fa settling AD843 op amp at its output, shown in Figure 3 will meet both of these settling time requirements f 1 V full-scale signals.

This means that valuable settling time overhead is save and the AD671 has the whole of the 50% of the CC clock period (see ϕ1A, ϕ2B, ϕ2A, ϕ1B of Figure 4) avai able for conversion. Note that the new pixel output sig nal must settle to the intermediate accuracy required b the AD671,and the ADC must complete its conversio within half a clock period.

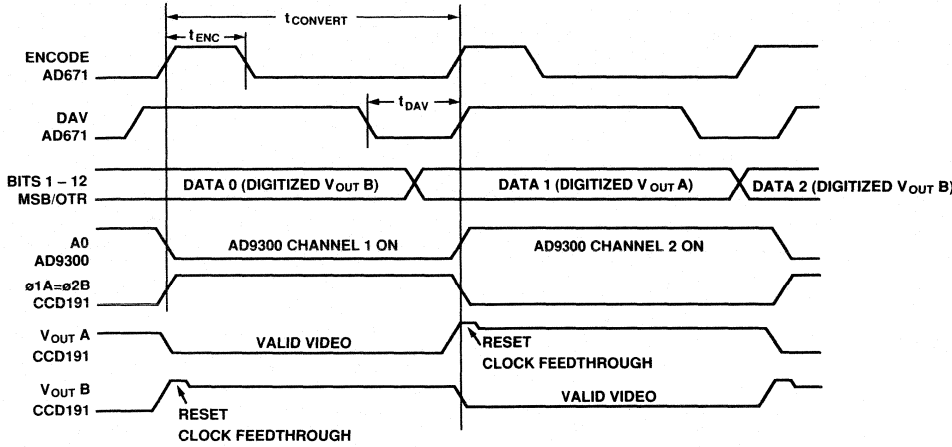

Figure 6. Linear Array CCD – Multiplexed Output Digitizing Timing Diagram

IIVING THE ANALOG INPUT TO THE AD671

achieve accurate conversion it is essential that a deband buffer or amplifier is connected at the input to ə AD671. The chosen amplifier should exhibit a low ss than 1 Ω) output impedance at frequencies in ex- ss of 10 MHz in order that its output can supply the pidly varying input currents to the AD671 while its ıtput voltage remains essentially constant.

ıe range of the CCD pixel signal at the output of the)9300 multiplexer is 0 V to approximately −1 V. A gain ctor of approximately −5 should be applied to this gnal in order to present the correct input signal range V) and polarity (+) to the AD671.

ırthermore, a high input impedance amplifier should ə used at the output of the AD9300 to avoid errors due multiplexer on-resistance and amplifier bias current oducts.

ıe AD843 FET input amplifier has low bias current nA) and wide bandwidth (34 MHz GBW), and when ɔnfigured for a gain of −5, fulfills all the above require- ıents of an input amplifier for the AD671.

EFERENCE INPUT TO THE AD671

ıe AD671 requires an external 5 V reference. The D586 provides a convenient 5 V reference. Its output ıould be decoupled with a 6.8 µF capacitor as shown in gure 3b.

GROUNDING AND DECOUPLING

It is essential that careful attention is paid to good printed circuit board layout and decoupling techniques in order to achieve the desired performance from the graphics image scanning system shown in Figures 3a, 3b and 3c. The power supply pins of each IC should be individually decoupled with a 4.7 µF tantalum and a 100 nF ceramic capacitor in parallel. The capacitors should be situated as close to the pins as possible with the ceramic one nearest to the device.

Ground planes should be used to isolate digital control signals from sensitive analog signals.

ACKNOWLEDGMENTS

I would like to express my appreciation to Tim Down of Optimum Vision (UK) and Tom Spencer of Loral Fairchild, Milpitas, California, for their invaluable comments and advice during the preparation of this application note.

REFERENCES

1989 Fairchild Weston CCD Imaging Databook
CCD191 Data Sheet (Fairchild Weston)
AD671 Data Sheet
AD9300 Data Sheet
AD843 Data Sheet

ANALOG DEVICES

AN-299
APPLICATION NOTE

NE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Simple Circuit Provides Ratiometric Reference Levels for AD782X Family of Half-Flash ADCs

by John Wynne

he 8-bit AD782X family of high speed ADCs comprises f the following:

AD7820; Single Channel, Unipolar, 1.36 µs conversion time.

AD7821; Single Channel, Unipolar or Bipolar, 660 ns conversion time (improved version of AD7820).

AD7824; Four Channel, Unipolar, 2.5 µs conversion time per channel.

AD7828; Eight Channel, Unipolar, 2.5 µs conversion time per channel.

ll of the converters run off of a single +5 V supply to onvert positive input signals. Additionally the AD7821 an convert negative input signals if its V_{SS} pin is tied to -5 V. Refer to the individual data sheets for more comlete information on the devices. This application note resents a circuit which can provide ratiometric refernce voltages suitable for use with the AD782X family of alf-flash ADCs.

he circuit in Figure 1 accepts a positive input reference oltage V_{REF} and, by implicit feedback, generates two additional reference voltages, $V_{REF}(+)$ and $V_{REF}(-)$, equally spaced above and below the original input reference level. The width of this window, determined by the ratio of R1 to R2, is a fixed percentage of the input reference voltage and remains a fixed percentage independent of the actual value of the reference level. Thus the window tracks the reference voltage. When used with ADCs such as the AD782X family, which have separate $V_{REF}(+)$ and $V_{REF}(-)$ inputs to impress a reference voltage across a resistor string, the circuit can generate programmable ratiometric reference levels to effectively increase the resolution of the ADC. With the added advantage of single supply operation the circuit can find uses in such applications as power supply monitoring, limit detection, head positioning servos' in disc drives, etc.

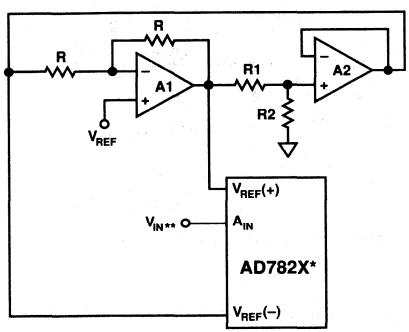

*ADDITIONAL CIRCUITRY OMITTED FOR CLARITY
**AD7820 & AD7821 – SINGLE CHANNEL, AD7824 – FOUR CHANNEL, AD7828 – EIGHT CHANNEL

Figure 1. $V_{REF}(+)$, $V_{REF}(-)$ Generator with AD782X ADC

From Figure 1 the expressions for $V_{REF}(+)$ and $V_{REF}(-)$ are as follows:

$$V_{REF}(+) = 2\,V_{REF}\frac{1}{1+X} \quad (1)$$

$$V_{REF}(-) = 2\,V_{REF}\frac{X}{1+X} \quad (2)$$

$$= V_{REF}(+)\,X$$

where: $X = \dfrac{R2}{R1 + R2}$

If a span voltage, V_{SPAN}, is defined as the difference between these levels then:

$$V_{SPAN} = V_{REF}(+) - V_{REF}(-)$$

$$= 2\,V_{REF}\frac{1-X}{1+X} \quad (3)$$

For instance, with X = 0 $V_{SPAN} = 2\,V_{REF}$
with X = 1/2 $V_{SPAN} = 2/3\,V_{REF}$
with X = 1 $V_{SPAN} = 0$

There is thus a nonlinear relationship between X and V_{SPAN}. This is shown very graphically in Figure 2 where $V_{REF}(+)$ and $V_{REF}(-)$ are plotted versus the resistor divider ratio X. Note that as X approaches unity the

generated voltages converge towards the reference voltage. Suggestions for making the ratio X digitally programmable are shown further in the text.

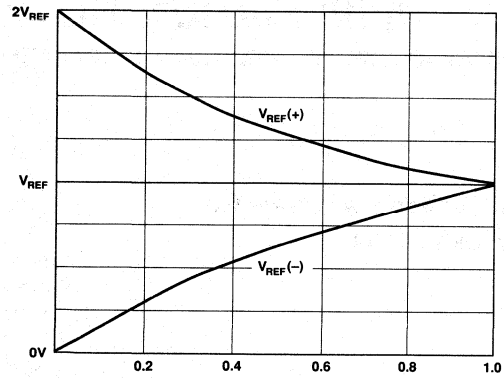

Figure 2. $V_{REF}(+)$, $V_{REF}(-)$ Voltages vs. Divider Ratio X

The plots indicate that the span voltage can vary from 2 V_{REF} to 0 V. In practice at either of these extremes, problems can arise. If used in a single-supply application, the op amp driving the bottom [$V_{REF}(-)$] of the ADC resistor string (e.g., A2 in Figure 1) will have difficulty in getting down to 0 V. Similarly, the op amp driving the top [$V_{REF}(+)$] of the ADC resistor string (e.g., A1 in Figure 1) might experience headroom problems if 2 V_{REF} approaches the op amp's V_{DD} level. At the other extreme the smallest usable span voltage is determined by the ADC. Typically the AD782X family, designed for a span voltage of 5 V, can convert with no missing codes down to a span voltage of 1 V.

Because of the transient currents which flow in the ADC resistor string during conversion it is necessary to decouple both $V_{REF}(+)$ and $V_{REF}(-)$ signals at the pins of the ADC. The op amps used for A1 and A2 should be considered for stability when driving capacitive loads of 0.1 μF or 0.2 μF.

dB or %

Dependent upon the application, the circuit can be viewed either as a means of increasing the effective resolution of the ADC or else as providing a voltage window (equal to full scale of the ADC) which can be placed around some input voltage to be monitored. For example, to increase the effective resolution of the converter from 8 bits to 10 bits in 6 dB steps (every 6 dB increase in dynamic range is equivalent to an extra bit of resolution), three different resistor divider ratios (X values) are required. Suitable values for X are found by choosing V_{REF} and V_{SPAN} and solving Equation 3 for X. Table I lists one possible set of values for X with their corresponding voltage levels computed for V_{REF} = 2.5 V.

X	$V_{REF}(+)$	$V_{REF}(-)$	V_{SPAN}	Effective Resolution
0.111	4.5 V	0.5 V	4.0 V	8 Bits
0.428	3.5 V	1.5 V	2.0 V	9 Bits
0.667	3.0 V	2.0 V	1.0 V	10 Bits

Table I. Possible Set of X Values to Increase AL Dynamic Range in 6 dB Steps

In applications such as power supply monitoring (when converting servo signals in HDD systems, it probably more useful to talk about the span range a being a percentage of the input voltage. In terms of th span voltage;

$$V_{REF}(+) = V_{REF} + V_{SPAN}/2$$

$$V_{REF}(-) = V_{REF} - V_{SPAN}/2$$

If V_{REF} is chosen to be equal to the nominal analog inpu voltage (V_{IN}) to be monitored, then the ADC can mea sure voltages which extend above and below nomina V_{IN} by $V_{SPAN}/2$. For instance, with X = 0.5; $V_{REF}(+)$ = 1.33 V_{REF} and $V_{REF}(-)$ = 0.667 V_{REF}. The span voltage i 0.667 V_{REF}. With V_{REF} chosen to be equal to the nomina input voltage V_{IN}, the full-scale analog input range to th ADC is V_{IN} ± 0.33 V_{IN} or V_{IN} ± 33%. Similarly wit X = 0.8, the full-scale range is V_{IN} ± 11%. Thus the AD result can be directly interpreted as a measure of th percentage variation of the input voltage from its nom nal value.

When monitoring voltages greater than +5 V, some in tial signal conditioning (i.e., attenuation) will be nece sary in order to avoid applying any input or referenc voltage higher than +5 V to the ADC. Figure 3 shows representative system to monitor a +15 V nominal sup ply having a tolerance of ±11%. It may be necessary t buffer the A_{IN} of the AD782X to avoid errors due to hig source impedance. To measure tolerances tighter tha

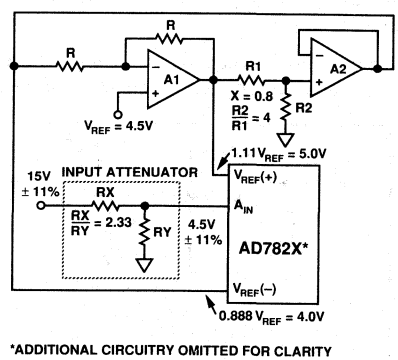

*ADDITIONAL CIRCUITRY OMITTED FOR CLARITY

Figure 3. Monitoring a +15 V ± 11% Supply

1% the span voltage will drop below 1 V since $V_{REF}(+)$ imited to +5 V maximum. To avoid this situation it be necessary to amplify the small tolerance voltage an acceptable level. One method of achieving this is own in Figure 4.

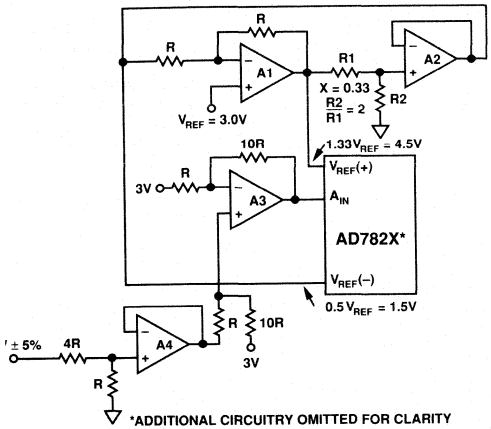

Figure 4. Monitoring a +15 V ± 5% Supply

e circuits presented so far all deal with a positive V_{REF}. n the circuits handle a negative V_{REF}? The answer is s, but only one ADC in the AD782X family can directly nvert negative voltages. This is the AD7821. By tying V_{SS} pin to −5 V (V_{DD} remains at +5 V), the AD7821 n convert negative voltages down to −5 V. However $V_{REF}(+)$ input must always be maintained more pos- ve than its $V_{REF}(-)$ input. This necessitates a switch- er of the $V_{REF}(+)$, $V_{REF}(-)$ connections when using a gative reference. For instance, in Figure 3, to monitor 15 V ±11%, the ADC must be the AD7821 with −5 V plied to its V_{SS} input; the 4.5 V reference now be- mes −4.5 V, the output of A1 now drives the $V_{REF}(-)$ put of the ADC; and the A2 output now drives $V_{REF}(+)$. ie op amps are now driven with bipolar supply volt- ges. All else remains the same.

ggested Circuits to Provide Programmable Divider atios

Figure 5 the divider ratio is dependent on whichever witch is closed, i.e., with S2 closed, divider ratio X is:

$$X = \frac{RC + RD}{RA + RB + RC + RD}$$

his circuit offers the advantage that the ON resistance f the switches does not play a part in determining the tio X.

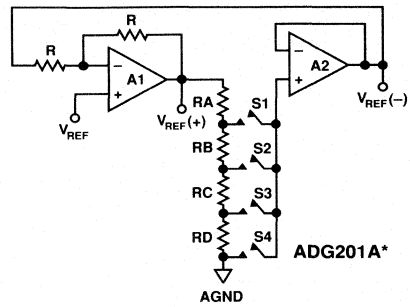

Figure 5. Programmable Divider Ratio X with Discrete Resistors

Figure 6 shows how an 8-bit DAC, the AD7524, can be used to determine whatever divider ratio X is required with much finer resolution than is possible with discrete resistors alone. The gain of A2 is given as:

$$G = 1 + \frac{256}{N}$$

where:

N is the decimal equivalent of the 8-bit DAC code.

Of the possible 255 gain levels programmable with the DAC, 128 of them cover the gain range from 1 to 2, 64 cover gains from 2 to 4, etc. Thus very fine resolution of low gains is possible. However since the noninverting gain of A2 is always greater than unity, RA & RB provide some fixed attenuation before the programmable gain stage. The overall divider ratio is:

$$X = G\frac{RB}{RA + RB}$$

The requirement that X must be always be less than unity [i.e., $V_{REF}(+) > V_{REF}(-)$] places an upper limit on the usable programmable gain range.

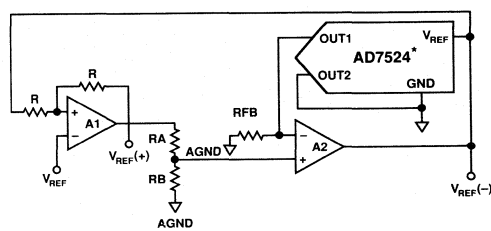

Figure 6. Programmable Divider Ratio X with 8-Bit CMOS DAC

ANALOG DEVICES

AN-215A
APPLICATION NOTE

NE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

3

Designer's Guide to Flash-ADC Testing

Part 1
Flash ADCs Provide the Basis for High Speed Conversion

by Walt Kester

Building high-performance circuits that take advantage of the high sampling rates of flash ADCs requires a knowledge of these converters' many intricacies. Part 1 of this 3-part series discusses the pitfalls of designing with flash ADCs, how to evaluate certain data-sheet specifications, and how to choose external components that complement your particular converter. Parts 2 and 3 will explore test and measurement methods that you can use to verify a converter's performance in your system.

To digitize analog signals whose bandwidths exceed 1 MHz, you'll probably need flash ADCs. Many flash converters with 4 to 10 bits of resolution are now available, thanks to recent advances in VLSI process technology and design techniques. However, to use these converters successfully at the high sampling rates that they provide, you must take into account and compensate for a variety of flash-converter characteristics.

The basic features of most flash converters are shown in **Fig 1**. A flash ADC simultaneously applies an analog-input signal to $2^N - 1$ latched comparators, where N is the number of the converter's output bits. A resistive voltage divider generates each comparator's reference voltage and sets each reference level 1 LSB higher than the level of the comparator immediately below it. Comparators that have a reference voltage below the input-signal level will assume a logic 1. The comparators with a reference voltage above the level of the input signal will produce a logic 0. A secondary logic stage decodes the thermometer code that results from the $2^N - 1$ comparisons. An optional output register latches the decoding stage's digital output for one clock cycle.

Timing is everything

One of the first difficulties you'll encounter when using flash converters is removing valid data from the converter. In practice, the comparator bank has two states controlled by a conversion-command signal. Various converters call this command the convert, the encode, or simply the clock command. When this signal is in its convert-command state, the comparators track the analog-input signal, and during this time the output data is invalid. When the command line changes state, it latches the comparator outputs. Valid output data is now available for transfer to an external register. You'll find most flash converters somewhat sensitive to the duty cycle and frequency of this command pulse. In other words, the performance of the converter, specifically its differential and integral nonlinearity performance, is related to the clock's duty cycle and frequency. Performance degradation is especially pronounced when you run the device at or near its maximum sampling rate.

Reprinted from EDN — January 4, 1990

Because of recent advances in VLSI processing and design techniques, many flash converters that have from 4 to 10 bits of resolution are available.

The way you handle the binary output depends on whether the converter has an internal output latch. Without a latch, the data will be invalid for a period equal to the sampling clock's pulse width. At high sample rates, the data-invalid time will impinge on the data-valid time, making it difficult for you to strobe the flash converter's output into an external register. For instance, if you operate a flash converter at 100M samples/sec with a 50%-duty-cycle sampling clock, the output data will be valid for only 5 nsec. When you consider the finite rise and fall time of the output binary bits, this short time doesn't leave you much leeway, even if you use the fastest external logic. In fact, you may ultimately lose data. The addition of an internal output latch simplifies clocking of the output data, because the output data is valid for approximately the entire clock cycle. In return for a longer data-valid time, you'll have to accept an inherent 1-cycle or more pipeline delay—an acceptable compromise in most systems applications.

Try to place an appropriate buffer register next to the flash converter. If you route the converter's digital output directly to a backplane data bus through a ca edge connector, signal coupling between the digi output signals and analog input will degrade the S ratio and harmonic performance.

In many high-speed data-acquisition designs, yo need a large and fast buffer memory to store the out data. A 500M-sample/sec converter can fill 1M byte memory in 2 msec. To reduce the required speed—a thus the cost—of the memory, you can demultiplex t high-speed data stream (**Fig 2**), which slows it to f quencies compatible with cost-efficient CMOS RAI **Fig 2**'s circuit clocks the two output registers at h the sample rate, and it latches data in each regis 180x out of phase from the other. Some flash convert that operate in excess of 200 MHz have onboard multiplexing for added convenience.

All that sparkles isn't gold

So far, these timing difficulties refer to how y deal with the converter's output data. But flash c verters can have internal problems as well. Low in frequencies can cause comparator metastability, a

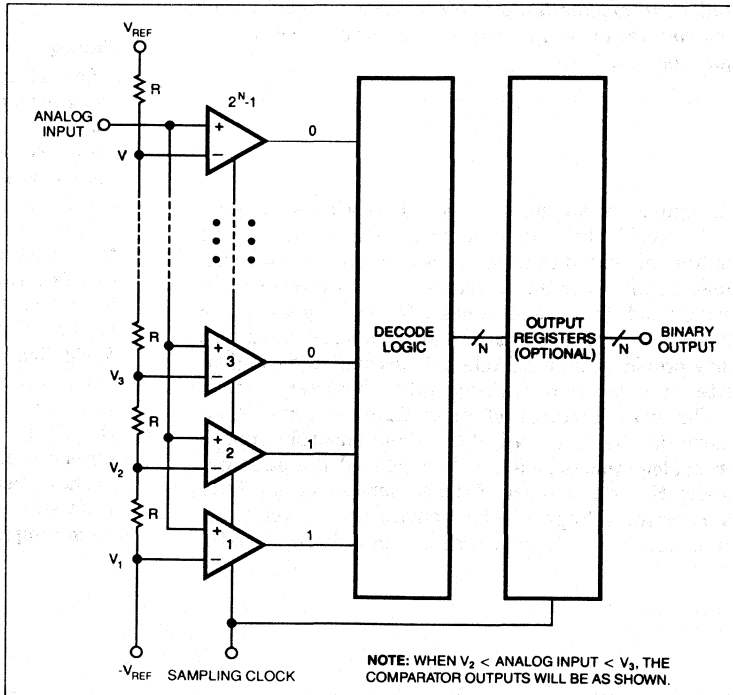

***Fig 1—Flash converters contain a bank** of $2^N - 1$ comparators, where N is the number of output bits. Decoding logic transforms the comparator outputs into the appropriate N-bit result. Timing and input characteristics of the converter, such as mismatched comparator delays, mismatched ladder resistors, and nonlinear input capacitances, are just a few sources of converter errors.*

igh input frequencies can lead to errors caused by ew-rate and delay mismatches. All of these errors ıay manifest themselves as sparkle codes in a poorly esigned flash converter.

Sparkle codes are random errors whose magnitude ıay approach the full-scale range of the converter. 'he term refers to the white dots or "sparkles" that ppear against a gray background when the ADC out- ut drives a video display. There are two sources of parkle codes: comparator metastable states and ther- ıometer-code bubbles.

A comparator metastable state occurs if the com- ıarator output falls between the logic-0 and logic-1 hreshold of the digital decoding logic. If the threshold ıncertainty region has a width of ΔV_L and the compara- ,or has a gain of A, then the error probability P_m is a ıniformly distributed value that's equal to

$$P_m = \frac{\Delta V_L}{Aq},$$

where q is the weight of the LSB.

A latched comparator in a flash converter has a regenerative gain of

$$A = A_0 e^{t/\tau}$$

when $t>0$, and

$$A = A_0$$

when $t \leq 0$. τ equals the regeneration-time constant, and t is the time after the application of a latch command. The probability of a metastable state, P_m, for a regenerative comparator bank driving decoding logic is

$$P_m = \frac{\Delta V_L}{A_0 q} e^{-t/\tau}.$$

The magnitude of the sparkle code depends on the location of the metastable comparator in the comparator bank and on the logic-decoding scheme. For instance, the 128th comparator determines the MSB of an 8-bit flash converter with straight binary decoding logic. If this comparator's output is in a metastable state, the decoding logic may mistakenly convert the input voltage whose correct binary representation is 01111111 to 11111111, producing a full-scale error. If the comparator bank's thermometer-code output is first decoded into Gray code, latched, and then converted into binary code, the metastable error is reduced to 1 LSB, regardless of the comparator in error. Flash converters, however, rarely use this scheme because of the ripple-through time and the increased logic density of the Gray-to-binary circuitry. Instead, converter designers often use "pseudo-Gray" decoding techniques to eliminate the delay time associated with traditional Gray-to-binary circuits.

Note that the probability of a metastable-state error increases as the time-after-latch, t, decreases (assum-

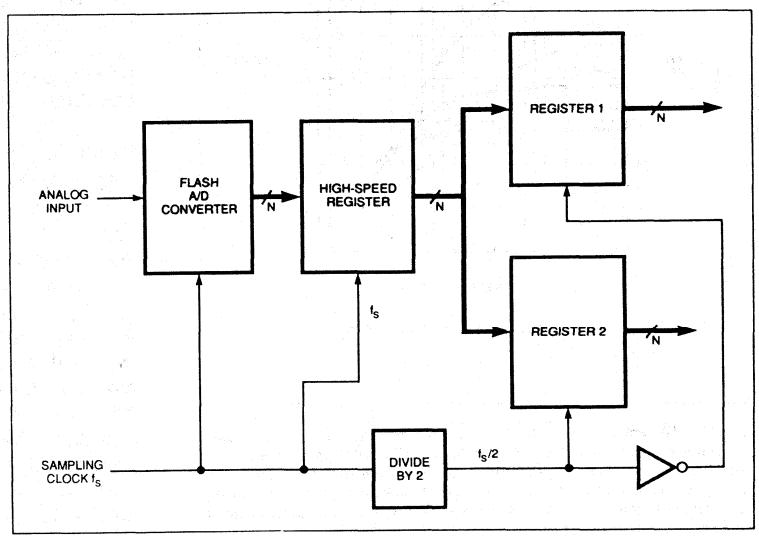

***Fig 2—Flash converters can operate** at extremely high sample rates. Thus, you must have an output register that can handle this fast data. By demultiplexing the converter's outputs, you can slow the data to a rate that's compatible with standard CMOS memory.*

One of the first difficulties you'll encounter when using flash converters is removing valid data from the converter.

ing a constant value of τ). This implies that the flash converter is more apt to produce metastable-state errors as you increase the sampling rate, because t must decrease correspondingly. Most manufacturers reduce metastable comparator states by minimizing the regeneration-time constant, τ. Lower regeneration-time constants result in higher power dissipation, which is one reason why many high-speed flash converters are power-hungry devices.

Thermometer-code bubbles are another potential source of sparkle codes. A well-behaved flash converter's comparator bank produces a specific sequence of ones up to a certain point in the input range of the converter, and it produces a sequence of zeros beyond that point. The decoding logic then assigns a binary number to the thermometer code. For low-frequency inputs, most flash converters' comparator banks are well behaved. At high speeds, however, delay mismatches among comparators may produce out-of-sequence ones and zeros in the thermometer code. The decoding logic then assigns an error binary code to these out-of-sequence points, or bubbles, which also result in sparkle codes. Again, proper comparator design and more sophisticated decoding-logic circuitry

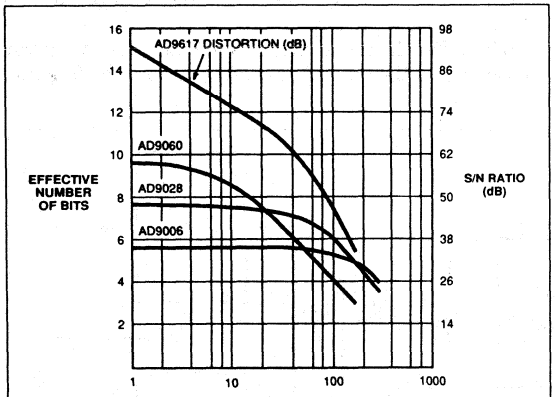

FLASH ADC	RESOLUTION (BITS)	MAXIMUM SAMPLING RATE (M SAMPLES/SEC)	TEST SAMPLING RATE (M SAMPLES/SEC)
AD9060	10	75	60
AD9028	8	300	250
AD9006	6	500	400

***Fig 3—Theoretically, a flash converter should maintain its performance** across the full Nyquist bandwidth. But as the curves in this figure illustrate, the S/N ratio starts to degrade well below each device's maximum sampling rate. (Note that these curves are based on test sampling rates that are somewhat lower than each device's maximum possible rate.)*

within the ADC itself can reduce these errors to acce able levels.

These comparator-timing errors degrade both t differential and integral linearity of a flash ADC as t input slew rate increases. In addition to static erro such as missing codes and sparkle-code errors, sle rate limitations manifest themselves as dynamic erro such as increased harmonic distortion and degradati in the S/N ratio. Ideally, a flash converter should mai tain its static performance specifications across the fi Nyquist bandwidth, and certain applications demar full performance beyond the Nyquist bandwidth. Th theoretical, rms S/N ratio for an N-bit ADC is give by the well-known equation

$$\text{S/N ratio} = 6.02N + 1.76 \text{ dB}.$$

However, as shown in a typical plot of S/N rati versus input frequency for actual flash converters (Fi 3), the S/N ratio degrades as the input frequency ir creases. This degradation starts well below these cor verters' maximum sampling rates. The left vertical axi of **Fig 3** shows the S/N ratio in another term: effectiv number of bits. The effective number of bits is simpl the value of N when you solve the above equation usin a specific value for the S/N ratio. Aperture jitter (sam ple-to-sample variations in the effective-sampling in stant) can also cause degradation in the overall S/N ratio for high-slew-rate inputs. Jitter can be interna or external to the converter. Part 3 of this series wil cover this topic in more detail. To minimize externall produced jitter components, you should always practic proper grounding, power-supply-decoupling, and pc board-layout techniques.

Watch for dangerous data-sheet territory

Most of the timing difficulties and error sources discussed so far are common to all flash converters. However, each converter features its own unique design and specifications; thus the data sheets require scrutiny. Don't be fooled into believing that sampling rate and input bandwidth are interchangeable; they're different specifications. It wasn't until recently that you'd even find input bandwidth specified for an ADC. Even now, no accepted industry-wide definition exists for a flash converter's full-power-bandwidth specification. Scrutinize the data sheet carefully and make sure you understand the manufacturer's definition and test method. The full-power bandwidth of a traditional op amp is the maximum frequency at which the amplifier

n produce the specified p-p output voltage at a speci-
d level of distortion. Another commonly used defini-
n calculates the full-power bandwidth by dividing
e amplifier's slew rate by $2\pi V_O$, where the output-
ltage range of the amplifier is $\pm V_O$.

When you apply traditional analog-bandwidth defini-
ns to flash converters, the results can be misleading.
ıe dynamic-error sources previously discussed may
ecome predominant long before the comparator front
ıd approaches its maximum bandwidth. If you use a
mmon definition of full-power bandwidth as the fre-
ıency at which the p-p reconstructed-sine-wave out-
ıt is reduced by 3 dB for a full-scale input, then the
fective number of bits (S/N ratio) at this input fre-
ıency may render the flash converter useless in your
stem. Thus, to get a true idea of a converter's per-
rmance, you must consider both the full-power band-
idth and the effective number of bits (S/N ratio) at
specific sampling rate.

Another definition you'll encounter occasionally for
ıll-power bandwidth is the maximum, full-scale input
gnal at a specified sampling rate that produces no
ıissing codes. Using this definition always gives the
ıost pessimistic number, so specifications based on
ıis definition appear on only a few data sheets. The
llowing is a recently proposed definition for full-
ower bandwidth (courtesy Chris Manglesdorf, a sen-
r scientist at Analog Devices): the frequency at which
he fundamental component of the reconstructed FFT
utput—excluding harmonics—is reduced by 3 dB from
ıll scale.

Just when you think you've mastered the intricacies
f flash ADCs, you'll realize that you have another
omponent to worry about: the input buffer amplifier.
Fortunately (or unfortunately depending upon your perspective), the flash converter—not the amplifier—usually limits the converter's dynamic performance. A flash converter typically digitizes a signal from a 50, 75, or 93Ω bipolar or unipolar source. If the input range of the flash converter is incompatible with the signal, then you'll clearly need a wideband op amp to generate the required gain and offset (**Fig 4**). In addition, the input capacitance of some flash converters may vary as a function of the analog input's signal amplitude. Therefore, so that the nonlinear capacitance doesn't produce undesirable harmonics in the digitized signal, you'll need to use a buffer amplifier for isolation. For certain flash converters, the input capacitance is so high that a buffer amplifier is needed just to preserve the signal bandwidth.

A good choice for the buffer is a high-speed transimpedance amplifier. These amplifiers have high bandwidths and flat frequency responses over a broad range of input frequencies. Also, many transimpedance amplifiers exhibit extremely low distortion. Pairing the right amplifier with your converter is important. For instance, **Fig 3** shows the S/N ratio of various converters plotted along with the harmonic distortion of the AD9617. Since the THD of the amplifier is better than the S/N ratio of the converters, the amplifier won't degrade the flash converters' performance over the major portion of their usable bandwidth.

Another factor to consider when driving the input of flash converters is the input-signal polarity. Positive input signals, which forward bias the substrate diode, can damage a converter that has a unipolar, negative input-voltage range. Installing an external Schottky diode provides effective protection.

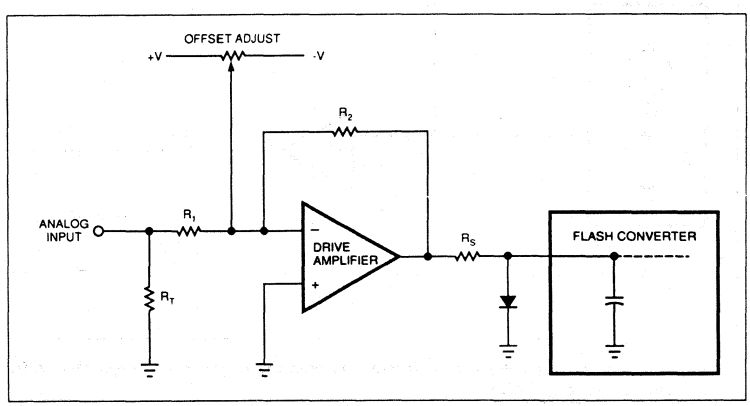

***Fig 4—If the voltage range of the analog signal** and the input range of the ADC aren't compatible, you'll have to adjust the gain using R_1 and R_2 and also adjust the input signal's offset. Because of the substantial value and the often nonlinear nature of a flash converter's input capacitance, you must choose an appropriate drive amplifier and value for R_S.*

Ideally, a flash converter should maintain its static-performance specifications across the full Nyquist bandwidth, but in reality ADCs fall far short of this ideal.

The flash converter's input capacitance and the drive amplifier's isolation resistor, R_S, form a lowpass filter. Typical series-resistor and input-capacitance values of 10Ω and 20 pF create a single-pole lowpass filter that has a bandwidth of 800 MHz. However, if the input capacitance changes from 20 to 15 pF over the input range of the converter, then an attenuation error of 1.4% occurs for a 50-MHz input signal. This 1.4% nonlinearity will produce 37 dBc of harmonics. (The unit dBc refers to the number of dB between the signal you're measuring and the carrier frequency). If you minimize the value of R_S and still maintain op-amp stability, you can reduce the attenuation error caused by these lowpass filtering effects. The signal dependence of input capacitance is rarely specified, but as converters move toward higher bandwidths, you can expect to see this parameter on more data sheets.

Few, if any, flash converters contain an internal voltage reference, so in addition to an external drive amplifier, you must also design your own voltage-reference generator. **Fig 5** illustrates a typical −2V, unipolar reference-voltage circuit for a flash converter. A buffer transistor is necessary because the resistance of the converter's ladder string is usually fairly low. The total reference-ladder resistance of a flash converter depends heavily on the fabrication process and may vary considerably from device to device. Also, the ladder's resistance may exhibit a large temperature coefficient.

If the flash converter allows bipolar operation, tl you'll have to generate two reference voltages. T circuit in **Fig 6** allows great flexibility in setting b the gain and the offset of a bipolar flash convert and it operates on ±5V power supplies. A few fl converters provide a sense pin for the voltage ref ence. You can use this pin to compensate for the v age drop caused by the package's pin and bond-w resistances, as shown in the bipolar-reference circ in **Fig 6**. In addition, some ADCs give you access one or more taps along the internal, reference-ladd resistor string. To achieve better integral linearit you can drive these taps from low-impedance source

Improve dynamics with T/H amplifiers

As previously discussed, the effective-sample tim delay variations among the latched comparators d grade the S/N ratio and the harmonic performanc You can visualize the individual comparators with an array as having variable delay lines in series wi their latch-strobe inputs. To understand the effect this delay on performance, consider an 8-bit, 100 sample/sec flash converter that's digitizing a full-scal 50-MHz sine-wave input. You can express the si wave as

$$v(t)=V_P\sin 2\pi ft.$$

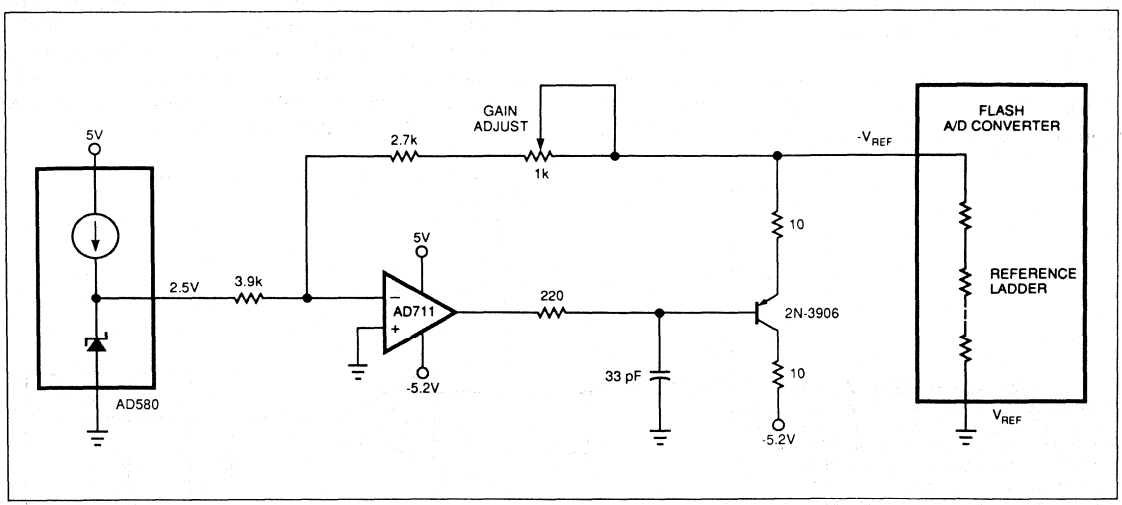

Fig 5—Flash converters don't have internal references, *so you must design them externally. This particular circuit provides a stable −2V reference for a unipolar converter.*

o improve flash-converter performance at ımpling rates as high as 25 MHz, you ın use front-end T/H amplifiers to imple- ıent a "track-and-slow-down" approach.

he maximum rate-of-change of this signal occurs at ıe zero-crossing point and is equal to

$$\left.\frac{dv}{dt}\right|_{max}=2\pi fV_P=\left.\frac{\Delta v}{\Delta t}\right|_{max}.$$

y solving this equation for Δt_{max}, you obtain

$$\Delta t_{max}=\frac{\Delta V}{2\pi fV_P}.$$

f the input-voltage range of the flash converter is 2V, r $V_P=1V$, then the LSB weight is 8 mV for an 8-bit ıDC. For the flash converter's error to be less than LSB, Δt_{max} must equal 25 psec. The effective-sample elay mismatch between comparators can't exceed this alue. If the mismatch is greater, a 50-MHz, full-scale ine-wave input will produce missing codes in the con- erter's output.

Placing an ideal track-and-hold (T/H) amplifier ahead f the flash converter theoretically would eliminate this ıroblem, because the flash converter basically would be digitizing a dc input. In actual practice, T/H amplifiers aren't ideal, especially at high speeds. The signal presented to the flash converter is still changing, although at a slower rate. Nevertheless, this "track-and-slow-down" approach can improve the flash-converter performance at sampling rates as high as approximately 25 MHz. At sampling rates above 25 MHz, the T/H circuit needs to be mounted on the same substrate as the flash converter in a suitable hybrid package. Monolithic T/H amplifiers in hybrid packages with 8-bit flash converters have successfully achieved 7 effective bits at Nyquist inputs and at sampling rates of 250 MHz. These hybrid packages exact a penalty of higher cost and power consumption, however.

You'll find it difficult to select an appropriate discrete T/H amplifier, because the interaction between the T/H amplifier and the flash ADC is hard to predict. You should evaluate key T/H amplifier specifications such as acquisition time, full-power bandwidth, slew rate, and harmonic distortion. Harmonic-distortion specifications typically are provided for the track mode. The T/H amplifier's performance may be considerably

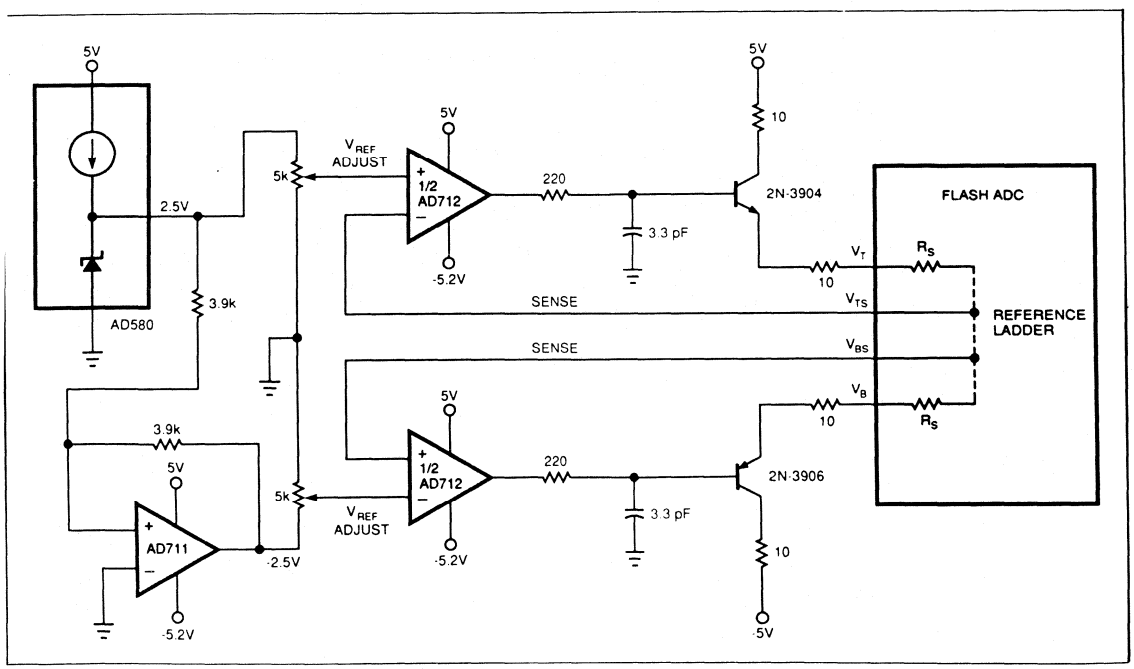

***Fig 6—This reference generator** for a bipolar flash converter uses the sense pins provided by the resistor ladder to compensate for the voltage drops in the package's pin and bond-wire resistances.*

worse than the stated specifications when the amplifier is in the hold mode and actually driving a flash converter. Also, the loading effects of the converter may degrade the T/H amplifier's performance. In addition, obtaining the optimum relationship between the various timing pulses that drive the T/H amplifier and the flash converter may require considerable experimentation.

Because of these many difficulties, the current goal of many ADC manufacturers is either to provide flash converters whose dynamic performance is acceptable without a T/H function, or to integrate the T/H function and converter on the same chip. In either case, manufacturers can fully specify the ADC for dynamic performance and spare you the somewhat difficult design problems associated with interfacing the T/H amplifier to the converter.

The knowledge of internal ADC features and t requirements these ADCs place on external circu should guide your initial design efforts. But once y finish the design and build your circuit, you'll want ensure that the combined performance of your cc verter and its support circuits meets your requi ments. Part 2 will discuss various DSP-based te methods that are particularly effective in flash-co verter testing.

Reference

1. Sheingold, Dan, *Analog-Digital Conversion Handbook* 3rd ed, Prentice-Hall, Englewood Cliffs, NJ, 1986.

ANALOG DEVICES

AN-215B
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Designer's Guide to Flash-ADC Testing

Part 2
DSP Test Techniques Keep Flash ADCs in Check

by Walt Kester

By testing your flash A/D converter, you can ensure that it's faithful to all the specifications listed on its data sheet. Part 2 of this 3-part series presents a number of methods, including sine-wave curve fitting and the FFT, that you can use to test flash converters. Readily available benchtop instruments or personal computers are the only equipment that you'll need to use these methods.

It's important to know how your flash A/D converter will perform in real-world applications. Therefore, you may want to perform any one of a variety of tests on your converter to determine its deviation from ideal performance. As Part 1 of this series discussed, flash ADCs exhibit errors due to static and dynamic nonlinearities, and these errors increase as the input signal's slew rate increases. Thus, the actual S/N ratio will fall short of the converter's theoretical value. Even if you don't apply these tests yourself, becoming familiar with them will help you evaluate data-sheet specifications more accurately, because many manufacturers use these same methods.

Another reason you may want to test your flash converter is to gain information that the manufacturer doesn't provide. Specifications such as S/N ratio and its related effective number of bits are key in all applications and are normally specified, but other specifications that are more important for your particular application may not be included on the data sheet. For example, video designs typically require that you know a converter's differential phase and gain (**Ref 1**). Communications systems may even depend on esoteric specifications such as the spurious-free dynamic range, which isn't available on many data sheets.

For a full-scale sine-wave input, the theoretical rms-signal to rms-quantization noise ratio is

$$\text{S/N RATIO} = 6.02N + 1.76 \text{ dB},$$

where N equals the number of bits (**Ref 2**). The rms quantization-noise voltage for an ideal ADC within the Nyquist bandwidth is $q/\sqrt{12}$, where q is the weight of the LSB expressed in volts.

The most popular method for extracting a flash converter's S/N ratio and effective number of bits is through discrete Fourier transforms (DFTs). Today, you can perform sophisticated DSP tests with PC-based test systems and standard software packages. The test system in **Fig 1**, for example, can execute a 1024-point FFT in less than one second. Most of the hardware you'll need is available as plug-in boards for the PC. However, you'll have to do a fair amount of work before you can begin to use a PC-based test system. First, you'll need to design a high-speed buffer-memory board to capture the data from the flash ADC. Typically, you'll need to use high-speed static CMOS or ECL RAMs. Second, plan to design an appropriate logic interface to connect this buffer memory to the digital I/O card of the PC.

Another hardware feature you might consider is an evaluation board, which certain manufacturers of video-speed ADCs supply to ease design testing. Many evaluation boards contain reference voltages, power-supply decoupling, timing circuits, output registers, and connectors. The evaluation boards usually have a

Reprinted from EDN – January 18, 1990

DSP test techniques determine your converter's deviation from ideal performance, and they even tell you certain specifications that the ADC's data sheet doesn't.

matching reconstruction DAC. In most cases, the manufacturer has optimized the design of these boards so that your ADC test won't be corrupted by faulty or poorly designed support circuits.

Your software must include a program to capture the data and then load it into the memory of the PC. If you plan to use FFT analysis, you must link a standard FFT software package to your test program. You may also have to generate a look-up table to store any special weighting functions required by your particular sampling scheme. Also, adding a coprocessor card will speed up the thousands of multiplications that FFT-based analysis requires.

If you don't have the time or the energy to build your own test system, consider one of the benchtop instruments available from a number of instrumentation manufacturers. These turnkey systems typically utilize a high-speed logic analyzer to capture data. Because menu-driven software allows you to select from a variety of tests, you incur practically no hardware or software development time.

To test a flash ADC using Fourier analysis, you must apply a spectrally pure sine wave to the converter and store a number of contiguous output data samples. Then, using DFT techniques, your test program calculates the rms-signal and rms-noise content and determines the ratio of the two. Noise calculations using DFT techniques include not only the converter's quantization noise but also the harmonics of the input sine wave. In addition, harmonics that fall outside the Nyquist bandwidth are aliased back into the Nyqui bandwidth because of the sampling process. Thus, t achieve accurate and repeatable results, the purity the sampling clock and the input sine wave is critical

You can use either coherent or noncoherent samplin to evaluate the ADC performance. Coherent samplin simply means that your record of samples contains a integer number of sine-wave input cycles. Alterna tively, noncoherent sampling produces a record tha contains noninteger multiples of the input. You mus choose between these sampling schemes based on th type of input data you expect. Coherent testing is mor suited to a laboratory environment when you kno the precise frequency content of an input signal, an it requires careful attention in the selection of the inpu and sampling frequencies. Noncoherent testing yield a better representation of ADC performance in a rea world application such as spectral analysis, becaus the precise frequency content of the signal being dig tized is a mystery.

However, whenever the number of time sample doesn't contain an integer number of input cycles (nor coherent testing), you'll have to time-weight the sam ples to reduce frequency side lobes. Without weight ing, discontinuities will cause the main lobe's energy– the fundamental—to leak into many other frequenc bins. The term "bins" refers to the spaces betwee spectral lines or spectral peaks. The number of bin for a particular spectrum equals the sampling fre

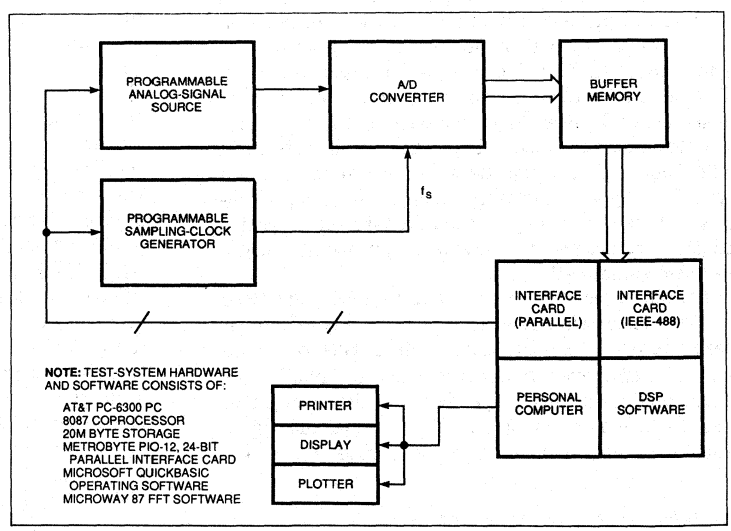

Fig 1—This DSP test system for a flash ADC *can execute a 1024-point FFT in less than one second, but the system requires a significant design effort. Hardware requirements include a high-speed buffer memory and logic interface between this memory and your PC. Software requirements include a program to capture the data and load it into the memory of the PC, as well as a link between your test program and a standard FFT software package.*

iency divided by the record length, or f_S/M. The leak-
:e of the signal from the central bin to side-lobe bins
akes accurate spectral measurements impossible—
)u simply can't distinguish the frequency bins that
ntain actual signal information from those that con-
in noise. Another reason to time-weight the samples
that the end user of your A/D-conversion system
ay be interested in the performance of the ADC using
1 identical or similar window.

Noncoherent sampling involves fewer input- and
ampling-frequency restrictions than coherent sam-
ing does, but it requires careful attention in the selec-
on and use of the weighting function. Also, to prevent
asking out harmonics of the fundamental, avoid using
puts that are integer submultiples of the sampling
equency. If your input frequency is an integer
ibmultiple of the sampling frequency, the quantiza-
on noise, $q/\sqrt{12}$, will be concentrated in the harmon-
s of the input frequency rather than uniformly distrib-

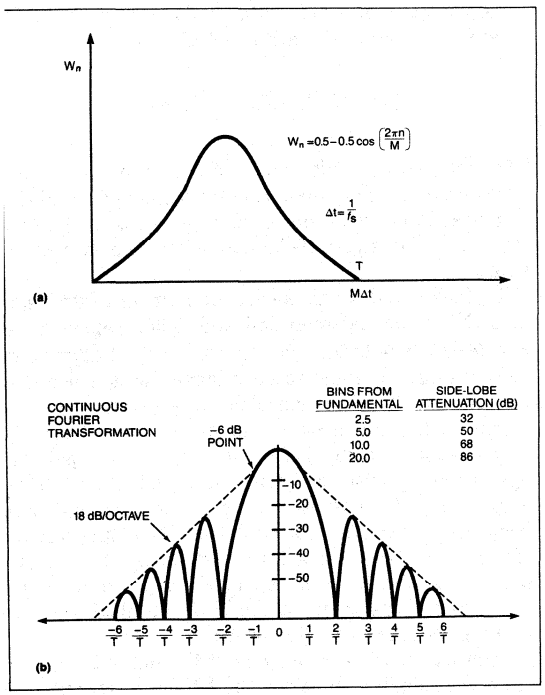

Fig 2—When the record of samples doesn't contain an integer number of input cycles*—that is, when you're using noncoherent sampling—you must precondition the data with a weighting function. The Hanning window shown here in the time domain* ***(a)*** *and the sampled-frequency domain* ***(b)*** *is a popular weighting function.*

uted across the Nyquist bandwidth. Ultimately, this condition leads to incorrect harmonic-distortion test results.

One popular weighting function is the Hanning window (**Ref 3**), which is described by the equation

$$W_n = 0.5 - 0.5 \cos\left(\frac{2\pi n}{M}\right),$$

where W_n is the weighting coefficient for the nth data sample, and M is the total number of samples. **Fig 2** graphically depicts the Hanning window in both the time and the frequency domains.

To calculate the S/N ratio, you have to decide the number of frequency bins to include in the fundamental and the number of bins to consider as noise. As **Fig 2** shows, you can correlate the amount of side-lobe attenuation with the lobe's distance, in terms of bins, from the fundamental bin. **Fig 2b** includes a table that lists some of these values. You'll have to make your decision based on the theoretical S/N ratio of the converter you're testing.

For example, an 8-bit converter has a theoretical, maximum S/N ratio of approximately 50 dB. In order to ensure that the side-lobe energy doesn't cause an artificially high noise measurement (and hence an artificially low S/N-ratio measurement), you should include at least 10 frequency bins on either side of the fundamental when calculating the signal level. (Simply take the square root of the sum of the squares of all 21 bins as your signal level.) Now, any side-lobe energy outside this region will be at least 68 dB below the fundamental signal level (18 dB below the theoretical, 8-bit quantization noise floor of 50 dB), and side-lobe leakage won't significantly affect the accuracy of your S/N-ratio measurement.

Other weighting functions may better suit your application. For example, **Fig 3** compares the popular Hanning window's spectral representation with the more sophisticated, minimum 4-term, Blackman-Harris type. For the same record length, the Blackman-Harris window provides better spectral resolution than the Hanning window, making it more suitable for critical spectral analysis, such as measuring 2-tone, third-order intermodulation-distortion products. The extra computations for the Blackman-Harris window don't lengthen processing time, because you calculate them only once and store them in a look-up table.

As previously stated, you can use coherent sampling if you know the characteristics of your input signal and if you choose the sampling rate accordingly. Coherent

Today, you can perform sophisticated DSP tests with PC-based test systems and standard software packages.

sampling eliminates leakage and the need for windowing (**Ref 4**); the spectral result of a coherently sampled signal is simply a single-frequency peak. Certain restrictions apply to the choice of the sampling rate and the sine-wave frequency, however. First, you must observe the following ratio:

$$\frac{f_{in}}{f_S}=\frac{M_C}{M}.$$

M_C equals the number of integer cycles of the sine wave during the record period. For a whole number of cycles, M_C must be an integer. To ensure that you don't take repetitive data, M_C should also be a prime number: 1, 3, 5, 7, 11, 13, 17, etc. By using prime numbers, you ensure that all samples during the record period are unique. When using coherent sampling, it's mandatory that the ratio M_C/M be constant. This requirement implies that you derive f_S and f_{in} from two locked frequency synthesizers.

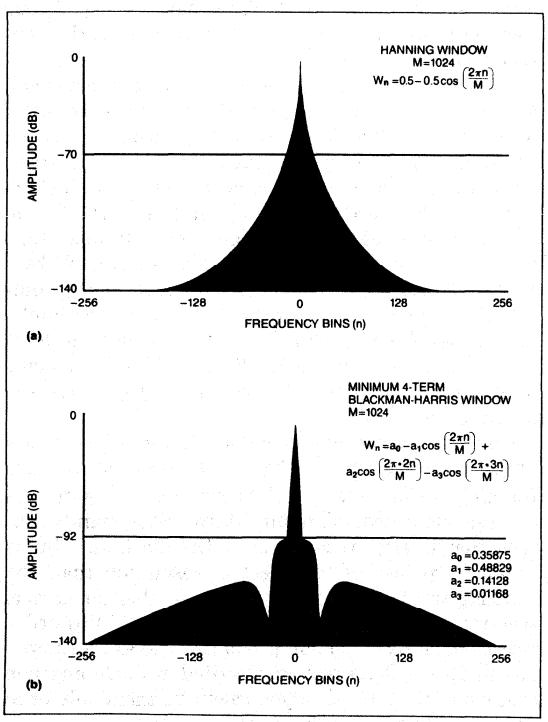

Fig 3—The Blackman-Harris windowing function (b) resolves closer peaks in a frequency spectrum than does the Hanning window (a). The mathematical expression for the Blackman-Harris window is more complex, but you only need to calculate the terms once and then store them in a look-up table.

Calculate the DFT

After selecting the record length and determini the weighting function (for noncoherent sampling), y must write your DFT test program. Your progra must find the DFT of the sequence of weighted da samples for M/2 frequencies (the Nyquist frequency Thus, the program should solve the following two equ tions for the kth frequency:

$$A_k=\frac{1}{M}\sum_{n=1}^{M} W_nD_n\cos\left[\frac{2\pi k(n-1)}{M}\right]$$

$$B_k=\frac{1}{M}\sum_{n=1}^{M} W_nD_n\sin\left[\frac{2\pi k(n-1)}{M}\right].$$

In these equations, A_k and B_k represent the magn tudes of the cosine and sine parts of the kth spectr line; n is the number of time samples; W_n is the weigh ing function; D_n is the amplitude of the time-functio data point; and k is the number of a spectral line. Th total magnitude of the kth spectral line is

$$\text{MAGNITUDE}_K=\sqrt{A_k^2+B_k^2}.$$

The program's results yield M/2 components, whic are the frequency-domain representation of the M tim samples. The resolution or spacing between the spe tral lines, Δf, equals f_S/M and is the bin size or bi width.

Typically, you should select the number of time sam ples (M) to be between 256 and 4096, depending o the desired resolution and the size of the buffer mem ory. M must be equal to an integer that's a power o two. If you're using noncoherent sampling, you ca compress leakage around the main lobe by using a larger record length, thereby leaving a larger percent age of the Nyquist spectrum uncontaminated. For ex ample, the Hanning weighting leakage is ±10 bins from the fundamental for 68-dB side-lobe suppression. If the record length is 256, then the leaky fundamental occupies 20 bins out of 128 spectral components, or 16% of the digital spectrum. When M equals 1024, the percentage reduces to 20/512, or 4%.

In practice, you can use one of the many FFT algorithms to simplify and speed the DFT calculations (**Ref 5**). An FFT algorithm will produce the same results as the DFT equations above, and the computation time is much faster.

he discrete Fourier transform is the most ɔpular method for determining a con- ːrter's true S/N ratio and effective num- ːr of bits.

ːrify the FFT

Consider the noise floor when verifying the FFT. ɜssuming that the round-off error contributed by the SP-noise calculation (the error caused by using a fi- te number of bits in the FFT multiplications and ldditions) is negligible, the rms-signal to rms-noise vel in a single frequency bin of width Δf is

$$\text{S/N RATIO}_{\text{FFT}}=6.02\text{N}+1.76\text{ dB}+10\text{ LOG}_{10}\left(\frac{\text{M}}{2}\right).$$

his equation represents the FFT noise floor. You ıould choose M so that any spurious components you ˈant to resolve lie at least 10 dB above this floor.

Basic software can easily generate an ideal N-bit ine wave by using the Integer (INT) function to trun- ate the value to the proper resolution. For instance, n input signal of frequency f_{in} is equal to

$$V_q=V_0\sin\left(\frac{2\pi nf_{in}}{f_S}\right),$$

vhere n is the nth time sample for an ADC that has nfinite resolution. You can calculate the corresponding ıuantized value using

$$V_q(n)=\text{INT}\left(\frac{V_0\sin\frac{2\pi f_{in}}{f_S}}{q}\right),$$

where $q=2V_0/2^N$. Substituting this expression for q n the above equation yields

$$V_q(n)=\text{INT}\left[2^{N-1}\sin\left(\frac{2\pi nf_{in}}{f_S}\right)\right].$$

The INT function simply truncates the fractional portion of $V_q(n)$.

To check the dynamic range of the FFT, calculate the S/N ratio by using 6.02N + 1.76 dB for increasing values of N and observing the point at which the S/N ratio no longer increases by 6.02 dB/bit. The sine-wave input to the weighting function and the FFT are more ideal as N approaches infinity. By making N arbitrarily large, you can greatly reduce quantization-error effects and analyze the true noise floor of the FFT. You can also examine the characteristics of the weighting function.

Match the sine wave to a curve

Another test method to use with flash ADCs is sine-wave curve fitting. You perform this test after the ADC digitizes the sine wave and after your test system stores the data in its memory. A record length of 1024 samples is usually sufficient. The software then calculates the best-fit, ideal N-bit sine wave to match the data points, based on the sine wave's amplitude, offset, frequency, and phase required to minimize the rms error between the actual and the ideal sine wave (**Refs 6** and 7). This method also requires that the input sine-wave frequency contains no subharmonics of the sampling rate. If you know the precise sine-wave frequency, the curve-fit algorithm is much simpler than the FFT method, and the probability that the algorithm will converge is higher.

After the software computes the rms error, Q_A, between the ideal sine wave and the actual sine wave, you can calculate the effective number of bits by using

$$\text{EFFECTIVE NUMBER OF BITS}=\text{N}-\text{LOG}_2\left(\frac{Q_A}{Q_T}\right),$$

where Q_T is the theoretical rms quantization error, $q/\sqrt{12}$. This measurement includes errors due to differential nonlinearity, integral nonlinearity, missing codes, aperture jitter, and noise, in addition to the quantization noise.

The effective number of bits that you calculate using the sine-wave curve-fitting method correlates with the value of the full-scale, FFT S/N-ratio measurement obtained using the equation

$$\text{EFFECTIVE NUMBER OF BITS}=\frac{\text{S/N RATIO}_{\text{ACTUAL}}-1.76\text{ dB}}{6.02}.$$

However, if you measure the effective bits of a sine-wave input signal whose amplitude is less than full scale, you must include the following correction factor in the above equation to achieve correlation between the two methods:

$$\text{EFFECTIVE NUMBER OF BITS}=\frac{\text{S/N RATIO}_{\text{ACTUAL}}-1.76\text{ dB}}{6.02}+\frac{\text{LEVEL OF SIGNAL BELOW FULL SCALE (dB)}}{6.02}.$$

One useful method for reducing the effects of the D/A converter in making gross back-to-back measurements on an ADC is the beat-frequency method. **Fig 4** illustrates a basic test setup. This test method stresses the converter with a near-Nyquist signal and drives the converter at its maximum sampling rate. Thus, the analog-input sine wave should be slightly lower in frequency than half the sampling frequency. The test system updates the registers that drive the DAC at an even submultiple of the sampling rate, f_S/N,

Coherent testing is more suited to a laboratory environment; noncoherent testing more closely represents ADC performance in the real world.

where N is a power of 2. (N is not the ADC's resolution.) The resulting signal from the DAC is a low-frequency sine wave whose exact frequency equals the difference between half the sampling rate and the analog-input frequency. As **Fig 4** shows, you should clock the DAC at a much lower rate, f_S/N—known as the decimation rate—thereby reducing the effects of glitches and other dynamic errors.

You can use the beat-frequency method to make signal-to-noise measurements over the Nyquist bandwidth, $f_S/2N$. You also can examine the low-frequency beat on an oscilloscope for missing codes and other nonlinearities. To measure the harmonic content of the beat frequency, you can use a low-frequency spectrum analyzer. The harmonics of the low-frequency beat are directly related to the harmonics of the analog-input frequency. A beat frequency of a few hundred kilohertz works well. To prevent jitter on the low-frequency beat signal, you must derive both the analog-input sine wave and the sampling frequency from frequency synthesizers or crystal oscillators.

This beat-frequency test is also effective in measuring the flash converter's performance for input signals near the sampling frequency. The performance under these conditions is useful for radar in-phase and quadrature-phase systems and in IF-to-digital conversion. To perform this test, set the ADC's analog-input frequency to slightly less than the sampling rate. The circuit generates a low-frequency beat even if the DAC updates at the sampling rate. However, updating the DAC at f_S/N reduces the effects of DAC dynamic errors on the measurements.

You can use DSP techniques and FFTs to analyze **Fig 4**'s decimated data for a wide range of input frequencies. You do have to remember the rules of a asing, however, to know where to expect the fund mental signal to show up in the FFT output spectrur You may think your FFT is sampling your signal at rate of f_S/N, but the converter is actually sampling a rate of f_S.

Once you understand how to use these various tecl niques, you can start to probe your particular converte to measure its real performance. Part 3 will discus how you apply these techniques to actually test a ADC in your system and determine a number of stat and dynamic specifications.

References

1. Kester, W A, "PCM Signal Codecs for Video Applications," *SMPTE Journal*, No. 88, November 1979, pg 770.
2. Bennett, W R, "Spectra of Quantized Signals," *Bell System Technical Journal*, No. 27, July 1948, pg 446.
3. Harris, Frederic J, "On the use of windows for harmoni analysis with the discrete Fourier transform," *IEEE Proceedings*, Vol 66, No. 1, January 1978, pg 51.
4. Coleman, Brendan, Pat Meehan, John Reidy, and Pat Weeks, "Coherent sampling helps when specifying DSP A/I converters," *EDN*, October 15, 1987, pg 145.
5. Ramirez, Robert W, *The FFT: Fundamentals and Con cepts*, Prentice Hall, Englewood Cliffs, NJ, 1985.
6. Peetz, B E, A S Muto, and J M Neil, "Measuring Wave form Recorder Performance," *HP Journal*, Vol 33, No. 11, November 1982, pg 21.
7. Frohring, B J, B E Peetz, M A Unkrich, and S C Bird "Waveform Recorder Design for Dynamic Performance," *HI Journal*, Vol 39, February 1988, pg 39.
8. "Dynamic performance testing of A to D converters," HP Product Note 5180A-2, 1982.

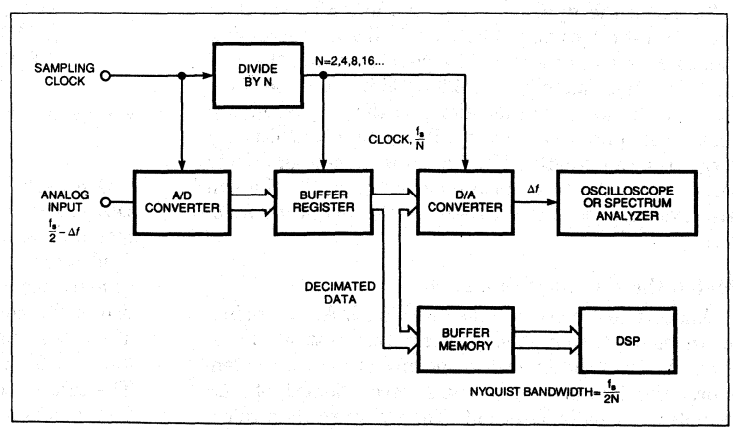

Fig 4—The beat-frequency test stresses the converter with a near-Nyquist input signal, *and it drives the converter at its maximum sampling rate. You can then examine the low-frequency beat on an oscilloscope and search for missing codes and other nonlinearities.*

ANALOG DEVICES

AN-215C
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Designer's Guide to Flash-ADC Testing

Part 3
Measure Flash-ADC Performance for Trouble-Free Operation

by Walt Kester

The first two parts of this series described the subtleties of flash A/D converters and the test methods used to evaluate these devices. Part 3 concludes the series with a discussion of the actual measurements you'll need to fully characterize flash A/D converters.

Although manufacturers have expanded the number of guaranteed specifications they put on their data sheets, the test conditions often won't match those of your system design. You can use the methods described in Part 2 of this series to test a flash A/D converter, but the measurements you need to perform depend on the converter's primary application. This final part of the series provides information on important measurements you'll need to characterize your converter's performance, including total harmonic distortion (THD), differential and integral nonlinearity, and noise power ratio. You'll probably want to start with the S/N ratio, a measurement that's common to most A/D converter applications.

The S/N ratio is the ratio of the rms fundamental to the rms quantization noise. As described in Part 2, you can measure this parameter by digitizing a pure sine wave and performing Fourier transformations on the data. The rms energy contained in the fundamental sine wave is equal to the square root of the sum of the squares of the peak value and the values of the appropriate number of samples, or bins, located on either side of the peak. The converter's resolution and its side-lobe roll-off characteristics determine the number of samples you'll need. For a detailed explanation of sampling requirements, see Part 2.

The rms energy in the remaining frequency bins represents the noise due to theoretical quantization, the converter's harmonic distortion and excess noise, and the FFT round-off error. Take the square root of the sum of the squares of the remaining samples (excluding the dc components) to determine the rms energy. The overall S/N ratio of the A/D converter is

S/N ratio = 20 log(rms signal level/rms noise level).

You can measure harmonic distortion in a similar manner. The test program (described in Part 2) examines the FFT frequency spectrum for the proper location of the desired harmonic (harmonics above $f_S/2$ will be aliased into the baseband) and determines the rms energy in that harmonic. The following equation calculates the harmonic distortion:

Harmonic distortion = 20 log(rms signal level/rms harmonic level).

Reprinted from EDN – February 1, 1990

The S/N ratio and harmonic distortion are key specifications in evaluating the performance of A/D converters.

The total harmonic distortion (THD) is the root-sum-square of the first five harmonics of the fundamental. Use this number in place of the rms harmonic level in the above formula.

Two-tone intermodulation tests using FFTs

In many applications, you don't have the simple case of a single input frequency. For example, in communication applications that multiplex several frequencies onto a single carrier, you need to measure intermodulation products. You determine this parameter by applying two sine waves of different frequencies (f_1 and f_2) to an A/D converter. You then measure the amplitudes of the third-order intermodulation products, which occur at frequencies $2f_1+f_2$, $2f_1-f_2$, $2f_2+f_1$, and $2f_2-f_1$.

Although it's possible to filter out most intermodulation distortion if the two tones are of similar frequencies, the third-order products will be very close to the fundamental frequencies and thus difficult to remove.

To avoid clipping-induced distortion, the amplitudes of the individual tones should be at least 6 dB below the full-scale range of the flash converter. In addition, the frequency separation of the two tones should be consistent with the resolution of the FFT. As discussed in Part 2, the spectral resolution of the FFT is a function of record length M, coherence vs noncoherence, and the properties of the windowing function that you choose.

In receiver applications, you often want to know the maximum ratio between the amplitude of a single-tone input signal and the amplitude of its maximum spurious component. For an ideal A/D converter, this ratio oc curs for a full-scale input sinusoid. In a practical A/I converter, however, spurious content is a function c slew rate. Therefore, the maximum spurious-free dy namic range for a given input frequency will probabl occur at a level somewhat below full scale. Becaus the spurious-free dynamic range is slew-rate depend ent, it's a function of input frequency and amplitude.

Fig 1 is a plot of the typical maximum spurious leve vs input signal level. Also shown is a plot of the corre sponding spurious-free dynamic range. The plot dem onstrates that the maximum spurious-free dynami range of 38 dB occurs for an input signal that's abou 3 dB below full scale.

The data you need to generate these plots is readil available from the family of FFTs calculated for th different input amplitudes. By knowing the input signa level that gives the highest spurious-free dynami range at frequencies close to the Nyquist frequency it's possible to set the gain of the system to take maxi mum advantage of the A/D converter's spectral charac teristics.

Histograms are helpful

Differential and integral nonlinearity are also important measurements of converter performance. Try a histogram test to obtain these measurements. To make a histogram analysis, digitize a known periodic input at a rate that's asynchronous relative to the input signal. To gather the sample data for the histogram, you'll need a buffer memory and a test system, as described

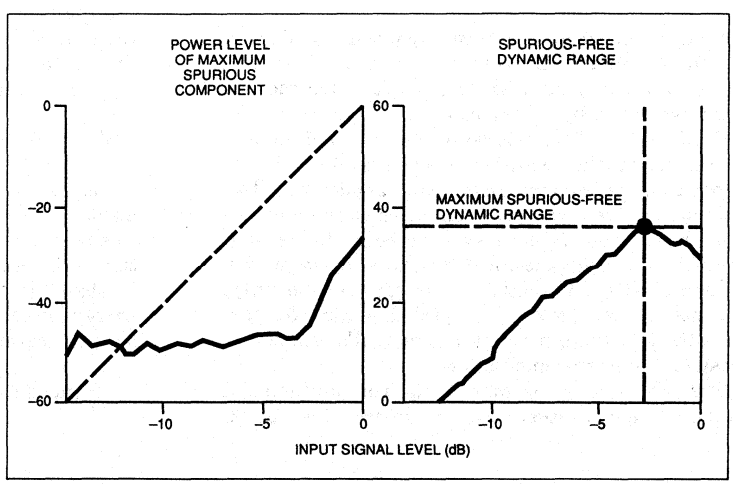

***Fig 1—These dynamic-range plots** show the power levels of spurious frequencies and the maximum spurious-free dynamic range. In this example, the maximum spurious-free dynamic range occurs at an input signal level that's 3 dB below full scale.*

n Part 2. The buffer memory will probably be too small to hold a statistically significant number of samples from a single run (several hundred thousand are usually required). For this reason, run several tests to acquire the data and load the contents of the buffer into the main memory of your test system after each run. Benchtop test systems from Hewlett Packard and Tektronix also provide histogram test capability.

After the test system accumulates a statistically significant number of samples, it can determine the relative number of occurrences of each digital code (the code density). This test routine then normalizes the data based upon the input signal and analyzes the results for linearity errors.

For an ideal A/D converter with a full-scale triangular-wave input, you'd expect an equal number of codes in each bin. The number of counts in the nth bin, H(n), divided by the total number of samples taken, M, is the bin width as a fraction of full scale. The ratio of the actual bin width to the ideal bin width, P(n), is the differential linearity. Ideally, this ratio should be unity. Subtracting 1 LSB gives you the differential nonlinearity.

You can determine integral nonlinearity with a cumulative histogram; the cumulative bin widths are the transition levels. However, the cumulative effects of errors can make the integral-nonlinearity measurement inaccurate. Histograms are used more often in evaluating differential nonlinearity.

High-speed, high-accuracy triangular waves are difficult to generate, so use a sine wave. All codes aren't equally probable with a sine-wave input, however, and you should normalize the histogram data using the probability density function for a sine wave, as shown in **Fig 2.**

To obtain accurate results, you need to take a large number of samples. For example, to determine the differential nonlinearity for an 8-bit flash converter to within 0.1 bit with 99-percent confidence, you'll need 268,000 samples. You can use hardware to count these samples, thus speeding up the software processing time. For high-speed sampling, decimate the output data to clock rates that are compatible with a slower-speed memory.

Using noise-power-ratio tests

You can use noise-power-ratio (NPR) tests to measure the transmission characteristics of frequency-division-multiplexed (FDM) communications links. In a typical FDM system, 4-kHz-wide voice channels are "stacked" in frequency for transmission over coaxial, microwave, or satellite equipment. At the receiving end, the FDM equipment demultiplexes the data and returns it to individual, 4-kHz baseband channels. In an FDM system that has 100 channels or more, Gaussian noise with the appropriate bandwidth approximates the FDM signal.

The test setup of **Fig 3** measures an individual 4-kHz channel for quietness by using a narrow-band notch (bandstop) filter and a tuned receiver (**Ref 4**), both of which measure the noise power inside this 4-kHz notch. The NPR measurements are straightforward. With the

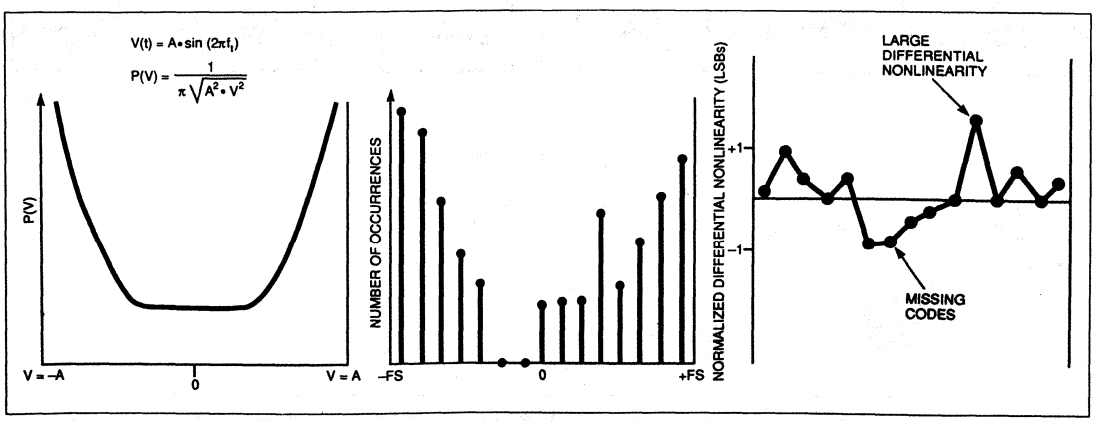

Fig 2—Histograms are often used to plot differential nonlinearity. *Shown here is a curve for the probability density function of a sine wave, which is used to normalize histogram data to produce a plot of differential nonlinearity.*

Where multiple frequencies exist on a single carrier, you need to measure intermodulation distortion as well as harmonic distortion.

notch filter out, the receiver determines the rms noise power of the signal inside the notch. The notch filter is then switched in, and the receiver determines the residual noise inside the 4-kHz slot. The ratio of the two readings, expressed in dB, is the NPR. You should test several slot frequencies across the noise bandwidth—low, midband, and high.

The NPR is usually plotted on an NPR curve as a function of rms noise level referred to the peak range of the system. For very low noise levels, the undesired noise is primarily thermal noise and is independent of the input noise level. Over this region of the curve, a 1-dB increase in the noise level causes a 1-dB increase in the NPR. As the noise level increases, the amplifiers in the system begin to overload, creating intermodulation products that cause the noise floor of the system to rise. As the input noise increases further, the effects of overload noise predominate, reducing the NPR dramatically. FDM systems are usually operated at a noise-loading level a few decibels below the point of maximum NPR.

In a digital system containing an A/D converter, the noise within the slot is primarily quantizing noise when low values of noise input signals are applied. The NPR curve is linear in this region. As the noise input level increases, the hard-limiting action of the converter causes clipping noise to dominate.

In a practical A/D converter, any dc or ac nonlinearities cause a departure from the theoretical NPR. Although the peak value of NPR occurs at a fairly low input noise level (rms noise = 1/4 V_O, where $\pm V_O$ is the range of the A/D converter), the broadband nature of the noise signal stresses the device, and the test provides a good indication of its dynamic performance.

Theoretically, NPR readings should be independent of any particular slot frequency. However, because of increased nonlinearities for the higher input frequencies, the NPR readings in the higher slots tend to be lower.

NPR testing using DSP techniques

Using FFT analysis techniques, you'll find NPR measurements a real challenge. Consider the case where the record length is 1024 and the sampling rate is 20 MHz. The FFT of 1024 contiguous time samples would place a spectral component every 19.53 kHz (20 MHz/1024). Because the notch-filter slot width is approximately 4 kHz, the probability of a spectral component falling within the notch is very low.

To achieve reasonable data stability in the FFT NPR analysis, a number of samples must fall within the notch. If ten samples are within the 4-kHz notch, then the resolution of the FFT would need to be 400 Hz, necessitating a record length of 50,000 for a sampling

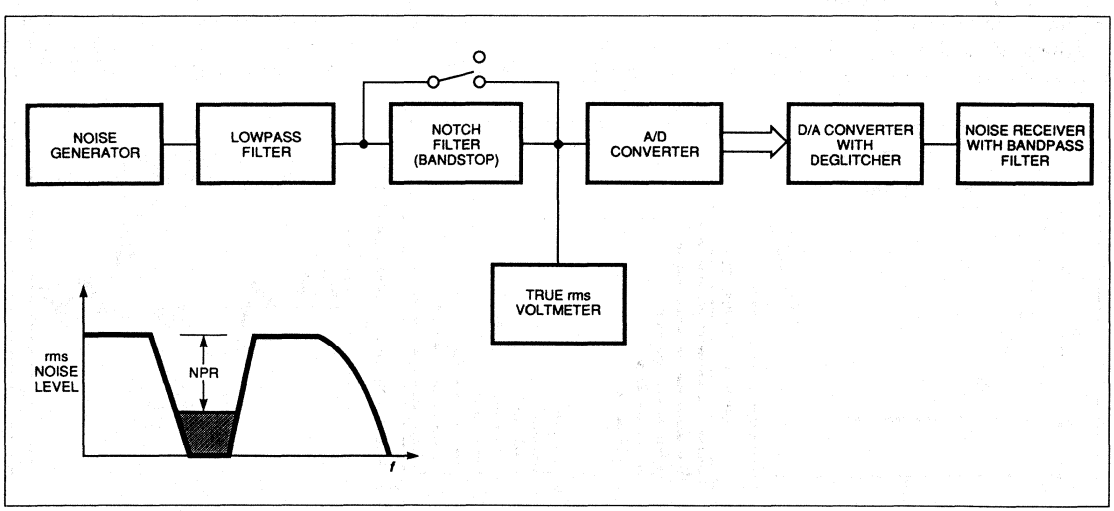

Fig 3—You can use this test setup to measure noise power ratio (NPR). *With the notch filter out, the receiver determines the noise power of the signal inside the notch. With the notch filter switched in, the receiver measures the residual noise inside the typical 4-kHz slot. The ratio of the two readings (in decibels) is the NPR.*

rate of 20 MHz. To avoid an extremely large buffer memory (and hence more demands on the FFT processor), you need to make the notch filter wider. For 20-MHz sampling and a 1024-word buffer memory, a notch filter that has a width of 200 kHz will provide ten frequency bins inside the notch. Even under these conditions, however, you should average the NPR calculations for several records to provide reasonable data stability.

Transient-response testing

The response of a flash converter to a transient input such as a square wave is often critical in radar applications. The major difficulty in implementing this test is obtaining a flat pulse that's commensurate with the converter's resolution.

A test setup for measuring the transient response of an A/D converter is shown in **Fig 4**. If you mount the Schottky-diode flat-pulse generator as close as possible to the analog input of the A/D converter, you can apply a signal to the A/D converter that's flat to at least 10-bit accuracy a few nanoseconds after it reverse biases the Schottky diodes.

You can use the same test setup to measure overvoltage recovery time. The amount of overvoltage is generally specified as a percentage of the A/D converter's range. For a converter with a 2V input range, 50% overvoltage corresponds to 1V above or below the nominal 2V input range. You make the starting point of the flat pulse correspond to the desired overvoltage condition. The actual recovery time is referenced to the time the input signal re-enters the A/D-converter input range. As in the transient-response test, you must consider the sampling (aperture) time delay when making this measurement.

The aperture-time and -jitter specifications of video A/D converters have probably been the least understood and most misused specifications in the entire field. The original concept of aperture time is centered around the classic S/H circuit of **Fig 5**. In an ideal S/H circuit, the switch has zero resistance when closed and opens instantly on receipt of an encode command. In practice, the sampling switch changes from a low to a high resistance over a certain finite time interval. An error occurs because the circuit tends to average the input signal over the finite time interval required to open the switch. As a result, the sampled voltage varies from the voltage at the instant the switch starts to open. The time required to open the switch is the aperture time. The error is determined by $E_a = t_a\ dV/dt$, where E_a is the aperture error, t_a is the aperture time, and dV/dt is the rate at which the input signal changes.

A simple first-order analysis, which neglects nonlinear effects, shows that no real error exists for such a switch. As long as the switch opens in a repeatable fashion, there is an effective sampling time that will cause an ideal S/H amplifier to produce the same hold voltage. The difference between this effective sampling point and the leading edge of the sampling clock is a fixed delay, which doesn't constitute an error. This effective aperture delay is the period from the leading edge of the sampling clock to the instant when the input signal equals the hold value. This specification

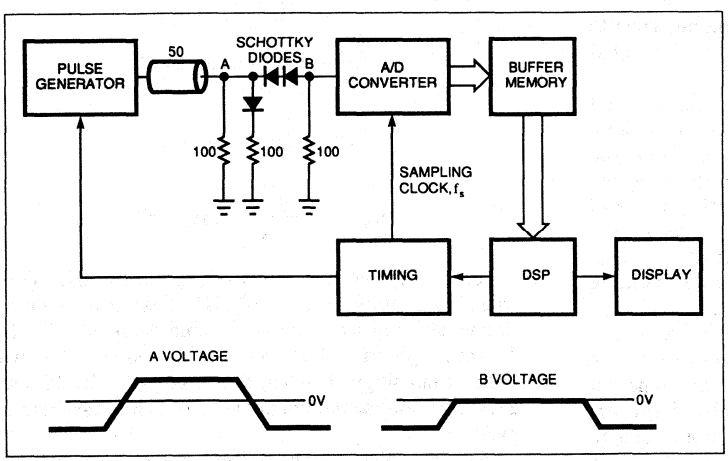

Fig 4—This test setup measures the transient response of an A/D converter. *The Schottky-diode network, located between points A and B in the circuit, generates a flat pulse for the input of the converter.*

In a practical A/D converter, the spurious-free dynamic range is a function of the converter's slew rate and can occur at a level below full scale.

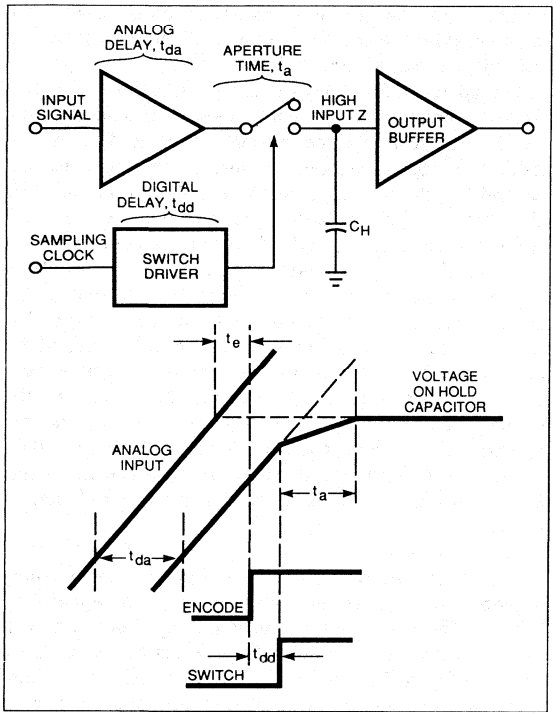

Fig 5—The concept of aperture time centers around the S/H circuit. *In practice, the sampling switch generates an error because of input-signal averaging over the finite time interval needed to open the switch. The aperture time is the time needed to open the switch.*

is important because it helps you determine when to apply the sampling clock with respect to the input signal timing.

The variation in effective aperture delay is important in simultaneous S/H applications. For example, in both I (in-phase) and Q (quadrature) radar receivers you may have to provide adjustable delays in the sampling clock to match the effective aperture delay times of several A/D converters. You should also consider delay-time tracking over a range of temperatures, especially in military systems where the specified operating temperature ranges from −55 to +125°C.

True aperture errors, however, do result from variable time delays. In a practical A/D converter, the sampling clock is often phase-modulated by some unwanted source; the source can be wideband random noise, power-line frequency, or digital noise due to poor grounding techniques. Phase jitter on the input sine wave can produce the same effect as jitter on the sampling clock. The resulting error is called aperture jitter. The corresponding rms voltage error caused by the rms aperture jitter qualifies as a valid aperture error.

The aperture-jitter specification is sometimes interpreted as a measure of the converter's ability to accurately digitize rapidly changing input signals. An A/D converter with an impressive aperture-jitter specification still may lose effective bits when digitizing a sine wave that has a maximum slew rate calculated from the aperture formula $E_a = t_a\ dV/dt$.

For example, assume that a 20-MHz, 8-bit flash converter has a bipolar input range of $\pm V_O$ ($2V_O$ p-p) and an aperture jitter specification of 20 psec rms. To calculate the maximum aperture-jitter error, convert the rms aperture jitter into a maximum value. If you consider that aperture jitter follows a Gaussian distribution similar to white noise, the rms aperture jitter, t_a, corresponds to the sigma (σ) of the distribution. The 2σ point on the distribution is a good place to set the maximum value, and the maximum aperture jitter becomes is $2t_a$.

If the corresponding maximum voltage error (ΔV) at the zero crossing of a full-scale sine wave is set to ½ LSB (½ LSB $= 2V_O/2^{N+1}$, where N equals the resolution of the A/D converter), then you can calculate the maximum full-scale sine-wave frequency, f_{max}, which will produce the ½ LSB aperture error, by using the following equations:

$$V(t) = V_O \cdot \sin(2\pi f_t),$$

$$\frac{dV}{dt} = 2\pi f V_O \cdot \cos(2\pi f_t),$$

$$\left.\frac{dV}{dt}\right|_{max} = \frac{\Delta V}{2t_a} = 2\pi V_O f_{max}, \text{ and}$$

$$f_{max} = \frac{\Delta V}{4\pi V_O t_a} = 2\pi t_a \cdot 2^{N+1}.$$

For $t_a = 20$ psec rms and N = 8, f_{max} is 16 MHz. These calculations imply that a 20-MHz flash converter can accurately digitize a full-scale sine wave of 16 MHz. In actual practice, however, the device may begin to suffer from skipped codes, decreased effective bits and S/N ratio, and ac nonlinearities at much lower frequencies.

You can calculate the effects of aperture jitter on

Histograms are useful in evaluating the differential nonlinearity of an A/D converter.

:he full-scale sine-wave S/N ratio as follows:

$$V(t) = V_0 \cdot \sin(2\pi f_t),$$

$$\frac{dV}{dt} = 2\pi f V_0 \cdot \cos(2\pi f_t), \text{ and}$$

$$\frac{dV}{dt}_{rms} = \frac{2\pi f V_0}{\sqrt{2}}.$$

For an rms error voltage, ΔV_{rms}, and an rms aperture jitter of t_a,

$$\frac{\Delta V_{rms}}{t_a} = \frac{2\pi f V_0}{\sqrt{2}}, \text{ and}$$

$$\Delta V_{rms} = \frac{2\pi f V_0 t_a}{\sqrt{2}}.$$

The rms-signal to rms-noise ratio, expressed in decibels, is

$$\text{S/N ratio} = 20 \log \left[\frac{V_0/\sqrt{2}}{\Delta V_{rms}}\right]$$

$$= 20 \log \left[\frac{1}{2\pi f t_a}\right] \text{dB}.$$

The S/N ratio that's due exclusively to aperture jitter in the above equation is plotted in **Fig 6** as a function of the full-scale input-sine-wave frequency for various values of aperture jitter.

Consider an 8-bit, 20-MHz A/D converter with an rms aperture jitter of 20 psec. For an 8-MHz full-scale input, the S/N ratio due only to aperture jitter is 60 dB, as calculated from the equation. The theoretical S/N ratio due to quantizing noise in an 8-bit flash converter is 50 dB. When you combine the S/N ratio of 60 dB with the S/N ratio of 50 dB, you obtain a theoretical S/N ratio of 49.6 dB, which encompasses both the ideal quantizing noise and the noise due to aperture jitter. A practical 8-bit device that has an rms aperture-jitter specification of 20 psec may, however, only achieve an S/N ratio of 40 dB under these conditions.

Therefore, to accurately evaluate the A/D converter's dynamic performance, you must carefully examine the S/N ratio, effective number of bits, and aperture-jitter specifications.

Try measuring the aperture jitter of an A/D converter using the test setup shown in **Fig 7**. The low-

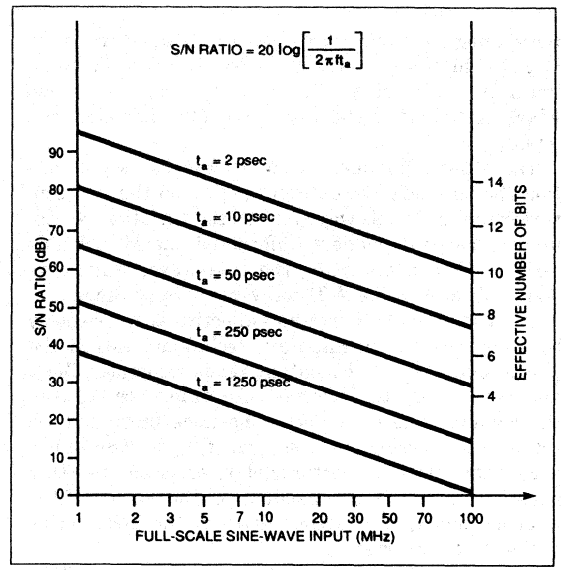

***Fig 6—This plot compares the S/N ratio** to the full-scale sine-wave input frequency for various values of aperture jitter.*

jitter pulse generator produces both the sampling clock and the analog input signal to minimize the phase jitter between them. Adjust the phase shifter until the A/D converter repetitively samples the sine wave at its point of maximum slew rate at midscale. Then take a histogram on the digitized A/D-converter output data.

An ideal A/D converter with no aperture jitter would have only one code present on the histogram. A practical converter gives a distribution of codes that you can fit to the normal distribution. The sigma (Σ) of the distribution corresponds to the rms error voltage, ΔV_{rms}, produced by the rms aperture jitter. Calculate the aperture jitter, t_a, from the formula

$$t_a = \frac{\Delta V_{rms}}{\frac{dV}{dt}}$$

where dV/dt is the rate-of-change of the sine wave at zero crossing.

If you sufficiently attenuate the input sine wave, any spreading of the distribution around the nominal code is due to intrinsic A/D-converter noise. As the input sine wave increases in amplitude, the slew rate, dV/dt, becomes proportionally greater, and the distri-

Noise-power-ratio tests are useful in determining the transmission characteristics of frequency-division-multiplexed communications links.

bution begins to spread because of the aperture jitter. Because high slew rates can affect the ac differential linearity of the converter, you should exercise caution when interpreting the histogram for high slew-rate inputs.

The offset adjustment shown in **Fig 7** lets you position the sine wave at different points on the A/D-converter range. In this way, you can see variations attributed to range-dependent differential-linearity characteristics. When offsetting the sine wave, make sure you don't exceed the A/D converter's input range.

It's also possible to measure effective aperture delay by using the locked-sine-wave technique. Adjust the phase shifter until the output reads midscale. Use a dual-trace scope to determine the difference between the leading edge of the sampling-clock pulse and the actual zero crossing of the sine-wave input. This difference is the effective aperture delay, which can be either negative or positive, depending on the values of the internal analog and digital delays in the S/H portion of the A/D converter.

At present, no industry standard exists for either the definition or the test for A/D-converter error rates. In flash converters, comparator metastable states can occur for low- or high-frequency input signals. At high frequencies, bubbles in the thermometer code of the comparator-bank output can also produce erroneous output codes.

Because error rates less than 1×10^{-16} are typical for well-behaved A/D converters, you need to take a large number of samples to properly measure the error rate. You must also take great care in the test-set layout, grounding, shielding, and power-supply decou pling so that 60-Hz, EMI, or RFI glitches don't creat erroneous errors.

Use the circuit in **Fig 8** to measure the error rat for low-frequency input signals. Apply a low-fre quency, full-scale sine wave (or triangle wave) to th A/D converter so that its rate of change is less than LSB/sample. This step ensures that the transitio zones between codes are all adequately exercised. A error amplitude of X LSBs is established as the lowe limit for the definition of a qualified error. Usually you select X to be several LSBs so that random nois doesn't produce errors. The software or hardware the examines the difference between each adjacent sampl and records the number of times this difference exceed the error threshold, X. If NQ is the number of qualifie errors that occur, and NT is the total number of sam ples taken, then the error rate, ER, is given by th equation $ER = NQ/2{\cdot}NT$.

As an example, consider an 8-bit, 100M-sample/se flash converter designed to take at least ten sample at each code level. For one slope of the triangle-wav input, the number of samples required i $10\times256=2560$ samples. The frequency of the triangl wave is

$$f_t = \frac{1}{2560 \cdot 2 \cdot 10\ \text{nsec}} = 19.5\ \text{kHz}.$$

At a 100-MHz sampling rate, the average time required to make an error for an error rate of 1×10^{-9} is 10 seconds.

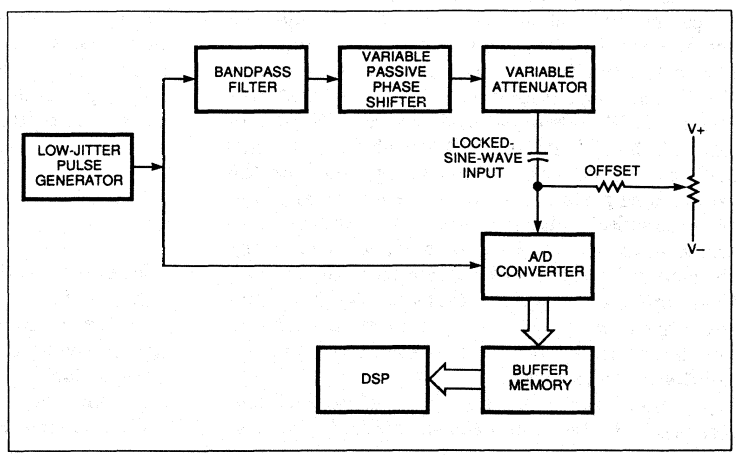

Fig 7—In this test setup for measuring aperture jitter, *you adjust the phase shifter until the A/D converter repetitively samples the sine wave at its point of maximum slew rate. You then take a histogram of the digitized A/D-converter output data. The offset adjustment lets you position the sine wave at different points on the converter's range.*

Aperture time and aperture jitter for A/D converters are probably the most misunderstood and misused specifications.

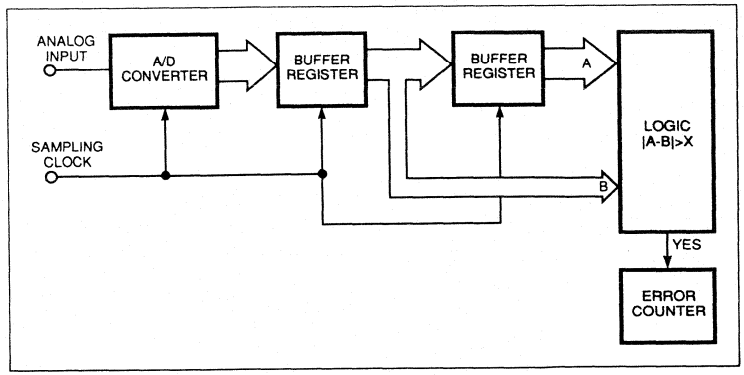

*Fig 8—**The effective aperture delay** is the time difference between the leading edge of the sampling-clock pulse and the actual zero crossing of the sine-wave input.*

In a similar manner, you can measure dynamic errors caused by fast input signals by using the beat-frequency approach. You choose the low-frequency beat frequency to give the proper number of samples per code level, and then you examine the decimated digital outputs for adjacent sample differences that exceed the allowable error amplitude.

In summary, determining appropriate error-rate criteria for an A/D converter depends upon both the application and the characteristics of the converter under consideration. Flash converters that use straight binary decoding with no additional correction logic are most subject to large metastable errors at midscale. For this situation, a low-amplitude dither signal centered on the midscale code transition might be an appropriate stimulus. In a more well-behaved flash converter, a full-scale signal that exercises all codes might be desirable.

If you plan to digitize composite video signals, you'll need to measure the differential-gain and -phase performance of the flash A/D converter. Differential gain is the percentage difference between the digitized amplitudes of two signals. Likewise, differential phase is the phase difference between the digitized values of the same two input signals. The input signals are typically a high-frequency low-level sine wave representing the color subcarrier frequency, superimposed on a low-frequency sine wave. Distortion-free processing of the color signal requires that the flash converter alters neither the amplitude nor the phase of the chrominance signal as a function of the luminance-signal level.

The best method for performing composite video tests is to use an A/D converter back-to-back with a D/A converter. Connect a TV test signal to the A/D converter and use the output of the D/A converter to drive a vectorscope. To ensure that the test accurately measures the A/D converter's performance, use a low-glitch D/A converter followed by a track-and-hold deglitcher. In addition, the dc accuracy of the D/A converter should exceed that of the A/D converter. When testing an 8-bit flash converter, use a D/A converter with at least 10 bits of accuracy.

References

1. Andrews, James R, Barry A Bell, Norris S Nahman, and Eugene E Baldwin, "Reference Waveform Flat Pulse Generator," *IEEE Transactions on Instrumentation and Measurement*, Vol IM-32, No. 1, March 1983, pg 27.

2. Schoenwetter, Howard K, "A Programmable Voltage Step Generator for Testing Waveform Recorders," *IEEE Transactions on Instrumentation and Measurement*, Vol IM-33, No. 3, September 1984, pg 196.

3. Gray, G A, and G W Zeoli, "Quantization and Saturation Noise Due to Analog-Digital Conversion," *IEEE Transactions on Aerospace and Electronic Systems*, January 1971, pg 222.

4. Tant, M J, *The White Noise Book*, Marconi Instruments, July 1974.

ANALOG DEVICES

AN-301
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Undersampling Techniques Simplify Digital Radio

by Richard Groshong and Stephen Ruscak

BY SAMPLING BELOW THE NYQUIST RATE WITH A NEW TYPE OF ADC, DESIGNERS CAN EXPLOIT THE BENEFITS OF DIGITAL RADIO.

Digital techniques offer some inherent advantages in communication-receiver design. These advantages, however, have been offset by both the cost and dynamic range limitations of analog-to-digital converters (ADCs). Consequently, designers continued to use analog circuitry despite its associated complexity and lack of flexibility. But recent introductions of low-cost, highly linear ADCs, combined with a technique known as undersampling, make digital processing the architecture of choice.

A review of receiver design shows how digital processing can improve performance. High-performance communication receivers invariably are superheterodyne designs. A typical example is a double-conversion unit that uses a broadband front-end and fixed-frequency bandpass filters after conversion to an intermediate frequency *(Fig. 1)*. The input filter ensures the selectivity needed to prevent strong signals outside of the receiver's normal tuning range from overloading the first mixer and generating spurious signals.

The first mixer upconverts the received signal to the first intermediate frequency (IF), which is high enough to simplify design of the broadband input filter but low enough to permit use of a low-cost crystal filter. The first IF is typically well above the highest receive frequency so the broadband input filter can attenuate undesirable signals. Otherwise, these signals can feed through the first mixer and prevent the receiver's IF rejection specification from being met. The input filter must also reduce the undesired first-mixer image response at

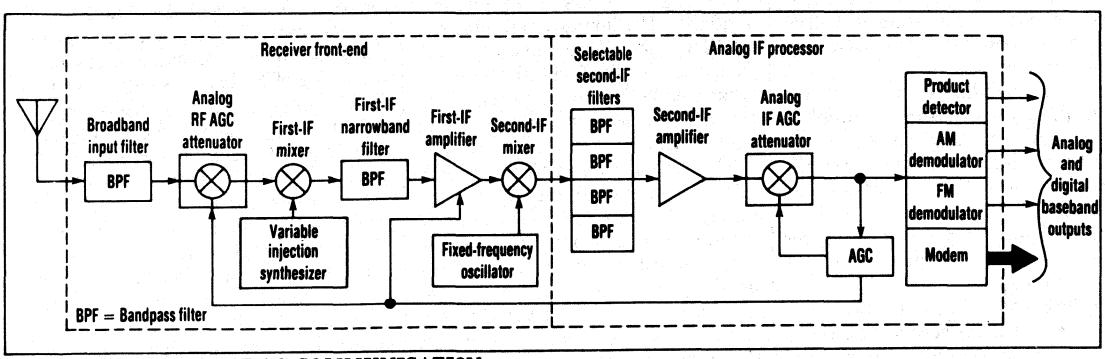

1. A TYPICAL ANALOG COMMUNICATION receiver may need several bandpass filters and demodulation circuits to support the bandwidth requirements and desired operating modes.

This article, the first of a two-part series, covers the fundamentals of digital radio and undersampling. The second part will describe a practical digital receiver, including the design of an analog front-end.

Reprinted from ELECTRONIC DESIGN – May 23, 1991

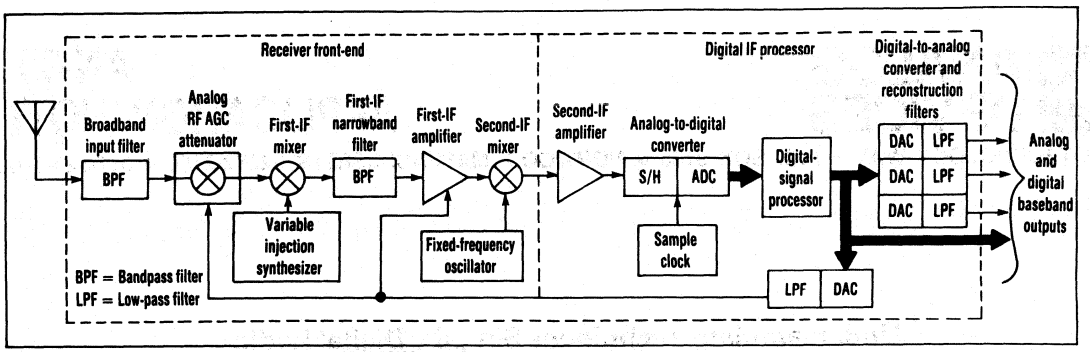

2. BY USING DIGITAL COMPONENTS (including a software-reconfigurable digital-signal processor) after the second IF amplifier, designers can simplify the receiver circuit and reduce costs.

frequencies twice the first IF away from the tuned frequency.

The first-IF bandpass filter protects the first-IF amplifier from overloading on signals that are within the receiver tuning range but outside of the desired receive-signal bandwidth. This filter also attenuates the second-mixer image response at twice the second-IF away from the desired signal frequency. The first-IF amplifier supplies just enough gain to allow a low receiver noise figure (for good sensitivity and a low noise floor) while maintaining a high dynamic range.

The second mixer downconverts the first IF to a frequency that permits the most receiver selectivity, gain, and signal processing at a reasonable cost. In the typical communication receiver, many second-IF bandpass filters may be needed to support the desired operating modes and bandwidth requirements. Additional frequency conversion to baseband may also be required, as well as multiple demodulation modes or embedded modems.

Analog-hardware costs, however, generally increase with the number of receiver operating modes. Consequently, using a software-reconfigurable digital-signal processor for the second-IF functions can reduce costs. Digital components can replace the second-IF and baseband functions *(Fig. 2)*.

Digital processing can also simplify the design of the receiver's analog front-end. For example, the digital-signal-processing algorithms can perform the fine tuning function, perhaps in 1-Hz steps over a 1-kHz range. This would allow the receiver's variable injection synthesizer to tune it's broadband input range in 1-kHz, rather than 1 Hz, steps. Also, the programmable digital filters can compensate the passband magnitude and phase responses of the analog front-end filter, so a lower cost filter can be used.

New Devices Help

Although the ADC in the second IF is a key component, it is only one element of a sampling system. Other elements include input conditioning amplifiers, a sample-and-hold amplifier, a voltage reference, digital interfacing circuitry, and clock conditioning circuitry. Older sampling systems were implemented in expensive, high-power hybrid devices. Recently, however, low-cost monolithic sampling systems have become available. For instance, an integrated ADC and sample-and-hold amplifier on one die form the core of what's called a "sampling ADC." In this article, the term ADC refers to such a system.

The sample rate of the second-IF ADC is an important factor. The digital-signal processor's efficiency is largely determined by the sample rate, regardless of the algorithms used by the processor. The higher the sample rate, the lower the efficiency. Likewise, nonlinearities in the sample-and-hold circuit will dominate the ADC's linear dynamic range, which consequently is affected by the second-IF. Higher IF frequencies result in lower dynamic range.

The digital-signal-processor designer, therefore, wants to use the lowest IF and sample rate possible to keep the receiver dynamic range and digital-signal-processor efficiency high. On the other hand, the analog front-end designer wants a high IF and sample rate to control image responses and aliasing in the ADC without a costly first-IF filter. Both designers can be satisfied by using undersampling techniques and sampling ADCs such as the AD679 and AD779, which have a built-in high-performance sample-and-hold circuit that is characterized above Nyquist frequencies.

Several operating considerations affect digital receiver design. For example, high frequency (HF) communication in the 2-to-30-MHz band generally consists of long-distance voice transmissions with bandwidths on the order of 3 kHz, in a dense interference environment. Undesired signals from anywhere in the world can interfere with small desired signals at the receiver. Spectrum crowding and the possibility of high-power transmitters close to the receiver make high dynamic range a must.

Specifically, a typical HF receiver must be sensitive enough to receive a 0.5-μV (–113 dBm) signal at the receiver input (50 Ω) with a 10-dB signal-to-noise ratio (SNR) in a 3-kHz

bandwidth. Because this bandwidth is small compared to the input sample rate required to digitize the IF signal, the ADC quantizing noise is not normally a limitation. However, the receiver must also be able to receive a 1-V signal (+13 dBm) without significant distortion. Consequently, in-band dynamic range must be at least 126 dB.

Unfortunately, a digital receiver's dynamic range is adversely affected by the fact that the bulk of the receiver gain must be ahead of the ADC and the second-IF bandpass digital filters. The high gain is needed to drive the ADC at a sufficient level. In an analog design, on the other hand, the gain occurs after the second-IF bandpass filters *(Fig. 1, again)*.

In the digital design, this positioning means that undesired out-of-band signals can pass through the wider-bandwidth receiver front-end and saturate the amplifier and ADC. This must be avoided because ADC saturation generates many aliased in-band harmonic and intermodulation products that can block a weak signal. The solution lies in automatic gain control (AGC) and the use of large-wordlength, highly linear ADCs.

The AGC performs the same overall function in a digital receiver as in an analog receiver. But in a digital receiver, more AGC must occur at the front-end to prevent overloading of the amplifier and ADC. The AGC signal is usually derived from the narrow-band IF filter output, which is implemented by the digital-signal processor. The designer must supply both analog and digital gain control, because the ADC can't possibly support the required 126-dB of dynamic range.

A system-level analysis determines the desired receiver gain distribution, or analog versus digital gain. The maximum analog gain must be high enough to prevent the ADC's quantizing noise from degrading the receiver noise figure or sensitivity below their desired limits at small input levels. As the signal increases above the sensitivity level, the AGC should not decrease the gain until an adequate SNR is obtained. Above that level, the analog AGC should hold the signal level at the ADC essentially constant.

Headroom Needed

In addition, the gain control must allow enough "headroom" at the ADC to avoid saturation during normally high peak-to-average voltage ratio periods of the desired signal. For normal speech signals, this headroom may have to be at least 20 dB above the average level.

When strong signals fall outside the digital filter's narrow bandwidth but inside the analog first-IF filter's bandwidth, the receiver must be able to detect any ADC overload. Then the receiver gain must be redistributed by reducing the analog gain and increasing the digital gain to maintain a constant output signal level. This action, however, will reduce the signal-to-quantization-noise ratio.

As noted earlier, a low sampling rate is a key factor in the overall performance of a digital receiver. A cursory analysis would conclude that the ADC's sampling rate must be at least twice the IF frequency. However, this is not true when sampling a bandpass signal.

In the time domain, sampling can be thought of as multiplying the analog input signal by an impulse train that has the same period as the sample clock *(Fig. 3a)*. In the frequency domain, the equivalent process is convolution of the analog signal spectrum with the impulse train spectrum. The result of the convolution is a set of images of the original spectrum at integer multiples of the sampling frequency, including dc *(Fig. 3b)*.

If the sampling rate decreases as the spectrum bandwidth remains constant, the adjacent images get closer together until they begin to overlap. At this point the signal spectrum is distorted, corrupting the information in the signal. This problem is called aliasing distortion and explains the traditional interpretation of the Nyquist criterion, which requires a sampling rate of more than twice the highest frequency of the signal being sampled.

In the case of a bandpass signal, however, designers can use undersampling, which means sampling at a rate less than twice the highest frequency being converted. To understand why undersampling (also

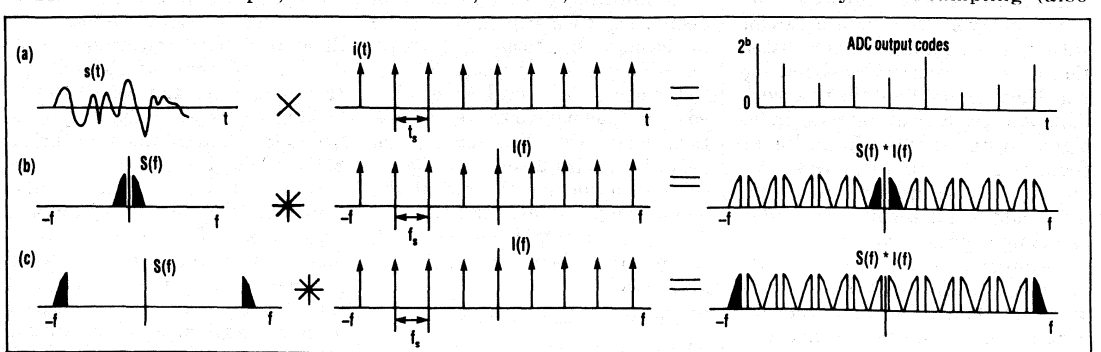

3. IN THE TIME DOMAIN, sampling is the multiplication of the analog signal and an impulse train with the same period as the sample clock (a). The equivalent in the frequency domain is convolution (b). If the analog signal is bandlimited to be less than one-half the sampling rate and doesn't cross a multiple of the sample clock, the converter can undersample the signal without causing aliasing (c).

called harmonic or bandpass sampling) is an ideal technique in radio applications, designers must study the spectrum of the typical radio receiver IF signal.

The IF signal, which is a bandlimited high-frequency signal, is applied to the input of the ADC, resulting in a frequency-domain convolution. Although the spectrum of the IF signal begins at a considerably higher frequency than that of a baseband signal containing the same information, the resulting convolution is theoretically identical *(Fig. 3c)*.

Now it should be evident that undersampling must also satisfy Nyquist's theorem. However, the sampling rate must be at least twice the bandwidth of interest, not twice the highest absolute frequency. Obviously, this makes a big difference in the case of an IF signal.

Proper undersampling presents some other restrictions. For one thing, the signal bandpass spectrum should not cross an integer multiple of one-half the sample rate. If it does, the convolution will create overlapping images and alias distortion. Also, designers must select an ADC with input circuitry that supports the highest signal frequency, otherwise the ideal impulse sampling model will not hold.

Another important factor is that ideal undersampling assumes that the sampling system behaves identically for both high-frequency and low-frequency inputs. At high frequencies, however, converters suffer from errors that manifest themselves as noise and harmonic distortion in the recovered baseband signal. These errors limit the receiver's ability to recover an undersampled signal in uncorrupted form. Imperfections that were insignificant for sub-Nyquist-frequency inputs require additional scrutiny in undersampling applications.

As a result, designers must examine ADC specifications that are very different than those that were important for conventional conversion applications. Traditional performance measures—such as integral nonlinearity (INL) and differential nonlinearity (DNL) are characterized with tests using dc inputs. But INL and DNL are only partly responsible for the nonlinearities that arise in the presence of ac signals.

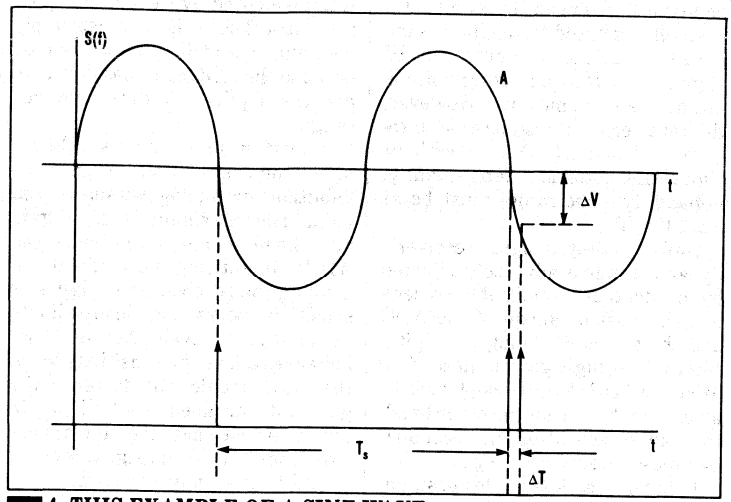

4. THIS EXAMPLE OF A SINE WAVE sampled by a clock of period T_s shows how the effects of aperture jitter increase with the input's slew rate. For a given jitter (Δt), the voltage error (ΔV) increases as the slew rate increases.

ADDITIONAL FACTORS

Radio designers must also be concerned with the converter's effect on the signal spectrum. This type of ac characterization specifies total harmonic distortion (THD), signal-to-noise-plus-distortion [S/(N+D)], intermodulation distortion (IMD), aperture jitter, and full-power bandwidth. Additionally, amplifier slew-rate limiting, phase jitter, and decreasing open-loop gain become increasingly significant in the presence of high-frequency signals.

The parameters may be interrelated. For example, converter resolution, which sets the quantization noise floor, also determines the receiver's maximum theoretical dynamic range and affects SNR. Assuming a full-scale sine wave signal, the maximum theoretical SNR of an ideal ADC is[1]

$$SNR(dB) = 6.02b + 1.76\ dB$$

where b = the number of bits.

This equation assumes that the measurement bandwidth spans the entire Nyquist bandwidth of one-half the sample rate ($f_s/2$). If the bandwidth is reduced after conversion to less than the full Nyquist-criteria bandwidth, the quantization noise power in this spectrum will be lower, resulting in a higher SNR. This condition is reflected in a generalized SNR equation:

$$SNR\ (dB) = 6.02b + 1.76\ dB + 10\log_{10}(f_s/2BW)$$

where f_s = the sample rate in Hz and BW = the information bandwidth in Hz.

For example, if f_s = 96 kHz, and the signal spans the entire Nyquist bandwidth (48 kHz), the maximum SNR for a 14-bit converter is 86.04 dB. But in the case of a bandlimited IF signal, which may span a 16-kHz bandwidth, the generalized equation predicts a maximum theoretical SNR of 90.81 dB.

Unfortunately, other noise sources, such as aperture jitter, also degrade an ADC's SNR. Aperture jitter or uncertainty is the variation associated with the exact moment that a sample-and-hold amplifier goes into hold mode. Both external (somewhat correctable) and internal (uncorrectable) sources cause aperture jitter. External jitter results primarily from the designer's inability

to deliver a jitterless sampling clock to the ADC. Internal jitter occurs when the sample-and-hold amplifier's sampling switch fails to open at precisely timed intervals.

A time-domain analysis shows that the effects of aperture jitter increase as the input signal's slew rate increases. An example is the case of a sample clock locked together with waveform A *(Fig. 4)*. In this example, if the sample clock has no variation (phase jitter), the output code from the ADC will always be zero. On the other hand, if the sample clock varies by an amount Δt, the ADC output will be the code that represents the error voltage, ΔV, whose amplitude is

$$\Delta V = (dV/dt) \times \Delta t$$

where dV/dt = the slew rate of the input signal in volts/s, at the sampling instant and Δt = the aperture uncertainty in seconds.

If the rms error voltage is well below the ADC's quantization noise, the resultant noise will be no worse than the quantization noise. As the error voltage increases, the converter's SNR, which previously was limited by quantization noise, will be dominated by noise related to aperture jitter. This noise results from phase modulation of the signal being sampled. If the aperture jitter is totally random, or "white," the noise produced by the phase-noise modulation will also be white. Regardless of the characteristics of the jitter-induced noise, aperture jitter will increase the overall noise floor of the ADC output spectrum.

Predicting Noise Levels

Designers can predict the theoretical limits of rms-signal-to-rms-aperture-jitter noise ratio (SN_jR) based on a full-scale sine wave input and a measurement bandwidth of $f_s/2$. The equation is[2]

$$SN_jR\,(dB) = 20\log_{10}(1/2\pi t_a f_a)$$

where t_a = the rms aperture uncertainty in seconds and f_a = signal frequency in Hz.

A graph of this equation shows the maximum SN_jR that designers can achieve based solely on the noise introduced by aperture jitter *(Fig. 5)*. Also shown is the maximum SNR that can be achieved if only quantization noise is taken into account. For a 14-bit converter with t_a = 150 ps, performance begins to degrade at about 53 kHz for a full-scale input.

When evaluating a converter's aperture-jitter performance, designers must consider the input signal's slew rate. Reducing the signal frequency or amplitude will result in a 1:1 reduction in jitter-related errors. That is, reducing the signal by a factor of 10 increases the undegraded bandwidth by 10 times.

The ADC's output contains undesired frequency components that were not present at the input. These components are at integer multiples of the input frequency. Consequently, designers should also evaluate harmonically related errors when trying to maximize dynamic range. THD, which describes the converter's linearity, is based on a ratio of the sum of the rms voltages of undesired frequency components to the rms voltage of the input. THD can be calculated with the equation

$$THD\,(dB) = 20\log_{10}[(h_1^2 + h_2^2 + \ldots h_n^2)/h_0]^{1/2}$$

where h_0 = the rms voltage of the fundamental and h_1 through h_n = the rms voltages of the integer harmonics of the fundamental.

At low frequencies, harmonic distortion generally results from unequal step sizes in the ADC transfer function (INL and DNL). At high frequencies, slew-rate limiting, signal feedthrough, nonlinear resistors and capacitors, and decreased open-loop gain in the sample-and-hold amplifier are the culprits. Like aperture-jitter-induced noise, THD increases when either amplitude or frequency increases.

Another important measure of linearity is IMD, which results when two tones applied simultaneously to an ADC are "mixed" as a result of nonlinearities. Like THD, at low frequencies IMD is a function of the converter's INL and DNL.

However, the ADC doesn't behave like a typical saturated amplifier at low frequencies. In the amplifier, an increase in amplitude causes a corresponding threefold increase in the third-order IMD products. But the ADC's transfer curve can be considered linear with small random quantizing-level errors distributed throughout the converter's full range. Therefore, increasing the in-

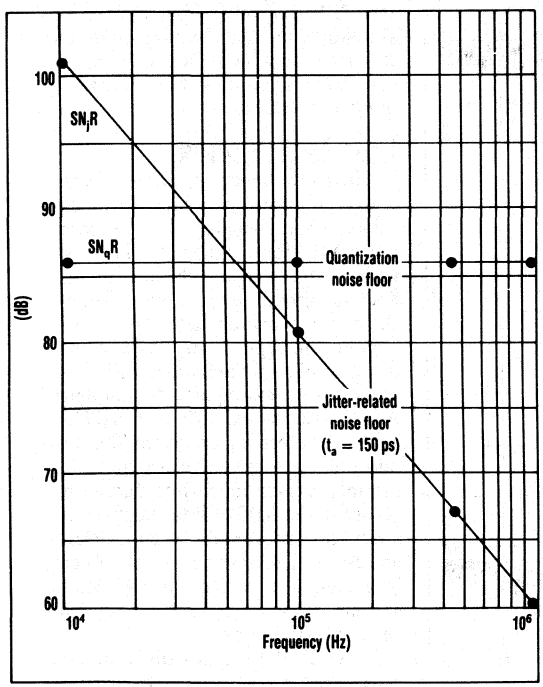

5. DESIGNERS CAN PREDICT the theoretical limits of the rms-signal-to-rms-aperture-jitter noise ratio, SN_jR. For an rms aperture uncertainty, t_a, of 150 ps, performance begins to degrade above 53 kHz, approximately. Also shown on this plot is the maximum SNR that can be achieved if only quantization noise is considered.

put signal's amplitude doesn't increase third-order IMD.

Instead, the distortion is fixed relative to the converter's full-scale range and, thus, independent the input amplitude. But this is true only if the third-order intercept of the ADC input amplifier is high enough that amplifier IMD isn't the dominant distortion source.

At high frequencies and high amplitudes, IMD resulting from INL and DNL is overshadowed by the effects of sample-and-hold amplifier errors. Typically, the amplifier cannot track large, high-frequency signals fast enough to maintain linearity, and IMD increases, as it does in any amplifier.

Another specification peculiar to undersampling converters is full-power bandwidth (FPB). Although no standard definition exists, FPB is often defined as the frequency at which the amplitude of the reconstructed digital output resulting from a full-scale input signal is reduced by 3 dB. A wide FPB ensures a flat output when the ADC undersamples high-frequency inputs. It's worth noting that the bandwidth increases at lower inputs.

Designers, however, should examine the tradeoffs between frequency response and SNR. For instance, although the converter's FPB may be well above the Nyquist rate, other specifications—such as S/(N+D) and THD—may vary as input frequencies approach the FPB. Also, smaller inputs reduce the SNR, but the resulting bandwidth increase reduces slew-rate-induced distortion and aperture-jitter-induced noise. So the FPB quoted in the data sheet shouldn't be the bottom line, although it's a good indication that the converter is intended for undersampling applications.

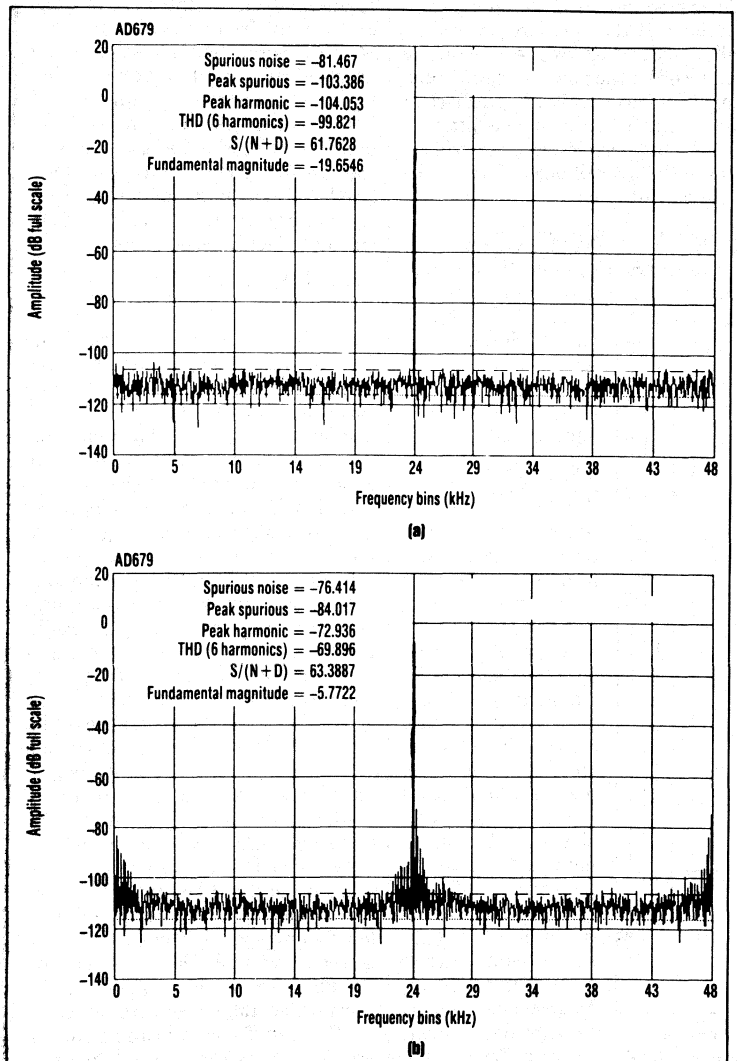

6. LOOKING AT A CONVERTER'S OUTPUT using fast Fourier transforms shows how noise and linearity specifications are affected by the magnitude of the input signal. These examples show the difference for inputs 19.6546 dB below full scale (a) and 5.7722 dB below full scale (b).

COST IS A FACTOR

Although linearity, sample rate, and resolution are prime considerations when choosing an undersampling ADC for digital radio, low cost and convenience are also important. And designers must also examine the device's performance for inputs beyond the typical sub-Nyquist frequency characterization. For instance, two examples of sampling ADCs are the AD679 and AD779. These converters are complete 14-bit, 100-ksample/s devices that differ only in their digital interfaces. Included in the AD679 data sheet are ac specifications for THD, FPB, IMD and S/(N+D). Graphs are specifically included showing the AD679's characteristics for signals beyond the Nyquist frequency.

S/(N+D) determines the dynamic range that the converter can realistically achieve. This figure, which is computed by analyzing the ADC's output using fast Fourier transforms (FFTs), is the rms summation of quantization noise, jitter-induced noise, circuit noise, and THD. The

2048-point FFTs for 456-kHz inputs that are 19.6546 and 5.7722 dB below the AD679's full-scale input range illustrate the dynamic range calculation *(Fig. 6)*.

As expected, the fundamental is at 24 kHz as a result of the undersampling convolution process. A number of specifications can be calculated directly from the FFT output data. For example, for an input 19.6546 dB below full-scale, the peak spurious noise is –103.386 dB relative to the converter's 10-V pk-pk full-scale range, and the integration of all the spurious noise in the Nyquist bandwidth is –81.467 dB *(Fig. 6a)*. The peak harmonic is –104.053 dB. Based on the first six harmonics, the THD is –99.821 dB, which is essentially buried in the noise floor.

Calculated Noise Floor

At 456 kHz, the theoretical SN_jR for a full-scale input is approximately 67.3 dB *(Fig. 5, again)*. Therefore, for a –20-dB input, the noise resulting from jitter is –87.3 dB. Adding the rms value of this noise to the ideal quantizing noise level of –86 dB results in an theoretical noise floor of –83.5 dB, which is relatively close to the measured level of –81.467 dB. The difference is due to external circuit noise and measurement error. This calculation shows that for inputs 20 dB below full-scale the ADC's dynamic range is limited by quantization noise rather than device nonlinearities.

A larger, high-slew-rate signal changes things. For an input 5.7722 dB below full-scale, the peak spurious signal is now –84.017 dB relative to full-scale, and all the spurious noise in the Nyquist bandwidth has increased by 5 dB to –76.414 dB *(Fig. 6b)*. THD has also increased but is still a respectable –69.896 dB. Consequently, for large high-frequency signals THD, not quantization noise, limits dynamic range.

Note that the AD679's wide input bandwidth not only allows the user to sample a high IF frequency, but also presents the disadvantage of sampling noise and other undesired signals over the same large bandwidth. Like the desired IF information band, the high-frequency noise is convolved with the sampling clock spectrum and can be aliased back to baseband, corrupting the signal. This is true for any undersampling ADC, unlike conventional converters, whose limited bandwidth filters out most of this noise. As a result, an ADC intended for undersampling is far more sensitive to the amount of noise at the input than a traditional converter, so board layout and decoupling considerations are even more critical.

Finally, convenience should not be overlooked in the choice of an ADC. The AD679 and AD779 are the only 14-bit monolithic ADCs that are suitable for undersampling. Both contain an integrated sample-and-hold amplifier, voltage reference, digital latches, and control logic, and have high-impedance analog inputs to simplify input buffering. The AD679 has a convenient 8-bit microprocessor interface and the AD779 has a parallel output that easily connects to digital-signal processors. The AD679 is also available in a pin-compatible 12-bit version, the 200-ksample/s AD678, which sacrifices dynamic range for speed. The AD678 is also specified at frequencies beyond Nyquist.

References:
[1]"Analog-digital conversion handbook," Analog Devices Inc., Wilmington, Mass., 1986, p. 542.
[2]"High-speed design seminar," Analog Devices Inc., Wilmington, Mass., 1989, pp. 1-86, 87.

ANALOG DEVICES

AN-302
APPLICATION NOTE

)NE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Exploit Digital Advantages in an SSB Receiver

by Richard Groshong and Stephen Ruscak

IMPROVEMENTS IN DIGITAL CONVERSION TECHNOLOGY HELP RADIO DESIGNERS TAKE ADVANTAGE OF DIGITAL ARCHITECTURES.

This article is the second of a two-part series on the design of digital radio receivers. The first part (ELECTRONIC DESIGN, May 23, p. 67) covered the fundamentals of digital radio and undersampling techniques, as well as selection criteria for the analog-to-digital converter. Part 2 covers the receiver design in greater detail.

The first step in digital-receiver design is to determine the radio's front-end specifications. The process is somewhat similar to that in a conventional analog design, but it must account for the characteristics of the digital IF processor used by the radio. After the front-end is completed, the designer can proceed to the IF processor, taking advantage of highly integrated devices to simplify the design.

The design example in this article is a 2-to-30-MHz SSB receiver employing undersampling techniques. A 16-kHz front-end bandwidth was chosen to accommodate up to four 3-kHz independent-sideband channels. This bandwidth also allows the digital-signal processor to perform fine tuning over at least a 1-kHz range, simplifying the variable-injection synthesizer *(Fig. 1)*.

The first IF is 100 MHz, which is high enough to ensure that the broadband input filter suppresses the IF and image response by 80 dB, but low enough to allow a low-cost crystal filter to be used as the first-IF narrowband filter. Because the DSP performs the fine tuning in 1-Hz steps, the variable injection synthesizer can tune from 102 to 130 MHz in 1-kHz steps. This "high-side" first-IF mixer injection causes a passband reversal or "flip" in the mixer, which will be corrected by the sampling process.

The first-IF amplifier makes up for losses in the first-IF mixer and filter, and maintains a front-end noise figure of 15 dB. After amplification, the signal is mixed with a fixed-oscillator signal of 99.544 MHz. The result is a second-IF of 456 kHz, which is high enough to allow the second-mixer image response at 99.088 MHz to be attenuated by 80 dB in the first-IF filter, but low enough to be sampled by the analog-

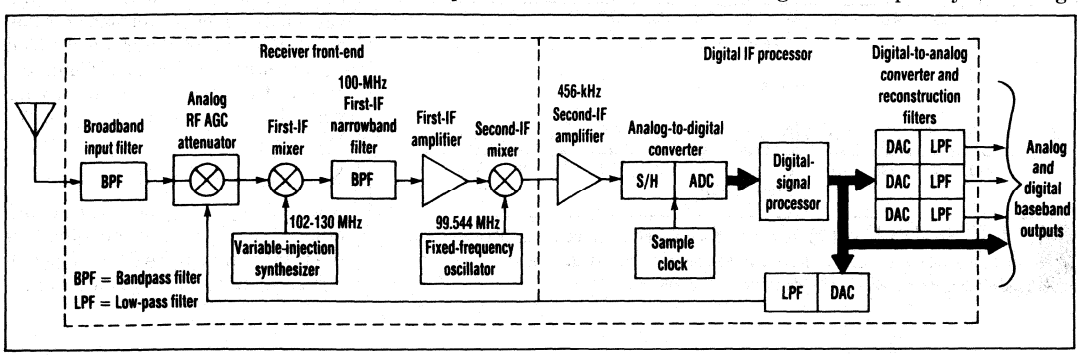

1. THE CHOICE OF 100 MHz for the first IF is high enough to ensure that the broadband input filter sufficiently suppresses the IF and image response but low enough to allow a low-cost crystal filter for the first-IF narrowband filter.

Reprinted from ELECTRONIC DESIGN – June 13, 1991

to-digital converter (ADC).

If the ADC's sample-and-hold bandwidth is greater than 456 kHz, as is the case with the AD779, the receiver can undersample the second-IF signal efficiently at a sample rate of 96 kHz. Because the analog-to-digital conversion is effectively the result of a convolution of the input signal and the sample-clock impulse spectrum, the sample clock must be as "clean" and jitter-free as the local oscillator or mixer injection signal. Any phase noise present on the clock will be impressed on the signal being sampled, effectively corrupting it. If isolation or buffering is inadequate, the spurious noise picked up by the sample clock can be disastrous.

When the second-IF spectrum is convolved with the sample-clock impulse spectrum, the IF signal is frequency translated to a lower IF of one-fourth the sample frequency, or 24 kHz *(Fig. 2)*. The convolution also creates a passband reversal that counteracts the reversal in the first-IF mixer. As expected, the sampled-IF spectrum repeats at every multiple of the sample frequency.

In this example, increasing the second-IF bandwidth will create aliasing distortion or overlap in the sampled IF spectrum because real-world filters can supply only finite stopband attenuation *(Fig. 2, again)*. But aliasing can be reduced to an acceptable level. In this design, the passband width of 16 kHz must be alias-protected to 60 dB.

This requirement, and a look at the IF sampling process, lead the designer to the bandpass specifications of the first-IF filter *(Fig. 3)*. Because of the sampling frequency and second-IF selected, the filter must have an 80-kHz-wide stopband. Given the 16-kHz-wide passband, the corresponding shape factor is 5 to 1. Attenuation of signals 40 kHz from the passband center of 100 MHz must be at least 60 dB. As noted, another 20 dB of attenuation is also needed at the image-response frequency of 99.088 MHz. A typical passband ripple specification might be 1 dB. A 4- or 5-pole crystal filter might satisfy these criteria.

SPREADSHEET OUTPUT

Receiver performance for thermal-noise-limited example

Digital-signal-processor receiver gain distribution and noise performance

Analog RF/IF noise figure	15 dB
ADC resolution	14 bits
ADC full-scale level	5 V peak (23.98 dBm)
ADC quantization level	−60.31 dBm
ADC sample frequency	96 kHz
ADC input bandwidth (BW1)	16 kHz
Information bandwidth (BW2)	3 kHz
Analog AGC threshold	−77 dBm

Antenna signal level (dBm)	Analog RF/IF gain (dB)	ADC signal level (dBm)	Noise at ADC input in BW1 (dBm)	Noise at ADC input in BW2 (dBm)	Quantizing noise of ADC in BW2 (dBm)	Total noise in BW2 (dBm)	Output SNR in BW2 (dB)	Digital processor gain (dB)	Output signal level (dBm)	Antenna overload level (dBm)
−113.00	57.00	−56.00	−59.96	−67.23	−74.11	−66.42	10.42	60.00	4.00	−33.02
−103.00	57.00	−46.00	−59.96	−67.23	−74.11	−66.42	20.42	50.00	4.00	−33.02
−93.00	57.00	−36.00	−59.96	−67.23	−74.11	−66.42	30.42	40.00	4.00	−33.02
−83.00	57.00	−26.00	−59.96	−67.23	−74.11	−66.42	40.42	30.00	4.00	−33.02
−73.00	53.00	−20.00	−63.96	−71.23	−74.11	−69.42	49.42	24.00	4.00	−29.02
−63.00	43.00	−20.00	−73.96	−81.23	−74.11	−73.34	53.34	24.00	4.00	−19.02
−53.00	33.00	−20.00	−83.96	−91.23	−74.11	−74.03	54.03	24.00	4.00	−9.02
−43.00	23.00	−20.00	−93.96	−101.23	−74.11	−74.10	54.10	24.00	4.00	0.98
−33.00	13.00	−20.00	−103.96	−111.23	−74.11	−74.11	54.11	24.00	4.00	10.98
−23.00	3.00	−20.00	−113.96	−121.23	−74.11	−74.11	54.11	24.00	4.00	20.98
−13.00	−7.00	−20.00	−123.96	−131.23	−74.11	−74.11	54.11	24.00	4.00	30.98
−3.00	−17.00	−20.00	−133.96	−141.23	−74.11	−74.11	54.11	24.00	4.00	40.98
7.00	−27.00	−20.00	−143.96	−151.23	−74.11	−74.11	54.11	24.00	4.00	50.98
17.00	−37.00	−20.00	−153.96	−161.23	−74.11	−74.11	54.11	24.00	4.00	60.98

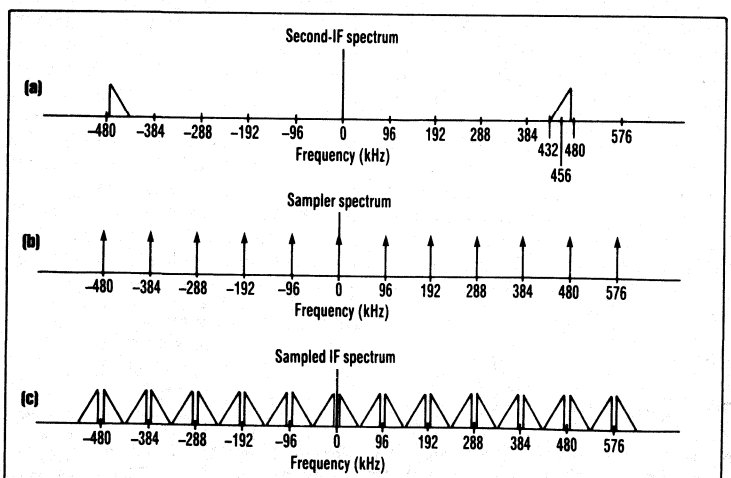

2. THE SAMPLING PROCESS IS a convolution of the second-IF spectrum (a) with the sampler spectrum (b). The resulting spectrum shows that if the second-IF bandwidth is increased, aliasing distortion will occur (c).

THE DIGITAL STAGE

The signal from the second-IF mixer must be amplified to a level sufficient to drive the ADC. The gain required depends on the ADC and various receiver performance trade-offs, and must be determined based on a system-level analysis. For example, the theoretical quantizing noise density of an ADC, N_{0q}, in W/Hz is:

$$N_{0q} = (V_{pk\text{-}pk}/2^b)^2/6f_sR$$

where $V_{pk\text{-}pk}$ = the peak-to-peak voltage range of the ADC in volts, b = the total number of ADC bits including sign, R = the assumed imped-

ance in ohms, and f_s = the sampling frequency in hertz.

For a 14-bit ADC with a 10-Vpk-pk range and a sampling frequency of 96 kHz, $N_{0q} = 1.2935 \times 10^{-14}$ W/Hz or –108.9 dBm/Hz, assuming 50-Ω impedance normalization. Other noise includes the thermally generated receiver noise. The level of this noise is the front-end noise figure, 15 dBm/Hz, minus the thermal noise generated in a 50-Ω resistor, –174 dBm/Hz, or –159 dB/Hz.

The receiver's minimum weak-signal gain must be high enough so that the weakest desired usable signal plus the receiver noise is greater than at least one ADC quantizing level. If not, the signal won't be recovered. The best way to ensure a large enough signal is to provide sufficient noise at the ADC input to dither across a quantizing level. To do so, a designer can amplify the in-band receiver thermal noise or add out-of-band noise, which the digital-signal processor can filter out later. Although it may not be the better method, this article discusses the simpler approach of amplifying the receiver thermal noise until it bridges a quantizing level.

For the AD779, one quantizing level is 610 μVpk-pk. The equivalent sine-wave amplitude is then 215.8 μV rms or –60.3 dBm. Therefore, amplifying the receiver noise so that its rms level at the ADC is also –60.3 dBm will guarantee adequate bridging of a quantizing level. With the front-end noise bandwidth (BW) of 16 kHz, the gain required to amplify the receiver noise to this level is:

Gain = quantizing level – thermal noise density – 10log BW
= –60.3 dBm + 159 dBm/Hz – 10log 16,000 Hz
= 56.7 dB

A spreadsheet program can perform these calculations, and others, to interactively examine receiver performance for different ADCs, sample rates, bandwidths, and gain distributions *(see the table)*.

Note that in this example, the analog gain is high enough to allow thermal noise to bridge a quantization level. As a result, the receiver's noise performance is dominated by thermal noise, not analog-to-digital quantization noise. At the sensitivity level, the thermal noise is 7 dB greater than the quantization noise when measured in the 3-kHz information bandwidth.

The output signal-to-noise ratio (SNR) is better than 10 dB at an antenna signal level of –113 dBm. Digital automatic gain control (AGC) holds the output signal level at 20 dB below full scale (+4 dBm) until the RF AGC threshold is reached. The output SNR increases with the antenna signal level until the level at the ADC reaches 44 dB below full-scale (–20 dBm). At this point, the analog AGC holds the ADC input and output levels constant. The output SNR continues to rise to 54 dB until the analog-to-digital quantization noise dominates the receiver noise.

Unwanted Signals

Another consideration is the antenna overload level, which is the antenna signal level at which an undesired received signal can pass unattenuated through the 16-kHz bandwidth front-end to the ADC and saturate it. Because this undesired signal would be outside of the 3-kHz information bandwidth, it would not affect the normal AGC setting.

As a result, the receiver must maintain some "headroom" to accommodate signal-level peaks. Only desired in-band signals should be present after the final narrowband filtering in the digital-signal processor. Consequently, 20 dB of headroom is sufficient at the digital-signal-processor output. At the ADC input, however, large out-of-band undesired signals may be present. Consequently, at least 40 dB of headroom is required.

The maximum headroom is limited by the SNR required for both noise- or distortion-limited cases. In the example design, the maximum SNR for large desired signals, without interference, is 54 dB. If an out-of-band undesired signal at the ADC has a level of +4 dBm, or –20 dB from the saturation level, the converter will generate harmonic and intermodulation distortion products that will alias and potentially fall inside the desired signal bandwidth. If these products are specified to be 70 dB down (–66 dBm) and the in-band desired signal level at the ADC is –20 dBm, the ultimate signal-to-noise-and-distortion, S/(N+D), ratio will be limited to 46 dB (–20 dBm + 66 dBm).

The next step in the design is the digital IF processor architecture, which will depend on the desired level of functionality and the choice of hardware. In a typical implementation of a digital SSB processor, the digitized IF signal is first processed by the IF translator *(Fig. 4)*. The

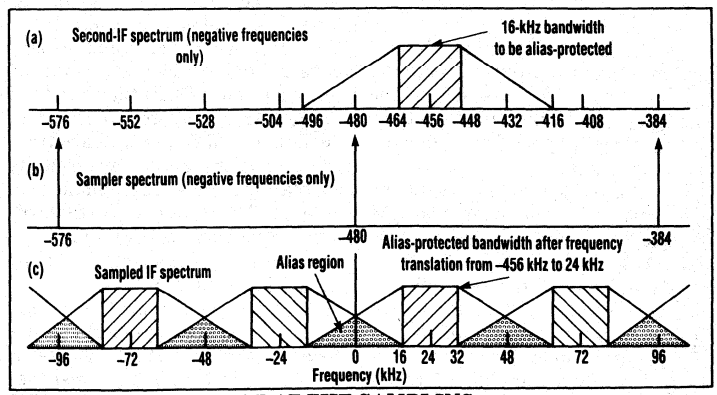

3. A DETAILED LOOK AT THE SAMPLING process helps in specifying the first-IF filter. The 456-kHz second IF (a) and 96-kHz sampling frequency (b) are the determining factors. To protect the desired 16-kHz bandwidth from aliasing, the filter must have no more than an 80-kHz-wide stopband, or a 5-to-1 shape factor (c).

translator shifts the SSB channel's center frequency to dc, or zero frequency, and at the same time converts the "real" IF signal to a complex baseband signal. The resulting real and imaginary components are often referred to as the inphase (I) and quadrature (Q) components, respectively. This representation of the signal as a rotating vector or phasor is commonly used in digital-signal-processing algorithms.[1]

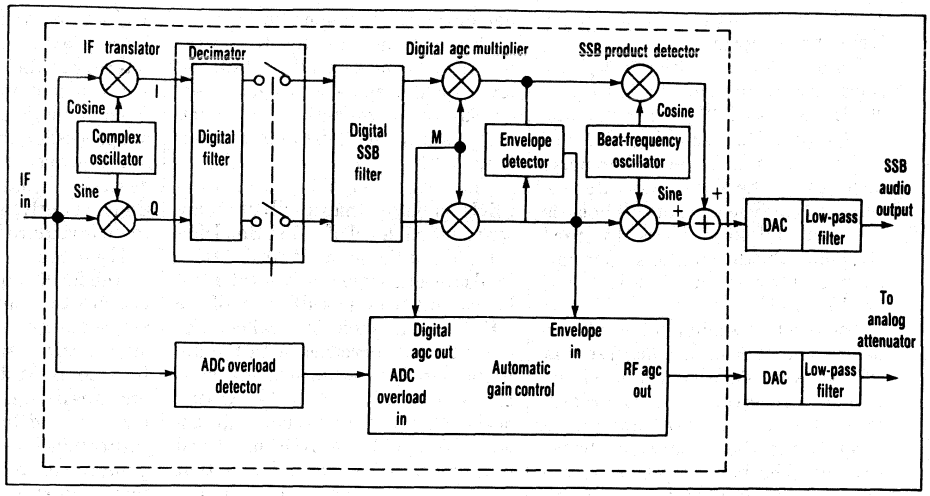

4. THE RADIO'S SSB FILTER, AGC, and demodulator are implemented by the digital-signal-processor chip. The system's sample rate is first reduced by a factor of 10 in the decimator to produce enough instruction cycles between samples to perform the required algorithms.

A complex oscillator and a multiplier perform the frequency translation and complex baseband conversion. Typically, the oscillator consists of a phase accumulator and sine and cosine algorithms *(Fig. 5)*. The phase accumulator is a modulo 2π adder that increments its output each sample period by adding a fixed constant, or phase increment, to its previous sample output. This process produces a continuously changing output phase angle ranging from 0 to 2π radians.

The algorithms then compute the phase angle's cosine and sine to generate the complex oscillator's real and imaginary outputs. Next, the real input signal is multiplied by the oscillator's outputs to produce the I and Q components of the complex baseband signal.

The next step is to implement the SSB filter, AGC, and demodulator in the digital-signal-processor chip. But the system's sample rate must be reduced as much as possible before the digital-signal processor can perform these computationally intensive algorithms. For example, if the DSP's instruction rate is 12 MHz and the sample rate remains at 96 kHz, only 125 instruction cycles will be available to run the algorithms on each point. This figure is insufficient. Therefore, the receiver decimates the sample rate by a factor of 10. The resulting 9600-Hz rate is enough to support the typical 3-kHz SSB bandwidth. It also enables 1250 instruction cycles between the decimated samples, enough to perform the required algorithms.

Filter The Input

Before the sample rate can be reduced, the input signal must be filtered to attenuate out-of-band signals. Otherwise, these signals can alias back in-band, causing distortion. This decimation filter usually consists of two identical finite-impulse-response (FIR) filters in the I and Q paths. Computer programs for designing FIR filters with arbitrary frequency response are readily available from many commercial sources. The most popular filter design algorithm, used in many of these programs, was distributed by the IEEE Press in 1979.[2]

After decimation, the signal passes through a high-performance receiver IF filter. This SSB filter is an FIR or infinite-impulse-response device. If necessary, the SSB filter's output level is adjusted by the digital portion of the AGC algorithm. The AGC controls the digital gain by multiplying the I and Q components by a positive scalar value, M *(Fig. 4, again)*.

The leveled I and Q outputs are then applied to the envelope detector and SSB product detector. The envelope detector computes the magnitude of the phasor represented by the I and Q components. The output envelope is simply the square root of the sum of the squared I and Q values. The AGC algorithm uses the envelope to adjust the receiver output level by a combination of analog (or RF) and digital gain control. To prevent saturation by strong out-of-band signals, the algorithm also senses any ADC overload and adjusts the RF and digital gain distributions accordingly.

The SSB product detector is a "half-complex" frequency translator that shifts the SSB carrier to its proper frequency. Just as in the IF translator, the circuit uses a complex oscillator, commonly called a beat-frequency oscillator. First, the I and Q signal values are multiplied by the cosine and sine outputs of the oscillator. Then the products are added to produce the receiver's audio output.

The hardware needed to implement the digital IF processor depends on several factors, but nearly

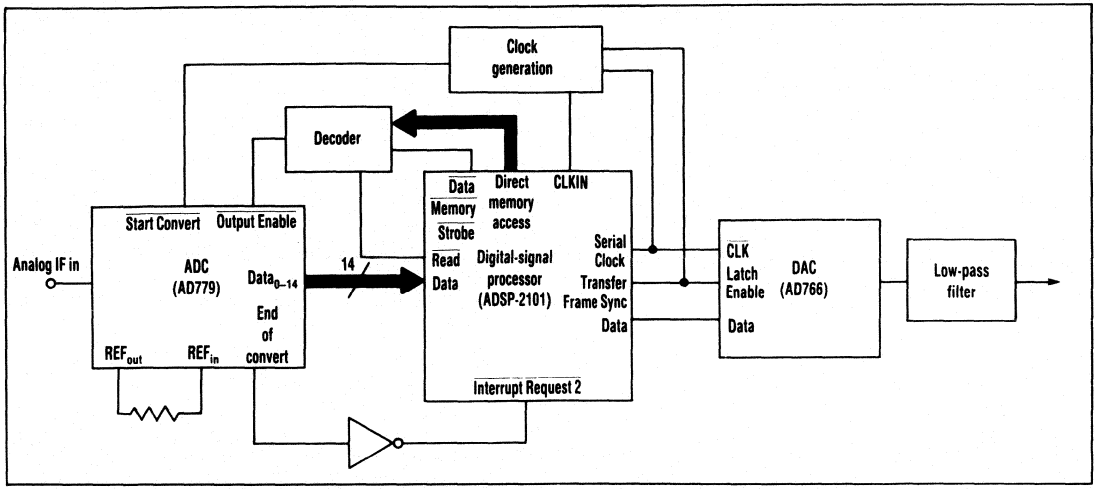

5. BECAUSE TYPICAL DIGITAL IF PROCESSOR architectures are highly integrated, most functions can be implemented in the DSP software. In this example circuit, only a few ICs are needed to perform the functions, as described in Figure 4.

all systems will include certain common components *(Fig. 6)*. The total number of ICs can be minimal, keeping overall system cost low.

Modern communication receivers would typically use a high-performance digital-signal processor, such as the ADSP-2101. This processor is well-suited to many of the operations of a digital receiver and can easily link to the 14-bit parallel output of the AD779. Conversion back to the analog domain is simplified by digital-to-analog converters (DACs), like the AD766. These low-cost, 16-bit, serial-input DACs require no additional glue logic or interface circuitry, have an on-board reference, and are completely tested and specified over temperature.

Additional circuitry performs decoding, clock generation, and output filtering. The decoder handles handshaking between the AD779 and the digital-signal processor. Depending on the number of peripherals connected to the processor, the decoder may be as simple as a few logic gates. Clock signals are typically generated by a high-stability crystal with low phase noise. The crystal can serve as the processor's master clock and can also be divided by counters or other means to generate the ADC and DAC sampling clocks.

On the falling edge of the sampling clock ($\overline{SC}$), the AD779 digitizes the incoming IF signal *(Fig. 6, again)*. Approximately 8 μs later, the conversion is complete and the AD779 asserts End of Convert (EOC). The falling edge of the inverted EOC signal interrupts the digital-signal processor. The control and address lines from the processor are decoded to generate the Output Enable ($\overline{OE}$) pulse for the ADC. The ADC then places the digitized output on the data bus to be read by the processor.

The ADSP-2101 offers several features tailored to communications systems. The ability to fetch two operands (typically a coefficient and a data point), multiply these operands, and then sum the results with previous products in one processor cycle makes FIR filter algorithms extremely efficient. And an internal 40-bit accumulator enables full-precision products to be calculated without round-off noise caused by truncation in the filter computations. Interprocessor communication between multiple DSPs, which is often needed in SSB equipment, is easily implemented with one of the processor's two integrated serial ports.

The digital IF processor's high integration level is indicative of the future of SSB equipment. As the performance of the ADC improves, it will be placed closer to the receiving antenna in the analog front-end. The ultimate goal is to digitize the incoming RF signal directly from the antenna and implement the remaining digital functions in custom ASICs.□

References:

[1] W.E. Sabin and E.O. Schoenike, *Single-sideband systems & circuits*, (McGraw-Hill Book Co., 1987).

[2] J.H. McClellan, T.W. Parks, and L.R. Rabiner; DSP Committee of the IEEE Acoustics, Speech, and Signal Processing Society, *Programs for digital signal processing*, (IEEE Press, 1979), Chap. 5.

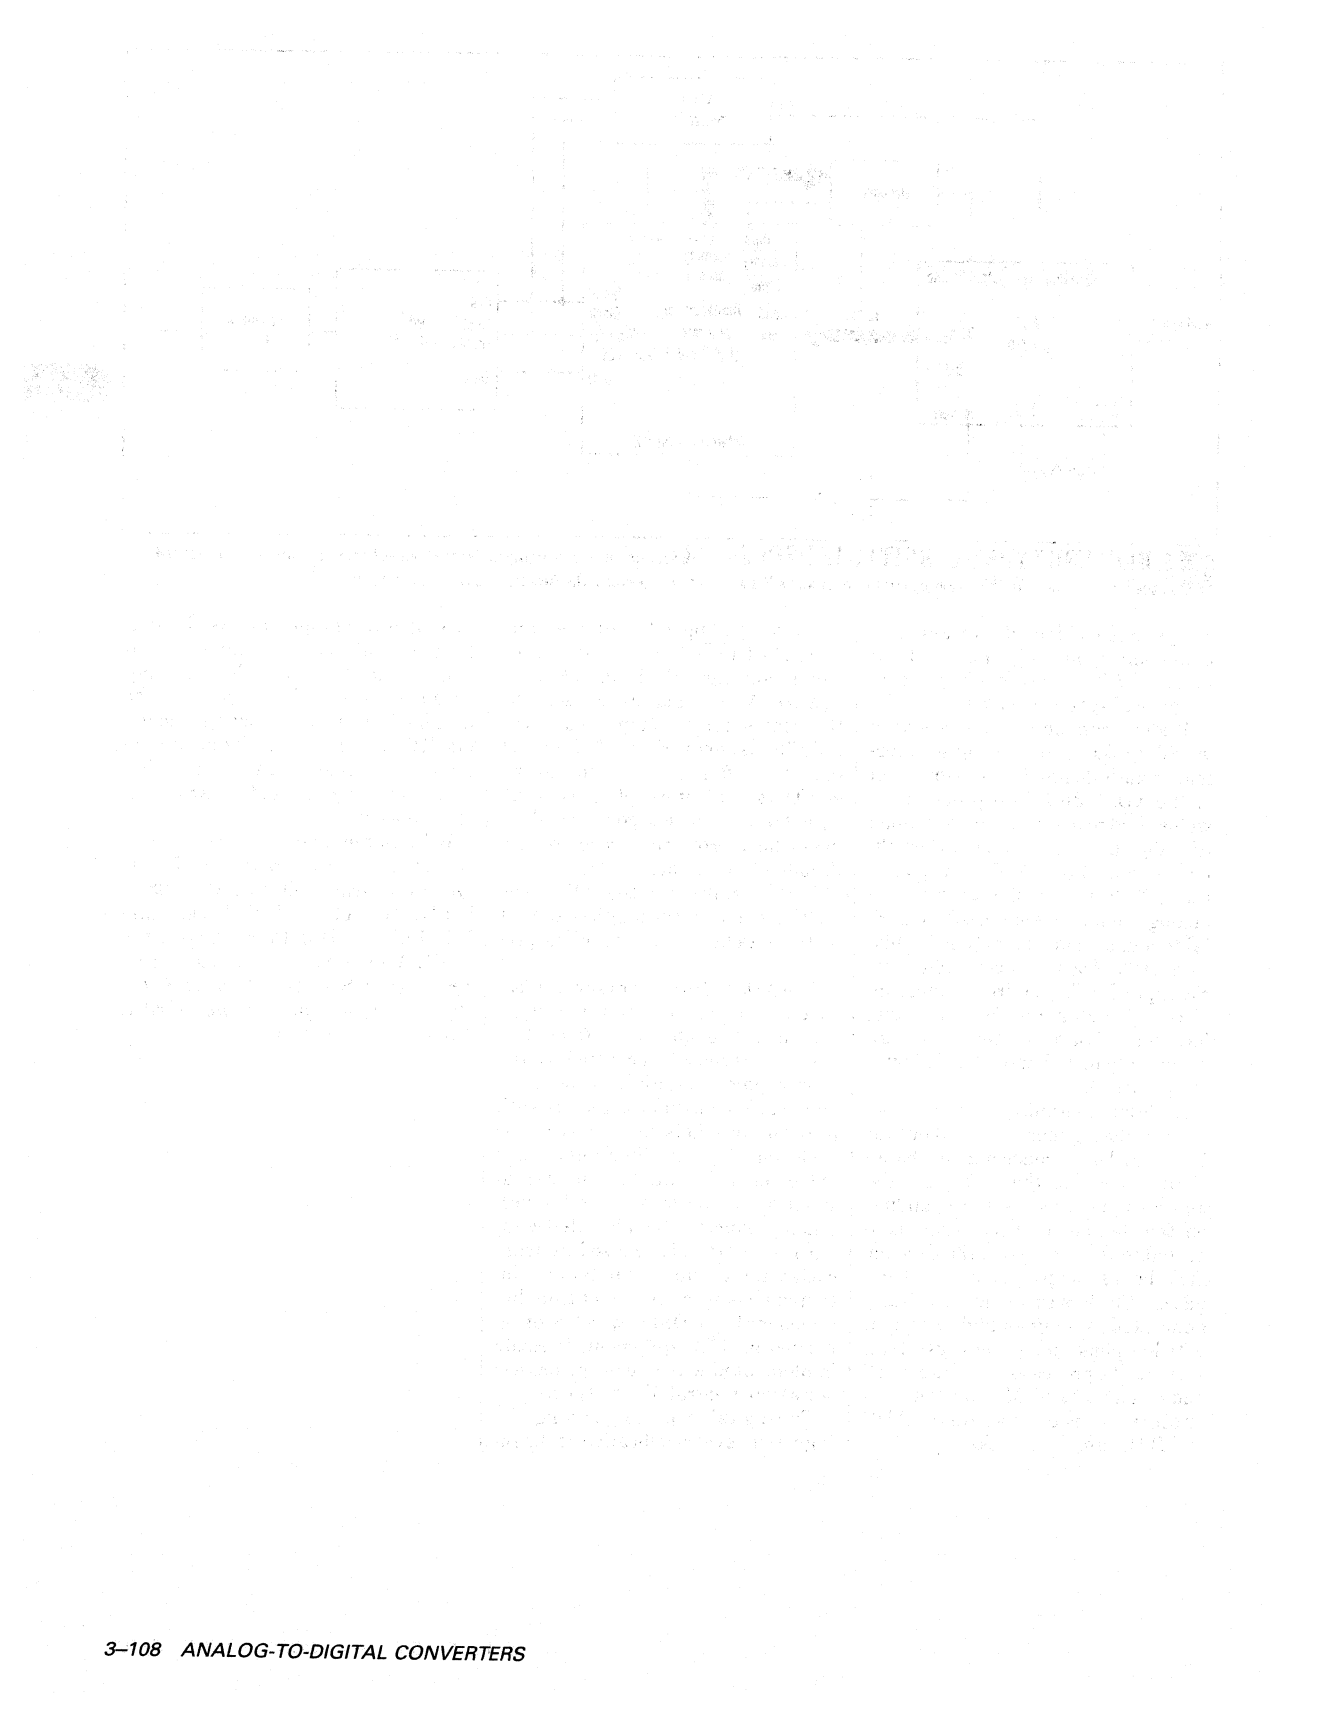

ANALOG DEVICES

AN-228
APPLICATION NOTE

NE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Evaluation Board for the AD7884 16-Bit A/D Converter

by Mike Curtin

NTRODUCTION

his application note describes the evaluation board for ıe AD7884 16-bit A/D converter. This is a 16-bit ıonolithic analog-to-digital converter with internal ample-and-hold and a conversion time of 5.3 µs. The ıroughput rate is 166 kSPS. It uses a 2-pass flash archi- ecture to achieve this speed and throughput. It has ıoth ac and dc specifications. Integral Linearity Error is :0.006% FSR and the Signal to (Noise + Distortion) Ra- io is 84 dB. Its fast 16-bit parallel output interface is ompatible with both DSPs (TMS320C25, ADSP-2101,)SP56000) and general purpose processors (MC68000, 80286, etc.). Full data on the converter is in the AD7884/AD7885 data sheet from Analog Devices. This should be consulted in conjunction with the application note when using the evaluation board.

The board operates from +15 and −15 volt power supplies. On-board components include the reference and op amp circuitry necessary for the analog front-end and output latches for interfacing to a 16-bit processor bus.

A full circuit diagram for the AD7884 board is shown in Figure 3.

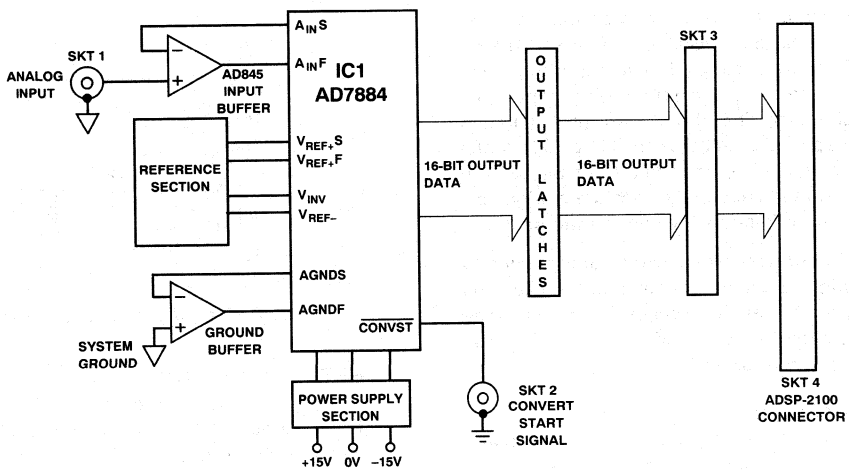

Figure 1. Evaluation Board Block Diagram

LINK OPTIONS

The evaluation board has several link options available and these are summarized in Table I below.

Table I. Link Options

Link No.	Function
LK1–LK4	Links 1–4 enable the user to select either the ±3 volt input range or the ±5 volt input range.
LK5	This allows $\overline{RD}$ to be tied permanently low.
LK6	This allows $\overline{CS}$ to be tied permanently low.
LK7	This link allows the two output latches to be enabled together (LK7 inserted) or separately (LK7 removed).
LK8	LK8 breaks the A_{IN} line so that the user can easily route it to extra signal conditioning such as antialias filtering or an extra gain stage. The board has a grid section where this external circuitry can be built.

POWER SUPPLIES, GROUNDING AND DECOUPLING

The board is powered from a ±5 V supply. These supplies drive two 5 V regulators (IC7 and IC8) which produce the ±5 V required for the AD7884. The AD7884 has one AV_{DD} pin and two V_{DD} pins. These are all driven from the same +5 V supply. The AV_{DD} pin is decoupled to signal ground with a 10 μF tantalum and a 0.1 μF ceramic capacitor. Both V_{DD} pins are decoupled with 0.1 μF capacitors only. The same decoupling arrangement is used for the negative supply pins. AV_{SS} is decoupled to signal ground with 10 μF and 0.1 μF while each of the V_{SS} pins is decoupled with 0.1 μF only.

The ±15 V supplies also drive the analog front-end circuitry which includes the AD586 reference and the op amps A1–A4. They are decoupled to signal ground with 10 μF tantalum and 0.1 μF ceramic disc capacitors. Power for the digital section of the board (IC9, IC10 and IC11) should be applied at either A31, B31 or C31 of the Eurocard connector, SKT 4, or at Pin 23 of SKT 3. The digital supply is routed separately from the analog section, and they are not joined anywhere on the board.

The evaluation board uses extensive ground planing to minimize any high frequency noise interference. The ground planing for the analog section is kept separate from that for the digital section, and they are joined only at Pin 30 (DGND) of the AD7884.

ANALOG INPUT SECTION

The analog input to the board is applied to the miniature BNC connector SKT 1, labelled A_{IN}. The input then goes to the link, LK8, and from there on to the input buffer, IC3. LK8 allows the user to divert the input to signal conditioning circuitry, which may be built on the grid section of the board. With LK8 inserted, A_{IN} goes directly to IC3, an AD845. This is a very fast JFET amplifier which combines very low input offset voltage and offset drift with fast settling and high slew rate. Consult the AD845 data sheet for full details.

The AD845 drives the AD7884 inputs through a series c links. These links allow the user to configure the ADC fc either a ±3 volt input or a ±5 volt input. Table II belo shows the necessary link settings for each range.

Table II. LK1–LK4 Settings

Input Range	LK1	LK2	LK3	LK
±5 Volts	AB	AB	AB	AB
±3 Volts	BC	BC	BC	BC

AGND BUFFER

IC2 is the force, sense amplifier for AGND. The AGNDS pin should be at zero potential. The board uses an OP-07 for IC2. The output of IC2 is decoupled with a 47 μF solic tantalum capacitor to AGND to deal with the fast curren transients on the AGNDS pin. The stability of this ar rangement is marginal, and if the user wishes to im prove the phase margin, the circuit given in Figure 2 may be used. A feedback capacitor (C_F) of 47 μF shoulc be used. This circuit compensates for the load capacitor by adding a low frequency zero and ensures an adequate phase margin.

MULTIPLEXER APPLICATIONS

When the AD7884 is used in multiplexer applications, it is possible for the analog input to see a full-scale step input. In these applications, it is better to use a very fast amplifier as the AGND buffer. Suitable op amps are the AD845 or AD847. With these, there is no need for the 47 μF capacitor to AGND at the output, and the step response time allows a throughput of 150 kHz to be achieved.

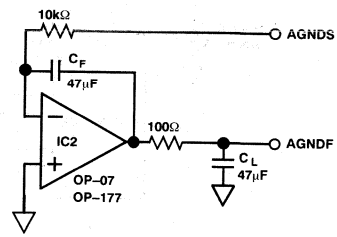

Figure 2. Compensation Circuit for AGND Buffer

REFERENCE SECTION

The required +3 V reference is derived from the AD586. The +5 V output is divided down to +3 V by R1 and R2 before being buffered by IC6. IC4 is a unity gain inverter which provides the −3 V negative reference. The gain setting resistors are on-chip and are factory trimmed to ensure precise tracking of V_{REF+}. Figure 3 shows IC4 and IC6 as AD845s. The very high slew rate of these amplifiers means that they can respond to the rapidly changing reference input impedance and input currents and hold the V_{REF+} and V_{REF-} inputs at the required dc levels.

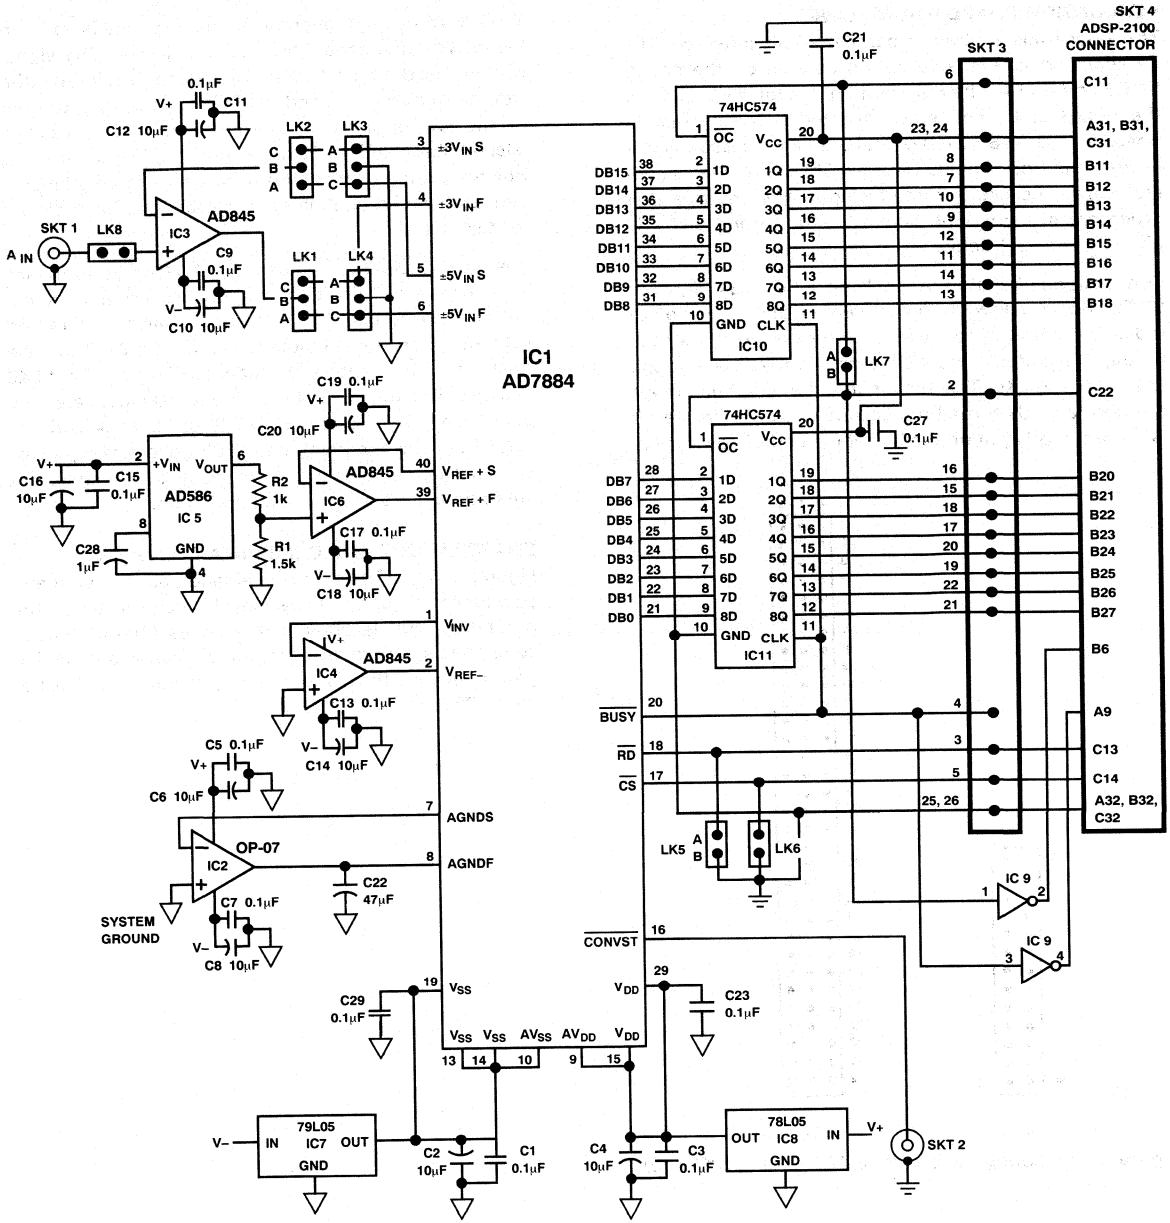

Figure 3. AD7884 Evaluation Board Circuit Diagram

EVALUATION BOARD INTERFACING

The board has a straight parallel 16-bit interface. The Parallel Port is available at two sockets on the board, SKT 3 and SKT 4. The Eurocard Connector, SKT 4 allows the board to be connected directly to the evaluation board for the ADSP-2100 Digital Signal Processor, which is available from Analog Devices. SKT 3 provides a general purpose socket for interfacing to any other parallel processor systems. The pinout for SKT 3 and SKT 4 is given in Figure 4 below.

Conversion is initiated by applying a $\overline{\text{CONVST}}$ signal at SKT 2. $\overline{\text{BUSY}}$ is low during conversion and goes high to indicate the end of conversion. In the simplest interface mode, $\overline{\text{CS}}$ and $\overline{\text{RD}}$ are tied permanently low through links LK5 and LK6. This means that the data bus drivers are permanently turned on. Each new conversion result is automatically put on the bus at the end of conversion.

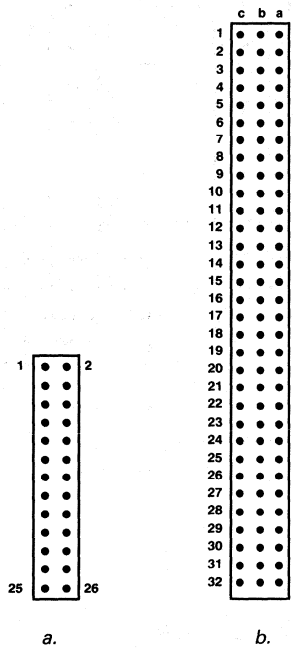

a. *b.*

Figure 4. Pin Configurations for SKT 3 (a) and SKT 4 (b)

The AD7884 design guarantees that the data is valid on the $\overline{\text{BUSY}}$ rising edge. This means that the $\overline{\text{BUSY}}$ signal can be used to latch the conversion result into the external latches, IC 10 and IC 11. These may then be read by the processor before the next conversion result is latched. For standard 16-bit interfacing, LK7 is inserted. However, if the user wishes to interface to an 8-bit system, LK7 can be omitted so that each 8-bit latch can be read separately.

It is also possible to evaluate the AD7884 performance when connected directly to the processor data bus and not isolated by latches. Take out the latches, IC10 and IC11 and short the D inputs to the Q outputs. Then use the $\overline{\text{CS}}$ and $\overline{\text{RD}}$ inputs in the standard fashion. Omit LK5 and LK6. This means that the $\overline{\text{CS}}$ and $\overline{\text{RD}}$ inputs are now controlled by Pins 5 and 3 (SKT 3) or Pins C14 and C13 (SKT 4). In this configuration, $\overline{\text{BUSY}}$ generates a processor interrupt at the end of conversion. The interrupt subroutine is a standard Read operation.

BOARD LAYOUT AND COMPONENT OVERLAY

Figures 5, 6 and 7 show the evaluation board layout and component overlay. These may be used to study the grounding and decoupling techniques for the various components. There is a separate analog and digital ground plane which is joined only at Pin 30 (DGND) of the AD7884.

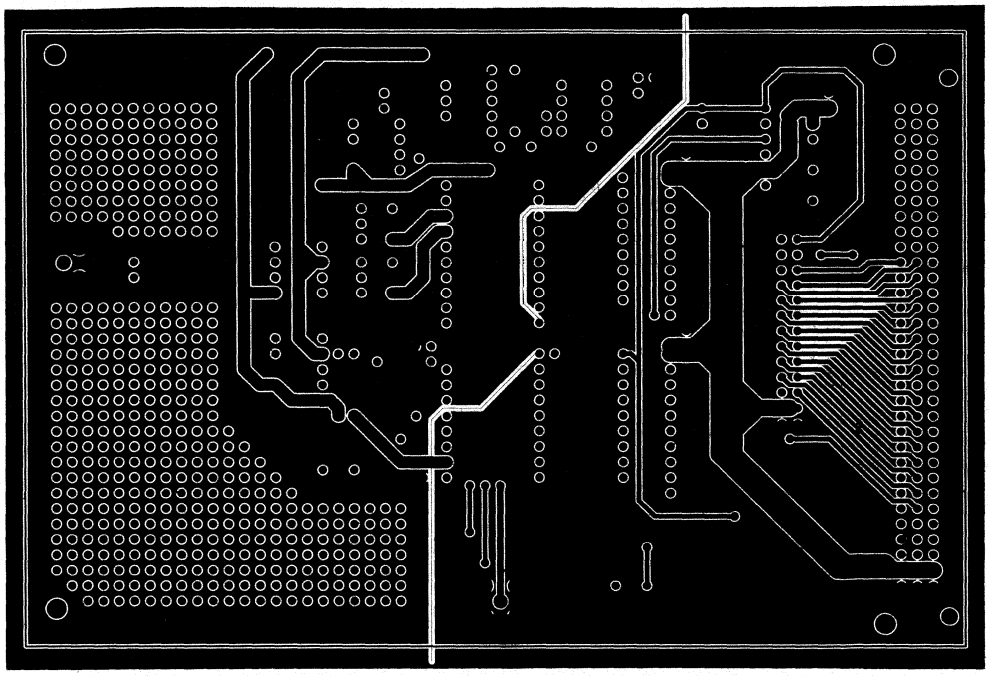

Figure 5. AD7884 Evaluation Board Component Side Layout

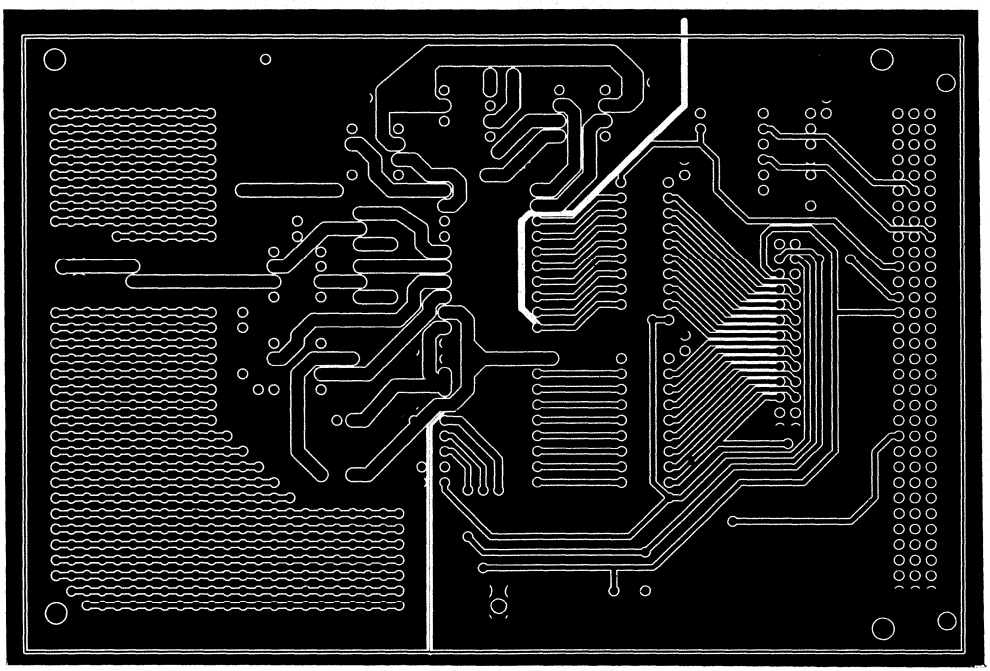

Figure 6. AD7884 Evaluation Board Solder Side Layout

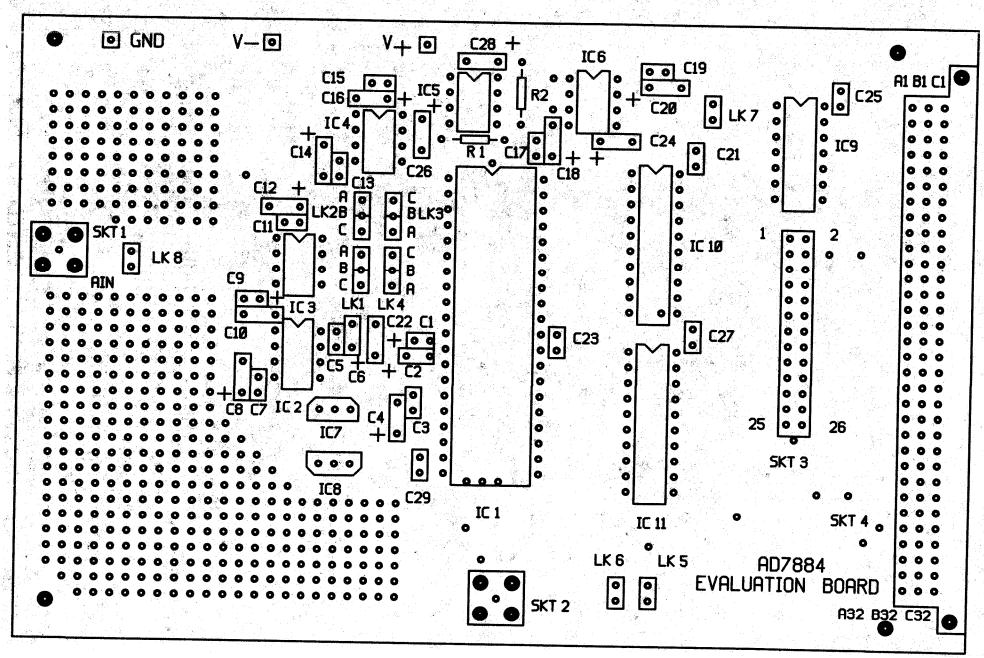

Figure 7. AD7884 Evaluation Board Component Overlay

ANALOG DEVICES

AN-229
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Evaluation Board for the AD7885 16-Bit A/D Converter

by Mike Curtin

INTRODUCTION

This application note describes the evaluation board for the AD7885 16-bit A/D converter. This is a 16-bit monolithic analog-to-digital converter with internal sample-and-hold and a conversion time of 5.3 μs. The throughput rate is 166 kSPS. It uses a 2-pass flash architecture to achieve this speed and throughput. It has both ac and dc specifications. Integral Linearity Error is ±0.006% FSR, and the Signal to (Noise + Distortion) Ratio is 84 dB. Its fast 8-bit byte output interface is compatible with both microcontrollers (8051, 68HC11, etc.) and general purpose processors (MC6800 Series, MC68008, 8088, Z80, etc.). Full information on the converter is in the AD7884/AD7885 data sheet from Analog Devices. This should be consulted in conjunction with the application note when using the evaluation board.

The board operates from +15 and −15 volt power supplies. On-board components include the reference and op amp circuitry necessary for the analog front-end and output latches for interfacing to a 16-bit processor bus.

A full circuit diagram for the AD7885 board is shown in Figure 3.

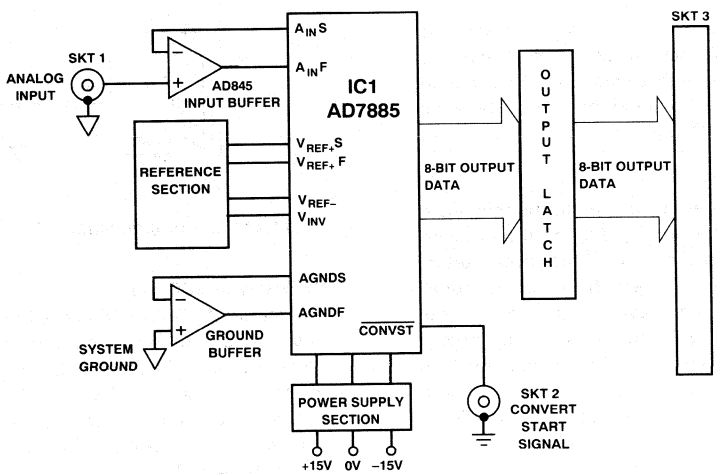

Figure 1. Evaluation Board Block Diagram

LINK OPTIONS

The evaluation board has a number of link options available and these are summarized in Table I below.

Table I. Link Options

Link No.	Function
LK1–LK4	Links 1–4 enable the user to select either the ±3 volt input range or the ±5 volt input range.
LK5	This allows $\overline{RD}$ to be tied permanently low.
LK6	This allows $\overline{CS}$ to be tied permanently low.

POWER SUPPLIES, GROUNDING AND DECOUPLING

The board is powered from a ±15 V supply. These supplies drive two 5 V regulators (IC7 and IC8) which produce the ±5 V required for the AD7885. The AD7885 has one AV_{DD} pin and one V_{DD} pin. These are driven from the same +5 V supply. The AV_{DD} pin is decoupled to signal ground with a 10 μF tantalum and a 0.1 μF ceramic capacitor. The V_{DD} pin is decoupled with a 0.1 μF capacitor only. The same decoupling arrangement is used for the negative supply pins. AV_{SS} is decoupled to signal ground with 10 μF and 0.1 μF while V_{SS} is decoupled with 0.1 μF.

The ±15 V supplies also drive the analog front-end circuitry which includes the AD586 reference and the op amps A1–A4. These are decoupled to signal ground with 10 μF tantalum and 0.1 μF ceramic disc capacitors. Power for the digital section of the board (IC9, IC10 and IC11) should be applied at Pin 23 of SKT 3. The digital supply is routed separately from the analog section and they are not joined anywhere on the board.

The evaluation board uses extensive ground planing to minimize any high frequency noise interference. The ground planing for the analog section is kept separate from that for the digital section, and they are joined only at Pin 21 (DGND) of the AD7885.

ANALOG INPUT SECTION

The analog input to the board is applied to the miniature BNC connector SKT 1, labelled A_{IN}. The input then goes to the link, LK8, and from there on to the input buffer, IC3. LK8 allows the user to divert the input to signal conditioning circuitry, which may be built on the grid section of the board. With LK8 inserted, A_{IN} goes directly to IC3, an AD845. This is a very fast JFET amplifier which combines very low input offset voltage and offset drift with fast settling and high slew rate. Consult the AD845 data sheet for full details.

The AD845 drives the AD7885 inputs through a series of links. These links allow the user to configure the ADC for either a ±3 volt input or a ±5 volt input. Table II shows the necessary link settings for each range.

Table II. LK1–LK4 Settings

Input Range	LK1	LK2	LK3	LK4
±3 Volts	AB	AB	AB	AB
±5 Volts	BC	BC	BC	BC

AGND BUFFER

IC2 is the force, sense amplifier for AGND. The AGNDS pin should be at zero potential. The board uses the OP-07 for IC2. The output of IC2 is decoupled with a 47 μF solid tantalum capacitor to AGND to deal with the fast current transients on the AGNDS pin. The stability of this arrangement is marginal, and the circuit shown in Figure 2 is used to improve phase margin. A feedback capacitor (C_F) of 47 μF should be used. This circuit compensates for the load capacitor by adding a low frequency zero and ensures an adequate phase margin.

MULTIPLEXER APPLICATIONS

When the AD7885 is used in multiplexer applications, it is possible for the analog input to see a full-scale step input. In these applications, it is better to use a very fast amplifier as the AGND buffer. Suitable op amps are the AD845 or AD847. With these, there is no need for the 47 μF capacitor to AGND at the output, and the step response time allows a throughput of 150 kHz to be achieved. Thus, if the user wishes to multiplex several channels into the AD7885, C3 and C26 should be removed, R4 and R5 should be shorted out and the OP-07 (IC2) should be replaced with either the AD845 and AD847.

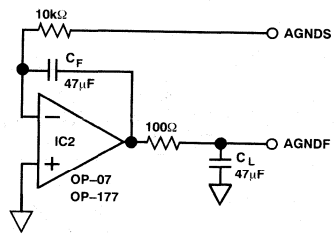

Figure 2. Compensation Circuit for AGND Buffer

REFERENCE SECTION

The required +3 V reference is derived from the AD586. The +5 V output is divided down to +3 V by R1 and R2 before being buffered by IC6. IC4 is a unity gain inverter which provides the −3 V negative reference. The gain setting resistors are on-chip and are factory-trimmed to ensure precise tracking of V_{REF+}. Figure 3 shows IC4 and IC6 as AD845s. The very high slew rate of these amplifiers have the ability to respond to the rapidly changing reference input impedance and input currents and hold the V_{REF+} and V_{REF-} inputs at the required dc input levels.

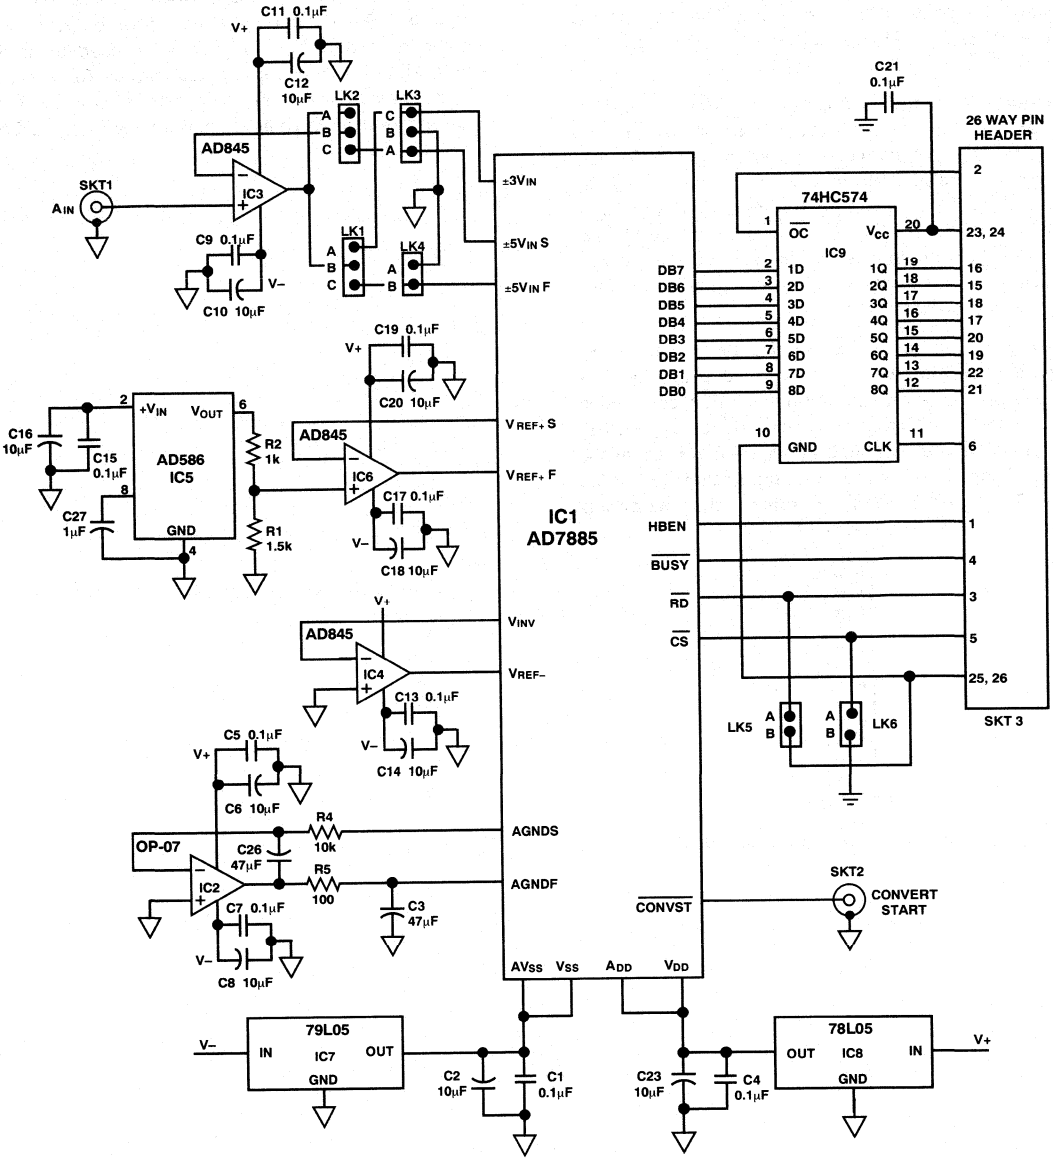

Figure 3. AD7885 Evaluation Board Circuit Diagram

EVALUATION BOARD INTERFACING

The board has an 8-bit byte interface. The 8-bit port is available at SKT 3. The pinout for SKT 3 is given in Figure 4.

Conversion is initiated by applying a $\overline{\text{CONVST}}$ signal at SKT 2. $\overline{\text{BUSY}}$ is low during conversion and goes high to indicate the end of conversion. When the $\overline{\text{BUSY}}$ line goes high, it is possible to read the conversion result by bringing $\overline{\text{CS}}$ and $\overline{\text{RD}}$ low. In the two-byte read operation, the HBEN input determines whether the most significant or least significant byte is read. When HBEN is high, then the 8 MSBs are put on the bus; when HBEN is low, then the 8 LSBs are put on the bus.

It is also possible to tie $\overline{\text{CS}}$ and $\overline{\text{RD}}$ permanently low through links LK5 and LK6. This means that the data bus drivers are permanently turned on. Each new conversion result is automatically put on the bus at the end of conversion. The state of HBEN determines whether the least significant byte or the most significant byte is on the bus. The AD7885 design guarantees that the data is valid on the $\overline{\text{BUSY}}$ rising edge.

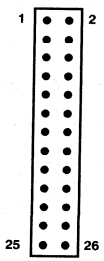

Figure 4. Pin Configuration for SKT3

BOARD LAYOUT AND COMPONENT OVERLAY

Figures 5, 6 and 7 show the evaluation board layout and component overlay. These may be used to study the grounding and decoupling techniques for the various components. There are separate analog and digital ground planes which are joined only at Pin 21 (DGND) of the AD7885.

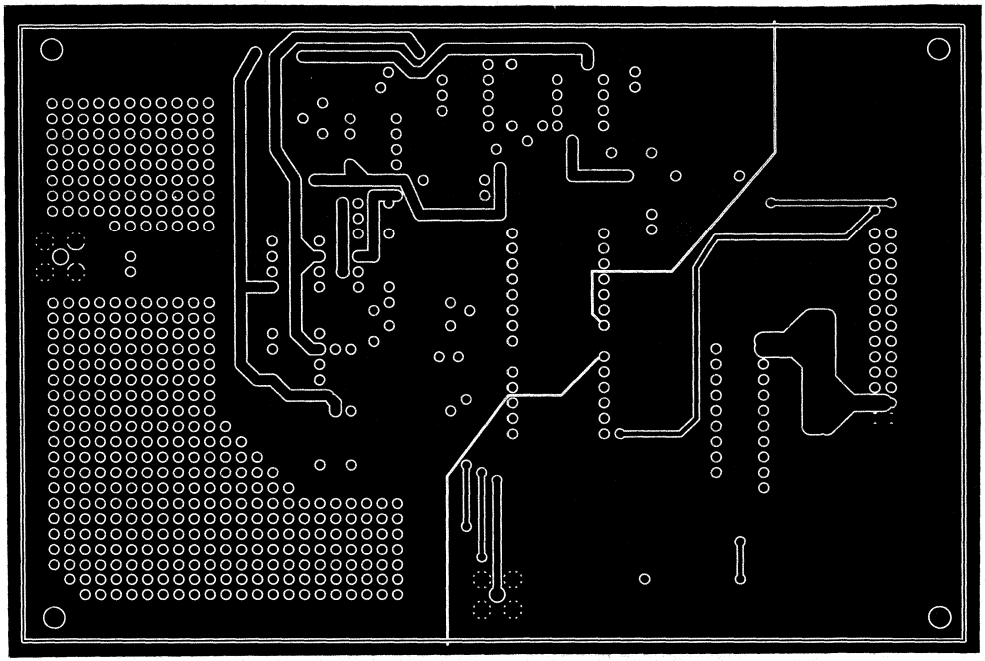

Figure 5. AD7885 Evaluation Board Component Side Layout

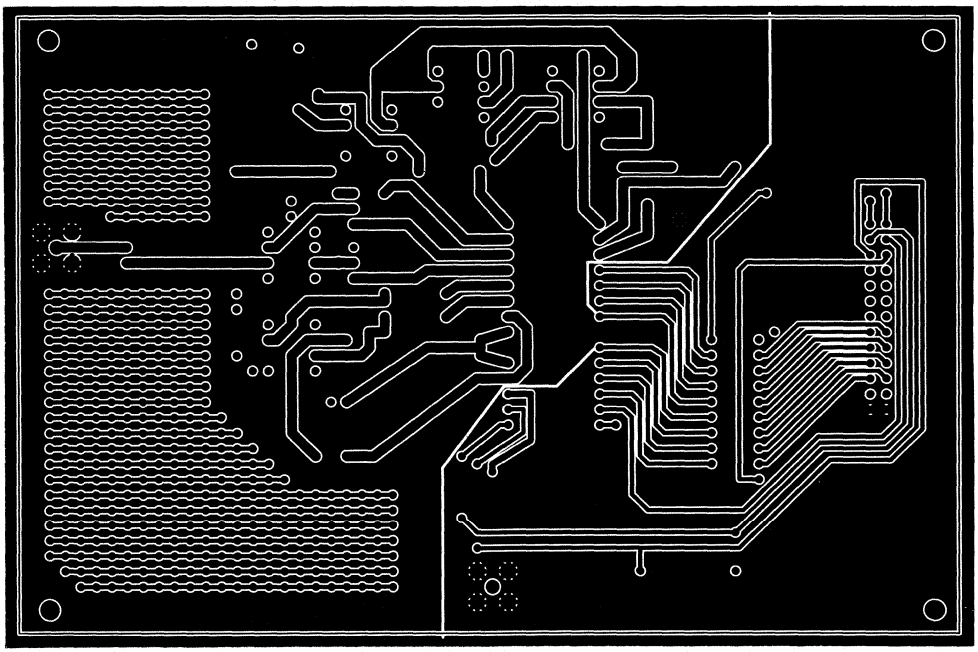

Figure 6. AD7885 Evaluation Board Solder Side Layout

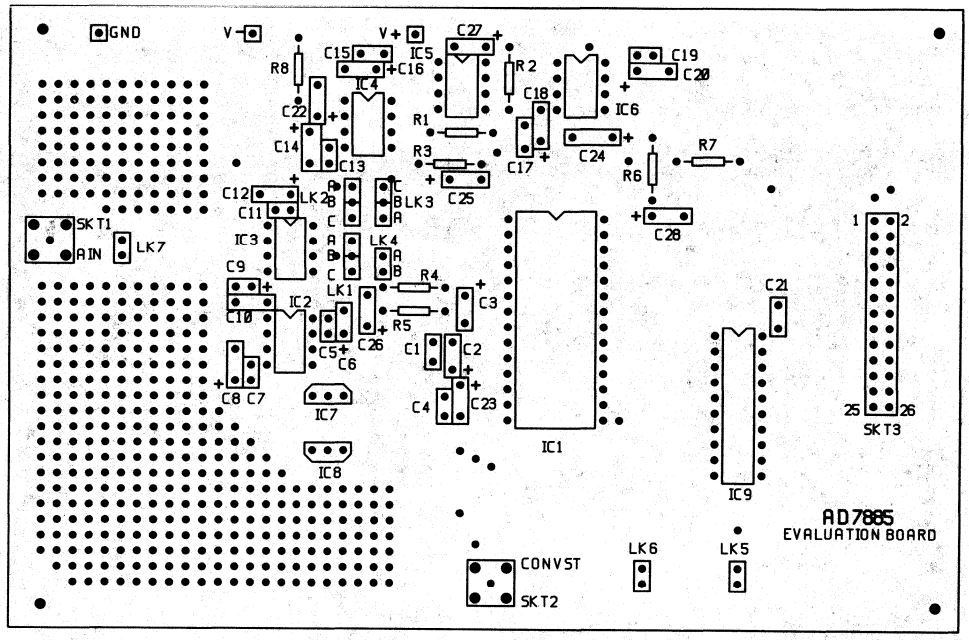

Figure 7. AD7885 Evaluation Board Component Overlay

ANALOG DEVICES

AN-368
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Evaluation Board for the AD7701/AD7703 Sigma Delta A/D Converters

by Mike Curtin

INTRODUCTION

This application note describes the evaluation board for the AD7701 and AD7703 A/D converters. These converters utilize sigma delta techniques to offer exceptional performance. The AD7701 is a 16-bit device with 0.0015% FSR INL error and no missing codes. The AD7703 is a 20-bit device with 0.0003% FSR INL error and no missing codes. Both parts (which are pin compatible) have a self-calibration mode which removes internal offset and gain errors and a system calibration mode which removes external circuit offset and gain errors. Both devices have a flexible synchronous serial interface which allows the ADCs to connect directly to digital signal processors and microcontrollers. In addition, the AD7701 has an asynchronous interface for use with UARTs. Full data on the AD7701 and AD7703 is available in data sheets from Analog Devices and should be consulted in conjunction with this application note when using the evaluation board.

The board has been designed to enable all the features of the converters to be fully evaluated. These features include analog input range, calibration mode, interface mode selection, master clock rate and output data rates.

The board operates from +5 volt and −5 volt power supplies. On-board components include a 2.5 V reference, clock generation circuitry, a decimation counter and parallel interface circuitry. The three different interface options are as follows: a parallel output port suitable for μC (microcontroller) and DSP (digital signal processor) interfacing, an asynchronous RS-232 interface (AD7701 only), a synchronous serial interface suitable for the serial ports of both microcontroller and DSP machines.

A full circuit diagram for the board is shown in Figure 5. The board layout and silkscreen are given in Figures 7, 8 and 9.

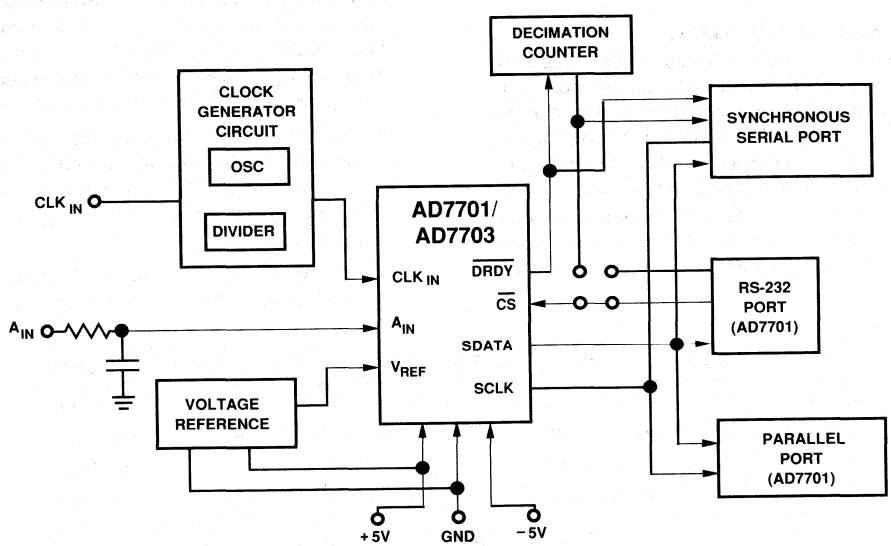

Figure 1. Evaluation Board Block Diagram

LINK OPTIONS

The evaluation board has several link options available and these are summarized in Table I below.

Link No.	Function
LK1	This link selects either the on-board oscillator (INT) or an external clock source (EXT) as the input to the clock generator section.
LK2	This allows any of the divider output rates to be used as the input clock to the AD7701/AD7703.
LK3	This is the baud rate selector when the AD7701 is in the AC (Asynchronous Communications) interface mode. Baud rates between 1200 and 19200 are available.
LK4	This link sets the divide ratio of the decimation counter.
LK5	This selects one of the three interface modes of the AD7701 and one of two modes of the AD7703.
LK6	This allows the calibration and sleep modes to be selected. It also selects the analog input range and determines whether the device does a self-calibration or a system calibration.
LK7	In the parallel interface mode for the board, this link configures the Data Acknowledge Signal (DACK) to be either active high or active low.
LK8	This link allows the $\overline{CS}$ input of the device to be controlled either by the output of the decimation counter or by the DTR input from the RS-232 interface.
LK9	This connects the on-board-generated baud clock to the device SCLK input.

Table I. Link Options

POWER SUPPLIES AND GROUNDING

The board is powered from a ±5 V supply. Both supplies are decoupled to ground with 10 μF tantalum and 0.1 μF ceramic disc capacitors. Power for the digital section of the board and the DV_{DD} pin of the AD7701/AD7703 is routed separately from the power for the analog section and AV_{DD} on the AD7701/AD7703. These analog and digital power lines are joined together only at one point which is at the input to the board. The four supply pins of the AD7701/AD7703 are each decoupled with 0.1 μF capacitors.

The evaluation board uses extensive ground planing minimize any high frequency noise interference from th on-board clocks or any other sources. Once again, th ground planing for the analog section is kept separa from that for the digital section and they are joined on at the board input.

BOARD INITIALIZATION

On power-up, the first step is to calibrate the device an set up the operating conditions. The evaluation boar offers two options here. The control pins (CAL, SC SC2, BP/$\overline{UP}$ and $\overline{SLEEP}$) can be set either by the DI switch, SW1, and the CAL pushbutton switch, SW2, c directly by LK6. If the board is operating in a stand-alon mode, the DIP switch, SW1, and the CAL pushbutto switch, SW2, set up the operating conditions and do th calibration. The truth table for the DIP switch settings i shown in Table II. Table III then shows the key for thes settings. When SW2 is activated, it will initiate a calibra tion cycle in accordance with the SW1-1 and SW1- settings. The CAL pin should be activated wheneve power is applied to the board or whenever the calibra tion mode is changed on SW1. If a processor is beinį used to control the board, LK6 can be driven by a port t set up the operating conditions and do the calibration The DIP switches should be in the off position if LK6 i being used to drive the control pins.

Switch	ON	OFF
SW1-1	SC1 = 0	SC1 = 1
SW1-2	SC2 = 0	SC2 = 1
SW1-3	Unipolar Input	Bipolar Input
SW1-4	Sleep Mode	Normal Mode

Table II. SW1 Settings

AD7701/AD7703 INTERFACE MODE SELECTION

The AD7701 may be set up for one of three serial inter face modes, and the AD7703 has two serial interface modes. These are controlled on the evaluation board by link option LK5. The modes are as follows:

1. The Synchronous Self-Clocking (SSC) mode which is intended for interfacing to digital signal processors and serial-to-parallel shift registers. In this case, data is clocked out on the falling edge of the internally generated SCLK.
2. The Synchronous External Clocking (SEC) mode, which is intended for interfacing to microcontrollers

CAL	SC1	SC2	CAL TYPE	ZERO REFERENCE	FS REFERENCE	SEQUENCE
⎍	0	0	Self-Cal	AGND	V_{REF}	One Step
⎍	1	1	System Offset	A_{IN}	–	1st Step
⎍	0	1	System Gain	–	A_{IN}	2nd Step
⎍	1	0	System Offset	A_{IN}	V_{REF}	One Step

Table III. Calibration Truth Table

such as the 8051 and 68HC11. In this case an external SCLK is applied to the AD7701/AD7703, and data is clocked out on its falling edge.

3. The Asynchronous Communications (AC) mode is present on the AD7701 only and is intended for interfacing to UARTs. In this mode, no clock is used. Data is clocked out in two serial bytes. Each byte has one start bit and two stop bits.

LK5 Position	Interface Mode
SSC	Synchronous Self-Clocking
SEC	Synchronous External Clocking
AC	Asynchronous Communication

Table IV. Interface Mode Selection

Table IV gives the link option positions for the various interface modes.

ANALOG INPUT SECTION

The analog input to the board is applied to the miniature BNC connector labelled A_{IN}. R12 and C22 provide filtering of any high frequency noise which may be present on the analog input. R12 also provides current limiting in the event of an input signal which exceeds the AV_{DD} supply.

The reference for the AD7701/AD7703 is provided by IC5 (the AD580, 2.5 V reference). This is decoupled with 10 μF and 0.1 μF to signal ground. The RC combination at the reference output act as a filter for any high frequency noise present. In addition, the AD580 output is buffered by IC12 (TLC271) and is available at V_{REFS} and V_{REFF} so that it can be used as a reference for other elements in the system.

CLOCK GENERATION

The evaluation board Clock Generator section is comprised of X1 (4.9152 MHz crystal), IC1, IC2 and associated passive components. Link option LK1 allows the user to choose between an externally supplied clock and that generated on the board. With LK1 in the EXT position, the input applied at the BNC connector labelled CLKIN is applied to the input of the counter IC2. This input should be +5 V CMOS compatible. Link option LK2 then selects a sub-multiple of this to drive the AD7701/AD7703 CLKIN pin. If LK1 is in the INT position, then the on-board 4.9152 MHz oscillator is selected as the counter input.

Counter IC2 (74HC4040) divides the input clock by 2^n where n = 0, 1, . . . 7. Table V shows the resulting CLKIN frequency for the various link positions when the on-board 4.9152 MHz oscillator is the counter input.

The AD7701/AD7703 has its own on-chip oscillator which needs only an external crystal or ceramic resonator to operate. Note however that ceramic resonators and low frequency crystals will need loading capacitors to operate correctly.

LK2 Position	CLKIN Frequency
0	4.9152 MHz
1	2.4576 MHz
2	1.2288 MHz
3	614.4 kHz
4	307.2 kHz
5	153.6 kHz
6	76.8 kHz
7	38.4 kHz

Table V. Choosing CLKIN Frequency for the AD7701/AD7703

The evaluation board allows operation with the on-chip oscillator. The desired crystal should be inserted in the X2 slot on the board. This crystal is now directly connected between the CLKIN pin and CLKOUT pin of the AD7701/AD7703. In addition, the wire link located directly above TP8 should be removed to disconnect the clock coming from IC2. Note that in addition to the X2 slot, there is also provision made for loading capacitors on the board. The space for these is located directly above X2.

DECIMATION COUNTER

Each time a data word is available for output from the AD7701/AD7703, the $\overline{DRDY}$ line goes low. If the $\overline{DRDY}$ line is tied directly to the device $\overline{CS}$ input, it will output data every time a data word is presented to the output pin. The output update rate is $f_{CLKIN}/1024$. Thus, for a CLKIN frequency of 4.096 MHz, the output update rate and output data rate will both be 4 kHz. Some applications will not require such a high output data rate and it is possible to reduce this by simply dividing down the $\overline{DRDY}$ signal before driving $\overline{CS}$. This function is known as decimation and is performed on the evaluation board by IC3, IC8a and Link Options LK4 and LK8. With a link inserted in LK8, Position A, the output of the counter IC3 drives the AD7701/AD7703 $\overline{CS}$ input. In this case, the counter accumulates 2^n counts (where n = 0 to 11) at which time the selected output enables the $\overline{CS}$ input of the AD7701/AD7703. The D input to flip-flop IC8a is enabled to a "1" at the same time as $\overline{CS}$ goes low. When $\overline{DRDY}$ returns high, IC8a is toggled and resets the counter IC3 which terminates the $\overline{CS}$ enable. Table VI gives the divide ratios for the LK4 settings and also gives the resulting output data rate, for a 4.096 MHz CLKIN.

EVALUATION BOARD INTERFACING

The board offers three separate interfaces: a synchronous serial port; an RS-232 port and a 16-bit parallel port.

Synchronous Serial Port

The synchronous serial port is available at the subminiature D connector, SKT1. The pin configuration for this is given in Figure 2. The port can be used in both the

LK4 Position	Divide Ratio	Output Rate (CLKIN = 4.096 MHz)
0	1	4 kHz
1	2	2 kHz
2	4	1 kHz
3	8	500 Hz
4	16	250 Hz
5	32	125 Hz
6	64	62.5 Hz
7	128	31.1 Hz
8	256	15.6 Hz
9	512	7.8 Hz
10	1024	3.9 Hz
11	2048	1.9 Hz

Table VI. Setting AD7701/AD7703 Output Data Rate

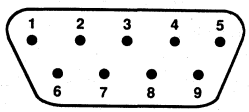

Figure 2. Pin Configuration for SKT1, Synchronous Serial Port

Synchronous Self-Clocking Mode and the Synchronous External Clocking mode. The port has six lines which are described below.

DGND Ground reference point

$\overline{CS}$ This buffered input controls the AD7701/AD7703 $\overline{CS}$ input

DR This buffered output is the $\overline{DRDY}$ signal from the AD7701/AD7703

SD This is the buffered SDATA output of the AD7701/AD7703

SCO Serial Clock Output. When the device is operating in the SSC mode, the serial clock from the AD7701/AD7703 is buffered and available at this pin. The output three state buffer is controlled by the MODE pin of the device. If the MODE pin is low (indicating SEC mode), the three state buffer is disabled.

SCI Serial Clock Input. When the device is operating in the SEC mode, the external serial clock is applied to this terminal and goes via the three state buffer to the SCLK pin. If the AD7701/AD7703 MODE pin is high (indicating SSC mode), then the three state buffer is disabled.

In the SSC mode serial data and serial clock are output from the AD7701/AD7703 whenever the $\overline{CS}$ line is activated. In order to control this by a processor port connected at SKT1, it is important to place the jumper in LK8 in position C (No Connection). This removes the decimation counter output from controlling the $\overline{CS}$ line.

In the SEC mode serial data is clocked out on the falling edge of the externally supplied serial clock whenever the $\overline{CS}$ line is activated. In order to have full control with a processor in this mode, the jumper in LK8 must be placed in position C and the jumper in LK9 must be placed in position B. These are both "No Connection" positions and disable the Decimation Counter control of $\overline{CS}$ and the Baud Rate Generator control of SCLK.

RS-232 Port (AD7701 Only)

The RS-232 port is available at the subminiature D connector, SKT2. The pin configuration for this is given in Figure 3.

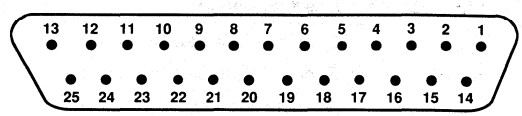

Figure 3. Pin Configuration for SKT2, RS-232 Port

When the AD7701 is set up for the AC (Asynchronous Communications) mode by tying the MODE pin to −5 V, data is transmitted in two serial bytes in UART-compatible format. An external SCLK sets the baud rate. On the evaluation board, this clock is derived from the counter IC2. Its frequency is set by the link option, LK3. The baud rates corresponding to the link positions are given in Table VII. Note that if the on-board baud clock is to be used, the jumper on LK9 should be in position A. This makes the connection between the output of IC2 and the AD7701 SCLK input. Alternatively, the baud clock can be controlled by a clock input to the SCI input at SKT1. In this case the jumper on LK9 should be in position B.

LK3 Position	Baud Rate Clock (CLKIN = 4.9152 MHz)
8	19.2 kHz
9	9.6 kHz
10	4.8 kHz
11	2.4 kHz
12	1.2 kHz

Table VII. Baud Rate Settings

The $\overline{DRDY}$ output from the AD7701 signals the CTS (Clear To Send) line of the RS-232 interface when data is available. The remote terminal then responds with a DTR (Data Terminal Ready) signal. When LK8 is set to position B, the DTR signal drives the $\overline{CS}$ input of the AD7701 and initiates the data transfer.

IC11, the MC145406 RS-232 Interface chip converts the TTL levels from the AD7701 to the required ±5 V signals to drive the remote terminal or takes the ±5 V signals and converts them into TTL levels. The RS-232 interface on the evaluation board is fully functional but does not comply with all the requirements of the EIA RS-232 standard. When the MC145406 receiver/driver chip is operated from ±5 V supplies rather than ±6 V, its driver output swing is reduced below the EIA specified limits. In practical applications this signal swing limitation only reduces the length of cable the circuit is capable of driving.

Parallel Port (AD7701 Only)

The 16-bit Parallel Port is available at two sockets on the board, SKT3 and SKT4. The Eurocard Connector, SKT3, allows this board to be connected directly to the evaluation board for the ADSP-2100 Digital Signal Processor, which is available from Analog Devices. SKT4 provides a general purpose socket for interfacing to any other parallel systems. The pinout for SKT3 and SKT4 is given in Figure 4.

When using the parallel interface, the AD7701 should be set up to operate in the SSC mode (LK5 set at SSC). Activation of the $\overline{CS}$ input on the AD7701 will determine the rate at which it attempts to update the output shift registers IC9 and IC10. $\overline{CS}$ can be controlled by the decimation counter output (LK8 in position A) or remotely at

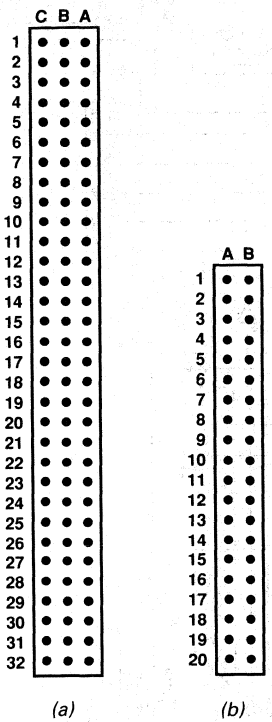

Figure 4. Pin Configurations for SKT3 (a) and SKT4 (b)

SKT1. If it is controlled remotely, LK4 must be in position 0 to ensure that the output register update and the remote $\overline{CS}$ occur at the same rate. When $\overline{CS}$ goes low, the 16-bit output data is clocked into the two 8-bit shift registers, IC9 and IC10. These are set up for shift right operation (S0 = 1, S1 = 0), and their three-state outputs are available at SKT3 and SKT4. A rising edge on Pin 2 of IC6 indicates that the output registers have been updated. This signal clocks the D-type flip-flop, IC8b to set its $\overline{Q}$ output low and indicate to the remote reading device that the registers have been updated and can be read. This is called the $\overline{PDR}$ (Parallel Data Ready) signal. The $\overline{PDR}$ signal also controls the S0 input of IC9 and IC10. When it goes low, the shift register outputs are held until a Data Acknowledge (DACK) signal is received from the remote device and resets flip-flop IC8b. In the ADSP-2100 interface, $\overline{PDR}$ drives an interrupt line in the processor. The service routine reads the 16-bit data by bringing $\overline{PCS}$ low and A0 high. When the data has been read it sends back a DACK signal to reenable the output shift registers. Note that LK7 allows the DACK signal to be either active high or low. If DACK is true when high, then LK7 should be in position A; whereas if DACK is true when low, LK7 should be in position B.

IC9 and IC10 are set up to be read in 16-bit parallel fashion with the low byte in IC9 and the high byte in IC10. They can also be configured to be read separately as two 8-bit bytes on an 8-bit bus. To do this, the Byte Wide Jumpers must be soldered in place. In addition, for proper "one byte at a time" address selection, the connections at Pin 10 of IC6 on the circuit board must be altered to control which register is to be read when A0 is a 1 and which is to be read when it is 0. This can be seen in Figure 5. For example, if the most significant byte is to be read when A0 is 0, IC10 should be enabled and IC9 disabled. The link between 4 and 5 on the board must be broken, and 1 must be joined to 4.

Link/ Switch	Position	Function
LK1	INT	On-board oscillator selected as input to clock generator section
LK2	0	IC2 counter divider ratio set at 1.
LK3	NC*	Baud rate selector not used.
LK4	0	Decimation counter divide ratio set at 1.
LK5	SSC	Synchronous self clocking interface mode selected.
LK6	NC*	Not used. SW1 sets the calibration mode.
LK7	A	DACK signal set to be active low.
LK8	A	Decimation counter output controls the $\overline{CS}$ input of the AD7701/AD7703..
LK9	B (NC*)	SCLK not controlled by the baud clock.
SW1-1	ON	Sets SC1 input at 0.
SW1-2	ON	Sets SC2 input at 0. With SC1, SC2 both set to 0, the device calibration mode is self-calibration of zero and full scale.
SW1-3	ON	Unipolar input range selected.
SW1-4	OFF	Normal mode (i.e., not low power Sleep mode) selected.

*NC = No Connection

Table VIII. Link and Switch Settings

SETUP CONDITIONS

In order to initially set up and verify that the evaluation board is functional, it is useful to set all the link and switch settings for a particular mode of operation. The settings, given in Table VIII, will set the board for the following: On-board 4.9152 MHz clock; unipolar analog input range; Self-calibration mode; 4.8 kHz output update rate; SSC Interface Mode. To obtain specified offset and full-scale performance, the CAL switch, SW2 should be depressed immediately after power-up. Figure 6 gives the link and switch positions on the board in order to ease setup.

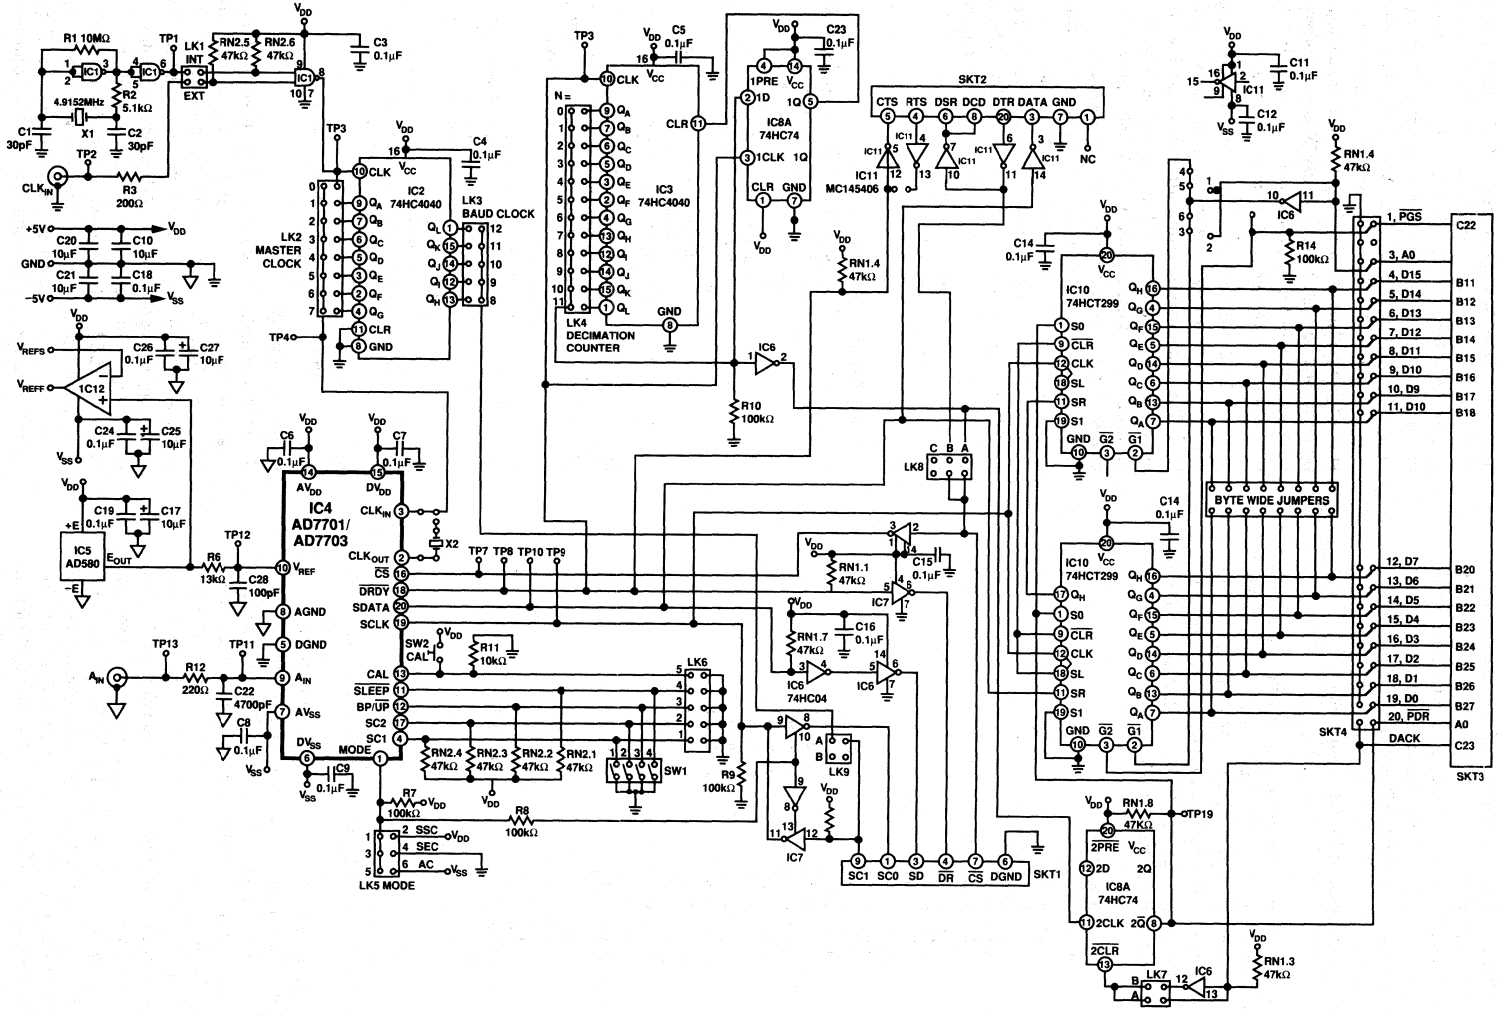

Figure 5. Evaluation Board Circuit Diagram

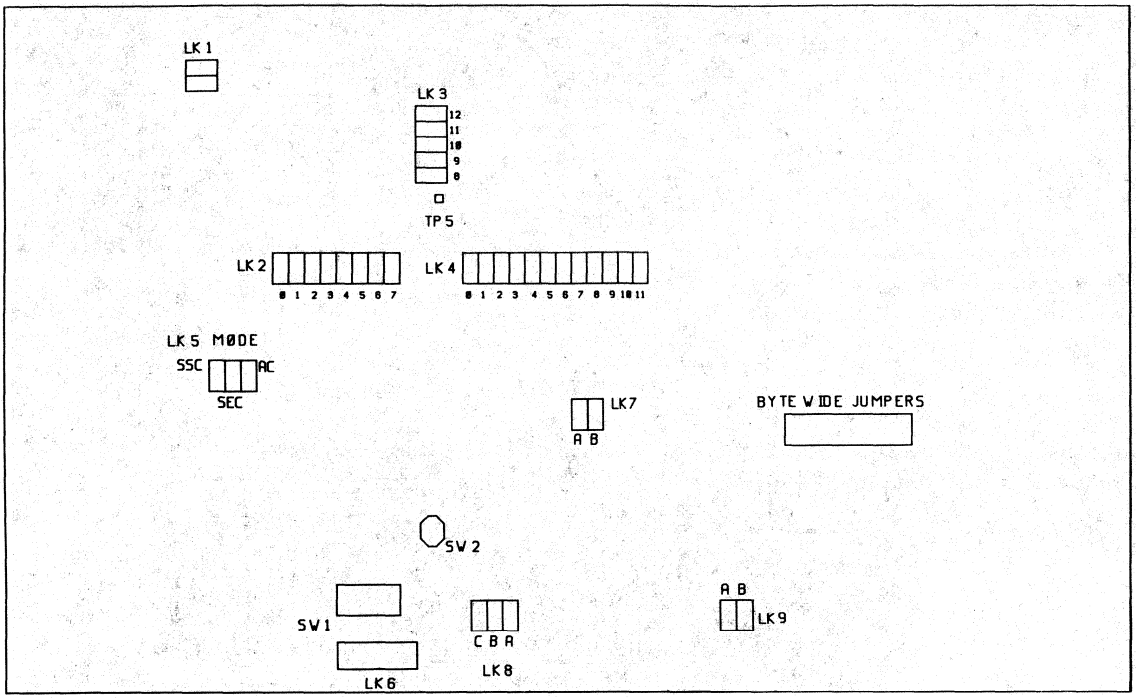

Figure 6. Positioning of Links and Switches on the Evaluation Board

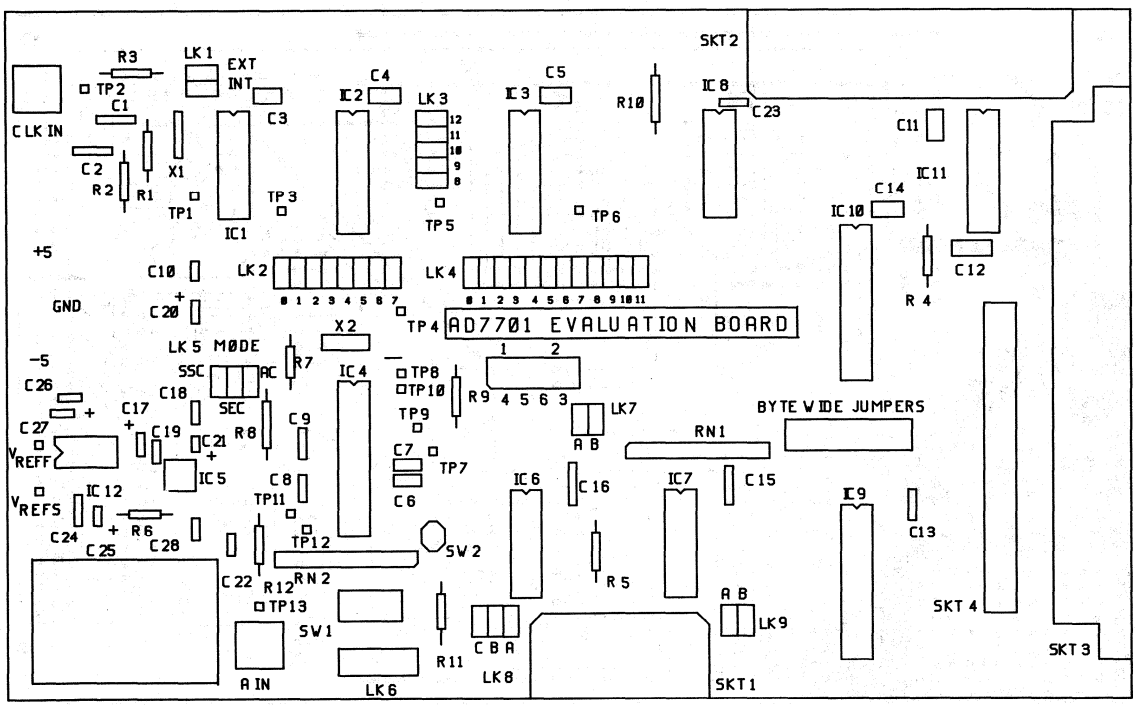

Figure 7. AD7701/AD7703 Evaluation Board Component Layout

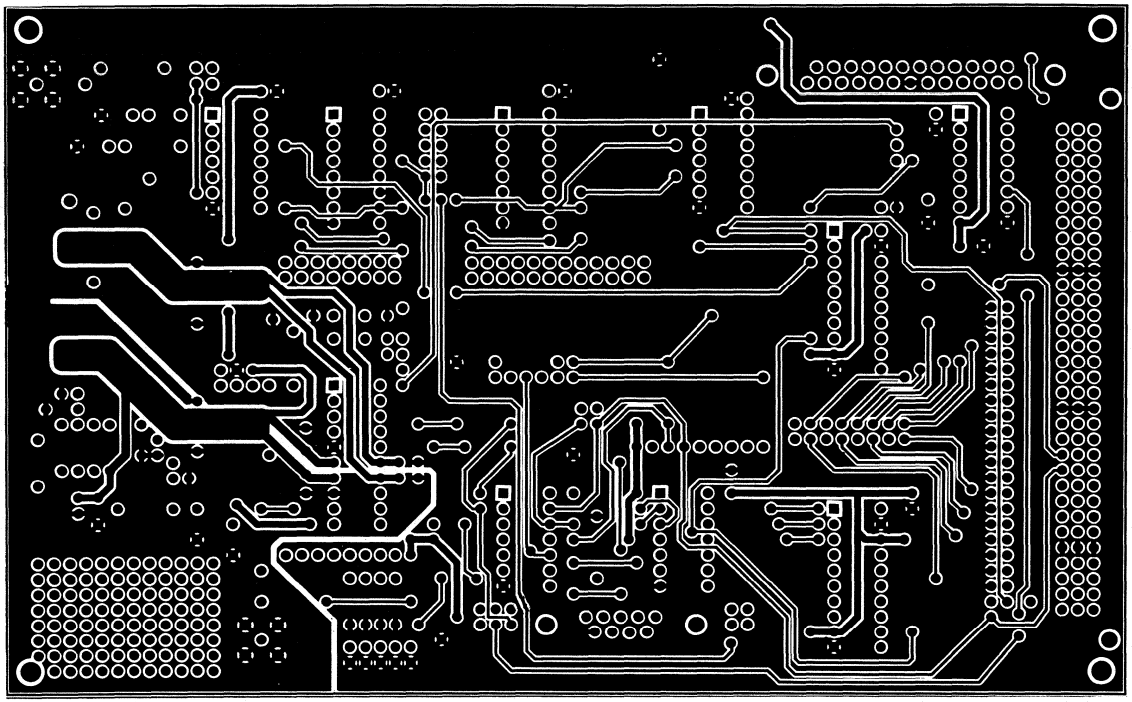

Figure 8. AD7701/AD7703 Evaluation Board Component Side Layout

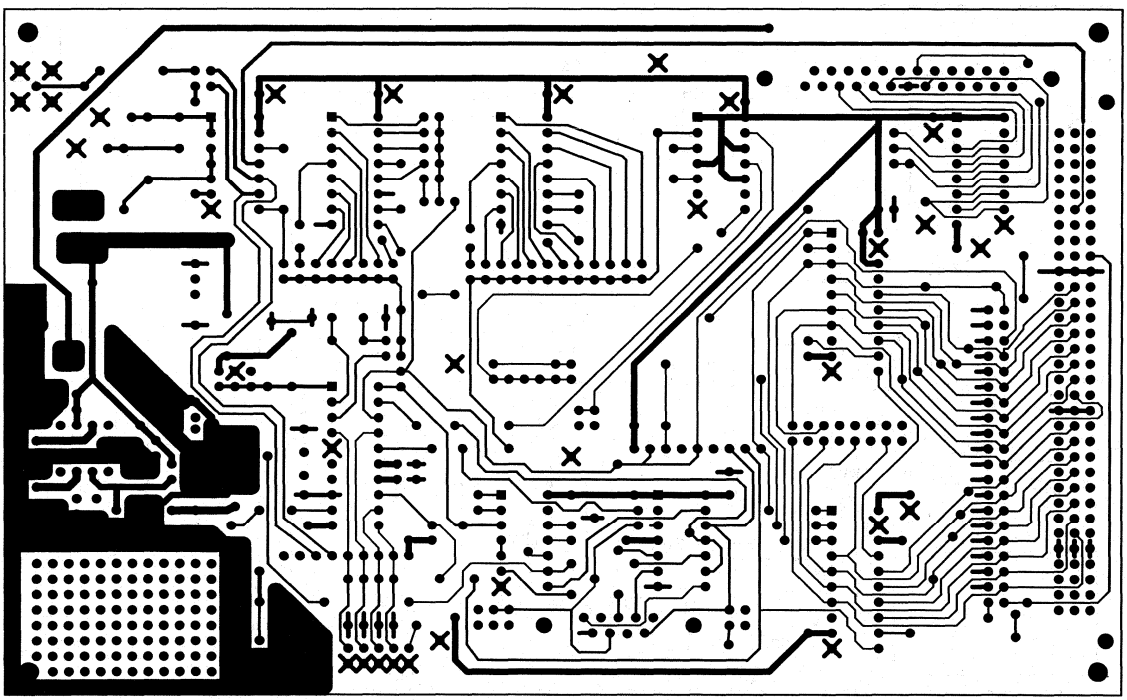

Figure 9. AD7701/AD7703 Evaluation Board Solder Side Layout

Audio Products
Contents

		Page
AN-111	A Balanced Summing Amplifier	4–3
AN-112	A Balanced Input High Level Amplifier	4–5
AN-113	An Unbalanced, Virtual Ground Summing Amplifier	4–7
AN-114	A High Performance Transformer-Coupled Microphone Preamplifier	4–9
AN-115	Balanced Low Noise Microphone Preamplifier Design	4–11
AN-116	AGC Amplifier Design with Adjustable Attack and Release Control	4–13
AN-121	High Performance Stereo Routing Switcher	4–15
AN-122	A Balanced Mute Circuit for Audio Mixing Consoles	4–19
AN-123	A Constant Power "Pan" Control Circuit for Microphone Audio Mixing	4–21
AN-124	Three High Accuracy RIAA/IEC MC and MM Phono Preamplifiers	4–23
AN-125	A Two-Channel Dynamic Filter Noise Reduction System	4–29
AN-127	An Unbalanced Mute Circuit for Audio Mixing Channels	4–31
AN-128	A Two-Channel Noise Gate	4–33
AN-129	A Precision Sum and Difference (Audio Matrix) Circuit	4–35
AN-130	A Two-Band Audio Compressor/Limiter	4–37
AN-131	A Two-Channel VCA Level (Volume) Control Circuit	4–39
AN-133	A High Performance Compandor for Wireless Audio Systems	4–41
AN-134	An Automatic Microphone Mixer	4–45
AN-135	The Morgan Compressor/Limiter	4–47
AN-136	An Ultralow Noise Preamplifier	4–55
AN-211	The Alexander Current-Feedback Audio Power Amplifier	4–57
AN-232	Bootstrapped IC Substrate Lowers Distortion in JFET Op Amps	4–73

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

A Balanced Summing Amplifier

The summing amplifier circuit shown in Figure 1 represents an excellent virtual ground summing amplifier using a balanced differential design that includes extremely low noise and wide bandwidth as featured by the SSM-2015. Any size audio mixing system can benefit from balanced virtual node mixing (summing). The low cost and exceptional performance of this design can be incorporated in any system with balanced or mixed balanced and unbalanced input sources.

IC_2, the PMI OP-41, serves as a DC servo-amplifier that is referenced to signal ground. The circuit functions as an integrator with a long time constant that retains the integrity of low frequency audio signals down to 5Hz, and keeps $e_{OUT} = 0V_{DC}$, ($\pm 10mV_{DC}$). The OP-41 is a FET input amplifier, with low input offset voltage (V_{OS}) and high input impedance. Although many low performance JFET/CMOS operational amplifiers can be employed, the summing output V_{DC} is a function of the servo's input offset voltage and its temperature coefficient ($\Delta V_{OS}/\Delta T$), which must be kept low for direct coupled summing applications.

In this design, the following facts predominate: $e_S = 0$, and $i_S = 0$. e_{IN} is the algebraic sum of the input(s) e_{IN1}, e_{IN2}, e_{IN3}, e_{INn} and etc. $e_{OUT} = [e_{IN1}(R_F/R_{IN1}) + e_{IN2}(R_F/R_{IN2}) + e_{IN3}(R_F/R_{IN3}) + e_{INn}(R_F/R_{INn})]$, etc. The input impedance therefore equals R_{IN1}, R_{IN2}, R_{IN3}, R_{INn}, etc. The overall gain of the circuit is set by R_F, and the gain of the individual channels can be adjusted independently by the values of R_{IN1}, R_{IN2}, R_{IN3}, R_{INn}, etc.

For individual source input(s), gain is $A_O = R_F/R_{IN}$.

The circuit configuration produces linear signal mixing at the summing nodes (IC_1 pins 10,11), whereas $e_S = 0$; therefore, no interaction occurs between the source inputs. Owing to the fact that the SSM-2015 is a bipolar transistor device, the noise is low ($1.3nV/\sqrt{Hz}$). The commonly used values of 10kΩ for R_F and R_{IN} are optimal for both minimum noise and previous stage loading, eliminating the need for buffer amplifiers and their noise contribution.

The input common-mode rejection for the SSM-2015 is typically 100dB as a result of true differential input topology. The differential thermal noise and DC offset drift is nearly eliminated by the common substrate construction employed. To exploit the high CMR of the SSM-2015, all signal resistors should be matched resistor networks or should employ 0.5% or better resistor tolerances.

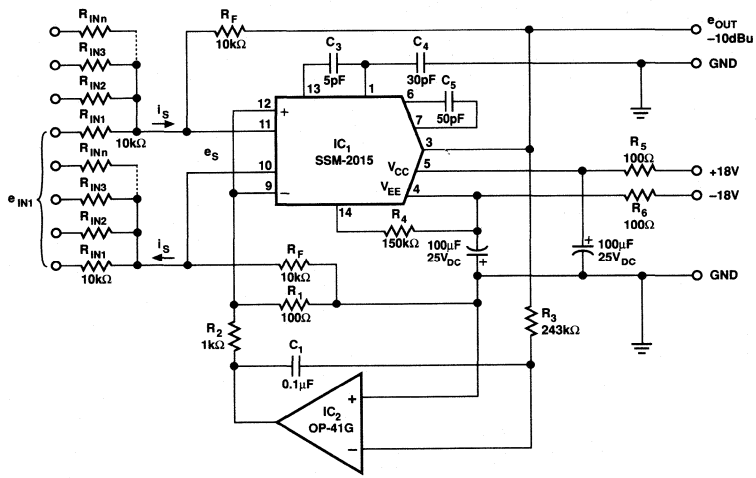

FIGURE 1

The output circuit topology of the SSM-2015 is complementary bipolar producing overall performance of 6V/µs slew rate, and is able to drive a 2kΩ unbalanced load. The circuit described can be directly coupled, eliminating coloration and distortion associated with coupling capacitors. The circuitry following this amplifier could be AC (capacitor) coupled if the DC servo IC_2 offset voltage of $\pm 10mV_{DC}$ is objectionable.

Audio performance challenges the best test equipment that might be used to measure high performance analog designs. For example: worst case THD for this circuit measures less than 0.008%, and IMD less than 0.02% over a band-width of 10Hz to 20kHz. See Table 1 for more performance details.

TABLE 1: Circuit Performance Specifications

Frequency Response (dB 20Hz to 20kHz)	±0.02
S/N Ratio @ +23dBu	103dB
TMD + Noise (@ +23dBu 20Hz to 20kHz)	0.008%
IMD (@ +23dBu SMPTE 60Hz & 4kHz, 4:1)	0.015%
CMRR (60Hz)	100dB
Slew Rate	6V/µs
Output Voltage (2kΩ load)	+23dBu or $11V_{RMS}$

)NE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

A Balanced Input High Level Amplifier

he balanced amplifier in Figure 1 utilizing the SSM-2015 feaures adjustable gain and can accept nominal audio signals om –27.5dBu to +0dBu with more than 30dB of headroom. he input terminals can tolerate common-mode voltages of 30 olts peak-to-peak. Common-mode noise rejection is greater an 100dB at 1,000Hz, while the EIN (Equivalent Input Noise) s a low –124dBu.

he IC_1 amplifier circuit is gain adjustable, and the design tilizes a 12-position switch with 2.5dB steps. Other resistor alues can be calculated to accommodate custom gain requirements.

C_1 is PMI's SSM-2015 true differential input IC amplifier. Its nput circuit utilizes two identical low noise bipolar transistors, vith access to the emitters that provide the gain adjustment. R_G R_{14} through R_{24}) sets the amplifier's gain using the equation:

$$\text{Gain} = 3.5 + \left(\frac{20 \times 10^3}{R_G}\right) \text{ for } R_8,\ R_{18} = 10.0\text{k}\Omega$$

The emitter feedback design exhibits both minimum noise and naximum common-mode rejection while retaining a very high input impedance. The output circuit topology is complementary bipolar producing 6V/μs slew rate, and is able to drive a 2kΩ unbalanced load. The circuit described can be directly coupled eliminating the distortion associated with coupling capacitors. Circuitry following this amplifier could be AC (capacitor) coupled if input normal-mode DC voltages are expected at the input of this circuit. Worst case THD measures less than 0.008%, and IMD less than 0.015%.

Input components C_1, C_2, R_1, and R_2 constitute a single pole low-pass filter that limits the input voltage slew rate, curbs interface transient intermodulation distortion, and keeps the amplifier from slewing. The input network has little effect on phase response within the pass band of 20Hz to 20kHz. To maintain high frequency common-mode performance, capacitors C_1 and C_2 should be matched for 1% tolerance.

For an output voltage of –10dBu, the balanced input amplifier circuit has an input sensitivity range of 0.0dBu to –27.5dBu. The common-mode voltage trim is included for maximizing application common-mode noise reduction and also allows the use of low cost components.

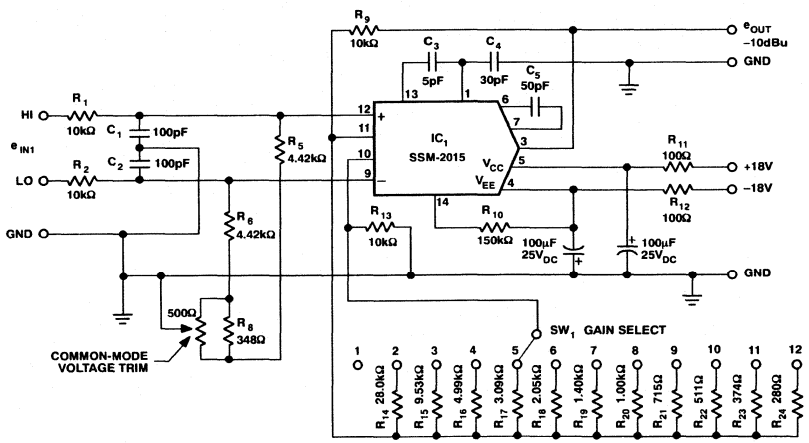

FIGURE 1

SW	G_{dB}	e_{IN}(dB)	R_G	VALUE (Ω)
1	10.0	0	R	ω
2	12.5	–2.5	R_{14}	28.0k
3	15.0	–5.0	R_{15}	9.53k
4	17.5	–7.5	R_{16}	4.99k
5	20.0	–10.0	R_{17}	3.09k
6	22.5	–12.5	R_{18}	2.05k
7	25.0	–15.0	R_{19}	1.40k
8	27.5	–17.5	R_{20}	1.00k
9	30.0	–20.0	R_{21}	715
10	32.5	–22.5	R_{22}	511
11	35.0	–25.0	R_{23}	374
12	37.5	–27.5	R_{24}	280

Specific gain can be calculated from the equation:

$$Gain_{dB} = 20 \log \left[3.5 + \left(\frac{20 \times 10^3}{R_G}\right)\right] \text{ for } R_8, R_{18} = 10.0k\Omega$$

TYPICAL APPLICATIONS

This design is ideal for use as the input amplifier in audio distribution amplifiers, for balanced input audio routing switchers, as the input buffer ahead of the A-to-D codec in digital recording and mixer equipment, or for the low noise high level input of mixing consoles.

TABLE 1: Circuit Performance Specifications

Frequency Response (dB 20Hz to 20kHz)	±0
S/N Ratio @ +23dBu	103
THD + Noise (20Hz to 20kHz) @ +23dBu	0.008
IMD (SMPTE 60Hz & 4kHz, 4:1) @ +23dBu	0.015
CMRR (60Hz)	100
Slew Rate	6V/
Output Voltage (2kΩ load)	+23dBu or 11V_{RI}

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

An Unbalanced, Virtual Ground Summing Amplifier

The summing amplifier circuit shown in Figure 1 represents a splendid unbalanced virtual ground summing amplifier. The design utilizes the SSM-2134, PMI's superior version of the popular NE5534 bipolar operational amplifier. This low noise amplifier can now be implemented where most equipment manufacturers use FET input operational amplifiers. The circuit described features reduction in noise, temperature, and input impedance effects on static condition output voltages, and elimination of unity gain instability.

The SSM-2134 helps reduce wide-band noise figures by 3dB to 10dB, while improving the frequency and phase response performance. Only minimal value compensation (C_2) is required for the SSM-2134. In the feedback loop, C_1 improves stability while keeping the slew rate at 10V/μs and bandwidth greater than 100kHz.

In this circuit, note the following design facts: e_{OUT} is the algebraic sum of the input voltage(s) e_{IN1}, e_{IN2}, e_{IN3}, e_{INn} and etc. $e_{OUT} = (-) [e_{IN1} (R_F/R_{IN1}) + e_{IN2}(R_F/R_{IN2}) + e_{IN3}(R_F/R_{IN3}) + e_{INn}(R_F/R_{INn})]$, etc. The individual input impedance therefore equals R_{IN1}, R_{IN2}, R_{IN3}, R_{INn}, etc. The overall gain of the circuit is set by R_F, and the gain of the individual channels can be adjusted independently by the values of R_{IN1}, R_{IN2}, R_{IN3}, and R_{INn}.

For individual source input(s), voltage gain = R_F/R_{IN}.

The circuit configuration produces linear signal mixing at the summing node (common tie point C_1, R_F, R_{IN1}, R_{IN2}, R_{IN3}, and R_{INn}), whereas e_S = 0, there is no interaction between the source inputs. Owing to the fact that SSM-2134 is a bipolar device, noise is low (2.8nV/$\sqrt{Hz}$). The commonly used values of 10kΩ for R_F and R_{IN} are optimal for both minimum noise and previous stage loading, eliminating the need for buffer amplifiers and their noise contribution.

This design produces maximum amplifier bandwidth with unconditional circuit stability for both input and output impedance (reactive or not) variations.

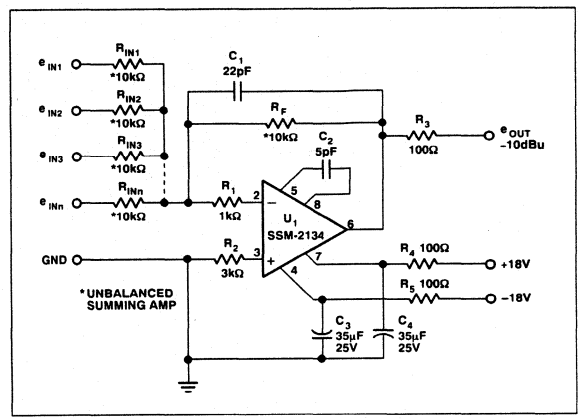

FIGURE 1

TABLE 1: Circuit Performance Specifications

Frequency Response (20Hz to 20kHz)	±0.02dB
S/N Ratio (@ +23dBu)	104dB
THD + Noise (@ +23dBu, 20Hz to 20kHz)	0.007%
IMD (SMPTE 60Hz and 4kHz, 4:1)	0.015%
Slew Rate	10V/μs
Output Voltage (2kΩ load)	+23.3dBu or 11.3V_{RMS}

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

A High Performance Transformer-Coupled Microphone Preamplifier

The SSM-2015 or SSM-2016 low noise differential amplifier is utilized in a transformer-coupled microphone preamplifier. The circuit shown in Figure 1 represents a microphone preamplifier with high performance, wide dynamic range, and ultra low noise. The design features a Jensen transformer-coupled preamplifier circuit with balanced/floating input, 1500Ω input loading, three step input attenuator, phantom microphone powering, and twelve amplifier gain choices. Although the design shown includes a twelve position gain selector, fixed gain applications can utilize the component value calculations and formula provided.

The design provides microphone input loading of 1500Ω. Input loading is capacitive reactive, and at higher input voltage frequencies, the low-pass network and transformer characteristics help attenuate unwanted normal-mode RF and ultrasonic voltages that might be present at the input terminals.

The input circuit contains a three-position input attenuator used to optimize source levels versus amplifier headroom. As usual, it's a compromise of headroom and preamplifier signal-to-noise. The attenuation is 0dB, –10dB, and –20dB while maintaining an input impedance of 1500Ω.

A phantom microphone powering circuit is included for condenser microphones that require 24 to 48 volts DC power.

The common-mode voltage range is limited only by the transformer's primary-to-shield breakdown voltage. Common-mode rejection is a product of the primary-to-secondary isolation and provides detachment of the microphone wiring environment. Although the balanced single-pole low-pass filter at the input terminals provides protection from radio frequency interference, this network, along with the capacitive effect of the primary winding to the grounded shield, plus the phantom powering resistors present a circuit path for external RF voltages to enter the preamplifier's circuit ground. A carefully planned single point (power supply) grounding, and the true balanced and differential input topology of the SSM-2015/2016 amplifier will eliminate unwanted external noise signals.

The network composed of R_4 and C_6 at the transformer secondary serves two functions. It minimizes transformer rising secondary winding signal amplitude with rising input frequency and deters secondary ringing, while helping to prevent amplifier input slewing. The SSM-2015/2016 differential input improves transformer performance substantially as compared with the conventional unbalanced design.

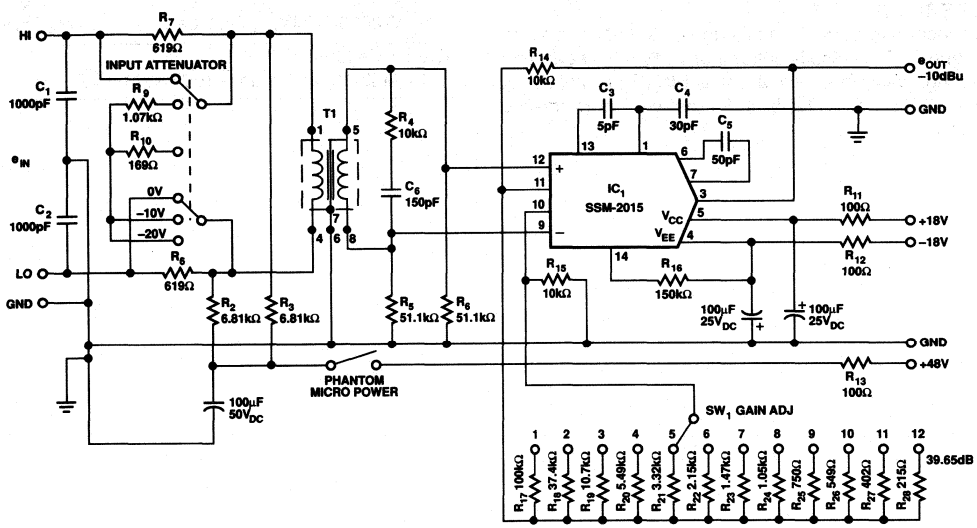

FIGURE 1

The circuit design incorporates a gain switch with twelve (12) calculated gain settings. The Jensen transformer, model JE-110K-HPC used in this application has a voltage gain of 17.9dB. For an output voltage of –10dBu, the microphone amplifier circuit has an input sensitivity range of –65dBu to –17.5dBu, with a typical output headroom of 33dB. The preamplifier circuit shown is gain adjustable from 9.6dB to 39.6dB in 2.5dB steps.

PMI's SSM-2015/2016 input circuit utilizes two identical low noise bipolar transistors, with access to the emitters, that provide the gain adjustment. The output circuit topology is complementary bipolar producing 6V/µs (2015) and 10V/µs (2016) slew rate into a 2kΩ unbalanced load.

R_G (R_{17} through R_{28}) sets the amplifier gain using the equation:

$$V_G = 3.5 + \left(\frac{20 \times 10^3}{R_G}\right)$$

for R_{14}, and R_{15} = 10.0kΩ.

SW	G_{dB}	*e_{IN}(dB)	R_G	VALUE (Ω)
1	9.6	–37.5	R_{17}	100k
2	12.1	–40.0	R_{18}	37.4k
3	14.6	–42.5	R_{19}	10.7k
4	17.1	–45.0	R_{20}	5.49k
5	19.6	–47.5	R_{21}	3.32k
6	22.1	–50.0	R_{22}	2.15k
7	24.6	–52.5	R_{23}	1.47k
8	27.1	–55.0	R_{24}	1.05k
9	29.6	–57.5	R_{25}	750
10	32.1	–60.0	R_{26}	549
11	34.6	–62.5	R_{27}	402
12	39.6	–65.0	R_{28}	215

*Input attenuator set to the 0dB position.

Unspecified overall circuit gain can be calculated from the equation:

$$G_{dB} = 20 \log\left[3.5 + \left(\frac{20 \times 10^3}{R_G}\right)\right] + 17.9$$

TYPICAL PERFORMANCE

Frequency response versus amplitude is ±0.2dB from 20 to 20,000Hz, and THD + noise is better than 0.03% over gain and frequency range described, with a typical EIN (Equivalent Input Noise) of –127dBu. See Table 1 for detailed performance specifications.

For applications where additional headroom is required, th SSM-2016 should be used. The SSM-2016 can be powere with up to ±36V_{DC} rails and drive 600Ω loads. If ±24V_{DC} rails ar used, headroom increases to 35.7dB (typically), while the EI remains at –127dB. As a consequence of the increased powe supply voltage, the SSM-2016 package power dissipation wi typically be 600mW with ±24V_{DC} rails (no signal), and will ris to 725mW with worst case signal conditions into 600Ω load.

For ±36V_{DC} power rails, although the headroom increases t 39.3dB, the SSM-2016 will dissipate 1.2 watts with no signa applied, and 1.5 watts worst case signal conditions into 600 load. Therefore, IC package cooling should be taken int consideration. Please see the SSM-2016 data sheet for IC pin out connections and recommended compensation capacito values. All other circuit component values shown here apply.

The transformer-coupled microphone preamplifier circuit de scribed above demonstrates robust, real-world usage refine ments, along with most operational features required by equip ment designers to deliver the highest performance. It will handle the most hostile microphone environments without distress t either the circuit or the user.

TABLE 1: Circuit Performance Specifications

Frequency Response (20Hz to 20kHz, –60dBu, 50dB gain)	±0.15dB
THD + Noise (20Hz to 20kHz, –60dBu, 50db gain)	0.045%
IMD (+23dBu, SMPTE 60Hz and 4kHz, 4:1)	0.05%
EIN (Equivalent Input Noise, 150Ω source)	–127dB
Input Impedance (20Hz to 5kHz)	1500Ω
Source Impedance	150Ω
CMR at 1kHz (common-mode rejection at 1kHz)	120dB
CMVR (common-mode voltage range)	±150V_{DC}
Slew Rate (overall circuit)	6V/µs
Gain Range (overall circuit)	17.5dB to 36dB
Output Voltage	
SSM-2015 (±18V_{DC}, 2kΩ load)	+23dBu or 11V_{RMS}
SSM-2016 (±24V_{DC}, 2kΩ load)	+25.7dBu or 15V_{RMS}
Output Headroom (SSM-2015, 2kΩ load, –10dBu nominal)	33dB

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Balanced Low Noise Microphone Preamplifier Design

The SSM-2015 differential amplifier is utilized in a transformerless, active-balanced input amplifier. The circuit shown in Figure 1 provides a microphone preamplifier design with excellent performance and low noise. The design features a transformerless preamplifier circuit with true-balanced input, 1500Ω input loading, phantom microphone powering, and high common-mode rejection. The design shown also includes a twelve position gain selector, or for fixed gain usage, component value calculations.

The design includes microphone input loading of 1500Ω, but the load resistor can be changed to accommodate other applications. Input loading is capacitive reactive at higher frequencies to attenuate unwanted RF and ultrasonic voltages at the input terminals.

The phantom microphone powering circuit provides power for condenser microphones that require 24 to 48 volts DC. The zener diodes CR_1, CR_2, CR_3, and CR_4 protect the input transistors of the SSM-2015 when connecting the microphone to the preamplifier circuit.

The common-mode voltage range is ±5.5 volts. Its common-mode rejection is optimized for most applications by the true-balanced and differential input topology of the SSM-2015. A balanced single pole low-pass filter at the input terminals provides protection for the circuit from radio frequency interference and prevents slewing of the SSM-2015 amplifier. The output circuit topology is complementary bipolar producing 6V/μs slew rate, and able to drive a 2kΩ unbalanced load.

The circuit design incorporates a gain switch with twelve (12) calculated gain settings. For an output voltage of –10dBu, the microphone amplifier circuit has an input sensitivity range of –65dBu to –27.5dBu, and an output headroom of 33dB. The overall circuit gain is adjustable from 27.5dB to 55dB in 2.5dB steps.

SW	G_{dB}	e_{IN}(dB)	R_G	VALUE (Ω)
1	27.5	–37.5	R_{15}	1.00k
2	30	–40	R_{16}	715
3	32.5	–42.5	R_{17}	511
4	35	–45	R_{18}	374
5	37.5	–47.5	R_{19}	280
6	40	–50	R_{20}	205
7	42.5	–52.5	R_{21}	154
8	45	–55	R_{22}	115
9	47.5	–57.5	R_{23}	86.6
10	50	–60	R_{24}	63.4
11	52.5	–62.5	R_{25}	47.5
12	55	–65	R_{26}	35.7

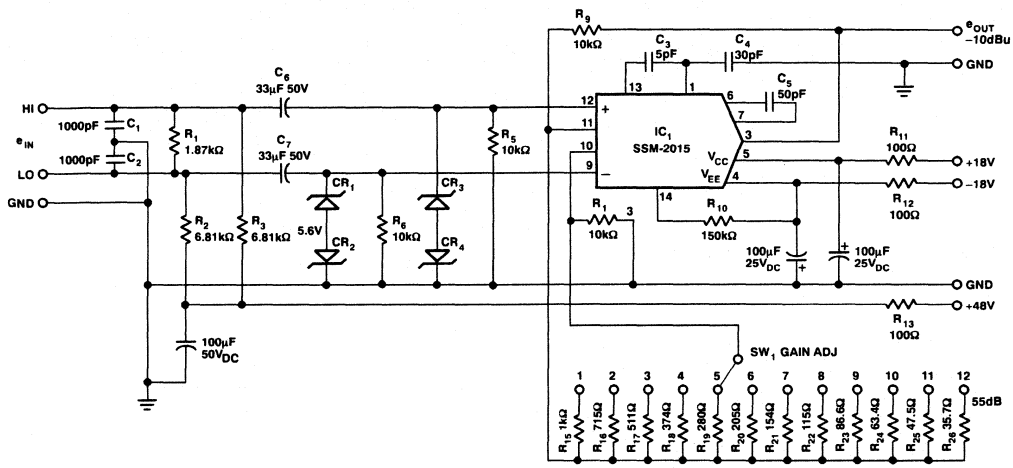

FIGURE 1

SSM-2015 input circuitry utilizes two identical low noise bipolar transistors, with access to the emitters that provide the gain adjustment. R_G (R_{15} through R_{26}) sets the amplifier's gain using the equation:

$$\text{Gain} = 3.5 + \left(\frac{20 \times 10^3}{R_G}\right) \quad \text{for } R_9, \& R_{13} = 10.0k\Omega$$

Unspecified gain can be calculated from the equation:

$$\text{Gain}_{dB} = 20 \log \left[3.5 + \left(\frac{20 \times 10^3}{R_G}\right)\right]$$

The frequency response amplitude is ±0.1dB from 20 to 20,000Hz, and THD + noise of better than 0.03% over the gain range described with a typical EIN (Equivalent Input Noise) of –124dBu.

The transformerless microphone preamplifier circuit described above demonstrates real-world usage refinements and includes most operational features required by equipment designers.

TABLE 1: Circuit Performance Specifications

Frequency Response (20Hz to 20kHz)	±0.1dB
THD + Noise (@ +23dBu, 20Hz to 20kHz)	0.03%
IMD (@ +23dBu, SMPTE 60Hz & 4kHz, 4:1)	0.05%
EIN (Equivalent Input Noise, 150Ω source)	–124dB
CMR (Common-Mode Rejection at 1kHz)	105dB
Slew Rate	6V/µs
Output Voltage (2kΩ load) +23dBu or	$11V_{RMS}$
Output Headroom (2kΩ load, –10dBu nominal)	33dB

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

AGC Amplifier Design with Ajustable Attack and Release Control

he automatic gain control (AGC) amplifier described below nd shown in Figure 1, features selectable gain reduction ompression ratios and time domain adjustable AGC attack nd release. This design employs the SSM-2013 VCA , SSM-110 precision level detector, two SSM-2134 low noise op mps, and an OP-215 FET input dual op amp.

he design features an inverting or noninverting input buffer mplifier, a voltage controlled amplifier with adjustable attack nd recovery characteristics, driven by a true RMS level detec-or. Additionally, it provides selectable gain reduction compres-ion, adjustable AGC output level, and maximum gain limit ontrols. Signal-to-noise ratio is better than 100dB and the tMS level detector allows the AGC amplifier to operate trans-arently throughout the audio spectrum. The gain recovery is near and time adjustable, and has maximum gain limiting gating) to preclude input source noise floor rise.

he input circuit includes a line level (–10dBu to 0dBu) buffer mplifier, that accepts inverting or noninverting inputs with reater than 10kΩ loading impedance. The buffer also isolates ne input source from the compressor gain reduction ratio elector, and limits step function slewing voltage.

The six-position gain reduction selector that follows the input amplifier provides adjustable compression that smoothes the AGC action. Six GAIN REDUCTION slope ratios of 2 to 22 can be selected, thus reducing the irritating *"hole producing and pumping"* character of most AGC circuits. The SSM-2013 VCA is chosen for its predictable behavior and its high performance. The dynamic range exceeds 94dB over the frequency range 20Hz to 20kHz. Over this frequency range, the amplifier achieves typically less than 0.01% THD + noise, and 0.03% IMD.

The SSM-2110's precision rectifier circuit produces the true RMS output that comprises a level detector. It results in a consistent and precise AGC action that retains good signal dynamics while leveling the input signal. It responds to the audio signal power density in a manner similar to human hearing.

Following the precision RMS rectifier is the VCA control voltage conditioning circuits. Constructed around U_6 (OP-215), the FET-input amplifier forms an integrator while the other amplifier provides the VCA control port buffer. The AGC output level is set by the rectified signal voltage compared to the reference voltage from the OUTPUT LEVEL control.

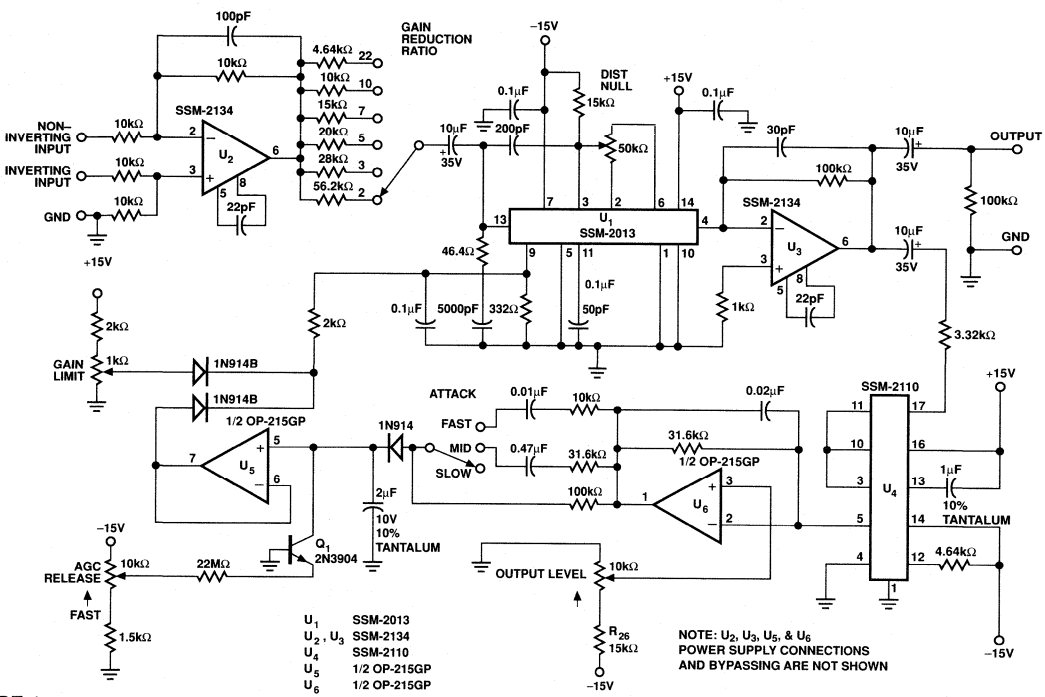

IGURE 1

The AGC attack and compression response is altered by adjusting the integrator charging time constant or integrator wave shape current. The three-position ATTACK switch allows selection of fast, medium, and slow compression and AGC response. When the slow position is selected, an insignificant amount of compression will take place, while fast and medium combine compression with the AGC action. The AGC release rate is controlled by a constant current discharge of the integrator capacitor. The recovery time constant is linear and adjusted by changing the integrator discharge current supplied by Q_1 and regulated by the RELEASE rate control.

The SSM-2134 has been selected for its low noise and high performance characteristics. The AGC circuit described is of the feedback class, that is, the level detecting rectifier follows the voltage controlled amplifier stage. This class of AGC circuit combined with the complementary gain reduction compression, driven by RMS level detection, and adjustable attack and release AGC action, allows this circuit to be as unobtrusive or as conspicuous as desired.

The flexibility and high performance of this design, along with the simplicity and cost effectiveness, allows this design to be suitable for incorporating in mixing console designs, or in stand-alone products.

TABLE 1: Circuit Performance Specifications

Input Voltage Range (Nominal for 0dBu Out)	–26dBu to +10dB (6mV to 2.45V_{RMS}
Rectifier Type	RM
AGC Amplifier Class	Feedbac
Attack Time	20 to 200m
Recovery Time (6dB)	3 to 32 SE
VCA Feedthrough (Trimmed)	–100d
Gain Limit Range (Gain Reduction 22)	–26dBu to –12dB
Frequency Response (20Hz to 20kHz)	±0.2d
S/N Ratio (@ ±10dB Gain)	106d
THD + Noise (@ +23dBu, 20Hz to 20kHz)	0.01
IMD (@ + 23dBu, SMPTE 60Hz & 4kHz, 4:1)	0.02
Output Voltage Slew Rate	6V/µ
Output Voltage (2kΩ Load)	+22dBu or 10V_{RM}

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

High Performance Stereo Routing Switcher

The SSM-2402 Dual Audio Switch comprises the nucleus for this 16 channels-to-one high performance stereo audio routing switcher, which features negligible noise and low distortion over the frequency range of 20Hz to 20kHz. This performance is achieved even while driving 600Ω loads at signal levels up to +30dBu.

The SSM-2402 affords a much simplified electrical design and printed circuit board layout, along with reduced manufacturing cost, when compared with discrete JFET circuits of similar performance. The electrical performance of the design described is vastly superior to CMOS switch designs, which are more prone to failure resulting from electrical static discharge.

The switching control of the SSM-2402 may be activated by conventional mechanical switches or 5 volt TTL or CMOS logic circuits. The application shown utilizes a simple mechanical control switch for illustration purposes only. Many diverse X/Y control schemes, destination control, or computer controlled designs can be utilized.

The "T" configuration of the SSM-2402 switch provides excellent ON-OFF isolation. The SSM-2402 also features 7ms ramped turn on and 4ms ramped turn off for click-free switching. Additionally, the switch has a break-before-make switching sequence. Both features become significant in large audio switching systems where the audio path can pass through multiple switching elements. Such controlled switching is very important in large systems used in broadcast program switching or in production work.

The application circuit design also employs the SSM-2015 balanced input amplifier (Figure 1). The input impedance is high (≈100kΩ), balanced or unbalanced. The input circuit incorporates a single pole RFI filter with a cutoff frequency set at 145kHz. In addition, the input circuit attenuates the signal by 25dB and extends the common-mode input voltage range to ±98 volts peak, with common-mode rejection greater than 70dB from 20Hz to 20kHz. The SSM-2015 is set to produce a 15dB gain. The signal drive level into the SSM-2402 switch is then

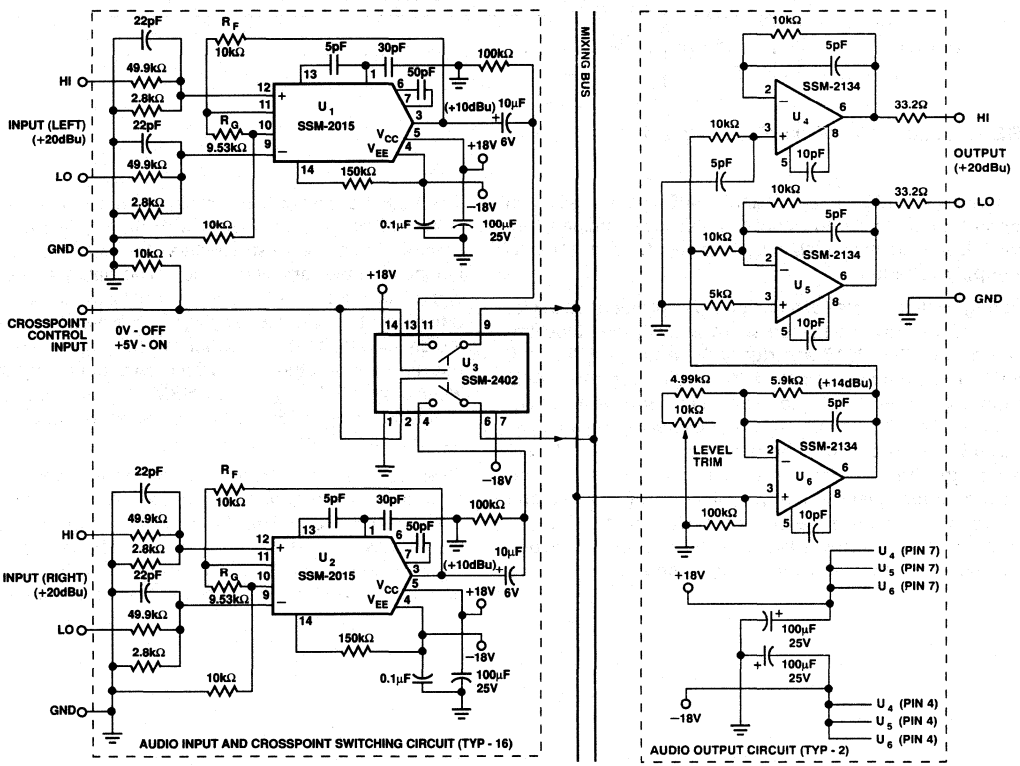

FIGURE 1: *Switcher Schematic*

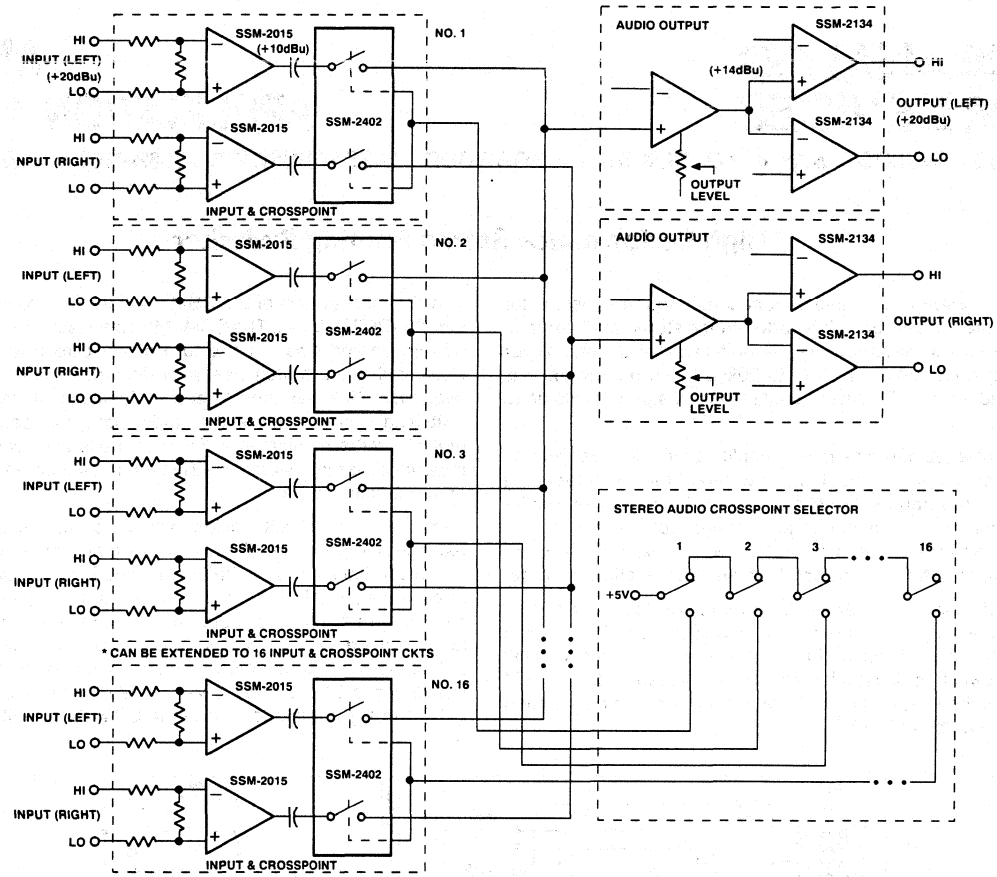

FIGURE 2: *Switcher Functional Block Diagram*

+10dBu with a +20dBu input level and +14dBu peak, well within ideal operating range. Good signal-to-noise is maintained, with generous head-room available by electing to use ±18VDC power supply voltages.

The routing switcher bus carries high level unbalanced audio, but is driven with low impedance sources. With the output impedance of the SSM-2015 at virtually 0Ω and the SSM-2402 switch ON, resistance is typically 60Ω. Bus-to-bus crosstalk is exceptionally low. For example, assuming 14pF coupling between buses and 20kHz signal, the crosstalk (isolation) exceeds 80dB. The 14pF would be representative for the 16 X 1 stereo design shown. Shielding of the buses with a printed circuit board ground plane and physically isolating the input and output circuits will reduce the crosstalk even further. The "T" configuration of the SSM-2402 switch virtually eliminates crosstalk between the various input signal sources.

The output amplifier incorporates a buffer amplifier that provides 4dB of gain (nominally), with adjustable output level trim control. The buffer also isolates the switching bus from the balanced output amplifier circuit. The balanced output is designed to drive 600Ω loads and utilizes two SSM-2134 IC amplifiers. The differential design increases drive capability, yet increases the heat dissipation surface area, and keeps IC package temperature well within safe operating limits, even when driving 600Ω loads. The SSM-2134 is recommended due to its low noise, wide frequency response, and output drive current capabilities.

Overall performance of the 16 X 1 stereo switcher is noteworthy. Input-to-output frequency response is flat to within 1dB over a 10Hz to 50kHz band. Total harmonic distortion plus noise is less than 0.03%, from 20Hz to 20kHz. SMPTE intermodulation distortion is less than 0.02%. The use of ±18VDC power supplies produces a +30dBm clip level, even when driving 600Ω loads.

TABLE 1: Circuit Performance Specifications

Max Input Level	+30dBu
nput Impedance, Unbalanced	100kΩ
nput Impedance, Balanced	200kΩ
Common-Mode Rejection (20Hz to 20kHz)	>70dB
Common-Mode Voltage Limit	±98V Peak
Max Output Level	+30dBu/dBm
Output Impedance	67Ω
Gain Control Range	±2dB
Output Voltage Slew Rate	6V/μs
Frequency Response (±0.05dB)	20Hz to 20kHz
Frequency Response (±0.5dB)	10Hz to 50kHz
THD + Noise (20Hz to 20kHz, +8dBu)	0.005%
THD + Noise (20Hz to 20kHz, +24dBu)	0.03%
IMD (SMPTE 60Hz & 4kHz, 4:1, +24dBu)	0.02%
Crosstalk (20Hz to 20kHz)	>80dB
S/N Ratio @ 0dB Gain	135dB

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

A Balanced Mute Circuit for Audio Mixing Consoles

The SSM-2402 Dual Audio Switch enhances the performance and simplifies the design of balanced high level switching (Mute) circuits used in audio mixing consoles. The use of the SSM-2402 and SSM-2134 creates a design that has negligible transient noise (as a result of signal switching), and exceptionally low signal distortion over a wide dynamic range. The balanced high level voltage switch then drives a virtual ground summing bus through 10kΩ resistors. Also included is a design for a virtual ground summing amplifier.

The SSM-2402 is a monolithic dual audio switch that improves electrical performance and eases printed circuit board layout design. The design reduces manufacturing cost when compared with discrete JFET designs of similar performance. Electrical performance is measurably superior to CMOS switch designs, and will be less prone to failure from electrical static discharge.

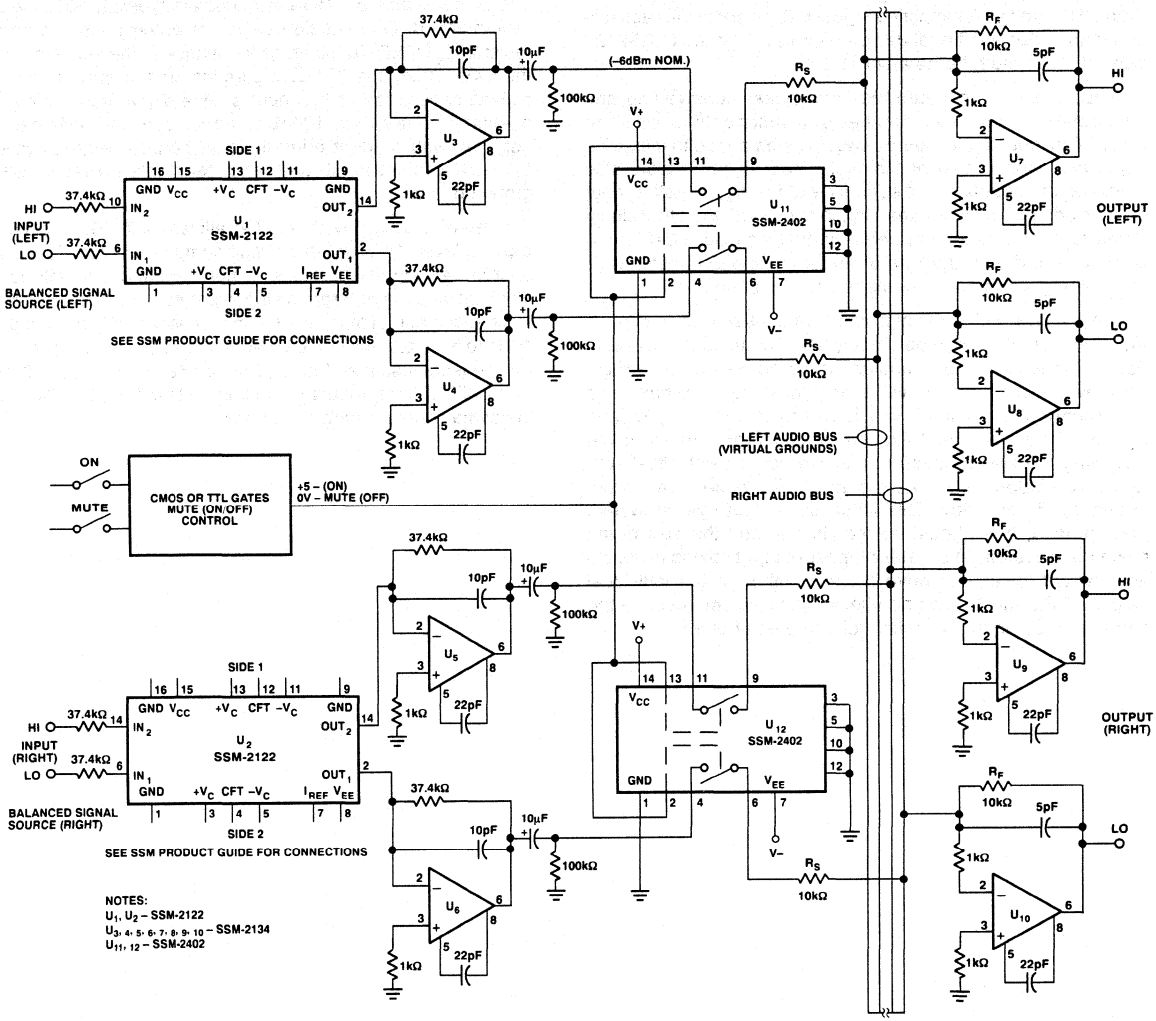

FIGURE 1: *Audio Mixer Channel Mute (On/Off) Circuit, a Balanced Design with High Level Bus Switching*

The "T" switch configuration of the SSM-2402 provides excellent ON-OFF isolation. The design shown further improves the ON-OFF isolation and left/right channel crosstalk figures by maintaining the common-mode rejection ratio of a fully balanced design. The switch features a 7ms ramped turn-on and 4ms ramped turn-off, and guaranteed break-before-make switching sequence for transient-free audio switching. The system performance is improved for large audio consoles that have multiple switches in the audio signal path.

The switch control ports are easily interfaced to conventional $5V_{DC}$ TTL or CMOS digital control circuits, further simplifying the control circuit design. Furthermore, product reliability and serviceability are improved by the simplified design. The application shown uses an elementary control circuit to functionally illustrate control voltage requirements. Customized logic gate control schemes or computer-controlled designs can be easily implemented.

The application circuit design employs dual audio switches driven by U_3, U_4, U_5 and U_6 inverting amplifiers. Their gain is controlled by two dual voltage controlled amplifier (VCA) elements U_1 and U_2. A simplified signal path is shown for application clarity. For additional design information of the SSM-2122 dual VCA, consult the data sheet.

The design shown in Figure 1 is signal phase noninverting, and incorporates a minimum number of components to minimize noise. The input signal source should be balanced to maximize separation and crosstalk isolation. PCB layout should also utilize equal inductance in each side of the signal path wiring, and include equal stray capacitance to ground and other signal paths to obtain maximum performance. The SSM-2122 VCA provides good gain tracking of the two audio channels, while maintaining accuracy in the balanced signal path.

U_{11} and U_{12} are utilized as high level switches so that other post-switching functions can be employed. A nominal drive level of 0dBu balanced (–6dBu unbalanced) is applied to the switch. This level is arbitrary, but will satisfy most signal-to-noise and headroom compromises. The output of amplifiers U_3, U_4, U_5, and U_6 are AC coupled prior to the switches to further minimize the switching transients caused by active component offset voltages. The balanced virtual ground mixing buses are current driven by R_S of 10kΩ. This value can affect overall system performance, and should be modified to suit the size of the mixing bus system. A greater number of input mixing channels will warrant lower bus drive current. Although the individual values of R_S and R_F can be altered, their values must be the same for a summing amplifier voltage gain of one (1).

TABLE 1: Circuit Performance Specifications

Max Input Level	+30dB
Input Impedence, Balanced	75k
Common-Mode Rejection (20Hz to 20kHz)	>70d
Common-Mode Voltage Limit	±12V Pea
Max Output Level	+30dB
Output Voltage Slew Rate	12V/µ
Frequency Response (±0.05dBu)	20Hz to 20kH
Frequency Response (±0.5dBu)	10Hz to 50kH
THD + Noise (20Hz to 20kHz, +8dBu)	0.005%
THD + Noise (20Hz to 20kHz, +24dBu)	0.03%
IMD (SMPTE 60Hz & 4kHz, 4:1, +24dBu)	0.02%
ON/MUTE Isolation (20Hz to 20kHz)	>85dB
S/N Ratio @ 0dB Gain	135dB

The active switches' ON resistances are typically 60Ω and are well matched. One should use 1% or better tolerance series resistor R_S (10kΩ) to minimize imbalance in the signal path. In the OFF state, the "T" configuration of the switch virtually eliminates leakage of the input source signal into the mixing bus(es). Greater than 100dB mute isolation at 1kHz can be obtained with prudent printed circuit board design since the SSM-2402 control inputs and switch terminals are separated by ground guards.

The use of $\pm 18V_{DC}$ power supplies allows a +30dBu (balanced) clipping level. All integrated circuit components mentioned will operate reliably at $\pm 18V_{DC}$, and noise contribution will be indiscernible, even in large mixing systems. The balanced input to balanced output frequency response is typically greater than 10Hz to 50kHz, within 1dB. Total harmonic distortion plus noise will measure less than 0.01%, from 20Hz to 20kHz at +30dBu with SMPTE intermodulation distortion less than 0.02% under the same measurement conditions.

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

A Constant Power "Pan" Control Circuit for Microphone Audio Mixing

The SSM-2134 permits the design of a constant power, transient-free "PAN" control circuit suitable for installation in the highest performance audio mixing consoles. The design incorporates unique and vital features. The PAN IN/OUT switch does not introduce transient type noise or interruptions in the audio when activated or deactivited, and when panning, an accurate constant power output is maintained between the sum of the two channels. The design allows "punching-in" and "punching-out" of the PAN circuit while mixing down or on-the-air, without transient clicks or holes in the mix.

The design utilizes conventional parts, e.g., a single SPST switch and a linear 10kΩ potentiometer. U_1 (SSM-2134) is used as a unity gain, inverting buffer with an input impedance of 37.4kΩ. The input source could be a VCA element or audio direct from the fader control. The values shown will allow a VCA, for example, the SSM-2013, to be used with only minor additions. The overall application circuit is noninverting from input to output.

The 15kΩ series input resistors R_S, plus the inverting input 15kΩ R_I in parallel with 5kΩ (1/2 of 10kΩ, with the PAN control in the center) forms an attenuator that has a 14dB loss. Rotating the PAN control in either direction decreases the attenuation to −11dB for one channel and maximum attenuation for the other.

$$\frac{1}{R_L} = \frac{1}{R_I} + \frac{1}{5k\Omega}, \qquad R_L = 3.75k\Omega$$

Attenuation is calculated as:

$$dB_{LOSS} = 20 \log \frac{R_L}{R_L + R_S} = 20 \log \frac{3.75k\Omega}{3.75k\Omega + 15k\Omega} = -14dB$$

Amplifier (U_2 & U_3) gain is:

$$dB_{GAIN} = 20 \log \frac{R_F}{R_I} = 20 \log \frac{75k\Omega}{15k\Omega} = +14dB$$

The frequency response is typically 10Hz to 50kHz, within 0.5dB. Total harmonic distortion plus noise will measure less than 0.007% from 20Hz to 20kHz, and SMPTE intermodulation distortion less than 0.01%. The amplifier clipping level is +24dBu with $\pm 18V_{DC}$ power supply rails. Headroom is nominally 30dB, and 27dB at full PAN for the operating channel.

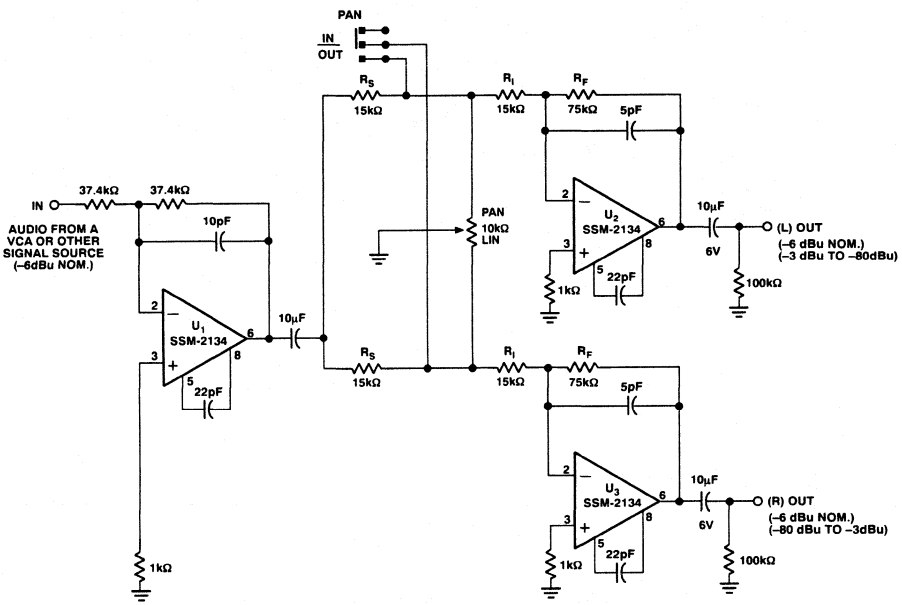

FIGURE 1: *Constant Power Type Control Circuit with Transient Free IN/OUT Switching*

TABLE 1: Circuit Performance Specifications

PAN Range, L ← C → R (L Out)	+3dB ← 0dB → –80dB
PAN Range, R ← C → L (R Out)	+3dB ← 0dB → –80dB
Max Input Level	+24dBu
Input Impedance, Balanced	37.4kΩ
Max Output Level (> 600Ω ±18V_{DC} PS)	+24dBu
Headroom	30dB
Output Voltage Slew Rate	<6V/μs
Frequency Response (±0.05dB)	20 Hz to 20kHz
Frequency Response (±0.5dB)	10 Hz to 50kHz
THD + Noise (20Hz to 20kHz, +8dBu)	0.005%
THD + Noise (20Hz to 20kHz, +24dBu)	0.03%
IMD (SMPTE 60Hz & 4kHz, 4:1, +24dBu)	0.02%
S/N Ratio	130dB

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Three High Accuracy RIAA/IEC MC and MM Phono Preamplifiers

Although the digital compact disk is rapidly supplanting the vinyl disk as the popular media method for professional and consumer audio entertainment, the electro-mechanical recording and reproduction of audio signals has many more years of life. The group of phono preamplifier application designs below will make the future years with vinyl more productive and pleasant. The applications employ solid engineering concepts, and dismiss "golden ear" discussions.

One design includes an input scheme for both moving coil (MC) and moving magnet (MM) – or variable reluctance – transducers. All designs employ extremely low noise circuit topologies, high accuracy active and passive RIAA (Recording Industries Association of America) equalization with selectable old RIAA or RIAA/IEC (International Electro-Technical Commission) curves. The applications incorporate both consumer and balanced output circuit configurations.

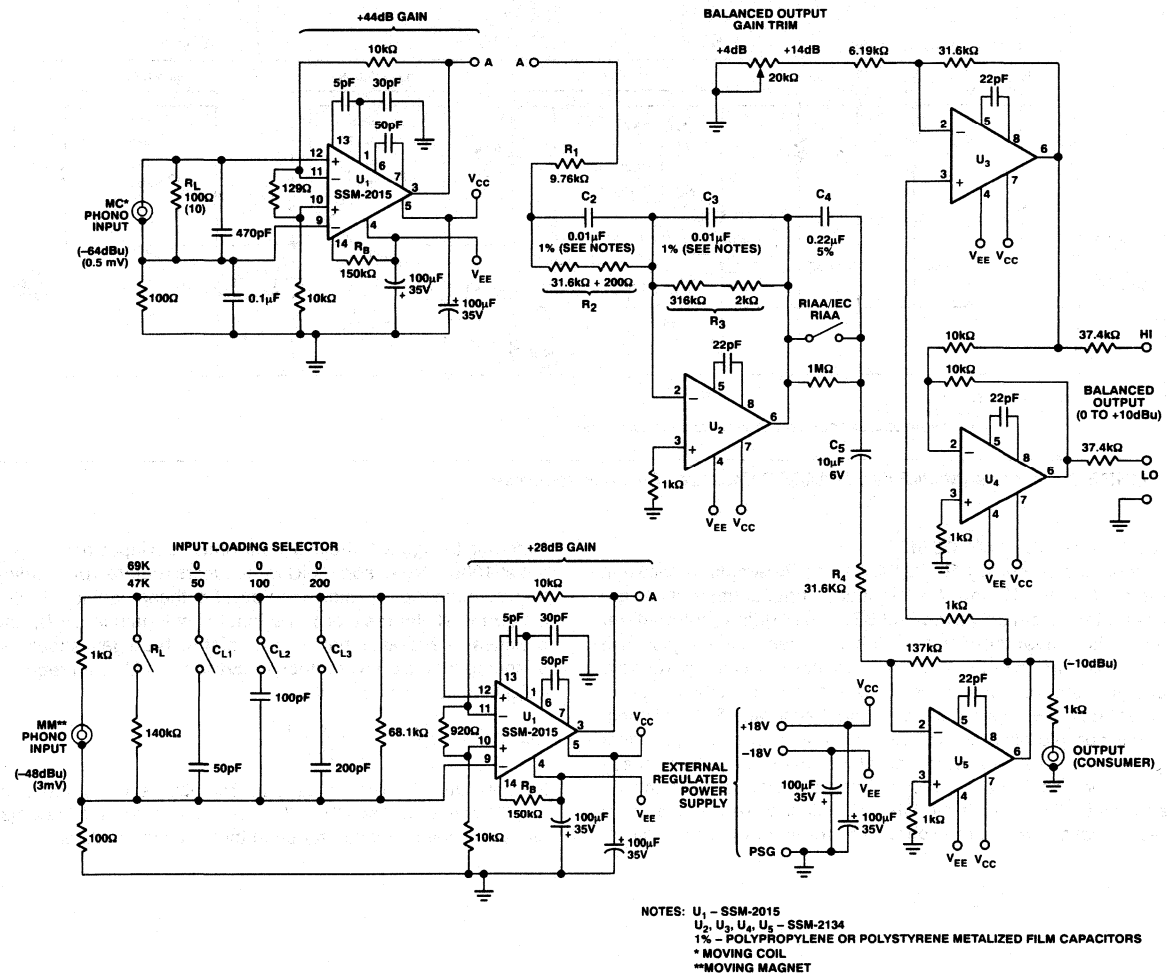

FIGURE 1: *High Accuracy RIAA/IEC MC or MM Phono Preamplifier*

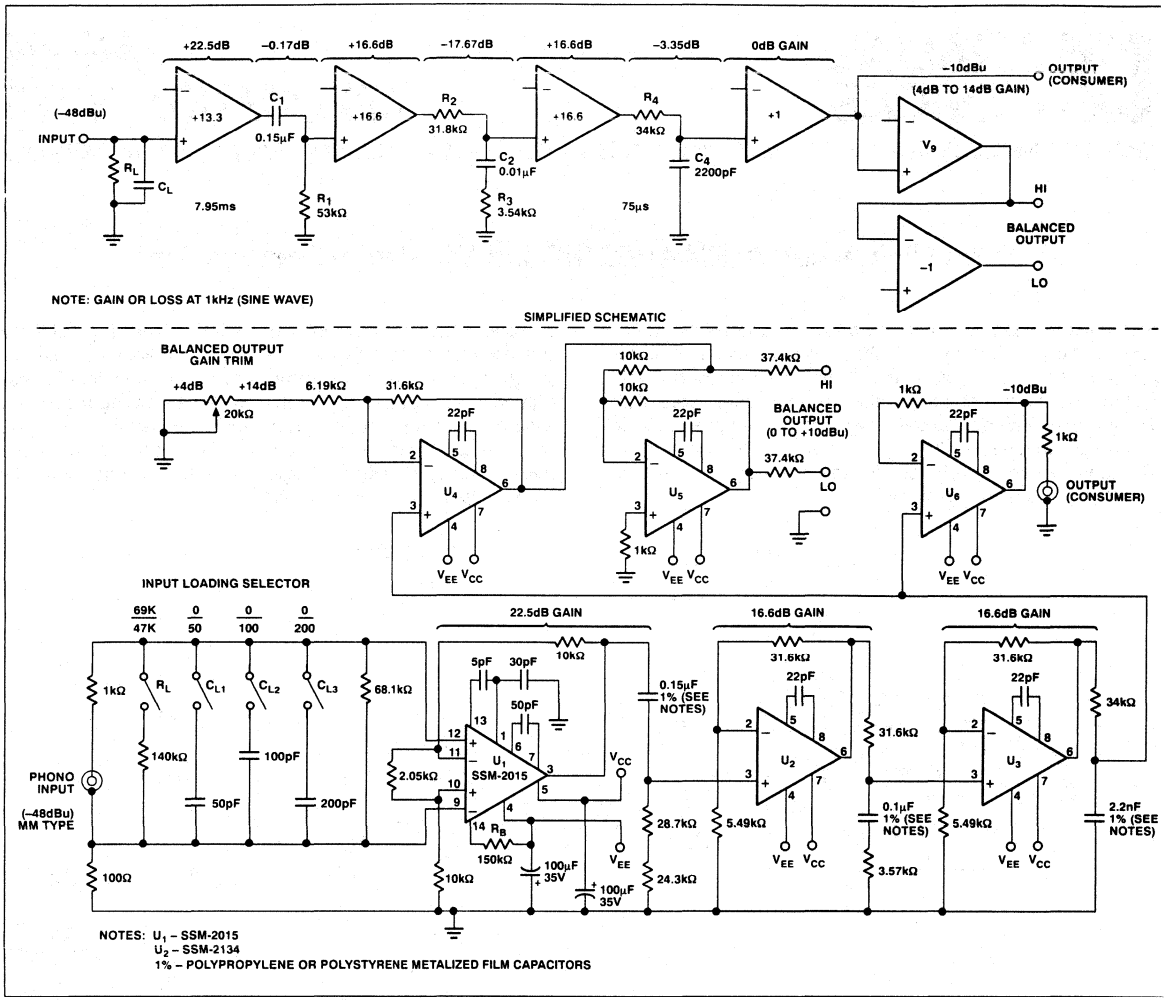

FIGURE 2: *Passive (Multi-Filter) RIAA/IEC Equalized Phono Preamplifier*

A HIGH ACCURACY DESIGN

In the High Accuracy RIAA/IEC Phono Preamplifier shown in Figure 1, both MC and MM transducer input configurations are presented. Both utilize the PMI SSM-2015 differential amplifier and take advantage of the high common-mode rejection it provides. The overall circuit structure *does not* incorporate any design compromises. It provides the lowest possible noise, adjustable MM input loading, highest accuracy RIAA filtering, and is completely devoid of transient and frequency dependent gain errors. The wide bandwidth stages minimize in-band phase shift, and provide exceptional phase and frequency response accuracy. This allows the RIAA/IEC filter to render the exact reciprocal of the recorded phase and frequency characteristics.

Referring to Figure 1, the MC input circuit has input loading (R_L set at 100Ω. (Note: some MC transducers require 10Ω loadin(for maximum reproduction accuracy. For these, replace R_L witl a 10Ω metal film resistor.) The input circuit gain is 44dB, an(provides a −20dBu signal level at point A. 44dB gain should b(adequate for most MC cartridges available. If U_1 gain require(adjusting, use the equation:

$$G_{dB} = 20 \log\left(3.5 + \frac{20 \times 10^3}{R_G}\right)$$

R_B sets the U_1 bias value and contributes to symmetrica amplifier slewing. The RIAA filter stage that follows U_1 stage(s all but eliminates noise produced by the input amplifier.

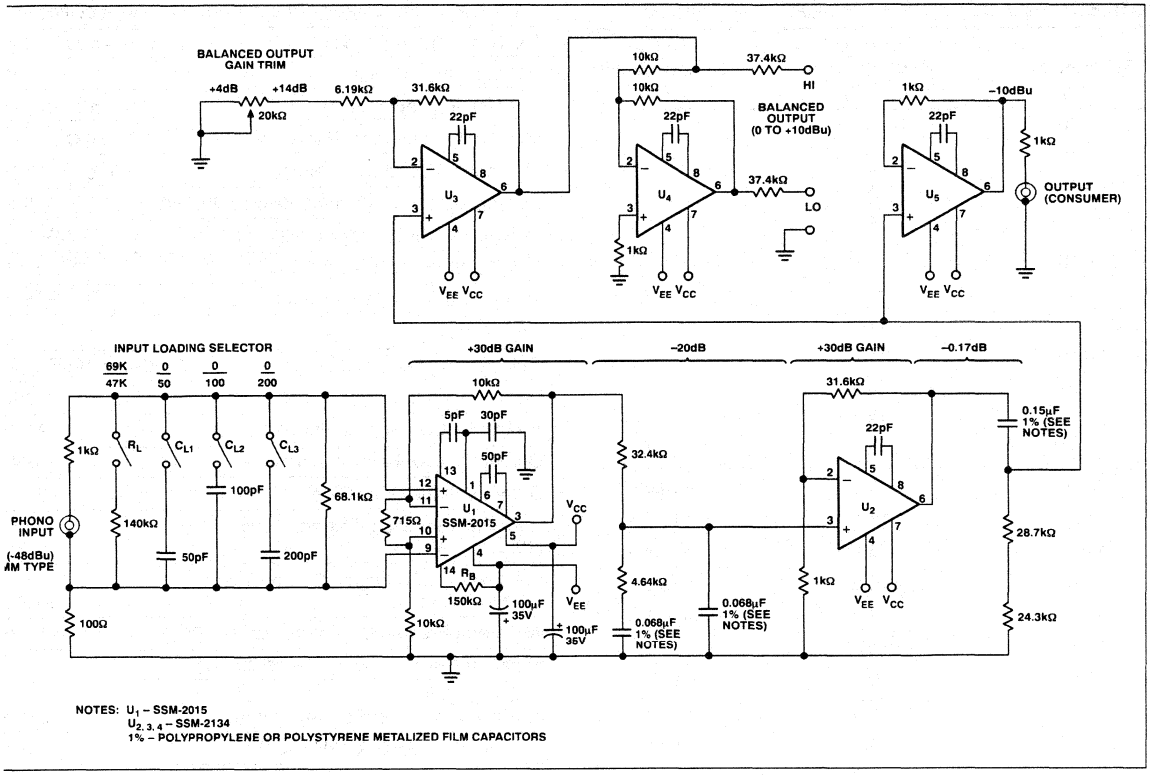

GURE 3: *Passive RIAA/IEC Equalized Phono Preamplifier*

ıe next stage contains the RIAA-RIAA/IEC equalization filter ıd is built around U_2, the SSM-2134 operational amplifier, and an active feedback type filter. The overall gain of this circuit : 1,000Hz is –2.5dB. RIAA equalization requires a gain of ɜ.3dB at 20Hz, and attenuation of 19.6dB at 20,000Hz. The ɔen-loop gain of U_2 is greater than 100dB at 20Hz, and 60dB : 20,000Hz, ensuring exceptional equalization accuracy.

hree filters make up the RIAA reproduction curve. The time ɔnstants are: 75µs, 318µs, and 3180µs, and a fourth time ɔnstant in the RIAA/IEC curve is 7960µs. The IEC filter was ıtroduced to minimize warp and infrasonic signal interference hile maintaining flat frequency response down to 40Hz.

The 75µs filter is formed by resistors R_1 (9.76kΩ) and R_2 (31.8kΩ) in parallel with capacitor C_2 (0.01µF).

The 318µs pre-emphasis filter is formed by R_2 (31.8kΩ) and C_2 (0.01µF).

The 3180µs filter is formed by R_3 (318kΩ) and C_3 (0.01µF).

The fourth pole, IEC 7960µs, a high pass filter is formed by R_4 (31.6kΩ) and C_4 (0.22µF), and provides 3dB attenuation at 20Hz rolling off at –6dB/octave thereafter.

Table 1 contains the complete RIAA and RIAA/IEC reproduction equalization characteristics. RIAA/IEC switching allows selection of either reproduction response curves. For the "audio purist," C_5 can be eliminated for a direct coupled design, thus reducing envelope and group delay distortions. All amplifier feedback circuits are direct coupled, and are referenced to circuit ground. The closed-loop gain is kept low to minimize input offset voltage. Therefore, only very small DC voltages can be expected at the output of the directly coupled version.

The high level amplifier, U_5, provides +12.7dB gain and feeds the unbalanced Consumer Output jack, with a nominal –10dBu level. U_5 is followed by balanced output buffer amplifier. The nominal output level is continuously adjustable from 0.0dBm to +10dBm at the balanced output terminals. The output source impedance is 75Ω, and will drive 600Ω loads to a maximum +30dBm clip point level. Table 2 shows circuit performance specifications.

TABLE 1: RIAA/IEC and RIAA Playback Characteristics

Frequency (Hz)	RIAA /IEC Relative Level (dB)	RIAA Relative Level (dB)
2.0	–0.2	
2.5	+1.8	
3.15	+3.7	
4.0	+5.7	
5.0	+7.6	
6.3	+9.4	
8.0	+11.2	
10.0	+12.8	
12.5	+14.1	
16.0	+15.4	
20.0	+16.3	+19.3
25.0	+16.8	+19.0
31.5	+17.0	+18.5
40.0	+16.8	+17.8
50.0	+16.3	+16.9
63.0	+15.4	+15.8
80.0	+14.2	+14.5
100	+12.9	+13.1
125	+11.5	+11.6
160	+9.7	+9.8
200	+8.2	+8.2
250	+6.7	+6.7
315	+5.2	+5.2
400	+3.8	+3.8
500	+2.6	+2.6
630	+0.8	+0.8
1,000	0.0	0.0
1,250	–0.8	–0.7
1,600	–1.6	–1.6
2,000	–2.6	–2.6
2,500	–3.7	–3.7
3,150	–5.0	–5.0
4,000	–6.6	–6.6
5,000	–8.2	–8.2
6,300	–10.0	–10.0
8,000	–11.9	–11.9
10,000	–13.7	–13.7
12,500	–15.6	–15.6
16,000	–17.7	–17.7
20,000	–19.6	–19.6

TABLE 2: High Accuracy Circuit Performance Specification

MC Nominal Input Level	–64dBu (0.5m
MC Input Impedance	100
MM Nominal Input Level	–48dBu (3.0m
MM Input Impedance, Resistive	69kΩ or 47k
MM Input Impedance, Capacitive	50pF to 350p
Common-Mode Rejection (20Hz to 20kHz)	>50c
Common-Mode Voltage Limit	±10V Pea
Nominal Output Level, Balanced	+8dBu/dB
Max Output Level, Balanced	+30dBu/dB
Output Impedance, Balanced	70
Gain Control Range, Balanced	0.0dBu to 10dBu/dB
Nominal Output Level, Unbalanced	–10dB
Max Output Level, Unbalanced	+24dB
Output Impedance, Unbalanced	1,000
Output Voltage Slew Rate	>6V/ µ
RIAA Reproduction Characteristics (20Hz to 20kHz)	±0.25d
RIAA/IEC Reproduction Characteristic (2Hz to 20kHz)	±1.0d
Wideband Frequency Response (±1.0dB)	0.0Hz to 70kH
Signal-to-Noise Ratio (20Hz to 20kHz)	>90d
THD + Noise (20Hz to 20kHz +8dBu, Any Output)	0.01%
IMD (SMPTE 60Hz & 4kHz, 4:1)	0.02%

A PASSIVE MULTI-FILTER DESIGN

The Passive Split Multi-Filter RIAA/IEC Preamplifier design shown in Figure 2, is intended for moving magnet (MM) inpu phono transducers. The design has an extremely low nois circuit topology, high accuracy passive RIAA/IEC equalizatio filters, and both unbalanced consumer and balanced outpu circuits. The input configuration utilizes the SSM-2015. It pro vides the lowest possible noise, adjustable resistive and ca pacitive input loading, and high accuracy passive RIAA filterin totally devoid of transient and frequency dependent gain errors

Referring to Figure 2, the following two stages contain the RIAA RIAA/IEC passive equalization filters. All high pass and lo pass filters are passive. The signal is amplified by U_2 and U_3 SSM-2134 op amps. The overall gain of the circuit at 1,000H is 38dB. RIAA equalization requires a gain of 19.3dB at 20Hz and attenuation of 19.6dB at 20,000Hz. Open-loop gain of U_2 and U_3 is greater than 100dB at 20Hz, and 60dB at 20,000Hz Closed-loop gain of U_1 is 22.5dB, and U_2, U_3 is 16.6dB ensuring an extensive gain margin for phase accuracy. Refer to Table 3 for complete circuit specifications.

TABLE 3: Passive Multi-Filter Circuit Performance Specifications

MM Nominal Input Level	–48dBu (3.0mV)
MM Input Impedance, Resistive	69kΩ or 47kΩ
MM Input Impedance, Capacitive	50pF to 350pF
Common-Mode Rejection (20Hz to 20kHz)	> 50 dB
Common-Mode Voltage Limit	±10V Peak
Max Output Level, Balanced	+30dBu/dBm
Nominal Output Level, Balanced	+8dBu/dBm
Output Impedance, Balanced	70Ω
Gain Control Range, Balanced	0.0dBu to 10dBu/dBm
Nominal Output Level, Unbalanced	–10dBu
Max Output Level, Unbalanced	+24dBm
Output Impedance, Unbalanced	1,000Ω
Output Voltage Slew Rate	>6V/μs
RIAA Reproduction Characteristic 20Hz to 20kHz)	±0.25dB
RIAA/IEC Reproduction Characteristic 2Hz to 20kHz)	±0.5dB
Wideband Frequency Response (±1.0dB)	0.0Hz to 70kHz
Signal-to-Noise Ratio (20Hz to 20kHz)	>90dB
THD + Noise 20Hz to 20kHz +8dBu, Any Output)	0.01%
IMD (SMPTE 60Hz & 4kHz, 4:1)	0.02%

TABLE 4: Uncomplicated Passive Circuit Performance Specifications

MM Nominal Input Level	–48dBu (3.0mV)
MM Input Impedance, Resistive	69kΩ or 47kΩ
MM Input Impedance, Capacitive	50pF to 350pF
Common-Mode Rejection (20Hz to 20kHz)	> 50dB
Commom-Mode Voltage Limit	±10V Peak
Max Output Level, Balanced	+30dBu/dBm
Nominal Output Level, Balanced	+8dBu/dBm
Output Impedance, Balanced	70Ω
Gain Control Range, Balanced	0.0dBu to 10dBu/dBm
Nominal Output Level, Unbalanced	–10dBu
Max Output Level, Unbalanced	+24dBu
Output Impedance, Unbalanced	1,000Ω
Output Voltage Slew Rate	>6V/μs
RIAA Reproduction Characteristic (20Hz to 20kHz)	±0.5dB
RIAA/IEC Reproduction Characteristic (2Hz to 20kHz)	±1.0dB
Wideband Frequency Response (±1.0dB)	0.0Hz to 70kHz
Signal-to-Noise Ratio (20Hz to 20kHz)	>90dB
THD + Noise (20Hz to 20kHz, +8dBu, Any Output)	0.01%
IMD (SMPTE 60Hz & 4kHz, 4:1)	0.02%

4

AN ECONOMICAL APPROACH

An Uncomplicated Passive RIAA/IEC Preamplifier is shown in Figure 3. It is a low cost, practical design for a passively equalized RIAA/IEC phono preamplifier. The design shown is for moving magnet (MM) input. It also is an extremely low noise input circuit design, and includes both unbalanced consumer and balanced output circuit configurations. The input circuit also utilizes the SSM-2015, and provides adjustable resistive and capacitive input loading. Wide bandwidth stages minimize in-band phase shift, and provide exceptional phase and frequency response accuracy. Table 4 details circuit performance data.

SUMMARY

For a phono transducer cartridge to deliver the performance as intended, it should be loaded with proper resistance and capacitance. The MM input circuits have adjustable transducer loading. Most transducers currently available will be accommodated with resistive loading of 69kΩ or 47kΩ, and capacitive loading of a few pF (input wiring dependent) to 350pF, in 50pF steps.

If greater input common-mode noise rejection is required, it can be obtained in all input designs by increasing the value of the 100Ω resistor and 0.1μF capacitor connected between the input RCA jack shield connection and the main circuit ground point. The values shown satisfy most requirements for 1 meter cables supplied with the newer tone arms.

All circuits described are signal noninverting, and constructed with bipolar IC amplifiers for lowest noise. They are compensated for widest bandwidth and circuit stability.

To achieve optimum trouble-free performance, a few construction and manufacturing tips should be observed. For grounding to be truly effective, all grounded components must return to a single point. This technique is effective in minimizing ground current loops that can cause excessive noise, signal cross-talk, AC power line noise, and circuit instability, and permit external noise spikes to enter. The ground center should be as close to the input amplifier (U_1) as possible. All grounded components of U_2, U_3, U_4, U_5, the output jack grounds, and the power supply ground lead should be tied to the same U_1 ground point.

As long as the power supply leads are kept short, and adequately filtered and bypassed with polyester film capacitor at the regulators, there is no need for individual decoupling capacitors at U_2, U_3, U_4, and U_5. The power supply voltages should be regulated for $\pm 18V_{DC}$.

All signal filter components should be of the highest quality, i.e., metalized polypropylene or polystyrene film, 1% tolerance capacitors (except for C_5, 5% tolerance is OK) and metal film resistors, 1% or better tolerance.

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

A Two-Channel Dynamic Filter Noise Reduction System

In this application, the SSM-2120 Dynamic Range Processor and SSM-2134 op amp are utilized in a dual-channel dynamic noise reduction circuit, where the input signal level and threshold control determine the corner frequency of a low-pass filter.

The SSM-2120 contains two class A VCAs (Voltage Controlled Amplifiers) that are used as the filter's control element, and two wide dynamic range full-wave rectifiers with control amplifiers. The VCA section or variable resistor is a current device con-

4

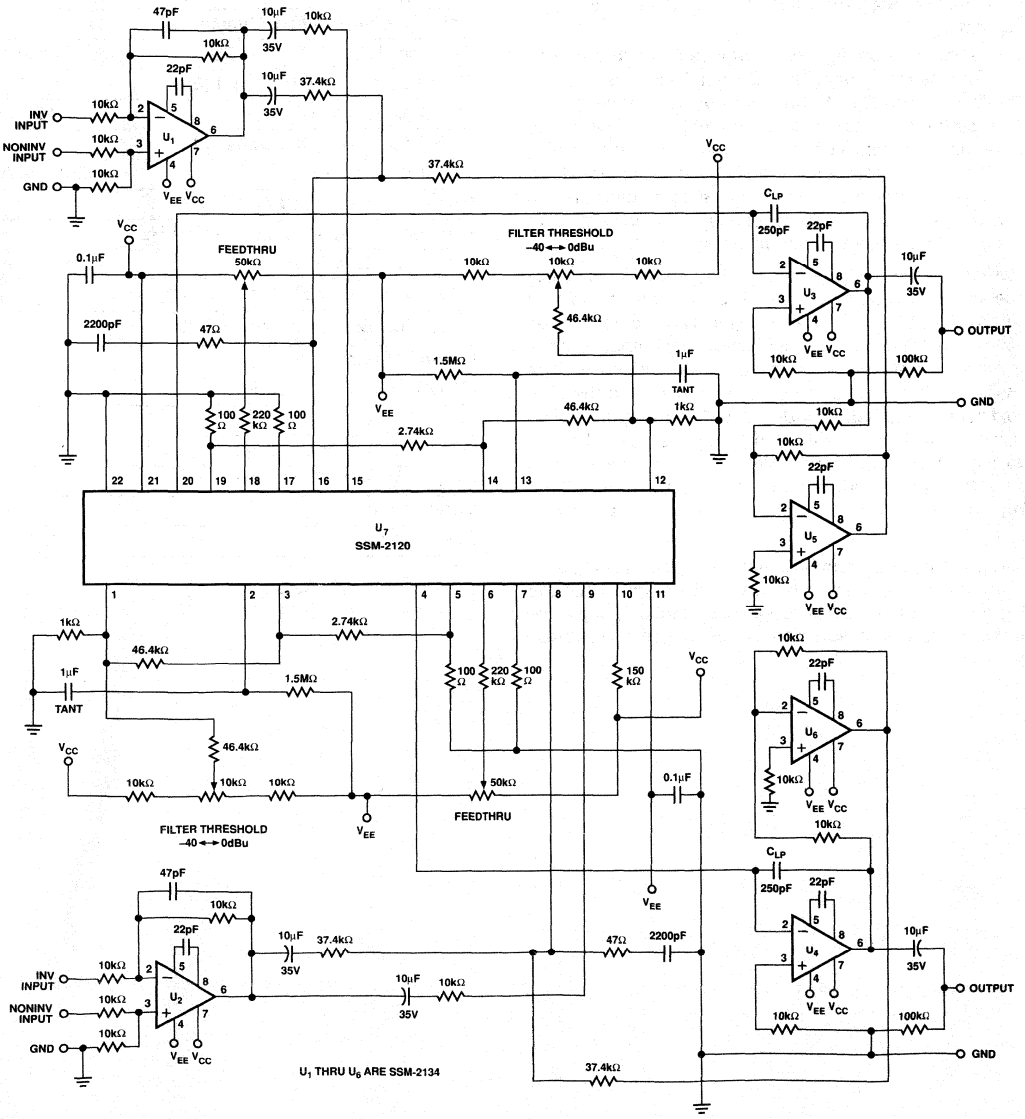

FIGURE 1: *Two-Channel Dynamic Filter Noise Reduction System*

trolled by the $+V_C$ voltage control ports. The VCAs are employed as variable resistor elements in a single-pole low-pass filter operating in a virtual ground configuration. The level detecting rectifier is a full-wave averaging type with more than 100dB dynamic range, followed by a LOG amp converter. The part also contains two operational amplifiers with PNP output transistors used to drive the $+V_C$ ports.

U_1 and U_2 (SSM-2134) are input amplifiers and source-load isolating buffers. They provide a choice of noninverting, inverting, or balanced inputs. Unbalanced loading is 10kΩ and balanced loading is 20kΩ. U_1 and U_2 gain is set at 0dB, with a –10dBu nominal input signal recommended using $\pm 18V_{DC}$ power supply rails. This configuration will provide an overall circuit headroom of more than 30dB. In less critical applications, the feedthrough trim controls and 220kΩ resistors can be eliminated.

DETAIL CIRCUIT DESCRIPTION FOR U_7 (SSM-2120), AND U_3 and U_4

The rectifier circuit is configured to provide a negative control voltage referenced to ground. The LOG amplifier's bias is set by the 1.5MΩ resistor. The 1.5MΩ resistor also provides the discharge path for the 1μF rectifier averaging capacitor. The discharge time constant controls the low-pass filter's action toward a lower corner frequency. The LOG amplifier provides a constant current charging for the 1μF averaging capacitor. It results in an attack (return to flat response) time constant T_{C0} of approximately 6ms, and a low-pass filter activation T_{C1} of 350ms.

The internal op amp of U_7 has the gain V_G set at 47. The potentiometer at the inverting input provides the adjustable threshold to activate the filter. The threshold adjustment ranges from –40dBu to 0dBu of input signal level. The output from the op amp drive transistor supplies only a negative control voltage to the VCA ($+V_C$) control port(s). For example, with the filter threshold control adjusted to 0dB and a –10dBu signal applied to the input, f_{C1} is ≈4kHz; or with –20dBu applied, f_{C2} is ≈1.2kHz, both rolling off at 6dB/Octave. With the input signal level exceeding the filter threshold setting (0V VCA control voltage present), the overall circuit frequency response is 20Hz to 16kHz, at ±1dB.

The VCA audio input current is limited by the 37.4kΩ resistor. The VCAs operate as current devices whose outputs feed the virtual ground of an amplifier loop. The feedback capacitor around the amplifier loop sets up a single-pole low-pass filter.

The SSM-2120 (U_7) inverts the signal current; therefore, U_5 ar U_6 are required to invert the output signal that is summed at tl input(s) of U_7.

The design is an effective single-ended noise reduction circ with low distortion and noise. When utilized on a noisy sign source, it will attenuate high frequency noise with inconspic ous operation.

TABLE 1: Circuit Performance Specifications

Nominal Input Voltage (–10dBu Out)	–10dE
Headroom (–10dBu Out)	+30dE
Input Voltage Range	–20dBu to +10dE
Input Type/Impedance, Balanced	20k
Input Type/Impedance, Unbalanced	10k
Dynamic Noise Reduction Class	Dynamic Low-Pas
Filter Activate Time Constant (6dB)	350m
Threshold Range (Level)	–40dBu to 0dB
Filter Deactivate Time Constant	6m
Signal Rectifier Type	Full Wave Averagin
Modulation Feedthrough, Trimmed	–100d
Frequency Response (20Hz to 16kHz)	±1d
Filter Type, Low-Pass	Single Pole, 6dB/O
Input 10dB Below Threshold Setting	f_{C1} = 3,800H
Input 20dB Below Threshold Setting	f_{C2} = 1,400H
Dynamic Range	
@ 0dB Gain (Ref. +22dBu)	106d
THD + Noise (20Hz to 20kHz)	0.02%
IMD (SMPTE 60Hz & 4kHz, 4:1)	0.05%
Output Voltage (2kΩ Load)	+22dB
Output Type	Unbalance
Power Supply	$\pm 18V_{DC}$ Regulate

NE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

An Unbalanced Mute Circuit for Audio Mixing Channels

s application note describes a dual channel unbalanced alog audio mute switch, for use in audio console mute cuits. The SSM-2402 dual audio switch, when used in the ual ground configuration, truly enhances any audio mute sign. The application, as shown in Figure 1, incorporates balanced stereo input buffers, dual stereo electronic virtual ound switches (with simplified control circuit), and virtual ound summing amplifiers.

THE AUDIO SWITCH AS IMPLEMENTED

The design utilizes the SSM-2402 (U_1 and U_2) dual audio switch in a virtual ground switching configuration. This method of operation improves linearity over a wide dynamic range. The SSM-2402 utilizes JFET switching, with internal wide bandwidth integrated amplifiers applied in a unique configuration. The result is low transient intermodulation distortion, low THD, and low IMD, while essentially eliminating all audio switching

4

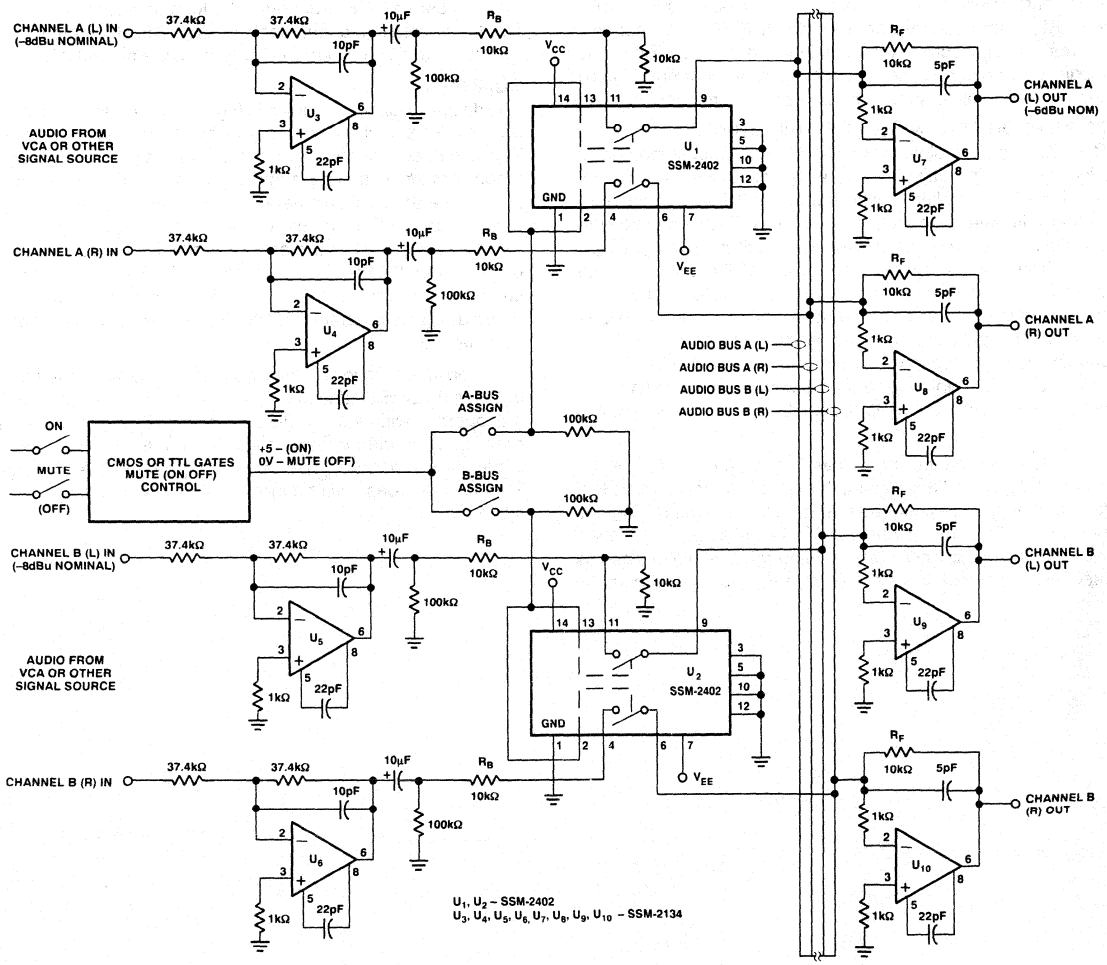

FIGURE 1: *Audio Mixer Channel Mute (On/Off) Circuit (Unbalanced Design with Virtual Ground Switching)*

transients. The SSM-2402 switch closed (ON) resistance is typically 60Ω in series with R_B (10kΩ). As shown in this mixing system, the tolerance of the 60Ω contributes to less channel imbalance than the 1% resistor tolerance, thus eliminating the need for level trim adjustments.

The SSM-2402 employs a "T" switching configuration that yields superior ON-OFF signal isolation. In the OFF state, the "T" configuration of the SSM-2402 virtually eliminates leakage of the input signal (down more than 100dB at 1kHz with guard pins 3, 5, 10 and 12 grounded) onto the mixing bus(es). The part also features a 7ms ramped turn ON and 4ms ramped turn OFF, for transient free audio switching even with signal applied. The switch also operates with a break-before-make switching sequence. These properties are significant when many remotely controlled electronic switches are connected in series and controlled by a single device, as in large audio systems managed by an automation computer.

CONTROL INTERFACE

In this application note, the bus assignment (selection) switches are shown functionally for clarity. The control ports of the SSM-2402 can easily be interfaced to conventional 5V TTL or CMOS logic control circuits. $+5V_{DC}$ (logic high) closes the switch (ON), and $0V_{DC}$ (logic low) opens the switch (OFF). The common interface levels improve the reliability and serviceability of any products it's designed into. Diverse logic gate control designs or computer controlled schemes can easily be implemented.

DRIVE REQUIREMENTS – THE INPUT CIRCUIT

The application employs two SSM-2402 dual audio switches in a four-bus configuration (two stereo buses) driven by U_3, U_4 and U_5, U_6 bipolar amplifiers. The buffer amplifiers are signal inverting, with their gain set to 0dB ($A_V = 1$). The input amplifiers also serve as source signal level clippers that prevent the input signal from exceeding the input range of the switches, thus preventing the switches from passing a distorted signal when overdriven in the open (OFF) state. A nominal input drive level of –10dBu is applied to the switch and will maximize the signal-to-noise ratio, and optimize headroom. The output of U_3, U_4, U_5, and U_6 are AC coupled to further minimize the switching transient noise caused by signal path DC voltages from previous origins.

The virtual ground mixing buses are current driven by R_B (10.0kΩ) resistors. Once again, this is a compromise value that can be changed to accommodate the extent of the mixing bus implemented. A greater number of input mixing channels will warrant a lower bus drive current. Although other values can be used, the resistance values of R_B and R_F should be the same.

As shown ($\pm 18V_{DC}$ power) R_B will apply approximately 1 peak current to the mixing bus. This is well within SSM-switching capabilities, as well as the SSM-2134 drive cap ties. The signal current is low enough to keep return gr currents low enough to prevent crosstalk resulting fron mechanical wiring constraints. Returning ground currents i pendently to the noninverting input of the summing amplif advised.

THE OUTPUT SUMMING AMPLIFIER

The design utilizes the SSM-2134, the PMI version of popular NE5534 bipolar operational amplifier. The circuit tures a significant reduction in summing amplifier nois decrease in temperature and bus impedance effects on static output voltage as a result of using a bipolar amplifier. design also balances the input circuit reflected source imp ance of the bipolar IC amplifier, alleviating the unity-gain in bility and eliminating the unbalanced input topology for inver summing designs that could cause output offset.

The SSM-2134 has a noise voltage of 2.8nV/$\sqrt{\text{Hz}}$, thus noise floor is reduced by 3 to 10dB. Additionally, frequency phase response performance have been improved. Only m mal compensation is required in the feedback loop of the SS 2134 to maintain unconditional stability. The slew rate rema greater than 10V/μs, with bandwidth exceeding 50kHz.

SUMMARY

The design application shown in Figure 1 is signal noninverti and utilizes a minimum number of noise generating elemer The circuit configuration produces linear signal mixing at virtual ground summing node ($e_S = 0V_{AC}$); therefore, no flected interaction occurs between the input sources. The sig input to any output frequency response is typically 10Hz 50kHz, ±0.5dB. Total harmonic distortion plus noise will mea ure less than 0.01%, from 20Hz to 20kHz. SMPTE intermoc lation distortion is less than 0.02%. With prudent printed circ board design, a greater than 100dB mute isolation @ 1kHz c be obtained.

The application shown employs $\pm 18V_{DC}$ power supplies produce a +24dBu audio output clip level. All SSM componer will operate with equal reliability at $\pm 20V_{DC}$, producing a proximately a 1dB increase in clip level. If the extra headroo is necessary, a $\pm 20V_{DC}$ power supply voltage is encourage The noise increase will be indiscernible, even in large mixir systems.

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

A Two-Channel Noise Gate

This application applies the SSM-2120 as a dual-channel noise gate or adjustable threshold downward expander. A noise gate is a type of noise reduction system that fully attenuates a VCA when no audio signal is present. The SSM-2120 contains two class A VCAs (Voltage Controlled Amplifiers) and two wide dynamic range full-wave rectifiers and control amplifiers. The VCA section is a current amplifier device whose gain is controlled by two gain ports that have dB/V scaling. The VCAs are employed as wide bandwidth amplifiers with the current inputs and outputs operating in a virtual ground configuration. The rec-

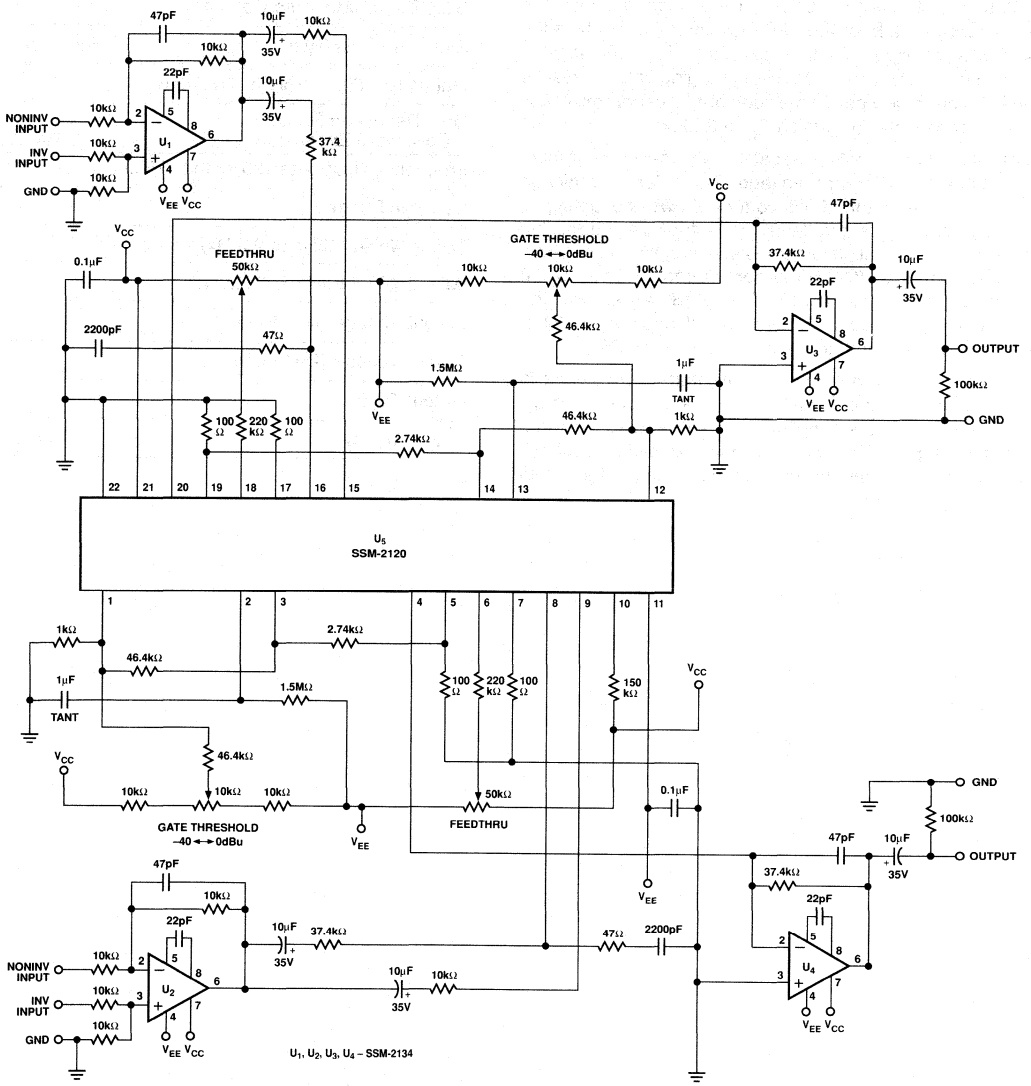

FIGURE 1: *Two-Channel Noise Gate*

tifiers are full-wave averaging type with 100dB dynamic range, followed by LOG converters. The part also contains two operational amplifiers with PNP output transistors connected in a common collector configuration.

Two SSM-2134s are used as input amplifiers, U_1 and U_2, to provide noninverting, inverting, or balanced inputs. Unbalanced loading is 10kΩ and balanced loading 20kΩ. U_1 and U_2 gain is set at 0dB, and with a –10dBu nominal input signal and $\pm 18V_{DC}$ power, will provide overall circuit headroom of 30dB. The VCA(s) could be DC coupled, although in this application they are AC coupled to reduce the dependence on trimming the side chain voltage modulation feedthrough. In critical applications the feedthrough trim controls and 220kΩ resistors should be added.

The SSM-2120's internal rectifier produces a negative DC voltage referenced to ground. The LOG amplifier bias is set by the 1.5MΩ resistor. The 1.5MΩ resistor also provides the discharge current path for the 1μF capacitor, that controls the gate's downward expansion time constant. The LOG amplifier provides a constant current capacitor charging value. It results in an attack (return 0dB gain) time constant T_C of approximately 6ms, and a downward expansion T_C of 350ms.

The internal op amp gain A_V is set at 47, with the inverting input also providing the reference voltage. The reference voltage range from the gate threshold control allows the gating to activate at any source signal level from –40dBu to 0dBu. The output from the op amp drive transistor supplies a negative control voltage to the VCA ($+V_C$) control port(s). The VCA(s) control ports have a sensitivity of 6mV/dB. As shown, the voltage divider provides a 2:1 downward expansion slope. Below the threshold level, the gain slope is $2dB_G/dB_{IN}$.

The VCAs are current output amplifiers that are designed to operate with virtual ground configurations such as U_3 and U_4. The VCA input current is supplied by the 37.4kΩ resistor and input voltage signal. The virtual ground amplifier feedback resistors are 37.4kΩ. With no VCA control voltage, the overall circuit voltage gain is 1 (0dB). Other non-gating gains can b attained by changing the output amplifiers feedback resisto value. The VCA input resistor should remain as shown fo maximizing the performance of the VCA(s).

TABLE 1: Circuit Performance Specifications

Nominal Input Voltage (–10dBu Out)	–10dB
Headroom (–10dBu Out)	+30d
Input Type/Impedance, Balanced Unbalanced	20k 10k
Downward Expander Class	Feedthroug
Threshold Sense Time Constant (6dB)	350m
Threshold Range (Level)	–40dBu to 0dB
Gate Deactivate Time Constant	6m
Signal Rectifier Type	Full-Wave Averagin
Modulation Feedthrough, Trimmed	< –60dB
Gain Reduction Ratio, Downward Expansion	1 to 2 (–2dB/dE
Frequency Response (20Hz to 20kHz)	±0.25d
Dynamic Range	100d
THD + Noise (20Hz to 20kHz)	0.02
IMD (SMPTE 60Hz & 4kHz, 4:1)	0.05
Output Voltage Slew Rate	6V/μ
Output Voltage (2kΩ Load)	+22dB
Output Type	Unbalance
Power Supply	$\pm$ 18V_{DC} Regulate

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

A Precision Sum and Difference (Audio Matrix) Circuit

When constructing an accurate sum and difference signal from stereo left and right sources, amplitude and phase (delay) errors can contribute substantial amounts of crosstalk in the reconstructed left and right audio channels. A minor 1dB difference, or 6° phase error, will result in only 25dB stereo channel separation. The design presented has essentially no phase or group delay in the sum or difference outputs as measured over the audio spectrum, 20Hz to 20kHz. This circuit utilizes matched (laser trimmed) resistor networks combined with high open-loop gain differential amplifiers to guarantee virtually no phase and amplitude error in the sum and difference channels.

Amplifiers U_3 and U_4 (SSM-2134) are utilized as input signal buffers that provide a low source impedance (0Ω) to the 10kΩ summing resistors that feed the virtual ground current summing

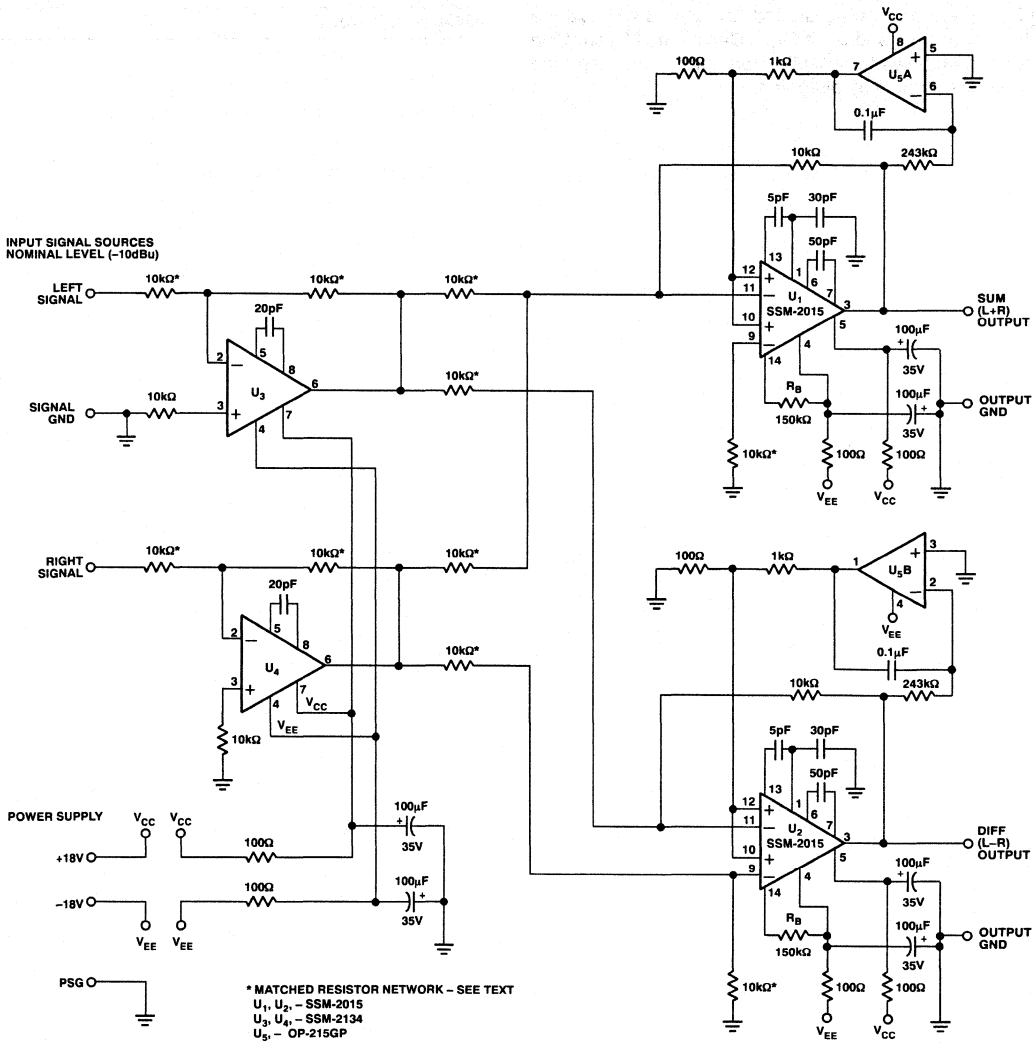

IGURE 1: *Precision Sum and Difference (Audio Matrix) Circuit*

nodes of U_1 and U_2. The overall gain of the buffer circuit is 0dB. The buffer amplifiers, U_3 and U_4, are compensated for a frequency response that extends to 100kHz at +24dBu. The buffers are configured as inverting amplifiers for lowest phase and group delay effects.

U_1 and U_2 (SSM-2015) are true differential input, high performance bipolar amplifiers. The current summing inputs have been employed for sum and difference operations. All bipolar and JFET op amps exhibit considerable propagation time differences between inverting and noninverting inputs, which result in phase and group delay errors. The SSM-2015 was selected because this device has practically equal propagation time between inverting and noninverting inputs (typically less than 10ns differential). This important characteristic produces high accuracy L + R and L – R signals. Because U_1 and U_2 are ultra-low noise audio preamplifiers, they contribute little noise and distortion to output signals.

U_5 (OP-215GP) is a dual JFET amplifier, and is utilized as a long time constant integrator, or DC servo amplifier. The noninverting input is referenced to ground ($0V_{DC}$) and will hold the output terminals of U_1 and U_2 at $0V_{DC}$. Capacitor (AC) coupling is not recommended as it would allow formation of envelope and low frequency group delay distortion.

CONSTRUCTION REQUIREMENTS

All 10kΩ resistors must be matched within 0.05% of each oth but can be 5% in value tolerance. The 0.1μF capacitor in t integrator circuit should be a metalized polyester film capaci with 10% tolerance. The power supply rails are regulated ±18V_{DC}.

TABLE 1: Circuit Performance Specifications

Frequency Response (± 0.02dB)	20Hz to 20k
Dynamic Range (20kHz Bandwidth)	104
THD + Noise (20Hz to 20kHz, +24dBu)	0.007
IMD (SMPTE 60Hz & 4kHz, 4:1, +24dBu)	0.015
Slew Rate	10V/
Nominal Signal Level	–10d
Maximum Output Voltage (2kΩ Load)	+23.3dBu or 11.3V_R
Amplitude Accuracy	0.05
Differential Error	<10

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

A Two-Band Audio Compressor/Limiter

The two-band audio compressor amplifier shown in Figure 1 features separate and adjustable signal compression ratios and threshold levels. This design employs the SSM-2120 Dynamic Range Processor, and SSM-2134 low noise op amps.

The design features: an inverting or noninverting input buffer amplifier; high-pass and low-pass filters; and a voltage-controlled amplifier driven by a log-average level detector with a full-wave rectifier. Additionally, there are fully adjustable gain reduction controls, and adjustable compression threshold level controls. Signal-to-noise ratio exceeds 100dB, while the level detector allows the compressor to operate transparently throughout the audio spectrum without interaction between high and low bands.

The circuit includes a line level input (–10dBu to 0dBu nominal) buffer amplifier, with inverting or noninverting inputs that have 10kΩ input impedance. The buffer also limits step-function slewing voltages from entering the next stage and isolates the input source from the high-pass and low-pass filter stages.

Both high-pass and low-pass filters are of the single-pole type. This filter class eliminates combing effects in the stop bands, and compensates for compressor artifacts when the two bands are summed back together in the output amplifier stage. Such artifacts include the compressor's effect on the amplitude of each independent band, and the filter's effect on the summing influence has been taken into account in the design.

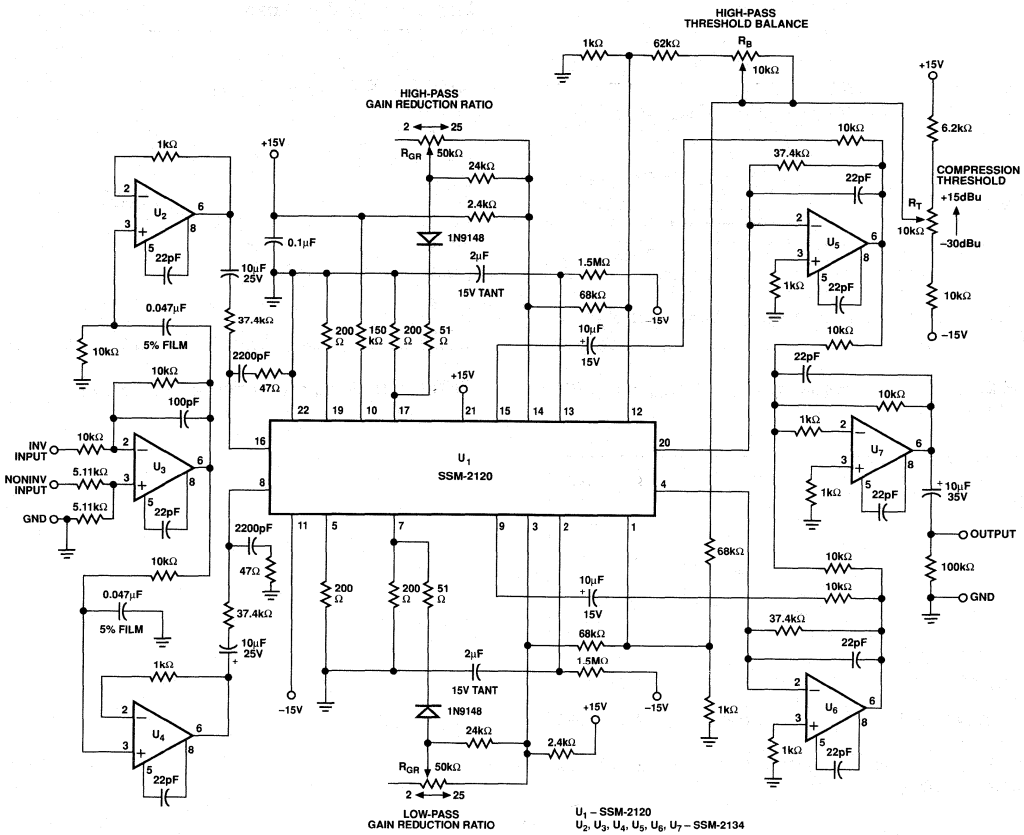

FIGURE 1: *Two-Band Audio Compressor*

Two continuously variable gain reduction controls (R_{GR}) in the VCA control circuit provide independent adjustment of compression gain slopes. The GAIN REDUCTION ratios are each adjustable from 2 to 25 for high-pass and low-pass bands. This range of adjustment provides mild compression to severe limiter/clipper action independently on each band. Thus the irritating "hole producing and pumping" character of single, wide-band compressor circuits can be reduced. The SSM-2120 provides a dynamic range of greater than 100dB over the frequency range of 20Hz to 20kHz with typically less than 0.02% THD + noise, and 0.05% IMD. Figure 2 shows THD+N vs. frequency, and Figures 3 and 4 provide compression and ratio characteristics.

The Compression Threshold control (R_T) allows the compressor to take effect from – 30dBu to +20dBu input levels. The SSM-2120's log-average precision rectifier is configured as a feedback-type level detector. The design produces consistent and precise compression profile of the input signal with no threshold level or compression drift over time and temperature.

The SSM-2120's full-wave log-averaging rectifier and control amplifier form an integrator and buffer circuit that isolates the low impedance VCA control port from the integrator timing circuit. This circuit senses the VCA output level and modifies its compression profile by feeding the averaged VCA signal plus the compression threshold control signal back into the VCA control ports. The control R_B balances the threshold amplitude between the two bands to pre-establish the compressor dynamics.

The small signal averaging time for the 2µF integration capacitor shown is 12ms. The attack time to 3dB of final value is about 6ms and is almost independent of signal level increase for level changes in excess of +10dB. The compression release rate is controlled by the 1.5MΩ discharge resistor in the integrator circuit. The recovery time constant is nearly linear since the discharge resistor returns to the $-15V_{DC}$ rail, and the rectifier produces a positive control voltage.

TABLE 1: Circuit Performance Specifications

Input Voltage Range (Nominal for 0dBu Out)	–10dBu to 0dBu 245mV to 755mV
Rectifier Type	Averaging
Compressor Amplifier Class	Feedback
Attack Time (+10dB or Greater Level Change)	6ms
Recovery Rate	1.67dB/ms
Feedthrough, Trimmed	–100dB
Gain Reduction Range	2 to 25
Frequency Response (20Hz to 20kHz)	0.2dB
Dynamic Range @ 0dB Gain	100dB
THD + Noise (20Hz to 20kHz)	0.02%
IMD (SMPTE 60Hz & 4kHz, 4:1)	0.05%
Output Voltage Slew Rate	6V/µs
Output Voltage (2kΩ Load)	+22dBu or $10V_{RMS}$

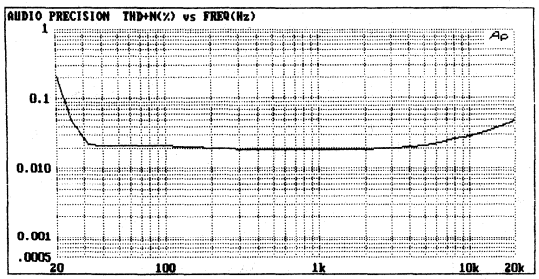

FIGURE 2: *THD+N vs. frequency (V_{IN} = $2V_{RMS}$,using 80kHz low pass filter)*

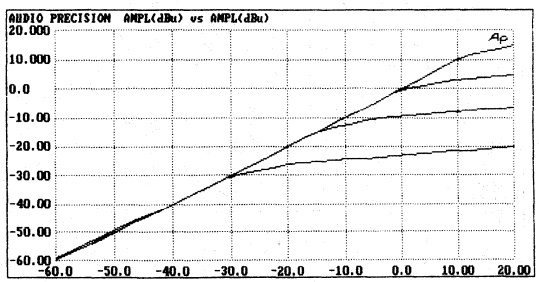

FIGURE 3: *Compression threshold variance characteristics (–30dBu to +15dBu)*

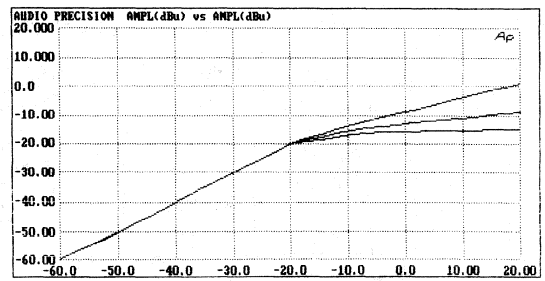

FIGURE 4: *Ratio characteristics (2:1 to 25:1)*

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

A Two-Channel VCA Level (Volume) Control Circuit

The dual-channel voltage-controlled amplifier (VCA) level control circuit describes a useful application of the SSM-2122 dual VCA, SSM-2134 low noise op amp, and PMI's OP-215BP JFET/bipolar op amp. This circuit is very handy when extremely close gain matching of a stereo audio source is desired, such as in ON-AIR and production audio consoles.

The design features a balanced input buffer amplifier and VCA driven by a level shifting amplifier which is controlled by a single 10kΩ linear potentiometer. Additionally, there are fully adjustable and independent gain limit and maximum attenuation trim controls. The VCA circuit has a nominal attenuation range greater than 95dB and has input overdrive protection. The signal-to-noise ratio exceeds 100dB with a gain of 10dB, and headroom of 32dB. The amplitude varies less than ±0.1dB over the frequency range 20Hz to 20kHz. Typical THD and IMD are less than 0.005% and 0.02%, respectively.

As shown, the circuit includes two line-level inputs designed for a –10dBu input signal level. The SSM-2134 (U_2 and U_4) input buffer amplifiers can be connected for balanced or unbalanced inputs with inverting or noninverting inputs. The input loading impedance is 10kΩ unbalanced and 20kΩ balanced. The input buffer amplifier also limits step function slewing voltages from entering following stages. Other input levels can be accommodated by adjusting the feedback resistor R_{F2}. For example: for a nominal input level of 0dBu, R_{F2} should be changed to 3.16kΩ, or for a nominal input level of +10dBu, R_{F2} changed to 1kΩ to provide the optimal current drive to the VCA. C_F should also be changed to 68pF and 220pF, respectively, for both U_2 and U_3.

For other input levels, R_{F2} can be calculated:

$$R_{F2} = 10 \times 10^3 \times EXP\left(\frac{10 + dBu}{-20}\right)$$

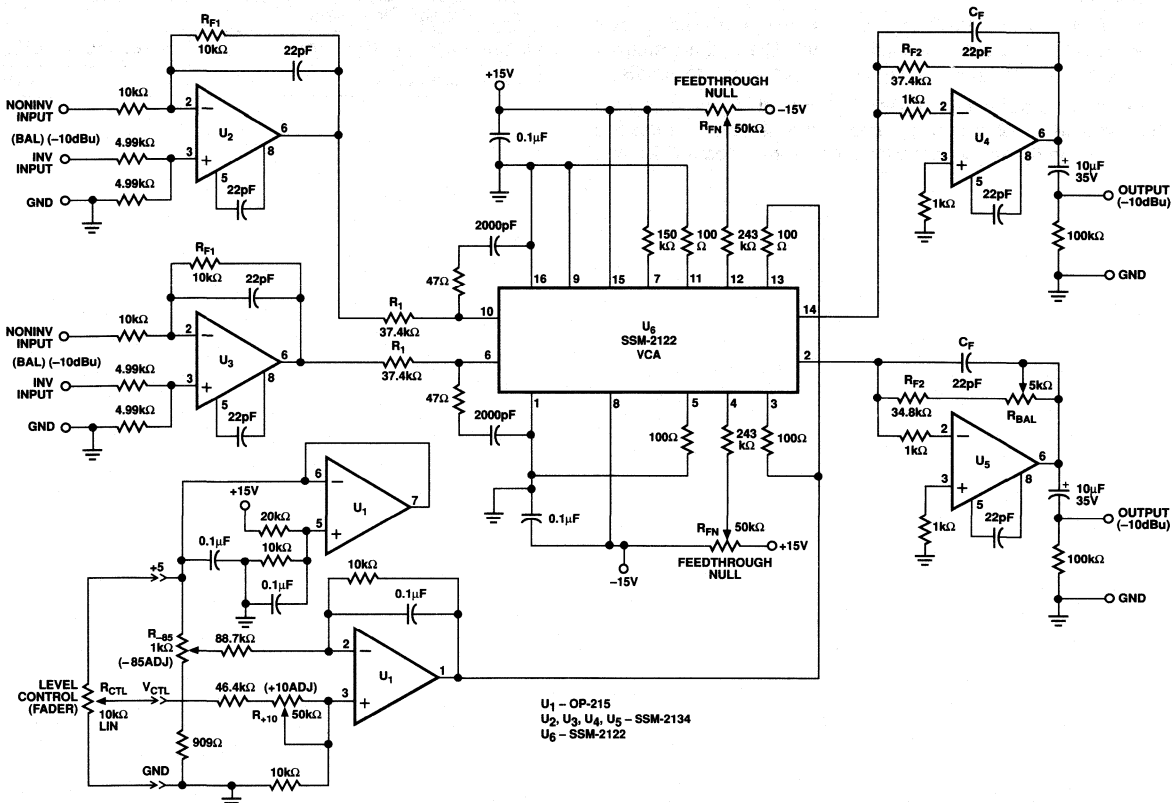

FIGURE 1: *Two-Channel VCA Level Control*

The SSM-2122 has a current-in and current-out structure. The input current is set by resistor R_1 and the virtual ground input of the SSM-2122. Similarly the transimpedance amplifier at the output converts the output current into a voltage. All devices in this design operate on $\pm 15V_{DC}$ power supply rails. The 37.4kΩ VCA input and output resistors optimize its dynamic range and minimize distortion. The SSM-2122 is a monolithic device, so the VCA gains remain uniform over a wide change in ambient temperature.

The SSM-2122 has two gain control ports that have a sensitivity of –6mV/dB. 0.0 volts at the gain control ports will yield 0dB overall gain; +60mV produces +10dB gain and –0.513V corresponds to a –85dB attenuation. The feedthrough trim null controls R_{FN} are not imperative for most applications. However, for very high performance requirements they will reduce attenuation control voltage feedthrough to less than 750µV. To adjust the null trim controls R_{FN}, inject a 100Hz sinewave into the control port through a 1kΩ resistor and a 100µF, 10V capacitor, and set the signal generator to 0.5 V_{RMS}. The control ports of U_6 are pins 3 and13. Adjust the level control R_{CTL} (fader) for 0dB gain, with the signal inputs shorted, then adjust R_{FN} for minimum 100Hz signal at the outputs.

The output amplifier(s) U_4 and U_5 are virtual ground connected current-to-voltage converters. The 37.4kΩ feedback resistors set the circuit voltage gain to 0dB with zero volts applied to the VCA control ports. Variable resistor R_{BAL} is used to balance the signal path gain of the two audio circuits. With the circuit gain set to 0dB and a test signal applied to the inputs, R_{BAL} is adjusted for equal output levels.

The VCA provides 10dB of additional gain at maximum level setting (+60mV at the VCA control ports). The THD is extremely low within the characteristic gain range of +10dB to –20dB.

The VCA control circuit is designed around U_1, the high input impedance OP-215 dual op amp. One half of U_1 is used to develop the 5V reference voltage for the level control element.

This is a fail-safe design – with no voltage applied or an ope connection at terminal V_{CTL}, the gain will descend to –85d Level control trimming is as follows: with the fader control set minimum (0V), the trim control R_{-85} is adjusted for maximu attenuation of –85dB or $-0.513V_{DC}$ at pin 1 of U_1. Then with th fader set to its maximum (5V), trim control R_{+10} is adjusted f maximum circuit gain of +10dB or +60mV. Since there is r interaction when adjusting R_{+10} for +10dB gain, the setting f R_{-85} will remain unaffected. Other max-attenuation values ca be used. R_{-85} has an attenuation range of –45dB(270mV) to 93dB (560mV).

TABLE 1: Circuit Performance Specifications

Input Voltage (Nominal for –10dBu Out)	–10dBu or 245mV$_{RM}$
Input Impedance, Unbalanced	10k
Imput Impedance, Balanced	20k
Headroom (Nominal for –10dBu In &Out)	32d
Feedthrough, Trimmed	<750µ
Gain Control Range (Nominal)	+10dB to –85d
Gain Control Voltage (+10dB to –85dB)	5V$_{DC}$ to 0V$_{D}$
Frequency Response (20Hz to 20kHz)	±0.1d
S/N Ratio @ 10dB Gain	100d
THD + Noise (20Hz to 20kHz, +22dBu)	0.005
IMD (SMPTE 60Hz & 4kHz, 4:1, +22dBu)	0.02
Output Voltage Slew Rate	6V/µ
Output Voltage (1kΩ Load)	+22dBu or 10V$_{RM}$
Output Impedance	<10

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

A High-Performance Compandor for Wireless Audio Systems

Wireless audio systems are finding increasing use in live performances, as well as in communications equipment where mobility is required. Designing such systems presents a difficult challenge: how to maintain adequate audio performance in view of power supply and current consumption limitations. To reduce transmission noise, the audio signal is usually compressed at

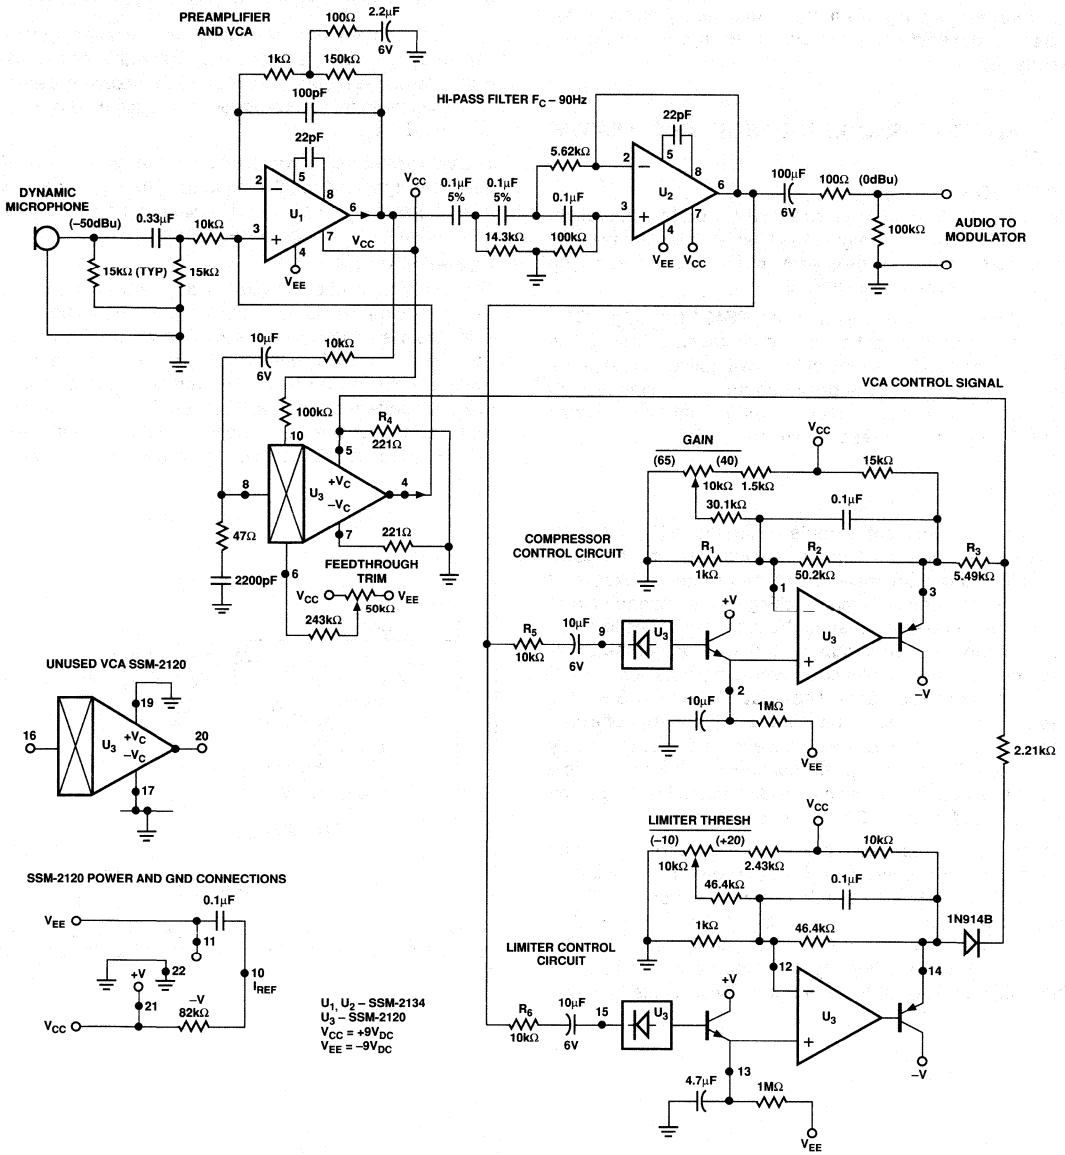

FIGURE 1: *Compressor-Limiter Circuit for Transmitter*

the transmitter and is expanded at the receiver by using a telecommunications industry-standard compandor IC, which has marginal audio performance as measured by professional standards. This application note describes a companding system utilizing the SSM-2120 Dynamic Range Processor, which permits considerable improvement over other techniques in terms of noise, distortion, feedthrough, and other key audio criteria.

Transmitters are battery powered and, hence, pose the most severe constraints on supply voltages and current consumption. Receivers are often AC powered, so bipolar supplies are more easily accommodated. Since the SSM-2120 requires split supplies, a voltage doubler circuit is necessary for the transmitter. In some cases, this may be considered unfeasible. In this event, however, the SSM-2120 is still very useful in the receiver expander circuit to complement any compressed signal, and to improve overall system performance. As a result, the compressor and expander sections of this application can be considered independently.

THE TRANSMITTER COMPRESSOR AND LIMITER CIRCUITS

COMPRESSOR

The design described is intended for $\pm 9V_{DC}$ battery operation, and includes a third-order high-pass filter for the elimination of subsonic noise and low frequency pops that would cause compandor overload or mistracking.

Figure 1 shows the connection of the SSM-2120 (U_3) VCA_1, rectifier, and control amplifier as a compressor. The VCA is connected in the feedback loop of the preamplifier U_1 to control the gain. The compressor is designed for a 2:1 compression characteristic. If the input rises 6dB, the output level will rise only 3dB. The gain compression expression is:

$$\text{Gain}_{\text{reduction ratio}} = \frac{R_2 \cdot R_4}{R_1 \cdot R_3}$$

as long as the rectifier input currents are limited by R_5, R_6 (10kΩ), and the rectifier has a ≈10μA reference current. The SSM-2120 rectifier and VCA have a dynamic range in excess of 100dB, resulting in exceptional tracking of the expander/compressor in the compandor system. High quality capacitors and resistors should be used to support the accuracy of the SSM-2120 elements. The small-signal averaging time for a 10μF integration capacitor is 25ms. The attack time to 3dB of final value is also about 26ms and is almost independent of signal level increases for level changes in excess of +10dB. The decay rate is 3ms per dB. The high-pass filter keeps frequencies below 90Hz from the input of the rectifier, reducing the low frequency distortion caused by the VCA control circuit.

DC and high-frequency feedback are provided for U_1 without sacrificing bandwidth or stability. The gain control is adjusted for 0dBu output with -50dBu applied to the microphone input terminals. The VCA is signal inverting. Its output current is summed, along with the microphone signal current, at the (virtual ground) noninverting input of preamplifier SSM-2134, U_1. The 10kΩ resistor at the input of the VCA limits the input current, iswhile the 2200pF, 47Ω network provides frequenc compensation for the VCA, keeping it stable.

The 100kΩ resistor from V_{CC} to pin 10 of U_3 establishes th operating current for the VCA. To minimize power supp current, all pins of the unused VCA should be returned t ground. The FEEDTHROUGH trim control is optional, and it ca be used to minimize the VCA control voltage from feedin through to the output.

PROTECTION LIMITER

The limiter uses the second rectifier and control amplifier fc separate and independent attack and decay times, along with steeper gain reduction slope. The limiter threshold control se the predetermined gain limiting point for high input signal level The gain reduction ratio is 4.6:1 as shown in Figure 1. Typicall the onset of gain limiting should be set to +10dBu at the outpu

As in the compressor control circuit, the rectifier input current i limited by R_6, 10kΩ, and the rectifier referenced to ≈10μA a well. Lower precision capacitors and resistors can be use here. Similar to the compressor, the attack time is much faste than the decay.

The VCA/Preamplifier was designed as a system. The VCA wa put in the signal feedback loop of the preamplifier principally t prevent preamplifier overload, while keeping the overall nois low, and minimizing component count.

POWER SUPPLY

The application circuit requires two power supply voltage: $\pm 9V_{DC}$. The power consumption, for the circuit shown in Figur 1, is less than 15mA from each supply. The design describe will operate properly with good dynamic range as the batter voltage begins to fall below the nominal $9V_{DC}$. It is assumed tha two 9 volt batteries would be used, but for the smaller hand-hel wireless microphones, a single 9 volt battery would be require Figure 2 depicts a DC-to-DC converter that will supply th $-9V_{DC}$.

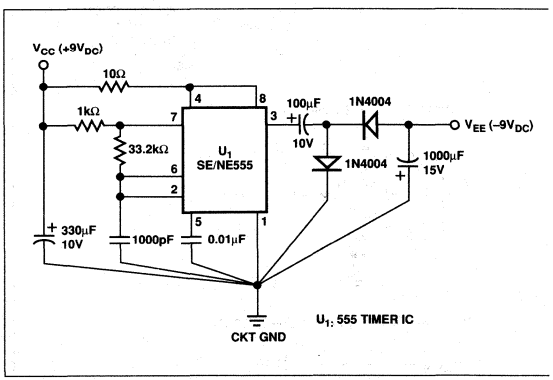

FIGURE 2: *DC-to-DC Converter for $+9V_{DC}$ to $-9V_{DC}$ at 15mA*

The converter circuit incorporates an astable oscillator running at 25kHz. It is followed by a capacitor-coupled level shifter and rectifier with a filter. A SE/NE555 timer is used in the "output sink" mode for maximum efficiency, and longest battery life.

THE RECEIVER EXPANDER SECTION

EXPANDER CONTROL

In Figure 3, the control connection of the SSM-2120 (U_3) VCA, rectifier, and control amplifier is shown. The control circuit connection to the VCA produces a 1:2 gain expansion curve. If the input rises 3dB, the output level will rise 6dB. The gain expansion ratio expression is:

$$\text{Gain}_{\text{expansion ratio}} = \frac{R_2 \cdot R_4}{R_1 \cdot R_3}$$

The rectifier input current is limited by a 10kΩ resistor connected to pin 9 of U_3, and the rectifier is biased at 10µA current through a 1.5MΩ resistor connected to V_{EE}. The SSM-2120 rectifier and VCA each have a 100dB dynamic range, resulting in accurate tracking of the compressor.

As with the compressor/limiter circuit, the small-signal averaging time for a 10µF integration capacitor is 26ms. The attack time to 3dB of final value is also about 26ms and is almost independent of signal level increases for level changes in excess of +10dB. The decay rate is 3ms per dB.

The control circuit gain values, as shown above, provide a control voltage to the VCA section $+V_C$ control port [U_3, pin 5(19)], which result in a 1:2 signal expansion characteristic. The LEVEL control sets the initial overall gain value, and is adjustable from –10dB to +20dB.

EXPANDER AUDIO

Input amplifier U_1 is a buffer between the input signal source (FM wireless receiver), the expander rectifier/control circuit, and the VCA audio signal input. If the signal source output impedance is below 100Ω, U_1 can be omitted. The nominal source signal level should be –10dBu. If signal gain or loss is required, U_1 gain structure should be modified to provide –10dBu to the VCA input current limiting resistor. The 37.4kΩ resistor ahead of the VCA input [pin 8,(16)] limits peak signal currents to avoid VCA distortion. The VCA signal input(s) are virtual ground current inputs. The 150kΩ resistor connected to V_{CC} and pin 10 of U_3 sets the VCA input/output current compliance range. A VCA input shunting capacitor shown from U_3 pin 8(16) to ground minimizes signal distortion and keeps the VCA stable by providing a high-frequency path to ground. The exact value is determined empirically. The output of the VCA feeds a virtual ground output amplifier, U_2. The overall audio path is signal noninverting, since the VCA is signal noninverting, and is combined with two inverting amplifiers U_1, and U_2.

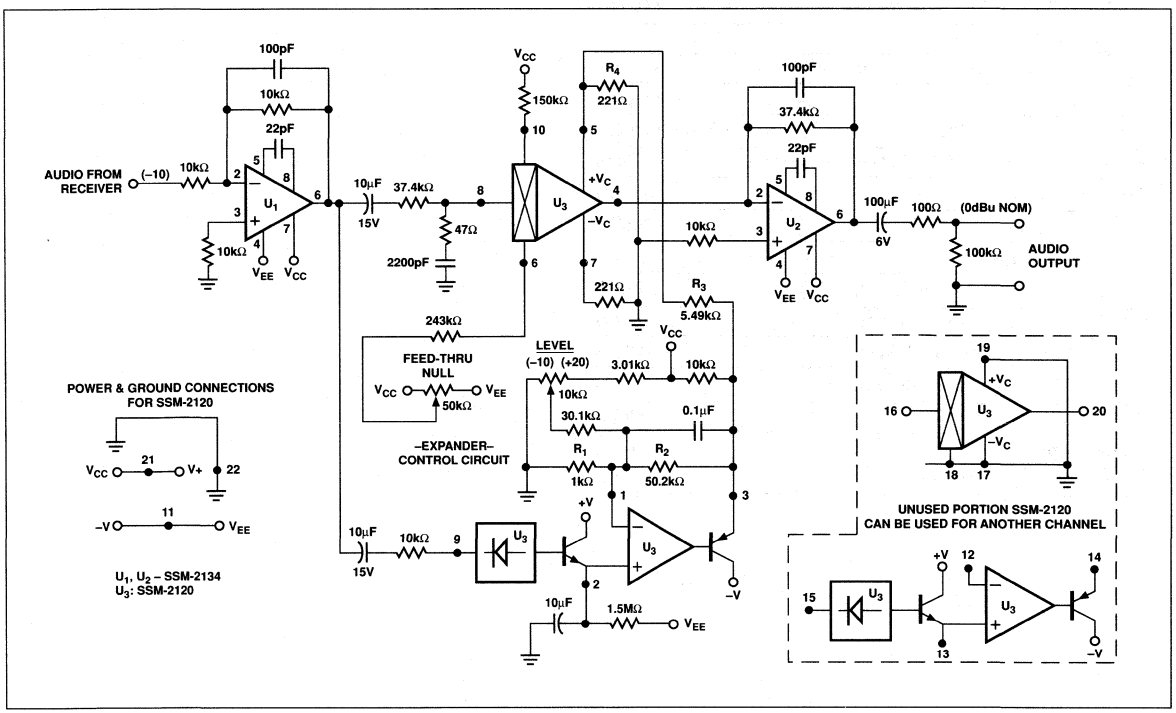

FIGURE 3: *Expander Circuit for Receiver Processor*

TABLE 1: Circuit Performance Specifications

Compressor/Limiter (Figure 1)	
Input Voltage Range (Nominal for 0dBu Out)	–65dBu to –40dBu
Compressor/Limiter Rectifier Type	Averaging
Compressor Amplifier Class	Feedback
Compressor Attack Time (to 3dB of Final Value)	26ms
Compressor Recovery Rate	3ms/dB
Feedthrough (Trimmed)	–70dBV
Compressor Gain Reduction Ratio	2:1
Limiter Attack Time (to 3dB of Final Value)	13ms
Limiter Recovery Rate	1.5ms/dB
Limiter Gain Reduction Ratio	4.6 to 1
High-Pass Filer, F_C	90Hz
High-Pass Filter Type	3rd Order Butterworth
Frequency Response (90Hz to 20kHz)	±0.5dB
S/N Ratio @ 0dB Gain	100dB
THD + Noise (% 1kHz to 20kHz)	0.05
IMD (% SMPTE 60Hz & 4kHz, 4:1)	0.05
Output Voltage Slew Rate	6V/µs
Power Supply, (Battery)	$\pm 9V_{DC}$
Output Voltage (2kΩ Load, $\pm 9V_{DC}$	+15dBu
Expander (Figure 3)	
Input Voltage Range (Nominal for 0dBu Out)	–10dBu
Expander Rectifier Type	Averaging
Expander Amplifier Class	Control Feed-Forward
Expander Attach Time (to 3dB of Final Value)	26ms
Expander Recovery Rate	3ms/dB
Feedthrough (Trimmed)	–70dBV
Gain Expander Ratio	1:2
Frequency Response (20Hz to 20kHz)	±1.0dB
S/N Ratio @ 0dB Gain (Dynamic Range)	100dB
THD + Noise (% 1kHz to 20kHz)	0.20%
IMD (% SMPTE 60Hz % 4kHz, 4:1)	0.25%
Output Voltage Slew Rate	6V/µs
Power Supply	$\pm 15V_{DC}$
Output Voltage (2kΩ Load)	+21dBu

AC coupling is not necessary but recommended. When the AC coupling is combined with feedthrough trimming, most of the unwanted sub-audible signals will be removed from the output signal. The unused portion of U_3 (SSM-2120) can be utilized in an identical second channel expander. The supply voltage should be held at ±15VDC to ±18VDC, to provide good dynamic range and circuit stability.

ANALOG DEVICES

AN-134
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

An Automatic Microphone Mixer

This application note describes an audio-signal activated microphone mixer, as shown, designed to accommodate eight input channels. The SSM-2120 Dynamic Range Processor is the nucleus of this design. The device includes two VCAs and two rectifier and control amplifier circuits. The application is designed for unattended microphone mixing functions, as would be used in a conference room that required sound reinforcement or conversation recording. The circuit provides automatic and transparent channel ON/OFF operation. The audio output automatically turns ON in less than 10ms and back OFF after 2 to 4 seconds of no audio. Each channel incorporates independent and automatic operation, with ON threshold sensitivity and level adjustment (trim) controls.

THE MICROPHONE PREAMPLIFIER CIRCUIT

For optimum circuit performance, the nominal output level from the microphone preamplifier or other audio source(s) should be –10dBu. The –10dBu level allows the SSM-2120 to provide the widest dynamic range and lowest noise, while optimizing the headroom of the automatic switch and the microphone preamplifier.

THE AUDIO SWITCH

Each audio signal is switched ON and OFF by a VCA (voltage controlled amplifier) element. By controlling the turn-ON and turn-OFF ramp time, the VCA produces transient-free switching. No distortion or "pops" are introduced using this technique. The turn-ON is ramped from a maximum of 90dB to 0dB attenuation in approximately 30ms. With the complete removal of the audio, the turn-OFF ramp of ≈100ms will begin after approximately three (3) seconds.

The VCA's input is a current input, virtual ground node. The design shown in Figure 1 assumes that $\pm 15V_{DC}$ will power the system. With this supply voltage, the 37.4kΩ input resistor(s) will keep the VCA operating at the optimum distortion and dynamic range. The 150kΩ resistor connected from V_{CC} ($+15V_{DC}$) to the reference current pin 10 of the 2120, sets the VCA bias operating point. The VCA's current output is then connected to a voltage by a transimpedance amplifier using a low noise SSM-2134 op amp. The VCA input(s) and output(s) are capacitively coupled to remove DC components from previous stages.

THE CONTROL CIRCUIT

The input signal is rectified and averaged before it is filtered by the integrator capacitor (10μF electrolytic). The small signal averaging time constant with this capacitor is approximately 60ms. The attack time is 30ms to 3dB of the final level, and is nearly independent of the magnitude of level increase. The discharge time is controlled by the 3.3MΩ resistor returned to V_{EE} ($-15V_{DC}$), which also sets the rectifier reference current.

The control circuit amplifier has a voltage gain of 217. The inverting input is used to set the ON/OFF threshold point, too, and allows for the ON Threshold level adjustment. The ON Threshold range is adjustable from -40dBu to 0dBu as referenced to the input of the VCA element. The control port +Vc of the VCA is used so that a negative control voltage applied will produce an attenuation effect. R_4 (221Ω) and R_5 (4.64kΩ) attenuate the control voltage by a factor of 22, resulting in a maximum attenuation value of 90dB with no audio signal present.

To minimize the effect of ON/OFF control voltage appearing in the output signal of the VCA, the Feedthrough null control circuit is recommended. It provides an external method for balancing the internal VCA currents and component values.

THE OUTPUT SECTION

The design incorporates a virtual ground current summing bus, that is fed by the individual mixing channel level control(s) and 10.0kΩ resistor(s). The Level control(s) provides 21dB attenuation range (0dB to –21dB), and is designed to balance the different inputs, but not turn them off fully. The 100kΩ (linear taper) level control(s), for a linear rotation produces a logarithmic attenuation curve.

The virtual ground summing amplifier establishes half of the balanced output circuit, with another inverting amplifier completing the balanced output circuit. The circuit is able to drive 600Ω loads to +24dBm levels with low distortion and high reliability.

TABLE 1: Circuit Performance Specifications

Input Voltage, without Preamplifier, (for +4dBu Out)	–10dB
Input Impedance, Unbalanced	≈1kΩ
Headroom (Nominal for –10dBu In and Out)	32dB
Turn ON Time (to 3dB of Final Value)	30ms
Turn OFF Time (No Signal)	≈3sec
Turn OFF Ramp Time	100ms
Feedthrough (Trimmed)	>1mV
ON/OFF Threshold Range (Nominal)	0dBu to –40dBu
ON/OFF Gain Extent	0dB to –90dB
Frequency Response for ±0.1dB	20Hz to 20kHz
S/N Ratio @ 0dB Gain	110dB
THD + Noise (from 20Hz to 20kHz)	0.005%
IMD (SMPTE 60Hz and 4kHz, 4:1)	0.02%
Output Voltage Slew Rate	12V/μs
Rated Output Level (600Ω Load)	+24dBu
Output Impedance	68Ω
Output Type	Balanced
Power Supply Requirements	$\pm 15V_{DC}$ Regulated

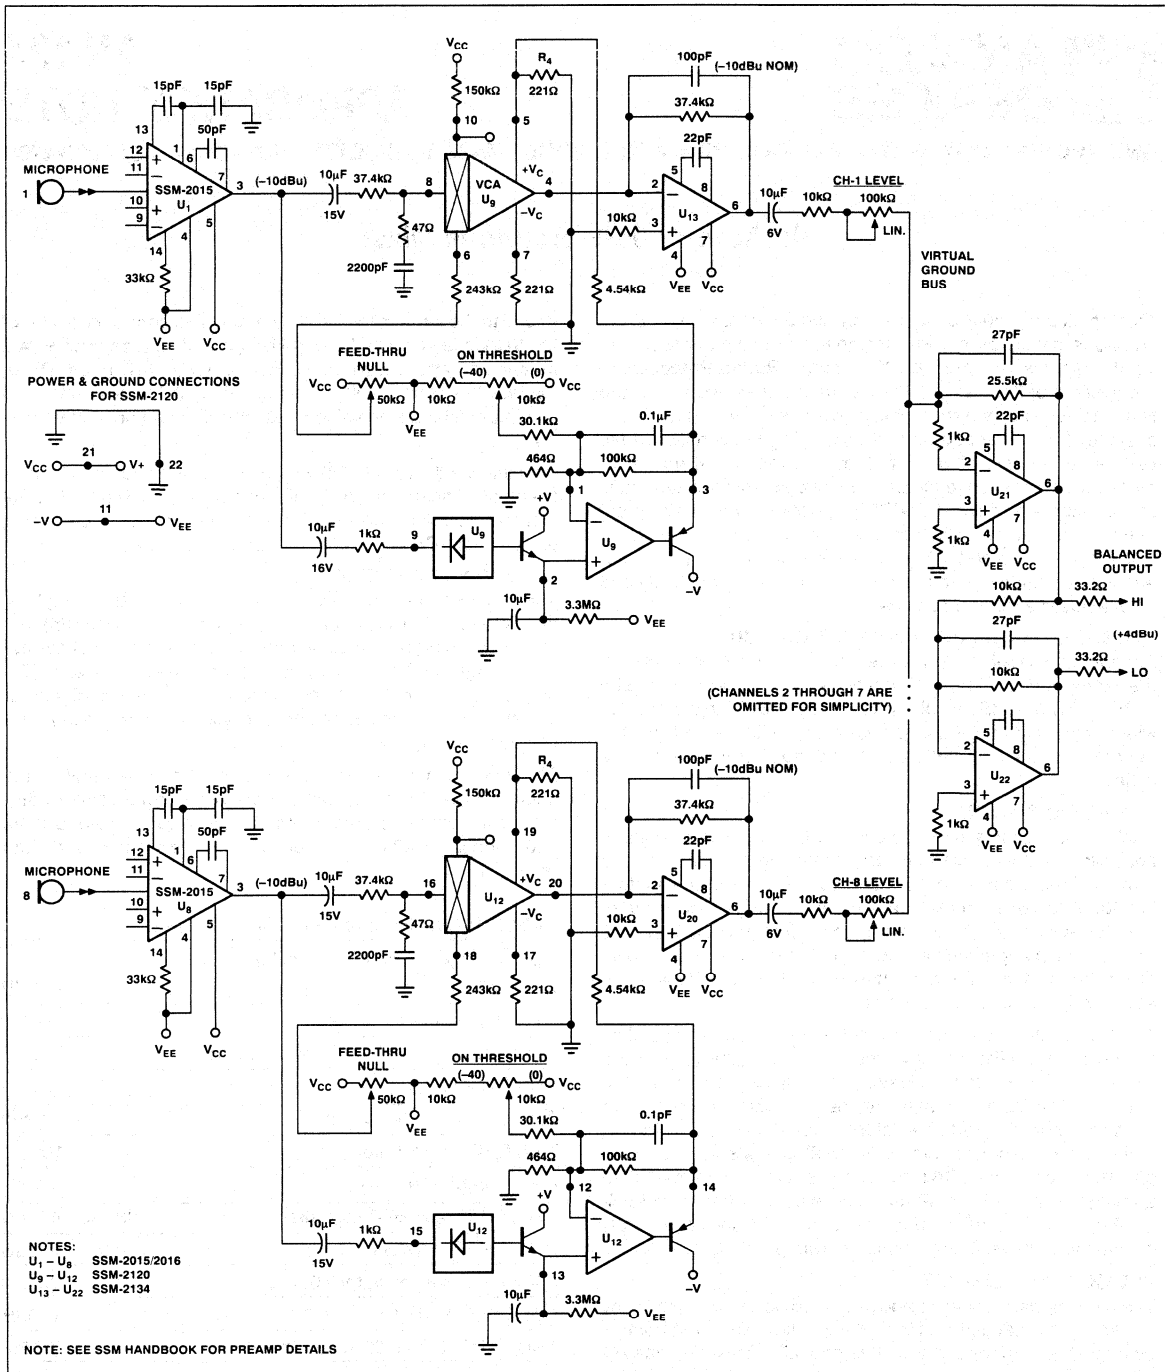

FIGURE 1: *Automatic Channel Activation Microphone Mixer Diagram Illustrates 8 Input Channels*

ANALOG DEVICES

AN-135
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

The Morgan Compressor/Limiter

The following application was written by Michael Morgan, a consultant to PMI with extensive experience in the design of professional audio equipment, including high-performance dynamic range processors. While currently a freelance consultant, Mr. Morgan spent nine years at Valley International as a principal designer of their products.

This application note describes the configuration of a low cost, high quality compressor with variable attack time and ratio control using the SSM-2110 Level Detector and SSM-2014 VCA. The discussion begins with an overview of compressor/limiter fundamentals, and is followed by a description of the unique attributes of the integrated circuits and their implications for the design engineer.

COMPRESSOR/LIMITER FUNDAMENTALS

The function of the audio compressor is, of course, to compress the dynamic range of the processed audio signal by altering the gain of its signal path in response to the relative level of the signal as compared to an arbitrary setpoint called the threshold, thus adding gain to low-level signals and reducing gain in the presence of high-level signals.

An audio compressor consists primarily of two functional sections, one of which derives a control signal by measuring and otherwise manipulating the audio signal to produce a voltage suitable for the second functional section, which is the gain control element. The gain control element is a device which can alter its attenuation, gain, or resistance in response to an external signal, such as a voltage or current.

The audio compressor differs from a similar device, called a limiter, in that a compressor exhibits a rotation point which is independent of its threshold setting. The rotation point is the locus on a graph of the compressor's transfer function at which the gain control element exhibits unity gain, and through which all lines derived from data describing the device's output level as a function of input level will pass. A limiter, in the purest sense, adds no gain and has no rotation point.

The compressor's ratio is defined as the increase in input level, in decibels, above the threshold which will result in an increase of output level equal to 1dB, and is a function of control circuitry gain. For example, each increase of 1dB in signal level above the threshold may cause a corresponding decrease in gain equal to 1dB, thus keeping the output level constant for a ratio of infinity:1, or it may cause a 1/2 dB decrease in gain, thus allowing the output level to rise at a ratio of 2:1. A compressor may have a very high ratio, and conversely, a limiter may have a very low ratio.

SHAPING THE RESPONSE

Figures 1a and 1b illustrate the transfer function of an ideal compressor as the ratio is varied with a fixed threshold, and as the threshold is varied with a fixed ratio, respectively. Note that in both cases, the rotation point is easily identified at 0dB. Note also that the compressor operates as a limiter when the threshold is equal to, or higher than, the rotation point. For this reason, the device described in this application note may be considered as a "compressor/limiter," but because it possesses a rotation point, we shall refer to the device as a compressor.

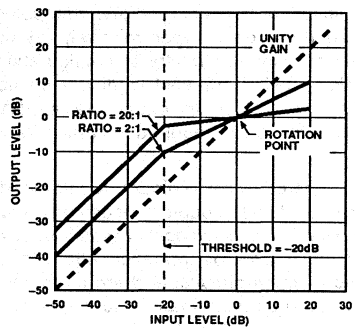

FIGURE 1a: *Output vs. Input Transfer Function of an Ideal Compressor*

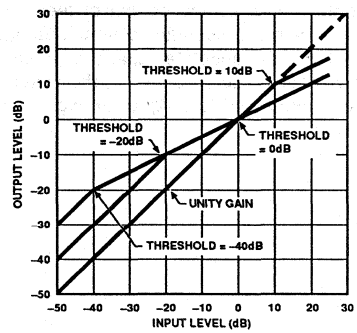

FIGURE 1b: *Output vs. Input Transfer Function of an Ideal Compressor Having a Fixed Ratio Showing the Effect of Threshold*

Audio compressors use two types of circuit topologies. In a feedback, or closed-loop configuration, the control signal is derived by measuring the output level of the gain control element. In a feedforward, or open-loop configuration, the control signal is derived by measuring the level of audio present at the input of the gain control device. Each topology has its typical advantages and disadvantages. The most common type of audio compressor uses the feedback topology. Among its advantages are: low parts cost; ease of configuration; ability to use simple, "linear" control circuitry and gain control elements. The disadvantages of the feedback topology are numerous: inability to realize continuously variable parameters accurately; heavy dependence upon circuit trimming to assure consistency in performance from unit to unit; tendency toward overshoot in either control signal or processed signal; virtual inability to configure circuitry for performing arbitrary dynamic functions, such as program control of release times, equalized sidechain functions, and interactive processing having more than one control function per gain control element.

The feedforward topology has long been considered by equipment designers to be the more versatile method of configuring audio dynamics processors. Among its advantages are: precise control of dynamics; ability to accurately and continuously vary processor parameters, such as ratio and attack-and-release time constants; easy circuit trimming for unit-to-unit consistency; possibility to realize arbitrary types of dynamics alteration; ease in configuration of interactive processing schemes using multiple control signals to operate a single gain control element; and relative freedom from control and signal overshoot. The disadvantages of feedforward topology have traditionally been: dependence upon relatively expensive and little understood logarithmic circuitry in configuration; difficulty in sourcing high-quality, low cost log/linear multipliers (dB/volt VCAs); dependence upon expensive log/RMS detection schemes to achieve the required accuracy for wide range of control.

By using integrated building blocks, feedforward control technology can be realized by equipment designers by virtue of their ease of application and low cost. These readily available integrated circuits deliver performance equal to or surpassing complicated discrete circuits, and are more cost effective for general use by equipment manufacturers.

THE SSM-2110 MONOLITHIC LEVEL DETECTOR

The SSM-2110 level detector IC represents a significant advancement in low cost, high quality converter circuitry. The device greatly simplifies the design of feedforward dynamic processors since it produces an accurate output that is proportional to the log of the absolute value of its input, and the log of the rms value of its input, in addition to the corresponding linear values. Such versatility is unique among detector/converter configurations.

VARIABLE TIME INTEGRATOR

In this application, the SSM-2110 is used in the design of feedforward compressor. The log of the absolute value of th input signal is extracted, then integrated by a E/I x C circu which corresponds roughly to an RC network in the linea domain (see Figure 2).

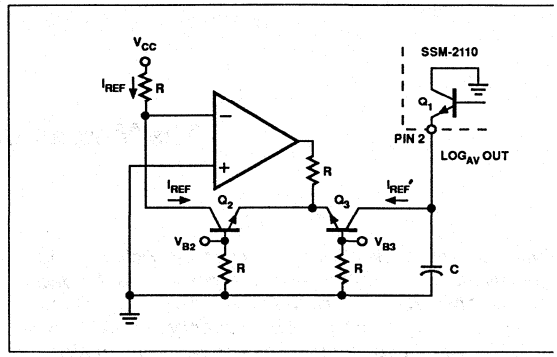

FIGURE 2: *Simplified Schematic of a Variable Time Constant Log of Average Integrator Using the SSM-2110*

The product of this operation is not, as one would expect, the a erage of the log of the absolute input value. During the integr tion process achieved by charging the integrator capacitor, (the charging current is proportional to the antilog of the voltag appearing at the base of Q_1. Since the voltage at the base of C represents the log of the absolute value of the input, the log ar anti-log terms cancel, thus leaving C to charge as a linear int grator with a current proportional to the absolute value of th input until the voltage across C approaches the voltage at th emitter of Q_1. In this manner, the order of the logging ar averaging operations are reversed. This is a very importa phenomenon which directly influences the audibility of the con pression process, and will be discussed at length later.

The integration time of the circuit in Figure 2 is varied by chan ing the current through the collector of Q_3. This is accomplishe by means of the multiplier circuit consisting of the operation amplifier, transistors Q_2 and Q_3, and their associated resistor

Current I_{REF} is forced to flow through the collector of Q_2. Th V_{be} of Q_2 is thus made to be proportional to the log of I_{REF} b virtue of the silicon transistor's intrinsic logarithmic propert idealized in the equation:

$V_{be} = kT/q \times \ln(I_c/I_s)$ where

k = Boltzman's constant (1.38×10^{-23}J/K)

T = Temperature in Kelvins

q = Charge on an electron (1.60×10^{-19}C)

I_c = Collector current

I_s = Reverse saturation current (extrapolated as $V_{be} \rightarrow$

When transistors Q_2 and Q_3 are closely matched, V_{be} of Q_2, which appears also at Q_3's emitter, causes a current equal to I_c of Q_2 to flow through the collector of Q_3. This transistor collector current, I_{REF}, may be used to charge or discharge a capacitor, to cause a voltage drop across a resistor, or may be converted to a voltage at the output of an operational amplifier. The collector current of Q_3 may be varied by applying a voltage at the bases of Q_2 or Q_3, or both bases simultaneously. As a rule of thumb, at 25°C, each 60mV change in V_b will cause a corresponding ten-fold change in Q_3's I_c. By using the "shorthand" log relationship for gain in which a ten-fold change in voltage (or current) equals 20dB, we can say that the collector current of Q_3 can be made to vary antilogarithmically at a rate of 1 dB/3mV (20dB/60mV). In effect, the circuit generates a voltage at the emitter of Q_2 which corresponds to the log of the input current, I_{REF}, adds the control voltage, then generates a current at the collector of Q_3 which is proportional to the antilog of the sum. Thus the portion of the circuit formed by the operational amplifier, Q_2, Q_3, and their associated passive components form a two-quadrant multiplier whose output is a high compliance current sink.

A positive voltage applied to the base of Q_2 will cause a corresponding decrease in Q_3's collector current, while a positive voltage applied to the base of Q_3 has the opposite effect, causing an increase in I_c of Q_3. Both bases may be controlled by bipolar voltages, but I_{REF} must flow in the direction indicated by conventional current flow through the transistors (must be sourced from a voltage more positive than the noninverting input of the operational amplifier for NPN transistors).

In operation, varying the current which discharges C also causes a varying offset voltage at the collector of Q_3 which equals the change in V_b, and must be compensated for in order to derive a useful control voltage. Figure 3 shows the response to a +10 volt pulse input having a repetition frequency of approximately 4 pps and a duty cycle of 50%. Note that the X-axis corresponds also to increasing integration time (decreasing I_c of Q_3). The illustration is a composite of several sampled waveforms, thus, scalar references in the X-axis are valid only for each pulse.

As can be seen in Figure 3, in the log average detection mode, the response of the device to large level changes is relatively fast, while the last 50 to 100mV of change occurs at the characteristic integration time determined by the status of the charge on the capacitor, C, as it is discharged by the collector current of Q_3. As the voltage across C approaches the voltage at the emitter of Q_1, the transistor behaves less as an antilog element, and more as a linear resistance proportional to V_{be}/I_{REF}.

These attributes determine the detector circuit's response to complex waveforms, and directly affect the audibility of the compression process. Consider the following explanation: Humans respond to changes in audio signal level by perceiving volume as being proportional to the log of the acoustic power emitted by a source, thus the human listener perceives a source emitting 10 watts of "sound," (if the reader will permit such simplifications) to be roughly twice as loud as the same source emitting only 1 watt. This implies that one should be able to control audio levels logarithmically for a natural "sound" in the processed output. That is generally the case, but the principle does not extend, in a strict sense, to the control of a compressor.

If one accepts the premise that the most common uses of the audio compressor are to enhance the "loudness" of the processed material, or to "level" the apparent volume of the processed material, one should be aware of the effect of waveform complexity upon perceived loudness. A simple example is found in the case of a musician playing an instrument: when called upon to perform a solo, in order to "stand out" from the background music, the instrumentalist produces more complex sounds, in addition to producing sounds at a higher relative level. The increase in complexity provides a psychoacoustic "cue" which translates to the human listener as increased perceived loudness.

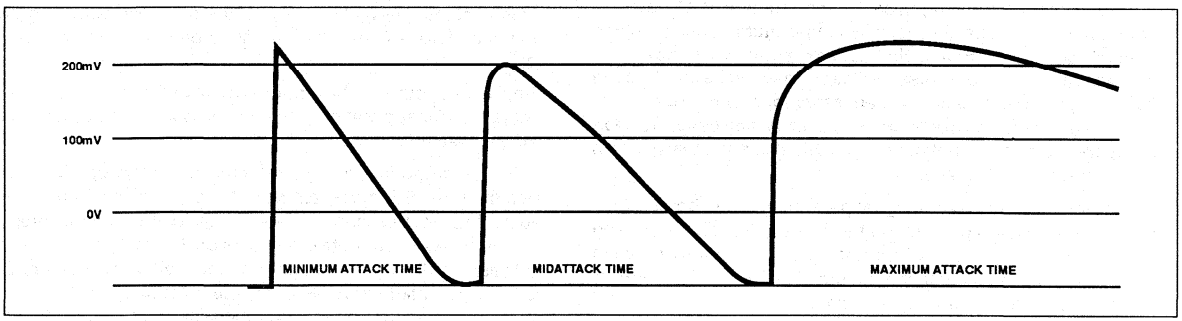

FIGURE 3: *Output Voltage Response of the Variable Time Constant Log of Average Integrator to a LF Square Wave Input*

During the compression process, if the detector circuitry produces a signal which calls for more gain reduction in response to the added harmonics in a sound which cause an increase in complexity, the gain control element will comply, thus making a solo instrumental exit the compressor at a lower level, foiling the intent of the performer. This is precisely what happens when using RMS detection – the detector circuit (correctly) assesses the increased complexity as an increase in sound energy, and calls for gain reduction.

The log averaging detector is relatively insensitive to increases in waveform complexity, and "ignores" the loudness cue thus provided. As a result, a complex waveform exits the compressor at a slightly higher level than it would if under the control of an rms detector. This rather unique "quirk" found in the log averaging process allows a solo instrumental or vocal to stand out in the processed signal, thus preserving the intent of the performer.

As a log averaging detector, the SSM-2110 exhibits remarkably little departure from an idealized log curve representing its input level throughout the entire range of the adjustable integration time, and offers superior performance to the equipment designer in this type of application. In addition, at higher input levels, the device does not compress the waveform at its output, thus under-reading the input value. In fact, the detector exhibits a gentle and quite predictable deviation from log conformity at high input current levels which results in the addition of a linear error term. This causes a slight over-reading of the input, and is quite useful for aficionados of "soft knee" limiting. It is unlikely that any real compressor design would require so wide a range of operation that this deviation might pose a problem (> 60 dB), but since the error term is so predictable and consistent, it can easily be corrected elsewhere in the control circuitry if necessary.

THE COMPRESSOR CONTROL CIRCUIT

Figure 4 illustrates a compressor control circuit incorporating the SSM-2110 as the detector element. Because the temperature compensation characteristics of the log recovery amplifier are not required in this application (control of a VCA having a complementary control sensitivity temperature coefficient) and to eliminate trimming of scale factor on a unit by unit basis, the log recovery amplifier is disabled by connecting its inputs to the IC's V_{REF} output. This step is necessary for proper operation of the IC when the log recovery amplifier is not used. The log recovery transistor is not used since offset is not a real concern in this circuit configuration.

The amplifier in the audio signal path, A_1, should be of a high quality, low noise type such as the SSM-2134. The remaining amplifiers may be general purpose types, preferably having FET input stages to minimize the effects of input bias currents on the accuracy of the multiplier circuits.

Amplifier A_2, and the SSM-2210 matched transistor pair Q_{1a}/Q_{1b} form the voltage-controlled current sink for the variable log averaging integrator. Amplifier A_4 boosts the output of the integrator to a usable level by increasing the nominal 6mV/dB scale factor of the signal at pin 2 to 1V/20dB, or 1V per decade. An offset corresponding to the change in voltage at pin 2 whic results from varying the integration time is applied via R_1; Variable resistor VR_2 allows the adjustment of the compressor' rotation point, or that input level at which the output of A_4 wi be 0V.

A pair of matching two-quadrant multipliers, which are confic ured using amplifiers A_6 and A_8 along with a four-transistc array Q_2 (MAT-04), allow adjustment of the compressor's rati and adds sufficient gain as a function of both the threshol setting from VR_4 and A_9, and the ratio, as determined by VR and A_5, to maintain the compressor's rotation point. Amplifier A converts the current output of the ratio multiplier from Q_{2b} int a voltage which charges the holding capacitor C_8 via Q_{3a} to voltage corresponding to the amount of gain reduction require of the VCA. Amplifier A_{12} converts the current output of th maintenance gain multiplier from Q_{2c} into a voltage correspond ing to the quiescent gain required for the VCA to maintain th compressor rotation point, and adds or reduces gain at the VC in response to the output gain control VR_7.

The release current sink is formed by amplifier A_{10}, and tw sections of monolithic transistor array Q_3. Compensation fc V_{be} of Q_{3a}, and for the quiescent change of voltage across C caused by varying the release current through Q_{3c} are applie via R_{31} to A_7.

Amplifier A_9 and diode D_3 form a precision half-wave rectifie whose output is a positive voltage equal to the gain reductio signal. This point may also source a gain reduction indicat with intrinsic scaling of +1V = –20dB. Since metering is a matt of preference for the design engineer, no attempt has bee made to include a gain reduction indicator as part of this circu discussion.

Since it is possible for the inputs of both A_6 and A_8 to be nega tive voltages, it is wise to include germanium diodes D_1 and D to prevent forward conduction of the internal base to emitter pr tective diodes in the MAT-04, thus eliminating the possibility reverse leakage coupling between the multipliers. The exis tence of these diodes also prevents application of the MAT-0 as the charging transdiode Q_{3a}, since a voltage more negativ than that across C_8 will frequently be present at the emitter Q_{3a}.

Amplifier A_{13} outputs the algebraic sum of the various gain co trol signals for application to the compressor VCA section whic will be described next.

Selection of the internal scale factor at 1V/decade, and incl sion of the –7.5V "math rail" are arbitrary choices made by th author in order to accommodate the use of the variable integr tor, and to skew the control markings on the front panel control indicated by the enclosed names associated with those variab resistor potentiometers. Placing the rotation point, as dete mined by R_2, at the nominal operating line level, e.g., +4dB, etc minimizes the effects of errors caused by the uncompensate temperature coefficient of the ratio multiplier, and any min deviations in log conformity inherent in the detector circuitry VCA sections of the compressor.

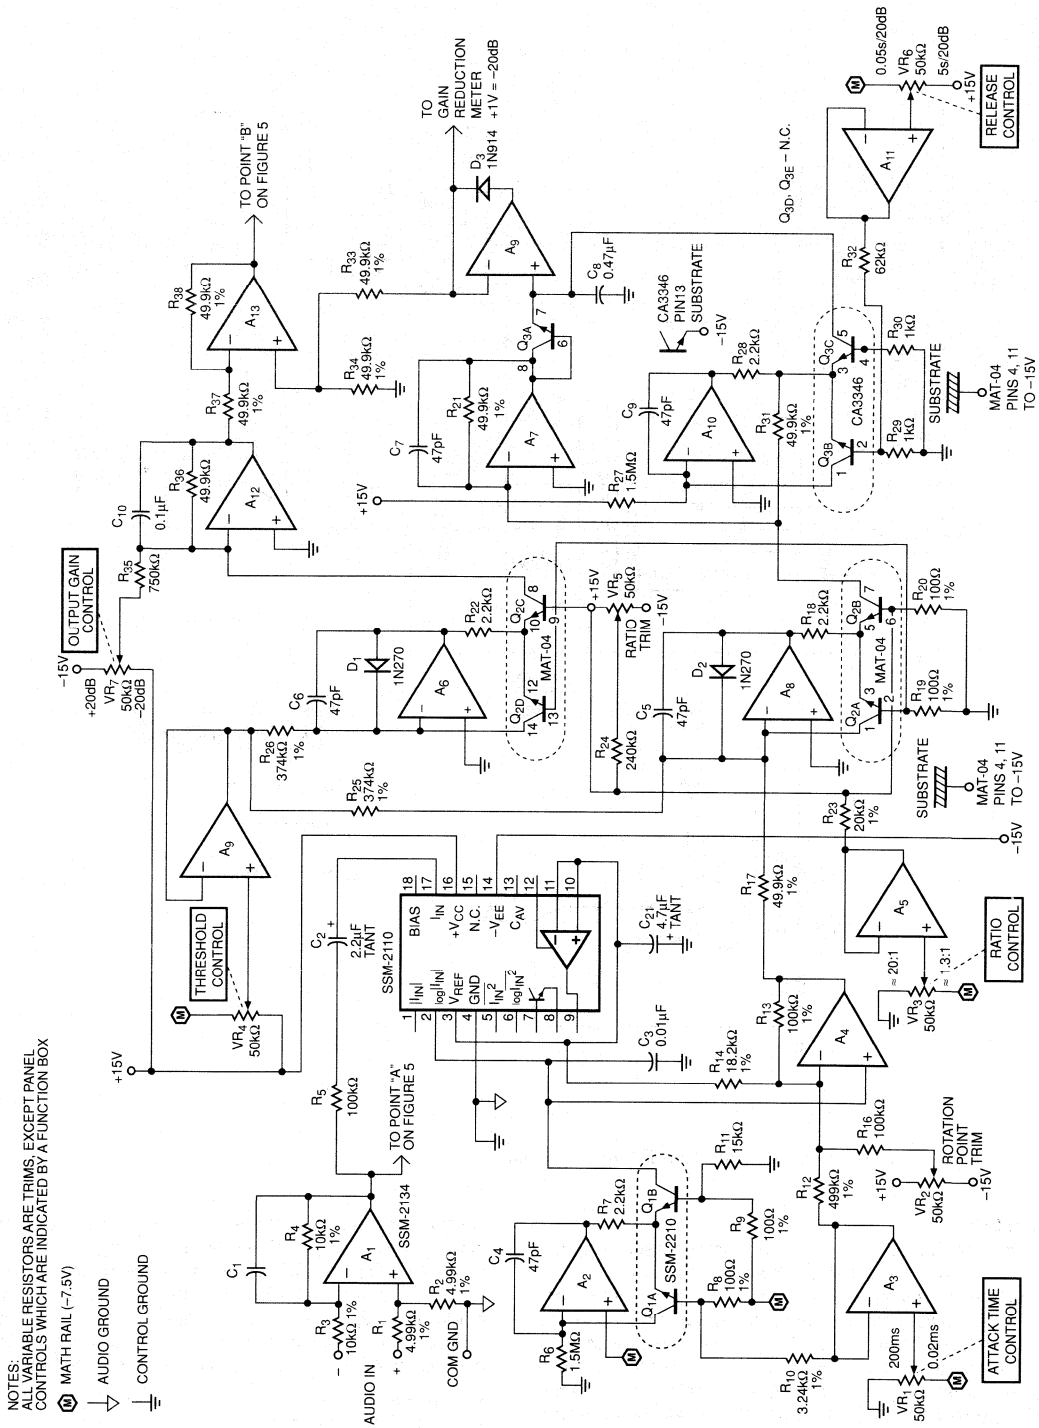

FIGURE 4: *Compressor Control Circuitry Schematic*

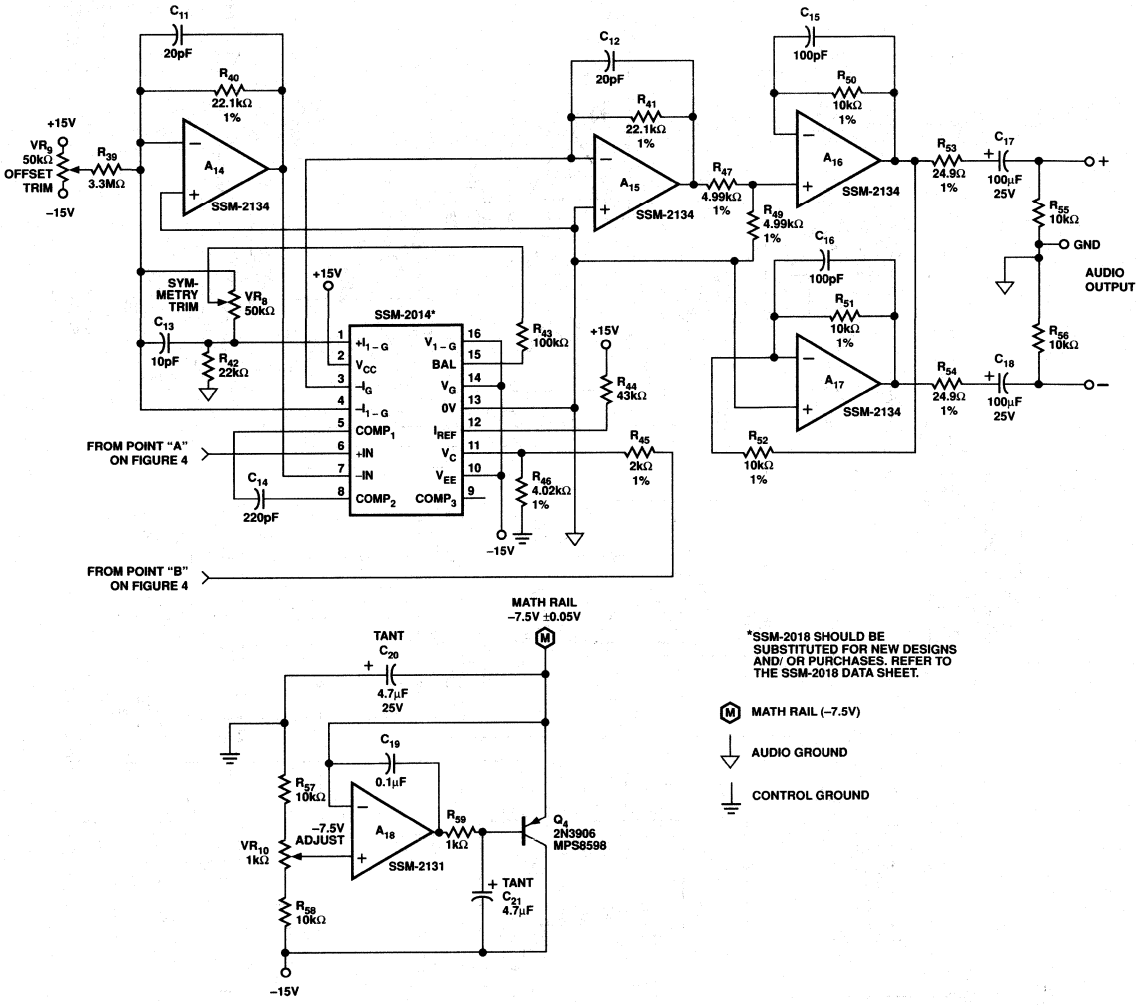

FIGURE 5: *VCA, Math Rail Regulator and Line Driver Schematic*

THE VCA SECTION

Figure 5 shows the VCA section of the compressor using the SSM-2014. The device is configured as an "outboard OVCE" in accordance with the literature supplied with the IC. Amplifiers A_{14}, A_{15}, A_{16}, and A_{17} should be high quality, low noise op amps such as the SSM-2134. A_{14} and A_{15} provide the necessary feedback and output buffers for this particular circuit. Resistive voltage divider R_{45}/R_{46} reduces the OVCE control port sensitivity to +1V/20dB attenuation to match the scale factor of the control circuitry. Variable resistor VR_8 is adjusted for lowe distortion products, preferrably at unity gain with a high lev input. VR_9 is adjusted by applying a low frequency (<50Hz) sign at about $2V_{p-p}$ to pin 11, and setting for least low frequen output at A_{15} under a no-input signal or shorted-input conditio The stability of both these trims will be a welcomed surprise the designer used to dealing with log/antilog VCAs.

A "quasi-balanced" line driver, consisting of A_{16}, A_{17}, and the associated passive components, completes the compress

circuitry. The unit is capable of driving a 600Ω load at a sustained output level of +21dBm, and has a maximum output of +26dBm.

The SSM-2014 provides the designer a degree of flexibility in configuration which is not easily available using other VCA topologies. Chief among its attributes is the ability to select the operating bias current by applying current to a single port on the chip. This allows the user to select an operating point which is optimized for best noise performance vs. distortion for a given application.

In considering the normal operation of the compressor, the most common scenario is that the unit is used to "track" instruments or vocals. Of secondary, but no less important, concern is use for compressing mixed program material to enhance apparent loudness. In both applications, the trade-off between the residual noise floor and distortion at high-signal levels, both a function of operating bias, is somewhat arbitrary. Since it is likely that the dynamic distortion inherent in the compression process in normal operation would be at least equally as noticeable as moderate distortion at high-signal levels, the "intermediate" bias setting as described in the literature accompanying the device was chosen. This places the device bias at approximately 300μA, with a value of 43kΩ for R_{44}. As a result, the noise floor for the VCA circuit is –84dB (ref. $0.775V_{RMS}$) in a 20 kHz bandwidth at 0dB gain. The 1kHz THD+Noise measurements yielded figures consistent with the published data, and SMPTE IMD measurements disclosed worst-case distortion products in the 0.2% range, which is acceptable in all but the most critical applications. Should the designer wish to implement the sliding bias scheme, as described SSM-2014 data sheet, the output of the absolute value at pin 1 of the SSM-2110 (see Figure 4), or the rms computing loop (pin 5) may be used to drive a comparator with the appropriate time constants in order to switch to class A operation in the presence of high level inputs. In practice, this makes little difference in the transparency of the compressor in normal operation. Listening tests of the compressor demonstrated the smooth, precise control expected of the feedforward circuit topology. As the attack time (integration time) control is advanced from fast to slow settings, the low frequency content in mixed material becomes more solid and better defined, but the tendency to "squash" the lows is relatively absent at faster attack times as compared to other compressors having adjustable attack times with comparable settings.

The log averaging detector scheme really shines on vocals and horns, bringing a soloist "up-front" with moderate attack times. This is a noticeable difference when compared to any RMS-type compressor used for comparison in the listening tests.

ADJUSTING FOR BEST PERFORMANCE

As in any compressor or limiter whose ratio must be trimmed in its initial setup (see Figure 4, VR_5), the unit is sensitive to incorrect adjustment. One of the most distressing sounds which can be produced by a compressor is "over-compression," in which the control circuitry causes too much gain reduction at high ratios. For this reason, the compressor ratio trim should be set with the **Threshold** control at 0dB (0V at the wiper of VR_4), and with an input of +20dB, the trim should be adjusted so that with the **Attack**, **Release**, and **Output** gain controls centered, the maximum **Ratio** setting, fully clockwise, produces an output level equal to or slightly greater than the rotation point.

When laying out circuitry using the SSM-2014 and SSM-2110, care should be taken to keep traces to virtual grounds as short as possible, and a single point audio ground should be used. The control ground should connect to the audio ground at one point, pin 4 of the SSM-2110, and supply traces should be heavily decoupled with high quality capacitors. Traces carrying audio signal should be kept well away from control circuitry, and the detector IC and VCA should be located away from heat sources such as regulators or power supply transformers.

Since all parameter control is derived from DC levels produced by the front panel controls, high quality potentiometers need not be used. All the front panel control scales may be marked in equal intervals, and follow the antilog law, i.e., equal spacing per dB of gain or threshold setting, equal spacing per decade of attack and release times, etc. The sole exception is the ratio control, whose scale is skewed so that 2:1 appears near the middle of the control, as one normally would expect of a traditional feedback compressor.

The compressor control circuit described in this application note was configured using only four quad op amps in addition to the SSM-2110 and three matched transistor arrays. By providing the basic building blocks for an audio dynamic range processor in monolithic form, the SSM audio chip set greatly simplifies the implementation of an otherwise complex processor.

MEASURED PERFORMANCE:

SMPTE IMD @ Unity Gain, 0dBv in	0.009%
SMPTE IMD @ Unity Gain, +20dBv in	0.11%
SMPTE IMD @ 20dB Gain Reduction, 0dBv in	0.06%
SMPTE IMD @ 20dB Gain Reduction, +20dBv in	0.025%
Residual Noise and Hum @ Unity Gain, 20kHz BW	–84dBv
Maximum Output Level into 600Ω, Balanced	+21dBm
Maximum Input Level Before Clipping	+21dBv
Usable Dynamic Range, Unweighted in 20kHz BW	103dB
Threshold Range Ref. Rotation Point	–40 to +20dB
Useful Range of Rotation on Point Adj.	–10 to +4dBv
Nominal Attack (Integration) Time Range	0.02 to 200ms
Nominal Range of Ratio Adjustment	1.3:1 to 20:1
Range of Release Time Adjustment	0.05 to 5s/20dB
Range of Output Gain Adjustment	–20 to +20dB

NOTE: 0dBv refers to 0.775 V_{RMS}

ANALOG DEVICES

AN-136
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

An Ultra Low Noise Preamplifier

by M. Jachowski

Achieving the maximum usable dynamic range from low output level transducers such as audio microphones, magnetic pick-ups, or low impedance strain gauges requires a preamplifier with very low input-referred voltage noise. The circuit shown in Figure 1 has extremely low noise, 0.5nV/$\sqrt{Hz}$, and can provide a gain of 1000 over a 200kHz bandwidth.

This amplifier's low noise characteristics are attributable to the SSM-2220's matched PNP transistor pair. Operating with 2mA collector current in each transistor, the SSM-2220 forms a differential input stage with a DC gain of 385, approximately 50μV of offset voltage, and only 0.5nV/$\sqrt{Hz}$ of broadband noise. When multiplied by the stage gain of 385, the input noise of the SSM-2220 appears as 192.5nV/$\sqrt{Hz}$ differentially at the inputs of the OP-27. This makes the 3.8nV/$\sqrt{Hz}$ of the op amp an insignifi-cant contribution to the overall noise of the circuit. In this example, the input stage compensation, C_1 and R_7, optimizes noise performance over the audio frequency range by allowing the differential pair to have a flat frequency response to 20kHz before being rolled-off for stability criteria. Input stage gain is reduced 20dB from 20kHz to 200kHz and then remains constant until the SSM-2220's gain-bandwidth limit is reached at about 8MHz. This compensation ensures the preamplifier's stability for gains from 100 to over 2000. Gain is set with resistors R_5 and R_6 where $A_{VCL} = 1 + R_5/R_6$. To limit the thermal noise contributed by the feedback loop impedance, R_6 should be no more than 10Ω (a 10Ω resistor creates about 0.4nV/$\sqrt{Hz}$ at +25°C).

The input stage current, 4mA, is established by the current source of Q_2, R_1, and a GaAsP LED. The LED is used as a 1.6V "zener" whose temperature coefficient is nearly identical to that of Q_2's base-emitter junction. This produces a temperature stable 1V drop across R_1 forcing 4mA to flow from Q_2's collector. The 4mA splits to 2mA in each side of the differential pair. With $h_{fe} = 150$ in the SSM-2220, input bias current will be about 13μA. Because the bias current is relatively large, the offset voltage created as it flows through unbalanced source impedances will quickly surpass the differential pair's offset, making necessary the offset trim, R_8. Low source impedances will reduce the offset drift as h_{fe} changes over temperature.

A low source impedance is also critical to maintain a low overall input noise. The 0.5nV/ $\sqrt{Hz}$ noise of the SSM-2220 input is equivalent to the thermal noise of a 15Ω resistor at +25°C.

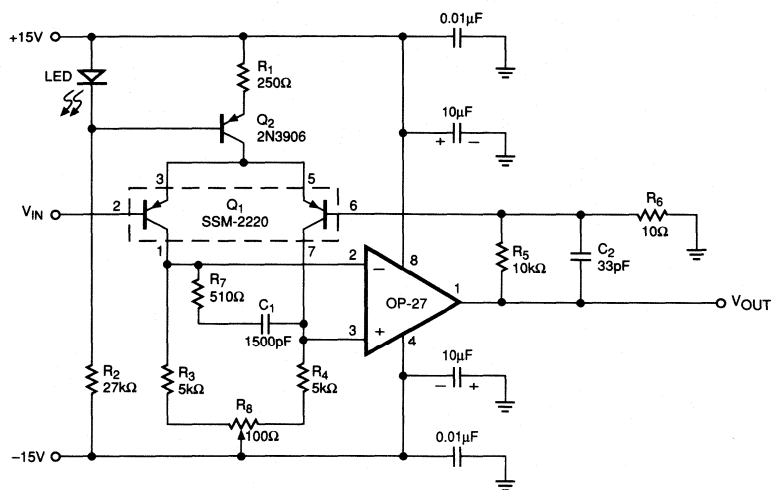

FIGURE 1: *This ultra low noise preamplifier shines new light on high-gain, low noise applications such as microphones, thermocouples, strain gauges, and magnetic pick-ups.*

Therefore, any transducer with a sourcing impedance greater than 15Ω will produce a noise which dominates that of the preamplifier. Figure 2 shows the total output noise of the preamplifier driven through a 10Ω source impedance. The analyzer displays total RMS noise voltage measured in a 0.03Hz bandwidth. The average broadband measurement is roughly 0.13μV on the vertical scale. Divided by the amplifier's closed-loop gain of 1000, this corresponds to 0.13nV at the preamp input, or expressed in nV/ Hz,

$$e_n = \frac{0.13nV}{\sqrt{0.03Hz}} = 0.75nV/\sqrt{Hz}$$

Taking into account the noise of two 10Ω source resistors, the noise attributable to the SSM-2220 is then,

$$0.75nV/\sqrt{Hz} = \sqrt{(e_{SSM})^2 + (0.4nV/\sqrt{Hz})^2 + (0.4nV/\sqrt{Hz})^2}$$

$$e_{SSM} = 0.49nV/\sqrt{Hz}$$

The 1/f noise corner frequency is also remarkably low, only about 0.25Hz. In the 20kHz audio bandwidth, the total RMS input-referred noise voltage contributed by the SSM-2220 differential pair is,

$$e_n = (0.5nV/\sqrt{Hz})\ (\sqrt{20kHz - 20Hz}) = 70.5nV_{RMS}$$

The thermal noise of a 10Ω source impedance in the same bandwidth is,

$$e_t = 1.28 \times 10^{-10}\sqrt{(10\Omega)\ (20kHz - 20Hz)} = 57nV_{RMS}$$

The total input referred noise of the preamplifier with 10Ω source impedances on each input is,

$$e_{total} = \sqrt{(70.5nV)^2 + (57nV)^2 + (57nV)^2} + 106nV_{RMS}$$

This is lower than the thermal noise of a single 50Ω resistor over the same bandwidth, $126nV_{RMS}$.

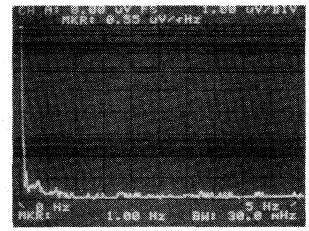

FIGURE 2: *The spectrum analyzer shows that, in a gain of 100(with 10Ω source impedances, the SSM-2220 preamplifier has less than 0.5nV √Hz broadband noise and a 1/f noise corne of about 0.25Hz. Total harmonic distortion is less than 0.005% of a 10V_{p-p} signal from 20Hz to 20kHz.*

ANALOG DEVICES

AN-211
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

The Alexander Current-Feedback Audio Power Amplifier*

by Mark Alexander

This application note was written by Mark Alexander, who received his BSEE from the University of Toronto in 1981. As a consultant to Analog Devices, Mark describes a unique power amplifier topology that is the result of his long standing interest in audio power amplifier design and critical listening of audio systems.

The current-feedback approach presented here meets the traditional audio requirements of power amplifiers, but also adds the additional benefit of very high speed and bandwidth (200 V/μs slew rate, 1 MHz bandwidth) that results in excellent dynamic performance, and hence, sound quality.

INTRODUCTION

The subject of power amplifier design is one of those controversial areas of audio engineering that continues to receive intense debate, despite the fact that there are literally dozens of papers available to guide the designer. Many different topologies have evolved from the relatively modest beginnings of solid state power amplifier design in the late 1950s and early 1960s, and this has lead to a few very unique and original designs. A substantial number of transistorized amplifiers that were built during these early years were little more than redesigns of vacuum tube circuits with lower voltage supply rails, and often had performance levels that left a great deal to be desired. Quite a few of them sounded significantly worse than their thermionic predecessors. The "real revolution" in audio power amplifier design actually occurred during the 1970s and introduced such new innovations as direct coupling, fully complementary design, pseudo Class A biasing, and current dumping; not to mention the discovery of the importance of dynamic intermodulation distortion testing and its relationship to slew rate. Unfortunately, the plethora of so called "new" amplifier designs that have proliferated since this period are often variations of older circuits that originated during the 1970s, and generally feature only slight modifications to the input, output or gain stages.

Some designers have demonstrated rail-commutated output stages, which allow them to improve the operating efficiency of a big amplifier to such an extent that the huge amount of output transistor heat sinking usually necessary is reduced to that of a much lower power design. These can suffer from "switch over" distortion caused by the output stage switching between different supply rails, and can be quite objectionable. Certainly, high output power should not be obtained at the expense of inferior operating specifications, but this is indeed the case with certain types of amplifiers. Some clever design techniques do achieve quite impressive performance, however, albeit at the expense of greatly increased circuit complexity. Still other amplifiers dispense completely with the familiar principles of negative feedback, and their creators claim that their circuitry provides a sound more "open and lifelike," even though the distortion performance is usually poor. On the whole, though, most audio power amplifiers are essentially discrete copies of monolithic voltage feedback op amps, such as the 4136, but are invariably simplified to reduce the transistor count.

The purpose of this technical note is to introduce the audio designer to a truly new power amplifier topology, not an adaptation of an existing design, that offers exceptional performance on a par with the best of the available solid state designs (voltage feedback or otherwise). This new topology completely dispenses with the principles of global voltage feedback, so commonly used in most amplifiers, in favor of a design based instead on the principles of current feedback. In addition, this note addresses many of the important practical aspects of successfully getting a design off the ground, aside from choosing the basic core amplifier design. In almost all cases, having a good basic amplifier topology is not enough to guarantee that the final piece of equipment will perform to the original design expectations. Consequently, additional topics such as board layout,

***Patent pending. The amplifier described herein is for informational purposes only with restricted use and no licenses implied. Readers are permitted to construct one stereo amplifier for their own personal, noncommercial use. For other uses, contact Analog Devices for licensing details.**

component selection, paralleling output devices, placement of high current wiring, and thermal design are considered as well.

A LITTLE BACKGROUND ON FEEDBACK

Before dissecting the new audio amplifier circuit in detail, some background on the differences in operational characteristics between voltage feedback and current feedback amplifiers is appropriate. Since it is likely that the reader may not have been previously exposed to the latter, an overview of voltage feedback followed by a look at the advantages of current feedback is necessary. This discussion will allow one to understand why circuits that make use of this relatively new topology are so important. Because the bandwidth of an audio amplifier is usually one of the most important specifications, a relatively simple equation for the upper −3 dB point is essential. Simplifying the current feedback amplifier and its attendant feedback network into a representative circuit model for nodal analysis provides the key to arriving at a compact, but reasonably accurate, expression for the frequency response. Appendix A has complete details of the circuit analysis.

The theoretical analysis of a voltage feedback circuit that often accompanies its frequent criticism has been well described in other works,[1] and thus will not be reiterated here. Since the original impetus behind the development of this new power amplifier design was a general dissatisfaction with the performance achievable by voltage feedback circuits, some discussion of their disadvantages is worthwhile. This will serve to set the stage for the in-depth discussion of current feedback amplifier analysis, in Appendix A. Although the analysis section can be skipped without disrupting the continuity of this note, the reader is encouraged to review it.

Constant gain bandwidth characteristics, resulting from the application of voltage feedback, present a problem if one requires reasonably high gain while simultaneously achieving wide closed-loop bandwidth. Some very high voltage power amplifiers may require gains as high as 50, for example, plus a bandwidth of several hundred kilohertz which obviously means that a gain bandwidth product in the range of 10 MHz to 20 MHz is needed. This is not easy to achieve, especially in a high voltage design. An additional problem with voltage feedback amplifiers is that their slew rate is usually limited by the transconductance stage which has a finite maximum output current, normally equal to the tail current of the differential input transistor pair, available to charge the compensation capacitor. High slew rate is very desirable in a large-signal audio power amplifier and mandates the use of large input-stage tail currents and small compensation capacitor values. Unfortunately, in the interests of amplifier stability, reducing the value of the compensation capacitor requires some degeneration of the input stage (to reduce its transconductance) which thus reduces the open-loop gain. This action reduces the loop gain available in the audio band and causes an increase in THD products, since it is the loop gain that serves to reduce the open-loop amplifier distortion, most of which originates in the highly nonlinear output stage. What all this boils down to is the fact that a difficult trade-off has to be made between stability, open-loop gain, and slew rate without compromising the overall ac performance and transient response. Clearly, a global voltage feedback scheme may not necessarily be the optimum choice for ultrahigh performance audio power amplifiers, and in some cases it will not even be possible to meet all the design goals using this topology.

Current feedback operational amplifiers were originally introduced because they overcame the bandwidth variation, inversely proportional to closed-loop gain, exhibited by voltage feedback amplifiers. They still show a slight variation of bandwidth, however, as the gain is increased above unity, but it is much less significant than with the latter. In fact, current feedback amplifiers don't begin to behave like voltage feedback amplifiers until the closed-loop gain is made quite large (~50). The simplified model of a current feedback amplifier in Figure 1 shows that it uses a unity gain input buffer whose output current is fed, via a bidirectional current mirror, into a transimpedance gain stage. The voltage generated here is then buffered and fed to the output terminal. Typical values for R_T are quite high, usually several hundred kilohms or even a few megohms. R_{INV} is the output resistance of the input buffer, and feedback resistors R_1 and R_2 set the input-to-output voltage gain in a fashion somewhat similar to that of a conventional op amp. Here, however, it is an error current I_1 that sustains the output voltage and not an error voltage.

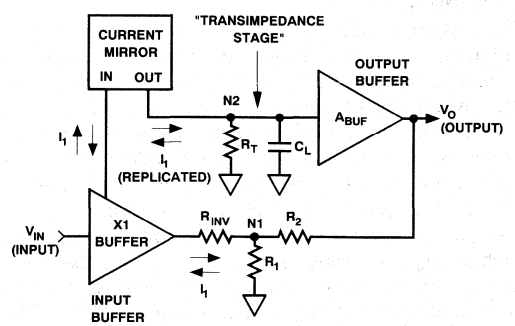

Figure 1. The Model of a Current Feedback Amplifier Shows that an Error Current, I_1, Determines the Overall Output Voltage.

The concept of a finite gain bandwidth product can also be applied to a current feedback amplifier as a measure of its performance, although it is only meaningful at high gains. Arguably, the most important attribute of this topology is that the amount of current available to charge the compensation capacitor during output slewing is proportional to the difference between the actual and final output voltages, just like a simple RC circuit. As such, there is theoretically no slew rate limit with this topology, which makes it very attractive for an audio power amplifier. Practical circuit limitations inevitably

impose a restriction on the maximum current level that can be handled in the gain stage of a current feedback amplifier, however, and it is this limiting that gives rise to a finite slew rate. Still, the slew rates achievable with these types of circuits are often higher by as much as a factor of 5 (or greater) than their voltage feedback counterparts, for a given quiescent supply current. Current feedback represents a much more logical choice for a power amplifier than voltage feedback, and this will be demonstrated.

POWER AMPLIFIER CIRCUIT TOPOLOGY

Prior to looking at the actual amplifier circuit, the simplified block diagram of Figure 2 will be considered to help understand how the overall design works on a system level. This will make the final amplifier circuit easier to follow. As may be gathered from Figure 2, this is a rather unconventional design, in which there are two op amp input stages feeding a single gain stage and power output buffer. By considering this design one block at a time, however, it is becomes easier to grasp the way in which each of the major sections interacts with one another.

The Input Stage

The input buffer used in this power amplifier is simply a conventional voltage feedback op amp chosen for its excellent audio characteristics, and reasonably high output current capability. This ensures that the limiting factor in terms of overall amplifier performance will be the current feedback gain block and not the input stage. The output current from input amplifier A_1 is taken from its power supply pins and fed to the emitters of a pair of common base cascode transistors that provide regulated dc voltages for the op amps. At first glance this might appear to be a very strange connection, because the power supply pins of A_1 are used as outputs and its output is used as an input. However, this is in accordance with the model shown in Figure 1 since the output current from the input buffer must be fed, via the bidirectional current mirror, into the transimpedance gain stage. It is here that the high output voltage is ultimately generated, prior to buffering by the unity gain output stage. The half-wave rectification action of A_1's output current, due to its class AB output stage, causes the two current mirrors to receive complementary input currents. When A_1 is sourcing output current, it causes a corresponding increase in the current of the upper mirror and a decrease in that of the lower mirror. This forces the voltage at the output of the transimpedance stage to swing positive. For cases where A_1 is sinking current, exactly the opposite is true. A current mode gain stage arrangement such as this is fully complementary and truly push-pull, which means it should exhibit low even-order distortion. Note that the quiescent supply current of A_1 conveniently serves to bias the two current mirrors that sit referenced to each power supply rail, thus providing an appropriate operating point for the transimpedance stage and bias voltage generator.

In most commercially available current feedback amplifiers, the input buffer stage has a gain of unity and is generally of an open-loop design. Here, an op amp is being used as the input stage instead and thus can be configured to provide some gain. This is extremely easy to do since it only involves tapping the shunt resistor to ground at the output of A_1. The overall amplifier midband gain is therefore:

$$A_V = \left(1 + \frac{R_7}{R_6}\right)\left(1 + \frac{R_8}{R_6 + R_7}\right) \qquad (1)$$

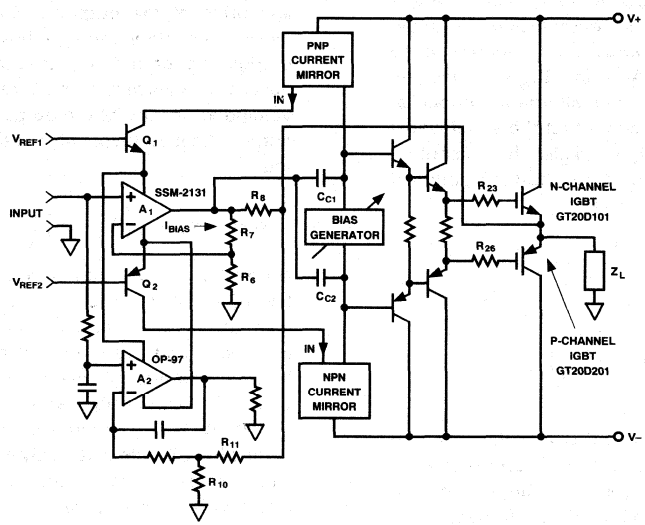

Figure 2. A Simplified Block Diagram of the Amplifier Shows that the Input Amplifiers, A_1 and A_2, Feed a Common Gain Stage and Output Buffer.

The Gain Stage and Frequency Compensation

The outputs of the two current mirrors that are connected to each supply rail feed an adjustable voltage bias generator which provides the necessary bias for class AB operation of the complementary MOS-IGBT (Metal-Oxide-Semiconductor Insulated-Gate-Bipolar-Transistor) output stage. The bias generator is designed to have very low output impedance over the operating frequency range of the amplifier. Compensation is provided by C_{C1} and C_{C2}; two capacitors are used instead of one to keep the structure of the gain stage symmetrical. Unlike the simplified current feedback model shown in Figure 1, this design has the compensation capacitors returned to the feedback summing node instead of ground. This alternate connection has a very beneficial effect on the amplifier step response when it is loaded by a fairly low value impedance such as a loudspeaker.

An IGBT emitter follower output stage, such as the one used in this amplifier, has a transfer function that contains two poles and a real zero plus the usual dc gain term of slightly less than unity. When the amplifier drives high values of load impedance, such as the feedback resistors alone, the two output stage poles are quite high in frequency (usually above 20 MHz), and contribute little excess phase shift within the amplifier's passband. Quite a different situation arises when a load is connected to the output of the amplifier. The two poles in the output stage now split apart, and the dominant one becomes sufficiently low in frequency that it contributes excess phase shift at lower frequencies within the amplifiers' passband. This can cause a considerable problem if the compensation scheme in Figure 1 is used since it may result in undesirable ringing on the edges of a square wave. The compensation scheme of Figure 2 overcomes this problem by inserting a high frequency closed-loop zero that tends to make the amplifier more stable. Also, this compensation arrangement allows the use of smaller capacitors than with the original scheme. Appendix A shows the complete response of the amplifier when this alternate compensation scheme is used. If we assume that the small signal transresistance, R_T, is quite high and also that the output buffer gain is near unity, then the closed-loop pole and zero will occur at frequencies given by:

$$f_{POLE} \cong \frac{1}{2\pi\left(2R_8 + \frac{R_8}{R_6 + R_7} R_{INV}\right) C_{C12}} \quad (2)$$

and

$$f_{ZERO} = \frac{\left(1 + \frac{R_8}{R_6 + R_7}\right)}{4\pi R_8 C_{C12}} \quad (3)$$

where C_{C12} is the sum of C_{C1} and C_{C2}. Notice that the frequency at which the zero occurs is approximately equal to the closed-loop bandwidth multiplied by the gain of the current feedback loop, if R_{INV} is fairly small in value. These equations, plus Equation (1), are the necessary design formulas needed to determine the gain and small signal bandwidth of this amplifier. Later on it will be demonstrated that the mathematical theory and actual measurements made on the circuit do indeed correlate very well with each other.

Driver and Output Stages

This part of the power amplifier design is quite conventional, relatively speaking, and no attempt was made to use error correction or pseudo class A biasing schemes to lower the output stage crossover distortion. Since the primary design objective for this amplifier was wide bandwidth and high slew rate, it was felt that any additional circuitry following the transimpedance gain stage might degrade the closed-loop stability. Besides, low crossover distortion can be achieved by running the output transistors at a sufficiently (but not excessively) high idling current. A simple double emitter follower driver stage, therefore, was chosen to buffer the voltage generated by the gain stage and feed it to the gates of the power IGBTs. This driver stage is capable of providing several hundred milliamps of charging current for the IGBT gate capacitances while the output is slewing, and is mandatory in a high speed design such as this.

DC Control Amplifier

The purpose of this additional input stage is to provide an accurate, low drift, dc gain path to the main output that is independent of the ac gain path and its poor dc characteristics. In the original version of this amplifier, expensive precision matched NPN and PNP dual transistors were used in the two current mirrors, but no dc control amplifier was used. It was incorrectly assumed that precise matching of the transistors in each mirror would result in very low output offset voltage, as long as the input buffer had reasonably low input offset voltage as well. As it happens, this is not the case with a current feedback amplifier. Any mismatch between the two current mirrors results in a finite amount of bias current appearing at the output terminal of the input buffer, which must flow through feedback resistor R_8 to the output. It cannot flow through R_6 and R_7 to ground, because the current in these resistors is set only by the voltage appearing at the output of the input buffer. The output offset voltage, without the dc control amplifier is thus:

$$V_{OOS} = V_{IOS\,(A1)}\left(1 + \frac{R_7}{R_6}\right)\left(1 + \frac{R_8}{R_6 + R_7}\right) + I_{BIAS} R_8 \quad (4)$$

Normally, V_{IOS} (A_1) can be made quite small by using a low offset op amp. Unfortunately, the output terminal bias current, I_{BIAS}, can be as large as 100 μA under static conditions and even larger if a thermal gradient exists between the two mirrors on the power amplifier driver board. This can easily lead to an output offset in excess of 100 mV, which changes as the amplifier warms up. A large offset like this is likely to cause an audible click when the relay that connects the loudspeakers to the amplifier is energized, and is generally undesirable.

The solution to these problems is a low frequency servo-loop that controls the dc output voltage, independently of any low frequency current or voltage fluctuations in

he main current feedback gain path. This is facilitated
y the use of a second low power precision op amp, A_2,
hat is configured as an integrator with very low cross-
ver frequency (less than 5 Hz). The low crossover fre-
uency ensures that the integrator will not have any
ffect on the performance of the overall amplifier in the
udio band. Voltage feedback is applied from the main
utput back to the input of the integrator through resis-
ors R_{10} and R_{11}, which set the closed-loop dc gain. This
ain is made equal to that given by equation (1). Since
A_2 drives a resistor connected to ground, as shown in
igure 2, it behaves as an operational transconductance
mplifier with the output current taken from its power
upply terminals. This compensating output current is
hen fed to the two common-base regulator transistors
here it is summed with the signal current from the
ower supply terminals of A_1. The output current of A_2 is
hus forced to cancel I_{BIAS} almost exactly because the dc
ain of the integrator, coupled with the additional gain
roduced by the transimpedance stage, is very high.
onsequently, the integrating control loop completely
verrides the current feedback loop at dc and the output
ffset is reduced from that given by Equation (4) to:

$$V_{OOS} = V_{IOS\,(A2)}\left(1 + \frac{R_{11}}{R_{10}}\right) \qquad (5)$$

his means that it can be made arbitrarily small through
he choice of a low offset amplifier for A_2. Here the cost
f an additional op amp is more than offset by not
aving to use expensive matched NPN and PNP dual
ransistors in the current mirrors.

MPLIFIER CIRCUIT DESIGN

he complete circuit diagram for one channel of the
mplifier is shown in Figure 3, and an accompanying
arts list is included in Appendix B. This design utilizes
IC op amps, 17 bipolar transistors in the gain and
river stages, and at least 2 complementary IGBT power
ransistors from Toshiba in the output stage. These re-
ently introduced devices are essentially similar to
ower MOSFETs in that they have a very high imped-
nce input terminal (the gate) and square-law transfer
haracteristics, but are manufactured using a slightly
odified double diffused MOS process. Unlike power
OSFETs, however, they feature consistently higher
urrent handling capability for N- and P-channel transis-
rs of a given die size. This allows one to get by with
smaller die size IGBT output stage than one using
OSFETs, thus providing a fairly substantial cost sav-
gs (especially on the P-channel transistors). The driver
tage in this amplifier can easily accommodate multiple
airs of power devices in the output stage, because of
s high peak current drive capability, but just a single
air of 250 V, 20 A IGBTs was used in the version that
as characterized here. Power supply voltages for the
river board and output stage may range from ±20 V to
75 V. Most of the components that mount on the com-
act driver board, the layout of which is shown in Figure
, are quite readily available and inexpensive.

An input filter with a cutoff frequency of approximately 2 MHz precedes the input stage. It was included to reduce the potential for RF interference problems, and to eliminate the possibility of the amplifier oscillating on power-up with the input left floating (something that was noticed during the original development of this topology). The filter is formed by the 100 Ω input resistor and 750 pF shunt capacitor. A 100 kΩ resistor is connected to ground at the input of A_1, and provides the necessary dc bias current path to ground if the input is inadvertently left open. The overall amplifier gain is set by R_6, R_7, and R_8, and substituting the values of these resistors into equation (1) yields a figure of 24.087 or 27.64 dB. If more gain from the circuit is desired, the values of R_6 and R_7 should be changed, but their sum should be kept approximately equal to 50 Ω so that the gain of the current feedback section stays constant (at about a factor of 16). By simply swapping the 16.5 Ω and 33.2 Ω resistors, for example, the gain of the input stage becomes approximately equal to 3, and the gain of the overall amplifier increases to a factor of 48.47 or 33.7 dB. In fact, the gain of the input stage can be made as large as 20 dB before its bandwidth drops below that of the rest of the amplifier.

The references for the two common-base regulator transistors (Q_1 and Q_2), which provide stable supply voltages for the op amps, are actually two pairs of standard NPN bipolar transistors (2N3904s) used as Zener diodes (Q_{14} through Q_{17}). They are connected in series (with their collector leads clipped off) to obtain a net breakdown voltage of around 15 V for the pair . There really is a good reason for using such an arrangement since it would obviously be easier to use a 15 V "Zener" diode, as opposed to this seemingly more complicated approach. In reality, the connection of two bipolar transistors in this manner exhibits significantly less low frequency noise than the 15 V "avalanche" diodes, as they are more appropriately called, and is actually more cost effective. The composite Zeners are bypassed with 10 μF 25 V tantalum capacitors, used mainly for reasons of economy and size, which filter out residual noise from the diodes as well as the power supply rails. Two resistors marked R_{BIAS} on the circuit diagram (R_1 and R_2), which are connected to each supply, serve to bias Zener connected transistors Q_{14} through Q_{17} and should be chosen such that with nominal power supply operating voltages (anywhere from 50 to 70 volts) about 1 mA of current will flow through them.

The two Wilson current mirrors connected to each rail, and fed from the collectors of Q_1 and Q_2, are formed from a low voltage transistor, a diode and a high voltage transistor (2N5551 or 2N5401). They are degenerated somewhat with 100 Ω 1% resistors to improve matching. Anti-saturation diodes (D_2 through D_5) have been included to prevent storage time problems with the cascode transistors (Q_4 and Q_6) in either of the two mirrors during clipping, and this results in extremely rapid recovery from overdrive. It should be noted that the onset of clipping in the transimpedance stage will occur at

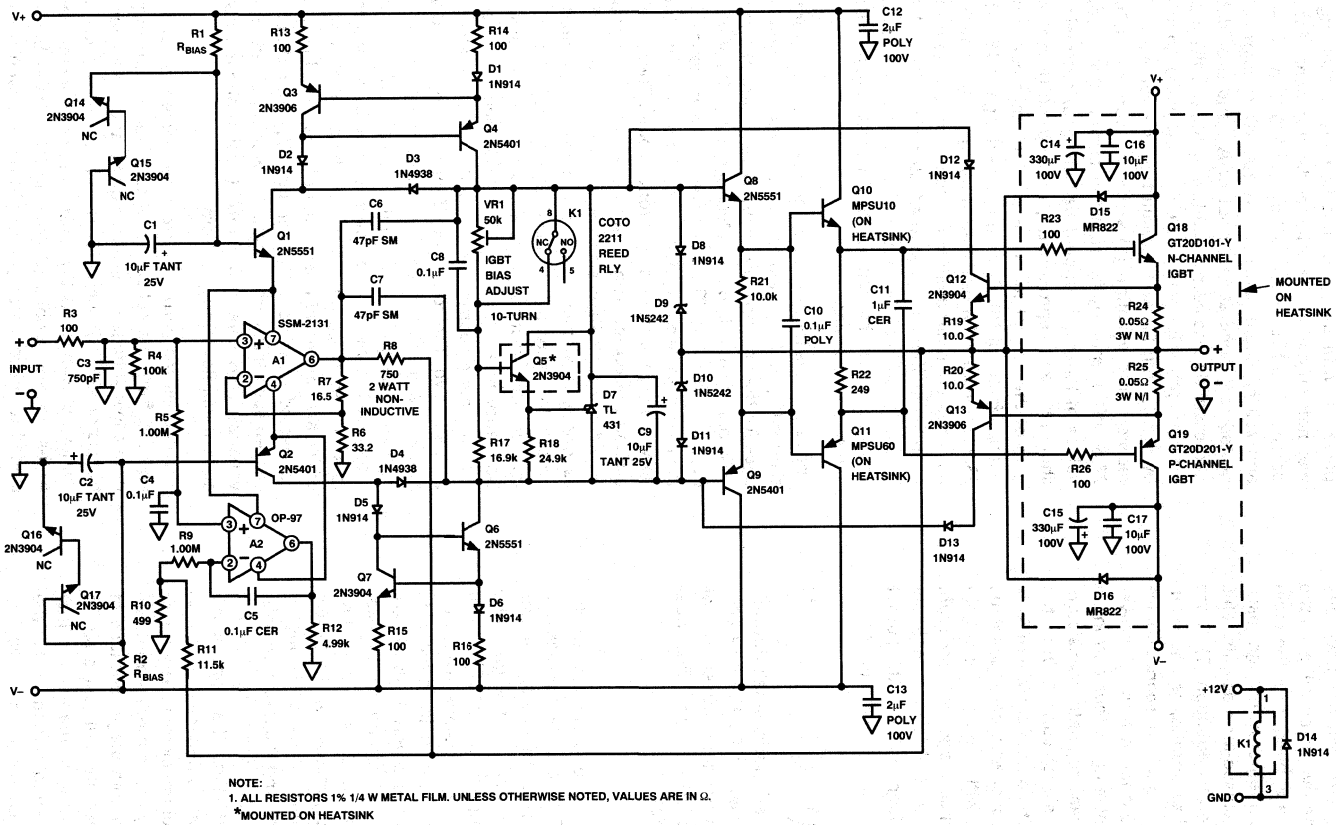

NOTE:
1. ALL RESISTORS 1% 1/4 W METAL FILM. UNLESS OTHERWISE NOTED, VALUES ARE IN Ω.
*MOUNTED ON HEATSINK

Figure 3. The Complete Amplifier Circuit Diagram Shows that Inexpensive Small-Signal Silicon Is Used Throughout to Minimize Overall Cost.

about 2 V from either power supply rail for very small overdrive conditions, but hard clipping in this stage will actually be dependent on the current limit of the input amplifier A_1. This occurs because, during hard clipping, the current summing network connected to the output of A_1 is no longer balanced and significant current can flow in its output stage. Consequently, the current in the mirrors increases very rapidly up to the value of A_1's maximum output capability (usually 30 mA to 40 mA), causing a corresponding voltage drop across the aforementioned 100 Ω resistors. The effect of this excessive current in the mirrors is such that it causes the clipped signal to "pull in" slightly from the rails, as the amplifier is driven harder and harder into its overload region. It is very important not to let the circuit stay in this condition for any significant period of time, since the power dissipation in Q_1 and Q_2 will increase far beyond their nominally rated value of a few hundred milliwatts. Peak dissipation in these transistors can reach as much as 1.5 W to 2 W, with typical rail voltages of 50 to 70 volts; therefore, very large dc input signals or low frequency square waves should be avoided. If these operating conditions are anticipated, however, clip-on heat sinks for Q_1 and Q_2 are mandatory.

Frequency compensation in this particular design is provided by two 47 pF compensation capacitors that are connected to the feedback summing node (C_6 and C_7), as mentioned previously. This results in a total value of 94 pF. The reason such a large value of capacitance was chosen is quite simple: it completely swamps out any nonlinear voltage dependent capacitances that are present at the high impedance gain node, resulting in constant amplifier bandwidth as the supply voltage is varied. Concerns about too low a slew rate, with such large compensation capacitors, are usually justified in a voltage feedback amplifiers, but here there is as much as 30 mA of current available to charge them and slew rate limiting will not normally be encountered.

A calculation of the expected frequency response of the amplifier is now in order, and can be accomplished quite easily by substituting the value of 94 pF for C_{C12}, and the values of 750 Ω for R_8, 16.5 Ω for R_7, and 33.2 Ω for R_6 into Equation (2). The value for R_{INV} is a little more difficult to determine since we must know *a priori* what the value of the closed-loop output resistance of A_1 is, at the overall −3 dB point of the amplifier. The solution to this problem actually involves a little bit of circular reasoning, but the motive behind it is rather easy to see. If Equation (2) is evaluated initially without considering the effect of R_{INV}, a closed-loop bandwidth of 1.12 MHz is calculated. Since the effect of a finite R_{INV} is to lower the bandwidth somewhat, a prediction of the final amplifier closed-loop bandwidth will allow an initial guess for this parameter to be made. In this case a prediction of a final closed-loop bandwidth of 1 MHz is made. If we now take the open-loop output resistance of A_1 from its data sheet (about 70 Ω) and divide it by one plus the value of its loop gain at the predicted −3 dB point of 1 MHz (about 7.68), a value of 9.11 Ω is obtained. When this estimate for R_{INV} is included in Equation (2), an overall closed-loop bandwidth of 1.034 MHz is the final result. This is really very close to the original guess of 1 MHz, and it seems that no further iteration will be necessary to get closer to an acceptable answer. It should now be plainly apparent that extraordinarily wide closed-loop bandwidth seems rather easy to come by in a current feedback power amplifier, even when the compensation capacitors are quite large. For this reason, careful board layout and wiring techniques are of tantamount importance in actually getting a design such as this to work properly without oscillating.

The output stage bias voltage generator, connected between the collectors of Q_4 and Q_6, is formed from a programmable shunt regulator (D_7), with an NPN emitter-follower buffer (Q_5) driving its control input. This buffer is not normally required in most applications because the control input bias current of D_7 (a TL431) is only a few microamps, but it is included here for thermal compensation of the output stage idling current. A common problem with biasing output stages that use vertical DMOS devices (MOSFETs and IGBTs) is that at moderately low current levels, the decrease in V_{TH} of approximately 3 mV/°C causes the collector current to increase for a fixed gate-to-emitter bias voltage. If transistor Q_5 is securely mounted on the same heat sink as the power IGBT output stage, its V_{BE} will decrease as the output transistors heat up. This decrease in V_{BE} of about 2 mV/°C, which is multiplied up in the bias generator by approximately a factor of three, thus helps to stabilize the quiescent current in the IGBT output stage. A form-C relay can also be included across the 50 kΩ bias adjustment pot (VR_1), as shown, to allow the amplifier to be powered up with zero bias voltage on the output transistors. This feature, when used in conjunction with resistive surge protection schemes for the main filter capacitors (and bridge rectifiers) during power up, will prevent any static voltage drop across the current limiting resistor due to the amplifier class AB idling current.

Some means must be provided, as well, to protect the output transistors from any condition that could cause their gate-to-emitter voltages to exceed the maximum allowed value of ±20 V. Thus, Zener diodes D_9 and D_{10} are connected from either side of the bias generator to the main output, and prevent the voltage seen between the gain stage and the emitters of the IGBTs from exceeding more than about 12 V.

The IGBT output stage is operated in a complementary emitter-follower configuration running at an idling current of about 100 mA, and series gate resistors R_{23} and R_{26} are included to limit the frequency response. This mitigates any tendency, in the fairly wideband output stage, towards parasitic oscillation. Current in the output stage is sensed across two low value resistors, R_{24} and R_{25}, connected in series with the emitters of the IGBTs. As the voltage drop across either of these two resistors increases towards 0.7 V, Q_{12} or Q_{13} will begin to conduct current away from the gain stage and thus limit the

output voltage. This is a convenient way to limit the current in the output stage to a safe value. Emitter degeneration resistors (10 Ω) must be used in conjunction with the two limiter transistors, Q_{12} and Q_{13}, because this circuit has quite a bit of gain when active and tends to oscillate slightly at high frequencies. Since these transistors must sink or source all the current from the transimpedance stage (up to the current limit of A_1) when the output current is being limited, the voltage across the 10 Ω resistors will increase slightly as the amplifier is driven into hard limiting. This causes a corresponding increase in the actual value of limited current, resulting in a somewhat "soft" limiting curve.

Of course, current limiting alone is not enough to guarantee power transistor integrity if short circuits to ground at the output are anticipated. This results from the fact that excessive power dissipation in the output stage will still occur if the current limit is set fairly high (actually a very desirable attribute in a modern amplifier). Fusing the power supply feed to the output stage will usually be necessary for protection of the power transistors.

PRACTICAL CIRCUIT CONSIDERATIONS

It is often mistakenly assumed that once a respectable topology has been chosen for a power amplifier, it is a simple task to construct a completed unit that meets all the original design goals. In fact, getting the physical details of an amplifier's construction properly sorted out can be just as time consuming as the actual design of the driver electronics themselves.

Circuit Board Layout

This is probably one of the most critical elements for a wide bandwidth audio power amplifier. The key to good board layout for this design is to keep trace lengths to an absolute minimum wherever possible, and to keep the overall layout very small in physical size. Figure 4 shows the layout of the board used to characterize this new topology, and as can be seen, the component packing density is reasonably high—it measures less than 9 cm on a side. The layout of the driver board actually follows the amplifier circuit diagram fairly closely in orientation, since it was begun on the left hand side where the input stage resides, and finished on the right where the output stage drivers are located.

Power supply busses travel along the top and bottom edges of the board, thus providing a convenient means of picking off power for the various stages. The two polyfilm bypass capacitors on the board actually have their own ground return paths that are independently isolated from the signal ground bus near the input stage. This may seem a like a subtle refinement, but on the original layout the bypass capacitors shared the same ground bus as the input stage, and strange low level oscillations were noticed on the first prototypes. It turned out that the oscillation was occurring as a result of the discharging of the bypass capacitors into the driver transistors (and ultimately the gates of the output

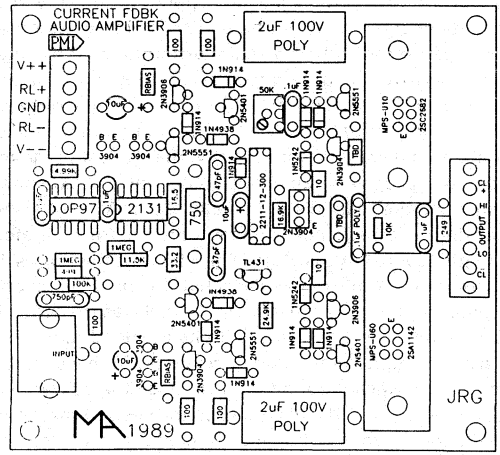

a. Topside Silkscreening

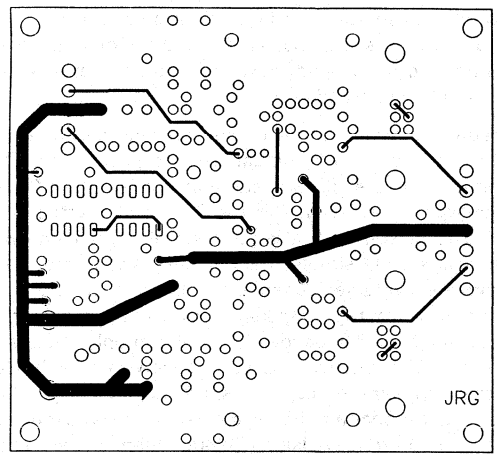

b. Topside Layout

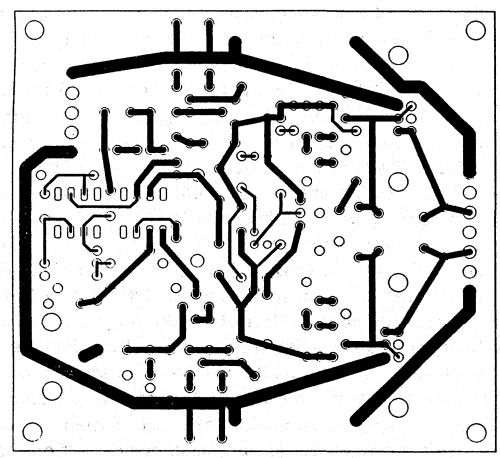

c. Bottomside Layout

Figure 4. A Compact Driver Board Contains All Amplifier Circuitry Except the IGBT Output Stage.

transistors) due to an initial perturbation in the circuit. This initial disturbance eventually lead to a self sustaining relaxation oscillation (of a few hundred hertz) because the ground bus surge, as the capacitors were discharging, was sufficiently large so as to be coupled back into the input stage of the amplifier. The improved layout of Figure 4 does not exhibit this anomaly.

Critical Component Selection

Some of the resistors in this design require great care in their selection, since the wrong type of resistive element will lead to unexpectedly poor performance. In particular, the 750 Ω feedback resistor R_8 should be an oversized completely noninductive metal film power resistor with a dissipation rating of at least 2 W (remember that the peak current in this component may be as high as 75 mA). Failure to use a resistor with a high enough power rating will very likely lead to thermal modulation of the actual resistance value and a corresponding increase in overall amplifier intermodulation distortion when large low frequency input signal components are present. Additionally, a low temperature coefficient of resistance is very desirable for this part. The current sensing resistors in the output stage (R_{24} and R_{25}) should also be of the low or noninductive variety. Since the short rise time of the amplifier (approximately 350 ns) means that a large di/dt in the load, and hence these resistors, can occur, any excessive inductance will cause the voltage across them to increase during fast edge transitions thus causing premature current limiting.

Input amplifier A_1 plays a significant role in the overall performance of the amplifier. It must possess all the desirable characteristics of a good line level audio op amp (namely low distortion, high slew rate and wide gain bandwidth product), plus it must have good output current capability as well. The SSM2131 BiFET audio op amp with a GBW of 10 MHz and slew rate of 40 V/μs more than meets the requirements for this design. Also, amplifier A_2 in the integrating dc control stage must have very low input offset current in addition to low offset voltage. This is because 1 MΩ resistors are used, in series with its input pins, to obtain the long time constant needed in this stage. Too large an input offset current would cause a sufficiently large differential dc error to appear across these resistors (many mV) and it would render a low input offset voltage op amp totally useless. The OP-97 adequately satisfies these requirements with an input offset current of only 30 pA and offset voltage of 30 μV.

Paralleling Output Transistors

This is an extremely important topic because most amplifiers will use more than one pair of output transistors per channel, so that low impedance loads can be accommodated without the output stage self-destructing. Since the maximum power dissipation in the output stage increases with decreasing load impedance, it is desirable to ensure adequate static and dynamic current sharing amongst all the output transistors. This will minimize the junction-to-case temperature rise in any one output device. Power MOSFET output stages can be effectively made to share current by means of tight thermal coupling between all transistors, and through the inclusion of appropriately valued series source-ballasting resistors. There is no reason to believe that power IGBT output stages, with their very similar square law transfer characteristics, will behave any differently if the same techniques are employed.

Typically for best current sharing in a MOSFET output stage, the value of the source resistors should be >>1/gm of each transistor over its desired drain current range. Since the transconductance is lowest at the output stage quiescent point, using this value of gm should guarantee sharing over the full output current range. Unfortunately, in practice this may lead to rather large resistance values and correspondingly large voltage drops when high values of load current are being delivered. A better solution is to do a limited amount of prescreening on the N- and P-channel IGBTs to eliminate any devices with larger than average characteristic deviations in V_{TH} and gm (at the idling point). Once this is done, it becomes feasible to use series emitter resistors in the range of 0.2/gm to 0.5/gm, which will help to minimize the voltage drop. For the Toshiba IGBT output devices used in this design, the typical gm at an emitter current of around 100 mA is close to 1S. For example, if an eight transistor output stage is needed that must have a total idling current of 400 mA, series emitter resistors in the range of 0.2 Ω to 0.5 Ω are acceptable along with some limited screening of the output transistors prior to installation.

Wiring Techniques

Some amplifier designers relegate power supply and output terminal wiring to the lowest level of the design phase. However, since these wires may carry large pulsating currents with a harmonic content well above the audio band, it pays to devote some attention to this task. Wiring is probably one of the most critical things that must be accomplished successfully, if the final design is to get anywhere close to the performance measured on a prototype breadboard (where the wires are normally quite short). Usually the layout of the power supply wiring is not particularly well controlled, but some very simple rules should be observed that will maximize the likelihood of success at first power up.

One of the most important rules in wiring layout is to use twisted pairs for the forward and return currents paths in any loop. This minimizes the series inductance of the conductors, since inductance increases with cross sectional loop area. Thus all power supply wires from the filter capacitors to the amplifier output stage(s) (and driver boards) should be twisted together, as shown in the system connection diagram of Figure 5. Fuses are placed in series with the power rails to protect the output stage in the event that an accidental short circuit in the load occurs. They should be of the fast blow type, and must be rated appropriately so that they will not

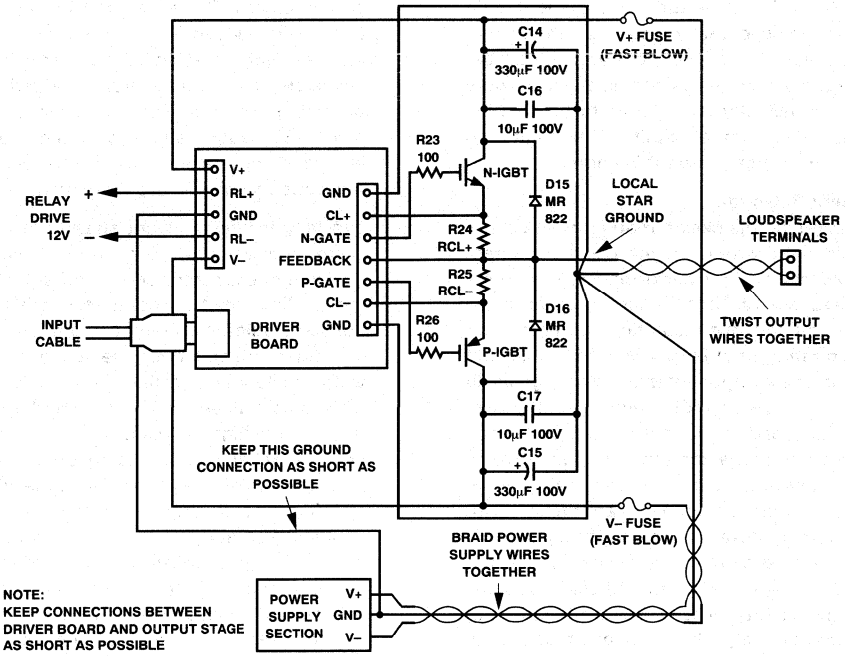

Figure 5. A Power Wiring Scheme Requires Proper Attention to Detail If Low Distortion Is to Be Achieved.

open up under peak output power levels. The wires that run from the output stage to the loudspeaker connectors should also be twisted together, as shown, to minimize their inductance. All interconnections between the driver boards and their respective output stages should be kept very short, in the interests of closed-loop stability. The series gate resistors for the IGBTs should be connected directly at the package terminals of these devices.

These tips are all a definite step in the right direction, but there is something else to consider that is decidedly not obvious. Since the positive and negative supply leads which feed the output stage(s) have half-wave rectified current waveforms, as shown in Figure 6, the harmonic current (occurring at even multiples of the fundamental output frequency) circulates in the loop formed between the power supply capacitors and the output transistors[2]. If there is any mutual inductance between these power supply leads and the output terminal loop, after the point at which negative feedback has been extracted, even order distortion components can be induced in the output that cannot be attenuated by the feedback action of the amplifier. For a typical amplifier with $R_L = 8\ \Omega$ and sinusoidal excitation, then at an output frequency f = 10 kHz, the induced second harmonic component in the output loop will be approximately 0.33% per µH of mutual inductance. It should be noted that the magnitude of the induced distortion components is proportional to the output frequency (i.e., they get larger as the frequency goes up), which can be minimized by keeping the power input and speaker wiring runs perpendicular to each other. Thus the output transistors should be physically connected to the power supply feed and output terminal cabling as shown in Figure 7. This approach minimizes the mutual coupling between the power input and output paths of the amplifier.

Heatsinking and Thermal Considerations

Heatsink selection should never be underestimated because it is one of those critical areas that, if neglected, will inevitably result in damage to the output transistors from excessive junction temperature. In most class AB power amplifiers, the total dissipation in the output stage is split equally between the two banks of output transistors (the N-channel units and the P-channel units). An equation that relates the power supply rail voltage and load impedance to the total maximum output stage power dissipation, under sinusoidal excitation, is given by:

$$P_{DISS}\,(max) = \frac{2\,V_{CC}^2}{\pi^2 |Z_L| \cos\theta} \qquad (6)$$

where θ is the phase angle of the load. As an example, consider the case of an amplifier with a two transistor output stage powered by ±60 V rails, and loaded by an impedance of 8 Ω with a phase angle of +30°. Under these conditions the maximum dissipated power will be 105.3 W. The Toshiba N- and P-channel IGBTs are rated for 180 W dissipation at a T_C of 25°C, but this is derated to zero at a T_C of 150°C. The junction-to-case thermal resistance ($R_{\theta JC}$) for these transistors is calculated by dividing the total difference in case temperature change (125°C) by that of the total change in power dissipation

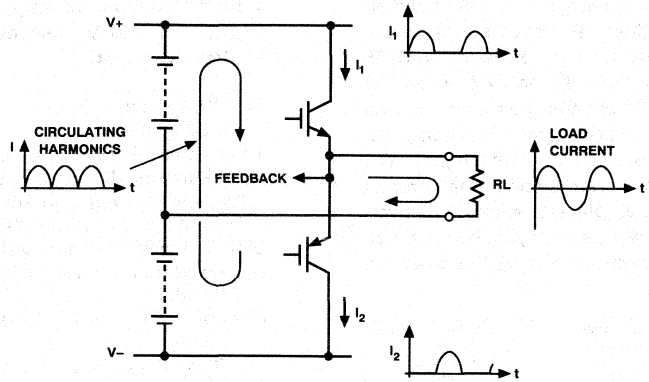

Figure 6. Harmonic Currents in a Power Amplifier Circulate Between the Supplies and the Class AB Output Stage.

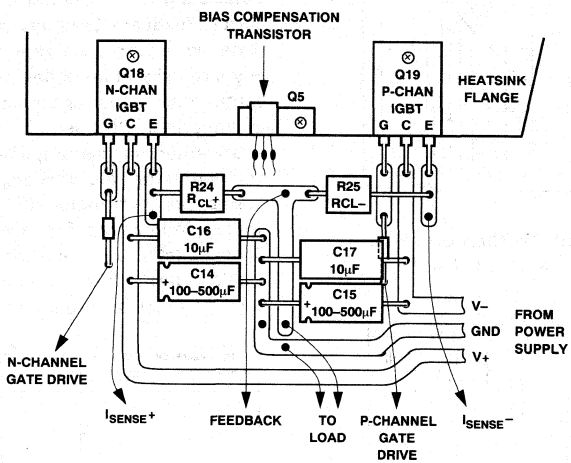

Figure 7. The Preferred Output Stage Layout and Component Placement

(180 W). This results in a figure of 0.694°C/W. Since the total power dissipated in the output stage is split equally between the two transistors, the effective $R_{\theta JC}$ is equal to 0.694/2 or 0.347°C/W. To ensure that the output stage transistors do not reach their maximum allowed junction temperature of 150°C, the total thermal resistance from junction-to-ambient (assuming T_A = 25°C) must not be greater than 125°C/105.3 W or 1.19°C/W. When the junction-to-case thermal resistance of the total output stage is subtracted from this number, we are left with the net allowed case-to-ambient thermal resistance ($R_{\theta CA}$) of 0.843°C/W. This value includes any thermal resistance due to the insulating washers that must be used to prevent the transistors from making electrical contact with the heat sink (often as much as 0.3°C/W per insulator). Thus in reality, some allowance for the interface materials must be made in the choice of the final extrusion which will provide heatsinking for the power transistors. In the example here, a large finned heatsink with a sink-to-ambient thermal resistance ($R_{\theta SA}$) of around 0.69°C/W is required. Of course, had two pairs of transistors been used in the output stage, the net $R_{\theta JC}$ would have been lower by a factor of two and a smaller extrusion could have been used for the heatsink. Thus there is a limited trade-off that can be made between the number of transistors and the size of the output stage heatsink, for a given power supply rail voltage and load impedance.

MEASURED PERFORMANCE

Table I provides a synopsis of the overall performance of the current feedback power amplifier using the new complementary IGBT output devices. Although this design does not achieve astoundingly low distortion levels typical of more complex topologies that employ linearization schemes in the output stage, the measurements made show that the THD and IMD generated by this circuit are still respectably low. Figure 8 shows that the overall harmonic distortion at 50 W output into an 8 Ω load is a minimal 0.001% at 1 kHz, rising to just under 0.009% at 20 kHz. This is a particularly good result considering that only one pair of output transistors has been used. Also, no low-pass LR isolation network has been used in series with the output that would tend to

attenuate the high frequency harmonics. This would artificially improve the amplifier performance in the vicinity of 20 kHz, and was deliberately excluded. SMPTE intermodulation distortion for 60 Hz and 7 kHz mixed 4:1 is plotted in Figure 9 as a function of the rms input level and, as the curve indicates, it is extremely low being just 0.0004% at 41.7 W into 8 Ω (0.92 V rms input). The absence of any significant upward slope in the curve of Figure 9, except where the amplifier is entering its overload region at about 0.95 V rms input, indicates a lack of thermal modulation effects on the 750 Ω feedback resistor.

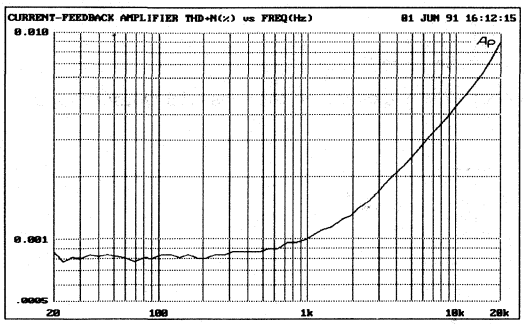

Figure 8. Amplifier THD Is Below 0.009% Throughout the Audio Band When Delivering 50 W to An 8 Ω Load.

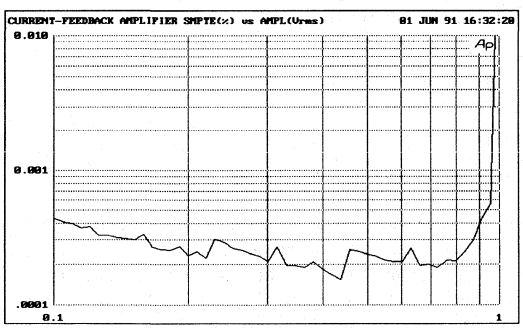

Figure 9. SMPTE Intermodulation Distortion (60 Hz/7 kHz 4:1, 40 W into 8 Ω) Is Exceptionally Low, Reaching Almost 0.0002% Before Rising as the Amplifier Enters its Overload Region.

Table 1. Summary of Amplifier Performance (V_{SUPPLY} = ±40 V, Current Limited to 2.5 A Average Per Rail, R_L = 8 Ω)

Sine Wave Power Output (Voltage Limited)	70 W
Total Harmonic Distortion at 1 kHz	0.001% at 50 W
Total Harmonic Distortion at 20 kHz (Depends Strongly on Idling Current Level in Output Stage)	0.009% at 50 W
SMPTE Intermodulation Distortion	0.0004% at 41.7 W
Dynamic Intermodulation Distortion (DIM-100)	0.0012% at 50 W
Frequency Response (−3 dB)	DC to 1 MHz
Slew Rate	>200 V/μs
Rise Time (Input Filter in Circuit)	400 ns
Total Quiescent Supply Current	130–150 mA

Static distortion measurements aside, what does put the current feedback topology into a class of its own is the dynamic performance. High slew rate is always critical in any large signal amplifier design, but proper waveform control during the reproduction of a square wave is just as important. Because of the nature of the gain stage arrangement in this amplifier, slew rate limiting occurs at a very large rate of change (typically 250 V/μs). Most normal program material is unlikely ever to cause slew limiting in this amplifier, even with large output swings. Consequently, the value measured for the DIM-100 dynamic intermodulation distortion test is a very low 0.0012% at 50 W output into 8 Ω, as shown in Figure 10.

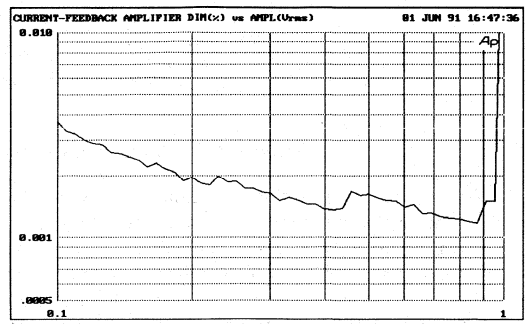

Figure 10. Low DIM-100 Transient Intermodulation Distortion (3.15 kHz/15 kHz 4.1, 50 W into 8 Ω) Results from the Clean Transient Response.

This is the lowest value of DIM-100 distortion that the author has ever seen reported for a solid-state power amplifier. In numerous listening tests, the "fast" sound of this amplifier and its tight LF performance have been commented upon. The large signal step response of the amplifier into an 8 Ω load at 100 kHz is shown in Figure 11, and the no load response with an 80 V p–p square-wave at the output is shown in Figure 12. Either photograph reveals that the amplifier is inherently stable and exhibits no trace of overshoot on fast edges.

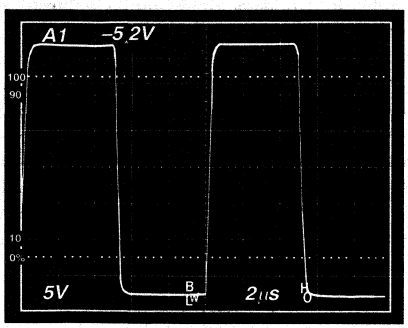

Figure 11. The Amplifier Exhibits Minimal Overshoot When Driving a High Frequency Square Wave into an 8 Ω Load.

Finally the frequency response, as shown in Figure 13, does indeed verify the somewhat overbearing calculations done earlier and proves that the closed-loop bandwidth extends all the way out to 1 MHz. Such a wide frequency response is definitely overkill for any audio power amplifier (200 kHz–300 kHz is probably more than adequate), but it does show what is achievable with a modern design.

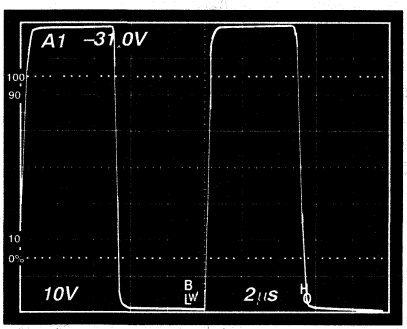

Figure 12. A Large Signal Square Wave at 100 kHz Shows that the IGBT Output Stage Is Inherently Stable Even Without a Load.

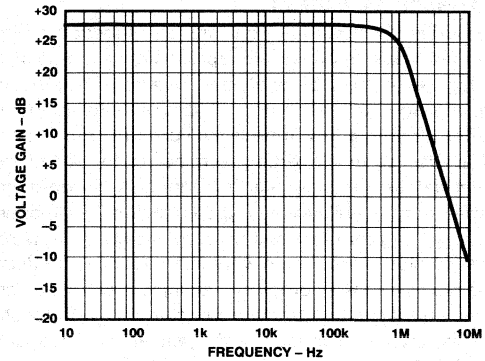

Figure 13. The Small-Signal Frequency Response Does Indeed Extend All the Way Out to 1 MHz, as Predicted by the Calculations.

CONCLUSION

Once in a while a new design comes along that offers a different way of doing an old job. The amplifier that has been presented here offers an evolutionary approach to the task of driving a loudspeaker. When proper attention is paid to all the details (and some of them are nontrivial indeed), current feedback amplifiers can offer superior sonic performance to all known topologies.

REFERENCES

[1]Mark Alexander, "A Current Feedback Audio Power Amplifier," 88th Convention of the Audio Eng. Soc., reprint #2902, March 1990.

[2]Edward M. Cherry, "A New Distortion Mechanism in Class B Amplifiers," Journal of the Audio Eng. Soc., Vol 29, No. 5, pp. 327-32 8, May 1981.

APPENDIX A: FREQUENCY RESPONSE OF CURRENT FEEDBACK AMPLIFIERS

To derive the input-to-output transfer function of a current feedback amplifier, the representative model shown in Figure 14 must be analyzed. Instead of a differential input stage, this topology utilizes a unity gain input buffer, driving a low impedance current summing node, which forces the inverting terminal to be at the same potential as the noninverting input. A nonzero input buffer output resistance, R_{INV}, is shown in series with the inverting terminal and must be included in the analysis of closed-loop gain versus frequency. Neglecting this resistance is a common oversight in simplified analyses, and leads to a transfer function that will not show any bandwidth variation with gain at all. Feedback is applied from the main amplifier output back to the inverting terminal through the current summing network that comprise of R_1 and R_2.

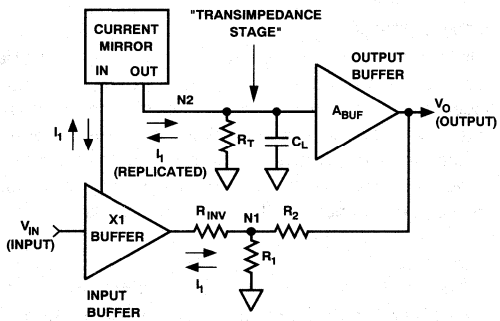

Figure 14.

The action of the input buffer is to force a finite current to flow through R_1 that must be balanced by an almost exactly equal but opposite current in R_2. Any difference between these two currents is an error current that flows into or out of the low impedance inverting terminal. This error current (as opposed to an error voltage in a conventional operational amplifier) is then mirrored and fed into a transimpedance stage, consisting of R_T and Cc, where current-to-voltage conversion takes place. The voltage generated here is buffered by another unity gain stage and is fed to the main amplifier output. Because the value of the small signal transresistance, R_T, is very high (normally several megohms) only minute error currents are needed to change the voltage at node 2 by several volts. Consequently, the amount of current that must flow into or out of the inverting terminal under steady state conditions is extremely small. The feedback network, even though it consists of fairly low value resistors, therefore presents a very light effective loac on the output of the input buffer. To derive a transfei function for this amplifier, nodal equations must be writ ten for nodes 1 and 2, and then combined in an appro priate way to obtain the final result:

$$\frac{V_O}{V_{IN}} = \frac{1+\frac{R_2}{R_1}}{1+\frac{R_2+\left(1+\frac{R_2}{R_1}\right)R_{INV}}{R_T A_{BUF}}+s\frac{\left(R_2+\left(1+\frac{R_2}{R_1}\right)R_{INV}\right)C_C}{A_{BUF}}} \quad (7$$

This relationship is actually very similar to that of a voltage feedback amplifier, and it can be seen that the dc closed-loop gain is nearly equal to $1 + R_2/R_1$ (assuming that the product $R_T A_{BUF}$ is also reasonably large). The low frequency gain term is something with which most users of IC op amps should already be familiar. At a first glance the frequency dependent term might seem to be quite similar to that of a voltage feedback amplifier, but it is in fact very different. This can be seen more easily if the expression for the closed-loop pole frequency is written down:

$$f_{POLE} \cong \frac{A_{BUF}}{2\pi\left(R_2+\left(1+\frac{R_2}{R_1}\right)R_{INV}\right)C_C} \quad (8)$$

Interestingly, this result shows that the pole frequency now depends predominantly on the value of feedback resistor R_2 and the input buffer output resistance R_{INV} multiplied by the closed-loop gain. Normally the value of R_{INV} is made as low as possible to minimize the change in pole frequency with gain, and is typically less than one tenth that of the minimum recommended feedback resistor value. At high gains, as mentioned before, the closed-loop bandwidth starts to become inversely proportional to the gain because the term in the denominator of Equation (8) due to R_{INV} starts to become dominant. The gain bandwidth product is thus:

$$GBW = \frac{A_{BUF}}{2\pi R_{INV} C_C} \quad (9)$$

The gain of the output buffer, A_{BUF}, also plays its part in determining the closed-loop pole frequency. As the main amplifier output is loaded, this gain drops well below unity, and causes a reduction in closed-loop bandwidth as dictated by Equation (8). This actually tends to make the amplifier more stable since the high frequency nondominant poles contribute less additional phase shift at the lower closed-loop −3 dB point. In fact, many commercial current feedback amplifiers show significant gain peaking with light loads, and don't begin to behave acceptably until loaded fairly heavily. Another thing to remember is that the minimum recommended

value for feedback resistor R_2 must be strictly adhered because too low a value will result in an excessively high closed-loop pole frequency. This can result in severe gain peaking due to the higher order poles becoming more dominant, and is especially a problem at low gains when the multiplicative effect of R_{INV} on the closed-loop pole time constant is minimal.

During early development of the current feedback power amplifier it was noticed that instability appeared on the edges of square waves, using the ground referenced compensation scheme. Some experimentation revealed that connecting the compensation capacitors to the feedback summing node made the instability disappear. An analysis of the amplifier response using this new arrangement was undertaken, since something must have changed to make it more stable. Indeed, when the compensation capacitors are returned to the feedback summing node instead of ground, the transfer function of the circuit changes quite significantly. This modified compensation arrangement also allows one to get by with smaller capacitors than before, but without compromising closed-loop stability. To see this, the current feedback model must be analyzed again but this time the compensation capacitor C_C is returned to the summing node instead of ground:

$$\frac{V_O}{V_{IN}} = \frac{\left(1 + \frac{R_2}{R_1}\right)\left(1 + s\left(\frac{2R_1R_2C_C}{R_1 + R_2}\right)\right)}{1 + \frac{R_2 + \left(1 + \frac{R_2}{R_1}\right)R_{INV}}{R_TA_{BUF}} + s\left(\frac{2R_2C_C + \left(1 + \frac{R_2}{R_1} + \frac{R_2}{R_T}\right)R_{INV}C_C}{A_{BUF}} - R_{INV}C_C\right)} \quad (10)$$

The major difference between Equations (7) and (10) is the appearance of a zero in the numerator determined by the parallel combination of R_1 and R_2, and some additional terms in the denominator. The zero tends to partially cancel the second pole of the amplifier due to the IGBT output stage, resulting in greatly improved stability. Probably the most interesting thing to notice about Equation (10) however, is that the R_2C_C time constant is now multiplied by a factor of two instead of unity as before. Since it is this time constant that predominantly determines the closed-loop pole frequency, the original compensation capacitor value can thus be scaled down by a factor of one half.

APPENDIX B: AMPLIFIER COMPONENT LIST FOR A SINGLE CHANNEL

	Quantity	Designator
Integrated Circuits		
SSM-2131P BiFET Audio Op Amp	1	A_1
OP-97FP Precision DC Op Amp	1	A_2
TL431CP Programmable Shunt Regulator	1	D_7
Transistors		
2N3904 NPN, 40 V (4 Are Used as Zener diodes)	7	Q_5, Q_7, Q_{12}, Q_{14}–Q_{17}
2N3906 PNP, 40 V	2	Q_3, Q_{13}
2N5401 PNP, 150 V (or 2SC2682 from NEC)	3	Q_2, Q_4, Q_9
2N5551 NPN, 160 V	3	Q_1, Q_6, Q_8
MPS-U10 NPN, 300 V[1]	1	Q_{10}
MPS-U60 PNP, 300 V[2]	1	Q_{11}
GT20D101-Y N-CHAN IGBT 250 V, 20 A (Toshiba)	1	Q_{18}
GT20D201-Y P-CHAN IGBT 250 V, 20 A (Toshiba)	1	Q_{19}
Diodes		
1N914 100 V, 100 mA Small Signal Diode	9	D_1, D_2, D_5, D_6, D_8, D_{11}, D_{12}, D_{13}, D_{14}
1N5242B 12 V, 500 mW Zener Diode	2	D_9, D_{10}
1N4938 200 V, 100 mA Low t_{RR} Diode	2	D_3, D_4
MR822 200 V, 5 A Low t_{RR} Rectifier	2	D_{15}, D_{16}
Resistors		
(All Values Are in Ohms, and Are 1/4W 1% Metal Film Unless Otherwise Specified)		
0.05 Ω, 3 W, 5% Noninductive (Shallcross LO-3 Series)	2	R_{24}, R_{25}
10.0	2	R_{19}, R_{20}
16.5	1	R_7
33.2	1	R_6
100 (2 Are Used as Gate Resistors for the IGBTs)	7	R_3, R_{13}–R_{16}, R_{23}, R_{26}
249	1	R_{22}
499	1	R_{10}
750, 2-5 W, 1% Noninductive Metal Film	1	R_8
4.99 k	1	R_{12}
10.0 k	1	R_{21}
11.5 k	1	R_{11}
16.9 k	1	R_{17}
24.9 k	1	R_{18}
100 k	1	R_4
1.00 M	2	R_5, R_9
R_{BIAS} (49.9 k with ±65 V Power Rails)	2	R_1, R_2
50 kΩ Multiturn Trimpot (Helitrim 68WR503 or Equivalent)	1	VR_1
Capacitors		
47 pF 5% Silvered Mica (or Ceramic) 200 V	2	C_6, C_7
750 pF 5% Silvered Mica (or Ceramic) 200 V	1	C_3
0.1 μF 10% Ceramic or Mylar 63 V	4	C_4, C_5, C_8, C_{10}
1 μF 10% Ceramic 100 V	1	C_{11}
2 μF 10% Polyfilm 100 V (Electrocube 230B1B205K)	2	C_{12}, C_{13}
10 μF 10% Tantalum Electrolytic 25 V	3	C_1, C_2, C_9
2 to 10 μF 10% Polyfilm 100 V	2	C_{16}, C_{17}
220 or 330μF 10% Aluminum Electrolytic 100 V	2	C_{14}, C_{15}
Miscellaneous:		
Form-C Reed Relay (Coto 2211-12-300)	1	K_1
Thermalloy 6100B Heatsink for the Driver Transistors	2	
Extra Large Finned Heatsink for the IGBT Output Stage	1	
Insulating Pads for the IGBTs	2	
5-Pin Molex Header 0.156 Inch Pin Spacing	1	
7-Pin Molex Header 0.156 Inch Pin Spacing	1	
3-Pin Molex Header 0.100 Inch Pin Spacing	1	
Right Angle RCA Jack	1	
Amplifier evaluation PC Board[3]	1	

[1]First choice substitution is NEC 2SC2682; second choice Toshiba 2SC2238B. Note correct pinouts.
[2]First choice substitution is NEC 2SA1142; second choice Toshiba 2SA968B. Note correct pinouts.
[3]Available to qualified OEMs. Contact local ADI sales office for details.

ANALOG DEVICES

AN-232
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Bootstrapped IC Substrate Lowers Distortion in JFET Op Amps

by Walt Jung

INTRODUCTION

Among the more popular IC op amps used today are standard process junction isolated FET types (JFETs). These amplifiers typically use a pair of differential P-channel FETs at the input, with remaining stages of bipolar NPN/PNP transistors. Specs typical for the better of these devices are largely good to excellent. For example, the input currents are only a few pA at room temperature, the common-mode rejection is quite good (albeit below that of the best all-bipolar op amps), and they can have excellent dynamics. On the whole, these attributes make for low overall dc errors and generally good ac performance. In terms of total harmonic distortion (THD) performance, it can be as low as 0.001% (10 ppm) over the audio range for some applications.

THE PROBLEM OF NONLINEAR INPUT CAPACITANCE

However, there is one factor which can cause rapid deterioration of JFET op amp THD performance, when operated from high source impedances. This is a phenomenon of distortion increase with increasing input signal rate-of-change, due to nonlinearity of the capacitance seen at the two inputs. Since junction isolated ICs use isolation "pockets" or "wells" to host the subject differential pair of P-channel FETs, there will necessarily exist parasitic capacitance from the three FET terminals to the IC substrate. Logically, the most troublesome of the three terminals is the gate, due to the high source impedances which can potentially be seen externally. A typical input capacitance at the inputs of a JFET IC is 3 pF–5 pF, and is seen from Pins 2 and 3 to the substrate (normally Pin 4).

This parasitic substrate capacitance has a nonlinear behavior as a function of applied CM voltage. Like the C/V characteristics of discrete FETs and bipolar transistors, the capacitance is minimum at higher bias voltages. Thus in an op amp application with a given CM input level, this suggests operating at the highest practical supply voltage within rating to minimize the effect. The input capacitance varies on an instantaneous basis with the applied ac CM voltage, and with a high impedance source, can generate high levels of distortion.

For JFET amplifier stages which operate as high source impedance followers, the CM input signal variation can give rise to excessive THD. Here, "high impedance" is qualified as 50k or greater, an approximate threshold where this distortion begins to be easily detectable at audio frequencies. For source impedances of 0.1 MΩ to 1 MΩ, the distortion can be as high as several tenths of a percent in the audio band.

The distortion is most easily seen as a 6 dB/octave rise in second harmonic distortion using a fixed level frequency sweep. The exact magnitude will vary with the device in use, and the specific combination of operating level, loading, supply voltages, and relative source impedance(s).

The designer has control over many of these circuit issues, and to some extent they can be optimized. One beneficial step which can lower frequency dependent distortion products is the use of *input impedance compensation*, a technique which ideally balances the source R/C components at the two amplifier inputs.[1] When used, this compensation serves to minimize the nonlinear effects of FET input amplifier common-mode capacitance.

To implement it, an optional feedback resistance "R_{S2}" is used in the circuit, where R_{S2} is simply made equal to the net nominal source impedance, "R_{S1}". In the Figure 1 example circuit, R_{S1} is a fixed 500k resistor inserted for testing purposes. . . . in a real application, R_{S1} could in fact be an inseparable part of the source. If the source is also capacitive, a compensatory feedback capacitance can be used in addition to R_{S2}. This technique can be effective in terms of lowering distortion due to nonlinear input capacitance.

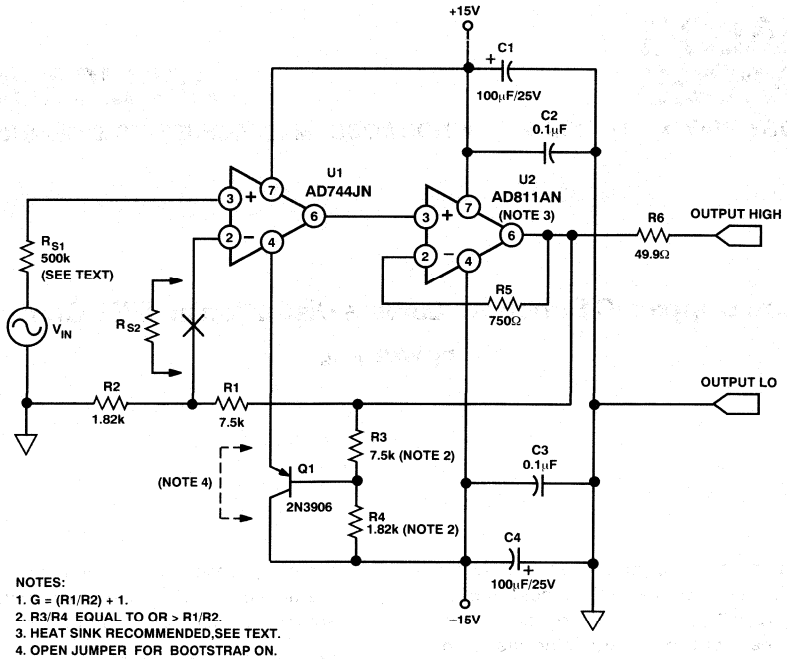

Figure 1.

MINIMIZING NONLINEAR INPUT CAPACITANCE EFFECTS

Unfortunately, there are still situations where it is desirable to operate a noninverting amplifier with unbalanced feedback/source impedances, but yet maintain minimal signal degradation in terms of noise/distortion. For example, the source impedance may simply be unknown. If it is higher than a few hundred kΩ, a JFET IC op amp is desirable from a noise standpoint, because of the low current noise. But, at the same time, it is undesirable from a distortion point of view due to the input C/V nonlinearity.

Figure 1 shows how a high performance JFET input op amp can be used in a circuit that reduces the nonlinear effects of CM input capacitance by bootstrapping the amplifier substrate. In this circuit stage U1 is an AD744, a low distortion JFET input op amp,[2] which determines the circuit's basic input characteristics. As will be shown, the AD744 is not unlike other JFET devices, and distorts with unbalanced high Z inputs due to nonlinear input C/V. When bootstrapped most effectively, however, the distortion due to input C/V nonlinearity can be reduced below the residual noise level.

With U1 loaded as shown by only the high impedance input (+) of buffer amp U2, it has virtually zero current drive requirements. This factor helps to maximize overall load dependent linearity, as well as to aid the bootstrapping action. Stage U2 provides primarily a high current output, which in this application is an AD811AN,[3] a high performance transimpedance amplifier used here as a voltage follower buffer. Designed primarily for video use, the AD811's relevant specs are a slew rate of 2500 V/μs, a bandwidth of 120 MHz, and 100 mA of output current. These factors enhance this circuit by providing both high and linear load drive capability, for loads of 600 Ω and below.

Because the buffer amplifier U2 eliminates the output stage current variations in U1, the bootstrap transistor Q1 can more effectively deliver bootstrap correction voltage to Pin 4 of U1. U1 is operated in a near ideal fashion with regard to loading, a factor which enhances linearity.

DESIGN FACTORS

The design as shown in Figure 1 operates at an overall voltage gain "G" set by R1 and R2 just as in a conventional noninverting amplifier, or:

$$G = \frac{R1}{R2} + 1 \qquad (1)$$

As is noted from the figure, there is also a separate resistor divider from the output of U2 to the V-rail, R3/R4. This divider feeds back a portion of the buffered U1 output signal from U2 to the negative supply pin of U1, Pin 4. With the two feedback dividers R1/R2 and R3/R4 matched in ratio, this has the net effect of driving the U1 substrate with a signal equal to that at the (+) input. As a result of reducing the voltage across the input capacitance, it is then effectively reduced. Note that for overall stability, the feedback via R3–R4 must be less than that of R1–R2. This will be so if this inequality is satisfied:

$$\frac{R3}{R4} \geq \frac{R1}{R2} \qquad (2)$$

It should be noted that this bootstrap technique is most effective at gains of 5 or more, because the R3/R4 divider tends to reduce the U1 supply voltage, and thus limits dynamic range. In the design process the standard gain expression is used for G, then R3/R4 are selected to satisfy Equation 2. Here, with G equal to 5.12, R3/R4 are simply selected for a ratio similar to this.

PERFORMANCE

With these design and device selection factors, the bootstrapped composite amplifier performance is remarkable considering its modest complexity. The following tests operated the Figure 1 circuit at ±15 V supplies and a 500 kΩ source for R_{S1}, and no compensation was used at R_{S2}.

For reference, the THD of a similar but nonbootstrapped circuit (i.e., Q1 jumpered) was about 0.1% at 10 kHz at a 3 V rms output level from U2, varying from a low of 0.01% at 1 kHz to as high as 0.2% at 20 kHz. At a fixed signal frequency of 10 kHz, Photo 1 shows this second harmonic distortion, at a level of 0.1%. The upper trace is the U2 output, while the lower trace is the distortion analyzer output set to a resolution of 0.3% full scale.

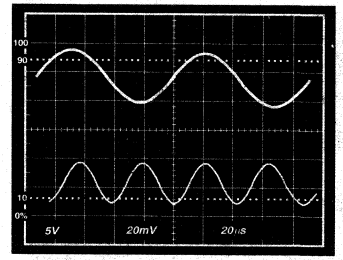

Photo 1. Output Signal (Top) and Distortion Components (Bottom, 0.3% Full Scale). Conditions: V_S ±15 V, Figure 1 Circuit, Without Bootstrapping, 3 V rms Out.

In contrast, Photo 2 shows the distortion under bootstrapped conditions, with a resolution of 0.03% full scale for the lower trace. In this case the harmonic distortion has essentially disappeared into the residual noise, indicating the potency of the bootstrap circuit.

4

UNBUFFERED BOOTSTRAP

It is worthwhile to note that the basic principle of this bootstrap will still operate to some degree without buffer amplifier U2, that is with U1 driving the load direct. However, the distortion null will not be as deep as it is shown in Photo 2, and it will also vary with loading.

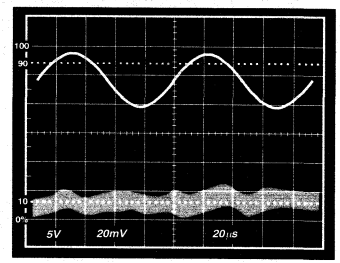

Photo 2. Output Signal (Top) and Distortion/Noise Components (Bottom, 0.03% Full Scale). Conditions: V_S ±15 V, Figure 1 Circuit, With Bootstrapping, 3 V rms Out.

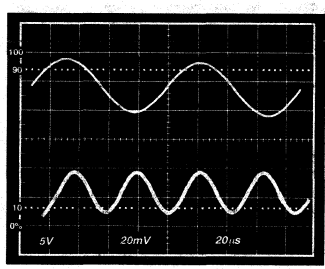

Photo 3. Output Signal (Top) and Distortion Components (Bottom, 0.1% Full Scale). Conditions: V_S ±15 V, Unbuffered U1 in Figure 1 Circuit, With Bootstrapping, 3 V rms Out.

Photo 3 shows this condition, where U1 drives the bootstrap divider directly, and the THD measures 0.03%, using a scale factor of 0.1% in the lower trace. For this example, this represents about a 10 dB distortion reduction over the same circuit conditions without the bootstrap.

The deterioration of distortion null is because U1 is called upon to deliver output current, which must pass through Q1. The V_{BE} of bootstrap transistor Q1 will then change with U1 load current, degrading the bootstrap tracking of the input drive to U1. The circuit selection choice is a trade-off between simplicity and less performance for the unbuffered case, versus the complexity of an additional IC, but much greater overall performance for the fully optimized Figure 1 version.

APPLICATION TIPS

For other operating hints, maximum output will be function of the power supplies, and will be 7 V rms c more with the ±15 V supplies used here. For suppl voltages of ±12 V or more a clip on heat sink is recom mended when the AD811AN is used for U2 (Aavid 5801 For low impedance loads, the supplies should be we bypassed with large electrolytics, returned to the loa common point.

Finally, it should be noted that the general principles o the nonlinear C/V input characteristic distortion mecha nism apply to virtually all JFET input IC op amps avail able today, independent of their source. Bipolar input o amps built on junction isolated processes can also b subject to the phenomenon, and may also benefit fro bootstrapping. However they are not as likely to be use with large source impedances where the distortion mag nitudes become a serious problem, as is true in the cas of JFET op amps.

References

[1]Analog Devices AD743 data sheet.

[2]Analog Devices AD744 data sheet, Figures 16 and 20.

[3]Dave Whitney and Walt Jung, "Applying a High-Performance Video Operational Amplifier," *Analog Dialogue*, Vol. 26-1.

Communications Products
Contents

		Page
AN-238	AD2xx Family for RS-232 Communications	5–3
AN-48	Designing Digital Repeaters with ICs	5–9
AN-50	A Variable-Frequency Clock-Recovery Circuit Using the RPT-82 or RPT-83	5–15
AN-119	Improved T148 (CEPT) PCM Repeater for #22AWG (0.7 mm) Unshielded Twisted-Pair Wire	5–17
AN-303	Understanding and Applying the AD7341/AD7371 Switched Capacitor Filters	5–19

ANALOG DEVICES

AN-238
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

AD2XX Family for RS-232 Communications

by Matt Smith

The AD230–AD241 is a product line specifically designed for RS-232 communications applications. All parts in the product line meet or exceed the standard requirements and offer superior performance in many areas. The present RS-232 requirements include conformance to the EIA-232D standard, low power consumption, low cost, high reliability and operation from a single +5 V power supply.

The AD2XX family of interface products meets this need by integrating RS-232 drivers, RS-232 receivers and a charge pump voltage converter onto the same chip. CMOS technology is used to keep the power consumption to an absolute minimum. In addition, some members of the family feature a shutdown or sleep mode which can be used to disable the devices thereby reducing the power consumption even further.

The AD2XX family is designed to meet the EIA-232D specifications while operating from a single +5 V power supply. This is achieved by the use of an on-chip voltage doubler.

Older generation RS-232 drivers required three separate power supplies: +5 V, +12 V and −12 V. This resulted in large bulky power supply units. Linear voltage regulators tend to be inefficient and are wasteful of power. This is especially a problem in today's portable equipment which operates with battery powered supplies. Ideally a single power supply should be used which can easily be derived from a battery pack. Switch mode regulators can achieve this and are efficient in terms of useful power but again can be quite bulky, a serious drawback in portable equipment. In addition to this, switch mode supplies generate severe electrical noise which requires careful screening in order to conform to strict EMI regulations.

The AD2XX solves all these problems by operating from a single +5 V power supply. An on-chip charge pump voltage converter generates ±10 V levels internally, thereby allowing the RS-232 output levels to be developed. The charge-pump technique is an extremely efficient method of stepping up the input voltage and is suitable for applications where the power requirements are modest. It is therefore ideally suited to RS-232-type applications.

CHARGE PUMP OPERATION

The charge pump uses a switched, floating-capacitor technique to double and then invert the input +5 V supply. This generates voltages of +10 V and −10 V.

The voltage doubler schematic is shown in Figure 1, while the inverter is shown in Figure 2.

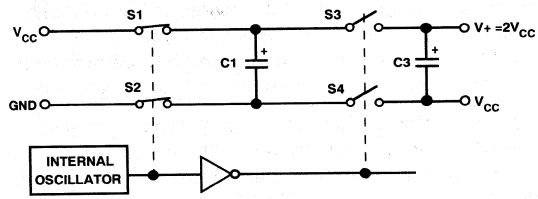

Figure 1. Voltage Doubler

The internal oscillator controls two phases of circuit operation.

During the first clock phase, switches S1 & S2 are closed causing capacitor C1 to charge up to V_{CC} (+5 V).

During the second phase, S1 & S2 are opened and S3 & S4 are closed. The negative terminal of C1 is connected to V_{CC} via S4. The voltage at V+ is therefore V_{CC} + V_{CC} = 2 V_{CC}. Capacitor C3 acts as a reservoir capacitor to maintain the voltage at 2 V_{CC} during clock phases when S3 & S4 are open. It should be noted that this reservoir capacitor is connected between V+ and V_{CC} for optimum performance.

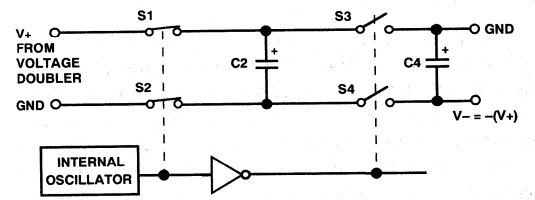

Figure 2. Voltage Inverter

The voltage at V+ is $2V_{CC}$ or +10 V, and this is then used to generate −10 V using a similar technique to that already described. Again, during the first clock phase S1 & S2 are closed thereby charging C2 up to 10 V. During the second clock phase, S1 & S2 are opened and S3 & S4 are closed. The positive terminal of C2 is connected to GND via S3. This forces the potential at V− to −10 V. Again the output reservoir capacitor (C4 in this case) maintains the output voltage relatively constant for clock cycles when S3 & S4 are open.

In order to conserve board space, the values of capacitors C1 to C4 can be reduced. Reducing C1 & C2 results in higher output impedance on the V+ and V− supplies, while reducing C3 & C4 causes increased ripple on the outputs. The increased output impedance & ripple is most noticeable at high temperatures, and if operation at extended temperatures is not required, then it is perfectly acceptable to reduce the component values.

The capacitors on the AD233 and the AD235 are integrated into the package, and so no external capacitors are required thereby reducing board space and saving on components.

Transmitter Outputs

The charge pump voltages (±10 V) are used internally to provide power for the RS-232 drivers. Under worst case conditions of high temperatures and maximum loading, the drivers are guaranteed to provide ±5 V levels on the RS-232 drivers. Typically the outputs provide ±9 V levels. This exceeds the minimum levels and permits operation well beyond the minimum RS-232 requirements.

Slew rate of each output is tightly controlled and is limited to less than 30 V/μs. This is achieved by internal slew limiting, and there is no need for external slew limiting capacitors as is the case with some bipolar designs.

Latch-Up Immunity

Because of the nature of the environment in which an RS-232 link may function, it is extremely important that the interface devices are capable of withstanding several forms of abuse. This can take the form of a user attempting to plug in the connector the "wrong way around." This can cause transmitter outputs from the peripheral device to become momentarily shorted to other transmitter outputs on the terminal. The transmitter outputs on the AD2XX product line are capable of withstanding shorting to ±15 V with the driver output at either polarity. This is the highest continuous voltage which can be present in an RS-232 link.

The parts must also be capable of withstanding signals applied even when power is removed. A peripheral device is often plugged into the terminal serial port before it is powered up. Again the AD2XX family contain internal protection to protect the device. In this case, the protection is on the receiver inputs and takes the form of passive resistive protection. Passive protection has the advantage that it continues to operate even when there are no power supplies to the device.

These protection schemes are also designed to protect the device if the interface cable is incorrectly wired. Due to the confusion which surrounds the RS-232 interface, this is an all-too-likely possibility. Some peripheral devices require signal crossing in the cable or connector, while others require a straight-through connection. This confusion is partly due to the fact that the initial specification applied to a terminal-to-modem interface. In practice the RS-232 port is used to communicate with a wide variety of peripheral devices and is not limited to modem interfaces.

Overvoltage Protection

The driver outputs are protected against damage by overvoltages greater than ±15 V on the outputs. This is achieved by internal series 300 Ω resistors on each transmitter output. This resistor also ensures full compliance with the EIA specification which requires a minimum power-down resistance of 300 Ω on each output. This resistor provides current limiting under fault conditions if a powered-up driver is inadvertently shorted to the powered-down driver.

If for any reason there is a requirement for even greater protection than is inherent in the device, then this may be applied externally. For protection against voltages in excess of ±15 V on the transmitter outputs, external series resistors may be used. A series 100 Ω resistor will provide protection up to ±20 V. This is the simplest scheme, but it causes a slight degradation of the output voltage swing due to the higher output impedance. This is not normally a problem as the output levels are well above the minimum RS-232 voltage level requirements.

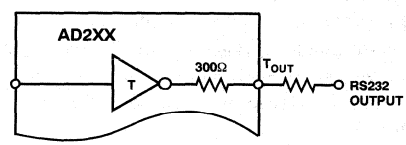

Figure 3. Resistor Protection to ±20 V for Transmitter Outputs

Another form of overvoltage protection uses TranZorbs.* These devices function as transient voltage suppressors and should be connected between the output and GND.

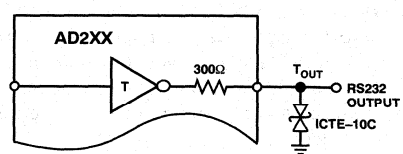

Figure 4. TranZorb Protection for Transmitter Outputs

*TranZorb is a registered trademark of General Semiconductor Industries, Inc.

As the outputs signal may swing either positive or negative, a bidirectional TranZorb should be used. This essentially contains two TranZorbs connected back to back. This scheme will provide voltage clamping at the TranZorb breakdown voltage for both positive and negative excursions. Effective spike suppression is also achieved due to the extremely fast TranZorb response time. A suitable TranZorb clamp rating for this application is ±10 V. This protection scheme does not degrade the output voltage swing as the previous scheme did.

The receiver inputs must also be capable of withstanding excessive input signal voltages. The AD2XX family are protected against overvoltages of up to ±30 V. This exceeds the RS-232 specification of ±25 V. If even higher levels of protection are required, then this can be provided externally by TranZorbs. A suitable clamp rating for this application is ±22 V.

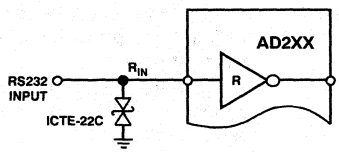

Figure 5. TranZorb Protection for Receiver Inputs

Noise Immunity

An RS-232 interface link is susceptible to noise pick up from the surrounding environment. The longer the link the more susceptible it becomes to outside interference. This is especially a problem in electrically hostile environments such as in a heavy industrial plant where large glitches can be injected or coupled onto the transmission line. These glitches can cause erroneous data reception by the RS-232 receiver. The AD2XX product line is designed to cope with noisy environments by the use of special on-chip filtering circuitry on the receiver inputs. These filters reject fast transient noise glitches up to 1 μs in duration. In addition, further noise immunity is achieved by the use of Schmitt trigger inputs having 0.5 V of hysteresis. The result is a much improved data link offering superior reliability and error-free communication even with severe external noise.

Using an External +12 V Supply

If an external +12 V power supply is available in a system as well as +5 V, then the internal voltage doubler may be made redundant thereby saving two capacitors on the charge pump voltage doubler. The power requirements on the +5 V supply will also be reduced. If this scheme is used, then it is important that the correct power-up sequence be observed. The +12 V supply should be powered up before the +5 V supply. This is generally the case in non-PC systems where the +5 V supply is derived from the +12 V supply using a 7805 type regulator.

If this sequence cannot be guaranteed, then a general purpose diode (1N914 or equivalent) should be connected in series with the +12 V supply. This is illustrated in Figure 6. The purpose of the diode is to prevent current backfeeding from V_{CC} and out the V+ line. Substantial fault-current could flow into a low impedance 12 V supply if this becomes a low impedance node when off. Under these conditions the diode will prevent the fault-current from flowing. Under normal conditions the diode is forward biased and has no effect.

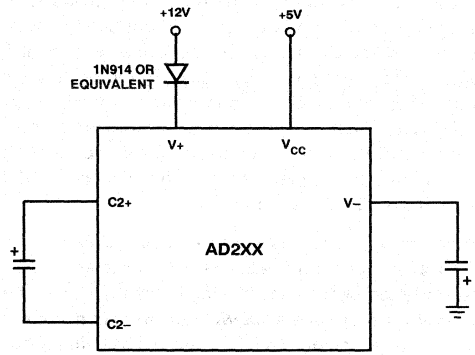

Figure 6. External +12 V Supply

Driving a Mouse

One of the most popular uses of a serial port on a personal computer is as a mouse interface. With the widespread acceptance of Windows,* a mouse has become an essential user interface tool. The older generation of mice interfaced to the PC using the serial communications port and generally required an external power supply to power the mouse hardware. The latest generation of mice use CMOS technology, and it is possible to power these devices directly from the communications port. In other words, the RS-232 transmitters should be capable of providing sufficient power to the mouse.

With a mouse connected, communication is in one direction only, i.e., from the mouse to the PC. Therefore the transmitter outputs on the serial port are not required for communication purposes. Instead, these are used as a power supply for the mouse. The mouse driver software sets up these transmitter outputs at permanently positive or negative levels as required by the mouse hardware. The transmitter outputs must therefore supply sufficient current drive to power the mouse. The power requirements are higher than those required by a standard RS-232 load. With the minimum RS-232 load, 3 kΩ, the driver must maintain ±5 V levels giving a current drive of ±2 mA.

In order to drive a mouse, however, greater current must be available. A typical mouse will require approximately 5 mA on a driver output. The AD2XX product line is designed to maintain ±5 V output levels with a current drain of 5 mA on each driver output.

There is a considerable variation in power requirements between mice from different manufacturers, and some mice may require even greater current drive. This may

*Windows is a registered trademark of Microsoft Corp.

be achieved by connecting two transmitters in parallel thereby doubling the effective current available. However, with the vast majority of mice this will not be necessary.

TYPICAL APPLICATIONS

The AD230–AD241 family contain a variety of transmitter/receiver combinations which will satisfy all communications needs. Some of the products are general purpose devices which are suitable for a wide range of differing applications. Others contain the exact combination of drivers/receivers for a particular function. One of the simplest RS-232 communication links uses 5 signal lines and may be implemented using a single AD232 device. This is illustrated in Figure 7.

More complex interfacing schemes use a greater number of transmitters and receivers. A complete RS-232 implementation designates 22 signal lines. This includes secondary data channels as well as provision for synchronous transmission. This has led to a certain amount of confusion and incompatibility problems between RS-232 terminals.

In practice, most of these lines are redundant. As a result, the industry has formed a de facto "standard" using eight signal lines. This has been shown to be perfectly adequate for most communication needs. Most modern personal computers have adopted this standard and contain a pair of 9-pin serial ports rather than the 25-pin DB25 connector which can be found on older machines. These 9-pin serial ports contain 8 signal lines and a ground terminal. The signal lines are made up from three transmitters and five receivers. This forms a small subset of the original 22 signal lines. Only the most important RS-232 signals are used. The signal designations and their functions are described below.

Pin No.	Function	Source	Abbreviation
1	Data Carrier Detect	DCE	DCD
2	Received Data	DCE	RXD
3	Transmitted Data	DTE	TXD
4	Data Terminal Ready	DTE	DTR
5	Signal Ground		
6	Data Set Ready	DCE	DSR
7	Request to Send	DTE	RTS
8	Clear to Send	DCE	CTS
9	Ring Indicator	DCE	RI

The function of each of these signal lines is as follows:

Function	Description
Transmitted Data	From the DTE to the DCE
Received Data	From the DCE to the DTE
Request to Send	A signal from the DTE to the DCE indicating that it wishes to transmit data
Clear to Send	A signal from the DCE, in response to a RTS, that it now can send the data
Data Set Ready	The DCE telling the DTE that it is connected to the telephone line
Data Terminal Ready	The DTE telling the DCE that it is ready to transmit or receive data
Data Carrier Detect	The DCE telling the DTE that it is receiving valid signals
Ring Indicator	The DCE telling the DTE that it has detected an incoming call
GND	Ground Reference

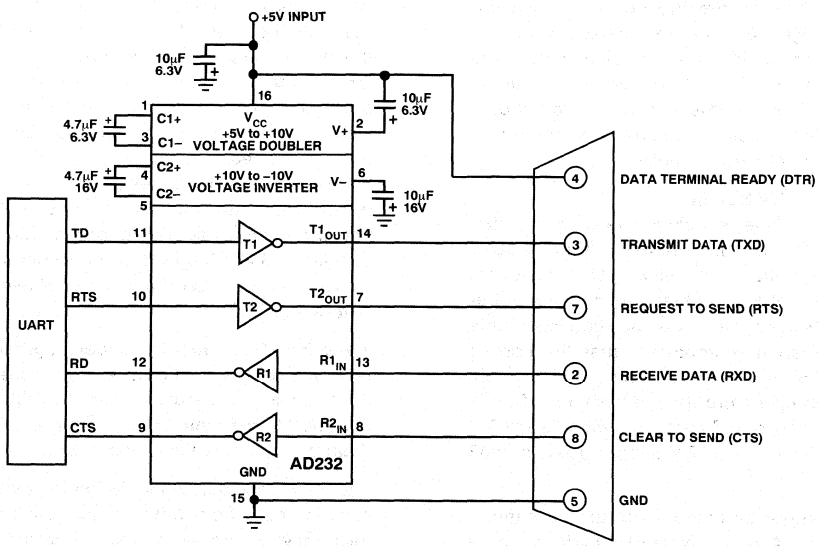

Figure 7. Minimum RS-232 Link Using the AD232

The AD241 is suitable for this 9-pin implementation as it contains sufficient drivers/receivers to meet the requirements in a single package. A typical application showing the AD241 is illustrated in Figure 8.

A typical communications sequence is as follows. The UART in the terminal (DTE) turns on the Request to Send (RTS) line. This signals the peripheral (DCE) that the terminal wishes to send data. The peripheral responds with a Clear to Send (CTS) signal. The terminal now starts transmitting data on its data (TXD) line. When transmission is complete, the terminal turns off the RTS line. In response to this, the peripheral turns off CTS and the link is terminated.

Figure 9 shows a different type of interface where a high resolution Analog to Digital Converter is interfaced to a personal computer using the RS-232 port. This circuit forms part of a remote data acquisition system. The remote ADC transmits serial data back to a personal computer for analysis. In this case therefore, the ADC acts as the peripheral device. This interface differs from the previous examples in that it shows the RS-232 interface device located at the peripheral end of the link. There will of course be a similar RS-232 arrangement located at the PC end.

The ADC chosen is a high resolution converter featuring an asynchronous UART compatible interface. Operation of the interface is as follows. The DRDY output from the AD7701 ADC signals the terminal that the conversion is complete and data is ready. In response to this, the remote terminal activates the DTR line. This indicates that the terminal is ready and this line is used as a chip select for the AD7701. This initiates the data transfer. The data is then transmitted as two serial bytes in UART compatible format.

5

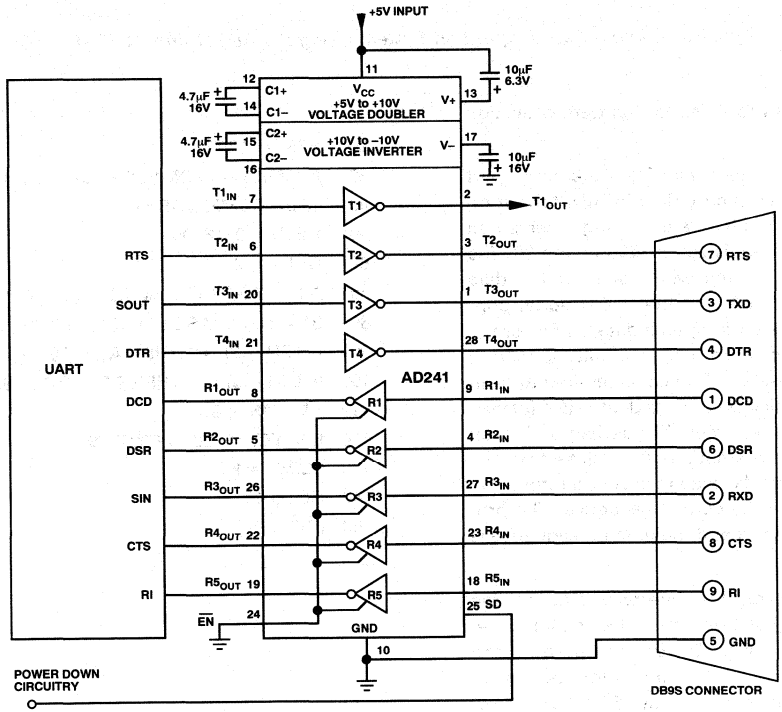

Figure 8. Complete RS-232 Link Using the AD241

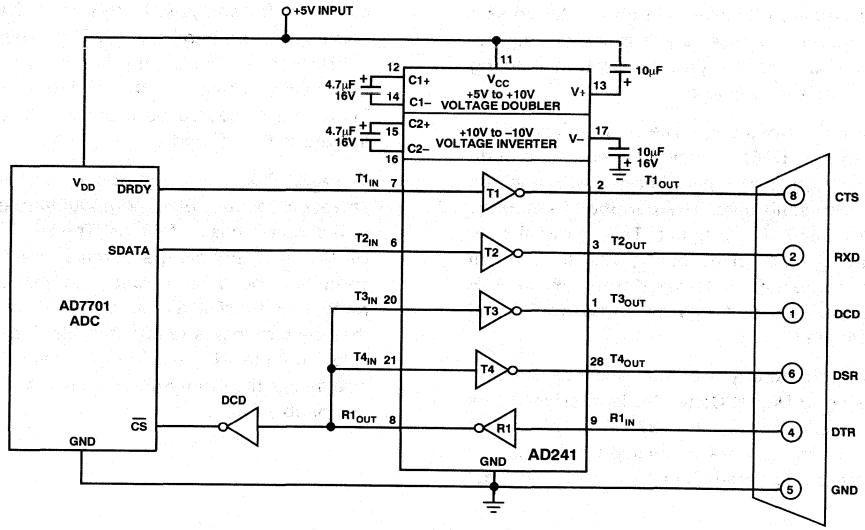

Figure 9. Remote Data Acquisition System Using the AD241 and AD7701

A simple Basic program to illustrate the data collection follows.

This program reads the input data from the serial port and displays the received binary data in HEX format on the screen. It assumes that the data is being received on COM1. The serial baud rate is set at 1200 in line 20. The next line defines a text string which sets up the data format and control line status of the COM1 serial port. The text string is then used to open a data buffer for the port into which two bytes are read. The data is read in two 8-bit bytes and these are then concatenated to form a single 16-bit result which is converted into HEX format. This is then printed on the screen. The serial port is then terminated by closing the data buffer. Note that the buffer must be opened and closed in order that the serial port control lines follow the correct sequence. The program loops back to the start and the reading sequence is repeated continuously.

For more information on the interface and on the AD7701 Analog to Digital Converter, the reader is referred to an application note "Evaluation Board for the AD7701/AD7703 Sigma Delta A/D Converter" published by Analog Devices, publication number E1483-15-1 2/90.

```
10  CLS
20  BR$ = "1200"
30  COMFIL$ = "COM1 : "+BR$+", N, 8, 2, RS, CS, DS, CD"
40  OPEN COMFIL$AS#1
50  FIELD 1, 2 AS D$
60  GET #1, 2
70  A$ = HEX$ (ASC ( MID$ (D$, 1, 1)))
80  B$ = HEX$ (ASC ( MID$ (D$, 2, 1)))
90  IF LEN (A$) = 1 THEN A$ = "0" + A$
100 IF LEN (B$) = 1 THEN B$ = "0" + B$
110 Y$ = A$ + B$
120 LOCATE 12, 20 : PRINT Y$
130 CLOSE #1
140 GOTO 40
150 END
```

ANALOG DEVICES

AN-48
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Designing Digital Repeaters with ICs

This note describes the use of the RPT-82 and the RPT-83, PCM carrier repeater integrated circuits, which perform all of the active functions required for a regenerative repeater operating at 1.544-2.048 Megabits per second (Mbps) data rates on PCM lines.

Most of the problems in digital voice transmission are caused by the transmission medium itself rather than by the transmit or receive hardware.

This is particularly true for the most commonly used medium, twisted-pair cable designed to carry analog voice-frequency signals in the 300- to 3700-Hz frequency range.

As more and more transmission routes are called upon to accommodate digital data, it becomes important, in terms of system cost and installation time, to use existing cable interchanges. As a result, audio-frequency-grade cable is called upon to transmit high-speed digital data. This was made possible through the development of digital regenerators.

A digital regenerator consists of two repeaters, one in each transmission direction. Repeaters receive digital data streams and retransmit them to the next receiver position. Current integrated-circuit repeaters eliminate the loading coils that were originally used to introduce a flat-band response over the voice band. A T1 repeater (the U.S. standard) can accept a degraded signal, regenerate it, retime the pulse stream, and transmit it over another 6000 feet of twisted-pair cable. The action taking place in the regenerative repeater stage is depicted within the dashed-line portion of Figure 1. When digital carrier repeater circuits are installed, the operating error rate for carrier systems with audio-grade cable pairs is well within the limits necessary for adequate voice communication.

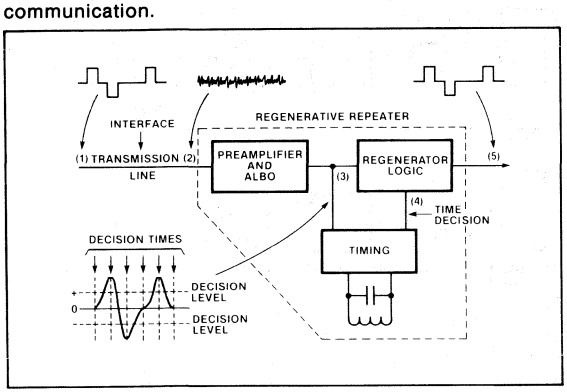

Figure 1: The functional elements of a regenerative repeater section are shown within the dashed line. The results being a re-shaping of the incoming pulses, which in turn enables a re-timed and regenerated data stream.

Several different integrated circuits are being used in the latest generation of repeater equipment. These circuits have cut the number of components required to build a digital regenerator while improving its capabilities and performance.

To grasp the concepts concerning use of the devices, it is helpful to first examine some of the basic theories relating to practical digital transmission.

The present T1 carrier uses bipolar or mark inversion digital format shown in Figure 2. There are several specific advantages provided by this coding system.

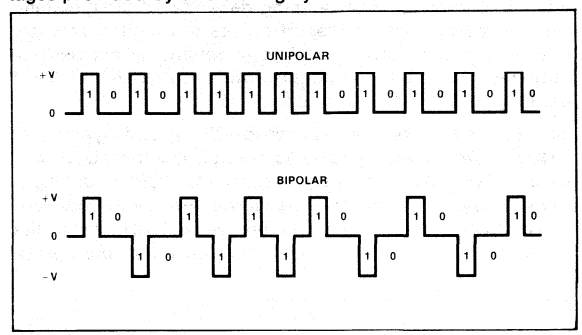

Figure 2: In the bipolar (alternate mark inversion) digital format, alternate 1 bits are inverted in polarity.

First, the maximum energy of the waveform is found at one-half of the bandwidth (data-transmission rate) of the system. For a binary code (see Figure 3), the most significant frequency is found at the zero level rather than the mid-frequency. This means that the energy spectrum encompasses double the bandwidth (3.088 MHz) of the transmission rate for the T1 rate. The bipolar scheme keeps the spectrum within the 1.544-MHz bandwidth of the transmission system.

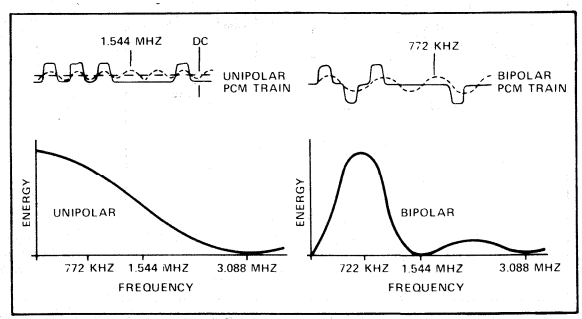

Figure 3: The (alternate mark inversion) bipolar digital format keeps the spectrum within the 1.544MHz bandwidth of the transmission system. The average dc component on the transmission pair is zero.

Another important characteristic of a bipolar code is that the average dc component on the transmission pair is zero. This is important in regenerator development because it allows the same channel to be used for data transmission and power supply. An advantage of this design approach is that extra pairs for powering a repeater location are not necessary with a simplex power arrangement.

Still another advantage provided by bipolar code is the capability to detect single-bit errors within the transmitted data. Since the sequence of pulses has alternating polarity, the absence of a pulse will be received as a bipolar violation. Though error bursts could produce incorrect error counts, it has been purported that the difference between the true error rate and measured error rate is negligible in high-error channels.

If the error rate is low enough to show differences for burst-error occurrences versus single-error counts, the channel will probably be performing well enough to make the difference unimportant.

There are also some problems associated with bipolar coding. For example, since retiming depends on the incoming pulse stream, a long series of zeroes will have a very low energy content. Some oscillator-resonant circuit designs may then damp out completely and fail to restart. This will depend on the Q of the circuit.

Some carrier designs have moved to different coding schemes to restrict the number of zeroes possible in a legal transmission. In European systems especially, the HDB3 code (high-density bipolar with no more than three zeroes in succession) has often been incorporated. In this discussion, however, the straight bipolar coding scheme will be considered when determining the relationship of the active components within the repeater.

The functional components of a repeater integrated circuit are illustrated in Figure 4. The repeater amplifies the incoming signal and equalizes the received signal to compensate for attenuation distortion and phase distortion. The analog amplifier also has a feedback loop that provides for a variable amplification dependent on the threshold detected for the incoming pulses (either positive or negative). Either of two methods can be used to obtain equalization: in the first approach, a fixed value of cable attenuation is provided; in the second approach, automatic equalization can be used with constant modification taking place to handle variable line lengths in that transmission section. The monolithic approach provides an equalization network and variable line matching circuitry.

The repeater must provide a timing-recovery circuit to extract the data transmission clock from the incoming pulse train. The recovery circuit will normally comprise a resonant circuit tuned to the peak energy level of the incoming signal. The resonant circuit is "energized" by the transitions at the received channel with the resonant frequency dependent on the Q of the circuit. A higher Q can allow the circuit to remain active through longer periods of zero signals levels. The recovered timing pattern is then used in the reproduction of the incoming series of pulses.

The repeater must be capable of differentiating between a data pulse and channel noise. The noise due to adjacent channel pairs or interference due to signals in adjacent time slots must not negate the reception of the proper incoming pulse sequence. Jitter due to timing variations from repeater section to repeater section can disrupt the signal detection

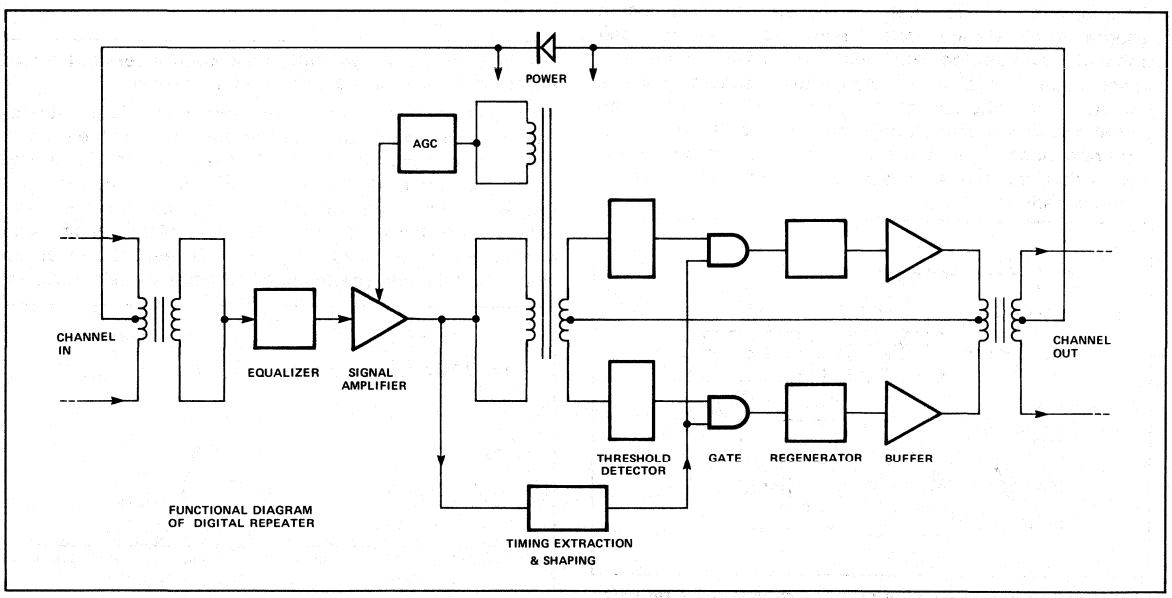

Figure 4: Functional diagram of a repeater integrated circuit shows the equalizer (to compensate for attenuation and phase distortion) and analog amplifier with feedback to provide variable amplification dependent on the detected threshold.

by causing pulse thresholds to be measured at the improper intervals. Pulse detection is accomplished by the detection of proper threshold levels and correct timing intervals. The repeater then will retime the detected levels to provide for a new pulse train. It is important when retiming the waveform that the input from the threshold detector precede the timing strobe. This ensures that the output pulse will follow the reconstructed timing, reducing the jitter components.

The final functional block necessary for the complete repeater function involves the regeneration of the output pulses. The buffer stages drive the output transformer to produce a posi-

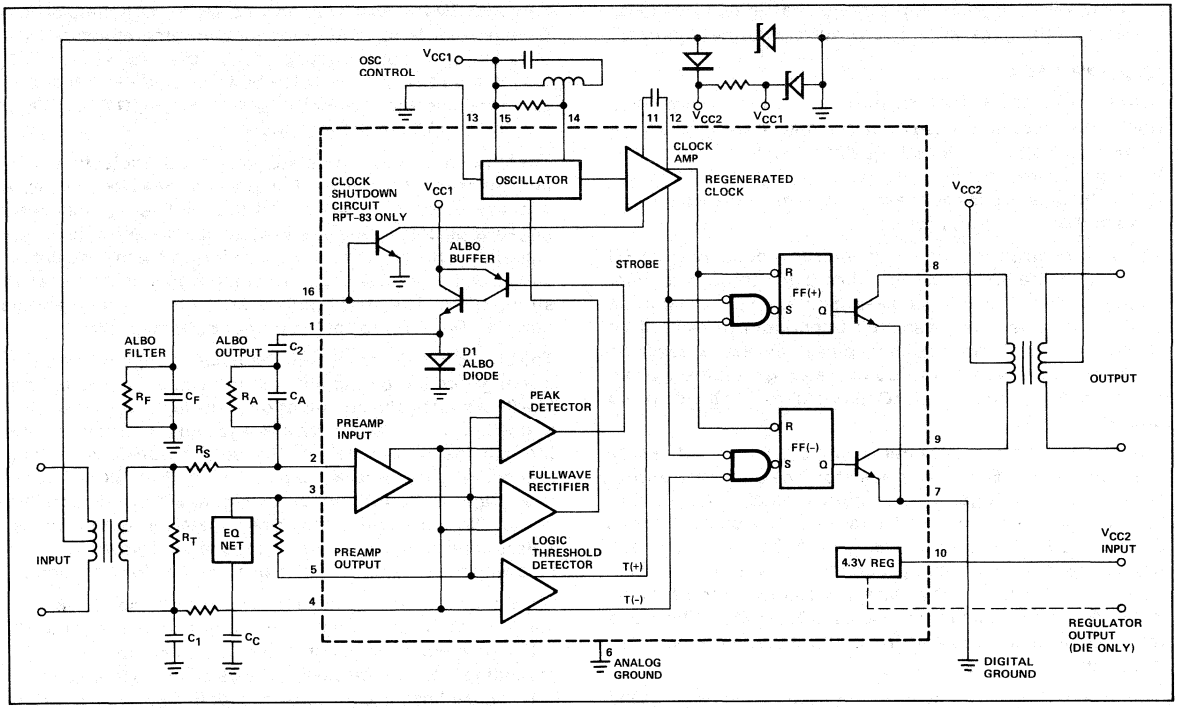

Figure 5: In ADI's RPT-82 and RPT-83, equalization and amplification of the incoming pulse train takes place through an amplifier with an external network controlled by feedback.

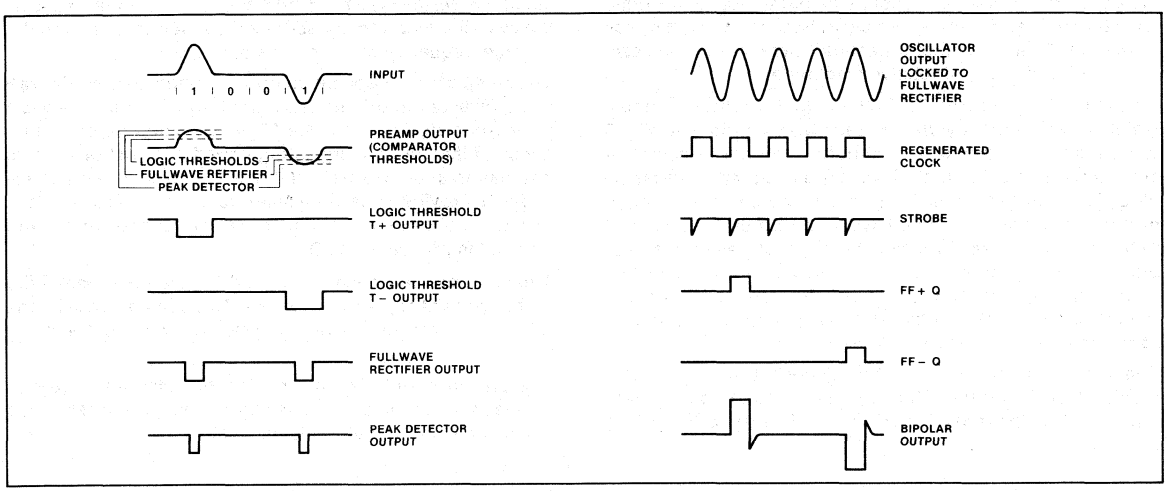

Figure 5A: Voltage waveforms at the outputs of the various functional blocks within the RPT-83 integrated circuit. The output pulses of the amplifier trip some combination of the level-detecting comparators.

tive or negative signal on the channel output pair. The design is such that, in most instances (U.S. and European), the generated signals are nominally square pulses. Different theories as to improving performance by band-limiting the output stage have been discussed and, in some instances, implemented. However, the majority opinion still seems to favor the square-wave components.

The repeater can be considered a combination of discrete yet interrelated functions. The overall function of the repeater is the shaping of incoming pulses, the retiming of these pulses, and the generation of an equivalent output stream.

IC APPROACH

In PMI's RPT-82 and RPT-83, the equalization and amplification of the incoming pulse train is achieved through a bipolar amplifier with an external equalization network controlled by feedback from the amplifier outputs (Figure 5). The result is that the pulses peak so the majority of the pulse amplitude is restricted to its own time slot.

The automatic attenuation of the pulses is done by an ALBO (automatic line build-out) circuit actuated through a buffer circuit on chip. The output pulses of the amplifier will trip one or all of the three level-detecting comparators. In terms of the ALBO feature, if the incoming pulses exceed a specified peak-reference value, the comparator generates current pulses that flow into the ALBO filter external to the chip and charge capacitor C_F.

The voltage at the ALBO filter is applied, in turn, to the base of the internal ALBO buffer transistor. The buffer will turn on when its base voltage exceeds approximately 1.7 volts. This 'ON voltage' is also specified in the electrical characteristics of both the RPT-82 and RPT-83. The ALBO diode acts as a series impedance, normally much larger, and therefore the dominant factor, than the shunt impedance of the external network. As the current increases at the buffer, the diode impedance decreases and the shunt impedance becomes dominant.

The overall effect is to increase the ALBO attenuation as the filter voltage level increases. The result is that the feedback loop will adjust itself until the pulses out of the preamplifier are equal to the fixed reference. The references for the other level detector and rectifier are set at fixed ratios of this peak reference. Therefore, their thresholds are fixed with respect to the pulse shape and relative amplitude.

On the RPT-83 (not the RPT-82), a clock shutdown circuit, which is activated by the current levels from the ALBO buffer, is also provided. This shutdown circuit turns off the clock amplifier at low input levels, thus neither the regenerated clock nor the strobe outputs are enabled at the output buffer flip-flops. The circuit was incorporated to prevent random-noise pulses from being reproduced as valid output pulses. The problem, however, can be that at long line lengths the correct signals are too low in energy to activate the clock circuit. Presently, PMI has customers using both schemes in their repeater designs.

The timing of the circuit is based upon the pulsing of a resonant tank circuit. The full-wave-rectifier comparator output pulses the tank network. The pulses try to "force" the oscillator circuit to phase lock to the incoming pulse waveform, while the tank attempts to resonate at its preset (with external components) frequency. The result of the two factors is a clock circuit that runs at an average bit rate for the incoming waveforms.

Again, the value of Q for the tank circuit will help determine the oscillator accuracy. If the Q is high, the resonant frequency dominates and it is more difficult to phase-lock to the incoming pulse rate. A high Q circuit also changes considerably with the temperature and long-term component drift. A Q value that is too low will cause adverse affects on the oscillator circuit by the jitter present in the pulse stream. This means the jitter will then also be transferred to the output pulse train as well. For both the RPT-82 and RPT-83, Qs of greater than 75 are recommended.

The logic-threshold comparator provides the detection function for incoming pulses. For positive received pulses, a negative pulse is generated on the T+ line; for incoming negative pulses, a pulse is sent on the T– line. The clock amplifier "squares" the timing waveform from the oscillation circuit and produces a square clock signal and a negative strobe pulse. The strobe pulses are 'ANDed' with the logic outputs (T+ or T–) to set the output buffer flip-flops.

The strobe is coincident with the positive edge of the clock signal generated by the tank circuit. Once the appropriate flip-flop is set by the combination of the logic output and the strobe pulse, the output driver stage causes current to flow through the proper half of the output transformer, thus regenerating the received bipolar pulse. The falling (negative) edge of the internal clock signal serves to reset the output flip-flops and thereby terminate the output pulse. This is important to prevent the regenerated waveform from following the data threshold instead of the clocking circuit.

One option is made available to users of both the RPT-82 and the RPT-83. The internal clock oscillator can operate in either an injection-locked mode or a pulsed-tank mode. By grounding a pin on the device, the oscillator is free-running but is phase-locked to the full-wave rectified output. With the pin open, the oscillator will only operate when pulsed by an incoming signal. The circuit then "rings" until the next incoming pulse is received from the rectifier. In both cases, the overall effect is to phase-lock the resonant circuit to the average frequency of the pulse train.

The completed T1 repeater may look something like Figure 6 in a typical configuration. The input transformer provides a two-to-one step-up from the 100-ohm characteristic impedance of the line. The matching impedance, in this case, is doubled to approximately 400 ohms. A fixed attenuation is added to provide a fixed line build-out of approximately 6dB. The automatic build-out then provides a varying attenuation for line lengths up to 36dB.

The equalization network is added to give the preamplifier higher gains at higher frequencies. This compensates for the roll-off characteristics of the preamp and the transmission line.

The oscillator tank circuit provides the resonant frequency for the oscillator. It is controlled by current pulses generated as the incoming waveform is received.

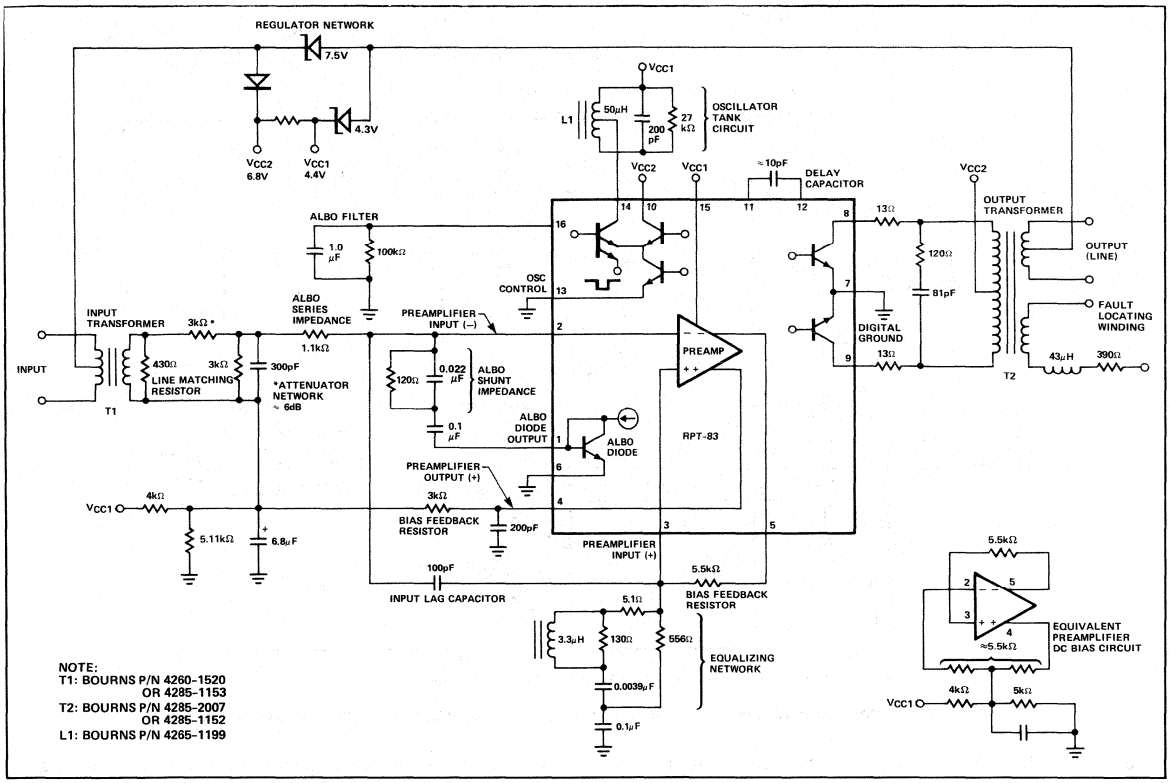

Figure 6: The complete T1 repeater system (1.544MHz) using the RPT-83 integrated circuit. The input transformer provides a 2:1 step-up from the 100 ohm line impedance.

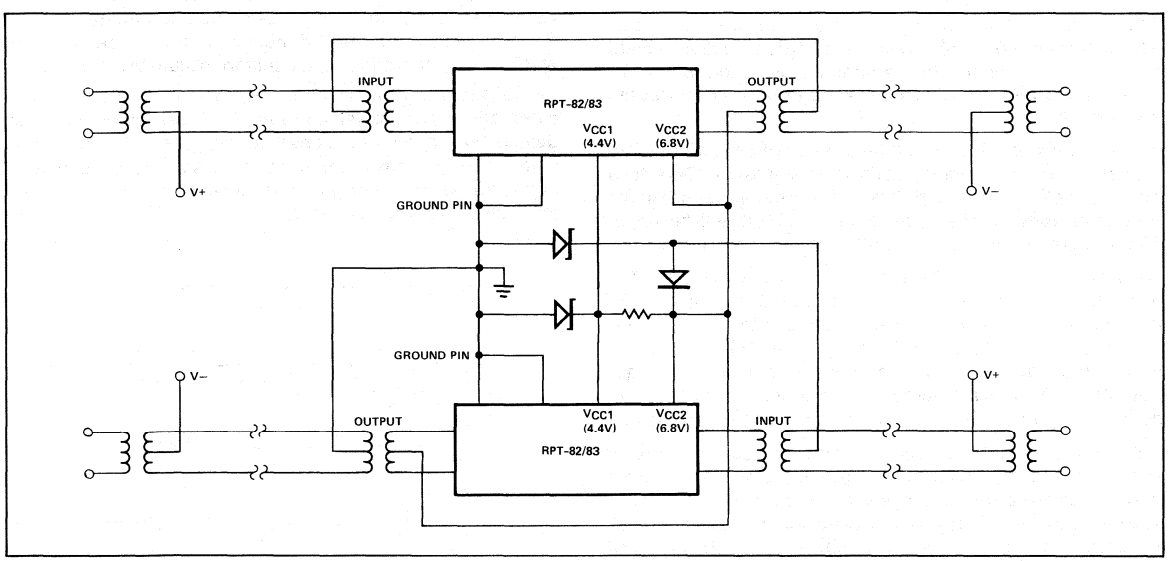

Figure 7: The simplex power supply design consists of two zener diodes and a diode float connected to the center taps of the line transformers. Nominal voltages are 4.4 and 6.8V.

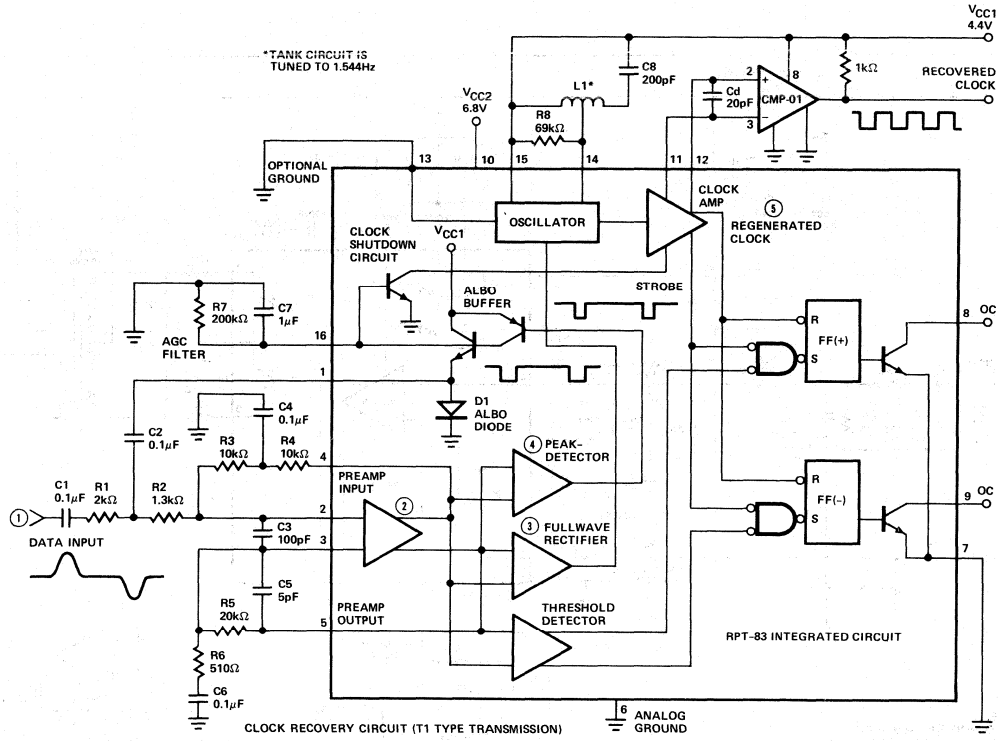

Figure 8: Integrated-circuit repeaters can also be used in clock-recovery circuits, as shown.

The ALBO filter aids in integrating the pulse output of the detector. As the voltage increases, more current flows through the ALBO diode, and the line build-out is characteristic of longer line lengths.

The power-supply current is available over the signal pair. The simplex power design consists of two zener diodes and a diode float connected to the center taps of the line transformers (Figure 7). The nominal voltage values required are 4.4 and 6.8V.

This design meets the specifications called for in a T1 carrier repeater. The integrated circuit uses less than 13mA total current (both voltage supplies). This means the maximum output current for the total repeater circuit will be under 50mA at worst-case conditions (all ones output signal).

A final application of the integrated circuits is the use of the repeater device in a clock recovery circuit. Data transmission is becoming more important in areas other than long-distance digitized audio. In these instances, the capability of recovering clocking from the data stream can be advantageous. This can mean single-pair connections that can transfer data without additional wiring for timing and clock signals.

Using the RPT-83 in conjunction with a precision comparator, such as the PMI CMP-01, will provide a recovery scheme capable of reproducing a clock waveform from input levels as low as 35mV peak-to-peak. Any system that requires clock retrieval from a data signal and synchronization to that signal can use a circuit similar to this design shown in Figure 8.

The test circuit was operated with an AMI incoming code at frequencies from 64kHz up to 1.544MHz. In the published design, the incoming data waveform is capacitively coupled to an input attenuator using fixed external components and the internal ALBO diode. Since the impedance at pin 1 of the device varies inversely with the amplitude of the input signal, the voltage at the preamplifier input (through resistor R2) will be held to less than 100mV peak-to-peak amplitude. This gives the design an input dynamics range of greater than 45dB while still producing a constant output waveform.

In the circuit shown in Figure 8, tests show none of the external component values to have critical tolerances. The design can be used to recover clocking from an incoming data stream—again a capability that has proven advantageous for designers of data interfaces in many areas other than telecom carrier exchanges.

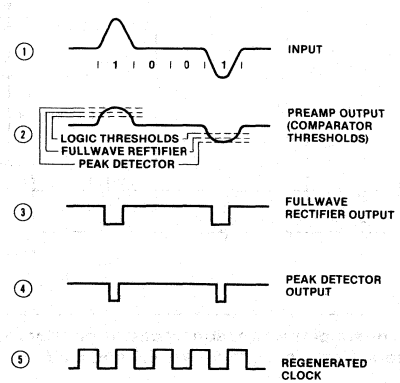

Figure 8A: Show the voltage waveforms within the circuit.

ANALOG DEVICES

AN-50
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

A Variable-Frequency Clock-Recovery Circuit Using the RPT-82 or RPT-83

This note describes a high-performance clock-recovery circuit, employing an RPT-83 repeater IC and a CMP-01 IC comparator, which helps derive a clock pulse from, and synchronizes it with, an incoming data signal. The circuit accepts a low-level bipolar pulse train (35mV peak-to-peak, minimum) while producing a usable recovered clock signal. Since the circuit also accepts high-level inputs (10V peak-to-peak, minimum), it can attain a 49-dB dynamic range.

The RPT-83 chip is conventionally connected, except that input and output transformers are not required and the (internal) output transistors are left as an open circuit (Figure 1). The recovered clock signal is picked off the RPT-83's clock amplifier (pins 11 and 12) by the CMP-01 comparator.

The comparator's output is basically a square wave at the data rate frequency (which, for this example, is the T1 transmission frequency of 1.544 Mbits/s).

Clock recovery can be best understood by referring to a more detailed diagram of the RPT-83's oscillator section and external tuned circuit (Figure 2). Grounding pin 13 creates a "locked oscillator" operating mode. Floating pin 13 creates a "pulsed tank" operating mode.

In the pulsed-tank mode, the external tuned circuit is stimulated each time the full-wave rectifier demands a current pulse. Between pulses, the tank circuit "rings" at its resonant frequency, damped only by R8 and the on-chip transistor collector resistor R.

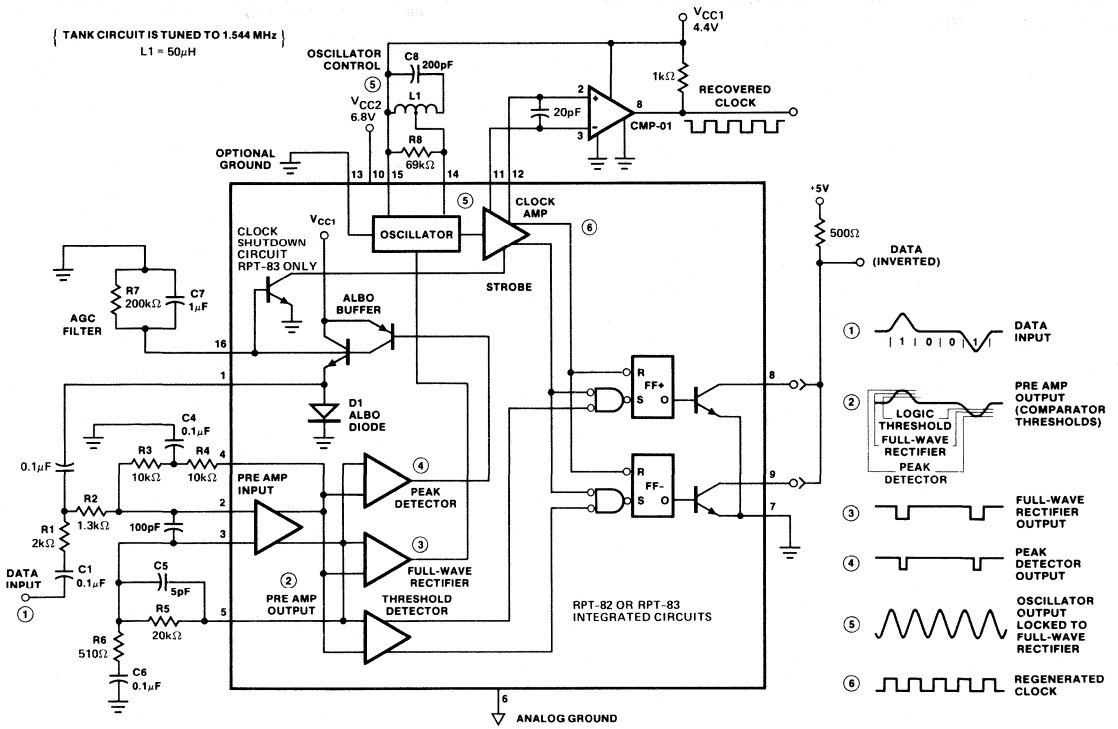

Figure 1: The clock-recovery circuit shown above will accept a 35mV peak-to-peak (minimum) bipolar pulse train and produce recovered clock pulses at the data rate frequency (1.544Mbits/s for this example). Key waveforms, shown in correct timing relationship, illustrate the circuit's operation. Comparator CMP-01 provides amplification and offset for a good TTL interface.

In the locked-oscillator mode, the oscillator runs continuously. It is injection-locked to the full-wave rectifier pulses by the second emitter of on-chip transistor Q. The oscillator circuit drives a buffer emitter follower, which in turn drives a clock amplifier. The clock amplifier has a differential output (pins 11 and 12) and squares the sinusoidal oscillations originating at pin 14.

Since the clock amplitude at pins 11 and 12 is only 1Vp-p around a common-mode dc level of 4.5V, the CMP-01 comparator provides a good interface to TTL or other logic. A small capacitor, C_4, may be connected across pins 11 and 12 to delay the clock relative to the data.

The circuit shown in Figure 1 does not make use of the data outputs of the repeater IC. However, if a data waveform is required in addition to the clocking pattern, the change in circuit configuration required is minor. By tieing the output leads of the two flip-flops together (pins 8 and 9) and adding a 500ohm pull-up resistor to +5V, an inverted data waveform is made available.

How "pulsed-tank" and "locked-oscillator" modes are selected is shown in Figure 2. Pickoff points for a regenerated clock signal (within the clock amplifier) are also evident in the diagram.

The circuit described in this note was designed and tested only with alternate-mark-inversion (AMI) codes. However, the system should work equally well with any return-to-zero code and would suit any of the many AMI-code variants. It can handle data rates up to 2.0 Mbits/s.

Designs have been tested at frequencies from 64kHz up to 1.544MHz by simply modifying the tank circuit. In all cases, performance met desired expectations.

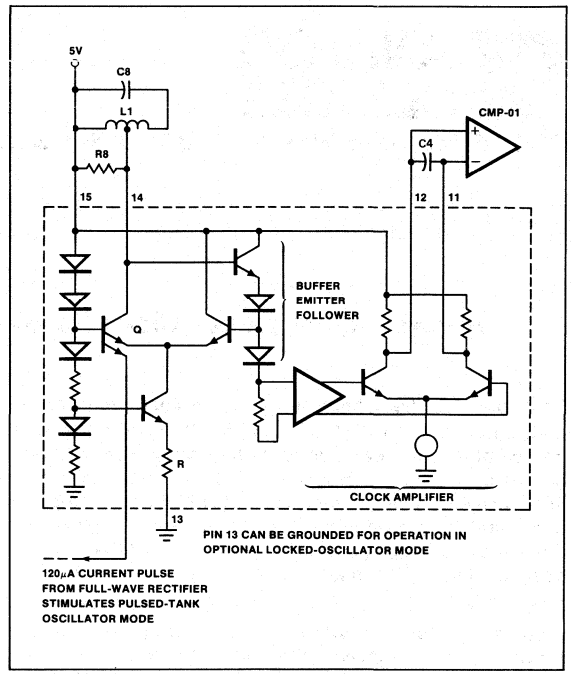

Figure 2: Detailed view of the RPT-83's oscillator and clock amplifier shows how "pulsed-tank" and "locked-oscillator" modes are selected. Pickoff points for a regenerated clock signal (within the clock amplifier) are also evident.

ANALOG DEVICES

AN-119
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Improved T148 (CEPT) PCM Repeater for #22AWG (0.7mm) Unshielded Twisted-Pair Wire

by Mike Jachowski

With only a few external components, the RPT-86/RPT-87 performs as a complete, high-speed PCM repeater. Configured as shown in Figure 1, it is capable of recovering and retransmitting AMI-coded data with an input amplitude dynamic range of from 0dB to –44dB at a rate of 2.048Mbit/s. This translates to a recovery distance of from 0 ft. to over 9000 ft. (2.7km) on #22AWG (0.7mm) unshielded, twisted-pair transmission line, 16pF/ft.

The RPT-86/87's dual ALBO diodes enable the circuit to automatically compensate for changes in signal amplitude by completing an AGC loop comprised of the diodes, preamplifier, peak threshold detectors, and the equalization network impedance. The diodes are variable impedance elements driven by the peak threshold detectors. They form divider networks with R_4 and R_5 capable of providing varying amounts of attenuation to the signal as diode impedance changes. If the signal grows in amplitude, the ALBO diodes become a lower impedance, thereby attenuating the signal before it enters the preamp. Likewise, if the signal is small, the diodes become high impedance allowing more signal to pass.

Being capacitively coupled, the impedance of the diodes can also change the filter response of the network to compensate for the frequency distortion imparted to the signal by the transmission line. The diodes control two poles which can be taken in or out of the network depending on diode impedance. The network also contains two zeroes, R_2C_1 and L_2C_5. For receiving on long transmission lines, the ALBO diodes remove the two poles from the network leaving an overall two zero network response. On short transmission lines, the diodes activate the poles, cancelling the zeroes and leaving the network with a flat frequency response.

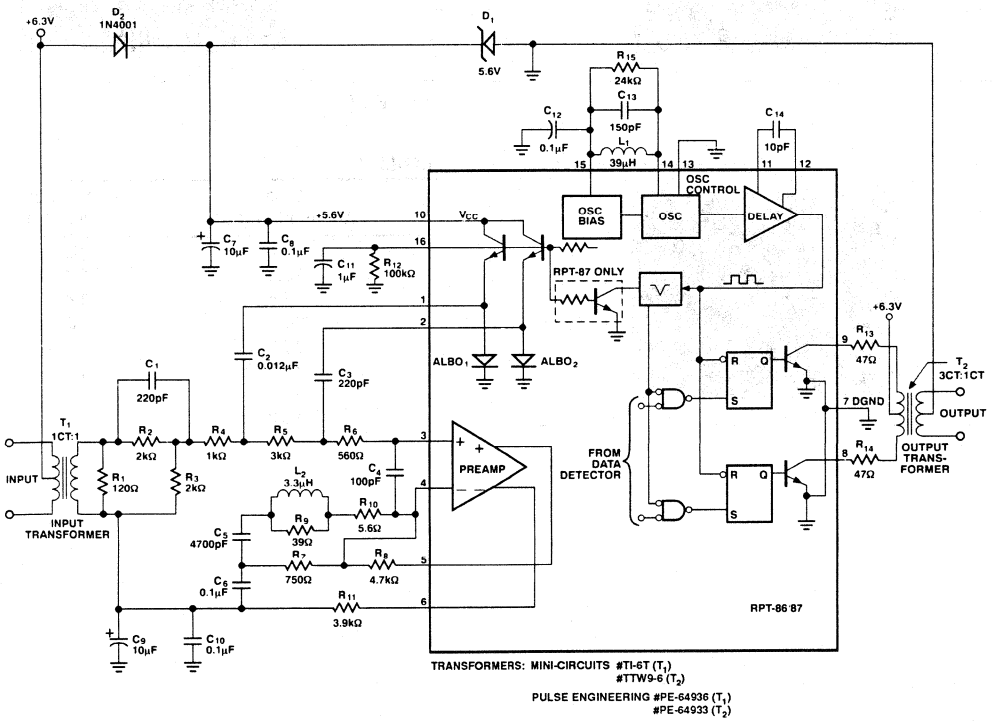

FIGURE 1: *A Complete 2.048Mbit/s Repeater Circuit Optimized for #22AWG (0.7mm) Unshielded Twisted-Pair Wire Transmission Line, 16pF/ft.*

In a typical 2.048Mbit/s repeater card application, two repeater ICs are used to transfer data in a bidirectional fashion using two twisted-pair lines. Power for the repeater card is derived from a 48mA direct current feed in one of the two lines. The line current passes through a 5.6V zener diode, D_1, establishing the supply voltage to the RPT-86/87. D_2 is added to create a temperature stable 6.3V supply for the output line drivers, stabilizing the magnitude of the transmitted data pulses over temperature variations. Furthermore, supplying the output stage with 6.3V allows the use of a center-tapped 3:1 drive transformer reducing the worst-case output drive current consumption to only 10mA per line. In this way, total repeater card quiescent current is a maximum of only 43mA with all ones transmitted in both directions and 0 ft. line length (about 2mA total ALBO current).

Figures 2, 3, 4 and 5 show typical RPT-86/87 waveforms in a 2.048Mbit/s application repeating data at the end of 6600 ft. (2km) of unshielded, twisted-pair wire. Figure 2 illustrates the $6V_{p-p}$ AMI signal, (a), as it enters the transmission line and as it appears after travelling through 6600 ft. of wire, (b). Figure 3 is the noninverting preamp output on pin #5 exhibiting proper data pulse reconstruction by the preamp and ALBO diode equalization network. Figure 4 shows the "eye-pattern" on pin #5 with a QRSS input signal demonstrating the repeater operating with a large signal-to-noise ratio after 6600 ft. (2km) of wire. The circuit's signal-to-noise tolerance, or interference margin, at several line attenuation levels is listed in Table 1. This test was made using a Sierra 415A-2 modified for 2.048Mbit/s testing. Figure 5 shows the retransmitted AMI data pulses created by the RPT-86/87's open-collector output line drivers as they pull on the opposing halves of the output transformer's center-tapped primary winding. The pulse lies within the normalized limits of the pulse mask recommended by ATT/CCITT specifications.

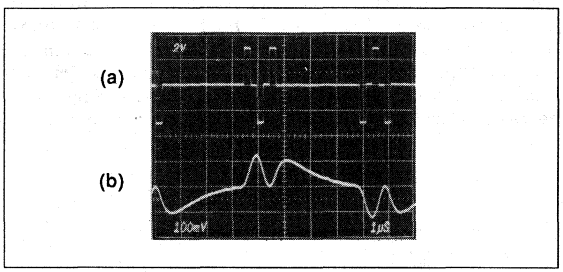

FIGURE 2

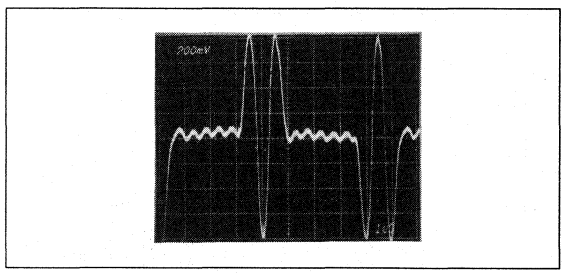

FIGURE 3

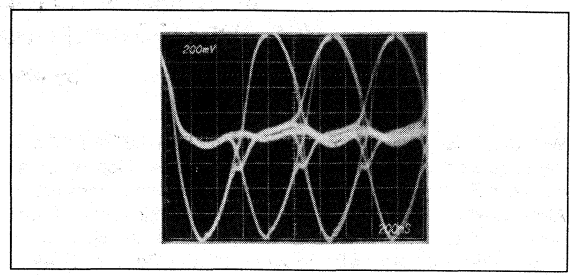

FIGURE 4

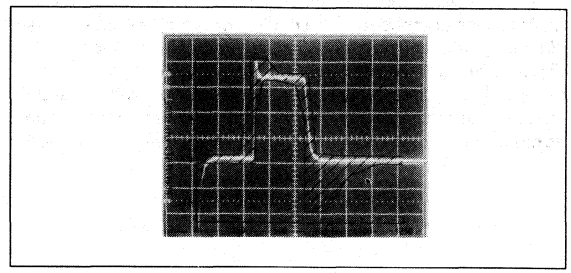

FIGURE 5

TABLE 1

Line Attenuation (dB)	Interference Margin (dB)
6.0	12
10.5	12
18.5	12
28.0	10
36.0	12
38.5	12
42.0	13
44.0	14

ANALOG DEVICES

AN-303
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Understanding and Applying the AD7341/AD7371 Switched Capacitor Filters

by John Reidy and Mike Curtin

INTRODUCTION

The AD7341 and AD7371 are high performance switched capacitor filters designed and specified for use in modems. The AD7341 is a reconstruction filter for the transmit channel and is used after a reconstruction D/A converter. It implements the filter function using a seventh order low pass switched capacitor filter and a second order low pass continuous time filter. For the specified CLKIN of 288 kHz, the cutoff frequency is 3.5 kHz. The AD7371 is an antialiasing filter for the A/D converter in the receive channel. It is made up of a second order low pass continuous time filter, a fourth order high pass switched capacitor section and a seventh order low pass switched capacitor filter. With the specified CLKIN frequency of 288 kHz, the lower cutoff frequency is 180 Hz and the upper cutoff frequency is 3.5 kHz. A complete description and specifications for both filters is available in the AD7341/AD7371 data sheet (Publication Number C1262-10-1/89), available from Analog Devices.

This application note explains some theory necessary to understand Switched Capacitor Filters (SCFs). It then covers two related aspects of the filters' operation. First, there is a look at the differences between synchronous and nonsynchronous operation from the accompanying A/D or D/A converter. Nonsynchronous operation means that the update rate of the DAC and the sampling rate of the ADC are not synchronous with the clock running the switched capacitor filters. This can cause a degradation in performance unless additional filtering is employed.

A feature of the filters which is not normally utilized is the ability to vary the cutoff frequency by adjusting the input clock. The specified cutoff frequency of 3.5 kHz applies with an input clock of 288 kHz. However, applications other than modems may require alternative cutoff frequencies which can be achieved by varying the input clock. When this is done over a wide frequency range it is necessary to take into account the on-chip continuous-time filters.

SAMPLING THEORY

Signal sampling is inherent in the filtering technique of switch capacitor technology. Consequently, it is helpful to review some theory on the practical aspects of sampling to understand switched capacitor filter operation. Equation 1 gives an expression for an arbitrary continuous time signal, g(t), that is sampled at time intervals of T_S:

$$s(t) = \sum_{n=-\infty}^{\infty} g(nT_S)\, h(t - nT_S) \quad \ldots\ldots\ldots\ldots\ldots\ldots (1)$$

where:

s(t) = Sampled Waveform

g(t) = Continuous Time Signal

$$h(t) = \begin{cases} 1, & 0 < t < T_S \\ 1/2, & t=0,\ t=T_S \\ 0, & \text{Otherwise} \end{cases}$$

T_S = Sampling Period

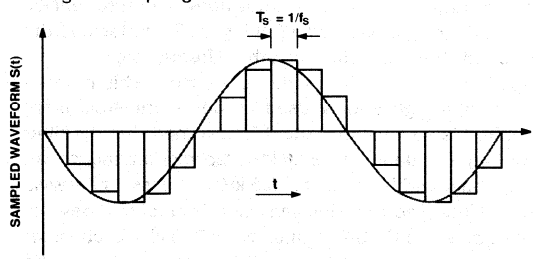

Figure 1. Time Domain Representation of a Sampled Waveform

The sampling process is depicted in Figure 1 for a single-tone sine wave. The signal waveform consists of a sequence of pulses which have a finite duration equal to T_S, the sampling period.

A corresponding expression for the frequency domain representation is given in Equation 2.

$$S(f) = \sum_{n=-\infty}^{\infty} G(f - nf_S)\, Sinc\,(\pi f/f_S) \quad \ldots\ldots\ldots\ldots\ldots\ldots (2)$$

where:

G(f) = Frequency Transform of g(t)

$Sinc(\pi f/f_S) = Sin(\pi f/f_S)/(\pi f/f_S)$ = Frequency Transform of h(t)

The frequency domain representation is the product of two functions. The first, $\Sigma\ G(f - nf_S)$, contains the original analog Fourier frequency spectrum, G(f), which is repeated at multiples of f_S. The term $Sinc(\pi f/f_S)$ is an attenuation profile due to the finite duration of the time domain pulses. Ideally, the function h(t) should be a series of impulse functions, in which case, the $Sinc(\pi f/f_S)$ term would be replaced by 1. In practice, however, the time domain pulses have a finite width resulting in a low-pass filtering type response. The resulting frequency domain profile for a single sine wave frequency is shown in Figure 2.

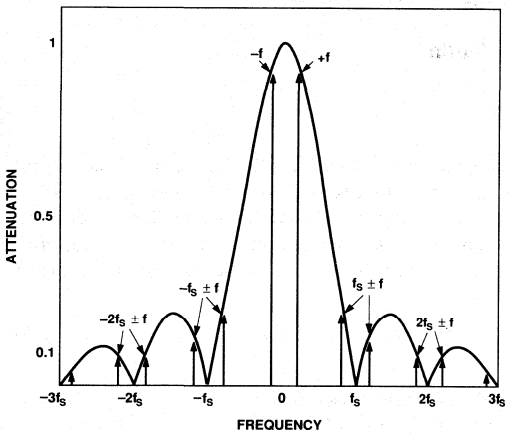

Figure 2. Frequency Domain Representation of the Sampled Signal, s(t), Shown in Figure 1

At this point it is worth mentioning some terminology used throughout this application note. The frequencies nearest the origin, labeled ±f in Figure 2, are termed the baseband frequencies. All other frequencies, $\pm f_S \pm f$, $\pm 2f_S \pm f$, etc., are called image frequencies. This nomenclature, although widely used, is not a universal standard. In this application note it is assumed that the filter sampling frequency is 57.6 kHz, corresponding to an upper bandwidth limit of 3.5 kHz, unless otherwise stated. Note also that this sampling frequency does not correspond to a CLKIN frequency of 57.6 kHz because of internal dividers. For further information consult the data sheet.

Figures 3 and 4 show simplified block diagrams which outline the filtering stages for the AD7341 and AD7371. Both filters have a second order continuous time filter in series with the signal path. It is located at the input for the AD7371 and at the output for the AD7341. Consequently, it contributes differently to the overall transfer function for the two filters. The AD7341 is designed as a reconstruction filter. Here, the continuous time section smooths the filter output so that it appears as an unsampled signal in the time domain. In the frequency domain this translates to attenuating the image frequencies. The AD7371 is an antialiasing filter. Here, the function of the continuous time section is to attenuate any frequencies outside the Nyquist limit to avoid aliasing in the switched capacitor filter itself. The output of the AD7371 is a sampled output with unfiltered image frequencies.

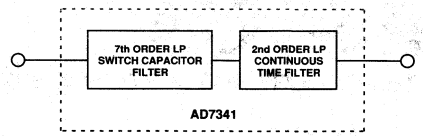

Figure 3. AD7341 Filtering Stages

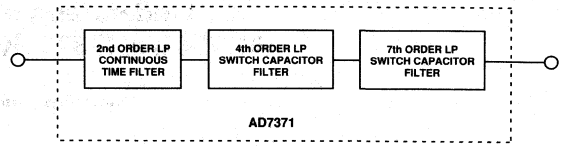

Figure 4. AD7371 Filtering Stages

The continuous time filter has to pass signal frequencies which are less than f_{MAX} (i.e., frequencies less than 3.5 kHz) and attenuate any frequencies above $f_S - f_{MAX}$, (54.1 kHz). There are almost four octaves between these frequencies allowing a large transition band for the filter roll-off. The continuous time filter response is shown in Figure 5. Examination of the plot shows that the cutoff frequency is not exactly at 3.5 kHz. The actual cutoff is at 4.76 kHz to allow some design margin for process variations. Though the filter is only second order, frequencies above 54.1 kHz are attenuated by 30 dB.

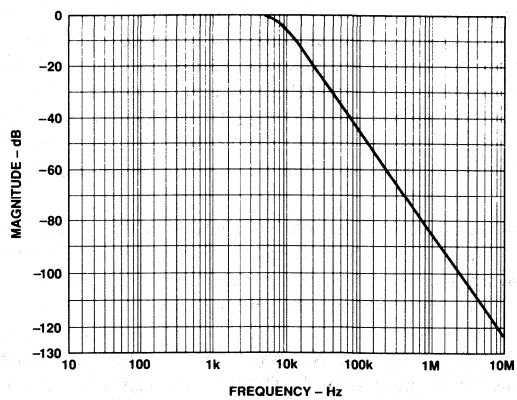

Figure 5. AD7341/7371 Continuous Time Filter Response

ACTUAL FILTER OUTPUT SPECTRUM

The next step is to apply the previous information to predict the filter output for a given input frequency. As an example, Figure 6 shows the actual output spectrum of the AD7371 when an input frequency of 1 kHz is applied to the input. As expected, the output spectrum consists of the input frequency of 1 kHz and image frequencies at multiples of f_S as described by Equation 2. From Equation 2, the magnitude of the image frequencies in dB relative to the fundamental is given by:

$20\ log\ [sin\ (\pi f/f_S)/(\pi f/f_S)]$

For signal frequencies which are much less than f_S, the magnitude of the base band frequency can be approximated to 0 dB. The magnitude of the image frequencies can be approximated to:

$$20 \log [f/(nf_S \pm f)]$$

where $nf_S \pm f$ represents the image frequency pairs that appear at each multiple of f_S. An alternative representation is given as:

$$20 \log \left[\frac{\text{Baseband Frequency}}{\text{Image Frequency}}\right]$$

Table I below gives the calculated values for the first three image frequency pairs. These values correlate very closely with the actual values shown in Figure 6.

n	Image Frequency (kHz)	Magnitude (dB)
1	56.6	35
1	58.6	35.3
2	114.2	41.2
2	116.2	41.3
3	171.8	44.7
3	173.8	44.8

Table I. Computed Magnitudes of the First Three Image Frequency Pairs in the AD7371 Output, for an Input Frequency of 1 kHz

Also present in the spectrum of Figure 6 are frequencies at odd multiples of $f_S/2$, i.e., 28.8 kHz, 86.4 kHz, 144 kHz. The presence of these components is due to the filter design and is not typical of sampling systems. They are directly related to the filter clock frequency and therefore do not vary with input signal frequency.

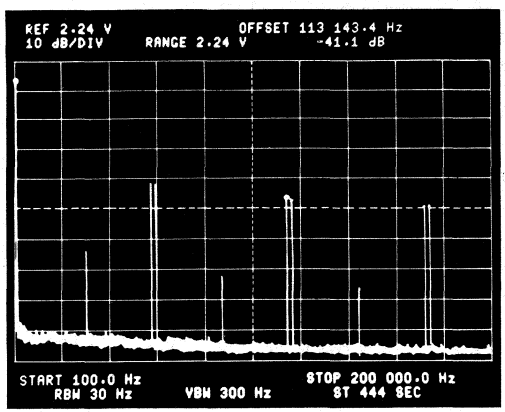

Figure 6. AD7371 Output Spectrum for an Input Frequency of 1 kHz

Other frequency components in the frequency spectrum are second and third order harmonics of the input frequency. These components are approximately 85 dB below the fundamental and therefore are too small to affect system performance. All other components are noise.

Figure 7 shows the output frequency spectrum of the AD7341 when a single tone sine wave is applied to the input. The attenuation of the continuous time filter results in smaller image frequencies and clock feedthrough components compared to the AD7371. The resulting magnitude of these images can be obtained by combining the frequency response of Figure 5 with the magnitude levels given in Table I. The resulting magnitudes of the first three image frequencies are shown in Table II.

n	Image Frequency (kHz)	Magnitude (dB)
1	56.6	65
1	58.6	65.3
2	114.2	83
2	116.2	83
3	171.8	94
3	173.8	94

Table II. Computed Magnitudes of the First Three Image Frequency Pairs in the AD7341 Output, for an Input Frequency of 1 kHz

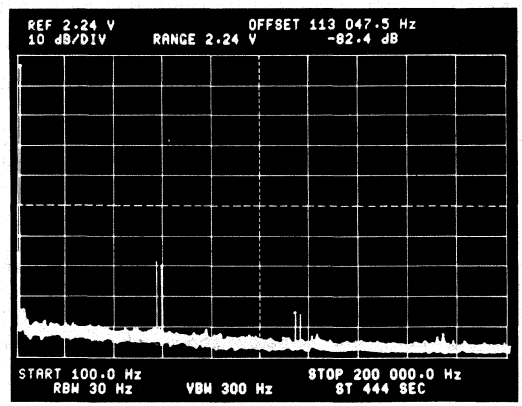

Figure 7. AD7341 Output Spectrum for an Input Frequency of 1 kHz

SYNCHRONOUS vs. NONSYNCHRONOUS SAMPLING

Interfacing the AD7371 to an ADC can cause aliasing problems because of the image frequencies and clock feedthrough components shown in Figure 5. One method of avoiding these difficulties is to use synchronous sampling as shown by the functional diagram of Figure 8. Synchronous sampling means that the ADC's $\overline{\text{CONVST}}$ input and the filter's update frequency are derived from the same source. In other words, the filter's update frequency is an integer multiple of the ADC's sampling frequency. The AD7371 has an on-chip clock divider which provides a $\overline{\text{SYNCOUT}}$ output, designed to give a variety of frequencies for sampling control, all of which are synchronous with the SCF clock.

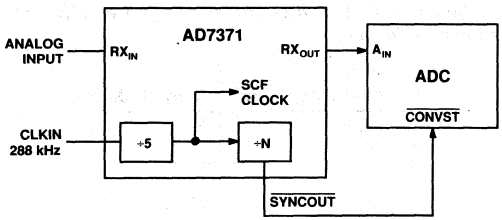

Figure 8. Synchronous Sampling Block Diagram

In synchronous interfacing, the image frequencies from the filter output map into the fundamental frequency position of the ADC's digital spectrum. This is easily verified by taking an example and comparing the aliased positions of the image frequencies for both a synchronous and a nonsynchronous ADC sampling frequency. Going back to the previous example, the switched capacitor frequency is 57.6 kHz and the signal frequency is 1 kHz. To complete the set up, consider a synchronous ADC sampling frequency of 57.6 kHz/8 = 7.2 kHz and a nonsynchronous sampling frequency of 8 kHz. The next step is to establish where the filter's image frequencies map into the ADC's digital spectrum. The digital frequency positions may be found by subtracting multiples of the ADC sampling rate from each of the image frequencies until they lie in the range ±1/2 of the ADC sampling rate. Table III below shows where the first three image frequency pairs from the switched capacitor filter map into the digital frequency spectrum.

Image Frequency (kHz)	ADC's Digital Frequency (kHz) f_{ADC} = 8 kHz	ADC's Digital Frequency (kHz) f_{ADC} = 7.2 kHz
56.6	0.6	1
58.6	2.6	1
114.2	2.2	1
116.2	3.8	1
171.8	3.8	1
173.8	2.2	1

Table III. Image Frequencies and Their Equivalent Frequency Positions in the ADC's Digital Spectrum. Signal Frequency = 1 kHz

It is clear from the above data that nonsynchronous sampling causes erroneous frequencies to appear in the passband. Synchronous sampling causes the image frequencies to overlap with the fundamental resulting in a single frequency spectrum for a single frequency input.

Referring back to Figure 6 there are also frequency components which appear at odd multiples of $f_S/2$. For synchronous sampling, these components will appear at either 0 or $f_S/2$ in the ADC's frequency spectrum. If the ratio of the ADC sampling frequency to the filter update frequency is even, these frequency components will appear at 0; if the ratio is odd, they will appear at $f_S/2$.

The above analysis can be applied to the DAC and reconstruction filter when generating analog output signals. In this case, the DAC output is oversampled by the switched capacitor filter. The lower order image frequencies generated from the DAC output will be attenuated by the switched capacitor filter. However, the higher order image frequencies will appear in the analog output spectrum of the filter. However, these components have a smaller magnitude than the lower terms and do not corrupt the spectrum as much as in the case of the ADC. Nonetheless, better results are achieved when synchronous sampling is employed.

DESIGNING A FILTER FOR NONSYNCHRONOUS APPLICATIONS

Inevitably, some applications will not be able to use synchronous operation. Then, additional filtering can be employed after the switched capacitor filter to improve performance. The function of the filter is to attenuate the image frequencies, such as those in Table III. The filter design must be based on the worst case, i.e., the image frequency with the largest magnitude. Neither of the values in Table III represents the worst case. Frequencies closer to $f_S/2$ will generate image frequencies which have a bigger magnitude and thus corrupt the spectrum more. Hence for a worst case analysis, Table III must be recalculated for the largest frequency in the passband spectrum, which is 3.5 kHz assuming an upper bandwidth limit of 3.5 kHz.

Table IV shows the recalculated image frequencies and their relative magnitudes.

Image Frequency (kHz)	ADC's Digital Frequency (kHz)	Magnitude (dB)
54.1	1.9	23.8
61.1	2.9	24.9
111.7	0.3	30.1
118.7	1.3	30.7
169.3	1.3	33.7
176.3	0.3	34.1

Table IV. Image Frequencies and Their Equivalent Frequency Positions in the ADC's Digital Spectrum. Signal Frequency = 3.5 kHz

For maximum attenuation of the image frequencies, the cutoff frequency should be as close as possible to the upper limit of the passband. Allowing approximately 20% headroom gives a value of 4.2 kHz for the desired cutoff frequency. The next step is to decide the stopband attenuation and transition band characteristics which

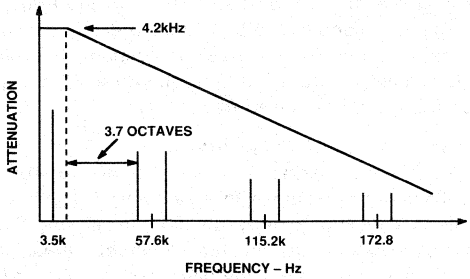

Figure 9. Image Frequency Spectrum and Required Filter Attenuation Profile

will ultimately determine the order of the filter. First, consider how a residue of an image frequency after attenuation will affect system performance. For example, a 12-bit system whose noise plus distortion is dominated by quantization noise has an ideal SNR of 74 dB. For the purpose of analysis assume that the value of the signal is 1 V rms. Then a value for the total rms noise in the system, e_Q, can be calculated as follows:

$$e_Q = log^{-1}(-74/20) = 199\ \mu V \quad \ldots\ldots\ldots\ldots\ldots (3)$$

Any additional noise, such as the residue from the image frequencies after attenuation, will add in an rms fashion to this noise, giving a degraded value for SNR:

$$SNR = 20\ log1/(e_Q^2 + e_F^2)^{1/2}$$

where:

SNR = System Signal to Noise Ratio
e_Q = Quantization Noise
e_F = rms Value of Additional Noise

A corresponding expression for the degradation, D, on system performance can now be written as follows:

$$D = 20\ log(1/e_Q) - 20log1/(e_Q^2 + e_F^2)^{1/2}$$

$$D = 10\ log(1 + e_F^2/e_Q^2)$$

A plot showing how additional noise degrades the system performance from 74 dB is shown in Figure 10. From the graph, if the system performance is not to degrade by more than 1 dB then any additional noise has to be lower than 80 dB relative to the signal.

This figure of 80 dB can now be taken as a design target for the filter. In other words, the rms residue from the filter after attenuation must be less than 80 dB. Looking at the image frequency magnitudes and locations shown in Table IV and Figure 9, the only image frequencies worth considering are the first two located at 54.1 kHz and 61.1 kHz. The other image frequencies which are smaller in magnitude (due to Sinx/x attenuation) will be further attenuated as they lie further up the frequency band. Taking the rms addition of the first two image frequencies gives a figure of approximately 21.3 dB. Subtracting this figure from the filter output target of

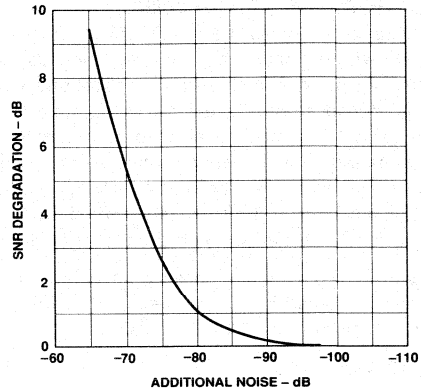

Figure 10. Degradation in SNR from the Ideal Value of 74 dB Due to Additional Noise

80 dB gives a value for the attenuation, A, required from the filter:

$$A = (80 - 21.3)\ dB = 58.7\ dB$$

The transition band in frequency octaves is given by:

$$log_2\ (54.1\ k\ /4.2\ k) = 3.7$$

The required transition band roll-off can now be written as 58.7/3.7 per octave or 15.9 dB per octave. This can be realized with a third order filter. One possible realization is shown in Figure 11 which is a third order Chebyshev configuration.

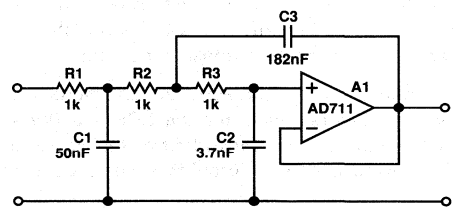

Figure 11. Chebyshev Filter Design for Nonsynchronous Applications

In this filter design example the frequency components which appeared at odd multiples of $f_S/2$ in Figure 6 were not considered. Taking the worst case component which lies at 28.8 kHz, its magnitude is approximately 60 dB down from the signal. The filter of Figure 11 attenuates this by a further 50 dB, making its effect on system performance negligible.

A similar exercise can be repeated for nonsynchronous operation of the AD7341 and accompanying DAC. However, in this analysis the on-chip continuous time filter will make the external filter requirement less stringent. A second order filter will suffice here.

VARYING THE CUTOFF FREQUENCIES

As already seen, the AD7341 and AD7371 filters are predominantly switched capacitor types, but they also contain continuous time filters. In any applications where the user wants to change the filter cutoff frequencies from the normal 3.5 kHz, the on-chip continuous time filter response must once again be considered. This filter response is the same for both the AD7341 and the AD7371. It is a second order Chebyshev which begins to roll off at 4.7 kHz and has a −3 dB point at 9.1 kHz.

The SCF sampling frequency necessary to obtain a particular filter cutoff can be determined by the following equations. For both the AD7341 and AD7371:

$f_S = f_{UFC}$ 57.6/3.5

where:

f_S = AD7341/7371 SCF Sampling Frequency
f_{UCF} = Upper Cutoff Frequency

For the AD7371, there is a ratio of 19.5 between the upper cutoff frequency (for the low pass section) and the lower cutoff frequency (i.e., the cutoff frequency for the high pass section).

Increasing the Cutoff Frequency

It is possible to increase the SCF sampling frequency by increasing the CLKIN frequency and get a corresponding increase in the cutoff frequency. However, for frequencies above 9 kHz the passband attenuation becomes significant and may be too large due to the on-chip continuous-time filter.

If a cutoff frequency of 10 kHz is required for the AD7371, for example, f_S needs to be 164 kHz and the CLKIN frequency needs to be 822 kHz. Figure 12, below, shows the filter frequency response for these conditions.

In cases where the AD7371 is operating in a nonsynchronous system (where the f_S for the SCF and the ADC sampling signal are not derived from the same clock source), it is necessary to provide a continuous time filter after the AD7371 to attenuate the image frequencies. This is accomplished by a third order Chebyshev filter as described in the previous section. This filter should be designed for a cutoff frequency of 1.2 × f_{MAX}, where f_{MAX} is the passband frequency.

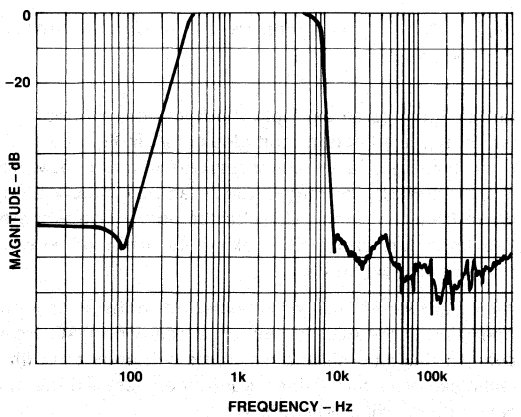

Figure 12. AD7371 Cutoff Frequency Set for 10 kHz

Decreasing the Cutoff Frequency

The AD7371 antialiasing filter for the receive ADC is itself a sampling system with a sampling frequency of 57.6 kHz and as such requires its own antialiasing filter. This is provided by the on-chip continuous time filter with characteristics as already described (see Figure 5).

If the cutoff frequency of the AD7371 is reduced in a synchronous system, then the on-chip antialiasing filter is no longer as effective. For example, halving the device CLKIN frequency automatically halves the upper cutoff frequency to 1.75 kHz. The SCF is now sampling at 28.8 kHz and the antialiasing filter needs to pass signals up to 1.75 kHz and attenuate signals above 27.05 kHz ($f_S - f_{MAX}$). The attenuation of ($f_S - f_{MAX}$) due to the on-chip filter is now reduced by a factor of about 14 dB (See Figure 5). A first order filter with a cutoff frequency at 2 kHz, positioned before the AD7371 will restore the attenuation level of signals above ($f_S - f_{MAX}$) to the required level. In fact, a first order filter is sufficient compensation down to cutoff frequencies of 1 kHz. Figure 13 gives an example of the AD7371 programmed for a cutoff frequency of 1 kHz with a simple first order filter providing the antialiasing.

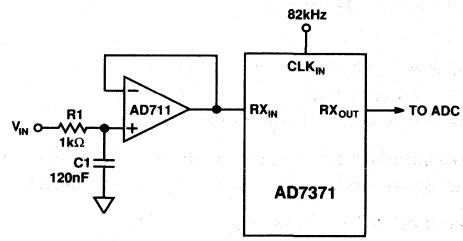

Figure 13. First Order Antialiasing Filter Used for the AD7371 When Set for a Cutoff Frequency of 1 kHz

Below 1 kHz, the on-chip continuous-time filter provides very little antialiasing and a second order Chebyshev with a cutoff frequency of about 1.2 × f_{MAX} is needed to complement it. An example is shown in Figure 14; this filter is designed for a 240 Hz cutoff which is suitable for an SCF cutoff frequency of 200 Hz.

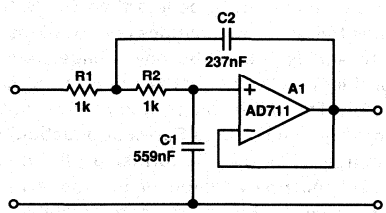

Figure 14. Second Order Chebyshev, Antialiasing Filter Used for the AD7371 When Set for a Cutoff Frequency of 200 Hz

When reducing the cutoff frequency of the AD7371 in a nonsynchronous system, there are two issues to be considered. The first is the antialiasing which has just been discussed. In addition to this there is the case of image frequencies appearing in the ADC passband as discussed earlier if the system is nonsynchronous (see synchronous versus nonsynchronous sampling). To attenuate these images, it is necessary to follow the AD7371 by a third order Chebyshev filter with an appropriate cutoff frequency. This should be set at 1.2 × f_{MAX}, the AD7371 passband frequency. Figure 15 shows the AD7371 in a nonsynchronous system with a cutoff frequency of 200 Hz and the necessary external continuous time filtering. The cutoff frequency of both external filters is 240 Hz.

In the AD7341, the continuous time filter occurs after the switched capacitor block. It is designed as the smoothing filter for the system. It eliminates the (f_S – f_{IN}) components due to the switched capacitor sampling. Reducing the cutoff frequency reduces the effectiveness of this filter. For synchronous applications, the compensation for this is the same as in the AD7371 case, except that here the filter comes after the AD7341. Again, a first order filter is sufficient down to programmed frequencies of 1 kHz, and below this a second order Chebyshev should be used.

In nonsynchronous systems, the output filter must be a third order Chebyshev to attenuate the image frequencies generated. The basic filter design is shown in Figure 11. The cutoff frequency should be chosen to be 1.2 f_{MAX}, where f_{MAX} is the chosen switched capacitor filter bandwidth.

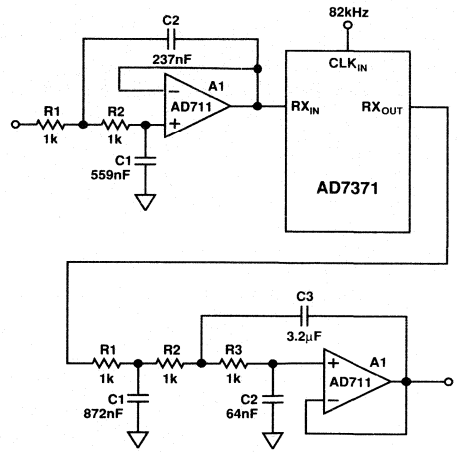

Figure 15. Extra Continuous Time Filtering in a Nonsynchronous System, AD7371 Cutoff Frequency = 200 Hz

Computational Products
Contents

		Page
AN-304	One-Chip "Slide Rule" Works with Logs, Antilogs for Real-Time Processing	6–3
AN-306	Synchronous System Measures μΩs	6–11
AN-307	Modem Circuit Techniques Simplify Instrumentation Design Techniques	6–13
AN-308	Commutating Amp Multiplies Precisely	6–27
AN-309	Build Fast VCAs and VCFs with Analog Multipliers	6–29
AN-212	Using The AD834 in DC to 500 MHz Applications: RMS-to-DC Conversion, Voltage Controlled Amplifiers and Video Switches	6–39
AN-216	Video VCAs and Keyers Using the AD834 and AD811	6–49

ANALOG DEVICES

AN-304
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

One-Chip "Slide Rule" Works with Logs, Antilogs for Real-Time Processing

by Lew Counts, Charles Kitchin and Steve Sherman

An analog computer IC multiplies and divides signals, takes powers, and also roots. A log-antilog section raises dynamic range during division.

An analog divider circuit seems to promise something for nothing. Just divide any number by zero and the result is infinity. But in theory, if numbers approaching zero are in the denominator, the divider must achieve infinite gain, and its input offset error must be smaller than the input signal.

In practice, analog divider chips fail to achieve even the modest 100-to-1 dynamic range vital to a wide-range circuit. A viable device calls for gains greater than 40 dB and input offset voltages under 100 μV.

But all is not lost. The AD538 multifunction IC's 1000-to-1 dynamic range surpasses the 50-to-1 limit of earlier one-chip dividers. Noise is held to only 25 μV—referenced to the input—over a 1000-Hz bandwidth, with its offset voltages under 100 μV.

A truly versatile analog computer on a chip, the 538 tackles one-quadrant multiplication and one- and two-quadrant division. It also calculates powers and roots of ratios. Like all analog computers, it runs in real time, making it the chip of choice when linearizing signals from transducers.

Its basic transfer function:

$$V_{out} = V_Y (V_Z/V_X)^M$$

makes the circuit simple to configure for a particular function. All that need be done is to connect specific pins and—in some cases—add one or two external resistors. Depending on the application, one, several, or all sections of the chip can be called into play.

The analog IC basically consists of an accurate 10-V/2-V reference and five precision op amps, all with offset voltages that are laser-trimmed to under 100 μV (Fig. 1). But the chip is more than the sum of its parts. It is designed as a complete analog computer whose system performance is specified for both voltage and current inputs. Its low input and output offset voltages, excellent linearity, and modest noise levels all contribute to its wide dynamic range—guaranteed from 1 mV to 10 V (80 dB).

The user's free access to the summing junctions of four of the chip's op amps gives it much of its versatility. External input resistors can change the pretrimmed offset or scaling voltages and allow multiple signals to be summed at each input terminal. The IC's power supply ranges from ± 4 to ± 18 V, letting

Reprinted from ELECTRONIC DESIGN – May 2, 1985

it run from standard split ±15-V supplies, as well as from ±12-V and even ±5-V units. The 2-V reference is particularly useful for driving the chip from ±5-V sources.

The op amps form three of the device's four major function blocks. The log ratio amplifier section harnesses three of them; the log and antilog section, one; and the output current-to-voltage stage, the last.

A logging operation

The V_Z and V_X inputs connect directly to the log-ratio section. This block furnishes an output voltage proportional to the difference between the natural logarithms, Ln, of the input voltages V_Z and V_X (Ln V_Z − Ln V_X). The transfer function between these inputs and the section's output pin (B) is given by

$$V_{out} = \frac{kT}{q} \text{Ln} \frac{V_Z}{V_X}$$

where k is Boltzmann's constant, T is the absolute Kelvin temperature, and q is the unit charge (1.60219×10^{-19}C).

The log ratio section may be used on its own (to compress the ratio of two signals) by temperature-compensating and scaling its output. To do so, its output (B) is joined to the summing junction of the output amplifier (I_{in}) through two external resistors. A 60-Ω 1% metal-film device is hooked in series with a 1-kΩ temperature-compensating resistor with a temperature coefficient of +3500 ppm/°C.

For most applications, however, the log ratio section is linked to the input of the antilog section (pin V_Y). That converts the signal from the logarithmic into the linear domain, according to the transfer function:

$$V_{out} = V_Y e^{(V_C q/kT)}$$

where V_C is the voltage at pin C. This section, like the

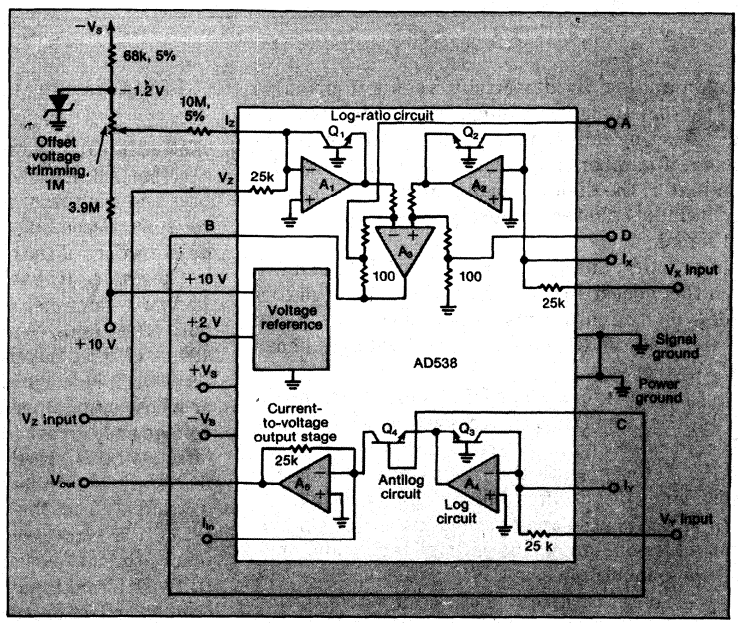

1. The transfer function of the AD538 one-chip analog computer is $V_{out} = V_Y (V_Z/V_X)^M$ when M is between 1/5 and 5. The flexibility of the transfer function allows the chip, which works in real time, to replace a microcomputer as well as several data converters.

Analog computer chip

log ratio block, may be used alone in order to expand a signal. Combining the transfer functions of both sections by tying together the B and C outputs results in the equation:

$$V_{out} = V_Y e^{\left(\frac{kT}{q} \times \frac{q}{kT} \times \text{Ln}\frac{V_Z}{V_X}\right)}$$

when $V_B = V_C$. This expression reduces to the multiplier-divider transfer function:

$$V_{out} = V_Y \frac{V_Z}{V_X}$$

Raising V_Z/V_X to the Mth power with an external resistor results in the analog computer chip's overall transfer function:

$$V_{out} = V_Y (V_Z/V_X)^M$$

where $\frac{1}{5} < M < 5$.

In most applications the V_Y input is used to set a convenient scale factor. The linear V_Y input signal is multiplied by adding the log of the V_Y input to the signal at C, which is already in the log domain. The chip's third section, the output current-to-voltage converter, buffers and scales the output from the antilog block.

Finally, the band-gap reference section supplies either +2 or +10 V, accurate to ±0.5% with a tempco of 25 ppm/°C. The 10-V reference is buffered and available at the ±10-V pin. The 2-V output is unbuffered. It is derived from the ±10-V pin by adding a resistive divider between the buffer's output and its summing point input.

As a result, any load on the 2-V output changes the feedback circuit and thus the 10-V reference voltage as well. To buffer the 2-V output, it is simply tied to the 10-V output.

Go forth and multiply

Taking the chip's sections singly—or even in pairs—only hints at its overall power. The device is easily configured as a high-performance, one-quadrant multiplier-divider (Fig. 1 again). As such, it handles only positive input voltages. Its transfer function:

$$V_{out} = V_Y (V_Z/V_X)$$

works with three input variables. That is, the chip carries out three-input calculations simultaneously—a task impossible with a conventional analog multiplier or even with a digital one.

What's more, the job is performed by the 538 in real time. When two input variables and a fixed reference are specified, the chip's dynamic range permits it to outclass such analog multipliers as the AD534. (The latter, though, performs four-quadrant multiplications—each of its two inputs may be a plus or minus voltage.)

In the circuit, if V_X is connected to either the 10- or the 2-V reference, a traditional one-quadrant analog multiplier results. (V_X, or any other input, may also be driven by a digital-to-analog converter, enabling the constant, or scaling, voltage, which establishes the output scale factor, to be set automatically or remotely.) With V_X tied to 10 V, the transfer function becomes:

$$V_{out} = (V_Y)(V_Z)/10\text{ V}.$$

Typically, the circuit achieves a bandwidth of 400 kHz when V_X varies over a range of 100 mV to 10 V. Maximum error with both inputs swinging

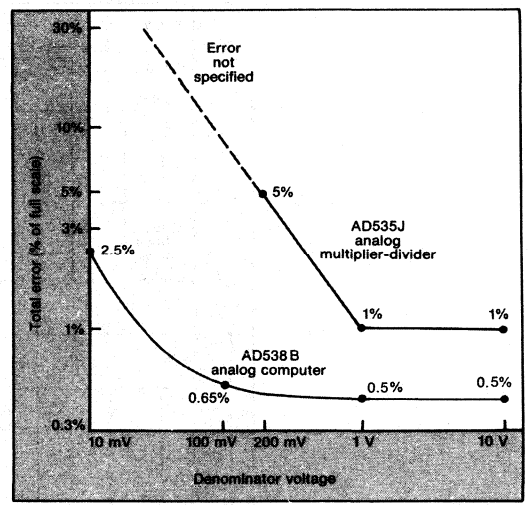

2. When used for division, the AD538 has a lower total error and a much wider denominator dynamic range than an analog multiplier-divider such as the AD535—which has even been specially trimmed to serve as a divider.

6

from 0 to 10 V is ±0.5% of the reading. Employing an external scaling voltage and the offset trim resistor (R_1) will halve the error.

Merely switching the 10-V reference from V_X to V_Y changes the circuit to a classical one-quadrant linear divider with 10-V scaling. V_X now serves as the denominator input terminal. The transfer function is specified as follows:

$$V_{out} = (10\ V)\,(V_Z)/(V_X)$$

The voltage applied to V_Z is divided by that applied to V_X. The internal trimming resistor reduces offset errors to less than 100 μV.

The typical −3-dB bandwidth of the one-quadrant divider is 370 kHz for inputs between 1 and 10 V, dropping slowly to 200 kHz for 2-mV signals. Its 74-dB dynamic range (from a 2-mV noise threshold at the input to a 10-V output clipping level) is much wider than that supplied by a conventional divider circuit using an analog multiplier.

As for harmonic distortion, the circuit's 10-V peak output is produced by a sine wave on the numerator input and 1 V on the denominator input. The second harmonic of the 10-V output is 70 dB below the fundamental.

Appearances are deceiving

The mathematics of multiplication and division makes multiplier circuits appear to have much lower input offset voltages than comparable divider circuits. Multiplying numbers less than one yields a result smaller than either of the original multiplicands. Multiplying the offset voltages at two different inputs, say 0.1 V each, produces an output offset voltage of 0.01 V. In addition, the product is generally divided by 10 (to increase dynamic range), further

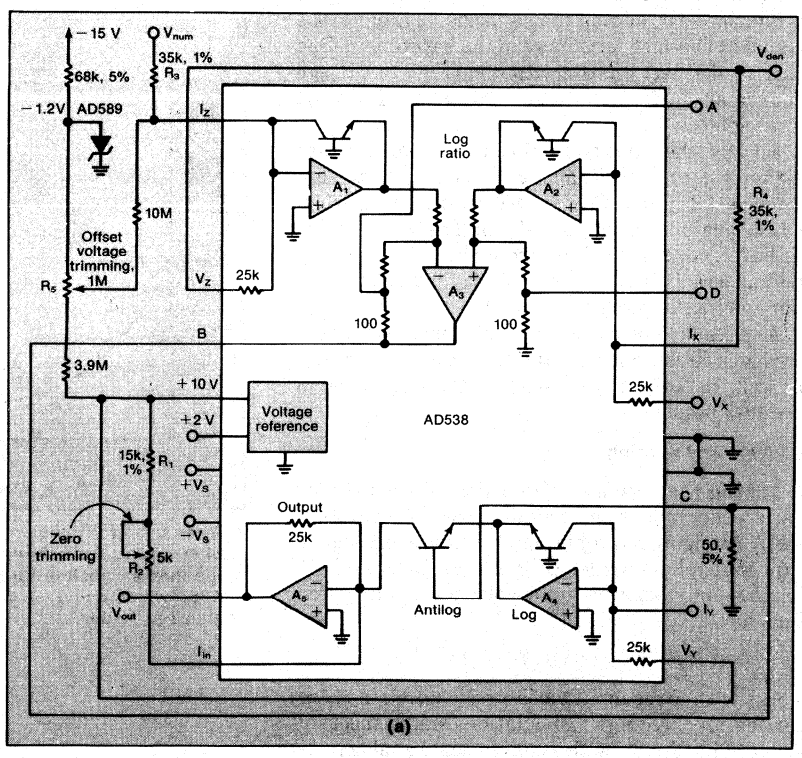

3. By implementing the transfer function $V_{out} = V_{num}/V_{den}$, the AD538 multifunction chip handles two-quadrant division. Thus the numerator is capable of handling both positive and negative voltages.

Analog computer chip

cutting output offset error to 0.001 V. Obviously, in multiplier circuits, output offset voltages matter more than do input offset voltages because the former add directly to the multiplier's output.

Divisive factors

In contrast, analog dividers usually run into difficulties over input offset voltages. The gain of an analog divider must approach infinity as the denominator moves closer to zero: Offset voltages or nonlinearities at the input are amplified by the $V_{scale}/V_{denominator}$ ratio of the divider. As the denominator voltage is reduced, input offset voltages in dividers quickly become intolerable, in turn severely degrading the device's dynamic range.

In multipliers, then, offsets are measured relative to full scale; in dividers, they must be gauged relative to the magnitude of the input signal. For any given accuracy, therefore, these offsets limit the minimum denominator voltage level.

The chip overcomes most of its input offset problems by employing high-quality op amps as input stages. These are quite different from the circuits found in the typical analog multiplier, such as the AD535, and the op amps hold offset errors under 100 μV.

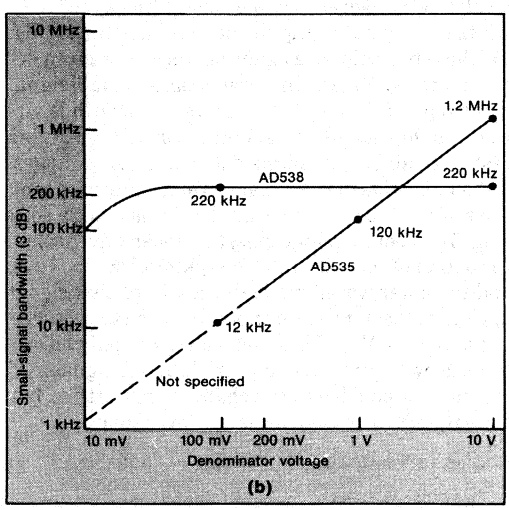

4. The bandwidth and dynamic range of the AD538 in the circuit of Fig. 3 are significantly greater than those of a conventional multiplier-divider IC, the AD535, in the same application.

A smaller total error

The device's superiority is readily seen by directly comparing total error from each chip for given denominator voltages (Fig. 2). The total error for the 535, configured as a multiplier-divider, is not even specified for denominator voltages below 200 mV. At that point it is, in fact, 5%. Alternatively, the 538's total error is only 2.5% at a denominator voltage of 10 mV—a level at which the typical multiplier-divider is unusable.

The term "total error" represents the sum of three elements: the inherent nonlinearity produced by the internal component errors; the numerator offset error multiplied by the set gain (V_Y/V_X); and the static offset error of the output amplifier. In other words, total error is expressed as:

$$\text{Nonlinearity} + (Z_{OS} \times V_Y/V_X) + \underset{\text{of output of output amp}}{V_{OS}}$$

where the subscript "OS" indicates offset.

Specifically, the chip's total error equals:

$$\pm 0.5\%\ V_{out} \pm 0.15\ \text{mV}(V_Y/V_X) \pm 0.15\ \text{mV}$$

When V_Y/V_X is small, the nonlinearity error dominates; when that ratio is large, the numerator offset error dominates; and when the output is small (under 10 mV), output offset dominates. Errors are defined in this manner so that the user may minimize total error by nulling out the error dominating a particular input condition.

Another quadrant heard from

Although the one-quadrant multiplier-divider tackles only positive input voltages, the chip can also be set up as a two-quadrant divider capable of handling bipolar numerator signals. That feat is accomplished by offsetting the nominal numerator input at V_Z by a voltage that tracks the denominator input signal V_{den}. The latter is made slightly larger than the former, by a ratio of 35 to 25 (Fig. 3). The ratio is set by feeding the numerator and denominator inputs through external 35-kΩ resistors (R_3,R_4) to the

summing node (current input) of their respective op amps, A_1 and A_2. The offsetting voltage V_{den} is summed with the numerator input, V_{num}, through the internal 25-kΩ resistor at V_Z. That offset changes the divider's transfer function from 10 V (V_Z/V_X) to:

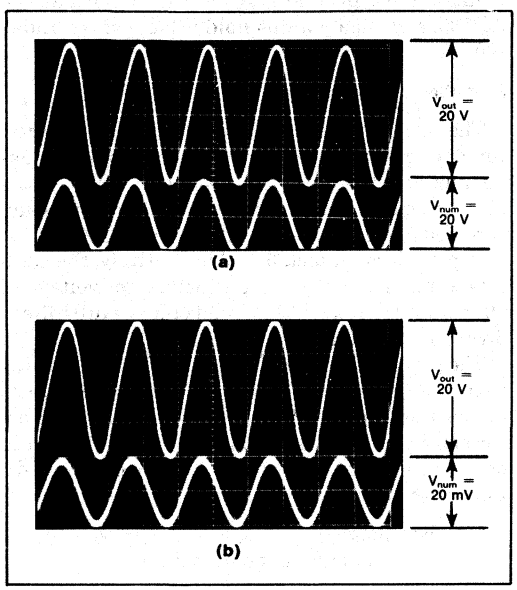

5. In a two-quadrant divider, the chip exhibits the same dynamic response with numerator inputs between 20 V pk-pk (a) and 20 mV pk-pk (b) with denominators of 10 V and 10 mV, respectively.

Table 1. Resistor values for powers less than one

Function	R_B (Ω)	R_C (Ω)	Power (M)
One-fifth root	162	40.2	0.20
One-quarter root	150	49.9	0.25
Cube root	100	49.9	0.33
Square root	100	100.0	0.50

Table 2. Resistor values for powers greater than one

Function	R_A (Ω)	Power (M)
Fifth power	48.7	5.0
Fourth power	64.9	4.0
Analog cube	97.6	3.0
Analog square	196.0	2.0
Arc tangent	931.0	1.21

$$V_{out} = +10\,V(V_{num} + V_{den})/V_{den}$$
$$= +10\,V(1+V_{num}/V_{den})$$
$$= 10\,V + 10\,V(V_{num}/V_{den})$$

As long as the denominator input equals or exceeds the numerator in magnitude, the circuit accepts bipolar voltages (such as sine waves) at the numerator input. However, its output now equals +10 V dc without any numerator voltage applied.

To make the circuit practical, the +10-V output offset is canceled by summing it with the 10-V reference; this is done in the output op amp section through resistors R_1 and R_2. The potentiometer trims the offset, so that there is zero out for zero in, leaving the simple transfer function:

$$V_{out} = 10\,V(V_{num}/V_{den})$$

At a denominator voltage of 100 mV, this log-antilog two-quadrant divider has 20 times the bandwidth of a typical multiplier-divider like the 535 (Fig. 4). These two chips are again compared for given denominator voltages, this time against small-signal bandwidth. The log-antilog device bandwidth is virtually independent of denominator voltage and is faster for all but the highest denominator voltages.

An indication of the dynamic range, precision, and linearity of the two-quadrant divider is readily seen (Fig. 5). In both photographs, the upper waveform is the output of the chip, a 20-V pk-pk 20-kHz sine wave, and the lower waveform is the source of that signal applied to the numerator input. In each case the output is 10 V × V_{num}/V_{den} and shows virtually no difference over a 1000-to-1 change in input signals. Of course V_{den} could have been changing at the time, but the outputs would then be harder to compare.

Getting to the root of the matter

The 538 takes roots in real time with an accuracy that might well appear unbelievable for an analog circuit (Fig. 6). When configured as a square-root (M = ½) circuit, it has a transfer function of:

$$V_{out} = 1\,V\,(V_Z/1\,V)^{1/2}$$

Its scale factor adjustment R_1 and offset trimming resistor R_2 let the curcuit reach an output voltage that is within 0.35% of the ideal over a range of 1 mV to 10 V (80 dB). For an input range of 10 mV to 10 V (60 dB), the output is within ±1 mV of the theoretically correct value. This is equivalent to 13-bit performance referred to the output or 17 to 20 bits referred to the input. Circuit bandwidth is typically

Analog computer chip

280 kHz with a 1-V pk-pk sine-wave input.

As in the divider circuit, the input signal at V_Z is divided by the level at V_X. Its 1-V scaling is obtained by halving the 2-V reference with resistors R_3 and R_4, applying roughly 1 V to both V_X and V_Y. V_X is actually set low, at 0.95 V, resulting in a V_Y value that is slightly high, providing a ±5% scale factor trimming range. The voltage terms in V_X and V_Y cancel, producing the dimensionless output required to determine a square function.

The square root is extracted by raising the quantity V_Z/V_X to the ½ power with resistors R_B and R_C. For maximum linearity the pair of resistors should be ratio-matched to better than ±1%. Changing the resistor values of R_B and R_C (Table 1) enables cube, quarter, and fifth roots to be determined according to the equation:

$$M = R_C/(R_C+R_B).$$

Finally, to raise the ratio V_Z/V_X to a power greater than one, it is only necessary to change the gain of the log-ratio subtracter op amp, A_3, by connecting a single resistor, R_A, between pins A and D. Varying the values of the resistor (see Table 2) creates M values from one to five according to the equation $R_A = 196/(M-1)$ and allows the arc-tangent function to be computed.

Since resistors of any value may be used, the chip is not limited to integral powers or roots. V_Z/V_X could be raised to the 3.17 power, say. It is only necessary that M stay between ⅕ and 5.

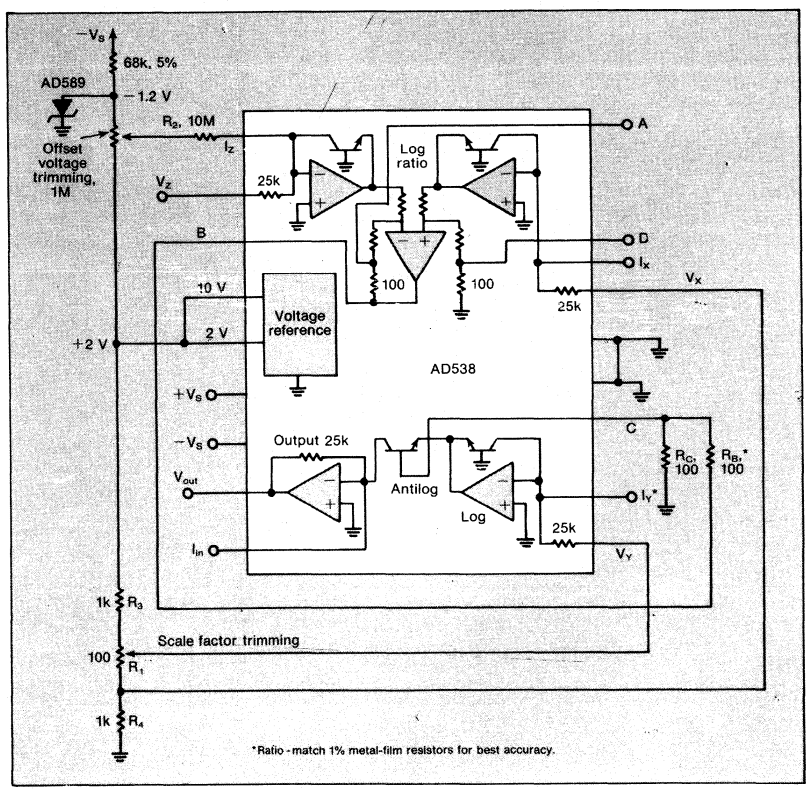

6. Resistors R_B and R_C set the chip to take the square root of the voltage at V_Z according to the transfer function $V_{out} = 1\ V\ (V_Z/1\ V)^{1/2}$.

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Synchronous System Measures μΩs

by Moshe Gerstenhaber and Mark Murphy

The circuit in **Fig 1** uses a synchronous-detection scheme to measure low-level resistances. Other low-resistance-measuring circuits sometimes inject unacceptably large currents into the system under test. This circuit synchronously demodulates the voltage drop across the system under test and can hence use extremely low currents while measuring resistance.

The 10V-pk, 1-kHz carrier generator injects a 1-mA reference current into unknown resistor, R_{TEST}. Instrument amplifier IC_1 and precision op amp IC_{2A} amplify the voltage across R_{TEST} by a gain of 100,000. Synchronous detector IC_3 demodulates this voltage, then op amp IC_{2B} acts as a lowpass filter on the demodulated voltage. The lowpass filtering will attenuate all uncorrelated disturbances, such as noise, drift, or offsets, while passing a dc voltage proportional to the unknown resistance.

The relationship between the output voltage and the unknown resistance is

$$V_{OUT} = 10 \times (2V/\pi) \times R_{TEST} \times 10^5/10\ k\Omega, \text{ or}$$
$$R = 0.0157 \times V_{OUT},$$

which is 15.7 mΩ/V at the circuit's output.

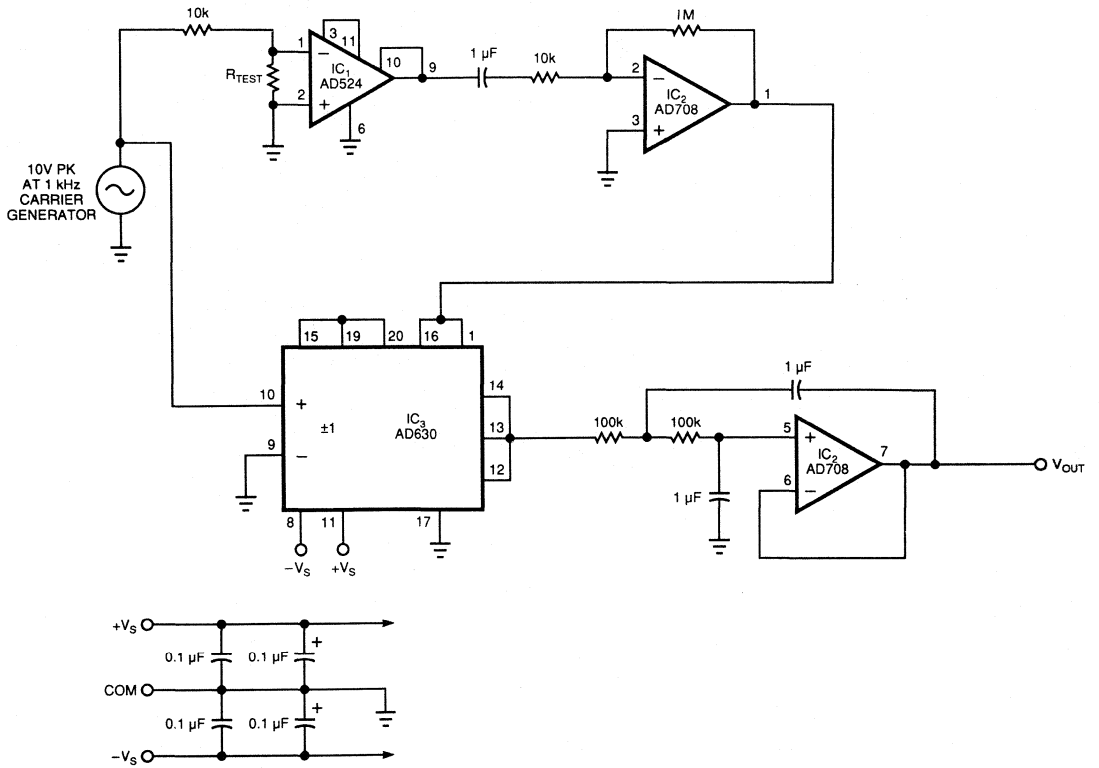

Fig 1—A synchronous demodulator helps this circuit *measure low-level resistances while rejecting uncorrelated disturbances, such as noise, drift, and offsets.*

Reprinted from EDN — May 9, 1991

ANALOG DEVICES

AN-307
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Modem-Circuit Techniques Simplify Instrumentation Designs

by Walt Jung and Moshe Gerstenhaber

A commutating modem circuit, whether an IC or pc-board version, can serve as a building block for many measurement-system functions, reducing noise and ensuring the integrity of transmitted data.

Although used extensively in communications, amplitude modulation/demodulation can also serve industrial-control instrumentation systems. In such applications, a commutating modulation/demodulation circuit is an extremely versatile building block that can perform several functions—such as amplification and clock generation—in addition to helping assure the integrity of measurement-data transmission. This article describes the operation and application of that type of modem in a variety of instrumentation situations.

Modulating and demodulating the output of an instrumentation system's sensors, a process often referred to as balanced mod/demod, provides a host of advantages:

- Ease of filtering
- Inherent noise discrimination
- Good linearity over a wide dynamic range
- Transmission of phase as well as amplitude information
- Relaxation of intermediate-stage amplification requirements.

Fig 1's representation of a typical control-instrumentation system helps demonstrate how to use a modem to realize these advantages. Transducers typically used in such control systems—such as linear and rotary variable differential transformers (LVDTs or RVDTs), ac bridges and photochoppers—measure occurrences that are usually static or slowly varying. This measurement

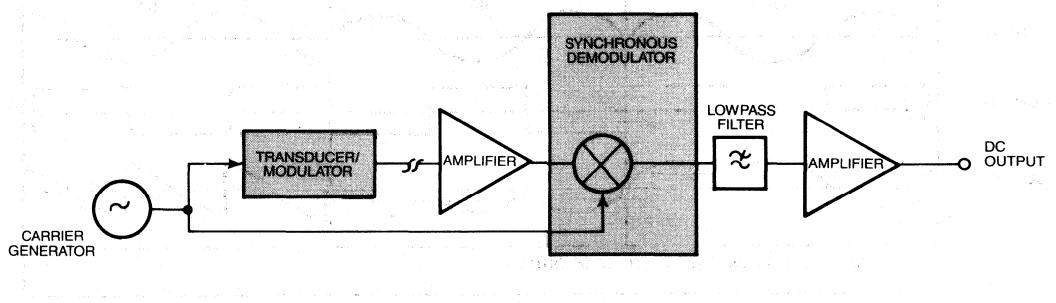

Fig 1—A typical instrumentation system *can include a synchronous demodulator that removes the carrier signal from a transducer's modulated output. Using balanced modulation/demodulation removes system offset, drift and noise too.*

Reprinted from EDN — August, 1984

Balanced modem circuit removes low-frequency noise

data can be imposed on a carrier signal by modulation; modulation capability may be inherent to the transducer, as is the case with an LVDT, or require use of an external modulator, as with some types of photochoppers.

The modulated signal is then typically transmitted to a central controller that extracts raw measurement data from the signal. If the carrier's frequency is well above the measured parameter's rate of change (typically, a carrier frequency of about 1 kHz gives sufficient carrier/data separation), simple post-demodulation filtering rejects all noise and dc errors but leaves the original signal intact.

The balanced modulation/demodulation scheme inherently rejects frequency disturbances—such as 1/f components, drift and power-line noise—that are either asynchronous to or in quadrature with the demodulating carrier. Thus, the combination of balanced modulation/demodulation and filtering eliminates many of the chief causes of measurement errors.

Consider, though, that the additional modulation/demodulation circuitry could introduce new sources of error. One type of modulation/demodulation circuit, the commutating modem, minimizes the sensitivity of the measurement signal to variations in the amplitude and waveform of the demodulating carrier signal, thus minimizing the effects of demodulation on measurement results.

To understand how a commutating modem minimizes noise sensitivity, consider two ac signals: a modem input that you wish to demodulate and a carrier signal that you can use to do the demodulating. When the carrier is positive, the commutating modem multiplies the input by +K (a scaling constant that you choose), and when the carrier is negative, by −K. Because the carrier's transitions between positive and negative states control the modulation process, the carrier's waveform and amplitude do not affect the output.

Fig 2 graphically shows the commutating modem's multiplication process. If the input and the carrier signal are in phase, the commutating modem's output is a positive, full-wave-rectified version of the input. When the input and the carrier are 180° out of phase, the output is a negative, full-wave-rectified version of the input. In either case, the average value of the output is proportional to the amplitude of input signal. But for an input signal in quadrature with the carrier, the output is a zero average value. In other words, the averaged output is proportional to the input amplitude as well as the phase difference between the input and the carrier.

Because modulation and demodulation are reciprocal processes, the **Fig 2** waveforms can also apply to a modulation scheme in which the input and output are reversed. That is, when the signal labeled output is mixed with the carrier at the modulator, that device creates the waveforms marked input.

This reciprocity is a powerful advantage of balanced modulation/demodulation because it allows you to use a common circuit for both modulation and demodulation. And most hardware implementations of the balanced modulation/demodulation circuit function are designed to be used either way.

Fig 3 shows a commutating balanced modulation/demodulation circuit that you can build with standard components. This circuit, commonly used in data acquisition and signal processing, is called an absolute-value circuit, a precision rectifier or a sign-programmable amplifier, depending on its specific use.

The **Fig 3** circuit is optimized for the commutating

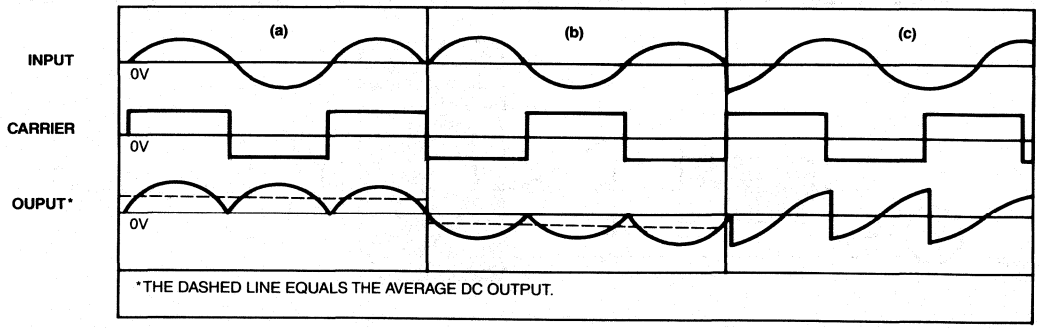

Fig 2—Demodulating an input signal *produces different outputs depending on the phase relationship between the signal and the modulating carrier. The average dc output is positive* ***(a)****, negative* ***(b)*** *or zero* ***(c)*** *when the input and carrier are in phase, 180° out of phase or 90 or 270° out of phase, respectively. For simplicity, the input is shown as a sine wave of the same frequency as the carrier, but a carrier whose frequency is an odd multiple of the signal frequency produces similar average-dc results.*

A commutating modem IC

The AD630 includes all the circuit blocks needed to implement balanced modulation and demodulation with high precision in a single IC. Shown functionally in the **figure**, it contains two uncommitted op-amp input stages, A and B. Depending on the state of the input to the IC's integral comparator, either stage A or B is connected to the chip's integrating output stage. When pin 10 is high relative to pin 9, input stage A is active; when the reverse is true, input stage B is active.

The chip's comparator is an internally latched stage with a specified switching window of ±1.5 mV (max), a window that includes not only comparator offset but a hysteresis of about 0.2 mV as well. The hysteresis is built in to make switching clean and unambiguous, even with slowly varying or noisy signals.

When switched, the AD630 can slew as fast as 45V/μsec between the output limits. In response to a 20V step, the output typically settles to within 0.01% of the final value in 3 μsec.

Inputs to stages A and B can be connected into conventional op-amp feedback loops, either with external components or internal resistors. On-chip resistors are ratio-trimmed for precise gains of ±1, ±2, +3 and +4 and are easily configured by appropriate pin strapping. Other gains are possible with external resistors.

The advantage of the internal resistors lies not only in their convenience, but also in their precision. Gain error is 0.05% (max) for B and K grade devices and 0.5% for other grades. These resistors also feature a 2-ppm/°C tracking specification. Resistors for bias-current compensation are also available at the A+ and B+ inputs for optional use in attaining the highest dc accuracy.

With its on-chip resistors, the AD630 can achieve submillivolt dc offsets and 0.1% or better dc accuracy without user trimming or gain-setting resistors. The dynamic specifications are such that the device is generally useful for modem uses to 100 kHz, and it is useful with high precision through the range of common instrumentation carrier frequencies—well above the limits of 0.5V/μsec-slew-rate devices.

All AD630 family devices come in plastic or ceramic 20-pin DIPs and operate from standard ±15V supplies, with conventional ±10V input/output ranges.

Many details of the AD630's design prove advantageous in various applications. For example, you can use the open collector output from the comparator (pin 7) as a logic status indicator. This status output can be used with an external pullup resistor to a positive voltage and can sink load current to $-V_S$ in the 0 state.

Another useful feature is that you can optimize device frequency compensation by strapping pin 12 to 13. Doing so reduces device bandwidth and phase at unity gain to its lowest value and reduces slew rate to 35V/μsec. With this pin open, the device speed is at its 45V/μsec maximum.

Two optional dc-offset trim connections use a pair of 10-kΩ trimmers: one between pins 5 and 6 and the other between pins 3 and 4. The wipers of both trimmers are tied to $-V_S$ (pin 8). The offset trimmer between pins 3 and 4 is first adjusted for a minimum-level square-wave output with a switching signal applied to the comparator and 0V dc input. The offset trimmer at pins 5 and 6 is then used to adjust the output offset to zero, with the comparator state fixed. These offset trims need only be used in the most demanding applications because of the low level of pretrimmed offset.

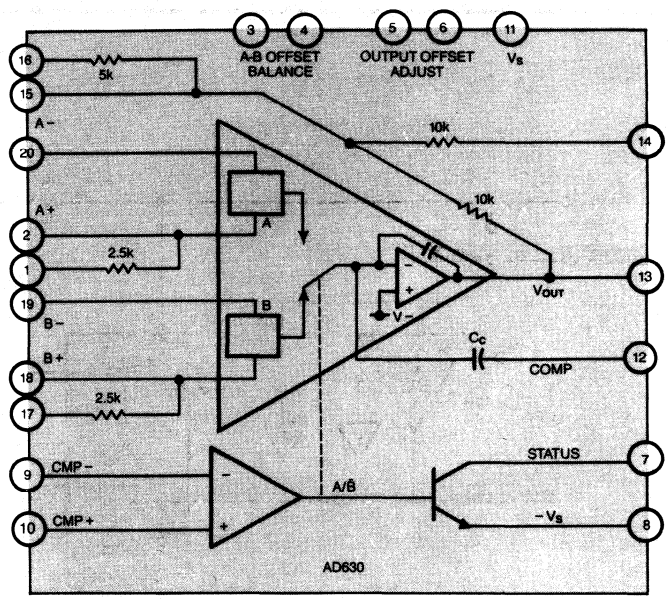

Uncommitted op-amp inputs A and B *allow the AD630 to operate as a commutating modem. An integral comparator selects between the two stages depending on the relative input levels at pins 9 and 10.*

function and minimizes errors from both the dc and ac sources. The carrier input controls output polarity with switch SW_1 set as shown, and the circuit functions as a modulation/demodulation circuit that exhibits the **Fig 2** waveforms. It might also be correctly referred to as a sign-bit amplifier, generally used in A/D or D/A conversion. In conversion applications, the carrier input would be called the sign-bit input.

When comparator stage A_3 is driven by the input signal rather than the carrier, the circuit is an absolute-value (or precision rectifier) circuit. Regardless of which signal does the driving, the output polarity can be changed by reversing the input of comparator A_3.

Although furnishing good performance, this circuit needs three separate active devices, a matched pair of precision resistors, board space, and design and debug time to build it. Furthermore, you have to optimize the amplifier for low noise, low offset and low drift and fast slew rate—not a trivial design problem. Integrating these components with an eye toward such performance objectives would be a logical improvement (**Ref 1**), and a chip that does just that is discussed in the accompanying **box** ("A commutating-modem IC"). This chip is used for simplicity in the examples that follow, but the discrete circuit of **Fig 3** could also be used.

Fig 4a shows a general-purpose modulation/demodulation circuit using the AD630. It is basically a switched-gain circuit with an absolute gain of 2; the V_{REF} input to the comparator determines the output polarity. Like the **Fig 3** circuit, this circuit can be used as a modulator or demodulator, a sign-bit amplifier, an absolute-value (precision rectifier) circuit or a phase detector applied to the device.

To understand the circuit operation, think of the AD630 as having three separate internal functions: a comparator stage and an output stage and two switched-input stages, only one of which is on at a time, depending on comparator state. The active input stage in conjunction with the output stage forms a single composite op amp, so you can regard the AD630 as an op amp that has two configurations depending on the comparator input.

For example, with input stage **A** on, the overall circuit is equivalent to that of **Fig 4b**. The circuit has a noninverting gain of 2 when V_{REF} is greater than V_{OV}, or in general, when the AD630's pin 10 has a higher input than its pin 9. Conversely, when input stage **B** is on, the overall circuit is equivalent to the circuit of **Fig 4c**. This circuit has an inverting gain of 2 when V_{REF} is less than V_{OV}, or in general, when pin 10 is lower than pin 9.

In this circuit, the feedback loop that sets the gain at 2 is defined using the chip's internal resistors that are

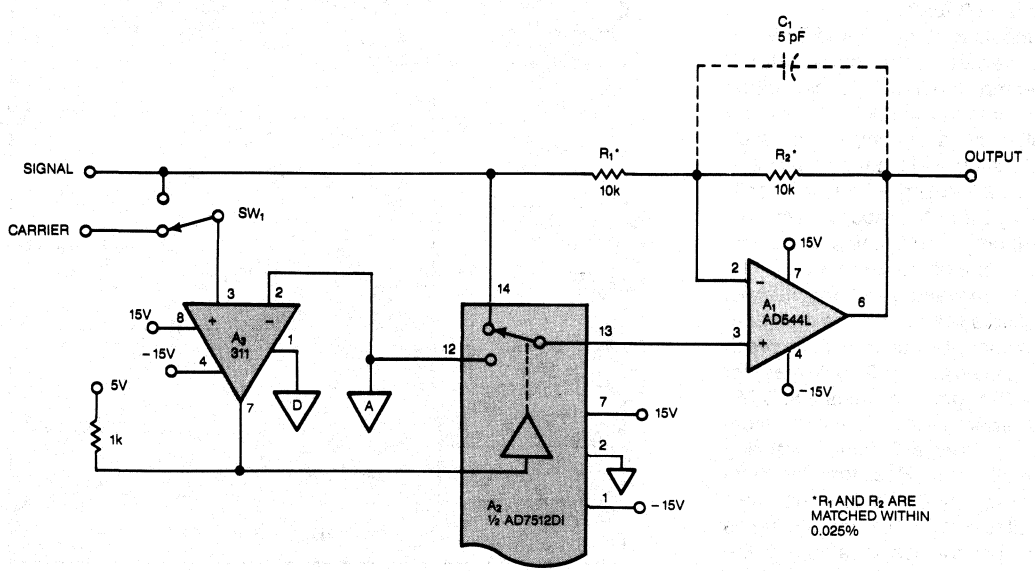

Fig 3—Known by a variety of names, *the basic commutating modulator/demodulator circuit is called a sign-bit amplifier when comparator A_3's carrier input, selected with switch SW_1 in the position shown, is a sign-bit signal. When the signal serves as A_3's input, the circuit is called an absolute-value or precision-rectifier circuit.*

connected to the circuit's pins 13, 14, 15 and 16. Including a variety of internal precision resistors with pin-accessible taps lets you use jumpers to simply select the gain state desired. If you should decide to use your own external feedback components, the sign and the gain magnitude might be different. Regardless of the feedback-loop approach you choose, the active state of inputs A and B (**Fig 4b** and **Fig 4c**) remain generally true.

In many applications, the A and B inputs are often interconnected, as A− and B− are in **Fig 4a**. This presents no output problem because the deselected input is switched off and presents no extra loading to the summing point.

However, the circuit presents different loads at its signal input depending on the input to the comparator. Note that the inverting configuration B in **Fig 4c** presents a load equal to R_{IN}, or 5 kΩ, when this input is selected, and the noninverting stage (**Fig 4b**) presents an input impedance intrinsically higher. Thus, the external input source faces a load that varies dynamically. To maintain the highest precision, you may have to add an input buffer to compensate for this dynamic loading.

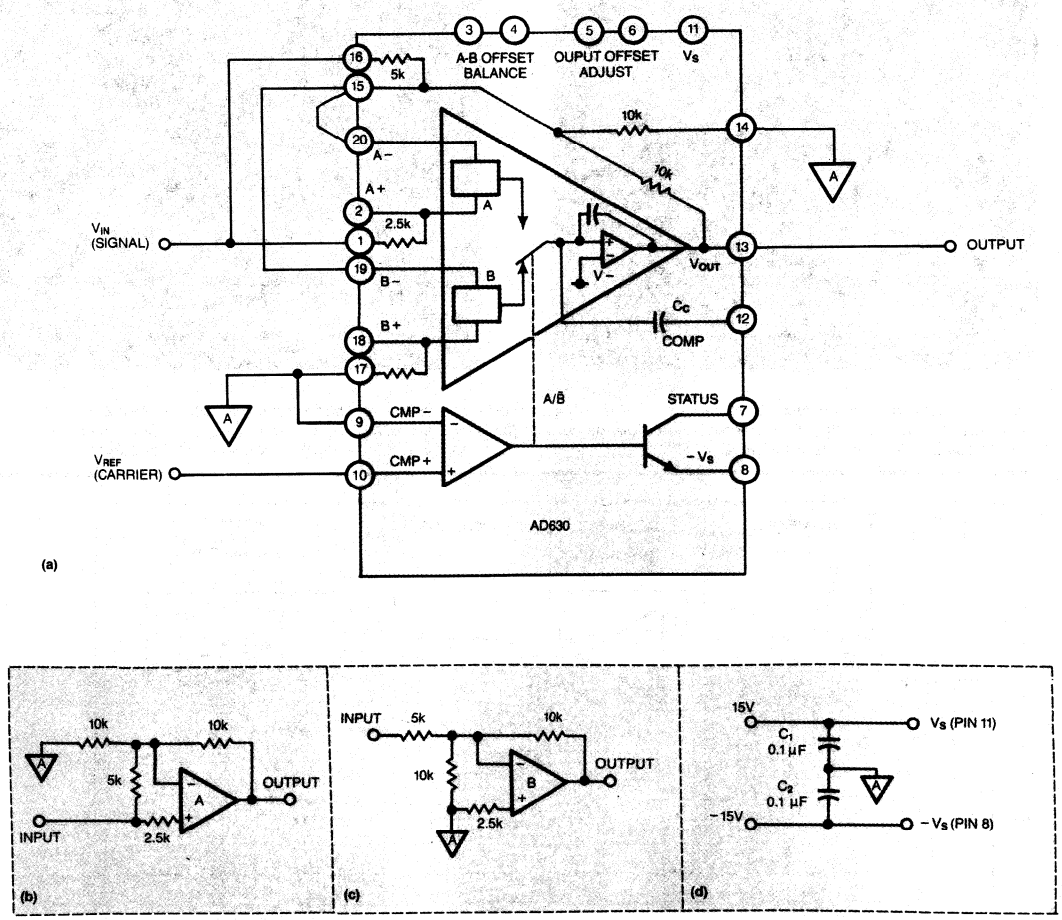

Fig 4—A building block IC *like the AD630 can be turned into a balanced modulator/demodulator* ***(a)*** *similar to that in* ***Fig 3.*** *Using the chip's precision resistors to set the gain (to a value of 2 with the configuration shown here) eliminates the need for trimming in most applications. The equivalent circuits seen when looking into the AD630's input when channel A or B is active are shown in* ***(b)*** *and* ***(c),*** *respectively, and the supply-conditioning circuitry for the chip appears in* ***(d).***

LVDTs' inherent modulation requires a precise carrier

For absolute-value applications, the AD630 has distinct advantages over the traditional op-amp approach. Even when implemented with high-performance amplifiers, traditional absolute-value circuits have classic problems: With high-frequency inputs of low amplitude, for example, the op amp must slew for high percentages of the time to accommodate the diode thresholds. As a result, the output waveform can be severely distorted. The balanced-modulator approach alleviates this problem simply by changing the sign of the forward signal path in synchronism with the zero crossings, with no diode thresholds to overcome. Using the **Fig 4** configuration as a gain-of-2 absolute-value circuit produces **Fig 5**'s results.

Such balanced modulation/demodulation techniques are highly useful with a variety of common transducers. What's more, a number of transducer types lend themselves to this approach because they perform the modulation of an ac carrier as an integral part of their transducing function. Examples of such transducers are those that fall in the LVDT and ac-bridge categories. Other types, such as some photochoppers, don't furnish modulation and require further interfacing to modulate their outputs. The **Fig 3** and **Fig 4** circuits can be built

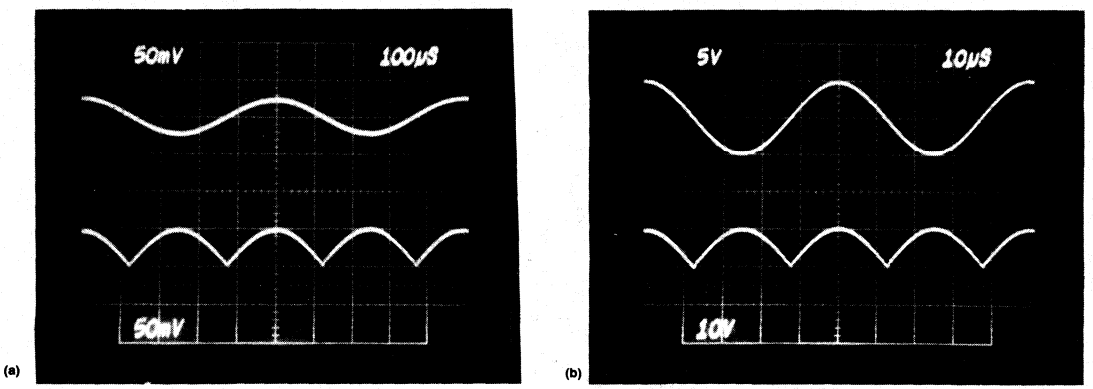

Fig 5—In an absolute-value circuit *with a gain of 2, the AD630 produces positive rectified outputs. In* **(a),** *the 2-kHz, 25-mV pk signal at the top produces the 50-mV pk rectified output below it. In* **(b),** *a high-amplitude (5V pk), high-frequency (20-kHz) signal at top produces the 10V pk output below it. In either case, the sign of the input is available from the device's status output.*

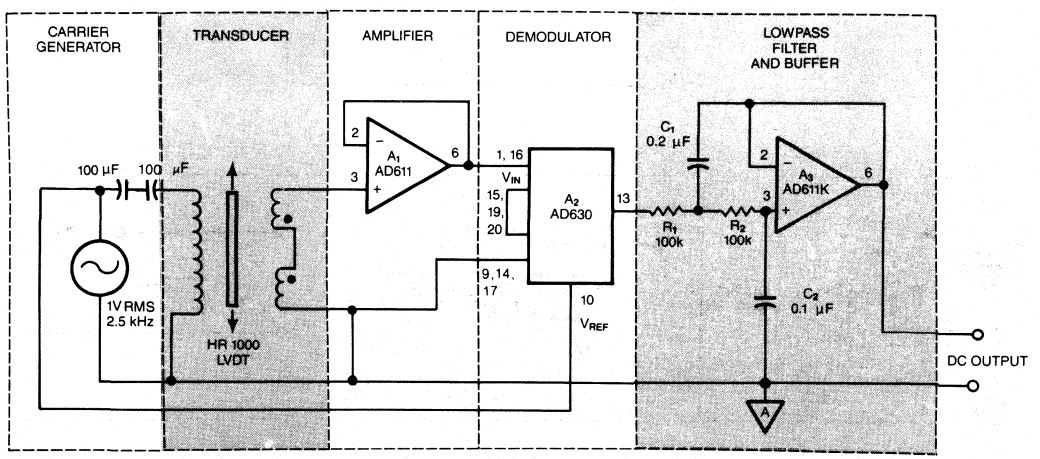

Fig 6—A linear variable differential transformer *whose measurement signal rides on a 1V rms, 2.5-kHz carrier can work with a commutating modulation/demodulation circuit such as the one shown in* **Fig 4.** *The lowpass filter at right, with a 10-Hz corner frequency, removes any high-frequency noise.*

into or added to the sensor/transducer to provide this modulation.

Among the transducers that have integral modulating capability, the LVDT measures linear displacement, and a related type, the RVDT, measures rotary displacement (**Ref 3**). Although the mechanical construction of these two types of transducers differs, they share many electrical similarities. Therefore, similar electrical considerations apply to the application of either type of device.

A transducer of either type is really a specialized transformer. The transformer is connected with two secondary windings in series opposition and exhibits a minimum electrical output at the LVDT's mechanical null point. The core motion, controlled by the object whose displacement is to be measured, varies the mutual inductance between the primary and the two secondaries. This variation produces a variable-phase and -amplitude ac output that is linearly proportional to the core displacement. A balanced demodulator can transform this variable ac output into a bipolar dc output that, when referred to the LVDT reference, or null, position, indicates the displacement.

The LVDT's physical advantages lie in the virtues of a low core mass, an essentially frictionless, no-hysteresis action, infinite resolution, and long mechanical life. Electrical advantages are the low sensitivity to external fields and the high common-mode isolation that a transformer inherently affords. The main disadvantage is the relatively high cost, which tends to be proportional to core length and accuracy. LVDTs come in a wide variety of mechanical and electrical configurations, with different nonlinearity and sensitivity specifications (**Ref 3**).

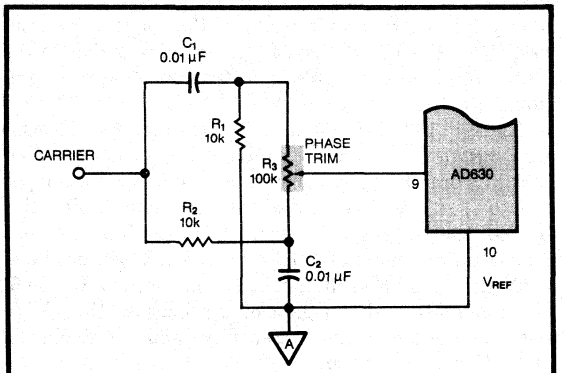

Fig 7—To trim a carrier signal *so that it's in phase with the input signal, you can use this lead-lag network. With the transducer at its maximum value, adjust resistor R_3 to maximize the circuit's average dc output.*

Fig 6 shows a representative application using a general-purpose LVDT, a Schaevitz Engineering (Pennsauken, NJ) Model HR 1000, whose core can travel distances of ±1 in. and whose output i

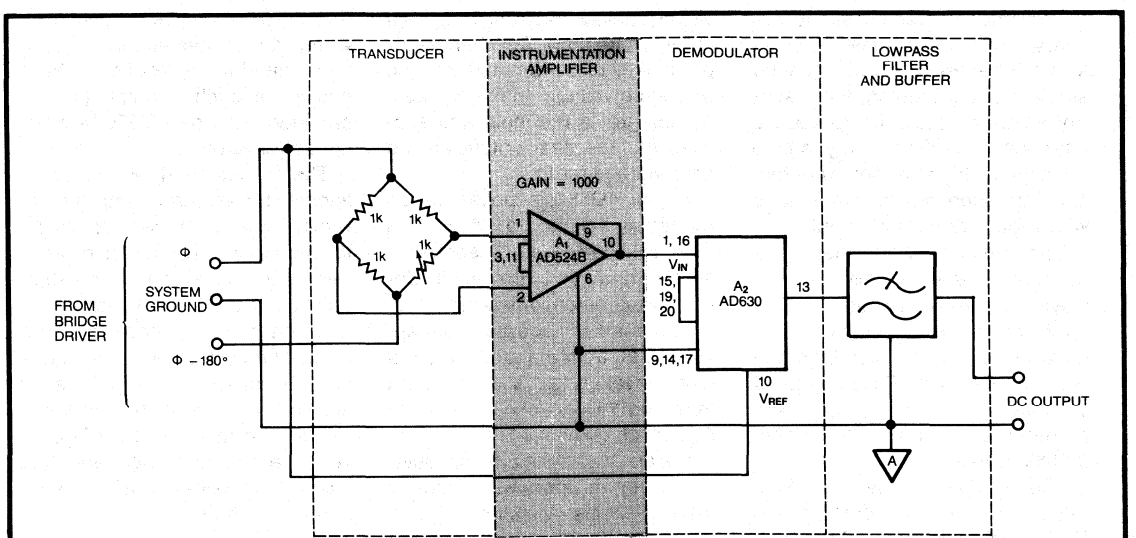

Fig 8—To measure changes in a bridge's variable arm, *this circuit uses a demodulator and filter similar to those in* ***Fig 6.*** *The bridge circuit's low-voltage output, however, requires use of a high-gain instrumentation amplifier, and the circuit requires a special bridge driver such as the one shown in* ***Fig 9.***

Carrier phase trimming may be needed for accuracy

within 0.25% FS. This circuit demonstrates some of the general principles of LVDT interfacing and can be used as a guide with other variable differential transformers. With the **Fig 6** circuit, system errors due to the electronics are negligible, and transducer errors predominate.

The HR 1000 can handle a 3V rms maximum excitation signal. An LVDT's sensitivity is typically specified in terms of millivolts of output per volt of excitation per mil (0.001 in.) of displacement (mV/V/mil). For this device, the rated sensitivity is 0.39 mV/V/mil, so with an excitation signal of 1.0V rms, the HR1000 produces a 390-mV rms FS (1-in. displacement) output. In **Fig 6**, a demodulator with a gain of 2 processes the LVDT output, doubling the sensitivity and producing a 780-mV rms FS output.

A Sallen-and-Key two-pole active filter with unity gain filters the demodulator output and buffers it to a low impedance. This circuit, using A_3 and the associated passive components, employs a maximally flat Butterworth alignment for minimum passband-amplitude errors. This filter's corner frequency is approximately 10 Hz, and the filter yields more than 80 dB of ripple attenuation for carriers as low as 1 kHz. Although this application is fairly straightforward, it serves to point out certain general rules you should know to use LVDTs effectively.

For maximum sensitivity and a minimum of susceptibility to carrier-frequency changes, the LVDT should see a high-impedance load. In **Fig 6**, a follower-connected FET-input op amp that acts as a buffer provides this load. The op amp also provides the balanced demodulator with a low source impedance that makes the circuit insensitive to dynamic load changes.

A stable oscillator

The performance of the carrier oscillator is critical to high-accuracy balanced modem signal processing. Although absolute accuracy and long-term frequency stability are not highly critical, amplitude stability and waveform purity are.

Amplitude stability is critical because the voltage level of the carrier signal applied to an LVDT or an ac bridge directly affects the overall system sensitivity. Therefore, the carrier amplitude must be stable against loading, time, supply-voltage and temperature changes. Waveform purity is also important—high harmonic content can introduce errors, particularly with regard to the LVDT null.

The circuit shown in the **figure** satisfies these amplitude and waveform requirements. It is a spot-frequency sine-wave oscillator whose frequency can be set to a value in the 1- to 5-kHz range. It produces a 1V rms output with harmonic distortion on the order of 0.05% or less.

The oscillator uses a Wien bridge configuration built around FET-input op amp A_{1A}. The Wien network consists of frequency-determining components R_1-C_1 and R_2-C_2. An oscillator frequency of 1 or 2.5 kHz can be selected by the choice of the timing components indicated in the **figure**.

Although the Wien bridge can provide a pure sine wave, automatic gain control (AGC) is required to regulate the output amplitude. AGC is provided by first having A_3, an AD536J rms/dc converter, detect and convert A_{1A}'s output. Op amp A_{1B} then compares A_3's rectified output to a dc reference voltage. In the process, A_{1B} integrates the input and supplies the resultant output as the gate voltage for Q_1.

A_4, an AD584, supplies the dc reference voltage to which A_3's output is compared. It applies a stable +2.5V dc to the amplitude-controlling potentiometer R_8. With a 2V control-voltage potential applied to R_6 from the arm of R_8, the loop develops a 1V rms output from the AD644 oscillator stage A_{1A}, which is available at TP_1.

So as not to disturb the automatic control loop with external loading effects, A_2 buffers the output. This design also lets you adjust the gain of the buffer stage to accommodate different output levels while the oscillator stage itself operates at a fixed low level for best stability and lowest distortion. This oscillator can operate with supply voltages ranging from ±5V to ±15V, but best performance occurs at ±15V.

Although the inherent stability of the AGC system nulls the differences in individual FET devices, there still may be applications where the carrier amplitude error might need to be monitored. With this circuit, such an output (TP_2) is provided from the AD536J's internal buffer at pin 6.

The dc voltage at TP_2 is proportional to the ac carrier amplitude to within 0.5%. It can be used simply for monitoring the carrier or as a reference signal for ratiometric-balanced detection. The latter use is much more effective than the former for system error minimization because an appropriately scaled/filtered version of this signal can be used as an A/D converter reference voltage, and thus the A/D converter provides ratiometric conversion.

This op amp can also serve to introduce gain scaling. If you opt for gain scaling here, low temperature-coefficient resistors with ratios that track closely yield the best gain stability. Although amplifier drift and offset at this point is not critical, gain stability is because it directly influences overall sensitivity.

Depending on the specific transducer used, a phase-trim network might be needed in the reference channel. The **Fig 7** combination lead-lag network provides a carrier reference phase trim for minimum dc output at null, and if used it should be connected between the carrier generator and the AD630 as indicated. The circuit should employ stable components, which suggest a high-resolution multiturn trimmer for R_3, and, for the other components, fixed-value capacitors and resistors with low temperature coefficients. Metal-film resistors and polystyrene capacitors offer this performance.

As previously noted, the carrier excitation voltage applied to the LVDT directly influences the overall system sensitivity. Therefore, the carrier generator should provide a stable, low-distortion sine wave (see **box**, "A stable oscillator") and should be ac coupled with the LVDT, as **Fig 6** shows, to eliminate any possible primary dc in the transducer; any dc voltage here could result in nonlinearity or catastrophic faults.

Use balanced modems with ac bridges

The oscillator requirements also apply to another widely used class of transducers: bridge-type sources. For bridge transducers that can accommodate ac excitation as well as dc, balanced demodulation of the amplified ac output suppresses a host of problems, including the drift and noise of the (usually) necessary preamplification.

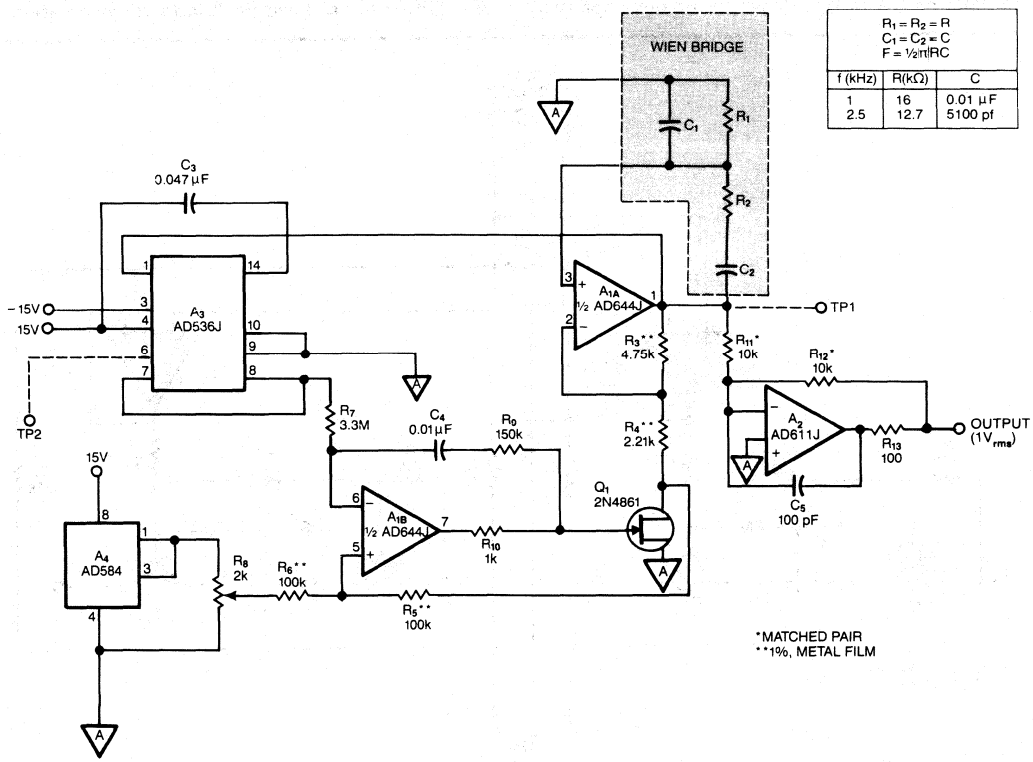

Based on a Wien bridge *that supplies a sine wave with little distortion, this oscillator circuit also employs automatic gain control to stabilize amplitude.*

AC bridge transducers work best with balanced inputs

Fig 8 illustrates an ac signal-conditioning system using a 1-kΩ/leg bridge, driven by a balanced 20V p-p, 1-kHz signal. An AD524B instrumentation amplifier (**Ref 4**) serves as a 1000× gain block (A_1), which is followed by a balanced demodulator with a gain of ±2 and by the 10-Hz lowpass filter used with the LVDT. Phase trim, if needed, is performed using **Fig 7**'s network.

Driving the bridge with balanced signals maximizes bridge sensitivity—enabling the system to resolve bridge unbalances of 1 ppm or less—and minimizes unbalance problems caused by stray capacitance. It also minimizes the common-mode voltage presented to the bridge preamp, thus enhancing input dynamic range. Remember that those preamp common-mode errors in quadrature with the carrier are nulled in the balanced demodulation process.

Fig 9 shows a push-pull circuit that provides balanced drive signals for the bridge. This driver develops two 10V p-p waveforms 180° out of phase when fed a 1V rms input from an oscillator. Because ac bridge excitation, like that of LDVTs, should be stable for best overall sensitivity, the **Fig 6** circuit's oscillator, described in the nearby **box**, should prove suitable for this application. The ratiometric option for the oscillator can be equally useful with bridge transducers.

The oscillator's output is fed to two halves of a dual op amp, with each half buffered by a 75-mA bipolar output stage. Although shown in **Fig 8** as driving a 1-kΩ/leg bridge (a load not directly suitable for typical IC op-amp loading), this buffer can drive even lower impedance bridges, down to 300Ω/leg. Note that the feedback resistors (R_1 through R_4) for the two driver amplifiers should be stable types with low temperature coefficients that track. This type of temperature-coefficient performance can best be realized using a single resistor-array device.

In addition to balanced modulation/demodulation op-

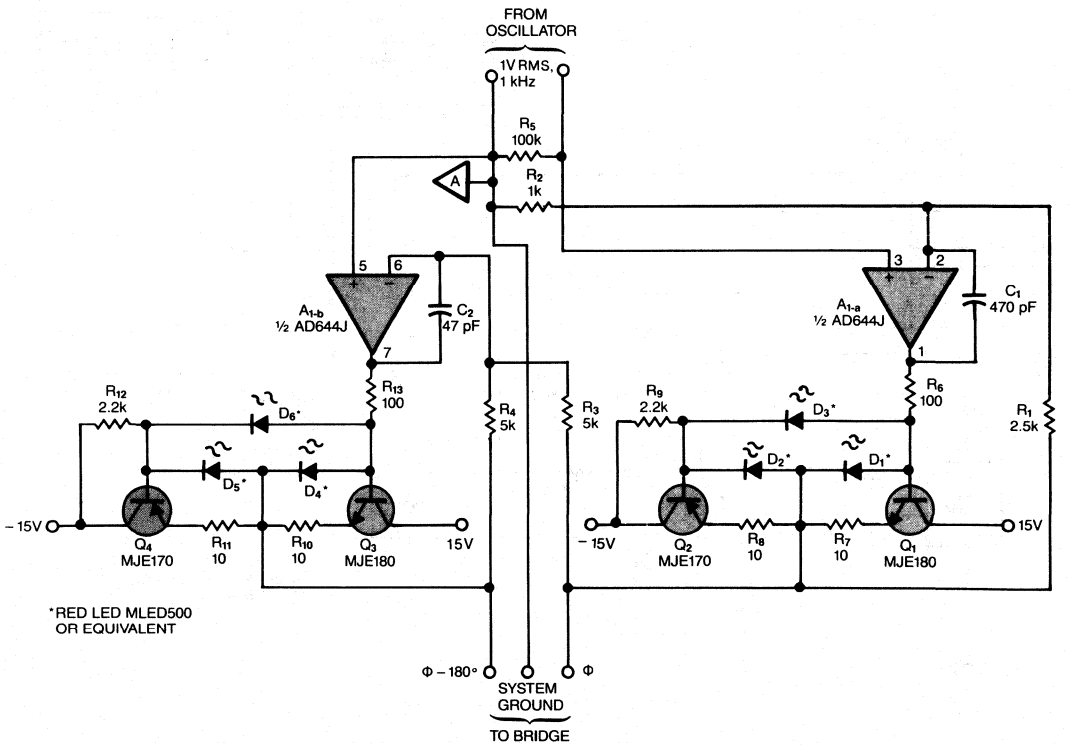

Fig 9—To drive the ac bridge *in Fig 8, this driver circuit produces two 5V peak sine waves that are 180° out of phase. Using these balanced signals maximizes system sensitivity.*

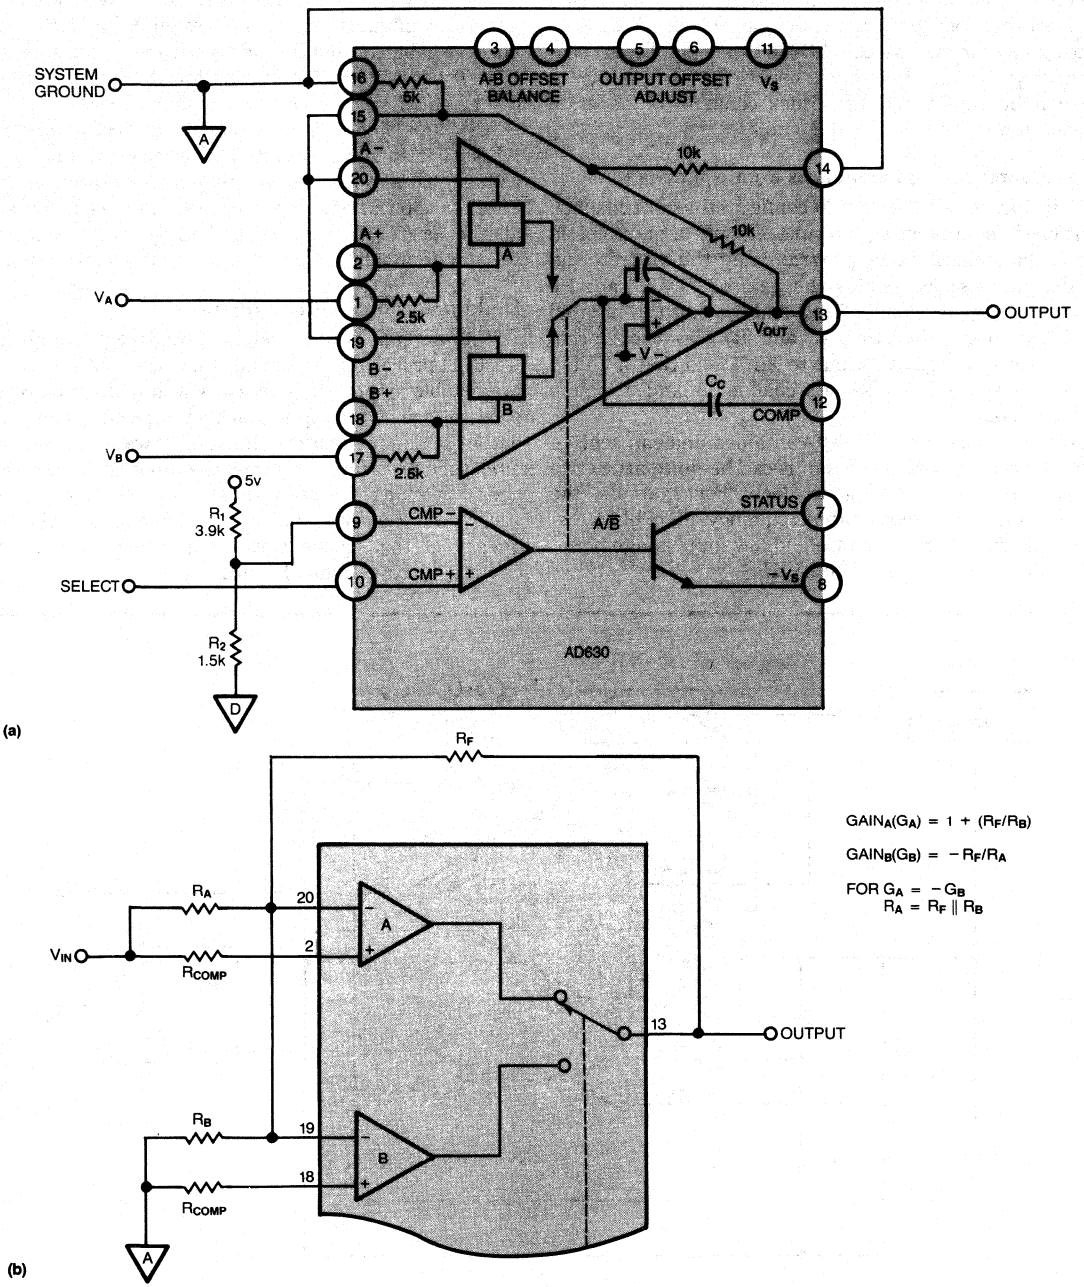

Fig 10—A dual-channel gain block *can be built using the AD630's internal resistors* ***(a)*** *or with external resistors* ***(b)*** *for a greater range of gains. In either case, you can select the desired port with a TTL-level input by adding the R_1-R_2 network* ***(a)*** *to the negative input of the chip's comparator.*

eration needed for transducers such as LVDTs and ac bridges, instrumentation applications typically need a general-purpose programmable-gain block. A commutating modem such as the AD630 can also fill this need, reducing the variety of parts needed for an application. It can be configured for either common or separate signal inputs to the A and B stages.

Use commutating modems as a gain blocks

Fig 10a shows the AD630 configured as a 2-channel, buffered, noninverting multiplexer with a gain of 4 on both channels. Jumper programming the on-chip gain resistors as shown yields the gain of 4; you can also jumper program gains of 1, 2 and 3 using the internal resistors only. Resistors R_1 and R_2 bias the on-chip comparator's negative input to +1.4V so that a TTL 1 at the positive input selects channel A and a logic 0 selects channel B.

Because this circuit is self-contained and can achieve both high ac and dc precision, it can be quite attractive as a basic "no-hassle" gain block. For example, the device makes a good audio-frequency gain block because it has a high slew rate and low input-stage noise, typically on the order of 10 nV/Hz$^{1/2}$.

With a few external resistors, a more extensive variety of gains is possible. **Fig 10b** shows this more general configuration. In this circuit R_F and R_B set a positive gain, G_A, and R_F and R_A set a negative gain, G_B. **Fig 10** notes the relationships for these gains, which can be equal on both channels or unequal.

Although it might not be obvious at first glance, the AD630 is also uniquely suited to precision timing applications that might be needed in instrumentation as well as other applications. It is so suited because it contains all the active circuitry needed to generate a square wave with amplitude and frequency characteristics.

Build precision timing circuits

With a commutating modem, a relatively straightforward approach to ensuring that the output has a predictable amplitude is to use a well-defined dc reference voltage as the input. Such an approach causes the output of the modulator (the square wave) to have an amplitude precisely ±K times the input reference voltage. Switching between these two states creates a square-wave output, and the switching rate (the frequency of the square wave) depends on the timing circuitry used to control switching.

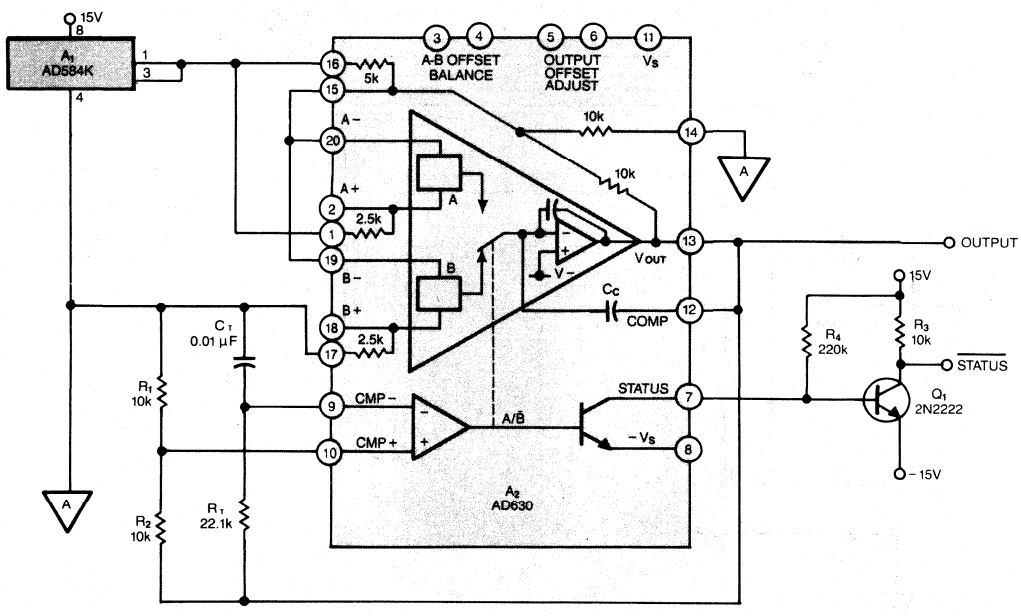

Fig 11—To produce a square wave *with a precise, 10V p-p amplitude, the AD630 switches the polarity of the AD584K's precision 2.5V dc output and multiplies it by a gain of 2. The square wave's frequency is set by C_T and R_T; in this case, it's about 2 kHz.*

In **Fig 11**, the AD630 is hooked up as a switched-gain amplifier, with an absolute gain, K, of 2 set as in **Fig 4**. With the gain set at 2, the peak output is 2 times V_{REF}, and the peak-to-peak output equals $4V_{REF}$. An AD584 precision reference-voltage source furnishes V_{REF}; connecting its pin 1 to pin 3 straps it for a +2.5V dc level, yielding the 5V pk, or 10V p-p, circuit output. In choosing a V_{REF} source, be certain that it has a low output impedance to minimize possible side effects from the dynamic loading of the switched-input AD630.

The applied reference voltage, V_{REF}, and the particular gain setting programmed into the AD630 determine the accuracy and stability of the output's amplitude. The basic tolerance applicable to an AD584K reference is 0.12%. Using an AD630AD (or AD630JN) strapped for a gain of 2, the output's amplitude error is only 0.1% greater than that of the reference used, indicating less than 0.25% overall (untrimmed) output-amplitude tolerance for the circuit. This level of precision allows the circuit to be useful as an amplitude calibrator.

The two networks at pins 10 and 9 (the comparator inputs) define the circuit's switching characteristics. The R_1-R_2 resistor network provides positive feedback, and the R_T-C_T network determines the timing delay in switching between states. The timing expression for the circuit is:

$$f = 1/(2.2R_TC_T).$$

As with all such RC-time-constant oscillators, the predictability and stability of this circuit are only as good as those of the components used, particularly those in the external timing networks. Both the R_1-R_2 and the R_T-C_T networks should have low temperature-coefficient components; metal-film resistors should be used for R_1, R_2, and R_T, and C_T should be a low-dielectric-absorption polystyrene or polypropylene film capacitor **(Ref 5)**.

Using the components shown, the frequency of the output square wave is approximately 2 kHz. The circuit operates with the best predictability and accuracy below 10 kHz, but it can be used with minor degradation to frequencies as high as 100 kHz.

A virtue of the timing scheme used in this circuit is that a comparison is done between a fraction of the output voltage and an exponential timing ramp derived from it. This basic scheme, popularized by the ubiquitous 555 timer **(Ref 5)**, provides high immunity to changes in output frequency with changes in output amplitude. The practical advantage of this scheme is that the reference voltage can be programmed for different amplitudes, and the AD630 can be strapped for different gains without disturbing the nominal operating frequency.

Although the main square-wave output from the AD630 provides ±5V in this circuit, the status output of the AD630's comparator is also available. This output, buffered by a discrete npn transistor as shown, provides an inverted status signal, which swings from +15 to −15V and can be interfaced easily to logic stages.

References

1. Gerstenhaber, M, and Brokaw, A P, "A Monolithic Balanced Modulator-Demodulator for Instrumentation," *ISSCC Digest of Technical Papers*, IEEE, New York, February 1982, pgs 44-45, 288 and 289.

2. Brokaw, P, Gerstenhaber, M, and Miller, S, "Fast, Flexible Switched Dual-Input Op Amp and Comparator," Analog Dialogue, Volume 17-1, Analog Devices, Norwood, MA, 1983.

3. Herceg, E E, *Handbook of Measurement and Control*, revised edition, Schaevitz Engineering, Pennsauken, NJ, 1976.

4. Wurcer, S, and Jung, W, "Instrumentation amplifiers solve unusual design problems," *EDN*, August 4, 1983.

5. Jung, W G, *IC Timer Cookbook*, 2nd Edition, Howard W Sams & Co, Indianapolis, IN, 1983.

6. Meade, M L, *Lock-in Amplifiers: Principles and Applications*, Peter Peregrinus Ltd, London, 1983.

ANALOG DEVICES

AN-308
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Commutating Amp Multiplies Precisely

by Moshe Gerstenhaber and Frank J. Ciarlone

By using a pulse-width-height modulation technique, the circuit in **Fig 1** implements a 0.015%-accurate multiplier. The circuit's output equals $V_X V_Y/10$. An AD581 voltage reference, an AD630 commutating amplifier, and an integrator comprising an AD707 op amp, 2000-pF capacitor, and 150-kΩ resistor first generate a precision triangle wave. For a given state of the AD630's output—$+V_{REF}$ at TP_1, for example—the integrator ramps until its output reaches $-11V$. Then, TP_1 changes state and the integrator begins ramping toward $+11V$. The triangle wave's period is 4.4RC or 1.32 msec, where R and C are the values of the integrator components.

The circuit uses a second AD630 driven by the variable V_X to compare the triangle waveform at TP_2 to the signal at V_Y. The duty cycle, T_1+T_2, at the output of this second commutating amplifier is as follows:

$$T_1 = 2RC(11 - V_Y)/10, \text{ and}$$

$$T_2 = 2RC(11 + V_Y)/10.$$

During T_1, the voltage at TP_4 equals $-1.1V_X$. During the remaining period, T_2, the pulse height will equal $+1.1V_X$. V_{OUT} is the average, obtained by lowpass filtering, of this T_1 and T_2 combined waveform and equals

$$V_O = \frac{-1.1\ V_X T_1 + 1.1 V_X T_2}{T_1 + T_2} = \frac{V_X V_Y}{10}.$$

You can use a higher bandwidth filter and a higher carrier frequency to build a faster multiplier.

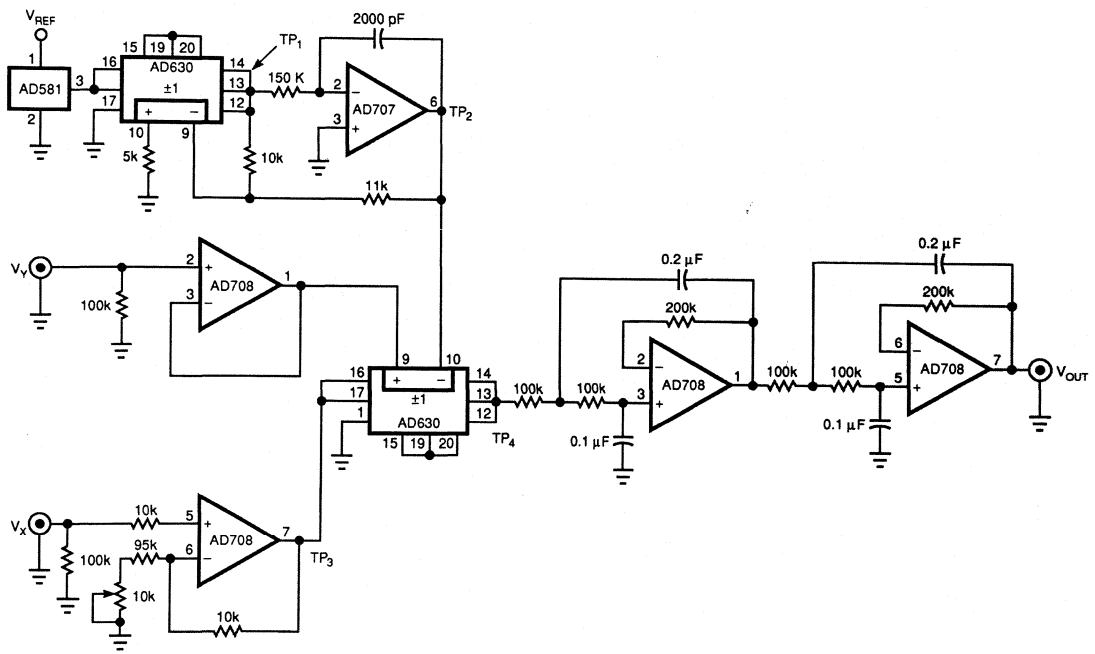

***Fig 1—Two commutating amplifiers join** a reference, an integrator, and a 4-pole filter to implement a 0.015%-accurate multiplier.*

Reprinted from EDN — October 25, 1992

ANALOG DEVICES

AN-309
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Build Fast VCAs and VCFs with Analog Multipliers

by Barrie Gilbert, Charles Kitchin and Ken Weigel

By using a high-speed analog-multiplier IC, you can accomplish signal-processing tasks that are beyond the reach of current digital technology. One such chip allows you to design megahertz-range voltage-controlled amplifiers and filters.

If you understand the circuitry and operating principles of a high-speed, monolithic, dual-channel analog multiplier, you can realize useful signal-processing functions that are difficult or impossible to accomplish when using digital techniques. For example, you can design a variety of voltage-controlled-amplifier and filter circuits based on the 60-MHz (typ) AD539 analog multiplier, and by carefully choosing companion operational amplifiers for such circuits, you can exploit to best advantage the multiplier IC's speed and accuracy.

Because the multiplier's signal paths are almost transparent (ie, they introduce little signal delay or noise), you can use the IC to build voltage-tunable filters that operate at much higher frequencies than were formerly possible. These filters exhibit wide-range, stable Q factors and avoid the clock noise associated with switched-capacitor filters. As an example, you can use the multiplier to configure a 0.12%-linearity voltage-controlled filter with center frequencies as high as 100 kHz.

A basic 2-quadrant multiplier

Before looking at specific applications for the multiplier, consider first how it performs its multiplication function. The AD539 has two signal inputs (V_{Y1} and V_{Y2}), a common control input (V_X), and a signal ground.

Reprinted from EDN — October 18, 1984

For the moment, however, treat the multiplier as if it had a single input, V_Y (**Fig 1**); the basic operation is the same for each multiplying channel of the device.

The common control input, V_X, controls a variable gain resistance, R_G, and the gain of the circuit is simply $-R_Z/R_G$. A fixed transresistance is formed by the op amp and R_Z, where R_Z is a feedback impedance built into the AD539. For $V_X>0$,

$$R_G = \frac{V_u}{V_w} \times R_z,$$

where V_W is the output voltage and V_U is a propor-

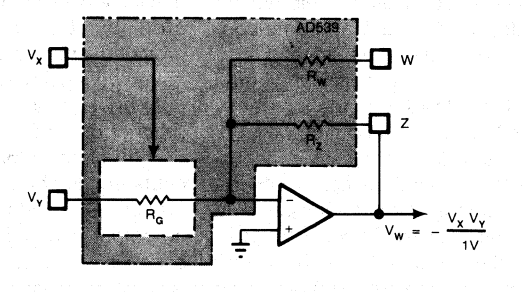

Fig 1—A basic 2-quadrant multiplier, *in conjunction with an op amp, multiplies two input voltages V_X and V_Y to produce output voltage V_W.*

tioning factor called the scaling voltage. The multiplier's transfer characteristic is

$$V_w = \frac{-V_x V_y}{V_u} \text{ for } V_x > 0.$$

When V_X is equal to or less than zero, R_G becomes an open circuit and the gain is zero. (The AD539 is designed in such a way that the operational amplifier in **Fig 1** is external to the multiplier; thus, you can select from a wide range of available devices the op amp that best suits your application. You need one op amp for each channel of the multiplier.)

The multiplier IC is not quite as simple as **Fig 1** shows, however. Input V_Y is buffered before it's applied to R_G, resulting in a constant input resistance of about 350 kΩ. The buffer has a limited signal range (specified as ±2V but usable to ±4V in many cases). Input V_X presents a relatively low input resistance of 500Ω, and it has a full-scale range from zero to +3V with a 10% overrange capability. Values of V_X below zero will not cause signal feedthrough, but they could lead to control feedthrough in some applications.

The other resistance in the gain equation, R_Z, has a nominal value of 6 kΩ. To improve versatility, an identical resistor, R_W, halves the gain (ie, it doubles the

The past and present of multipliers

Inexpensive monolithic multipliers first appeared around 1970, following the discovery of the translinear principle, which exploits the precise logarithmic properties of bipolar transistors. The first generation of such multipliers included simple diffused circuits of low accuracy (2 to 5% after manual trimming), and they required many external components.

As the art of multiplier design matured, it became possible to put the complete function—including input interfaces, output amplifier and scaling reference—on one chip, using stable thin-film resistors. With such devices, you could strap pins to perform multiplication and division and to calculate squares and square roots. It was still necessary, however, to compensate manually for three offset voltages and to calibrate the scale, an awkward and time-consuming chore. The limitations of the supporting circuitry prevented the full utilization of the inherently high bandwidth (about 50 MHz) of the multipliers' cores.

These limitations faded with the third generation of multipliers in the mid-1970s. These were individually calibrated by laser-trimming the thin-film resistors while the circuits were still in the wafer stage. On-line instrumentation assured high accuracy by minimizing the total error of each device. The assembled circuit met the specified accuracy in all its basic configurations. However, the bandwidth was still only 1 MHz, and the distortion and noise performance were not outstanding.

The latest generation of multipliers addresses real-world signal-processing requirements more closely than its predecessors. Static accuracy—conformity to an ideal transfer characteristic under dc conditions—is now more than 100 times higher than that of early multipliers. More important, such neglected matters as bandwidth, differential gain and phase linearity and noise are now receiving more emphasis.

Today, analog multipliers are available in many types, differentiated by the actual function provided and the application. All have two input ports and one output port, and they all generate the product of the two inputs. An important class of these circuits accepts inputs of either polarity and preserves the signs to generate an output. When both inputs have essentially linear amplitude response the device is a 4-quadrant multiplier.

A second class of multiplier accepts a bipolar input at its linearly responding input port but responds to only one polarity at its control port. In this 2-quadrant device, only the signal input can change sign. The AD539 is a 2-quadrant multiplier: Its control input (X) can be positive only (0 to 3V full scale), whereas the signal inputs (Y) can be bipolar (±2V full scale).

As is often the case, a tradeoff in performance in one area results in worthwhile improvements elsewhere. For example, in voltage-controlled amplifiers, it is unnecessary and undesirable for the multiplier to respond to both polarities at the control input because the primary objective is to control only the magnitude of the output, not its phase. Two-quadrant multipliers provide greater gain accuracy and lower noise than comparable 4-quadrant devices in such applications.

Finally, some applications—mostly in analog computing—require operation in only one quadrant. One-quadrant devices sometimes appear as log-antilog circuits and usually have high static accuracy at the expense of speed.

value of V_U to 2V) when placed in parallel with R_Z. This gain-adjustment capability is useful when you want to maintain the output swing at a value that's near the same level as the input. R_W can also be used to add an additional signal to the output.

The output from the AD539 is actually a current of nominally ±1 mA (±2.2 mA peak), which can drive a grounded resistive load as high as 500Ω at the multiplier's full bandwidth, conservatively rated at 30 MHz min. Balance between the two channels of the chip is excellent: Using precision 50Ω loads, the gains track within 0.025 dB, and phase matches within ±0.1° from dc to 10 MHz.

The use of external op amps provides other design benefits. For instance, without external op amps, you must use low-resistance loads to profit from the full bandwidth of the AD539's signal channels, but such loads limit outputs to a few hundred millivolts. Also, the loading affects the scaling; it becomes less exact, because it involves the ratio of the external load to the internal thin-film resistors (which have an absolute tolerance of about 20%).

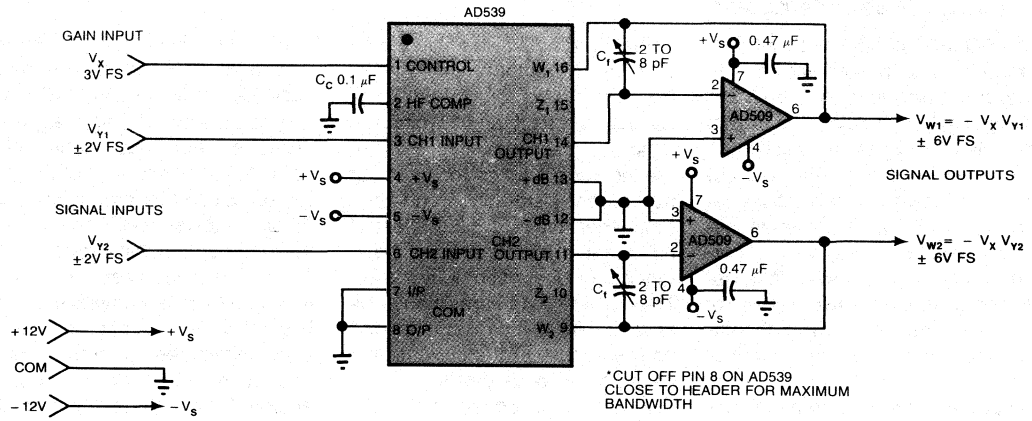

Fig 2—This dual-channel voltage-controlled amplifier *(VCA) exhibits 0.05% total harmonic disortion at 10 kHz.*

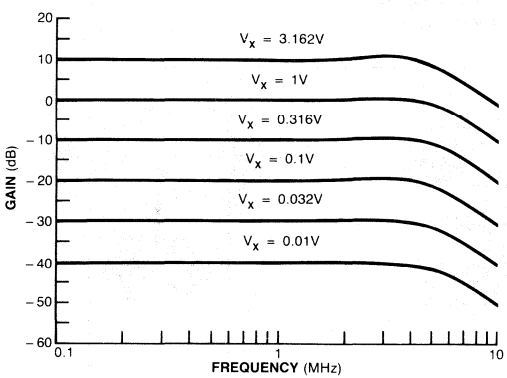

Fig 3—The frequency response of the *Fig 2* dual-channel VCA incorporating an AD509 op amp demonstrates that signal bandwidth is virtually independent of gain.

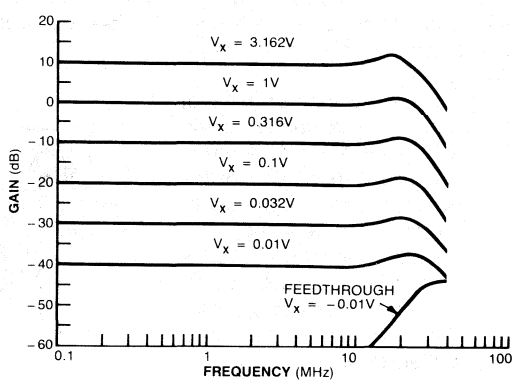

Fig 4—The −3-dB bandwidth exceeds 25 MHz *when you use the ADLH0032 hybrid op amp in the* **Fig 2** *dual-channel VCA.*

Chip multiplies signal input by magnitude of control voltage

However, when you use external op amps and the IC's internal feedback resistors R_Z and R_W to provide current-to-voltage conversion, the ratio of these matched internal resistors with respect to R_G determines the scaling and thereby guarantees high accuracy. Furthermore, the performance of the op amp primarily determines the bandwidth capability of the circuit. For example, using an AD509 op amp, the −3-dB bandwidth is about 6 MHz; using an ADLH0032 hybrid op amp, 25 MHz is possible. And in situations

Controlled-cascode design

The AD539 employs a controlled-cascode design that allows it to operate at frequencies to 60 MHz. In contrast, a typical 2-quadrant multiplier **(Fig A)** consists of a translinear cell (Q_1 through Q_4) and a voltage-to-current converter (Q_5, Q_6 and R_Y). Translinear circuits operate entirely in the current mode; their behavior is largely dependent on the ratios of currents, which result from the linear dependence of transconductance on current.

Input voltage V_Y determines the value of a pair of complementary currents through Q_1 and Q_2. Because the ratio I_{Q3}/I_{Q4} is equal to I_{Q1}/I_{Q2} and the magnitude of the output is proportional to I_X, the overall circuit behaves as a controlled transresistance of value $R_Y I_Y/I_X$.

There are several sources of distortion and noise in this circuit:

- V_{BE} mismatch in Q_1 through Q_4 generates even-order distortion, the magnitude of which is independent of I_X.
- Junction resistance causes odd-order distortion, which is zero at only one I_X value.
- Junction resistances also generate a noise component proportional to I_X.
- Collector-base capacitances at Q_3 and Q_4 leak high-frequency signals to the output even when the dc gain is zero.

A controlled-cascode design addresses these problems by changing the location of the output **(Fig B)**. Although the differential signal current is still applied to Q_1 and Q_2, the output is now taken from their collectors. When $I_X=0$, Q_3 and Q_4 are off, and the circuit is simply a cascode circuit. There are no mechanisms for even- or odd-order distortion generation other than the negligible effect of nonlinear alpha. The base resistances no longer contribute noise. The effective transresistance is now simply R_Y; it does not depend on the accuracy of a current ratio as in the previous circuit.

When I_X is equal to I_Y, the source currents are completely diverted from Q_1 and Q_2, so the output shuts off. The high-frequency feedthrough is now very small because C_{BC} of Q_1 and Q_2 no longer connect the input to the output. For values of I_X between zero and I_Y, the effective transresistance is $R_Y/(1-(I_X/I_Y))$. Consequently, the circuit behaves like a cascode that has controllable gain. Analysis shows that the distortion caused by area mismatches and resistances is much lower than that of the first circuit.

The controlled cascode used in the AD539 is a derivative of **Fig B**'s circuit. It's a configuration that also eliminates a troublesome peak in high-frequency response arising from C_{BC} of Q_3 and Q_4. Furthermore, the design of the cell transistors achieves a balance between the conflicting requirements of low- and high-frequency accuracy.

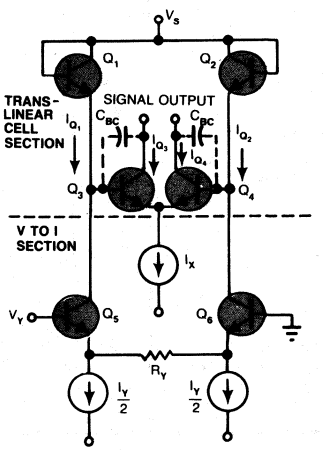

Fig A—A typical 2-quadrant multiplier cell *has several sources of distortion and noise.*

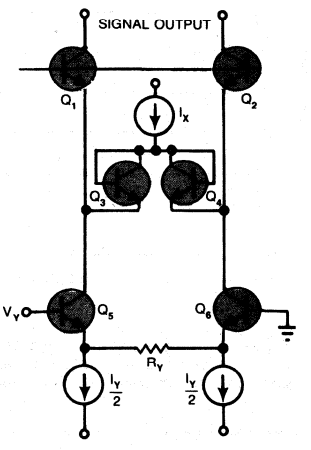

Fig B—The controlled-cascode multiplier cell *minimizes distortion by furnishing a signal output at the collectors of Q_1 and Q_2.*

where bandwidth is not an issue, many low-cost op amps are sufficient. An AD611, for example, provides low distortion and a full ±10V output at frequencies as high as 200 kHz.

Design a voltage-controlled amplifier

Fig 2 shows a dual-channel amplifier that has voltage-controllable gain. This gain is numerically equal to the control voltage V_X in the range 0 to +3.3V. The signal bandwidth is essentially independent of gain, as shown in **Fig 3.** In addition, the control-channel bandwidth is about 5 MHz (for a V_X rated from 300 mV to 3V), using a 3-nF HF compensation capacitor (C_C). When rapid gain control is not required, C_C should be 0.1 μF; this value reduces high-frequency distortion in the signal channels. In addition, the ±12V supplies allow the circuit to realize a full-scale output swing of ±6V with some overrange. To avoid control breakthrough, V_X should not vary substantially from the nominal control range.

Looking at the performance of this voltage-controlled amplifier (VCA), the total harmonic distortion for an input of 1V rms is typically below 0.05% at 10 kHz, for a V_X above 300 mV. The output noise for a V_X of 1V is 50 μV rms from 10 Hz to 10 kHz and 550 μV from 10 Hz to 5 MHz. Furthermore, the signal leakage when V_X is deliberately made slightly negative (say, −10 mV) is −75 dB at 5 MHz.

If you need a higher bandwidth, you can use a hybrid op amp—for example, the ADLH0032. **Fig 4** shows the high-frequency response with C_F adjusted for 1 dB of peaking; under these conditions, the −3-dB bandwidth exceeds 25 MHz. Careful board layout and supply decoupling keep signal leakage below −90 dB at low frequencies and below −60 dB at 20 MHz when V_X is −10 mV.

To meet more stringent design requirements, you can employ an ultralow-distortion single-channel VCA **(Fig 5).** In this design, the input connects directly to channel 1 and to amplifier A_1. This amplifier inverts the

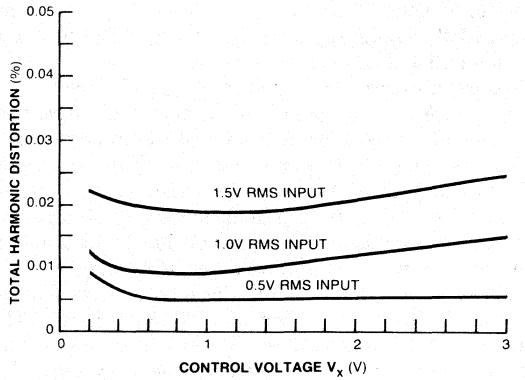

Fig 6—Less than 0.01% total harmonic distortion *is achievable with the* ***Fig 5*** *ultralow-distortion VCA.*

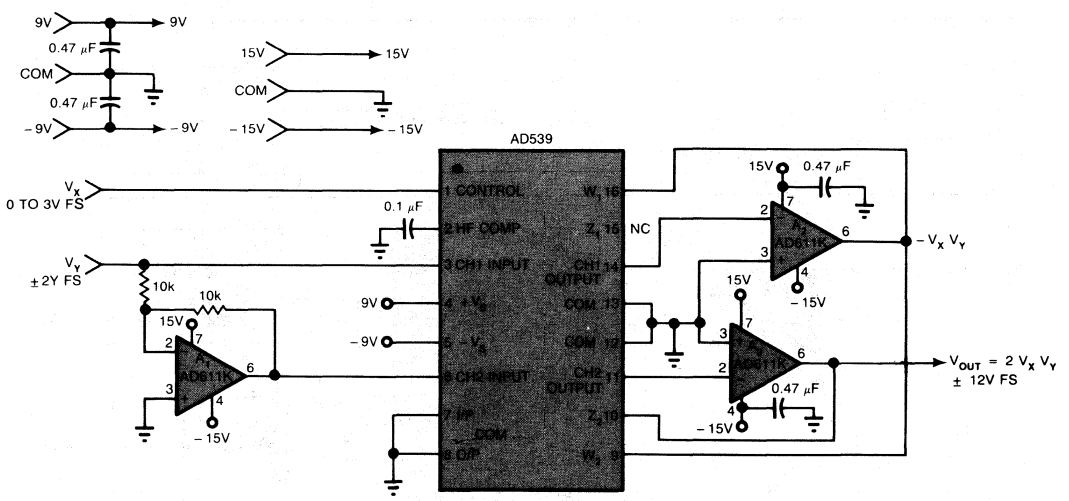

Fig 5—An ultralow-distortion VCA, *whose output appears at A_3's output, uses both of the AD539's multiplier channels to achieve distortion cancellation.*

signal and drives the AD539's channel 2. Driving the multiplier in symmetrical fashion minimizes distortion.

Note that this circuit provides two outputs. The output of channel 1's amplifier, A_2, is $-V_XV_Y$; it can be used directly and has the same characteristics as the **Fig 2** VCA. This output is also used to obtain the low-distortion A_3 output, which is equal to $2V_XV_Y$.

For the low-distortion output, A_2's output connects to one of channel 2's scaling resistors, W_2; the other channel 2 scaling resistor, Z_2, provides the feedback path for channel 2's output amplifier, A_3. Because these internal resistors are matched, A_3 simply inverts the output of A_2. At the same time, it performs current-to-voltage conversion for channel 2 of the multiplier and adds the result to the inverted A_2 output. Therefore, the total output of A_3 is $-(-V_XV_Y)$ from channel 1 and A_2 plus $-(-V_XV_Y)$ from A_1 and channel 2; this amounts to $2V_XV_Y$. Amplitude limits on this low-distortion output arise from the supply voltages of A_2 and A_3.

Fig 6 shows that the total harmonic distortion can remain below 0.01% over a wide gain range. The increase in distortion for small values of V_X apparently stems from noise in the measurement system. The balanced configuration of this VCA eliminates control feedthrough when V_X is outside its nominal range.

Many variations of this design are possible, with gains as high as 60 dB. What's more, a series connection of the two channels creates a circuit that has a square-law gain response. This function is useful, for example, in swept-gain applications; it compensates for inverse-square-law propagation-path losses.

Two-quadrant multipliers also find profitable application in the construction of voltage-tunable filters. In such a design, a loop consisting of two integrators and

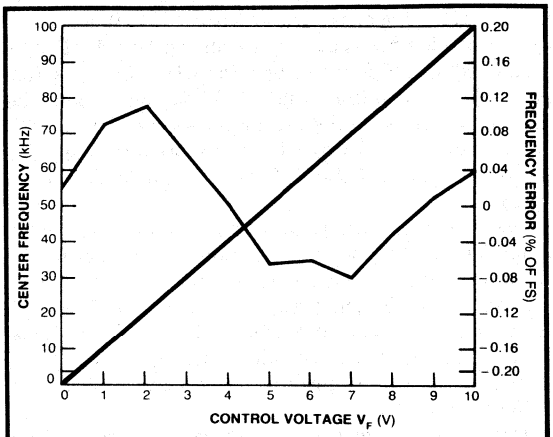

Fig 8—In response to a 0 to 10V control-voltage input, *the* ***Fig 7*** *voltage-tunable filter exhibits a 0- to 100-kHz center frequency. Frequency error remains less than 0.12%.*

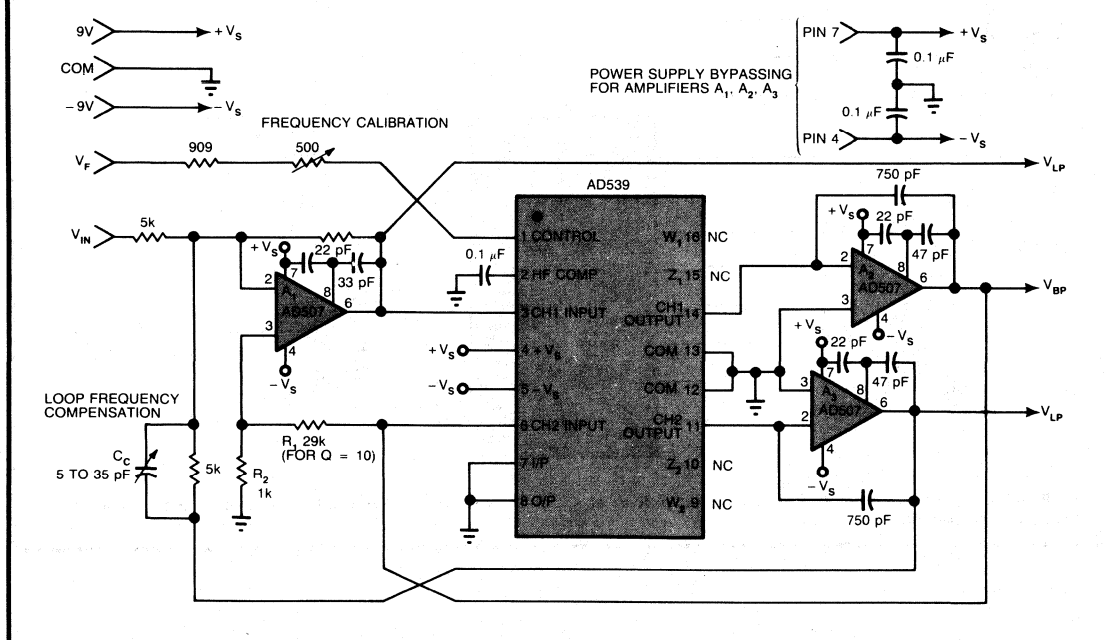

Fig 7—A voltage-tunable state-variable filter *changes center frequency in proportion to variations in the control voltage* V_F.

an inverter (used also as a summing amplifier) forms the basic filter. The filter has lowpass, highpass and bandpass responses, and it features a variable Q over a wide range. Furthermore, the low offset voltage of the control channel ensures accuracy over a much wider range of control voltages than is possible using a typical 4-quadrant multiplier.

In the integrators, pairs of capacitors and resistors determine the resonant frequency. To control this frequency, it's customary to vary the values of the resistors. Because the channels of an AD539 behave like voltage-controlled resistances, they allow you to create voltage-tunable filters.

A 2-pole voltage-tunable filter section (**Fig 7**) provides a center frequency of nominally 0 to 100 kHz, easily scaled to other ranges by the choice of capacitors (which should, however, always be equal). Each channel of the AD539 behaves like a voltage-controlled transresistance, having a value of

$$R_G = \frac{6\ k\Omega}{V_x / 1V}.$$

The absolute value of R_G is not exact; only the ratios of the on-chip resistors undergo trimming during manufacturing. Accurate, temperature-stable calibration of the frequency requires the addition of an external trim resistor, placed in series with the (resistive) control input. When the circuit is thus adjusted to compensate for the tolerance of the thin-film resistors, the resonant frequency is $f = 10 \times V_F/1V$ (f is in kilohertz, V_F in volts).

Fig 8 shows the measured center frequency of the voltage-tunable filter and the deviation from the ideal value. The deviation is less than 0.12% of the full range. With the values shown for R_1 and R_2, the nominal Q is 10, and it varies by only a few percent over the 1- to 100-kHz range. In general,

$$Q = \frac{R_1 + R_2}{3R_2}.$$

The minimum Q, 0.33, occurs when the value of R_2 approaches infinity; the maximum Q is limited in practice by dynamic-range considerations. Capacitor C_C helps to maintain loop stability and is adjusted for minimum variation of Q with V_F. If a higher bandwidth is necessary, high-speed op amps make feasible a

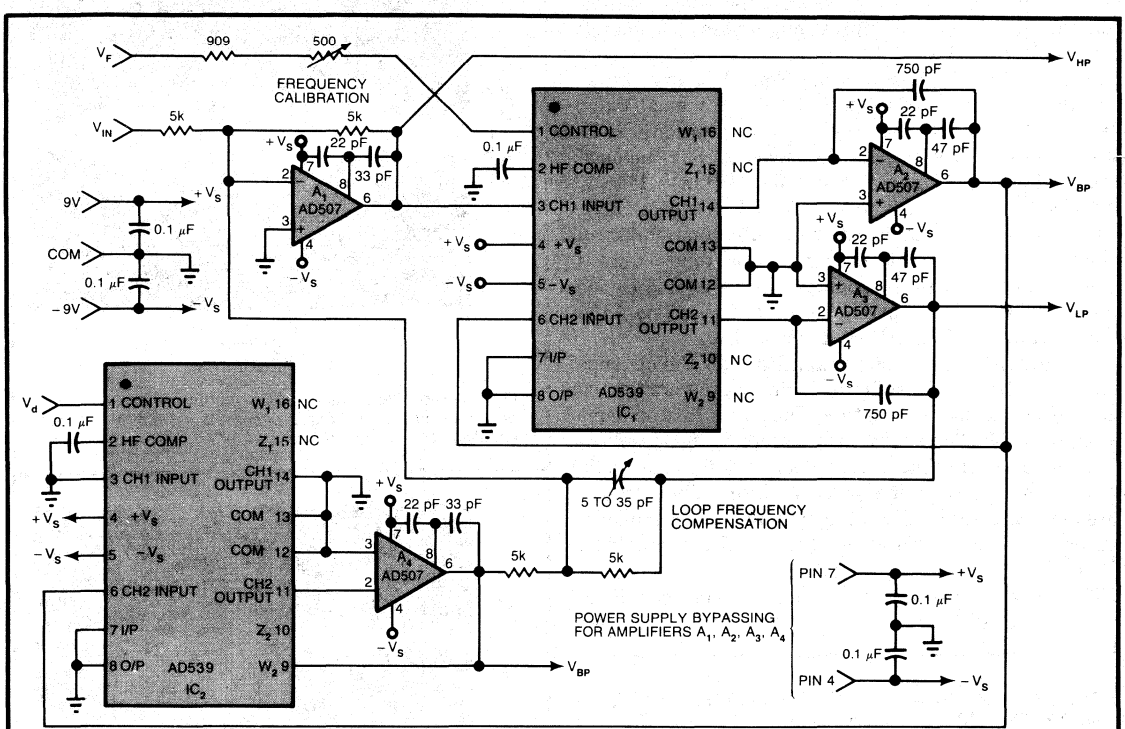

Fig 9—This variable-Q voltage-tunable filter *uses one multiplier to vary center frequency in response to V_F and a second one to vary the Q in response to changes in V_d.*

Distortion for 1V rms input is typically below 0.05% at 10 kHz

full-scale frequency of 1 MHz.

In any case, there are limits to the input amplitude. Note that the output of A_2 is Q times larger than V_{IN}. This output drives the channel 2 input of the AD539, whose signal capacity is about ±4V. Consequently, V_{IN} must be less than ±4V/Q to avoid clipping.

A more sophisticated voltage-tunable filter provides voltage control of Q. The **Fig 9** circuit incorporates two AD539s, one of which provides this function. The frequency scaling remains 10 kHz/V of V_F, but the damping (d) is now equal to the numerical value of the control voltage, V_d. Consequently,

$$d = V_d,\ Q = \frac{1}{V_d}\ (V_d \text{ in volts}).$$

If you need to alter the scaling of Q, connect a resistor in series with pin 1 of IC_2. You can use the bandpass output from the output of either A_2 (in which case the gain is equal to Q) or A_4. This second connection has the advantage of providing a gain that doesn't vary with the value of Q.

For the **Fig 9** filter design, **Fig 10** shows the ac output for V_d values of 100 mV, 300 mV, 1V and 3V (for Q=10, 3.3, 1 and 0.33, respectively). In addition, **Fig 11** shows the stability of Q as V_F varies over its nominal range. Here, Q=10 and V_F=100 mV, 300 mV, 1V, 3V and 10V (for f=1, 3, 10, 30 and 100 kHz, respectively).

Op-amp offset voltages can be troublesome for small values of V_F, when the dc behavior tends toward open-loop conditions. In applications requiring a wide range of frequency control, it's necessary to trim the offset voltages of amplifiers A_1 through A_3.

In addition, compensation capacitors for the op amps must provide a safe stability margin without adding too much phase to the integrators. Such very fast op amps as the ADLH0032 require more complex high-frequency compensation, but filters designed for the audio range can use such inexpensive, internally compensated op amps as the AD611. Nominally, the frequency-compensation capacitor at pin 2 of the AD539 should be 0.1 μF.

The output of the voltage supplies in these circuits is somewhat flexible. **Figs 7** and **9** show 9V supplies for the AD539s and the AD507s. In these applications, there is no advantage in using higher voltages. You can, however, use 15V supplies.

Finally, digital control of the center frequency or Q or both requires only the addition of a D/A converter. The AD558 is a good choice for this application, having a 2.5V full-scale output. Moreover, it's capable of using the same power supplies.

The AD539-based filter design stacks up favorably against a typical switched-capacitor filter (**table**). In comparison, the state-variable filter shown in **Fig 7** exhibits 30 times the maximum center frequency, 25 times the dynamic range and five times the maximum Q at 20 kHz. The obvious tradeoff between these approaches is in current consumption: The AD539-based

TABLE—COMPARISON OF AD539-BASED FILTER WITH TYPICAL SWITCHED-CAPACITOR FILTER

PARAMETER	SWITCHED-CAPACITOR FILTER	AD539 ANALOG FILTER
MAXIMUM Q AT 20 kHZ	10	50
MAXIMUM CENTER FREQUENCY	30 kHz	1 MHz
SIGNAL DYNAMIC RANGE	52 dB	80 dB
POWER-SUPPLY CURRENT	4 mA	28 mA

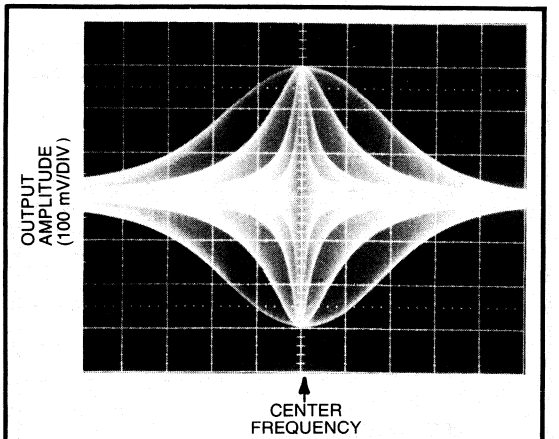

Fig 10—The wide variation of Q *apparent in these superimposed outputs of the* ***Fig 9*** *circuit stems from V_d values equaling 100 mV, 300 mV, 1V and 3V.*

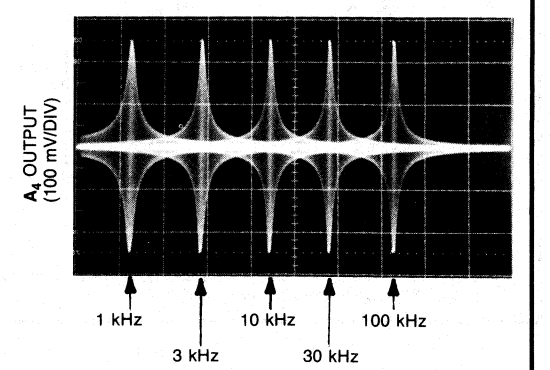

Fig 11—The stability of Q as V_F varies *is illustrated by the uniform shapes of the* ***Fig 9*** *circuit's ac response.*

filter consumes seven times the current. Nevertheless, the simplicity, accuracy and high bandwidth of the circuits using the AD539 can outweigh this disadvantage in most applications.

References

1. Meyer, R G and Sansen, W M, "Distortion in Bipolar-transistor Variable-gain Amplifiers," *IEEE Journal of Solid State Circuits,* Vol SC-8, August 1973, pg 275.

2. Gilbert, B, "A New Wideband Amplifier Technique," *IEEE Journal of Solid State Circuits,* Vol SC-3, December 1968, pg 353.

3. Tow, J, "Step-by-step Active Filter Design," *IEEE Spectrum,* December 1969, pg 64.

4. *Nonlinear Circuits Handbook,* D Sheingold, ed, Analog Devices, 1974, pg 138.

ANALOG DEVICES

AN-212
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Using the AD834 in DC to 500 MHz Applications: RMS-to-DC Conversion, Voltage-Controlled Amplifiers, and Video Switches

by Mark Elbert and Barrie Gilbert

NTRODUCTION

The AD834 is the fastest four quadrant multiplier available, having a useful bandwidth of 800 MHz, compared to the 60 MHz bandwidth of the AD539 two-quadrant multiplier, the 10 MHz bandwidth of the AD734 four-quadrant multiplier, or the 1 MHz bandwidth of the industry-standard AD534 four-quadrant multiplier. Its monolithic construction and high speed makes the AD834 a candidate for such HF applications as balanced modulation-demodulation, power measurement, gain control, and video switching, at frequencies that were previously beyond the scope of analog multipliers.

The AD834 does not sacrifice accuracy to achieve its speed. In common with all of the Analog Devices multipliers, laser trimming is used during manufacture to null input and output offsets and to establish precise scaling. In typical applications the total static error can be held to less than ±0.5%.

It is available in 8-pin plastic DIP, SOIC, and ceramic packages for the commercial, industrial, and military temperature ranges and operates from ±5 V supplies.

The main challenge in using the AD834 arises from its current-mode output stage. In order to maintain the highest possible bandwidth, the AD834's outputs are in the form of a pair of differential currents from open collectors. This is an inconvenience when a more conventional ground-referenced voltage output is needed. Thus, this application note discusses methods for the accurate conversion of these currents to a single-sided ground-referenced voltage.

These applications include a wideband mean-square detector, an rms-to-dc converter, two wideband voltage-controlled amplifiers, a high-speed video switch, and transformer-coupled output circuits. These applications provide the user with a complete and proven solution, in many cases including recommended sources for critical components.

OVERVIEW OF THE AD834

The AD834, shown in block schematic form in Figure 1, is the outcome of Analog Devices' continuing dedication to high-accuracy analog signal processing. In particular,

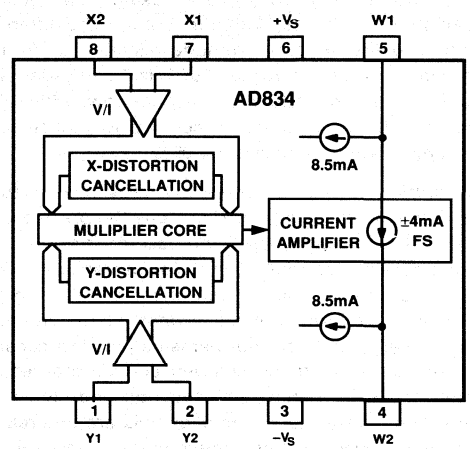

Figure 1. AD834 Block Diagram

it incorporates the experience gained in twenty years of manufacturing analog multipliers. The part is constructed on a 3 GHz epitaxial bipolar transistor process using laser-trimmed thin-film resistors. Attention to many subtle details has resulted in unusually low distortion and noise. Figure 2 shows a more detailed, but still simplified, circuit schematic.

The X- and Y-inputs are applied to high-speed voltage-to-current (V/I) converters, having a transresistance of 285 Ω and a small-signal input resistance of about 25 kΩ. The full-scale input voltage is ±1 V for both inputs. The input bias currents are typically 45 μA. Therefore, the dc resistance seen by both inputs of a differential pair must be equal to minimize offset voltages, just as for an op amp. Resistors at the inputs also minimize the risk of

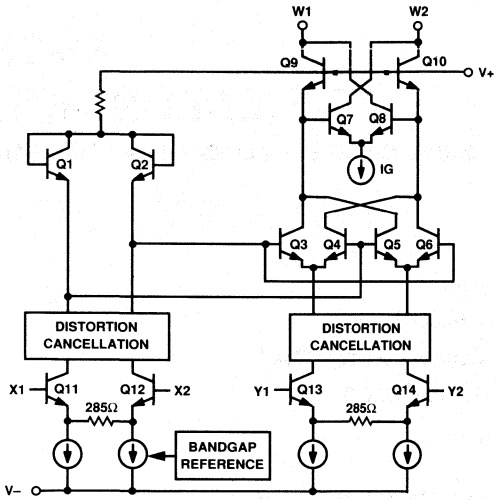

Figure 2. Simplified AD834 Schematic

high frequency oscillations. The V/I converters have a common-mode range of ±1.2 V, using the recommended supply voltages. Within that range, the differential inputs exhibit a common-mode rejection of 70 dB, conservatively specified for f < 100 kHz. Even-order distortion in the V/I converters is inherently low, while distortion cancellation circuitry is included to reduce odd order nonlinearity to typically ±0.05%.

The *multiplier core* is a well-known translinear circuit. The translinear principle [Ref. 1] exploits the precise logarithmic relationship between the base-emitter voltage (V_{BE}) and collector current (I_C) of a bipolar transistor. The input and output signals of translinear circuits are always in current form. Voltage swings at the internal nodes are very small, so that parasitic junction capacitances do not have to be charged and discharged, a common cause for bandwidth reduction and slew-rate limiting. Thus, translinear multiplier cells are inherently fast; they are also readily implemented in monolithic form. However, they can introduce distortion if not carefully designed.

This distortion is due primarily to emitter area mismatches and ohmic resistances in the core transistors (Ref. 2). Using the traditional convention in naming the channels, as shown in Figure 2, the X channel is susceptible to these effects, while the Y signal-path is essentially linear (the four output devices, Q3 through Q6, behaving in many respects like common-base stages, or cascodes). Therefore, the signal requiring the lowest possible distortion should always be handled by the Y channel. For example, in a balanced modulator application, the carrier (local oscillator voltage) should be applied to the X input and the baseband signal to the Y input.

The output from the core is in the form of a pair of differential currents. Now, the scaling of these currents is customarily controlled by adjustment of the bias currents in the V/I converter used on the X-input, which also determines the currents in the diode-connected transistors, Q1 and Q2.

In classical voltage-output multipliers, the range of adjustment needed to absorb the inevitable resistor mismatches is small, and this method of trimming the scaling factor is acceptable. In the AD834, however, the transfer function involves the two input voltages V_X and V_Y, the scaling voltage (generated in the band-gap reference circuit, and trimmed to an accurate value which is assumed here to be 1 V) and the output current, I_W:

$$I_W = \frac{V_X V_Y}{1V} \cdot \frac{I}{R} \qquad (1)$$

In this expression, the value of a resistance, R, determines the calibration of the output current. As fabricated, thin-film resistors have an initial uncertainty which can be as large as ±20%, and the customary methods of trimming the scale factor would result in other compromises (for example, erosion of the available signal range in the X-input V/I converter).

Therefore, the AD834 uses a "Gilbert gain-cell" [Reference 3] after the core to provide the needed adjustment of the effective value of R, which, in fact, is achieved by varying the current gain of this cell through trimming the current I_G. R, after the I_G trim, has an effective value of 250 Ω, resulting in a full-scale output current of ±4 mA when both inputs are at their full-scale value of ±1 V. The typical current-gain is 1.6, and because this type of amplifier is very fast and buffers the core outputs, the overall bandwidth of the multiplier is actually enhanced over that which would be obtained using the core outputs directly.

The bias currents from the core, and the gain-setting current I_G, result in a fairly large standing current—typically 8.5 mA—which flows into the outputs W1 and W2 (Pins 4 and 5). Only the *differential* output is precisely specified to be ±4 mA.

The output currents can be converted back to voltages in a variety of ways. In the simplest case, load resistors connected to the positive supply might be used, but these do not convert the (two) differential outputs to a single-sided voltage.

For the AD834 to operate properly, the output Pins (4 and 5) must be pulled above V+ to avoid saturation of Q7–Q10. To avoid using a separate supply to do this, several of the circuits included here use a voltage-dropping resistor in series with the positive supply Pin (6) of the AD834; this is a higher value than necessary for decoupling purposes.

This dropping resistor lowers the voltage at Pin 6 to provide an extra margin of bias for the output transistors. For example, in the mean-square circuit in Figure 3, 11 mA of quiescent current across the 169 Ω dropping resistor creates 1.86 V of headroom. The decoupling re-

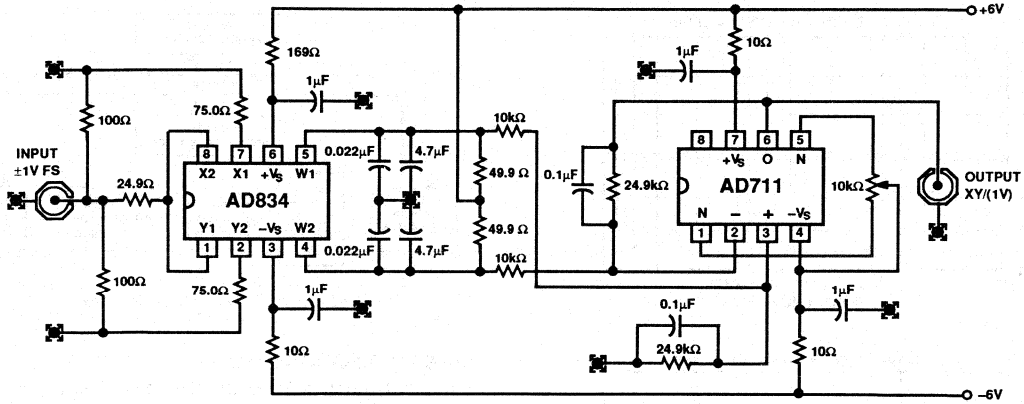

Figure 3. A DC to 500 MHz Mean Square Circuit

sistor in series with the negative supply to Pin 3 is only 10 Ω, since it is included just to decouple the supplies.

Much of this application note, however, is concerned with *more effective ways* of loading the outputs. For example, because they are fully calibrated, the outputs of two or more AD834's can be accurately summed by simply connecting them in parallel, as is done in the rms application discussed later in this application note.

MEAN-SQUARED DETECTOR

We will begin with a discussion of a mean square detector (Figure 3), whose output is a dc voltage proportional to the input power. This circuit is useful in that it requires only a calibrated signal generator and a dc voltmeter to demonstrate the very high speed of the AD834.

The input signal is applied to the X- and Y-inputs connected in parallel. The *instantaneous* output current is thus proportional to the square of the input voltage. The square of a sinusoidal input voltage of amplitude A is an offset cosine at twice the frequency:

$$A\,(sin\omega t)^2 = A^2\,(1 - cos\,2\,\omega t)/2 \qquad (2)$$

If the input to the AD834 has this sinusoidal form, then the instantaneous output current (using Equation 1) is simply

$$I_W = 2A^2\,(1 - cos2\omega t)\,mA \qquad (3)$$

the average value of which is just 2 mA for the maximum 1 V amplitude sinusoid.

The full-scale differential voltage which would be measured across Pins 4 and 5 of the AD834 is, therefore, 2 mA × (50 Ω + 50 Ω), or 200 mV. This average is extracted by the low-pass filter formed by the 4.7 μF 0.022 μF (AVX part #SR505E475MMAA and #SR505a223JAA) capacitors in conjunction with the 50 Ω collector load resistors, having a −3 dB frequency of about 650 Hz.

Two capacitors are used in parallel since the 4.7 μF capacitor uses the compact but lossy Z5U dielectric material while the 22 μF capacitor uses a high Q NPO dielectric which ensures good filtering at the highest frequencies. Note that the 4.7 μF capacitors have a −20% to +80% tolerance, so their −3 dB frequency is not accurate, nor does it usually need to be. Further filtering is performed by the capacitors in shunt with the feedback resistors of the AD711 operational amplifier, configured to have a −3 dB frequency of 65 Hz.

Due to finite averaging of the circuit, there will be some ripple for low frequency inputs. For the circuit shown, a 1 kHz input will produce the mean-square plus a −42 dB 2 kHz ripple; for 100 kHz input, the ripple will be only −80 dB. Since the output is band limited, we can use a generic low speed op amp with ample common-mode range, obviating the need for level shifting. The differential gain of the amplifier can be chosen to provide a convenient scale factor.

The full-scale gain of the circuit in Figure 3 is calculated as follows. The average output current is ±2 mA for 1 V (peak) sinusoidal input, which creates ±100 mV across each 50 Ω output load resistor or 200 mV differential. The amplifier is configured for a differential gain of 2.5 (feedback resistance over source resistance), yielding a circuit gain of 0.5 V dc output for 1 V rms input.

The bandwidth of this circuit is limited by package capacitance and inductance. In the 8-pin cerdip, the multiplier's response normally starts to rise at 500 MHz due to package resonance and peaks at 800 MHz before rolling off. A 24.9 Ω resistor at the input dampens the resonance yielding an essentially flat response out to 800 MHz. (The package inductance will be different for a surface mount AD834.) Figure 4 shows the results over frequency for three different power levels using the test configuration shown in Figure 5.

Neglecting the 24.9 Ω in series with the high impedance inputs, the input resistance to the mean square circuit in Figure 3 is 50 Ω. Since the full-scale input range is ±1 V, the maximum measurable power with a 50 Ω input load is 10 mW (20 dBm), assuming a sinusoidal input.

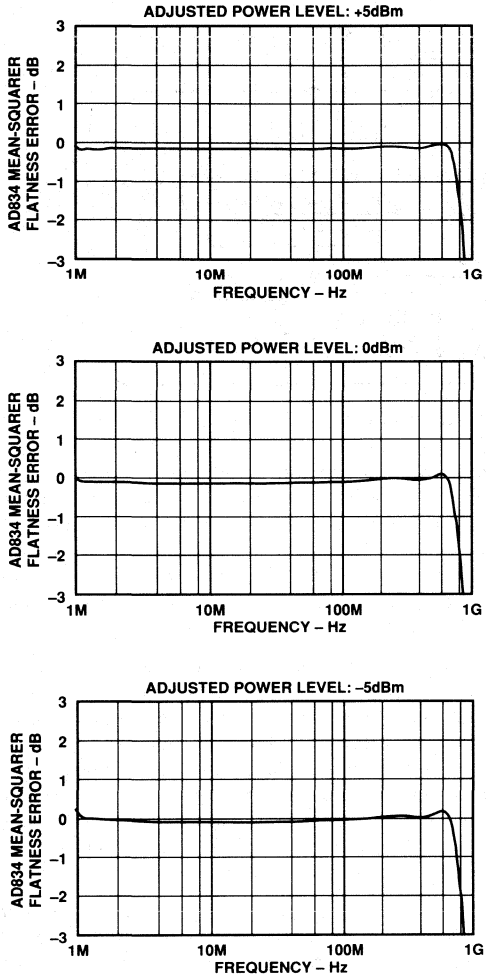

Figure 4. Frequency Response of Mean Square Circuit for Input Power Levels of −5 dBm, 0 dBm, and +5 dBm

For greater input ranges, a voltage divider with a series resistance of 50 Ω at the input will scale down the voltage seen by the AD834 while maintaining a proper termination resistance. For example, if the input signal is applied to a 45 Ω resistor in series with a 5 Ω resistor to ground, then taking the AD834's input from the middle node of the voltage divider provides 20 dB attenuation of the input signal, while maintaining a termination resistance of 50 Ω (45 Ω + 5 Ω).

Detection of low power signals is limited by dc offsets and the common-mode rejection of the op amp. For example, a −20 dBm signal, corresponding to 22.4 mV rms across 50 Ω, would result in a 4.5% error in the presence of only 1 mV of offset in the op amp. A 10% error can occur if the AD834's X channel offset is just 2 mV.

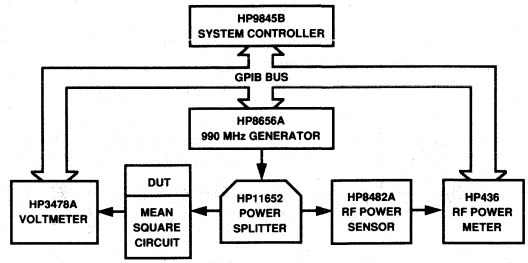

Figure 5. Test Configuration

RMS-TO-DC CONVERTER

The root mean square (rms) circuit (Figure 6) is little more than the mean square detector circuit described above followed by a square root circuit. The frequency response is determined by the front end squarer and output filter. From the mean-square discussion, the squarer functions well past 500 MHz, while the lower −3 dB frequency response is 340 Hz (100 Ω and 4.7 μF). Note that a resistor divider network at the input determines the full-scale input voltage to be ±2 V peak.

The square root function is performed by a squaring AD834 in the feedback loop of an AD711 operational amplifier. The 2N3904 transistor functions as a buffer. The resistive divider network (two 100 Ω) between the buffered output and the X and Y channel inputs of the AD834 used in the square root section determines the output scaling to be ±2 V full scale.

The outputs of the two AD834s are current-differenced. Accurate output differencing and summing is possible owing to the precision of the laser trimmed AD834 output signal current scaling. The AD711 forces the difference between the two AD834 signal current to zero. Any error in the nulling generates a voltage across the two 100 Ω pull-up resistors.

After additional filtering and level shifting by the 15 kΩ, 85 kΩ, and 0.1 μF network, the residual error is amplified by the full AD711 open loop gain. The amplified error signal forces the AD834 in the feedback loop to match its output to the mean-squaring AD834's output. The error is nulled when the rms circuit's output is equal to the square-root of the circuit's input mean-squared, hence the rms function.

The accuracy of the circuit at small signal levels is limited by inevitable offset voltages. While a nominal 0 V input with a 1 mV error to a mean-square function generates a 1 μV output error, the same input error generates a 31.6 mV output error through a square root circuit.

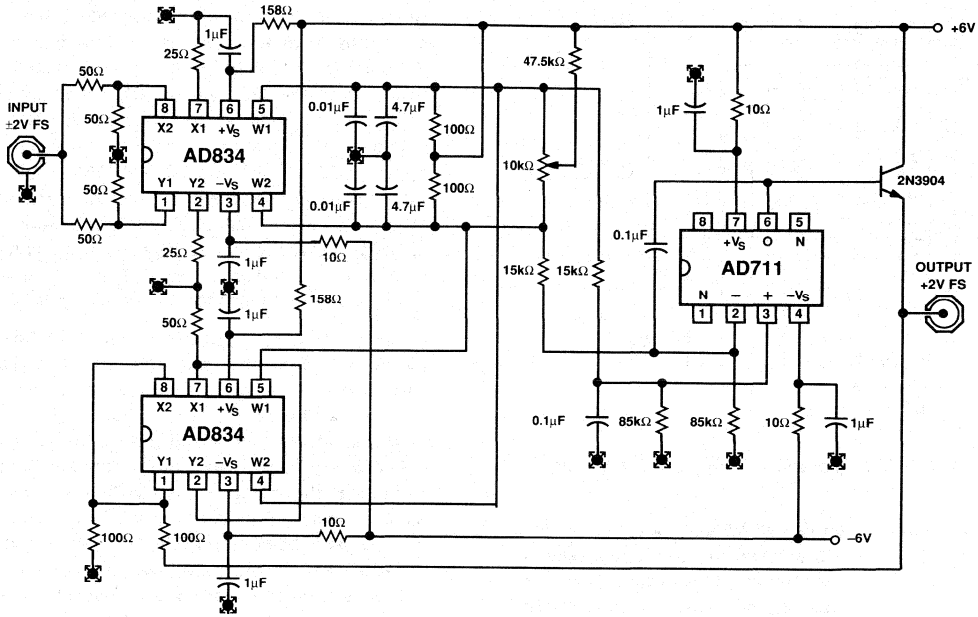

Figure 6. DC to 500 MHz RMS-to-DC Converter

DC COUPLED VCA APPLICATIONS

Where the dc response of the AD834 cannot be discarded, some form of level shifting, either passive or active must be employed, since high speed op amps often have inadequate common-mode range. The following applications show the use of active and passive level shifting circuits in the implementation of wideband voltage-controlled amplifiers.

A DC TO 60 MHz VOLTAGE-CONTROLLED AMPLIFIER USING PASSIVE LEVEL SHIFTING

Figure 7 shows the schematic of a circuit employing a passive network as a level shifter. The op amp chosen here is the AD5539.

The AD5539 is built on the same process as the AD834 and provides a 2 GHz gain-bandwidth product at high closed-loop gain. Unlike most op amps, the AD5539

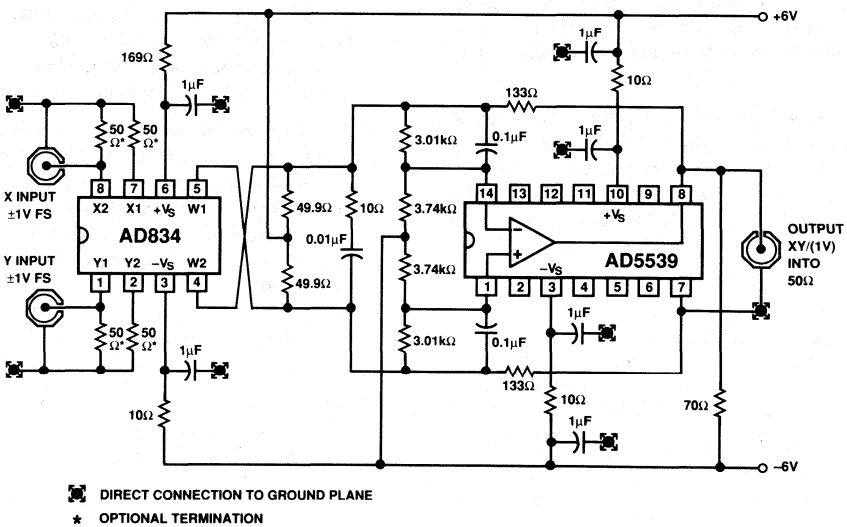

Figure 7. DC to 60 MHz Voltage-Controlled Amplifier Using Passive Level Shifting

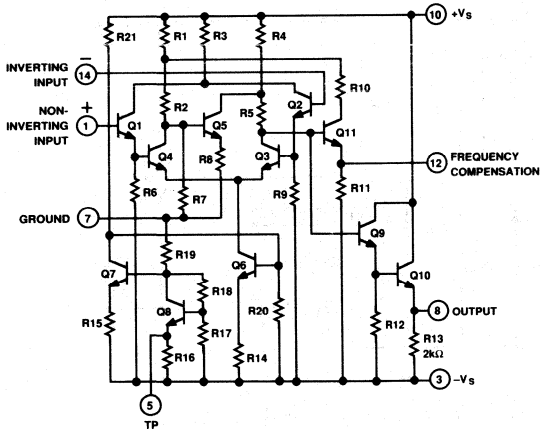

Figure 8. AD5539 Operational Amplifier Simplified Schematic

features a ground pin and an all-NPN output stage which operates in "Class A" to achieve the part's high speed (see Figure 8). Closer examination shows that there is a limited amount of "headroom" between the output node and the inputs, and between these voltages and ground. This, its high speed, and other unusual attributes of the AD5539 require special care in its use.

First, consider the consequences of its Class A output stage. In most op amps, the output can both "pull up" and "pull down" on the load, but the NPN emitter-follower output stage can only pull up. The AD5539 has an internal pull-down resistor (R11) of 2 kΩ, which can only supply two or three milliamps. A general-purpose high-speed multiplier must be able to swing to at least ±1 V while driving the minimum likely load resistance of 50 Ω. At this output level, the load current will be ±20 mA, which must therefore be supplied by an external pull-down resistor. In fact, the pull-down current must be considerably more than this, and requires careful consideration.

Figure 9 shows how the calculation is done. The 425 mV voltage sources are just "$I_B R_C$," that is, the standing current of 8.5 mA at the AD834 multiplied by the load resistor R_C, which we have here set to 50 Ω. The 200 mV sources in Figure 9 (a) are the "$I_W R_C$" generators when the full-scale output current is +4 mA. From here, we calculate V1 = 5.375 V and V2 = 5.775 V.

Next, we calculate the voltage at W2. Because the input current to an ideal op amp is zero, there is no loading at W2 and the voltage is simply V2 multiplied by the attenuation ratio 125/(125 + 50), or 4.125 V. Because the input voltage to an ideal op amp is zero, W1 is at the same voltage, so we can now calculate the current in the upper 50 Ω resistor as (5.375–4.125)/50 mA or 25 mA. Again, there is essentially no current at the input of the op amp, so the 25 mA all flows in the feedback resistor of 125 Ω, resulting in a voltage drop across it of 3.125 V. Finally, we calculate the output as the voltage at W1 (4.125 V) minus this drop; that is, the output is at +1 V.

Notice a somewhat surprising result at this point: although a current of 20 mA flows into the load, a larger current, 25 mA, flows in the feedback resistor! This unusual state of affairs is due to the very low value of the feedback resistor needed to reduce the scaling factor to the desired value, and the relatively large voltage needed at the output of the AD834 to ensure proper biasing of its outputs W1 and W2. Thus, even though the load needs to be sourced 20 mA, we still need to provide at least 5 mA in the pull-down resistor R_P to bias the output emitter-follower in the AD5539. The situation gets more severe when the output current of the AD834 is reversed, because we now need to sink 20 mA in the 50 Ω load and the voltage across the feedback resistor is now even higher.

This situation is shown in Figure 9(b). The calculation is exactly as before, and we discover that the current in the feedback resistor is now 39.7 mA. So R_P needs to provide the load current of 20 mA and an additional 40 mA or so in the feedback path, while the voltage across it is 5 V. This would require R_P = 83 Ω. In practice, it should be slightly lower to prevent slew rate limiting the fall time. Also, the feedback resistor will be raised from 125 Ω to 133 Ω to make up for the finite gain of the AD5539 under these heavily-loaded conditions. If we take the parallel sum of the 50 Ω load, the 70 Ω pull-down and about 150 Ω effective feedback resistance, the actual load on the amplifier is only 24 Ω!

The AD5539 is stable for uncompensated gains of greater than 5, and the AD5539 in this circuit is operating at a gain of just over 3. The 0.01 μF and 10 Ω network compensates by throwing away enough open loop gain to be stable when driving a 50 Ω load. For higher impedance loads, the 10 Ω compensation resistor may need to be reduced.

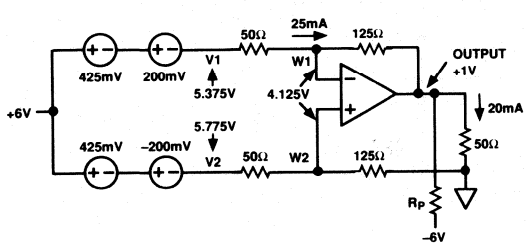

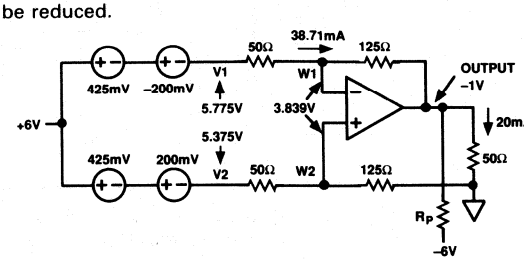

Figure 9. Equivalent Circuits for Calculating the Value of the Pull-Down Resistor

A level-shifting network is included between the nodes W1 and W2, whose average voltage is about +4 V, to the input of the AD5539 which must be close to ground. With the values shown, the op amp inputs are set slightly below ground (about −460 mV). This network halves the low frequency open-loop gain, which has some effect on the dc accuracy in the presence of offset voltages at the input to the AD5539. If output offset is important, a 500 Ω potentiometer should be inserted in series with the 3.74 kΩ resistors and its slider taken to −6 V. It is then adjusted for zero output with both X and Y inputs set to zero.

Note also that the "inner" Pins X1 and Y2 on the AD834 are grounded to minimize HF feedthrough; the resulting phase-reversal at the X input is corrected by swapping W1 and W2.

Figure 10 shows the pulse response with the input pulse applied to the X input and the Y input set to +1 V, indicating a rise time of 6 ns.

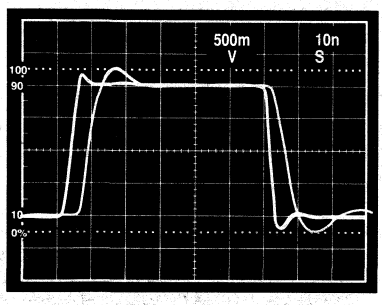

Figure 10. Pulse Response of the DC to 60 MHz Voltage-Controlled Amplifier

Figure 11 shows a set of frequency responses taken on an HP8753B network analyzer for Y inputs of +1 V, 316 mV, +100 mV, and 0 V. In the case of 0 V, the Y input is adjusted to null the input offsets. Note that the high frequency feedthrough is less than −65 dB of full scale (f < 3 MHz).

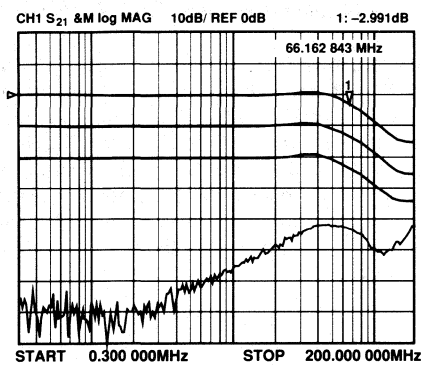

Figure 11. Frequency Response of the DC to 60 MHz Voltage-Controlled Amplifier

A DC TO 480 MHz VOLTAGE-CONTROLLED AMPLIFIER USING ACTIVE LEVEL SHIFTING

Figure 12 (a) shows an active level shifter, using a PNP transistor as a common base stage or cascode. Here, the AD834 is modeled by three ideal current sources, two for the 8.5 mA bias currents and one for the ±4 mA differential signal current. The transistors' bases are tied to +5 V, setting the emitter potential stays at 5.7 V resulting in a voltage of 3.3 V across the resistors R1 and R2 in the absence of signal. Figure 12 (b) shows an equivalent circuit.

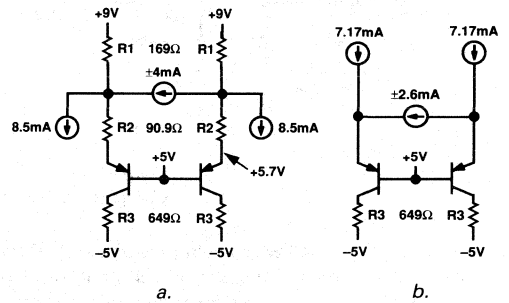

Figure 12. An AD834 Output Stage Using Active Level Shifting

The equivalent dc bias current of 7.17 mA is found by solving for the current flowing into the emitter when the signal current generator is zero. In the ac domain, the signal current generator sees R1 and R2 both tied to low impedance nodes. By inspection, the original signal current has been scaled by:

$$\pm\ 2.6mA\ =\ \pm 4\ mA\ \times\ \frac{R1}{R1\ +\ R2} \qquad (4)$$

Since AD834's outputs have very high output impedances, the equivalent series resistance can be ignored. The entire 7.17 mA flowing into the cascode's emitter flows out the cascode's collector, assuming a good α, and across R3. The voltage across R3 is:

$$4.65\ V\ =\ 7.17\ mA\ \times\ 649\ \Omega \qquad (5)$$

The operational amplifier's inputs are 350 mV below ground and are within the common-mode range of a wideband amplifier.

The bandwidth of a transistor configured as a cascode is the unity gain frequency (f_T) of the transistor, provided that the user does not create any spurious poles. Choosing an R1 and R2 such that their parallel sum is too large for the transistor's parasitic emitter-base capacitance or an R3 too large for the transistor's parasitic collector-base capacitance will create unwanted poles that lower the frequency response of the circuit.

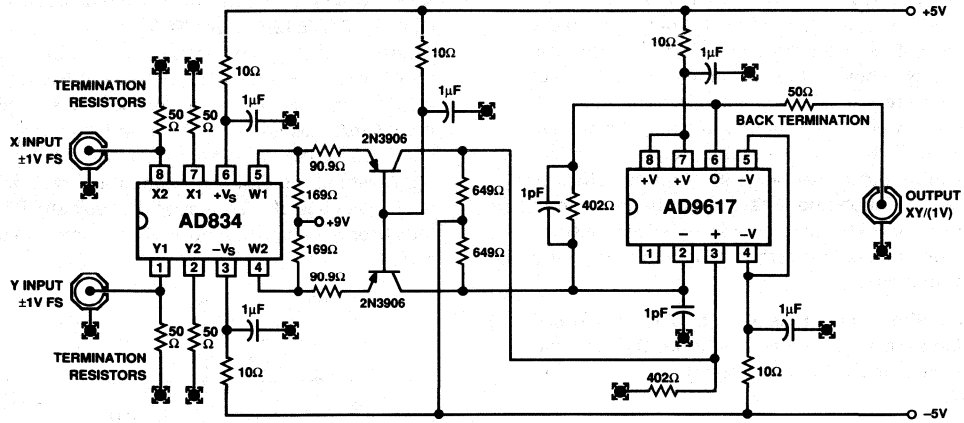

Figure 13. A DC to 480 MHz Voltage-Controlled Amplifier Using Active Level Shifting

Another potential pitfall when using the active PNP level shifter is oscillations at the cascode's emitter. The input impedance of a bipolar junction transistor's emitter is inductive at frequencies approaching its gain-bandwidth product (f_T), while the AD834's output is capacitive. Due to the high bandwidth of the system, these impedances can lead to oscillation.

To prevent such oscillations, the emitter in Figure 12 has been isolated from the AD834's output by R2. This prevents oscillations while providing signal attenuation (gain control) as related in Equation 4. The 2N3906s provide wideband level shifting without resonance or oscillation. Care must be taken when using alternative transistors.

The signal current at the cascodes' collectors is now fed to a wideband amplifier in a differential current to voltage converter configuration as shown in Figure 13. This configuration is similar to an op amp driven current-to-voltage converter which typically follows a current output multiplying digital-to-analog converter.

The AD9617 makes an excellent choice to drive the current to voltage converter. The AD9617 is a second-generation transimpedance amplifier (also known as a current feedback or TZ amplifier) with a fully complementary output stage (unlike the AD5539), and optimized for use with a 400 Ω feedback resistor.

The AD9617 inputs are tied directly to the collectors of the cascodes. The op amp creates a virtual short between the input nodes, forcing all the signal current to flow in the feedback paths. The differential transresistance of the convertor is 400 Ω. The desired scaling can be attained by means of the R1 and R2 attenuation network described above. The full-scale gain of the circuit (X = Y = 1 V) at the AD9617's output is calculated as:

$$2 \times 2.6\,mA \times 400\,\Omega = 2.08\,V \qquad (6)$$

or 1.04 V after the reverse termination resistor. The actual circuit shows a full-scale gain closer to unity.

Figure 14 shows the full-scale step response (−1 V to +1 V) applied to the X input and the Y input set to +1 V, demonstrating the circuit's capabilities with a rise time of under 2 ns while exhibiting some overshoot, but no ringing. Note that the output slews at over 500 V/μs.

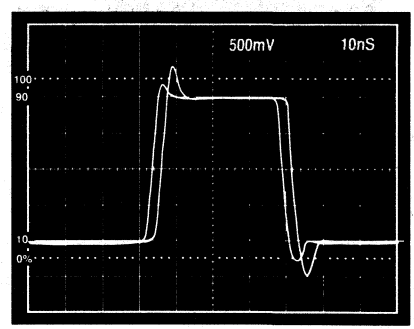

Figure 14. Step Response of the Wideband VCA

Figure 15 shows a set of frequency responses taken on the HP8753B network analyzer for Y inputs of +1 V, 316 mV, +100 mV, and 0 V. The Y input is actually adjusted to null the input offsets. Note that the circuit has a small-signal bandwidth of 500 MHz (at an input power level of 0 dBm). This bandwidth is possible with the two 1 pF capacitors at the inverting node. The high frequency feedthrough is less than −80 dB of full-scale (f < 2 MHz).

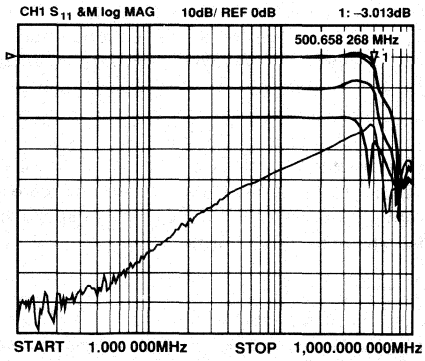

Figure 15. Frequency Response of the Wideband VCA

THE AD834 AS A VIDEO SWITCH

With 0 V or +1 V applied to the X channel as gate control and the video signal to the Y channel, the AD834 becomes a high-speed video switch. Figure 16 illustrates this idea with a high speed current switching circuit centered around an ECL switch. The current flows through either Q1 or Q2, depending on the input voltage. Current switching ensures fast and clean switching to determined levels (+1 V and ground), and allows the user to over- or under-drive the gate input.

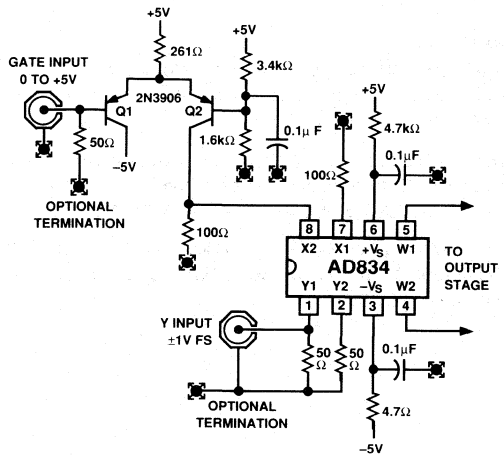

Figure 16. The AD834 as a High-Speed Video Switch

The AD834 switches on as the gate input rises from +1 V through +2 V at the gate circuit input. Below 1 V, Q1 absorbs almost all of the current from the 216 Ω resistor; the 2N3906 transistor is turned off. In this state, the 100 Ω resistor from the X2 input to ground accurately shuts the Y channel off, with Y channel feedthrough to the output measured at −50 dB. With the base of Q2 held at 1.6 V, the transistor's emitter potential is 2.35 V. A steady 10.2 mA (minus base current) from the 261 Ω resistor generates +1 V across the 100 Ω resistor at the X2 input independent of the exact high level of the gate input.

Figure 17 shows a scope photograph of a 1.5 ns rise-time pulse gating a 200 MHz signal. The resulting envelope rise time is 2.7 ns; it has a fall time of 3.0 ns. Although the switched signal may be much slower, the output stage from the AD834 should have a bandwidth greater than 100 MHz in order to maintain an envelope rise time of 3.5 ns.

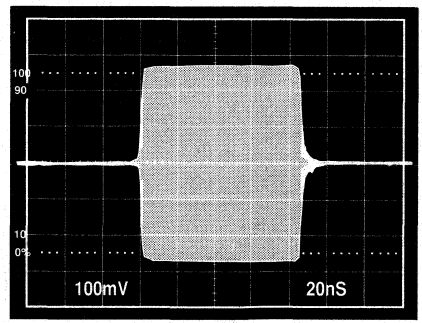

Figure 17. Rise Time of the Video Switch

AC OUTPUT-COUPLING METHODS

In many applications, the dc component at the output can be discarded. In such cases, a wideband buffer can easily ac couple to the AD834 output. The circuits below show the use of simple transformers and baluns for passive, ac coupled output circuits.

TRANSFORMER-COUPLED OUTPUT

Figure 18 shows the use of a center-tapped output transformer, which provides the necessary dc load condition at the outputs W1 and W2, and is designed to match into the desired load impedance by appropriate choice of turns ratio. The specific choice of the transformer design will depend entirely on the application. Transformers may also be used at the inputs. Center-tapped transformers can reduce high frequency distortion and lower HF feedthrough by driving the inputs with balanced signals. Suitable center-tapped transformers include the Coilcraft WB2010PC, which the manufacturer specifies for 0.04 MHz to 250 MHz operation.

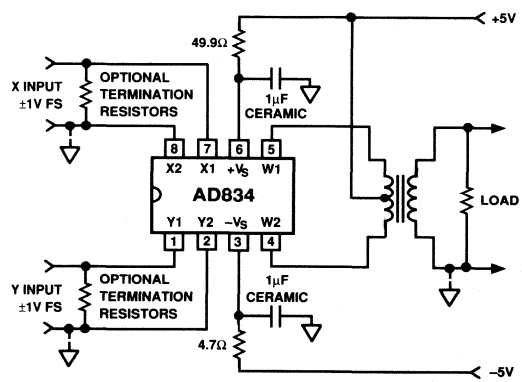

Figure 18. The AD834 with Transformer-Coupled Output

BALUN-COUPLED OUTPUT

Figure 19 shows a circuit which uses blocking capacitors to eliminate the dc offset, and a balun, a particularly effective type of transformer, to convert the differential (or balanced) signal to a single-sided (or unbalanced) output. A balun consists of a short length of transmission line wound on to a toroidal ferrite core, which converts the 'bal'anced output to an 'un'-balanced one (hence the use of the term).

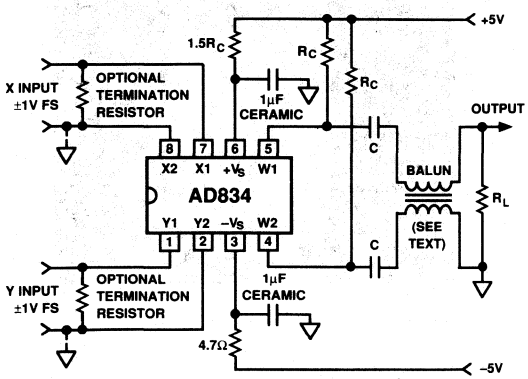

Figure 19. The AD834 with Balun-Coupled Output

Although the symbol used is identical to that for a transformer, the mode of operation is quite different. In the first place, the load should now be equal to the characteristic impedance of the line, although this will usually not be critical for short line lengths. The collector load resistors R_C may also be chosen to reverse-terminate the line, but again this will only be necessary when an electrically long line is used. In most cases, R_C will be made as large as the dc conditions allow, to minimize power loss to the load. The line may be a miniature coaxial cable or a twisted pair.

It is important to note that the upper bandwidth limit of the balun is determined only by the quality of the transmission line; hence, it will usually exceed that of th multiplier. This is unlike a conventional transformer where the signal is conveyed as a flux in a magneti core, and is limited by core losses and leakage induc tance. The lower limit on bandwidth is determined b the series inductance of the line, taken as a whole, an the load resistance (if the blocking capacitors C are suf ficiently large). In practice, a balun can provide excellen differential-to-single-sided conversion over much wide bandwidths than a transformer.

IMPLEMENTATION

Building these circuits requires good high frequenc techniques. The circuit schematics suggest suitable lay out. **Ground plane is essential for all of the circuit described in this applications brief.** It should cover a much of the component side of the PCB as possible, bu not directly underneath the IC or encircling any individ ual pins. Sockets add to the pin capacitance and induc tance, and should be avoided. If sockets are necessary use individual pin sockets such as AMP p/n 6-330808-3 They contribute far less stray reactance than the molde socket assemblies. Each power trace should be decou pled at the IC with a 0.1 μF low inductance ceramic capacitor, in addition to the main decoupling capacitor All lead lengths should be kept as short as possible. Fo lead lengths longer than an inch, stripline technique should be used.

REFERENCES

[Ref. 1]Gilbert, Barrie, "Translinear Circuits: A Proposed Classification," *Electronic Letters,* Vol. 11, No. 6, pp. 126 127.

[Ref. 2] Gilbert, Barrie, "A Precise Four-Quadrant Multiplier with Subnanosecond Response," *Journal of Solid State Circuits*, Vol. SC-3, No. 4, pp. 365-373.

[Ref. 3] Gilbert, Barrie, "A New Wideband Amplifie Technique," *Journal of Solid-State Circuits*, Vol. SC-3 No. 4, pp. 353-365.

ANALOG DEVICES

AN-216
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Video VCAs and Keyers Using the AD834 and AD811

by Eberhard Brunner, Bob Clarke, and Barrie Gilbert

NTRODUCTION

Voltage-controlled amplifiers (VCAs) built from analog multipliers take one of two forms. In the first, the multiplier acts as a voltage-controlled attenuator ahead of a fixed-gain amplifier. This type of VCA is used in applications where only a moderate maximum gain, but a fairly high maximum loss, are needed. In the second, the variable attenuation is placed in the feedback path around an op amp, which, in fact, implements an analog divider, more suitable for applications requiring high gains.

This application note describes practical circuits in which the wide bandwidth of the Analog Devices AD834 Four-Quadrant Multiplier and the AD811 Current-Feedback Op Amp are exploited to provide a video-quality VCA with a maximum gain of 12 dB (×4) or 20 dB (×10), based on the first of the above methods. A slightly modified form of this VCA, using two multipliers whose outputs are summed, provides the first of two video keyer designs; a second design uses global negative feedback around the multipliers to achieve improved accuracy and some simplification.

A VIDEO-QUALITY VCA

The VCA is shown in Figure 1. The AD834 multiplies the signal input by the control voltage. Its outputs are in the form of differential currents from a pair of open collectors, ensuring that the full bandwidth of the multiplier (which exceeds 500 MHz) is available for certain applications. In this case, more moderate bandwidth is obtained using current-to-voltage conversion, provided by the AD811 op amp, to realize a practical amplifier with a single-sided ground-referenced output. Using feedback resistors R8 and R9 of 511 Ω the overall gain ranges from −70 dB for V_G ~0 to +12 dB (a numerical gain of four) when V_G = +1 V.

The −3 dB bandwidth is 90 MHz (Figure 2) and is essentially independent of gain. The response can be maintained flat to within ±0.1 dB from dc to 40 MHz at full gain (Figure 3) with the addition of an optional capacitor of about 0.3 pF across the feedback resistor R8. The circuit produces a full-scale output of ±4 V for a ±1 V input, and can drive a reverse-terminated load of 50 Ω or 75 Ω to ±2 V. Figure 4 shows the typical pulse response.

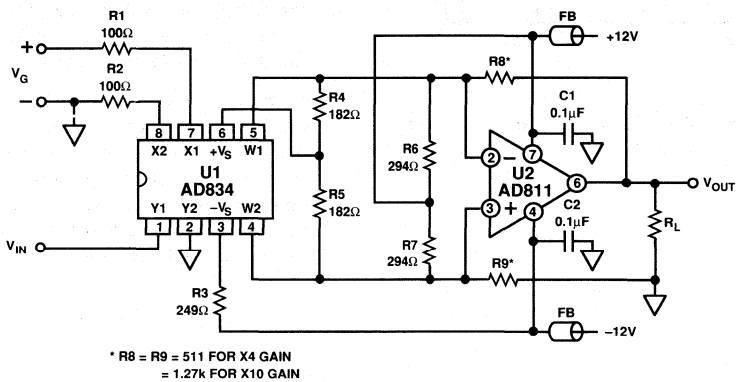

Figure 1. Complete VCA Provides Up to 20 dB of Gain (G = BW = 25 MHz) and a Bandwidth of Over 90 MHz (G = 12 dB)

The gain can be increased to 20 dB (×10) by raising R8 and R9 to 1.27 kΩ, with a reduction of the −3 dB bandwidth to about 25 MHz (also shown in Figure 2) and a maximum output voltage of ±9 V using the ±12 V supplies. It is not necessary to alter R6 and R7 for the high gain version of the amplifier, although an optimized design would raise these slightly to restore the common-mode voltage at the input of the AD811 to +5 V.

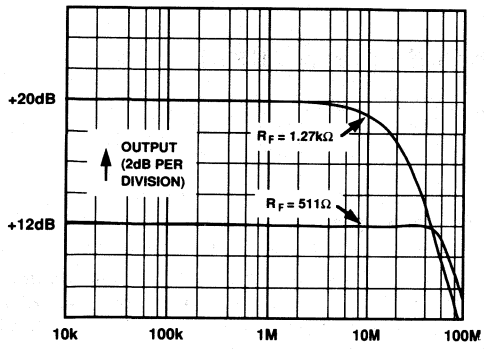

Figure 2. Small-Signal Response of the VCA Shows a −3 dB Bandwidth of 90 MHz for the 12 dB Version and 25 MHz for the 20 dB Version

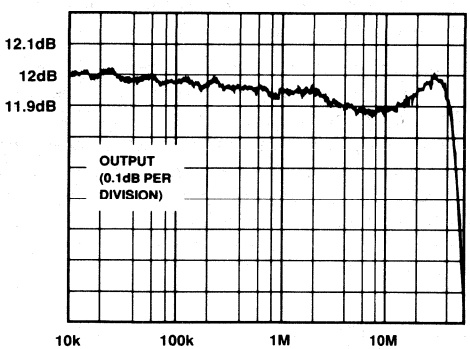

Figure 3. AC Response Can Be Held Flat to Within ±0.1 dB from DC to 40 MHz by Addition of a 0.1 pF Capacitor Across R8

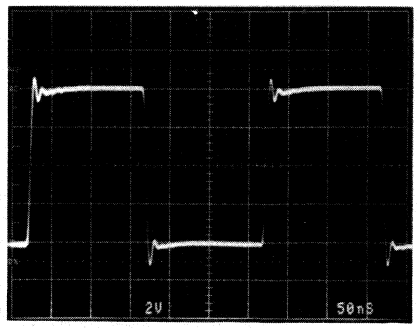

Figure 4. Full-Output Pulse Response for the 12 dB Amplifier

The gain-control input may be a positive or negative ground-referenced voltage, or fully differential, depending on the user's choice of connections at Pins 7 and 8. As shown, a positive value of V_G results in an overall noninverting response. Reversing the sign of V_G simply causes the sign of the overall response to invert. In fact, although we have called this a voltage-controlled amplifier, it can just as well be used as a general-purpose four-quadrant multiplier with good load-driving capabilities and fully symmetrical responses from X- and Y-inputs.

We have used the Y-input of the multiplier for the signal, since this port is slightly more linear than the X-input, and have shown X2 and Y2 grounded. These inputs each draw about 45 μA of bias current, so the grounded (unused) inputs should be terminated preferably in the same resistance as the source, in each case, to minimize offset voltages. The resistance of the signal source may in some cases be essentially zero (as in the case of a transformer-coupled input, or certain signal generators); note that a doubly terminated cable line of impedance Z_O will present a dc resistance of $Z_O/2$ at the input. Resistors R1 and R2 have been included in Figure 1 to minimize the likelihood of small aberrations arising in the signal path in those cases where V_G is derived from a source having poor HF characteristics; they may be omitted in the four-quadrant multiplier application.

High-frequency circuits such as those described herein are sensitive to component layout, stray capacitance, and lead lengths. Use a ground plane and make short, direct connections to ground. Bypass the power-supply connections—inductance in the power-supply leads can form resonant circuits that produce response peaking or even sustained oscillations.

Circuit Analysis

To understand the operation of the VCA, we need first to consider the scaling properties of the AD834, which is actually an accurate nonlinear (two-input) voltage-controlled current source. Figure 5 shows a simplified schematic of the whole VCA.

The exact transfer function for the AD834 would show that the differential voltage inputs at X1, X2 and Y1, Y2 are first multiplied together, divided by the scaling voltage of 1 V (determined by the on-chip bandgap reference) and the resulting voltage is then divided by an accurate 250 Ω resistor to generate the output current. A simplified form of this transfer function is

$$I_W = (X_1 - X_2)(Y_1 - Y_2) \times 4mA \quad (1)$$

where I_W is the differential current output from W1 to W2 and it is understood that the inputs X_1, X_2, Y_1, and Y_2 are expressed in volts. Thus, when both differential inputs are 1 V, I_W is 4 mA; this current is laser-calibrated to close tolerance, which simplifies the use of the AD834 in many applications. Note carefully the direction of this current in determining the correct polarity of the output connections.

It is easy to show that the output of the AD811 is

$$V_{OUT} = 2 \times I_W \times R_F \quad (2)$$

where R_F is the feedback resistor. For R_F = 500 Ω (499 Ω is the nearest standard resistor value), the overall transfer function of the VCA becomes

$$V_{OUT} = 4(X_1 - X_2)(Y_1 - Y_2) \quad (3)$$

which reduces to $V_{OUT} = 4\,V_G\,V_{IN}$ using the labeling conventions shown in Figure 1. As noted, the phase of the output reverses when V_G is negative. A slightly higher value of R_F is used to compensate for the finite gain of the AD811.

Both the AD811 and the AD834 can operate individually from power-supply voltages of ±5 V. However, to ensure proper operation of the AD811's input stage, the common-mode voltage at W1 and W2 must be within the common-mode range of these inputs. There are several ways to do this. We can use separate supplies of ±5 V for the AD834 and ≥ ±9 V for the AD811. Here, we have chosen to show how the VCA can be biased from one dual supply of nominally ±12 V. Figure 5 also helps to understand the dc biasing design.

We begin by deciding to place the AD834's outputs at about +5 V (a little higher than they operate in the other published applications of this product). Under dc conditions, the high open-loop gain of the op amp forces W1 and W2 to assume the same potential. We calculate the values of R_A to introduce the required 7 V drop, by considering the components of the total current in each of these resistors for the zero-signal condition.

First, when V_{OUT} = 0, the current in resistors R_F must be 10 mA (5 V/ 500 Ω). Second, the standing current into W1 and W2, due to the AD834's internal biasing, is 8.5 mA per side. Third, in this application we provide the positive supply voltage for the AD834 (at Pin 6) via resistors R_B which each carry one-half of the total supply current of 11 mA. Thus, the total current in resistors R_A is 24 mA (10 + 8.5 + 5.5 mA) and a value of 294 Ω is chosen (the closest standard value to 7 V/24 mA) for these resistors. Finally, we choose R_B to set the voltage at Pin 6 to +4 V, which is high enough to ensure accurate operation of the AD834 over the full signal and temperature ranges; the nearest standard resistor value is 182 Ω (1 V/ 5.5 mA).

The presence of these resistors (whose parallel sum is 112 Ω on each side) at the input of the op amp causes it to operate at a "noise gain" of 4.45 (499 Ω/112 Ω), but this neither has any significant effect on the dc scaling of the system, nor does it lower the closed-loop bandwidth (as would be the case for a conventional voltage-feedback op amp).

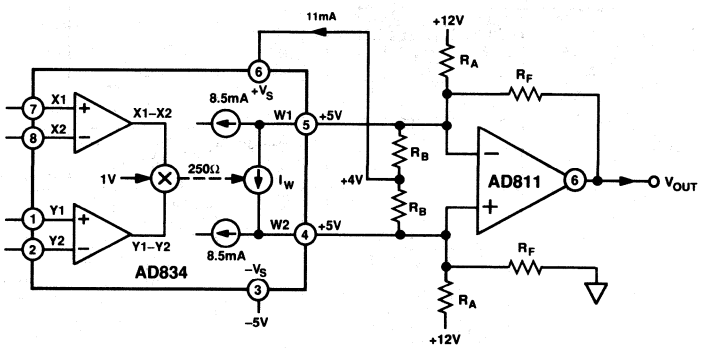

Figure 5. Simplified Schematic of the VCA for Analysis Purposes

A VIDEO KEYER BASED ON THE VCA

Using two AD834s and adding a 1 V dc source, a special form of a two-input VCA called a video keyer (Figure 6) can be assembled. Keying is the term used in reference to blending two or more video sources under the control of a further signal or signals to create such special effects as dissolves and overlays. The circuit described here is a two-input keyer, with video inputs V_A and V_B, and a control input V_G. The output at the load is given by

$$V_{OUT} = GV_A + (1 - G)V_B \quad (4)$$

where G is a dimensionless variable (actually, just the gain of the "A" signal path) that ranges from 0 when $V_G = 0$, to 1 when $V_G = +1$ V. Thus, V_{OUT} varies continuously between V_A and V_B as G varies from 0 to 1.

The operation is straightforward. Consider first the signal path through U1, which handles video input V_A. Its gain is clearly zero when $V_G = 0$ and the scaling we have chosen ensures that it is unity when $V_G = +1$ V; this takes care of the first term in Equation 4. On the other hand, the V_G input to U2 is taken to the *inverting* input X2 while X1 is biased at an accurate +1 V. Thus, when $V_G = 0$, the response to video input V_B is already at its full-scale value of unity, whereas when $V_G = +1$ V, the differential input $X_1 - X_2$ is zero. This generates the second term in Equation 4.

To generate the 1 V dc needed for the "1–G" term, an AD589 reference supplies 1.225 V ±25 mV to a voltage divider consisting of resistors R2 through R4. Potentiometer R3 should be adjusted to provide exactly +1 V at the X1 input.

In this case, we have shown an alternative arrangement using dual supplies of ±5 V for the AD834 and ±12 V for the AD811. Also, the overall gain in this case is arranged to be unity *at the load*, when it is driven from a reverse-terminated 75 Ω line. This means that the "dual VCA" has to operate at a maximum gain of ×2, rather than ×4 as in Figure 1. However, this cannot be achieved by lowering the feedback resistor, since below a critical value (not much less that 500 Ω) the AD811 will become unstable. This is because the dominant pole in the closed-loop ac response of a current-feedback amplifier is controlled by this feedback resistor. It would be possible to operate at a gain of ×4 and then attenuate the signal at the output. Instead, we have chosen to attenuate the signals by 6 dB at the input to the AD811; this is the function of R8 through R11.

The −3dB bandwidth is about 85 MHz and the gain is flat within ±0.1 dB to 30 MHz (Figure 7). Output noise and signal isolation with either channel fully off and the other fully on is about −60 dB to 20 MHz. The feedthrough at 100 MHz is limited primarily by board layout.

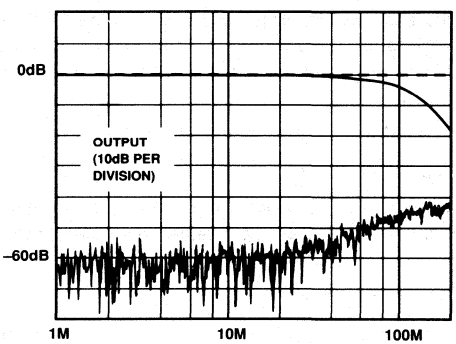

Figure 7. AC Response of the Video Keyer, at V_G = Zero and +1 V; Feedthrough Is About −60 dB

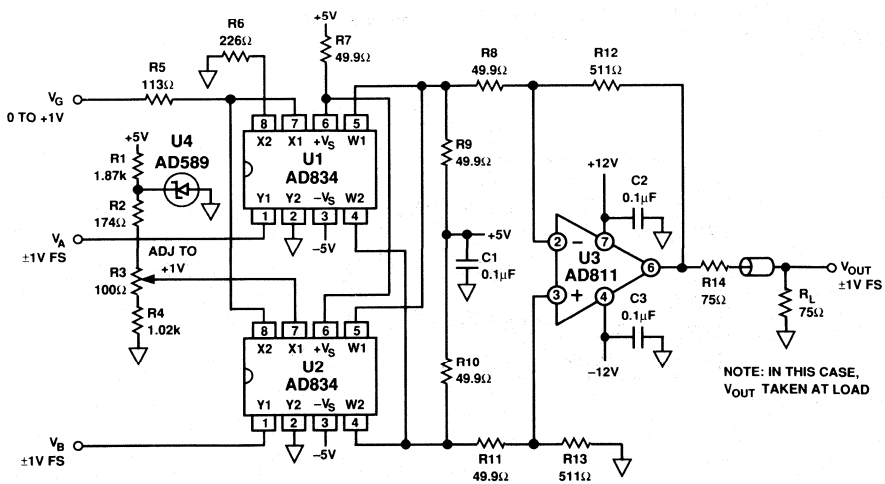

Figure 6. A Two-Input Video Keyer Based on the VCA

A FEEDBACK KEYER

The gain accuracy of the "VCA-based" keyer is dependent on the feedback resistor, R_F. Also, any nonlinearity in the multipliers will show up as a differential gain error. Using an alternative technique, in which the feedback is routed back to unused signal inputs on the AD834s, we can eliminate the feedback resistor and achieve higher accuracy. In the design shown here, we have also used a level-shifting network between the AD834 and the AD811 that eliminates the need for separate power supplies for the two ICs. (In fact, this technique can also be used in the VCAs.)

The basic idea is shown in Figure 8. Note first that V_{OUT} is returned to the inverting inputs Y2 of the multipliers and that their outputs are added. The sum is forced to zero by the assumed high open-loop gain of the op amp. Multiplier M1 produces an output $G(V_A–V_{OUT})$, while M2 produces an output $(1–G)(V_B–V_{OUT})$, where G is $V_G/(1\ V)$ and ranges from 0 to 1. Therefore, the complete system is described by the limiting condition

$$G(V_A–V_{OUT}) + (1–G)(V_B–V_{OUT}) \rightarrow 0 \qquad (5)$$

which requires that

$$V_{OUT} = GV_A + (1–G)V_B \qquad (6)$$

exactly as required for a two-input keyer. The summation of the differential current-mode outputs of the two AD834s is simply achieved by connecting together their respective W1 and W2 nodes. The resulting signal—essentially the loop error represented by the left-hand side of Equation 5—is forced to zero by the high gain of an AD811 op amp.

Figure 9 provides a practical embodiment of these ideas. The gain-control details to provide G and (1–G) terms

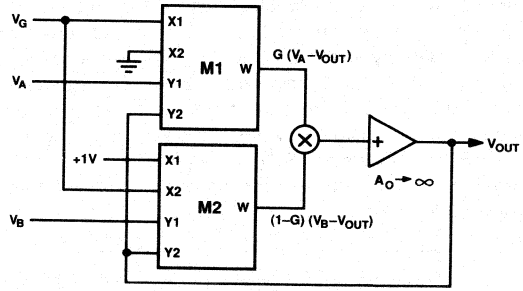

Figure 8. Elements of a Feedback Keyer

are identical to those used previously. The bias currents required at the output of the multipliers are provided by R8 and R9. A dc-level-shifting network comprising R10/R12 and R11/R13 ensures that the input nodes of the AD811 are positioned within an acceptable common-mode range for this IC. At high frequencies, C1 and C2 bypass R10 and R11, respectively.

R14 is included to lower the HF loop gain, and is needed because the voltage-to-current conversion in the AD834s, via the Y2 inputs, results in an effective value of the feedback resistance of 250 Ω (see Figure 5); this is only half the minimum value of 500 Ω required for HF stability of the AD811. (Note that this resistance is unaffected by G: when G = 1, all the feedback is via U1, while when G = 0 it is all via U2.) Resistor R14 reduces the fractional amount of output current from the multipliers into the current-summing inverting input of the AD811, by sharing it with R8. This resistor can be used to adjust the bandwidth and damping factor to best suit the application.

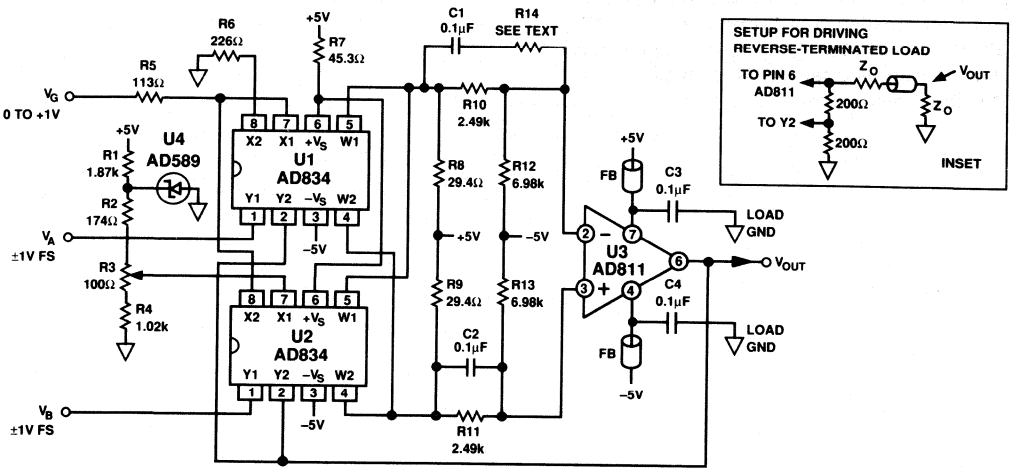

Figure 9. A Practical Embodiment of a Feedback Keyer. The Inset Shows the Feedback Configuration (Gain of ×2) for Driving a Reverse-Terminated Load.

Figure 10 shows the small-signal ac response of this system of the "A" channel at unity gain and zero gain; as is inevitably the case, there is a small amount of feedthrough at the highest frequencies. Two representative values of R14 are shown; using 402 Ω, the pulse response is considerably overdamped, resulting in a −3 dB bandwidth of 15 MHz, while a value of 107 Ω provides a maximally flat response with a −3 dB bandwidth of 70 MHz.

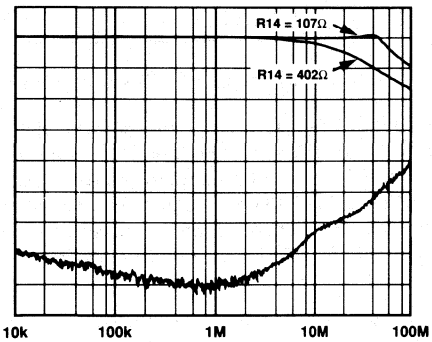

Figure 10. AC Response of the Feedback Keyer. For V_G = +1 V, the −3 dB Bandwidth Is 15 MHz Using R14 = 402 Ω and 70 MHz with R14 = 107 Ω. For These Measurements, R_L = 50 Ω

Figure 11 shows the pulse response at unity gain: in (a) R14 = 402 Ω, while in (b) R14 = 107 Ω. The frequency and pulse responses of the "B" channel, and of the gain-control input are the same, being limited by the output amplifier rather than the AD834s. Likewise, the differential gain and phase behavior will be determined primarily by the AD811; the data sheet should be consulted for more information. The feedthrough at 1 MHz is about −80 dB and −64 dB at 10 MHz and, as before, is eventually limited by board layout. All of these results used a 50 Ω load at the output.

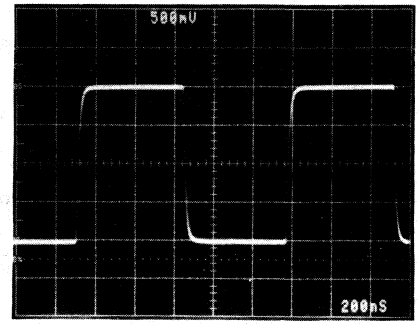

a.

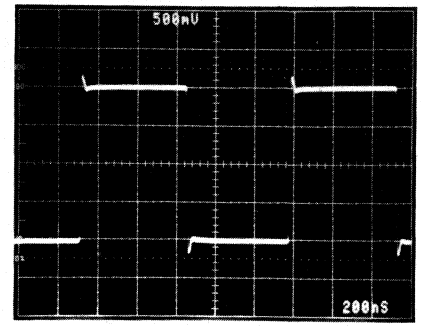

b.

Figure 11. Pulse Response of the Feedback Keyer. In (a), R14 = 402 Ω While in (b), R14 = 107 Ω. For These Measurements, R_L = 50 Ω

Unlike Figure 6's circuit, this keyer provides unity-gain operation. In applications where a reverse-terminated line (50+50 Ω or 75+75 Ω) is to be driven, the gain can be doubled by the inclusion of a resistive divider between V_{OUT} and the Y2 pins; equal resistors of 200 Ω can be used (see the inset in Figure 9). This halving of the feedback voltage also lowers the bandwidth, which can now be restored by reducing, or even eliminating, R14. Figures 12 and 13 show the modified circuit's performance when driving a 50 Ω reverse-terminated line.

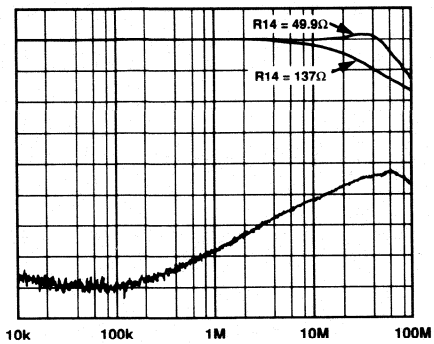

Figure 12. AC Response of the Feedback Keyer, Now Configured for a Gain of ×2. For V_G = + 1 V, the −3 dB Bandwidth Is 15 MHz Using R14 = 137 Ω and 70 MHz with R14 = 49.9 Ω. For These Measurements, R_L = 50 Ω

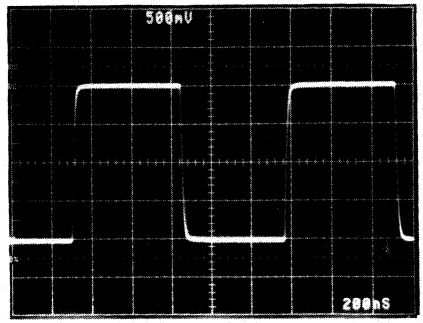

a.

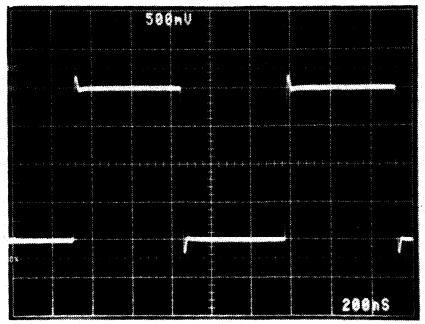

b.

Figure 13. Pulse Response of the Feedback Keyer Now Configured for a Gain of ×2. In (a), R14 = 137 Ω While in (b), R14 = 49.9 Ω. For These Measurements, R_L = 50 Ω

Data Acquisition Subsystems
Contents

Page

AN-312 DAS Packs Memory for Easy μP Interface 7–3

7

ANALOG DEVICES

AN-312
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

DAS Packs Memory for Easy μP Interface

by Hans Tucholski and John Wynne

An 8× 8 dual-port RAM makes a CMOS data-acquisition system—combining an 8-bit a-d converter and an eight-channel multiplexer—transparent to μPs like the 8085 and 6800.

For the first time, a microprocessor-compatible data-acquisition IC eliminates the need to wait for a-d conversion data. The AD7581 from Analog Devices not only contains dual-port, 8-byte RAM, it also holds an eight-channel multiplexer, an 8-bit successive-approximation a-d converter, address latches, interface logic, and three-state buffers for direct connection to a microprocessor bus—all on a single 147 × 125-mil CMOS chip.

With all conveniences, including memory, right on board, there is no need for interrupt software, and data are simply taken from the a-d converter in a normal microprocessor data-read cycle. The upshot of all this? The AD7581 can be treated by a microprocessor as a fast memory, which makes interfacing to processors as easy as adding another memory chip to the acquisition system.

As Fig. 1 implies, information flow within the device is unidirectional from input multiplexer to a-d converter, from the converter to memory, and from memory to data bus. Communication between different parts of the device takes place through internal bus structures. The successive-approximation conversion takes place on a continuous channel-sequencing basis. Data are automatically transferred to proper locations in the 8 × 8 RAM at the end of the conversion.

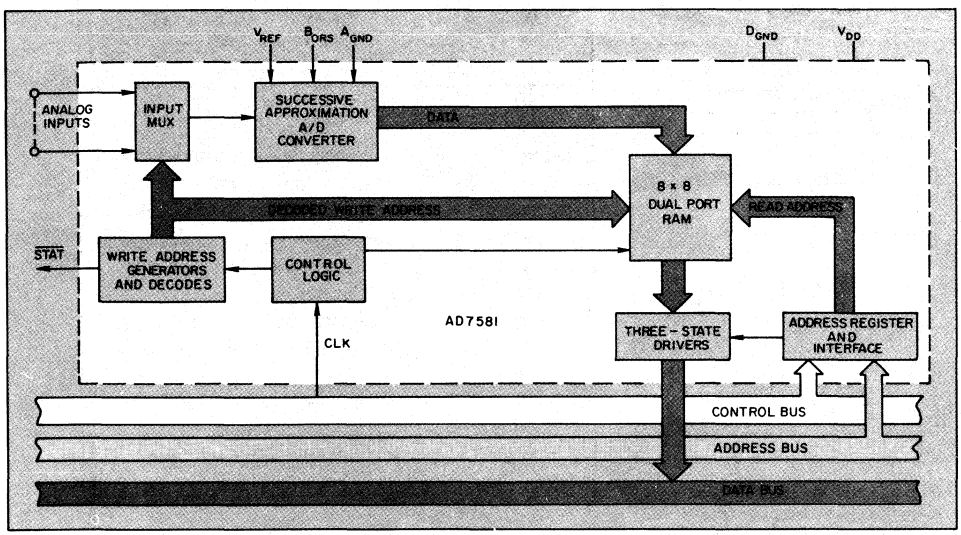

1. The AD7581 monolithic CMOS data-acquisition component includes an eight-channel multiplexer, an 8-bit successive-approximation a-d converter, three-state output drivers, and an 8 × 8 dual-port RAM.

Reprinted from ELECTRONIC DESIGN – September 13, 1980

Easy μP Interface

Because the AD7581 is viewed as memory (rather than as an a-d converter) by a microprocessor, it can be treated as an interleaved DMA device to allow easy μP interfacing. Access to the RAM is allowed at any time for any channel simply by applying the proper address to address pins (A0, A1, A2) and by activating chip select ($\overline{CS}$). The $\overline{CS}$ enables the read-address decoder, which transfers data bits from the addressed location to the input of the three-state output drivers, and thus onto the DB0 through DB7 data-output pins.

Methods of accomplishing this will vary according to the microprocessor used: For processor systems featuring a multiplexed address/data bus such as the 8085, the proper address must be latched into the AD7581 from the address lines using the address-latch enable (ALE). However, for systems having separate address and data buses, such as the 6800, the address latches can be made transparent simply by tying ALE high. The address decoder for the AD7581 is enabled with VMA (for 6800-type systems) and $\overline{RD}$ (for 8085-type systems).

The clock inputs, similarly, vary according to the microprocessor selected. The master clock is tied to ø1 in 6800-type systems, while other timing signals such as ALE can be used in 8085-type systems. Figure 2 (top) shows a minimum system configuration for the AD7581 interfaced with a 6800, while Fig. 2 (bottom) shows a minimum configuration for the 8085.

A general data-move program for use with the 6800 assumes that the destination address is within 256_{10} bytes of the source (AD7581 RAM) address (Fig. 3, top). In this example, the AD7581 has been assigned locations 8000_{16} to 8007_{16}. The index register generates both addresses, initially pointing directly to the source buffer and, by indexing using the displacement OF_{16}, subsequently pointing to the destination address. The move distance of this program is limited to 256_{10} bytes.

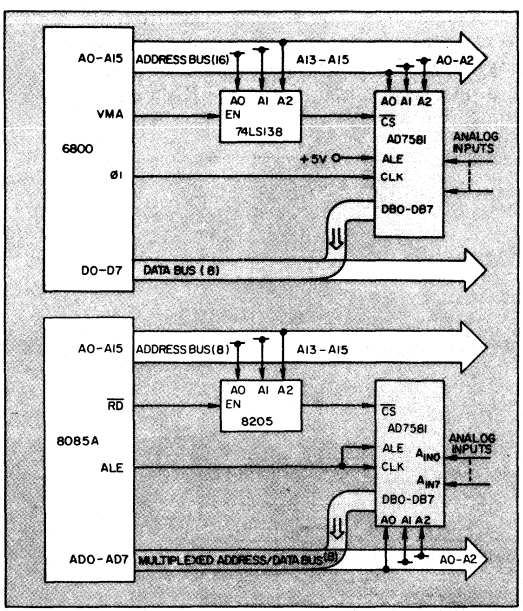

2. Interfacing with 6800-type microprocessors (top) is accomplished simply by tying the phase-output of the processor to the clock input of the AD7581, and making the ALE Input pin high at +5 V. With the 8085 (bottom), the ALE output of the processor is tied to the AD7581's ALE and clock.

```
      LDX    #   8000     LOAD SRCE ADDRESS
      LDA  B #   8        LOAD BUFFER LENGTH
LOOP  LDA  A  0,  X       LOAD CHANNEL DATA
      STA  A  0F, X       STORE CHANNEL DATA
      INX                 INCREMENT SOURCE POINTER
      DEC  B              DECREMENT BUFFER LENGTH
      BNE     LOOP        RETURN IF NOT EMPTY

      MVI     B      08   LOAD B WITH #8
      LXI     H,     8000 LOAD SRCE ADD INTO HL
      SHLD           2050 STORE SRCE ADD IN SRCE POINTER
      LXI     H,     2060 LOAD DST ADD INTO HL
LOOP  SHLD           2052 STORE DST ADD IN DST POINTER
      LHLD           2050 LOAD SRCE ADD FROM SRCE POINTER
      MOV     A,     M    LOAD CHANNEL DATA INTO A
      INX     H           INCREMENT SRCE ADD
      SHLD           2050 STORE SRCE ADD IN SRCE POINTER
      LHLD           2052 LOAD DST ADD FROM DST POINTER
      MOV     M,     A    STORE CHANNEL DATA IN DST ADD
      INX     H           INCREMENT DST ADDRESS
      DCR     B           DECREMENT REGISTER B
      JNZ            LOOP RETURN IF NOT EMPTY
```

3. Moving data with 6800 (top) and 8085 micros: Source and destination can be anywhere in memory (bottom).

For a more general system, in which the source and destination buffers can be anywhere in memory, the program for the 8085 is applicable (Fig. 3, bottom). Register HL assumes the role of index register and performs all the necessary addressing. The AD7581 is again assigned addresses 8000_{16} to 8007_{16} while destination addresses are 2060_{16} to 2067_{16}. Source and destination addresses are held during program execution in locations 2050_{16} and 2052_{16}.

Data transfer without μP interface

Although the AD7581 has been deliberately designed to appear as RAM to a microprocessor, it is equally at home in *nonmicroprocessor*-based, or hardware-oriented, systems. In some real-time applications, it may be required to provide an interrupt signal when a particular channel receives updated data. To achieve this, it is necessary to identify which channel is currently under conversion.

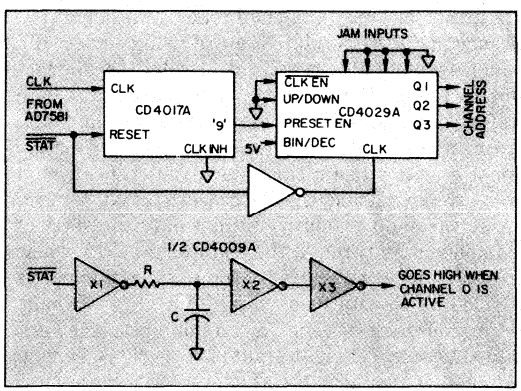

4. Without a microprocessor interface, the AD7581 can be harnessed to a write-address generator (top). Input-channel selection without microprocessor control is helped by a channel-0 identification circuit (bottom).

In a microprocessor-based system, channel identification is straightforward, accomplished by monitoring the $\overline{\text{STAT}}$ output and then using software techniques to identify the input multiplexer address. In the stand-alone mode, of course, this is not possible and a hardware solution is necessary. The hardware solution must, in effect, make the internal address controlling the input multiplexer available externally. To achieve this, the nature of the status signal, $\overline{\text{STAT}}$, changes from its usually passive role of indicating the a-d converter's operating state to the more important role of providing an external synchronizing clock.

Figure 4 (top) shows one way to derive the address of the channel under conversion from the $\overline{\text{STAT}}$ signal. The CD4029A counter is incremented every eight input clock periods via the inverted $\overline{\text{STAT}}$ signal. The BCD output of the 4029A provides the required channel address. When STAT remains low during channel-0 conversion, the counter is reset from the CD4017A.

For applications requiring an interrupt signal only after each complete scan, the circuit of Fig. 4 (bottom) should prove useful. The time constant RC is chosen so that X_2 ignores the normal $\overline{\text{STAT}}$'s low pulse width (eight clock periods) but responds to the much wider $\overline{\text{STAT}}$ pulse width (72 clock periods) that occurs during channel-0 conversion. Typically for a 1-μs clock period, C= 0.022 μF, R = 1.8 kΩ.

Figure 5 shows a simple stand-alone configuration for the AD7581 in which channel selection is controlled by a thumbwheel BCD switch. Since the channel-selection address is always present, ALE can be tied high, making the input address latches transparent. The $\overline{\text{STAT}}$ pin is directly connected to the $\overline{\text{CS}}$ pin to provide output data after every conversion is complete. A data latch is required to avoid losing the data when the three-state bus drivers return to their high-impedance state. The rising edge of $\overline{\text{STAT}}$ strobes the data into the 8212 latch and, after some propagation delay, floats the data bus.

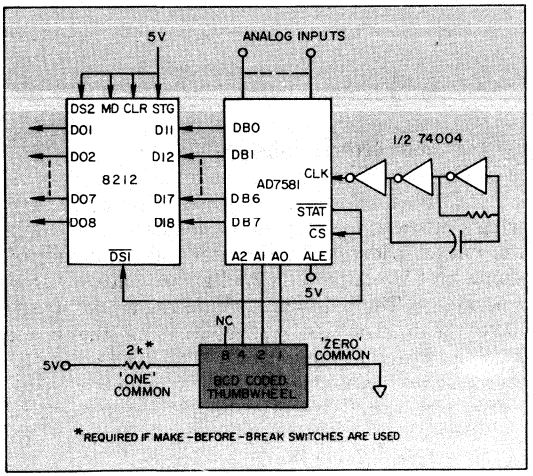

5. Outboard BCD switches provide address decoding in the absence of a microprocessor.

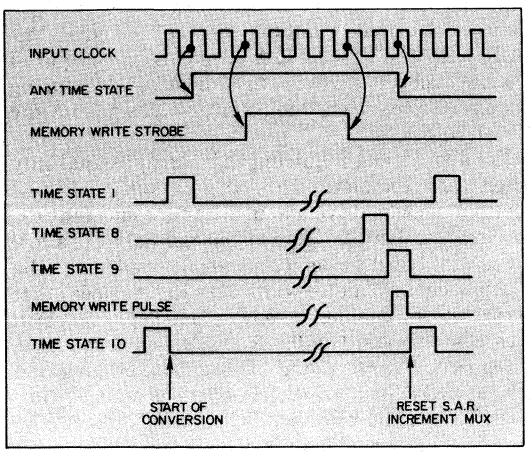

6. Ten separate timing states are required to cycle data through the AD7581. The first eight of these states cycle each channel of the DAS through successive-approximation conversions; the ninth cycle clears the successive-approximation register and loads the memory. The tenth state enables the write-address counter. The entire sequence requires 640 clock pulses.

Whatever the application, successful operation depends on the internal logic sequence.

For starters, all timing for the analog-to-digital converter is derived from the input clock frequency. Start-up logic is included on the device to establish the correct sequences on power-up. A maximum 800 clock pulses are required for this period. Internal timing relationships are shown in Fig. 6. Timing logic inside the AD7581 provides ten separate time states, or T-states, for each channel from the input clock frequency. Each T-state extends for eight input clock periods. Thus, 640 clock pulses are required for a full scan through all eight channels, which gives a conversion rate of 80 μs per channel for a 1-MHz clock rate.

The first eight of these T-states are required for loading 8 complete conversions into the 8 × 8 RAM, while the successive-approximation register is cleared during the tenth T-state. This T-state also increments the a-d converter's write-address counter, whose output is decoded to provide a 1-of-8 output to drive the input multiplexer as well as to select the corresponding RAM location.

Figure 6 shows memory being updated during time state 9. Access to a storage element in the 8 × 8-bit memory array is via complementary transmission gates, which have separate read and write data lines (see "Fast On The Inside"). Although the memory updating is related to the rising edge of the memory-write pulse and is completed in approximately 100 ns, the memory-write pulse stays high for four input clock periods. This keeps the write operation independent of clock waveforms. When the memory-write pulse returns low, the transmission gates are closed and the data are latched.

A status signal, identifying not only the start of a complete scan through the eight channels but also the status of the converter through each channel, is provided by the status-output signal, $\overline{\text{STAT}}$. During the tenth T-state of each conversion cycle, the $\overline{\text{STAT}}$ output goes low for eight input clock periods. During channel-0 conversion, the $\overline{\text{STAT}}$ output remains low for an additional 64 clock periods over normal. This is the only signal related to the operating state of the a-d converter that is externally available. Figure 7 shows this $\overline{\text{STAT}}$ signal and its relationship to the a-d converter's conversion cycle.

The on-chip multiplexer consists of eight thin-film resistors, each in series with a MOS SPDT switch, and not the usual one-thin-film-resistor-with-eight-MOD-SPST scheme. This allows the inputs to exceed

Fast on the Inside

There are several internal reasons why the AD7581 can be treated simply as a fast memory device.

The heart of the swift, microprocessor-compatible device is its storage element: an 8 × 8-bit static CMOS memory connected over complementary transmission gates to separate data lines for writes and reads.

High-speed performance stems from a ten-transistor cell design, which also permits direct-read-lines to be connected to the three-state buffers without sense amplifiers. The static design also gives virtually zero power dissipation. And the double-layer metal interconnect used in the fabrication process of the device allows a dense layout.

Memory update, which occurs during part of the memory-write time state, is related to the rising edge of the clock. Although it is completed in approximately 100 ns, a long pulse of four clock periods is used because this makes the write operation independent of clock waveforms and allows simple logic to be used. At the end of the write-enable pulse, WE, the data are latched. Since the memory always contains the most recent conversion data and sports a dual-port structure, any channel can be read at any time.

Memory words are selected by placing the appropriate data on the address inputs. An address-latch-enable (ALE) input is available to latch the data internally, and becomes necessary when the AD7581 is used in 8085 systems and the address is obtained from the shared address/data bus, AD_8 - AD_{11}. In other cases, the latch-enable input can simply be tied to V_{DD}.

The address data are not decoded completely at this point, since this would cause unnecessary transients at the memory latch. The decoder requires a read-strobe signal, which then selects the word location, transferring its bits to the input of the three-state buffers. The strobe signal is designated "Chip Select" rather than "Read" because it also enables the three-state buffers.

The chip-select logic has been designed to meet the timing requirements when interfaced directly to an 8085 bus. Memory access is fast, with valid data available approximately 250 ns after the falling edge of Chip Select, $\overline{CS}$. Furthermore, the AD7581 clears the bus about 80 ns after "reading" at the rising edge of $\overline{CS}$—well before the processor places another address onto the bus.

Interleaved DMA operation is achieved simply by using the processor system clock as a clock for the a-d-converter system and the processor read as chip select. In this way, internal memory writing does not coincide with the processor's data reading.

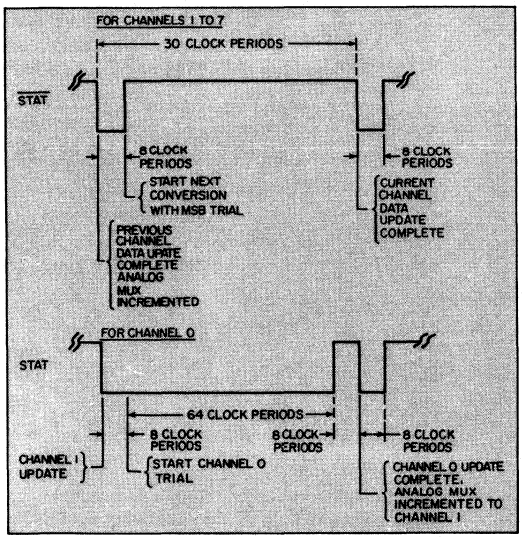

7. The STAT signal, which identifies the start of a scanning/conversion cycle, goes low during the tenth T-state, and remains low while the zero-channel update is completed.

the supply voltage without causing latch-up and reduces feedthrough between the channels, while providing constant input resistance.

All channels are switched through their series 20-kΩ resistors to analog ground, AGND, when they are not under conversion. The channel under conversion is summed with the output from the reference-driven R-2R ladder network to produce an error voltage at the comparator's summing junction. This error voltage is successively reduced to 0 V by the successive-approximation technique to give a correct digital representation of the analog input. The conversion speed of the a-d converter is limited mainly by the settling time of the error voltage at the comparator's summing junction and is determined by the switch capacitances of the multiplexer and the d-a-converter network. The comparator, consisting of a differential-to-single-ended input stage, a gain stage, and a logic driver, provides adequate voltage gain.□

Digital-to-Analog Converters
Contents

		Page
AN-313	Getting the Most from High Resolution D/A Converters	8–3
AN-342	Analog Signal-Handling for High Speed and Accuracy	8–15
AN-109	Understanding and Preventing Latch-Up in CMOS DACs	8–23
AN-210	Adding Additional Input Channels to the AD7773/AD7775	8–31
AB-6	Single-Supply Operation of the DAC-08 and DAC-20	8–33
AN-225	12-Bit Voltage-Output DACs for Single-Supply 5 V and 12 V Systems	8–35
AN-310	Single-Supply 12-Bit DAC Circuits Using the AD7568	8–39
AN-314	14-Bit DACs Maintain High Performance Over Extended Temperature Range	8–43
AB-3	Digital Nulling of Precision Op Amps	8–49
AN-17	DAC-08 Applications Collection	8–51
AN-19	Differential and Multiplying D/A Converter Applications	8–63
AN-343	Circuit Ideas for IC Converters	8–71
AN-21	4–20 mA Digital to Process Current Transmitter	8–75
AN-23	D/A Converter Generates Hyperbolic Functions	8–77
AN-208	Understanding LOGDACs	8–79
AN-137	A Digitally Programmable Gain and Attenuation Amplifier Design	8–83
AN-142	Voltage Adjustment Applications of the DAC-8800 TrimDAC, An Octal, 8-Bit D/A Converter	8–85
AN-219	Electronic Adjustment Made Easy with the TrimDAC	8–99
AN-316	AD7224 Provides Programmable Voltages Over Varying Ranges	8–103
AN-317	Circuit Applications of the AD7226 Quad CMOS DAC	8–107
AN-318	AD7528 Dual 8-Bit CMOS DAC Application Note	8–113
AN-319	Simple Interface Between D/A Converter and Microcomputer Leads to Programmable Sine-Wave Oscillator	8–121
AN-320A	CMOS DACs and Op Amps Combine to Build Programmable Gain Amplifiers – Part I	8–127
AN-320B	CMOS DACs and Op Amps Combine to Build Programmable Gain Amplifiers – Part II	8–139
AN-321	3-Phase Sine Wave Generation Using the AD7226 Quad DAC	8–149
AN-209	8th Order Programmable Low-Pass Analog Filter Using Dual 12-Bit DACs	8–153
AN-322	Improve Function Generators with Matched D/A Converters	8–161
AN-323	Microstepping Drive Circuits for Single-Supply Systems	8–167
AN-324	Simple DAC-Based Circuit Implements Constant Linear Velocity (CLV) Motor Speed Control	8–175
AN-325	12-Bit Analog I/O Port Uses AD7549 and 8051 Microcomputer	8–181
AN-326	Interfacing the AD7549 to the MCS-48 and MCS-51 Microcomputer Families	8–185
AN-327	DAC ICs: How Many Bits Is Enough?	8–189
AN-206	Analog Panning Circuit Provides Almost Constant Output Power	8–193
AN-207	Interfacing Two 16-Bit AD1856 (AD1851) Audio DACs with the Philips SAA7220 Digital Filter	8–197
AN-328	Generate Four Channels of Analog Output Using AD7542 12-Bit D/A Converters and Control the Lot with Only Two Wires	8–201
AN-329	Dynamic Performance of CMOS DACs in Modem Applications	8–207
AN-37	Eight-Channel Codec Applications	8–213
AN-38	Four-Channel Shared Codec	8–221
AN-39	Companding D/A Converter	8–227
AN-41	Improved Shared-Channel Codec Design with Analog Devices' Companding DACs	8–239
AN-42	A 1 kHz, 0 dBm0 Standard Signal Generator	8–247
AN-101	Cross-Plot Generator Allows Quick A/D Converter Evaluation	8–249
AN-205	Video Formats and Required Load Terminations	8–253

		Page
AN-330	Animation Using the Pixel Read Mask Register of the ADV47x Series of Video RAM-DACs	8–257
AN-331	Changing Your VGA Design from a 171/176 to an ADV471	8–265
AN-332	Improved PCB Layouts for Video RAM-DACs Can Use Either PLCC or DIP Package Types	8–267
AN-333	Design and Layout of a Video Graphics System for Reduced EMI	8–269
AN-237	Choosing DACs for Direct Digital Synthesis	8–285

ANALOG DEVICES

AN-313
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Getting the Most from High Resolution D/A Converters

by Scott Wayne

Here's a close look at D/A converter specs, requirements, error sources, and test methods—and how they can affect your circuit designs

High resolution digital-to-analog converters have many varied applications that can be loosely grouped into two categories: instrumentation and waveform reconstruction. Instrumentation applications include raster scan, process control, automatic test equipment, robotics, and others. Waveform reconstruction includes digital audio, sonar, and telecommunications, as well as specialized waveform generation. Each of these two classifications has widely different specifications and requirements for D/A converters. For example, instrumentation applications require traditional D/A specifications such as good linearity and high stability, while waveform reconstruction applications require low total harmonic distortion (THD) and a high signal-to-noise ratio (S/N).

D/A converter for instrumentation

A high-resolution D/A converter for instrumentation must have low differential nonlinearity (DNL) and integral nonlinearity (INL) over temperature. It also must have low offset and gain drift, a high power supply rejection ratio (PSRR), and low noise. This explanation of the more important instrumentation specifications — as well as techniques for measuring and improving them — should prove helpful in applying high-resolution D/As.

As the digital input to a D/A converter changes by one bit, the analog output changes by some amount. Ideally, a one bit digital input change should produce a constant 1 LSB (least significant bit) output change anywhere along the input-output transfer function. Differential nonlinearity, usually the most important specification for an instrumentation converter, is a measure of the deviation between the actual analog output change and the theoretical change of 1 LSB. It is specified at room temperature in LSBs or as a percentage of full scale range (FSR). Differential nonlinearity is a function of time and temperature. Its drift is given in ppm 1000 hours or ppm/°C.

Reprinted from ELECTRONIC PRODUCTS MAGAZINE – December 12, 1983

Monotonicity is another measure of differential nonlinearity. If a converter is monotonic, the analog output will remain constant or increase as the digital input is increased. Nonmonotonicity implies a differential linearity error of greater than 1 LSB. To compute the monotonic temperature range of a converter, subtract the initial differential nonlinearity of the D/A from 1 LSB. Divide the result by the DNL drift temperature coefficient. This gives the minimum temperature deviation around room temperature for which the converter will remain monotonic. For example, given a 16-bit device with an initial linearity error of ½ LSB (30 ppm) and a linearity drift of 1 ppm/°C, the monotonic temperature range is 25°C to ±30°C or −5 to +55°C. Monotonicity is a very important specification for process control applications because a nonmonotonic converter will cause the control loop to oscillate endlessly.

Differential nonlinearity can easily be measured by directly comparing the analog output produced by one digital input and the analog output produced by the next sequential digital input. The effects of differential nonlinearity and nonmonotonicity errors are shown in Fig. 1.

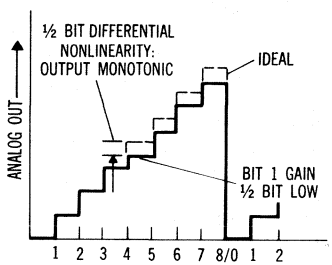

(a) Differential nonlinearity

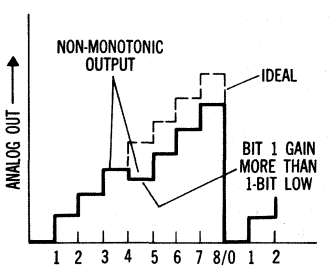

(b) Non-monotonicity

Fig. 1. Differential linearity and monotonicity errors

Measuring integral nonlinearity

Integral nonlinearity (INL), also referred to as nonlinearity or relative accuracy, is the deviation of the actual converter output from a straight line drawn between the end points of the converter's input-output transfer function (see Fig. 2). INL is very difficult to measure as the measurement involves finding the difference between two large numbers. A 6½-digit voltmeter that is accurate to better than 0.0002% would be necessary to accurately measure the integral nonlinearity of a 16-bit converter. For this reason, most integral nonlinearity measurements are made by comparing the converter under test to a reference converter of known accuracy. The errors are then read with a null meter. A precision 18-bit converter can be used to test converters with up to 16-bit accuracy. A precision divider, traceable to the National Bureau of Standards, must be used to properly test an 18-bit converter.

Summation errors, or superposition errors, occur whenever the analog output due to any digital bit combination does not equal the algebraic sum of the analog outputs due to each digital input bit alone. The numerous sources of summation error in D/A converters depend primarily upon their internal architecture. For example, in a typical R-2R ladder configuration, if both outputs are not at the same potential, summation errors will occur. This can easily be seen in the case of the 2-bit converter shown in Fig. 3. Also, if the switch resistance is not the same in both the on and the off positions, summation errors will again occur.

Another cause of summation errors is the feedback resistor in a voltage-output converter. As the analog output increases from zero to full scale, the power dissipated by the feedback resistor increases, and the resistor heats up. This causes a change in resistor value, and a corresponding gain change along the transfer function. Since this apparent gain is different for a sum of bits than it is for the individual bits alone, a summation error occurs. Gain changes are especially troublesome in hybrid and monolithic converters, where the power dissipation of the feedback resistor is limited by its small size.

The simplest method for measuring differential nonlinearity is to

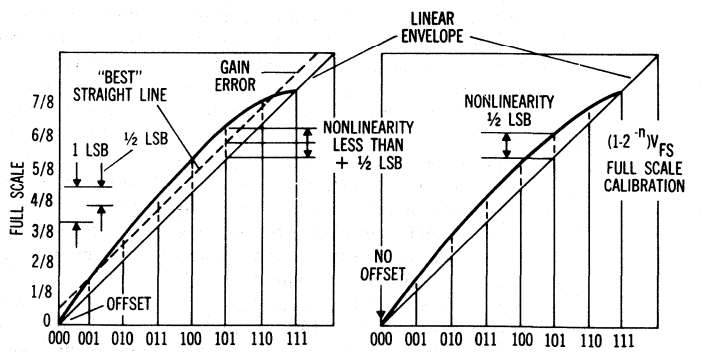

(a) ½ LSB nonlinearity achieved by arbitrary location of "best straight line."

(b) Nonlinearity reference is straight line through end points.

Fig. 2. Comparison of linearity criteria for 3-bit D/A converter (straight line through end points is easier to measure and gives a more conservative specification)

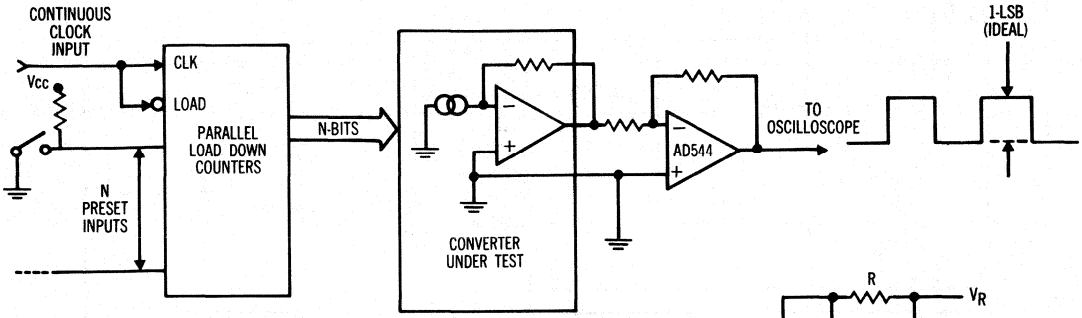

Fig. 4. A simple DNL tester

connect a parallel-load down counter to the device under test (see Fig. 4). The inputs to the counter are first preset to the initial desired digital input value. When the clock goes low, the digital input is asynchronously loaded into the counter, and presented to the converter. The counter counts down by one on the rising edge of the clock pulse, and a digital input value one bit less than the preset input is presented to the converter. Ideally, with a continuous clock signal applied to the counter, the resulting analog output will be a square wave with an amplitude of 1 LSB for any digital input. The deviation of this amplitude from 1 LSB is the differential nonlinearity at the digital input being observed. Amplification of the analog output may be added if desired.

For 18-bit converters, there are 2^N or 262,144 possible digital input combinations. It would be quite tedious to set this many input combinations with toggle switches. Even with an automated tester that required as little as 10 ms per test, the entire sequence would take almost one hour. Fortunately, it is not necessary to test all input combinations to characterize a high resolution device with no summation errors — a total of only N tests is required. If the summation errors are found to be reasonably small, a total of 2N tests may be sufficient. These tests are usually performed at the major carries, the transitions between a test bit "ON" and all of the bits less significant than the test bit ON, and the test bit OFF.

A simple method for measuring integral nonlinearity is to connect the converter to a parallel-load right-left shift register (see Fig. 5). The register is preset with 00 . . . . 01. As long as the MSB is a zero, zeroes are shifted in from the right. When the high bit reaches the MSB, ones are shifted in from the left. When the register reaches 11 . . . . 11 it is again preset and the process continues. A reference converter of at least two bits more accuracy than the one under test is connected to the same register. The two analog outputs are then subtracted and amplified by an instrumentation amplifier. The output of the amplifier is ideally zero volts. The deviation of the output from zero volts is the integral nonlinearity of the converter under test. Zero and gain servos can also be added to enhance this test.

Program the measurements

A practical method for simultaneously measuring both integral and differential nonlinearity is to program the desired digital inputs into a PROM (see Fig. 6). As the PROM is cycled, the converter under test undergoes a transition from a single input bit ON alone to all of the lesser significant bits ON, while the input of the reference converter sees just the single bit ON. Each time the lesser significant bits are ON, one LSB of current is added to the output of the converter under test. The two analog outputs are differenced, amplified, and displayed. Ideally, the output of the amplifier is zero volts. The deviation of the output voltage due to each input is

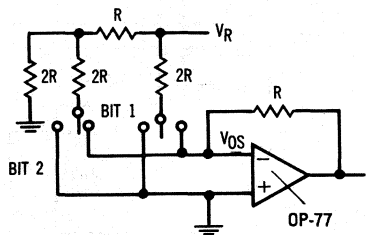

(a) Two bit voltage output R-2R ladder D/A converter

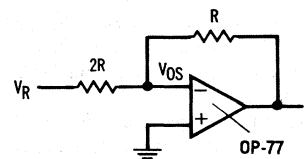

(b) Bit 1 ON, $V_O = -V_R/2 + 3/2\,V_{OS}$

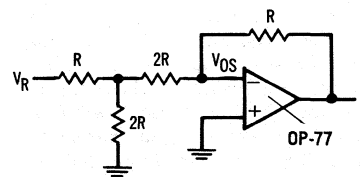

(c) Bit 2 ON, $V_O = -V_R/4 + 11/8\,V_{OS}$

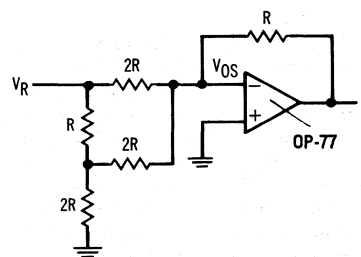

(d) Bits 1 and 2 ON, $V_O = -3/4\,V_R + 15/8\,V_{OS} \neq \sum 1+2$

Fig. 3. Summation errors in a 2-bit D/A converter

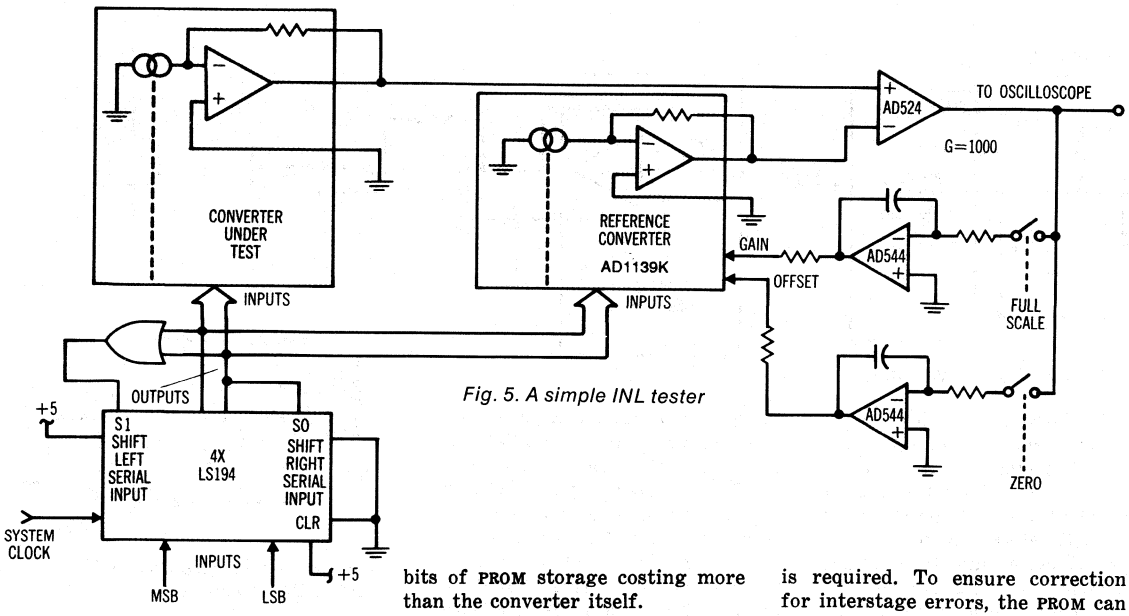

Fig. 5. A simple INL tester

the integral nonlinearity. The difference in output voltage due to adjacent digital inputs is the differential nonlinearity.

This test can easily be automated by replacing the PROM with a microprocessor and using a bus-interfaceable voltmeter in place of the oscilloscope. For each digital input, the output voltage of the amplifier is measured and recorded. This voltage represents the integral nonlinearity. The meter is then nulled, and the digital inputs switched. Output voltage is again measured and recorded to determine the differential nonlinearity. Digital inputs are again switched and the process is repeated until the test is complete.

To improve the linearity of a D/A converter, the error voltage due to each digital input must first be measured. A corresponding correction signal can then be added by a lower resolution converter with a full scale value worth only a few LSBs of the device to be corrected. The input to the correction converter is stored in a PROM that is addressed by the LSBs of a common input bus. While this approach works theoretically, correcting an 18-bit resolution converter from 16-bit to 18-bit accuracy would require 256K x 8-bits of PROM storage costing more than the converter itself.

For a converter with no summation errors, the worst case integral linearity error will be less than or equal to half of the worst case differential linearity error. Therefore, if the summation errors and the differential linearity errors of a converter are corrected, the integral linearity errors will be corrected also. This means that instead of correcting every possible input, relatively few corrections can be made giving the same net result.

Suppressing summation errors

Many high resolution converters are composed of several independent internal stages of 4, 8, or 12-bits. The internal architecture can guarantee that these stages don't interact. Therefore, they cannot produce summation errors. Summation errors in the less significant stages may be suppressed with respect to full scale to the extent that they may be neglected. Information about the size and location of the summation errors for any given converter must be determined experimentally.

The only summation errors that are greater than ¼ LSB at 18 bits occur in the four most significant bits. To correct for all of the summation errors, only a 16 x 8 PROM is required. To ensure correction for interstage errors, the PROM can be increased to 32 x 8 — still many orders of magnitude less than a full correction scheme. The only remaining task is to correct for the differential linearity errors in the 13 lesser significant bits. The net result is equivalent to an 18-bit accurate D/A converter.

A semiautomated calibration scheme is shown in Fig. 7. The analog output due to each digital input to be corrected is compared to the output due to the digital input one bit smaller. One LSB of current is added to the output during the smaller input. The correction converter's input is incremented until the correct analog output is obtained. This input value is stored in RAM. The lesser significant bits are corrected in a similar manner, substituting trimming potentiometers for the correction converter.

Instrumentation applications

Instrumentation applications for high resolution D/A converters are numerous and varied. Many of these applications use the converter as a programmable voltage source. If the converter has a current output, it can also be used as a programmable current source. Careful note must be taken of the voltage com-

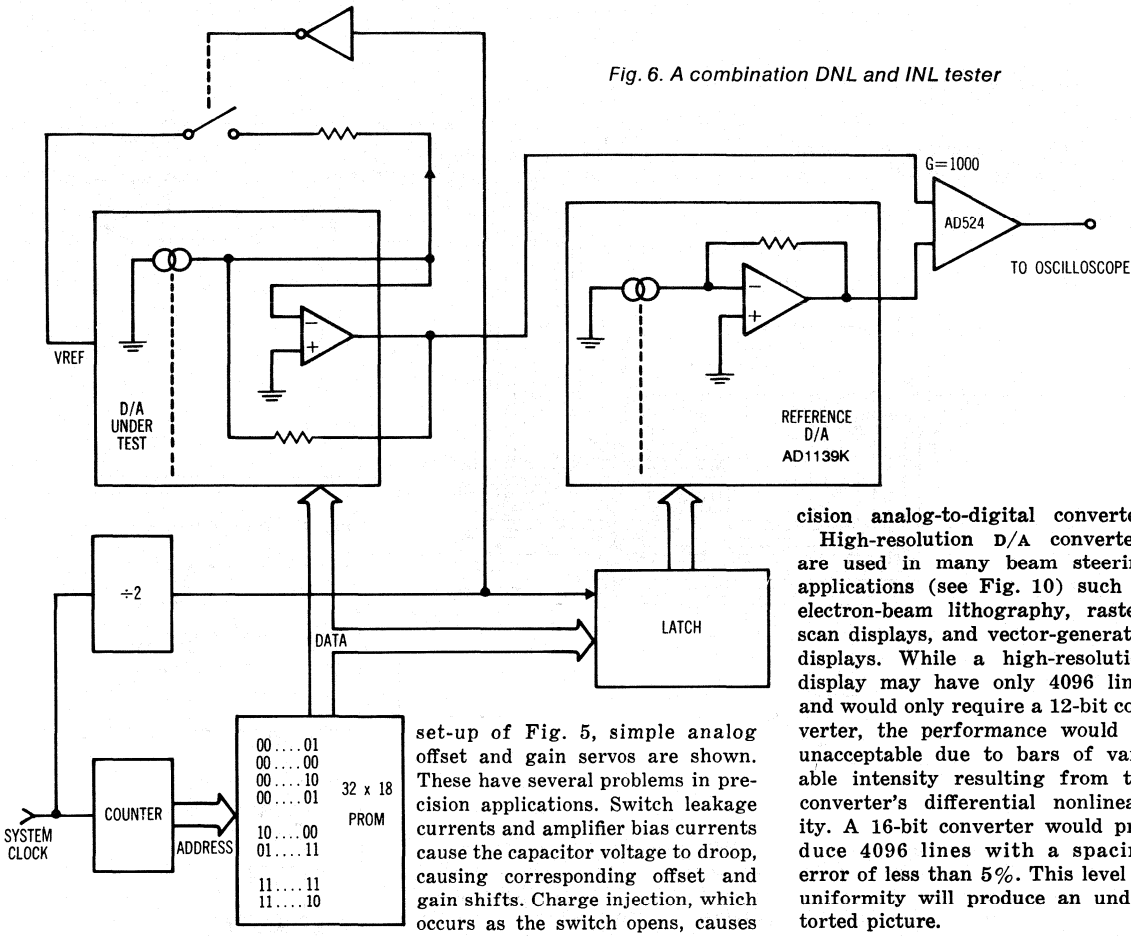

Fig. 6. A combination DNL and INL tester

pliance — the maximum voltage that can appear at the current output terminal while maintaining specified linearity. A more useful feature of a current output device is that the output amplifier can be custom tailored to the application. For example, a low drift amplifier could be used for precision applications, or a power amplifier could be used to provide a large output drive for a programmable power supply (see Fig. 8). The programmable power supply can be used to control the magnetic fields produced by the electromagnets in a cyclotron, as a voltage reference, or to test A/D converters in ATE systems.

A high resolution D/A converter is often used as a digitally controlled potentiometer for use in an auto zero or gain calibration circuit. In the integral linearity test set-up of Fig. 5, simple analog offset and gain servos are shown. These have several problems in precision applications. Switch leakage currents and amplifier bias currents cause the capacitor voltage to droop, causing corresponding offset and gain shifts. Charge injection, which occurs as the switch opens, causes a step at the amplifier output resulting in similar offset and gain errors. The amplifier output constantly ramps up and down due to time delays throughout the circuit, mostly due to the RC time constant of the integrator. This, too, produces gain and offset shifts. All of these problems can be eliminated by replacing the amplifier and switch with a D/A converter and comparator (see Fig. 9). The offset converter is incremented until the converter under test is zeroed, and the gain converter is incremented until the converter under test is adjusted to full scale. The digital inputs are latched, and the correction voltages remain constant and jitter-free until the next auto zero or auto gain cycle. This same technique can be used to make a zero droop sample and hold, or a precision analog-to-digital converter.

High-resolution D/A converters are used in many beam steering applications (see Fig. 10) such as electron-beam lithography, raster-scan displays, and vector-generated displays. While a high-resolution display may have only 4096 lines and would only require a 12-bit converter, the performance would be unacceptable due to bars of variable intensity resulting from the converter's differential nonlinearity. A 16-bit converter would produce 4096 lines with a spacing error of less than 5%. This level of uniformity will produce an undistorted picture.

D/A converters speed assembly

A growing area where high resolution converters are used is in manufacturing and automated assembly, for positioning of robot arms or precision machining (see Fig. 11). A computerized lathe or milling machine can produce a component of up to three feet in length with 0.0005-in. tolerances. A converter with at least 16 bits of resolution is needed to control the depth of cut as well as the horizontal position of the cut. Under microprocessor control, the horizontal positioning converter would be incremented to position the object, while the depth control converter would be adjusted to accurately and repeatedly set the depth of cut.

In multiplying D/A converters, the voltage reference can be varied.

Fig. 7. Improving the linearity of a D/A converter

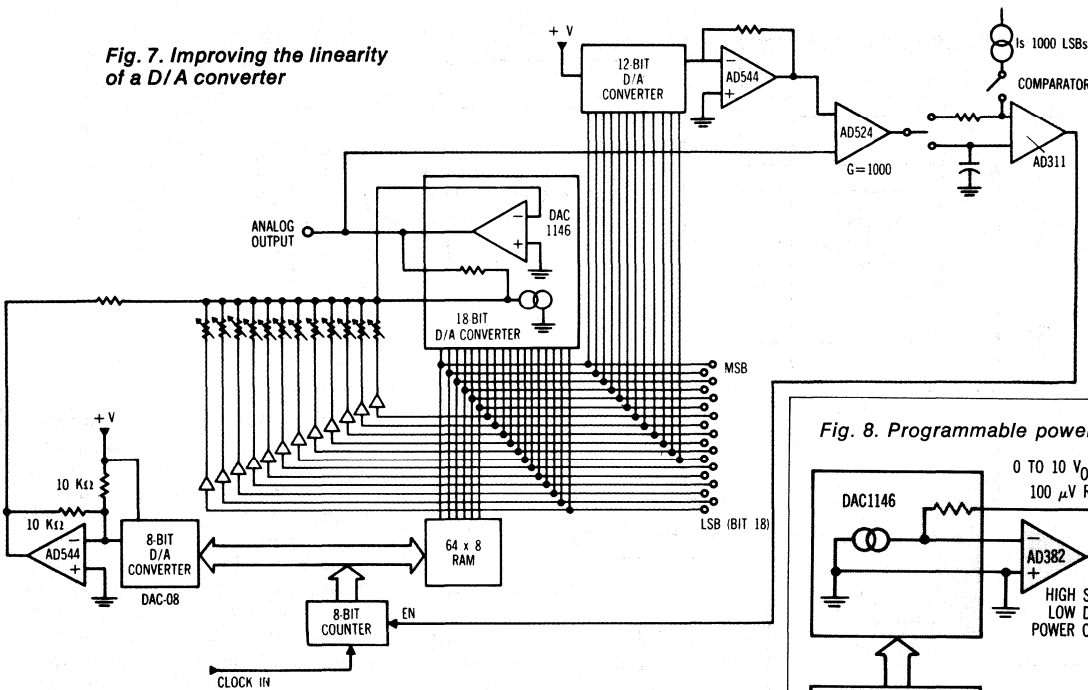

Fig. 8. Programmable power supply

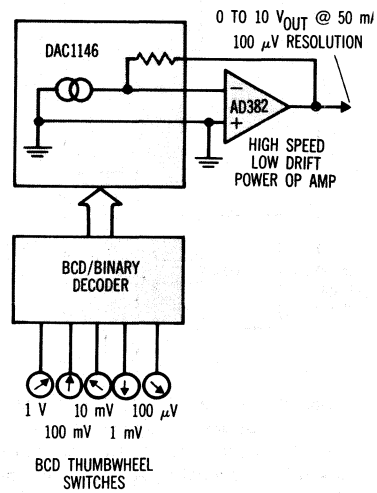

The analog output is the product of the reference input and the digital input. Such converters can multiply in one, two, or four quadrants, depending on the allowed polarity of the voltage reference and the digital inputs.

An interesting application for a high-resolution multiplying D/A converter is a digitally programmed resistor for use as a resistance temperature detector (RTD) simulator (see Fig. 12). For resolution and accuracy of 0.1°C over a −220 to +850°C range, an RTD simulator must be capable of varying from 10 to 400 Ω with 10 mΩ of resolution. This requires a 16-bit converter. The excitation current from the RTD meter may vary by ±5%, so a multiplying converter must be used to maintain a constant resistance. The analog output is proportional to the excitation current and the digital input. As the excitation current increases, the reference voltage increases. This causes the output and the excitation voltage to increase. The simulated resistance is just the excitation voltage divided by the excitation current; therefore, it's dependent only upon the digital input code value. The system simulates a resistance programmed via a digital source at the converter's inputs.

Waveform reconstruction specs

Users of high-resolution D/A converters in waveform reconstruction applications are generally not concerned with differential nonlinearity or any of the other traditional specifications. Instead, a new and highly specialized set of specifications are required. This is largely because many converters used for waveform reconstruction are part of a larger system where a dynamic waveform is digitized and reconstructed. The user's only interest is that this process be completed with a minimum of error, exhibited in several ways, the most significant being total harmonic distortion. Other sources of error are intermodulation distortion, noise, limited dynamic range, poor settling time, and aliasing.

Dynamic range is the ratio of the smallest output signal (1 LSB) to the largest output signal (full scale) that a converter can produce. For an N-bit converter, full scale is equal to 2N LSBs. Theoretically, the dynamic range of an N-bit converter is 6N dB. Converter noise, coupled with inaccuracies in the LSB weight, can reduce this theoretical dynamic range slightly.

Signal-to-noise ratio (S/N) is the

Fig. 9. Offset calibration

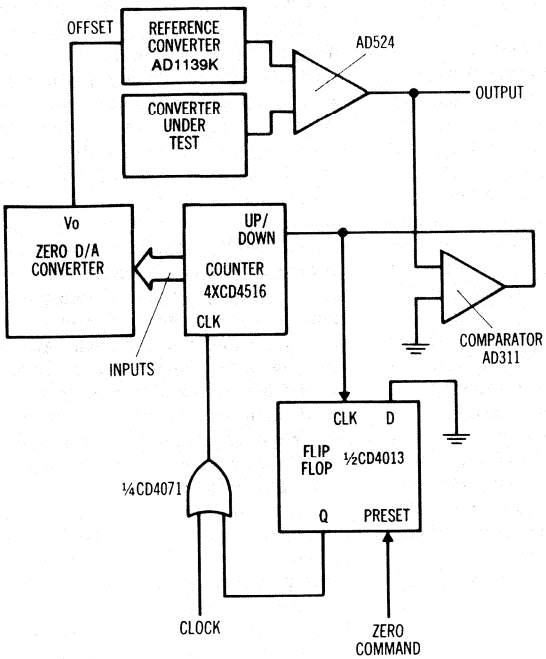

Fig. 10. Beam steering

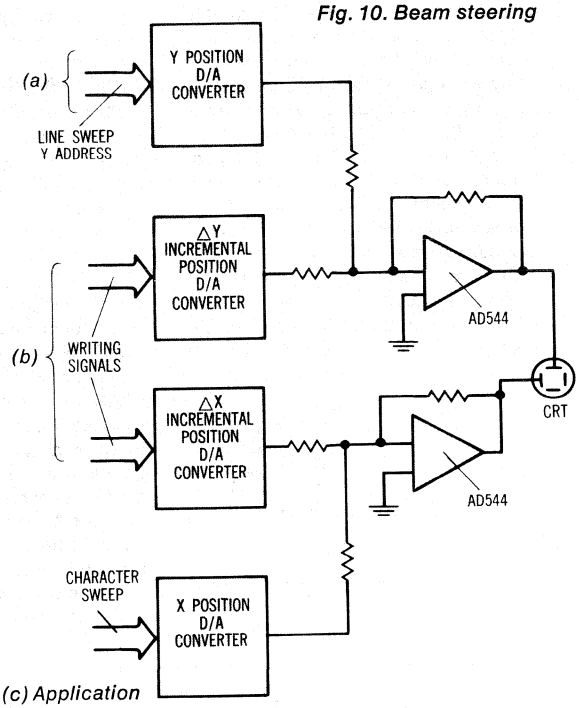

ratio of the maximum RMS signal to the RMS quantization error (see Fig. 13). The RMS value of the maximum sine wave that a converter can produce is the peak output divided by the square root of two, or $Q \cdot 2^{N-1}/\sqrt{2}$, where Q is defined to be the quantization interval. The quantization error increases linearly from $-Q/2$ to $+Q/2$ and then abruptly returns to $-Q/2$. The RMS value of this sawtooth wave is the peak output divided by the square root of three, or: $Q/\sqrt{12}$. The signal-to-noise ratio, then, is: $2^N\sqrt{1.5}$. This can be expressed in dB as S/N = 6.02N + 1.76 dB.

Waveform reconstruction is a dynamic process and requires that the selected converter perform well in a dynamic sense. This means that the settling time must be faster than the time allotted by the process. (Settling time is the length of time between the switching of the digital inputs of the converter and the time when the output reaches and remains within a specified error band around its final value.)

Total harmonic distortion (THD) is the most important specification for a converter used for waveform generation or reconstruction. Roughly speaking, this is the difference between an ideal sine wave and a reconstructed version at the converter's output. THD is defined as the ratio of the square root of the sum of the squares of the RMS values of the harmonics to the RMS value of the fundamental. This means that the RMS energy of each harmonic must be squared and added together. The square root is taken and divided by the RMS energy of the fundamental. The result is the THD of the converter. For a converter with a finite number of digital inputs, N, and associated output voltages, THD can be calculated from the formula:

$$\text{THD} = \frac{\text{RMS error}}{\text{RMS signal}} = \frac{\sqrt{\frac{1}{N}\sum_{i=1}^{N}[E_L(i) + E_Q(i)]^2}}{E_{RMS}} \times 100\%$$

(a) 11 lines of 4096 positioned with 12 bit accuracy. Note dark and light bars

(b) 11 lines of 4096 positioned with 16 bit accuracy. Note uniform bars

where E_L (i) is the linearity error of the converter at sampling point i, and E_Q (i) is the quantization error at sampling point i. Intermodulation distortion is caused by the additional error products produced when the ideal output is composed of two sine waves of different frequency.

When a D/A converter changes state, not all of its switches open and close simultaneously. Internal voltages and currents require finite times to reach their final values. In a CMOS converter, a large amount of charge is stored on the gate-to-source and gate-to-drain capaci-

tances of each switch. This charge is pumped to or from the analog output as the switches open and close. All of these effects combine to produce an output voltage or current spike, commonly known as a glitch. The amplitude of the glitch is nonlinearly related to the converter's input; thus, the glitch produces energy at harmonics of the fundamental frequency and causes harmonic distortion. When the analog output is asymmetric about zero volts, the second harmonic is generated. Asymmetric behavior of increasing or decreasing signals generates third harmonics. Other asymmetric behavior will cause other harmonics (see Fig. 14).

Eliminating the glitches

A deglitcher (see Fig. 15) is a special type of track-and-hold amplifier that holds the output constant from just before the input code changes, to just after the glitch has died out. It then acquires the new, stable output. The deglitcher generates its own output glitch, caused by charge injection from its switching circuits. But this glitch is usually small, of constant amplitude, and independent of the digital input code. It produces energy at the sampling frequency, not at the signal frequency or any of its harmonics, and therefore does not cause harmonic distortion. Band-limiting the deglitcher output with a fixed time constant suppresses distortion caused by slew rate limiting of the output amplifier.

Audio signals are inherently bipolar, and so the converter used to reconstruct them must be bipolar (see Fig. 16). One possible method for configuring such a device is to offset a standard binary converter by half scale. A second method is to sense the polarity of the signal and switch the output amplifier to either an inverter or a buffer as required. The remainder of the input determines the magnitude of the output voltage. Offset binary converters have a simple, clean architecture, and so can be made smaller, more reliable, and less expensive than a sign magnitude converter with the same performance. Alternatively, extra money can be spent on better resistor networks to yield an offset binary converter with improved performance over a sign magnitude device of the same price. For audio applications, however, a substantial amount of signal is around zero volts. An offset binary converter exhibits its worst temperature performance around zero volts because it counts on the MSB to track with the sum of all lesser significant bits, and with the bipolar offset resistor. A sign magnitude converter exhibits its best performance around zero volts because all bits are OFF. Good performance around zero is important because of the effect on S/N ratio. Large signals tend to mask error-induced noise, while small signals will be buried by the same noise. A binary converter with a current output can be converted to sign magnitude

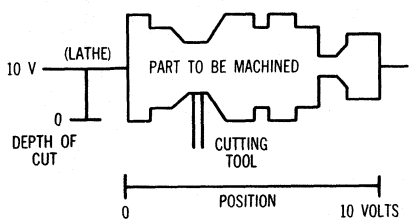

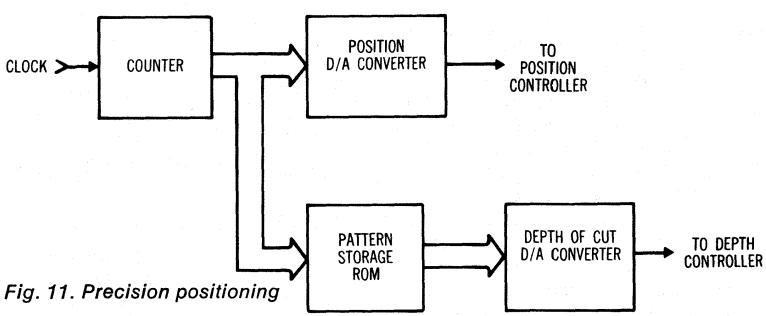

Fig. 11. Precision positioning

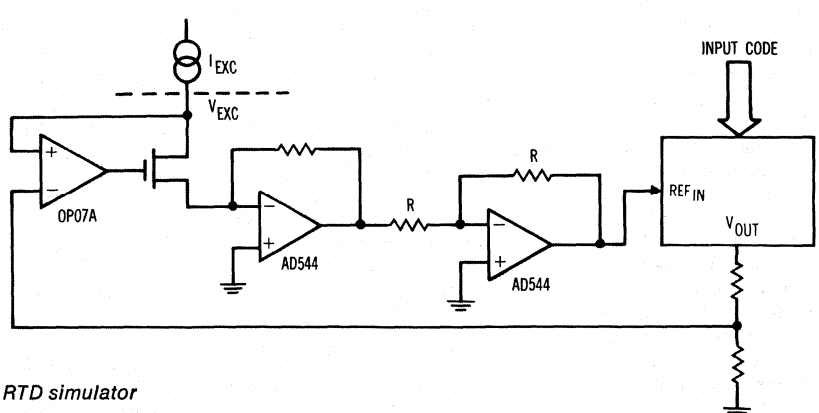

Fig. 12. RTD simulator

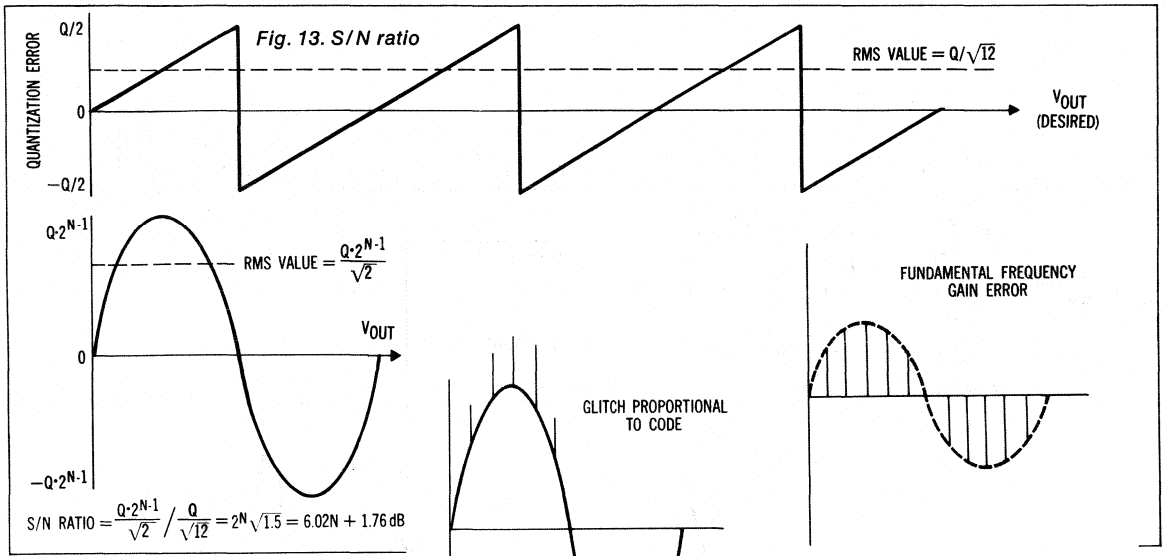

Fig. 13. S/N ratio

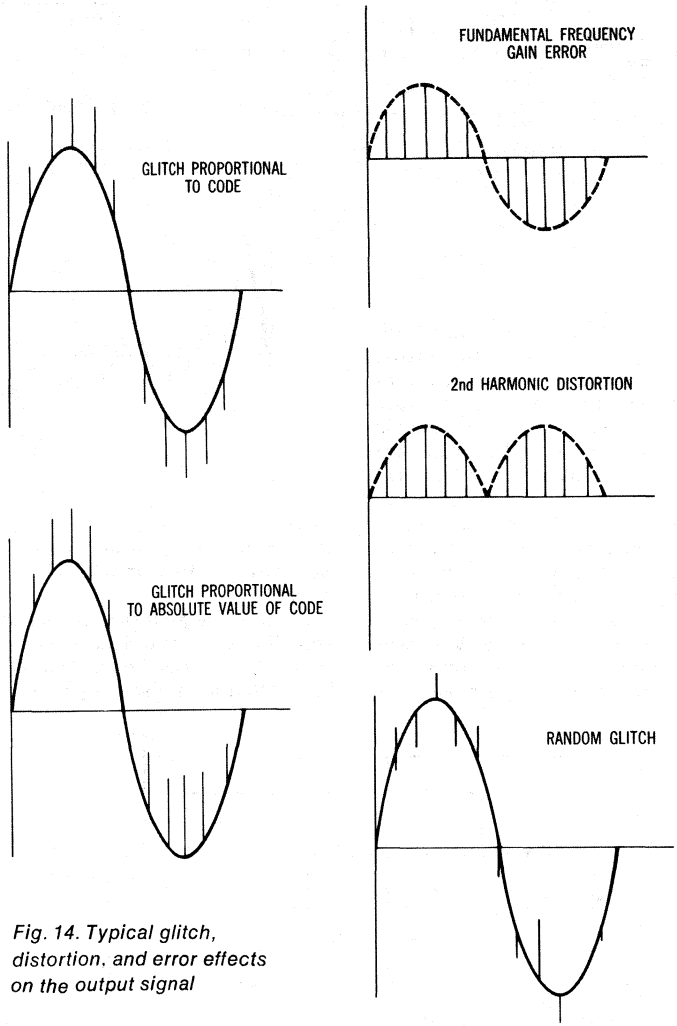

Fig. 14. Typical glitch, distortion, and error effects on the output signal

architecture by the addition of a low-drift amplifier, an SPDT CMOS switch, and an output buffer (see Fig. 17). The result is maximum DNL drift of ±½ ppm/°C for ±⅛ full scale range and ±1 ppm/°C over the full scale range. Typical drifts are ±¼ ppm/°C around zero and ±½ ppm/°C over the full range.

Check total harmonic distortion

Total harmonic distortion can be tested using the circuit of Fig. 18. In this test set up, the PROM contains one cycle of a computer generated sine wave. Frequency select switches are used to program the adder with the number of codes that it should skip on each count. This series of sinusoidally-related digital codes is fed to the converter to generate a staircase approximation of an analog sine wave. If the adder is set to increase the PROM address by one on each count, 4096 inputs will be presented to the converter on each cycle through the PROM. If the adder is set to increase the PROM address by 1024 on each count, only four inputs will be presented on each cycle. In this way, any of 2048 discrete frequencies between 12 Hz and 25 kHz can be generated with a constant 50 kHz sampling rate. The D/A output (see Fig. 19) is deglitched and displayed on the spectrum analyzer. Total harmonic distortion can be computed by comparing the amplitude of the fundamental frequency with the amplitudes of the harmonics as discussed.

A high resolution D/A converter with good dynamics can be used as a precision function generator. Through microprocessor control, the converter can produce the standard output functions — sine waves, square waves, pulses, triangle

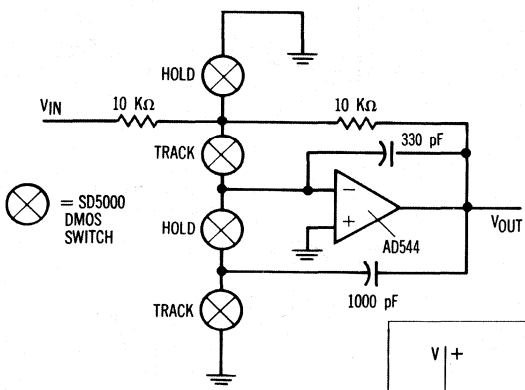

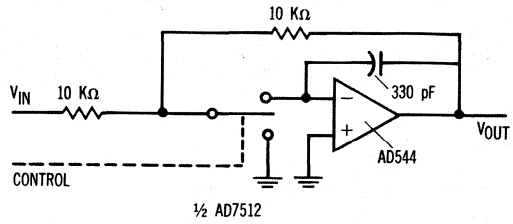

Fig. 15. Deglitcher circuits for D/A converters

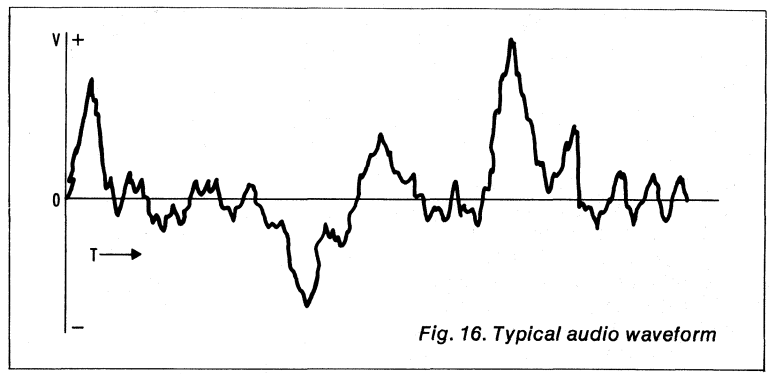

Fig. 16. Typical audio waveform

offset, and inversion are controlled by appropriately scaling the inputs. Frequency and duty cycle are controlled by processor timing. More complex functions such as double pulses, bursts and pulse trains are also easily implemented. The real advantage of using a high resolution device as a function generator is that arbitrary waveforms can be generated. These can be used to simulate complex waveforms, to program test patterns, or to generate hyperbolic or other transcendental functions. A 16-bit converter with a settling time of less than 5 μs can produce outputs at up to 200 kHz with distortion of less than 0.002%.

A very high performance, reasonably low cost, two-channel digital audio system can be implemented by combining a fast, high-resolution converter with a track-and-hold amplifier and two deglitchers. This implementation allows a stereo pair of analog outputs to be simultaneously updated at 50 kHz. The track and hold stores the right-channel output while the left chan-

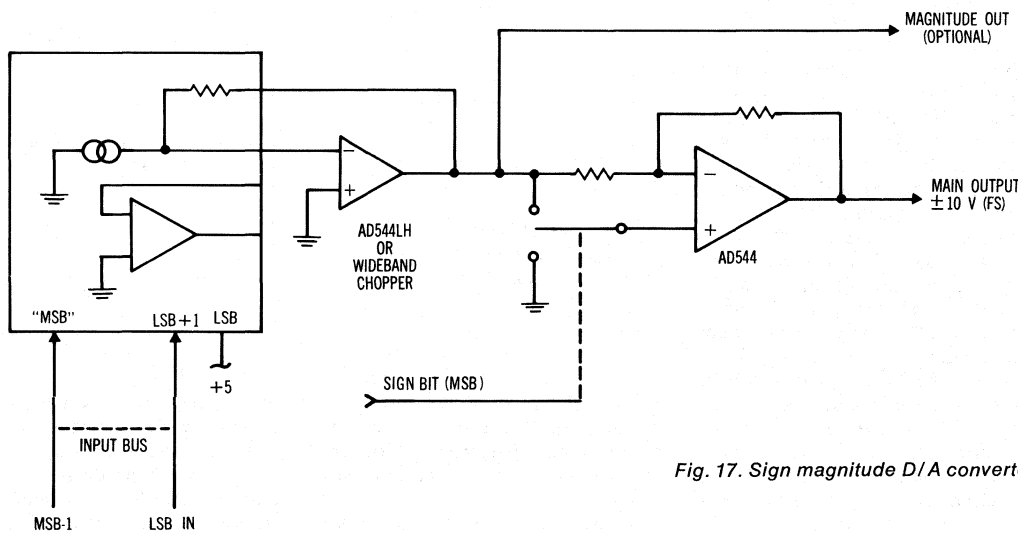

Fig. 17. Sign magnitude D/A converter

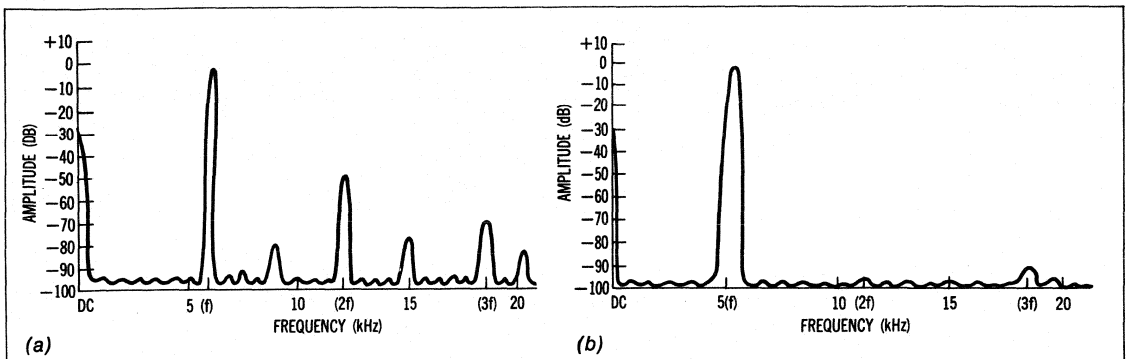

Fig. 19. Spectral distribution of D/A output signal (a) before deglitching and (b) with deglitching circuitry added

nel is presented to the input of the converter. The two deglitchers are then switched from the hold mode to the follow mode, and their outputs are switched from the previous sample to the new sample. The deglitcher response is bandlimited to eliminate distortion caused by slew rate limiting of the amplifier. A time constant of 3.4 μs can ensure the passage of full-power 20 kHz sine waves without distortion. A 16-bit converter used along with a track-and-hold amplifier and deglitchers can produce an audio signal with distortion of less than 0.005%.

Telecommunications provides another major area of application for high resolution converters. Data is digitized for transmission over telephone, broadcast, or satellite communications links and then reconstructed on the receiving end. Again, the significant parameters are THD and settling time, which limits the rate of transmission.

Other uses for high resolution converters in waveform reconstruction include sonar and seismic research. In sonar, a signal is transmitted and the time until it is reflected back to the receiver is measured. The time, stored in digital format, is presented to the converter. A picture of the ocean bottom is subsequently generated. In seismic research, a similar process occurs. An explosive is detonated, and the vibrations are measured using A/D converters at specified locations. Data are processed by a computer, and presented to a D/A converter which is used to generate a seismic profile.

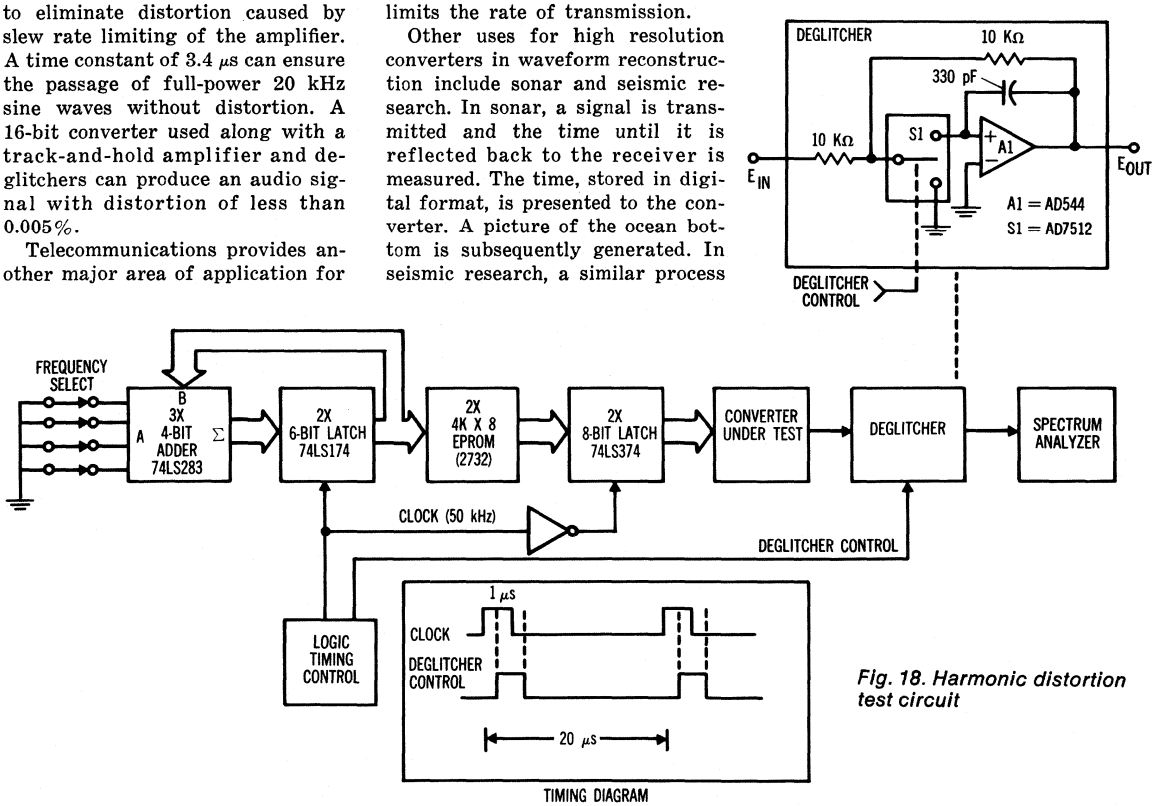

Fig. 18. Harmonic distortion test circuit

ANALOG DEVICES

AN-342
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Analog Signal-Handling for High Speed and Accuracy

by A. Paul Brokaw

You've bought an IC a/d or d/a converter that's specified for 10-bit-and-better resolution and accuracy. Or, you've bought a current-ouput DAC with submicrosecond settling to 1/2 LSB. Much design effort, technological development, and process competence have been expended to solve the hardest part of your interface problem. But . . . you aren't out of the woods yet! Here are some of the issues that you will have to come to grips with to preserve speed, resolution, and accuracy:

1. If your DAC is a current-output type and you want voltage, the use of an op amp requires that you deal with the dynamic and steady-state signal-interfacing problems.
2. You will have to minimize interference introduced via common power-supply connections.
3. You will have to decide where "ground" should be and how to keep it there.
4. If "ground" is remote, you will have to couple to it without reduced accuracy or succumbing to interference.
5. If your analog signal is being converted by a successive-approximations converter, you may have to buffer the source from fast transients incidental to conversion.

To become aware of these potential problems is to have taken the first step towards solving them. Since all circuits and systems differ in important little ways, there are no"cook-book" solutions that can be blithely employed for satisfactory results in all cases. However, a little thought will go a long way towards solving them. The purpose of this Brief is to remind you of some of the things you should be thinking about.

DAC's AND OP AMPS – DYNAMIC PROBLEMS

A current-output DAC is usually connected to the summing point of an inverting op amp, and then the feedback loop is closed via the internal "span" resistor, R_F, as Figure 1 shows. The output impedance of the DAC can generally be treated as a parallel combination of resistance and capacitance. The shunt capacitance, C_O, combines with R_F to add a pole to the open-loop response, which may result in poor closed-loop response.

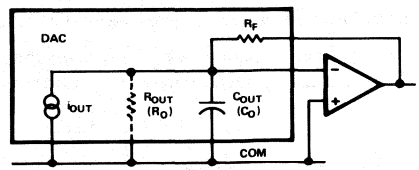

Figure 1. Equivalent circuit of current-output DAC.

Figure 2 shows how the open-loop amplitude and phase response might appear if the spurious pole due to C_O is below the undisturbed system-crossover frequency. Not only will the closed-loop bandwidth be reduced, but –more seriously– excess phase shift will be introduced. The extra phase shift reduces the system frequency stability margins and may cause ringing (and perhaps even oscillation).

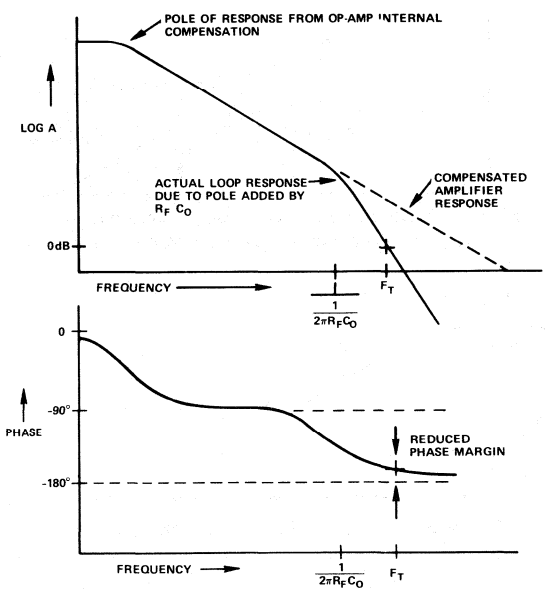

Figure 2. Amplitude and phase response of the circuit of Figure 1. The additional pole increases settling time by reducing bandwidth and increasing both overshoot and ringing.

As Figure 3a shows, the loop-stability margins can be restored by connecting a feedback capacitor, C_F, in parallel with the feedback resistor. This capacitance creates a zero in the open-loop transfer function, which can be adjusted to correct the phase margin. However, if R_{OUT} is very large (as is often the case with current-output DAC's), the large pole-zero mismatch remaining (Figure 3b) may result in slow settling.

Even with finite values of R_{OUT}, a small residual pole-zero mismatch (Figure 3c) may result in long-settling "tails"; the DAC output voltage may appear to settle quickly, but then it slowly changes —by a significant amount— to its final value, over the course of tens of microseconds, or even milliseconds.[1]

The residual mismatch will be eliminated when the DAC-output circuit and the feedback network form a frequency-compensated voltage divider, i.e., when $R_OC_O = R_FC_F$. This condition can usually be satisfied, but sometimes it requires large values of C_F. Unfortunately, C_F—which introduces an open-loop zero—also produces a closed-loop pole, which reduces the overall bandwidth and results in increased settling time.

R_F is generally fixed by the desired DAC gain; the minimum value of C_O is a property of the converter not under the system-designer's control. Therefore, C_F and R_O are the only tw parameters that can be manipulated (reduced). As R_O' (tl effective value of R_O) is reduced by shunting the DAC outp with a resistor, the required value of C_F is reduced, and tl closed-loop bandwidth is increased (Figure 4). The unity-ga bandwidth of the op amp, b, limits the open-loop system ban width, which, in turn, limits the realization of closed-loo bandwidth. As R_O' is reduced, the *open-loop* bandwidt obtainable for a fixed op-amp bandwidth, b, is also reduced.

A compromise can be reached by adjusting R_O' to provide th same open- and closed-loop bandwidth. For a fixed C_O an R_F, the values of R_O' and C_F can be determined from:

$$R_O'C_O = R_FC_F = \frac{1 + \sqrt{1 + 8b\pi R_FC_O}}{4b\pi} \qquad (1)$$

The resistive component at the DAC output also influences th effect ot the amplifier offset (V_{OS}) and noise on the overal output voltage. Both are magnified by $(1 + R_F/R_O)$.

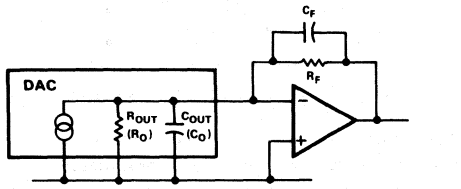

Figure 3a. Improving loop stability by the use of feedback capacitance, C_F.

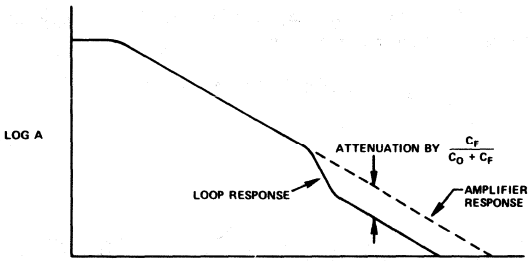

Figure 3b. Response of circuit 3a, neglecting R_{OUT}. Pole- Zero mismatch may yield poor transient response.

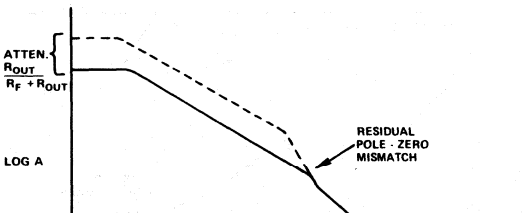

Figure 3c. Response of circuit 3a with finite R_{OUT}.

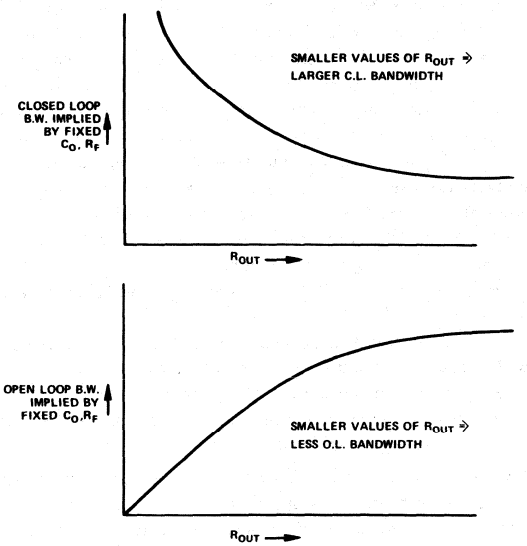

Figure 4. Effect of varying R_{OUT} (R_O') on open-loop and closed-loop bandwidth.

DAC's AND OP AMPS - NULLING PROBLEMS

Perhaps the best way to control V_{OS} in an op amp used with a DAC is at the source—to choose an op amp with sufficiently low offset over the temperature range (such as the AD510). The next-best way is to null the op-amp's offset by the standard V_{OS} trim, taking pains to connect the pot wiper to the appropriate supply terminal *at the device.*[2] The amplifier's offset-trim adjustment should be used *only* for V_{OS} nulling; if it is used to compensate for offsets caused by the flow of bias current through the feedback resistor, as well as for offsets

[1] This process is discussed in some detail, with waveforms, in the Appendix to an article, "Settling Time of Operational Amplifiers," by Bob Demrow, appearing in ANALOG DIALOGUE 4-1 (1970).

[2] The reasons for this are well-documented in the Application Note, "An IC-Amplifier User's Guide to Decoupling, Grounding, and Making Things Go Right for a Change," by the author, available from Analog Devices. A heavily edited version appeared in EDN Magazine, October 5, 1975.

occurring in external circuitry, the amplifier input stage will have to be unbalanced, which will cause its V_{OS} tempco to be degraded.

If the amplifier lacks offset-adjust terminals, or if it is necessary to compensate for the additional sources of offset mentioned above in one convenient place, there are two commonly used ways of providing the trim; they are shown in Figure 5. The more-desirable approach is shown in 5a; the correction is applied to the amplifier's positive input terminal, as a voltage. Since it is effectively in series with V_{OS}, the V_{OS} correction is unaffected by changes of R_O'.

The less-effective way is to introduce a current at the summing point, as shown in 5b. If the resistances in the circuit (including R_O') are constant, there is no problem. However, if R_O' can vary, the output offset will change. If the change of R_O' is a function of the applied *digital code*, the result can be increased differential nonlinearity.

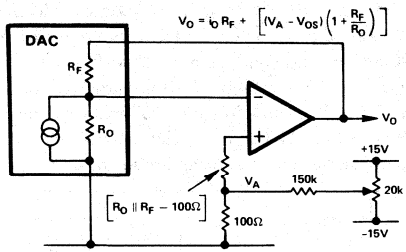

(a) Nulling offset with voltage applied to op-amp reference input.

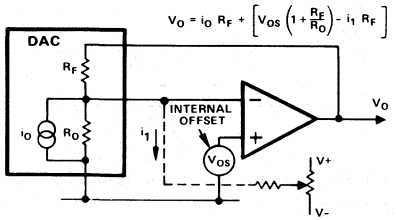

(b) Nulling offset with current added at op-amp summing point.

Figure 5. External offset-null methods.

For example, if the DAC is an inverted R-2R-ladder type, as shown in Figure 6, the output resistance, R_O, approaches R for codes containing many 1's, 3R for codes containing a single 1, and ∞ for all-0's. If R = 10kΩ, the resistance looking back into the network is about 10kΩ for more than four 1's and 30kΩ for a single 1. Thus, for the one-bit transition from 0011111111 to 0100000000, the error voltage, V_{OS} (1 + R_F/R_O), changes from 2 V_{OS} to (4/3)V_{OS}. If the offset had been nulled at all-0's (1 + R_F/R_O = 1, since $R_O \to \infty$), the offset error will be +V_{OS} at the first code and (+1/3)V_{OS} at the second code; the incremental change of error will be (-2/3)V_{OS}. If V_{OS} is not much smaller than the voltage equivalent of the least-significant bit, a tangible error will result. It will be especially pernicious in the case of a multiplying-DAC application with small analog inputs. The solution is simple: use Figure 5a instead of 5b.

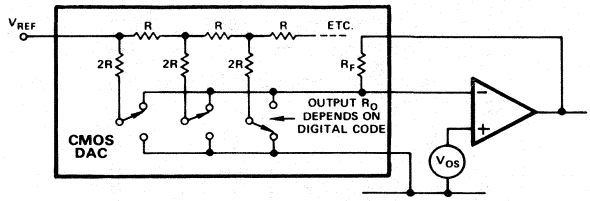

Figure 6. Variable output resistance of inverted R-2R ladder in CMOS and voltage-switching DAC's.

"Foreign" currents in common ground and power lines can introduce offset, noise, and other errors that will be amplified in the same way as V_{OS} errors. It is important to refer the amplifier circuit (and its external V_{OS} trim), the load across which the output voltage is developed, and the DAC's reference input – all of these – to the DAC terminals, in the manner shown in Figure 7.

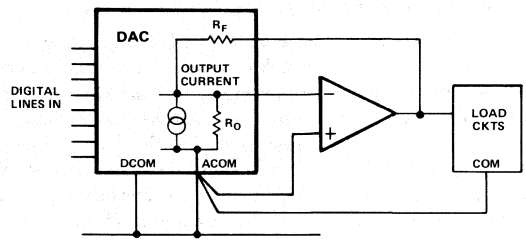

Figure 7. Referring buffer amplifier and load circuits to analog common.

BYPASSING AND DECOUPLING

In "virtual-ground" systems, such as an op amp, driven by a current-output DAC, the DAC output current doesn't actually return to ground, but to one of the power supplies, by way of the op amp's output stage (Figure 8). To reduce the impedance in the high-frequency current path, the bypass capacitor should be connected so as to return the currents from one (or both) power terminals to ground *at the DAC*. If the DAC output is active, it may require bypassing of its own supplies for the same reason.

WARNING: You and your drafting department may have conflicting objectives. Your objective is to design circuits that work and to communicate the important details to whoever assembles them. Your drafting department (or so it may seem) has the objective of drawing nice, neat, squared-off diagrams, in which the lines representing conductors are nicely equipotential. You may have noticed that, in Figures 7 and 8, these niceties have been avoided. The lines are configured to resemble closely the job that the wires perform, converging at the common analog connection. Again, the bypass-capacitor lead, in Figure 8, wends its way purposefully around the op amp's acute angle to its power-supply terminal, rather than shooting straight up to meet the power-supply line (a sure recipe for costly debugging). If you think your drafting depart-

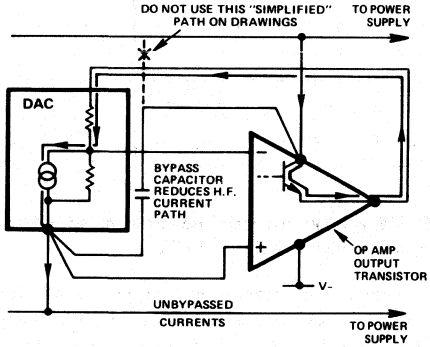

Figure 8. Bypassing power supplies for virtual-ground applications. Arrows show unbypassed current flow.

ment may have a mind of its own, you may want to include a special message for the person who builds the circuit to be sure that it gets built the way you want it built.

Figure 9a shows an example of ineffective decoupling. Here, the op amp drives a load, which connects to a long ground line (returning to the power-supply terminal), and the supply-decoupling for the amplifier returns to the power supply through another long line. The return path for the load current is as long as, or longer than, the supply lines powering the op amp. The "local" decoupling is not only ineffective; it may actually contribute to noise on the power-ground bus.

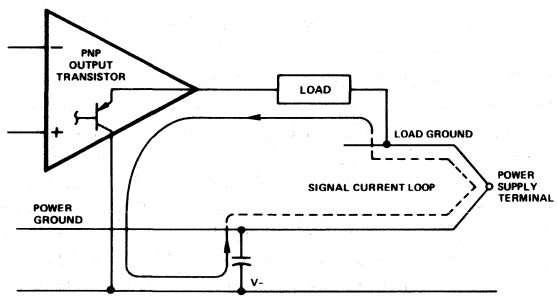

(a) Decoupling for negative supply ineffective.

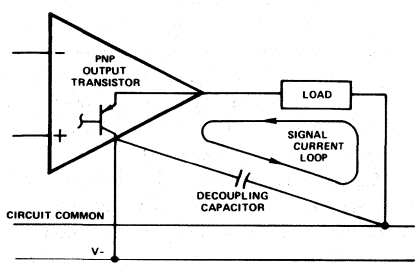

(b) Decoupling negative supply optimized for "grounded" load

Figure 9. Effective and ineffective decoupling.

The cardinal rule of decoupling is: *Make it easy for the current to get back by the shortest path.* Figure 9b shows a more-effective scheme, in which the decoupling capacitor connects by the shortest path between the load return and the load-voltage control element. Here, an op amp, swinging a resistive load-circuit negative, drives the load from an internal PNP transistor, connected to V–. Decoupling the V– pin of the op amp to the low side of the load provides the most direct return path for high-frequency currents, and bypasses them around ground and power buses.

GROUNDING

Great amounts of effort, and many decoupling components, are spent in the attempt to correct problems created by poor ground-current management. In large systems, and in systems which deal with both high-level and low-level signals, "ground" (or common bus) management becomes an important aspect of design. The worst sin– allowing low-level analog signals to share conductors with logic returns or power connections– is an invitation to trouble.

Figure 10a shows an example of a path, shared by digital and analog signals, between the common connections of a converter and an analog system. If the least-significant bit is worth 2.5mV, and digital on-off current fluctuations of the order of 100mA are flowing an a lead with 0.1Ω resistance, the resulting 4-LSB uncertainty (not including spikes and glitches) suggests that it would be folly to waste money on a 12-bit converter, when 10-bit resolution is the very best that can be hoped for, because of wiring limitations.

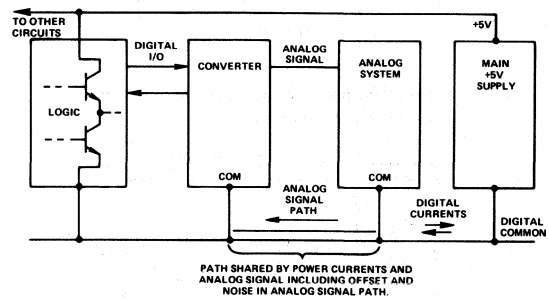

(a) Shared path produces interference and errors.

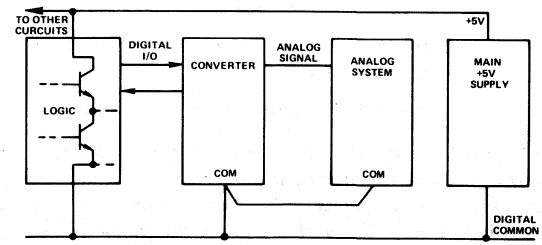

(b) This connection minimizes common impedance between analog and digital (including converter digital currents).

Figure 10. Proper and improper grounding.

As Figure 10b shows, in concept, an analog subsystem can be locally interconnected, with a single-wire connection to the digital common. This signal connection carries only the digital currents required for the *converter's* digital interface. Moreover, analog signals are not forced to share a conductor, even with those currents. The analog subsystem should be powered by a supply with a local common return, which may be connected to the digital common but does not share any cur-

rent-carrying conductors. Ideally, there are no "foreign currents" flowing between the analog system and the digital system, except for those within the converter. If the two systems are joined only at the converter, the foreign currents share the shortest path, and their effect is minimized.

In practical systems, it is often impossible to avoid multiple foreign-current paths. In systems which include several d/a and a/d converters, for example, each converter is a path for digital currents, yet it must have access to the analog signal common. Frequently, the ground problem in such systems can be treated by using an analog common which handles *only* analog signal returns - and a separate system of returns for all digital or high-level signals (Figure 11). Occasionally, a third system of analog power commons may be used to advantage. Since the analog common must be connected to the digital common at no more than one point, safety diodes should be added to any modularized system. These diodes prevent large voltages from developing between ground systems if the key grounding unit, or "Mecca", should be removed from the system.

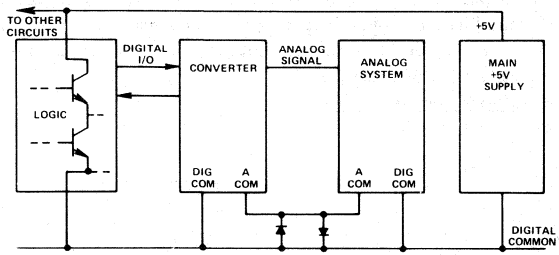

Figure 11. Improved ground current management (analog and digital common must be joined in either converter or analog system. Diodes are fault protection if this connection is broken.)

WHEN COMMON GROUND IS IMPRACTICAL

In large systems, it is often impractical to rely on a single common point for all analog signals. In these cases, some form of differential (or even *isolation*) amplifier is required to translate signals between ground systems. For the inveterate op-amp user, a simple subtractor, or "dynamic bridge" circuit may come to mind. These circuits translate a signal which is referred to one ground system into a similar or amplified signal, referred to a different ground system (Figure 12). The common-mode rejection of the amplifier and a resistance-ratio match are used to eliminate the effects of voltage differences between the two grounds, or common points.

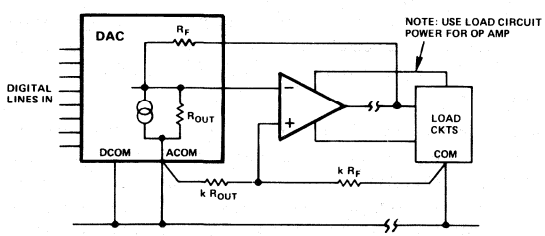

Figure 12. Use of differential amplifier to eliminate the effects of common-mode voltage.

It is generally wise to power the op amp from the power available at the *load* side of the circuit, and/or to decouple it with respect to the *load* common. The reason for this can be deduced from the circuit architecture of the most-common types of op amps (Figure 13).

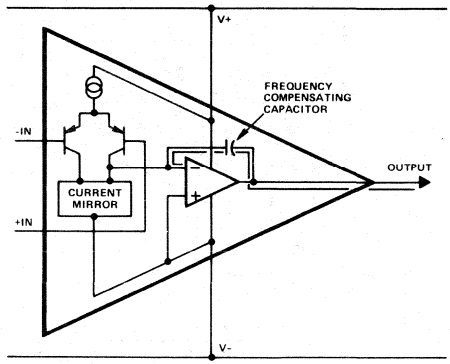

Figure 13. Typical op-amp circuit architecture. Reference for output integrator is V−.

An op amp converts a differential input signal to a single-ended output signal. In many popular op amps, the differential-to-single-ended conversion is done with respect to V− (some use V+), and the resulting signal drives an integrator.[3] The integrator characteristic is used to frequency-compensate the amplifier, and the integrator input is referred to the single-ended output, at V-. The integrator acts as a unity-gain follower for fast signals applied to its non-inverting (or reference) input. As a result, signals applied to the V- terminal have their high-frequency components conveyed directly to the output. Signals having frequency components above the amplifier *closed-loop* bandwidth will be transmitted from V- to the output with little or no attenuation.

As Figure 14a shows, if the op amp used as a subtractor amplifier is powered from or bypassed to the same common line as the input signal, any high-frequency signals associated with that common will appear as part of the output signal. If the ground-noise includes appreciable high-frequency noise (such as logic currents produce), the common-mode rejection will be defeated.

If, on the other hand (14b), the op-amp supply terminals are referred to the *output* signal common, no extraneous signals are coupled into the integrator. Any ground noise appears as a common-mode input signal and is reduced by the common-mode rejection of the amplifier (which is typically very much better than the negative-supply-voltage rejection at high frequencies).

Since noise-rejection performance of the subtractor depends on carefully matched source and feedback resistance ratios, it cannot be used in all situations. Whenever the source impedance cannot be controlled, or is exceptionally high, the subtractor (or dynamic bridge) becomes impractical. In this situation, ground noise and other remote-grounding difficulties can often be avoided by the use of an *instrumentation amplifier.*

[3] The reference mentioned in footnote 2 provides considerable detail regarding the integrator-reference and compensation schemes of some 32 device families.

IC instrumentation amplifiers, such as the AD521, accept differential input signals at high impedance, provide a fixed gain (which can be selected without introducing overall feedback that joins the input and output circuitry), and

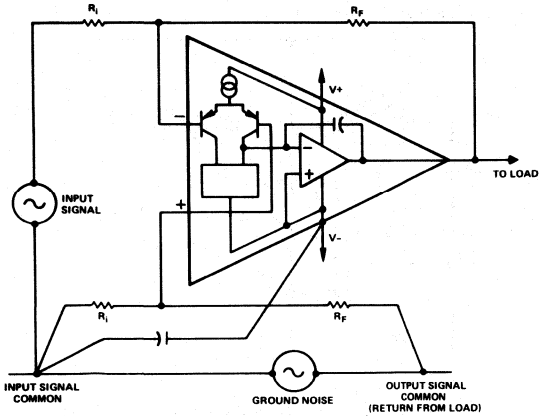

(a) Decoupling to *input* common includes ground noise in the path from the load to the integrator driving the output.

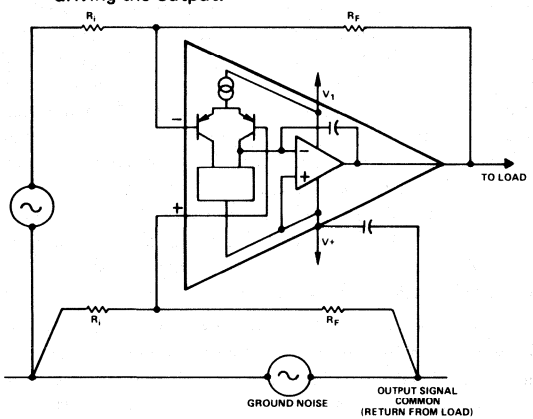

(b) Decoupling to *output* common eliminates ground noise from integrator reference path. Ground noise is minimized in output signal.

Figure 14. Proper and improper decoupling of subtractors using op-amp with integrator referred to V–.

develop the output voltage with respect to a reference terminal, which may be connected to the input common of a remote load-circuit (Figure 15).

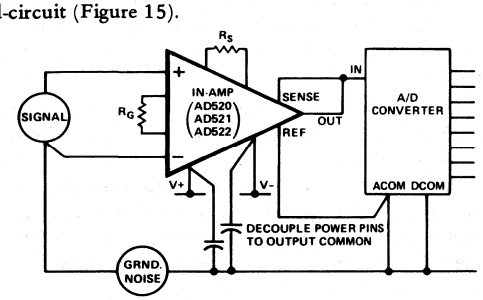

Figure 15. Use of instrumentation amplifier to interface separate ground systems.

Some instrumentation amplifiers are quite versatile and can provide additional functions, while isolating the common returns. For example, the output-reference terminal can be used to add fixed or variable bias voltages to the output.

If the common-mode voltages are very large, or if galvanic isolation is essential for safety, isolation amplifiers, such as the 286 (described elsewhere in these pages), or amplifiers powered by dc-to-dc converters may be highly desirable.

A/D CONVERTERS

The input impedance of many analog-to-digital converters changes during the conversion process and can affect the performance of an amplifier furnishing the input signal.

For example, in successive-approximation converters, the input current is compared to a trial current (Figure 16). The comparison point is diode-clamped, but it may swing plus-and-minus several hundred millivolts. This gives rise to a modulation of the input current. The output impedance of a feedback amplifier is made artifically low by the loop gain. At high frequencies, where the gain is low, the amplifier output impedance rises to its open-loop value. In the case of most IC amplifiers, the open-loop output impedance is a minimum of 25 ohms, due to current-limiting resistors (but is more typically 100-200 ohms).

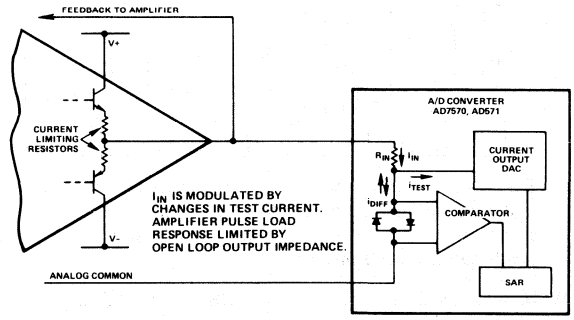

Figure 16. Relationship between successive-approximations A/D converter and op amp that is the source of the input signal.

Even a few-hundred microamperes, reflected from the change in converter loading, can introduce errors in instantaneous input voltage. If the conversion speed and the bandwidth of the amplifier are compatibly fast, the output may return to the nominal voltage before the converter makes its comparison, so that little or no error is introduced. However, many precision amplifiers have relatively narrow bandwidth. This means that they recover very slowly from output transients. Naturally, precision amplifiers are more likely to be used in high-resolution systems, where small errors are less tolerable. As a result, fast, high-resolution systems may suffer from amplifier output-transient errors.

There are a number of ways of solving the problem. Perhaps the easiest is to use an a/d converter that has its own on-board buffer, e.g., the AD572, or most modular types. In other cases, a sample-hold, with low output impedance, may serve as a buffer, as well as to provide the sampling function. Another solution is to use (carefully) a wideband op amp, such as the AD509, which does not include output current-limiting resistors. Finally, it is not difficult to construct an

nside-the-loop buffer that can stiffen the output of a slow, ıccurate amplifier.

'igure 17 shows a simple unity-gain buffer, constructed from ın NPN and a PNP transistor in a compound connection. The ɔutput impedance of this buffer remains low at high frequen- :ies. A good rule-of-thumb for selecting transistor types to use ın complementary-compound is that the input device – in this :ase, the NPN– should be a high-frequency transistor, and :he output device– the PNP– should be a relatively slower :ransistor.

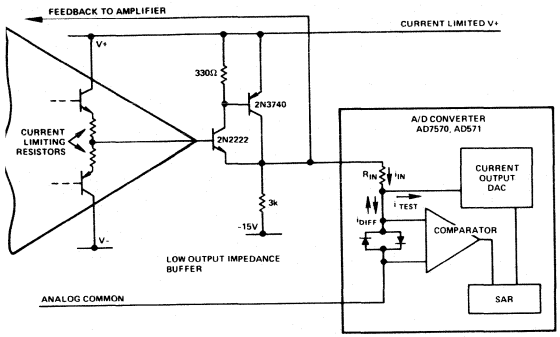

Figure 17. Inside-the-loop buffer provides stiff drive for unipolar ADC.

Since the buffer is not current-limited, a small power-device, capable of pulling down a 200-300mA current-limited supply without damage, has been suggested as the PNP. If the system is definitely safe from overloads, a smaller PNP can be used. This buffer is intended for positive unipolar signals; the 3kΩ resistor provides ample bias to keep the output impedance low over the active range.

A more-complex, protected buffer, for better performance with bipolar input signals is shown in Figure 18. An AD580 voltage reference can be used as a constant-current load to keep the buffer active over the bipolar range. This buffer also includes a bypassed resistor to limit the available output current without pulling down the power bus.

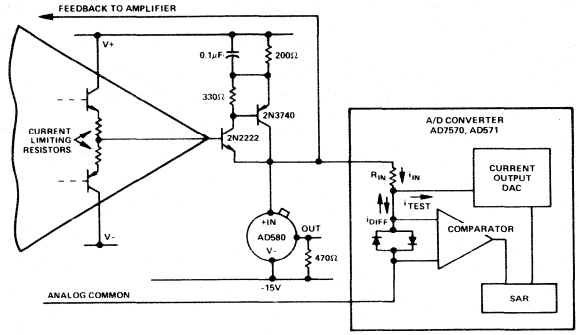

Figure 18. Protected buffer capable of driving bipolar signals into ADC.

TO CONCLUDE

As we told you, our objective was to make you aware of some of the analog problems of implementing interface circuitry, to start you thinking about how to solve them, and to give you some concrete ideas (but not "cookbook remedies"). We hope that they will help make you next system startup somewhat less painful.

ANALOG DEVICES

AN-109
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Understanding and Preventing Latch-Up in CMOS DACs

by Mark Alexander

INTRODUCTION

Many designers now use analog and digital CMOS ICs in their designs to conserve power and increase board functionality. An extensive range of CMOS logic, A-to-D converters, D-to-A converters and op amps are currently available; however, some engineers using these devices for the first time are either unaware of, or are intimidated by the tendencies of some junction isolated CMOS ICs to latch up under certain conditions. Latch-up is defined as the generation of a low impedance path between the power supply rails by the triggering of a parasitic four-layer bipolar structure commonly referred to as an SCR (Silicon Controlled Rectifier). This parasitic SCR is inherent in all CMOS input and output circuitry, and applies to both analog and digital integrated circuits. This application note examines the mechanisms that cause SCR action in CMOS integrated circuits, discusses the pitfalls that are commonly encountered in circuit design that can lead to latch-up, and finally presents techniques to eliminate it. The reader can skip the sections on theory if desired and proceed straight to the sections on Identifying the Causes of Latch-up, and Latch-up Prevention Techniques.

SCR OPERATION

Prior to discussing latch-up in CMOS ICs, it is useful to review the basic theory of SCR operation to gain an understanding of the underlying mechanisms that cause latch-up. An SCR has the PNPN structure and characteristics shown in Figures 1 and 2, respectively. The SCR is a normally-off device that looks like a reverse-biased diode until triggered. Then it latches and conducts a high current until the current falls below a minimum holding value. The four-layer SCR can be modelled with two bipolar transistors, one NPN and one PNP, connected as shown in Figure 3. The resulting device has three main terminals – an anode, a cathode and a gate, plus a secondary gate which is connected to the base region of the PNP transistor. Conduction of current through the SCR is initiated by injecting sufficient current into the base of Q_2 to turn it on. When this is done, Q_2 draws collector current through the base-emitter junction of Q_1. Consequently, Q_1 turns on, which causes additional current to be injected into Q_2's base. This in turn causes Q_2 to turn on even harder, thus supplying more base current to Q_1. The positive feedback arrangement formed by Q_1 and Q_2 sustains SCR conduction even when the gate current is removed altogether. Once triggered, the device will remain in this low impedance state indefinitely until one of two things happen: if the voltage applied across the SCR is reduced to the point where Q_1 or Q_2's emitter-base junction turns off, then Q_1 or Q_2 will cease to conduct because it has no base current and the SCR will turn off; alternatively if the current through the SCR is reduced below the minimum holding current required to sustain conduction, it will also turn off.

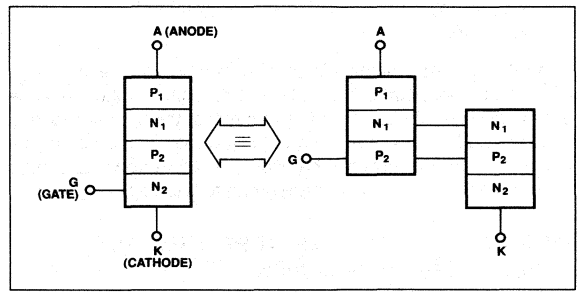

FIGURE 1: *An SCR has a four-layer structure as shown, but can also be thought of as two bipolar transistors merged together.*

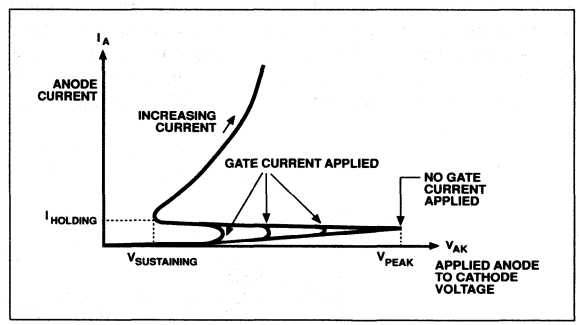

FIGURE 2: *The typical characteristics of an SCR indicate that once triggered, the device conducts a high current generally limited only by the power supply and load.*

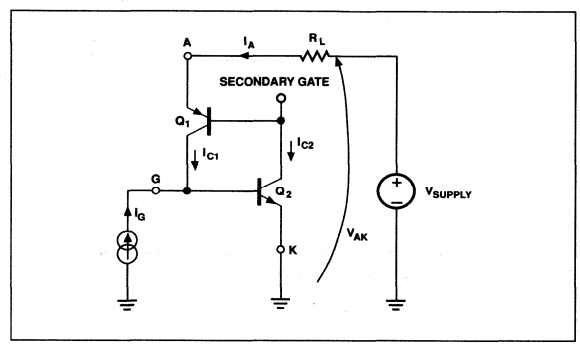

FIGURE 3: *The bipolar transistor model of an SCR shows the positive feedback loop formed when the collectors of Q_1 and Q_2 feed the base terminals of each other.*

There are three ways in which the SCR can be triggered into a conducting state. The first and most direct way is to externally inject a sufficiently large current into either of the gate terminals.

A second and less obvious way is to increase the applied anode to cathode voltage (V_{AK}) until avalanche breakdown occurs in the N_1-P_2 junction. This effectively causes current injection into the gate region of the SCR and will cause it to latch on and continue conducting even when the applied voltage is reduced below the peak sustaining voltage of the device. The third way to trigger the device is to apply a sufficiently high $\Delta V_{AK}/\Delta t$, or slew rate,so that current is injected into the gate region through the depletion layer capacitance of the N_1-P_2 junction. Although the junction depletion region capacitance decreases with increasing V_{AK}, thus reducing the amount of current injected, high enough slew rates can still cause sufficient current injection into the gate region to make the SCR latch on.

It should be noted that the effects of breakdown voltage and $\Delta V_{AK}/\Delta t$ on the trigger point of the SCR are intensified by the presence of any gate current flow. That is to say that the peak sustaining voltage and the maximum allowable slew rate both decrease as the applied gate current increases. These effects are often noticeable in CMOS ICs that are prone to latch up.

DIGITAL INPUT AND OUTPUT RELATED SCR STRUCTURES IN CMOS DACs

A typical unbuffered CMOS inverter is shown in Figure 4 for simplicity, and contains one N-channel MOSFET with its source and body tied to DGND, and one P-channel MOSFET with its source and body tied to V_{DD}. This input and output structure, while shown together in an inverter for convenience, is characteristic of all the digital inputs and outputs on PMI's CMOS DACs. The diodes shown from the output terminal to DGND and V_{DD} are body diodes inherent in the construction of the MOSFETs themselves,

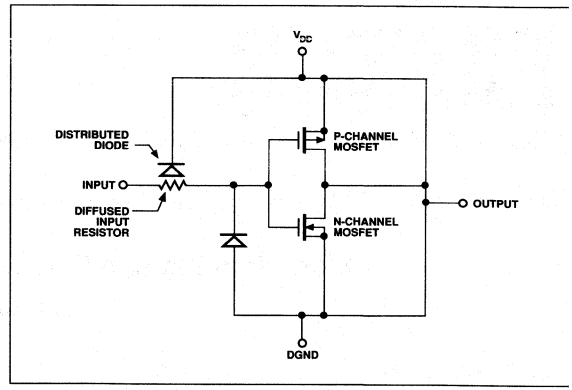

FIGURE 4: *A simple CMOS inverter includes input protection diodes to improve the ESD sensitivity of the device. The output diodes are formed because the bodies of the MOSFETs are electrically connected to V_{DD} and DGND.*

while the input resistor, input diode to DGND and the distributed diode to V_{DD} are deliberately included to clamp the gates at the power supply rails. This input protection circuitry greatly reduces the sensitivity of the MOSFET gates to electrostatic discharge (ESD) damage due to oxide rupture, and is now common on almost all CMOS integrated circuits.

Forward biasing any one of the four diodes shown in Figure 4 can lead to SCR action if the current through them is high enough. Figure 5 depicts a cross sectional view of the CMOS inverter, showing how the N- and P-channel MOSFETs are fabricated. The P-channel device is formed directly in the N– type substrate, while the N-channel device sits in a P– type well. A parasitic four-layer

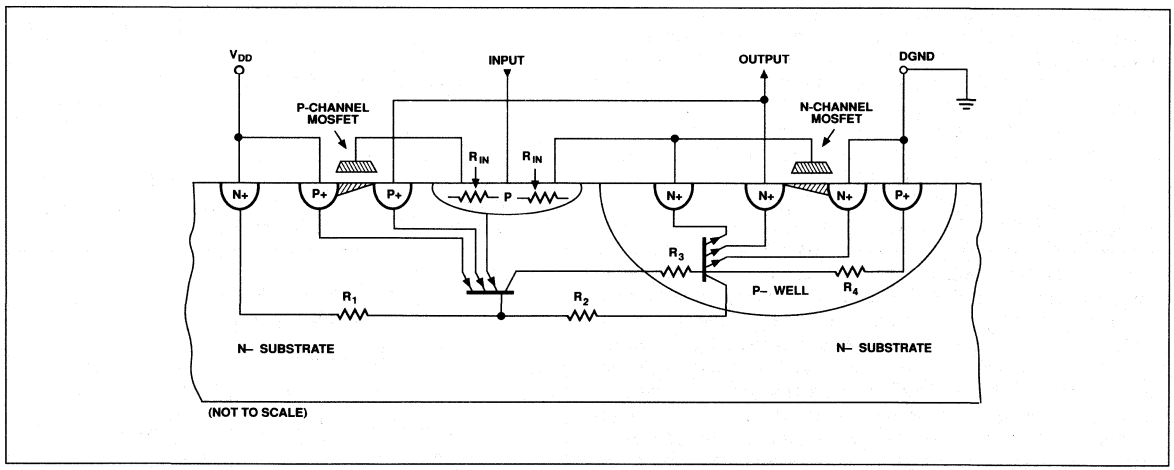

FIGURE 5: *This cross sectional view of a simple CMOS inverter with its input protection diodes clearly shows the presence of parasitic PNP and NPN transistors.*

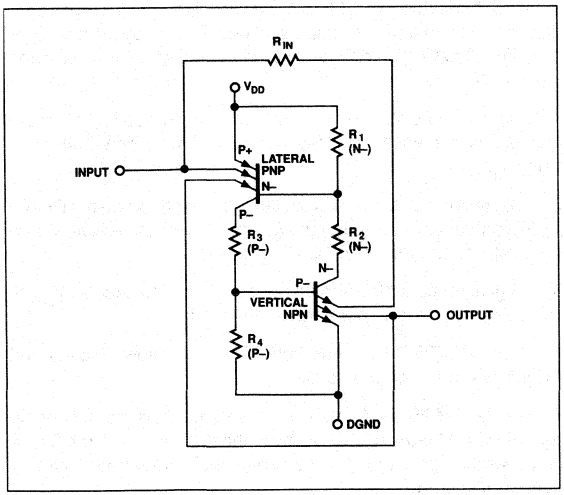

FIGURE 6: *Rearranging the parasitic bipolar transistors shown in Figure 5, gives the easily recognizable circuit of an SCR.*

SCR, shown schematically in Figure 6, is formed when the N and P-channel MOSFETs are located in close proximity to each other. The multi-emitter vertical-NPN transistor results from the fabrication of the N-channel MOSFET in the P– well, and the multi-emitter lateral-PNP transistor results from the fabrication of the P-channel MOSFET in the N– substrate. Fortunately, the PNP transistor has a wide base region and consequently a low current gain, β_F, usually much less than one. Thus, it typically takes hundreds of milliamps of current to trigger SCR action. Notice that the parasitic SCR is connected directly across the power supply rails of the CMOS device, and when triggered, presents a low impedance path that will cause excessive current to flow. This condition is potentially destructive if allowed to persist, since it will result in permanent damage to the bond wires or chip metallization due to localized overheating.

In a CMOS DAC no explicit gate terminal is available, but triggering the parasitic SCR can still be accomplished in a number of ways. The most common cause is when the input voltage exceeds V_{DD} by about 0.3 to 0.7 volts. Then the parasitic PNP transistor's emitter-base junction turns on, thus causing SCR action. The input undervoltage situation is different, however, because there is an input resistor in series with the protection diode connected to DGND. The input can thus be swung further below DGND (usually by as much as 2-3 volts) before the diode will conduct. This means that under normal operating conditions, there is more of a margin of protection against input undervoltage than input overvoltage. Similarly, if the output terminal is allowed to exceed V_{DD} by about 0.3 to 0.7 volts, then conduction of the lateral-PNP transistor initiates SCR action. Finally, pulling the output below DGND by more than 0.3 to 0.7 volts causes the vertical-NPN transistor's emitter-base junction to turn on and also causes the SCR to trigger.

ANALOG GROUND AND OUTPUT RELATED SCR STRUCTURES IN CMOS DACs

8

There are two other ways in which latch-up can be triggered in a CMOS DAC, in addition to the digital input/output overvoltage situation already mentioned. CMOS DACs such as those manufactured by PMI include analog circuitry on the same chip as the digital logic, as shown in Figure 7. Currents flowing from the thin-film resistor R-2R ladder network are switched either to AGND or I_{OUT} (a summing node) using N-channel transistors. These are contained in an isolated P– well, and have one drain connection and two sources. This second P– well is connected to analog ground and *not* digital ground.

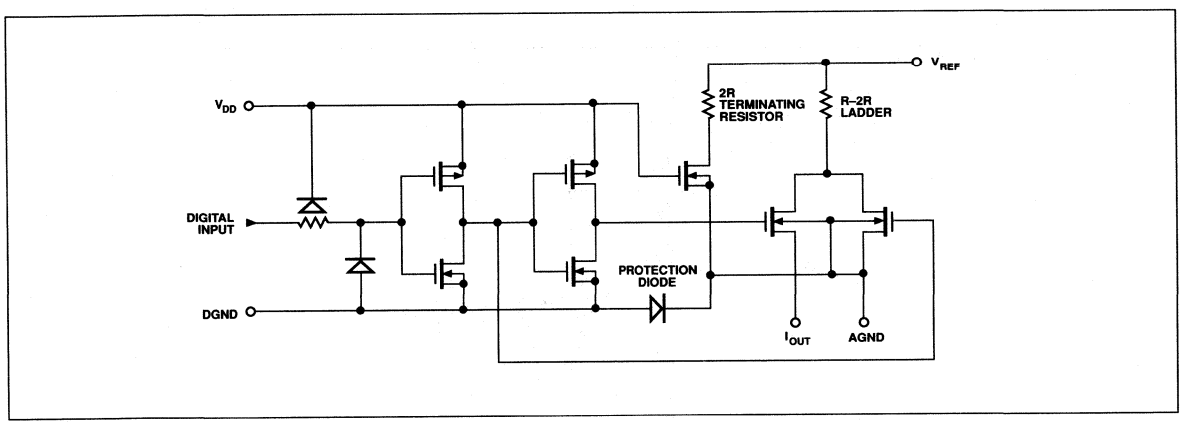

FIGURE 7: *This simplified diagram of a CMOS DAC with a thin film R-2R ladder shows the digital input circuitry plus the NMOS current switches. A protection diode is included to improve the ESD tolerance.*

Figure 8 shows how an additional parasitic SCR structure is formed by the NMOS current switches. If the I_{OUT} terminal is pulled below AGND by more than 0.3 to 0.7 volts then the parasitic NPN transistor associated with the analog P– well is turned on. This may result in SCR latch-up when the power supplies are turned on and sufficient current is flowing out of I_{OUT}.

Note that an additional clamping diode is included from DGND to AGND as shown in Figure 9, being formed by an N+ diffusion in a separate P– well. The purpose of this diode clamp is to reduce the ESD sensitivity of N-channel current switches. However, the protection diode from DGND to AGND adds another parasitic SCR to the device structure. If DGND is allowed to exceed AGND by more than 0.3 to 0.7 volts, then the parasitic NPN transistor is forward biased, and the probability that the device will latch up when the power is turned on is high. Both of these situations should be avoided.

IDENTIFYING THE CAUSES OF LATCH-UP

Now that the mechanism of SCR latch-up has been understood, potentially dangerous situations can be identified that could cause catastrophic failure of a CMOS IC. A list of latch-up inducing conditions to be avoided is an important checklist for system designers. The following guidelines should be followed for all designs using CMOS ICs.

1. Digital inputs should not be allowed to exceed V_{DD} by more than 0.3 volts at any time. This includes a power-down situation when V_{DD} = zero.

2. Digital inputs should also not be allowed to go below DGND by more than about 2 to 3 volts (many data sheets usually give –0.3 volts as the maximum limit, however).

3. Digital outputs also should not be allowed to exceed V_{DD} or DGND by more than ±0.3 volts at any time.

4. For a CMOS DAC, DGND should not be allowed to exceed AGND by more than 0.3 volts.

5. For a CMOS DAC also, I_{OUT} should, in general, not be allowed to fall below AGND by more than 0.3 volts. Some DACs can tolerate significant I_{OUT} current flow, however, without any danger of latch-up.

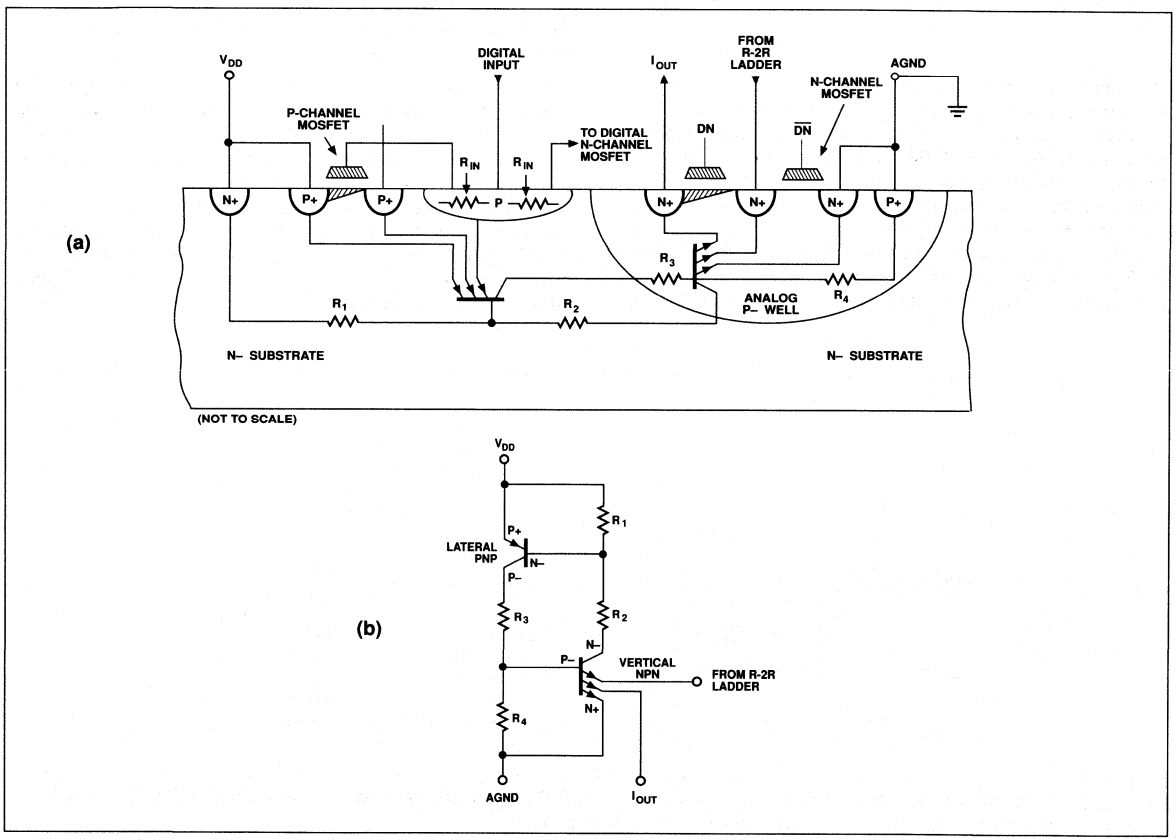

FIGURE 8: *The NMOS current steering transistors of a CMOS DAC shown in (a) form an additional parasitic SCR structure shown in (b) that can cause latch-up when I_{OUT} is pulled more than a base-emitter drop below AGND.*

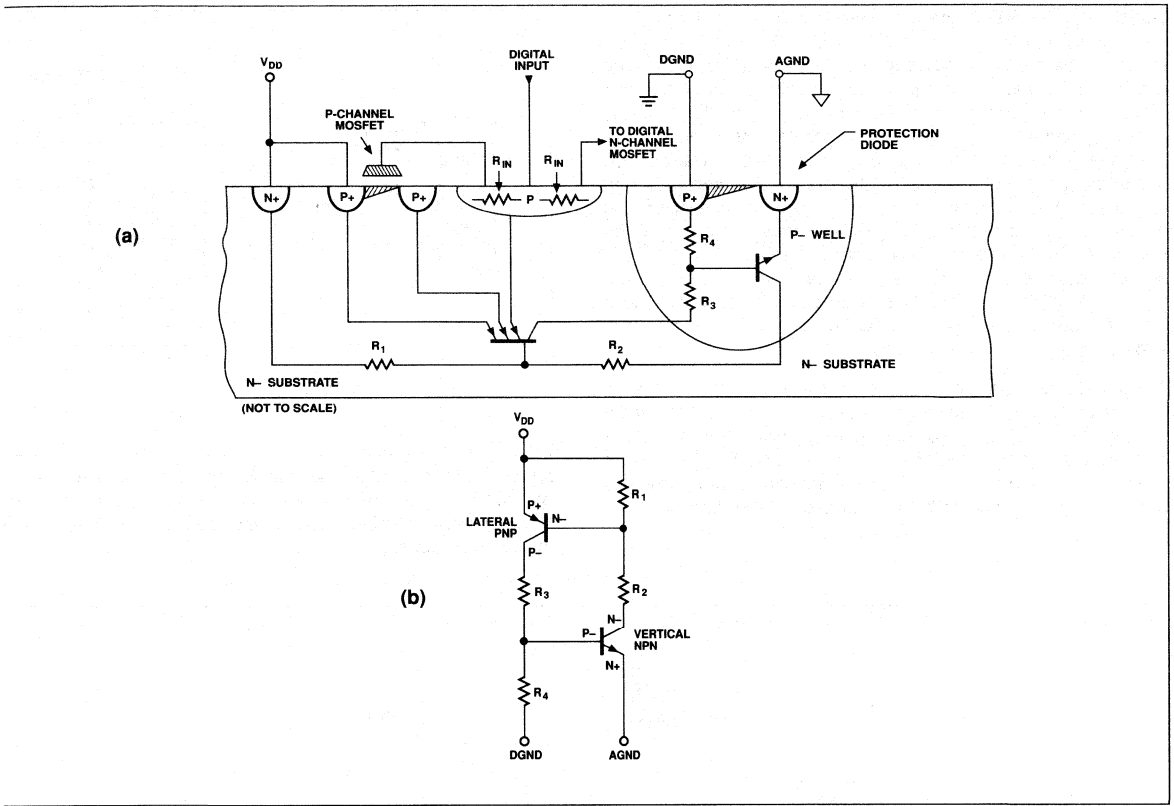

FIGURE 9: *A clamp diode shown in (a) from DGND to AGND reduces ESD sensitivity, but unfortunately adds an additional parasitic SCR to the DAC structure as shown in (b).*

Even a normally benign system design can sometimes unknowingly violate one or more of the aforementioned rules. Consider the case where a CMOS logic IC, with its power supply at zero volts (i.e., turned off), has its inputs driven by another logic IC that is powered from a different power supply. Here the upper input protection diodes of the IC being driven are forward biased, and if enough current flows through them then SCR action can occur. This situation is more common than most engineers would care to admit. It is especially noticeable in large systems that use local on-card regulators to power the ICs on board, but distribute signals amongst many boards through a backplane interconnection scheme. When raw power is applied to all boards, some of the on-card regulators start up before others, thus allowing signals to be applied to ICs on boards that have not yet powered up. Figure 10 shows how this can happen. This violates rule No.1 and could definitely cause some devices to latch up and conduct high currents after their power supplies came up to the normal operating voltages.

Other problems can occur even on single board systems that use multiple power supplies for the analog and digital circuitry.

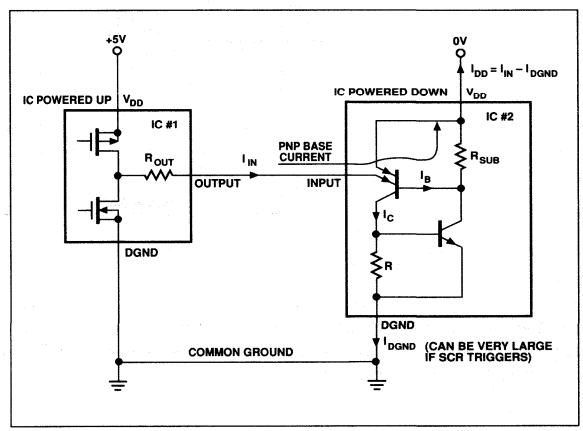

FIGURE 10: *When an unpowered CMOS IC has any inputs driven more than +0.3 to +0.7 volts above DGND, the parasitic PNP transistor will begin to conduct. Note that before the SCR triggers, the base current for the lateral-PNP transistor flows out of the V_{DD} pin of IC #2.*

Consider a CMOS DAC that is powered from a +12 volt supply, but is driven by on board +5 volt CMOS logic. If the +12 volt supply comes up before the +5 volt supply, then there will be no forward biasing of the digital input protection diodes on the DAC, and latch-up will not occur. On the other hand, if the +5 volt supply comes up first and one or more of the digital inputs to the DAC are high, then the upper input protection diodes are forward biased and latch-up will probably occur if there is enough current flowing into these inputs. With more and more data acquisition and analog output boards being designed for personal computers, the lack of a guaranteed turn-on sequence for the +5 volt and ±12 volt supplies that most PCs typically use can be a hazardous situation. Indeed, improper power supply sequencing is the cause of almost all cases of latch-up in CMOS ICs no matter what they are – logic devices, DACs or op amps.

Additional problems can also be encountered when CMOS logic outputs are used to drive highly reactive loads or unterminated transmission lines. Either of these can cause sufficient overshoot and undershoot to forward bias the parasitic input or output diodes, and can also lead to latch-up induced failures, if there is sufficient current flow for the time duration of the ringing.

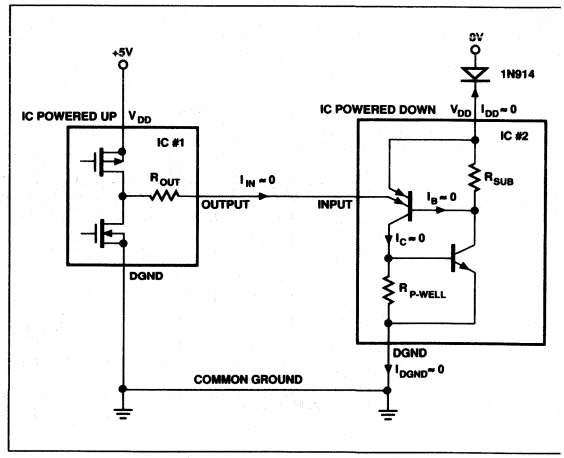

FIGURE 11: *Adding an inexpensive silicon diode in series with the V_{DD} pin of the unpowered IC, effectively prevents the parasitic lateral-PNP transistor's base current from flowing and inhibits SCR action.*

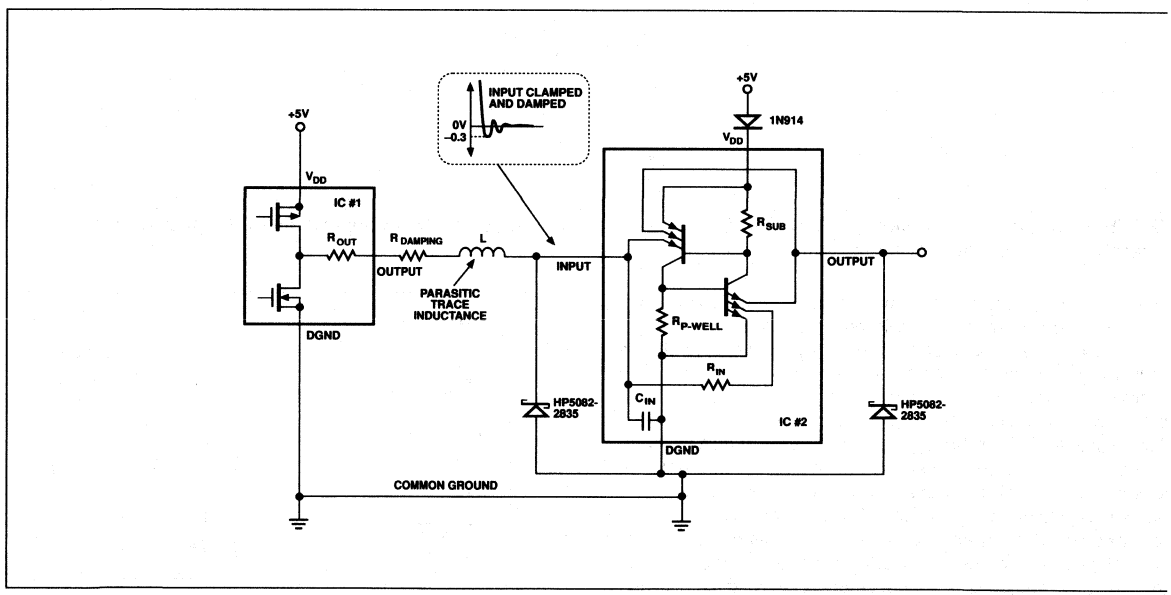

FIGURE 12: *Adding Schottky diodes from the inputs and outputs of a CMOS IC to DGND, protects against undervoltages from causing conduction of the parasitic NPN, thus inhibiting SCR action. The series damping resistor makes ringing due to long PC board traces die out more quickly.*

Fortunately, latch-up problems with modern CMOS ICs are becoming rarer and rarer due to special design techniques that are being incorporated to raise the required input or output current trigger levels to very high values (hundreds of milliamps). However, rugged ICs are still no substitute for good design practice.

LATCH-UP PREVENTION TECHNIQUES

The following recommendations should be implemented in general, for all CMOS designs that violate one or more of the previously discussed rules:

1. If the digital inputs or outputs of a device can go beyond V_{DD} at any time, a diode (such as a 1N914) connected in series with V_{DD} will prevent SCR action and subsequent latch-up. This works, because the diode prevents the base current of the parasitic lateral-PNP transistor from flowing out of the V_{DD} pin, thus prohibiting SCR triggering. This is shown in Figure 11.

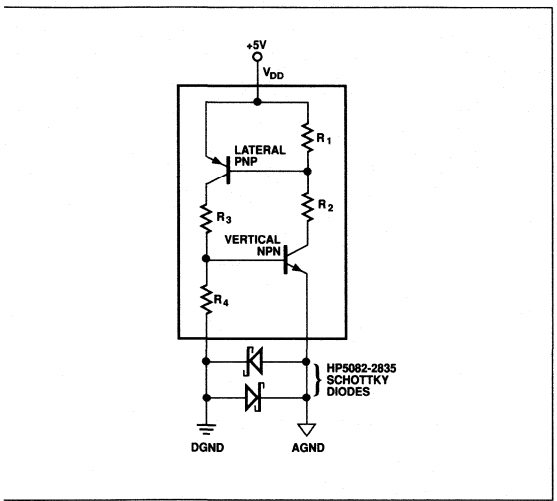

FIGURE 13: *Connecting Schottky diodes between DGND and AGND prevents conduction of the parasitic NPN transistor, and helps to minimize injected noise from DGND to the analog output.*

2. If the digital inputs or outputs of a device can go below DGND at any time, a Schottky diode (such as an HP5082-2835) connected from those inputs or outputs to DGND will effectively clamp negative excursions at –0.3 to –0.4 volts. This prevents the emitter-base junction of the parasitic NPN transistor from being turned on, and also prevents SCR triggering. Figure 12 shows the connections for the Schottky diodes.

3. If the DGND potential can occasionally exceed AGND by more than 0.3 volts, a Schottky diode placed between the two pins of the device will prevent conduction of the associated parasitic NPN transistor. This provides additional protection against latch-up as shown in Figure 13. An extra diode connected in inverse parallel with the one just mentioned provides clamping of DGND to AGND in the other direction and will help to minimize digital noise from being injected into the DAC.

4. In circuits where the I_{OUT} pin of a CMOS DAC can be pulled below AGND, another Schottky diode clamp between these two terminals will prevent sensitive ICs from latching up. This condition sometimes occurs with high-speed bipolar operational amplifiers that are used as current-to-voltage converters following the DAC. During power-up or power-down transitions, the op amp's inverting input presents a low impedance from I_{OUT} to the negative supply rail. An unprotected DAC may not be reliable without the requisite Schottky diode clamp to AGND.

5. Finally, in designs that have long digital PC board traces between components and are prone to inductive ringing problems, a series damping resistor of 10–100Ω often will be beneficial. This resistor increases the damping factor of the equivalent series RLC network and causes the ringing to decay more quickly. This will help to prevent conduction of the input or output protection diodes.

Keeping these relatively simple and inexpensive procedures in mind when designing with CMOS ICs will prevent commonly encountered problems due to latch-up, resulting in trouble-free operation. The old proverb that says, "An ounce of prevention is worth a pound of cure," is very appropriate here.

ANALOG DEVICES

AN-210
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Adding Additional Input Channels to the AD7773/AD7775

by John Wynne

The AD7773 and AD7775 accept differential input signals through the differential input pins, $V_{IN}(+)$ and $V_{IN}(-)$. However, a special mode exists which disconnects the output of the rectifier from the $C_{INT}(+)$ pin and allows single-ended input signals to be applied via the $C_{INT}(+)$ pin. Locations CR7, CR8 and CR9 of the 10-bit wide Control Register are decoded to provide a number of different functions and modes of operation within the AD7773 and AD7775. CR7 determines whether a signal will be acquired via the synchronous detector's differential inputs or direct from the $C_{INT}(+)$ pin. $\overline{CR7}$ is ANDed with the internally generated integrate signal INT to make or break the signal path from the rectifier output to the $C_{INT}(+)$ pin. With CR7 low the rectifier output drives the external integrating capacitor on the C_{INT} pins, and all input signals are acquired through the $V_{IN}(+)$ and $V_{IN}(-)$ differential input pins. With CR7 high the synchronous detector stage is bypassed, and all input signals are now acquired through the single-ended $C_{INT}(+)$ pin. Table I shows the different options/modes available and their respective addresses.

Table I. Functional Address Decoding for the AD7773 and AD7775

CR9	CR8	CR7	Function
0	0	X	Soft Reset
0	1	X	Power Down
1	0	0	Not Allowed
1	0	1	Calibration Mode
1	1	0	Normal Mode
1	1	1	Bypass Mode

X = Don't Care

Bypass Mode: In this mode the synchronous detector, rectifier and integrator are bypassed and the single-ended input signals to be captured are now applied via the $C_{INT}(+)$ pin, the C_{INT} capacitor having been removed. To select the bypass mode locations, CR7, CR8 and CR9 of the control register are loaded with logic highs. A simplified timing diagram of the channel operating in the bypass mode is shown in Figure 1.

When CTRL goes high at the start of a signal capture, the input begins to be tracked by one of the four T/H amplifiers. On the falling edge of CTRL, the tracking T/H amplifier is put in the hold mode, and the voltage on its hold capacitor remains held for subsequent A/D conversion. Unlike the synchronous detector mode the discharge switch, SW3, remains open between successive CTRL signals. When the number of bypass signals captured equals the preprogrammed number expected, the capture sequence ends and the held voltages are sequentially applied to the ADC and converted. From here on, channel operation is identical to the Synchronous Detector mode as described in the AD7773/AD7775 data sheet. Note that locations SR1–SR4 of the Status Register now convey no meaningful information and should be ignored for Bypass Detector operation.

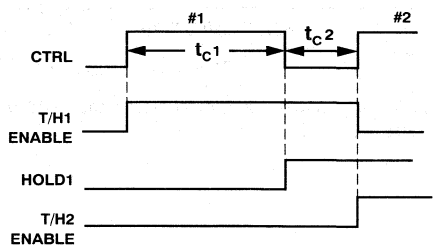

Figure 1. Channel Timing Wafeforms in Bypass Mode

There are a number of considerations which should be followed when changing between modes. The first is that no mode change be attempted before the burst capture and conversion sequence is complete, i.e., not until location SR0 of the Status Register returns low. This will avoid any inadvertent corruption of a conversion in progress. The second consideration involves the delay between writing to the Control Register and starting a new burst capture sequence. This time is defined under the Demodulator Timing Characteristics as the $\overline{WR}$ rising edge to CTRL rising edge and is specified as

200 ns minimum. It is required to ensure that the correct conditions have been set up internal to the device.

A final consideration involves allowing sufficient time for the integrating capacitor, C_{INT}, to discharge when changing from the Bypass mode to one of the other operating modes. This is necessary since C_{INT} is not discharged by the internal discharge switch, SW3, either between successive CTRL pulses or even on completion of the burst capture sequence. A discharge time of 300 ns—equivalent to t_C2, the CTRL Low time in Figure 1—is adequate after transferring out of the Bypass mode. This discharge time and the previous set up time of 200 ns must be added together to arrive at a final overall delay of 500 ns.

Adding Extra Channels

Figure 2 shows a recommended circuit for adding one extra input channel to the AD7773 and AD7775. Figure 3 shows the recommended circuit for multiple input channels. Op amps A1 and A2 can be an OP-292 from Analog Devices, a dual op amp designed for single +5 V operation. Op amp A3 is a transconductance amplifier such as the CA3080 from RCA. This amplifier must be gated OFF via its bias input pin when the Bypass mode is not selected. The analog input voltage signal(s) to both circuits is expected to range from 0 V to 1.2 V.

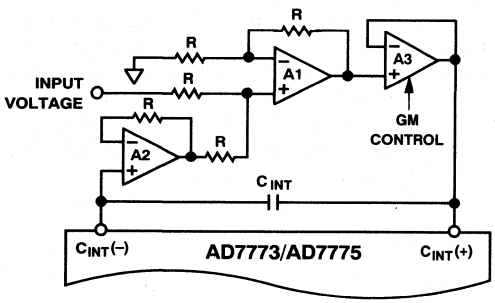

Figure 2. Adding One Additional Input Channel to the AD7773 and AD7775

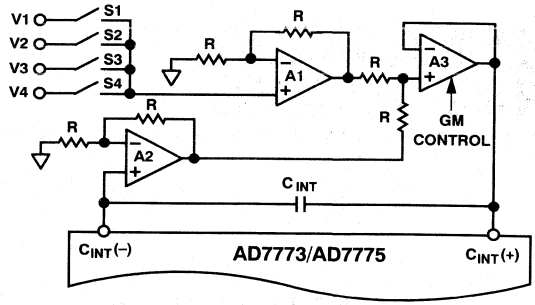

Figure 3. Adding Four Additional Input Channels to the AD7773 and AD7775

ANALOG DEVICES

AB-6
APPLICATION BRIEF

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Single Supply Operation of the DAC-08 and DAC-20

FEATURES

- Simple Interface
- Compatible with CMOS and Open Collector TTL
- No Degradation in Performance

GENERAL DESCRIPTION

The DAC-08 may be operated from a single supply when properly biased. This circuit will allow the use of a single power supply, or battery, and still realize the premium performance of these high speed DACs.

The resistive voltage divider inputs to V_{LC} and logic inputs provide the necessary voltage levels for operation from CMOS and Open Collector TTL logic.

8

CIRCUIT

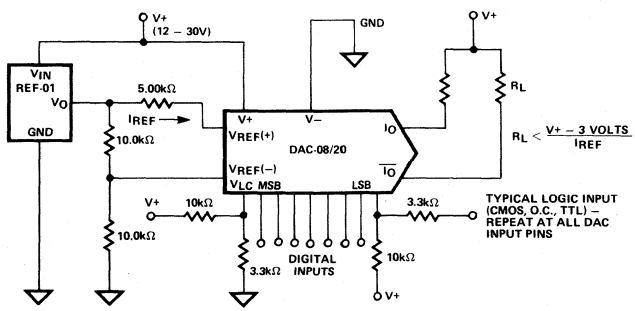

ANALOG DEVICES

AN-225
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

12-Bit Voltage-Output DACs for Single Supply 5 V & 12 V Systems

by John Wynne

The basic CMOS Multiplying DAC, pioneered by Analog Devices, is a voltage-in, current-out device. Intended as a two-quadrant multiplier, the output signal current, which is proportional to the product of the digital input word and the input voltage reference, is steered into a virtual earth summing junction of an external op amp, the loop being closed around the op amp by the on-chip feedback resistor RFB. This circuit inherently provides signal inversion, a positive input signal is translated into a negative output voltage. The AD75XX family of CMOS multiplying DACs all behave like this.

In single supply applications CMOS DACs are operated in what is known as the voltage-switching mode. Figure 1 shows an AD7545A, a 12-bit CMOS DAC, configured for such operation. Reference voltage V_1 is applied to the OUT1 terminal, reference voltage V_2 is applied to the AGND terminal, and the output voltage is available at the V_{REF} terminal. No signal inversion occurs from input to output hence the attraction of the circuit for single-supply applications. However, some linearity degradation should be expected when using the AD75XX DACs in the voltage-switching mode. This is due to the imbalance of the drive voltages on the N-channel switches in the R-2R ladder. With increasing voltage levels on V_1 & V_2, Integral Linearity or Relative Accuracy degrades much more rapidly than Differential Nonlinearity. In a typical HDD servo loop it is more important for the DAC driving the voice coil motor to be monotonic rather than be accurate to 12-bits. Appendix 1 shows some typical linearity performance curves for the AD7243 and AD7545A for various V1, V2 combinations.

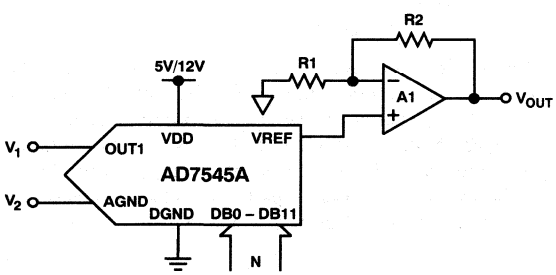

Figure 1. Voltage-Switching Operation of the AD7454A with Buffer Amplifier

Another family of CMOS DACs available from Analog Devices, the AD72XX family, produce a positive output voltage from a positive input voltage. Normal operation of these DACs is in fact the voltage-switching mode. The AD7237, a dual 12-bit DAC with 8+4 loading, and the AD7243, a serial loading 12-bit single DAC, are both suitable for use in the circuits presented here. Note that the AD7237 and AD7243 are not suitable for 5 V operation. In order to allow easy comparisons to be made between the various circuits in this application note the same V_1, V_2 nomenclature is used for all DACs. Thus for the AD7545A, V_1 is applied to OUT1, while for the AD7243 it is applied to REFIN.

8

Voltage-Switching Mode Operation

There are several points worth noting about voltage-switching mode operation:

1. For the AD75XX DACs the output signal is a voltage at a constant output impedance equal to the ladder impedance. The circuit has no significant gain error as long as the current load at the V_{REF} output is small, i.e., an output buffer amplifier is needed.

   For the AD7237 and AD7243 DACs the voltage output is already buffered by an on-board amplifier obviating the need for an additional external amplifier.

2. For the AD75XX DACs the reference voltage inputs V_1 and V_2 do not see constant input impedances but see input impedances which change with digital input code. Hence V_1 and V_2 must be low impedance sources. A somewhat similar situation exists with the AD7237 and AD7243 DACs, although with these devices the V_1 input is internally buffered by an op amp allowing a high source impedance for V_1. However V_2 must still be a low impedance source.

3. V_1 must always be more positive than V_2. If V_1 goes more than 0.3 V negative with respect to V_2, internal diodes and/or parasitic transistors will turn on leading to heavy current flow and possible device destruction. System power-on and power-off are the most obvious occasions when this reversal can take place, and these events should be carefully characterized.

4. For the AD75XX DACs the feedback resistor RFB is not generally used in the voltage-switching mode. For the circuits where it is not required tie it to V_1.

5. The output voltage expressions for Figures 2 to 4 assume that the DAC ladder termination resistor is at the same potential as V_2. This assumption is valid for the AD7545A and AD7548, single 12-bit DACs, and AD7537, AD7547 & AD7549, dual 12-bit DACs.

6. In the output voltage expressions for Figures 2 to 5, D represents the fractional equivalent of the DAC digital input code in decimal. For a 12-bit DAC, D can vary from 0 to 4095/4096, i.e., D can almost, but never quite, equal unity. However, in order to allow easy comparisons between the different circuits the output voltage expression for each circuit is evaluated with D = 0, D = 1/2 and D = 1. It should be remembered that the output voltage expressions evaluated with D = 1 are in fact 1 LSB higher than is theoretically possible.

Figure 2 shows the classic voltage-switching mode circuit with a CMOS multiplying DAC and two reference voltage sources V_1 and V_2. When V_2 = 0 V the output voltage ranges from 0 V when D = 0 to $(1+R2/R1)V_1$ when D = 1. An obvious problem with having 0 V as one end of the output signal range in a single supply system is that, although the DAC output will get to 0 V, the single-supply op amp output almost certainly will not, especially if it is required to sink any appreciable load current. By applying a second reference voltage V_2 to the AGND pin the all 0s code output voltage can be raised above 0 V. For 5 V systems V_2 must be kept quite small (≤0.5 V) since it reduces the drive on the DAC switches directly ($V_{DD} - V_2$) and generally increases the onset of linearity errors for a fixed value of V_1. For 12 V systems V_2 can be increased to 3 or 4 V, although for Figure 2 this is unlikely to be the case since V_2 is multiplied by the gain factor (1+R2/R1) to produce the all 0s end of the output signal range. This limits the usefulness of the circuit to applications where either V_2 is small or the required output signal range is small, i.e., a small gain factor. Note also that two low-output impedance voltage sources are required if the DACs are from the AD75XX family. With the AD7237 and AD7243 DACs, V_1 is buffered by on-chip buffer amplifiers obviating the low source impedance requirement on V_1.

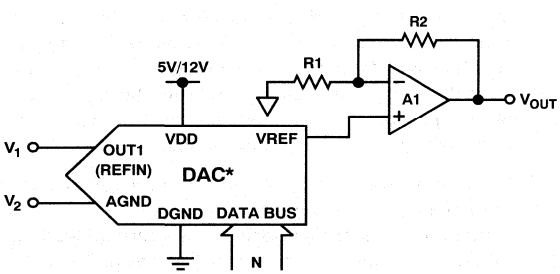

*AD7545A, AD7237 OR AD7243
ADDITIONAL CIRCUITRY OMITTED FOR CLARITY

Figure 2. Classic Voltage-Switching Mode Operation

For the circuit of Figure 2 the output voltage is given as:

$V_{OUT} = V_2(1+R2/R1) + (V_1-V_2)(1+R2/R1)D$
where 1 LSB = $(V_1 - V_2)(1+R2/R1)(1/4096)$.

When D = 0,
$V_{OUT} = V_2(1+R2/R1)$.

When D = 1/2,
$V_{OUT} = V_2(1+R2/R1)+(V_1-V_2)(1+R2/R1)(1/2)$.

When D = 1,
$V_{OUT} = V_2(1+R2/R1)+(V_1-V_2)(1+R2/R1)$
$= V_1(1+R2/R1)$

Supplying two reference voltages, V_1 and V_2, to the circuit of Figure 2 is not the only method of moving the all 0s code output voltage above 0 V. Relevant only to the AD75XX family, Figure 3 shows the feedback resistor RFB, found on these DACs, wired to its own V_{REF} output. This has the effect of moving the all 0s code output voltage to $(V_1/2)$ (1+R2/R1). Note that this is achieved without the need for a second reference voltage V_2. An interesting feature of this circuit is its voltage doubling nature—the all 1s code output voltage is double the all 0s code output voltage. Figure 3 is suitable for both 5 & 12 V systems. For 12 V systems, V_1 should be no more than 1.25 V above AGND in order to maintain 12-bit relative accuracy. For 5 V systems it is recommended that V_1 be no more than 0.25 V above AGND.

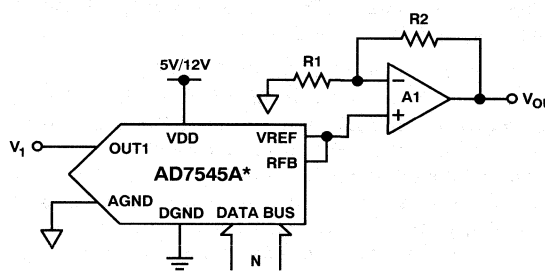

*ADDITIONAL CIRCUITRY OMITTED FOR CLARITY

Figure 3. Using the Feedback Resistor R_{FB} to Implement a Voltage Doubling Function

For the circuit of Figure 3 the output voltage is given by:
$V_{OUT} = (V_1/2)(1+R2/R1)+(V_1/2)(1+R2/R1)D$
where 1 LSB = $(V_1/2)(1+R2/R1)(1/4096)$.

When D = 0,
$V_{OUT} = (V_1/2)(1+R2/R1)$.

When D = 1/2,
$V_{OUT} = (V_1/2)(1+R2/R1)+(V_1/2)(1+R2/R1)(1/2)$.

When D = 1,
$V_{OUT} = V_1(1+R2/R1)$.

Figure 4 is a variation on Figure 2 and provides an all 0s code output voltage equal to V_2 and not $V_2(1+R2/R1)$ which Figure 2 provides. Hence, by appropriate choice of the gain setting resistors, R1 & R2, this circuit is suitable for producing wide output voltage ranges. The circuit requires two low output impedance voltage sources and is suitable for both 5 and 12 V systems.

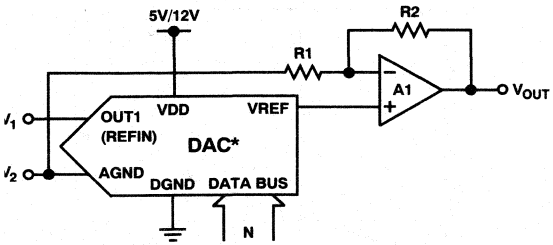

*AD7545A, AD7237 OR AD7243
ADDITIONAL CIRCUITRY OMITTED FOR CLARITY

Figure 4. Avoiding Any Gain Factor Multiplication of V_2

For the circuit of Figure 4 the output voltage is given by:

$V_{OUT} = V_2 + (V_1 - V_2)(1+R2/R1)D$

where 1 LSB = $(V_1 - V_2)(1+R2/R1)(1/4096)$.

When D = 0,

$V_{OUT} = V_2$.

When D = 1/2,

$V_{OUT} = V_2 + (V_1 - V_2)(1 + R2/R1)(1/2)$.

When D = 1,

$V_{OUT} = V_2 + (V_1 - V_2)(1 + R2/R1)$.

Figure 5 is a particular implementation of Figure 4 using an AD7243. This is a serial input, voltage-output 12-bit DAC which includes the output buffer amplifier A1 and matched, equal value, gain setting resistors. It is specified to operate from 12 V to 15 V V_{DD}. For 12-bit performance V_1 should be between 1.25 V and 3 V above V_2. With the equal value gain resistors giving a gain factor of 2, the output voltage range of this circuit is $2(V_1 - V_2)$ with one end of this voltage range being anchored at V_2 when D = 0. Input voltage V_1, which is applied to the REFIN input of the AD7243, is buffered by an on-chip buffer amplifier. This means that V_1 can be generated by a resistive divider from some higher voltage and does not need to be a low impedance source. Dependent upon the value of load resistance to ground, the output op amp in the AD7243 can experience headroom problems if the output voltage is required to swing within several volts of the V_{DD} line. With V_{DD} = 12 V ± 10%and a load of 2 kΩ to ground headroom problems, causing signal limiting, can occur above an 8 V output signal. Increasing the value of load resistance defers the onset of headroom problems.

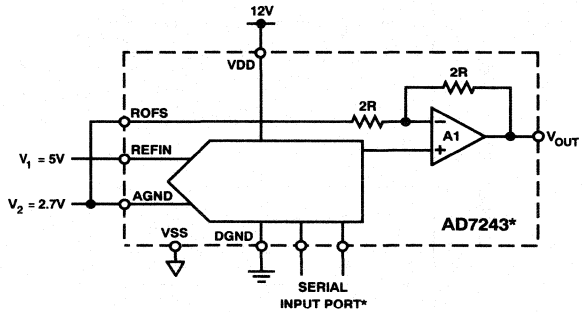

*ADDITIONAL CIRCUITRY OMITTED FOR CLARITY

Figure 5. Voltage-Switching Operation with the AD7243

For the circuit of Figure 5 the output voltage is given by:

$V_{OUT} = 2.7 + (2.3)2D$

where 1 LSB = 4.6/4096 = 1.1 mV.

When D = 0,

V_{OUT} = 2.7 V.

When D = 1/2,

V_{OUT} = 5 V.

When D = 1,

V_{OUT} = 7.3 V.

Another feature of the circuit (in common with Figure 4) is that the output voltage at the half-scale code, i.e., D = 1/2, is equal to V_1 and will follow V_1 as V_1 changes. Thus V_1 behaves as a quasi-ground and can be used by the signal conditioning circuitry to refer the analog output signal to a point other than 0 V. This is more obvious if the output voltage expression for Figure 5 is rewritten as shown

$V_{OUT} = V_1 + (V_1 - V_2).(2D-1)$

$= 5\text{ V} + (2.3).(2D-1)$.

Appendix 1

Figures A1–A4 show typical linearity plots of the AD7243 while Figures A5–A9 show linearity plots for the AD7545A.

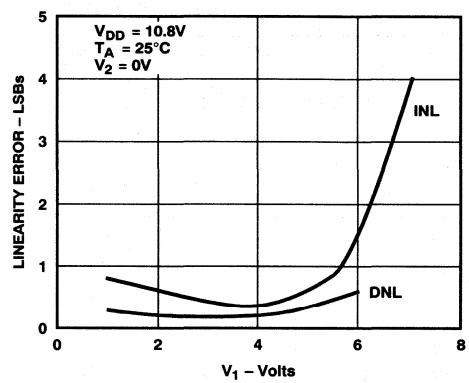

Figure A1. AD7243 Linearity Error, V_2 = 0 V

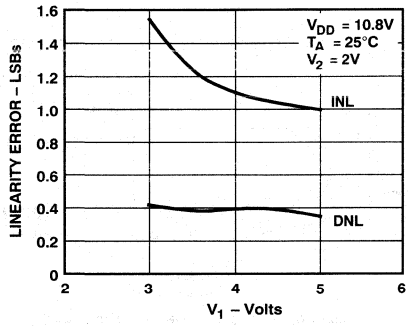

Figure A2. AD7243 Linearity Error, V_2 = 2 V

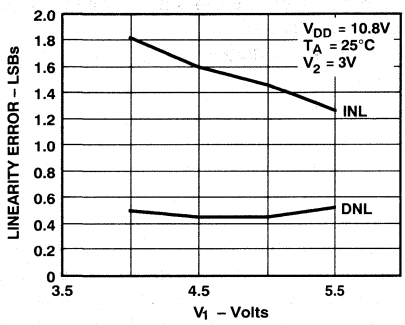

Figure A3. AD7243 Linearity Error, V_2 = 3 V

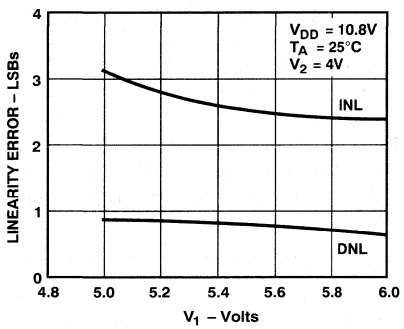

Figure A4. AD7243 Linearity Error, V_2 = 4 V

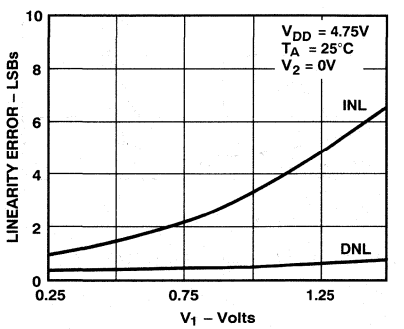

Figure A5. AD7545A Linearity Error, V_{DD} = 4.75 V, V_2 = 0 V

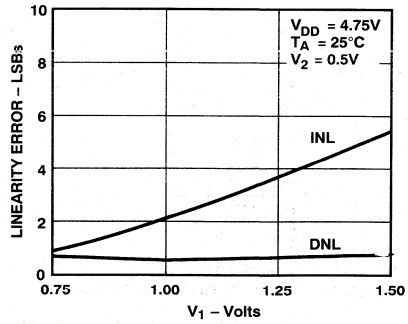

Figure A6. AD7545A Linearity Error, V_{DD} = 4.75 V, V_2 = 0.5 V

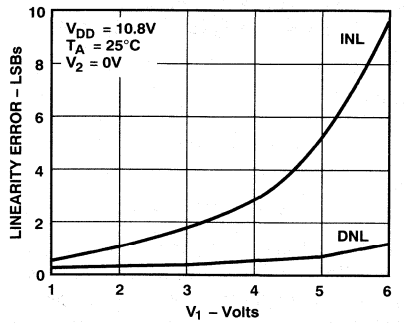

Figure A7. AD7545A Linearity Error, V_{DD} = 10.8 V, V_2 = 0 V

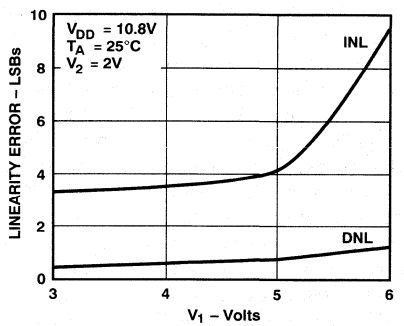

Figure A8. AD7545A Linearity Error, V_{DD} = 10.8 V, V_2 = 2 V

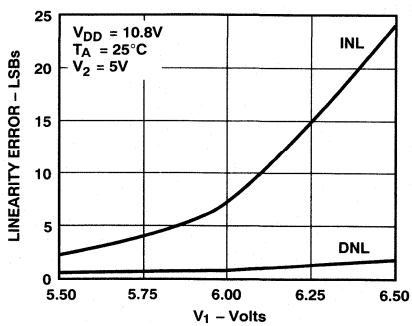

Figure A9. AD7545A Linearity Error, V_{DD} = 10.8 V, V_2 = 5 V

ANALOG DEVICES

AN-310
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Single Supply 12-Bit DAC Circuits Using the AD7568

by John Wynne

This application note suggests a number of different circuit configurations in which a CMOS D/A converter can be used to generate a variety of popular output voltage ranges. All of the circuits work off of positive supplies only and feature the AD7568, an octal 12-bit current-out CMOS DAC. Figures 1–4 show the DACs being operated in their current-steering mode, while Figures 5–7 show operation in the voltage-switching mode. A previous application note, AN-225, covered the issues involved in using current-out DACs in the voltage-switching mode, and the reader is referred to this for additional information.

In the output voltage expressions for Figures 1 to 7, the term D represents the fractional equivalent of the DAC digital input code in decimal, N. For a 12-bit DAC, D = N/4095 and can vary from 0 to 4095/4096, i.e., D can almost, but never quite, equal unity. However, in order to allow easy comparisons between the different circuits, D is assumed to range from 0 to 1. It should be remembered that the output voltage expressions evaluated with D = 1 are in fact 1 LSB higher than is theoretically possible.

Although the AD7568 is a +5 V only device and is specified in the normal current-steering mode with the I_{OUT1} and I_{OUT2} outputs at 0 V, the DACs still work well with their current outputs biased above 0 V. For instance with I_{OUT1} and I_{OUT2} raised above ground by up to 1.23 V, a typical bandgap value, the DACs remain monotonic to 12 bits.

In the voltage switching mode configuration I_{OUT1} can be raised above I_{OUT2} by up to 1.23 V and 12-bit DAC monotonicity is still assured. This performance is maintained even if the I_{OUT2} output is itself at 1.23 V above ground.

Appendix 1 shows some typical linearity performance curves for the AD7568 under various biasing conditions.

Current-Steering Mode

Figure 1 shows the simplest circuit in this mode. The DAC reference input, V_{REF}, is shown tied to 0 V, but it can be biased to any positive voltage up to twice the value of V1 to change the full scale output voltage.

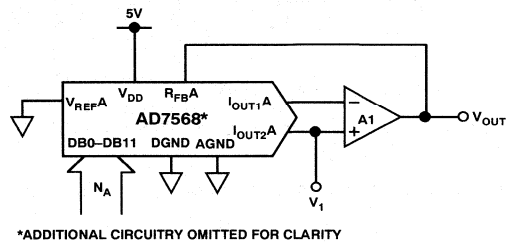

Figure 1. Simple Current Steering Mode Circuit

For the circuit of Figure 1, the output voltage is given by:

$$V_{OUT} = V_1 (1 + D)$$

For example, if V_1 = 1.23 V, the output voltage goes from 1.23 V at D = 0 to 2.46 V at D = 1. This circuit is very limited in its output voltage range, and consequently gain is generally added around the op amp to extend the performance. A circuit to multiply the output voltage range by some gain factor is shown in Figure 2.

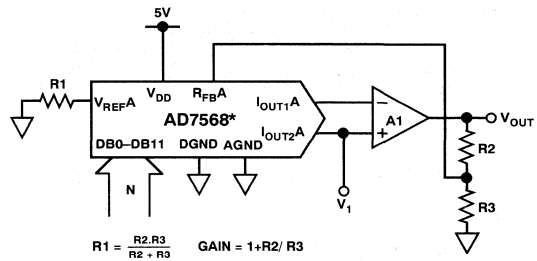

Figure 2. Adding Gain to Circuit of Figure 1

Using three external resistors is the recommended way (Reference 1) of adding gain to the standard configuration without requiring a large gain adjustment range or compromising the circuit's temperature coefficient. The resistors R1, R2 and R3 should all have similar temperature coefficients, but they need not match the temperature coefficient of the DAC.

For the circuit of Figure 2, the output voltage is given by:

$$V_{OUT} = V_1(1 + D)(1 + R2/R3)$$

Similar to the circuit of Figure 1, the interesting feature of this circuit is its voltage doubling nature regardless of the gain factor—the all 1s code output voltage is double the all 0s code output voltage. For instance, with V_1 = 1.23 V and R2 = R3 (Gain Factor = 2), the output voltage varies from 2.46 V at D = 0 to 4.92 V at D = 1.

However simple voltage doubling may not be to everyone's requirement. If R3 in Figure 2 is returned to V_1 instead of 0 V, then the all 0s code output voltage avoids being multiplied by the added gain factor and is simply V_1. Figure 3 shows this configuration.

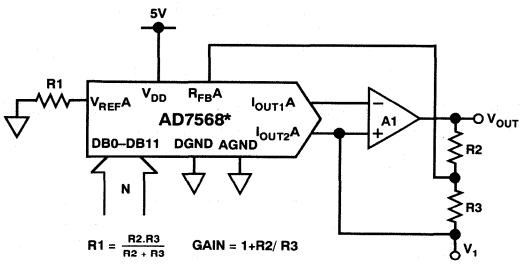

Figure 3. Avoiding Gain Factor Multiplication of All 0s Code Output Voltage

For the circuit of Figure 3, the output voltage is given by:

$$V_{OUT} = V_1 + V_1 \cdot D\,(1 + R2/R3)$$

In this instance, with V_1 = 1.23 V and R2 = 2R3 (Gain Factor = 3), the output voltage range is from 1.23 V at the all 0s code to 4.92 V at the all 1s code.

The facility to apply a second voltage to the V_{REF} input of the DAC allows the output voltage range to be closely tailored to requirements. As an example, Figure 4 generates an output voltage of 2.5 V to 7.5 V, i.e., 5 V ± 2.5 V.

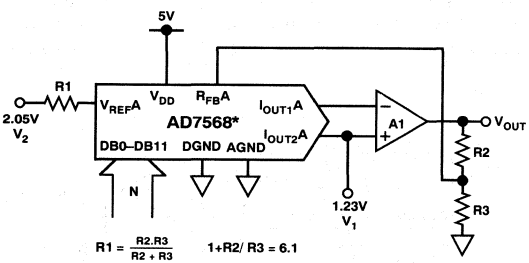

Figure 4. Circuit to Generate 5 V ± 2.5 V Output

For the circuit of Figure 4, the output voltage is given by:

$$V_{OUT} = V_1\,(1 + R2/R3) + (V_1 - V_2)\,D(1 + R2/R3)$$

The gain factor is chosen to produce the high-side 7.5 V output with V_1 again equal to 1.23 V. This value is then used in the above expression to solve for V_2 when the output voltage is at its low-side value of 2.5 V. A feature of this particular circuit is that V_2 is greater than V_1, and hence the output voltage vs. code transfer function is inverted, i.e., all 0s code produces 7.5 V output; all 1s code produces 2.5 V output.

Voltage-Switching Mode Circuits

Figure 5 is discussed in the previously mentioned AN-225 but is reproduced here to show the suitability of the AD7568 in voltage-switching mode applications for 5 V-only systems.

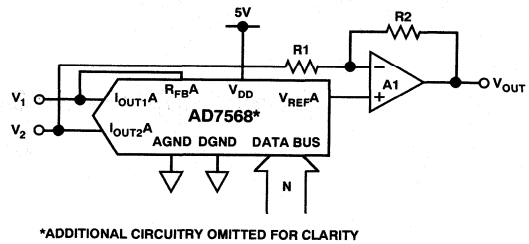

Figure 5. Voltage-Switching Mode Operation

For the circuit of Figure 5, the output voltage is given by:

$$V_{OUT} = V_2 + (V_1 - V_2)\,(1 + R2/R1)\,D$$

If R1 = R2 and the voltage sources are two bandgaps where the V_1 bandgap is with respect to V_2, i.e., V_1 is 2.46 V above ground, the output voltage over the code range is as follows:

When D = 0,

$$V_{OUT} = 1.23\ V$$

When D = 1/2,

$$V_{OUT} = 2.46\ V$$

When D = 1,

$$V_{OUT} = 3.69\ V$$

The output voltage is of the form 2.46 V ± 1.23 V. In this circuit the output voltage swings 1.23 V up and down around an effective bias point of 2.46 V. This type of output is commonly termed a V_{BIAS}/V_{SWING} output and is very popular in the hard disk and optical disk drive industry. Using bandgaps for V_1 and V_2, the minimum output voltage of Figure 5 is 1.23 V. If it is important to have 0 V as the minimum output voltage, this can be easily achieved by returning R1 in Figure 5 not to $I_{OUT2}A$, but to $I_{OUT1}A$. In addition R1 must be made equal to R2. The output voltage expression for this configuration is:

$$V_{OUT} = V_2 + (V_1 - V_2)\,(2D - 1)$$

This is of the form 1.23 V ± 1.23 V, i.e., now when D = 1/2, V_{OUT} = 1.23 V and not 2.46 V.

Figure 5 can also be rearranged to generate a V_{BIAS} level higher than either 1.23 V or 2.46 V by tying the DAC feedback resistor to $V_{REF}A$ instead of $I_{OUT1}A$. The output voltage expression under these conditions is given by;

$$V_{OUT} = V_2 + (1/2)(V_1 - V_2)(1 + R2/R1)(1 + D)$$

If R2 = 3R1, the above expression simplifies to;

$$V_{OUT} = 1.23\ V + 2.46\,(1 + D)$$

This is of the form 4.92 V ± 1.23 V.

This highlights the fact that although the DAC itself is limited to a 5 V power supply, it can control output voltages higher than this.

Programmable V_{BIAS} Level

Some applications require a variable V_{BIAS} level and a constant V_{SWING} level. For the circuit of Figure 6 the output voltage is given by;

$$V_{OUT} = V_2 (1 + D_B)(1 + R2/R1) - V_1(2D_A - 1)(R2/R1)$$

With R1 = R2 the output voltage simplifies to:

$$V_{OUT} = 2V_2 (1 + D_B) - V_1(2D_A - 1)$$

where D_A and D_B are fractional representations of the digital codes to DACs A and B respectively. When $V_1 = V_2 = 1.23$ V, the output voltage swings 1.23 V (determined by D_A) around a bias level which is programmable (via D_B) from 2.46 V to 4.92 V. Note that there is a code inversion effect in this circuit also. For a given V_{BIAS} level the output voltage is at its minimum with all 1s in DAC A and vice versa.

Using bandgaps for V_1 and V_2, the minimum output voltage of Figure 6 is 1.23 V. If 0 V is the minimum output voltage required, this can be easily achieved by tying the DAC B feedback resistor, $R_{FB}B$, to $V_{REF}B$. This has the effect of halving the voltage at the noninverting input of A2. The output voltage now swings ±1.23 V (determined by D_A) around a bias level which is programmable (via D_B) from 1.23 V to 2.46 V.

Also note in Figure 6 that if DAC B is removed and V_2 is applied directly to the noninverting input of A2, the output voltage will be of the form 2.46 V ± 1.23 V. In comparison with Figure 5 which has a similar output voltage signal this configuration requires one more op amp and two matched resistors but only one bandgap reference ($V_1 = V_2$).

Programmable Full-Scale Level

Another common requirement is for a circuit which can have programmable full-scale V_{SWING} levels while retaining the same V_{BIAS} value. Such a circuit is shown in Figure 7.

With R1 = R2 the output voltage for Figure 7 is given by:

$$V_{OUT} = V_2 + D_A (V_1 - V_2) (2D_B - 1)$$

When $V_1 = V_2 = 1.23$ V and the V_1 bandgap is with respect to V_2, i.e., V_1 is 2.46 V above ground, the output voltage swings ±1.23 V maximum around a bias level of 1.23 V. The full-scale V_{SWING} voltage is set by D_A to be between 1.23 V and 0 V theoretically.

8

REFERENCES

1. Brokaw P., "Input Resistor Stabilizes MDACs Gain," *EDN*, January 7, 1981, pp. 210–211.

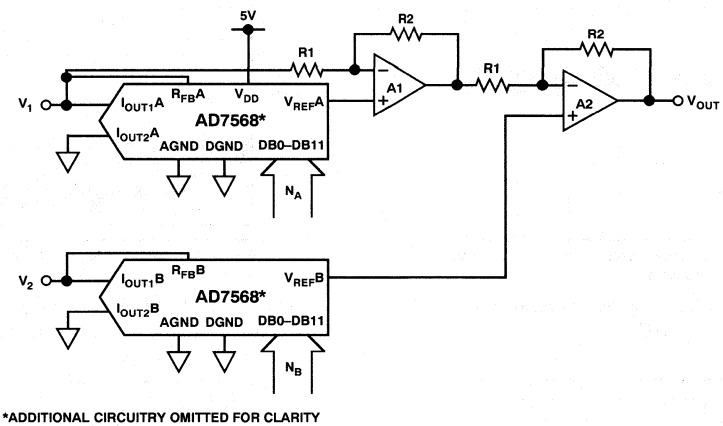

Figure 6. Programmable V_{BIAS} Level

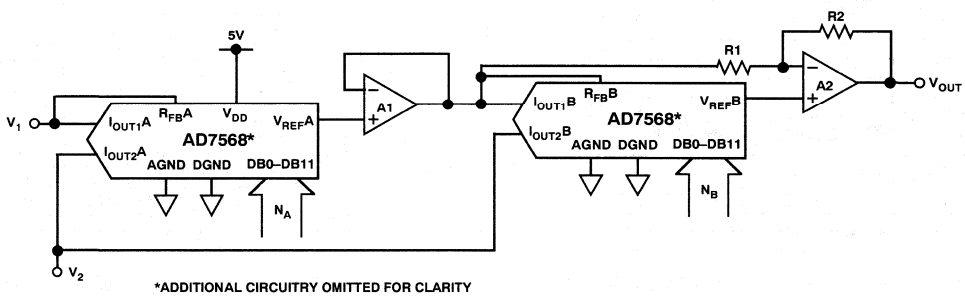

Figure 7. Programmable Full-Scale V_{SWING} Level

Appendix 1

Figures A1–A6 show typical linearity plots of the AD7568 tested in both the Figure 1 and Figure 5 configurations of this application note. The plots of Figures A1–A2 were generated using the Figure 1 circuit and use the same DAC channel within the AD7568 to show temperature effects on the linearity performance. The plots of Figures A3–A4 were generated using the Figure 5 circuit and again the same DAC channel within the AD7568 is used to show temperature effects on the linearity performance. Similarly for Figures A5–A6, the same channel is recorded. For every plot the DAC V_{DD} is +4.75 V and 1 LSB = 300 microvolts. The worst case conditions, in linearity terms, occur at low power supplies and high temperatures.

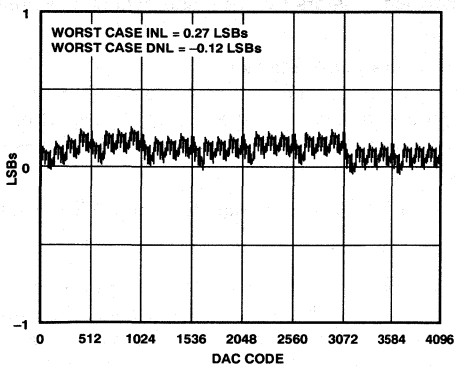

Figure A1. Current-Steering Mode Linearity with V_1 = 1.23 V and T_A = +25°C

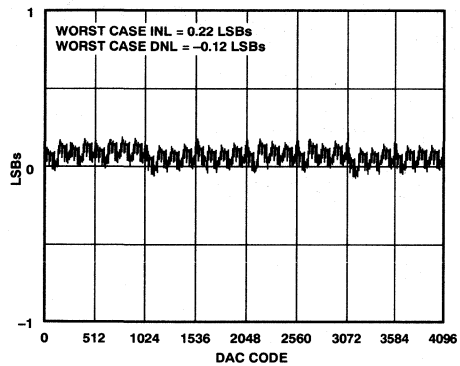

Figure A2. Current-Steering Mode Linearity with V_1 = 1.23 V and T_A = +85°C

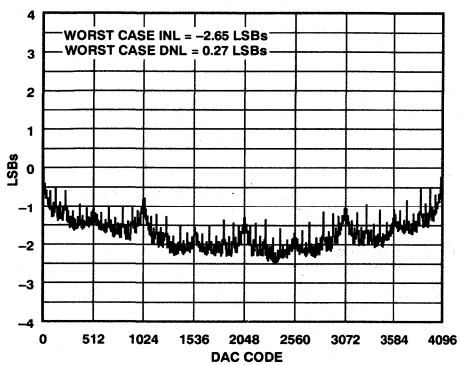

Figure A3. Voltage-Switching Mode Linearity with V_1 = 1.23 V, V_2 = 0 V and T_A = +25°C

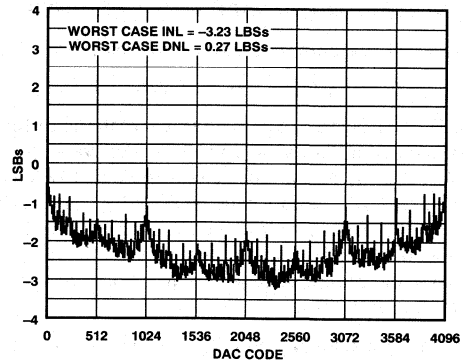

Figure A4. Voltage-Switching Mode Linearity with V_1 = 1.23 V, V_2 = 0 V and T_A = +85°C

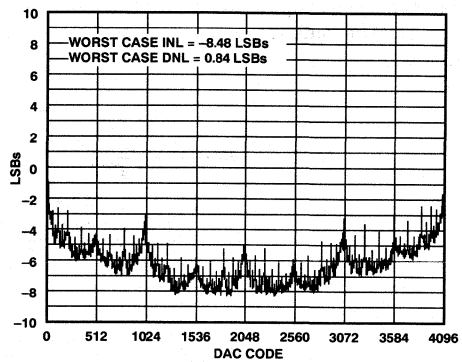

Figure A5. Voltage-Switching Mode Linearity with V_1 = 1.23 V, V_2 = 1.23 V and T_A = +25°C

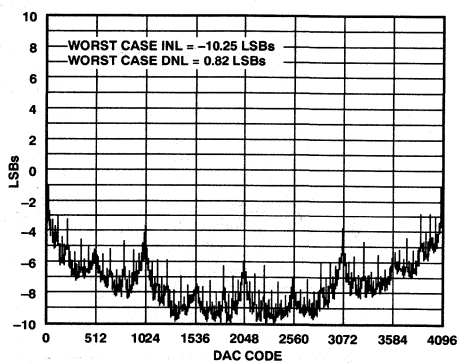

Figure A6. Voltage-Switching Mode Linearity with V_1 = 1.23 V, V_2 = 1.23 V and T_A = +85°C

ANALOG DEVICES

AN-314
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

14-Bit DACs Maintain High Performance Over Extended Temperature Range

by Mike Curtin and Robert Stakelum

INTRODUCTION

As circuit designers look for higher performance d/a converters in instrumentation and control applications, the AD7534/AD7535 14-bit DACs from Analog Devices offer the unique combination of low power consumption and high performance over the full temperature range. What makes this possible is a proprietary low leakage configuration which keeps output leakage current very low even at high temperatures. Consequently, linearity error and gain error are much less susceptible to temperature drift. This application note shows that the devices work very well even when operated above their specified temperature range. Test results are given for devices at 200°C. To show that reliable operation is maintained at this high temperature a 1000 hour life test was carried out at 200°C. Results of this life test and a failure analysis are presented. Though the 200°C results are impressive, Analog Devices does not guarantee performance outside the specified temperature range. However, the results give an indication of the exceptional reliability of the devices in the range up to 125°C and also show the consistency of important specifications up to 125°C and beyond.

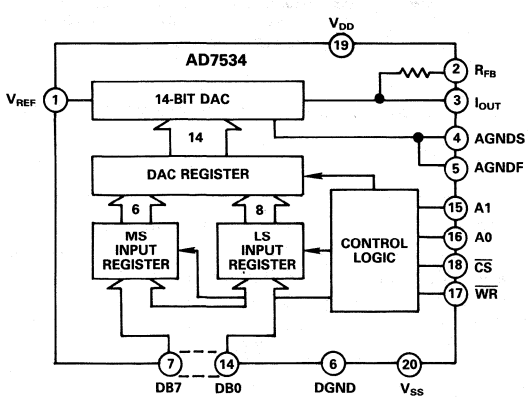

Figure 1. AD7534 Block Diagram

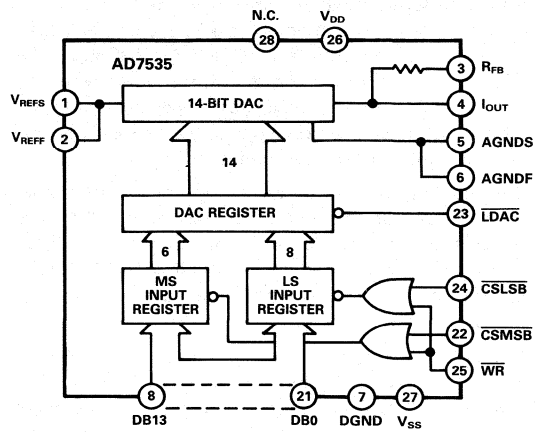

Figure 2. AD7535 Block Diagram

LOW LEAKAGE CONFIGURATION IN AD7534/AD7535

Leakage current in CMOS DACs has two components:

1. Leakage from the V_{DD} supply to the I_{OUT} terminal. This is independent of input code.
2. Leakage from the R-2R ladder through off-switches. This is a maximum for all zeros at the input and is a minimum with all ones at the input, i.e., all switches on.

To eliminate this leakage current flowing into the I_{OUT} line, the AD7534/AD7535 uses a novel configuration (patent pending). By holding the V_{SS} pin at $-0.3V$ through a simple resistor divider (see Figure 3), the leakage current is virtually eliminated. This means that any leakage current effects on linearity and gain error (usually significant above 100°C) are also eliminated.

If V_{SS} is held at 0V instead of $-0.3V$, the device will exhibit all the normal specifications degradation due to leakage. Figure 4 is a typical linearity plot of a device at 125°C, showing the effect of V_{SS} on end-point linearity.

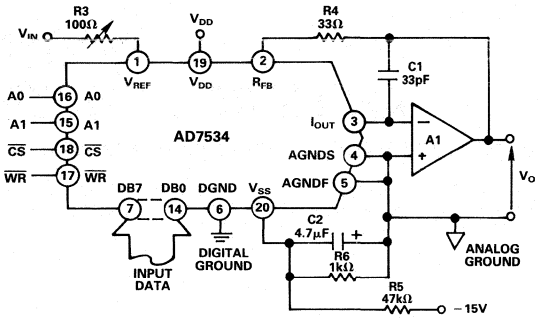

Figure 3. Unipolar Binary Operation of the AD7534

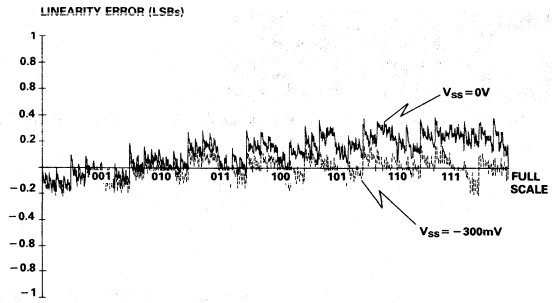

Figure 4. Effect of V_{SS} on Ad7534/AD7535 Linearity at 125°C

OPERATION AT VERY HIGH TEMPERATURES

With its excellent performance up to 125°C, and its special low leakage configuration, the AD7534/AD7535 should be considered for very high temperature operation (i.e., in the 200°C region). Before looking at the performance level at such a high temperature, it has to be realized that several factors will influence the device reliability.

1. Chemical interaction on the surface of the die can be a hazard at high temperatures. This can result in wire bond strength degradation. To ensure that this degradation did not occur with the AD7534/AD7535, a wire bond pull strength test was conducted after the high temperature life testing.
2. The aging rate of a device approximately doubles for every 10°C rise in T_J (junction temperature). The fact that the AD7534/AD7535 DACs are in CMOS is a major advantage in this respect. Table I is a comparison of the AD7534/AD7535 with a typical bipolar DAC, which has a much higher power consumption. At 200°C, T_J for the bipolar DAC exceeds that for the AD7534/AD7535 by 20°C. Thus, its life expectancy is considerably less.

In order to establish a good confidence level for 200°C operation, 7 devices were tested at the elevated temperature. Figures 5 to 8 are a record of these results.

The graph of Figure 5 shows how leakage current increases dramatically above 100°C with V_{SS} = 0. When V_{SS} = −300mV, there is a significant improvement, with leakage at 200°C limited to 10nA. The effects of leakage on

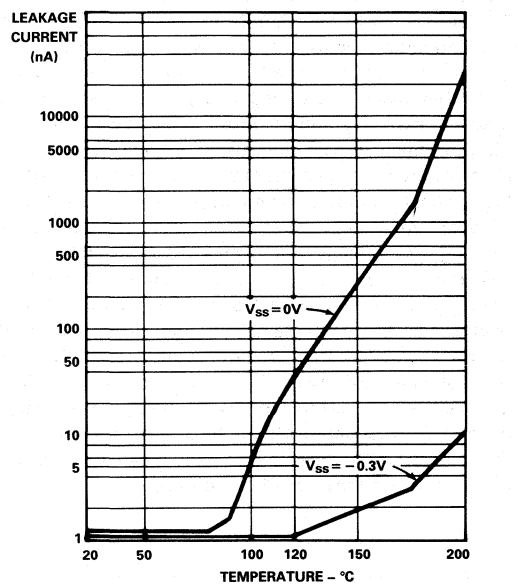

Figure 5. Leakage Current vs. Temperature for AD7534/AD7535

DEVICE	AD7534	AD7535	TYPICAL BIPOLAR DAC
Package	20-Pin Side Braze	28-Pin Side Braze	24-Pin Side Braze
θ_{ja}	80°C/W	60°C/W	65°C/W
P (Typical Power Consumption)	47.5mW	47.5mW	350mW
T_J ($T_A + \theta_{ja} \cdot P$)	203.8°C	202.8°C	222.7°C

Table I. Comparison of Junction Temperatures for T_A (Ambient Temperature) of 200°C

linearity error and gain error are reflected in Figures 6 and 7. Without the low leakage configuration ($V_{SS}=0$) the magnitude of the linearity and gain errors make the parts unusable. By using the low leakage facility, typical linearity error is less than 1LSB and typical gain error is less than 2LSBs.

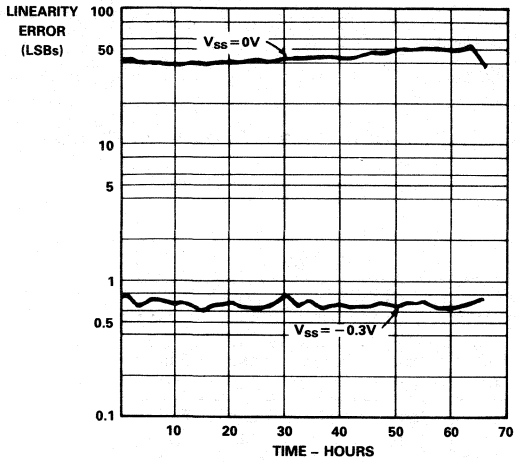

Figure 6. Linearity at 200°C

The stability of the devices over 65 hours of continuous operation is impressive. Linearity drift during this time is ±0.1LSBs while gain error drift is ±0.3LSBs. Linearity stability with temperature is shown in Figure 8. This is a typical all-codes linearity plot for the AD7534/AD7535 at two temperatures; 25°C and 200°C. Linearity error is well within 1LSB over the extended temperature range.

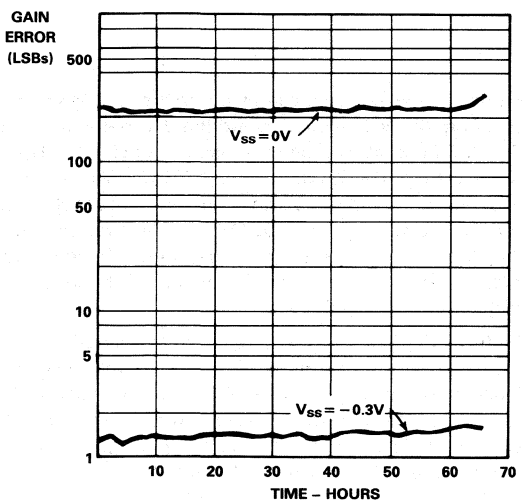

Figure 7. Gain Error at 200°C

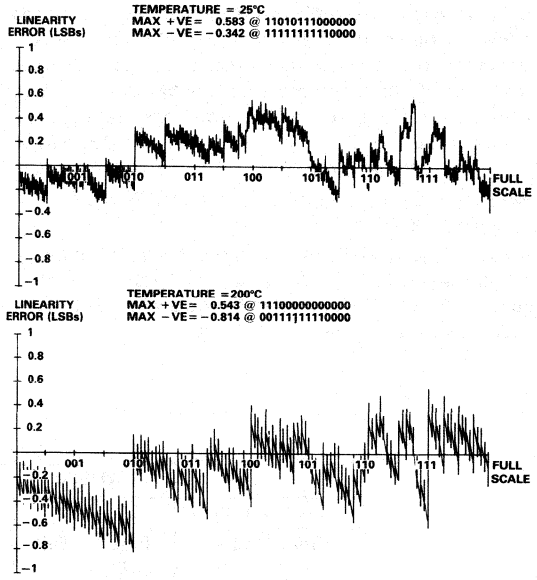

Figure 8. All-Codes Linearity Plot for the Same Unit (AD7534) at 25°C (Top) and 200°C (Bottom)

8

LIFE TEST

The initial operation and testing of seven devices at 200°C for 65 hours gives a good indication of device functional performance. However, at this high temperature, the devices age much faster and may exhibit some temperature-related failure. To examine the reliability of the AD7534/

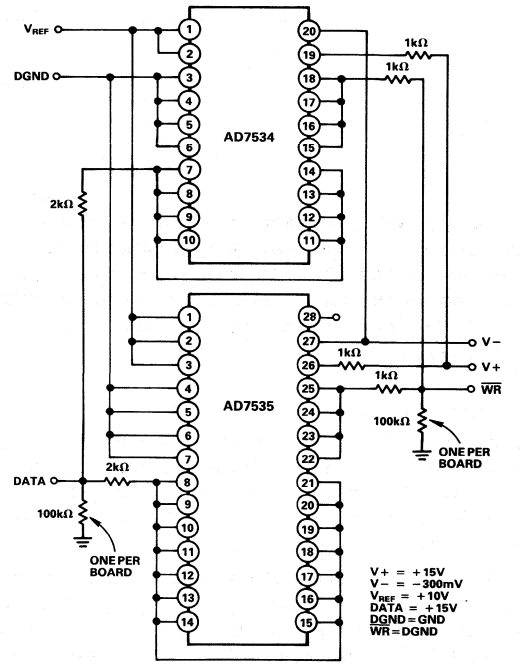

Figure 9. AD7534/35 200°C Burn-In Diagram

AD7535 at 200°C, a sample of the devices was subjected to a 200°C, 1000-hour life test. This sample contained devices taken from 2 separate production lots for each model. The life test required a special high temperature burn-in board, made from a glass polymide material. The resistors, sockets and solder used on the board had to be capable of withstanding 200°C. The life test was conducted with the devices powered up and voltage levels applied as in Figure 9.

There was a parallel life test carried out at 125°C to correlate failures. Parts from the same production lots were used in both tests. If a lot-related failure occurred at 200°C, it could be verified in the 125°C life test.

device being misapplied during testing, causing the bond wire at the AGNDS pin to be blown open and rendering the part unusable. The other device was a linearity failure.

Time (Hrs)	0	72	333	500	732	1000
Linearity Err. (LSBs)	1.82	1.82	1.87	1.87	1.89	2.04
Specification Limit	2LSBs					

Table III. AD7534 Which Failed after 1000 Hours

	T = 125°C				T = 200°C			
	AD7534		AD7535		AD7534		AD7535	
Initial Sample	30		29		28		24	
After 72 hrs	Pass 30	Fail 0	Pass 29	Fail 0	Pass 28	Fail 0	Pass 24	Fail 0
After 142 hrs	Pass 30	Fail 0	Pass 29	Fail 0	Pass 28	Fail 0	Pass 24	Fail 0
After 333 hrs		–		–	Pass 28	Fail 0	Pass 22	Fail 2
After 500 hrs	Pass 30	Fail 0	Pass 28	Fail 1	Pass 28	Fail 0	Pass 22	Fail 0
After 732 hrs	Pass 30	Fail 0	Pass 28	Fail 0	Pass 28	Fail 0	Pass 21	Fail 1
After 1000 hrs	Pass 30	Fail 0	Pass 28	Fail 0	Pass 27	Fail 1	Pass 21	Fail 0

Table II. Life Test Results

Table II shows the complete life test results. From an initial sample of 28, one AD7534 failed between 732 and 1000 hours at 200°C. The AD7535 failure rate was higher with three failures from an initial sample of 24. At 125°C, the AD7535 failure rate was one as opposed to none for the AD7534's.

FAILURE ANALYSIS

Table II shows that there were no AD7534 failures at 125°C. There was one AD7535 failure after 500 hours, which was due to an isolated masking defect on an aluminium line contacting to a current steering switch in the DAC ladder.

At 200°C, there was one AD7534 life test failure. This was caused by marginal drift outside its linearity specification. Table III charts this drift through life test. The total change of 0.22LSB over the 1000-hour life test is not significant. There were 3 AD7535 failures at 200°C. Two occurred after 333 hours of life test. One of these was as a result of the

Probing revealed that the fault was in the DAC MSB latch cell. Both the MSB switches and the decoder were functional. Further failure analysis was inconclusive. There was a further AD7535 failure after 730 hours. This was due to a gate oxide rupture on an N-channel current steering switch. The nature of the 200°C life test failures suggests that there was no special temperature related failure mechanism. They are random failures of the kind normally seen in 125°C life tests.

Of further interest is the way in which certain critical specifications drift through the life test. Table IV shows four of these specifications (I_{DD} at logic level inputs of 0.8V and 2.4V, Linearity Error and Gain Error). The average, $\bar{X}$, and standard deviation, σ, of each sample is given at three stages of life test (0, 500 and 1,000 hours).

Not only that, but it is also possible to compare at a glance the drift at 125°C and 200°C. The figures show that the device specifications examined are very consistent through both the 125°C and 200°C tests.

		AD7534												AD7535											
		I_{DD} (0.8V) (mA)			I_{DD} (2.4V) (mA)			LIN ERR (LSBs)			GAIN ERR (LSBs)			I_{DD} (0.8V) (mA)			I_{DD} (2.4V) (mA)			LIN ERR (LSBs)			GAIN ERR (LSBs)		
		$\bar{X}$	σ	SPEC.	$\bar{X}$	σ	SPEC.	$\bar{X}$	σ	SPEC.	$\bar{X}$	σ	SPEC.	$\bar{X}$	σ	SPEC.	$\bar{X}$	σ	SPEC.	$\bar{X}$	σ	SPEC.	$\bar{X}$	σ	SPEC.
125°C	0 HRS	0.21	0.071	3	0.81	0.241	3	0.10	0.68	2	−1.26	2.2	8	0.31	0.12	4	1.6	0.6	4	0.23	0.49	2	0.8	1.3	8
	500 HRS	0.20	0.03	3	0.81	0.243	3	−0.04	0.69	2	−1.0	2.2	8	0.29	0.095	4	1.6	0.67	4	0.19	0.49	2	1.13	1.36	8
	1000 HRS	0.20	0.02	3	0.84	0.28	3	−0.09	0.71	2	−1.0	2.1	8	0.32	0.13	4	1.4	0.8	4	0.15	0.51	2	0.99	1.39	8
200°C	0 HRS	0.20	0.025	3	0.91	0.226	3	−0.16	0.85	2	−1.8	2.7	8	0.24	0.046	4	1.5	0.54	4	0.16	0.437	2	0.70	1.5	8
	500 HRS	0.20	0.026	3	0.91	0.231	3	−0.10	0.85	2	−1.8	2.8	8	0.24	0.038	4	1.6	0.57	4	0.23	0.41	2	0.75	1.6	8
	1000 HRS	0.21	0.035	3	0.92	0.22	3	0.15	0.90	2	−1.8	2.8	8	0.24	0.037	4	1.6	0.57	4	0.19	0.37	2	0.98	1.7	8

Table IV. AD7534 and AD7535 Drift through Life Test

WIRE BOND STRENGTH

To confirm that there was no bond strength degradation during the 200°C life test, a pull strength test was performed on 54 of the wire bonds after the 1000-hour life test. Table V compares the pull strengths obtained after the 1000-hour, 200°C life test with those obtained after a standard 1000-hour, 125°C life test. As can be seen, there is no substantial difference between both sets of figures. In fact, in this particular case, the 200°C figures are better than those for 125°C, and confirm that no bond strength degradation took place.

Pull Strength after 1000 hours (grams)			
125°C		200°C	
$\bar{X}$	σ	$\bar{X}$	σ
8.1	0.79	8.9	0.63

Table V. Wire Bond Strength Comparison

CONCLUSIONS

The AD7534 and AD7535 offer excellent high temperature performance. This is due to the novel low leakage configuration which minimizes leakage current. The 200°C operation of the devices demonstrates clearly how effective the configuration is. In addition, the life test data at 200°C shows that there is no apparent extra failure mechanism. The experimental results (both functional and life test) should be taken as a measure of exceptional device performance and reliability over the specified temperature range (−55°C to +125°C).

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Digital Nulling of Precision Op Amps

FEATURES

- Digitally-controlled offset nulling is achieved by imbalancing the first stage collector currents of a precision op amp.
- Greater than 1.5mV of offset voltage may be nulled to zero with 5μV resolution at 25° C.
- This application is especially useful in microprocessor-controlled systems where stringent error budgets exist.
- Circuit uses the nulling terminals with a DAC-08 substituted for the conventional nulling potentiometer.

GENERAL DESCRIPTION

The input offset voltage of a precision op amp (OP-05, OP-07, OP-77) may be nulled to <5μV using the complementary current outputs of a DAC-08 to change the ratio of collector currents in the first stage. With V_{OS} being defined as the voltage which must be applied between the input terminals to force V_{OUT} to zero and assuming all errors to be in the first stage, V_{OS} may be expressed as:

1) $$V_{OS} = \frac{kT}{q} \log_e \frac{I_{C1}}{I_{C2}} \cdot \frac{I_{S2}}{I_{S1}}$$

where:

k = Boltzmann's constant = 1.38×10^{-23} joules/° K
T = Absolute temperature, ° K
q = Charge of an electron = 1.6×10^{-19} coulomb
I_S = Theoretical reverse-saturation current
I_C = Collector Current

Changing the ratio I_{C1}/I_{C2} over a ±3% range results in an input offset voltage nulling range of greater than 1.5mV at 25° C.

CIRCUIT

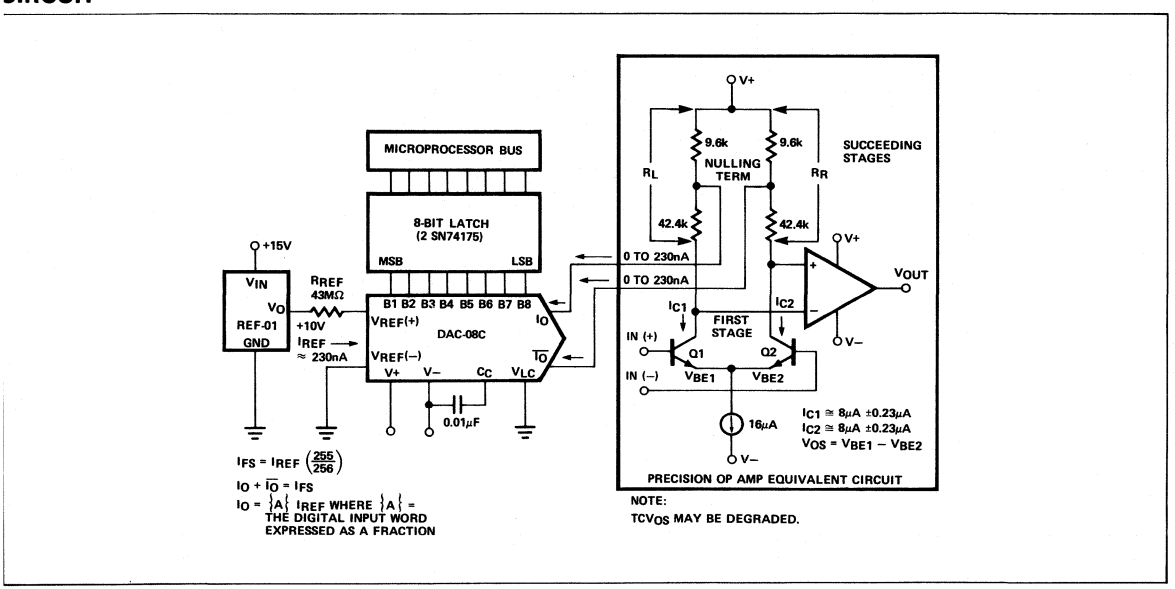

ANALOG DEVICES

AN-17
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

DAC-08
Applications Collection

GENERAL DESCRIPTION

There has been a trend in recent years toward providing totally dedicated digital-to-analog converters with limited applications versatility. This application note describes a new type of monolithic DAC designed for an extremely broad range of applications, the Analog Devices DAC-08.

Several unique design features of this low-cost DAC combine to provide total applications flexibility. Principal among them are: dual complementary, true current outputs; universal logic inputs capable of interfacing with any logic family; 85ns settling time; high-speed multiplying capability; and finally, the ability to use any standard system power supply voltages. A description of these features is given followed by specific applications using each feature.

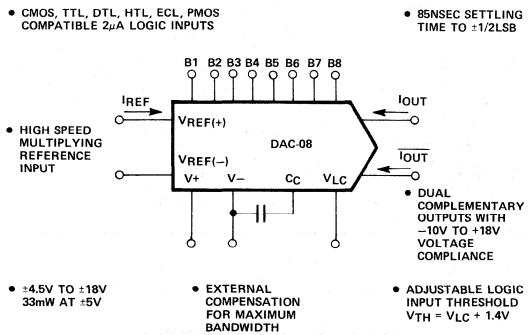

Figure 1. The Flexible D/A Converter

OUTPUT

HIGH VOLTAGE COMPLIANCE CURRENT OUTPUTS

Many older current-output DACs actually have resistive outputs which must be terminated in a virtual ground. The DAC-08, however, is a true digitally-controlled current source with an output impedance typically exceeding 20MΩ.

Its outputs can swing between −10V and +18V with little or no effect on full-scale current or linearity. Some of the applications that require high output voltage compliance include:

1. Precise current transmission over long distances.
2. Programmable current sources.
3. Analog meter movement driving.
4. Resistive termination for a voltage output without an op amp.
5. Capacitive termination for digitally-controlled integrators.
6. Inductive termination with balanced transformers, transducers and headsets.

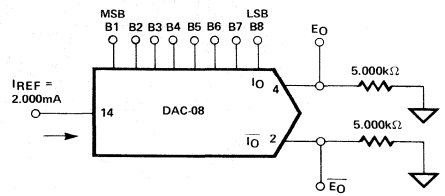

	B1	B2	B3	B4	B5	B6	B7	B8	I_O mA	$\overline{I_O}$ mA	E_O	E_O
FULL SCALE −1LSB	1	1	1	1	1	1	1	1	1.992	0.000	−9.960	0.000
FULL SCALE −2LSB	1	1	1	1	1	1	1	0	1.984	0.008	−9.920	− 0.40
HALF SCALE +LSB	1	0	0	0	0	0	0	1	1.008	0.984	−5.040	−4.920
HALF SCALE	1	0	0	0	0	0	0	0	1.000	0.992	−5.000	−4.960
HALF SCALE −LSB	0	1	1	1	1	1	1	1	0.992	1.000	−4.960	−5.000
ZERO SCALE +LSB	0	0	0	0	0	0	0	1	0.008	1.984	−0.040	−9.920
ZERO SCALE	0	0	0	0	0	0	0	0	0.000	1.992	0.000	−9.960

Figure 2. Basic Unipolar Negative Operation

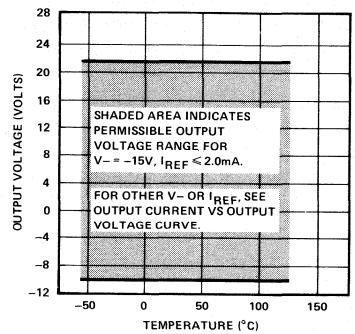

Figure 3. Output Voltage Compliance vs Temperature

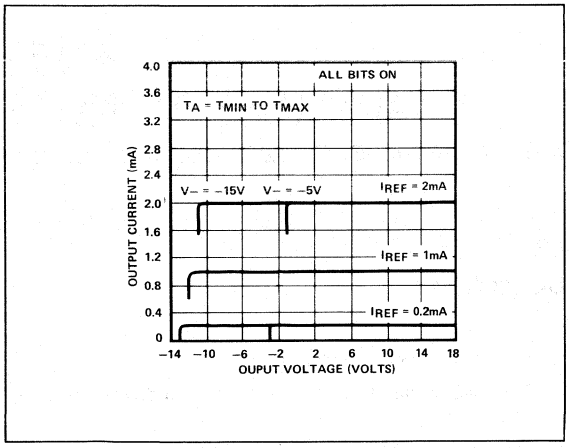

Figure 4. Output Current vs Output Voltage (Output Voltage Compliance)

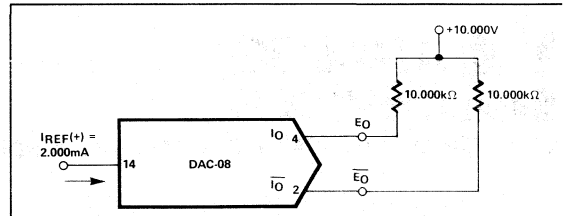

	B1	B2	B3	B4	B5	B6	B7	B8	E_O	$\overline{E_O}$
POS FULL SCALE –LSB	1	1	1	1	1	1	1	1	– 9.920	+10.000
POS FULL SCALE 2LSB	1	1	1	1	1	1	1	0	– 9.840	+ 9.920
ZERO SCALE +LSB	1	0	0	0	0	0	0	1	– 0.080	+ 0.160
ZERO SCALE	1	0	0	0	0	0	0	0	0.000	+ 0.080
ZERO SCALE –LSB	0	1	1	1	1	1	1	1	+ 0.080	0.000
NEG FULL SCALE +LSB	0	0	0	0	0	0	0	1	+ 9.920	– 9.840
NEG FULL SCALE	0	0	0	0	0	0	0	0	+10.000	– 9.920

Figure 5. Basic Bipolar Output Operation

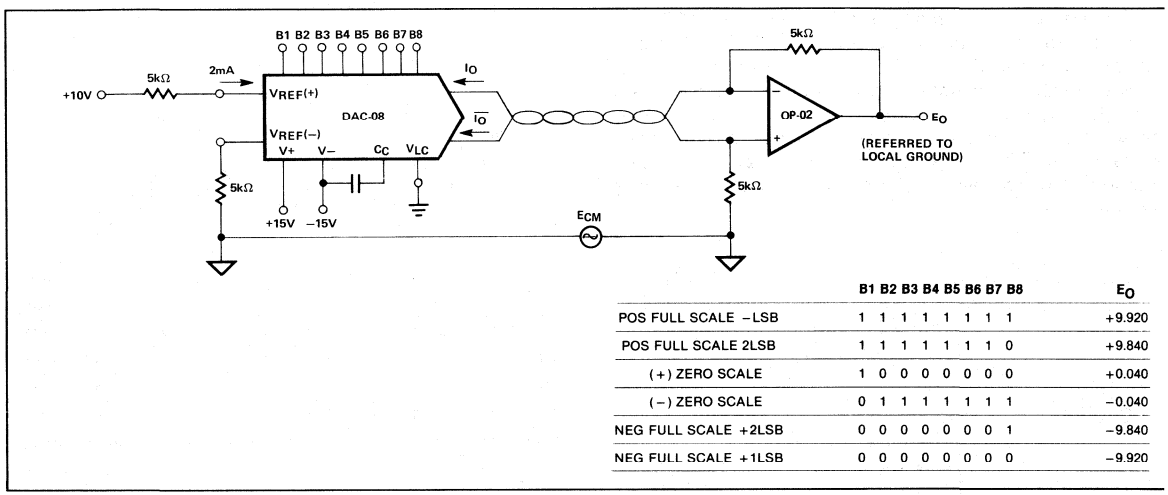

	B1	B2	B3	B4	B5	B6	B7	B8	E_O
POS FULL SCALE –LSB	1	1	1	1	1	1	1	1	+9.920
POS FULL SCALE 2LSB	1	1	1	1	1	1	1	0	+9.840
(+) ZERO SCALE	1	0	0	0	0	0	0	0	+0.040
(–) ZERO SCALE	0	1	1	1	1	1	1	1	–0.040
NEG FULL SCALE +2LSB	0	0	0	0	0	0	0	1	–9.840
NEG FULL SCALE +1LSB	0	0	0	0	0	0	0	0	–9.920

Figure 6. High Noise Immunity Current to Voltage Conversion

DUAL COMPLEMENTARY OUTPUTS

Convertional DACs have a single output, so they cannot drive balanced loads and are limited to a single input code polarity. The DAC-08 was designed to overcome these limitations.

Input coding of positive binary or complementary binary is obtained by a choice of outputs, I_O for positive-true or I_O for negative-true. In many applications both are used either independently or in combination. Dual comlementary outputs allow some very unusual and useful DAC applications:

1. CRT display driving without transformers.
2. Differential transducer control systems.
3. Differential line driving.
4. High-speed waveform generation.
5. Digitally controlled offset nulling of op amps.

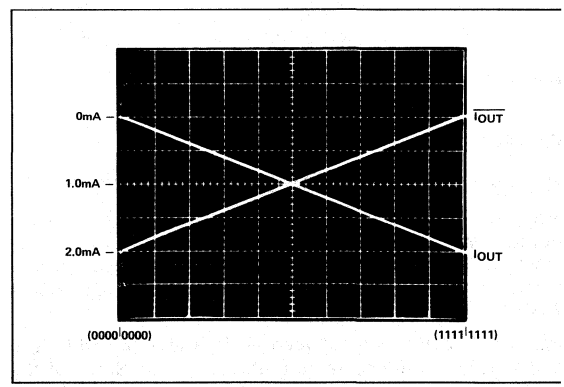

Figure 7. True and Complementary Output Operation

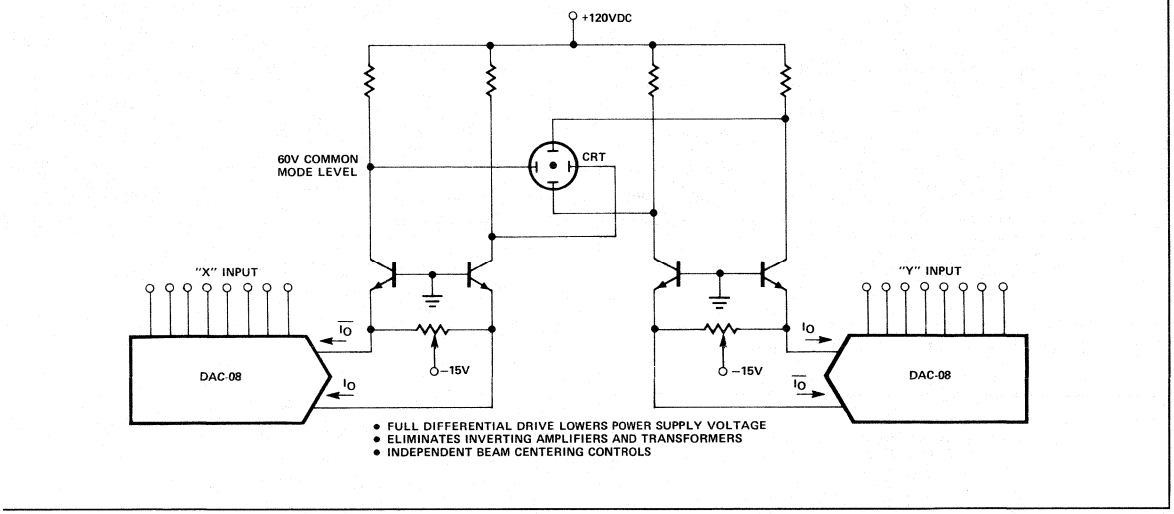

Figure 8. CRT Display Driver

8

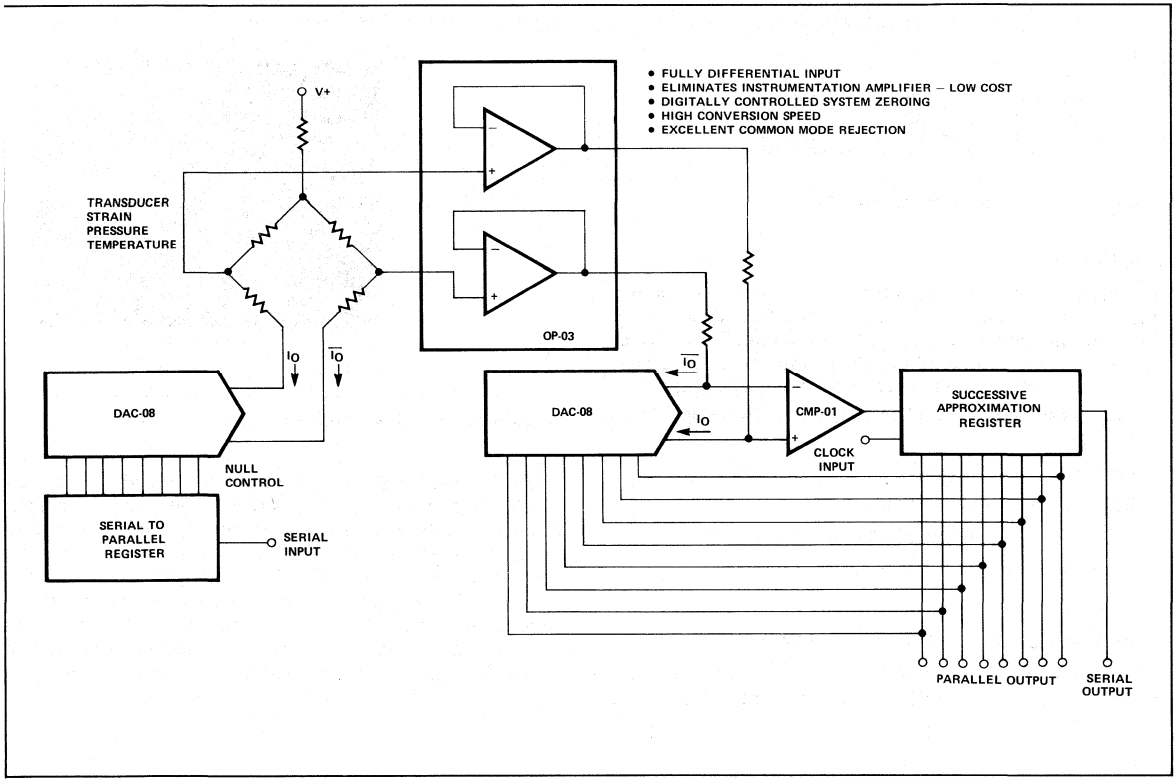

Figure 9. Bridge Transducer Control System with Full Differential Input

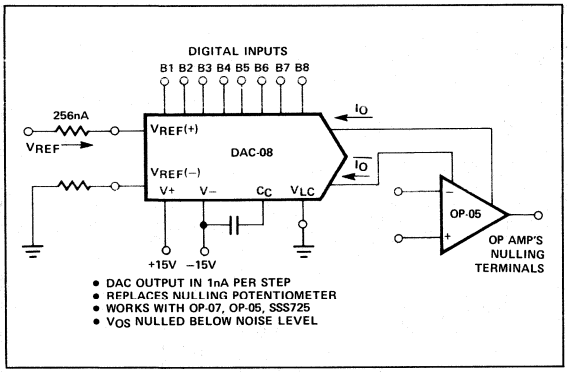

Figure 10. Digitally Controlled Offset Nulling

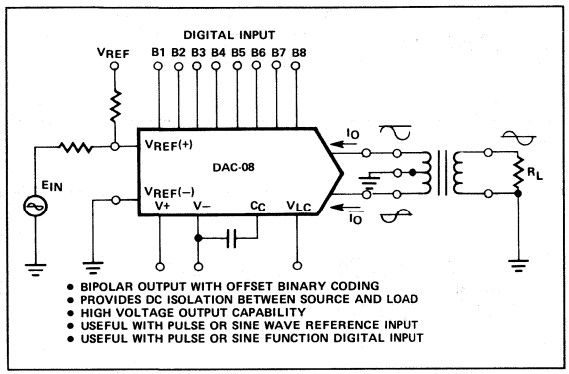

Figure 11. Balanced Transformer Drive

HIGH SPEED

Sub-microsecond settling times are common in current-output DACs. Many DACs settle in 500ns; 300ns is not unusual. But 85ns settling time for a low-cost DAC is exceptional, and this characteristic allows the use of the DAC-08 in formerly difficult and expensive-to-build applications:

1. 1μs, 2μs and 4μs A/Ds. (These are completely described in AN-16.)
2. 15MHz Tracking A/Ds.
3. ECL compatible applications.
4. Video displays requiring a low-glitch DAC.
5. Radar pulse height analysis system.

LOGIC INPUTS

ADJUSTABLE INPUT LOGIC THRESHOLD

Most DACs have TTL or CMOS compatible inputs which require complicated interfaces for use with ECL, PMOS,

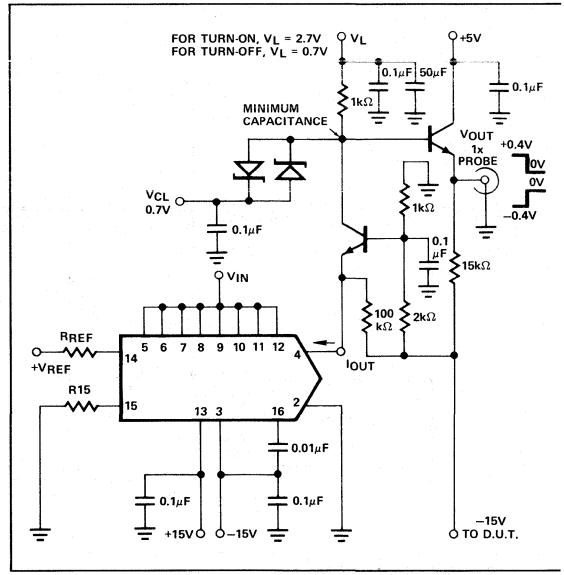

Figure 12. Settling Time Measurement Circuit

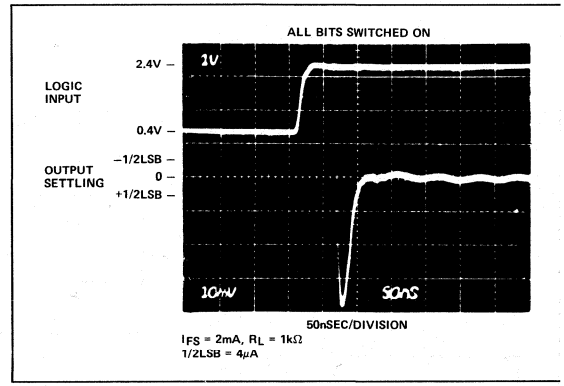

Figure 13. Full Scale Settling Time

NMOS or HTL logic. By contrast, the DAC-08, with typica logic input current of 2μA and an adjustable input logic threshold, interfaces easily with any logic family in use today. The logic input threshold is 1.4V positive with respect to Pin 1; for TTL Pin 1 is therefore grounded; for other families Pin 1 is connected as shown in the interfacing figure. An adjustable threshold and a −10V to +18V input range greatly simplify system design especially with other than-TTL logic. The circuits shown in c and d provide a $2V_{BE}$ V_{LC} compensation to minimize temperature drift.

ECL applications without level translators.

Direct interfaces with Hi-Z RAM outputs.

CMOS applications without static discharge considerations.

HTL or HNIL applications without level translators.

System size, weight, and cost reduction.

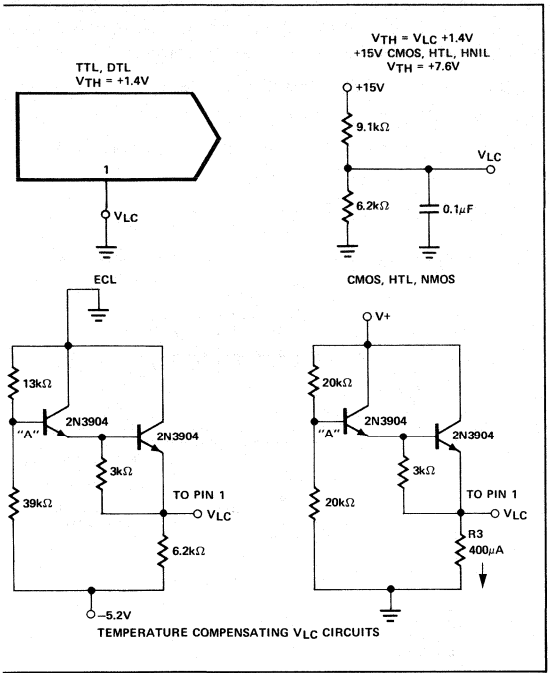

igure 14. Interfacing with Various Logic Families ($V_{TH} = V_{LC} + 1.4V$)

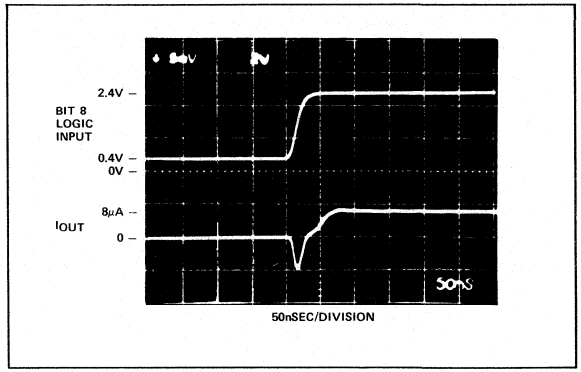

Figure 16. LSB Switching

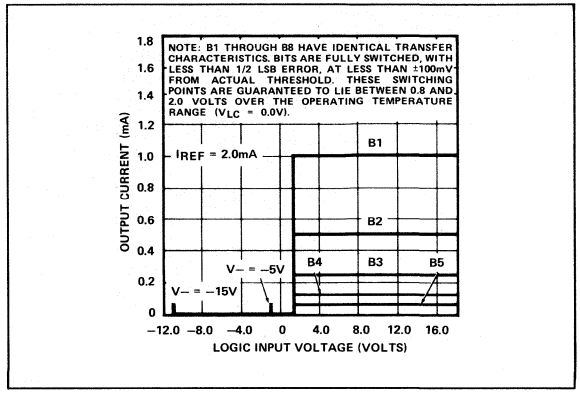

Figure 17. Bit Transfer Characteristics

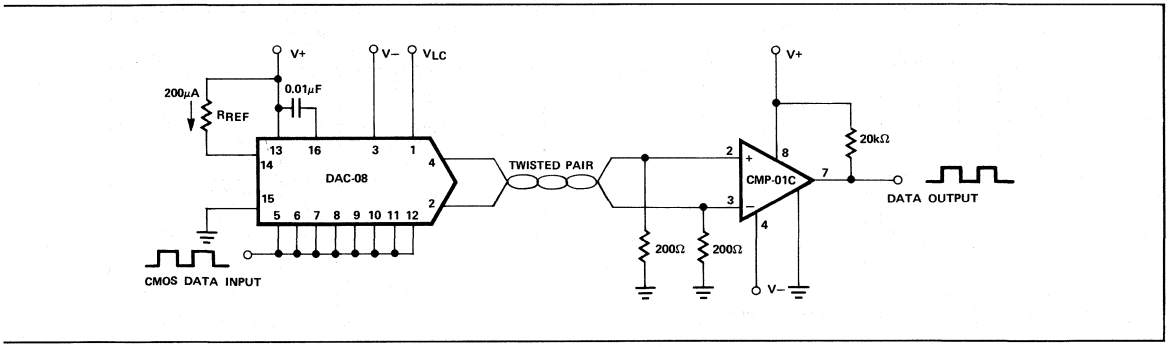

igure 15. CMOS Differential Line Driver/Receiver

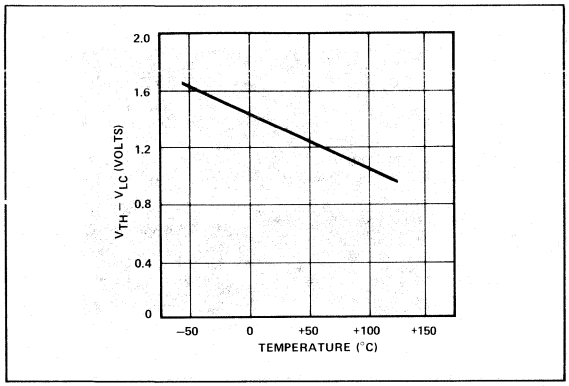

Figure 18. $V_{TH} - V_{LC}$ vs Temperature

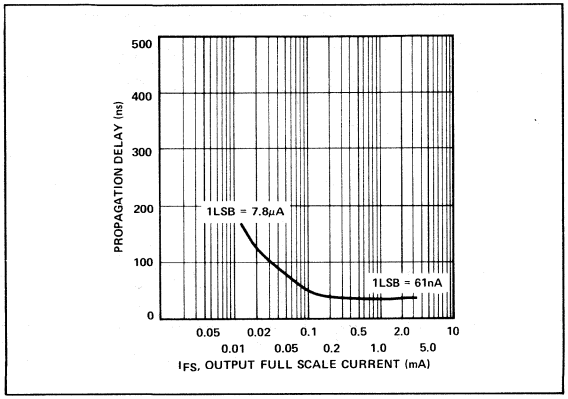

Figure 19. LSB Propogation Delay vs I_{FS}

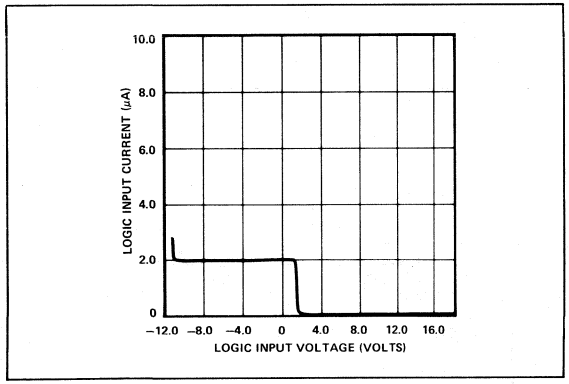

Figure 20. Logic Input Current vs Input Voltage

REFERENCE INPUTS

MULTIPLYING CAPABILITY

Fixed internal references are included in many DACs, bu they limit the user to non-multiplying, single-polarit reference applications and do not allow a single-syster reference. To achieve the design goals of low cost and tota applications flexibility, the DAC-08 uses an externa reference. Positive or negative references may be applie over a wide common-mode voltage range. In addition, th full-scale current is matched to the reference curren eliminating calibration in most applications.

1. Digitally controlled full-scale calibration.
2. 8 x 8 multiplication of two digital words.
3. Digital Attenuators/Programmable gain amplifiers.
4. Modem transmitters to 1MHz.
5. Remote shutdown and party line DAC applications.

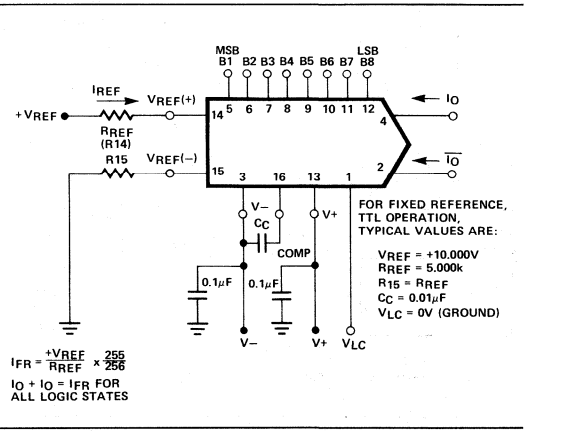

Figure 21. Basic Positive Reference Operation

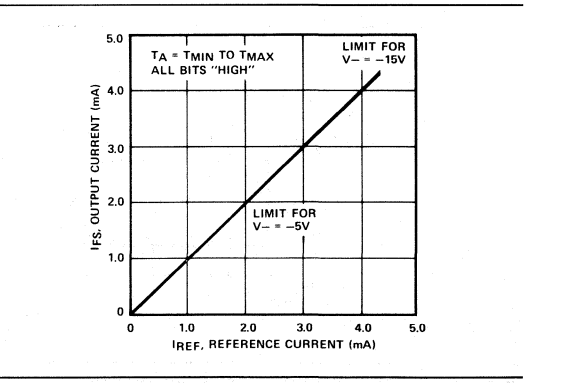

Figure 22. Full-Scale Current vs Reference Current

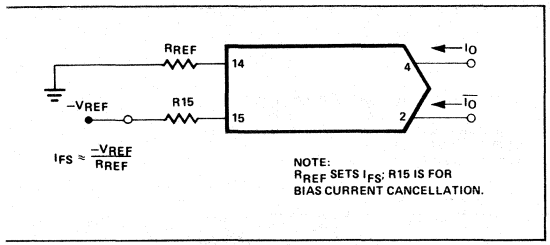

gure 23. Basic Negative Reference Operation

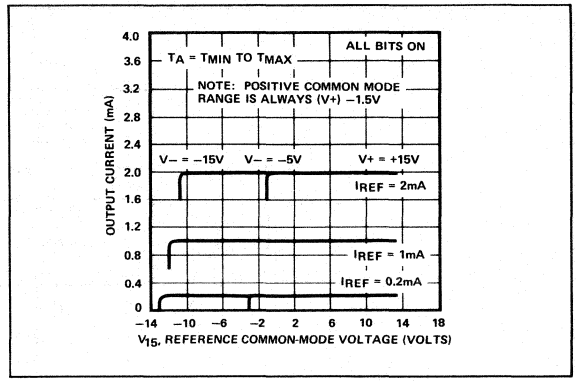

Figure 24. Reference Amp Common-Mode Range

IGH SPEED

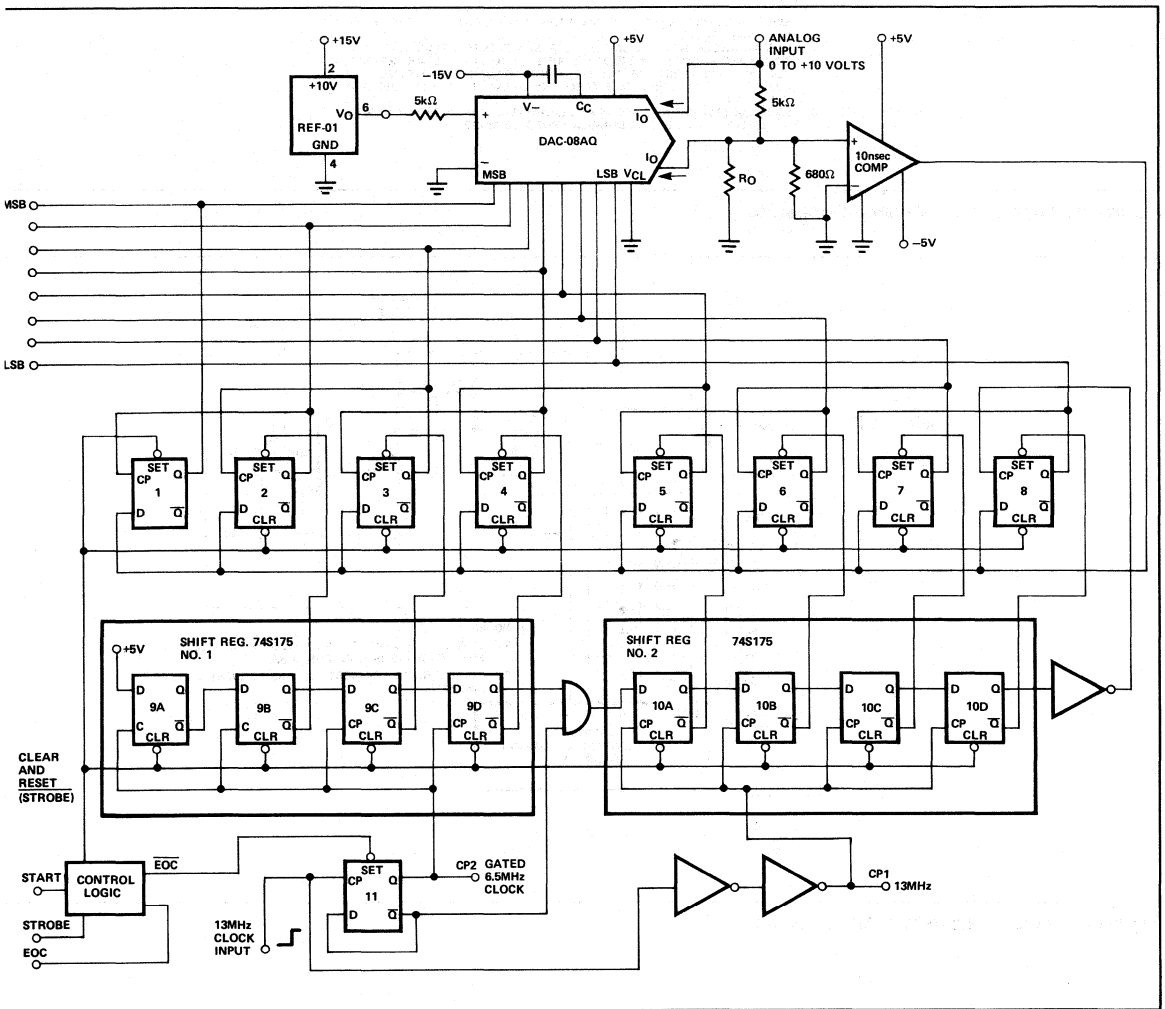

Figure 25. Simplified Schematic 1μs A/D

8

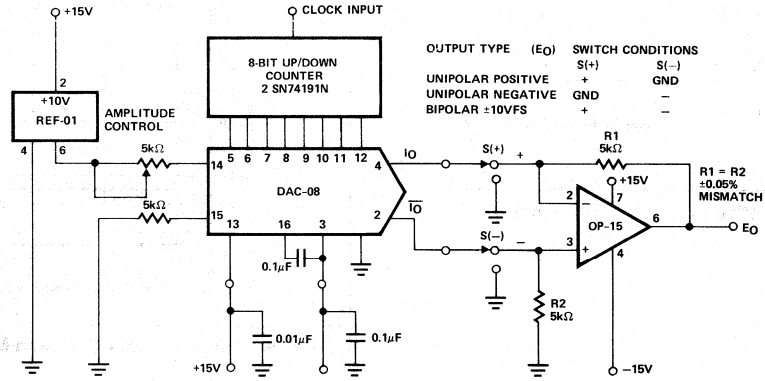

NOTES:

1. BIPOLAR OUTPUT IS SYMMETRICAL AROUND ZERO. ADJUSTABLE PEAK-TO-PEAK AMPLITUDE.
2. FOR TRIANGLE WAVE, COUNT UP TO FULL REVERSE AND COUNT DOWN.
3. FOR POSITIVE-GOING SAWTOOTH, COUNT UP TO FULL CLEAR, REPEAT.
4. FOR NEGATIVE-GOING SAWTOOTH, COUNT DOWN. CLEAR REPEAT.
5. FOR OTHER WAVEFORMS, USE A ROM PROGRAMMED WITH THE DESIRED FUNCTION.
6. 85nSEC SETTLING TIME PERMITS WAVEFORM PERIOD OF 25.6μSEC OR 39kHz REPETITION RATE IN THE UNIPOLAR POSITIVE MODE.

Figure 26. High-Speed Waveform Generator

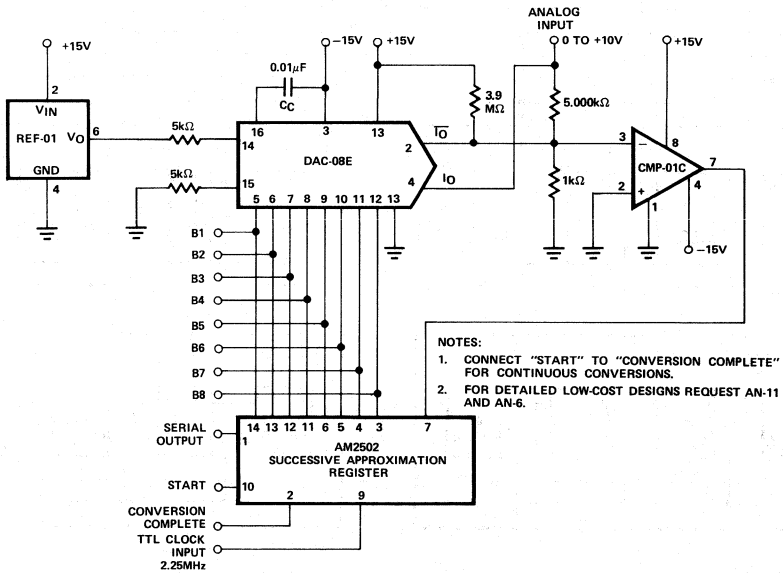

NOTES:

1. CONNECT "START" TO "CONVERSION COMPLETE" FOR CONTINUOUS CONVERSIONS.
2. FOR DETAILED LOW-COST DESIGNS REQUEST AN-11 AND AN-6.

Figure 27. Four IC Low-Cost A/D Converter

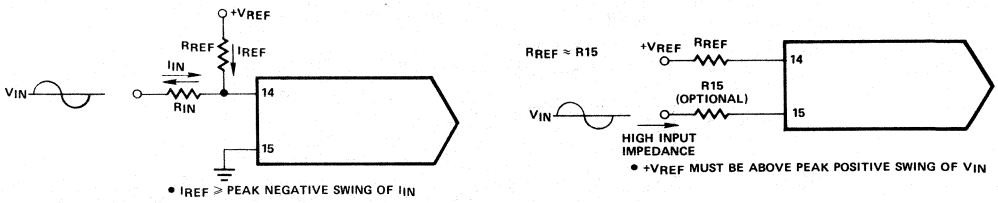

gure 28. Accommodating Bipolar References

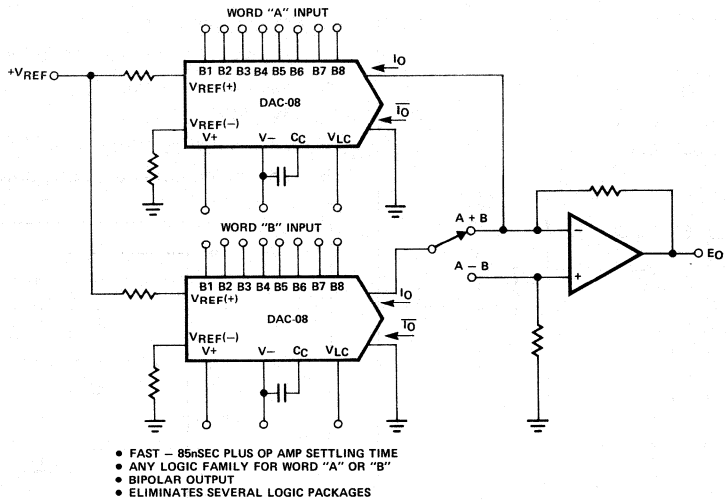

gure 29. Digital Addition or Subtraction with Analog Output

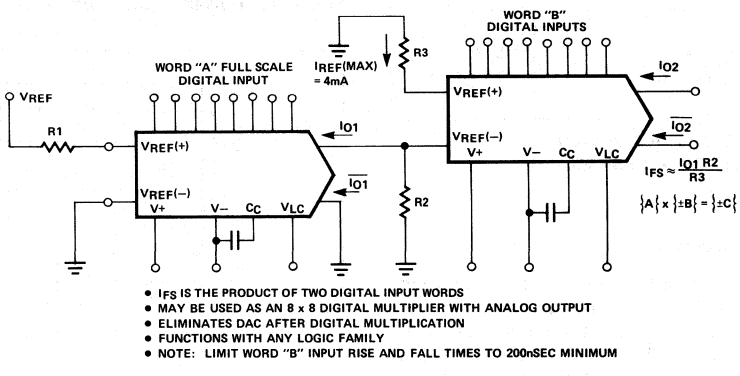

gure 30. Digitally Controlled Full-Scale Calibration (Multiplier)

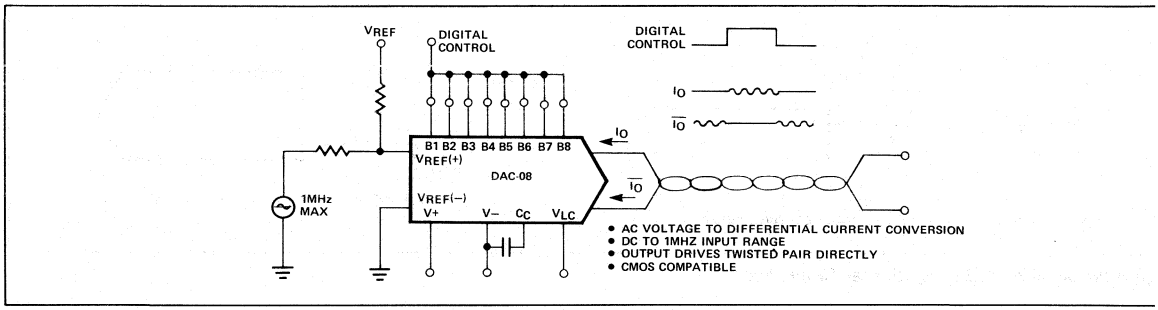

Figure 31. Modem Transmitter

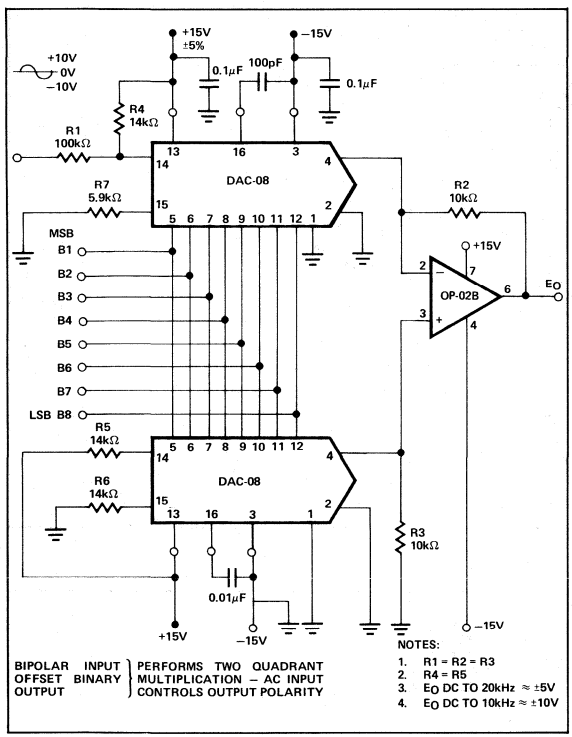

Figure 32. DC-Coupled Digital Attenuator/ Programmable Gain Amplifier

POWER SUPPLIES

POWER SUPPLY REQUIREMENTS

The DAC-08 works with ±4.5V to ±18V supplies allowing use with all standard digital and analog system supply voltages plus most battery voltages. With only 33mW of power dissipation at ±5V and 85ns settling time, it has a lower speed power product than CMOS DACs. Power dissipation is almost constant over temperature, and bypassing is accomplished with 0.01μF capacitors — no large electrolytics are required. These power supply requirements allow:

1. Battery operation.
2. Use of unregulated or poorly regulated power supplies
3. Use in space-limited areas due to small bypas capacitors.
4. Use in constant power dissipation applications.
5. Common digital and analog power supplies.

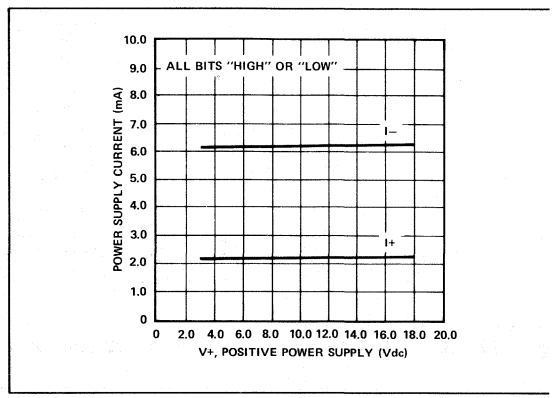

Figure 33. Power Supply Current vs V+

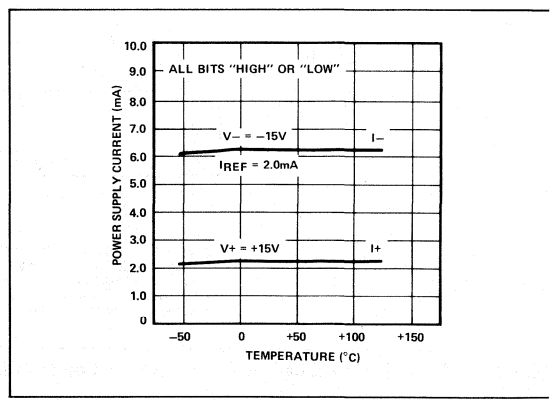

Figure 34. Power Supply Current vs Temperature

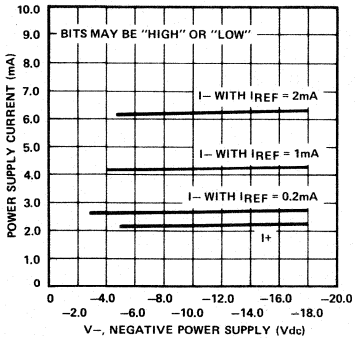

igure 35. Power Supply Current vs V–

THER APPLICATIONS

ICROPROCESSOR APPLICATIONS

he ability to use μP power supply voltages and the ability to nterface with any logic family make the DAC-08 especially useful in μP applications:

1. Tracking A/D converters.
2. Successive approximation A/D converters.
3. Direct drive from Hi-Z MOS RAM outputs.

By programming the ROMs with the successive approximation or the tracking A/D algorithm, all of the logic for A/D conversion is contained in the μP. This is a very inexpensive approach, since there is no need for the usual A/D conversion logic packages.

CONCLUSION

High voltage compliance complementary current outputs, universal logic inputs and multiplying capability make the Precision Monolithics DAC-08 the most versatile monolithic high-speed DAC available today.

8

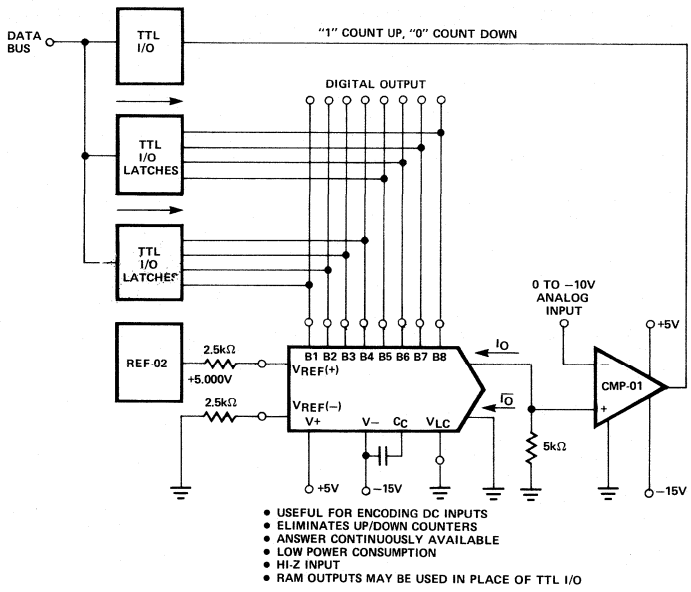

igure 36. Microprocessor Controlled Tracking A/D Converter

)NE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Differential and Multiplying D/A Converter Applications

NTRODUCTION

he introduction of low-cost monolithic D/A converters has implified data acquisition and control system design. This pplication note describes several new applications using ıe multiplying capability and dual complementary current utputs of the Analog Devices DAC-08.

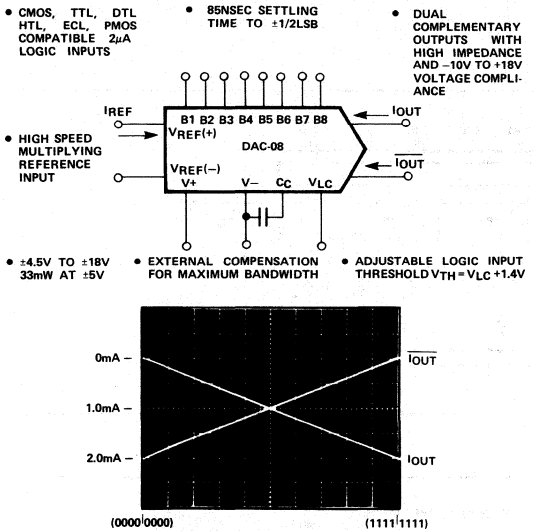

Figure 1. The Universal DAC

MULTIPLYING DAC BASICS

multiplying DAC has an analog output which is the product of a digital input word and a reference voltage and can e expressed as:

$$1.)\ E_O = E_{REF}\left(\frac{A1}{2}+\frac{A2}{4}+\frac{A3}{8}+\frac{A4}{16}+\frac{A5}{32}+\frac{A6}{64}+\frac{A7}{128}+\frac{A8}{256}\right)$$

or a current reference, current output DAC, the expression ecomes:

$$2.)\ I_O = I_{REF}\left(\frac{A1}{2}+\frac{A2}{4}+\frac{A3}{8}+\frac{A4}{16}+\frac{A5}{32}+\frac{A6}{64}+\frac{A7}{128}+\frac{A8}{256}\right)$$

he DAC-08 has complementary/differential current oututs, and I_O has a complement expressed as:

3.) $\overline{I_O} = I_{FS} - I_O$ for all input logic states.

The relationship of I_{REF} to I_O and $\overline{I_O}$ is illustrated in Figure 2 and in Figure 3, the basic DC reference connections. References may be either positive or negative, and a bipolar output voltage may be achieved using the high compliance current outputs alone or with an output operational amplifier. The simplest form of a multiplying DAC accepts a unipolar varying reference input.

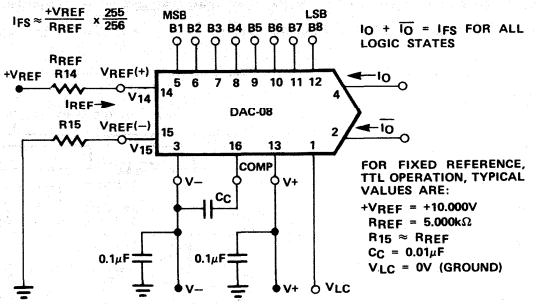

NOTES:
1. R_{REF} SETS I_{FS}. R15 IS FOR BIAS CURRENT CANCELLATION.
2. PINS 14 AND 15 ARE OP AMP INPUTS, SO V14 ≈ V15.

Figure 2. Positive Reference Connection

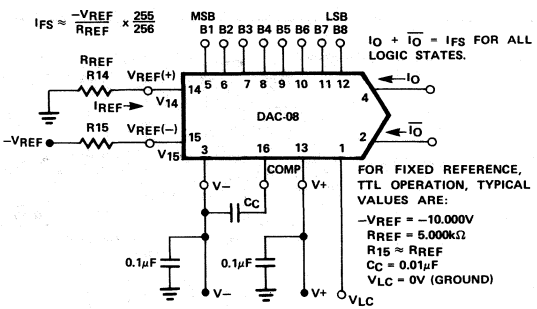

NOTES:
1. R_{REF} SETS I_{FS}. R15 IS FOR BIAS CURRENT CANCELLATION.
2. PINS 14 AND 15 ARE OP AMP INPUTS, SO V14 ≈ V15.

Figure 3. Negative Reference Connection

BIPOLAR REFERENCES

Operation with bipolar references is achieved by modulating I_{REF} as shown in Figure 5. To aid in understanding bipolar operation, see the equivalent circuit in Figure 4. The reference inputs of the DAC-08 are op amp inputs — $V_{REF}(+)$ being the inverting input and $V_{REF}(-)$ being the noninverting input. Excellent gain linearity of the reference amplifier allows multiplying operation over a range of I_{REF} of 4μA to 4mA with monotonic operation from less than 100μA to 4mA.

REFERENCE AMPLIFIER COMPENSATION

AC reference applications will require the reference amplifier to be compensated using a capacitor from Pin 16 to V−. The value of this capacitor depends on the impedance presented to Pin 14: for R_{IN} values of 1.0, 2.5 and 5.0kΩ, minimum values of C_C are 15, 37, and 75pF. Larger values of R_{IN} require proportionately increased values of C_C for proper phase margin.

FAST PULSED OPERATION

For fastest multiplying response, low values of R_{IN} enabling small C_C values should be used. For R_{IN} = 1kΩ and C_C = 15pF, the reference amplifier slews at 4mA/μs enabling a transitic from I_{REF} = 0 to I_{REF} = 2mA in 500ns. If R_{IN} or the parall equivalent resistance at Pin 14 is less than 200Ω, no compe sation capacitor is necessary, and a full-scale transitic requires only 16ns.

TWO-QUADRANT MULTIPLICATION

There are two forms of two-quadrant multiplication: bipol digital, where the digital input word controls output polarit and bipolar analog, where the analog reference input co trols output polarity.

Bipolar digital two-quadrant mutliplication is shown Figure 6 with the output polarity being controlled by a offset-binary-coded digital input word.

Bipolar analog two-quadrant multiplication is shown Figure 7. A bipolar reference voltage is connected to the u per DAC-08 and modulates the reference current by ±1.0m around a quiescent current of 1.1mA. The lower DAC-08 als has a reference current of 1.1mA; due to the parallel digit inputs, the lower DAC-08 effectively subtracts out the quie cent 1.1mA of the upper DAC-08's reference current at all i

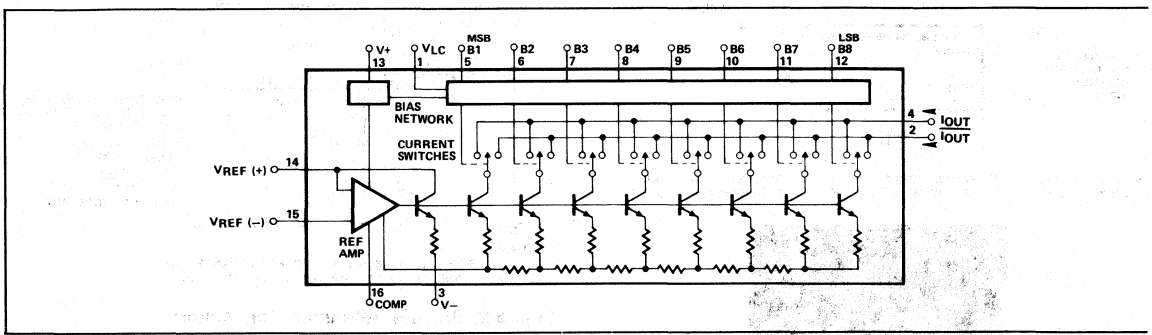

Figure 4. DAC-08 Equivalent Circuit

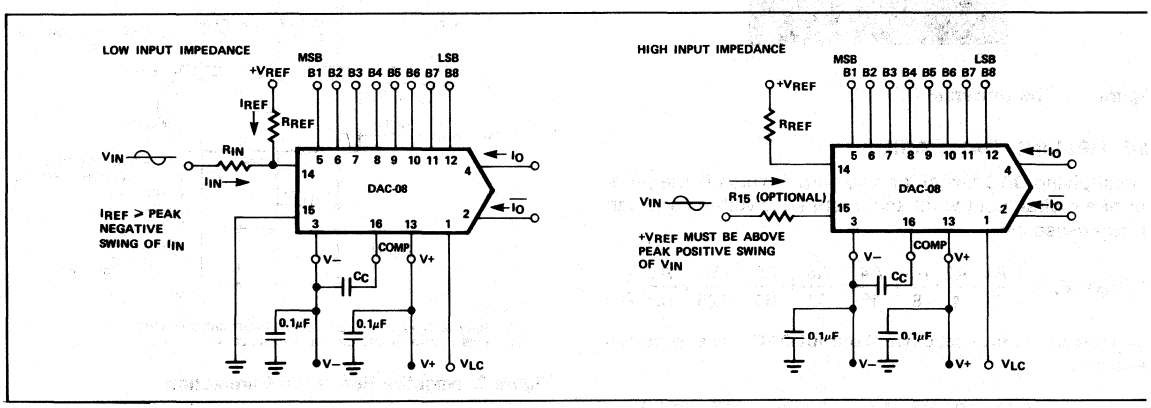

Figure 5. Bipolar Reference Connections

ıt codes, since the voltage across R3 varies between 10V and 0V. Thus, the output voltage, E_O, is a product of a gital input word and a bipolar analog reference voltage.

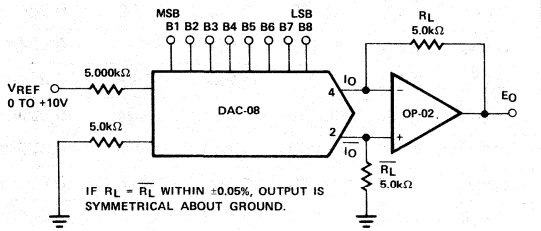

	B1	B2	B3	B4	B5	B6	B7	B8	I_O mA	$\overline{I_O}$ mA	E_O (V)
POS. FULL SCALE	1	1	1	1	1	1	1	1	1.992	0.000	+9.96
POS. FULL SCALE −LSB	1	1	1	1	1	1	1	0	1.984	0.008	+9.88
(+) ZERO SCALE	1	0	0	0	0	0	0	0	1.000	0.992	+0.040
(−) ZERO SCALE	0	1	1	1	1	1	1	1	0.992	1.000	−0.040
NEG. FULL SCALE +LSB	0	0	0	0	0	0	0	1	0.008	1.984	−9.88
NEG. FULL SCALE	0	0	0	0	0	0	0	0	0.000	1.992	−9.96

igure 6. Bipolar Digital Two-Quadrant Multiplication (Symmetrical Offset Binary)

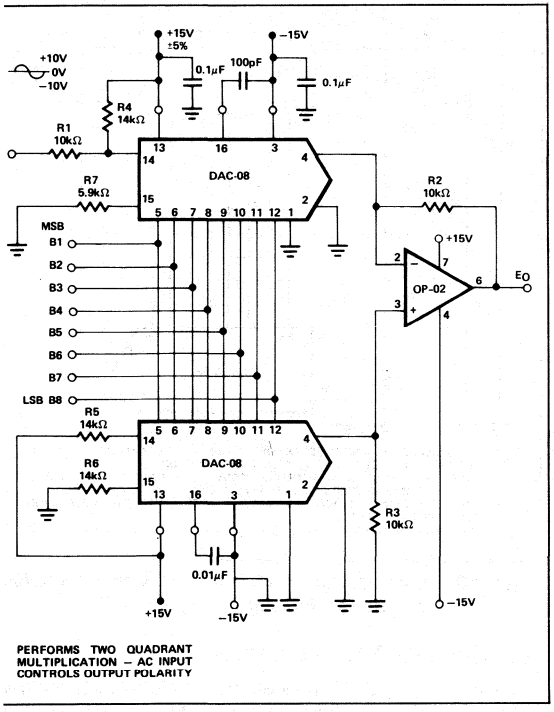

igure 7. Bipolar Analog Two-Quadrant Multiplication (DC-Coupled Digital Attenuator)

FOUR-QUADRANT MULTIPLICATION

Four-quadrant multiplication combines the two forms of two-quadrant multiplication. Output analog polarity is controlled by either the analog input reference or by the offset binary digital input word. One implementation of this function with the DAC-08 is shown in Figure 8 with output current values listed in Table 1.

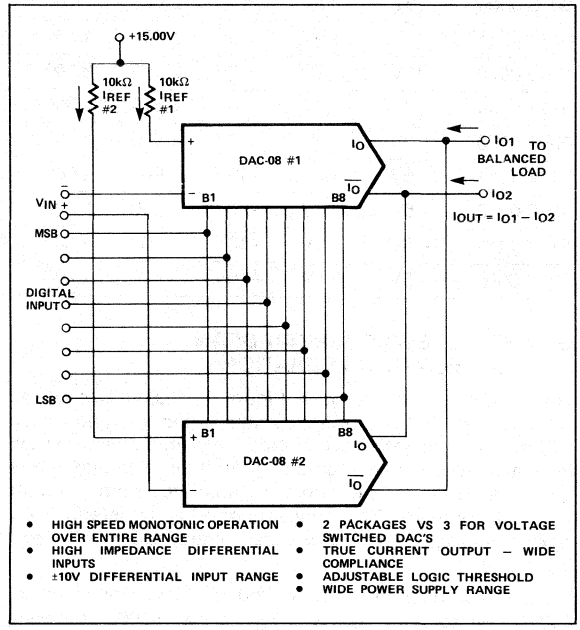

Figure 8. Four-Quadrant Multiplying DAC with with Impedance Input

The four-quadrant multiplying DAC circuit shown accepts a differential voltage input and produces a differential current output. An output op amp is not shown because it is not always required; many applications are more suited for high output compliance (−10V to +18V) differential current outputs. Typical balanced loads include transformers, transducers, transmission lines, bridges and servos.

Operation of the four-quadrant multiplier may be more easily visualized by considering that if either $V_{IN} = 0V$ or the offset binary digital input code is at midscale (corresponding to zero), then a change in the other input will not affect the output. Zero multiplied by any number equals zero.

A common mode current will be present at the output and must be accommodated by the balanced load. A pair of matched resistors may be used at the outputs to shunt most of the common mode current to ground, thus reducing the common mode voltage swing at the output.

8

Table 1. Four-Quadrant Multiplying Current Values in Figure 8.

DIGITAL INPUT	$V_{IN}(+)$	$V_{IN}(-)$	V_{IN} DIFF.	I_{REF} #1 (mA)	I_{REF} #2 (mA)	I_O#1 (mA)	I_O#2 (mA)	I_{O1} (mA)	I_O#2 (mA)	I_O#1 (mA)	I_{O2} (mA)	I_{OUT} DIFF.
1111 1111	+5V	−5V	+10V	2.000	1.000	1.992	0	1.992	0.996	0	0.996	0.996m
1000 0000	+5V	−5V	+10V	2.000	1.000	1.000	0.496	1.496	0.500	0.992	1.492	0.004m
0111 1111	+5V	−5V	+10V	2.000	1.000	0.992	0.500	1.492	0.496	1.000	1.496	−0.004m
0000 0000	+5V	−5V	+10V	2.000	1.000	0	0.996	0.996	0	1.992	1.992	−0.996m
1111 1111	0V	0V	0V	1.500	1.500	1.494	0	1.494	1.494	0	1.494	0.000m
1000 0000	−10V	−10V	0V	2.500	2.500	1.250	1.240	2.490	1.250	1.240	2.490	0.000m
0111 1111	+10V	+10V	0V	0.500	0.500	0.248	0.250	0.498	0.248	0.250	0.498	0.000m
0000 0000	0V	0V	0V	1.500	1.500	0	1.494	1.494	0	1.494	1.494	0.000m
1111 1111	−5V	+5V	−10V	1.000	2.000	0.996	0	0.996	1.992	0	1.992	−0.996m
1000 0000	−5V	+5V	−10V	1.000	2.000	0.500	0.992	1.492	1.000	0.496	1.496	−0.004m
0111 1111	−5V	+5V	−10V	1.000	2.000	0.496	1.000	1.496	0.992	0.500	1.492	0.004m
0000 0000	−5V	+5V	−10V	1.000	2.000	0	1.992	1.992	0	0.996	0.996	0.996m

HIGHEST SPEED FOUR-QUADRANT MULTIPLYING CONSIDERATIONS

The configuration shown in Figure 10 makes use of the DAC-08's ability to operate in a fast-pulsed reference mode without compensation capacitors. This technique provides lowest full-scale transition times. An internal clamp allows quick recovery of the reference amplifier from a cutoff ($I_{REF} = 0$) condition. This connection yields a reference slew rate of 16mA/μs which is relatively independent of R_{IN} and V_{IN} values.

Input resistances are not limited to 10kΩ. For example, 100kΩ resistors for R_{IN1} and R_{IN2} allow ±100V reference voltage inputs making this connection especially useful in high common mode voltage environments. Except for different reference treatment, operation and digital input coding are identical in the circuits shown in Figure 8 and in Figure 10; both have the transfer function shown in Figure 9.

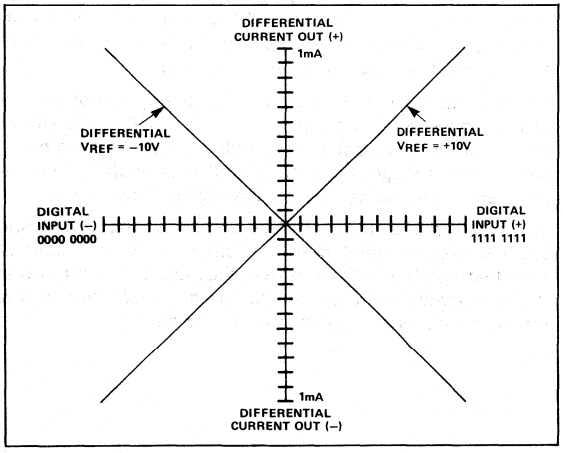

Figure 9. Four-Quadrant Multiplying DAC Transfer Function

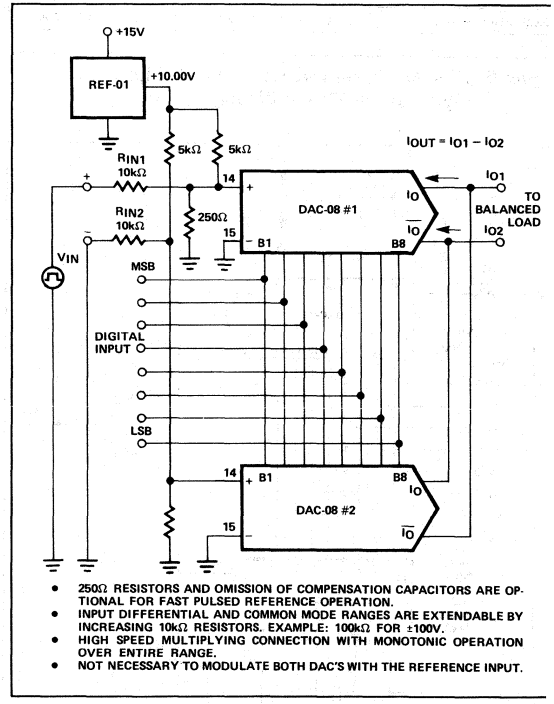

- 250Ω RESISTORS AND OMISSION OF COMPENSATION CAPACITORS ARE OPTIONAL FOR FAST PULSED REFERENCE OPERATION.
- INPUT DIFFERENTIAL AND COMMON MODE RANGES ARE EXTENDABLE BY INCREASING 10kΩ RESISTORS. EXAMPLE: 100kΩ FOR ±100V.
- HIGH SPEED MULTIPLYING CONNECTION WITH MONOTONIC OPERATION OVER ENTIRE RANGE.
- NOT NECESSARY TO MODULATE BOTH DAC'S WITH THE REFERENCE INPUT.

Figure 10. Four-Quadrant Multiplying DAC with Extendable Input Range and Highest Speed

AC-COUPLED MULTIPLICATION

Some multiplying DAC applications are more easily achieve with AC coupling. At the same time, a high impedance inpu is often required to avoid loading a relatively high source in pedance. Both requirements are met by the circuits show in Figure 11 and Figure 12 which use the compensatic

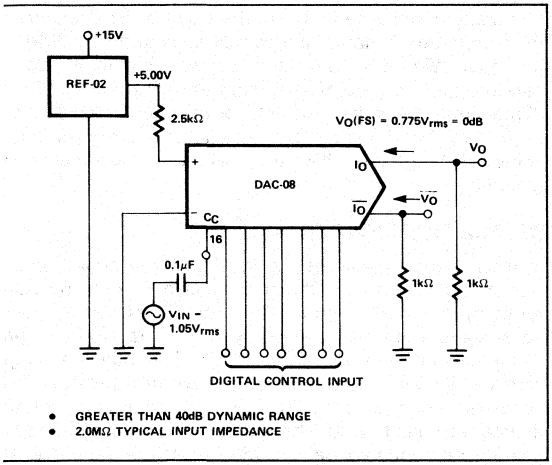

igure 11. High Input Impedance AC-Coupled Multiplication (Audio Frequency Digital Attenutor)

apacitor terminal (C_C) as an input. This is possible because 'C is the base of a transistor whose emitter is one diode rop (0.7V) away from the R-2R ladder network common aseline internal to the DAC-08.

Vith a full-scale input code the output, V_O, is flat to ›200kHz and is 3dB down at approximately 1.0MHz making ıis type of multiplying connection useful even beyond the udio frequency range. Such a connection is illustrated in igure 12 operating at 455kHz, the highest recommended perating frequency in this connection.

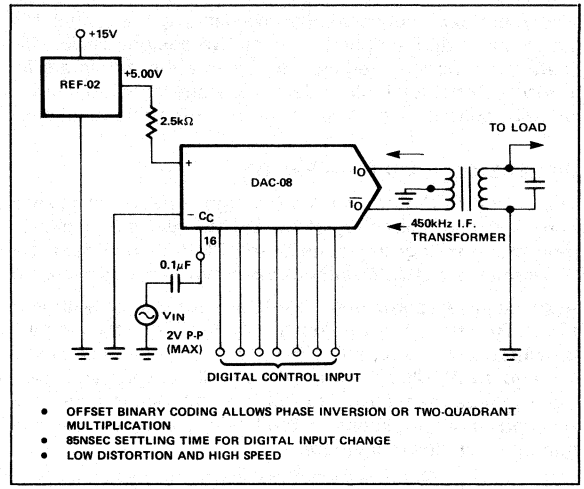

Figure 12. High Input Impedance AC-Coupled Multiplication (I.F. Amplifier/Digital Attenuator)

DIFFERENTIAL AND RATIOMETRIC A/D CONVERSION

8

Complementary/differential current-source outputs and multiplying capability allow the DAC-08 to be used in differential and ratiometric A/D converter designs directly without signal conditioning amplifiers. This group of applications begins with the basic differential A/D converter and ratiometric A/D converter connections followed by more specific applications.

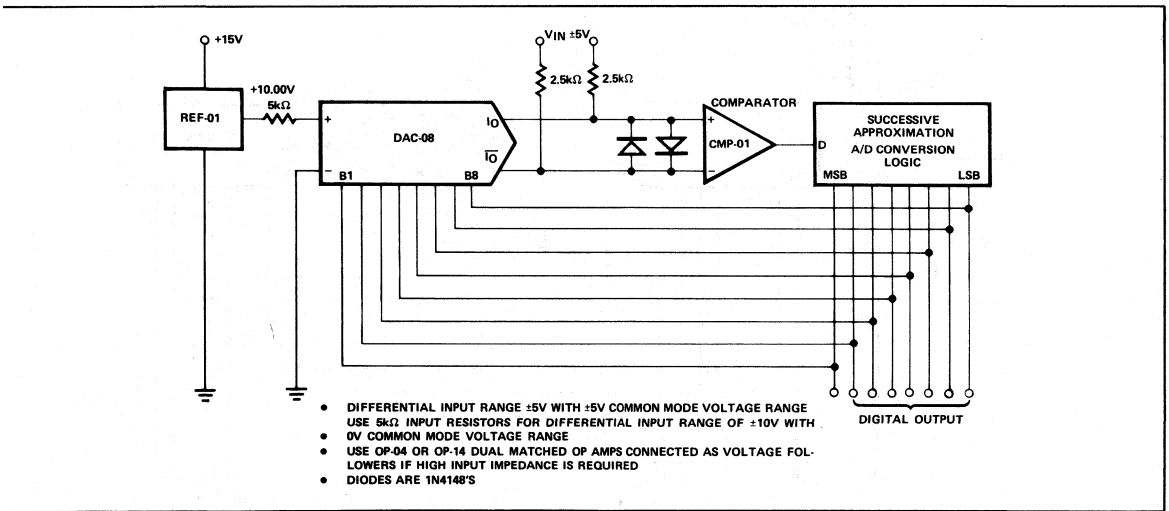

Figure 13. Differential Input A/D Conversion Basic Connections

These are extremely cost-effective designs due to their low parts count and simplicity. Alternative designs performing identical functions require instrumentation amplifiers for the differential-to-single-ended input signal conditioning and analog multipliers or dividers for the arithmetic functions.

DIFFERENTIAL A/D CONVERSION

The circuit in Figure 13 uses the high voltage compliance current output capability of the DAC-08 and the high common mode voltage rejection of the CMP-01 to construct a differential input ADC without input signal conditioning.

A successive approximation ADC is constructed with four ICs: a REF-01 +10V reference, a 2502-type successive approximation register, a CMP-01 precision voltage comparator, and a DAC-08. As shown, the circuit converts an analog input in less than 2.0μs. For lower speed requirements, the A/D conversion logic can be the tracking or servo type consisting of up/down counters.

FOUR-QUADRANT RATIOMETRIC A/D CONVERSION

Ratiometric A/D conversion with fully differential X and Y inputs is accomplished with the circuit in Figure 14. Here, one set of inputs, V_X, is connected in a manner similar to the circuit in Figure 13, and the other set of inputs V_Y, is connected in a multiplying fashion. Operation is as follows: I_{REF} for both the upper and the lower DAC-08 is modulated between 1mA and 3mA; and the resulting output currents are differentially transformed into voltages by the 5kΩ resistors a the comparator's inputs and compared with the V_X differer tial input. When the conversion process is complete (com parator inputs differentially nulled to less than 1/2 LSB) digital output is available which corresponds to the quotier of V_X/V_Y. Thus, four-quadrant ratiometric A/D conversion i achieved with four ICs and without instrumentatio amplifiers.

BRIDGE TRANSDUCER NULL

In many control systems, bridges must be nulled, and digital representation of the bridge's error must be provide for computer monitoring and control. The circuit in Figur 15 accomplishes both tasks by using the DAC-08 com plementary/differential current outputs to null the bridg with the DAC-08 connected in a tracking differential A/I converter configuration. The REF-02 reference voltag source provides both the bridge excitation voltage and th positive reference voltage for the DAC-08. Some of the ac vantages of this circuit are listed at the bottom of Figure 1

POWER MONITOR

Another differential current-input ADC is shown in Figure 1 with a transformer-coupled input. An up/down counter, precision high-speed comparator, and the DAC-08 form tracking A/D converter which continuously monitors th analog input. Two precautions must be observed: the com mon mode voltage at the comparator's inputs must not e

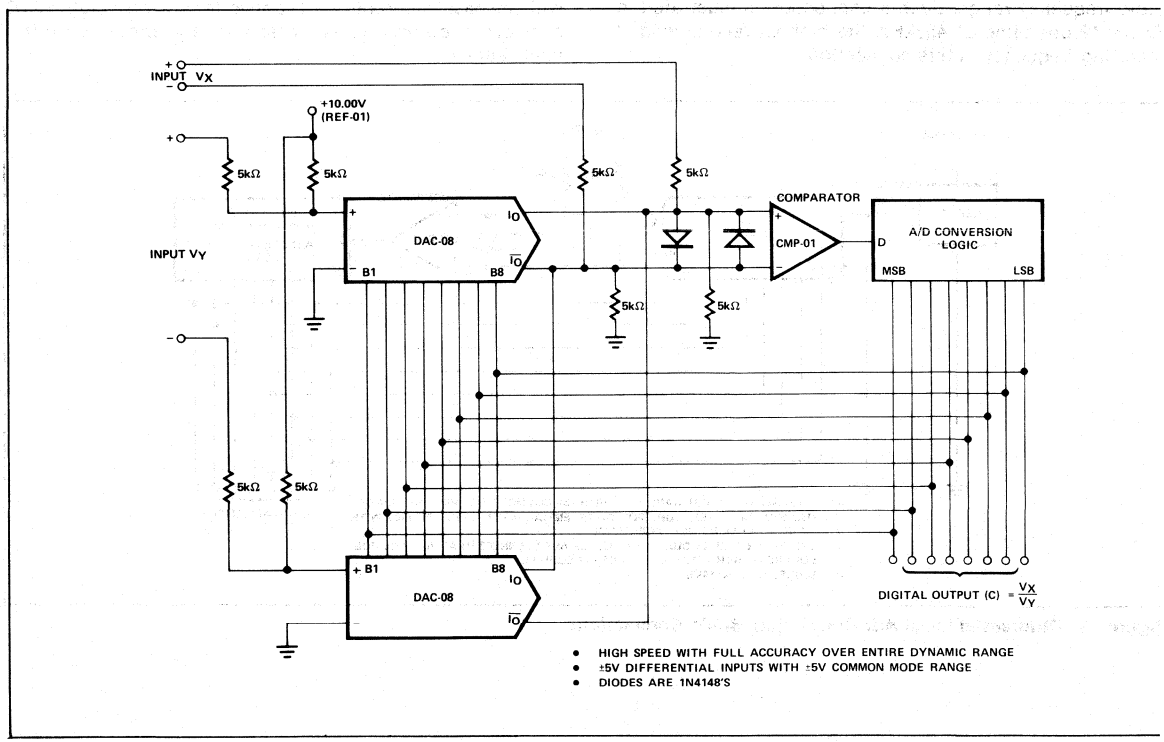

Figure 14. Four-Quadrant Ratiometric A/D Conversion Basic Connections

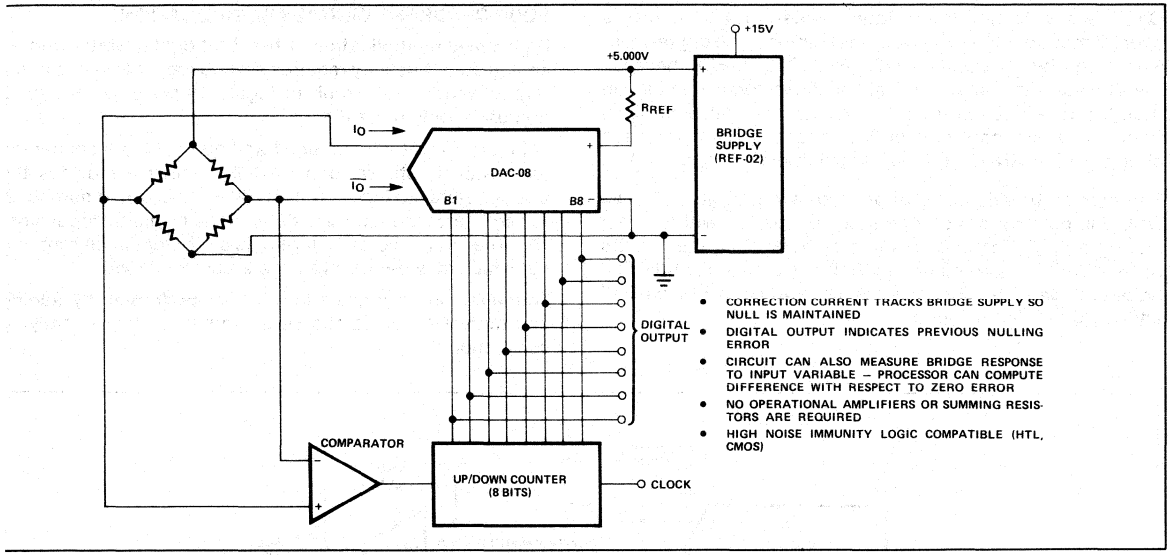

Figure 15. Bridge Transducer Null

8

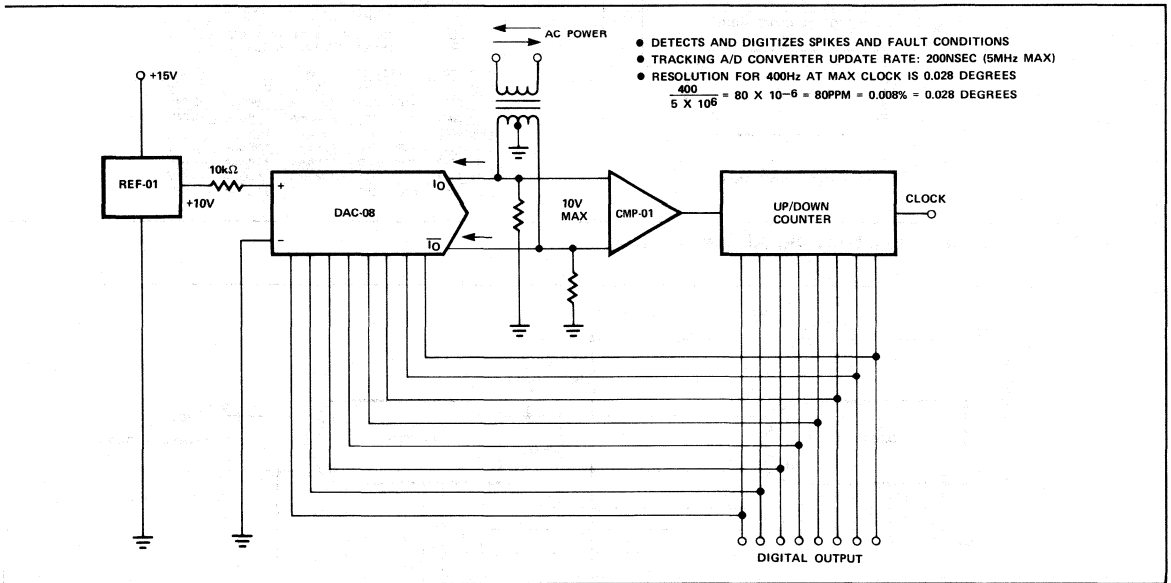

Figure 16. Power Fault Monitor and Detector

ceed ±10V; and the differential voltage must not exceed 11V. Voltage-limiting resistors at the comparator's inputs are recommended.

ALGEBRAIC DIGITAL COMPUTATION

Frequently, a digital arithmetic operation (addition, subtraction, multiplication, or division) must be performed, and an analog output must be provided. Traditionally, the arithmetic operations are performed with several ICs, and the output drives a D/A converter. This section decribes applications of the DAC-08 as an arithmetic building block, new design approaches that reduce the number of packages required in many applications. Today's low cost, versatile DACs merit a designer's consideration as arithmetic elements.

One benefit is not immediately apparent and deserves special mention. In all of these applications, the digital input words can be CMOS, TTL, DTL, NMOS, or MECL, because the DAC-08 interfaces with all of those logic families. In fact, the two input words may even be from different logic families to eliminate special level translators or interface circuitry. (See AN-17 "DAC-08 Applications Collection.")

The first arithmetic application is shown in Figure 17. Two DAC-08s perform a fast algebraic summation with a direct analog output. The circuit works by paralleling the outputs of two DAC-08s and summing their currents while driving a balanced load. The output is the algebraic sum of word "A" and word "B" in all four quadrants.

FOUR-QUADRANT DIGITAL MULTIPLICATION

High-speed multiplication of two 8-bit digital words with ar analog output usually requires several logic packages and a D/A converter. The circuit in Figure 18 performs this func tion using only three ICs.

In Figure 18 DAC-08 number 1 and number 2 are connectec as previously shown, and DAC-08 number 3 provides the analog reference inputs to DAC-08 number 1 and number 2 Those reference inputs are determined by digital input worc "A." The circuit's output, I_{01}-I_{02}, is a differential current out put which may be used to drive a balanced load.

Four-quadrant multiplication is thus performed by addinç one more DAC-08 to the basic four-quadrant multiplyinç connection.

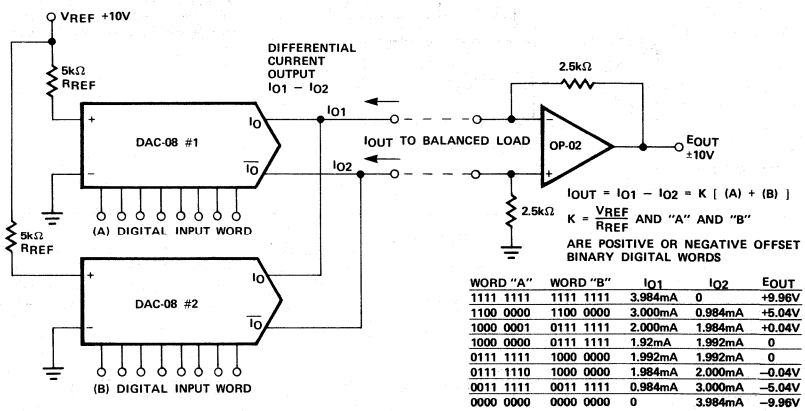

WORD "A"	WORD "B"	I_{O1}	I_{O2}	E_{OUT}
1111 1111	1111 1111	3.984mA	0	+9.96V
1100 0000	1100 0000	3.000mA	0.984mA	+5.04V
1000 0001	0111 1111	2.000mA	1.984mA	+0.04V
1000 0000	0111 1111	1.92mA	1.992mA	0
0111 1111	1000 0000	1.992mA	1.992mA	0
0111 1110	1000 0000	1.984mA	2.000mA	−0.04V
0011 1111	0011 1111	0.984mA	3.000mA	−5.04V
0000 0000	0000 0000	0	3.984mA	−9.96V

Figure 17. Four-Quadrant Algebraic Digital Computation

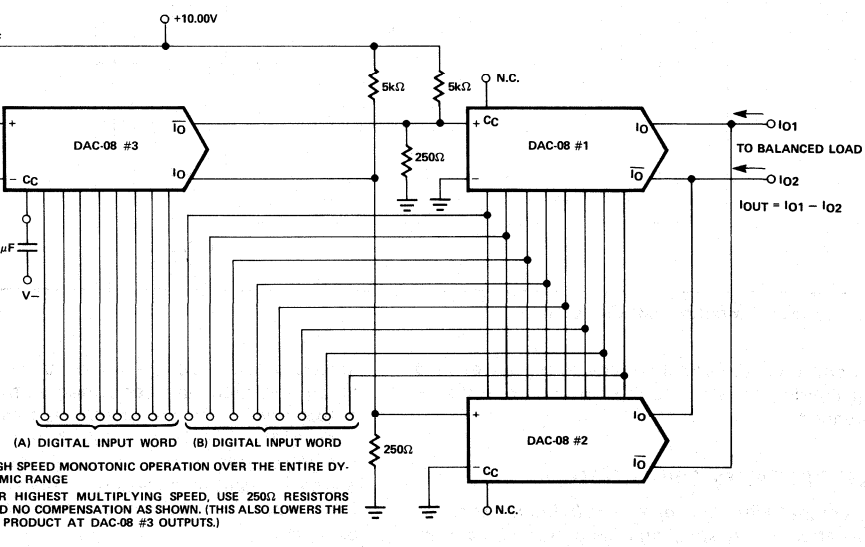

Figure 18. Four-Quadrant 8-Bit × 8-Bit Digital Multiplier

ANALOG DEVICES

AN-343
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Circuit Ideas for IC Converters

by Walter Jung,[1] Jeff Riskin, and Lew Counts

The alert circuit designer is constantly on the lookout for new devices and new ways to use existing devices to realize needed functions more efficiently, at lower cost, or in ways not previously practical. Some recently introduced analog-digital-conversion integrated circuits fit in this class. In these pages, we offer a few circuit ideas, either for direct application, or to stimulate the Reader's thinking about related possibilities.

Included in the discussion are such devices* as the AD537 V/f converter, the AD1408, AD561, and AD7520 d/a converters, and the AD581 reference, as well as some older devices, in a variety of circuits suited for instrumentation, data-acquisition, and process-control applications. All these ideas are workable; a few of them are ready to hook up and use as they stand. The others are useful to illustrate concepts and are ripe for adaptation and further modification for specific applications.

OHMS-TO-FREQUENCY CONVERTER

Ohms-to-volts conversion is a familiar property of many digital voltmeters. However, ohms-to-*frequency* conversion provides added flexibility, since it facilitates remote measurements, averaged measurements, and optional a/d or f/V conversion at the destination.

In the circuit of Figure 1, the 1V reference voltage available at the AD537 is unloaded by buffer amplifier A1, which drives a reference current into the resistor under test in the feedback circuit of amplifier A2. The output voltage, proportional to resistance, develops a current at the input of the V/f converter, which generates a square wave at a frequency proportional to current, and hence to R_X. Since the reference for the measurement is the same as the reference for the conversion, ratiometric operation minimizes the effects of variation of the AD537's reference with temperature.

A counter can be used to read resistance directly. Typical laboratory counters have more than adequate resolution; models with adjustable gate time permit the decimal place to be located as appropriate for the resistance range being measured. For example, a gate time of 1s will provide a readout in Hz, and the central measurement range will provide a direct readout, 1Ω/Hz, or 1kΩ/kHz, up to the 100kHz full-scale range.

In this application, we are taking advantage of the typically wide dynamic range of V/f conversion to provide a readout of the most-frequently used resistance values on a single range. After calibration at 100kHz full-scale, with a 100kΩ standard, (R2 is adjusted), and at 100Hz low-scale, with a 100Ω standard, (R4 is adjusted), the linearity error will typically be no more than ±0.06%.

R_S and R1-R2 should be stable precision types, and C_t an NPO ceramic, for best stability and repeatability. C1 and C2 serve as noise bypasses, but C2 should have low leakage (polystyrene), since it is effectively in parallel with R_X.

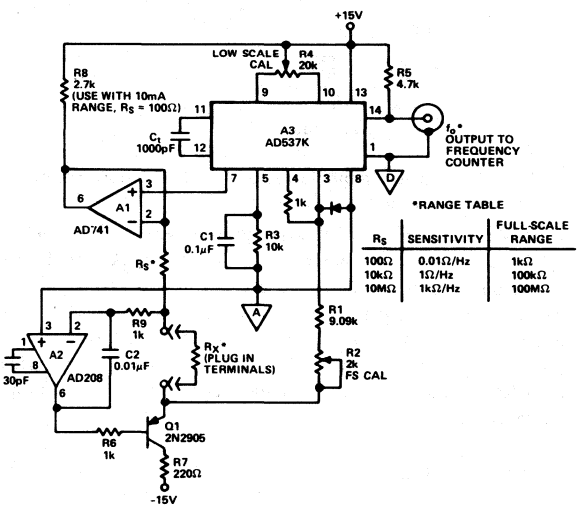

Figure 1. Ohms-to-Frequency Converter

As the chart notes, two additional ranges are suggested. The 0.01Ω/Hz range has greater resolution and accuracy, for R_X < 1kΩ, with a 1kΩ (= 100kHz) full-scale limit. Resistances

[1]This article is adapted from portions of Walter Jung's *IC Converter Cookbook*, published by Howard W. Sams & Co., Indianapolis, Indiana (1978).

less than 0.1 ohm can be resolved on this scale. A pullup resistor, R8, should be used on this range, to minimize loading on A1, since R_S (= 100Ω) will draw 10mA.

The highest scale range (1kΩ/Hz) allows resistances in the tens of megohms to be read. A low-bias-current amplifier, such as the AD208 (or the AD517, or a FET-input amplifier) should be used to minimize errors due to the flow of bias current in R_S.

ALGEBRAIC MANIPULATIONS – QUOTIENTS OF DIGITAL INPUTS

Since d/a converters multiply analog inputs by digital numbers, devices that permit a wide range of analog variation can perform a variety of algebraic manipulations involving multiplication or division of analog and digital quantities.[2,3] An example of the technique can be seen in Figure 2, a circuit that produces an analog quotient of two digital words, multiplied by a constant or variable reference.

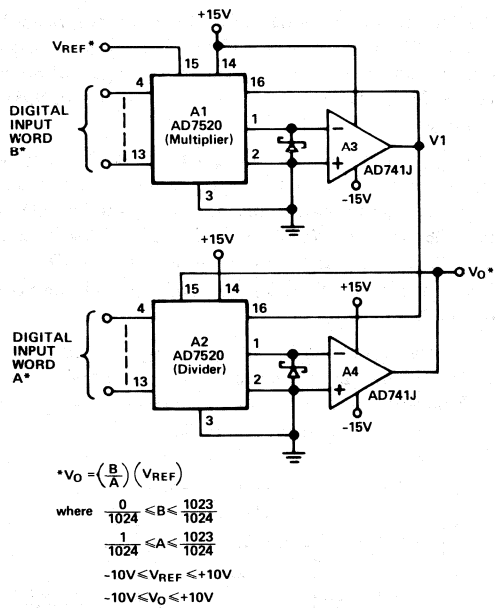

Figure 2. Algebraic manipulations – Analog quotient of two digital words.

In this circuit, two CMOS d/a converters are used. Converter A1 is connected in the forward path of op amp A3, producing an output, $V_1 = -B\ V_{REF}$, where B is the fractional binary value corresponding to the input code. Converter A2 is connected in the *feedback* path of op amp A4, producing an output, $V_O = -V_1/A$, where A is the fractional binary value associated with A2's input code. The overall relationship, therefore, is

$$V_O = \frac{B}{A} V_{REF} \qquad (1)$$

V_{REF} may be of any value in the range ±10V, B may be any number from 0 to 1023/1024, in steps of 1/1024, and A may be any such number from 1/1024 to 1023/1024. Naturally the ratio is limited to values for which the output, V_O, is within bounds. V_{REF} may be positive or negative, ac or dc, and the output will be of the same polarity.

Like analog division circuits, this circuit has an output error characteristic inversely proprotional to the denominator, A.

8-BIT-PROGRAMMABLE SQUARE-WAVE OSCILLATOR

Programmability is an important new degree of freedom in analog circuit and system design. Virtually any circuit parameter can be made digitally controllable with little difficulty using a/d and d/a conversion devices. It is important to be aware that "digitally controllable" doesn't necessarily mean that programmed circuits *must* interface with computers processors, or even digital systems. In many cases, the digital input can be provided by manually operated switches, which need not be fancy, since they need only to switch binary levels This circuit and those that follow illustrate a variety of practical examples of programmable circuits.

Figure 3 shows an 8-bit (255-frequency) programmable oscillator with square-wave output. The circuit comprises a current output d/a converter (AD1408 family) and a current-to frequency converter (AD537 family). The digital input produces a linearly related current from the DAC; this current driven directly to the input of the VFC, produces a square wave that has a frequency proportional to the numerical value of the digital input word.

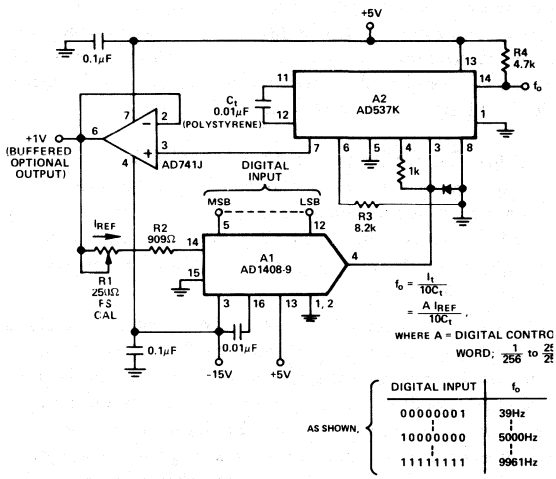

Figure 3. 8-bit programmable oscillator, square-wave output.

The AD1408-9 (9-bit-linearity) DAC is scaled for 1mA full scale current output, to match the 1mA full-scale input of the AD537K. The 1mA reference current for the DAC is derived from the 1V reference output of the AD537, buffered by the AD741 follower-connected op amp. Since the basic reference source is common to both devices, errors due to its drift tend to cancel out.

[2] *Analog-Digital Conversion Handbook*, D.H. Sheingold, ed. [3rd edition (1986), published by Prentice Hall. Available from Analog Devices, Inc., Norwood MA, 02062 P.O. Box 796.]

[3] "Application Ideas for Multiplying DACs," *Analog Dialogue* 12-1, 1978.

polystyrene capacitor is used for C_t, and its tempco is compensated for by loading the AD537's V_T output with R3[4]. R3 can be adjusted to trim the overall system tempco. The circuit, as shown, has a nominal full-scale frequency of 10kHz (9961Hz for all-1's), with C_t = 0.01μF. Worst-case nonlinearity of the specified DAC-VFC combination is 0.16%. The output is a TTL-compatible square wave.

PROGRAMMABLE GAIN TRIMMING OR CALIBRATION

It is usual, in the design of devices such as converters or amplifiers, to concentrate design attention on linearity, since gain and offsets are considered to be reducible errors. Nevertheless, in high-precision applications, the gain must eventually be calibrated. D/A converters are improved substitutes for potentiometers, if the gain of the device-to-be-calibrated is set at a value near the nominal value, and the programmed converter provides the difference. Calibration can be performed automatically, under software control, with the required incremental value retained in a counter or latched; or it can be performed manually, using a thumbwheel switch.

In the circuit of Figure 4a, an AD1408-9 (256 adjustment steps) provides the incremental adjustment range for the scale factor of an AD562 12-bit DAC. When pin 5 of the AD562 is connected to a 10V reference, the gain will be 0.2% high. In this circuit, a programmable 0 to -2μA current applied to the summing point will provide a ±0.2% range of gain change in 15.6ppm/LSB increments (1/16 of an LSB in the AD562).

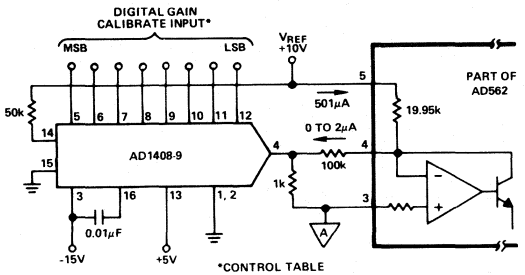

(a) D/A gain calibration.

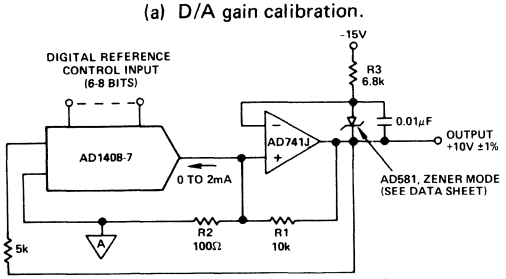

(b) Direct reference voltage calibration.

Figure 4. Gain calibration methods.

The performance of the components used to achieve this function is not highly critical, since their contribution to overall gain error is reduced by their small weighting. The use of this scheme with an AD562 DAC is a simple example, but it is applicable wherever automatic calibration to high absolute accuracy is required.[5] Coarser steps (fewer bits) could have been used (6 bits of a 1408-7) if appropriate.

In Figure 4b, a related scheme is used to calibrate the output of a buffered reference circuit. The basic reference is an AD581 10V bandgap reference, connected as a 2-terminal "Zener diode", in the feedback path of an op amp. The 1% positive feedback increases the output voltage to 10.1V, and the 2mA full-scale output from the AD1408-7 DAC, flowing in the 100Ω resistor, can reduce the output voltage to about 9.9V. Thus, the adjustable range is 10V ±0.1V, in increments of about 780μV/bit, for 8-bit control.

Amplifier gain can also be trimmed by using a DAC to set incremental gain values in the neighborhood of nominal gain. A typical scheme for programming inverting-amplifier gain would employ a CMOS DAC, with its input attenuated, in shunt with the input resistor of an inverting operational amplifier.

PROGRAMMABLE OFFSET

A programmed constant offset (or offset-zeroing voltage) can be introduced at the *reference* input of an instrumentation amplifier, to provide an output offset, independent of gain. Figure 5 shows how an AD521 instrumentation amplifier might operate in conjunction with an AD561 10-bit d/a converter. In this case, the nominal full-scale output range of the AD561 is ±1V, when loaded by 2.5kΩ. Larger offset ranges than ±1.67V would be available by using a follower-with-gain between the DAC output and the amplifier's *reference* input, or by providing a portion of the AD521 gain via *sense* feedback,[5] the offset would be amplified by the same amount. Smaller offset voltages are obtained by simply reducing R_X.

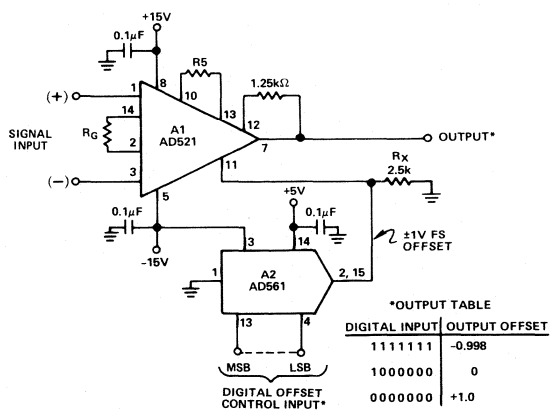

Figure 5. Programmable offset instrumentation amplifier.

4-20mA CURRENT CONTROLLER

A common requirement in industry is for transmission of analog data in the form of a 4-20mA current, to minimize the effects of ground-potential differences, series resistance, and voltage-noise pickup. 4mA corresponds to zero, 20mA to full scale.

[5] For another example, see the AD572 12-bit ADC data sheet, Figure 11.

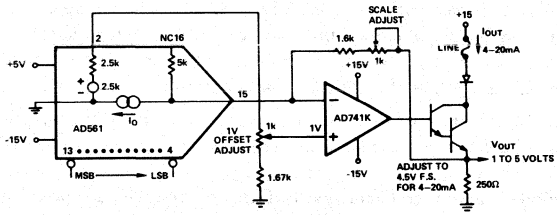

Figure 6. Process control current source.

Figure 6 shows a circuit to accomplish this with 10-bit resolu tion. An AD561 is used, in conjunction with an op amp and a Darlington transistor. With an all-0's digital input, the 1kΩ *offset* pot is adjusted for 4mA of output current. With all 1's the *scale-adjust* pot is set for 20mA (or 19.98mA) of out put current.

Although the load is shown here as being referred to a +15V supply, it may –in general– be returned to any positive voltage within the breakdown rating of the transistor used. The diode protects against reverse-polarity faults, the fuse against shorts.

)NE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

4-20 mA Digital to Process Current Transmitter

his application note describes a three IC, 4-20mA process urrent, digital-to-analog converter that can be constructed t low cost. It operates from a −5V ±1V negative power supply nd a +23V ±7V positive power supply, has 24V output vol- ıge compliance, and occupies less than 4 square inches of rinted circuit board space. Other significant features include TL logic input compatibility, 8-bit binary coding, 0° to 70°C operation, and 5μs full scale settling time into a 500Ω oad.

HEORY OF OPERATION

, fixed current of 0.5mA is added to a DAC's output current arying between 0 and 2.0mA and the resulting total current s multiplied by a factor of 8 to produce an output current of .0 to 20mA.

ı the schematic, first note the REF-01CJ, a +10V ad- ıstable reference. Its output goes to the noninverting input f one half of A3, a dual precision op amp. The inverting in- ut is within a feedback loop forcing +10V to appear at the op of R4, a 20kΩ resistor; a 0.5mA current will flow in R4 through Q1, a high h_{FE} transistor. The same +10V is applied to R3, the reference input resistor of a multiplying IC D/A converter, the DAC-08. Full scale output current of the DAC will be the difference in voltage between the +10V reference and Pin 14 of the DAC divided by R3; Pin 15 will be at the same voltage as Pin 14 because it is a high impedance point, the noninverting input of an op amp internal to the DAC. After calibration a current of 0 to 2mA (depending on the digital input code) will flow into the DAC's output, Pin 4.

Both the DAC's output current and the fixed 0.5mA flow in R5, a 800Ω precision resistor. The voltage developed by that current is applied to the noninverting input of the other half of A3 and will also appear across R6, a 100Ω precision resistor. Thus, eight times the 0.5 to 2.5mA current in R5 flows in R6, or 4 to 20mA. Almost all of this current appears at the output because the 2N6053 is a high h_{FE} device, a power darlington transistor.

Some other components need explanation. C1 provides frequency compensation of the DAC's reference amplifier; C2

CHEMATIC DIAGRAM

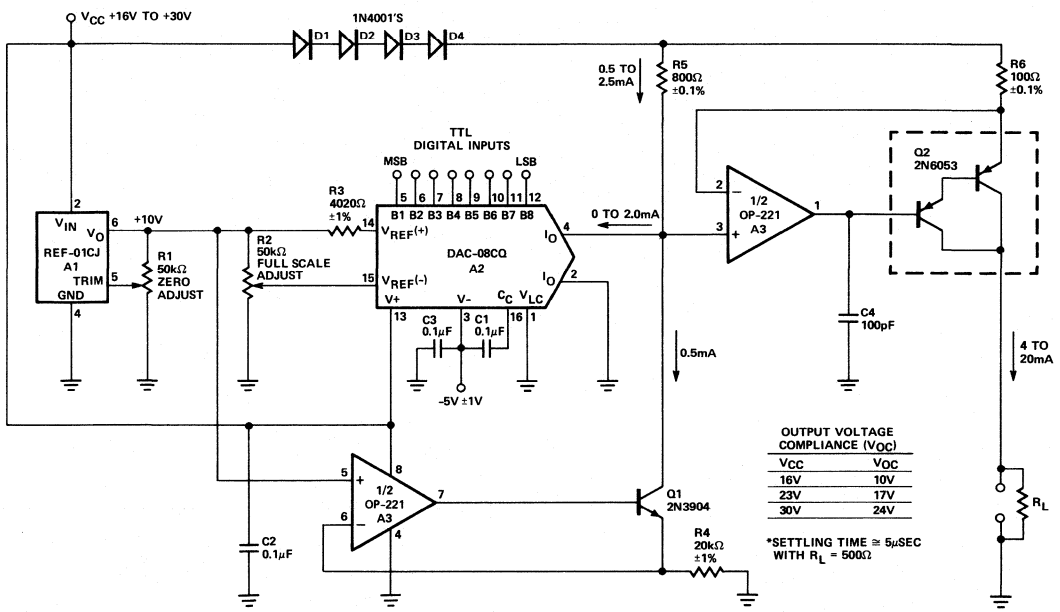

and C3 are power supply decoupling (bypass) capacitors; C4 prevents high frequency oscillations. D1 through D4 insure at least 2.5V differential between the op amp's inputs and its positive power supply under all conditions. R1 and R2 are zero scale and full scale adjustments respectively.

CALIBRATION PROCEDURE

Apply +23V ±7V and −5V ±1V to the converter with a current-measuring meter connected between the output and ground. Make the digital inputs all zeros, < +0.8V. Adjust R1 until the output current is 4.0mA. Now change the digital inputs to all ones, > +2.0V. Adjust R2 until the output current is 20mA. Calibration is now completed.

OUTPUT VOLTAGE COMPLIANCE

Output voltage compliance is V_{CC} −6V. For example, at V_{CC} = +16V, the output may go to a maximum of +10V without affecting output current. Thus, a 500Ω resistor would be the maximum load resistor at V_{CC} = +30V, V_{OC} = 24V, and R_L Maximum = 1.2kΩ.

SCALE MODIFICATION

Although the values shown are for the more common 4-20mA requirement, operation at 1-5mA or 10-50mA may be achieved by changing some components. For 10-50mA, change R6 to 40Ω; this makes the multiplying factor 20 instead of 8. For 1-5mA, replace the 2N6053 with a 2N5087, and change R6 to 400Ω.

CONCLUSION

A simple, low-cost process current converter has been shown with wide application in the controls industry. The design is tolerant of wide power supply variations, has hig voltage compliance, and is easily calibrated. Reliability ar cost are optimized by using only three integrated circuit the Precision Monolithics DAC-08, REF-01, and OP-14, plu a few readily available discrete components.

REFERENCE

Crowley, B., "Circuit Converts Voltages to 4-20mA For In dustrial Control Loops," **Electronic Design**, Jan. 5, 197(page 116.

PARTS LIST

Circuit Symbol(s)	Description
A1	+10V Reference, PMI REF-01CJ
A2	8-Bit DAC, PMI DAC-08CQ
A3	Dual Op Amp, PMI OP-221
C1-C3	0.1μF +80%/−20% 50V, Type CK-104
C4	100pF ±5% Mica, DM100ED101J03
D1-D4	Power Diode, 1N4001
Q1	NPN Transistor, 2N3904
Q2	PNP Power Darlington, Motorola 2N605
R1-R2	50kΩ Potentiometer, Bourns #3006P-1-503
R3	4020Ω ±1%, RN55C4021F
R4	20kΩ ±1#, RN55C2002F
R5	800Ω ±0.1%, GR#8E16D800
R6	100Ω ±0.1%, GR#8E16D100

ANALOG DEVICES

AN-23
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

D/A Converter Generates Hyperbolic Functions

Measurement and control systems frequently require fine resolution around a setpoint with wide dynamic ranging capability. This can be satisfied by systems designs which use a high resolution, strictly linear approach; but this is costly and often unnecessary. Nonlinear function fitting using multiplying digital-to-analog converters (DAC's) offers a desirable alternative by being both simpler and more cost-effective. This application note describes how extended range hyperbolic functions of the type A/X or —A/X (where "A" indicates an analog constant, while "X" represents a decimally-expressed digital divisor) are easily generated by just two lowcost I.C.'s; an operational amplifier and a multiplying DAC. Circuit configurations are provided for each polarity output along with dynamic performance photographs and general design guidelines for either binary or BCD-coded divisors.

THEORY OF OPERATION (A/X)

Figure 1a shows the A/X function circuit which uses a two-digit BCD-coded DAC, the DAC-20EX, and a decompensated, wide-bandwidth op-amp, the OP-17. A constant current, I constant, equal to the value of one least significant bit (LSB), flows into the DAC output terminal, I_O. Simultaneous adjustment of the scale factor and output amplifier offset voltage is enabled by a multi-turn, low tempco potentiometer, R5, which adjusts current $-I_R$ producing voltage $-V_R$ across R2. The LSB value (scale factor) equals $-V_R/R1$.

Zener diode, D_Z, provides a stable reference voltage source. Because feedback for the op amp is through the DAC, capacitors C1 and C2 are added to provide proper phase compensation. Reference resistor R3 is determined by the scale factor and the maximum current allowed into the DAC reference input V_R+. Bias

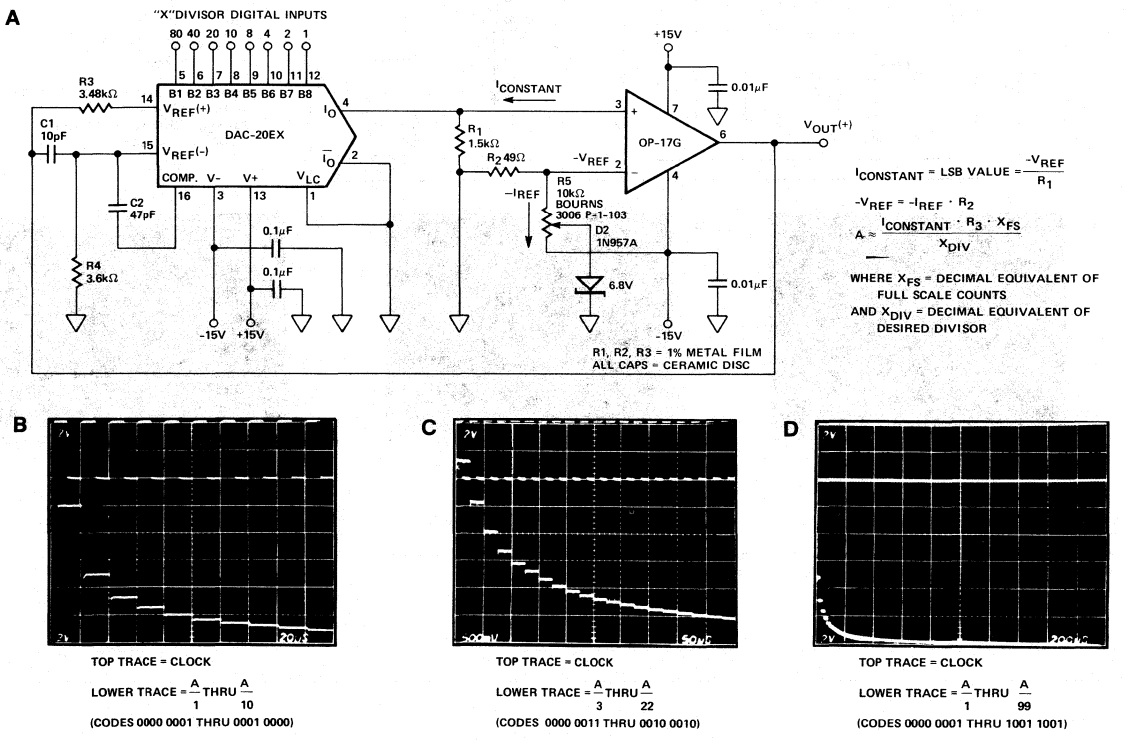

Figure 1. A/X Function Generator

current compensation for the DAC reference amplifier is accomplished by R4.

Figures 1b, 1c, and 1d show dynamic performance of circuit 1a when the digital inputs are swept by an external BCD up-counter with codes of 0000 0001 through 1001 1001 (division by zero is not allowed).

THEORY OF OPERATION (—A/X)

The circuit configuration for the —A/X function is shown in Figure 2a. It is quite similar to that of Figure 1a with both the DAC reference amplifier and output amplifier terminals reversed. Capacitors C1 and C2 provide phase compensation. Figures 2b, 2c, and 2d show dynamic performance of circuit 2a.

DESIGN CONSIDERATIONS

1. Circuit speed and settling time are dictated by output op amp slew rate, scale factor, and compensation. Use of slower amplifiers considerably increases the illustrated settling times. Effective slew rate of circuit 1a is 3V/μs, wh circuit 2a slews 0.6V/μs.
2. Layout and breadboarding of high gain, wide-bandwic devices necessitates considerable care with a ground pla with single point grounding being highly desirable. D coupling capacitors located close to the devices' supp inputs are essential.
3. Accuracy of the circuit is within 1% over the 0°C +70°C temperature range with 1% metal film resisto R1, R2 and R3. DAC linearity becomes an importa factor as the divisor decreases; for this reason 1/4 LS linear DACs are recommended.
4. Binary coding may be accomplished by substituting an 8- binary-coded DAC-08EX for the two-digit BCD-coded DA 20EX. In addition to adjusting circuit values however, higher performance op amp such as the OP-17F is desirat because the output amplifier's input offset voltage di becomes a more significant error source for overall sca factor stability over temperature. This is due to the increas resolution of the binary coding.

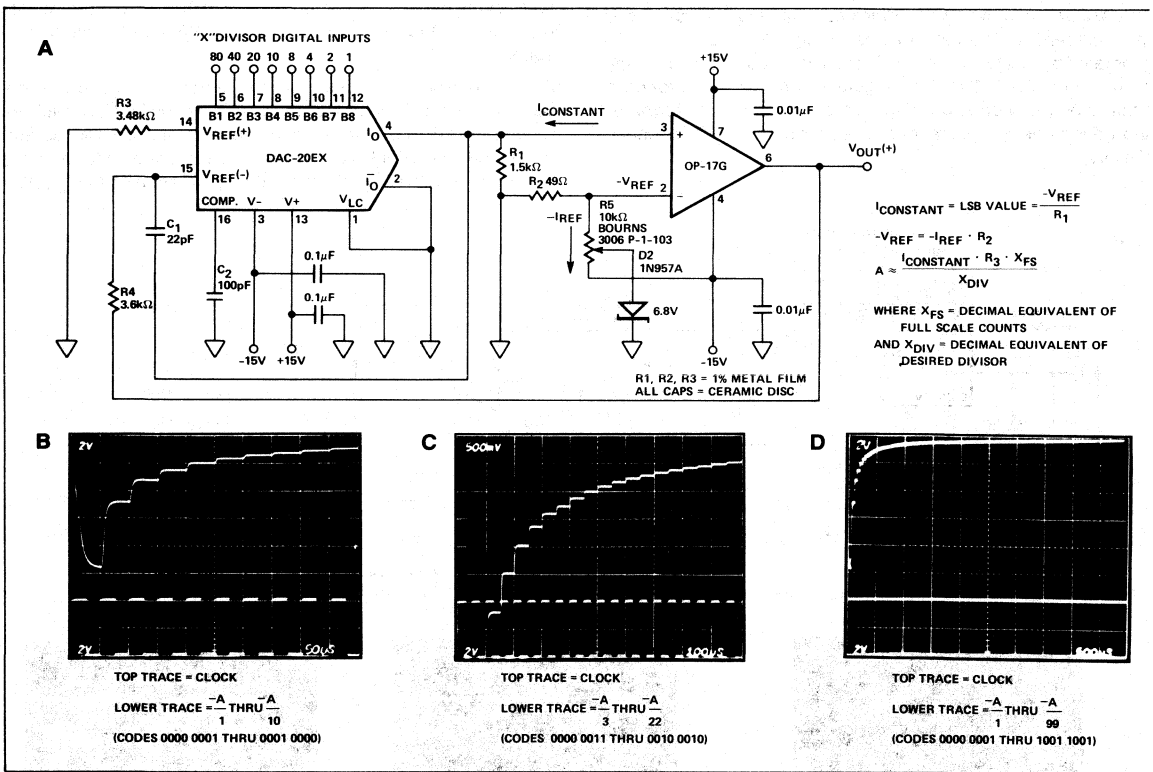

Figure 2. –A/X Function Generator

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Understanding LOGDACs

INTRODUCTION TO THE ANTILOG D/A CONVERTER

Analog Devices' AD7100 Series LOGDACs are CMOS multiplying DACs characterized by an exponential (antilogarithmic) digital-to-analog transfer function.

Perhaps the easiest way to visualize what a LOGDAC® does is to compare it to two well-known circuits—the classic 3-terminal potentiometer and a CMOS multiplying DAC (digitally controlled potentiometer). As shown in Figure 1a through Figure 1c, the transfer function of the three circuits is of the form:

$$V_{OUT} = \alpha V_{IN} \qquad \text{EQN1}$$

WHERE:
α = attenuation factor
$0 \leq \alpha \leq 1$

In each case shown in Figure 1, α is a dimensionless number which can range from 0 (maximum attenuation) to approximately 1 (minimum attenuation). Additionally, each circuit has an analog input (V_{IN}), an analog output (V_{OUT}) and a mechanism for controlling the attenuation factor α.

The above in conjunction with Figure 1a through Figure 1c illustrates the similarity of the pot, M-DAC and LOGDAC functions. How, then does the LOGDAC *DIFFER* from the linear M-DAC?

The answer is resolution. The basic differentiating feature of the LOGDAC versus the linear multiplying DAC is the way resolution is specified.

RESOLUTION OF LOGDACs VERSUS LINEAR DACs

From Figure 1b, the attenuation factor α of a linear DAC is:

$$\alpha = \left(\frac{N}{2^n}\right) \qquad \text{EQN2}$$

WHERE:
n = Number of digital input bits to the DAC
N = Integer value of DACs digital input

NOTE: Many treatments of multiplying DACs label the attenuation factor "D" where the digital input "D" is:

$$D = \alpha = \left(\frac{N}{2^n}\right) = \frac{\text{Bit 1}}{2^1} + \frac{\text{Bit 2}}{2^2} + \ldots \frac{\text{Bit n}}{2^n}$$

WHERE: Bit 1 through Bit n = 1 or 0
n = Number of bits

Since N was postulated to be an integer, the smallest possible change of N is plus or minus one count. The resolution of α is, therefore, ± 1 part in 2^n or 1 part in full scale. Important to realize is that the voltage resolution of a linear DAC is the same (barring differential nonlinearity effects) at all points on its transfer function.

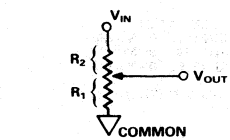

$V_{OUT} = \left(\frac{R_1}{R_1 + R_2}\right)(V_{IN})$

FORM: $V_{OUT} = \alpha V_{IN}$

WHERE: α = ATTENUATION FACTOR = $\left(\frac{R_1}{R_1 + R_2}\right)$

Figure 1a. Three Terminal Pot

The resolution of a LOGDAC is different, however. The following discussion shall endeavor to show that a LOGDAC's voltage resolution is different at all points on its transfer function, i.e., barring differential nonlinearity effects, the resolution expressed as a *percent of reading* (not percent of full scale) is constant throughout the LOGDACs range.

LOGDAC is a registered trademark of Analog Devices, Inc.

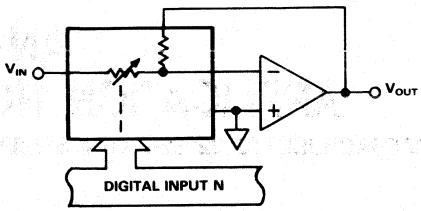

$V_{OUT} = -\left(\frac{N}{2^n}\right)(V_{IN})$

FORM:
$V_{OUT} = -\alpha V_{IN}$

WHERE:
α = ATTENUATION FACTOR = $\left(\frac{N}{2^n}\right)$
n = NUMBER OF DIGITAL INPUT BITS
N = DIGITAL INPUT ($0 \leq N \leq 2^n - 1$)

EXAMPLE: FOR 8-BIT DAC
$n = 8$, $2^n = 256$
$0 \leq N \leq 255$

$0 \leq \alpha \leq \left(\frac{255}{256}\right)$

Figure 1b. Linear Multiplying DAC (Digitally Controlled Pot)

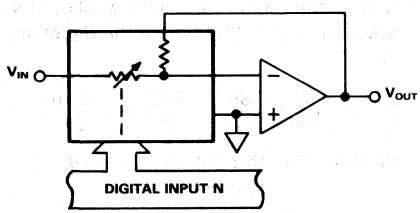

$V_{OUT} = -V_{IN}(10)^{-0.05rN}$

FORM:
$V_{OUT} = -\alpha V_{IN}$

WHERE:

α = ATTENUATION FACTOR = $10^{0.05rN}$

N = DIGITAL INPUT
FOR AD7111: $0 \leq N \leq 239$
FOR AD7118: $0 \leq N \leq 59$
FOR AD7115: $0 \leq N \leq 199$

r = LOGDAC RESOLUTION IN dB

Figure 1c. LOGDAC (Digitally Controlled Pot)

From Figure 1c, the LOGDACs attenuation factor α is an exponential function (antilog) of the basic form:

$$y = a^{-x} \qquad \text{EQN3}$$

If base number 10 is chosen (other base numbers can be used, incidentally), the attenuation factor α for the LOGDAC of Figure 1c becomes:

$$\alpha = 10^{-\left(\frac{rN}{20}\right)} \qquad \text{EQN4}$$

WHERE:
N = Integer value of DACs digital input
for AD7115: $0 \leq N \leq 199$
for AD7111: $0 \leq N \leq 239$
for AD7118: $0 \leq N \leq 59$
r = LOGDAC resolution in dB
for AD7115: r = 0.1
for AD7111: r = 0.375
for AD7118: r = 1.5

Taking the LOG of both sides of EQN4 gives:

$$\log_{10}\alpha = -\left(\frac{rN}{20}\right)$$

$$20\log_{10}\alpha = -rN$$

$$\therefore \quad \alpha_{dB} = -rN \qquad \text{EQN5}$$

From EQN5, it is readily apparent that a plus one count change of N causes an attendant $-(r)$dB change in the attenuation factor α (and thus also a $-(r)$dB change in the DAC's output voltage).

Figure 2 is a graph of the general LOGDAC transfer function for the circuit of Figure 1c. It shows quite simply that increasing the digital input N causes a decrease in the output voltage V_{OUT}. Additionally, it shows the nonlinear relationship of V_{OUT} relative to N. Figure 3 is an expanded section of the transfer function shown in Figure 2. It illustrates the fact that the ratio of any two adjacent LOGDAC output voltage levels is the same throughout the transfer function. To further amplify the significance of this point, consider that a + one count change in the digital input N causes the output voltage to decrease in amplitude by a *fixed ratio relative to where it was before the change.* (At all points on its transfer function . . .) Thus, we have a DAC with *percent of reading* resolution as opposed to the linear DAC which defines resolution in terms of *percent of full scale.*

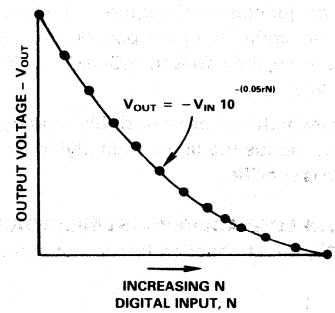

Figure 2. LOGDAC D/A Transfer Characteristic

To summarize, a linear DAC's voltage resolution is fixed throughout its transfer function. However, the LOGDAC

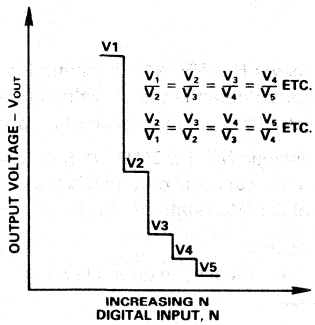

Figure 3. Expanded LOGDAC Transfer Function Illustrating the Concept of % of Reading Resolution

xhibits a continuously variable output voltage resolution ıroughout its transfer function range. The LOGDAC's oltage resolution is coarsest at or near full scale (0dB) nd finest at or near 0 scale (mute). Table I shows the equialent percent of reading resolution for various Analog evices LOGDACs.

Model	dB Resolution ($\Delta N = \pm 1$ Count)	% of Reading Resolution ($\Delta N = +1$ Count)	($\Delta N = -1$ Count)
AD7118	±1.5dB	−15.9%	+18.9%
AD7111	±0.375dB	−4.2%	+4.4%
AD7115	±0.1dB	−1.1%	+1.2%

Table I.

ASIC CIRCUIT CONFIGURATIONS

NTILOG DAC (Exponential with Negative Exponent)

ıe circuit of Figure 4 generates output voltage levels as ɜtermined by the equations:

$$V_{OUT} = -V_{IN}10^{-\left(\frac{rN}{20}\right)} \quad \text{EQN6}$$

and/or

$$V_{OUT} = -V_{IN}e^{-(0.11512rN)} \quad \text{EQN7}$$

WHERE:

r = LOGDAC resolution in dB
- for AD7118: r = 1.5
- for AD7111: r = 0.375
- for AD7115: r = 0.1

N = Integer equivalent of digital input
- for AD7118: $0 \le N \le 59$
- for AD7111: $0 \le N \le 239$
- for AD7115: $0 \le N \le 199$

V_{IN} = ac or dc input voltage (nominal range ± 10V)

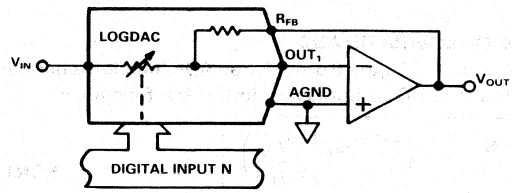

Figure 4. ANTILOG D/A Converter (Negative Exponent)

Features of the circuit of Figure 4 include:

1. It provides dB attenuation of V_{OUT} relative to V_{IN} as determined by the digital word N. (i.e., output range is 0dB to − dB)
2. The circuit provides % of *reading* resolution.
3. The analog input can be voltage or current, ac or dc, positive or negative polarity – i.e., the circuit is basically a CMOS multiplying DAC.

ANTILOG DAC (Exponential with Positive Exponent)

The circuit of Figure 5 is analogous to a multiplying DAC divider circuit. It provides signal gain of V_{OUT} relative to V_{IN} as determined by the equations:

$$V_{OUT} = -V_{IN}10^{+\left(\frac{rN}{20}\right)} \quad \text{EQN8}$$

and/or

$$V_{OUT} = -V_{IN}e^{+(0.11512rN)} \quad \text{EQN9}$$

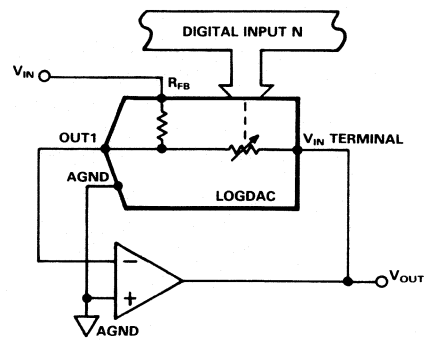

Figure 5. ANTILOG D/A Converter (Positive Exponent)

Basically, the analog input or reference voltage is applied to the on chip feedback resistor (RFB) and the amplifier output is connected to the V_{IN} terminal of the LOGDAC. The LOGDAC then ends up in the amplifier's feedback loop, thus the circuit provides *dB* gain of V_{OUT} relative to V_{IN} as determined by the digital input N (i.e., V_{OUT} range is 0dB to positive dB). As does the negative exponential DAC of Figure 4, this circuit provides % of *reading* resolution.

LOG OR LOG RATIO ADC

The circuit of Figure 6 provides an ADC function while performing a LOG compression. Its transfer function is:

$$N = \left(\frac{1}{-r}\right)\left(20\log_{10}\left|\frac{-V_{IN}}{V_{REF}}\right|\right) \qquad \text{EQN10}$$

OR

$$N = \left(\frac{8.68659}{-r}\right)\left(\ln\left|\frac{-V_{IN}}{V_{REF}}\right|\right) \qquad \text{EQN11}$$

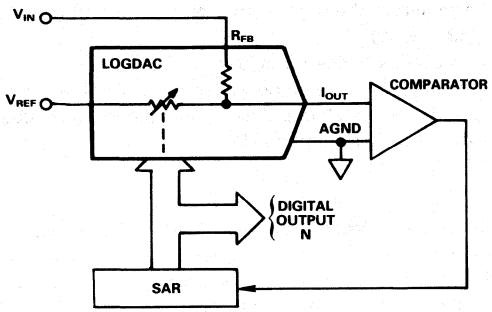

Figure 6. Log or Log Ratio A/D Converter

If the digital answer N (of EQN10) is multiplied by $-r$, the numerical value obtained is the dB value of the absolute value of V_{IN} relative to V_{REF} (answer 0dB to $-$dB).

If the digital answer N (of EQN11) is multiplied by $-r/8.68659$, the numerical value obtained is the natural log of the absolute value of V_{IN} relative to V_{REF}.

Circuit Constraints:

1. V_{IN} and V_{REF} must be of opposite polarity
2. $V_{IN} \leq |V_{REF}|$

ANALOG DEVICES

AN-137
APPLICATION NOTE

)NE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

A Digitally Programmable Gain and Attenuation Amplifier Design

by James Wong

8

ly adding two resistors to the output amplifier feedback loop of current output D/A converter, you can get gain control in ddition to attenuation control. Figure 1 shows a complete igitally programmable amplifier that is capable of producing ain as well as attenuation in the range of 1/64 to 64. The circuit erives its range by using a 12-bit CMOS D/A converter.

he design is based on the fact that the transfer function from ıe input to the output of the D/A converter is purely voltage ttenuation. Connecting the two resistors R_1 and R_2 in a "T" onfiguration inside the feedback loop of the output amplifier roduces a voltage gain from the resistor junction to the output. R_1 is much smaller than R_{FB} (in this case R_{FB} is 11kΩ), then ıe gain produced is approximately equal to $1 + R_2/R_1$, or 64. he result is a programmable gain amplifier that has a transfer ınction of:

$$A_V = -\left(\frac{D}{4096}\right)(64),$$

/here D represents the binary weighted digital code of the D/A onverter.

)f course, the added gain of the T-network does increase the oise gain of the circuit. Be sure to choose a low noise amplifier ɔ begin with.

ly using a low noise, high-frequency op amp such as the OP-1, besides keeping noise level down, it gives the circuit a wide andwidth performance even at high-gain settings.

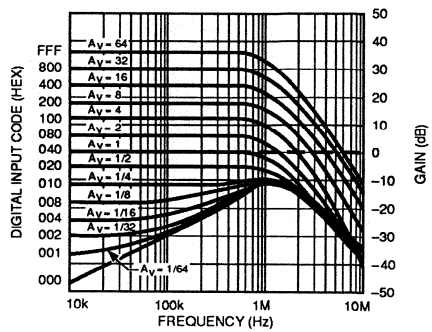

FIGURE 2: *Gain vs. Frequency Response for the Twelve Binary Gain Settings*

Figure 2 shows the frequency response of the programmable gain circuit at various gain settings. At high gains, the amplifier has 1MHz bandwidth. At gains below 1/4, the D/A converter's stray capacitance feedthrough limits the amplifier bandwidth, while still achieving 20kHz.

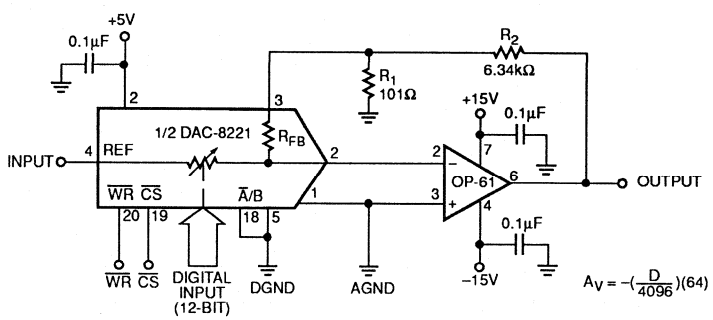

FIGURE 1: *Two Resistors R_1 and R_2 Add a Gain of 64 to the D/A Converter, Resulting in a Simple Digitally Programmable Gain Amplifier*

ANALOG DEVICES

AN-142
APPLICATION NOTE

)NE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Voltage Adjustment Applications of the DAC-8800 TrimDAC® An Octal, 8-Bit D/A Converter

by Joe Buxton

The DAC-8800, a monolithic octal 8-bit digital-to-analog converter, s a digitally-controlled voltage adjustment device. The DAC's lesign makes it ideal for replacing trimming potentiometers in nany applications. Not only does it replace potentiometers, but he DAC has many advantages over them, such as solid state eliability, very low drift over temperature and time, elimination of shifts due to vibrations, and automating the adjustment process.)uring manufacture of complex electrical systems, potentiometers must be manually adjusted taking considerable time and ost for labor, or expensive robotic systems must be developed or the same purpose. However, the DAC-8800 can automate the ystem's voltage adjustments so that a computer can now control he calibration.

This application note first describes the basic architecture and operating modes of the DAC-8800, including the reference input ange limits, the load that the DAC places on the references, single upply operation, and serial interfacing. The last half of this note hows many basic circuits for using the DAC in a wide variety of pplications, such as two wire interfaces and stand-alone operaion for systems not based on digital controllers. Also included are techniques for adjusting the offset of operational amplifiers; using two DAC outputs together for coarse and fine control of a voltage; digitally changing the gain of a voltage-controlled amplifier; and trimming voltage references.

BASIC ARCHITECTURE

As the functional diagram shows in Figure 1, the DAC-8800 has eight individual DACs divided into two groups of four, each group having its own high and low reference inputs. Each DAC's output is independently controlled by a serial interface through which the 8-bit data word and 3-bit address are loaded.

Each of the DACs contains an R-2R ladder connected between the high and low reference inputs as shown in Figure 2. The output voltage is set by the position of the switches according to the formula below:

$$V_{OUT}(D) = D \times (V_{REF}H - V_{REF}L)/256 + V_{REF}L$$

where D is the digital code.

As this equation shows, the output can vary from $V_{REF}L$ to $V_{REF}H$ in 256 steps. It is significant that, while the output voltage can vary over this range, the DAC-8800's output impedance is always equal to a constant R_{OUT}, the characteristic resistance of the ladder. R_{OUT} is typically 12kΩ but can vary between 8kΩ and 16kΩ from device to device. The DAC's accuracy depends not on the absolute value of the ladder resistors but rather on the relative resistor matching. Thus, variations in output impedance do not

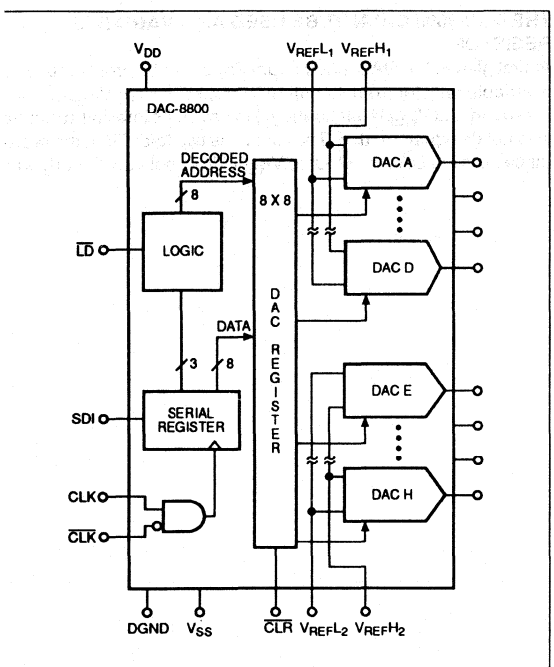

FIGURE 1: *DAC-8800 Block Diagram*

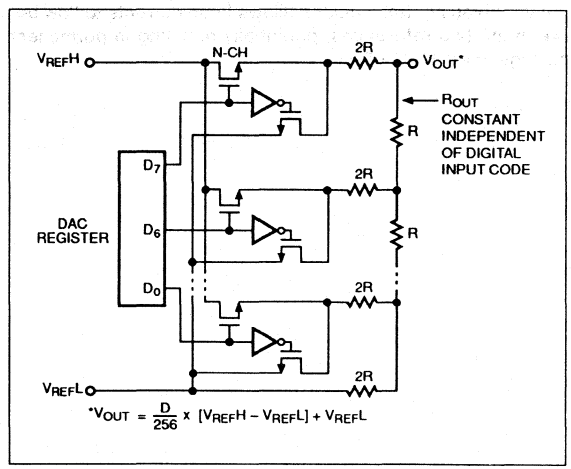

FIGURE 2: *DAC-8800 R-2R Ladder Network*

TrimDAC is a registered trademark of Analog Devices, Inc.

affect the linearity of the DAC. To easily understand the DAC, each output can be thought of as a Thevenin equivalent circuit of a voltage source in series with R_{OUT} as in Figure 3, where R_{OUT} is 12kΩ. The digital code then varies the voltage source between $V_{REF}L$ and $V_{REF}H$.

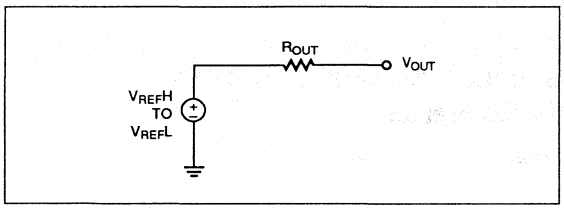

FIGURE 3: *Thevenin equivalent of each DAC output. R_{OUT} is typically 12kΩ.*

REFERENCE INPUT LIMITS

The switches in the R-2R ladder are N-channel enhancement MOSFETs with extremely low ON resistance. To ensure the DAC's linearity, the MOSFETs' gate-to-source voltage (V_{GS}) needs to be greater than the switches' intrinsic threshold voltage, which for the DAC-8800 is 2.5V. When the voltage falls below 2.5V the MOSFETs' ON resistance increases, which causes resistance mismatching in the R-2R ladder. Any mismatching degrades the precise R-2R ratios and thus decreases the linearity of the DAC.

In the DAC-8800, the gate-to-source voltage is equivalent to the voltage difference between V_{DD} and $V_{REF}H$, respectively. Figure 2 shows that $V_{REF}H$ is connected to the drain of the MOSFET switches, and, when the switches are on, the drain voltage is basically equivalent to the source voltage. The gate voltage is driven by CMOS logic, and when the switch is on, the logic connects the gate to V_{DD}. Thus, $V_{REF}H$ must be at least 2.5V below V_{DD}, as shown in Figure 3 of the DAC-8800's data sheet. However, to guarantee the data sheet error specifications over –55°C to +125°C, the gate-to-source voltage needs to be at least 4V. There is no similar limitation between the reference input and the negative supply, V_{SS}. Thus, the reference inputs can go to V_{SS}. An important note: because of internal protection diodes in the DAC, $V_{REF}L$ should not be allowed to go higher than $V_{REF}H$. Forward biasing these diodes allows large currents to flow between the two references, potentially resulting in permanent damage to the DAC-8800.

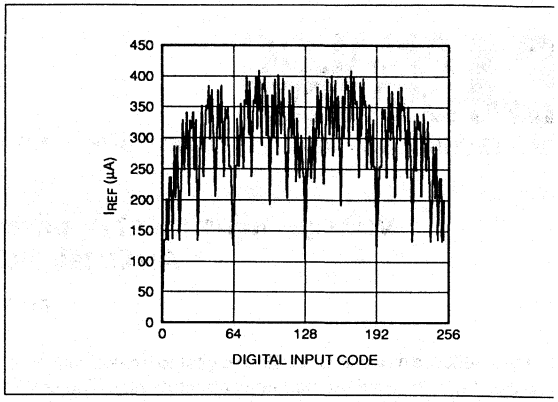

FIGURE 4: *$I_{REF}H$ variation versus digital code. One of four DACs connected to $V_{REF}H$. The other 3 DACs are loaded with zero code.*

REFERENCE INPUT CURRENT CHANGES WITH DIGITAL CODE

As the digital code changes, the resistance looking into the reference input changes significantly. Figure 4 shows the current demand into the $V_{REF}H$ pin as a function of the digital code for one of the four DACs referenced from that pin. This graph was generated with the following conditions: V_{DD} = +12V, V_{SS} = 0V $V_{REF}H$ = +5V, and $V_{REF}L$ = 0V. As can be seen, the load on the reference varies from zero to 400µA. With all four DACs operating, the load current can go up to a maximum of 1.6mA. It is important to keep in mind that the current changes in abrupt steps. Thus, in applications where speed is important, any device driving the reference pin must be capable of handling these step current changes. A fast recovery op amp (such as the OP-42) or reference is recommended.

THE DAC-8800 CANNOT BE USED AS A VARIABLE RESISTOR

At first glance the DAC-8800 might appear to be ideal for use as a variable resistor from its output to $V_{REF}L$, where $V_{REF}L$ is tied to ground and $V_{REF}H$ left floating. However, its internal structure was not designed for this. The reason is two fold. First, the resistance from the DAC's output to $V_{REF}L$ does not vary linearly with

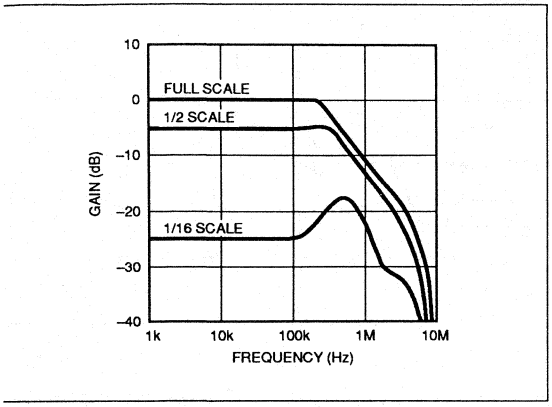

FIGURE 5: *DAC-8800 Bandwidth Under Different Gains*

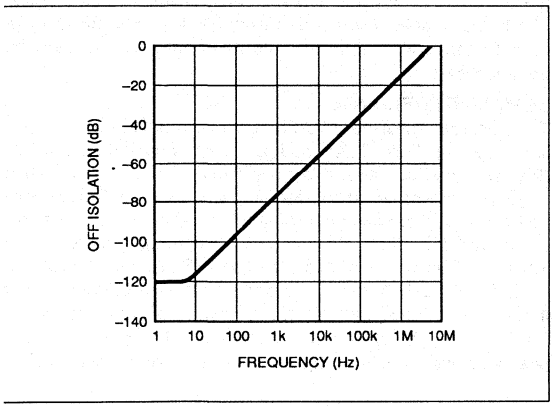

FIGURE 6: *DAC-8800 OFF Isolation*

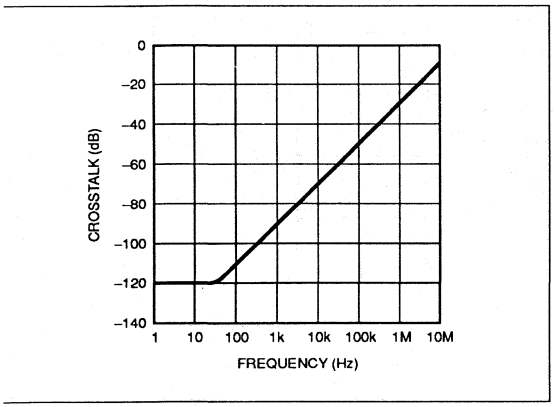

FIGURE 7: *DAC-8800 Crosstalk*

the digital code. Rather, it changes erratically, similar to the way the reference current changes in Figure 4. These seemingly random changes are due to various switches connected to $V_{REF}H$ turning on in binary fashion rather than sequentially and creating alternate current paths from the output to $V_{REF}L$.

Second, the NMOS switches are not bi-directional. In this variable resistor configuration, the switches connected to $V_{REF}H$ would actually have current flowing in reverse direction from the source to the drain. They maintain their low ON resistance only when current is flowing normally from the $V_{REF}H$ side (the drain) towards the output (the source). In backwards operation the source voltage causes changes in the ON resistance. Thus, any change of the voltage on the DAC's output will change the ON resistance and ultimately change the resistance to ground, even at the same digital code. Obviously, the DAC-8800 was designed to work as a voltage attenuator, and not as a variable resistor.

AC MULTIPLYING MODE OPERATION

The DAC-8800 is designed primarily as a DC adjustment device. However, it can also be used in multiplying mode by applying an AC signal to the reference input. In such applications, bandwidth, off-isolation, and crosstalk are important to the circuit's performance. The bandwidth of the DAC-8800 is limited by the ladder resistance and the internal capacitance, which are both specified in the data sheet. The typical resistance of 12kΩ, combined with the reference capacitance of 75pF, limits the bandwidth to 177kHz. Figure 5 shows actual network analyzer measurements of the –3dB bandwidth, which for this particular part occurs at 360kHz. The fact that the measured bandwidth is twice the typical points out how the bandwidth can vary due to varying capacitance and resistance from device to device. The worst case bandwidth is approximately 100kHz based on worst case resistor and capacitor values of 16kΩ and 100pF, respectively. Remember, as mentioned in the reference input limits section, $V_{REF}H$ cannot go below $V_{REF}L$. Any AC signal into $V_{REF}H$ must be biased to avoid this condition.

The off isolation of the DAC-8800, shown in Figure 6, was measured using an AC signal for $V_{REF}H$ and measuring an associated DAC output with all the bits off. The off isolation reveals how much of the input signal will feed through to the output. An interesting correlation can be made between this graph and the bandwidth graph of Figure 5, for the 1/16 scale measurement. The 1/16 scale shows a 10dB rise in the gain above 100kHz. This is actually due to the capacitive feedthrough of the DAC-8800.

The crosstalk versus frequency graph in Figure 7 was measured as the crosstalk from one set of four DACs to the other set of four DACs in the package. In other words, DACs A through D were set to full scale, and a frequency dependent signal was injected into their $V_{REF}H$ input. The crosstalk was then measured on the outputs of DACs E through H. The graph shows DC crosstalk of –120dB rising up to –50dB at 100kHz, revealing excellent performance for DC and low frequency AC signals.

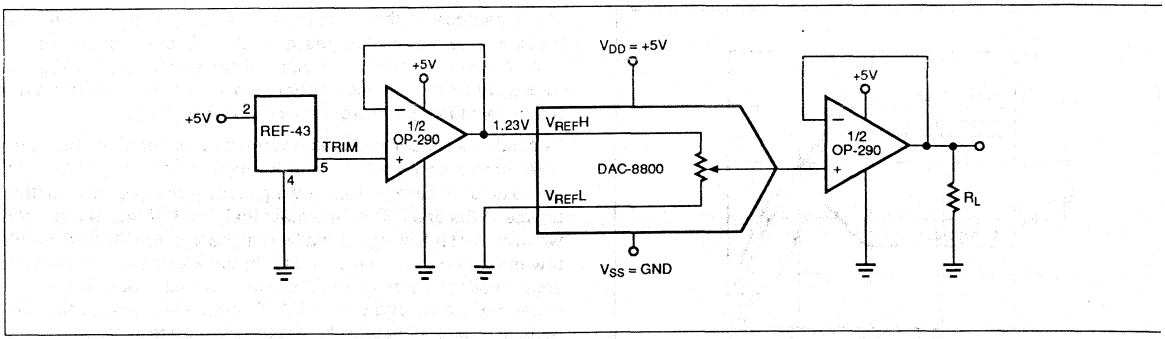

FIGURE 8: *DAC-8800 Single +5V Operation*

OUTPUT NOISE

The DAC-8800 exhibits basic broadband white noise of typically 18nV/$\sqrt{Hz}$. Its dominant noise source is the resistor ladder. Thus, the noise does not exhibit any measurable 1/f noise.

SINGLE +5V SUPPLY OPERATION OF THE DAC-8800

The DAC-8800 is ideal for single supply applications because its output range includes ground. In fact, the DAC-8800 can work well with a single +5V only. Even with this low of a supply voltage, $V_{REF}H$ can be connected to a 1.23V bandgap reference (Figure 8). Although the 1.23V bandgap reference violates the 4V of headroom requirement, the DAC is still within ±1/2 LSB of total unadjusted error. The 4V below the positive supply limit was set with a safety margin of about 0.5V to account for operation over the full operating temperature range.

The 1.23V bandgap reference voltage is derived from the TRIM pin of a precision 2.5V reference device, the REF-43. The buffer amplifier is needed because the TRIM pin's impedance is 50kΩ. The DAC-8800's reference inputs characteristically range from 12kΩ to 40kΩ depending on the digital code, which would load the trim pin excessively. The OP-290's low offset voltage of 75μV and low temperature drift characteristics maintain the reference's accuracy. The OP-290 also has the ability for its output to operate to ground with the addition of a load resistor; 10kΩ works well. The output amplifier is needed to buffer the DAC-8800's high output impedance when the output is connected to a low impedance load.

SERIAL INTERFACING

The digital control of the DAC-8800 is a standard three-wire serial interface with clock (CLK), load ($\overline{LD}$), and serial data input (SDI) (Figure 9). Additionally, an inverted $\overline{CLK}$ input pin is available for negative edge triggered data loading. Either CLK or $\overline{CLK}$ can also be used as a chip select pin. When loading data, 3 address bits are loaded, MSB first, followed by 8 bits of data, again MSB first. Thus 11 bits in all are loaded through the SDI pin to control each DAC. The DAC-8800 can run on a clock as fast as 6.6MHz making it possible to load all eight DACs in as little as 14 microseconds. Furthermore, the DAC-8800 maintains TTL compatibility for positive power supply voltages greater than or equal to +5V.

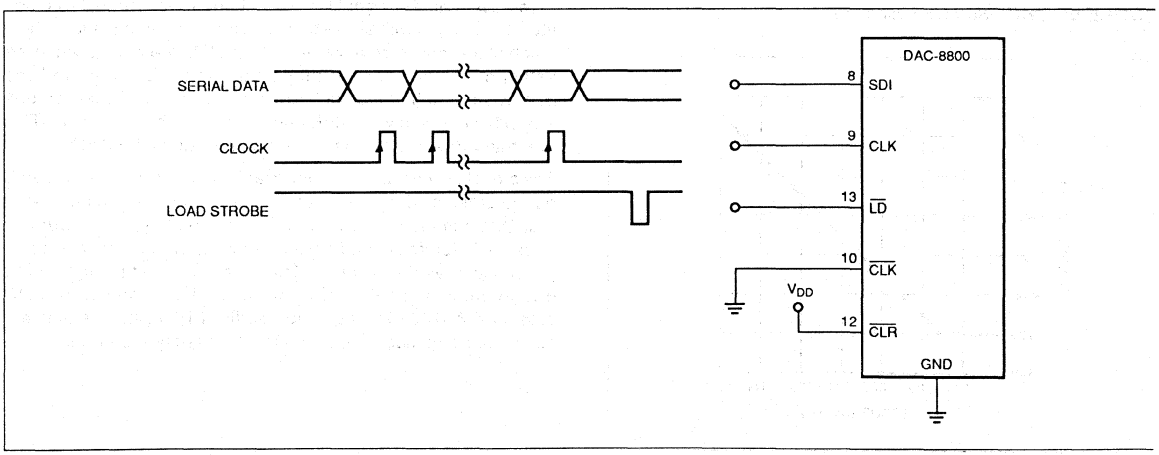

FIGURE 9: *DAC-8800 Serial Interfacing*

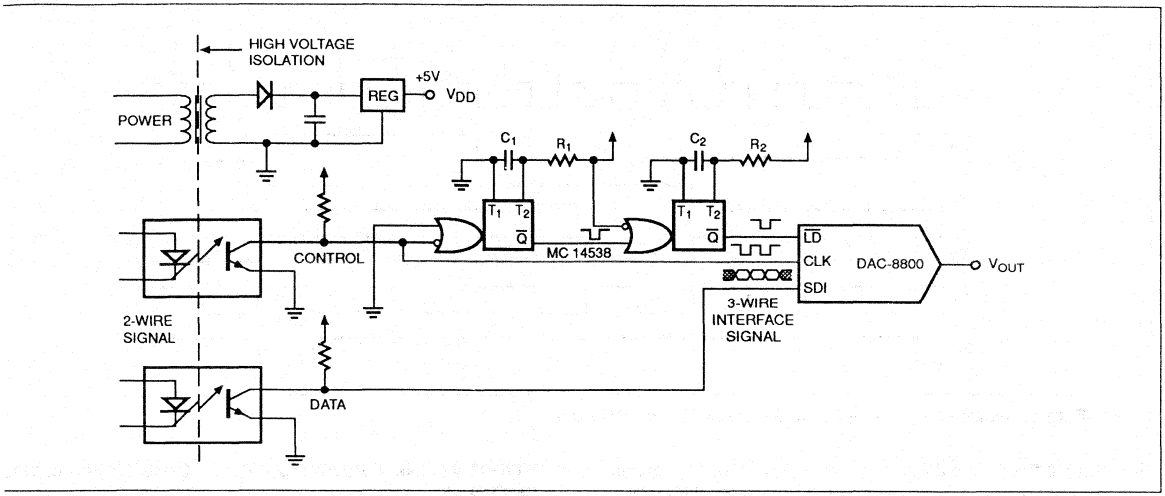

IGURE 10: *Isolated Two-Wire Serial Interface for the DAC-8800*

'WO WIRE INTERFACES FOR PROCESS ENVIRONMENTS

ligh voltage isolation using opto-couplers is often necessary for erial interfaces found in process control applications. In these nd other applications where minimizing the number of data lines ; desirable, two-wire signal interfaces can be used (Figure 10). 'his simple circuit translates the two-wire interface into the three ata lines required to load the DAC-8800. The $\overline{\text{LOAD}}$ signal is enerated using two retriggerable one-shots. The first one-shot's meout should be set longer than the clock period. Each suc-eeding clock pulse will retrigger the one-shot until all 11 bits are)aded into the DAC. Then the clock must pause long enough to llow the one-shot to time out. When the first one-shot's output goes low, it triggers the second one-shot, which produces the $\overline{\text{LOAD}}$ pulse, and finishes the loading cycle.

There are some common pitfalls when using one-shots. For example, the timeout set by the external resistor and capacitor can vary over temperature and from part to part. Even more significant is the variation due to resistor and capacitor tolerances. A typical capacitor can vary by ±10% which will cause an equivalent ±10% variation in the timing of the one-shot. To avoid the problems of one-shots, a second method using a counter is recommended (Figure 11a). The counter keeps track of the number of clock cycles and, when all the data has been input to the DAC, the external logic creates the $\overline{\text{LOAD}}$ pulse.

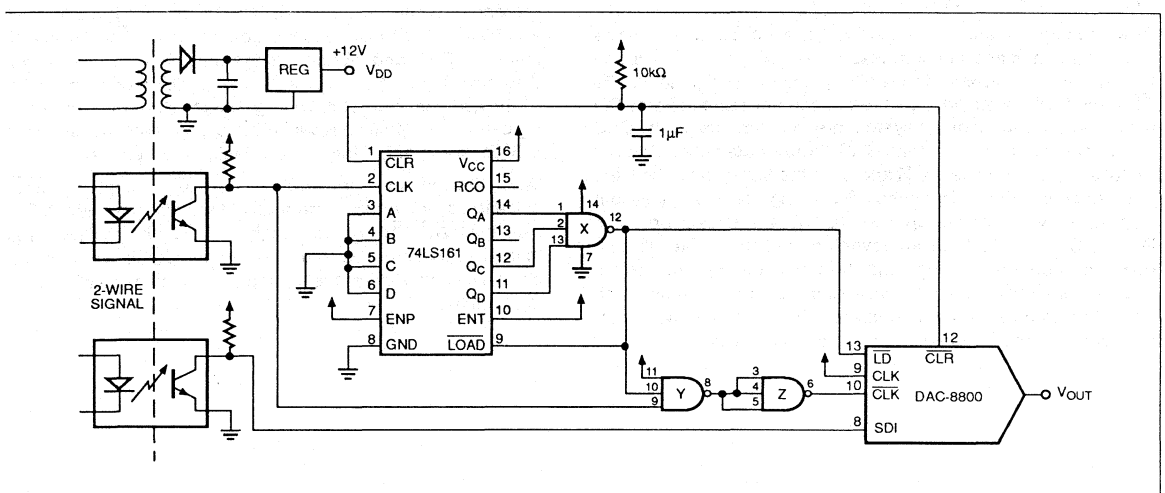

IGURE 11a: *Isolated Two-Wire Serial Interface Using a Counter*

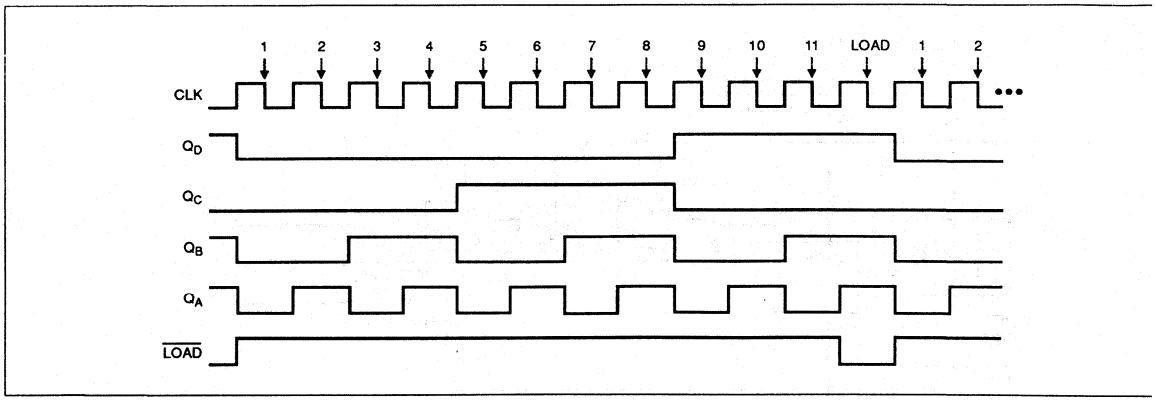

FIGURE 11b: *Isolated Two-wire Serial Interface Timing Diagram*

Referring to the timing diagram (see Figure 11b), the counter is incremented on every rising edge of the clock. Additionally, the data is loaded into the DAC-8800 on the falling edge of the clock by using the $\overline{CLK}$ input instead of the CLK input. The reason for using the $\overline{CLK}$ input becomes apparent after considering the $\overline{LOAD}$ pulse. The timing diagram shows that after the eleventh bit has been clocked, the output of the counter is binary 1010. On the following rising CLK edge the output of the counter changes to binary 1011, upon which NAND gate 'X' goes low to generate the $\overline{LOAD}$ pulse. The $\overline{LOAD}$ signal is connected to both the DAC's $\overline{LD}$ and the counter's $\overline{LOAD}$ pins. Since the counter has a synchronous clear, the $\overline{LOAD}$ pulse remains low until the next CLK pulse. NAND gates 'Y' and 'Z' prevent the twelfth falling CLK edge (labelled 'LOAD' in the timing diagram) from clocking the DAC, which would load false data into the DAC. Using the $\overline{CLK}$ input allows sufficient time from the CLK edge to the $\overline{LOAD}$ edge, and from the $\overline{LOAD}$ edge to the next CLK pulse, to satisfy the timing requirements for loading the DAC-8800.

After loading one address of the DAC, the entire process can be repeated to load another address. If the loading is complete then the CLK must stop after the twelfth pulse of the final load. The $\overline{CLK}$ input will be pulled high and the counter reset to zero. The timing requirements of the system are the same as for the DAC alone, and can be found in the DAC-8800's data sheet. Another feature of this circuit is the R and C on the CLR pins of both the DAC and the counter. This simple RC timing circuit will clear both chips upon system power-up. The 74LS161 was chosen because, like the DAC-8800, it has an asynchronous clear. The RC time constant should be set longer than the power supply turn-on time. The values shown in the circuit give a time constant of 10ms, which should be adequate for most systems. This same two-wire interface can be used for most three-input serial DACs.

STAND-ALONE OPERATION PROVIDING NONVOLATIL SETTINGS

Whenever a system with a DAC-8800 is powered on, the DA(8800 needs to have all eight of its data words loaded to set tl proper DC output voltages. In a system with a microprocessor microcontroller, this is a straightforward operation. However, some systems the DAC-8800 may be the only part with a digit interface. In this case, the circuit shown in Figure 12a will aut matically load the DAC on system power-up. The core of the circ is a serial input/output EEPROM device (U_2), preprogramme with the appropriate data for the DAC. The counter labelled U counts through 8 addresses, which are serially shifted into tl EEPROM by U_3, a parallel to serial shift register. The EEPRO shifts out a 16-bit word associated with each address. Only 11 the 16 bits are actually shifted into the DAC before the $\overline{LOAD}$ puls arrives.

The second counter, U_7, in combination with the flip-flop U_6, coun the loading of the bits into the EEPROM and into the DAC-880 When all the bits are loaded the logic sends a $\overline{LOAD}$ pulse whic loads the DAC and increments the address on U_4. The timir diagram in Figure 12b gives a detailed description of the loadir of one address. The CLK INH logic inhibits the shift register du ing certain CLK pulses because 9 bits need to be loaded into tl EEPROM and only 8 bits are available in the register.

When the system is powered-up, R_1 and C_1 create a $\overline{PWRU}$ pulse to asynchronously clear all of the counters and the DA(After the $\overline{PWRUP}$ pulse goes high, the free running clock begir to load all 8 addresses. After the eighth address is loaded, tl clock is disabled to remove any digital switching noise in the an log circuitry.

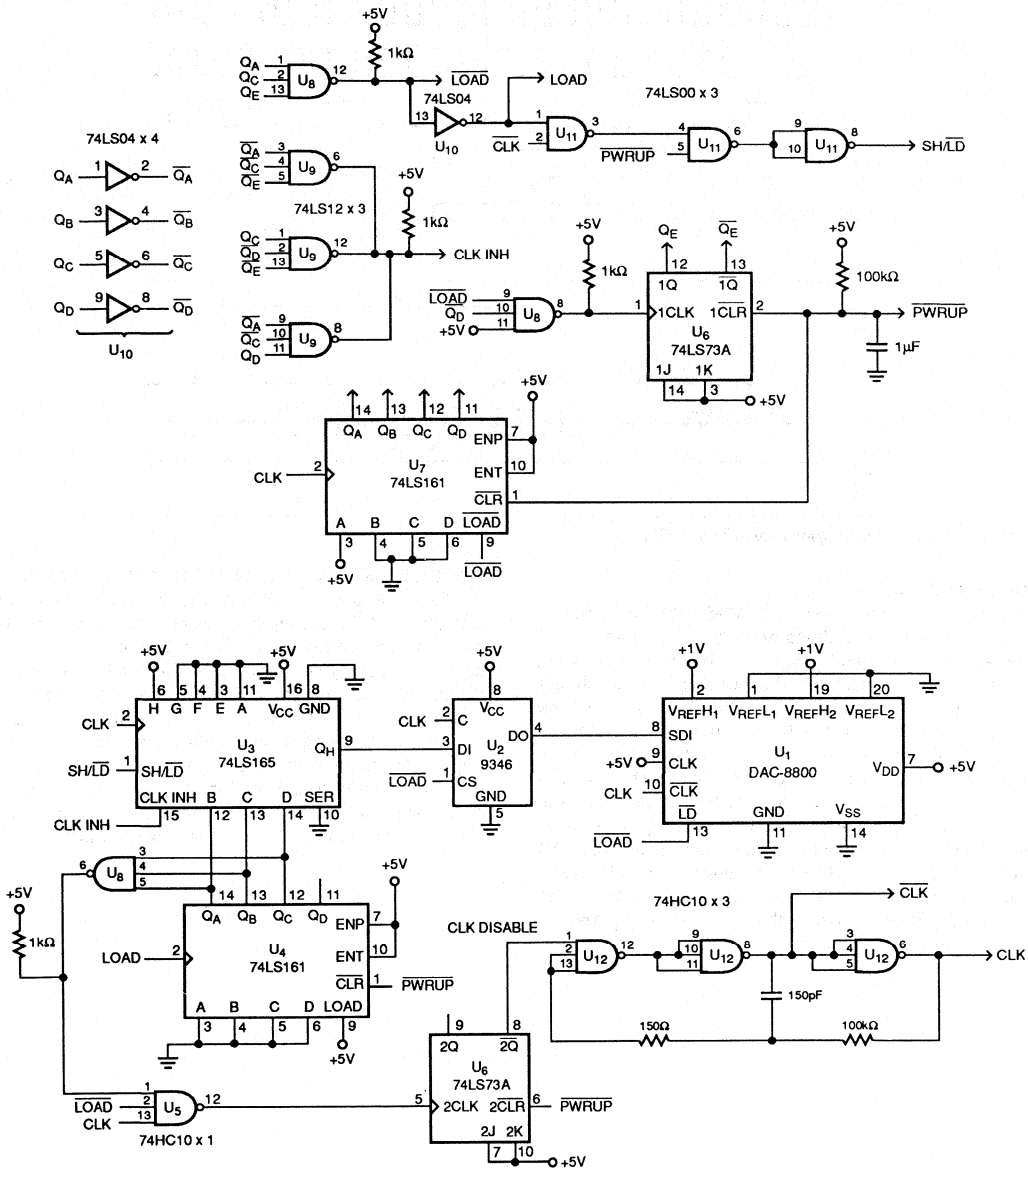

FIGURE 12a: *Stand-alone operation of the DAC-8800. The EEPROM stores data to set the DAC's output voltages on system power-up.*

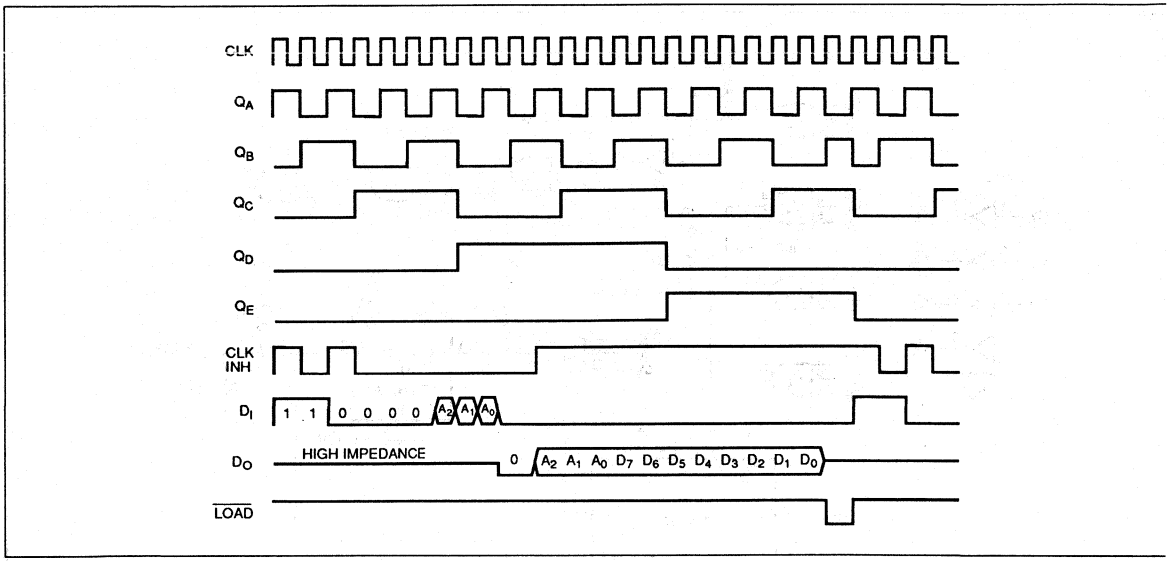

FIGURE 12b: *Stand-Alone Operation Timing Diagram*

DAC-8800 TEST FIXTURE

Figure 13 shows a fixture to test the DAC-8800. The circuit includes switches to set the address and data so that each DAC can be loaded with any digital word. After the switches are set, pressing the push-button activates the monostable multivibrator which generates the clock signal to load the DAC register. The counter selects the successive MUX channels, which switch the bits in proper loading order. After all eleven bits are loaded, the $\overline{LOAD}$ switch is manually toggled to generate a $\overline{LOAD}$ pulse. The $\overline{CLR}$ switch should be high at all times except to clear all eight DACs, in which case $\overline{CLR}$ needs to be switched low and then back to high. The CS switch should always be set low to keep the DAC selected at all times. The DAC-8800's outputs are buffered by OP-400 operational amplifiers; however, the outputs can be configured many different ways for the actual tests required.

CONDITIONS TO AVOID IN USING THE DAC-8800

The DAC-8800 is very resistant to the most common latch-up conditions.[1] For example, it does not latch-up when the digital inputs go high before the DAC is powered up. Nor is the DAC's own power supply sequencing significant. However, a few conditions still exist that are potentially destructive. In order to prevent damage from latch-up and ESD, the internal DAC-8800 design includes large body diodes placed from many of the pins to ground or either of the supplies, as shown in Figure 14. Looking at the placement of the diodes, it is easy to understand what conditions need to be avoided. The voltage across these diodes should be less than 0.3V, or they will start to turn on.

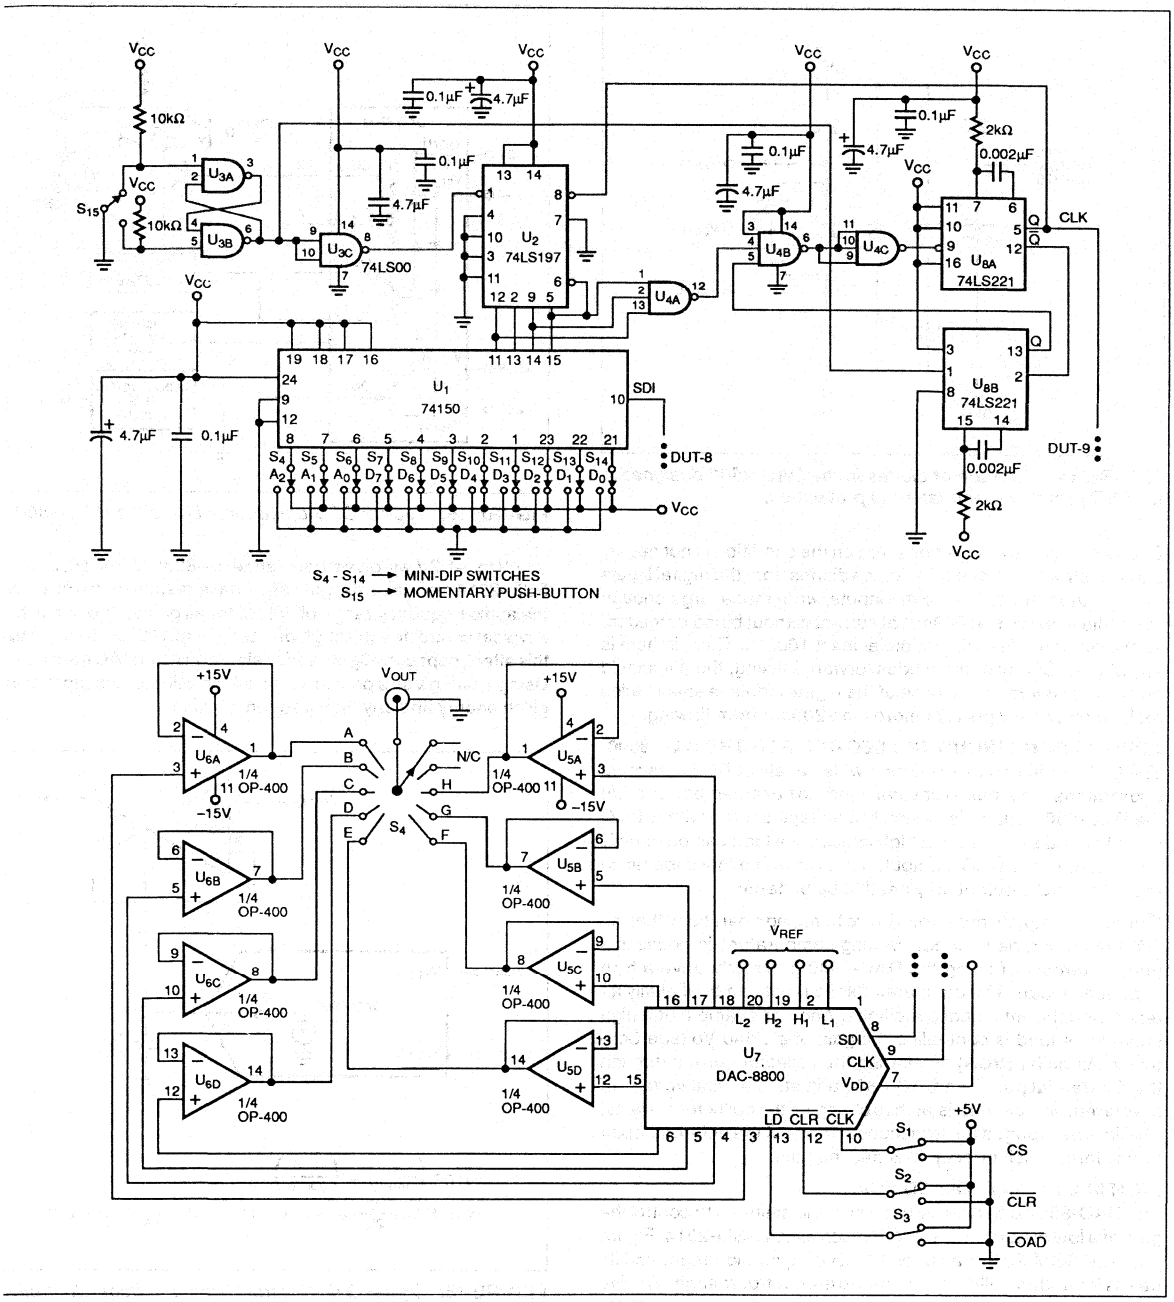

FIGURE 13: *DAC-8800 Test Fixture*

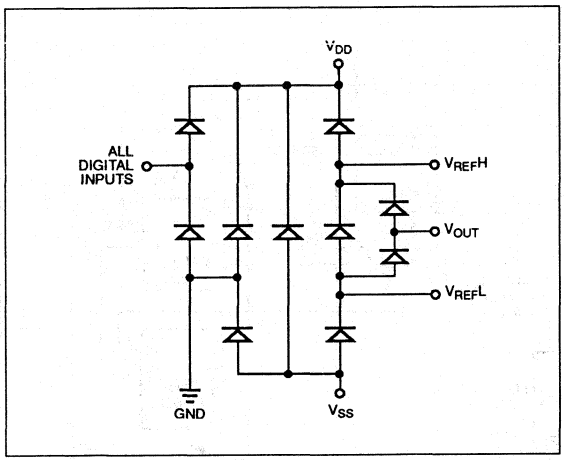

FIGURE 14: *Diagram of diodes in the DAC-8800 designed for ESD protection and latch-up prevention.*

Even if one of the diodes does turn on the condition is not necessarily destructive. For example, the diodes from the digital inputs to V_{DD}, and from ground to the inputs, were made large enough to handle in excess of 200mA of current without being damaged. All the other diodes can handle at least 100mA. Thus, if there is any chance of any of the diodes forward biasing, the pin should be current limited. In the case of the digital inputs, a small series resistor can easily prevent more than 200mA from flowing.

APPLICATION CIRCUIT COLLECTION FOR THE DAC-8800

The DAC-8800 can be used for a wide variety of DC adjustment applications. The main point that needs to be remembered is that the DAC-8800 output is basically a voltage source with a 12kΩ output impedance. Thus, a high impedance load can be directly connected to the DAC's output, however with a low impedance load, the DAC's output may need to be buffered.

Figure 15 suggests numerous basic trimming operations that the DAC-8800 can be used for. Setting comparator trip points is a prime example of using the DAC-8800 to directly drive a high impedance load. The comparator's trip point can be digitally altered for different signal conditions. Another example of a high impedance load is controlling the gain of a video Voltage Controlled Amplifier (VCA) by altering the collector current through the differential pair. The DAC-8800 adjusts the base voltage of the current source transistor thus changing the collector currents. This in turn changes the transconductance of the differential pair transistors, which directly changes the gain.

DIGITALLY-CONTROLLED VCA

The DAC-8800 can also be used in audio systems to control the gain of a low distortion audio VCA such as the SSM-2014 (Figure 16). The SSM-2014 has over 100dB of dynamic range, and its gain is logarithmically proportional to the control voltage, V_C. The DAC can be connected directly to the control port of the VCA, which has a gain sensitivity of –30mV/dB. A reference range from

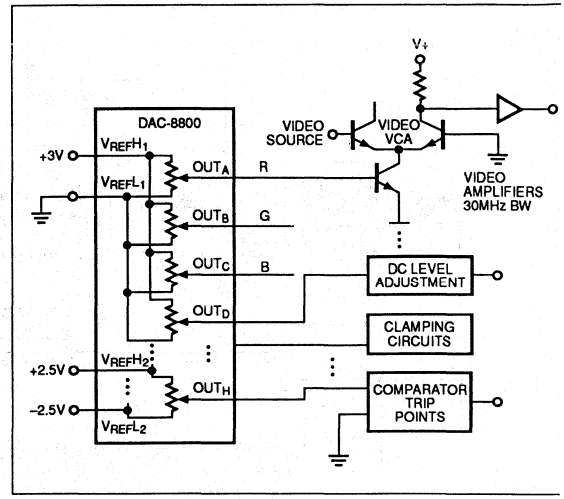

FIGURE 15: *Typical DC Adjustments Using the DAC-8800*

+2.5V to –1.2V will give a gain range of –80dB to +40dB, and the SSM-2014 maintains flat gain and phase response to well above the audio frequency range of 20kHz for all gains. The circuit has a typical control feedthrough of 1.3mV/V at 100Hz. To minimize this effect, capacitor C_S is used to slow down the DAC transitions. Using 1.0µF gives a pole at 13Hz, which will filter out most of the glitch energy and any high frequency noise.

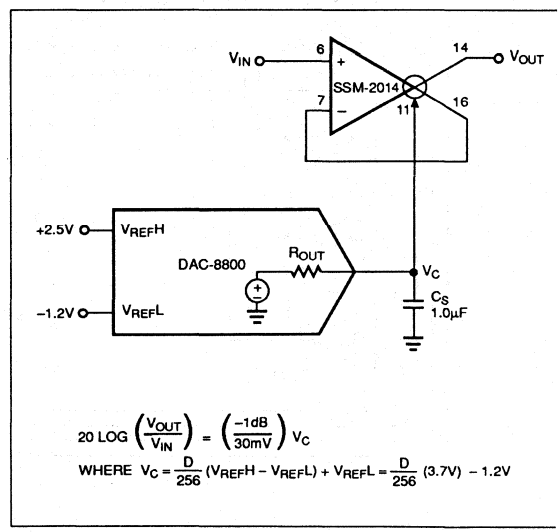

FIGURE 16: *Digital Gain Control Using a Voltage-Controlled Amplifier*

TRIMMING OP AMP OFFSET VOLTAGE

A frequently encountered DC adjustment application is trimming op amp offsets. There are many straight forward methods for trimming; one of which is connecting the DAC-8800 to the op amp's null pins. The DAC-8800 can directly null op amps to the negative supply, provided the supply is −12V or less in magnitude (Figure 17). This limit is because of the maximum 20V limit across the DAC. Since the positive supply needs to be at least +5V for logic interfacing, the DAC-8800's negative supply is safe to around -12V. The references used need to be near the voltage of the op amp's trim pins, which is typically a couple hundred millivolts above the negative rail. The figure shows reference values that work well for trimming the OP-42 over a ±40mV range. The actual voltages can be created using resistor dividers from the negative supply to ground and buffering the reference inputs with op amps. In cases where the op amp is adjusted from the positive rail, one of the alternative methods in the following paragraphs is needed. The reason for this is that the DAC-8800's output would need to be able to go up to the positive supply. However, $V_{REF}H$ is limited to 4V below the positive supply. Thus, the offset cannot be directly adjusted around the positive supply.

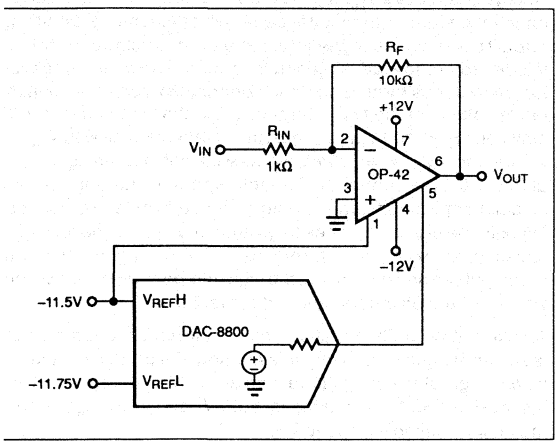

FIGURE 17: *Using the DAC-8800 for offset nulling directly on the op amp's null pins.*

A simpler method of offset nulling that gets around the supply voltage limitation is to connect the DAC-8800's output in series with a resistor to the summing node of the amplifier (Figure 18). This adds a small current that cancels the op amp's offset voltage. The series resistor should be large to provide a fine adjustment. The noise of the DAC and the series resistor might at first appear to be a problem, but it is actually attenuated by the 1kΩ input resistor. Therefore, the 1kΩ noise dominates. For the values in Figure 18, the adjustment range is ±50mV on the output. Figure

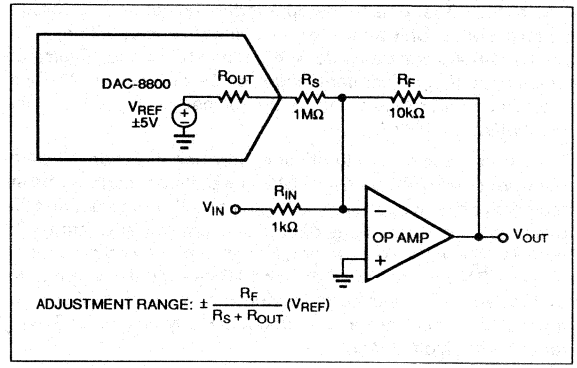

FIGURE 18: *Offset nulling by connecting the DAC-8800 to the summing node of an amplifier.*

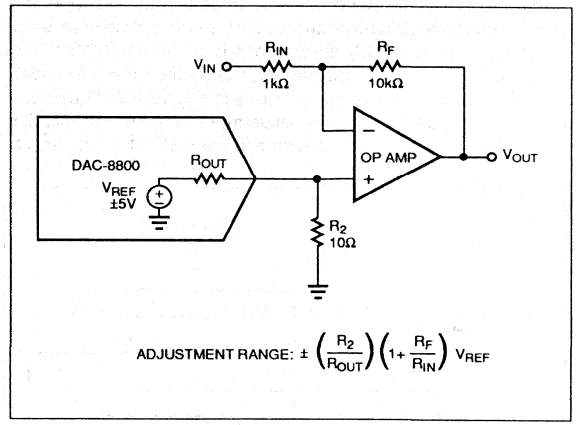

FIGURE 19: *Offset nulling by connecting the DAC-8800 to the noninverting node of an amplifier.*

19 shows an alternative method of offset nulling by adjusting the voltage at the amplifier's noninverting input. The resistor divider is recommended to provide for fine control of the offset. The adjustment range for the values in Figure 19 is ±42mV. The small resistor-to-ground also reduces the DAC's output noise to a point where it is insignificant compared to the op amp's own noise. With R_2 = 10Ω, the input noise caused by the DAC reduces to 15pV/$\sqrt{Hz}$; the noise of R_2 is much larger than this.

In both nulling applications shown, a positive and negative reference is required; however, in certain systems, only one reference may be available. Thus, the DAC-8800 can only adjust the offset in one direction. If this is the case, the amplifier can be forced to offset in either the positive or negative direction by connecting

one of the offset pins to the appropriate power supply voltage. Then the offset only needs to be adjusted in one direction, which the DAC-8800 can easily do with just one reference. Typical op amps have V_{OS} adjustment ranges on the order of ±5mV. Thus, the adjustment range of the DAC-8800 needs only to adjust the input offset by 10mV.

If the op amp is forced to offset in only one direction, the temperature coefficient of V_{OS} (TCV_{OS}) will almost certainly be affected. For a typical op amp, TCV_{OS} is directly proportional to the offset voltage. Thus, if the offset voltage is larger to start with, then TCV_{OS} will also be larger. For example, an op amp with a simple NPN input stage will exhibit a TCV_{OS} equal to V_{OS} divided by the temperature in degrees Kelvin at which that offset was measured. This translates to approximately 3µV/°C of TCV_{OS} for every millivolt of V_{OS}.

TRIMMING VOLTAGE REFERENCES

Figure 20 shows the DAC-8800 being used to trim a voltage reference such as PMI's REF-01. The output of the DAC is connected to the TRIM pin of the reference just as the wiper of a potentiometer would be connected. This entire circuit can easily be used in single supply applications because the DAC-8800's output can go to V_{SS}. Furthermore, this method can be used for all of PMI's references provided there is at least 4V of headroom between V_{DD} and V_{OUT}. The adjustment range of this circuit is the same as the ±300mV specified in the REF-01's data sheet.

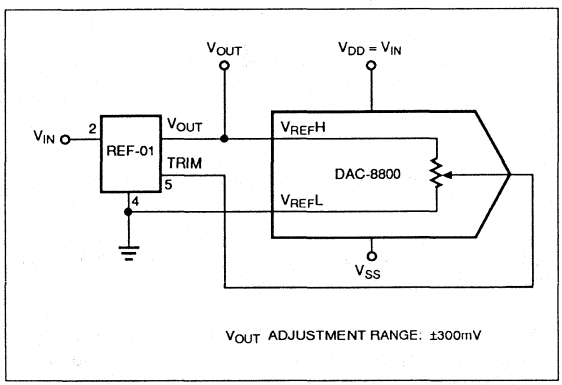

FIGURE 20: *Reference Trimming Using the DAC-8800*

COARSE-FINE CONTROL

Two of the DAC-8800's outputs can be connected together, and the resulting output is the average of the two unconnected outputs (Figure 21). This can easily be seen by thinking of the Thevenin equivalent circuit in Figure 3. The two output resistances in the same package are well matched so they form an accurate resistive divider, which averages the two DAC voltages. Such a circuit could be useful for performing a coarse-fine control, where one of the references is set to 1/10 the other reference. For example, setting $V_{REF}H_1$ to 1.0V, $V_{REF}L_1$ to –1.0V, $V_{REF}H_2$ to 100mV, and $V_{REF}L_2$ to –100mV gives an output adjustment range of ±0.5V ±0.05V.

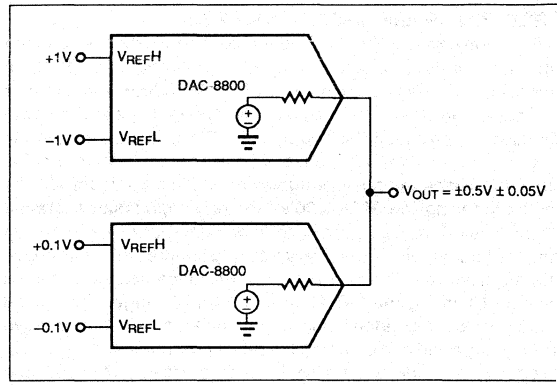

FIGURE 21: *Coarse-Fine Control by Averaging the DAC-8800 Outputs*

This application is limited by the voltage difference between th references. If $V_{REF}1$ is too much larger than $V_{REF}2$, the outpu voltage goes above $V_{REF}2$, and the switches in $V_{REF}2$ start to tur on independent of the digital code, which causes a significar error. It can even cause $V_{REF}2$ to increase by forcing current bac into the reference, causing nonlinearities. The actual threshol for the switches turning on varies depending upon the commo mode voltage of the two references, but the worst case is 1.5 between $V_{REF}1$ and $V_{REF}2$. Thus, the above example with $V_{REF}H$ = 1.0V and $V_{REF}H_2$ = 100mV works well, but increasing $V_{REF}H$ above 1.6V may cause the switches to start turning on. Fo unipolar applications, this problem can be avoided by connect ing both the high references to ground and setting the low refer ences to, for example, –10.0V and –1.0V. Configured this way the output can never go above the high reference, and the switche will never turn on independent of digital code.

An alternative method that avoids the aforementioned probler is shown in Figure 22. Using this method, the output voltage ca never be greater in magnitude than any of the references. To ge the maximum output voltage, set both V_1 and V_2 to $V_{REF}H$, the the equation simplifies as follows:

$$V_O = V_{REF}H\,[\,(2R_{OUT} + R_S) / (2R_{OUT} + R_S)\,] = V_{REF}H$$

Thus, the maximum output voltage is equal to the reference c the DACs. Another advantage is that only half as many reference are needed. One thing to be careful of is that the percentag adjustment range of each DAC output will vary with changin output resistances from device to device. In Figure 22, with R_{OU} = 12kΩ, R_S = 96kΩ sets the output to be 10% of V_1 and 90% of V_2 However, if R_{OUT} changes to 8kΩ, then the percentages becom 7% and 93%, respectively. If this is unacceptable, then R_S need to be variable from at least 64kΩ to 128kΩ to cover the entire outpu resistance range of the DAC.

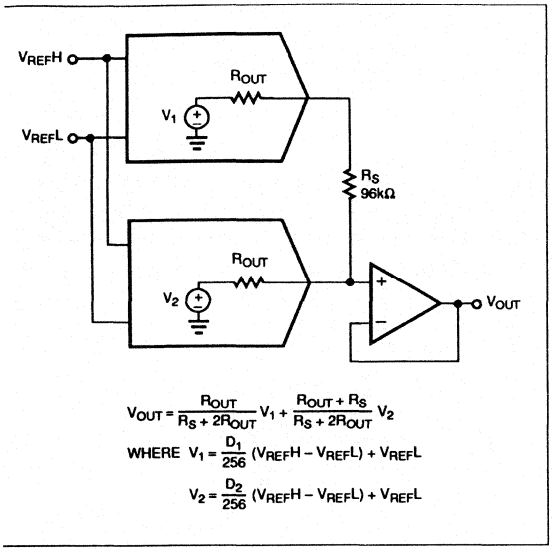

FIGURE 22: *Course-fine adjustment using the same reference for both DACs.*

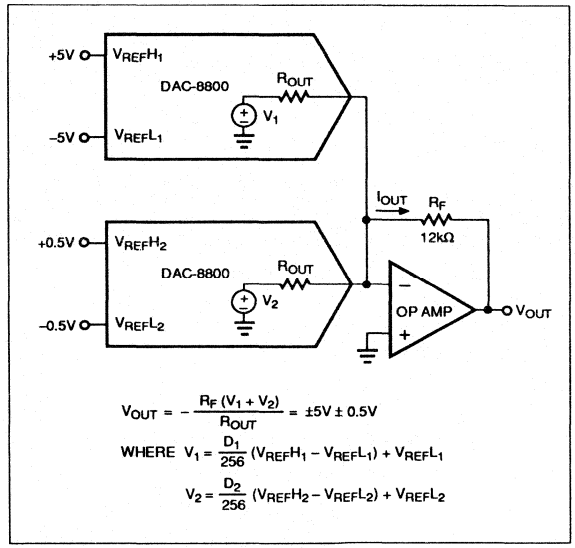

FIGURE 23: *Voltage summer by connecting the DAC-8800 outputs to the summing node of an amplifier.*

8

Another method, shown in Figure 23, is current summing. In this case, the DAC outputs are connected to the virtual ground of the op amp, avoiding the problem of the switches being forced on. The feedback resistor should be 12kΩ to match the output impedance of the DAC-8800. As in the above application, the feedback resistor may need to be varied from 8kΩ to 16kΩ depending on variations in the DAC-8800's output resistance. Also remember that the op amp inverts the DAC's reference voltages, so a 5V high reference gives –5V at the output of the op amp.

AN ADJUSTABLE REFERENCE FOR ANALOG-TO-DIGITAL CONVERTERS

In an analog-to-digital conversion circuit the DAC-8800 works well as a digitally-controlled reference (Figure 24). Using the DAC, the reference voltage can be adjusted for different ADC sensitivities. The DAC-8800 output may need to be buffered by an op amp because of the typical ADC's low reference input impedance. For flash converters the typical input impedance is usually well below 1kΩ, which is much too low for the DAC-8800's output impedance of 12kΩ. A separate DAC in the same package can provide the negative reference as well, but it too has to be buffered with an op amp.

As can be seen by the many different application circuits shown in the above section, the DAC-8800 is a very versatile device. Of course, the application examples shown here are only a small selection of the many different ways the DAC-8800 can be used.

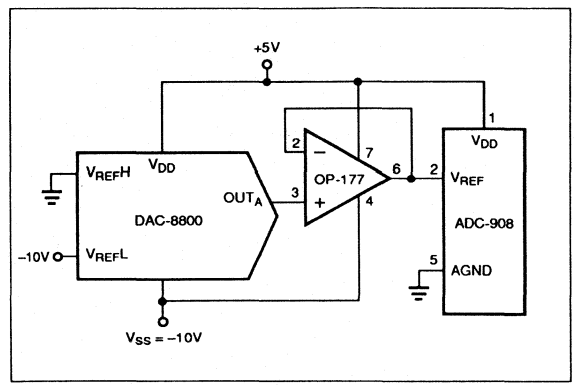

FIGURE 24: *The DAC-8800 as a digitally-controlled variable reference for ADCs. The DAC's output needs to be buffered by an op amp.*

REFERENCES

1. AN-109, "Understanding and Preventing Latch-up in CMOS DACs." Analog Devices, Inc., January 1989.

ANALOG DEVICES

AN-219
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Electronic Adjustment Made Easy with the TrimDAC

by Walter Heinzer and Joe Buxton

The TrimDAC® is a multi-channel d/a converter designed specifically for adjusting gains and dc levels in electronic circuits digitally and without moving parts. It combines many of the properties of the adjusting pot(entiometer) with the prospect of hands-off automatic adjustment and high reliability. The highly desirable attributes of TrimDACs include small package size, many devices per package, serial interface (reduces pin count) and low power dissipation. TrimDACs in electronic adjustment reduce cost in two ways: the higher speed of adjustment under software control saves time and capital investment; and the device itself is quite cheap.

Most designers of new circuit designs would like to avoid the once-ubiquitous variable resistor because of its mechanical sensitivity, relatively wide absolute tolerances, and high labor cost. But there is generally a need for factory adjustments and calibrations in electronic equipment. Even digital products need a power supply adjusted or calibrated to a specified tolerance. And many electronic systems are connected to real world sensors or output devices in *systems* that need calibration. A key issue facing engineers who design such systems is cost reduction of factory calibration and field maintenance.

Electronic factory-calibration of chips by semiconductor manufacturers is already widely used; calibration and adjustment problems are solved on (and with) integrated circuits by autozero, self-calibration, Zener-zap, fuse link, EPROM and laser trim. Something akin to this in larger-scale real-world systems is highly desirable.

Recognizing this problem, we sought to design products to fill the need for digitally adjustable electronic devices to automate, speed up, and eliminate manual and mechanical adjustments.

For example, consider the CRT display; curvature aberrations in the manufacturing of glass tubes require that the elements of focus (convergence & color purity) be individually adjusted, especially for high-resolution displays. Since the convergence adjustment of CRT display systems with resolutions of more than 1,000 lines requires that 6 to 8 variable-resistor adjustments be made in high-volume production—currently by robot-controlled screwdrivers—displays are ideal candidates for electronically controlled adjustment devices. The TrimDAC offers an attractive alternative to this mechanical adjustment approach. Previously a labor-(human or robot) intensive process taking minutes, the operation now can be done in seconds.

TrimDAC is a registered trademark of Analog Devices, Inc.

VOLTAGE ADJUSTMENT, THE FIRST GENERATION

The first TrimDAC, the DAC-8800, is a monolithic CMOS IC with all the ingredients necessary for general-purpose dc voltage setting. It contains eight unbuffered voltage-output d/a converters in a 20-pin skinny DIP package (Figure 1). The output voltage range, unipolar or bipolar, can be independently set for each group of four DACs. Output voltage ranges are established by the choice of high- and low- external-reference inputs. The 8-bit DACs provide 256 voltage levels within each range.

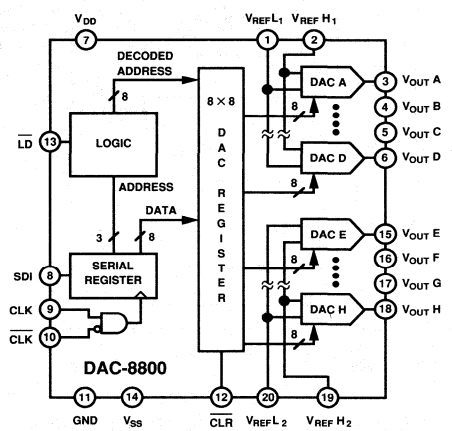

Figure 1. DAC-8800 Block Diagram. Shared References Determine Output Voltage Range.

A TTL-compatible 3-wire serial interface loads the contents of the eight internal DAC registers. These can all be set to zero by an asynchronous Clear (CLR) input, very handy for system power-up. An internal regulator provides TTL compatibility over a wide range of V_{DD} supply voltages. Single-supply operation is available by connecting V_{SS} to GND. The device achieves its performance and flexibility with a low 24 mW of dissipation.

The output voltage of each DAC is changed by clocking an 11-bit word (3 address bits, 8 data bits) into the serial shift register. The internal logic decodes the three address bits to establish which internal DAC register will receive the 8 bits of data from the serial register during the Load (LD) strobe. One DAC is updated with each LD strobe. At the maximum clock rate of 6.6 MHz, all eight d/a converters can be loaded in as little as 14 microseconds.

The output voltage range is determined by the external input voltages applied to V_{REFH} and V_{REFL} (Figure 2). If V_{SS} is negative, V_{REFL} may be set to a negative value; this results in a programmable bipolar range of output voltages. The relationship between V_{OUT} and V_{REFH}, V_{REFL}, and the digital input, D (a base-10 integer between 0 and 255), is:

$$V_{OUT}(D) = (D/256)(V_{REFH} - V_{REFL}) + V_{REFL}$$

The DAC-8800 is tested for operation with V_{DD} = 12 V and V_{SS} = 0 V or –5 V. However, it was designed to operate from a wide variety of available supply-voltage combinations. Here are some typical pairings: V_{DD}, V_{SS} = *+15 V, 0 V; +12 V, 0 V; +12 V, –5 V; +5 V,–5 V; +5 V, –12 V; +5 V, 0 V.*

The primary application of the DAC-8800 is with fixed reference inputs for dc voltage setting. Outputs may be applied directly to high-impedance circuits—or to external op amps for buffering.

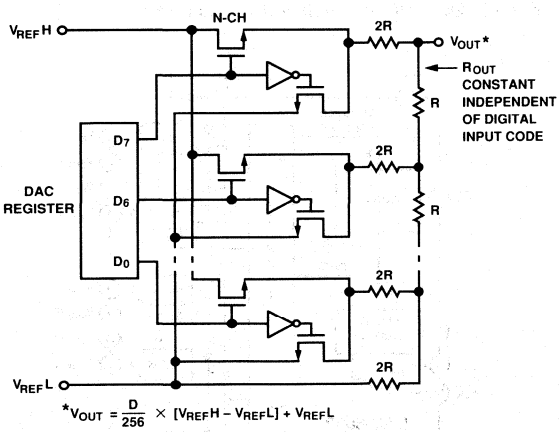

Figure 2. Simplified Equivalent Voltage-Switching DAC Circuit. The Output Resistance Remains Constant at a Nominal 11 kΩ.

ADDING GAIN: THE SECOND GENERATION

Second-generation TrimDACs, such as the DAC-8840 and DAC-8841, solve the problem of replacing variable resistors for adjusting ac or varying dc voltages—for example, in audio volume control. Other common applications where ac signals must be attenuated include some circuits found in video displays, projection-TV displays, instrumentation, oscilloscopes, medical gear, modulation circuits, modems, and so on.

The DAC-8840 contains a multiplying DAC structure with four-quadrant multiplying capability. Figure 3 shows the connection of one of the eight independent channels of the DAC-8840. This multiplying channel has a 1-MHz

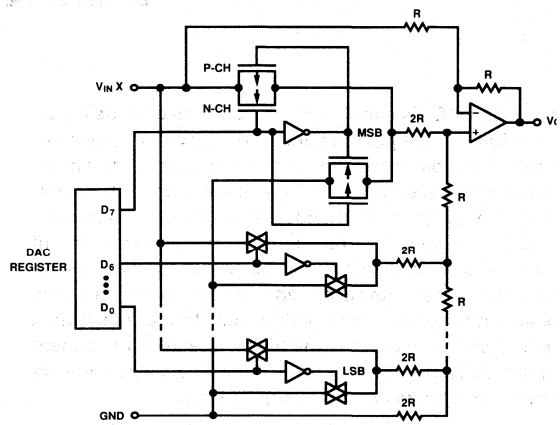

Figure 3. One Channel of the Four-Quadrant Multiplying DAC-8840.

bandwidth for ±3-V input signal levels while operating from ±5-V supplies. A typical signal channel has 0.01% total harmonic distortion and can slew at 2.5 V/μs. Because the output amplifier is connected in a differencing (push-pull) configuration, the gain for signals applied to V_{IN} can range from full-scale positive to full-scale negative, depending on the applied digital (offset binary) word. The magnitude of the binary word corresponds to the wiper position of a pot, with zero output at half-scale; a Preset control input sets all DACs to this "zero" position. Figure 4 describes this serial input CMOS octal D/A converter in greater detail.

The gain transfer function of a DAC-8840 channel is:

$$V_{OUT}(D) = (D/128 - 1)\ V_{IN}$$

where D is the value of the binary input, a decimal integer between 0 and 255. At full-scale, V_{OUT} = 127/128 × V_{IN}; when D = 128 (also the Preset condition), V_{OUT} equals zero volts; and when D = zero, $V_{OUT} = -V_{IN}$.

In the DAC-8840, eight DAC registers store the output state; they are updated from an internal serial-to-parallel shift register loaded from a standard 3-wire serial-input

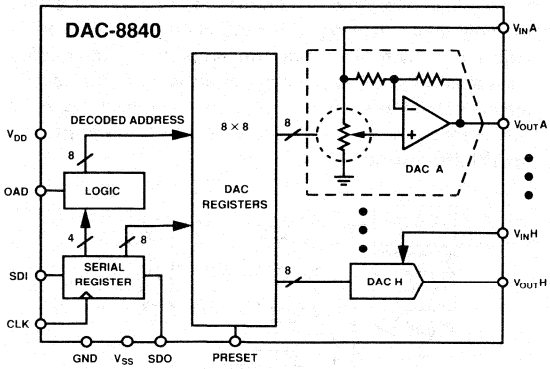

Figure 4. DAC-8840 Block Diagram. Note the 3-Wire Input and Serial Data Output Pin (SDO) for Daisy-Chaining Additional Packages.

digital interface. The data word clocked into the serial-input register (SDI) consists of 12 bits; the first four determine the address of the DAC register to be loaded with the 8 data-bits. A serial data output pin at the other end of the shift register (SDO) allows simple daisy-chaining in multiple DAC applications without additional external decoding logic (Figure 5). The fourth address bit, which decodes as a NOP for the package, makes it possible to select a single DAC in one of the packages to be updated when all the packages receive the common LD DAC strobe signal.

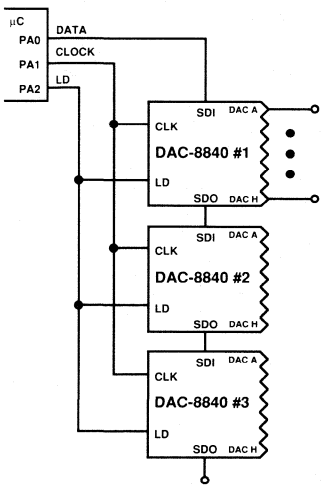

Figure 5. DAC-8840s in a Serial Daisy Chain Minimize Chip Decoders.

The DAC-8841, a mask option of the DAC-8840, offers an ideal octal DAC for +5-V single-supply applications. The DAC and amplifier of each channel are configured as shown in Figure 6, with the amplifier connected for a non-inverting gain of two. This configuration is a 2-quadrant multiplying arrangement with a 1-MHz bandwidth. AC signals applied to the V_{IN} terminal can be attenuated to zero or amplified by a factor of up to two, with 256 possible level settings from zero to 2 × (255/256) V_{IN} :

$$V_{OUT}(D) = 2 \times (D/256) \times (V_{IN} - V_{REFL}) + V_{REFL}$$

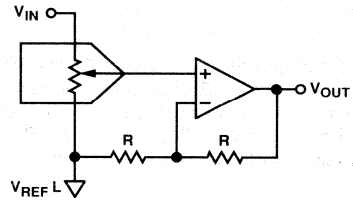

Figure 6. Internal Connections of the +5-V-only DAC-8841.

VARIABLE RESISTORS VERSUS TRIMDACS

From the above overview of the DAC-8800 and DAC-8840/41 TrimDACs, we can compare them to mechanically variable resistors (pots), reviewing the strengths and weaknesses of each.

Advantages of TrimDACs over Potentiometers: Better mechanical stability, improved product life, improved temperature coefficients, smaller size; computer control can eliminate technician costs; remote operation, constant output resistance, and low output resistance with low power dissipation.

Advantages of Potentiometers over TrimDACs: Voltage range usually much greater, no separate power supply required, simple human interface, no memory required, no "zipper noise" (the sound heard when using a DAC to adjust audio levels).

Another useful advantage of the potentiometer at present is a nonvolatile memory. That is, in a vibration-free environment, the wiper of the potentiometer stays where it was last set, even with the power off. The TrimDAC devices described here do not contain nonvolatile memory; for them, the required memory is generally supplied by system E^2PROM. Since in many of today's systems a low-cost high-density E^2PROM holds system set points for current time, date, mode, parameters and so on, it is an easy matter to share this nonvolatile memory with the TrimDAC calibration set points; they are reloaded at system powerup.

TYPICAL APPLICATIONS

In professional audio equipment, voltage-controlled amplifiers (VCA) are used to set gain, fade, pan and mix signals. The dc control inputs of these VCAs are ideally controlled by the DAC-8800 (Figure 7). The addition of the capacitor at the VCA voltage control port, VC, helps to limit the slew rate, reducing the clicking to a subaudible level. One DAC-8800 can control 8 channels of logarithmically set gain and attenuation levels.

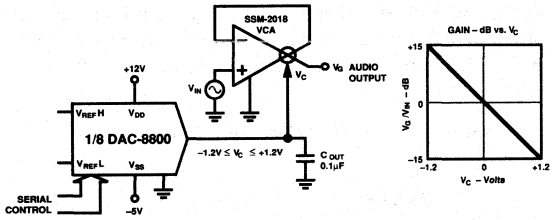

Figure 7. Setting Gain of a Voltage-Controlled Amplifier in Professional Audio Equipment. One DAC-8800 Can Serve 8 Channels. The Damping Capacitor at the Voltage-Control Point Minimizes Zipper Noise by Slowing Rates of Gain Change to Subaudio Frequencies.

Figure 8 shows a selection of output configurations of a DAC-8800, including simple buffers, summing circuits with coarse/fine control, and adding gain for increased output swing. A DAC-8800 can be used in system offset nulling by connecting its output to the summing node of any convenient op amp, using an appropriate value of summing resistance or a T-network.

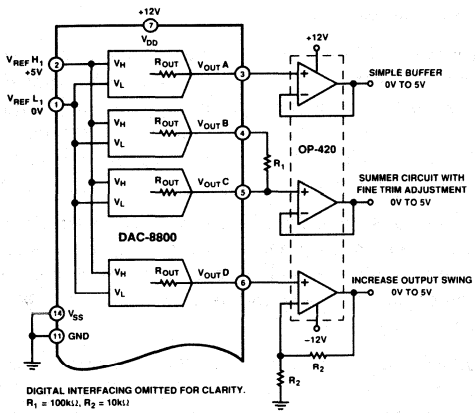

Figure 8. Some Ways of Buffering the DAC-8800 Output.

For video convergence and deflection control, especially in multi-sync displays, the DAC-8840 can be used to adjust the sawtooth waveform amplitudes, their reference bias voltages, and the parabolic waveforms used to linearize them as they are summed together to drive the CRT deflection. Figure 9 shows a block diagram of a typical arrangement.

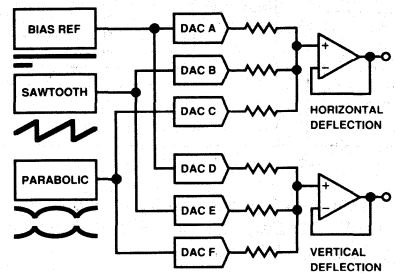

Figure 9. DAC-8840's Four-Quadrant Multiplying Capability Simplifies Amplitude Adjustment of Waveform Components in Video Deflection.

Availability

The 20-pin DAC-8800, and the 24-pin DAC-8840 and DAC-8841 are available for two temperature ranges—extended industrial (−40°C to +85°C) and military (−55°C to +125°C). Packaging includes plastic and ceramic DIPs and SOL surface-mount packages.

ANALOG DEVICES

AN-316
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

AD7224 Provides Programmable Voltages Over Varying Ranges

by Mike Byrne

This application note discusses some uses of the AD7224 in providing programmable output voltages over varying ranges. The AD7224 is a monolithic, voltage-mode CMOS DAC with output buffer amplifier housed in a 0.3 inch wide, 18-pin DIP.

The circuit of Figure 1a shows a common configuration for a potentiometer with the output voltage, V_{OUT}, varying from V_A to V_B. This output voltage is varied using the variable resistor VR1. The function performed by this circuit can be implemented using the circuit of Figure 1b. This circuit uses the AD7224 and just three external components to give a programmable output voltage varying from V_X to V_Y.

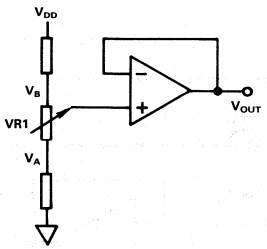

Figure 1a.

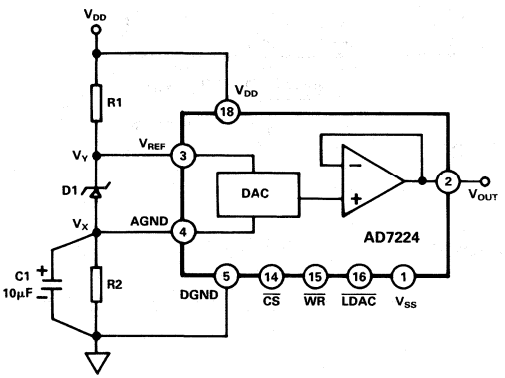

Figure 1b.

The operation of the circuit depends upon the principle that all the current which flows into the V_{REF} input of the AD7224 will flow out of AGND. Therefore, although the current flowing into the V_{REF} input will vary (because of varying input impedance with code) the current flowing through R2 will always remain constant. In practice some current flowing into the V_{REF} input will "leak" away from AGND. This leakage current varies with digital code. This variation is typically of the order of 3µA which represents an error of only 0.05% since the typical current which flows through the zener is 6mA. For example, if V_X is biased at 2V, with R2 = 330Ω the variation of V_X with digital code is typically less than 1mV.

An additional separate current, which does not flow into the V_{REF} input, flows from V_{DD} through internal circuitry and out of AGND. This current is independent of code and therefore does not cause variations in V_X over the input code range. It is in the range 50µA to 150µA and is dependent upon the $(V_Y–V_X)$ voltage; the lower the $(V_Y–V_X)$ voltage the smaller the current. It does mean that V_X, and hence V_Y, will be slightly larger than one would expect based on the current flowing through R1. For example, with R2 = 330Ω, V_{DD} = +12V and $(V_Y–V_X)$ = +6.1V it was found that V_X was 43mV larger than expected. Under the same conditions except $(V_Y–V_X)$ = 3.3V, V_X was 22mV larger than expected. In either case the increased V_X can easily be adjusted for by reducing R2.

The zener diode, D1, must be biased so that the current variation through it will not significantly affect its zener voltage. The worst case variation for the circuit of Figure 1b is 1mA over the input code range but typically this is less than 500µA. The total variation on V_Y, including the effect of the variation on V_X, was found to be less than 2mV. This was for various values of D1 with V_X = 2V and R2 = 330Ω.

The circuit of Figure 1b is not as flexible as that of 1a and a number of limitations exist. V_Y must be at least 4V below the V_{DD} supply voltage to ensure correct operation of the AD7224. Additionally, V_Y must always be positive with respect to V_X and in turn, V_X must always be positive with

respect to DGND. The entire circuit can, however, be operated in single supply with the V_{SS} pin of the AD7224 tied to DGND. When operating in single supply, the AGND should be biased over 2V above DGND to ensure specified linearity performance over the range of reference voltages. This limitation does not apply to differential linearity since the part remains monotonic over the range of reference voltages regardless of whether the part is operated in single or dual supplies. Table I shows different values for R1 and R2 for different zener diode voltages for a V_{DD} supply voltage of +12V. Capacitor C1 is used to eliminate spikes on AGND. These spikes occur when the digital input is changed because the zener diode presents a non-zero impedance to the V_{REF} input. Even without C1 they are not really a problem since they are masked out from the output by the slew rate of the output amplifier.

D1	R1	R2	V_X	V_Y
3.3V	820Ω	390Ω	2.8V	6.28V
3.9V	680Ω	680Ω	4V	8V
4.4V	680Ω	390Ω	2.78V	7.26V
5.6V	680Ω	330Ω	2.16V	7.9V

Table I.

An alternative method for generating a programmable output voltage over a varying range is shown in Figure 2. In this case the AD7224 is used with the LM10, which contains a reference, reference amplifier, A1, and output amplifier, A2, on one chip. The circuit contains just two narrow DIPs and four resistors. The technique involved is one of implicit feedback. For further information on implicit feedback circuits see page 39, "Nonlinear Circuits Handbook" published by Analog Devices.

The on-chip reference on the LM10 is a 200mV reference, V_R. Amplifier A1 is used as a buffer for this reference and has an output current source capability of 1mA. Note that the output voltage from the circuit is taken from the A2 output (which is in fact the reference voltage for the AD7224) and not from the V_{OUT} of the AD7224. The gain placed on amplifier A1 determines the lower end of the output voltage range at V_O. For example, if the gain of this amplifier is 5 the lower end of the V_O range will be 1V while if the gain is unity the lower end will be 200mV. The ratio of R4 and R3 then determines the span of the output voltage. For example, (with a gain of 5 on A1), if R4/R3 = 9 then the output voltage range, at V_O, will be 1V to 10V over the digital input code range of the AD7224. If R4/R3 = 4 then the output voltage range will be 1V to 5V.

The output voltage does not vary in a linear fashion with digital input code. The relationship is shown in Figure 3 for an R4/R3 ratio of 9 and a gain of 5 on A1. The expression for output voltage, V_O, is as follows:

$$V_O = V_R \cdot (1 + G1) \cdot \frac{(1 + G2)}{(1 + G2 \cdot D)}$$

where G1 = R2/R1
G2 = R4/R3

and D is a fractional representation of the digital word in the DAC register ($0 \leq D \leq 255/256$).

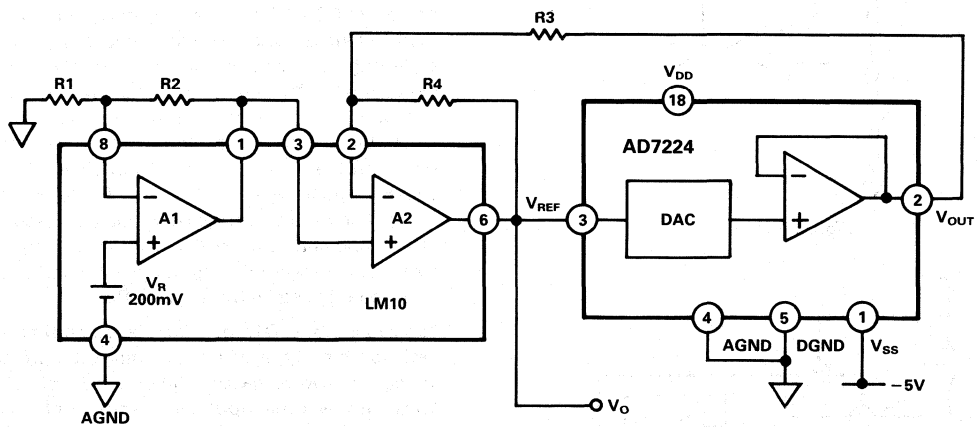

Figure 2. AD7224 with LM10

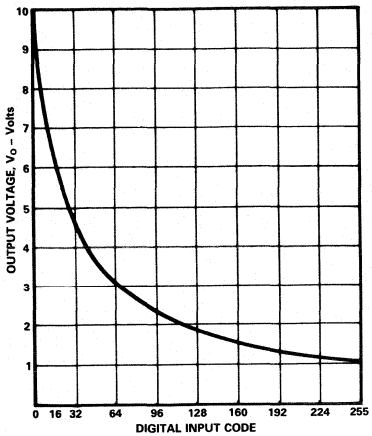

Figure 3. V_O Versus Digital Code

Another configuration, using the AD7224 with the LM10 and giving a similar output voltage range, is outlined in Figure 4. In this case A1 is used to drive the V_{REF} of the AD7224 directly. A2 is then used in conjunction with a series pass transistor, T1, to provide current boosting with up to 100mA being provided across R_L. The output voltage range with the component values given in Figure 4, is, once again, 1V to 10V. The positive supply and return paths for the transistor and load have to be kept separate from the supply paths for the rest of the circuit to avoid errors caused by resistive drops with large currents flowing. Varying the ratios of R6 to R5 and R8 to R7 changes the output voltage range with the expression for the output voltage, V_O, being

$$V_O = V_R \cdot \frac{(1 + G1) \cdot (1 + G2)}{(1 + G1 \cdot D)}$$

where G1 = R6/R5
G2 = R8/R7

and D is a fractional representation of the digital word in the DAC register.

In both the circuit of Figure 2 and Figure 4 the AD7224 is shown operating from dual supplies (i.e., $V_{SS} = -5V$). The reason for this is that the AD7224 loses its output sink capability with output voltages near AGND when operated in single supply. In dual supply operation it maintains a 400μA output sink capability over the entire output voltage range. Resistor values for R3 in Figure 2 and R5 in Figure 4 should be chosen so that this 400μA output sink current is not exceeded. If single supply operation is required, the V_{SS} of the AD7224 can be tied to DGND with the AGND of the AD7224 (and all other AGND points in the circuit) biased to 2V. In this case the output will again have its full sink capability over the output range.

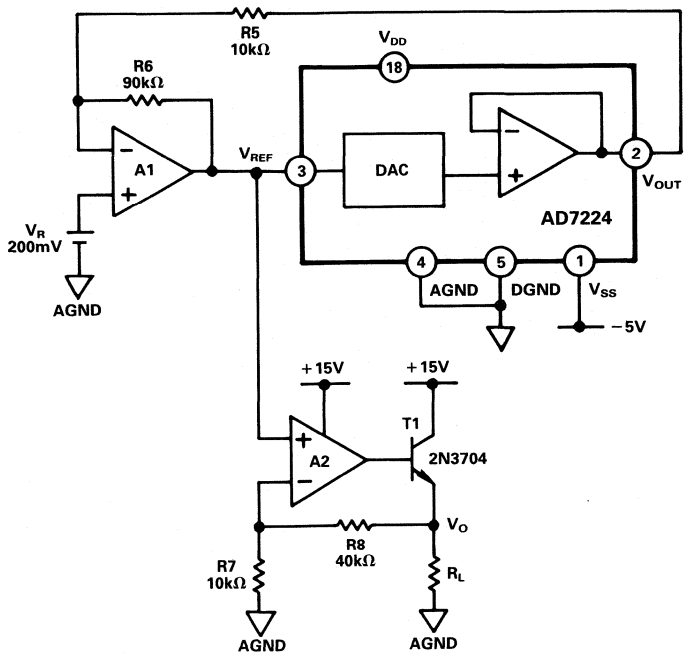

Figure 4. Current-Boost Configuration

ANALOG DEVICES

AN-317
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Circuit Applications of the AD7226 Quad CMOS DAC

by Mike Byrne

The AD7226 is a monolithic quad 8-bit CMOS DAC packaged in a 20-pin DIP. Each DAC output is buffered by a CMOS amplifier which is capable of developing +10V across a 2kΩ resistor. Data is loaded from a common 8-bit data bus into one of the on-chip latches provided for each individual DAC (see Figure 1).

The AD7226 has certain features which make the part a unique and extremely useful device. Firstly, housing four DACs plus interface logic and output buffer amplifiers in a 20-pin package allows for substantial savings in circuit board space requirements and complexity. Additionally, the converters are operated in the voltage-mode which allows single supply operation for the part, with a unique DAC switch pair arrangement allowing an extended reference range not previously available with voltage-mode converters. Since all four DACs are fabricated on the same chip, precise matching and tracking between them is inherent.

This application note discusses some uses of the AD7226 in dc or voltage setting type applications. Operation of some of these circuits relies on the inherent DAC-to-DAC matching provided by the AD7226. Other circuits benefit from the circuit board space saving offered by the AD7226 and from its ability to operate with a single power supply.

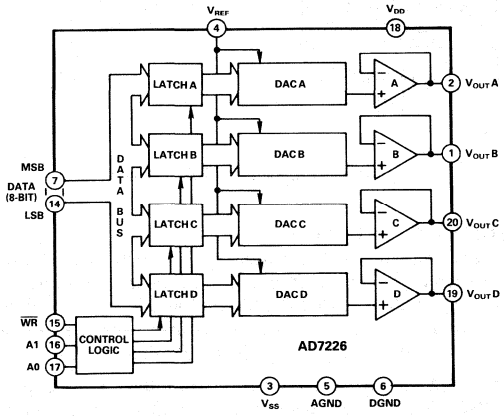

Figure 1. AD7226 Functional Diagram

This application note does not discuss the basic operation of the AD7226; consult the data sheet for this information.

AD7226 APPLICATIONS DISCUSSED IN THIS NOTE

1. Programmable Offset Adjust of Operational Amplifiers using one-channel of the quad DAC per op amp.
2. Set-Point Controller Circuit which allows fine adjust over a wide voltage range.
3. Self-Programmable Reference Voltage using one DAC of the AD7226.
4. Staircase Window Comparator Circuit for Measurement of Threshold Values for a TTL device.
5. V_{SS} Generation Circuits to allow dual supply operation of the AD7226 from a single power supply.
6. 5V Single Supply results showing excellent Differential Nonlinearity performance.

PROGRAMMABLE OFFSET ADJUST

The AD7226 can be used to provide programmable input offset voltage adjustment for operational amplifiers. The circuit configuration used to achieve this is shown in Figure 2. Each output of the AD7226 can be used to trim the input offset voltage of one operational amplifier. This means that programmable offset adjustment can be achieved on four operational amplifiers by the addition of just one 20-pin device and some extra resistors.

The circuit configuration uses the input offset voltage nulling pins provided on most operational amplifiers. Resistor R2, tied to +10V, provides a fixed bias current to one offset node. The output of the D/A converter is connected via R1 to provide a variable bias current to the other offset node. Therefore, changing the code on the D/A converter provides offset adjust for the operational amplifier. For symmetrical adjustment, the bias current through R2 should equal the current in the other offset node with the half-full scale code (i.e. 10000000) on the D/A converter.

Resistors R1 and R2 are chosen such that enough current variation over the DAC code range is given to provide the required range of offset adjustment for the op amp in question. Reducing the values of R1 and R2 increases the range of offset which can be trimmed, with a corresponding reduction in resolution.

The method of programmable offset adjustment can be used with most operational amplifiers which have offset

nulling pins provided. Table I shows some results achieved with four popular op amps. The table shows the typical range of offset values which can be trimmed using the given values of resistors R1 and R2. It also gives typical figures for final values of offset achieved after using the method outlined above.

The circuit configuration of Figure 2 ensures that, for increasing code on the D/A converter, the op amp offset goes more positive when using the first three op amps of Table I. For the TL091 and other op amps of the family (TL061, TL071, etc.) the trim terminals must be swopped (as indicated in Figure 3) to ensure that for increasing code the output offset will go in the positive direction.

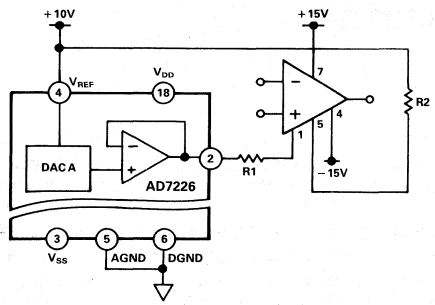

Figure 2. Offset Adjust Using AD7226

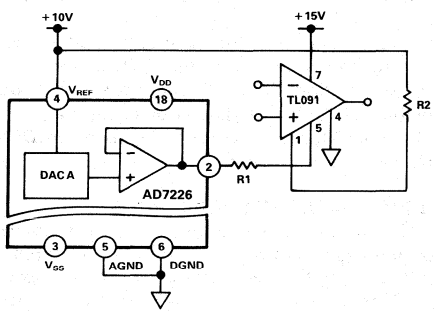

Figure 3. Offset Adjust for TL091

The trimming of op amp offsets in this manner increases the offset drift temperature coefficient of the op amp. For example, using the AD544 this increase will be 3μV/°C per millivolt of offset adjustment assuming 0ppm/°C temperature coefficient for the external resistors R1 and R2. However, the same drift would have been introduced had an external trimpot been used to trim the offset. The method outlined above has the advantage of being programmable and, therefore, any drift over temperature can be adjusted out during a periodic calibration cycle.

The programmable input offset voltage adjust can be used to deliberately introduce input offset voltage into the op amp. This could be useful in a system context where programmably introducing offset at a node in the system could give the desired value at the output of the system.

The first three op amps shown in Table I operate from dual supplies. The TL091 op amp is specified to operate at 5V single supply. The AD7226 can operate at 5V single supply and will remain monotonic to 8-bits under these conditions. This is basically all that is required for the configurations of Figure 2 and Figure 3 to function. The 5V single supply operation of the AD7226 is discussed later in this application note.

Op Amp	R2 (kΩ)	R1 (kΩ)	Range (mV)	Final Offset (μV)
AD741	1200	1000	±6.75	−14.5
AD544	620	500	±2.75	3.3
AD542	470	360	±2.4	4.6
TL091*	1000	500	+9.0	2.5

*Operating in Single Supply V_{SS} = 0V V_{DD} = +15V.
All other op amps dual supply ±15V.

Table I. Typical Op Amp Offset Results

Some operational amplifiers, especially duals and quads, do not have trim terminals available to the user. The AD7226 can still be used to provide offset adjustment by programmably varying the voltage at the required op amp input terminal. One such configuration is outlined in Figure 4. The noninverting input of the op amp is offset in a negative direction via R4 to −15V. In a similar fashion to the previous method, for symmetrical adjustment the current through R4 should equal the current through R3 with the half-full scale code on the D/A converter. Once again increasing the digital code will vary the offset in a positive direction. The resistor configuration will be seen as a low impedance by the noninverting input of the op amp to prevent noise injection. The circuit configuration does not affect the gain or transfer function of the op amp. Using the TL044 quad op amp, with the components values given in Figure 4, the typical range of offset which could be trimmed using the method outlined for Figure 3 was ±5.5mV. A typical figure for the final offset value achieved was 10μV. Similar techniques can be used for other op amp configurations.

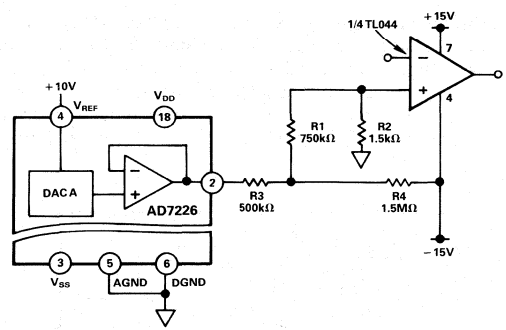

Figure 4. Alternative Offset Adjust

SET-POINT CONTROLLER

The set-point controller circuit of Figure 5 allows programmable fine and coarse adjust of the output voltage, V_{OUT}, over a 200V range. The circuit has a fine adjust resolution of 8mV which means that an output voltage can be set in the range −100V to +100V to within ±4mV.

The circuit uses DAC A as a programmable coarse adjust and DAC B as a programmable fine adjust of the output voltage. DAC A has an effective output voltage range of

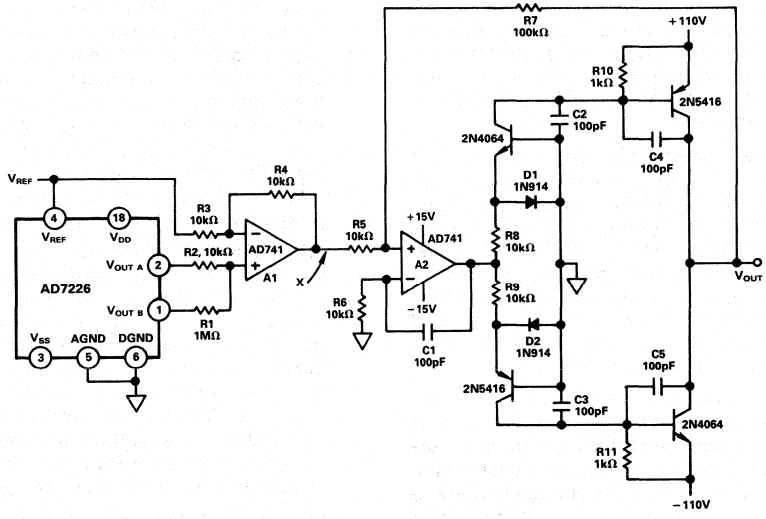

Figure 5. Set-Point Controller Circuit

−100V to +100V over the digital input code range giving an LSB size of 800mV for this coarse adjust DAC. DAC B has an effective output voltage range, at V_{OUT}, of 2V giving an LSB size of 8mV for this fine adjust DAC.

Amplifier A1 sums the output of both DACs and provides a bipolar output voltage, V_X, at point X. Amplifier A2, with the additional level shifting circuitry, provides a gain of 10 between V_X and V_{OUT}. In general, the output voltage V_{OUT} can be expressed as:

$$V_{OUT} = \frac{R_7}{R_5}\left\{ 2 \cdot V_{OUTA} \cdot \left(\frac{R_1}{R_1 + R_2} \right) + 2 \cdot V_{OUTB} \cdot \left(\frac{R_2}{R_1 + R_2} \right) - \left(V_{REF} \right) \right\}$$

For the component values given in the circuit of Figure 5 this can be simplified to give the expression for V_{OUT} as:

$$V_{OUT} = 10\left\{ V_{OUTA} \cdot \left(\frac{2000}{1010} \right) + V_{OUTB} \cdot \left(\frac{20}{1010} \right) - \left(V_{REF} \right) \right\}$$

The resolution of the fine adjust circuitry can be changed by varying R2. This is done without any significant effect on the overall output voltage range. Adjusting R7 will change the output voltage range, varying the resolution of both the fine adjust and coarse adjust DACs.

This circuit of Figure 5 is capable of developing ±100V across a 2.7kΩ load. It is useful in set-point controller applications or in a programmable power supply. The other two channels of the AD7226 can be used to perform normal D/A converter functions or configured as above to provide a second set-point controller.

SELF-PROGRAMMABLE REFERENCE

The circuit of Figure 6 shows how one D/A converter of the AD7226, in this case DAC A, may be used in a feedback configuration to provide a programmable reference voltage for itself and the other three converters. The relationship of V_{REF} to V_{IN} is dependant upon digital code and upon the ratio of resistors R1 and R2. It can be expressed by the formula

$$V_{REF} = \frac{(1+G)}{1+G \cdot D_A} \cdot V_{IN}$$

where $G = R2/R1$

and D_A is a fractional representation of the digital word in latch A ($0 \leq D_A \leq 255/256$).

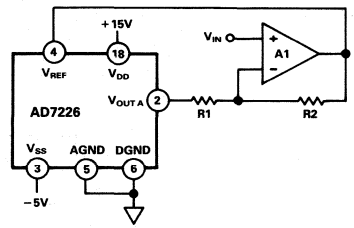

Figure 6. Self-Programmable Reference

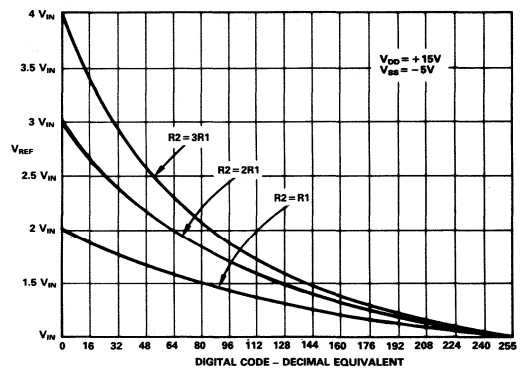

Figure 7. Variation of V_{REF} with Feedback Configuration

Figure 7 shows typical plots of V_{REF} versus digital code for three different values of R2. With $V_{IN} = 2.5V$ and $R2 = 3R1$

the voltage at the output of A1 (i.e. reference voltage for the AD7226) will vary between +2.5V and +10V over the digital input code range. This gives an effective 10-bit dynamic range to the other three D/A converter outputs with the minimum LSB size being 10mV over a potential full scale output range of 0 to 10V.

The circuit of Figure 6 should only be used when the AD7226 is operated from dual supplies (i.e. $V_{SS} = -5V$). One reason for this is that the AD7226 is specified at +10V reference voltage only for single supply operation. More importantly, however, is the fact that the AD7226 has reduced current sink capability at output voltages near 0V when used in single supply (see AD7226 data sheet). This means that the circuit would not operate correctly at lower values of digital input code. For correct operation with dual supplies R1 must be greater than 6.8kΩ.

Different configurations of A1 and resistors can be devised to give other ranges of reference voltage. It must be noted that varying the voltage at the V_{REF} pin varies the reference voltage for all the converters. Whatever configuration devised must ensure that the voltage at the V_{REF} pin of the AD7226 must never go negative with respect to either AGND or DGND.

THRESHOLD TESTING

In test systems and other applications it is necessary to determine whether some unknown voltage lies within certain limits. In these cases the AD7226 can be used with external comparators or op amps to provide a programmable staircase window comparator. The voltage levels on the comparator and the window sizes are determined by the digital code in the latches of the D/A converters. Figure 8 shows one such application where the AD7226 can be used to measure the threshold levels of a TTL device under test.

V_{OUTA} and V_{OUTB} of the AD7226 in addition to the six external comparators, form a staircase window comparator. Each adjacent pair of comparators forms a window of programmable size. When a voltage lies within a window, the output from that window goes high . Window 1 is set with an upper limit of $V_{CC} = 5V$ and a lower limit of $V_{OUTA} = 2.4V$. Window 2 has an upper limit of 2.4V and a lower limit of 0.4V, while Window 3 has an upper limit of 0.4V and a lower limit of AGND = 0V. These levels can be programmably set and varied as required.

The staircase window comparator circuit is used in testing the output threshold levels of the device under test (DUT1). When D8 goes high, SW1 is set so that 400μA is sourced from the output of DUT1. The V_{OH} of DUT1 is applied to the staircase window comparator. For the part to pass its V_{OH} test, the output from Window 1 must go high. Any other window going high indicates that the device fails. When D8 goes low, SW1 is switched so that the output of DUT1 will sink 16mA and the V_{OL} of DUT1 is measured. Window 3 of the staircase window comparator must go high for the device to pass its V_{OL} test.

Window 2 is not really necessary in the application shown. However, it has been included to demonstrate the non-overlapping staircase comparator configuration. V_{OUTC} and V_{OUTD} of the AD7226, along with additional comparators, can be used to extend the staircase window comparator to give five programmable non-overlapping windows. The window structure for these five non-overlapping windows is shown in Figure 9 with the upper limit of the staircase now at V_{REF}. If overlapping windows are required, the configuration can easily be adapted as shown in Figure 10a. In this case all four outputs are used with the external comparators to provide three overlapping windows. The window structure for this overlapping window configuration is shown in Figure 10b.

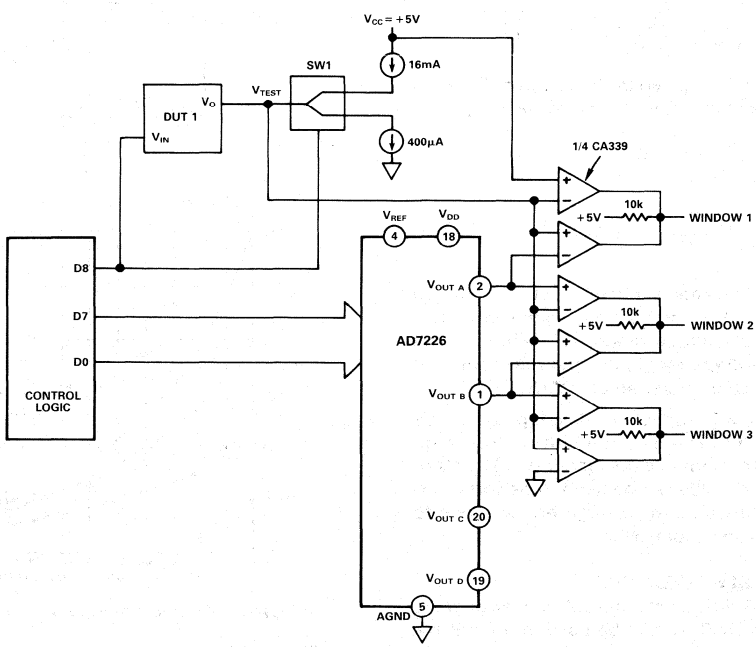

Figure 8. Threshold Testing

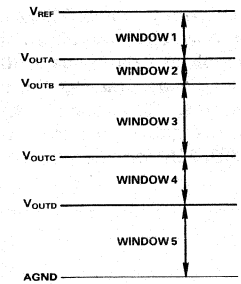

Figure 9. Window Structure

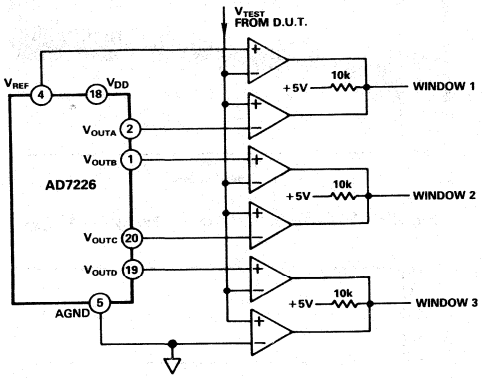

Figure 10a. Overlapping Windows

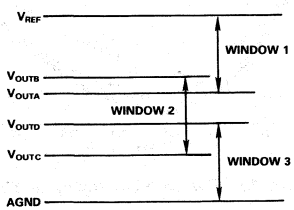

Figure 10b. Window Structure

V_{SS} GENERATION

Operating the AD7226 from dual supplies results in enhanced performance over single supply operation on a number of parameters. The negative V_{SS} gives additional headroom to the output amplifier which results in improved negative-going settling-time, improved zero code error performance and an extended input reference range. Some applications may require this enhanced performance but may only have a single power supply rail available. The following circuits show some methods of generating a negative V_{SS} from a single power supply rail for the AD7226.

Figure 11 shows one such method of generating a negative supply using one CD4049, operated from a V_{DD} of +15V. Two inverters of the hex inverter chip are used as an oscillator. The other four inverters are paralleled and used as buffers for higher output current. The square-wave output is level translated to a negative-going signal, then rectified and filtered. The circuit configuration shown will provide an output voltage of −5.1V for current loadings in the range of 0.5mA to 8mA. This will satisfy the AD7226 I_{SS} requirement over the commercial operating temperature range. Noise spikes which are generated on this V_{SS} line from the clock can be considerably reduced by decoupling the V_{DD} supply line to the CD4049.

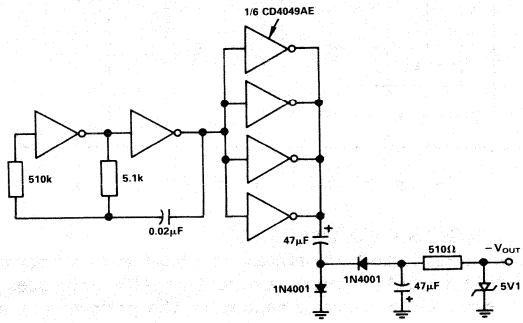

Figure 11. V_{SS} Generation Circuit

An alternative method of generating a negative supply from a positive rail is to use Analog Devices AD7560, a DC-DC voltage converter. This can provide a −5V supply from a +5V rail. The circuit configuration used to achieve this is outlined in Figure 12.

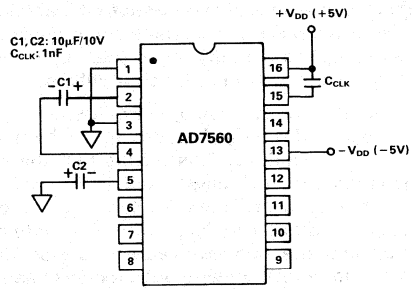

Figure 12. AD7560 Giving −5V

Some applications may require a +5V reference but may only have a single supply rail available. The AD7226 is specified at +5V reference when used with dual supplies only. The circuit of Figure 13 could prove useful in such applications. It provides a V_{SS} of −5V, the +5V reference and V_{DD} of +11.4V to +16.5V from a single +16.4V to +21.5V power supply rail (in battery applications two 9V batteries in series). The AD584 is a pin-programmable precision voltage reference. The AD380 op-amp buffer establishes the "ground" for the AD7226 midway between 0 and +10V. Hence, pin 1 of the AD584 can be used as the +5V reference for the AD7226. The V− of the input signal is used to provide a V_{SS} of −5V. If the input voltage can exceed the limits above (i.e., +16.4V to +21.5V) a zener diode may be required to provide a regulated supply voltage to the V_{DD} pin of the AD7226. The voltage at the V_{DD} pin must never exceed +17V with respect to the AD7226 "ground".

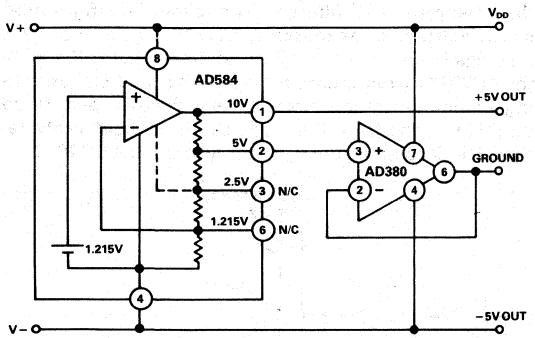

Figure 13. Voltage Splitting Using AD584

5V SINGLE SUPPLY OPERATION

The AD7226 can be operated from a single +5V power supply rail, but because the output amplifier now loses a considerable amount of headroom, the performance of the part is degraded. However, one important parameter which retains its specified performance is differential nonlinearity. At a single +5V supply this remains within ±1LSB which ensures that the AD7226 will remain monotonic over the output voltage range.

This monotonic operation makes the AD7226, at single +5V supply, suitable for applications where the absolute value (or accuracy) of the output voltage is not important but where it is essential that the output increases as the digital code increases. An example of this type of application has already been outlined in this application note: trimming the offset of operational amplifiers.

The required overhead voltage of 4V between the reference voltage and the V_{DD} input must still be maintained under these conditions. This means that the reference voltages must remain below 1V to ensure monotonic operation. At these low reference voltages, the output offset voltage is obviously very large with respect to the input reference and the relative accuracy is also degraded. However, the output voltage does increase as the input digital code varies from 0 to 255. This can be seen from the plot of Figure 14 which shows differential nonlinearity for single +5V supply at a reference of 600mV. Figure 15 shows a plot of relative accuracy for the same part under the same conditions. Additionally, the digital input threshold levels and digital input currents are not affected by operating the AD7226 from the single +5V supply rail.

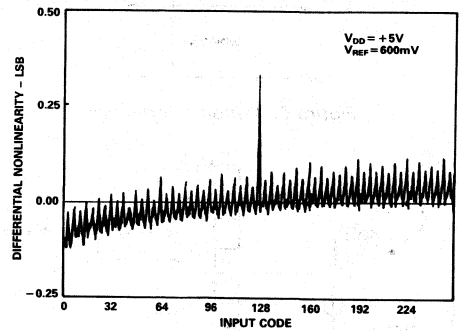

Figure 14. Differential Nonlinearity at 5V V_{DD}

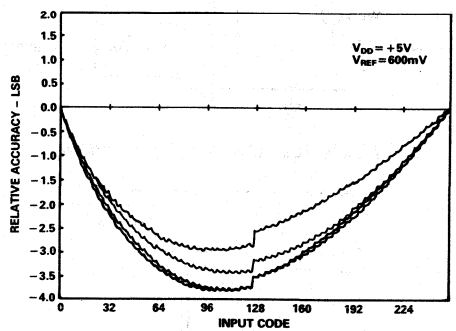

Figure 15. Relative Accuracy at 5V V_{DD}

ANALOG DEVICES

AN-318
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

AD7528 Dual 8-Bit CMOS DAC Application Note

by Paul Toomey and Bill Hunt

INTRODUCTION

The AD7528 is a monolithic dual 8-bit CMOS DAC packaged in a 20-pin DIP. Each DAC has its own 8-bit data latch which loads data from a common 8-bit data bus (see Figure 1). Since both DACs are fabricated on the same chip, precise matching and tracking between DACs is inherent. This property of the AD7528 dual DAC, along with the P.C. board space saving it allows, makes the AD7528 a unique and extremely useful device.

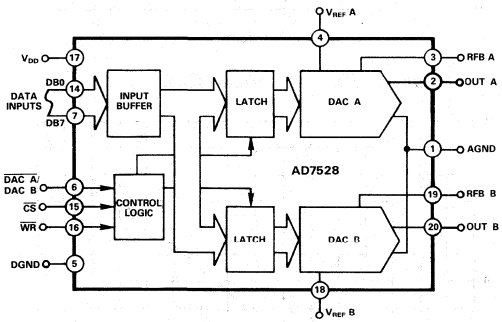

Figure 1. AD7528 Functional Diagram

This note discusses the AD7528 applications circuits listed below. Several of these circuits rely on the DAC to DAC matching provided by the AD7528. All of the circuits benefit from the high packing density the AD7528 allows, especially when used with dual and quad op-amps such as the AD644 or TL074. Not discussed in this note are basic details of AD7528 operation, consult the data sheet for this information.

AD7528 APPLICATIONS DISCUSSED IN THIS NOTE

1. State-variable filter (S.V.F.) with programmable center frequency, selectivity and gain.
2. Programmable sine wave oscillator with linear control.
3. Function fitting sine wave synthesizer with amplitude control facility and programmable phase shift.
4. Programmable voltage/current source, unipolar and bipolar circuits.
5. Programmable gain amplifier with no trimpots.
6. Programmable waveform generator for vector scan CRT displays.
7. AD7528 single-supply operation circuits for low budget applications requiring multiple analog outputs

STATE VARIABLE FILTER WITH PROGRAMMABLE CENTER FREQUENCY, SELECTIVITY (Q) AND GAIN

The state variable filter (or universal filter as it is often called) is a convenient 2nd order filter block. It provides simultaneous low-pass, high-pass and bandpass outputs. All filter parameters can be readily adjusted. Figure 2 shows a typical filter circuit with expressions for center frequency, Q and gain for the bandpass output.

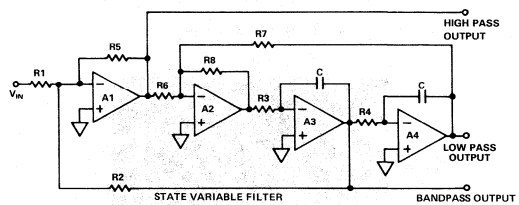

Figure 2. State Variable Filter

BANDPASS TRANSFER FUNCTION

$$\frac{V_{OUT}}{V_{IN}}(f) = \frac{A_O}{1 + jQ\left[\frac{f}{f_o} - \frac{f_o}{f}\right]}$$

$$f_o = \frac{1}{2\pi R3\,C} \cdot \sqrt{\frac{R8}{R7}} \quad (\text{For } R3 = R4)$$

$$Q = \frac{R6}{R8} \cdot \frac{R2}{R5} \cdot \sqrt{\frac{R8}{R7}}$$

$$A_O = -\frac{R2}{R1}$$

Where f = frequency of V_{IN}

A_O = gain at f = f_o

Q = circuit Q factor, i.e., $\frac{f_o}{\text{3dB Bandwidth}}$

f_o = resonant frequency.

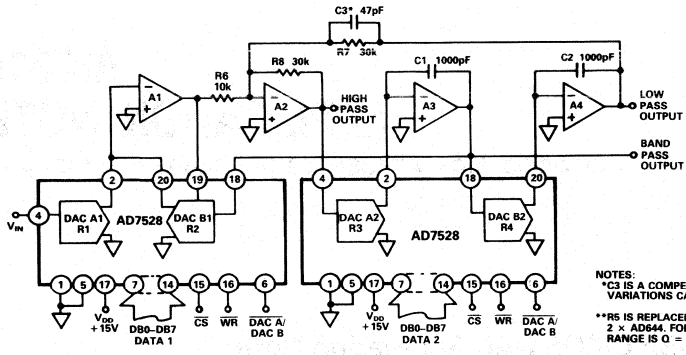

CIRCUIT EQUATIONS:

$C1 = C2,\ R3 = R4,\ R7 = R8$

$$f_o = \frac{1}{2\pi R3C1}$$

$$Q = \frac{R6}{R8} \cdot \frac{R2}{R5^{**}}$$

$$A_O = -\frac{R2}{R1} \text{ For Bandpass Output}$$

DAC EQUIVALENT RESISTANCE EQUALS

$$\frac{256 \times (\text{DAC LADDER RESISTANCE})}{\text{DAC DIGITAL CODE (DECIMAL)}}$$

NOTES:
*C3 IS A COMPENSATION CAPACITOR TO ELIMINATE Q AND GAIN VARIATIONS CAUSED BY AMPLIFIER GAIN BANDWIDTH LIMITATIONS

**R5 IS REPLACED BY DAC B1 INTERNAL RFB = 11kΩ, OP-AMPS ARE 2 × AD644. FOR COMPONENT VALUES SHOWN PROGRAMMABLE RANGE IS Q = 0.3 TO 4.5, f_o = 0 to 15kHz.

Figure 3. Digitally Controlled State Variable Filter

Introducing the DACs as Control Elements:

By replacing R1, R2 and R3, R4 with matched DAC pairs the filter parameters can be made programmable as shown in Figure 3. DAC A1 and DAC B1 control filter gain and Q, while DAC A2 and DAC B2 control center frequency (f_o). For the component values shown the programmable Q range is from 0.3 to 4.5 and is independent of f_o (see Figure 4). Center frequency (f_o) is programmable from 0 to 15kHz (see Figure 5) and is independent of Q.

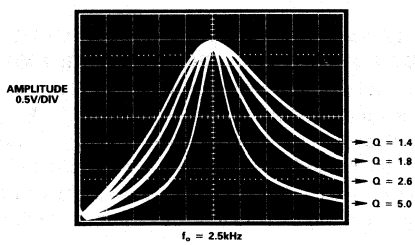

Filter 4. Filter Q Variation

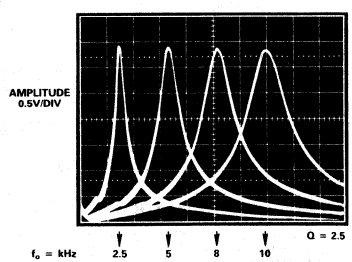

Figure 5. Filter f_o Variation

Programming

The graph in Figure 6 shows how the circuit Q varies with DAC B1 code and Figure 7 shows how the center frequency varies with DAC 2 (A and B) code for the component values given in Figure 3. Gain variation alone is accomplished by changing DAC A1 code. Unity gain occurs when the data in DAC A1 and DAC B1 latches is identical. Since the AD7528's logic inputs are TTL or CMOS compatible, the DACs are readily interfaced to most microprocessors, (see data sheet for hookups) thus providing an ideal microprocessor-to-filter interface.

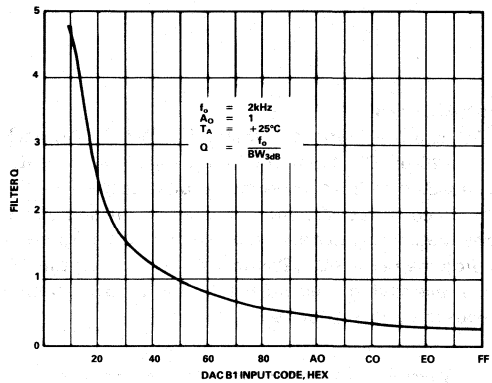

Figure 6. Filter Q Variation with Code

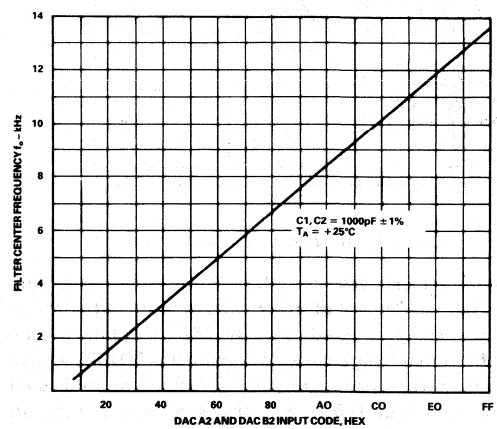

Figure 7. Filter f_o Variation with Code

PROGRAMMABLE SINE WAVE OSCILLATOR WITH LINEAR CONTROL

Frequency control of many oscillator circuits can be accomplished using two ganged potentiometers. However, the two potentiometers must track precisely over their full temperature range if a linear response is required. Figure 8 shows a high performance sine-wave oscillator realized using a state-variable filter. The frequency of oscillation is set by ganged potentiometers P1 and P2.

igure 9 shows the same circuit with P1 and P2 replaced y the AD7528 matched pairs.

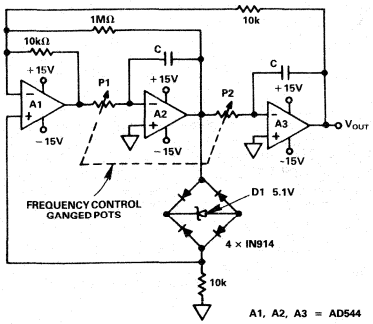

Figure 8. Sine Wave Oscillator Using a State Variable Filter

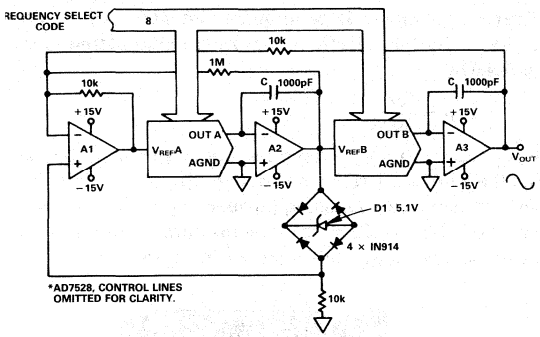

Figure 9. Programmable Sine Wave Oscillator Using a State Variable Filter and a Dual DAC

The equivalent resistance of each DAC, as seen by op-amps A2 and A3 varies with input code from infinity at code 00 Hex (0000 0000) to a minimum of $\approx$ 11kΩ (DAC ladder resistance) at code FF HEX (1111 1111).

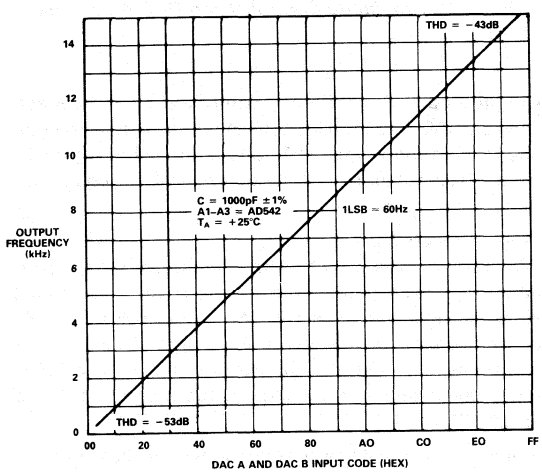

Figure 10. Frequency vs. DAC Code for Programmable Sine Wave Oscillator (Figure 9)

Loading each DAC latch with the same code provides a linear code versus frequency relationship as shown in Figure 10. The frequency of oscillation can be expressed as:

$$\text{Output Frequency} = \frac{N}{256\,(2\pi RC)}\ \text{Hz}$$

Where R = DAC ladder resistance i.e. V_{REF} input resistance.
C = is as shown in Figure 9.
N = decimal representation of digital input code. For example, N = 128 for input code 10000000.

For the component values given in Figure 9, output frequency is variable from 0 to 15kHz. Output amplitude is controlled by the zener diode D1. Total harmonic distortion for the circuit shown is −53dB at low frequencies (1kHz) and −43dB at higher frequencies (14kHz). Note that a cosine output is also available at the output of op-amp A2.

FUNCTION FITTING SINE WAVE SYNTHESIZER

In this application the multiplying capabilities of the two CMOS DACs are used to synthesize a sine wave based on a function fitting technique. This allows very low frequency, highly stable sine waves to be generated.

8

Function Fitting:

Function fitting is a technique for translating a mathematical or empirical relationship from one medium (such as a mathematical formula) to another medium (usually a physically realizable device or system). This application uses the dual DAC to implement a one quadrant sin X approximation in the form of the quadratic polynomial.

$$Y = 1.828N - 0.828N^2 \text{ where } 0 \le N \le 1 \text{ and } N = \frac{2}{\pi} X$$

The graph of Figure 11 shows the relationship between sin X and its quadratic approximation given above. The

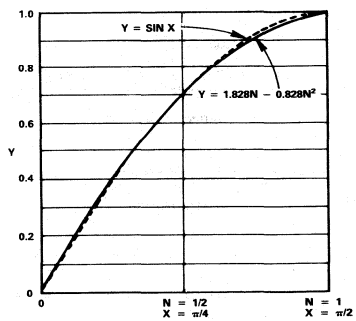

Figure 11. Relationship Between Sin X and its Quadratic Approximation

circuit of Figure 12 implements the function by ramping N up and down using an up/down counter, and switching the circuit output polarity. This generates sin X in four stages (see Figure 13).

Circuit Operation: (Figure 12)

An input clock drives the up/down counter in real time. The counter is connected so that it counts up and down continuously, providing an output pulse at "borrow" every time it reaches the all zeros count.

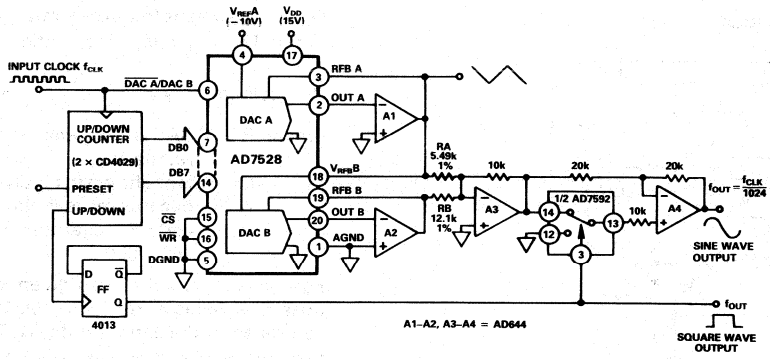

Figure 12. Function Fitting Sine Wave Generator

DAC A produces a triangle waveform consisting of two ramps of opposite slope, each generated in 256 steps at op-amp A1 output. This is the N variable.

DAC B is driven with the same digital word as DAC A. Its reference input is driven by op-amp A1, thus DAC B multiplies the digital version of N by the analog version of N to produce an output from op-amp A2 of $-N^2$ (negative sign is due to inversion through A2).

Since the N and $-N^2$ signals are of opposite polarity, the $Y = 1.828N - 0.828N^2$ expression is implemented by summing N and $-N^2$ signals in the correct ratios determined by RA and RB.

The analog switch, in conjunction with op-amp A4, changes the circuit's output polarity at one-half the triangle wave frequency, thus producing both positive and negative halves of the sine wave.

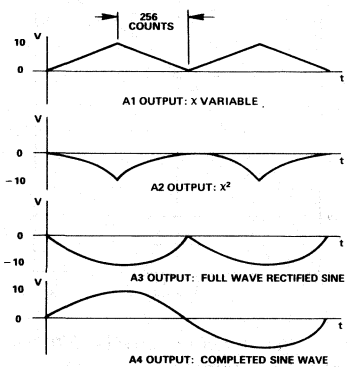

Figure 13. Sine Wave Synthesis Using Function Fitting

Distortion:

Distortion in the output sine wave is a function of the quadratic approximation fit to the sine curve. Errors in the values of R_A and R_B will, therefore, contribute directly to distortion. If R_A is made adjustable over a small range, it can be trimmed to minimize distortion. Distortion was measured for the circuit of Figure 12 at –33dB and was constant over sine-wave frequency 0–2.5kHz.

The circuit of Figure 12 generates constant amplitude sine waves in the 0–2.5kHz frequency range. Output frequency is given by

$$f_{OUT} = \frac{f\,clk}{1024}$$

It is possible to obtain rapid frequency sweeping by varing the input clock rate. The counter up/down output provides a useful zero crossing pulse. By applying an ac signal to the DAC A reference input, the output sine wave can be amplitude-modulated as shown in Figure 14.

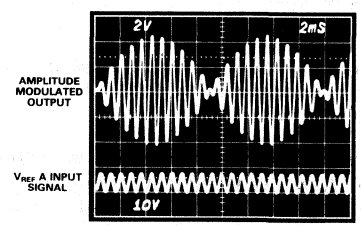

Figure 14. Amplitude Control Using DAC A Reference Input

Output amplitude is directly controlled by the voltage level on DAC A reference input. Ac or dc signals may be used within the range ±10 volts, 0 to 10kHz. Figure 14 shows the amplitude of a 1kHz sine wave being controlled by a 55Hz sine wave.

Programmable Phase

Two sine waves with 0 to 360°programmable phase relationship can be generated by running two of the circuits, shown in Figure 12, from the same input clock source. By allowing one circuit to start running from zero count a preset number of clock cycles before the other, a phase difference between their output sine waves is introduced (see Figure 15).

Since each complete cycle is generated by 1024 clock cycles, phase steps of 360/1024 degrees are programmable. Note also that the phase difference is independent of output frequency.

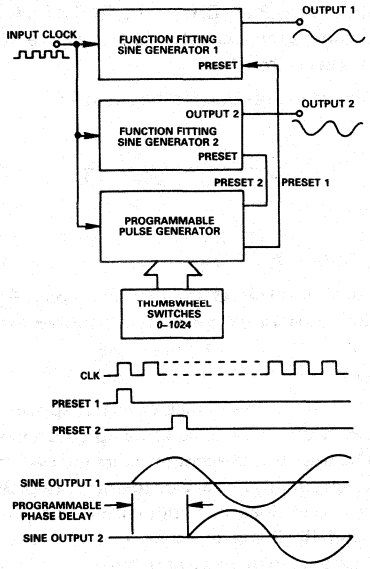

Figure 15. Programmable Phase Relationship

PROGRAMMABLE VOLTAGE/CURRENT SOURCE USING DUAL DAC AD7528

The circuit in Figure 16 is that of a unipolar V/I source. A negative reference is required for a positive output voltage.

The circuit provides;

a) A programmable output voltage

$$V_{OUT} = +V_{REF}A \frac{N_A}{256}, \text{ for } N_A = 0 \text{ to } 255.$$

provided the current limit is not exceeded.

b) In the voltage mode, a programmable load current limit given by:

$$I_{OUT(max)} = V_{REF}B \cdot \frac{R_1}{R_2} \cdot \frac{1}{R_{S2}} \cdot \frac{N_B}{256}$$

for $N_B = 0$ to 255.

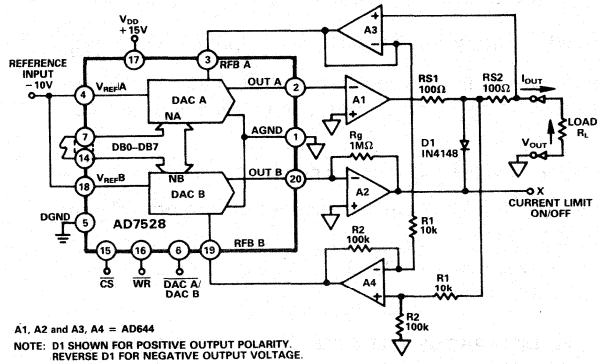

Figure 16. Programmable Voltage/Current Source V_{OUT} = 0 to +10V, I_{OUT} = 0 to +10mA

(c) A constant current feature by setting $N_A = 255$ i.e., maximum output voltage capability, and limiting the load resistance value R_L such that

$$I_{OUT(max)} \cdot R_L < \frac{255}{256} V_{REF}A$$

$$\text{with } I_{OUT\ (max)} = V_{REF}B \cdot \frac{R_1}{R_2} \cdot \frac{1}{R_{S2}} \cdot \frac{N_B}{256}$$

as in (b).

A useful feature of the circuit is the possibility of load current "readback" in the voltage mode (or load voltage readback in the current mode). By monitoring point X in Figure 16, as the current limit value is reduced, a state change will take place when current limit is attained. The set current limit value will correspond to the load current.

In the circuit DAC A with amplifier A1 and buffer A3 acts as a standard programmable voltage source when $V_{REF}A$ is held constant. The voltage drop across resistor R_{S2} provides a voltage proportional to load current with R_{S1} acting as a current limit on amplifier A2. Amplifier A4 with resistors R_1 and R_2 references the voltage across R_{S2} to ground and also provides gain $\left(\frac{R_2}{R_1}\right)$. The output of A4

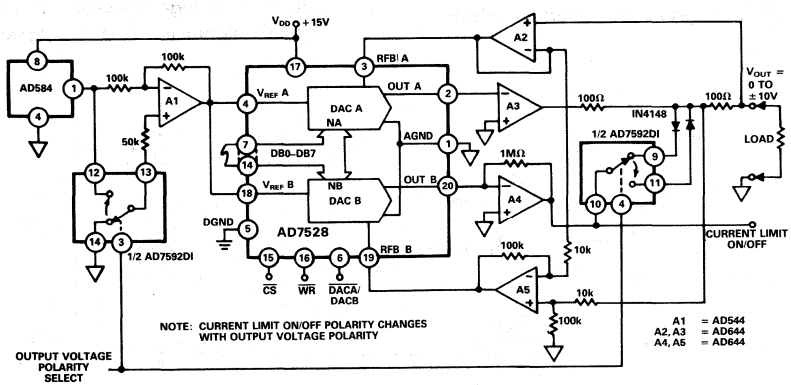

Figure 17. Programmable Voltage/Current Source with Bipolar Output

is compared with a proportion $\frac{N_B}{256}$ of $V_{REF}B$ (usually $V_{REF}A = V_{REF}B$) by amplifier A2. If $V_{OUT4} > V_{REF}B \frac{N_B}{256}$ then current limit is required and the output of A2 via diode D1 draws load current to maintain a constant load current.

If $V_{OUT4} < V_{REF}B \frac{N_B}{256}$ then the current limit is not required and amplifier A2 output is disconnected as diode D1 is reversed biased.

Figure 17 shows a similar circuit for bipolar operation, i.e., 0 to ±10V at 0 to ±10mA.

DUAL DAC PROGRAMMABLE GAIN AMPLIFIER WITH NO TRIMPOTS

A unique advantage of the matched DACs available in the AD7528 is utilized in the programmable gain/attenuation circuit shown in Figure 18. The equivalent resistance of each DAC from its reference input to its output is used to replace the input and feedback resistors in the standard inverting amplifier circuit. By loading DAC's A and B with suitable codes, programmable gain/attenuation over the range −48dB to +48dB can be achieved.

In the circuit of Figure 18, the DAC equivalent resistances are given by:

$$R_{DAC}A = \frac{256\,R_{LD}A}{N_A} \quad \text{and} \quad R_{DAC}B = \frac{256 R_{LD}B}{N_B}$$

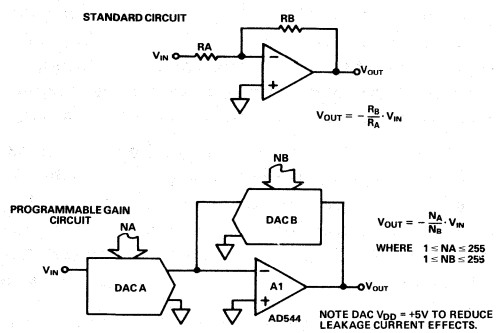

Figure 18. Dual DAC Programmable Gain Amplifier

Where: $R_{LD}A$ and $R_{LD}B$ are DAC A and DAC B R-2R ladde resistances respectively, N_A and N_B are the DAC codes in decimal (1–255).

The resultant gain expression for the circuit is

$$\frac{V_{OUT}}{V_{IN}} = -\frac{256\,R_{LD}B}{N_B} \cdot \frac{N_A}{256\,R_{LD}A}$$

This simplifies to:

$$\frac{V_{OUT}}{V_{IN}} = -\frac{R_{LD}B \cdot N_A}{R_{LD}A \cdot N_B}$$

But since DAC A and DAC B are a matched pair, $R_{LD}A = R_{LD}B$. This simplifies the expression even further to give:

$$\frac{V_{OUT}}{V_{IN}} = -\frac{N_A}{N_B} \qquad 1 \le N_A \le 255 \qquad 1 \le N_B \le 255$$

Notice that DAC ladder resistance does not appear in the expression. Previous PGA circuits using DACs have always had to be trimmed to accommodate the DAC ladder resistance which usually has a wide tolerance (typically 8k to 20k). This circuit does not suffer from this problem since $R_{LD}A$ and $R_{LD}B$ are matched to better than 1%. Notice also that the circuit has a constant input resistance of $R_{LD}A$. The two unused feedback resistors, $R_{FB}A$ and $R_{FB}B$ are also precisely matched and could be used to provide other DAC code vs. gain relationships.

PROGRAMMABLE WAVEFORM GENERATOR FOR VECTOR SCAN CRT DISPLAYS

Figure 19 shows the dual DAC in a triangle/rectangle wave generator in which the period of each half cycle can be programmed. Such a circuit is useful for vector scan CRT displays to generate variable rate sweep signals (depending upon whether a long or short vector is to be drawn). DAC A determines the ramp rate for the positive going ramp of the triangle while DAC B determines the ramp rate for the negative going ramp. The integrator output voltage is sensed by comparators A4 and A5. When this voltage reaches +10V or −10V the comparators drive the R-S flip-flop G1 and G2 which selects the output of the appropriate DAC via the double-pole ganged analog switch SW1, SW2.

The switching arrangement shown has the advantage that high speed switches (such as CD 4016 or AD7519) can be used to change between the two DACs without introducing significant output glitches at the changeover.

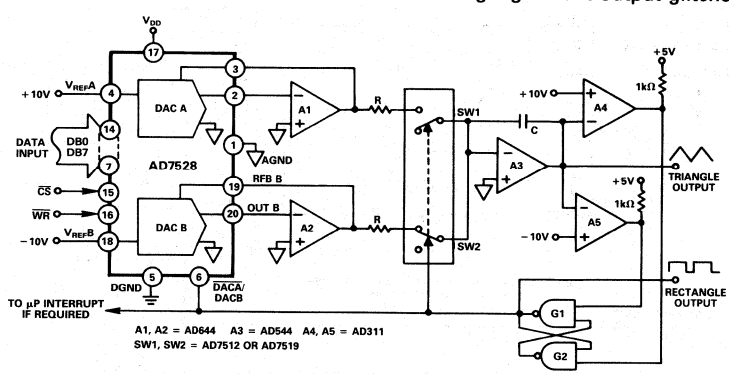

Figure 19. Digitally Programmable Waveform Generator

Furthermore, one DAC can be updated from the data bus and allowed to settle while the output of the other DAC is being used to generate the ramp signal. The output of flip-flop G1 and G2 automatically connects the "unused DAC" to the data bus for further data update if necessary. The output of the flip-flop can be used to drive the interrupts of a microprocessor if required.

Selecting Waveform Parameters:

The period (t) of the waveforms generated by the circuit is given by:

$$t = 512\,RC\left[\frac{1}{N_A} + \frac{1}{N_B}\right]$$

where N_A and N_B are the DAC A, DAC B codes in decimal (1–255) respectively.

If DAC A and DAC B latches contain the same codes, the expression simplifies to:

$$t = \frac{1024\,RC}{N_A} \text{ i.e., output frequency } f = \frac{N_A}{1024\,RC} \text{ Hz}$$

The mark-to-space ratio of the rectangle wave output is dependent on the ratio of N_B to N_A

$$\text{Mark to Space Ratio} = \frac{N_A}{N_B}$$

A special case, exists when the code in either DAC is zero. In this case the circuit will stop oscillating as the integrator input voltage will be zero. If the all zeros condition can occur, it is advisable to connect a 10MΩ resistor from the V_{REF} terminal to the output terminal of each DAC, i.e, V_{REF}A to OUT A and V_{REF}B to OUT B. This provides sufficient bias current to keep the circuit oscillating, and does not affect frequency calculations significantly as the 10M resistor introduces only 1/4 LSB of additional error into each DAC output.

AD7528 SINGLE SUPPLY OPERATION

In low budget digital designs requiring analog outputs, the cost of adding an extra power supply rail for the DAC circuits can be a limiting factor. The AD7528 in the single supply configurations shown below provides an ideal cost effective solution for such applications (especially where multiple analog outputs are required).

Single Supply, Voltage Switching Mode:

In this mode, the normal DAC R-2R ladder is inverted. The reference voltage is applied to the DAC OUT A or OUT B terminal and the output voltage is taken from the DAC V_{REF}A or V_{REF}B terminal. For the DACs to retain their specified linearity, the reference voltage range must be restricted as follows:

For $V_{DD} = +15V$, V_{REF} max $= +2.5V$

For $V_{DD} = +5V$, V_{REF} max $= +0.5V$

Figure 20 shows a circuit for use with a +15 volt power supply giving four separate 0 to +10V outputs. The op-amps used have a Class A output configuration for small signal levels, thus allowing their outputs to go to zero volts for zero volts input. At higher signal levels, the outputs convert to Class B.

Single Supply, Current Steering Mode:

This mode of operation is described in the AD7528 data sheet, and is suitable for single +10 volt to +15 volt supply operation. This is achieved by biasing the AD7528 analog ground (AGND) +5 volts above the power supply ground. Unlike the previous circuits the available drive for the DAC switches is now V_{DD} −5 volts so the 5 volt specifications apply for linearity. Figure 21 shows how a +2 volt to +8 volt analog output may be obtained using two op-amps per DAC. The two DAC reference inputs are tied

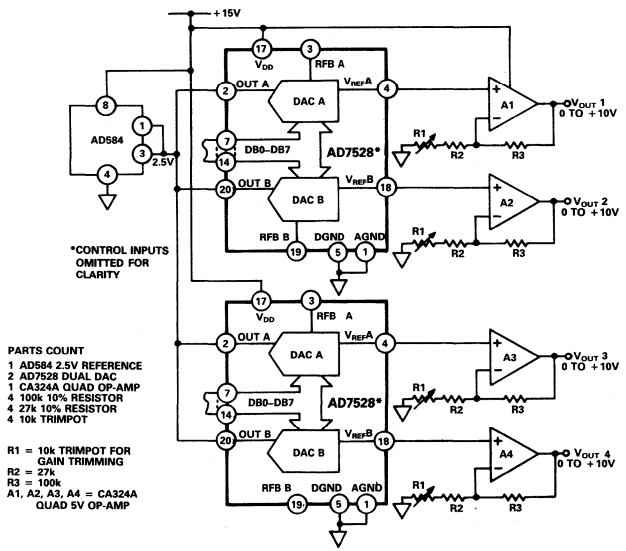

Figure 20. Four Channel Analog Output Circuit

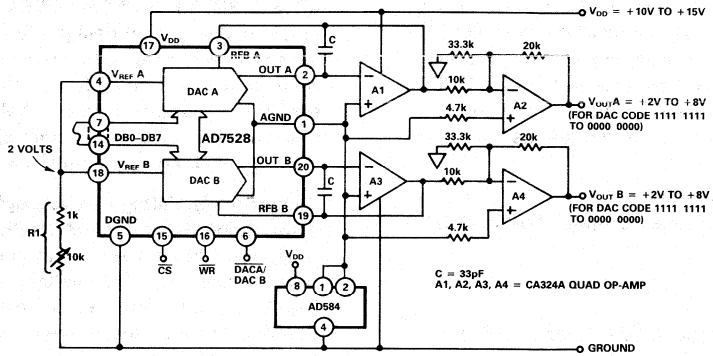

Figure 21. AD7528 Single Supply Operation with AGND Biased to +5 Volts

together and a reference input voltage is obtained without a buffer amplifier by making use of the constant and matched impedances of the DAC A and DAC B reference inputs. Current flows through the two DAC R-2R ladders into R1; R1 is adjusted until V_{REF}A and V_{REF}B inputs are at +2 volts. The adjustment is independent of either DAC code.

Each analog output channel has a +2 to +8 volt range for DAC codes 1111 1111 to 0000 0000.

Reference

Dan H. Sheingold, "Nonlinear Circuits Handbook," available from Analog Devices.

D. P. Burton, Application Note "Methods For Generating Complex Waveforms And Vectors Using Multiplying D/A Converters" Analog Devices Publication Number: E671–15–9/81.

Application Guide to CMOS Multiplying D/A Converters, Analog Devices Publication Number: G479–15–8/78.

ANALOG DEVICES

AN-319
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Simple Interface Between D/A Converter and Microcomputer Leads to Programmable Sine-Wave Oscillator

by John Wynne

This application note outlines a very simple interface between the AD7542, a 12-bit CMOS D/A converter, and one of the most popular industry building blocks available, the 8-bit MCS-48 microcomputer family. The interface is used in a circuit which generates programmable low distortion sine waves. The circuit is based upon repeatedly adding a known constant to an accumulator which in turn scans a sine look-up table. The data from this table is presented to the DAC for conversion into an analog waveform. Changing the constant changes the generated output frequency.

THE INTERFACE

The AD7542 accepts data in three 4-bit nibbles over a 4-bit wide data bus. The data is stored in three 4-bit data registers until, under program control, the new 12-bit word is transferred to the DAC register to update the analog outputs (see Figure 1). Internal register selection is achieved by decoding the address inputs A0 and A1 as shown in Table I. All data loading or data transfer operations are synchronized to the rising edge of the control signal $\overline{WR}$.

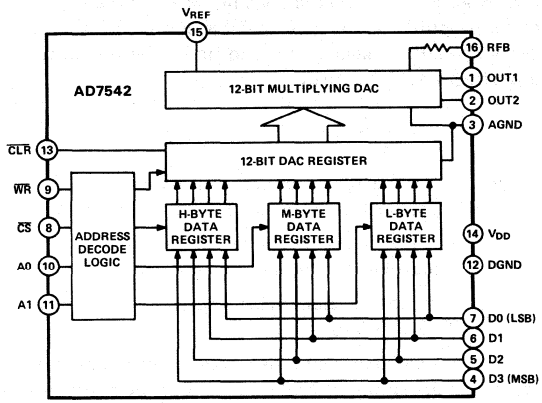

Figure 1. AD7542 Functional Block Diagram

A1 (Pin 11)	A0 (Pin 10)	Register Selected
0	0	Low Byte Register
0	1	Mid Byte Register
1	0	High Byte Register
1	1	DAC Register

Table I. AD7542 Internal Register Decoding

8

On all MCS-48 microcomputers the lower half of Port 2, P20–P23, can be used for I/O expansion with a dedicated I/O expander device, the 8243. The 8243 contains four 4-bit I/O ports which serve as extension of the on-chip I/O and are addressed as ports 4–7 with their own MOV, ANL and ORL instructions.

All communications between the MCS-48 microcomputer and the 8243 occurs over P20–P23 with timing provided by on output pulse on the PROG pin of the processor. The AD7542 is interfaced to the MCS-48 microcomputer over P20–P23 and is accessed by the processor as if it were an 8243. However, only the Transfer Accumulator to Port instructions are used in the interface. Each transfer consists of two 4-bit nibbles, the first containing the "op-code" and port address and the second containing the actual 4-bits of data.

Referring to Figure 2, a high to low transition of the PROG line indicates that address and op-code information are present on P20–P23. The port address on lines P20 and P21 must be externally latched at this point in time. The MCS-48 timing ensures that the falling edge of PROG coincides with the falling edge of ALE. Thus ALE is used to latch the port address information into an external 7475-type latch. The op-code information is not used in this application. When PROG returns high, the 4-bit wide bus contains valid data. By driving the $\overline{WR}$ input of the AD7542 with the PROG line, valid data is loaded into the AD7542 at this time.

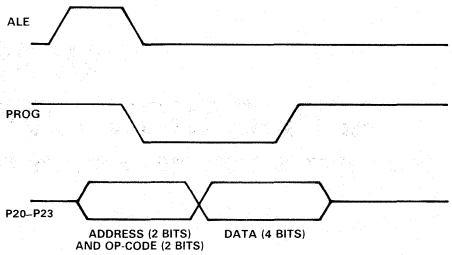

Figure 2. Expander I/O Interface Timing

Nibble 1				Nibble 2			
P23	P22	P21	P20	P23	P22	P21	P20
X	X	A	A	D	D	D	D
Op-Code Address				(MSB) DATA (LSB)			

A (P21)	A (P20)	Port Selected
0	0	Port 4
0	1	Port 5
1	0	Port 6
1	1	Port 7

Table II. Expander I/O Multiplexed Bus (P20-P23)

Table II shows the two nibbles of a transfer instruction. Comparing Table II with Table I, it can be seen that the address of the Low Byte register is the same as the Port 4 address in the I/O expansion mode. Thus a TRANSFER ACCUMULATOR TO PORT 4 instruction will in fact load the Low Byte register of the AD7542. A similar correspondence exists between the remaining AD7542 registers and the MCS-48 expansion ports. Only four instructions are thus required to load the AD7542:

```
MOVD  P4, A   Load Low Byte register
MOVD  P5, A   Load Mid Byte register
MOVD  P6, A   Load High Byte register
MOVD  P7, A   Load DAC register
```

The basic hookup is shown in Figure 3.

The advantage of this interface lies in its simplicity. In either single or multi-DAC applications, mapping the AD7542's as I/O devices greatly simplifies both the software and the chip select decoding over what would be required if the devices were memory mapped (in external data memory).

Multiple AD7542s can be connected to P20–P23 if additional chip select lines are supplied for the additional devices. For instance using the top half of Port 2 (P24–P27) directly as chip selects allows a maximum of four AD7542s on the I/O port; using P24–P27 with a 4-to-16 line decoder (such as the 8205) allows a maximum of sixteen AD7542s on the port.

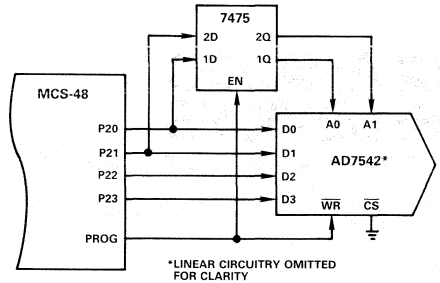

Figure 3. Basic AD7542 to MCS-48 Interface

PROGRAMMABLE SINE-WAVE GENERATOR

One method of generating sine waves is to drive an up-counter at a given clock rate, the output of the counter being applied to a look-up table containing phase-angle information. This data in turn is fed to a DAC which generates the sampled waveform. A low-pass filter is then used to smooth the output waveform. The easiest way to change the output frequency is to increase the rate at which the counter is driven. However, changing the fundamental sampling frequency has severe implications for the output filter. If the software that loads the counter, instead of incrementing it by one per loop, increases the counter by two in the same period of time, the counter fills and rolls over twice as quickly. Thus, although the rate of counter loading–the sampling frequency–remains the same, doubling the loaded value–the frequency constant–doubles the output frequency. In the same way, each integer multiple of the initial frequency constant produces a corresponding multiplication of the output frequency. This is the basis of the technique used here and is described in more detail in References 1–3. The references also include software examples for using the circuit to generate touch-tone frequencies for telecommunication applications.

The number of bits in the accumulator (or counter), m, is determined by the resolution required in output frequency variation. The number of possible output signal frequencies is 2^{m-1}, and the spacing between them is inversely proportional to 2^m. A 14-bit accumulator is sufficient for the 1Hz spacing required here. The generated output frequency is given by

$$f_s = \frac{N.fc}{2^m}$$

where: N is the frequency constant which establishes the output frequency.

f_c is the frequency of executing the loop, i.e., the sampling frequency.

m is the number of bits in the accumulator

Choosing a base frequency constant of 1 to produce a 1Hz output signal ensures that all other frequency constants N always equal the output frequency. Thus for a 14-bit counter to increment in single steps and roll over to zero once per second, requires a sampling rate of 2^{14} samples/sec (a sampling frequency of 16,384Hz). Thus each timing

oop should take 61.035μs to execute. For the MCS-48 family, the external crystal frequency is divided by fifteen to provide the internal machine cycle time. A crystal of 4.9155MHz gives a machine cycle time of 3.056μs which means 20 machine cycles per timing loop. Any branch in the program must be tailored so that a complete loop through the program will always take 20 machine cycles.

THE SINE TABLE

The MCS-48 architecture limits the external look-up table to 256 entries. Thus, 256 bytes encode the sine functions 360° in 1.406° steps (360/256). Placing the full 360° in the look-up table simplifies the software at the expense of increased distortion in the output. The synthesis program of Reference 2 stores a quarter period (first quadrant) of the output waveform in the look-up table and gives the software required to generate the additional quadrants.

The high byte of the 14-bit accumulator is used as a pointer to reference the look-up table. The complete look-up table is shown in Table III along with the BASIC program used to generate the table on an AIM-65. The AD7542 in Figure 5 is arranged for bipolar operation with offset binary coding. To ensure that the output waveform is symmetrical around 0V, the entries in the look-up table are adjusted so that the value FFH is assigned to the greatest positive amplitude equal to (V_{REF} − 1LSB) for 90°, and the value OIH to the greatest negative amplitude equal to– (V_{REF} − 1LSB) for 270°. This allows the output for 0° and 180° to fall exactly on 80H–equivalent to 0V output for bipolar operation.

0700	80	83	86	89	0780	80	7D	7A	77
0704	8C	90	93	96	0784	74	70	6D	6A
0708	99	9C	9F	A2	0788	67	64	61	5E
070C	A5	A8	AB	AE	078C	5B	58	55	52
0710	B1	B3	B6	B9	0790	4F	4D	4A	47
0714	BC	BF	C1	C4	0794	44	41	3F	3C
0718	C7	C9	CC	CE	0798	39	37	34	32
071C	D1	D3	D5	D8	079C	2F	2D	2B	28
0720	DA	DC	DE	EO	07A0	26	24	22	20
0724	E2	E4	E6	E8	07A4	1E	1C	1A	18
0728	EA	EB	ED	EF	07A8	16	15	13	11
072C	F0	F1	F3	F4	07AC	10	0F	0D	0C
0730	F5	F6	F8	F9	07B0	0B	0A	08	07
0734	FA	FA	FB	FC	07B4	06	06	05	04
0738	FD	FD	FE	FE	07B8	03	03	02	02
073C	FE	FF	FF	FF	07BC	02	01	01	01
0740	FF	FF	FF	FF	07C0	01	01	01	01
0744	FE	FE	FE	FD	07C4	02	02	02	03
0748	FD	FC	FB	FA	07C8	03	04	05	06
074C	FA	F9	F8	F6	07CC	06	07	08	0A
0750	F5	F4	F3	F1	07D0	0B	0C	0D	0F
0754	F0	EF	ED	EB	07D4	10	11	13	15
0758	EA	E8	E6	E4	07D8	16	18	1A	1C
075C	E2	E0	DE	DC	07DC	1E	20	22	24
0760	DA	D8	D5	D3	07E0	26	28	2B	2D
0764	D1	CE	CC	C9	07E4	2F	32	34	37
0768	C7	C4	C1	BF	07E8	39	3C	3F	41
076C	BC	B9	B6	B3	07EC	44	47	4A	4D
0770	B1	AE	AB	A8	07F0	4F	52	55	58
0774	A5	A2	9F	9C	07F4	5B	5E	61	64
0778	99	96	93	90	07F8	67	6A	6D	70
077C	8C	89	86	83	07FC	74	77	7A	7D

Table IIIa. Look-Up Table for 360°of Output Waveform

```
L       = 3840
FOR X   = 0 to 255
LET A   = Sin [(X * 1.40625)/57.2958]
LET Y   = INT (128.5 + 127 * A)
POKE L, Y
LET L   = L + 1
NEXT X
END
```

Table IIIb. Basic Program for AIM-65 Used to Generate Look-Up Table

SOFTWARE AND HARDWARE

Figure 4 shows the arrangement of registers within the microcomputer. The accumulator is implemented with two 8-bit registers, Reg 0 and Reg 1. Register 0 is the least significant. The frequency constant N is stored in Reg 2 and Reg 3. Register 2 is the least significant. Both the 14-bit accumulator and the 14-bit wide frequency constant occupy the most significant 14-bits of their respective register pairs. The high byte of the 14-bit accumulator is used as the pointer for accessing the look-up table.

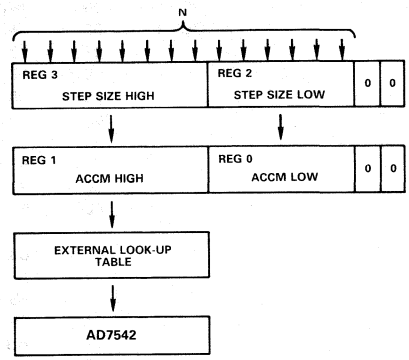

Figure 4. Register Arrangement for Program

The sequence of actions for the synthesis of a sine-wave output is shown in Table IV with the actual program listing in Table V. A circuit diagram of the synthesizer is shown in Figure 5. The 2716 EPROM contains both the program and the sine look-up table for the 8035. Program memory is all external and starts at location 000H, the address to which the processor is vectored after a Reset. The look-up table is treated as external data memory and occupies the highest page in the 2716–starting at 700H–for ease of decoding.

Bit 7 of Port 2 (P27) is used to differentiate between the high and low byte of the right-justified frequency constant N being loaded in through Port 1. A low level on P27 indicates the low byte (8LSB's) is present on P17–P10, a high level on P27 indicates the high byte (6LSB's) is present on P15–P10. For the program to work properly, the low byte of the frequency constant–with P27 low–must be present on Port 1 before a Reset is applied. Almost immediately after a Reset, the low byte is read and the program enters

a loop to allow the high byte of the frequency constant to be presented at Port 1. The program exits the loop–when P27 goes high–by reading the high byte. The frequency constant is next shifted left two places to position it properly in Registers 2 and 3 before the program enters the frequency synthesis loop properly.

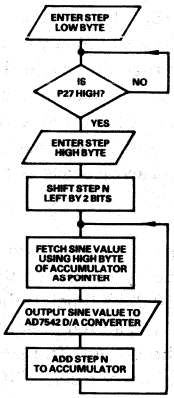

Table IV. Sequence of Action for Sine-Wave Generation

	Location	Op-Code	Mnemonic		Statement
	0000	27	CLR	A	
	01	A8	MOV	RO,A	Clear 14-Bit Accumulator
	02	A9	MOV	R1,A	
	03	3C	MOVD	P4,A	Clear 4 LSB's of DAC
	04	09	IN	A,P1	Get Low Byte of Frequency Constant
	05	AA	MOV	R2,A	
LOOP 1	06	0A	IN	A,P2	Test for High Byte
	07	F7	RLC	A	
	08	E6	JNC	LOOP 1	
	09	06			
	0A	09	IN	A,P1	Get High Byte of Frequency Constant
	0B	AB	MOV	R3,A	
	0C	97	CLR	C	
	0D	FA	MOV	A,R2	Shift Frequency Constant Two Places to the Left
	0E	6A	ADD	A,R2	
	0F	AA	MOV	R2,A	
	10	FB	MOV	A,R3	
	11	7B	ADDC	A,R3	
	12	AB	MOV	R3,A	
	13	97	CLR	C	
	14	FA	MOV	A,R2	
	15	6A	ADD	A,R2	
	16	AA	MOV	R2,A	
	17	FB	MOV	A,R3	
	18	7B	ADDC	A,R3	
	19	AB	MOV	R3,A	
LOOP 2	1A	81	MOVX	A,R1	Fetch Sine Value
	1B	3D	MOVD	P5,A	Output Sine Value to DAC
	1C	47	SWAP	A	
	1D	3E	MOVD	P6,A	
	1E	3F	MOVD	P7,A	
	1F	97	CLR C		
	20	F8	MOV	A,R0	Update Low Byte of Accumulator
	21	6A	ADD	A,R2	
	22	A8	MOV	R0,A	
	23	F9	MOV	A,R1	Update High Byte of Accumulator
	24	7B	ADDC	A,R3	
	25	A9	MOV	R1,A	
	26	00	NOP		Filler to Obtain Desired Sampling Frequency
	27	00	NOP		
	28	04	JMP	LOOP 2	
	0029	1A			

Table V. Program Listing for Sine-Wave Generation

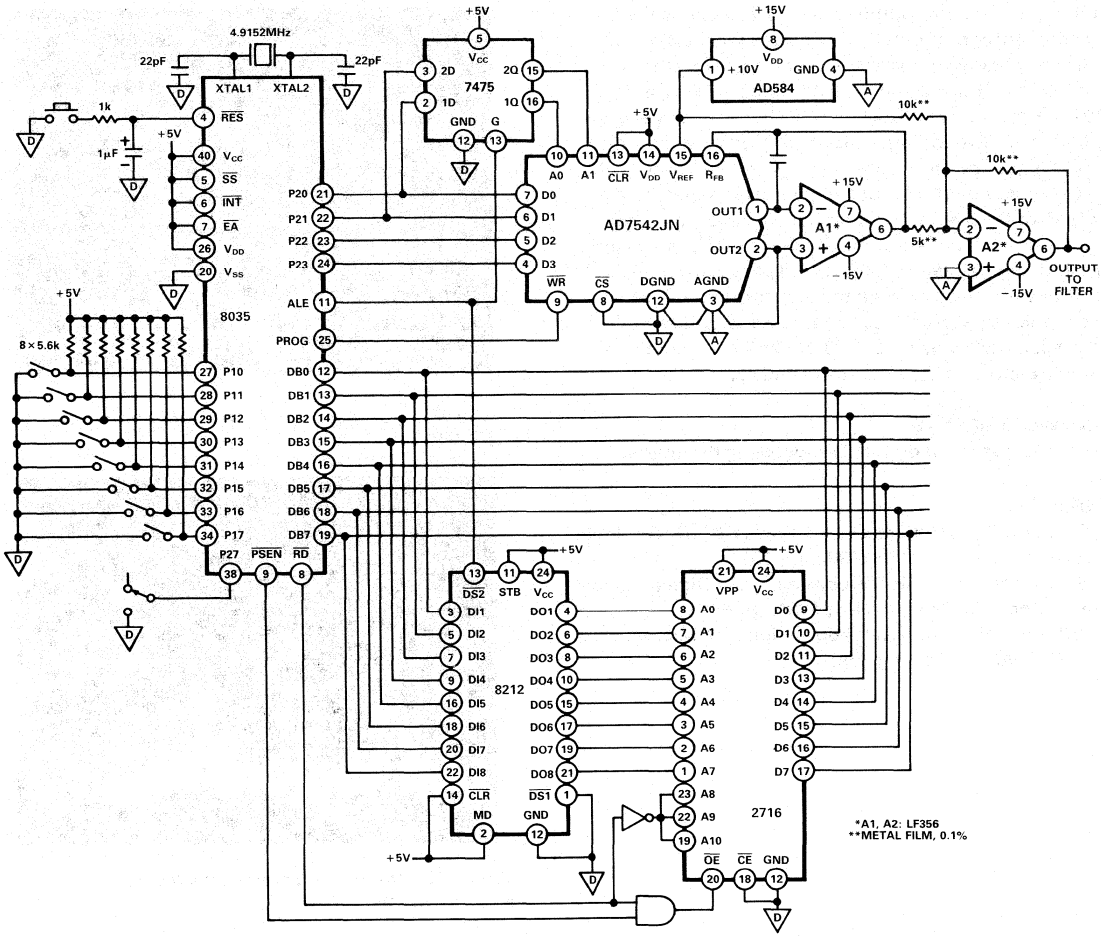

Figure 5. Circuit Diagram of Frequency Synthesizer

EXPERIMENTAL RESULTS

The filter used to smooth the output waveform from the DAC is a sixth-order active RC low-pass filter (Reference 3) and is shown in Figure 6. The cut-off frequency is 4kHz. The response is down 58dB at 8kHz, half the sampling frequency. Experimental results are shown in Table VI.

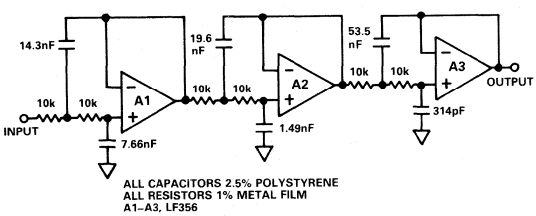

Figure 6. Low-Pass Output Filter, 4kHz Bandwidth

Input Frequency Constant N in Decimal	Output Frequency	Total Harmonic Distortion (dB)
0	0	–
1	1	–
16	16	−48
17	17	−48
32	32	−52.6
33	33	−48
64	64	−54.8
65	65	−43
128	128.01	−56.2
129	129.01	−44.7
256	256.02	−54.6
257	257.02	−45.2
512	512.05	−55.4
513	513.05	−44.1
1024	1024.1	−55
1025	1025.1	−45
2048	2048.2	−52.4
2049	2049.2	−44.8
4096	4096.4	−57
4097	4097.4	−47

Table VI. Experimental Results from Frequency Synthesizer

An interesting feature of the results in Table VI is the increased distortion of output waveforms whose frequencies are not integer power's of 2. For instance at f_{OUT} = 1024Hz (2^{10}Hz) total harmonic distortion is −55dB; at f_{OUT} = 1025Hz the output distortion is −45dB, an increase of 10dB. This is due to harmonics being generated at the "carry" frequency, the frequency at which the low byte of the 14-bit accumulator (in Reg 0) generates a carry to the high byte of the accumulator (in Reg 1). When synthesizing an output frequency of 1025Hz a carry between Reg 0 and Reg 1 is generated every 64 passes through the loop. Since one pass through the loop takes 61μs a series of harmonics spaced at 256Hz intervals is added to the output spectrum causing the 10dB deterioration in total harmonic distortion. See Figures 7a and 7b.

ACKNOWLEDGMENTS

To Norm Bernstein for originally suggesting the interface.

REFERENCES

1. J.E. Galbraith, "Generate Sine Waves by Direct Table Look-Up," EDN, April 28, 1982, pg 101.
2. T.S. Kinsel and J.H. Wuorinen, "A Digital Signal Generator," IEEE Micro, Nov, 1981, pg 6.
3. H. Chamberlin, "Musical Applications of Microprocessors," Hayden, 1980, pg 369.

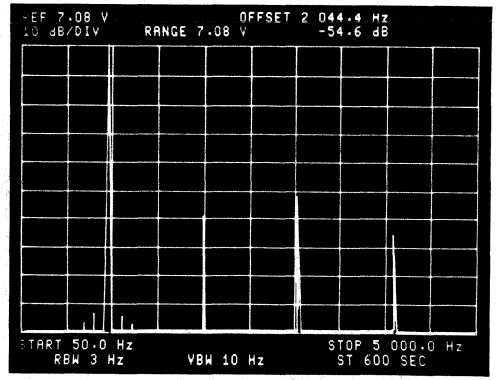

Figure 7a. Output Waveform Spectrum for F_{OUT} = 1024Hz

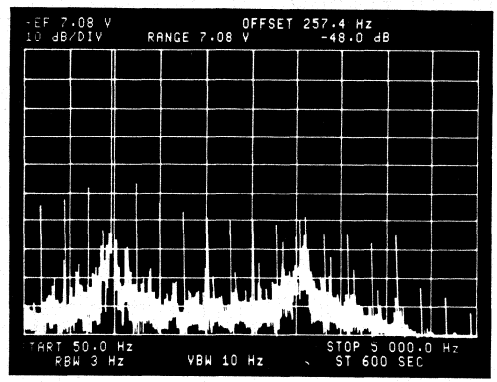

Figure 7b. Output Waveform Spectrum for F_{OUT} = 1025Hz

ANALOG DEVICES

AN-320A
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

CMOS DACs and Op Amps Combine to Build Programmable Gain Amplifiers Part I

by John Wynne

The ability to software program the gain of an amplifier can be very useful to designers. Allied to this ability is the requirement for a large number of programmable gain levels and not simply 3 or 4 user selectable gain levels. CMOS DACs are a natural choice for control of a PGA circuit – not only do they satisfy the above two requirements but they are low cost and can be highly accurate. The traditional way of using CMOS DACs to build a PGA is to place the DAC in the feedback loop of a simple inverting op amp configuration. The DAC behaves as a programmable resistance and thus allows the circuit gain to be digitally controlled. High circuit gains require large values of effective feedback resistance. However, as the effective feedback resistance increases, its precise value becomes less and less defined. This results in a circuit whose gain accuracy decreases with increasing gain. This application note is published in two separate parts. Part I looks at the PGA circuits based upon a single DAC and analyzes the gain errors arising from such a configuration. It also includes an analysis of the dc error sources which limit the dc accuracy of the PGA. Two recently released CMOS DACs from Analog Devices, the AD7534 and AD7538, combine a number of features which make them suitable for use in a PGA circuit. The AD7534 is in a 20-pin, 0.3" wide package and loads data in an 8+6 format. The AD7538 is in a 24-pin, 0.3" wide package and has a 14-bit parallel loading format. Both have similar specifications. A detailed comparison is made between the performance of a PGA circuit based on the 14-bit AD7534 and one based on a 12-bit AD7545. In Part II of the application note, PGA circuits built with dual DACs are investigated. These circuits offer the advantage of greater accuracy over a wider dynamic range in comparison with a single DAC solution. Monolithic dual 8-bit (AD7528) and 12-bit (AD7537, AD7547 and AD7549) DACs are now available which make this a cost effective solution.

THE BASIC EQUATIONS

Turning a DAC into a PGA element simply involves placing the network in the feedback path as opposed to the input path of an inverting op amp circuit configuration. Both configurations are shown in Figure 1. The transfer function of Figure 1a is:

$$\frac{V_{OUT}}{V_{IN}} = -\frac{R_{FB}}{R_{EQ}} \quad (1)$$

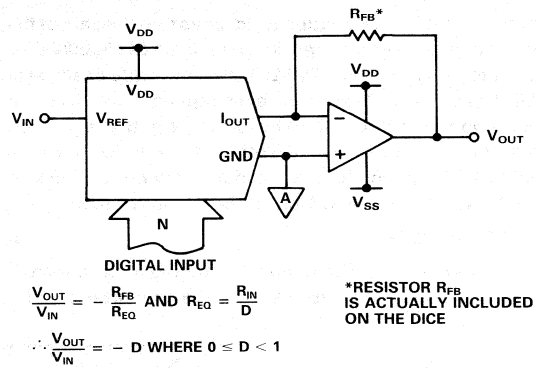

Figure 1a. Standard DAC Configuration to Provide Attenuation

R_{EQ} is the equivalent transfer impedance of the DAC from the V_{REF} pin to the I_{OUT} pin and can be expressed as

$$R_{EQ} = \frac{2^n R_{IN}}{N} \quad (2)$$

where: n is the resolution of the DAC

N is the DAC input code in decimal

R_{IN} is the constant input impedance of the DAC ($R_{IN} = R_{LAD}$ for an R-2R DAC)

Substituting this expression into Equation 1 and assuming zero gain error for the DAC ($R_{IN} = R_{FB}$) the transfer function simplifies to

$$\frac{V_{OUT}}{V_{IN}} = -\frac{N}{2^n} \quad (3)$$

This ratio is commonly represented by D and, as such, is a fractional representation of the digital input word, i.e.

$$\frac{V_{OUT}}{V_{IN}} = -D \quad (4)$$

When the DAC network and the feedback resistor R_{FB} are swopped around as in Figure 1b, the ideal transfer function is obviously the inverse of Equation 3 or

$$\text{GAIN IDEAL} = \frac{V_{OUT}}{V_{IN}} = \frac{-2^n}{N} = \frac{-1}{D} \quad (5)$$

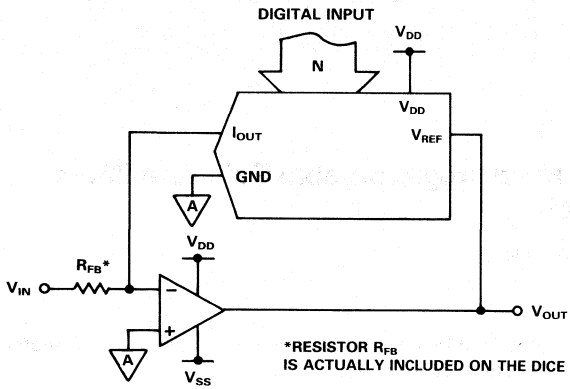

Figure 1b. Transposed DAC Configuration to Provide Gain

With an 8-bit DAC, Equation 5 indicates that the gain of the circuit can be varied from 256 (an input code of 00000001) down to unity (actually 256/255, an input code of all 1s) in 255 steps. The all 0s code is never applied. This avoids an open-loop condition thereby saturating the amplifier. With the all zeros on (all bits off) code excluded there remains $2^n - 1$ possible input codes allowing a choice of $2^n - 1$ output levels. In dB terms, the dynamic range is

$$20 \log_{10} (2^n - 1) \quad (6)$$

For an 8-bit device the dynamic range is approximately 48dB. Increasing the resolution of the converter by 1 bit increases the dynamic range by 6dB. Thus a 12-bit device will have a maximum dynamic range of 72dB. Closely allied to the number of programmable gains available is their distribution over the gain range from minimum gain to maximum gain. Figure 2 is a graphical representation of Equation 5 with the circuit gain plotted in dB for an 8-bit DAC. The graph indicates that one half (128) of the total number of steps or LSBs available cover the 0 to +6dB range (gains from 1 to 2), one quarter (64) of the steps cover the next 6dB of gain (+6 dB to +12dB or gains from 2 to 4), one eighth(32) of the steps cover the next 6dB (+12dB to +18dB or gains of 4 to 8), etc. Figure 2 illustrates quite clearly the decreasing number of gain steps available as the circuit gain increases in 6dB bands. From Figure 2 it can be seen that the last 6dB gain band (+42dB to +48dB or gains of 128 to 256) is covered in just one step. This occurs when the DAC digital code changes from decimal 2 to decimal 1. Thus most of the programmable gain steps are concentrated at the lower end of the gain range offering tremendous resolution between adjoining gain settings (less than 0.05dB). Figure 2 might suggest that this step "bunching" is the only drawback of the PGA circuit. This is not so. When the various error sources are taken into account, the actual performance falls quite short of the ideal.

DEFINING THE ERRORS

The main culprit in the performance shortfall is the DAC Integral Linearity which causes a rapid deterioration in system accuracy with increasing gain. The reason is not hard to see. For example, with an 8-bit DAC and a required gain of 16, the DAC input code is decimal 16. If the DAC is 8-bit linear, i.e. ±0.2% F.S. or ±1/2LSB, then the effective input code can vary from 15.5 up to 16.5. This means the system gain of equation 5 can vary from 256/15.5 (or 16.51) down to 256/16.5 (or 15.51) – an error of ±3% even

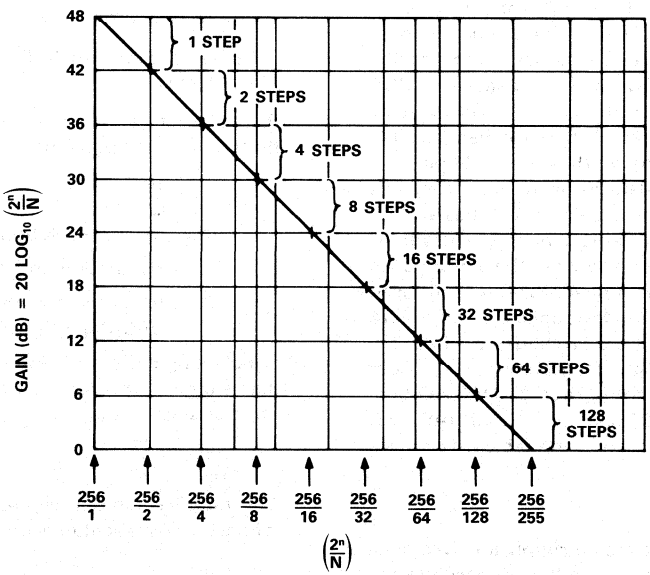

Figure 2. Step Distribution in dB vs. Digital Code for an 8-Bit DAC

though the DAC itself has a maximum error of ±0.2%. Higher gains produce correspondingly higher errors.

The transfer function of Equation 2 assumed that there was no gain error in the DAC. DAC gain error does exist, however, and it results in the ideal gain expression of Equation 5 being multiplied by a term close to unity. The smaller the gain error then the closer this multiplier term is to unity. When both linearity errors and gain errors are included, the system gain expression becomes

$$\text{GAIN ACTUAL} = \frac{V_{OUT}}{V_{IN}} = -\frac{2^n}{(N+X)(1+\Delta)} \qquad (7)$$

Where X is the linearity error in LSBs

$$\text{and } (1+\Delta) = \frac{R_{FB}}{R_{IN}}$$

The difference between the ideal gain and the actual gain can be expressed in percentage terms as

$$E(\%) = -\left[\frac{\Delta}{(1+\Delta)} + \frac{1}{(1+\Delta)}\left(\frac{X}{N+X}\right)\right] \cdot 100\% \qquad (8)$$

Since the gain error of a DAC can be trimmed to zero whereas nothing can be done to reduce the integral nonlinearity, it is instructive to consider the percentage gain error due to DAC nonlinearity alone and then to add the DAC gain error term.

Quick Rule-of-Thumb Helps

With zero DAC gain error equation 8 simplifies to

$$E(\%) = -\left(\frac{X}{N+X}\right) \cdot 100\% \qquad (9a)$$

This expression applies for all gain settings except unity when an additional error term must be added to Equation 9a. For any R-2R ladder network the maximum output current (with all 1s applied to the DAC) is always 1LSB less than the input current. The "missing" LSB worth of current flows through the ladder termination resistor to signal ground. Therefore, at a gain setting of unity, a unique error term equal to 1LSB in percent must be added to the percentage error expression of Equation 9a.

By rearranging Equation 9a, the percentage gain error can be usefully expressed in terms of the programmed gain and DAC linearity (accuracy) as

$$E(\%) = -\left(\frac{2^n}{N+X}\right)\left(\frac{X}{2^n} \cdot 100\%\right) \qquad (9b)$$

or, in words, the maximum percentage gain error is equal to the required gain times the DAC linearity in percent. This is a quick rule-of-thumb for estimating output error for gains greater than unity. The solid line in Figure 3 is a graphic representation of Equation 9a plotted for a 12-bit accurate, 12-bit resolution DAC and shows the increasing gain error with increasing gain. For comparison purposes, a plot for a 14-bit accurate, 14-bit resolution DAC is represented by the dotted line in Figure 3. Because the linearity error X can be either positive or negative, the vertical axis of Figure 3 is graduated in terms of ±E(%). Figure 3 is, however, an approximation, albeit a good one, of Equation 9a. The straight line representation of error versus gain in Figure 3 suggests that the maximum positive and negative gain errors are symmetrical over the entire gain range. In fact at very high gains the maximum gain errors become noticeably asymmetric and continue to be increasingly asymmetric with increasing gain.

This is due to the ±X linearity term in the denonimator of Equation 9a having an increasing impact with increasing gain (N reducing). The magnitude of the asymmetry can be seen from the results in Table A1 of Appendix 1 which shows computed values of Equation 9a for the AD7545LN

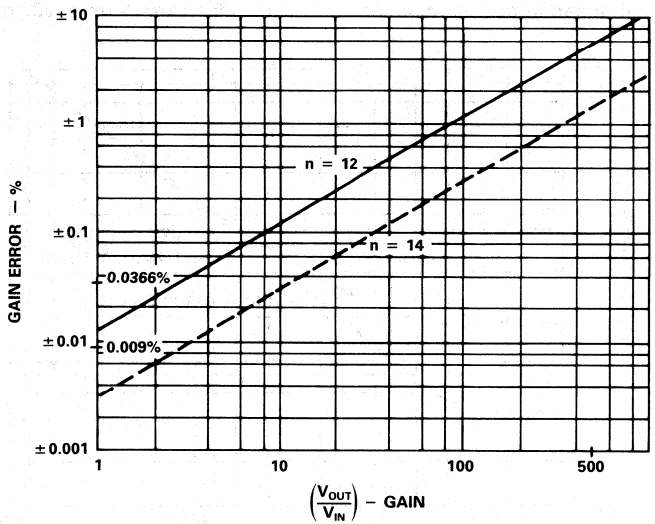

Figure 3. Comparison of Worst Case Gain Errors Between PGA Systems Based on a 12-Bit Resolution, 12-Bit Accurate DAC and a 14-Bit Resolution, 14-Bit Accurate DAC. DAC Gain Error is Zero in Both Cases.

using both positive and negative linearity terms. For instance, at a gain of 128, the maximum positive gain error is +1.587% and the maximum negative gain error is −1.538%, an imbalance or asymmetry of 0.049%. At a gain of 16 the gain error asymmetry is 0.008%, a vanishingly small number. For low gain settings, Figure 3 is thus a good approximation of the gain errors to be expected from a DAC-based PGA circuit with zero DAC gain error. At the all 1s or unity gain setting, the effect of the LSB worth of signal current "lost" in the ladder termination is to always increase the gain, albeit only slightly, over unity. For the computed AD7545LN values shown in Table A1, the X1 gain error figures include the additional term of 1LSB or +0.0244%.

DAC Gain Error Causes Skew

When a DAC with non-zero gain error is used, the more complete Equation 8 must now be used to compute the total error. From inspection of Equation 8, non-zero DAC gain error has two effects; the first is to multiply the error results of the previous analysis by a term $\frac{1}{(1+\Delta)}$ which is close to unity. The overall effect of this is small and serves only to increase or reduce the gain asymmetry error at any given gain setting. (Positive DAC gain error reduces the asymmetry.) The second and major effect of DAC gain error is to add an error term which is independent of gain setting. Since the term is constant, it has the greatest impact at low gain settings and serves to skew the transfer curve of Figure 3 either in a positive or negative direction. (Positive DAC gain error obviously causes a negative skew.) Table A2 in Appendix 1 shows computed values of Equation 8 for all combinations of DAC gain error and linearity error. The unity gain figures include an additional +0.0244% error term due to the ladder termination. The worst case values from this table are used to plot Figure 4. Shown dotted in this figure is the transfer curve of Figure 3. Thus even with a 12-bit resolution, 12-bit accurate DAC (X = ±1/2LSB) having a tight gain error specification of ±5 LSBs, the worst case system gain error can be as large as 0.159% for a gain of 1.

The gain error analysis leading to Figure 4 assumed that the loop gain of the system is sufficiently high so as to cause no appreciable error. With the AD OP-07 this is a valid assumption at dc and low frequencies if extreme gain settings are avoided. Exactly how the loop gain affects the gain error can be seen from the first term in expression A7 of Appendix A2. At high signal frequencies and high gain settings the finite loop gain can contribute to the gain error of the system. This topic is dealt with in greater detail in Part II of this application note.

ADDITIONAL ERROR SOURCES CAUSE DC OFFSET ERRORS

Gain errors such as these are not the only errors which affect the PGA circuit of Figure 1b. Additional error sources result in a gain-dependent dc offset voltage at the op amp output. The additional DAC-related error sources are code-dependent output resistance and output leakage current. Op amp-related error sources are input offset voltage, input bias current and finite open loop gain. Appendix 2 outlines the contribution of each individual error source to the output. Assume the requirement is for a PGA with gains from 1 to 64. Assume also that the components available are the AD7545LN and the AD OP-07E. The relevant specifications for these components at 25°C are shown in Tables I and II. When the min/max figures from Tables I & II are substituted into Equation A7, the worst case output voltage (occurring for $N = 64_{10}$, i.e. Gain of 64) is equal to

$$V_{OUT} = -64.58\,V_{IN} \pm 29.1\text{mV} \qquad (10)$$

Gain error is approximately 0.9%. Only 6.5mV of the total dc offset error term of 29.1mV is due to the input offset

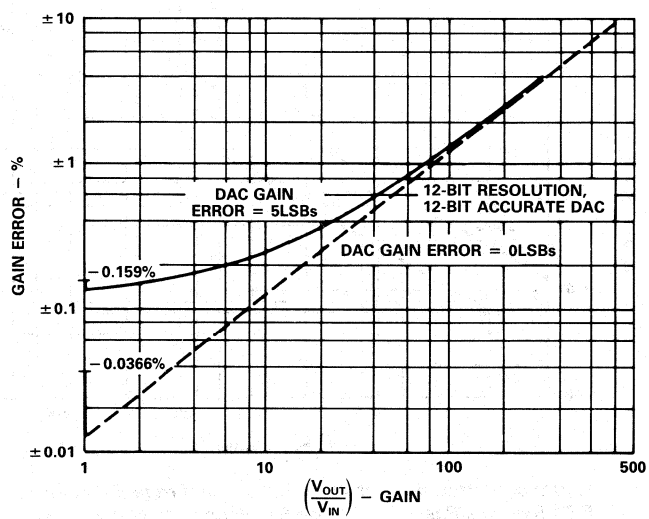

Figure 4. Theoretical Worst Case Gain Errors for AD7545-Based System with and without DAC Gain Error

voltage V_{OS}; the remainder is due to both DAC leakage current I_{LKG} (the major contributor at 16.1mV) and op amp input bias current $I_B(-)$. For the purpose of calculating the op amp input offset voltage error term, a value of $R_0 = 3$ R_{FB} was used in A7. If V_{IN} is a dc signal, then the dc offset error of 29.1mV will be indistinguishable from the output signal. Referring the dc offset error term to the system input results in an equivalent input error of 29.1/64mV or 0.455mV. The smallest input signal V_{IN} for less than, say, 1% error at the X64 gain setting is thus 45.5mV. This error is in addition to the gain error of 0.9%.

Errors Increase with Temperature

If the PGA is expected to perform accurately over a wide temperature range, the drift performance of the error sources with temperature will be important. To a greater or lesser extent all of the terms in Equation A7 are temperature sensitive. A similar analysis of the circuit at, say, 70°C using the data specifications indicates that the gain error will remain essentially unchanged but that the dc offset error can increase dramatically. Over this temperature range, the input offset voltage can increase by nearly 80%, the input bias current by 40%, but the already dominant component, DAC leakage current I_{LKG}, can increase by 400%. When these worst case percentage increases are translated into output voltages, the input offset voltage term can add 5.2mV; the input bias current term can add 2.5mV, but the DAC leakage current term can add 64mV. However, experience shows that the DAC leakage current really only begins to increase at temperatures greater than 85°C. A representative plot of DAC output leakage current versus temperature for a CMOS DAC is shown by the solid line in Figure 6. If the PGA circuit must operate over the military temperature range of −55°C to +125°C, then AD7545UD and AD OP-07H grade devices must be used in place of the previous commercial grades. Tables I and II show the relevant specifications of these high-temperature devices. Using these figures in Equation A7 the output voltage of the PGA at +125°C is

$$V_{OUT} = -64.59\,V_{IN} \pm 349.6\text{mV} \qquad (11)$$

Parameter	AD OP-07E $T_A = +25°C$	AD OP-07E $T_A = +70°C$	AD OP-07H $T_A = +125°C$
Open Loop Gain, A_{OL}	2.10^6 min	$1.8.10^6$ min	$1.5.10^6$ min
Input Bias Current, $I_B(-)$	4nA max	5.6nA max	6nA max
Input Offset Voltage, V_{OS}	75μV max	134μV	200μV max

NOTE: $V_{DD} = +15V$, $V_{SS} = -15V$

Table I. AD OP-07 Specifications at $T_A = +25°C$, +70°C & +125°C

Parameter	AD7545LN $T_A = +25°C$	AD7545LN $T_A = +70°C$	AD7545UD $T_A = +125°C$
Resolution, n	12-bits	12-bits	12-bits
Relative Accuracy, x (Integral Linearity)	±1/2LSB max	±1/2LSB max	±1/2LSB max
Gain Error	±5LSBs max	±6LSBs max	±6LSBs max
Output Leakage, I_{LKG}	10nA max	50nA max	200nA max
Input Resistance, R_{IN}	25kΩ max	25kΩ max	25kΩ max

NOTE: $V_{DD} = +5V$

Table II. AD7545 Specifications at $T_A = +25°C$, +70°C & +125°C

Programmed gain of the circuit is again 64. The gain error is virtually unchanged from its 25°C value as shown in equation 10. The dc offset errors, though, have increased by an order of magnitude. Of the 349.6mV, the op amp bias current contributes 9.7mV; the input offset voltage contributes 17.4mV, while DAC leakage current accounts for the remaining 322.5mV. Referring the 349.6mV to the system input results in an equivalent input error, for the X64 gain setting of 349.6/64mV or 5.5mV. The smallest dc input signal V_{IN} for less than 1% error at the X64 gain setting is thus 550mV. This is at +125°C and is an order of magnitude greater than the minimum input signal at +25°C. Table IVa lists the gain and dc offset errors derived from Equation A7 at +25°C, +70°C and +125°C when using the 12-bit AD7545.

REDUCING THE ERRORS

Gain Errors

From the graphical representation in Figure 4 a somewhat surprising conclusion can be drawn. In precision applications, DAC gain error is as significant a contributor to system error as is DAC integral linearity. This is because any DAC gain error results in a constant error term having the greatest impact on the most accurate gain settings. For example, at a gain setting of 1, the error from Figure 4 is ±0.159%. This is for a 12-bit resolution, 12-bit accurate DAC with a DAC gain error of ±5LSBs. If DAC gain error did not exist – dotted line in Figure 4 – the maximum system error would be ±0.036%, one fifth of its previous value. It is of course possible to trim the DAC gain error to zero by means of a potentiometer. However, the temperature coefficients of the trim components will differ from that of the thin-film ladder resistors of the DAC resulting in a change in DAC gain error over temperature. This topic of trimming CMOS DACs for minimum drift is dealt with extensively in Reference 1.

As the programmable gain setting is increased (N reducing), DAC integral nonlinearity becomes the dominant contributor to the system gain error. For example, at a gain setting of 64, the system error from Figure 4 is approximately ±0.9%, of which ±0.78% is due to DAC integral nonlinearity. In place of the single gain stage, two series-connected gain stages can be used to provide the same overall gain. Since the product of the gain settings of the individual stages is equal to the overall gain, individual stages can now have much lower gain settings and hence higher accuracy. If two stages are available, each having a similar performance to that shown in Figure 4, then for the same overall gain setting of 64 (individual gain settings of 8), the system gain error is approximately ±0.44%. More than half of this error is due to DAC gain error. In order to keep overall system cost, complexity and size to a minimum, monolithic dual 8-bit and 12-bit DACs are now available from Analog Devices which are well suited for this application. The AD7528 is a dual 8-bit DAC with a parallel loading structure. The AD7537, AD7547 and AD7549 are dual 12-bit DACs with 8+4, full-parallel and nibble (4-bits) loading structures respectively. The performance of these circuits is investigated in Part II of the application note.

DC Offset Errors

Depending on the application, the dc offset errors may or may not be significant. The most obvious solution to eliminating these dc voltages is to ac couple the output signal. However, this may not always be possible even in those applications where dc levels are unimportant. If the frequency of interest extends down to, say, a few Hertz, there will be a conflict between choosing a value (and hence, a physical size) for the coupling capacitor and suffering an increase in gain errors due to low frequency fall-off. Depending on the circuit performance required and on the specific application a designer can identify the most important circuit parameters and accordingly choose components for the most cost-effective solution.

DAC output leakage current is very much the major contributor to dc offset errors – especially at high temperatures. To date CMOS DACs with low output leakage currents over the military temperature range have not been available.

A SINGLE DAC SOLUTION

Two recently introduced high resolution CMOS DACs from Analog Devices boast specifications which are close to the ideal for a DAC operating in the PGA mode. The AD7534 and AD7538 are low-cost 14-bit resolution DACs with extremely tight DAC gain error specifications, ±4LSBs maximum. Coupled with 13-bit accuracy specifications and a patented low-leakage technique (US Pat No. 4,590,456) which maintains a very low-leakage current at high temperatures, both gain and dc offset errors are much reduced over the previous case. Table III shows the relevant specifications for the AD7534KN at +25°C. Figure 5 is a graphical representation of Equation 8 plotted for the AD7534KN. Shown dotted in the figure is the graphical representation of Equation 9 – where DAC gain error is assumed to be zero. These curves should be compared with those of Figure 4. Tables A3 and A4 in Appendix 1 show computed values of Equations 9 and 8 respectively for all combinations of DAC gain error and linearity error.

A dedicated pin (V_{SS}) is available on the AD7534 which can be tied either to 0V (normal DAC output leakage) or tied to −300mV (reduced DAC output leakage) to implement the patented low-leakage scheme. Figure 6 is a plot of leakage current versus temperature for both conditions. It clearly shows the improvement gained by using the low-leakage scheme. If the relevant specifications from Table III are substituted into equation A7 (using AD OP-07E specifications from Table I), the worst case output voltage at a gain of 64 (occurring for $N = 256_{10}$) is equal to

$$V_{OUT} = -64.264\,V_{IN} \pm 11.1mV \qquad (12)$$

Gain error is approximately 0.41%. This result is for 25°C, and it should be compared with the result in Equation 10. Both gain and dc offset errors have been at least halved.

Parameter	AD7534KN $T_A = +25°C$	AD7534KN $T_A = +70°C$	AD7534TD $T_A = +125°C$
Resolution, n	14-bits	14-bits	14-bits
Relative Accuracy, x (Integral Linearity)	±1LSB max	±1LSB max	±1LSB max
Gain Error	±4LSBs max	±4LSBs max	±4LSBs max
Output Leakage, I_{LKG}	±5nA max	±10nA max	±20nA max
Input Resistance, R_{IN}	10kΩ max	10kΩ max	10kΩ max

NOTE: $V_{DD} = +15V, V_{SS} = -0.3V$

Table III. AD7534 Specifications at $T_A = +25°C$, $+70°C$ & $+125°C$

For the purpose of calculating the op amp input offset voltage error term, a value of $R_0 = 15\,R_{FB}$ was used in A7. Table IVb shows the individual components of the dc offset error term. The analysis is repeated at +70°C and +125°C, and these results are also included in Table IVb. Comparing the results of Table IVa with Table IVb the effectiveness of the low leakage scheme on the AD7534 is immediately apparent – dc offset errors at +125°C are an

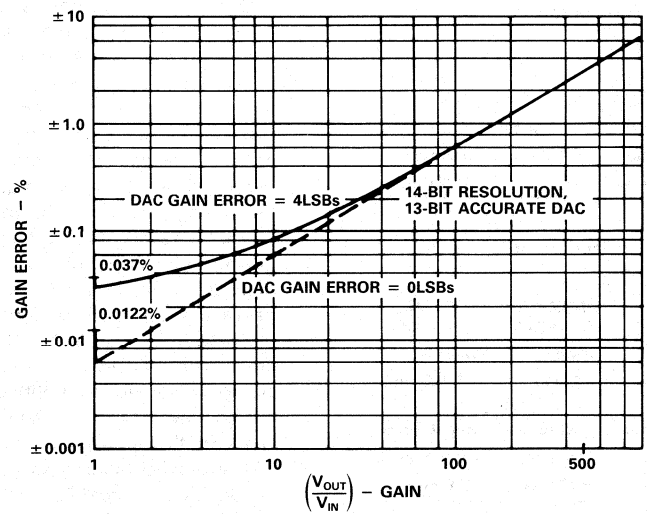

Figure 5. Theoretical Worst Case Gain Errors for AD7534-Based System with and without DAC Gain Error

order of magnitude less with the AD7534. The effectiveness of the low-leakage scheme is further demonstrated by the fact that, for identical gain settings, the equivalent input error voltage for the AD7534 at +125°C is approximately the same as the equivalent input error voltage for the AD7545 at +25°C.

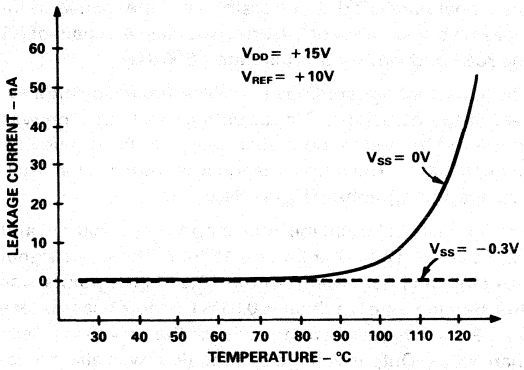

Figure 6. Graph of Typical Leakage Current vs. Temperature

NOISE GAIN

A fundamental requirement of any PGA circuit is that it be monotonic, i.e. if the DAC digital code is changed to increase the gain, then the circuit gain does in fact increase or, at the very least, the circuit gain does not decrease. If the linearity of the DAC is concomitant with the resolution, i.e. 12-bit resolution DAC with 12-bit accuracy, then the DAC is inherently monotonic, and ideally all possible input codes are usable for gain setting. However, if the linearity is less than the resolution, so to speak, i.e. 11-bit accurate but 12-bit resolution DAC, then adjoining codes may overlap causing nonmonotonic operation. In the case of the 11-bit accurate, 12-bit resolution DAC every second code should be avoided, i.e. adjacent codes should not be used. If the DAC specification is more extreme, i.e. 10-bit accurate but still 12-bit resolution, only every fourth code should be used. Thus the number of programmable (and guaranteed monotonic) gain settings is now one quarter of the total number of steps.

The AD7534KN is a 14-bit resolution DAC with 13-bit relative accuracy. Additionally, it is guaranteed monotonic to 14-bits over temperature. This means that ideally all 16, 384 codes are usable for gain setting. However, there is another error source which must be taken into account. The contribution of the input offset voltage, V_{OS}, to the dc offset errors (via the circuits noise gain) has already been mentioned. This noise gain results in a dc term being added to or subtracted from the op amp output. The greatest likelihood of the system becoming nonmonotonic occurs when the DAC input code changes from one code setting to the next adjacent code setting. Adjacent code transitions which are the most critical are those which change the DAC output impedance, R_O, from a low value to a higher value. However, even at the most critical code transition, a nonmonotonic step can still be avoided if the analog input signal voltage, V_{IN}, is sufficiently large with respect to the op amp input offset voltage, V_{OS}. Analysis of the AD7534 over the X1 – X64 gain range in integer steps shows that the worst case code transition is the last code transition. The output impedance R_O changes from $R_O = 6.7\ R_{FB}$ at X63 gain setting ($N = 260_{10}$) to $R_O = 15\ R_{FB}$ at X64 gain setting ($N = 256_{10}$). Simultaneously solving expression A7 at this worst case code transition yields a minimum value for V_{IN} which is the threshold value for possible nonmonotonic operation. This value is 6 V_{OS}. Input signal levels greater than this value will avoid any nonmonotonic transitions occurring. For integer gains which are not integer powers of 2 e.g. X3, X5, X6 etc., the nearest 14-bit code was chosen in each case. Table A5 in Appendix 3 lists both the digital input code used and the resulting $\left(1 + \frac{R_{FB}}{R_O}\right)$ ratio for each integer gain setting in the X1 – X64 gain range.

In some applications, such as closed-loop servo systems, the requirement is for very fine gain adjustment rather than integer gain increments. As Figure 2 suggests, half of the total number of codes, or 2^{n-1}, cover the gain range from X1 to X2. For the AD7534 this means that 8192 codes are available for gain setting between X1 and X2, allowing extremely fine gain adjustment. Between the X1 and X64 gain settings there are just over 16,000 possible gain settings. However, in return for having so many gain settings available, there is a price to be paid in terms of the minimum V_{IN}/V_{OS} ratio required for monotonicity. Computer analysis shows that in order to avoid a nonmonotonic transition between any adjacent code pair over the full X1

		DC ERRORS			
TEMPERATURE	GAIN ERROR	I_{LKG}	$I_B(-)$	V_{OS}	TOTAL
T_A = +25°C	0.91%	16.1mV	6.5mV	6.5mV	29.1mV
T_A = +70°C	0.91%	80.1mV	9mV	11.7mV	100.8mV
T_A = +125°C	0.92%	322.5mV	9.7mV	17.4mV	349.6mV

Table IVa. Error Analysis Over Temperature with AD7545, Gain = 64

		DC ERRORS			
TEMPERATURE	GAIN ERROR	I_{LKG}	$I_B(-)$	V_{OS}	TOTAL
T_A = +25°C	0.41%	3.2mV	2.6mV	5.3mV	11.1mV
T_A = +70°C	0.41%	6.4mV	3.6mV	9.5mV	19.5mV
T_A = +125°C	0.41%	12.8mV	3.8mV	14.1mV	30.7mV

Table IVb. Error Analysis Over Temperature with AD7534, Gain = 64

– X64 code table, the minimum analog input voltage, V_{IN}, should be greater than 1350 V_{OS}. At first glance this is a surprisingly large ratio but, on reflection, should be expected. When adjacent 14-bit codes are used, the percentage increase in gain is a tiny fraction of a percentage point. This is especially true at the lower gain settings. However, the percentage change in noise gain for the same code transition can be orders of magnitude greater. For instance, for the code transition $16{,}379_{10}$ to $16{,}378_{10}$, the signal gain increases by 0.006% whereas the noise gain changes by 0.775%. The ratio of V_{IN} to V_{OS} required to ensure monotonicity at this code transition is 400.

Noise gain is usually assumed to contribute a dc error term only; acting only on the op amps input offset voltage, V_{OS}. However, expression A1 in Appendix 2 shows that the noise gain also contributes to the gain error of the system. The very high open-loop gain of the AD OP-07, however, ensures that any contribution is at a negligible level.

TEST RESULTS

DC measurements were taken on a number of AD7534KN/AD OP-07E combinations at +25°C and a representative result is indicated in Figure 7. The "starred" points indicate the measured errors at the selected gain settings between X1 and X256. At each gain setting the dc input signal level was adjusted to provide a 5V output signal level. As a comparison, the theoretical worst-case error curves from Figure 5 are reproduced again in Figure 7. No adjustments were made to either DAC gain error or op amp input offset voltage. Since the measurements taken were dc, the representative error curve in Figure 7 includes the offset errors at +25°C. However, these were measured at approximately 40μV and were considered negligible in comparison to the input test signal level.

With a gain of 64 the 3dB bandwidth of the system is typically 7kHz. In order to avoid gain errors due to op amp roll-off the signal bandwidth should be restricted to an appropriate limit. For instance, restricting the signal bandwidth to below 100Hz (1/70th of 3dB frequency) ensures that additional gain errors due to roll-off are below 0.01% (a simple, single pole roll-off is assumed). Wider bandwidth op amps than the AD OP-07 can be used in order to increase the signal bandwidth but possibly at the expense of increasing one or more of the error sources. At a gain of X1 the 3dB bandwidth is approximately 500kHz.

The output voltage settling time will be limited by the slew rate of the AD OP-07. Output voltage settling time was measured for a step input change at two fixed gain settings (X1, X64). The output response is shown in the oscilloscope photographs of Figure 8a and 8b.

For Figure 8a (X1 Gain) the input step size is ±200mV, and for Figure 8b (X64 Gain) it is ±154mV. The input signal rise time in either case from the 10% to 90% points was 1μs. Settling time to within ±0.01% for the X1 case is less than 30μs, whereas settling time for the X64 case is less than 300μs. Output voltage settling time was also measured for a change in gain settings (i.e. DAC code) with a steady input signal. The oscilloscope photograph of Figure 8c shows the output response when changing from a X1 setting to a X64 setting. The input signal level was a constant −154mV. The output voltage settling time to within ±0.01% is approximately 300μs going from X1 to X64, while going from X64 to X1 it is approximately 60μs. The dynamic response of the AD7545LN & AD OP-07 PGA combination was similarly measured and found to be indistinguishable from the above.

When ac signals are to be amplified the nonlinearity or distortion of the system becomes important. CMOS DACs are manufactured with high quality, thin-film resistors having very low noise and very small voltage coefficient. In addition, very little of the input signal voltage is developed across the signal-steering switches of the ladder network; hence most of the distortion in the output signal

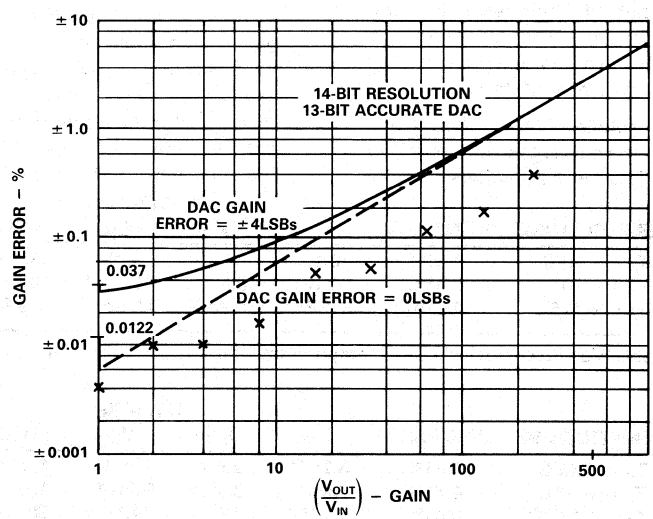

Figure 7. Measured vs. Theoretical Worst Case Gain Errors for AD7534-Based System

s due to the op amp. Table V compares the performance of the 12-bit AD7545LN with the 14-bit AD7534KN in terms of harmonic distortion. In both cases, at any particular gain setting, the input signal level was adjusted to provide a 6V rms output signal level. The test frequency was 200Hz, and distortion levels were measured using a Hewlett Packard HP339A Distortion Measurement Set with it's 3rd order, 30kHz low pass filter function switched into the signal path. Results for the two PGA systems are very similar. Table VI compares the two systems in terms of voltage noise performance over a bandwidth from 22Hz to 22kHz. To help the comparison, the same low noise AD OP-27 op amp was used in both circuits. The superior performance of the AD7534-based system is due to the lower value of DAC ladder resistance compared with the AD7545.

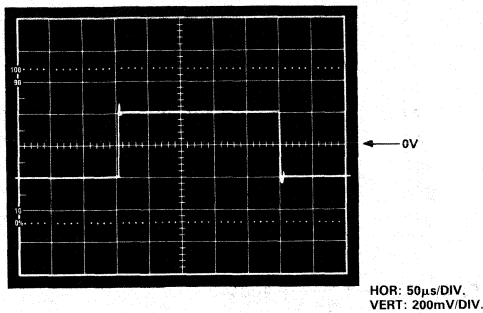

Figure 8a. Gain of 1. Output Response to ±200mV Step Input AD7534 System

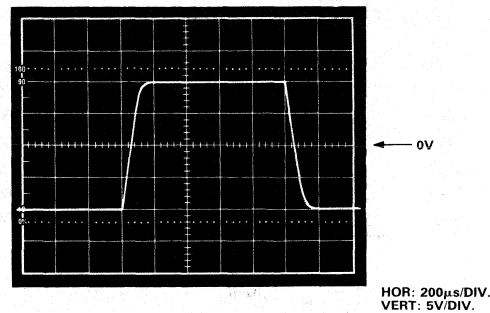

Figure 8b. Gain of 64. Output Response to ±154mV Step Input, AD7534 System

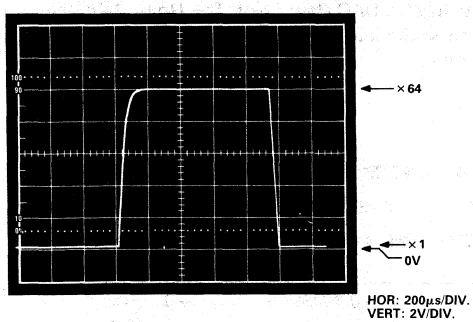

Figure 8c. Output Response when Gain Switching between ×1 and ×64. Constant Input Signal of −154mV. AD7534 System

GAIN	AD7545LN & AD OP-27	AD7534KN & AD OP-27
X 1	5.5µV	4.5µV
X 2	9µV	7µV
X 4	18µV	12µV
X 8	35µV	23µV
X 16	72µV	45µV
X 32	145µV	89µV
X 64	285µV	175µV

Table VI. Output Voltage Noise vs. Gain Settings. Readings are rms, 22Hz to 22kHz.

GAIN	AD7545LN & AD OP-07	AD7534KN & AD OP-07
X 1	<−90dB	<−90dB
X 2	<−90dB	<−90dB
X 4	<−90dB	<−90dB
X 8	−89dB	−88dB
X 16	−86dB	−86dB
X 32	−82dB	−83dB
X 64	−76dB	−79dB

Table V. Total Harmonic Distortion Levels vs. Gain Settings for a Constant 6V rms Output Signal

REFERENCES

1. Application Note; "Gain Error and Gain Temperature Coefficient of CMOS Multiplying DACs" by Phil Burton. Available from Analog Devices, Publication No. E630c-5-3/86.
2. Handbook; "CMOS DAC Application Guide." Available from Analog Devices, Publication No. G872a-15-4/86.
3. Clayton, G.B., "Operational Amplifiers, Second Edition," Butterworths (1979).
4. Transaction Brief; "Expression for the Output Resistance of a Switched R-2R Ladder Network" by E. David Erb and Gregory M. Wierzba. IEEE Trans. Circuits & Systems, Vol CAS-30, No 3, March 1983, p167-169.

APPENDIX 1

SYSTEM GAIN ERROR

GAIN	CODE N	RELATIVE ACCURACY X = +0.5	RELATIVE ACCURACY X = −0.5
1	4095	+0.0122%	+0.0366%
2	2048	−0.0244%	+0.0244%
4	1024	−0.0488%	+0.0489%
8	512	−0.0976%	+0.0978%
16	256	−0.195 %	+0.1957%
32	128	−0.389 %	+0.3922%
64	64	−0.775 %	+0.7874%
128	32	−1.539 %	+1.587 %
256	16	−3.03 %	+3.226 %
512	8	−5.88 %	+6.666 %

$$\text{Error (\%)} = -\left(\frac{X}{N+X}\right)\cdot 100\% \qquad (9a)$$

Table A1. Computed Values for Equation 9a for 12-Bit Resolution (n = 12), 12-Bit Accurate (X = ±0.5LSBs) DAC with Zero DAC Gain Error. The Unity Gain Values Include an Additional +0.0244% Error Term as Discussed in the Text.

SYSTEM GAIN ERROR

GAIN	CODE N	RELATIVE ACCURACY X = +1	RELATIVE ACCURACY X = −1
1	16383	0%	+0.0122%
2	8192	−0.0122%	+0.0122%
4	4096	−0.0244%	+0.0244%
8	2048	−0.0488%	+0.0488%
16	1024	−0.0976%	+0.0978%
32	512	−0.195 %	+0.1957%
64	256	−0.389 %	+0.3922%
128	128	−0.775 %	+0.7874%
256	64	−1.539 %	+1.587 %
512	32	−3.03 %	+3.226 %

$$\text{Error (\%)} = -\left(\frac{X}{N+X}\right)\cdot 100\% \qquad (9a)$$

Table A3. Computed Values for Equation 9a for 14-Bit Resolution (n = 14), 13-Bit Accurate (X = ±1LSBs) DAC with Zero DAC Gain Error. The Unity Gain Values Include an Additional +0.0061% Error Term as Discussed in the Text.

SYSTEM GAIN ERRORS

		DAC GAIN ERROR, Δ = +0.0012		DAC GAIN ERROR, Δ = −0.0012	
GAIN	CODE N	REL. ACCURACY X = +0.5	REL. ACCURACY X = −0.5	REL. ACCURACY X = +0.5	REL. ACCURACY X = −0.5
1	4095	−0.109%	−0.085%	+0.134%	+0.159%
2	2048	−0.146%	−0.098%	+0.096%	+0.145%
4	1024	−0.171%	−0.073%	+0.071%	+0.169%
8	512	−0.219%	−0.024%	+0.022%	+0.218%
16	256	−0.317%	+0.074%	− 0.075%	+0.318%
32	128	−0.511%	+0.270%	−0.270%	+0.513%
64	64	−0.896%	+0.665%	− 0.656%	+0.908%
128	32	−1.658%	+1.464%	−1.420%	+1.709%
256	16	−3.148%	+3.1 %	−2.914%	+3.349%
512	8	−5.977%	+6.537%	− 5.769%	+6.795%

$$\text{Error (\%)} = -\left[\left(\frac{\Delta}{1+\Delta}\right)+\left(\frac{1}{1+\Delta}\right)\left(\frac{X}{N+X}\right)\right]\cdot 100\% \qquad (8)$$

Table A2. Computed Values for Equation 8 for 12-Bit Resolution (n = 12), 12-Bit Accurate (X = ±0.5LSBs) DAC with DAC Gain Error of ±5LSBs (Δ = ±5/4096). The Unity Gain Values Include an Additional +0.0244% Error Term as Discussed in the Text.

SYSTEM GAIN ERRORS

		DAC GAIN ERROR, Δ = +0.000244		DAC GAIN ERROR, Δ = −0.000244	
GAIN	CODE N	REL. ACCURACY X = +1.0	REL. ACCURACY X = −1.0	REL. ACCURACY X = +1.0	REL. ACCURACY X = −1.0
1	16383	−0.025%	−0.012%	+0.024%	+0.037%
2	8192	−0.037%	−0.012%	+0.012%	+0.037%
4	4096	−0.049%	0 %	0 %	+0.049%
8	2048	−0.073%	−0.024%	−0.024%	+0.073%
16	1024	−0.122%	+0.073%	−0.073%	+0.122%
32	512	−0.219%	+0.171%	−0.171%	+0.22 %
64	256	−0.413%	+0.368%	−0.365%	+0.417%
128	128	−0.799%	+0.763%	−0.751%	+0.812%
256	64	−1.562%	+1.563%	−1.514%	+1.612%
512	32	−3.054%	+3.2 %	−3.007%	+3.251%

$$\text{Error}(\%) = -\left[\left(\frac{\Delta}{1+\Delta}\right) + \left(\frac{1}{1+\Delta}\right)\left(\frac{X}{N+X}\right)\right] \cdot 100\% \qquad (8)$$

Table A4. Computed Values for Equation 8 for 14-Bit Resolution (n = 14), 13-Bit Accurate (X = ±1LSBs) DAC with Gain Error of ±4LSBs (Δ = ±4/16,384). The Unity Gain Values Include an Additional +0.0061% Error Term as Discussed in the Text.

APPENDIX 2

The DAC-related error sources, in addition to linearity error, are code-dependent output resistance, leakage current and gain error. Op amp-related error sources are input offset voltage, input bias current and finite open loop gain. Putting all the error sources together results in a complete but somewhat unwieldly expression for the output voltage. For a more complete treatment of the error sources the reader is referred to References 2, 3. If the op amp of Figure 1b has an open loop gain of A_{OL}, then the output voltage of the circuit is

$$V_{OUT} = -V_{IN} \cdot \frac{R_{EQ}}{R_{FB}} \cdot \left\{ \frac{1}{1 + \left(\frac{1}{A_{OL}}\right)\left(1 + \frac{R_{EQ}}{R_{FB} \parallel R_0}\right)} \right\} \qquad (A1)$$

The $\left(1 + \frac{R_{EQ}}{(R_{FB} \parallel R_O)}\right)$ term in this expression is called the noise gain G_N or closed loop gain $\frac{1}{\beta}$ of the circuit. R_O is the code-dependent output resistance of the DAC which appears between the op amp summing junction and AGND.

The output voltage due to the DAC leakage current I_{LKG}, op amp input offset voltage V_{OS} and input bias current $I_B(-)$ is conveniently found by combining these separate error sources into a single input offset error source E_{OS}. The output voltage due to this error source is thus

$$\left(\frac{1}{\beta}\right) \cdot E_{OS} = G_N \cdot E_{OS}$$

$$\text{Where } E_{OS} = \pm V_{OS} \pm I_{LKG} \cdot R \pm I_B(-) \cdot R \qquad (A2)$$

R is the effective source impedance seen by the inverting input to AGND, i.e.,

$$R = R_{FB} \parallel R_O \parallel R_{EQ}$$

or

$$R = \frac{R_{EQ}}{G_N} \qquad (A3)$$

The output voltage due to E_{OS} is thus

$$\pm V_{OS} \cdot G_N \pm I_{LKG} \cdot R_{EQ} \pm I_B(-) \cdot R_{EQ}$$

or

$$\pm V_{OS} \cdot G_N + (I_{LKG} + I_B(-)) \cdot R_{EQ} \qquad (A4)$$

The noise gain G_N (or closed loop gain $\frac{1}{\beta}$) is a function of the DAC output rsistance. This resistance is quite a complex function of the digital input code, roughly increasing from a minimum value to a maximum value as the DAC code changes from all ones to all zeros respectively. [Reference 4].

The DAC output resistance is also influenced by the type of ladder network used in the DAC design. The two DACs in comparison here use two different networks; the AD7545 is a straight R-2R ladder design while the AD7534 uses a segmented design. This means that the noise gain term will be different for both DACs.

When all of these components are summed together the output voltage expression becomes

$$V_{OUT} = -V_{IN} \cdot \frac{R_{EQ}}{R_{FB}} \left\{ \frac{1}{1 + \left(\frac{1}{A_{OL}}\right)\left(1 + \frac{R_{EQ}(R_{FB} + R_0)}{R_{FB} \cdot R_0}\right)} \right\}$$

$$\pm (I_{LKG} + I_B(-)) \cdot R_{EQ}$$

$$\pm V_{OS} \cdot \left(1 + \frac{R_{EQ}(R_{FB} + R_0)}{R_{FB} \cdot R_0}\right) \qquad (A5)$$

When R_{EQ} is replaced by $\frac{2^n R_{IN}}{N + x}$ and the DAC gain error is taken into account by letting

$$\frac{R_{FB}}{R_{IN}} = (1 + \Delta)$$

$$\text{where } \Delta = \frac{R_{FB} - R_{IN}}{R_{IN}} \qquad (A6)$$

the output expression becomes

$$V_{OUT} = -V_{IN} \cdot \frac{2^n}{(N+x)(1+\Delta)} \left\{ \frac{1}{1+\left(\frac{1}{A_{OL}}\right)\left\{1+\left(\frac{2^n}{(N+x)(1+\Delta)}\right)\left(1+\frac{R_{FB}}{R_0}\right)\right\}} \right\}$$

$$\pm (I_{LKG} + I_B(-)) \cdot \frac{2^n R_{IN}}{N+x}$$

$$\pm V_{OS}\left\{1+\left(\frac{2^n}{(N+x)(1+\Delta)}\right)\left(1+\frac{R_{FB}}{R_0}\right)\right\} \qquad \text{(A7)}$$

Expression A7 can be simplified since the effect of DAC linearity and DAC gain error on the error sources is almost negligible. Similarly, for the system gains under consideration, the gain error factor in expression A1 can reasonably be assumed to be unity, i.e.,

$$\frac{1}{1+\left(\frac{1}{A_{OL}}\right)G_N} \approx 1 \qquad \text{(A8)}$$

When these assumptions are made, expression A7 simplifies to

$$V_{OUT} = -V_{IN}\frac{2^n}{(N+X)(1+\Delta)}$$

$$\pm \{I_{LKG} + IB(-)\} \cdot R_{IN} \cdot (\text{SYSTEM GAIN})$$

$$\pm V_{OS}\left\{1+(\text{SYSTEM GAIN})\left(1+\frac{R_{FB}}{R_O}\right)\right\} \qquad \text{(A9)}$$

As the system gain is increased, so too are the dc offset errors due to leakage currents and input offset voltage.

APPENDIX 3

SYSTEM GAIN	$CODE_{10}$	$(1+R_{FB}/R_O)$	SYSTEM GAIN	$CODE_{10}$	$(1+R_{FB}/R_O)$	SYSTEM GAIN	$CODE_{10}$	$(1+R_{FB}/R_O)$
1	16383	1.08	24	683	1.31	47	349	1.27
2	8192	1.00	25	655	1.25	48	341	1.30
3	5461	1.31	26	630	1.25	49	334	1.25
4	4096	1.00	27	607	1.24	50	328	1.20
5	3277	1.28	28	585	1.27	51	321	1.21
6	2731	1.31	29	565	1.28	52	315	1.26
7	2341	1.21	30	546	1.21	53	309	1.28
8	2048	1.00	31	529	1.14	54	303	1.25
9	1820	1.20	32	512	1.06	55	298	1.26
10	1638	1.25	33	496	1.13	56	293	1.27
11	1489	1.25	34	482	1.20	57	287	1.22
12	1365	1.31	35	468	1.22	58	282	1.24
13	1260	1.22	36	455	1.25	59	278	1.24
14	1170	1.24	37	443	1.26	60	273	1.22
15	1092	1.18	38	431	1.26	61	269	1.24
16	1024	1.03	39	420	1.22	62	264	1.15
17	964	1.18	40	410	1.25	63	260	1.15
18	910	1.23	41	400	1.17	64	256	1.07
19	862	1.23	42	390	1.22			
20	819	1.28	43	381	1.23			
21	780	1.20	44	372	1.22			
22	745	1.27	45	364	1.23			
23	712	1.21	46	356	1.22			

Table A5. Digital Input Codes and Resulting $(1 + R_{FB}/R_O)$ Ratio vs. System Gain for the AD7534

ANALOG DEVICES

AN-320B
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

CMOS DACs and Op Amps Combine to Build Programmable Gain Amplifiers Part II

by John Wynne

This Application Note investigates the performance of dual CMOS DACs as gain determining elements in a PGA system. It details how greater accuracy over a wider dynamic range can be achieved with a dual DAC circuit as opposed to a single DAC solution. A dual DAC approach spreads the required system gain over two stages. This results in an overall system gain which is the product of the individual gain stages, but the overall gain error is essentially only the sum of the individual gain error terms.

Part I of this Application Note[1] looked at the error sources which exist when a CMOS DAC is used as a programmable resistance in the feedback loop of an op amp. A detailed comparison was also made between the performance of a PGA circuit based on a 14-bit DAC, the AD7534, (AD7538) and a 12-bit DAC, the AD7545.

DUAL DAC SOLUTION

With two DACs available in a single package – and, of course, two op amps also available in a single package – it becomes possible to use two simple PGA circuits in series without taking up too much extra printed circuit board area over that required by a single PGA stage. Figure 1 shows such a circuit. Each section contributes its own gain errors and offset errors to the overall error. Gain errors are due to DAC integral nonlinearity and to DAC gain error. Offset errors are due to DAC leakage currents and op amp input bias currents and input offset voltages.

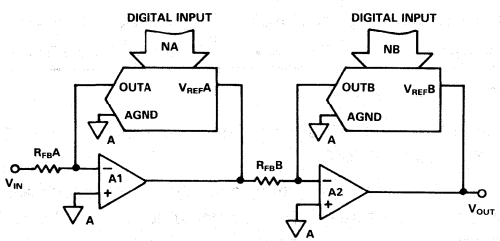

Figure 1. Dual DAC PGA Circuit

The dual DAC circuit does however exaggerate the dc offset error terms. The offset errors of the first stage are multiplied by the gain setting of the second stage. This can lead to quite large dc offset voltages for large gain settings, one third of a volt or more, at the output of the second stage. For this reason, dual DAC PGA systems are only considered suitable for ac signals, and the analysis presented here emphasizes this.

Since DAC gain error can be trimmed to zero (whereas nothing can be done to reduce the integral nonlinearity) it is instructive to first consider both the system gain and the percentage gain error due to DAC nonlinearity alone and then to add the DAC gain error terms. Following this, comparisons are made between the dual DAC performance and the performance of both the single 12-bit and 14-bit DAC systems which were the subject of Part I of this Application Note.

THE BASIC EQUATIONS FOR A DUAL DAC PGA

When only linearity errors are considered, the gain of a single stage is:

$$\frac{V_{OUT}}{V_{IN}} = -\frac{2^n}{(N+x)} \qquad (1)$$

Where n is the resolution of the DAC
x is the DAC linearity error in LSBs
and N is the DAC code in decimal.

With two similar stages in series the system gain is obviously:

$$\frac{V_{OUT}}{V_{IN}} = \left(\frac{2^n}{N_A+X_A}\right) \cdot \left(\frac{2^n}{N_B+X_B}\right) \qquad (2)$$

Where A and B suffixes refer to DAC A (1st stage) and DAC B (2nd stage) respectively.

The percentage gain error of a single stage, again assuming zero DAC gain error, can be expressed as:

$$E(\%) = -\left(\frac{X}{N+X}\right) \cdot 100\% \qquad (3)$$

With two stages in series the overall gain error can generally be taken as the sum of both stages. If the first stage has a percentage gain error of E_A% and the second stage has E_B%, then the total system gain error can be expressed as:

$$\text{Error}(\%) = \left(E_A + E_B + \frac{E_A \cdot E_B}{100}\right)\% \quad (4)$$

For precision PGA systems these individual gain error terms, E_A and E_B, will normally be held to 1% or less. Under these conditions Equation 4 simplifies to:

$$\text{Error}(\%) = (E_A + E_B)\% \quad (5)$$

or using Equation 3:

$$\text{Error}(\%) = -\left(\frac{X_A}{N_A + X_A} + \frac{X_B}{N_B + X_B}\right) \cdot 100\% \quad (6)$$

Equation 2 shows the overall system gain is the product of the individual gain stages, while Equation 6 shows the overall percentage gain error is effectively the sum of the individual gain error terms. Hence, greater accuracy over a wider dynamic range is possible with a dual DAC PGA system over any single DAC solution. Also from Equation 2 the only constraint on the codes to both DACs, N_A and N_B, is that the overall system gain is achieved. Equation 6, however, indicates that in order to keep the gain error as small as possible, the codes to both DACs must be chosen to be as equal to each other as possible. Simultaneously meeting both these requirements results in the optimum system.

Equation 3, which gives the percentage gain error of a single stage, is correct for all gain settings except unity when an additional term must be added to the simple equation. This additional error term is due to an LSB worth of signal current being "lost" in the R-2R ladder termination resistor thereby making an ideal X1 gain impossible.

Therefore, for each DAC with a gain setting of unity (all 1s setting) an error term equal to 1LSB (expressed as a percentage) must be added to the total error term. For a dual-DAC system Equation (6) applies for all combinations of gain settings except those cases where an individual gain stage is set to unity gain. An extra error term must be added for each unity gain setting.

When both linearity errors and DAC gain errors are included, the gain of a single stage is:

$$\frac{V_{OUT}}{V_{IN}} = -\frac{2^n}{(N + X)(1 + \Delta)} \quad (7)$$

$$\text{where } (1 + \Delta) = \frac{R_{FB}}{R_{DAC}}$$

With two similar stages in series, the system gain is simply the product of the two stage gains or:

$$\frac{V_{OUT}}{V_{IN}} = \left(\frac{2^n}{(N_A + X_A)(1 + \Delta_A)}\right) \cdot \left(\frac{2^n}{(N_B + X_B)(1 + \Delta_B)}\right) \quad (8)$$

For an individual stage the percentage gain error is:

$$\text{Error}(\%) = -\left\{\frac{\Delta}{(1+\Delta)} + \frac{1}{(1+\Delta)}\left(\frac{X}{N+X}\right)\right\} \cdot 100\% \quad (9$$

Using similar reasoning to that used before in conjunctio with Equations 4 and 5, the overall percentage gain erro is again the sum of the individual error terms or:

$$\text{Error}(\%) = -\left[\left\{\frac{\Delta_A}{(1+\Delta_A)} + \frac{1}{(1+\Delta_A)}\left(\frac{X_A}{N_A+X_A}\right)\right\} + \left\{\frac{\Delta_B}{(1+\Delta_B)} + \frac{1}{(1+\Delta_B)}\left(\frac{X_B}{N_B+X_B}\right)\right\}\right] \cdot 100\% \quad (10$$

Again, Equation 10 is valid for all combinations of gai settings except those cases where an individual gai stage is set to unity gain. For each unity gain setting a additional 1LSB (expressed as a percentage) must b added to the output error term.

COMPARING THE ERRORS

It is interesting to compare the dual DAC performance with both single 12-bit and 14-bit DACs on the basis o DAC nonlinearity alone and then to add the DAC gain error terms. Table A1 in Appendix 1 shows computed values of equation 6 at selected combinations of gain settings. The effect of DAC gain error is included in Table A2 which shows computed values of Equation 10. Also included in the tables are the individual codes for each DAC at the selected gain settings. The error terms in both tables are calculated for a PGA system based on the AD7547LN. This is a dual 12-bit DAC with full parallel loading, housed in a skinny 24-pin package. Relevant specifications for this device are shown in Table I. The gain error analysis leading to Equations 6 and 10 assumed that the individual loop gains of the dual DAC system were sufficiently high so as to cause no appreciable error. This is a valid assumption at dc and low frequencies since extreme gain settings are not used in practice. This topic is dealt with later in greater detail.

Parameter	AD7547LN $T_A = +25°C$	AD7547LN $T_A = +70°C$	AD7547UQ $T_A = +125°C$
Resolution, n	12 bits	12 bits	12 bits
Relative Accuracy, X (Integral Linearity)	± 1/2LSB max	± 1/2LSB max	± 1/2LSB ma
Gain Error	± 1LSB max	± 1LSB max	± 2LSB max
Output Leakage, I_{LKG}	10nA max	150nA max	250nA max
Input Resistance	20kΩ max	20kΩ max	20kΩ max

NOTE: $V_{DD} = +12V$ to $+15V$

Table I. AD7547 Dual DAC Specifications at $T_A = +25°C$ $+70°C$ and $+125°C$

Figure 2 plots the results due solely to DAC nonlinearity errors of the three PGA circuits under comparison. As might have been expected, the dual DAC circuit produces the highest errors at the lowest gains. For system gains of 1 to 4 the dual 12-bit DAC circuit produces the highest errors. At a gain of 4, however, its performance equals

that of a single 12-bit DAC circuit and thereafter is increasingly superior with increasing gain. At a gain of 64 its performance equals that of a single 14-bit DAC circuit, and thereafter is increasingly superior with increasing gain. The straight line representation in Figure 2 of percentage error versus gain for the dual DAC circuit, however, is an idealization. Because the total gain error now depends on two DAC codes, a number of code combinations exist which will give the same gain (within some chosen tolerance) but produce widely different system gain errors. For instance, if the required gain is 64, the obvious and best distribution of gain settings for the dual DAC system is X8 and X8. From Figure 2 this will give a system error of approximately ±0.2%. If gain settings of X1 & X64 were

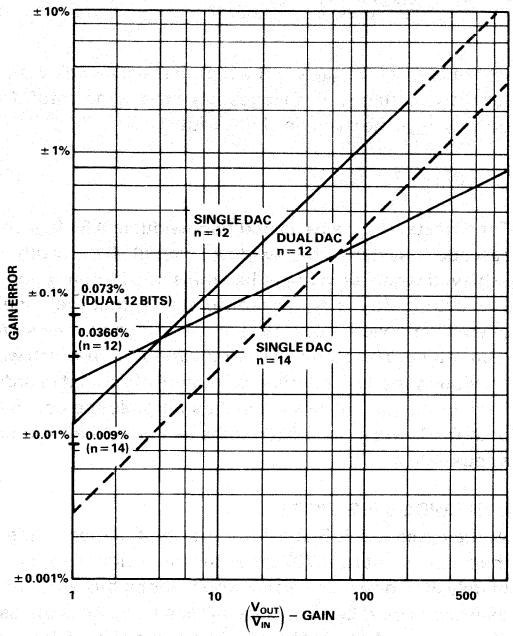

Figure 2. Comparison of Worst Case Gain Errors between PGA Systems Based on Single and Dual 12-Bit Resolution, 12-Bit Accurate DACs and a 14-Bit Resolution, 14-Bit Accurate DAC. DAC Gain Error Is Zero in All Cases

chosen, the system error would be slightly greater than ±0.8%, a four-fold increase. For this example the optimum codes are obvious. However, for many other system gains which are not integer powers of two and which may not even be integers, choosing the optimum gain distribution will require some attention.

The total error results of Table A2 are plotted in Figure 3 along with results from Part I of this Application Note, of a single 12-bit DAC (AD7545LN) and a single 14-bit DAC (AD7534KN). The very tight gain error of the dual DAC results in a system where the gain error, at all gain settings, can actually be less than that of a single 12-bit DAC.

DYNAMIC PROBLEMS
STABILITY AND COMPENSATION

For the purposes of investigating the dynamic performance of DAC-based PGA systems, an equivalent circuit for a single DAC and op amp stage is shown in Figure 4. Programmable circuit gain is ideally set by resistors R_{FB} and R_{EQ} around the op amp. A real DAC, however, also adds both a code dependent resistance R_O and code dependent capacitance C_O between the op amp summing junction and ground. Capacitance C_2 represents the sum of both stray capacitance and any added capacitance across the equivalent feedback resistor. Since this equivalent resistance can normally be quite large, the effect of C_2 on the frequency response can be very important.

With an internally compensated op amp, whose open loop gain rolls off due to a single dominant pole, the 3dB bandwidth for an ideal PGA system is very easily found from the constant gain-bandwidth product concept. For instance, if the gain-bandwidth product of A1 in Figure 4 is 1.10^6, then for a X1 inverting gain ($1/\beta = 2$), the ideal closed-loop bandwidth is 500kHz; whereas, for a X64 gain, it is reduced to approximately 15kHz. Thus, signal bandwidth is inversely proportional to gain.

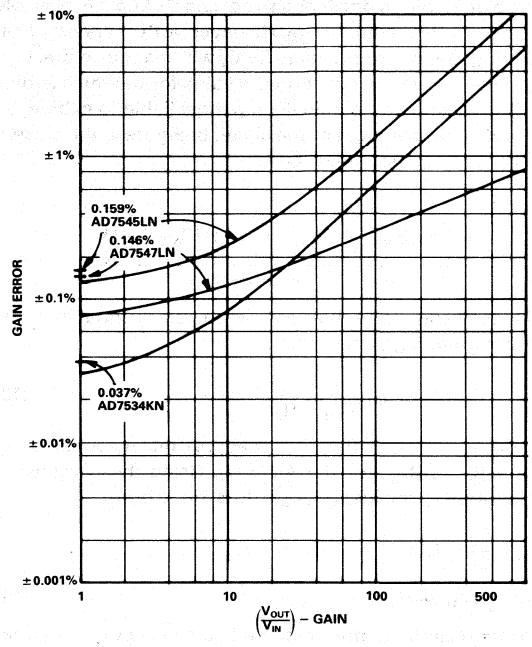

Figure 3. Comparison of Theoretical Worst Case Gain Errors between the Three PGA Systems when DAC Gain Errors Are Included

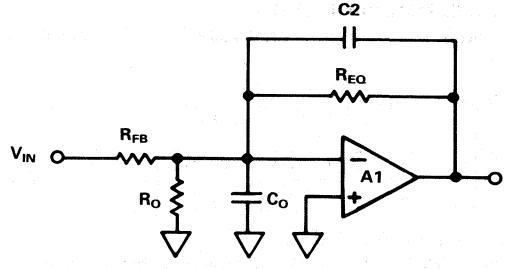

Figure 4. Equivalent Circuit for a Single Stage

The output capacitance of the DAC, C_0, combines with R_{EQ} to add an extra pole to the closed-loop response. This pole upsets the unconditional stability characterized by a loop gain slope of −6dB/octave and a −90° maximum phase shift by adding extra phase shift. The amount of extra phase added depends on the position of the spurious pole with respect to the undisturbed unity loop gain frequency. If the phase margin is reduced by too much, the system stability is upset and gain peaking becomes increasingly evident. For a PGA system intended to amplify ac signals, gain peaking can be a serious source of error. Loop stability margins can be restored and gain peaking eliminated by placing a phase lead capacitor C_2 across the feedback resistor R_{EQ}.

The value of C_2 plays an important part in the system performance. If C_2 is too small (with respect to C_0), gain peaking may occur. If it is too large, the bandwidth of the amplifier will be unnecessarily reduced. Ideally, the ratio of C_0 to lead capacitance C_2 should be equal to the closed-loop gain minus one. For a fixed gain setting it is possible to satisfy this requirement. However, with every new gain setting, the optimum value for C_2 will change. In practice, the best solution is to choose a value for C_2 which avoids gain peaking at the gain setting most liable to exhibit it – the X1 gain setting. For the circuit of Figure 4, the closed-loop gain can be expressed as:

$$\frac{1}{\beta} = \frac{\left\{1+R_{EQ}\left(\frac{1}{R_{FB}}+\frac{1}{R_0}\right)\right\}\left\{1+S(C_0+C_2)R\right\}}{1+SC_2R_{EQ}} \qquad (11)$$

At low frequencies the closed-loop gain is equal to the dc noise gain, G_N, of the circuit:

$$G_N = 1 + R_{EQ}\left(\frac{1}{R_{FB}}+\frac{1}{R_0}\right) \qquad (12)$$

At the X1 gain setting, R_{EQ} is approximately equal to R_{FB} and $R_0 = 3\ R_{FB}$ (for the AD7547). Using these values in Equation 12 the ideal C_0/C_2 ratio is found from:

$$\frac{C_0}{C_2} = G_N - 1 = \frac{1}{\beta} - 1$$

$$\text{or } C_2 = 0.75\,C_0 \qquad (13)$$

At gain settings other than unity the value of C_2 is larger than necessary to ensure no gain peaking occurs. The drawback to this is that the overall bandwidth is reduced and the settling time is increased.

SMALL-SIGNAL BANDWIDTH

For Figure 4, the output signal voltage can be expressed as:

$$V_{OUT} = -V_{IN} \cdot \frac{R_{EQ}}{R_{FB}}\left(\frac{1}{(1+SC_2R_{EQ})}\right) X \left[\frac{1}{1+\frac{1}{A(\omega)}\left\{\frac{G_N(1+S(C_0+C_2)R)}{1+SC_2R_{EQ}}\right\}}\right] \qquad (14)$$

where $R = R_{FB} \| R_0 \| R_{EQ}$
and $A(\omega)$ is the open-loop gain of the amplifier, a complex quantity.

The closed-loop signal bandwidth is set at $1/2\pi\, C_2\, R_{EQ}$, by the value chosen for C_2. In terms of the digital input code to the DAC the 3dB signal bandwidth can be expressed as:

$$f_{3dB} = \frac{D}{2\pi C_2 RLAD} \qquad (15)$$

and since system gain is inversely proportional to D:

$$f_{3dB} = \frac{1}{GAIN} \cdot \left(\frac{1}{2\pi C_2 RLAD}\right) \qquad (16)$$

When two stages are cascaded the overall bandwidth is:

$$1.1\sqrt{\left(\frac{0.35}{f_{3dBA}}\right)^2+\left(\frac{0.35}{f_{3dBB}}\right)^2} \qquad (17)$$

where f_{3dBA} and f_{3dBB} are the high frequency 3dB cutoffs for the two stages. If n stages have the same cutoff frequency, f_{3dB}, the cascaded bandwidth is:

$$\sqrt{2^{1/n}-1} \cdot f_{3dB} \qquad (18)$$

For two stages, the cascaded bandwidth is 0.64 f_{3dB}. In a cascaded system, the widest overall bandwidth is achieved when each stage has the same high frequency 3dB cutoff. This is equivalent to saying that the widest overall bandwidth occurs when each stage has the same gain setting. This statement complements that which was previously made in relation to Equation 6 viz. that in order to keep the gain error as small as possible the codes to both DACs must be chosen to be as equal to each other as possible.

DYNAMIC GAIN ERRORS

At the cascaded 3dB cutoff frequency, the magnitude o the output signal is 0.707 times the input signal – approxi mately 30% down. In certain applications this amount o gain error could be excessive. In these instances the use able bandwidth should be considered in terms of the addi tional gain error resulting from the fall-off in closed-loop signal gains. The reduced bandwidth for the cascadec system is determined by the gain equation:

$$\text{Amplitude} = \frac{1}{1+\left(\frac{f}{f_{3dB}}\right)^2} \qquad (19$$

Both stages are assumed to have the same high fre quency 3dB cutoff. For example, the reduced bandwidth necessary to restrict additional gain errors to below 0.1% or the total amplitude to 0.999, can be found from Equa tion 19 to be:

$$f = 0.032\, f_{3dB}$$

$$\text{or } f = 0.032 \cdot (f_{CASC}/0.64)$$

$$= 1/20^{th} \text{ of } f_{CASC} \qquad (20$$

This means that signal frequencies up to 1/20 of the cas caded bandwidth will have less than 0.1% of additiona gain error.

nite loop gain can also contribute additional gain error. ıe square bracketed term in Equation 14 is called the ain error factor. Ideally this factor is equal to unity; the nount by which it differs from unity is equivalent to the ddditional gain error due to finite loop gain. Appendix 2 ontains both an analysis of the gain error factor and a orked example of its magnitude in a typical application. ı comparison with the gain errors introduced by the sig- al bandwidth, the analysis indicates that additional gain rrors due to finite loop gain can be ignored.

OISE AND DISTORTION

loise does not cause gain errors in the sense of introduc- ıg constant multiplicative terms. Rather, it acts to reduce ıe signal to distortion ratio of the PGA system. This topic ; well covered by standard op amp texts and will not be iscussed here. Analog Devices' CMOS DACs are man- factured with high quality, thin-film resistors and exhibit ery little excess noise over that expected from an equiva- ent Johnson noise source. In addition, since these thin- lm resistors have a very small voltage coefficient, any listortion through the R-2R ladder results from R_{ON} mod- lation of the signal-steering switches. In practice, how- ver, distortion is usually limited by the op amp itself.

\C COUPLING THE OUTPUT

n a dual DAC PGA system the dc offset errors of the first tage are multiplied by the gain of the second stage. This an lead to quite large dc error voltages at the second tage output. AC coupling the output of the second stage bviously eliminates the problem but adds a low fre- uency pole to the response. This low frequency pole will ontribute gain errors to low frequency input signals. The ituation is exactly analogous to the gain error caused by igh frequency fall-off discussed previously. Placing the ac coupling capacitor after the second stage also allows a single trim potentiometer to correct for DAC gain errors of both DACs. However, the temperature coefficient of the rim potentiometer will not match the temperature coeffi- ient of the thin-film resistors of the DACs resulting in a change in DAC gain error over temperature. Taking into account the fact that the DAC gain error of the AD7547LN s specified to remain within ±1LSB from 0 to +70°C, better circuit performance over temperature may result if no trim potentiometer is used.

TEST RESULTS

AC measurements for a PGA system based on an AD7547LN dual-DAC and two AD OP-27E op amps are plotted in Figure 5. The "starred" points indicate the measured errors at the selected gain settings between X1 and X512. At each gain setting the input signal level was adjusted to provide a 6V rms output signal level. The test frequency was 200Hz. As a comparison, the worst-case theoretical curves from Figures 2 and 3 for a dual DAC PGA system area also included in Figure 5. As mentioned previously in the text, these solid-line curves are in fact idealizations since in a dual DAC system any desired gain can usually be realized by different code combinations.

Parameter	AD OP-27E $T_A = +25°C$	AD OP-27E $T_A = +70°C$	AD OP-27A $T_A = +125°C$
Open-Loop Gain, A_{OL}	1.10^6 min	$0.75.10^6$ min	$0.6.10^6$ min
Input Bias Current, I_B (–)	40nA max	60nA max	60nA max
Input Offset Voltage, V_{OS}	25μV max	50μV max	60μV max

NOTE: $V_{DD} = +15V$, $V_{SS} = -15V$

Table II. AD OP-27 Specifications at $T_A = +25°C$, +70°C and +125°C

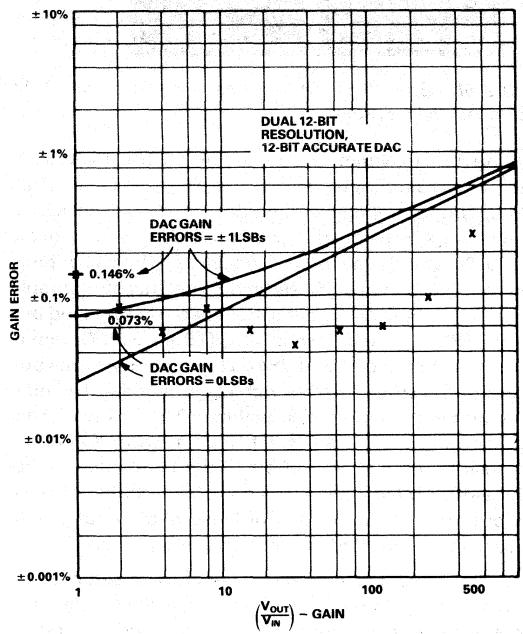

Figure 5. Measured vs. Theoretical Worst Case Gain Errors for AD7547-Based System

Output voltage settling time was measured for a step input change at two fixed gain settings (X1, X64). The output response is shown in the photographs of Figure 6a and 6b. For Figure 6a (X1 Gain) the input step size is ±200mV, and for Figure 6b (X64 Gain) it is ±154mV. The input signal rise and fall time in either case from the 10% to 90% points was 400ns.

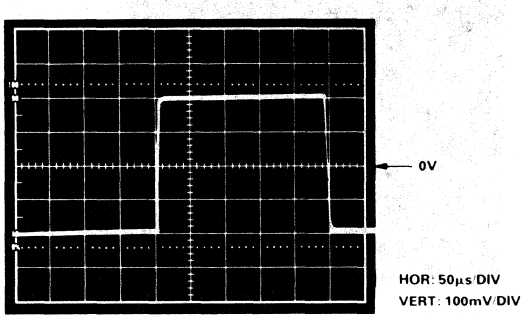

Figure 6a. Gain of 1. Output Response to ±200mV Step Input, AD7547 System

8

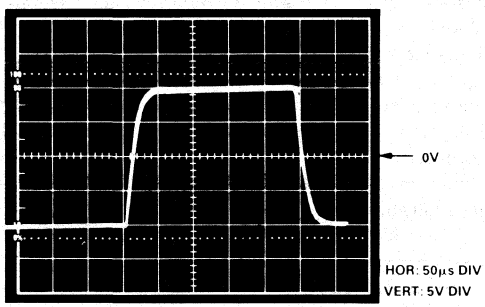

Figure 6b. Gain of 64. Output Response to ±154mV Step Input, AD7547 System

Settling time to within ±0.01% for the X1 case is less than 15μs; whereas, settling time for the X64 case is less than 100μs. Output voltage settling time was also measured for a change in gain settings with a steady input signal. Figure 6c shows the output response when changing from a X1 setting to a X64 setting. To get this photograph the codes in both DACs are simultaneously switched between 4095_{10} (X1 and X1) and 512_{10} (X8 and X8). The input loading structure of the AD7547 allows simultaneous updating of the DAC registers with common data. The output voltage settling time to within ±0.01% is less than 70μs going from X1 to X64, while going from X64 to X1 it is less than 15μs. The effect of the high-speed, wide bandwidth AD OP-27 op amps on voltage settling time is obvious from a comparison of these photographs with those of Figure 8 in Part I of this Application Note.

Total harmonic distortion levels versus gain settings are shown in the right hand column of Table III. For comparison, the harmonic distortion levels measured previously for the AD7545 – and AD7534 – based PGAs are also included in this table. In all three cases, for any particular gain setting, the input signal level was adjusted to provide a 6V rms output signal level. The test frequency was 200Hz. The bandwidth of the Hewlett Packard HP339A distortion measuring set on which the measurements were taken was purposely limited by the set's 3rd order, 30kHz low pass filter to minimize the noise bandwidth.

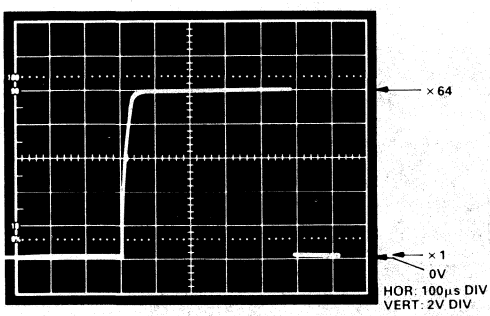

Figure 6c. Output Response when Gain Switching between X1 and X64. Constant Input Signal of +154mV, AD7547 System.

GAIN	AD7545LN & AD OP-07	AD7534LN & AD OP-07	AD7547LN & AD OP-27 (2)
X 1	< −90dB	< −90dB	< −90dB
X 2	< −90dB	< −90dB	< −90dB
X 4	< −90dB	< −90dB	< −90dB
X 8	−89dB	−88dB	< −90dB
X 16	−86dB	−86dB	< −90dB
X 32	−82dB	−83dB	< −89dB
X 64	−76dB	−79dB	< −84dB
X 128	Not Measured	Not Measured	< −79dB
X 256	Not Measured	Not Measured	< −74dB
X 512	Not Measured	Not Measured	< −68dB

Table III. Total Harmonic Distortion Levels vs. Gain Settings for a Constant 6V rms Output Signal

The small signal bandwidth at a number of gain setting is shown in Table IV. These measured frequencies follo closely the values predicted from Equations 15 – 18. value of 47pF was used for the phase lead capacitor, C in each of the two stages. With large signal levels and hig signal frequencies distortion becomes severe due to op amp limitations. For example, for the AD OP-27 with a output signal level of 6V RMS the limit before distortio rapidly increases is typically 40kHz. The input signal level used for the bandwidth measurements were varied wit gain setting in order to maintain total harmonic distortio below 75dB. Thus at high gain settings (small band widths), signal levels were high in order to maximize th signal-to-noise ratio; at low gain settings (wid bandwidths), signal levels were low in order to avoid o amp induced distortion.

System Gain	F_{CASC} Measured	F_{CASC} Computed
X 1	135kHz	135.5kHz
X 4	60kHz	67.7kHz
X 8	44kHz	43kHz
X 16	32kHz	34kHz
X 64	15.4kHz	16.9kHz
X 256	7.8kHz	8.4kHz

Table IV. Small Signal Bandwidth vs. Gain Setting fo Dual-DAC PGA System

Table V compares the voltage noise performances of the AD7547-based PGA and the single DAC-based systems.

GAIN	AD7545LN & AD OP-27	AD7534KN & AD OP-27	AD7547LN & AD OP-27
X 1	5.5μV	4.5μV	6.2μV
X 2	9μV	7μV	11.2μV
X 4	18μV	12μV	17μV
X 8	35μV	23μV	31.5μV
X 16	72μV	45μV	63μV
X 32	145μV	89μV	127μV
X 64	285μV	175μV	260μV

Table V. Output Voltage Noise vs. Gain Settings. Readings are rms, 22Hz to 22kHz.

To help the comparison the same AD OP-27E op amp was used in both of the single DAC circuits and in the first stage of the dual DAC circuit. No phase lead compensation capacitors were used in any of the circuits. The results indicate that there is little difference between the single 12-bit and dual 12-bit PGAs in terms of output voltage noise.

FIXED GAIN CIRCUIT

In addition to programmable gain in a system, a certain amount of fixed gain can also be useful. It is possible to combine both functions around a single DAC plus op amp. The circuit is shown in Figure 7. Resistors R1 and R2 determine the fixed amount of gain. Normally, the DAC input resistance would be included in the gain expression making circuit gain a function of that resistance. This problem is eliminated (Reference 2) by adding an input resistor in series with R_{FB}. This input resistor, R_{IN}, has a value equal to the parallel combination of R1 and R2. The output voltage is:

$$V_{OUT} = -\frac{V_{IN}}{D} \cdot \left(1+\frac{R1}{R2}\right) \qquad (21)$$

Resistors R1, R2 and R_{IN} should have similar temperature coefficients, but they need not match the temperature coefficient of the DAC. Minor gain adjustments can be made by varying the attenuator ratio as adjustment sensitivity is almost unaffected by R1.

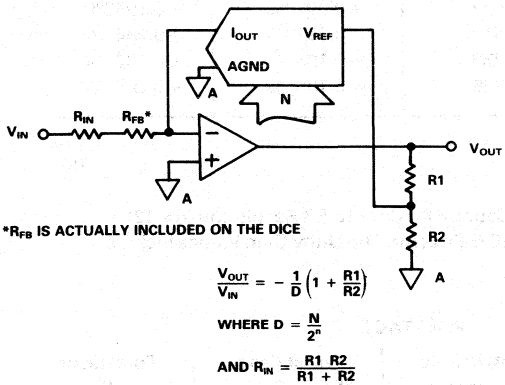

Figure 7. Adding Some Fixed Gain to the Basic Stage

REFERENCES

1. "CMOS DACs and Op Amps Combine to Build Programmable Gain Amplifiers, Part I" by John Wynne. Available from Analog Devices, Publication No. E1037-15-1/87.
2. "Input Resistor Stabilizes MDAC's Gain" by Paul Brokaw, EDN, January 7, 1981, pp. 210-211.
3. Transaction Brief; "Expression for the Output Resistance of a Switched R-2R Ladder Network" by E. David Erb and Gregory M. Wierzba. IEEE Trans. Circuits & Systems, Vol. CAS-30, No. 3, March 1983, pp. 167-169.

APPENDIX 1

	1st STAGE		2nd STAGE		
GAIN	**DAC A Code N_A (Decimal)**	**Worst Case Error, E_A**	**DAC B Code N_B (Decimal)**	**Worst Case Error, E_B**	**Total Error $E_A + E_B$**
1	4095	+0.0366%	4095	+0.0366%	+0.0732%
2	2896	+0.0173%	2896	+0.0173%	+0.0345%
3	2365	+0.0211%	2365	+0.0211%	+0.0423%
4	2048	+0.0244%	2048	+0.0244%	+0.0488%
8	1024	+0.0489%	2048	+0.0244%	+0.0733%
16	1024	+0.0489%	1024	+0.0489%	+0.0978%
32	512	+0.0978%	1024	+0.0489%	+0.1467%
64	512	+0.0978%	512	+0.0978%	+0.1956%
128	256	+0.196 %	512	+0.0978%	+0.294 %
256	256	+0.196 %	256	+0.196 %	+0.392 %
512	128	+0.392 %	256	+0.196 %	+0.588 %

$$\text{Error}(\%) = -\left\{\frac{XA}{NA+XA} + \frac{XB}{NB+XB}\right\} \cdot 100\% \qquad (6)$$

Table A1. Computed Worst Case Gain Error for Equation 6 for Dual 12-Bit Resolution (n = 12), 12-Bit Accurate (X = ±0.5LSBs) DACs with Zero DAC Gain Error. The Unity Gain Values Include Additional Error Terms as Discussed in the Text.

	1st STAGE		2nd STAGE		
System Gain	**DAC A Code N_A (Decimal)**	**Worst Case Error, E_A**	**DAC B Code N_B (Decimal)**	**Worst Case Error, E_B**	**Total Error $E_A + E_B$**
1	4095	+0.073 %	4095	+0.073 %	+0.0146%
2	2896	+0.0417%	2896	+0.0417%	+0.083 %
3	2365	+0.0456%	2365	+0.0456%	+0.091 %
4	2048	+0.049 %	2048	+0.049 %	+0.098 %
8	1024	+0.073 %	2048	+0.049 %	+0.122%
16	1024	+0.073 %	1024	+0.073 %	+0.146%
32	512	+0.122 %	1024	+0.073 %	+0.195 %
64	512	+0.122 %	512	+0.122 %	+0.244 %
128	256	+0.22 %	512	+0.122 %	+0.342 %
256	256	+0.22 %	256	+0.22 %	+0.44 %
512	128	+0.417 %	256	+0.22 %	+0.637 %

$$\text{Error}(\%) = -\left[\left\{\frac{\Delta_A}{(1+\Delta_A)} + \frac{1}{(1+\Delta_A)}\left(\frac{X_A}{N_A+X_A}\right)\right\} + \left\{\frac{\Delta_B}{(1+\Delta_B)} + \frac{1}{(1+\Delta_B)}\left(\frac{X_B}{N_B+X_B}\right)\right\}\right] \cdot 100\% \qquad (10)$$

Table A2. Computed Worst Case Gain Error for Equation 10 for Dual 12-Bit Resolution (n = 12), 12-Bit Accurate (X = ±0.5LSBs) DACs with DAC Gain Errors of ±1LSB (= ±1/4096). The Unity Gain Values Include Additional Error Terms as Discussed in the Text.

PPENDIX 2

ne gain error factor from equation 14 is:

$$\frac{1}{1+\frac{1}{\beta A(\omega)}} = \frac{1}{1+\frac{1}{A(\omega)}\left\{\frac{G_N(1+S[C_0+C_2]R)}{1+SC_2R_{EQ}}\right\}} \quad \text{(A1)}$$

r

$$\frac{1}{1+\frac{1}{\beta A(\omega)}} = \frac{1}{1+\frac{1}{A(\omega)}\left\{\frac{G_N(1+j\omega/\omega_1)}{1+j\omega/\omega_2}\right\}} \quad \text{(A2)}$$

where $\omega_1 = \frac{1}{(C_0+C_2)\cdot R}$

and $\omega_2 = \frac{1}{C_2 \cdot R_{EQ}}$

quation A2 can be rewritten emphasizing the phasor ualities of its terms:

$$\frac{1}{1+\frac{1}{\beta A(\omega)}} = \frac{1}{1+\left|\frac{G_N \cdot r_1}{A_{OL}(\omega)\cdot r_2}\right| \angle\phi} \quad \text{(A3)}$$

where $r_1 = \sqrt{1+(\omega/\omega_1)^2}$

$r_2 = \sqrt{1+(\omega/\omega_2)^2}$

$A_{OL}(\omega) = A_{OL}/\sqrt{1+(\omega/\omega_P)^2}$

A_{OL} is the dc value of the open loop gain and ω_P is the break frequency of the op amp

Also, $\angle\phi = \theta_1 - \theta_2 + \theta_3$ (A4)

where $\theta_1 = t_{an}{}^{-1}\left(\frac{\omega}{\omega_1}\right)$

$\theta_2 = t_{an}{}^{-1}\left(\frac{\omega}{\omega_1}\right)$

$\theta_3 = t_{an}{}^{-1}\left(\frac{\omega}{\omega_P}\right)$

The magnitude of the gain error factor can now be written as:

$$\left|\frac{1}{1+\frac{1}{\beta A(\omega)}}\right| = \frac{1}{\sqrt{1+\left|\frac{G_N \cdot r_1}{A_{OL}(\omega)\cdot r_2}\right|^2 + 2C_{OS}\theta\left|\frac{G_N \cdot r_1}{A_{OL}(\omega)\cdot r_2}\right|}} \quad \text{(A5)}$$

8

since the cosine of an angle θ between 90° and 180° is negative, then it is possible for the gain error factor of equation A5 to be greater than unity causing gain peaking. From Equation A4 the angle θ is the sum of θ_3, due o the op amp pole at ω_P; θ_1 due to the closed-loop zero it ω_1 and θ_2, due to the closed-loop pole at ω_2. The relative positions of ω_1 and ω_2 determine the overall response of he system.

Equation A5 is the gain error factor for a single stage. For a dual DAC PGA the overall gain error factor can be taken o be the sum of both stage factors.

From the AD7547 data sheet, $C_0 = 140pF$ max and $R_{LAD} = 20k\Omega$ max. The value of compensation capacitor, C_2, required to prevent gain peaking depends upon the minimum value of programmable gain to be used. For a X1 gain setting $C_2 = 0.75\ C_0$ or $C_2 = 100pF$ (from Equation 13). If the overall gain required is X16, then the widest system bandwidth occurs when both stages have gain settings of X $\sqrt{16}$ i.e., X4. The signal bandwidth for a single stage is found from Equation 16 to be:

$$f_{3dB} = \frac{1}{GAIN}\left(\frac{1}{2\pi C_2 R_{LAD}}\right)$$

$$= 20kHz$$

The cascaded bandwidth is:

$$f_{CASC} = 0.64\, f_{3dB}$$

$$= 12.7kHz$$

If the additional gain error due to signal roll-off is to be kept below, say, 0.1%, then from Equation 19, the useable system bandwidth reduces to 1/20 of the cascaded bandwidth, i.e., 600Hz. If AD OP-27 op amps are used which have an open-loop gain of approximately 80dB at 600Hz, then, from Equation A5, the additional gain error due to a single stage is found to be less than 0.004%. Two similar stages will contribute less than 0.01% of additional gain error. This error is an order of magnitude below that caused by signal roll-off. Hence, in comparison with signal roll-off, nonideal gain error factors can be ignored as a source of additional gain error.

A value of 5.5 for G_N was used in Equation A5 to determine the additional gain error. Appendix 3 shows how this value was determined and gives a software listing to help determine G_N for any other gain setting.

APPENDIX 3

From Equation 12

$$G_N = 1 + R_{EQ}\left(\frac{1}{R_{FB}} + \frac{1}{R_O}\right)$$

This can be rewritten as

$$G_N = 1 + \frac{R_{EQ}}{R_{FB}}\left(1 + \frac{R_{FB}}{R_O}\right)$$

$$\text{or } G_N = 1 + (\text{Stage Gain}) \cdot \left(1 + \frac{R_{FB}}{R_O}\right)$$

Table A3 lists a program which solves the $(1 + R_{FB}/R_O)$ term for a straight R-2R ladder network. It is an extention of the expression contained in Reference 3. It is written for a HP85 computer. The resolution of the DAC is entered first followed by the code in decimal for which the output resistance is required. The answer, R_{OUT}, is equal to

$$R_{OUT} = 1 + \frac{R_{FB}}{R_O}$$

This is now used to find the noise gain, G_N.

```
  4 ! ************************
    **
  5 ! * PROGRAM TO PLOT DAC ROUT
     *
  6 ! * FOR ANY NUMBER OF BITS
     *
  7 ! ************************
    **
  8 CLEAR
 10 DISP "NO OF BITS" @ INPUT N
 15 DIM B(20)
 18 DISP @ DISP "ENTER CODE IN D
    EC"
 20 INPUT F
 25 D=F
 27 CLEAR @ DISP @ DISP @ DISP "
          COMPUTING"
 30 GOSUB 1000
 40 GOSUB 4000
 50 DISP @ DISP
 52 CLEAR @ DISP @ DISP @ DISP "
    NUMBER OF BITS = ";N
 54 DISP @ DISP @ DISP "CODE IN
    DEC = ";F
 56 DISP @ DISP
 60 DISP "ROUT = ";E
 70 END
1000 REM DEC TO BIN
1010 FOR C=N-1 TO 0 STEP -1
1020 D=D-2^C
1030 IF D<0 THEN D=D+2^C @ B(N-C
     )=0 ELSE B(N-C)=1
1040 NEXT C
1050 RETURN
4000 REM CALCULATE ROUT
4005 T=0 @ S=0
4010 FOR C=1 TO N STEP 1
4020 IF B(C)=1 THEN 4040
4030 GOTO 4070
4040 Y=1+2^(1-2*C)
4050 T=T+Y
4070 NEXT C
4080 FOR C=2 TO N-1
4090 FOR J=C+1 TO N
4100 IF B(C)=1 AND B(J)=1 THEN 4
     120
4110 GOTO 4140
4120 X=2^(2-C-J)*(2^(2*C-2)-1)
4130 S=S+X
4140 NEXT J
4150 NEXT C
4160 E=T-S @ E=3/E
4165 E=1/E
4170 E=E+1
4175 E=E-F/2^N
4190 RETURN
```

Table A3. HP-85 Listing for Solving $(1+R_{FB}/R_O)$ Ratio

ANALOG DEVICES

AN-321
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

3-Phase Sine Wave Generation Using the AD7226 Quad DAC

by Mike Byrne

This application note discusses the use of the AD7226 in waveform generation. It discusses how to generate 3-phase sine waves which can be used in driving small 3-phase motors. The AD7226 is a quad, 8-bit, CMOS D/A converter packaged in a 20-pin DIP (consult the datasheet for information on the part).

The circuit of Figure 1 shows one method of generating 3-phase sine waves using the AD7226. The circuit consists of a 2716 EPROM, in which sine wave values are stored in digital form; a counter, which counts through the 2716 addresses, and control logic which controls the counter and the loading of the AD7226. The proper codes for synthesizing three sine waves are stored in sets of three. Each set of three code values results in output sine wave values which have a 120° separation between them. For example, if Address X contains a code corresponding to 0° then Address X+1 contains a code corresponding to 120°and X+2 a code corresponding to 240°. The next address, X+3, is the start of another set of three codes and it contains a value corresponding to (0°+1.4°) (i.e., 360°/256, assuming a 256-element sine wave table is used). Address X+4 contains a value corresponding to (120°+1.4°) and so on. Therefore, for a 256-element sine wave, 768 (3×256) addresses in the EPROM are used.

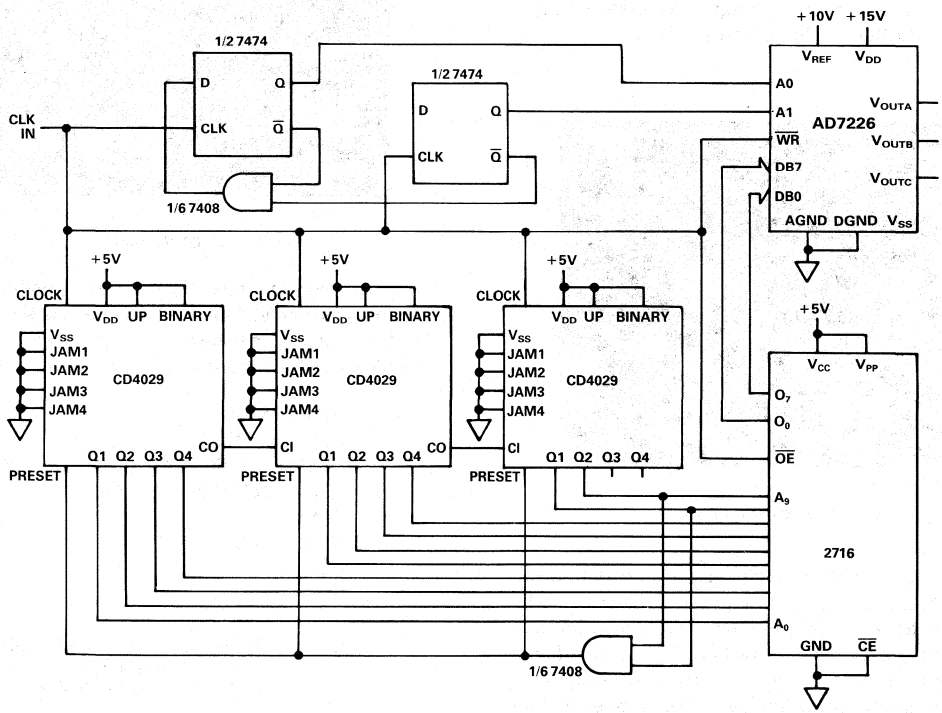

Figure 1. 3-Phase Sine Wave Circuit

The logic in the circuit controls the output of data from the 2716 and the loading of the D/A converters in the AD7226. It ensures that successive values from the EPROM are loaded to the D/A converters in turn. Referring to the example already given, data from Address X is loaded to DAC A, Address X+1 to DAC B, Address X+2 to DAC C, Address X+3 to DAC A again and so on. A full cycle for the three sine waves is generated by stepping through the full look-up table in the EPROM. When data from the last address has been loaded to the AD7226, the counter is reset and the cycle begins again.

With a 256-element sine table the output frequency is 1/768th of the input clock frequency. Varying the input clock frequency over the range 4kHz to 1MHz will vary the output frequency from a few Hz to 1.2kHz. The matching and tracking between the D/A converters of the AD7226 ensures a very good match between the output amplitude of each of the sine waves. Figure 2 shows some typical resulting waveforms, after filtering with a simple RC, for a 1.2kHz output frequency. This was found to be the maximum frequency at which the circuit could be run without appreciable change in the distortion figures for the output waveforms. The limiting factor on the frequency range is the 450ns data access time specification of the 2716 EPROM. To increase the frequency range, the 2716 should be replaced by an EPROM with a faster access time. Alternatively, if a sine wave look-up table with less elements is used then the frequency of the output sine waves will be increased since the time to cycle through a full look-up table is decreased. For example, if an 85-element sine table is used, with the same clock frequency, then the maximum output frequency (again using the 2716 EPROM) is increased to 3.6kHz.

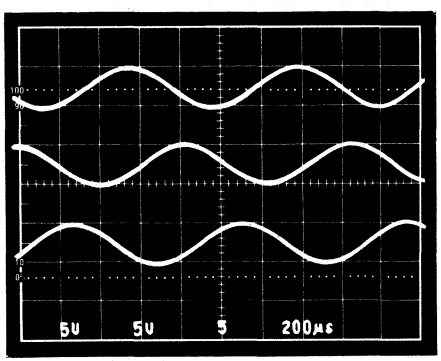

Figure 2. 3-Phase Sine Wave Output

Figure 3 shows typical harmonic distortion results using the circuit of Figure 1 after the output has been filtered by a single pole RC with a 3dB frequency slightly larger than the output frequency. The RC will not have a significant effect on the portion of the spectrum shown but will however remove harmonics of the sampling frequency. The results show that the largest harmonic of the sine wave is typically −55dB below the fundamental. The total harmonic distortion figure is typically −48dB. These distortion figures were achieved up to a frequency of 1.2kHz, using a 256-element sine wave table. When an 85-element table is used and the output frequency increased to 3.6kHz, the total harmonic distortion figures are not appreciably worse after the output has again been filtered with a simple RC circuit.

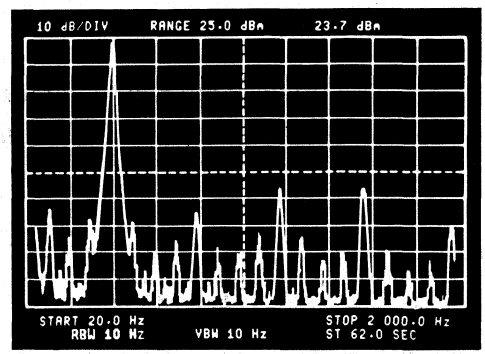

Figure 3a. Output Spectrum (400Hz)

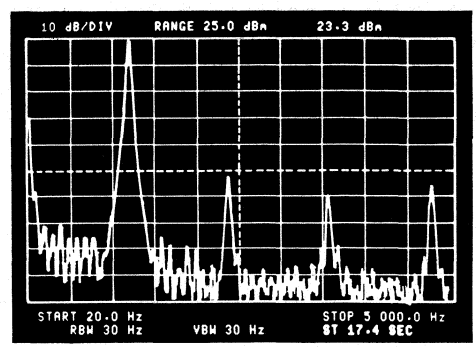

Figure 3b. Output Spectrum (1.2kHz)

A basic program for generating the code values for a 256-element sine wave table on the AIM-65 is shown in Table I. The program has to be run through a total of three times to fill all 768 locations of the sine wave look-up table.

```
INPUT "P";P
FOR X=0 TO 255
LET A=SIN [(X * 1.40625+P)/57.2958]
LET Y=INT (128.5+127 * A)
POKE L,Y
LET L=L+3
NEXT X
END
```

NOTE: [L=3328 and P=0 for 1st Sine Wave
L=3329 and P=120 for 2nd Sine Wave
L=3330 and P=240 for 3rd Sine Wave]

Table I. Basic Program

In the program the numerical values which the SIN function acts upon are in radians. For lesser element sine tables line 2 in the program should be modified. For example, to get a 3 × 85-element table, STEP 3 should be added to the end of the line. The decoding to reset the counter would also have to be changed since the counter would now only count through 255 states.

In the circuit of Figure 1 the fourth D/A converter of the AD7226, DAC D, is not addressed. The circuit can easily be adapted to address this fourth DAC. In this case DAC D can be used in a feedback configuration to provide a programmable reference voltage for itself and the other three converters. This configuration is shown in the circuit of Figure 4. It means that the amplitude of the output sine waves can be programmed by varying the digital code in latch D. Note that varying the digital code in latch D varies the output amplitude of all three sine waves. The relationship of V_{REF} to V_{IN} is dependent upon digital code and upon the ratio of resistors R1 and R2. It can be expressed by the formula:

$$V_{REF} = \frac{(1+G)}{1+G \cdot D_D} \cdot V_{IN}$$

where G = R2/R1

and D_D is a fractional representation of the digital word in latch D ($0 \leq D_D \leq 255/256$).

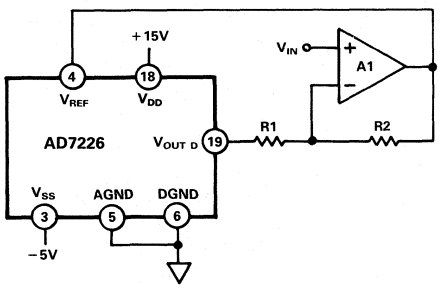

Figure 4. Self-Programmable Reference

Figure 5 shows typical plots of V_{REF} versus digital code for three different values of R2. With V_{IN} = 2.5V and R2 = 3R1 the voltage at the output of A1 (i.e., the reference voltage for the AD7226) will vary between +2.5V and +10V over the digital input code range. The circuit of Figure 4 should only be used when the AD7226 is operated in dual supplies. This is because the AD7226 has reduced current sink capability at output voltages near 0V when used in single supply (see AD7226 data sheet). As a result the circuit would not operate correctly at lower values of digital input code. For correct operation with dual supplies R1 must be greater than 6.8kΩ.

The circuit of Figure 1 need not necessarily be used to generate 3-phase sine waves. Since the codes in the EPROM determine the shape of the output waveforms then the D/A converters can be made to produce various different outputs. For example, the circuit can be used to generate a square wave output from DAC A, a triangular wave output from DAC B and a sine wave output from DAC C, with all waveforms at the same frequency. The method for storing the codes in EPROM is as outlined before with the codes stored in sets of three. The resulting waveforms are shown in Figure 6. Once again, only three of the AD7226 outputs are used. This is because the circuit of Figure 1 is designed for selecting only three of the D/A converters. It can easily be adapted so that the fourth D/A converter can be used to provide a fourth output waveform. Alternatively, the fourth D/A converter can be used to program the amplitude of the output waveforms as before.

8

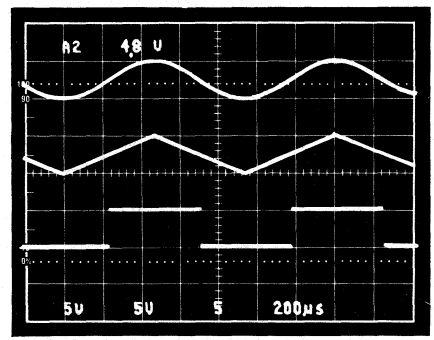

Figure 6. Output Waveforms

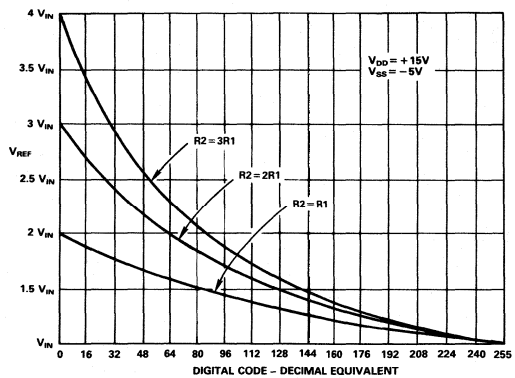

Figure 5. Variation of V_{REF} with Feedback Configuration

ANALOG DEVICES

AN-209
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

8th Order Programmable Low Pass Analog Filter Using Dual 12-Bit DACs

by Bill Slattery

INTRODUCTION

This application note describes the design of a low pass analog filter whose cutoff frequency can be programmed from 100Hz to 50kHz. The filter is designed as a plug in expansion board for IBM PC AT/XT* or compatibles. A high order filter function is implemented, giving a very fast roll-off in the transition band. This design realizes an 8th order function with a roll-off equalling 48dB/octave. The note also discusses some of the tradeoffs and practical limitations which must be considered when designing a filter.

The design is based on a 2nd order universal active filter, as shown in Figure 1. The required performance is achieved by cascading four of these 2nd order stages.

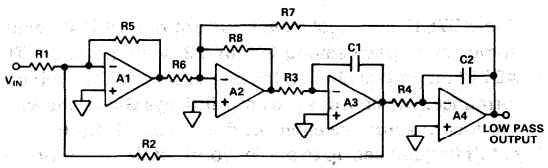

Figure 1. Universal Active Filter

The cutoff frequency of the filter is determined by R3, R4, C1 and C2. Digital control of the cutoff frequency is achieved by replacing resistors R3 and R4 in each stage by CMOS Multiplying Digital to Analog Converters (DACs). The DAC is in effect configured as a digitally programmable resistance.

To have accurate control of cutoff frequency, R3 and R4 within each stage must be closely matched. This is best achieved by replacing these two resistors with a monolithic dual 12-bit DAC. Analog Devices produces a range of suitable dual 12-bit DACs, the AD7537, AD7547 and AD7549. Since these have two DACs on one chip, DAC resistance matching will be in the order of 0.5%. Additionally, the use of 12-bit DACs ensures excellent resolution over the wide range of cutoff frequencies which can be programmed to the filter.

Applications for this filter include Industrial Process Control, Automatic Test Equipment (ATE), Sonar Signal Processing, Instrumentation, Audio Systems and Data Acquisition Systems. In Digital Signal Processing (DSP) applications, it can be used as the front end, low pass, anti-alias filter.

THE FILTER FUNCTION

A 2nd order low pass filter function is given by

$$\frac{V_{OUT}(s)}{V_{IN}(s)} = \frac{A_O \omega_O^2}{s^2 + \frac{\omega_O}{Q} s + \omega_O^2} \quad (1)$$

where $s = j\omega$

ω_O = 3dB bandwidth (cutoff frequency)

Q = circuit Q factor

A_O = gain at $\omega = \omega_O$

The universal active filter shown in Figure 1 has a 2nd order low pass transfer function given by

$$\frac{V_{OUT}(s)}{V_{IN}(s)} = \frac{\frac{R5\,R8}{R1\,R6}\left(\frac{1}{C1\;R3}\right)^2}{s^2 + \frac{R5\,R8}{R2\,R6}\left(\frac{1}{C1\;R3}\right)s + \left(\frac{1}{C1\;R3}\right)^2} \quad (2)$$

when $R3 = R4$

$R7 = R8$

and $C1 = C2$

By comparing coefficients between Equations (1) and (2) we see that

$$A_O = \frac{R5\,R8}{R1\,R6} \quad (3)$$

$$Q = \frac{R2\,R6}{R5\,R8} \quad (4)$$

$$\omega_O = \frac{1}{C1\,R3} \quad (5)$$

hence $$f_O = \frac{1}{2\pi\,C1\,R3}$$ (filter cutoff frequency) (6)

IBM PC AT/XT is a trademark of International Business Machines Corp.

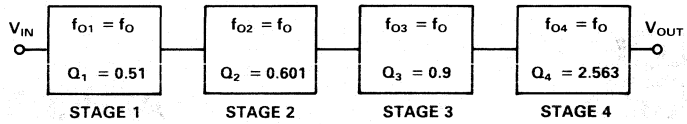

Figure 2. Block Diagram of 8th Order Butterworth Filter

8th Order Butterworth Filter Function

In realizing a particular 8th order filter function, each 2nd order stage is individually programmed to a specific cutoff frequency, f_O and Q. Values of f_O and Q for different filter types can be found using tables or software routines which are widely available (References 1 and 2).

The design discussed in this application note implements a unity gain, Butterworth filter function. The cutoff frequency which must be programmed to each stage is the same as the overall filter cutoff frequency f_O. The Qs of each stage, however are not the same. Figure 2 shows the required values of f_O and Q for each stage.

Since the gain of each stage is unity (A_O = 1), Equations (3) and (4) can now be solved using the values of Q given in Figure 2, to yield the required resistor values. Table I lists the required resistor values of each stage.

	Stage 1	Stage 2	Stage 3	Stage 4
R1	39k	150k	120k	12k
R2	20k	82k	82k	33k
R5	12k	82k	82k	3.9k
R6	10k	1.8k	2.7k	10k
R7	33k	3k	3k	33k
R8	33k	3k	3k	33k

1% tolerance resistors should be used.

Table I. Resistor Values for Each Stage of This 8th Order Butterworth Filter

DIGITAL CONTROL OF CUTOFF FREQUENCY

The cutoff frequency of the filter is determined by C1, C2, R3 and R4. However, since C1 = C2 and R3 = R4, the cutoff frequency f_O can be expressed in the form shown in Equation (6).

Digital control over cutoff frequency is achieved by replacing R3 and R4 with one of the following configurations:

1. The AD7537 DAC.
2. The AD7537 DAC in series with a padding resistance R_{PAD}.

The principal difference between the AD7537, AD7547 and AD7549 is in their loading structure. The AD7537 is chosen for this design, see Figure 3. Its 2 byte (8+4) loading structure makes it ideal to use with a microprocessor based system which has an 8-bit data bus. The IBM PC AT/XT or compatible is just such a system.

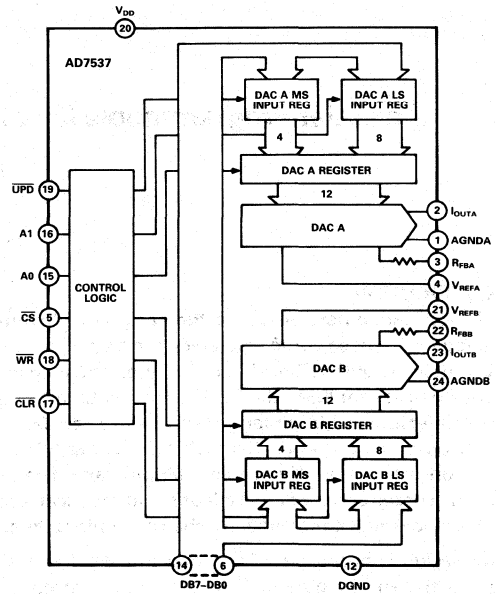

Figure 3. Functional Block Diagram of AD7537

The AD7537 is a dual 12-bit current output DAC whic comes in a space saving 24-pin 0.3" wide package. Th AD712 op amp is chosen for this design. It is a high pe formance, low-cost dual op amp. Its large Gain Bandwidt Product ensures excellent performance of the filter fc cutoff frequencies up to and beyond 50kHz. (For more d tailed information on the AD7537 and AD712 consult th relevant data sheets.)

Design (1)

Figure 4 shows how R3 and R4 are replaced by the AD753 dual DAC. The AD7537 is now in effect a pair of prc grammable resistors.

The equivalent resistance of a DAC in this configuratio is given by

$$R_{EQ} = \frac{R_{DAC}}{D} \quad (7$$

where R_{DAC} is the DAC ladder resistance

and D is the digital fraction programmed to the DAC D is given by

$$D = \frac{N}{2^n} \quad (8$$

where n = resolution of DAC in bits (in this case n = 12)

N = DAC digital code (decimal 1 to 4095)

Since R3 is now replaced by R_{EQ} we can write

$$f_O = \frac{1}{2\pi C1 R_{EQ}} \quad (9)$$

Choosing C1 = 220pF and R_{DAC} = 14k

and solving between Equations (7), (8) and (9), the cutoff frequency can be written as a function of the decimal code N.

$$f_O = (0.01262 \cdot N)\ kHz \quad (10)$$

In practice the cutoff frequency can now be programmed between 100Hz and 50kHz with a resolution of 13Hz.

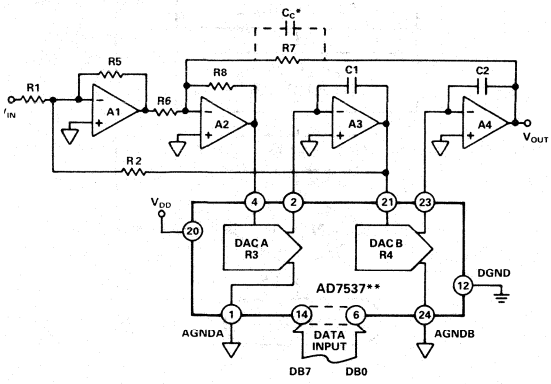

*See "GAIN BANDWIDTH PRODUCT SECTION".
**ADDITIONAL CIRCUITRY OMITTED FOR CLARITY.

Figure 4. Single Stage of Programmable Filter Using an AD7537

For a known DAC ladder resistance, a stage can be accurately programmed to a specific cutoff frequency. The two frequency controlling resistors are effectively matched to within 0.5%, since both are in one monolithic package. However, to achieve an overall accurate filter cutoff frequency, the DACs in each of the four stages must have similar DAC ladder resistances. DAC ladder resistance can vary between 9k and 20k with a typical value of 14k, for different AD7537s.

Design (2)

One solution that can be adopted to avoid using specially selected AD7537s is to use a padding resistance R_{PAD} in series with the DAC ladder resistance as shown in Figure 5. A large value of R_{PAD}, 100k, is used. This has the effect of desensitizing the variations in DAC ladder resistance. This means that the specified variation in DAC ladder resistance now has an insignificant effect on the overall filter performance.

The equivalent resistance of a DAC in this configuration is given by

$$R_{EQ} = \frac{R_{DAC} + R_{PAD}}{D} \quad (11)$$

Choosing C1 = 22pF, R_{PAD} = 100k and R_{DAC} = 14k

and solving between Equations (8), (9) and (11), the cutoff frequency can be written as a function of the decimal code N.

$$f_O = (0.01549 \cdot N)\ kHz \quad (12)$$

The cutoff frequency can now be programmed between 100Hz and 50kHz with a resolution of 16Hz.

It should be noted that since we are using DACs whose ladder resistance can vary between 9k and 20k, the above expression for f_O will be accurate to within ±4.5%. If the padding resistance R_{PAD} was not used, f_O could vary by up to ±30% of the programmed value. Design 1 does not have this problem since DACs whose ladder resistance is known to be 14k are used. Cost and accuracy tradeoffs must be considered when choosing either of the two design options.

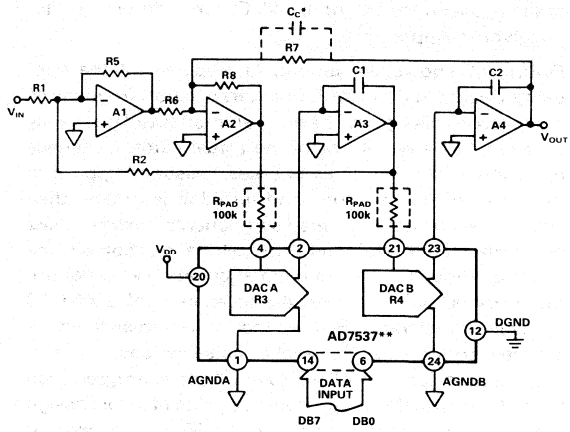

*See "GAIN BANDWIDTH PRODUCT SECTION".
**ADDITIONAL CIRCUITRY OMITTED FOR CLARITY.

Figure 5. Single Stage of Programmable Filter Using an AD7537 and Series Resistances R_{PAD}

GAIN BANDWIDTH PRODUCT

A compensation capacitor C_C must be used on the 4th Stage of the filter. This is required to reduce effects caused by Gain Bandwidth Product limitations. The value of capacitance depends on the type of function implemented and the design option choosen.

For this 8th Order Butterworth Filter, the value of the capacitor can vary between 15pF and 47pF

Design (1) – Use 15pF

Design (2) – Use 47pF.

DIGITAL INTERFACE

The filter circuit can be built on a standard IBM prototype card and plugged directly into the expansion slot on the main board of an IBM PC AT/XT or compatible. The prototype board basically consists of buffers and some circuitry which generates an enable line $\overline{EN}$. Figure 7 shows the interface layout used. Each AD7537 is addressed ($\overline{CS}$) using a 74LS138 (3 to 8 line decoder). Address lines A3 to A5 as well as the $\overline{EN}$ and $\overline{IOW}$ lines are used to decode each address. The valid addresses which can be used are HEX300 to HEX31F.

Since the AD7537 is a dual 12-bit DAC, data is loaded to each DAC with a 2-byte (8+4) loading instruction. Address lines A0 to A1 select DAC A and DAC B, see Figure 3, as well as selecting the low and high bytes of data. The address decoder also generates the $\overline{UPD}$ line which is used to update the data in all the DACs.

An IBM Basic program for this Butterworth filter is given in Table A1, Appendix 1. The program is run by inputting the required cutoff frequency in kHz as well as the values of R_{DAC}, R_{PAD}, and C1. The program calculates and outputs the relevant digital code to each of the DACs. Figure 6 gives a flow diagram representation of the program used.

PERFORMANCE

Plots of amplitude and phase response as well as a plot of the noise spectrum for the 8th Order Butterworth Filter are given in Appendix 2.

Figure A1 shows the amplitude response of the filter using Design 1. Plots for four different cutoff frequencies are given in this figure. The programmed cutoff frequency is found to be within 3% of the actual cutoff frequency achieved, for each of the four cases. Passband ripple is in the order of 0.2dB. Stop-band rejection is greater than 100dB, except for frequency components between 2kHz and 4kHz, and at 41kHz and 1MHz, where rejection is in the order of 90dB. The reduction in attenuation at lower frequencies is mainly due to the increased value of DAC equivalent resistance R_{EQ}. As the value of resistance increases, so too does the level of noise increase. In practice, frequencies below 500Hz should not be programmed to the filter with the component values used in this design option. Cutoff frequencies of less than 500Hz can however be achieved if larger values of capacitance C1 are used. This effectively reduces the range of selectable cutoff frequencies but allows lower cutoff frequencies with a reduced noise floor to be programmed to the filter. At high frequencies, greater than 10MHz, it is found that there is less attenuation; attenuation reduces below 90dB. This is due both to gain bandwidth limitations of the op amp and feedthrough across the system.

Figure A2 shows the amplitude response for various cutoff frequencies, using Design 2. Again we can see that the pass band ripple is less then 0.2dB. The accuracy of the cutoff frequency however will vary by about 4.5% from the programmed cutoff frequency. Also, it should be noted that in using Design 2 stopband rejection will be reduced to about 70dB. This is due to increased noise caused by the large value of padding resistance R_{PAD}. This can be reduced somewhat by using low noise resistors.

Figure A3 is a plot of the noise spectrum at the output of the filter, using Design 1. The cutoff frequency is programmed at 50kHz, and the input is grounded.The largest noise component in our system occurs at about 41kHz and has an amplitude of −54.8dBm* (0.4mV). Since the filter is present in a noisy environment, i.e., the board of an IBM PC AT/XT or compatible, it is not surprising that there should be noise present in the system. In practice, there is a minimum input signal level that can be applied to the

*Measured with respect to 50Ω.

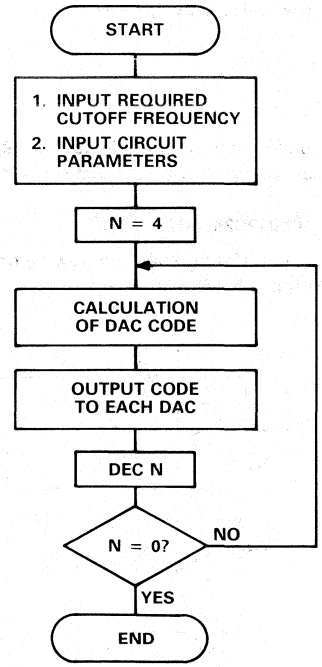

Figure 6. Flow Diagram Representation of Control Algorithm

filter to achieve a specific Signal-to-Noise Ratio (SNR) ir the filter's pass band. In this design, using the noise component at 41kHz, an SNR greater than 90dB can be achieved by applying input signal levels of greater thar 1.3V rms.

Figure A4 is a plot of the filter's phase response. The phase response of the system determines the Group Delay of the filter.

$$\text{Group Delay} = \frac{1}{2\pi}\frac{d\phi}{df}$$

Thus a linear phase response implies a constant group delay. The phase response plot shown exhibits some slight nonlinearity. This is as expected for a Butterworth filter function. If phase response is an important design consideration, a different filter function should be considered, eg., Bessel. (References 1 and 2.)

PERFORMANCE EXTENSION

1. The frequency range of this filter can be extended to 100kHz, using Design 1, by simply replacing the 220pF capacitors in each stage by 110pF capacitors. Stop-band attenuation of 85dB and passband ripple of less than 0.2dB is achieved over the full range.
2. A Chebychev response can be achieved by using different values of R1 and R2 in each stage. R1 and R2 can be evaluated by solving Equations (3) and (4) for the values of Q given in Figure 8. Figure 8 gives the cutoff frequency f_O and Q for each stage of a 0.2dB ripple 8th Order Chebychev Filter. It should be noted that for a Chebychev filter the cutoff frequency of each stage is

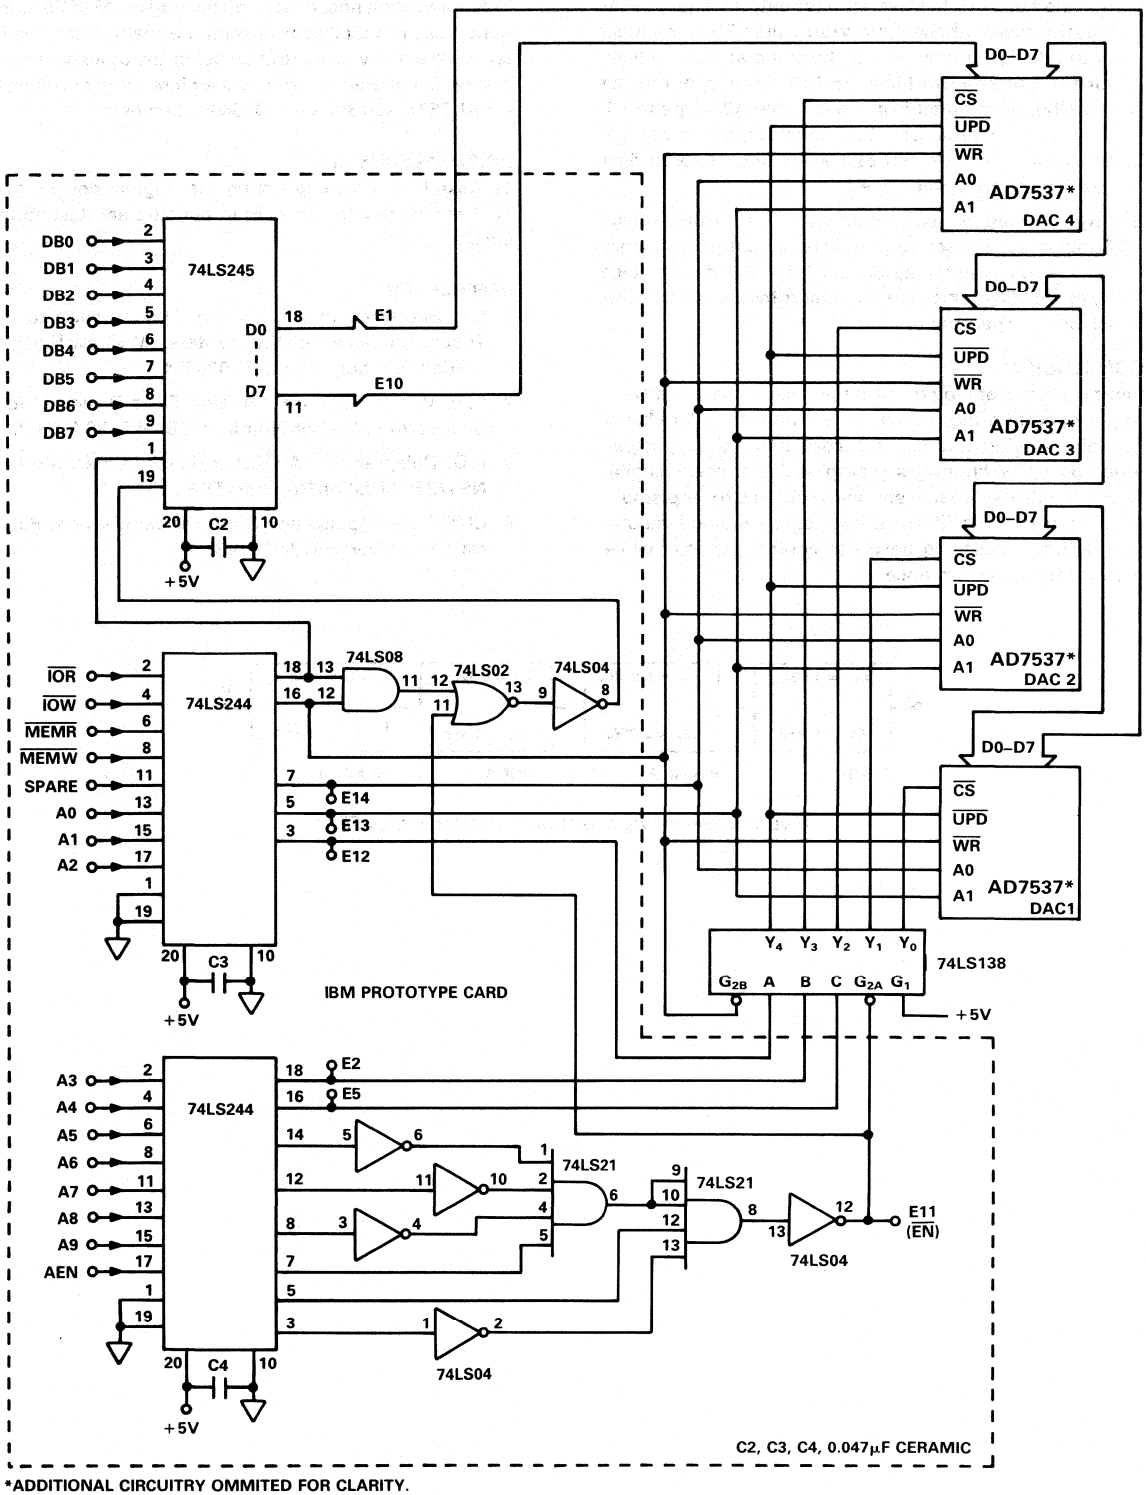

*ADDITIONAL CIRCUITRY OMMITED FOR CLARITY.

Figure 7. IBM PC AT/XT Digital Interface

not the same as the overall filter cutoff frequency, as was the case with the Butterworth filter. The remainder of the design proceedure is the same as that outlined for the Butterworth filter. An IBM Basic program for this Chebychev filter is given in Table A2, Appendix 1. This program tunes each individual stage to its particular cutoff frequency so that a specified overall filter cutoff frequency is achieved.

3. Programmability over filter type (i.e., Butterworth, Chebychev, Bessel, etc.) as well as cutoff frequency can be achieved by replacing R1 and R2 of each stage by an AD7537. This gives us total software control over all the filter parameters.

CONCLUSION

Filtering in the analog domain has many inherent advantages over filtering in the increasingly popular digital domain. Analog filtering is a continuous time process whereas digital filtering is a sampled data process. Digital filters require extensive software routines to implement such algorithms as FIR and IIR filters. The analog filter utilizes digital processing power to control the filter while giving analog performance.

This application note deals with the IBM PC AT/XT or compatible as the digital controller for the analog filter. The filter could equally be controlled using any other microprocessor. Suggested microprocessor interfaces are given in the AD7537, AD7547 and AD7549 data sheets.

ACKNOWLEDGEMENTS

To Mike Curtin for suggesting the original design idea
To Sean Morley for his help in building and testing the circuit.

REFERENCES

1. L.P. Huelsman and P.E. Allen, "Introduction to the Theory and Design of Active Filters." McGraw-Hill Publication Number: ISBN 0-07-030854-3.
2. M.E. Van Valkenburg, "Analog Filter Design." Holt-Saunders Publication Number: ISBN 4-8338-0091-3.
3. A.D. Delagrange, "An Active Filter Primer, Mod 1" NSWC Publication Number: TR 82-552.
4. CMOS DAC Application Guide, Analog Devices Publication Number: G872a-15-4/86.

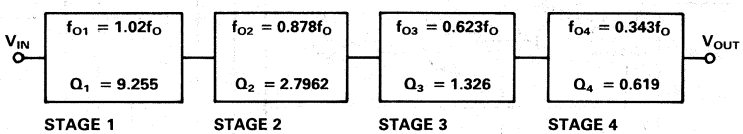

Figure 8. Block Diagram of 8th Order Chebychev Filter

```
10  CLS
20  PRINT "8th-Order Low-Pass Butterworth Filter"
30  PRINT
40  INPUT "F0 (kHz) (500Hz . . . . 48kHz)                    = " ;FO
50  F0 = F0*1000
60  PRINT
70  INPUT "RDAC(Kohms) (14K TYPICALLY)                       = " ;RDAC
80  RDAC = RDAC*1000
90  PRINT
100 INPUT "RPAD(Kohms) (0 FOR DES.1, 100K FOR DES.2)         = " ;RPAD
110 RPAD = RPAD*1000
120 PRINT
130 INPUT "C(pFarads) (220pF FOR DES.1, 22pF FOR DES.2)       = " ;C
140 C = C*1E-12
150 REQ = 1/2/22*7/FO/C
160 N = 4096*(RPAD + RDAC)/REQ
170 PRINT                                                     = " ;HEX$(N);"H"
180 PRINT "CODE
190 N = INT (N)
200 N1 = INT(N/256)'   MSB
210 N2 = N-N1*256 '   LSB
220 ADDR = &H300
230 FOR C = 0 TO 3
240 OUT ADDR,N2
250 OUT ADDR + 1,N1
260 OUT ADDR + 2,N2
270 OUT ADDR + 3,N1
280 ADDR = ADDR + 4
290 NEXT C
300 OUT &H310,0 'UPDATE DACS
310 END
```

Table A1. IBM Basic Program for 8th Order Butterworth Filter

```
10  CLS
20  PRINT "8th-Order Low-Pass Chebychev Filter"
30  PRINT
40  INPUT "F0 (KHz)                                          = " ;FO
50  F0 = F0*1000
60  PRINT
70  INPUT "RDAC (Kohms)  (14K TYPICALLY)                     = " ;RDAC
80  RDAC = RDAC*1000
90  PRINT
100 INPUT "RPAD(Kohms)  (0 FOR DESIGN 1)                     = " ;RPAD
110 RPAD = RPAD*1000
120 PRINT
130 INPUT "C(pFarads)  (220pF FOR DES. 1)                    = " ;C
140 C = C*1E-12
150 FOR I = 0 TO 3
160 READ X
170 FC = X*FO
180 REQ = 1/2/22*7/FC/C
190 N = 4096*(RPAD + RDAC)/REQ
200 PRINT
210 PRINT "CODE" ;I;"                                        = ";HEX$(N);"H"
220 N = INT(N)
230 N1 = INT(N/256) '   MSB
240 N2 = N-N1*256 '   LSB
250 ADDR = &H300
260 ADDR = ADDR + I*4
270 OUT ADDR,N2
280 OUT ADDR + 1,N1
290 OUT ADDR + 2,N2
300 OUT ADDR + 3,N1
310 NEXT I
320 OUT &H310,0'UPDATE DACS
330 END
340 DATA 1.02,.878,.623,.343
```

Table A2. IBM Basic Program for 8th Order Chebychev Filter

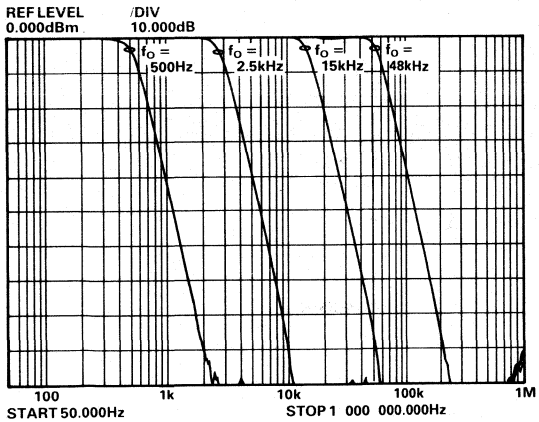

Figure A1. Amplitude Response of Butterworth Filter for Various Cutoff Frequencies Using Design 1

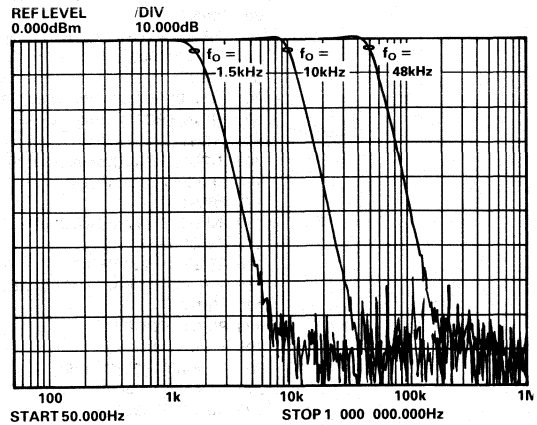

Figure A2. Amplitude Response of Butterworth Filter fo Various Cutoff Frequencies Using Design 2

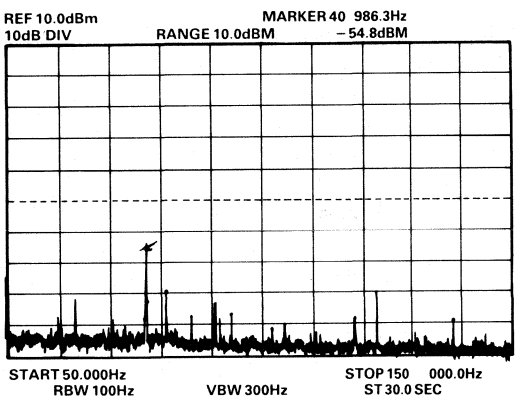

Figure A3. Noise Spectrum of Filter

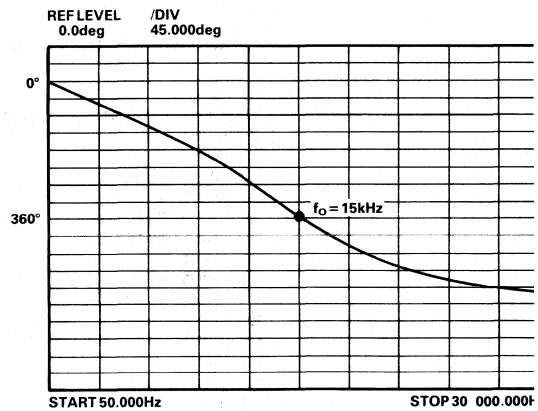

Figure A4. Typical Phase Response of Butterworth Filter

ANALOG DEVICES

AN-322
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Improve Function Generators with Matched D/A Converters

by Paul Toomey

With a pair of precisely matched 8-bit CMOS DACs in one IC, you can create inexpensive circuits that digitally control high-precision analog waveforms.

Using the AD7528, a low-cost IC containing two 8-bit CMOS multiplying DACs, you can design a variety of signal-generation circuits for audio, computer-graphics and process-control applications. These designs, which require a pair of closely matched, tightly tracking D/A converters, have heretofore proved impractical to produce because the selection and testing necessary to match the DACs has made them quite expensive.

Three useful circuits—a programmable sine-wave oscillator, a function-fitting sine-wave synthesizer and a triangle/rectangle-wave generator—clearly demonstrate the 7528's design potential. Each circuit is digitally controlled; therefore, before considering the specific designs, you should examine interfaces for widely used microprocessors.

To that end, note that the AD7528's two on-chip 8-bit latches and control inputs ($\overline{CS}$, $\overline{WR}$ and DAC Select) are logic-level and speed compatible with most μPs. Data-hold time, often a problem in applications using members of the 6800 family, equals zero. **Fig 1a** shows an interface to a 6800-based system, while **Fig 1b** details the connections for an 8085-based system.

Use dual DACs as ganged potentiometers

With this background information, turn to the dual DAC's applications. You can control the frequency of many oscillator circuits by using two ganged potentiometers, provided that the pots track precisely over their full resistance range. **Fig 2**'s high-performance sine-wave oscillator is a state-variable filter design in which an AD7528 acts as a pair of digitally controlled matched resistors to replace the ganged potentiometer. The equivalent resistance of each DAC, as seen by op amps A_2 and A_3, varies from infinity at input code 00_H to a minimum of 10 kΩ (the DAC ladder resistance) at FF_H.

Because the same code is loaded into each DAC, the expression for the output frequency (in hertz) becomes $f=N/256(2\pi RC)$, where N is the DAC code (base 10), R

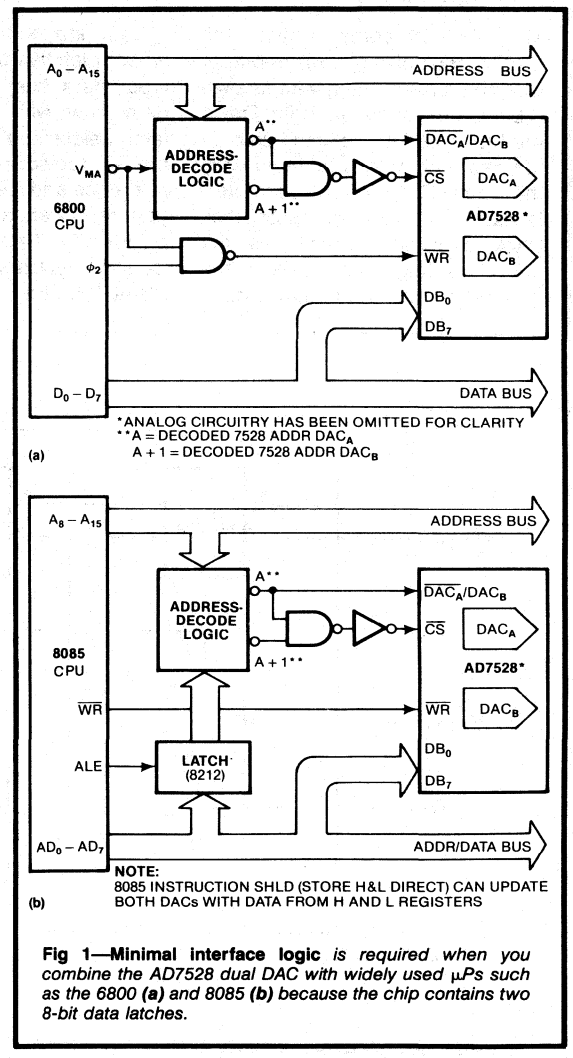

Fig 1—Minimal interface logic *is required when you combine the AD7528 dual DAC with widely used μPs such as the 6800* ***(a)*** *and 8085* ***(b)*** *because the chip contains two 8-bit data latches.*

Reprinted from EDN — May 12, 1982

Summing resistors are critical to function fitter's performance

is the DAC ladder resistance and C is the feedback capacitance for A_2 and A_3. For the circuit values shown in **Fig 2**, output frequency varies from 0 to 15 kHz in steps of approximately 60 Hz, with an amplitude of about 20V p-p. Total harmonic distortion measures −53 dB at 1 kHz and −43 dB at 14 kHz.

Combine multiplication and function fitting

Function fitting translates a mathematical or empirical relationship from one medium (such as a mathematical formula) to another (usually a physically realizable device or system). By taking advantage of the multiplying capabilities of the AD7528's two CMOS DACs, you can use this technique to synthesize extremely low-frequency, highly stable sine waves. For example, in **Fig 3a**'s design, the IC implements a 1-quadrant sin X approximation in the form of the quadratic polynomial $Y=1.828N-0.828N^2$, where $0\leq N\leq 1$ and $N=(2/\pi)X$. (See **Appendix** for the derivation of this equation.) By ramping N up and down and strategically switching the

Identical twins

The AD7528 contains two 8-bit CMOS multiplying current-output D/A converters, each with its own 8-bit latch. Because both DACs are fabricated at the same time on the same chip, they have virtually the same R/2R-ladder characteristics, and these characteristics track extremely closely as the supply voltage varies and the ambient temperature changes.

Data inputs to the AD7528's on-chip latches share a common input data bus; you control the data flow with the $\overline{DAC_A}/DAC_B$ latch-select input. Each section of the IC provides its own analog reference and feedback inputs, so you can operate the DACs independently. Channel-to-channel isolation typically measures −77 dB at 100 kHz.

Settling time varies with supply voltage. After a digital input change, the analog output settles to within 90% of its final value in 80 nsec with a 15V supply; with a 5V supply, settling occurs in 270 nsec. Regardless of the supply voltage, though, the AD7528 typically draws just 1 mA.

For more information on the AD7528, **Circle No 750.**

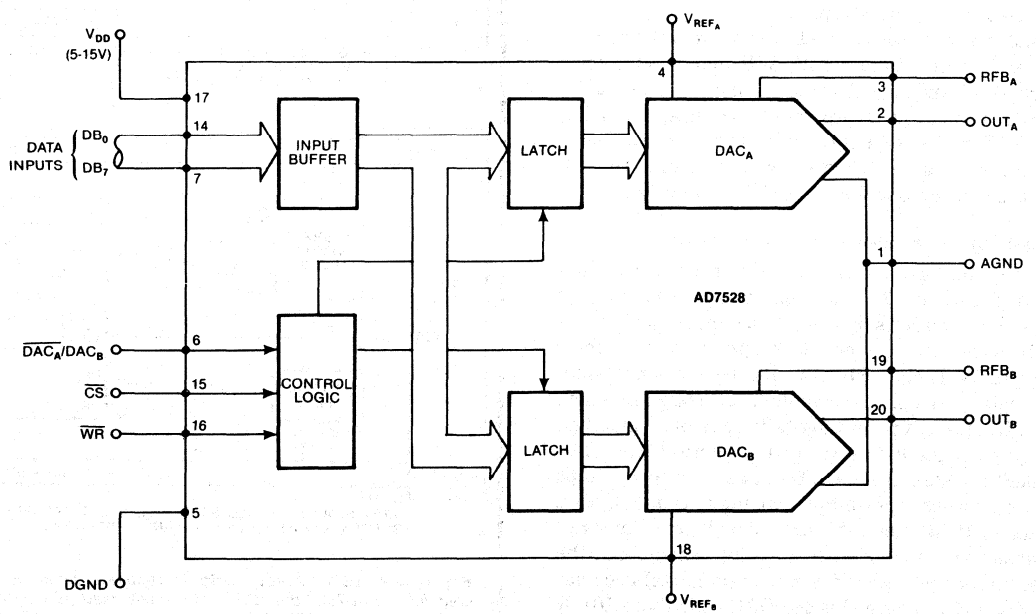

Simultaneous fabrication *of the AD7528's two CMOS multiplying DACs yields virtually identical R/2R-ladder characteristics.*

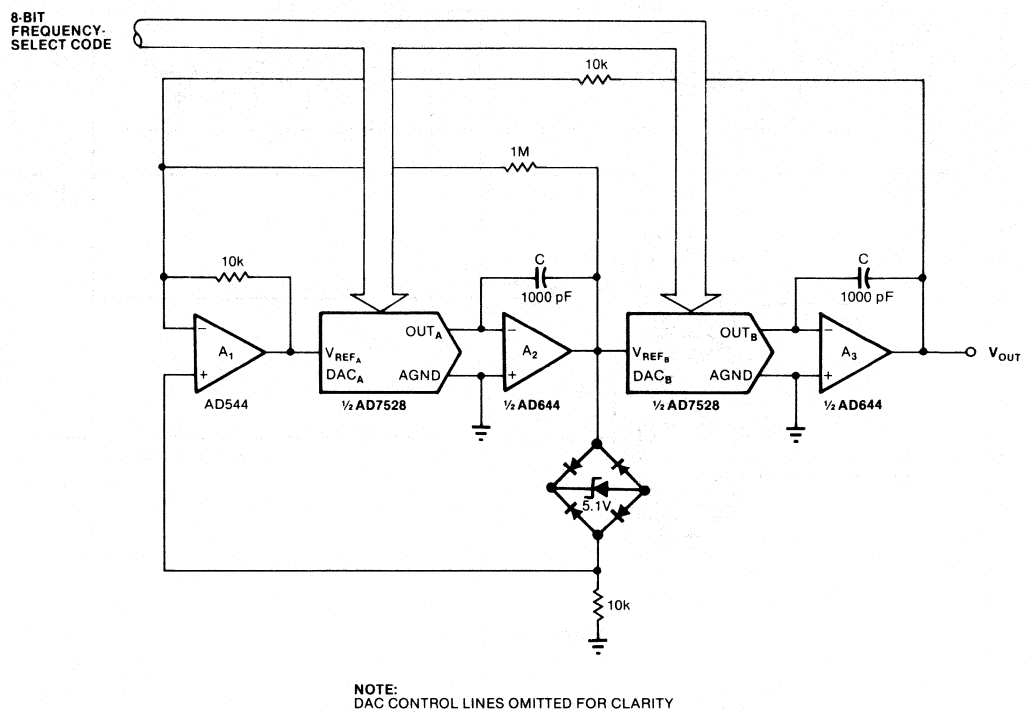

Fig 2—DAC_A and DAC_B act as matched, digitally controlled resistors, *replacing a ganged precision potentiometer, in this state-variable-filter-based programmable sine-wave oscillator.*

output polarity, the circuit generates sin X in four stages (**Fig 3b**).

Specifically, an input square-wave clock drives a counter that counts up and down continuously, providing a Borrow output pulse every time it reaches the all-ZEROs count. In turn, DAC_A produces a triangle wave consisting of two ramps of opposite slope, each generated in 256 steps of op amp A_1's output. This is the analog N variable.

To produce the analog N^2 variable, DAC_B multiplies N's digital version by its analog version. Op amp A_2 inverts the product signal, producing $-N^2$. Then amplifier A_3 sums the two variables in the correct 1.828 and 0.828 weighting, as determined by R_A and R_B, to produce the Y signal. A smooth sine wave results when op amp A_4 switches the circuit's output polarity to generate properly timed positive and negative signals.

Distortion in the output sine wave is a function of the quadratic-approximation fit to the sine curve. Therefore, the values of R_A and R_B are critical to proper circuit operation. Measured distortion for **Fig 3**'s circuit values is −35 dB; it remains constant over the 0 to 2500-Hz range, as does the output amplitude.

If you want to sweep the circuit's output frequency, you can do so rapidly by varying the input clock rate. The counter's Borrow output provides a valuable zero-crossing pulse in such use. Additionally, by applying an ac signal to DAC_A's reference input, you can amplitude-modulate the output sine wave.

Generating triangles and rectangular pulses

Certain applications require triangular or rectangular waveforms in which you can program the period of each half cycle. For example, the former serve as sweep signals in vector-scan CRT displays to generate variable-length vectors. The design depicted in **Fig 4** serves such requirements extremely well.

In operation, DAC_A defines the ramp rate for the triangle's positive slope; DAC_B, the rate for the negative slope. Comparators A_4 and A_5 sense the integrator's output, driving gates G_1 and G_2 when the level reaches +10 or −10V. These gates act as an RS

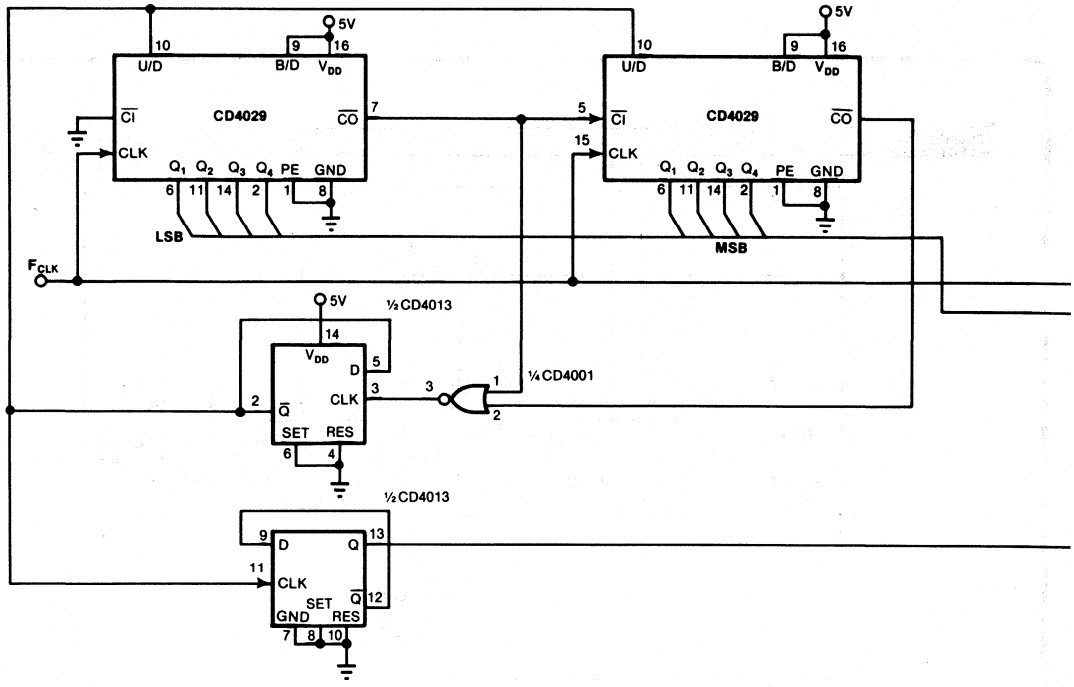

(a)

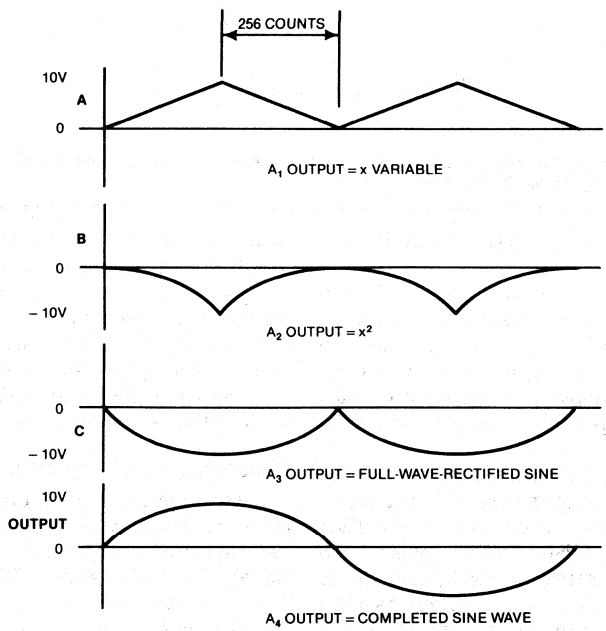

Fig 3—The multiplying capability *of the AD7528 gets used to advantage in this function-fitting sine-wave generator* ***(a)****. The digital representation of N is* DAC_B*'s input word, while the analog representation of N drives the DAC's analog reference input. As a result,* DAC_B *outputs an analog signal equal to* N^2*. Resistors* R_A *and* R_B *relate to the quadratic polynomial that approximates the 1-quadrant sin X function; they must be stable and accurate to achieve minimum distortion* ***(b)****.*

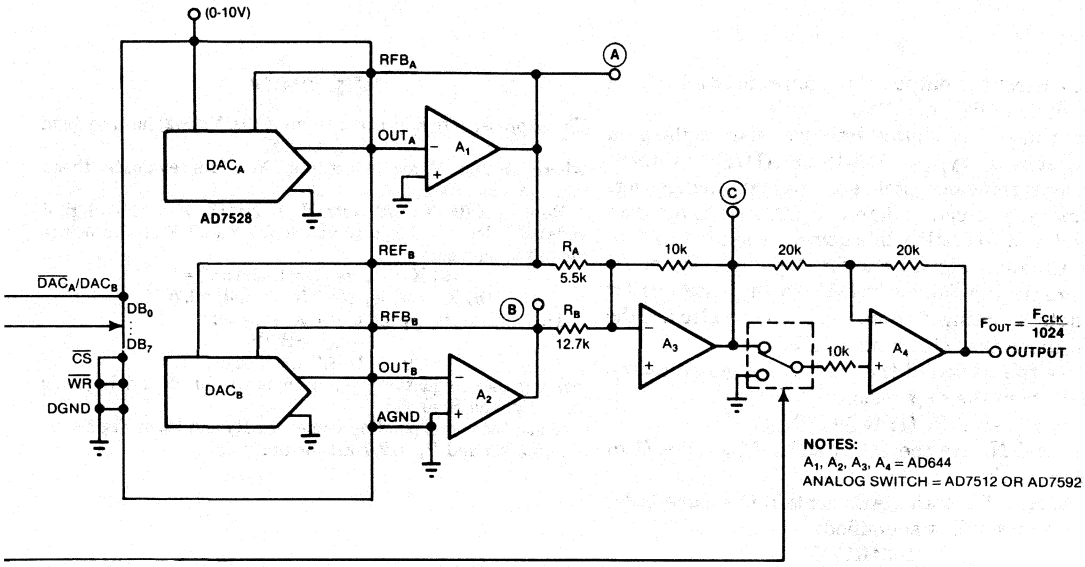
(0-10V)
RFB_A
A
OUT_A
A1
DAC_A
AD7528
C
REF_B
R_A
5.5k
10k
20k
20k
DAC_A/DAC_B
DB_0
DB_7
RFB_B
B
R_B
12.7k
A3
F_OUT = F_CLK / 1024
OUTPUT
CS
WR
DGND
OUT_B
DAC_B
A2
10k
A4
AGND
NOTES:
A1, A2, A3, A4 = AD644
ANALOG SWITCH = AD7512 OR AD7592

8

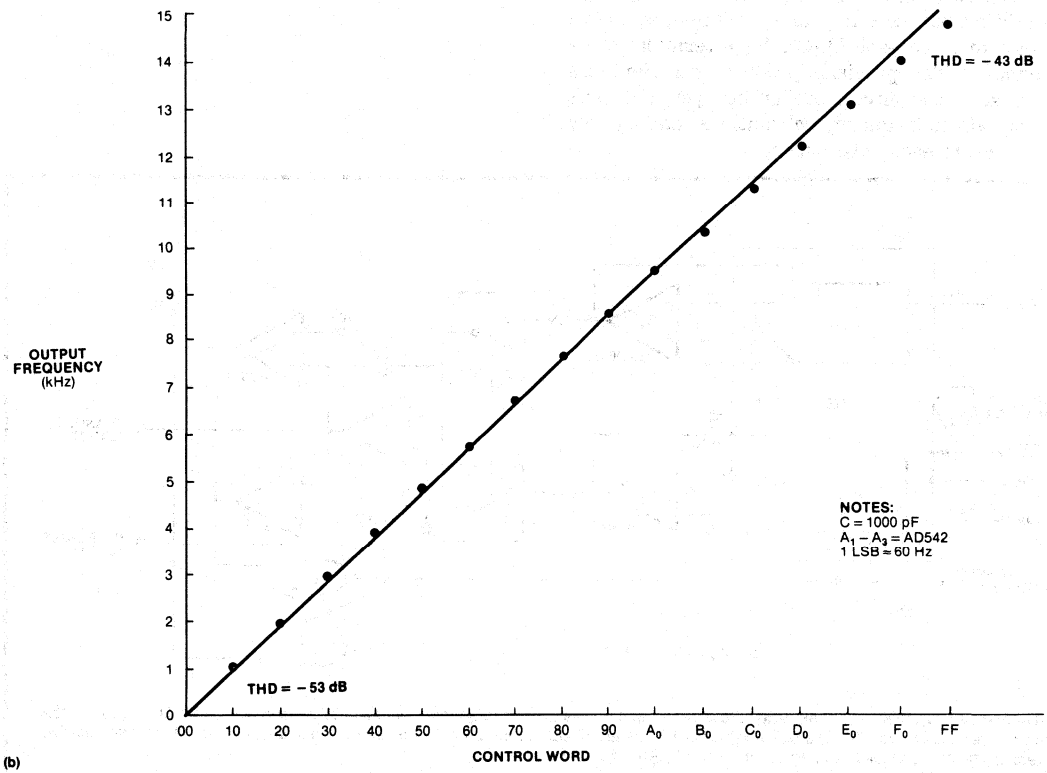
OUTPUT FREQUENCY (kHz)
15
14
13
12
11
10
9
8
7
6
5
4
3
2
1
0
THD = −43 dB
THD = −53 dB
NOTES:
C = 1000 pF
A1 − A3 = AD542
1 LSB = 60 Hz
00
10
20
30
40
50
60
70
80
90
A0
B0
C0
D0
E0
F0
FF
CONTROL WORD

(b)

flip flop to select the output of the appropriate DAC via analog switches SW_1 and SW_2.

The arrangement shown lets you use high-speed switches, such as types CD4016 or AD7592, without introducing significant glitches at the DAC changeover. Furthermore, you can update one DAC from the data bus and allow it to settle while using the second DAC to generate the ramp signal. The flip-flop output automatically connects the "unused" DAC to the data bus for further data updating if necessary. You can also use the output to drive the interrupts of a μP.

Calculate the period of the waveforms generated by this circuit from the expression

$$T=512RC((1/N_A)+(1/N_B)),$$

where N_A and N_B are the DAC_A and DAC_B codes (1 to 256), respectively.

If the latches for both DACs contain the same code, the expression for T is simplified:

$$T=1024RC/N$$

or $f=N/1024RC$ (in hertz).

The mark-to-space ratio of the rectangular wave equals N_B/N_A.

Note one caution regarding **Fig 4**'s circuit. If the code in either DAC becomes all ZEROs, the integrator input voltage becomes zero, and the circuit stops oscillating. If this condition can occur in your application, connect a 10-MΩ resistor from each DAC's V_{REF} terminal to its output terminal. This provision produces sufficient bias current to keep the circuit oscillating, yet it doesn't significantly affect frequency calculations, adding only ¼-LSB error into each DAC output.

Appendix

The expression to approximate sin X ($0 \le X \le \pi/2$) has the form

$$Y=AX+BX^2,$$

where A and B are constants. You can evaluate these constants as follows:

Replace X ($0 \le X \le \pi/2$) with N ($0 \le N \le 1$), where N=(digital code)/256. Next, substitute values for Y and N at two points along the sine curve:

(a) $X=\pi/2$, ie, $N=1$: $\sin(\pi/2)=1$

(b) $X=\pi/4$, ie, $N=0.5$: $\sin(\pi/4)=0.7071$.

This substitution produces two equations:

$$1=A(1)+B(1)^2$$

$$0.7071=A(½)+B(½)^2.$$

Solving these equations simultaneously for A and B yields $A=1.828$ and $B=0.828$.

Thus, the corresponding values for R_A and R_B in **Fig 3a** are $R_A=5.5$ kΩ and $R_B=12.7$ kΩ for unity gain.

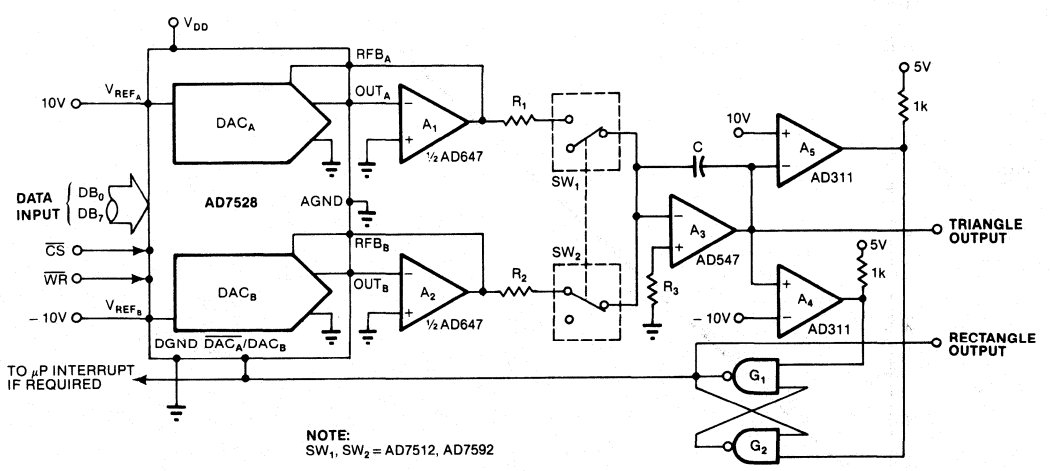

Fig 4—You can update one DAC and allow it to settle *while using the second DAC to generate a ramp in this programmable waveform generator. Oscillation stops if either DAC has an all-ZEROs input code, however. If this is a problem in your application, connect a 10-MΩ resistor between each DAC's V_{REF} and output terminals.*

ANALOG DEVICES

AN-323
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Microstepping Drive Circuits for Single Supply Systems

by John Wynne

This application note presents a brief review of two-phase permanent magnet stepper motors and outlines some of the issues involved in microstepping. The static position errors introduced by the DACs and by load torque are described and the merits of closed-loop versus open-loop control discussed. Two practical circuits are presented which use positive-only power supplies (+5V and +12V), a situation common in both floppy and low density hard disk drives.

SOME BASICS OF THE TWO-PHASE PM STEPPER MOTOR

In a permanent magnet stepper motor torque is generated by the interaction of two magnetic fields, the stator field and the rotor field. The stator field is generated by current flowing in the phase windings; the rotor field is due to the permanent magnet pole pairs arranged radially on the rotor. A two-phase stepper motor has two phase windings, A and B, separated by 90 electrical degrees, each phase having two possible current directions, positive and negative. Assuming only one phase is turned on at any one time, the stator field can have one of four possible orientations. The rotor field, and hence the rotor, will align itself with respect to the stator field at the one position where the field vectors are in equilibrium.

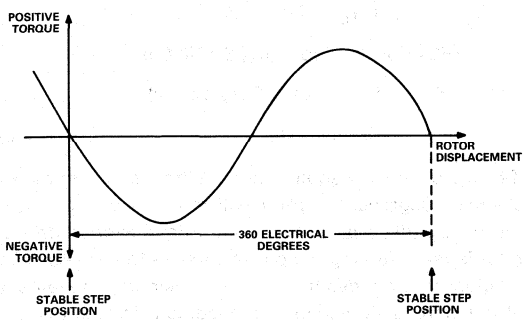

Figure 1. Static Torque vs. Rotor Displacement for an Ideal Stepper Motor

Exactly how the magnetic fields interact to produce torque can easily be seen from the ideal static torque versus rotor displacement curve of Figure 1. This is generated by energizing a single phase and rotating the rotor by hand.

The convention used in Figure 1 is that a movement to the right corresponds to a clockwise rotation of the rotor while a positive torque in the figure represents a clockwise torque developed on the rotor. The torque output can be seen to be a sinusoidal function of the electrical angle ϕ between the field vectors or:

$$T = -K\ I\ Sin\phi \qquad (1)$$

where K is a motor constant and I is the phase current. The maximum torque T_{max} exerted on the rotor occurs when the angle between the rotor and stator field vectors equals 90 electrical degrees. With maximum rated current I_R flowing, this maximum torque is known as the peak static torque or simply the holding torque, i.e., $T_{HMAX} = KI_R$ which is a constant for a given motor. The minus sign in the torque expression indicates that the position of stable equilibrium occurs on the negative slope of the torque curve. If the rotor is moved in a clockwise direction, a counterclockwise torque is developed to return the rotor to its stable or step position. A similar situation exists if the rotor is moved in a counterclockwise direction. Thus a constant phase excitation provides one stable rotor position every 360 electrical degrees. If the phase excitation is reversed, a new stable position will be established midway between these original positions, i.e, at 180 degrees. In a two-phase motor, the two-phase windings are physically separated by 90 electrical degrees, hence stable rotor positions will occur at every 90 electrical degree rotation of the stator field. A two-phase PM stepper will therefore move four full steps for every 360 electrical degree rotation of the stator field vector.

The number of pole pairs N determines the relationship between a stator field vector shift of α electrical degrees

and the resulting rotor movement of θ mechanical degrees;

$$\alpha = N\theta \quad (2)$$

The full step mechanical angle is therefore

$$\theta_{FS} = \pi/2N \quad (3)$$

For instance a rotor magnetized with 25 pole pairs will have a mechanical step angle of 3.6°. Figure 2 shows the complete static torque versus rotor position curves for a two-phase permanent magnet stepper motor with each phase excited in turn. Stable equilibrium positions for the rotor exist at the full step positions, $\pi/2N$, π/N, $3\pi/2N$ and $2\pi/N$. A full step rotation of the rotor results from a shift in the stator field vector of 90 electrical degrees. The basic sequence for a two-phase motor of 4 full steps for one full revolution (360 electrical degrees) of the stator field vector is obvious from Figure 2.

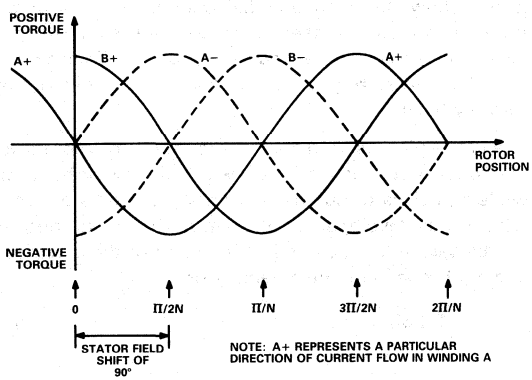

Figure 2. Static Torque vs. Rotor Position Curves for a Two-Phase, PM Stepper Motor

SOME MICROSTEPPING BASICS

In Figure 2 at rotor position 0, phase A current is positive and equal to its maximum rated value I_R; phase B current is zero. At rotor position $\pi/2N$, phase B current is positive and equal to its maximum rated value I_R while phase A current is zero. By correctly choosing and controlling the ratio of phase A current to phase B current it is possible to generate a stable equilibrium position anywhere between the stable full step positions. Figure 3 is a vector representation of the two-phase currents, I_A and I_B, where only the first step from the sequence of four in Figure 2 is represented.

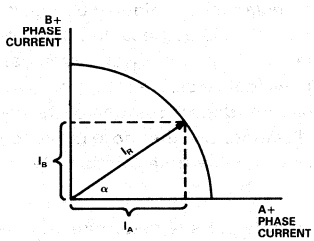

Figure 3. Vector Representation of Phase Currents

From Figure 3 the angular position of the stator field vector, and hence the rotor position, is proportional to the phase currents:

$$\alpha = \text{Arc tan } (I_B/I_A) \quad (4)$$

In a practical circuit dual digital-to-analog converters (DACs) set the required phase current magnitudes while pulse width modulator (PWM) controllers ensure the set levels are met and maintained in the windings. The maximum number of current ratios or microsteps possible is determined by the resolution n of the DACs, i.e., maximum number of microsteps = 2^n. In practise the number of microsteps used is usually an integer power of 2, i.e,. 8,16,32, etc. There are advantages to using DACs whose resolution exceeds the minimum resolution necessary to generate the required number of microsteps.

Not alone is it necessary to choose the correct ratios of phase currents for each new microstep, but it is also necessary to ensure that the holding torque due to the combined phases remains constant and equal to the holding torque obtained with full rated current in a single phase, i.e., from Figure 3

$$\sqrt{I_A^2 + I_B^2} = I_R = \text{Constant} \quad (5)$$

If this condition is not met, it is possible, with a purely friction load, to have a different position error at each microstep. This is discussed later.

The torque output generated by the two motor phases is:

$$T_A = -KI_A \sin\phi \quad (6)$$

$$T_B = -KI_B \sin(\phi - \pi/2)$$

$$\text{or } T_B = KI_B \cos\phi \quad (7)$$

To move the rotor by θ mechanical degrees (or the stator field vector by $N\theta$ electrical degrees) it is necessary to vary the phase currents in a cosinusoidal fashion:

$$I_A = I_R \cos N\theta \quad (8)$$

$$I_B = I_R \sin N\theta \quad (9)$$

In an ideal motor both torque contributions can be added to provide the total torque:

$$T_{TOT} = T_A + T_B$$

$$= -KI_R \cos N\theta \sin\phi + KI_R \sin N\theta \cos\phi$$

$$= -KI_R (\sin\phi \cos N\theta - \cos\phi \sin N\theta)$$

$$T_{TOT} = -KI_R \sin(\phi - N\theta) \quad (10)$$

If the stable torque equilibrium position is moved by $N\theta$ electrical degrees, the rotor will follow it by moving θ mechanical degrees to where the torque exerted on the rotor is zero. Varying the phase currents in a sine-cosine relationship provides the required microstep movement of the stator field vector and ensures a constant stator field vector magnitude (or holding torque) at each microstep position.

STATIC POSITION ERRORS

Microstepping accuracy depends on how closely equations 6–10 can be realized in practise with a real motor. The perfectly sinusoidal static torque vs. rotor displacement curve exhibited by an ideal motor can be upset by a number of factors such as the presence of detent torque, a nonlinear relationship between torque and energizing current, by nonidentical static torque vs. angular displacement curves for Phase A and Phase B, etc. It is worthwhile analyzing the positional errors which can be caused by the PWM controllers and the DACs.

PHASE GAIN ERRORS

One of the most common causes of inaccuracies in microstepping is unequal maximum torque contributions from both phases. This can be due to the motor itself, to mismatch in the PWM controllers (for instance, due to different values of sense resistors) or to mismatch in the DAC outputs (for instance, different gain errors). The torque expressions (Equations 6 and 7) can be rewritten to include a phase gain error:

$$T_A = -KI_A \sin\phi \quad (11)$$

$$T_B = (1+M)\, KI_B \cos\phi \quad (12)$$

where $I_A = I_R \cos N\theta$ (8)

$$I_B = I_R \sin N\theta \quad (9)$$

and M is the mismatch between phases. The stable rotor position occurs when $T_A + T_B = 0$:

$$-KI_R \cos N\theta \sin\phi + (1+M)\, KI_R \sin N\theta \cos\phi = 0 \quad (13)$$

$$\text{or } \tan\phi = (1+M) \tan N\theta \quad (14)$$

$$\text{or } \phi = \text{Arc tan}[(1+M) \tan N\theta] \quad (15)$$

For an ideal motor at its stable rotor position $\phi = N\theta$. The difference between ϕ and $N\theta$ at the stable rotor position indicates positional error:

$$\phi - N\theta = \text{Arc tan}[(1+M) \tan N\theta] - N\theta \quad (16)$$

Equation 16 is plotted in Figure 4 for two phase gain error values of 1% (M = 0.01) and 5% (M = 0.05). The horizontal axis of Figure 4 is the rotor position in electrical degrees. Maximum error occurs at the half-step position. For example, in a system with sixteen microsteps per full step, a phase gain mismatch of 5% results in a microstep position error of nearly 25%.

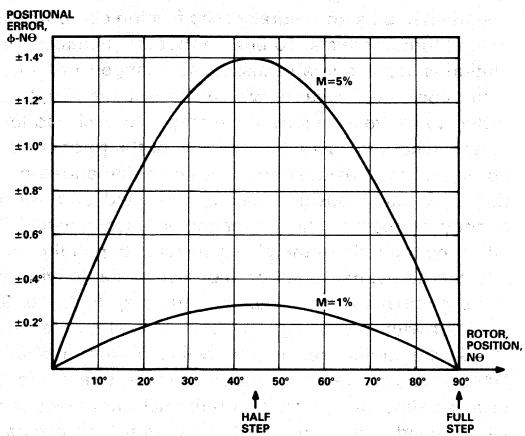

Figure 4. Microstepping Positional Error Due to Phase Gain Errors

ERRORS DUE TO DAC RESOLUTION AND ACCURACY

From before, the angular position of the rotor is proportional to the ratio of phase currents

$$\alpha = \text{Arc tan}\,(I_B/I_A) \quad (4)$$

The locus of ideal phase current ratios required to microstep the rotor between two full-step positions lies along the quarter circle in Figure 3. However, due to the finite resolution of the DACs in a practical system, it is not possible to choose the ideal current ratios required at each microstep to exactly track the circular locus. Consequently, the DAC finite resolution introduces an angular error in the different positions of the current vector. The angular error at each microstep can be expressed as

$$\delta_\alpha = \text{Arc tan}\,(I_B/I_A) - \alpha_{MICROSTEP} \quad (17)$$

where $\alpha_{MICROSTEP}$ is the ideal microstep size in any given system. For example, Figure 5 compares the percentage angular errors generated by two different resolution DACs (6-bits and 8-bits) for a system with 8 microsteps per full step (i.e., $\alpha_{MICROSTEP} = 90°/8 = 11.25°$). Figure 5a shows that the limited resolution available with a 6-bit DAC can contribute almost 5% angular error to the microstep position; whereas with an 8-bit DAC, Figure 5b, the errors are approximately an order of magnitude less.

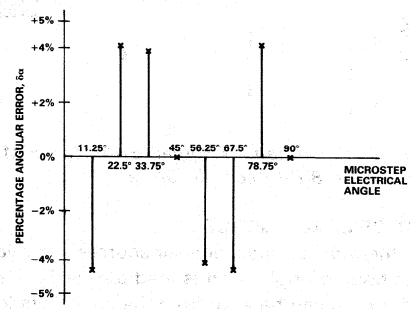

Figure 5a. Percentage Angular Error for 8-Microstep System Using Ideal 6-Bit DACs

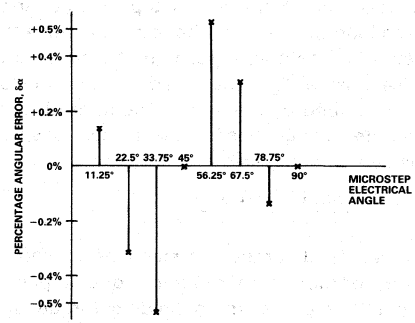

Figure 5b. Percentage Angular Error for 8-Microstep System Using Ideal 8-Bit DACs

The graphs of Figure 5 assume the DACs have no integral linearity errors. If the DACs are assumed to be accurate to their resolution level, i.e., the 6-bit DACs are 6 bits accurate (±0.8%) and the 8-bit DACs are 8 bits accurate (±0.2%), the errors increase over the previous case. Figure 6 repeats the comparsion using such DACs. Maximum angular error using a 6-bit accurate DAC can now be up to 10%.

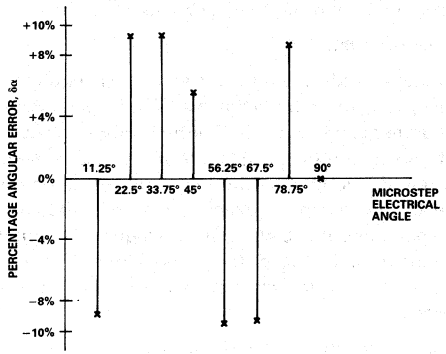

Figure 6a. Percentage Angular Error for 8-Microstep System Using 6-Bit Accurate DACs

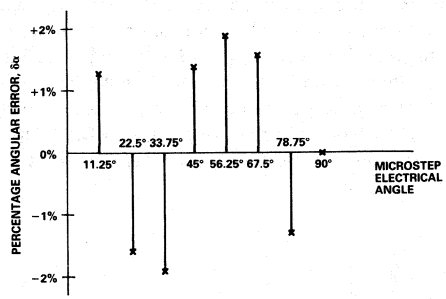

Figure 6b. Percentage Angular Error for 8-Microstep System Using 8-Bit Accurate DACs

EFFECT OF LOAD TORQUE

One of the most significant error sources in a microstepper positioning application is load torque T_L. The effect of load torque can be seen from the typical static torque vs. rotor position curve in Figure 7.

With no load, the rotor is at rest at position 0. Applying a clockwise load torque T_L to the rotor will move the rotor to a new rest position, a position at which the applied clockwise torque T_L is balanced by the anticlockwise torque developed by the motor. The rotor displacement between the two rest positions is the step position error θe. This can be expressed as

$$\theta e = \frac{\text{Arc Sin } (-T_L/T_{HMAX})}{N} \qquad (18)$$

For a given load torque the step position error is inversely proportional to the number of pole pairs on the rotor and to the maximum holding torque available. Increasing the number of pole pairs decreases the full step size thereby increasing the slope of the static torque

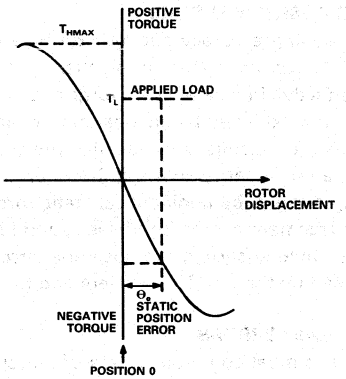

Figure 7. Static Position Error Due to Load Torque

vs. rotor displacement curve over that obtainable with a bigger step size motor. For two motors with equal load torque and holding torque but with different step sizes the steeper slope will result in smaller step position errors. However, not alone do motor costs obviously increase with smaller step size, but so too does the static torque vs. displacement curve depart from the ideal sinusoidal relationship. Specialist motors are available from manufacturers which have been deliberately optimized for microstepping during manufacture.

The other option for reducing step position error indicated by Equation 18, that of increasing the holding torque (by increasing the phase currents), is not available when microstepping since the cosinusoidal phase currents ensure the exact opposite occurs, i.e., that the holding torque remains constant at each microstep. Thus, when microstepping, the maximum holding torque is deliberately restricted to equal the one-phase on torque, not the higher torque which normally would be available with two-phase-on operation.

CLOSED-LOOP VS. OPEN-LOOP CONTROL

Microstepping improves the positioning resolution possible in any control application. However, the positional accuracy can be significantly worse than that offered by the original full-step accuracy specification due primarily to load torque effects. To ensure that the increased resolution is useable it is necessary to change from a normally open-loop system, where the rotor is simply directed to move so many microsteps, to a closed-loop system where the actual movement of the rotor is monitored and any residual error between demanded position and actual position can be detected and subsequently corrected. The correction is accomplished by adjusting the ratio of the phase currents to pull the rotor into the demanded position. For two systems with the same number of microsteps per full step (e.g., 16, 32 etc.) but with one using 6-bit DACs and the other 8-bit DACs, the higher resolution available from the 8-bit DACs offers a greater number of current ratios to be selected allowing a more accurate final rotor position to be achieved. Thus, fine positioning with high accuracy can be achieved with microstepping but only in a closed loop system.

PRACTICAL CIRCUITS

Two typical microstepping drive circuits are presented in this application note. Both use dual 8-bit DACs to control the phase currents and a single 8-bit ADC for closed-loop position control. The first circuit, Figure 9, uses an AD7628 dual DAC and an AD7820 ADC. The AD7628 is identical to the industry standard AD7528 but offers TTL compatibility (and is fully specified) at a +12V power supply. It is beyond the scope of this application note to deal in any detail with the types of position/velocity transducers which are used to monitor movement of the head assembly in disk drives. The AD7820 half-flash 8-bit ADC is widely used for converting these transducer output signals and it is shown in block diagram form in Figure 9. For example, a typical transducer might be an incremental optical encoder with quadrature triangular output waveforms although stepper motors are now becoming available (such as the Portescap P750) which include, in addition to the phase windings, integral velocity windings which can be used directly to obtain rotor velocity and hence position. The second circuit, Figure 11, uses a recently released device from Analog Devices, the AD7669, which contains both a dual 8-bit DAC and an 8-bit ADC.

Typical DAC output voltage waveforms and associated direction control information for clockwise rotation are shown in Figure 8.

The "sine/cosine" data is usually held in look-up tables in the controller. The rate at which the data is loaded to the DACs obviously determines the microstepping rate. The controller also supplies the direction control information required by the H-bridge power driver. The output driver used in both circuits is the UDN2998W, a dual H-bridge driver available from Sprague Electric. It is packaged in a 12-pin single-in-line power tab package

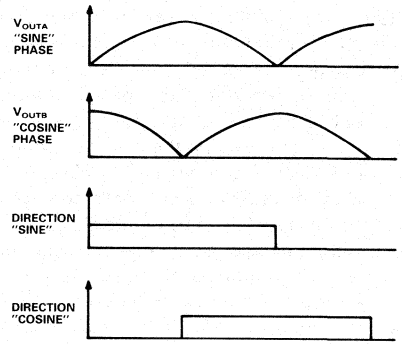

Figure 8. Typical DAC Output Voltages for Microstepping. Direction Signals Are for CW Rotation with UDN-2998W.

for high current capabilities (+2A continuous). In practice, it was found necessary to ground the power tab to a quiet ground in order to avoid parasitic oscillations.

The circuits of Figures 9 and 11 also differ from each other in that one is based on a fixed frequency PWM technique while the other uses a frequency modulation PWM approach. The fixed frequency approach allows a synchronized chopping frequency which can help to reduce noise on the power supply lines and ground return paths. The frequency modulation approach provides a fixed-ripple current characteristic leading to a higher efficiency drive.

8

FIXED FREQUENCY PWM DRIVE

The complete circuit is shown in Figure 9. The NE555 timer produces short (2.5μs) low level pulses at a repetition rate of approximately 45kHz. This is high enough to avoid audible noise and low enough to avoid unnecessary switching losses. These negative pulses reset

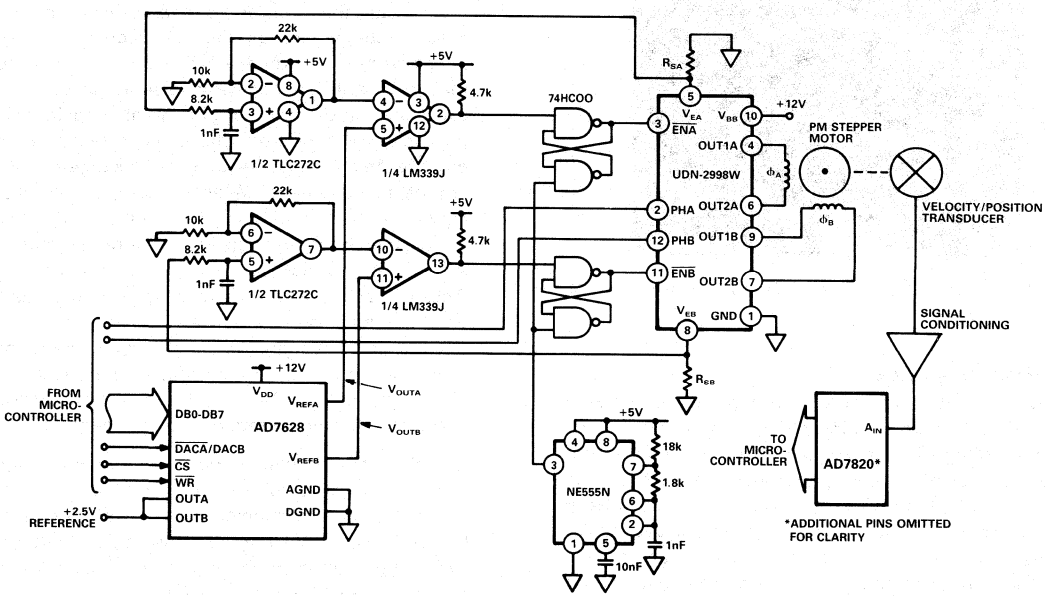

Figure 9. Fixed Frequency PWM Drive

the flip-flops and enable the phase currents via the UDN2998W dual H-bridge driver. The actual value of the phase currents are monitored by means of the sense resistors R_{S1} and R_{S2} between the emitters of the output sink drivers and ground. The voltage signals across the sense resistors – representing the phase currents – are filtered by the 8.2kΩ/InF combination to eliminate noise spikes, amplified by X3.2 across I_{C1} a/b and compared with the DAC control voltages by I_{C2} a/b. For either channel, when the amplified sense voltage exceeds the DAC control voltage, the comparator output is driven low setting the flip-flop and turning off the source driver. The phase current now decays slowly, the recirculation path being through the on-chip ground-clamp diode, phase winding and output sink transistor. The current will continue to fall until the next pulse resets the flip-flop and the sequence is repeated.

The lower trace of Figure 10 shows a typical output voltage waveform from one of the DACs. The waveform covers two full steps with eight microsteps per full step. The top trace is the amplified sense voltage at the output of the corresponding sense amplifier. The vertical axis is 1V/division. The horizontal axis is uncalibrated for photographic purposes, but the time for eight microsteps is 12.8ms. The motor used is a Portescap P530 two-phase stepper motor with a full step size of 3.6°. It is intended for 0.45° microstepping – eight microsteps per full step – resulting in a total of 800 microsteps per revolution. With a reference voltage of +2.5V on the DACs and a sense resistor of 0.5Ω, the peak current is (2.5/3.2)/0.5 or 1.56A which is equal to the nominal rated current for one-phase on operation.

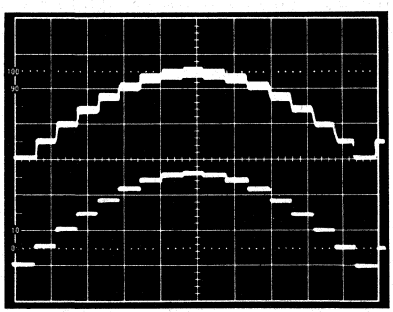

Figure 10. Circuit Waveforms for Fixed Frequency PWM Circuit of Figure 9

VOLTAGE-MODE DACs

The AD7628 DACs are operated in what is known as a voltage-mode or voltage-switching configuration. The reference voltage is applied to OUTA and OUTB terminals of the AD7628, and the output is taken from the V_{REFA} and V_{REFB} terminals. A positive reference voltage gives a positive output voltage – hence the output signals are voltages (not currents) at a constant source impedance equal to the ladder resistance and can be expressed as

$$V_{OUTA} = D_A \cdot V_{REF} \quad (19$$

$$V_{OUTB} = D_B \cdot V_{REF} \quad (20$$

where $D_A = N_A/256$ and $D_B = N_B/256$, with N_A and N_B i hexadecimal format being the codes supplied to DAC A and DAC B respectively. This mode of operation is virtually free of gain error, hence DAC to DAC output voltag matching will be better than 0.4%. Close matching i important in open-loop micropositioning application since inequalities between the maximum values o phase currents will cause positioning errors.

The circuit of Figure 9 responds to DAC output voltage extending from V_{REF} down to and including 0V. An operating range which includes 0V is important in applications such as disk drive read/write head movemen which require very fine positioning indeed. Such systems can have 32, 64 or more microsteps per full ste producing very small voltage steps around 0V. Thes occur when the rotor position is close to a full ste position. Also at this position the phase with the leas current in it has the biggest impact on the positionin error.

FREQUENCY MODULATION PWM DRIVE

By removing the flip-flop from Figure 9 which synchronizes the phase-on times with a master clock frequency, the resultant circuit will free-run at a rate determined by the response of the loop components and the motor parameters. This circuit is shown in Figure 11.

For any one phase assume that the comparator output is high indicating that the phase current is less than it should be. The high on the comparator output produces a low on the enable input of the H-bridge and turns on the phase current. The phase current is monitored by the sense resistor, filtered and amplified before being compared with the control voltage. When the amplified sense voltage exceeds the DAC control voltage, the phase current is turned off. The comparator output remains low until the decaying current falls sufficiently to cause the comparator output to go high again, restarting the cycle. Although both phases have separate loops, in practice the chopping frequencies lock together avoiding any beat phenomena. With the components shown in Figure 11, and again using the Portescap P530 stepper motor, the chopping frequency is approximately 25kHz. It should be possible to reduce the component count of Figure 11 by replacing the dual op amp and quad comparator (half of which is used) by one of the modern quad op amps designed for single supply operation. Due to the slower response time of the op amps when used as comparators the chopping frequency will drop but the circuit should still operate due to its regenerative nature.

It is also very simple to add an extra control signal to allow the phase currents to be turned off independent of the controller. By replacing the inverter gates shown in Figure 11 with NAND gates, the spare input on each NAND gate can now act as a control, high for normal operations; low for emergency stop.

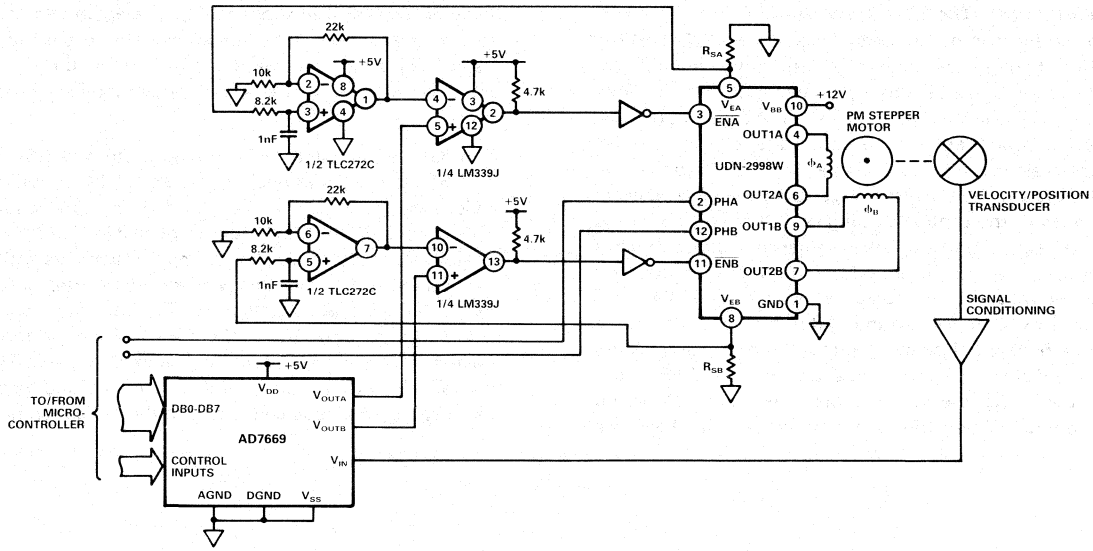

Figure 11. Frequency Modulation PWM Drive

8

The top trace in Figure 12 shows the amplified sense voltage at the output of one of the sense amplifiers. The lower waveform shows the DAC control voltage for that phase. To allow a direct comparsion with the previous waveform of Figure 10 the vertical axis of Figure 8 is again 1V/division with an identical microstepping rate of 1 microstep per 1.6mS or 8 microsteps per 12.8mS. The fixed-ripple current characteristic of the drive is very obvious.

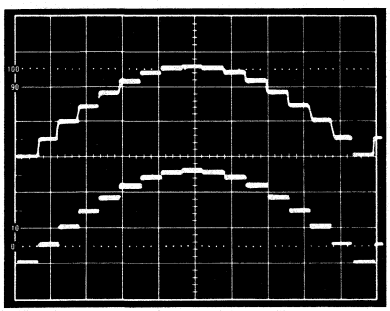

Figure 12. Circuit Waveforms for Frequency Modulation PWM Circuit of Figure 11

VOLTAGE OUTPUT DACs

The two DACs used in this circuit are 8-bit voltage DACs contained on the monolithic AD7669 die. An on-chip bandgap voltage reference of +1.23 (not available to users) drives both DAC inputs. The output buffer amplifiers of each DAC are configured with internal X2 gains and provide unipolar 0 to +2.5V outputs. DAC-to-DAC full-scale error matching is better than ±0.4%. As mentioned previously, the AD7669 also contains an 8-bit ADC. The ADC input voltage range matches the DACs' output voltage range of 0 to +2.5V with a typical DAC-to-ADC gain match of 1%.

MICROSTEPPING RATE

Outside of disk drive applications where microstepping is used to position the read/write head, its primary attraction is the extremely smooth single step response which can be achieved by replacing a single excitation change with a number of smaller amplitude excitation changes. Such applications are inherently low speed – below 100 rpm – since at higher speeds the motor inertia makes the rotor rotate smoothly regardless of whether the drive is microstep, full step or half step. In head positioning applications in disk drives motor speed will depend on whether the drive is in the seek or track mode. In the seek mode where the head is moving some distance across the disk surface, highest possible speeds are required to meet the velocity profile. In the track mode where the head is tracking a selected data track, the microstepping rate is lower. For the fixed frequency PWM circuit of Figure 9, the maximum microstep rate before any individual microstep disappeared was found to be 1400 microsteps/second. With 0.45° microsteps this rate translates into 105 rpm. The rate can be much increased above this before the motor stalls, but the sense voltage waveforms increasingly depart from the DAC control waveforms. With 0.225° microsteps (16 per full step) the maximum microstep rate which retains all of the microsteps in the sense waveforms was found to be 1850 microsteps/sec or approximately 70 rpm. Similar results were found for the frequency modulation PWM circuit of Figure 11. With 0.45° microsteps, maximum microstep rate retaining all eight levels was found to be 1330 microsteps/sec or 100 rpm. With 0.225° microsteps the rate was 1660 microsteps/sec or approximately 62 rpm.

NEW PWM MOTOR DRIVER AVAILABLE

Sprague Electric has recently released the UDN2917, a new dual full-bridge stepper motor driver suitable for microstepping applications. Output current in both windings of a bipolar stepper motor is sensed and controlled independently in each bridge by an external sense resistor, internal comparator and monostable multivibrator. When the bridge is turned ON, current increases in the motor winding, and it is sensed by the external sense resistor until the sense voltage reaches the level set at the comparator's input. The comparator then triggers the monostable which turns OFF the source driver of the bridge. After turn-off, the motor current decays, circulating through the ground clamp diode and sink transistor. The source driver's OFF time is determined by the monostable's external RC timing components. When the source driver is re-enabled, the winding current (the sense voltage) is again allowed to rise to the comparator's threshold. This cycle repeats itself, maintaining the average motor winding current at the desired level. This is a frequency modulation PWM drive.

The outputs are capable of driving motors at 1.5A and 45V. Maximum combined (source and sink) output saturation voltage is 2.6V at 1A and 3.1V at 1.5A. A lower power version, the 2916, is also available. The 2917 is available in two packages for applications versatility. One option is a 28-pin DIP with the two center pins on each side bonded together (commonly known as a batwing) and connected to the die attach pad to form a heat tab. For surface mount applications, a 44-pin PLCC is available with two sides of this quad package formed into heat sinking tabs connected to the die attach pad.

ANALOG DEVICES

AN-324
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Simple DAC-Based Circuit Implements Constant Linear Velocity (CLV) Motor Speed Control

by John Wynne

Most hard-disk and floppy-disk magnetic drives operate with the disk spinning at a constant speed and the READ/WRITE head transferring information to or from the disk at a constant rate. This type of recording is called constant angular velocity (CAV) recording since each data-bit call is given the same angular velocity regardless of its position on this disk surface. However, the size of the data bits will vary depending on their location. Closest to the hub the bit-cell density is maximum thinning out to a minimum nearest the disk circumference. This non-uniform areal bit density (number of bits in a given area) is obviously wasteful of disk space.

CD-ROM drives—and CD players—employ a different type of recording technique in order to achieve a constant areal bit density over the operating disk surface. These drives use constant linear velocity (CLV) recording, a technique which actually varies the speed of the spinning disk according to where the data-bit cell is located. Disk speed is highest for cell locations closest to the hub with decreasing speed for cell locations towards the disk circumference. Varying the rotational speed ensures that the distance between data-bit cells is constant regardless of cell location, and it also maintains a constant data transfer rate from the disk. CLV recording can boost disk capacity by up to 35% over that achievable with CAV recording. However, one of the drawbacks to this technique is the relatively long access times due to the required speeding-up or slowing-down of the disk.

It is somewhat ironic that the first generation of optical CD-ROM drives on the market in the main ignored the capacity boost possible with CLV (an ideal drive for optical drives because of the continuous spiral track on which data is recorded) and concentrated instead on reducing the seek times in order to compete with high density hard-disk drives. Some manufacturers have introduced drives which go some way towards combining the high density offered by CLV with the fast access times offered by CAV by using a technique called banding or zoning. In this approach the surface of the disk is divided into a small number of concentric bands or zones. Within each band the recording is CAV. From band to band a degree of CLV is achieved, not by changing the rotational speed of the disk, but by changing the data recording rate for each new band. This complicates the design of the controller somewhat since a variable frequency phase-locked loop (PLL) will be required for the data-separator circuit. If access times are the only yardstick by which different drives are to be compared then CAV-recorded disks will inherently be faster than CLV-recorded disks. However many applications exist where disk capacity is as important as, if not greater than, data access times; archival applications being an obvious one. CLV recording with its substantial capacity advantage will be the drive of choice in these areas.

8

The circuit presented here produces a transfer characteristic which mimics exactly the relationship between data-bit cell position (i.e., head position) and disk rotational speed required to achieve CLV recording. In mimicking this relationship the DAC digital input data represents the radius at which the READ/WRITE head is to be positioned, while the circuit output voltage represents the correct rotational speed for that head position. By employing such a circuit in a servo loop controlling the speed of the spindle motor it should be possible to directly program the correct disk speed for any data location on the disk resulting in faster data access times.

SPEED VS. POSITION

For any rotating system, spindle speed and linear velocity are related by the simple expression (see Figure 1).

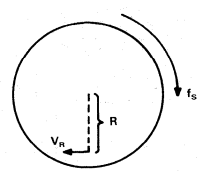

$V_R = 2\pi f_S R$

V_R = LINEAR VELOCITY

f_S = SPINDLE SPEED

R = RADIAL POSITION

Figure 1. Pictorial Explanation of Terms

$$V_R = 2\pi f_S R \quad (1)$$

where V_R is the linear velocity

f_S is spindle speed

R is radial position

Essentially V_R is the velocity of the disk passing under the READ head. This linear velocity must remain constant over the operational area of the disk in order to ensure a constant data transfer rate. For this to happen the spindle speed must be inversely proportional to the radial position of the head or

$$f_S = \frac{V_R}{2\pi R} \quad (2)$$

An optical drive platter has maximum and minimum radii between which the data is stored. The radial position of the READ head can be expressed in terms of these limits as

$$R = R_{MIN} + X(R_{MAX} - R_{MIN}) \quad (3)$$

where R_{MIN} is the minimum operating disk radius,

R_{MAX} is the maximum operating disk radius,

X is a fractional representation of the required head position, $0 \le X \le 1$.

Substituting this expression into the previous one and normalizing the result to the highest spindle speed produces a ratioed expression as

$$f_{RATIO} = \frac{f_S}{f_{S\,MAX}} = \frac{1}{1+X\left(\frac{R_{MAX}}{R_{MIN}} - 1\right)} \quad (4)$$

The READ head at the minimum radius corresponds to X = 0 and has the highest spindle speed. As X increases towards the outer edge of the disk the spindle speed decreases. The shape of this curve is shown in Figure 2.

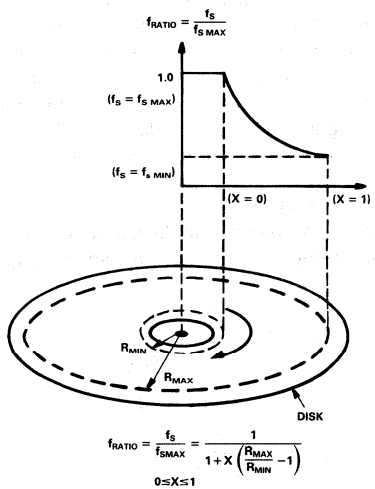

Figure 2. Relationship Between Disk Spin Speed and Head Position for a Constant Linear Velocity

The minimum value of f_{RATIO} depends on the disk diameter. For example, with a disk of diameter 12 inches R_{MAX} could be 5 inches and R_{MIN} be 2 inches. This means f_{RATIO} varies from 1.0 to 0.4, i.e., the minimum speed is 0.4 f_S max.

THE MIMICKING CIRCUIT

The circuit in Figure 3 has a transfer function equal to

$$V_{RATIO} = \frac{V_{OUT}}{V_{REF}} = \frac{1}{1+D1\frac{R_{FB}}{R_{DAC1}}}$$

where D1 is a fractional representation of the DAC digital input data and can vary from 0 to almost unity and R_{DAC1} is the ladder resistance of DAC 1.

Note that

$$D1 = \frac{N1}{2^n}$$

where N1 is digital input data in decimal and

n is resolution of DAC.

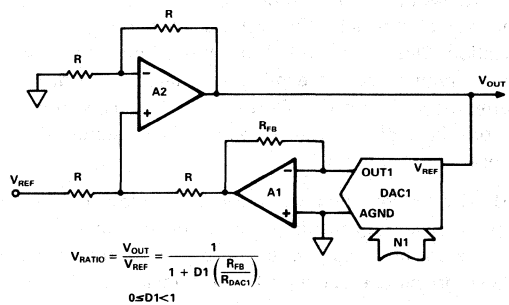

Figure 3. Circuit to Mimic the Speed/Position Curve of Figure 2

The V_{RATIO} transfer function of Expression (5) mimics the required f_{RATIO} relationship of Expression (4) exactly. The value of the R_{FB}/R_{DAC1} ratio in Figure 3 is dependent on the DAC resolution and the R_{MAX}/R_{MIN} ratio which is fixed for a given disk diameter and can be found by equating the denominators of Expressions (4) and (5) as:

$$\frac{R_{FB}}{R_{DAC1}} = \frac{X}{D1}\left(\frac{R_{MAX}}{R_{MIN}} - 1\right) \quad (7)$$

Using the R_{MAX} and R_{MIN} values from the previous example and a 12-bit resolution DAC and solving Expression (7) when X = 1 and D1 = 4095/4096 yields an R_{FB}/R_{DAC1} ratio of 1.5004. Due to the manufacturing spread in DAC ladder resistance, implementing the required R_{FB}/R_{DAC1} ratio on a production basis is itself an interesting exercise and will be dealt with separately [see later].

Using The Circuit

A block diagram of a possible motor speed controller incorporating the DAC circuit is shown in Figure 4. The effect of the DAC circuit is to augment the existing motor speed control circuitry (shown in dashed box) by providing a coarse feedback signal which very quickly ensures the disk is speeded up or slowed down to the proper speed after a new SEEK command is received. When the READ head approaches its correct new position, motor control is handed back to the fine-speed control elements within the dashed box.

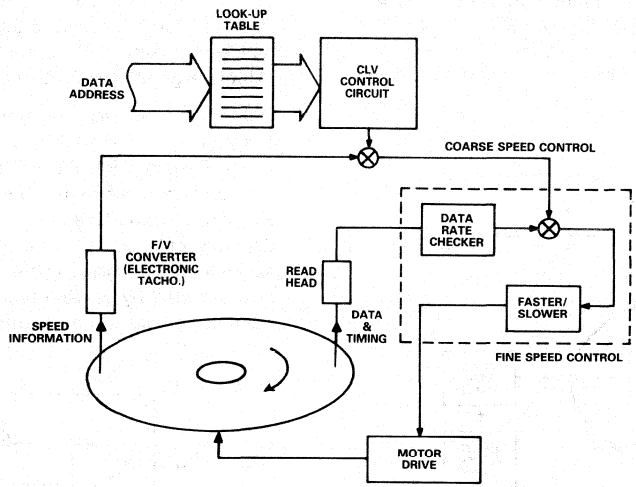

Figure 4. Block Diagram of CLV Motor Speed Controller with Coarse and Fine Speed Control Loops

8

The disk look-up table in Figure 4 performs the task of changing absolute data addresses into corresponding fractions of the disk radius. Figure 5 shows the "transfer curve" for the digital-in, digital-out look-up table. The curve is nonlinear reflecting the fact that with constant density recording more data locations are available as the spiral track circumference increases, i.e., towards the edge of the disk. With CAV recording this relationship would be linear reflecting the fact that each track across the disk surface has the same amount of data on it.

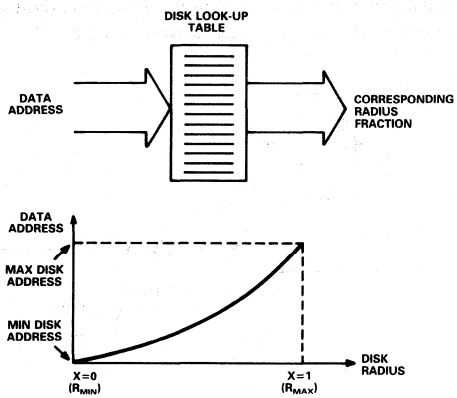

Figure 5. Relationship Between Disk Data Address and Corresponding Location on Disk Surface

The number of look-up values in the table is equal to or less than the resolution of the DAC being used. With a 12-bit DAC such as the AD7545, 4096 different disk speeds are programmable, hence the look-up table need be no bigger than 4096 entries. Each of these different speeds corresponds to a block of contiguous data addresses. The fine speed control circuitry can differentiate (if required) between data addresses and hence disk speed in any one block. Using a higher resolution DAC obviously increases the number of programmable speeds available, e.g., a 14-bit DAC such as the AD7538 allows 16,384 speeds to be programmed. As an example of the system in operation in a CD-ROM drive, consider that the READ head is positioned at the minimum disk radius R_{MIN}. At this radius the disk will be revolving at its maximum speed f_S max, typically 900 rpm. A SEEK command to read data from an address corresponding to a position 3/4 of the way across the operating disk area requires that the disk speed be reduced by 53% to 423 rpm, an f_{RATIO} of 0.47 (from Equation 4 and using an R_{MAX}/R_{MIN} ratio of 2.5). The error signal to achieve this reduction is generated by loading the digital equivalent (or the closest code possible) of the fractional 3/4 distance into the DAC. In this example the most suitable code, N, is 3072_{10} which is loaded from the look-up table. From Equation (5) the output voltage of the circuit will now be 0.47 V_{REF}, a V_{RATIO} of 0.47. The signal is compared with the output of the linear frequency/voltage converter whose output voltage at f_S max is

V_{REF}.. The difference between these two signals is the coarse error signal which drives the disk to its correct speed of 423 rpm. The same code, 3072_{10}, from the look-up table is loaded into another DAC producing a voltage which will drive the head to the required position, 3/4 of the way across the disk surface. Figure 6 shows this coarse head positioning loop with a voice coil motor driving the laser head assembly.

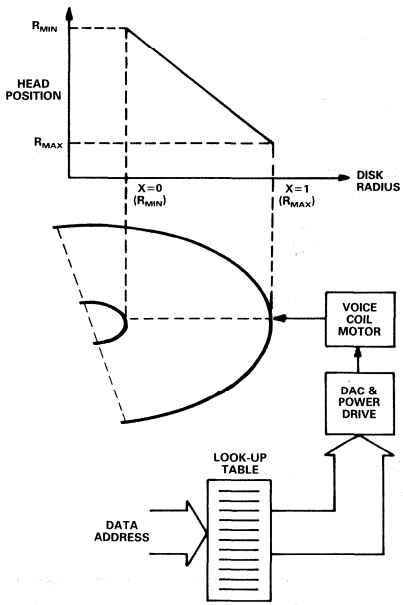

Figure 6. Coarse Control Loop for Positioning Head Assembly

IMPLEMENTING THE R_{FB} /R_{DAC1} RATIO

Standard CMOS DAC manufacturing strives to make the on-chip feedback resistor R_{FB} as equal as possible to the DAC ladder resistance R_{DAC}. The circuit of Figure 3 relies on being able to choose an R_{FB}/R_{DAC1} ratio which is not unity but dictated by the system parameters of disk size and number of programmable speeds required, i.e., DAC resolution. Equation 7 gives the exact relationship between these parameters. Since the required R_{FB}/R_{DAC1} ratio is always going to be greater than unity, the obvious solution is to add a resistor R1 in series with the feedback resistor, Figure 7.

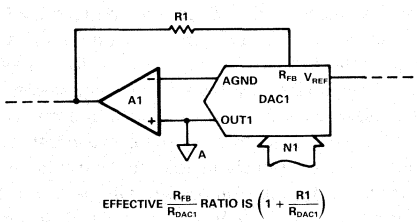

Figure 7. Incorrect Way to Implement R_{FB}/R_{DAC1} Ratio

There are two drawbacks to this solution. The first is tha R1 would have to be a select-on-test resistor due to th manufacturing spread in absolute values of the DA ladder resistance. Typically the spread is ±20% around mean value of 11kΩ. A second drawback is that fo significant changes in gain, R1 will be large relative t the DAC resistors and its temperature coefficient is un likely to match the temperature coefficient of the DAC typically −300ppm/°C. As a result such a gain trim woul exhibit large temperature coefficients.

One solution to realizing nonstandard, repeatable gai factors uses three additional resistors as shown in Fig ure 8. Resistors R1, R2 and R3 should all have simila temperature coefficients, but they need not match th temperature coefficient of the DAC. The overall gai depends only on the values of external resistors an excludes the absolute value of the ladder itself. Gain i best adjusted by varying the attenuator ratio, as adjust ment sensitivity is almost unaffected by R1.

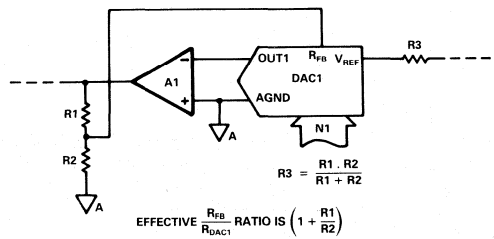

Figure 8. Recommended Way to Implement Fixed R_{FB}/R_{DAC1} Ratio

A second method of realizing a repeatable gain is t replace the feedback resistor R_{FB} in Figure 3 with a pro grammable feedback resistor. A second CMOS DA placed in the feedback loop behaves as a programmabl resistance and allows the circuit gain to be digitally pro grammed. Such a circuit is shown in Figure 9.

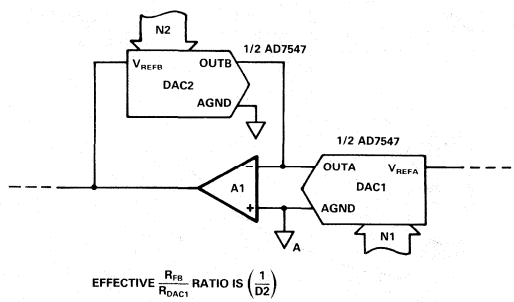

Figure 9. Recommended Way to Implement Programmable R_{FB}/R_{DAC1} Ratio

ıe effective feedback resistance has a value equal to

$$R_{FB} = \frac{R_{DAC2}}{D2}$$

here R_{DAC2} is the DAC ladder resistance of the feed-ıck DAC and D2 is the fractional representation of the gital input data to the feedback DAC. The V_{RATIO} ex-·ession of Equation 5 now becomes

$$V_{RATIO} = \frac{V_{OUT}}{V_{REF}} = \frac{1}{1+D1\left(\frac{R_{DAC2}}{R_{DAC1} \cdot D2}\right)}$$

here D1 and R_{DAC1} relate to the primary DAC. If both ACs are on the same monolithic chip, then they will ıve identical ladder impedances. This simplifies the RATIO expression to

$$V_{RATIO} = \frac{1}{1+D1\left(\frac{1}{D2}\right)}$$

oth DACs will generally have the same resolution, n, so ıat the V_{RATIO} expression is shown finally to be propor-ɔnal to the codes in the two DACs

$$V_{RATIO} = \frac{1}{1+N1\left(\frac{1}{N2}\right)}$$

here N1 and N2 are the codes applied to DAC1 and AC2 respectively. Thus the R_{FB}/R_{DAC1} ratio is now ɔlely determined by the code N2 in the feedback DAC.

To find this code, the denominators of Equation 4 and the final V_{RATIO} expression above are equated for the condition when the spindle speed is at its maximum. Using the previous example with $R_{MAX}/R_{MIN} = 2.5$, $X = 1$ and $N1 = 4095_{10}$ the correct code for DAC2 is found to be $N2 = 2730_{10}$. This code is loaded into DAC2 and is not changed thereafter. Suitable monolithic dual 12-bit DACs for this application are the AD7537/AD7547 family from Analog Devices which are available in skinny 24-pin DIP packages or in standard surface mount packages. The AD7537 with its 8+4 bit loading structure is suitable for 8-bit bus systems, while for 16-bit bus systems the AD7547, with its parallel 12-bit loading structure, is more suitable.

The programmable solution to realizing the R_{FB}/R_{DAC1} ratio raises the intriguing possibility of building a universal drive capable of reading disks of different diameters. The disks would need to have the same R_{MIN} dimension and the same maximum spindle speed f_S max. Recorded on the innermost sectors of such disks would be R_{MAX} information and the number of programmable speeds required. The drive would read this information at power-on at the standard maximum spindle speed and set the circuit parameters accordingly, a universal drive.

8

ANALOG DEVICES

AN-325
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

12-Bit Analog I/O Port Uses AD7549 and 8051 Microcomputer

by Mike Curtin

INTRODUCTION

In the process control industry, many slowly varying analog signals need to be measured and controlled. Examples of these are temperature, pressure, position, etc. This application note describes the design of an analog Input/Output port based on the AD7549 dual DAC and the 8051 microcomputer which will meet this requirement. The I/O port measures analog signals and also provides an analog output voltage which may be used in various system control loops (e.g., control voltage on a hydraulic servo valve).

HARDWARE DESCRIPTION

The two main components in the I/O port are the AD7549 dual DAC and the 8051 microcomputer. The AD7549 is a dual 12-bit DAC. Figure 1 illustrates the block diagram. For further information consult the AD7549 Data Sheet, available from Analog Devices. One DAC of the AD7549 provides the analog output voltage while the other performs the D/A function in a Successive Approximation ADC. The 8051 provides the interfacing signals to load DACB with the data for the analog output. It also performs the successive approximation routine with DACA to measure the analog input, A_{IN}.

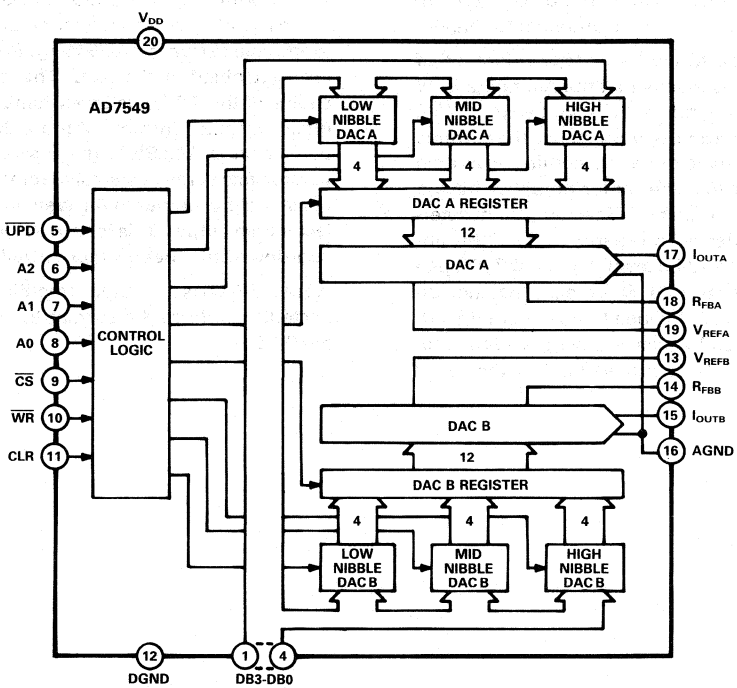

Figure 1. AD7549 Functional Block Diagram

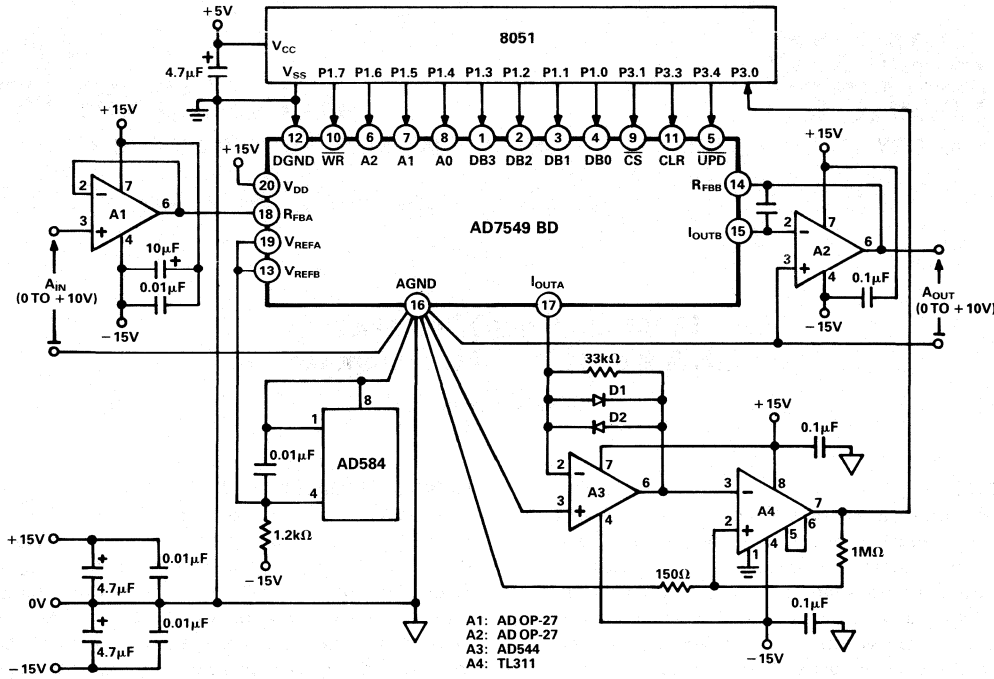

Figure 2. Analog I/O Port Circuit Diagram

The I/O port circuit diagram of Figure 2 also incorporates the AD584 voltage reference, an input buffer for A_{IN}, comparator circuit for the ADC and output amplifier for A_{OUT}.

The AD584 is connected to give a voltage reference of −10V to both DACs. This means that the A_{IN} range is 0V to +10V and the A_{OUT} range is also 0V to +10V. A1 (AD OP-27) buffers the analog input. Since the currents in the DAC are being switched at something less than 100kHz, it is important that the buffer amplifier have a high loop gain at this frequency. The loop gain determines the output impedance and so the ability of the amplifier to maintain A_{IN} at a 12-bit level. The AD OP-27 open loop gain at 100kHz is 40dB (100). This means that the buffer output impedance is 0.7Ω (Open Loop Output Impedance divided by loop gain) which is driving a 10kΩ load. Error due to this is 0.007% and is well below the 12-bit level.

The ADC comparator section has two stages. A3 is an AD544 with a gain resistor of 33kΩ and back-to-back diodes (HSCH-1001) to reduce settling time constraints. It is possible to trim the AD544 V_{OS} (Input Offset Voltage) to take out offset in the ADC. The deciding factor in the choice of the AD544 is its extremely low input bias and offset currents. Currents of more than 100nA can contribute errors of 1/2LSB to the answer. A4 is a TL311 comparator with some hysteresis. Note that there is no pull-up resistor at the output of A4 even though it is an open collector type output. This is because P3.0, when configured as an input, provides an internal pull-up resistor.

A2 (AD OP-27) is the output amplifier for DACB and is connected to perform the current-to-voltage conversion for the DAC.

```
0100 AOUT:  CLR   P3.3      Disable the CLR line
102         SETB  P3.4      Set UPD high
104         CLR   P3.1      Bring CS low and select the AD7549
106         MOV   R2, #04   Load Register R2 with 04. This
                            will be used to set the device
                            address lines
108         ACALL ADDRS     Register address loaded to AD7549
10A         MOV   R0, #21
10C         MOV   A, @R0    Load data (W) into DAC B low nibble
10D         ACALL DATA      register
10F         INC   R2        Set up next register address and
110         ACALL ADDRS     load to the AD7549
112         MOV   A, @R0
113         SWAP  A         Load data (V) into DACB mid nibble
114         ACALL DATA      register
116         INC   R2        Set up next register address and
117         ACALL ADDRS     load to the AD7549
119         DEC   R0
11A         MOV   A, @R0    Load data (U) into DACB high nibble
11B         ACALL DATA      register
11D         INC   R2        Set up DACB Register address and load
11E         ACALL ADDRS     to the AD7549
120         CLR   P1.7      Strobe the WR line to load data
122         SETB  P1.7      (UVW) to DAC B
124         SETB  P3.1      Bring CS high to deselect AD7549
126         RET             Return to main program

0140 ADDRS: MOV   A,R2      This subroutine takes the register
141         SWAP  A         address in R2, formats it and loads
142         ORL   A, #80    it out to the AD7549. It then returns
144         MOV   P1,A      to AOUT routine
146         RET

0150 DATA:  ANL   A, #OF    This subroutine transfers the data
152         ORL   P1,A      nibble in the lower half of A to
154         CLR   P1.7      the AD7549 data bus and strobes the
156         SETB  P1.7      WR line low to load the appropriate
158         RET             register, before returning to AOUT

0200 AIN:   MOV   R0, #21
202         SETB  P3.0      Set up port line P3.0 as an input
204         CLR   P3.3      Disable the CLR line
206         SETB  P3.4      Set UPD high
208         CLR   P3.1      Bring CS low and select the AD7549
20A         MOV   P1, #00   Load DAC A low nibble register with
20D         SETB  P1.7      all 0's
20F         SETB  P1.4
211         CLR   P1.7      Load DAC A mid nibble register with
213         SETB  P1.7      all 0's
215         CLR   P1.4
217         SETB  P1.5
219         CLR   P1.7      Load DAC A high nibble register with
21B         SETB  P1.7      all 0's
21D         SETB  P1.4
21F         CLR   P1.7
221         SETB  P1.7      DAC A is now loaded with all 0's
223         MOV   P1, #0A0  Set up address for DAC A high nibble

226         SETB  P1.3      Set DAC 1st (5th,9th) MSB
228         CLR   P1.7      Bring WR low
22A         CLR   P3.4
22C         SETB  P3.4      Strobe UPD pin
22E         SETB  P1.7      Bring WR high
230         JNB   P3.0,290  Test comparator A4 output. If 0,
                            jump to routine to clear bit
233         SETB  P1.2      Set DAC 2nd (6th,10th) MSB
235         CLR   P1.7      Bring WR low
237         CLR   P3.4
239         SETB  P3.4      Strobe UPD pin
23B         SETB  P1.7      Bring WR back high
23D         JNB   P3.0,294  Test comparator A4 output. If 0,
                            jump to routine to clear bit
240         SETB  P1.1      Set DAC 3rd (7th,11th) MSB
242         CLR   P1.7      Bring WR low
244         CLR   P3.4
246         SETB  P3.4      Strobe UPD pin
248         SETB  P1.7      Bring WR back high
24A         JNB   P3.0,298  Test A4 output. If 0, jump to
                            routine to clear bit
24D         SETB  P1.0      Set DAC 4th (8th,12th) MSB
24F         CLR   P1.7      Bring WR low
251         CLR   P3.4
253         SETB  P3.4      Strobe UPD pin
255         SETB  P1.7      Bring WR back high
257         JNB   P3.0,29C  Test A4 output. If 0, jump to
                            routine to clear bit
25A         CLR   P1.7      Bring WR low
25C         CLR   P3.4
25E         SETB  P3.4      Strobe UPD pin
260         SETB  P1.7      Bring WR back high
262         INC   R0
263         MOV   A, #0F
265         ANL   A,P1      Read nibble from port, and place
267         MOV   @R0,A     result in address specified by R0
268         JNB   P1.5,271
26B         CLR   P1.5      Set up address for DAC A mid nibble
26D         SETB  P1.4
26F         AJMP  226
271         JNB   P1.4,278
274         CLR   P1,4      Set up address for DAC A low nibble
276         AJMP  226
278         MOV   A,23      Take the 2 least significant nibbles and
27A         SWAP  A         combine them in data memory location
27B         ORL   A,24      23
27D         MOV   23,A
27F         SETB  P3.1      Bring CS high to deselect AD7549
281         RET             Return to main program

290         CLR   P1.3      These instructions clear the AD7549
292         AJMP  233       data bits and return to the
294         CLR   P1.2      successive approximation routine
296         AJMP  240
298         CLR   P1.1
29A         AJMP  24D
29C         CLR   P1.0
29E         AJMP  25A
```

Table I. 8051 Routines for Programming the I/O Port

SOFTWARE DESCRIPTION

Table I lists the complete analog I/O port software subroutines. The I/O port should be considered as part of a larger control system. Whenever an analog input is to be measured or an analog output to be delivered, the program jumps to the appropriate subroutine. These subroutines are A_{OUT} and A_{IN}. A_{OUT} takes the 12 bits of data UVW contained in data memory locations 20, 21 and loads this data to DACB. So, the output of A2 (A_{OUT}), is the analog value of the digital word, UVW.

A_{IN} is the successive approximation routine for converting the analog signal, A_{IN}, into its digital value XYZ and placing the result in data memory locations 22 and 23. The routine initializes port outputs, clears the contents of DACA and then proceeds into the successive approximation routine proper. In this, it makes extensive use of the bit-handling instructions available on the 8051. Individual port bits may be cleared or set with a single instruction (CLR or SETB). Also, a single instruction (JNB) can test the state of port bits and jump to another location depending on the bit state. The use of these instructions simplifies the complete successive approximation routine. Table II shows the memory organization for the Analog I/O Port.

DATA MEMORY	CONTENTS
020	0U
021	VW
022	0X
023	YZ

Table II. Memory Organization for Analog I/O Port

PERFORMANCE

To perform the Analog Output function, the user jumps from the main program to subroutine A_{OUT}. This occupies 55 bytes of memory and has an execution time of 74μs. This means that within 74μs of jumping to the subroutine, A_{OUT} has reached the analog equivalent of UVW (Data in 20, 21). Since A2 (AD OP-27) has excellent input offset voltage characteristics, A_{OUT} specifications will match those on the AD7549 data sheet. When the AD7549KN (BD, TD) is used, the integral linearity error is 1/2LSB. Differential linearity is less than 1LSB, ensuring guaranteed monotonicity over temperature. Full-scale error (gain error) is 3LSBs max which corresponds to 0.073% F.S.R.

A_{IN} is contained in 145 bytes of memory, and has an execution time which varies between 140μs and 180μs, depending on the value of A_{IN}. This is the ADC conversion time. For the slowly varying signals which occur in process control systems, this speed is adequate. However, for a user who needs to measure higher frequency signals, increased bandwidth can be obtained by using an AD585 Sample/Hold amplifier instead of the input buffer, A1. This allows sampling of signals up to 2.7kHz.

Figure 3 shows how the output of A3 (top trace) varies during the A/D conversion cycle. At the end of the cycle, DACA is loaded with the digital value of A_{IN}, causing the currents flowing into A3 inverting terminal to balance and bringing the voltage at the output of A3 to zero. The bottom trace is the start of conversion signal. This particular conversion is completed in approximately 160μs.

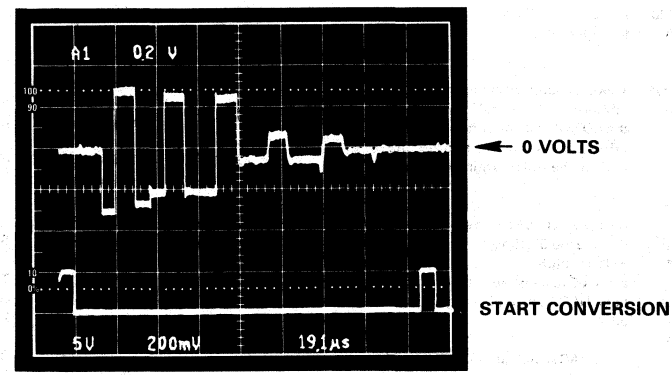

Figure 3. Voltage at A3 Output During the Conversion Cycle

ANALOG DEVICES

AN-326
APPLICATION NOTE

NE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Interfacing the AD7549 to the MCS-48 and MCS-51 Microcomputer Families

by Mike Curtin

ne AD7549 is the industry's first monolithic dual 12-bit /A converter. Not only that, but because of the data load- g structure, it is housed in a space-saving 20-pin 0.3" ide package. The data for each DAC is contained in three bit nibbles. Table I shows how the digital control inputs ad the internal registers. The dedicated CLR line is espe- ally useful in system reset and autocalibration routines hile $\overline{UPD}$ allows both DACs to be updated together from eir respective input registers. The AD7549 data sheet ives further information on the device and it is recom- ended that the user consult this where necessary.

he applications of the AD7549 are many and varied. They clude volume and balance control in audio systems, utoff frequency and Q factor control in programmable fil- ers and frequency and amplitude control in programma- le oscillators. As well as these audio based applications, the AD7549 can be used in general to replace two discrete DACs with a dual DAC in a small package. One particular example of this is when the DACs are used as analog output ports in microcomputer based control systems. This application note examines such systems and deals with the hardware and software interfacing needs.

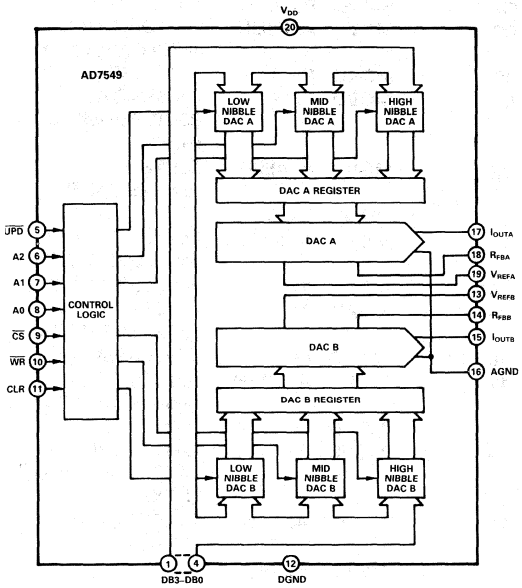

Figure 1. AD7549 Functional Diagram

CLR	$\overline{UPD}$	$\overline{CS}$	$\overline{WR}$	A2	A1	A0	FUNCTION
0	X	X	1	X	X	X	No data transfer.
0	1	1	X	X	X	X	No data transfer.
1	X	X	X	X	X	X	All registers cleared.
0	1	0	⌐	0	0	0	DAC A LOW NIBBLE REGISTER loaded from Data Bus.
0	1	0	⌐	0	0	1	DAC A MID NIBBLE REGISTER loaded from Data Bus.
0	1	0	⌐	0	1	0	DAC A HIGH NIBBLE REGISTER loaded from Data Bus.
0	1	0	⌐	0	1	1	DAC A Register loaded from Input Registers.
0	1	0	⌐	1	0	0	DAC B LOW NIBBLE REGISTER loaded from Data Bus.
0	1	0	⌐	1	0	1	DAC B MID NIBBLE REGISTER loaded from Data Bus.
0	1	0	⌐	1	1	0	DAC B HIGH NIBBLE REGISTER loaded from Data Bus.
0	1	0	⌐	1	1	1	DAC B Register loaded from Input Registers.
0	0	1	⌐	X	X	X	DAC A, DAC B Registers updated simultaneously from Input Registers.

NOTE: X = Don't Care

Table I. AD7549 Truth Table

Microcomputers differ from microprocessors in that they contain not only an on-chip CPU but also various combinations of RAM, ROM, I/O ports, timers, etc., which enables the user to configure a complete control system with a minimum amount of hardware. It is, therefore, hardly surprising that the number of microcomputers used in instrumentation and control is ever increasing. Two of the most popular microcomputer families are the MCS-48 and the newer MCS-51, both of which are manufactured by Intel. Interfacing the AD7549 to these systems requires little external hardware and keeps system component count to a minimum.

MCS-48 INTERFACE

The various members of the MCS-48 family differ only in their internal memory capabilities. In the following discussion the 8048 microcomputer is representative of the

complete family. The 8048 contains a set of specific instructions for communicating with the 8243 I/O expander. These are the MOVD, ANLD, ORLD instructions.

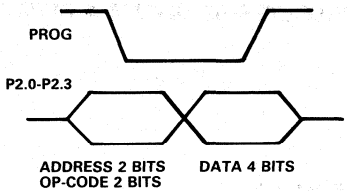

Figure 2. MCS-48 Input/Output Expander Timing

MOVD reads or writes data directly between the microcomputer's accumulator and one of the four output ports on the 8243 I/O expander. ORLD and ANLD perform the logical OR and logical AND functions between the accumulator and an output port, placing the result in the output port. Figure 2 shows the timing for the Output Expander instructions. All communication with the 8048 occurs over the lower half of Port 2 (P2.0–P2.3). Timing control is provided by PROG. Each transfer consists of two 4-bit nibbles. The first contains the op code and port address and the second contains the actual 4-bits of data. A high to low transition of the PROG line indicates that address is present, while a low to high indicates the presence of data. Because of the similarity in structure between the AD7549 and the 8243 it is possible to use the 8243 instructions to program the AD7549.

P2.0	P2.1	8243 Port	AD7549 Register
0	0	P4	LOW NIBBLE REGISTER
0	1	P5	MID NIBBLE REGISTER
1	0	P6	HIGH NIBBLE REGISTER
1	1	P7	DAC REGISTER

Table II. AD7549 Register Selection

Table II shows how the registers are selected. As an example, the instruction "MOVD P4, A" sets P2.0 and P2.1 to zeros in the first nibble of the instruction (Figure 2). The negative going edge of PROG latches 0 0 to A0 and A1 causing the LOW NIBBLE REGISTER to be loaded with the bottom 4 bits of A. Note that the AD7549 A2 pin is not shown in Table II. This effectively selects DAC A or DAC B and would be set appropriately prior to the MOVD instruction. Figure 3 shows the interface circuit for the MCS-48 to the AD7549.

Data inputs DB3–DB0 connect directly to P2.3–P2.0. Address lines A0, A1 are obtained by latching P2.0, P2.1 on the negative-going edge of PROG. All other AD7549 control lines are port outputs on the 8048. Table III is a listing of OUTLD, a routine for loading the dual DAC. To load DAC A with UVW (HEX) the routine initially selects the device, and clears all registers. Then, using the MOVD instruction, each nibble register is loaded with its correct 4 bits of data. Finally the DAC A Register is loaded with the contents of the 3 nibble registers. The routine then goes through a similar procedure to load DAC B with XYZ (HEX). When the high speed (11MHz) version of the 8048 is used, it takes 57μs to load both DACs.

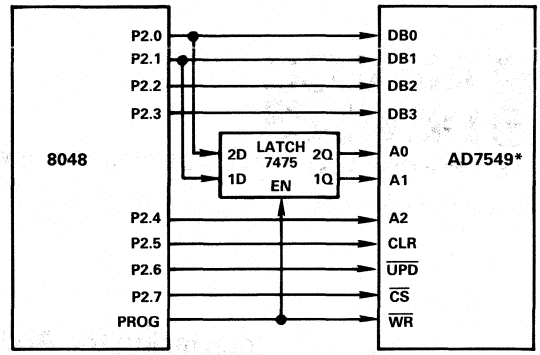

*LINEAR CIRCUITRY OMITTED FOR CLARITY

Figure 3. Interface Circuit for the MCS-48 to AD7549

```
0000 OUTLD:  MOV   A, #60H    Bring CS low to select
02           OUTL  P2, A      device and set CLR
                              high to clear all DAC
                              registers
03           MOV   A, #40H    Disable CLR, Set
05           OUTL  P2, A      A2 = 0 to select DAC A
06           MOV   A, #0VWH   Load DAC A LOW
                              NIBBLE REGISTER
08           MOVD  P4, A      with W (HEX)
09           SWAP  A          Load DAC A MID
                              NIBBLE REGISTER
0A           MOVD  P5, A      with V (HEX)
0B           MOV   A, #0UH    Load DAC A HIGH
                              NIBBLE REGISTER
0D           MOVD  P6, A      with U (HEX)
0E           MOVD  P7, A      Load DAC A
                              REGISTER with UVW
                              (HEX)
0F           MOV   A, #50H    Set A2 = 1
                              to select DAC B
11           OUTL  P2, A
12           MOV   A, #0YZH   Load DAC B LOW
                              NIBBLE REGISTER
14           MOVD  P4, A      with Z (HEX)
15           SWAP  A          Load DAC B MID
                              NIBBLE REGISTER
16           MOVD  P5, A      with Y (HEX)
17           MOV   A, #0XH    Load DAC B HIGH
                              NIBBLE REGISTER
19           MOVD  P6, A      with X
1A           MOVD  P7, A      Load DAC B
                              REGISTER with XYZ
                              (HEX)
1B           MOV   A, #C0H    Set CS high to
                              deselect DAC
1D           OUTL  P2, A
```

Table III. 804B Program Listing

MCS-51 INTERFACE

In the following text, the term "8051" is used to generically refer to all members of the MCS-51 family. Like the MCS-48 family, the various members differ mainly in their internal memory capabilities (the 8052/8032 also has an extra 16-bit timer/counter). In comparison to the 8048, the 8051 offers more speed, more input-output pins, and more memory in addition to a full duplex serial I/O port. There is no facility on the 8051 to expand I/O ports in the same manner as the 8048. However, it does have features which make an AD7549 interface very simple. In particular, the 8051's Boolean processing instructions provide direct bit handling (i.e., single bit addressing and updating). Individual bits can be set, cleared or complemented with the 2-byte instructions SETB, CLR or CPL. In addition, bits can be moved to and from the carry flag with the MOV instruction, and logical ANL and ORL functions can be performed between the carry and the addressed bit. With this bit handling facility, it is possible to interface to the AD7549 without any external hardware.

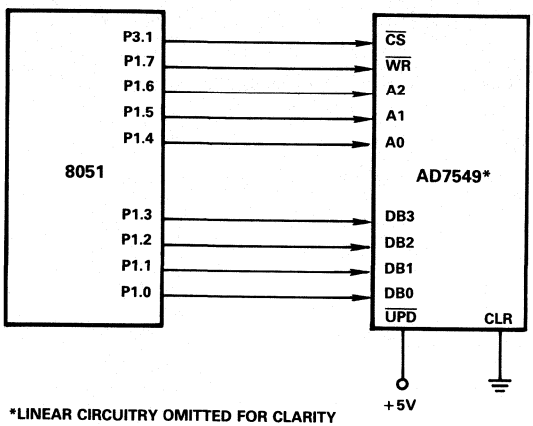

*LINEAR CIRCUITRY OMITTED FOR CLARITY

Figure 4. MCS-51 to AD7549 Interface

Figure 4 shows the interface circuit and Table IV lists the software routine DACLOAD which loads both DACs. The data to be loaded into the DACs is stored in data memory locations 20 to 23. Table V shows how this is organized. DACLOAD programs DAC A with UVW (HEX) and DAC B with XYZ (HEX). P3.1 is the AD7549 $\overline{CS}$. This is activated at the beginning of the routine and brought high at the end. Port lines P1.4–P1.6 control the register addressing. The routine initially sets up the appropriate nibble register address in 8051 register R2. Then, the ADDRS subroutine formats this and loads it to the device address lines. The DATA subroutine places the correct data on the device lines DB3–DB0 and loads the appropriate register by strobing $\overline{WR}$ (P1.7) low. The next AD7549 nibble register address is then set up by incrementing R2. When the three nibble registers are loaded the 12-bits of data are transferred to the DAC regiser by setting up the DAC register address and strobing $\overline{WR}$ low.

```
0100  DACLOAD:  CLR   P3.1     Bring CS low to select device
102              MOV   R2, #00  Load 1st register address into R2
104              LCAL  ADDRS    Register address loaded to DAC
107              MOV   R0, #21
109              MOV   A, @R0
10A              LCAL  DATA     Data (W) loaded into register
10D              INC   R2       Next register address setup
10E              LCAL  ADDRS    Register address loaded to DAC
111              MOV   A, @R0
112              SWAP  A
113              LCAL  DATA     Data (V) loaded into register
116              INC   R2       Next register address setup
117              LCAL  ADDRS    Register address loaded to DAC
11A              DEC   R0
11B              MOV   A, @R0
11C              LCAL  DATA     Data (U) loaded into register
11F              INC   R2       DAC A register address setup
120              LCAL  ADDRS    DAC A register address loaded to DAC
123              CLR   P1.7     Bring WR low
125              SETB  P1.7     Bring WR high. Data (UVW) has
                                now been loaded to DAC A
127              INC   R2       Next register address setup
128              LCAL  ADDRS    Register address loaded to DAC
12B              INC   R0
12C              INC   R0
12D              INC   R0
12E              MOV   A @ R0
12F              LCAL  DATA     Data (Z) loaded into register
132              INC   R2       Next register address setup
133              LCAL  ADDRS    Register address loaded to DAC
136              MOV   A, @R0
137              SWAP  A
138              LCAL  DATA     Data (Y) loaded into register
13B              INC   R2       Next register address setup
13C              LCAL  ADDRS    Register address loaded to DAC
13F              DEC   R0
140              MOV   A @ R0
141              LCAL  DATA     Data (X) loaded into register
144              INC   R2       DAC B register address setup
145              LCAL  ADDRS    DAC B register address loaded to DAC
148              CLR   P1.7     Bring WR low
14A              SETB  P1.7     Bring WR high. Data (XYZ) has
                                now been loaded to DAC B
14C              SETB  P3.1     Bring CS high to deselect device
0200  ADDRS:    MOV   A, R2    This subroutine takes the
201              SWAP  A        appropriate register address held in
202              ORL   A, #80   R2, formats it and loads it out to
204              MOV   P1, A    the device.
206              RET
0210  DATA:     ANL   A, #0F   This subroutine transfers the data
212              ORL   P1, A    nibble in A to the device data bus
214              CLR   P1.7     and strobes the WR line low to load
216              SETB  P1.7     the appropriate register.
218              RET
```

Table IV. 8051 Program Listing

When the 12MHz version of the 8051 is used, the program loads the AD7549 in 140μs. The amount of program memory used is 93 bytes. Note that the device architecture allows interfacing using only nine of the 8051's thirty-two input/output lines, so that in large control systems, the DACs do not unduly overburden the 8051 I/O capability. The system is easily expanded by taking extra I/O lines and using them as $\overline{CS}$ inputs for successive devices.

Memory Location	Contents
20	0U
21	VW
22	0X
23	YZ

Table V. Data to be Loaded to the AD7549

ANALOG DEVICES

AN-327
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

DAC ICs: How Many Bits is Enough?

by Robert Adams

The DACs designed for use in digital audio playback circuitry are available in a diverse choice of architectures, and each have their own performance tradeoffs.

At the heart of every digital audio playback system lies the single-most critical component for high-fidelity audio: the digital-to-analog converter (DAC). These converters handle the delicate task of translating the 16-bit binary words encoded on the disc or tape into corresponding analog signals worthy of amplification and, ultimately, of the human ear.

Magnavox's first CD player (circa 1983) employed dual 14-bit converters, and by 1989 many models heralded 20-bit converters. Today however, the words "bitstream" and "MASH" — displayed prominently on the face of many CD players and other digital reproduction devices — refer to a 1-bit converter architecture that is revolutionizing the digital audio industry and attracting the attention of many.

Just how many bits are necessary? Given that CDs are encoded in a 16-bit PCM format, why are 18- and 20-bit converters showing up?

This article examines the DACs designed and optimized for use in digital audio playback circuitry. It compares the diverse choice of architectures available for the job and discusses the design and performance tradeoffs of each.

Basic theoretical concepts

Although you don't have to be an expert to understand audio D/A conversion, there are two basic concepts that are fundamental to the understanding of this topic: *sampling* and *quantization*.

Sampling is the act of taking an analog signal waveform, which has a value at every instant of time, and sampling it at regular intervals. This supplies us with information about the value of the waveform only at the sampling instants. The Nyquist Theorem states that, as long as the original analog signal contains no frequencies higher than half of the sampling frequency, the signal can be re-constructed completely from the sampled values.

For most digital audio applications, we have two common sampling frequencies: 44.1kHz (the consumer standard) and 48kHz (the professional standard). Both of these frequencies are a little above twice the highest frequency of interest (20kHz).

If we imagine that we have a "perfect" sampler, free of real-world non-idealities, we can say that the signal can be recovered completely, without any added noise or distortion. We assume that the output of this perfect sampler is either a noiseless analog voltage (held on a capacitor, for example) or a digital word with infinite precision. To store the sampled values for later playback, one must use a finite number of bits to represent the sampled value. This is where *quantization* comes into play.

By using a finite digital word-length to represent the sampled value, we have added an uncertainty to the sampled value, one which can no longer be completely recovered on playback. Quantization can be thought of as a *many-to-one mapping* (a continuous range of input values results in a single output value). For consumer digital audio equipment, the sampled values are stored in a 16-bit format, with a theoretical dynamic range of 98.1dB. (It's more complicated than just 6dB/bit×16 bits=96dB, and it is commonly misunderstood.)

Errors introduced by quantization are heavily dependent on the signal being quantized. High-amplitude signals with a complex spectrum (pop music, for example) generally cause quantization errors that are like white noise. Very soft signals that are not much larger than one quantization level can be severely distorted unless dither noise (a random noise uncorrelated with the signal) is added to the signal before quantization. The importance of using dither to de-correlate the quantization error from the input signal has received wide attention, and most manufacturers now use it where appropriate in their designs.

A typical DAC circuit

Over the years, many modifications in D/A circuitry have improved performance and reduced cost. For example, early CD players generally used a single, parallel-input DAC, alternated between both right and left channels. They employed two output sample-and-hold circuits — one per channel — to reduce the signal-dependent "glitch" that they produced when switching from one sample to the next, and to divide the output samples into separate left and right signals. The sample-and-hold amplifier's distortion and the inter-channel time delay drew criticism from audiophiles, who were convinced that such degradations were audible. Modern designs usually employ one DAC for each channel, and special design techniques are used in the DAC itself to reduce the glitch to the point where de-glitching circuitry is no longer necessary.

Figure 1 shows a typical D/A circuit, which might be found inside a modern CD player. The most important IC is, of course, the D/A chip itself. In this case, our AD1860, an 18-bit DAC, is being used. Like most modern DACs, this device uses a serial interface to reduce pin count. This serial interface consists of three signals: serial data input, bit clock input and word clock input. This serial bit stream is converted internally to a parallel word, where it drives the individual current-steering switches. One advantage of using an internal latch is that the timing of when each bit actually switches can be carefully controlled, which reduces the signal-dependent glitch problem found in older designs.

The other key function is the digital oversampling filter chip, which takes in data at 44.1kHz (or 48kHz) and produces an output at a higher rate (two, four, eight

Adams is with the Semiconductor Division of Analog Devices, Wilmington, MA.

Reprinted from Sound and Video Contractor — February 20, 1991

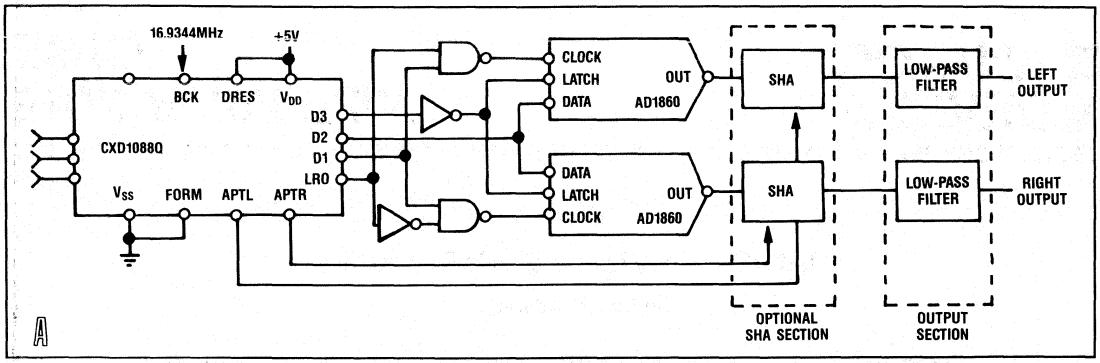

Figure 1. This typical D/A circuit uses an 18-bit DAC, which uses a serial interface to reduce pin count.

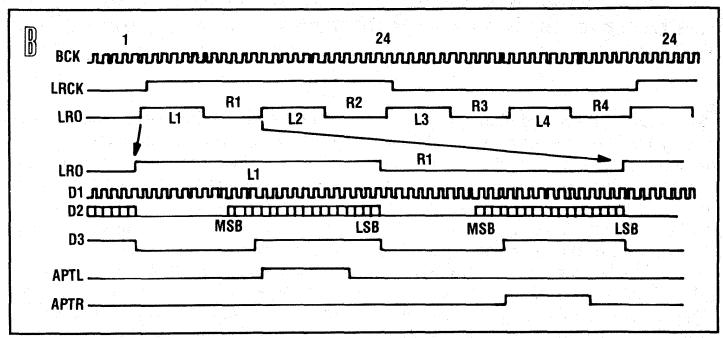

or even 16 times the input sample rate). Next we will look into the theory behind DAC filters, and why analog brick-wall filters have almost universally given way to digital interpolation filters.

D/A reconstruction filters: Is digital better?

On the playback side, sampling theory predicts that there will be "images" of the original signal around multiples of the sampling frequency. For example, a 1kHz signal that was originally sampled at 44.1kHz will have components at 43.1kHz, 45.1kHz, 87.2kHz, 89.2kHz, etc., as shown in Figure 2.

If we want perfect reconstruction of the original continuous waveform, we must filter out or otherwise eliminate these unwanted look-alike or image components before they are fed to the rest of our audio system. However, this filtering is, strictly speaking, not necessary. All of the unwanted components are in the ultrasonic range, and the only reason to remove them is to ensure that the analog circuitry in a typical power amp or pre-amp will not become non-linear in the presence of large ultra-high-frequency signals.

Figure 3 shows two methods for removing these unwanted image signals. The first method is to follow the DAC with a *brick-wall* analog filter, which has flat response over audio frequencies and sharply cuts off just above 20kHz. The second method is to do most of this filtering digitally and use a simple analog filter at the DAC output.

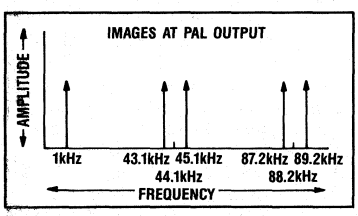

Figure 2. Sampling theory predicts that there will be "images" of the original signal around multiples of the sampling frequency. Perfect reconstruction of the original continuous waveform requires that these unwanted look-alike or image components be eliminated.

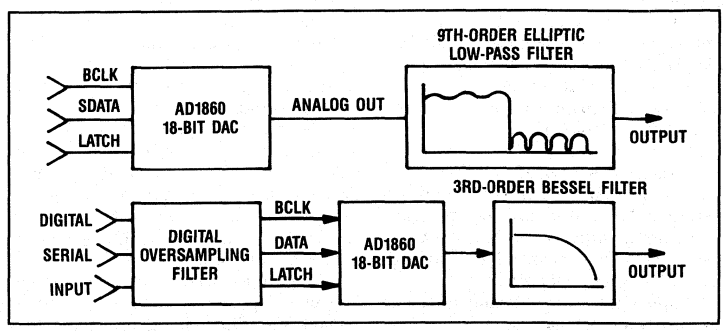

Figure 3. There are two methods for removing unwanted image signals. Follow the DAC with a brick-wall analog filter, which has flat response over audio frequencies and sharply cuts off just above 20kHz, or do most of this filtering digitally and use a simple analog filter at the DAC output.

The analog reconstruction filter: Fast disappearing

In the early days of digital audio, the reconstruction filter was typically a complex, analog, active (or passive) filter. This filter had to have very flat frequency response over the audio band and then drop precipitously just above 20kHz, with an ultimate attenuation of at least 80dB.

These filters became known as *brick-wall* filters, and they required exceptionally high-tolerance components to meet audio standards. These filters also introduced large phase shifts near the edge of the audio band. Whether this phase shift introduced audible artifacts became a controversial subject, and early detractors of digital audio often blamed the so-called "gritty" sound of digital recordings on the high phase shift introduced by this filter, when in fact they were probably hearing the effects of poor low-level linearity in the ADC or DAC (more on this later).

In any event, the combination of high cost, high power consumption and large phase shifts drove the development of digital oversampling filters that have effectively replaced brick-wall analog filters in most applications.

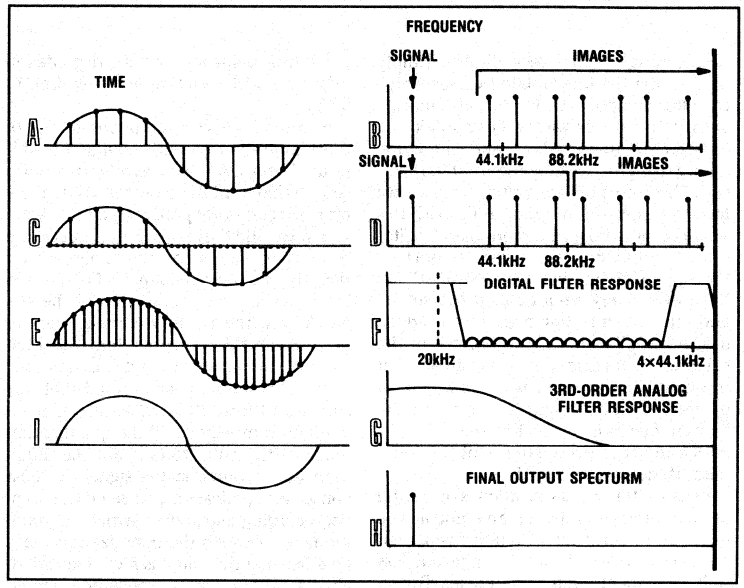

Figure 4. *For digital oversampling, we start with a sampled signal at 44.1kHz (Figure 4a), which has images in the frequency domain (Figure 4b). The sample rate of the digital signal is increased by zero-valued samples inserted between the existing samples (Figure 4c), resulting in the spectrum shown in Figure 4d. This zero-stuffed signal is applied to a digital low-pass filter. The filter response for a 4× interpolator is shown in Figure 4f, and the smoothed time-domain signal is shown in Figure 4e. A simple analog filter finishes the job by removing the high-frequency image, as shown in Figures 4g and 4h, resulting in the reconstructed waveform shown in Figure 4i.*

To the rescue ...

Digital oversampling filters take the burden off the analog filters by filtering out most of the undesired information (images) in the digital domain before the signal is applied to the DAC. There must still be an analog output filter, but it can be a simple third-order design with minimal phase shift.

Figure 4 shows how oversampling filters work in both the time and frequency domains. We start with a sampled signal at 44.1kHz (Figure 4a), which has images in the frequency domain (Figure 4b). The next step in the process is to increase the sample rate of the digital signal by inserting zero-valued samples between the existing samples (Figure 4c), resulting in the spectrum shown in Figure 4d.

This may seem like a trivial operation, because we have not really changed the spectrum of the signal. But it is, in fact, an important conceptual step; the signal in Figure 4a is sampled at 44.1kHz and has images around multiples of 44.1kHz, but the signal in Figure 4c is sampled at four times 44.1kHz and has images around four times 44.1kHz.

Even though the spectrum looks the same, we now regard the 44.1kHz images as the actual signal itself, which just happens to have repeated regions at multiples of 44.1kHz. This conceptual step is necessary because any digital filter that we build will have its own image response at multiples of the input sample rate. Because we want our filter response to attenuate frequencies beyond 44.1kHz, it has to run at a faster sampling rate than 44.1kHz.

The third step in the process is to apply this zero-stuffed signal to a digital low-pass filter. The filter response for a 4× interpolator is shown in Figure 4f, and the smoothed time-domain signal is shown in Figure 4e. Note that the filter response has an image around the 4× oversampling rate. Therefore, an image of the baseband signal appears at the DAC output around this frequency and must be removed by the external analog filter. But the digital filter has done most of the hard work, and a simple analog filter can now finish the job by removing this high-frequency image, as shown in Figures 4g and 4h, resulting in the reconstructed waveform shown in Figure 4i.

One advantage of using a finite impulse response (FIR) digital filter is that the phase response can be made linear. (See "The Performance Aspects of Digital Oversampling" and "The Saori Signal Processor" in the October 1989 and September 1990 issues, respectively for more on FIR filters.) This solves the problem (if it is a problem) of too much phase shift in the audio band, although we should bear in mind that the digital filter is every bit as much of a "brick-wall" filter as the old analog filter was. It does not provide relief for those who can hear beyond 20kHz.

Another more tangible advantage of the digital approach is that we have now replaced a fussy analog circuit requiring high-precision components and lots of circuitry with a monolithic IC that is inexpensive and whose performance does not change significantly from one unit to the next. It is this combination of economic incentive and improved performance that has lead virtually all manufacturers to incorporate digital oversampling filters into their designs.

Today it is common to find digital audio devices that use 18- or even 20-bit DACs. The number of bits used in the DAC is often prominently displayed on the front panel, with more expensive units using more bits. The technically savvy person will often ask the obvious question: If there are only 16 bits in the original signal, why would you need an 18- or 20-bit converter?

There are several answers to this question, some of which are technical in nature and some that have more to do with marketing issues. (How does a manufacturer differentiate a low-end CD player from a high-end unit if they both sound the same?) The technical answer is that the digital oversampling filter produces a digital output with much more than 16 bits. This does *not* mean that the digital filter has somehow magically removed noise from the signal or increased the dy-

namic range of the recording. It simply means that the interpolated values which the filter computes do not usually coincide with 16-bit integer values. They usually fall in between 16-bit values, and, therefore, a longer digital word is required to represent the interpolated value. So now we face the question of what to do with this long digital word (often more than 30 bits).

If we simply drop the lower bits and feed the upper 16 bits into a 16-bit DAC, we have effectively re-quantized the signal, and some quantization noise will be added to the signal that was not present in the original digital signal. If we postulate that the original recording was made with a perfect 16-bit ADC with a theoretical 16-bit dynamic range of 98.1dB, then we have added another noise source of the same magnitude to the signal.

Because this re-quantization is occurring at the interpolated rate and not at the original 44.1kHz rate, its noise is spread out over a wider frequency range, and less of it appears at audio frequencies. For an 8:1 interpolation ratio, the total noise will be about 1.5dB higher than theoretical 16-bit performance, resulting in a noise floor of 96.6dB. This assumes that the truncation noise is not correlated with the signal, which may be an optimistic assumption for low-level signals. A more pessimistic value for the noise degradation might be 3dB, resulting in a S/N ratio of 95.1dB.

If instead we save the upper 18 bits of the digital filter output signal and feed it to an 18-bit DAC, we are adding a noise source that is 12dB below the original 16-bit theoretic noise produced by the ADC. When we add these two noise sources, the result is only a fraction of a decibel higher than the theoretic 16-bit noise floor, resulting in almost no degradation of the signal. This is the primary technical reason that 18- and 20-bit converters are often used in digital audio playback systems.

But one should bear in mind that this slim 1.5dB (or 3dB worst case) noise advantage is mostly lost if the quantization noise of the converters is not the dominant noise source in the signal. In most consumer applications, noise on the original recording and in the playback environment will exceed the quantization noise by a fair margin, and it is highly questionable whether this improvement could realistically be audible.

Another reason for using converters of higher resolution is that the distortion performance (linearity) of higher-resolution converters is typically better than that of lower-resolution products. For example, a 16-bit DAC, such as our AD1856, typically exhibits full-scale distortion of 0.002%; a 20-bit part (such as the AD1862J) has distortion numbers of less than 0.0012%. Yet another advantage of high-resolution converters is that the thermal noise performance (noise with no digital codes changing) is typically better than the noise performance of their 16-bit counterparts.

The requirements for the professional audio market are quite different from those of the consumer market. When you are recording live, uncompressed material, there is a legitimate need for very-high-resolution ADCs and DACs. The dynamic range of a closely miked percussion instrument using a low-noise condenser microphone can exceed 120dB without much trouble!

The professional recording industry is in the midst of a major push toward 18- and even 20-bit converters to meet these needs. At present, few product offerings provide this level of performance, and most of these products are expensive hybrids. But in the near future, the same performance will be available in inexpensive monolithic ICs that use a radically different design approach known as sigma-delta conversion.

In Part II, we'll cover design approaches for D/A conversion and discuss the specifications in product literature.

ANALOG DEVICES

AN-206
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Analog Panning Circuit Provides Almost Constant Output Power

by John Wynne

In audio recording and playback it is often required to split or "pan" a single signal source into a two-channel signal for stereo effects. To locate the signal source as desired in the sound stage, the overall signal power is maintained constant while the relative levels in the derived channels are adjusted. This application note describes a circuit which limits variations in the total output power to ±0.1dB. It uses two CMOS Multiplying D/A Converters to control the signal levels in the derived channels. CMOS DACs are ideal in this application because of their low distortion. The DACs function as resistive attenuators using thin-film resistors which have low noise and very low voltage coefficient. Converters which are especially suitable for this application are Dual-DACs which contain two CMOS DACs on a monolithic substrate. These are available from Analog Devices with 8-bit resolution (AD7528) or 12-bit resolution (AD7537/47/49).

The simplest and most obvious circuit for panning is shown in Figure 1. The digital data N_B fed to DAC B is the 2's complement of the data N_A fed to DAC A. For n-bit converters the digital input codes, N_A & N_B, can be represented by fractional values, D_A & D_B respectively, where $D_A = N_A/2^n$ and $D_B = N_B/2^n$, with N_A and N_B in decimal format. Because of the 2's complement arrangement between the DACs, the all 0's code or full mute condition is not allowed – in theory at least. The relationship between the two fractional representations, D_A & D_B, is given as:

$$D_A + D_B = 1 \qquad (1)$$

The output voltage expressions for the two channels in Figure 1 are as follows:

$$V_{OUTA} = -D_A \cdot V_{IN} \qquad (2)$$

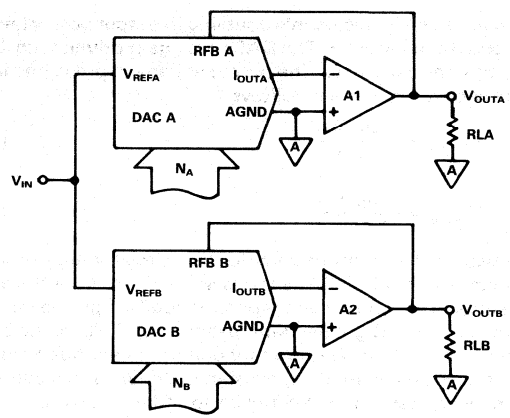

Figure 1. Simple Power Splitter Circuit

$$V_{OUTB} = -D_B \cdot V_{IN} \qquad (3)$$

The performance of a 12-bit system is shown in Figure 2 where the total output power level is plotted versus N_A. The 0dB output power level used as the reference level in Figure 2 is the total output power available when $D_A = D_B = 0.5$, the balanced condition. With this simple panning circuit, the total output power level at either extreme of the allowable input code range has increased by 3dB (a doubling of the power) over the power output level at the balanced condition. Load impedances, RLA & RLB, are assumed equal for both channels.

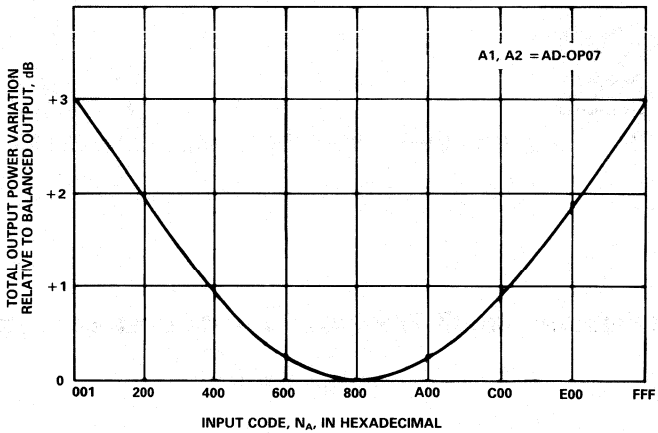

Figure 2. Response of Figure 1 Using 12-Bit DACs

A circuit which avoids this doubling in output power level is shown in Figure 3. The DACs are again driven with 2's complementary data. The output voltage expressions for the two channels are as follows:

$$V_{OUTA} = \frac{-D_A \cdot V_{IN}}{1 + D_A} \quad (4)$$

$$V_{OUTB} = \frac{-D_B \cdot V_{IN}}{1 + D_B} \quad (5)$$

Power splitter performance for this circuit when implemented with 12-bit DACs is shown in Figure 4. At the extremes of the input code range the total output power is now only 0.5dB greater with respect to its level at the balanced condition. This amount of output variation would be acceptable in many applications. Certain applications, however, may demand better performance than this.

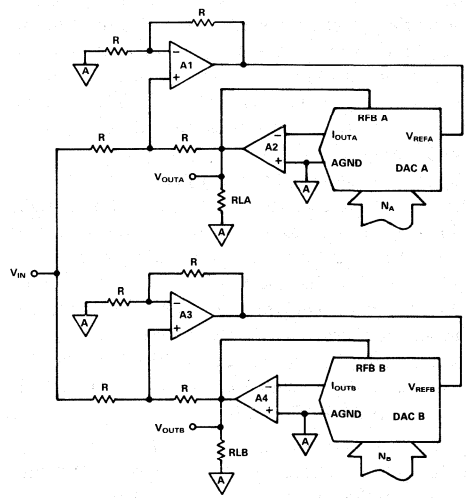

Figure 3. Improved Power Splitter Circuit

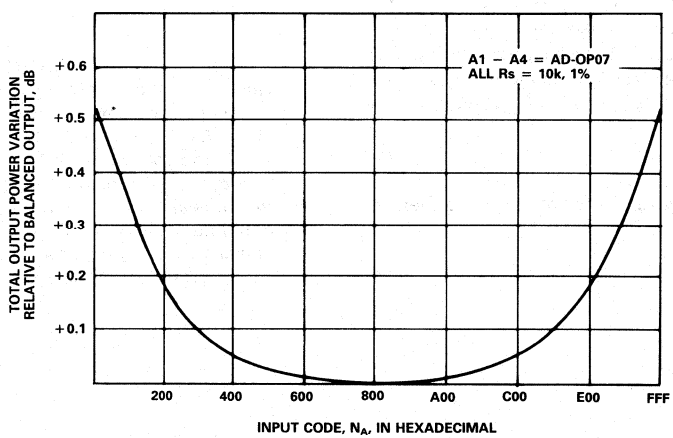

Figure 4. Response of Figure 3 Using 12-Bit DACs

HIGHER PERFORMANCE POWER SPLITTER

A higher performance power splitter circuit can be built by adding some gain around the CMOS DACs of Figure 3. The gain factor required is equal to $\sqrt{2}$.

In order to provide a DAC output voltage which is $\sqrt{2}$ times greater than normal, the effective value of the feedback resistor must be made equal to $\sqrt{2}$ times the DAC ladder impedance R_{DAC}. Reference 1 outlines how additional gain can be added to the standard configuration without requiring a large gain adjustment range or compromising the circuit's temperature coefficient. The circuit configuration of Figure 5 provides the additional $\sqrt{2}$ gain factor. Resistors R1, R2 and R3 should have similar temperature coefficients, but they need not match the temperature coefficient of the DAC. The three resistors are precision (0.1%) metal film resistors with standard EIA/MIL values. When both DAC circuits of Figure 3 are changed to include the gain resistors of Figure 5, the output voltage expressions for the two channels become:

$$V_{OUTA} = \frac{-\sqrt{2} \cdot D_A \cdot V_{IN}}{1 + \sqrt{2} \cdot D_A} \quad (6)$$

$$V_{OUTB} = \frac{-\sqrt{2} \cdot D_B \cdot V_{IN}}{1 + \sqrt{2} \cdot D_B} \quad (7)$$

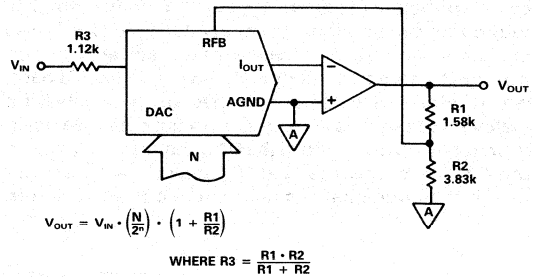

Figure 5. Additional Resistors to Provide $\sqrt{2}$ Gain Factor

The performance of the revised circuit is shown in Figure 6. The total power output at either extreme of the input code range is equal to the total power output at the balanced condition. The output power level remains constant within a ±0.1dB error band. With the exception of R1, R2 and R3, all resistors are metal film, 10kΩ, 1% tolerance.

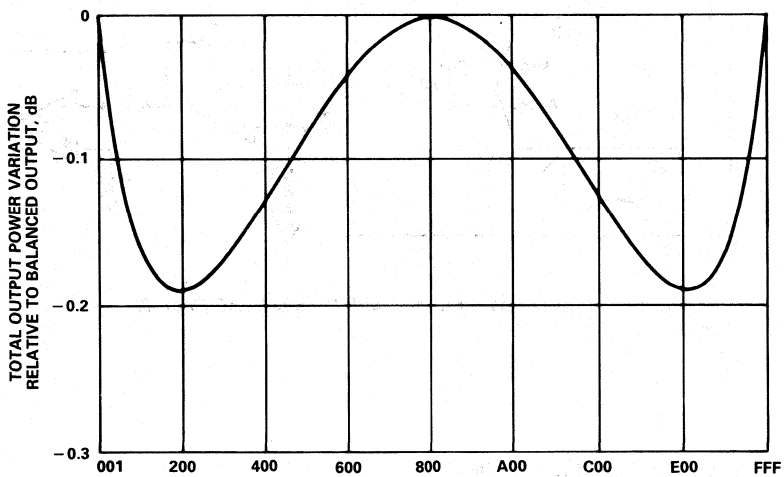

Figure 6. Response of Figure 3 Using $\sqrt{2}$ Gain Factor

STAND-ALONE APPLICATIONS

Many applications which are microprocessors-based will have no difficulty in generating the 2's complement data required for the previous power splitter circuits. It is also possible in non-microprocessor-based systems to perform the 2's complement operation with dedicated hardware. If either of these approaches are unsuitable, for whatever reason, it is still possible to build a power splitter circuit by rewiring one of the DAC channels and driving both DACs with identical data. Figure 7 shows the rewiring and the additional circuitry required to achieve this.

The output expressions for the two channels are now:

$$V_{OUTA} = \frac{-D \cdot V_{IN}}{1 + D} \quad (8)$$

$$V_{OUTB} = \frac{-(1-D) \cdot V_{IN}}{2 - D} \quad (9)$$

The performance of Figure 7 is identical to that of Figure 3. Additionally, the circuit performance can be upgraded in a similar fashion as previously described for Figure 5.

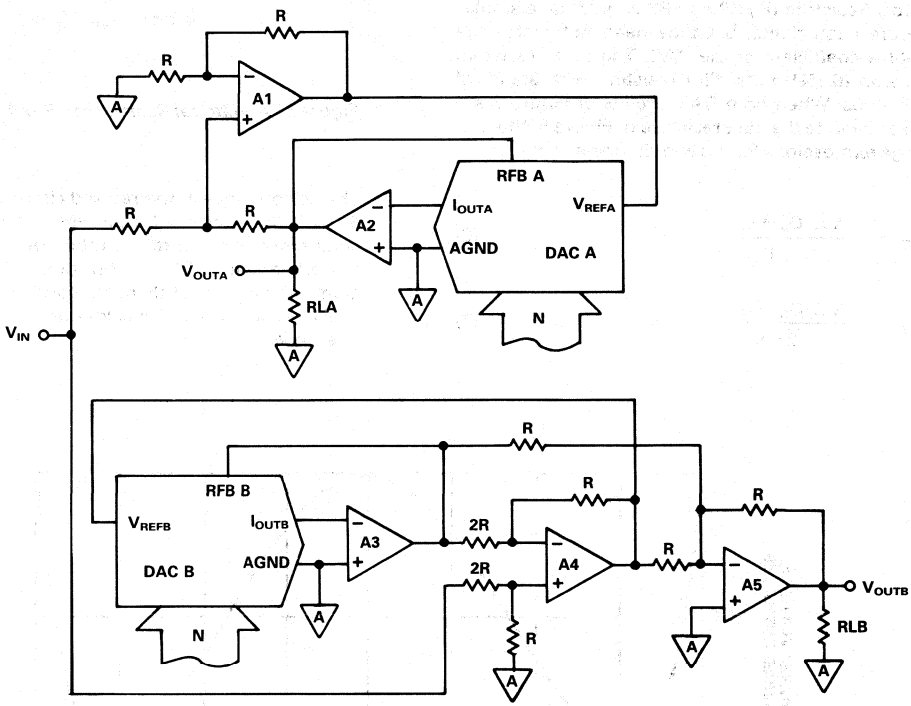

Figure 7. Power Splitter Circuit with Identical DAC Data, N

REFERENCES

1. Brokaw P., "Input Resistor Stabilizes MDACs Gain." EDN, January 7, 1981, pp. 210-211.

ANALOG DEVICES

AN-207
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Interfacing Two 16-Bit AD1856 (AD1851) Audio DACs with the Philips SAA7220 Digital Filter

by Kevin Greene

INTRODUCTION

The AD1856 is a complete 16-bit DAC used primarily for digital audio applications. The AD1851 is a lower noise, second generation version of the AD1856. Each device provides a voltage output amplifier, 16-bit DAC, 16-bit serial-to-parallel input register, and voltage reference. The AD1856 is specified to operate with ±5 V to ±12 V supplies and achieves a *maximum* of 0.0025% total harmonic distortion (THD). The AD1851 operates with ±5 V supplies and has a *maximum* of 0.004% total harmonic distortion + noise (THD + N), and a signal-to-noise ratio (SNR) of at least 107 dB. Their performance and ease of use make the AD1856/AD1851 popular choices for 16-bit audio designs.

The Philips SAA7220 is a 4× oversampling digital interpolating filter. Some listeners prefer the sound quality of this filter over other digital filters on the market. The SAA7220 features attenuation correction in the audio passband. Other digital filters typically do not incorporate this feature making it difficult, if not impossible, to achieve a flat frequency response using a Butterworth or Bessel filter. These features make the SAA7220 a popular choice for system designers.

The combination of the AD1856/AD1851 with the SAA7220 yields a system with very low THD + N and an excellent SNR. If a Bessel (also known as Thompson) filter is used, the transient response will be exceptional as well. Figure 1 shows a block diagram of the system. Unfortunately, the output interface of the SAA7220 is compatible with I²S format, while the AD1856/AD1851 has a standard 3-line interface.

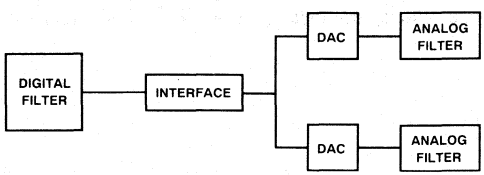

Figure 1. Block Diagram of System

INTERFACE

The interface needs to take the SAA7220 output, which conforms to I²S format (see Figure 2a), and modify it to match the input requirements of the AD1856/AD1851.

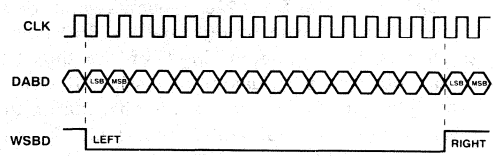

Figure 2a. Output of SAA7220

Three signals must be present for proper operation of the AD1856/AD1851: Data, Clock, and Latch Enable. These signals are shown in Figure 2b. Bringing Latch Enable low after the least significant bit of any word latches the previous 16 bits and starts the conversion.

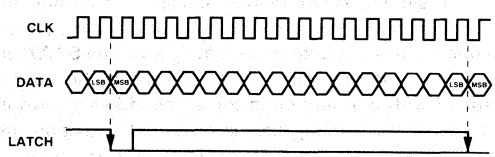

Figure 2b. Signal Requirements for AD1856 (AD1851)

In order for the AD1856/AD1851 to be compatible with the SAA7220, the interface must delay the WSBD line by one clock cycle, gate the clock on the left channel, and simultaneously latch both DACs to eliminate any phase shifts between channels. A quad NAND gate and a dual D flip-flop accomplish this as shown in Figure 3.

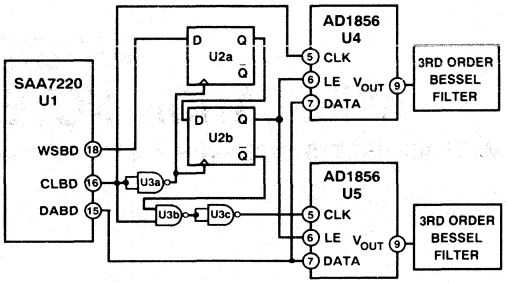

Figure 3. Interface of the SAA7220 and AD1856 (AD1851)

CIRCUIT DESCRIPTION

Data bits are valid on the negative edge of the clock. The D flip-flops are positive-edge triggered, therefore one NAND gate (U3a) is used to invert the clock to the flip-flops. Referring to Figure 2a, WSBD goes low one clock cycle prior to the LSB. Figure 2b shows LATCH going low after the LSB. The flip-flops delay WSBD by one clock cycle, similar to a shift register, to correct this mismatch. The other two NAND gates (U3b and U3c) are used to form an AND gate to gate the clock to the left channel DAC (U5). The SAA7220 transmits serial data, left channel first, followed by right data. WSBD is low during left data transmission. Using $\overline{Q}$ of U2b, the AND gate is "on" during the left data transmission. Therefore both DACs clock in the 16 left channel bits. During the transmission of right data, WSBD goes high ($\overline{Q}$ goes low), turning "off" the clock to U5. Only the right channel (U4) clocks in the right data. After the LSB of right data is clocked into U4, both DACs can be simultaneously latched from the Q output of U2b.

PASSBAND ATTENUATION

All ripple-free low-pass filters have attenuation in the passband similar to Figure 4a. Typically, the corner frequency is approximately 30 kHz for many audio circuits. This results in signal attenuation at 20 kHz. The SAA7220 provides 1.0 dB of digital pre-emphasis in the 20 kHz passband with attenuation in the stopband as shown in Figure 4b. If the analog filter is designed such that the magnitude is 1 dB down at 20 kHz, the combined response will theoretically be flat in the passband as shown in Figure 4c.

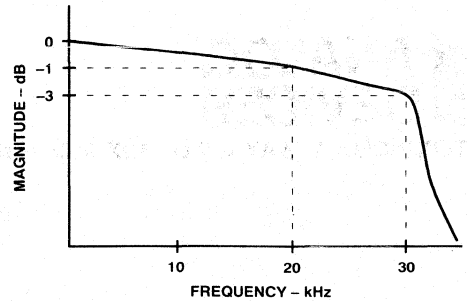

Figure 4a. Analog Filter Response

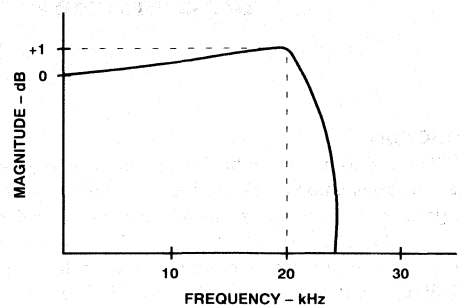

Figure 4b. Digital Filter Response

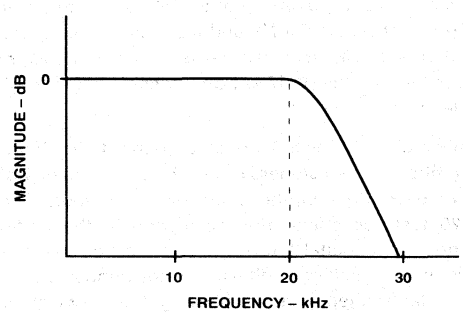

Figure 4c. Combined Response

Normally, Chebyshev or Elliptic filters, with very low ripple, would be required to achieve a flat response. However, these filters inherently have a poor transient response. A Bessel filter, which is optimized for linear phase, has essentially no overshoot or ringing associated with its step response. Additionally, the impulse response lacks oscillatory behavior. The Elliptic, Chebyshev, and even Butterworth filters all suffer from these shortcomings. The improvement in the transient response reduces the distortion of the overall circuit.

Figure 5 shows one such Bessel filter architecture. It's a 3rd order Sallen and Key filter with the −1 dB point of about 20 kHz. An advantage of this filter is that it uses only one op amp which reduces the component cost. A disadvantage is that the sensitivity (the change in magnitude to variations in the component values) increases dramatically as the order of the filter increases. The sensitivity of a third order filter is usually acceptable. Additionally, the capacitor values need to be modified to match available values. This causes the filter to roll-off slightly more than 1 dB at 20 kHz. Another design possibility is to use two op amps, cascading a two-pole filter and a one-pole filter. This would reduce the sensitivity of the filter and possibly achieve a more accurate −1 dB point.

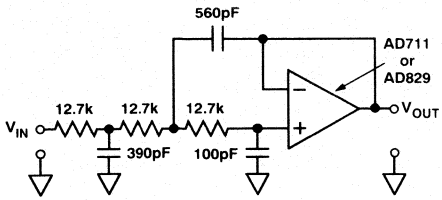

Figure 5. 3rd Order Bessel Filter

RESULTS

By combining the SAA7220's digital filter, the AD1856 (or even better, the AD1851), and a Bessel filter, an excellent system can be achieved. The result is a circuit with THD + N and SNR specifications comparable to the highest quality systems using Butterworth or Bessel filters, but additionally offering the flat frequency response of a system using a Chebyshev or Elliptic filter.

ACKNOWLEDGMENTS

To Bill Thompson for his input. To Steve Ruscak for his input and help throughout the project.

REFERENCES

1. L.P. Huelsman and P.E. Allen, *Introduction to the Theory and Design of Active Filters,* McGraw-Hill Publication Number: ISBN 0-07-030854-3.
2. Arthur B. Williams, *Electronic Filter Design Handbook,* McGraw-Hill Publication Number: ISBN 0-07-070430-9.

ANALOG DEVICES

AN-328
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Generate 4 Channels of Analog Output Using AD7542 12-Bit D/A Converters and Control the Lot with Only Two Wires

by John Wynne

This application note describes how a Universal Asynchronous Receiver Transmitter (UART) can be configured to drive four AD7542 12-bit D/A converters. The AD7542 is ideal in this application because of its data loading structure. The functional diagram of Figure 1 shows the AD7542 to consist of three 4-bit data registers, a 12-bit DAC register, address decoding logic and a 12-bit multiplying DAC. Data is loaded into the data registers in three 4-bit nibbles, and subsequently transferred to the 12-bit DAC register. These data move operations are controlled by the address inputs A0 and A1. Each character received by the UART contains one 4-bit nibble of data and four bits of associated control information.

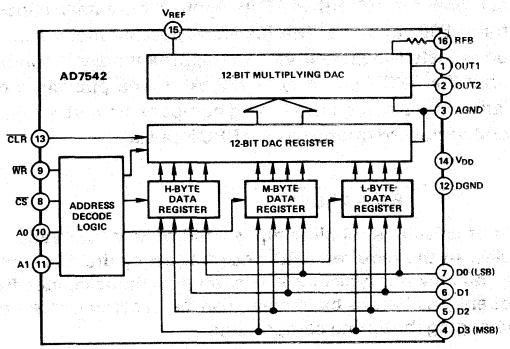

Figure 1. AD7542 Functional Block Diagram

A block diagram of the circuit is shown in Figure 2. All data and control signals pass over one opto-coupler allowing very simple wiring between a host processor and remote analog output channels. The power consumption of the complete circuit is kept very low due to the fact that every active component in the circuit is constructed with either CMOS or low power Bifet technology.

SERIAL DATA FORMAT

The format of the serial data used by the receiving IM6402 is shown in Figure 3. This consists of one start bit, eight data bits (LSB first) and two stop bits. Transmission may be either from another UART or directly from a microcomputer. Many modern microcomputers have dedicated serial input/output ports which can be used for transmission. However, small variations between microcomputers do exist and so it will be necessary to tailor the receiving UART to allow different formats to be accommodated.

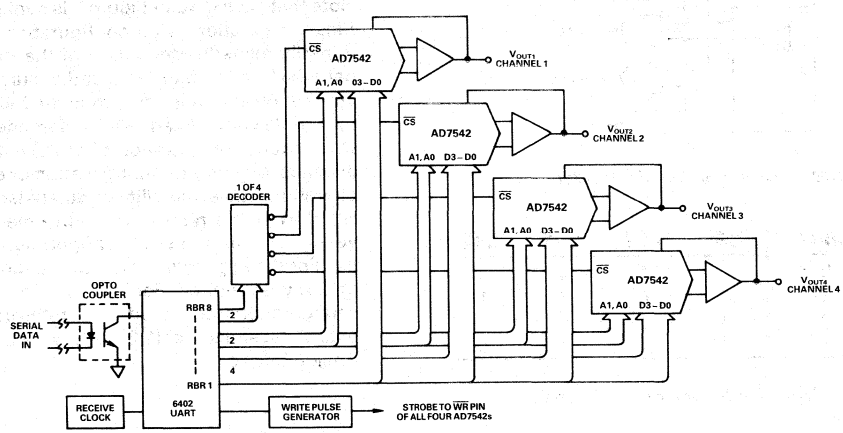

Figure 2. Block Diagram of Circuit

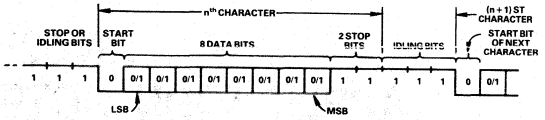

Figure 3. Serial Data Format

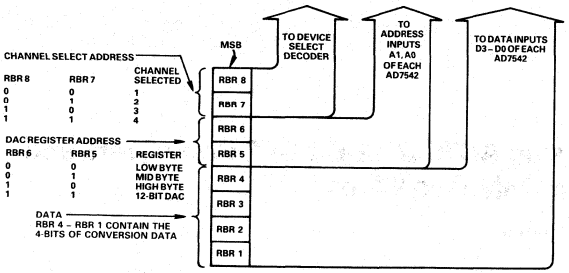

Figure 4. Breakout of the 8-Bit Character Received by the UART

As a character is received at the UART, the start and stop bits are stripped off and the data is transferred to the Receiver Buffer Register. A high level on the Data Received (DR) pin indicates that a new character has been received. Each 8-bit wide character contains 4-bits of conversion data and 4-bits of control data. A detailed functional breakout of each bit in the 8-bit wide character is shown in Figure 4.

The two most significant bits go to the decoder to select one of the four DACs to be updated. Character bits RBR6 and RBR5 go to address inputs A1, A0 respectively of each DAC to select one of the internal registers for a data move operation. The four least significant bits (RBR4 – RBR1) contain the data to be loaded to one of the 4-bit nibble registers.

The simple 1-of-4 decoder is shown in Figure 5. The truth table is shown in Table I. The four decoded outputs go to the Chip Select (CS) inputs of the DACs. This signal must be low for a data move operation to occur.

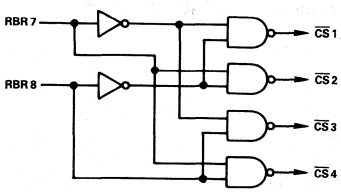

Figure 5. Channel Select Decoder

RBR8	RBR7	$\overline{CS1}$	$\overline{CS2}$	$\overline{CS3}$	$\overline{CS4}$
0	0	0	1	1	1
0	1	1	0	1	1
1	0	1	1	0	1
1	1	1	1	1	0

Table I. Truth Table for Figure 5 Decoder

The final signal which the DACs require is a Write ($\overline{WR}$) pulse tocontrol DAC loading. This signal is common to all four DACs and is generated by means of the Data Received (DR) output (see Figure 6). As previously mentioned, a high level on the DR pin indicates that a new character has been received by the UART. The combination of the CMOS inverters and the resistor-capacitor networks of Figure 6 produces a delayed low-going signal on the Data Received Reset ($\overline{DRR}$) input. This clears the DR output low which in turn drives the $\overline{DRR}$ input high again. This low-going pulse on the $\overline{DRR}$ input (equal to the width of the delay) is used as the $\overline{WR}$ pulse for the DACs.

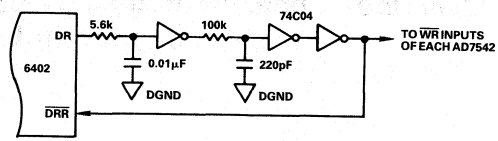

Figure 6. Write Pulse Generator Circuit

The complete circuit is shown in Figure 7. The optocoupler used is a high speed HCPL-2530 manufactured by Hewlett Packard. The Receiver Register Clock (RRC) is generated by the ICM7555, a low power CMOS version of the popular 555 timer. Note that this clock rate must be 16 times the receiver baud rate. The values of the timing components, RA, RB and C in Figure 7, have been chosen for a 9600 baud rate. The 10k ohms potentiometer should be adjusted (adjusted value is approximately 9k ohms) until the RRC frequency is 153.6kHz. The published relationship between the timing components and the generated clock frequency on the RRC input is

$$f_{CLK} = \frac{1.46}{(RA + 2RB)C}$$

In practice this relationship is valid only at low frequencies. At the required RRC clock frequency the measured frequency was typically 30% lower than the calculated frequency. Different baud rates can be accommodated by changing the timing components.

DAC OUTPUT CIRCUITRY

Note that each DAC of Figure 7 is configured for Unipolar Binary Operation. This configuration provides straight (binarily-related) attenuation of the input reference signal, which can be either a.c. or d.c., current or voltage. The application examples shown in the following section use this mode with a fixed −10V reference to the four DACs. Other modes of operation of an AD7542 are possible, for instance, 4-quadrant multiplication or even single supply operation. The availability or otherwise of suitable power supplies at the remote location determines the output voltage ranges and DAC configurations which are possible. For the application examples a remote power supply of ±15V is required. For information on these other DAC configurations the reader is referred to the AD7542 data sheet and to Ref [1,2], all of which are available from Analog Devices.

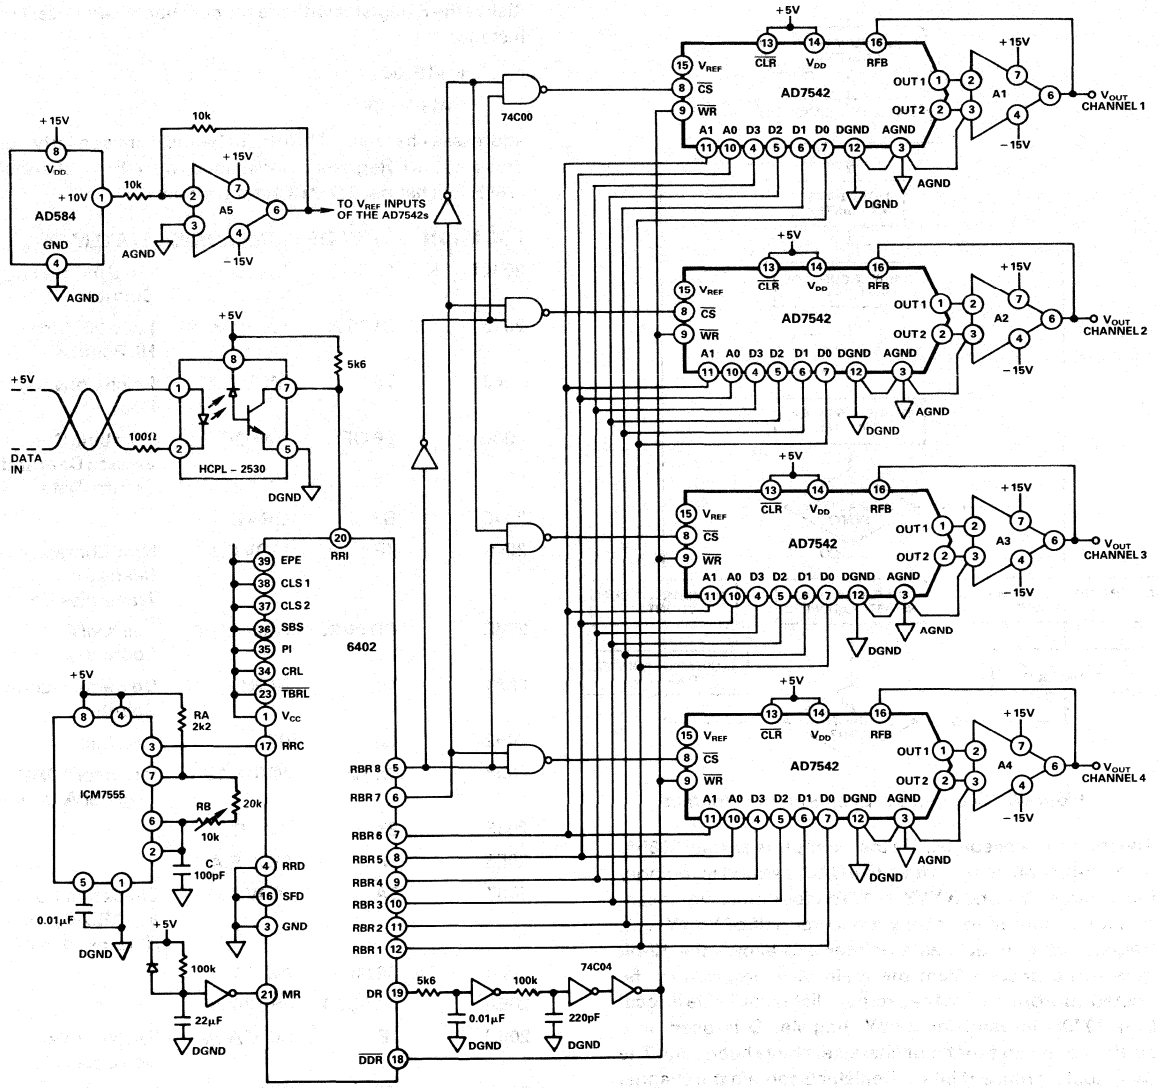

Figure 7. Detailed Circuit Diagram

THE SOFTWARE INTERFACE

Before data transmission can occur the 4-bits of conversion data must be packed with the proper 4-bits of control data to form an 8-bit character. This character is then passed to the transmitter for serial transmission. If the transmitter is a UART similar to the receiving one, the 8-bit character need only be written into the UARTs transmit register and the UART commanded to transmit. Start and stop bits are automatically added by the UART itself. If the transmitter is based upon a microprocessor then some extra software is required to supervise the serial transmission. An example is given which loads 12-bits of data to one of the four DACs and immediately updates the analog output. This involves transmitting four characters sequentially to the addressed DAC, i.e., the 3 data nibbles and the load DAC register command. The software is written in the form of nested subroutines; the first, called PACK, assembles the 8-bit character prior to transmission; the second, called XMIT, transmits this character with the necessary start and stop bits; the final one, called DELAY, determines the baud rate. The sequence of events in the PACK subroutine are shown in Figure 8 with the actual program listing shown in Table II. The other routines are discussed in Appendix A1. The PACK routine of Figure 8 ensures four characters are sequentially sent. It can be considerably simplified to allow single 4-bit data loading to any AD7542 register and also to defer updating of a DAC's analog output until required.

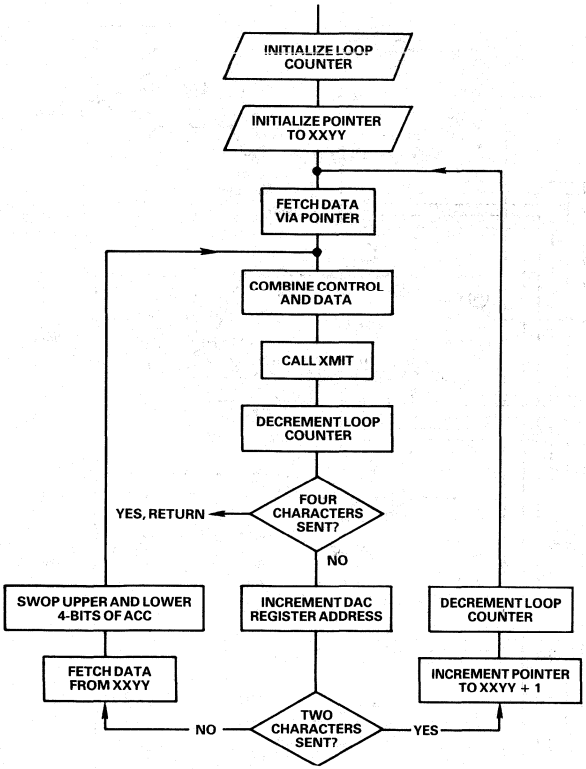

Figure 8. Sequence of Events in Pack Routine

The microprocessor used in the example is an Intel 8085A. New 12-bit data to be written to a DAC is stored in memory locations XXYY and XXYY + 1. The four most significant bits are assumed to occupy the lower half of XXYY + 1. Register pair HL acts as a pointer and allows the 12-bit data to be passed from the main user program to the nested subroutines. In the program listing of Table II, location 2070H is used for XXYY. Register D is used as a counter to ensure that four characters have been sent. The two most significant bits of Register E contain the channel address. The channel address is loaded into Register E before the program enters the nested subroutines and does not change during the program execution.

The next two most significant bits of Register E contain the AD7542 register address. Accumulator A is used to combine the conversion data loaded from XXYY (or XXYY + 1) with the control data of Register E. The combined word is then moved to Register C. Register C is used to pass the 8-bit character from the PACK subroutine to the XMIT subroutine. To completely update the four channels requires 16 characters to be transmitted. Table III shows the truth table for the control portion of the 8-bit character. The two most significant bits of the character determine which channel is to be updated. Before calling the PACK subroutine to transmit new data, the user must first initialize the E Register with the proper channel address. For instance

```
MVI E, 00
CALL PACK
```

addresses channel 1. The other channels are similarly addressed; the E Register being loaded with 40H for channel 2, 80H for channel 3 and COH for channel 4.

LOCATION	OP CODE	MNEMONIC	STATEMENT
2044 PACK	16 05	MVI D, 05	Load 05 into Loop Counter
2046	21 70 20	LXI HL 2070	Load 2070 into HL Pointer
2049	7E	MOV A,M	Fetch Conversion Data
204A	E6 OF	ANI OF	Combine Conversion Data and Control Data
204C	B3	ORA E	
2040	4F	MOV C,A	New Character Ready for Transmission
204E	CD 00 20	CALL 2000	Call XMIT Routine
2051	15	DCR D	Decrement Loop Counter
2052	C8	RZ	Finished?
2053	3E 10	MVI A, 10	Increment DAC Register Address
2055	83	ADD E	
2056	5F	MOV E,A	
2057	7A	MOV A,D	Check that LOW- and MID-Bytes Have Been Sent
2058	C6 00	ADI 00	
205A	EA 65 20	JPE 2065	
205D	7E	MOV A,M	Fetch Conversion Data
205E	OF	RRC	Swop Upper and Lower 4-Bits of ACC
205F	OF	RRC	
2060	OF	RRC	
2061	OF	RRC	
2062	C3 4A 20	JMP 204A	
2065	23	INX H	Increment Pointer
2066	15	DCR D	
2067	C3 49 20	JMP 2049	

Table II. Program Listing for PACK Routine

CHARACTER TO BE TRANSMITTED								HEX EQUIVALENT	REGISTER ADDRESSED
CONTROL				DATA					
MSB							LSB		
0	0	0	0	X	X	X	X	0X	DAC 1 Low-Byte Register
0	0	0	1	X	X	X	X	1X	DAC 1 Mid-Byte Register
0	0	1	0	X	X	X	X	2X	DAC 1 High-Byte Register
0	0	1	1	X	X	X	X	3X	DAC 1 DAC Register
0	1	0	0	X	X	X	X	4X	DAC 2 Low-Byte Register
0	1	0	1	X	X	X	X	5X	DAC 2 Mid-Byte Register
0	1	1	0	X	X	X	X	6X	DAC 2 High-Byte Register
0	1	1	1	X	X	X	X	7X	DAC 2 DAC Register
1	0	0	0	X	X	X	X	8X	DAC 3 Low-Byte Register
1	0	0	1	X	X	X	X	9X	DAC 3 Mid-Byte Register
1	0	1	0	X	X	X	X	AX	DAC 3 High-Byte Register
1	0	1	1	X	X	X	X	BX	DAC 3 DAC Register
1	1	0	0	X	X	X	X	CX	DAC 4 Low-Byte Register
1	1	0	1	X	X	X	X	DX	DAC 4 Mid-Byte Register
1	1	1	0	X	X	X	X	EX	DAC 4 High-Byte Register
1	1	1	1	X	X	X	X	FX	DAC 4 DAC Register

Table III. Truth Table For Control Data

8

APPENDIX A1

A simple latch is all that is required with the 8085A to generate the serial out stream (see Figure A1). In practice the least significant bit of a spare I/O port (Port A) on an 8355 was used as the serial out pin. This port was designated Port 00. The sequence of events for the XMIT routine is shown in Figure A2 and the corresponding program appears in Table AI.

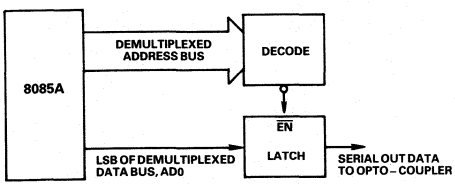

Figure A1. Serial Out Generator

The XMIT subroutine is entered with the character to be transmitted contained in register C. A bit counter is set to 11_{10} and is decremented as each bit is transmitted. The 11 bits are made up of 1 start bit, 8 data bits and 2 stop bits. The DELAY subroutine is called after each bit is transmitted (see Table AII). By varying the length of the delay different transmission rates can be achieved. For a rate of 9600 bits/second, "XX" in the delay subroutine was set to 11H (17_{10}). With this baud rate 872.7 11-bit characters per second are transmitted giving a true data rate of 6981.6 bits/second. At this speed an analog output can be updated in 4.5mS. Table AIII shows the baud rates obtained for different values of "XX" with an 8085A clock frequency of 3.1 MHz. Some provision should be made to allow the UART to clock itself out of an error condition. This could arise if the serial transmission link is broken or on power-up of the remote circuit. This requirement is met if, after transmitting four characters to one DAC, transmission ceases for a period equal to the time taken

LOCATION	OP CODE	MNEMONIC	STATEMENT
2000 XMIT	06 B	MVI B, B	Set Bit Counter to 11
2002	3E FF	MVI A, FF	Initialize Port 00 as Output
2004	D3 02	OUT 02	
2006	79	MOV A,C	Fetch Character
2007	B7	ORA A	Clear Carry
2008	17	RAL	
2009	D3 00	OUT 00	Send to Port
200B	CD 20 20	CALL 2020	Call DELAY Routine
200E	1F	RAR	Rotate Next Bit
200F	37	STC	
2010	05	DCR B	
2011	C2 09 20	JNZ 2009	Finished?
2014	C9	RET	Return

Table AI. Program Listing for XMIT Routine

to transmit a single character. During this time the serial data line must be held high to hold the RRI input high. This happens automatically in fact, since at the end of each character the line goes high for the stop bit and remains high until the low start bit of the next character occurs.

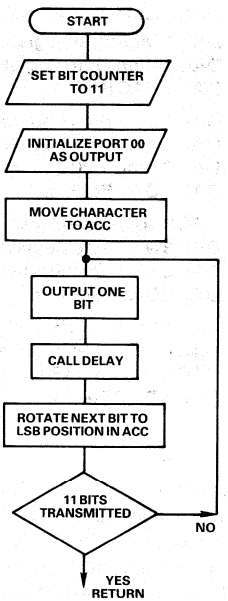

Figure A2. Sequence of Events in XMIT Routine

LOCATION	OP CODE	MNEMONIC	STATEMENT
2020 DELAY	D5	PUSH DE	Save DE on Stack
2021	1E XX	MVI E, XX	Load Required Baud Rate
2023	1D	DCR E	Delay
2024	C2 23 20	JNZ 2023	Finished?
2027	D1	POP DE	Pop DE from Stack
2028	C9	RET	Return

Table AII. Program Listing for DELAY Routine

BAUD RATE	"XX"
9600 bits/sec	11H
4800 bits/sec	27H
2400 bits/sec	56H
1200 bits/sec	BOH

Table AIII. Baud Rate Variation vs. XX for Delay Routine

ACKNOWLEDGEMENTS

To Phil Burton for originally suggesting the idea and also to Ray Speer for his help.

REFERENCES

1. Phil Burton, "CMOS D/A Converter Circuits for Single +5V Supplies," PUB. NO. E828-10-4/84
2. CMOS DAC Application Guide, PUB. NO. G872-30-10/84

ANALOG DEVICES

AN-329
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Dynamic Performance of CMOS DACs in Modem Applications

by Mike Curtin and Matt Smith

INTRODUCTION

In the new high-speed modems manufactured to meet the V.32 and V.33 standards, it is of prime importance to be able to produce a high-quality carrier signal. The D/A converter used to produce this needs excellent dynamic characteristics; harmonic distortion must be typically less than −70dB. This application note evaluates three Analog Devices CMOS DACs when used in this application. It explains how to get the best performance from each DAC and looks at the requirements for deglitchers. The note is intended to provide the information necessary for modem designers to evaluate these DACs and assess their suitability for particular systems.

TEST CIRCUIT AND CONDITIONS

In both the V.32 and V.33 standards, the carrier signal frequency is 1800Hz. The D/A converter digitally constructs a composite signal with update rates of 9.6kHz in V.32 systems and 14.4kHz in V.33 systems. One of the fall back rates for both standards is 7.2kHz. Figure 1 shows the test circuit used to evaluate the distortion performance of the DACs. The system generates an 1800Hz sine wave using the three DAC update rates already mentioned (14.4kHz, 9.6kHz and 7.2kHz). To do this the HP300 Series computer generates the digital values for the sine wave based upon the output frequency and the update frequency. These are then loaded into RAM. When the timing control logic block is activated, it sequences through the RAM at the predetermined update rate and loads the digital words into the DAC to produce the 1800Hz sine wave.

This is then fed into the spectrum analyzer via the deglitcher which may or may not be used depending on the glitch performance of the DAC under test. The spectrum analyzer shows the spectral content of the output sine wave, and output distortion can be calculated from this.

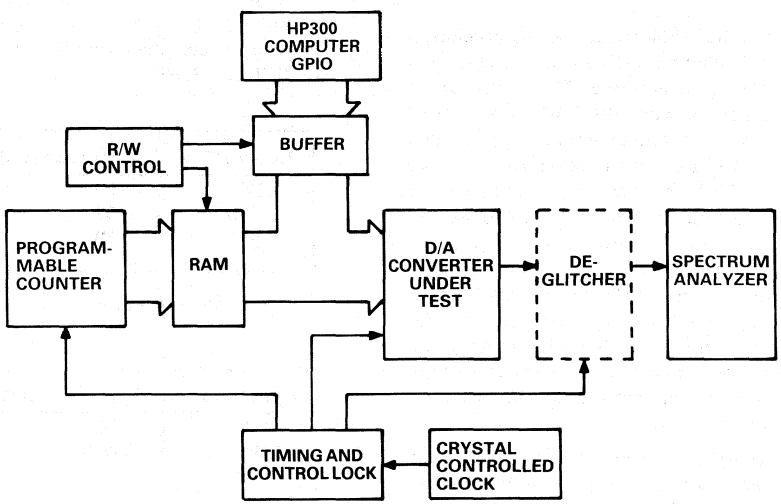

Figure 1. Harmonic Distortion Test Circuit

DEGLITCHER

When the code changes in a current-steering CMOS DAC, there is a capacitive coupling of charge across the switches in the DAC. This causes an injection of charge into the I_{OUT} line which in turn causes a voltage spike or glitch to appear at the output of the current to voltage amplifier. When the DAC is being used in sine wave construction, the presence of these output glitches results in increased harmonic distortion. This harmonic distortion increases as the glitches increase.

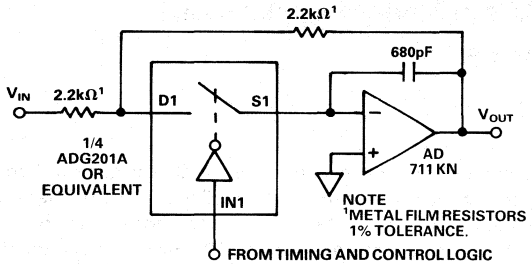

Figure 2. Deglitcher Circuit

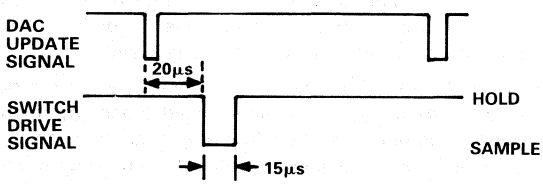

Figure 3. Deglitcher Timing

In order to overcome the problems caused by glitches, it is possible to use a deglitcher. Such a circuit is shown in Figure 2 with its timing in Figure 3. The digital value loaded to the DAC changes when the DAC update signal goes low. The switch in the deglitcher is turned on 20µs after this, causing V_{IN} to be sampled. This delay allows plenty of time for the DAC output to settle. The deglitcher samples for 15µs before going back into hold and waiting for the next DAC update. At the slowest sampling rate of 7.2kHz this means that the signal must be held for up to 125µs. This presents no problems for the circuit. The control signal for the deglitcher can be derived from the DAC update signal with a dual monostable. Alternatively, if a microcomputer with on-chip timer-counter is used to load the DAC, then the deglitcher control signal can be set up as an output from the controller.

The deglitcher of Figure 2 is used in some of the following circuits with excellent results. It is up to the system designer to decide if he can achieve acceptable performance without the deglitcher. The results given in this application note will help in making this decision.

AD7537/AD7547 In Current-Steering Mode

The AD7537/AD7547 are dual 12-bit DACs, packaged in narrow 0.3", 24-pin DIPs or 28-terminal LCCCs and PLCCs. Power consumption is low (100mW typical). The only difference between the two devices is in their loading structures. The AD7547 has a 12-bit parallel structure while the AD7537 has (8+4) loading. Figure 4 shows the AD7537/AD7547 set up in the standard current mode for ±5V output. This circuit is used with the system of Figure 1 to produce an output sine wave with 10V pk-pk amplitude and 1800Hz frequency.

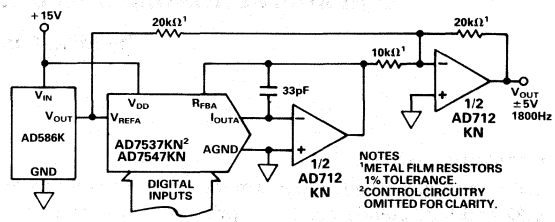

Figure 4. AD7537/AD7547 in Bipolar Current-Steering Mode

The distortion figures for this circuit at the three update rates (14.4kHz, 9.6kHz, 7.2kHz) vary from −69dB to −73dB. This is shown in Table I. To improve on this you can use the deglitcher circuit of Figure 2. This cleans up the output waveform and gives improved distortion performance which is shown in Table II. The improvement gained by using the deglitcher is about 10dB which shows that the glitches in the DAC have a considerable effect on the ac performance. APPENDIX A contains a selection of spectral responses obtained with different circuits. All of those shown are for the 14.4kHz sampling frequency, making comparisons between the circuits more meaningful. Figures A1 and A2 show the spectral responses for Figure 4 with and without the deglitcher. When the deglitcher is used, the level of the 2nd harmonic drops from −70dB to −90dB.

Update Rate	Distortion (dB)	2nd Harmonic (dB)
7.2kHz	−73	−73
9.6kHz	−69	−71
14.4kHz	−69	−71

Table I. Distortion Performance of Circuit in Figure 4

Update Rate	Distortion (dB)	2nd Harmonic (dB)
7.2kHz	−85	−85
9.6kHz	−78	−86
14.4kHz	−76	−90

Table II. Distortion when Using Deglitcher of Figure 2

AD7537/AD7547 in Voltage-Switching Mode

Another way of getting improved distortion performance from the AD7537/AD7547 is to use it in the voltage-switching mode. Figure 5 shows the circuit diagram. The DAC connections have now been reversed with the reference voltage applied to the I_{OUTA} terminal, AGNDA grounded and the V_{REFA} terminal as the output. R_{FBA} is not used and is tied to I_{OUTA} to prevent stray pickup. The glitches at the output of this circuit are much smaller than in Figure 2. The point where the glitches appear (I_{OUT}) is now connected to a low impedance point (AD580 output) which can absorb the current glitches without producing voltage spikes. The resulting major-carry glitches from Figure 5 are typically 20nV-secs compared to 200nV-secs for Figure 4. When operating in the voltage-switching mode, the DAC linearity degrades as the reference voltage increases. For this reason, the reference input in Figure 5 is limited to +2.5V, giving an output signal of ±2.5V.

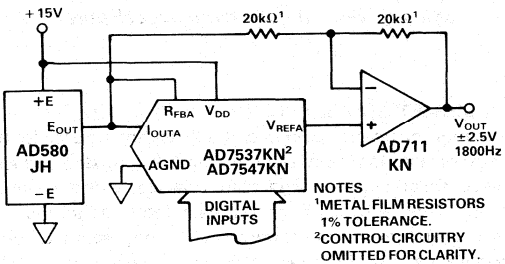

Figure 5. AD7537/AD7547 in the Voltage Switching Mode

Table III shows the distortion figures for the voltage-switching circuit. The results are the same both with and without the deglitcher and Figures A3 and A4 are the spectral responses when sampling at 14.4kHz. The deglitcher makes no difference to the performance of this circuit, verifying the absence of glitches in the outputs. However, if you compare the distortion results directly with those for the standard current-steering circuit plus deglitcher, you can see there is a slight degradation. Though the glitches are much smaller, the linearity performance in the voltage mode is degraded. (See the AD7537/AD7547 data sheets for typical performance.) This degraded linearity will increase distortion in the output signal. The distortion figures for the voltage-mode circuit are still better than −70dB making it suitable for many modem applications.

Update Rate	Distortion (dB)	2nd Harmonic (dB)
7.2kHz	−75	−75
9.6kHz	−74	−82
14.4kHz	−72	−79

Table III. Distortion Performance of Figure 5 with and without Deglitcher

AD7245/AD7248

The AD7245 and AD7248 are voltage-switching, 12-bit DACPORTs®. Each contains a reference, 12-bit DAC and output amplifier. They differ only in their loading structure; the AD7245 is a 12-bit parallel load device while the AD7248 has a byte loading structure. When connected as in Figure 6, the AD7245/AD7248 can be programmed for a sine wave output as previously discussed. The output signal range is ±5V.

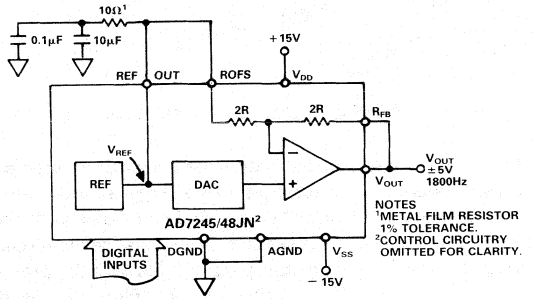

Figure 6. AD7245/AD7248 Connected for Bipolar ±5V Output

Table IV shows the distortion results obtained from this circuit, and Figure A5 is the spectral response. The distortion is caused by slew rate distortion in the AD7245/AD7248 output amplifier. If the deglitching, circuit of Figure 2 is cascaded with Figure 6 and used as a sample/hold amplifier, there is a marked improvement in performance. The output is now sampled at a time when the slew rate effects are over and distortion is determined by the SHA output amplifier (AD711). The distortion results with this setup are shown in Table V, and Figure A6 shows the spectral response for the 14.4kHz – sampled signal.

Update Rate	Distortion (dB)	2nd Harmonic (dB)
7.2kHz	−76	−76
9.6kHz	−57	−71
14.4kHz	−61	−69

Table IV. Distortion of AD7245/AD7248 in Figure 6

Update Rate	Distortion (dB)	2nd Harmonic (dB)
7.2kHz	−82	−82
9.6kHz	−74	−83
14.4kHz	−75	−82

Table V. AD7245/AD7248 Distortion Performance Using Circuit of Figure 2

DACPORT is a registered trademark of Analog Devices, Inc.

AD7538 in Current-Steering Mode

The AD7538 is a 14-bit CMOS DAC and is of interest to the designer who needs somewhat better than 12-bit system performance. The device is packaged in the same narrow 24-pin package as the previous DACs.

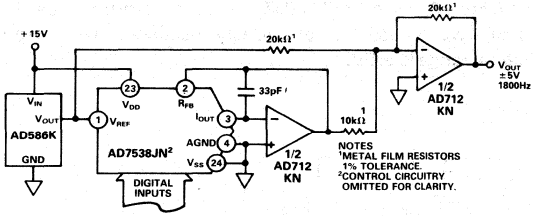

Figure 7. AD7538 in Bipolar Current-Steering Mode

Figure 7 shows the AD7538 in the standard bipolar current-steering configuration. Table VI gives the dynamic performance of Figure 7 when used on its own without a deglitcher. The performance is a reflection of the larger D/A glitch impulse which this part exhibits over the previous 12-bit devices. When the deglitcher of Figure 2 is used there is a major improvement in the performance as can be seen from Table VII. THD is now down to a level of −86dB for the 14.4kHz−sampled signal. Figure A7 and A8 are the spectral responses for Figure 7 with and without the deglitcher.

Update Rate	Distortion (dB)	2nd Harmonic (dB)
7.2kHz	−56	−56
9.6kHz	−57	−65
14.4kHz	−54	−57

Table VI. Distortion Performance of Figure 7

Update Rate	Distortion (dB)	2nd Harmonic (dB)
7.2kHz	−88	−88
9.6kHz	−84	−89
14.4kHz	−86	−89

Table VII. Distortion Performance Using Deglitcher

AD7538 in Voltage-Switching Mode

When the AD7538 is operated in the voltage mode, the reference driving I_{OUT} must have extremely good dynamic characteristics (i.e., response to a changing load). The nominal impedance which it is driving is 6kΩ and this changes with the DAC code. As well as this, the glitches which the reference must absorb at I_{OUT} as the code changes are large. The AD580 is not capable of delivering this performance and must be buffered by the AD712 shown in Figure 8. The distortion results are in Table VIII. The spectral responses for Figure 8 are shown in Appendix A. Figure A9 is the response without deglitching and Figure A10 is the response with deglitching.

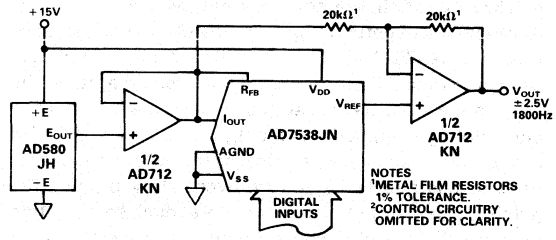

Figure 8. AD7538 in the Voltage Switching Mode

Update Rate	Distortion (dB)	2nd Harmonic (dB)
7.2kHz	−75	−75
9.6kHz	−75	−79
14.4kHz	−72	−78

Table VIII. Distortion Performance of Figure 8

DISCUSSION OF RESULTS

The method used to calculate the distortion figures for the various circuits was to feed the digitally constructed sine wave directly into the spectrum analyzer without any filtering. Then, the distortion was calculated as the rms sum of the distortion components. This may seem to be a somewhat cumbersome approach. However, if a distortion meter was used, there would be very severe output filtering requirements. For example, when sampling at 7.2kHz, even if we wanted to measure only up to the 2nd harmonic (3.6kHz), a filter with cutoff at 3.6kHz and sufficient attenuation at 7.2kHz to eliminate the clock frequency would be necessary. The best way of practically evaluating the circuit distortion is to use the spectrum analyzer. For each of the update rates, components up to half this rate were summed to calculate the circuit distortion. In this way, all harmonics of interest are included.

In practical modem systems, the output signal from the D/A converter will be followed by a filter section. This low-pass filter will nominally have a cutoff frequency of 3.5kHz for V.32 and V.33 systems. It will remove clock components and other unwanted noise from the carrier signal. The output filter will also attenuate all but one of the harmonics at the output. The one that does not get attenuated is the 2nd harmonic (3.6kHz) and so it is of special interest to the modem designer. For this reason the 2nd harmonic in each of the circuits tested has been listed separately in the tables.

In analyzing the performance of the various circuits, the best results were obtained from the AD7537/AD7547 and AD7538 when both were operated in the current-steering mode and the output was deglitched. The 2nd harmonic level was equal for both of these (−89dB), but in terms of THD the AD7538 was superior (−86dB versus −76dB for the 14.4kHz sampled circuit). In practical terms this means that if the system designer is using a 14.4kHz clock-rate, his filter requirements will be less if he uses the AD7538.

Attenuation of the higher harmonics isn't as critical as with the AD7537/AD7547 and so a lower order filter is possible. If the designer wants to eliminate the deglitcher, then the best circuits to use are the AD7537/AD7547 or the AD7538 in the voltage-switching mode. Both of these give good results with a THD of −72dB and 2nd harmonic level of −78dB in the 14.4kHz sampled circuit.

The AD7245/AD7548 achieves good performance when the circuit of Figure 2 is used with it as a sample/hold. It also has the advantage of an on-board reference. THD for this circuit is −75dB and the 2nd harmonic level is −82dB.

CONCLUSIONS

Each of the CMOS DACs discussed in this application note is capable of delivering better than −70dB distortion performance when synthesizing an 1800Hz signal with the stated update frequencies. Some, but not all, of the DAC configurations require deglitching to achieve this. There are varying levels of performance among the devices which deliver better than −70dB THD, ranging from −72dB for the AD7537/AD7547 in the voltage mode without deglitching to −86dB for the AD7538 in the current-steering mode with a deglitcher.

It should be remembered that the results contained in this application note are typical and were obtained from a range of devices taken randomly. Several production/fabrication lots were sampled. However, the results are meant to show typical performance only and do not guarantee that this will be met in all cases.

APPENDIX A

Spectral Responses of Figures 4, 5, 6, 7 and 8. F_{OUT} = 1800Hz. Update Rate = 14.4kHz. Responses are shown both with and without the deglitcher circuit of Figure 2.

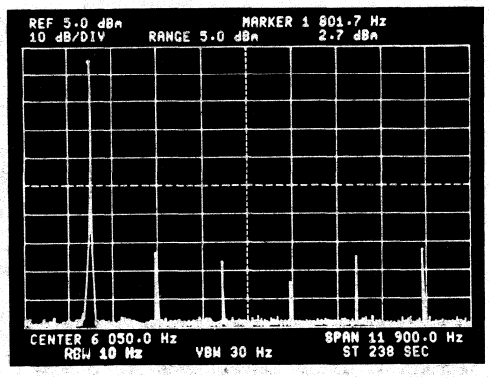

Figure A1. Spectral Response of Figure 4 (AD7537/AD7547 in Current-Steering Mode) without Deglitcher

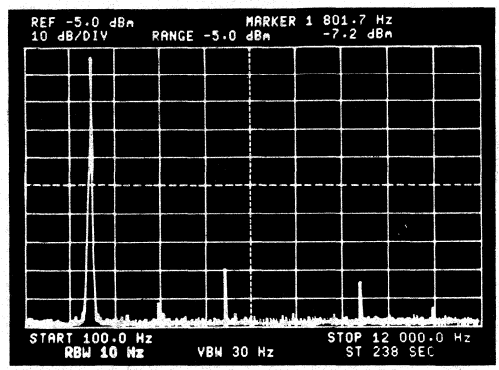

Figure A2. Spectral Response of Figure 4 with Deglitcher

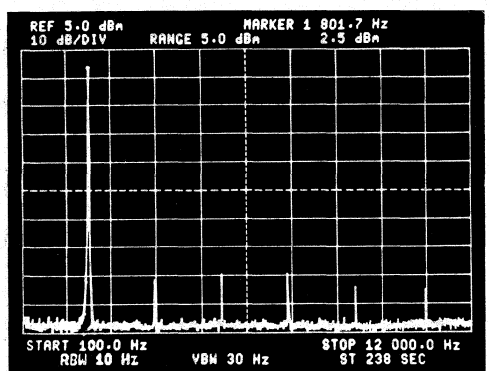

Figure A3. Spectral Response of Figure 5 (AD7537/AD7547 in Voltage-Switching Mode) without Deglitcher

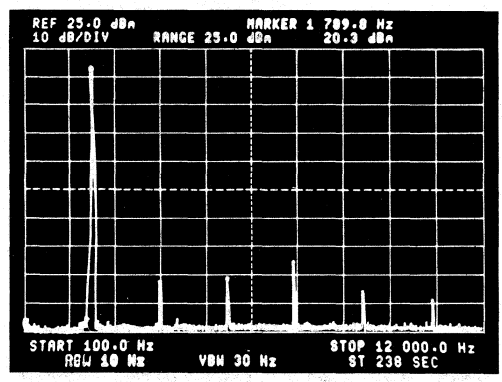

Figure A4. Spectral Response of Figure 5 with Deglitcher

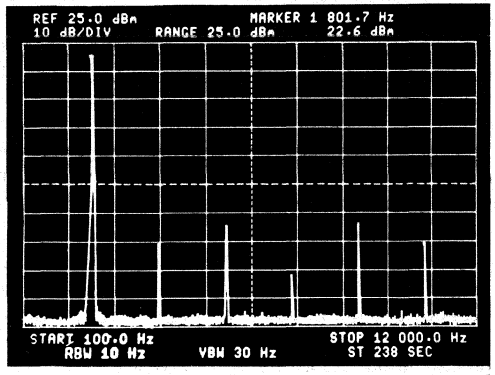

Figure A5. Spectral Response of Figure 6 (AD7245/ AD7248 in Bipolar Output Configuration) without Deglitcher

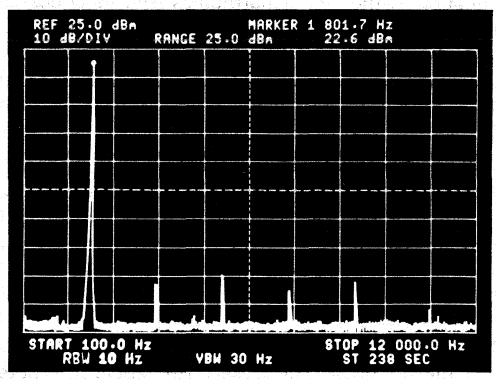

Figure A6. Spectral Response of Figure 6 with Deglitcher

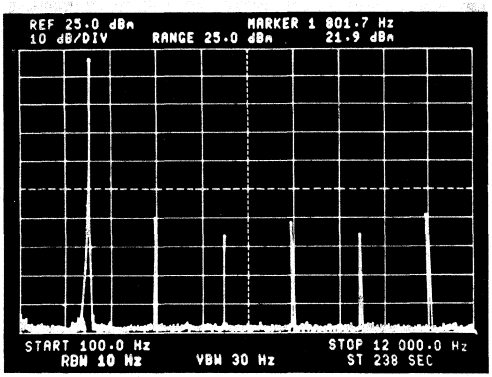

Figure A7. Spectral Response of Figure 7 (AD7538 in Current-Steering Mode) without Deglitcher

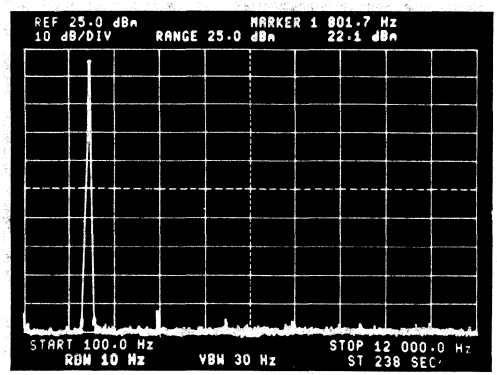

Figure A8. Spectral Response of Figure 7 with Deglitcher

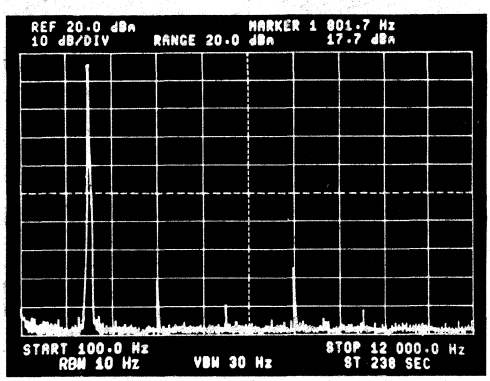

Figure A9. Spectral Response of Figure 8 (AD7538 in Voltage-Switching Mode) without Deglitcher

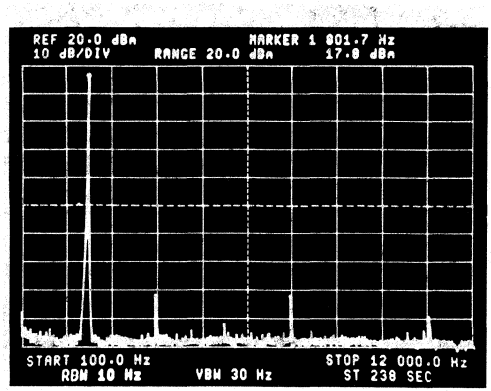

Figure A10. Spectral Response of Figure 8 with Deglitcher

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Eight-Channel CODEC Applications

Analog Devices Inc. has developed a CODEC System for the purpose of demonstrating its use in a shared-channel system. These circuits provide a working digital transmission system incorporating eight analog input channels digitally interfacing to eight output channels. The circuit design was completed so that a single analog printed circuit board could be used either for encoding eight channels or decoding eight channels. A single digital timing and interface board is used to provide system clocks and the parallel digital data bus from transmitter to receiver. The design uses pluggable connections so that, if desired, the encoding portion or the decoding portion of the system can be tested separately. The user need only provide three voltage supplies (±15VDC, +5VDC), the appropriate filters and any applicable transmission test equipment.

HARDWARE

The entire eight-channel system consists of twenty-one integrated circuits mounted on three printed circuit boards. Two of the boards, as mentioned, are a common layout which is used for the analog functions of the system. The analog board used for the encoder (analog-to-digital conversion) requires the addition of a COMDAC® (DAC-86, 89), a MUX-88, a SMP-81, a CMP-01, and a REF-02. The other analog board is used for the decoder (digital-to-analog conversion) with the addition of a DAC-86, MUX-88, REF-02, and an OP-16. The general schematic and parts list for the analog boards are shown in Figures 1 and 2.

The control board is a TTL circuit of twelve devices designed to provide three basic functions: a) the seccessive approximation register with interface circuit to the DAC-86, b) the

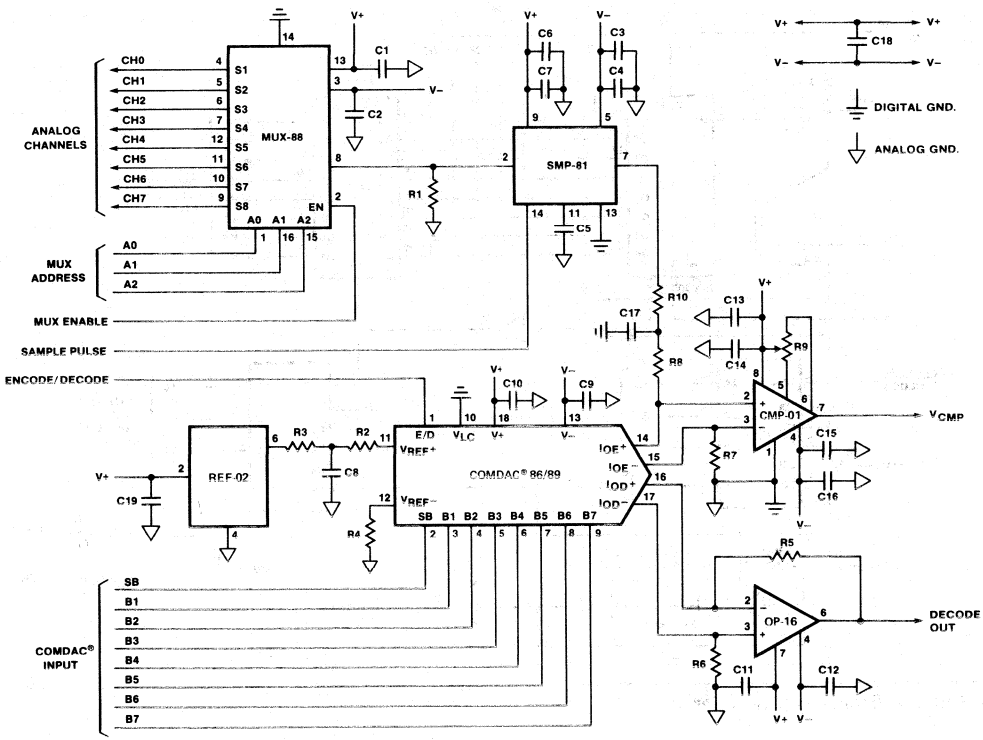

Figure 1. COMDAC® Encoder/Decoder Analog Circuit Board

COMDAC is a registered trademark of Analog Devices, Inc.

TRANSMIT (Eight-Channel Encoder)	RECEIVE (Eight-Channel Decoder)
MUX-88	MUX-88
SMP-81FY	COMDAC® DAC-86EX
COMDAC® DAC-86EX (89)	(89)
	OP-16F
CMP-01EJ	REF-02
REF-02	

RESISTORS	
R1 — 20kΩ	R6 — 2.49kΩ
R2 — 330Ω	R7 — 2.49kΩ
R3 — 9.1kΩ	R8 — 1.5kΩ
R4 — 10kΩ	R9 — 2kΩ (POT)
R5 — 2.49kΩ	R10 — 1.0kΩ

CAPACITORS	
C1, C2, C4, C7, C8, C9, C10, C11, C12, C13, C15, C18, C19	0.1μF
C3, C6, C14, C16	10μF
C5	5000pF
C17	100pF

Figure 2. Parts List — Analog Board

clock for the encoder circuit, and c) the digital interface between the encode and decode sections. The encoder clock design is a multiple frequency clock, the usefulness of which is treated in a later section, generated from a programmable read-only memory frequency divider. The PROM was used to provide flexibility in changing the clock waveforms if the user so wishes. The resultant clock waveforms are also described later; the circuit schematic is shown in Figures 3a and 3b. The PROM data is listed in Figure 4. The digital interface is provided by using a standard 8-bit, parallel-in, parallel-out latch updated as the successive approximation process is completed for each input channel. The remainder of the circuit consists of the SAR and the multiplexer address counters. A parts list for the circuit is shown in Figure 5. A complete circuit schematic for the digital board is shown in Appendix A.

The boards are interconnected by use of four mini-dip connectors and cables, a 16-pin and 14-pin from each analog board to the controller. The lead designations of the two connectors are shown in Figure 6. The input and output channels are accessible through "banana"-type plugs; this allows optional connections from the transmission line in order to try different types of filters and line interface circuits. The two analog boards require ±15VDC and "banana" plugs are provided to interface to the appropriate supplies. The control board requires +5VDC only. The entire system layout is shown in block diagram in Figure 7.

SYSTEM OPERATION AND DESIGN

To achieve a workable system configuration, the encoding and decoding operations were approached as two separate designs. The entire transmission link was then connected to complete the end-to-end tests.

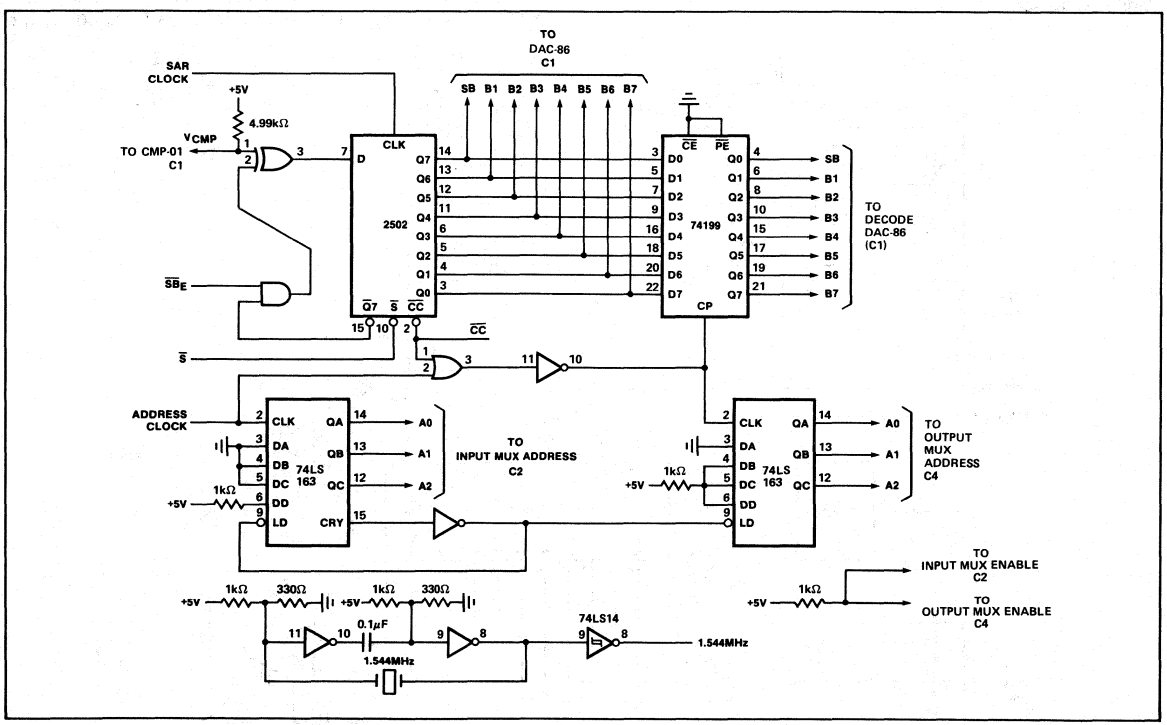

Figure 3A. Encoder/Decoder Controller

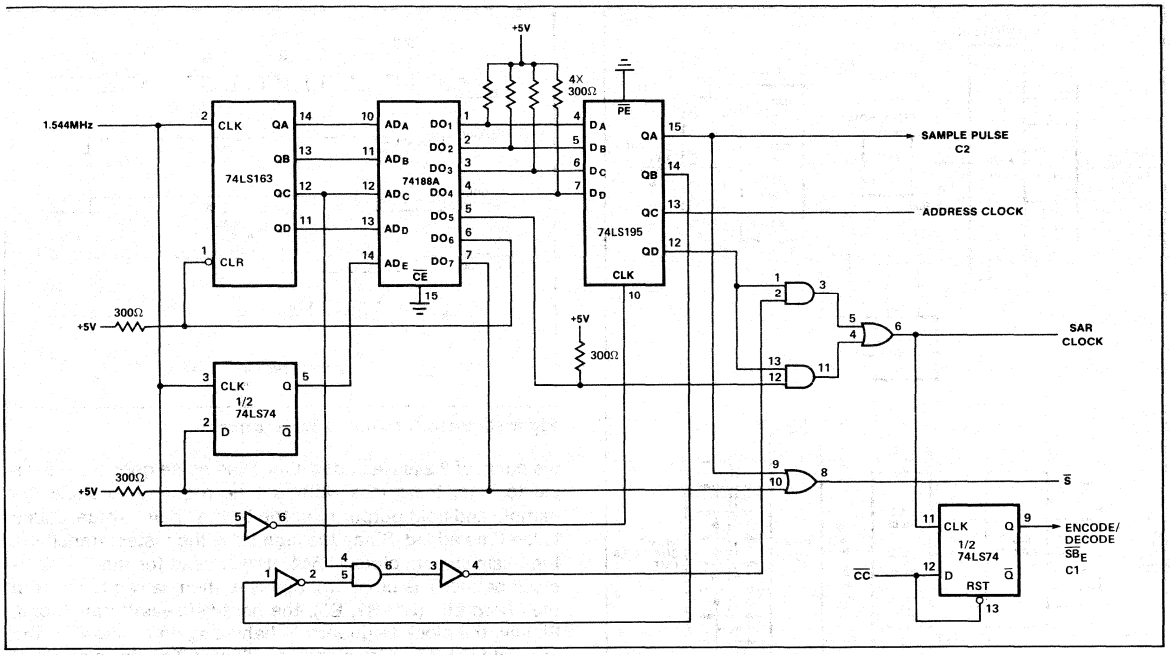

Figure 3B. Encode Clock

ADDRESS	DATA (MSB—LSB)	ADDRESS	DATA (MSB—LSB)
00	20	0C	27
01	28	0D	27
02	20	0E	2F
03	21	0F	6F
04	39	10	67
05	2B	11	67
06	23	12	6F
07	2B	13	6F
08	23	14	67
09	2B	15	67
0A	23	16	6E
0B	2F	17	0E

$MSB = D0_8$
$LSB = D0_1$
Address 18 — 1F — Unused.

Figure 4. PROM-Based Clock

PARTS			
74LS04	2	74LS163	3
74LS08	1	74188A	1
74LS14	1	74LS195	1
74LS32	1	74199	1
74LS74	1	2502	1
74LS86	1	1.544MHz Crystal	

Figure 5. Parts List — Controller

CONNECTORS		
PIN #	C1 (C3)	C2 (C4)
1	+5	+5 GND
2	+5	Sample Pulse
3	+5	MUX Enable
4	VC MP	A0 — Address
5	+5 GND	+5 GND
6	+5 GND	A2 — MUX Address
7	+5 GND	A1 — MUX Address
8	B7 — LSB	+5 GND
9	B6	+5 GND
10	B5	+5 GND
11	B4	+5 GND
12	B3	+5 GND
13	B2	+5 GND
14	B1	+5 GND
15	SB — MSB	+5 GND
16	Encode/Decode	+5 GND

*C2 is 14 Pin — C1 is 16 Pin

Figure 6. Pin Designations — System Connectors

The accuracy of the encoder, the analog-to-digital converter, is dependent upon several factors. The first, and most significant, is the speed (settling time) of the companding DAC in conjunction with the comparator. Other factors are switching time of the multiplexer, acquisition time of the sample-and-hold, hold-step settling time of the sample-and-hold, and output noise of the sample-and-hold and

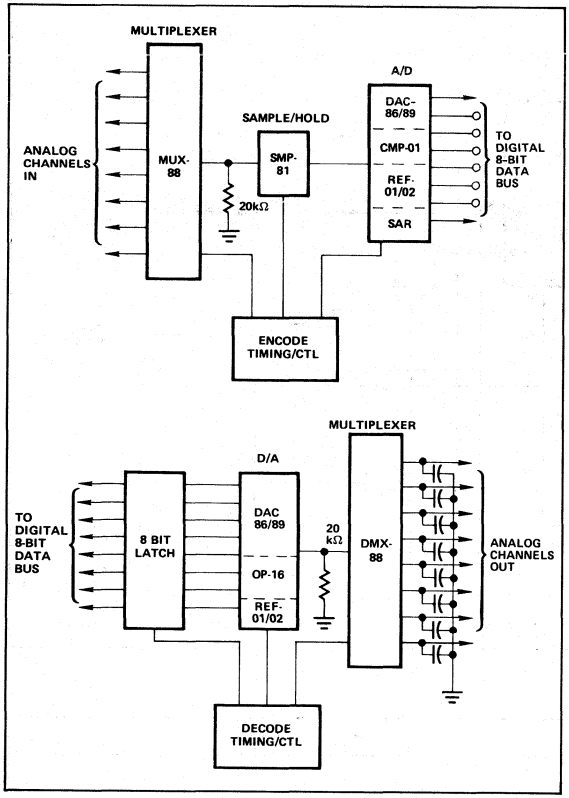

Figure 7. Eight-Channel System Layout

multiplexer. All of these characteristics had to be considered while developing the encoder circuit.

In a companded A/D conversion, using a successive approximation method, different digital bits require different settling times. There is a constant increase in the settling time required as one goes from the sign bit to the chord bits and then to the step bits. This increase in settling time is partially caused by the scaled current sources used to design a companding D/A and is most affected by the magnitude of the voltage level being converted. So as the output currents of the DAC-86 become smaller (at smaller voltage levels), the comparator is less responsive. Also as the less significant bits are clocked in, the situation becomes more critical since progressively smaller current changes are provided to force the comparator to change output state. The encoder clock waveforms, shown in Figure 8, depict a system approach to accommodate these timing characteristics. The governing design criteria was that a limited amount of time is available to complete the successive approximation of the analog signal (for eight channels, with a sampling frequency of 8kHz, this means 15.6μs) and the optimum clock must fit within this period. All functions (including sampling, encoding, switching) must therefore be completed with an acceptable accuracy, within this time constraint. To achieve accurate data acquisition with the sample-and-hold, a sam-

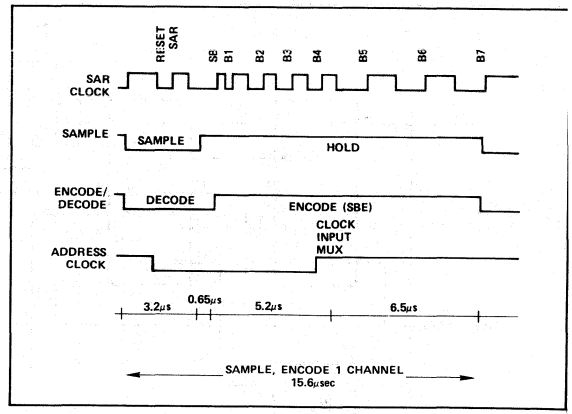

Figure 8. Encoder Timing Waveforms

ple pulse of 3.2μs was used. Once the pulse goes from sample to hold, there is a waiting time required to allow the sample-and-hold output to settle to the "held" value, 650ns is the time added. Since the sign bit is the fastest transition, the basic system clock (1.544MHz) is used for the first SAR clock period. The clock frequency is then halved to clock in the chord bits (B1, B2, B3), the next "slowest" transitions. Finally, the clock frequency is halved again to allow for the step bits (B4-B7), a frequency of 386kHz. The times required for the different bit conversions are shown in Figure 8. The SAR is reset during the sampling period and once the sample circuit has entered the hold mode, the input multiplexer's address is changed. This provides ample time for the MUX to switch to the next analog input with little effect on the adjacent waveforms. In addition, to further minimize interchannel crosstalk, a 20kΩ resistor to ground is added to the multiplexer output.

To provide sufficient noise immunity from the sample-and-hold to the comparator, a simple filtering circuit using a 100pF capacitor is added. A 4.9kΩ resistor to +5V from the comparator output aids in increasing the DAC-86 and comparator speed. In order to test the encoder without adding a decoder, a high-precision voltage source (DC levels) was used as an input to one or more channels of the A/D circuit. The data output (one byte every 15.6μs) was sampled and stored in a logic analyzer. By reviewing the samples for each channel, and comparing the data from all channels of the eight channel system, the effects of changes in the clock on the accuracy of the design could be observed. By inputting a steady-state value, the data out would not only demonstrate the "consistency" of the A/D conversion but also make obvious any adjacent channel effects that were produced.

The analyzer sampling also presents a method of initializing the system components. By grounding all inputs and digitally sampling the data bus output, the zero level can be set so as to produce an alternating series of data bytes equivalent to 0 volts (10000000 and 00000000).

A similar design approach was used in developing an eight-channel decoder. In this case, the response time of the DAC-86 and op amp offers little problem since an eight-

channel decoder is relatively slow. However, the accuracy of the devices does become important. The design considerations become, therefore, the nonlinearity of the DAC-86 in conjunction with the switching time and charge injection of the multiplexer. The DAC-86 is manufactured to strict linear specifications, which assures both excellent decoding linearity and absolute accuracy. The multiplexer is used both as a switch to demultiplex the output waveforms and as a holding circuit by adding capacitance to the output leads. One effect seen in the multiplexer is that as a channel is switched off, there is a charge injected onto the output. This charge is not normally obvious in a MUX design, however, when working into a high impedance (such as a filter) and with capacitors on the source leads (outputs), the charge can add an offset to the waveform. To minimize the effect, a large capacitor (0.1μF or greater) is added. Since the charge pulse is a short duration signal, the signal on the larger capacitor will be less affected by the charge than if a smaller component was used.

In terms of a system decoder design, this circuit could be used for more than eight channels. The op amp and DAC-86 are both capable of responding to more output channels. An eight-channel system was incorporated to remain compatible with the analog board that is used in the demonstrator. In order to decode more channels (>12), the output capacitance on the demultiplexer would need to be reduced. The other circuit components would remain the same.

The decoding circuitry can be tested separately by adding a series of data bytes and monitoring the output channels. The CCITT recommendation for testing PCM systems includes a method of testing a decoder by introducing a standard sequence of digital data words in order to produce a 1kHz sinusoid at a nominal level of 0dBm0.* This method proved useful in "debugging" the circuit design prior to attempting the end-to-end tests.

SYSTEM TESTS

Once the encoder and decoder were functioning separately, the entire system was connected with the appropriate interfacing to allow for full-system transmission tests. The measurements taken were the standard set of PCM specifications observed in the majority of data sheets for telecommunication oriented products. These tests included signal-to-total distortion, gain tracking, intelligible crosstalk, and idle channel noise.

There are two different methods of performing the tests. For the U.S. standards, a sinusoid signal is used as the channel input and the measurements are taken around the base (test) frequency. For the European tests, a pseudo-random or "white" noise source is preferred as the input signal. The test results discussed here were obtained with the U.S. testing procedures; similar test results have been achieved using European test methods.

The test set-up for signal-to-total distortion and gain tracking measurements is shown in Figure 9. The results are plotted in Figures 10 and 11. As is seen, each test was performed with both the C-Message Weighting and the 3kHz Flat response terminating configurations. The results using the μ-law parts (DAC-86) are represented. In terms of signal-to-total distortion, the system exceeds the recommended standard at all input levels by 2dB or greater. The system is also well within the recommended gain tracking limits for both terminations.

The crosstalk and idle channel noise measurements are given in Figure 12. One consideration of the system design, in terms of performance, was that the most difficult characteristic for a shared-channel system to minimize is the intelligible crosstalk specification. The design is directed

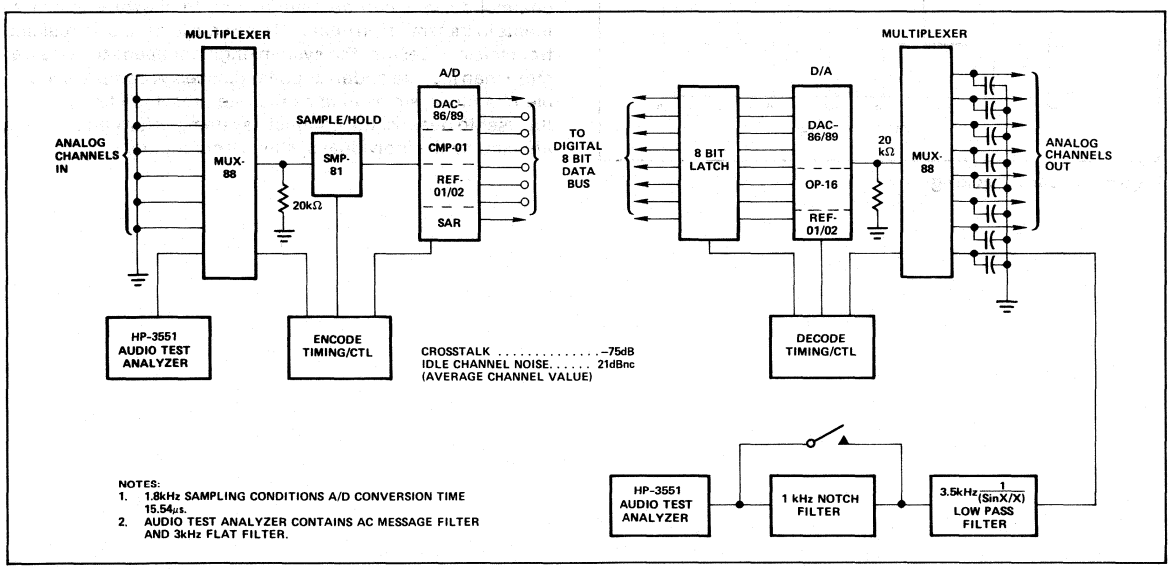

Figure 9. Eight-Channel Test Configuration

*Reference — CCITT Sixth Plenary Assembly (1976), Orange Book Vol. III-2

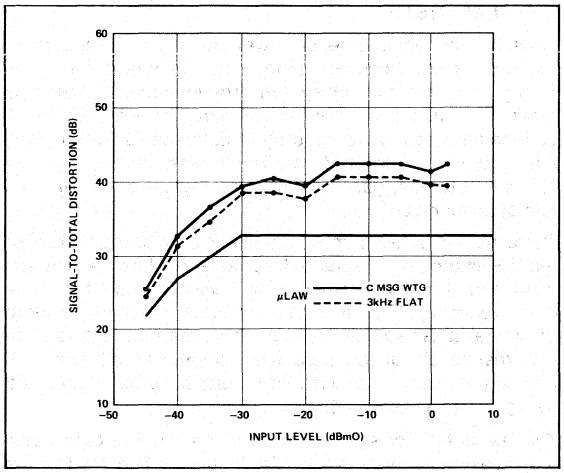

Figure 10. Signal-To-Quantizing Distortion vs. Input Level

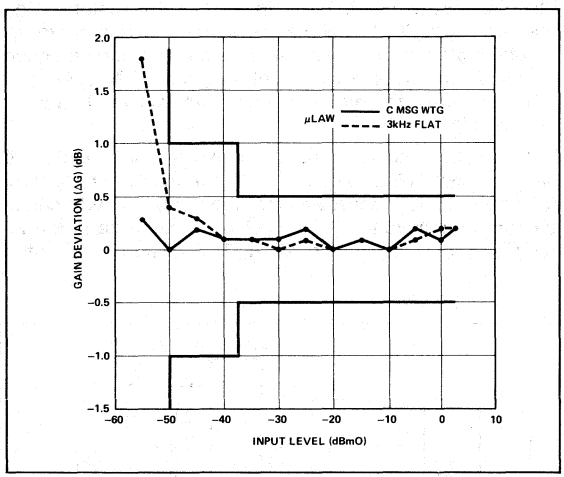

Figure 11. Gain Tracking

IDLE CHANNEL NOISE

CHANNEL	NOISE (dBm0)	CHANNEL	NOISE (dBm0)
1	−66.9	5	−62.6
2	−67.6	6	−65.3
3	−67.0	7	−67.0
4	−67.2	8	−61.8

CROSSTALK

FREQUENCY	INTELLIGIBLE CROSSTALK
300 — 2900Hz	≤ −78dBm0
2900 — 3400Hz	≤ −70dBm0

Figure 12. Idle Channel Noise and Crosstalk

toward optimizing this measurement. The idle channel noise is at the system recommended level when measured without output filtering. It is further reduced by adding a PCM receive filter on the decoder multiplexer.

CONCLUSIONS

The circuitry just discussed is meant to represent one approach to designing an eight-channel, shared CODEC system. It is not meant to be the only design, but provides a working system upon which to base further engineering development. It should be noted that in terms of transmission testing, the design is an end-to-end system. The configuration presents the users with a complete circuit enabling them to observe the individual device characteristics significant in producing a shared-channel design.

The design provides a starting point from which most characteristics important to both shared-channel and single-channel designs can be manipulated to allow for improvements in transmission quality. To be able to develop a realistic transmission design, the system engineer needs to consider more than just the coder/decoder devices. A complete multiple-channel system, such as the one presented here, allows the user to observe the complete system performance as it is affected by the individual system components.

APPENDIX A

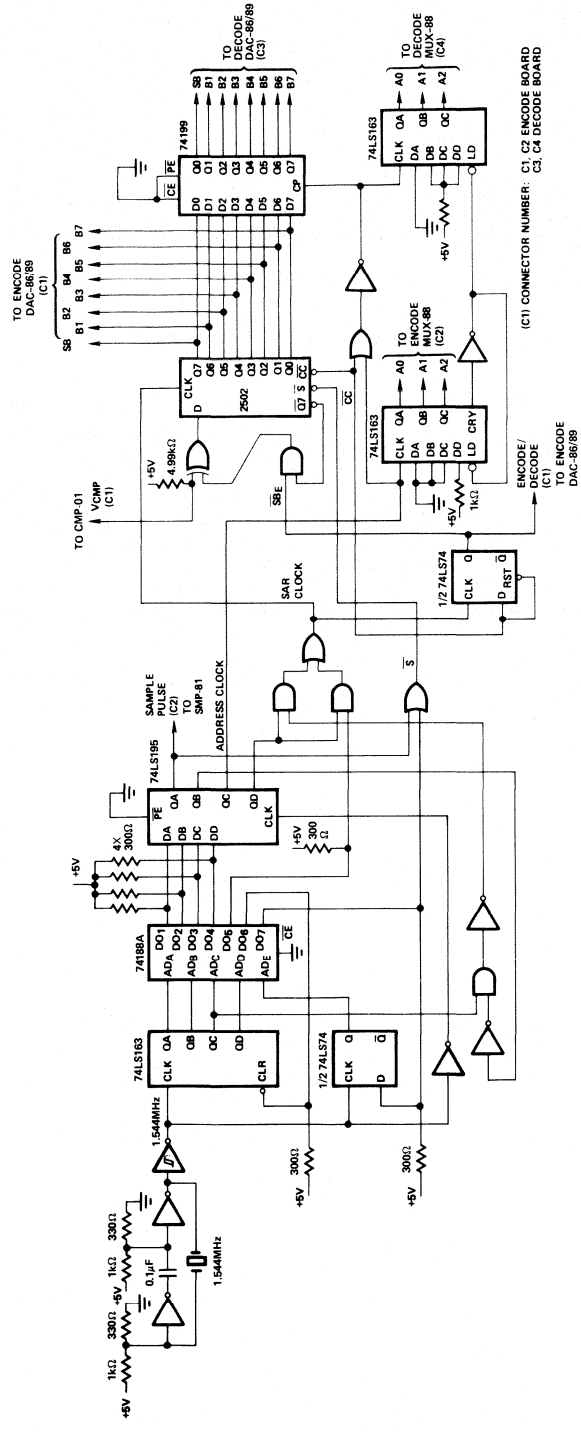

Figure A-1. Circuit Schematic

8

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Four-Channel Shared CODEC

FOUR-CHANNEL SHARED CODEC

A four-channel CODEC assembled from LSI components is a cost-effective digital transmission system requiring a relatively small number of devices. The system makes use of a single COMDAC® companded DAC-86 or DAC-89 digital-to-analog converter for both encoding and decoding (see Figure 1). The timing of the circuitry is compatible with ATT and CCITT system specifications.

Each channel is sampled at the standard 8kHz rate. With four channels this allows approximately 31.2μs to encode the sampled analog input and to decode the received digital signal for the same channel. To simplify the timing system requirements equal amounts of time were allowed for encoding and decoding, thus permitting 15.6μs for the more critical encode portion of the cycle, The encode/decode clocking scheme for this CODEC was incorporated directly from a successful eight-channel CODEC system which has been published as ADI Application Note 37. One original feature of the four-channel design was the use of dual eight-channel multiplexer ICs to switch the four channels. This results in a system whose interchannel crosstalk is practically negligible. Crosstalk figures of −85dB have been observed. This article describes the design procedure and reviews the transmission characteristics of the completed system.

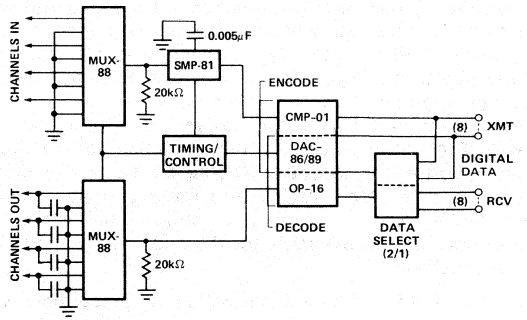

Figure 1. Four-Channel CODEC

CIRCUIT DESIGN

The analog circuitry required for a four-channel system is shown in Figure 2. The circuit uses the same printed circuit card as the transmit or receive sections of the eight-channel CODEC design (with the addition of a second multiplexer). The analog inputs all connect at the input multiplexer(MUX-88) using alternating inputs, the output (drain) of the MUX

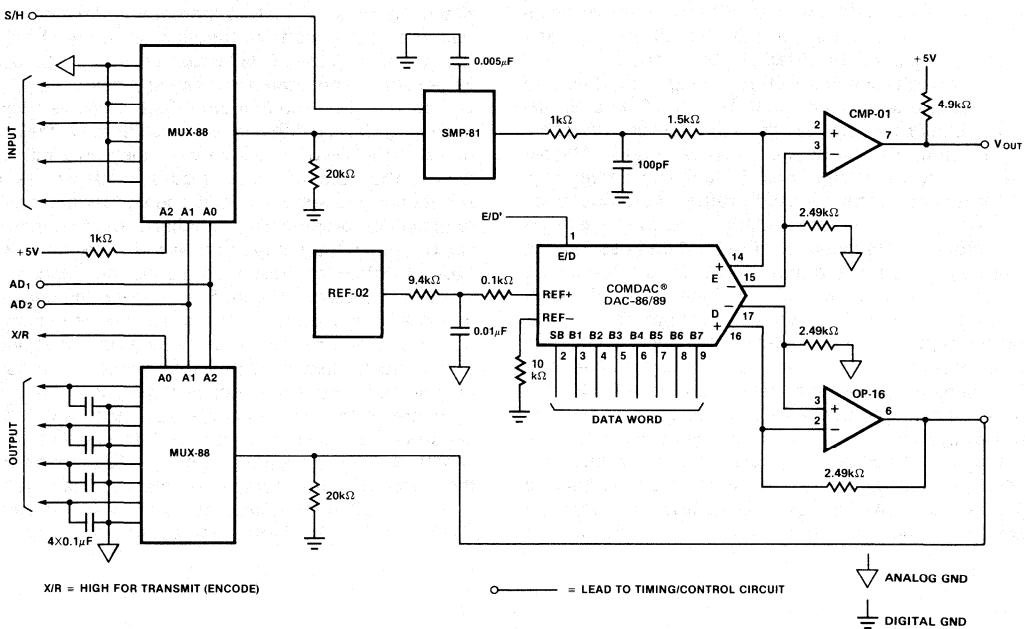

Figure 2. Four-Channel CODEC — Analog Board

COMDAC is a registered trademark of Analog Devices, Inc.

drives the SMP-81 sample-and-hold device. Once the input level is held, the COMDAC® (DAC-86 or 89), in conjunction with the comparator (CMP-01), begins the analog-to-digital conversion sequence. A successive approximation encode procedure is used; this generates, within the allotted conversion time (15.6μs), an eight-bit digital approximation of the analog level. The data byte is available at the successive approximation register output for the next 15.6μs time frame. During this time the CODEC decodes the incoming digital signals.

The decode cycle begins as soon as the encode cycle is completed. The DAC-86 is switched to the decode mode and an eight-bit data word is presented at its input leads, presumably from some distant analog-to-digital converter through a switching matrix. As the DAC-86 is switched to decode, the operational amplifier (OP-16) converts the output current of the DAC-86 into the appropriate voltage level. The output multiplexer is switched to the proper analog port as the decode procedure is initiated. On the output leads of the "de"-multiplexer, a hold capacitor is used to provide an output holding function. The "staircase" waveform is then available for filtering and the final subcriber interface.

As shown in the Figure 2, several circuit precautions were taken to reduce the internal noise levels. Foremost among these is the ample use of grounding throughout the analog circuit. All power supply inputs to the ICs are bypassed with capacitance (0.1μF) to ground. In addition, any spare land area of the board is filled with ground paths. The various voltage return paths are kept separate except for one common location on the board at the supply input leads. For further noise protection, a 20kΩ resistor to ground is connected to the drain leads (the common output or input) of both multiplexers. This reduces the multiplexer output noise and any crosstalk voltage feedthroughs. Also, in terms of the multiplexers, the output MUX is addressed to a grounded terminal (source connection) in between the active channels. This aids in minimizing the mutiplexer injection phenomena (discussed in detail in Application Note 37, describing the eight-channel CODEC design) and helps to further reduce the device crosstalk. The significance of using two multiplexers in this design is to allow unused channels of the output MUX to be connected to ground. The active channels are then alternated, in terms of addressing, with the grounded terminals. Addressing the multiplexers requires only two of the address leads to be controlled by a binary counter for data selection. The third address input is either held active (as in the input MUX), or can follow the system transmit and receive control signal (lead X/R, as in the output MUX). The address sequence for the input MUX is repeated through ports 4 to 7 (100 to 111), the most significant bit is held high and the two lower bits are counter outputs. This scheme allows maximum settling time for the MUX since the next channel to be encoded is selected more than 16μs prior to the conversion. For the output MUX, the two **most** significant bits of the address are the counter output leads, the least significant bit is the transmit/receive lead (X/R). As is shown in Figure 3, the address sequence alternates from active output port (even addresses) to

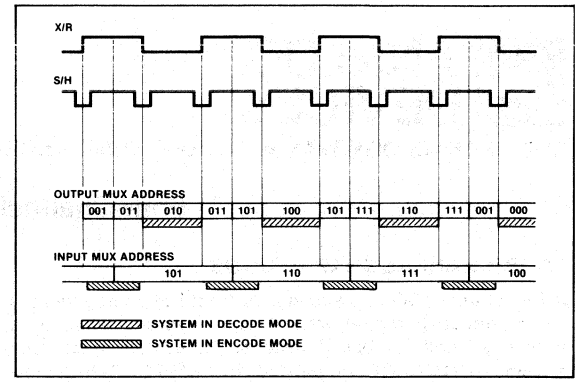

Figure 3. Four-Channel CODEC — Multiplexer Sequencing

grounded ports (odd addresses). The counter is changed while the MUX is selecting an unused (grounded) channel. This type of sequencing reduces the interchannel interference of the MUX and greatly adds to the system's measured performance.

To minimize sample-and-hold noise, a simple filter circuit is added to the output terminal of the device. A similar approach was used in the eight-channel design. Another feature in common with the eight-channel system is the use of a 4.9kΩ resistor pull-up from +5V to the output of the comparator; this decreases the switching time of the device for the encode procedure.

The timing waveforms generated for the four-channel system are based on the encoder clock used in the eight-channel CODEC. This clock circuit is shown in Figure 4. In a companded A/D conversion, using a successive approximation method, different digital bits require different settling times. There is a constant increase in the settling time required as one goes from the sign bit to the chord bits and then to the step bits. This increase in settling time is partially caused by the scaled current sources used to design a companding D/A and is most affected by the magnitude of the voltage level being converted. So that, as, the output currents of the DAC-86 become smaller (at smaller voltage levels), the comparator is less responsive. As the less significant bits are clocked in, the situation becomes more critical since progressively smaller current changes are provided to drive the comparator. A multi-frequency clock will take advantage of these timing variations; such a clocking scheme is shown in Figure 5. The governing design criteria is that a limited amount of time is available to complete the successive approximation of the analog signal and the optimum clock must fit within this period. All functions (including sampling, encoding, switching) must be completed, therefore, with an acceptable accuracy, within this time constraint. To achieve accurate data acquisition with the sample-and-hold, a sample pulse of 3.2μs was used. Once the pulse goes from sample to hold, there is a waiting time required to allow the sample-and-hold output to settle to the "held" value; 650ns is the time added. Since the sign bit is

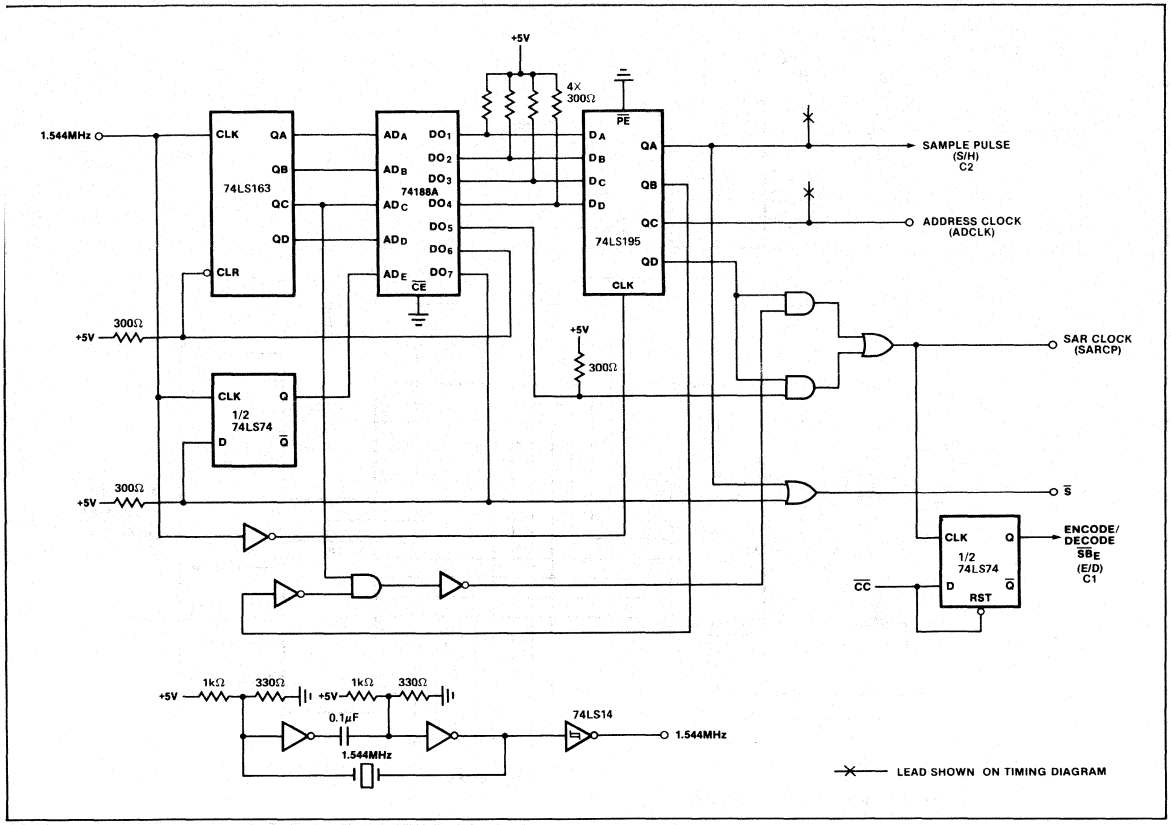

Figure 4. PROM Based System Clock

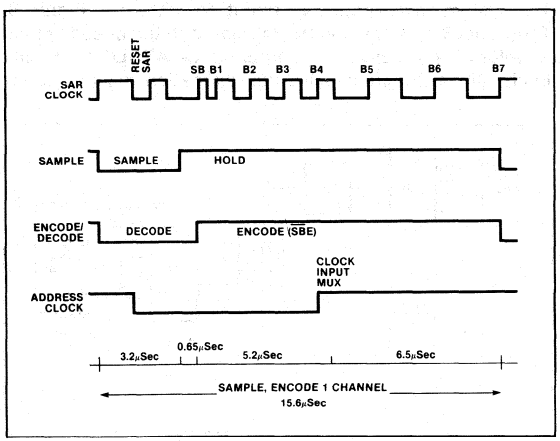

Figure 5. Four-Channel CODEC — Encoder Timing

the fastest transition, the basic system clock (1.544MHz) is used for the first SAR clock period. The clock frequency is then halved to clock in the chord bits (B1, B2, B3), the next "slowest" transitions. Finally, the clock frequency is halved again to allow for the step bits (B4-B7), a frequency of 386KHz. The SAR is reset during the sampling period and once the sample circuit has entered the hold mode, the input multiplexer's address is changed. This provides ample time for the MUX to switch to the next analog input with little effect on the adjacent waveforms.

Using this timing pattern as the starting point, the original eight-channel system was converted to a four-channel bi-directional design. The only additional control functions to be added were the timing signals needed to switch the DAC-86 between the encoding and decoding modes, to operate the output multiplexer, and to select the proper data inputs for the DAC-86. The DAC-86 mode select and the multiplexer address leads, as mentioned previously, are generated from a single system transmit/receive control (shown as X/R in Figure 6). The lead is the system monitor of the mode in which the CODEC is operating. When active (logic "1"), the DAC-86 and associated parts are in the encode mode. In the encode mode: the successive approximation register clock is enabled, the encode/decode lead to the DAC-86 (E/D) is enabled, the data input selector directs the SAR output back toward the DAC-86 for the feedback needed in successive approximation, and the output MUX connects the OP-16 to a grounded (unused) channel. The

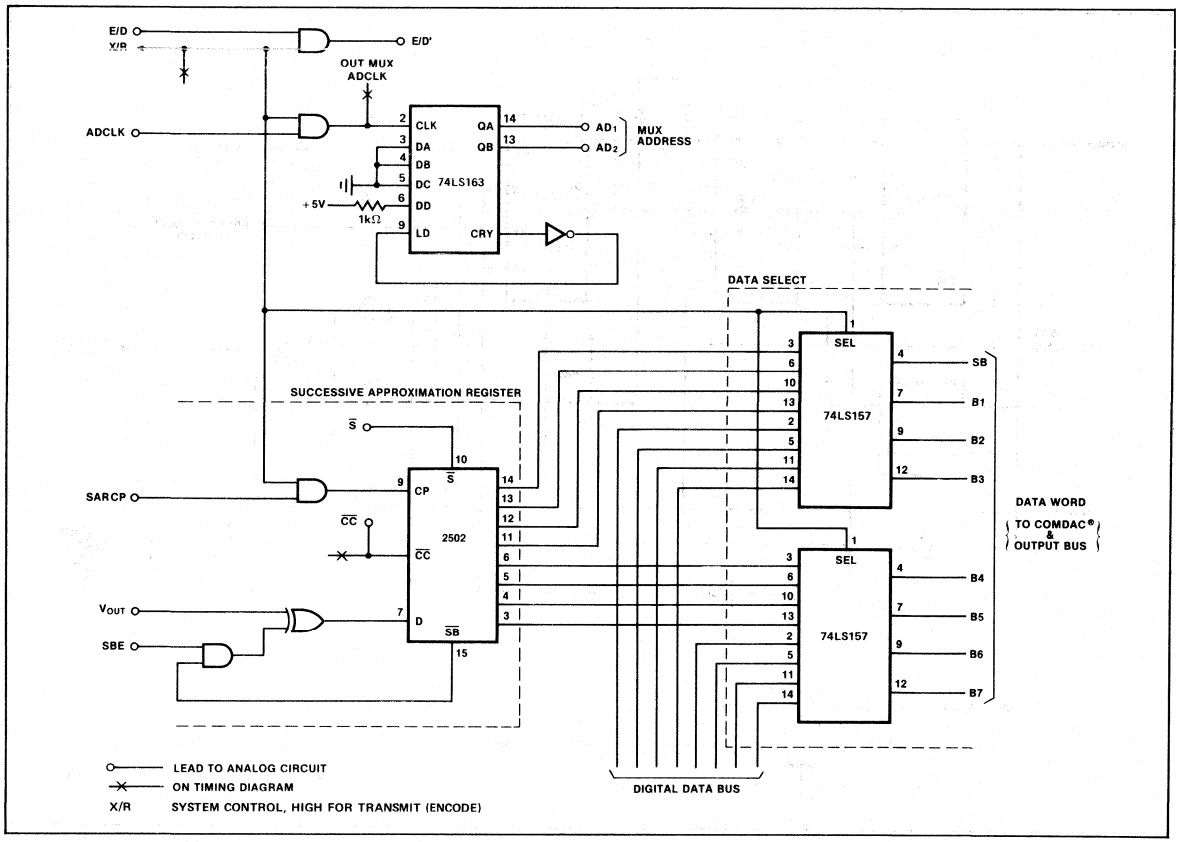

Figure 6. Four-Channel CODEC — Control Board

X/R lead remains at logic "1" until encoding is completed, then goes to ground (logic "0"), the decode state. To decode a data byte, the DAC-86 is held in the decode mode (E/D is low), the output MUX is addressed to an active output port, and the SAR clock is disabled (this register will hold the last encoded data word throughout the decode cycle — it is not cleared until a new input signal is to be encoded). The data selector is directed to the digital system bus and the decoding of the byte on the bus begins. As described, the address leads of the multiplexers are programmed such that the input MUX will always be directed toward the next active channel, once the previous analog sample has been held. But the output MUX does not connect to an active channel until the decode cycle begins; during the encode cycle only unused (grounded) ports are addressed.

SYSTEM TESTS

The system as configured in the block diagram (Figure 7) is a complete four-channel CODEC. To perform the transmission tests, it was decided to use a single CODEC circuit and transmit data in a "loopback" configuration. As was mentioned earlier, the only signal required from an external controller is the X/R lead. This is generated for the test circuit by halving the inverted ADCLK lead from the prom-based clock.

The system waveforms that result are shown in Figure 8. These clock patterns can be correlated to the encode clock waveforms in Figure 5 by comparing the ADCLK or the S/H leads. Since the successive approximation register is not

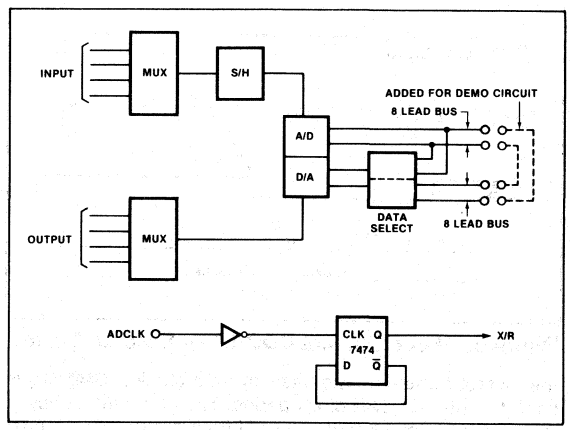

Figure 7. Four-Channel CODEC — Demonstrator Layout

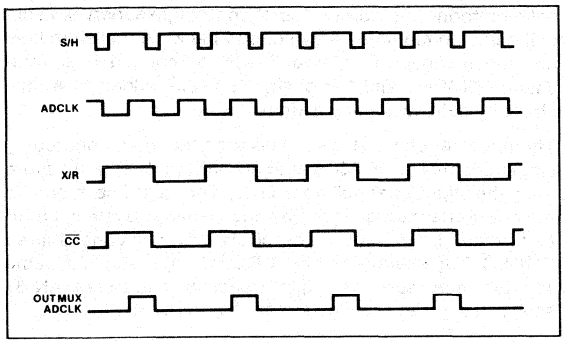

Figure 8. Four-Channel CODEC — Demonstrator Timing

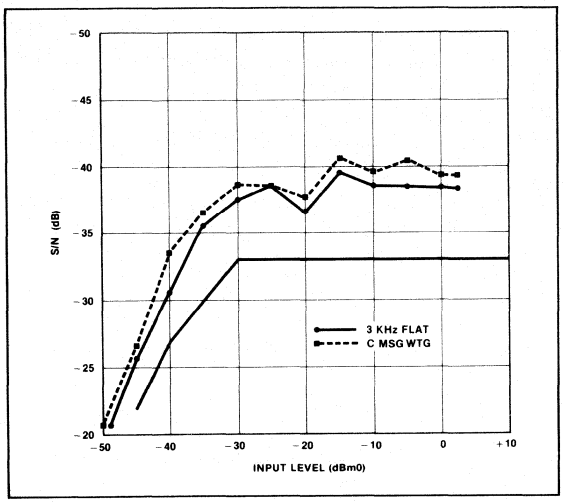

Figure 9. Signal-To-Total Distortion (Four-Channel)

cleared until after the decode cycle, an external register is not necessary to hold the data for the decoding process.

The transmission tests that were completed were the typical telephone network tests as described in the eight-channel application information. The tests include signal-to-total distortion, gain tracking, intelligible crosstalk and idle channel noise. Again the test method used for the first two tests was based on a sinusoidal input signal, as is common in the AT&T specifications, at a frequency between 400 and 3400Hz using a frequency-selective wave analyser. The results of all testing are shown in Figures 9 through 11.

It is of some interest to compare this data with the test results of the eight-channel CODEC design. In particular, the idle channel noise and the crosstalk measurements are improved. This can be partially explained by the different manipulations of the output multiplexer. This does however, tend to point to the fact that the output MUX can be a significant source of noise and cross channel interference. Further data is certainly necessary, but these results do point out an area of concentration for the system designer wanting to improve system performance.

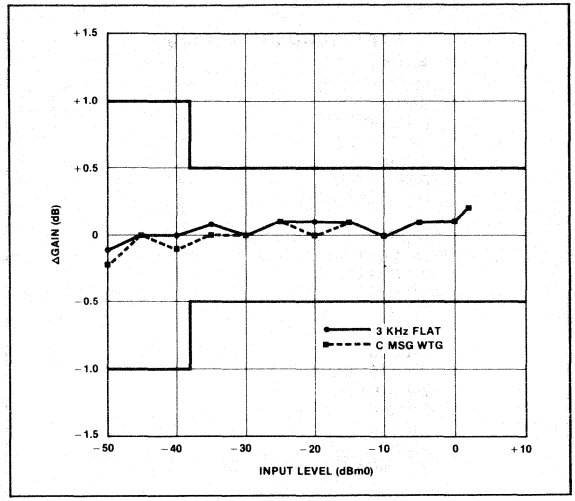

Figure 10. Gain Deviation (Four-Channel)

IDLE CHANNEL NOISE

Channel	
1	
2	no channel
3	had noise
4	level > 2dBrnc

INTELLIGIBLE CROSSTALK

f_{input}	Level
400-3400Hz	≤ −85dBmO

Figure 11. Transmission Measurements (Four-Channel)

In terms of signal-to-total distortion and gain tracking, the four-channel results compare favorably with the eight-channel data and both systems exceed the AT&T requirements. Overall, the transmission tests point out that using a single DAC-86 for four-channel transmitting and receiving is a realistic approach and can comply with all "system" standards.

CONCLUSIONS

The testing described in the preceding pages demonstrates the feasibility of encoding and decoding four channels with a single DAC-86. The system has several advantages: 1) a smaller number of devices are required to complete the CODEC function than were necessary for the eight-channel design, 2) the clock circuitry (PROM-based timing generator) is common to all encoders, so only a single such circuit is needed for multiple CODECs. Both of these factors contribute to reduced printed circuit board area for multiple transmission channels. The devices needed for a four-channel CODEC are listed in Figure 12. In terms of package sizes, only one device is larger than sixteen pins (the DAC-86/89 is 18 pins) and three of the components are only eight pins. This should make system layout fairly simple and allow relatively dense component packing. If chan-

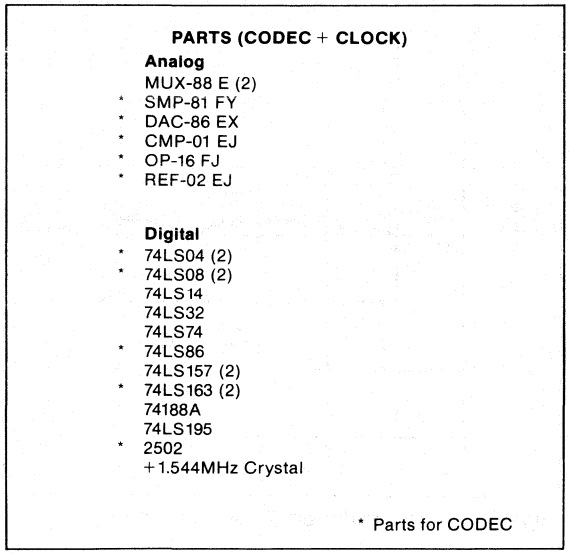

Figure 12. Four-Channel CODEC — Parts

nel monitoring is incorporated, then a single supervision circuit could administer several circuit packs in a system line-up. The number of external leads is reduced in a four-channel CODEC and the design is easily added to a bus-structure data switching system.

The price per channel is still less than that being quoted by single-channel CODEC designers although slightly more than the eight-channel approach. The sacrifice made in price-per-channel is offset by the gains in system architecture and board layout offered by a four-channel shared CODEC. The shared channel CODEC approach is a viable solution to producing a digital transmission system. AN-37 shows designs at even lower per channel cost.

ANALOG DEVICES

AN-39
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Companding D/A Converter

INTRODUCTION

A companding digital-to-analog converter (DAC) is the key component in PCM CODEC systems. (CODEC is an acronym for coder-decoder.) A CODEC performs the coding functions which consist of an analog-to-digital conversion (ADC) of the input analog (voice) signal and decoding, which consists of a digital-to-analog conversion (DAC) of the received digital input.

The DAC is used for both encoding and decoding; it is in a feedback loop to generate the ADC functions. Voice signals in telephony require a system with a very large dynamic range. The dynamic range (DR) of a CODEC is defined as the ratio of the largest resolvable signal to the smallest signal which can be encoded. The dynamic range of the CODEC is the same as that of the DAC used in either the decode mode or in the feedback loop of the successive approximation type ADC. The dynamic range of a DAC is simply the ratio of its output for a linear input of one least significant bit (LSB) to that of the largest, all "1s," input. This ratio is usually expressed in decibels using the equation:

$$DR = 20 \log_{10} \frac{I_{MAX.}}{I_{LSB}}$$

where for a current output DAC I_{MAX} is the output current for all "1s" input and I_{LSB} is the output current for one LSB input. Using this equation a linear bit DAC can be shown to resolve a ratio of 2^n:1 therefore:

$$DR = 20 \log_{10} \frac{2^n}{1} \approx 6^n$$

The wide dynamic range requirements of a telephone system require the equivalent dynamic range of a 12-bit system or 72dB. However, this system would not be satisfactory for telephone voice transmission because of its excessive bandwidth requirements. With present day T1 type transmission systems a 64kbits/sec data rate is required to transmit each voice channel. The use of the linear system would increase this bit rate to 96kbits/sec. This would provide more accuracy than is needed at the expense of excessive bandwidth.

For voice systems the most important criterion is the signal-to-noise ratio. In a PCM system noise is due almost entirely to quantizing distortion. Thus, a non-linear DAC has a non-linear transfer characteristic to compress the analog signal into a digital word and a complementary transfer characteristic to expand the digital words into analog signals with a wide dynamic range. For a telephone system a CODEC requires a fairly uniform signal-to-distortion ratio over its entire dynamic range. Achieving this uniform signal-to-distortion ratio over a wide dynamic range requires the use of non-uniform coding. A non-uniform CODEC is a coder-decoder pair whose input amplitude range is divided into steps of unequal widths, such that the width of the quantizing steps increase in proportion to the amplitude of the signal. To achieve uniform signal to distortion performance a logarithmic transfer function is required. The word compand, (compand is an acronym for compress — expand) was borrowed from analog systems to describe this non-uniform coding system where quantizing and coding is such that step size depends on the input amplitude.

COMPANDING PRINCIPLES

Companding requirements differ for different signal distributions. As mentioned above, voice signals require constant S/D performance over a wide dynamic range. In order to accomplish this the distortion must be proportional to the signal level. This feat is best achieved by the use of a logarithmic compression law. However, a truly logarithmic assignment of code words is not physically possible since this implies an infinite number of codes. Two methods for generating practical implementations of logarithmic transfer functions have been derived which have become industry standards. These methods are generally known by their transfer functions which are called μ-law and A-law respectively. Both of these transfer functions are normally implemented with eight-bit non-linear DACs to achieve a 72dB dynamic range. This is the equivalent dynamic range of a twelve-bit linear DAC. The μ-law and the A-law transfer functions are described by the following equations:

$$\mu\text{-law} \quad Y = \frac{\ln(1 + \mu|X|)}{\ln(1 + \mu)} \operatorname{sgn} X \quad \text{for } -1 \le X \le +1$$

$$\text{A-law} \quad Y = \frac{1 + \ln A|X|}{1 + \ln A} \operatorname{sgn} X \quad \text{for } 1/A \le X \le 1$$

$$Y = \frac{A|X|}{1 + \ln A} \operatorname{sgn} X \quad \text{for } 0 \le X \le 1/A$$

These laws have unique signal-to-distortion characteristics for each value of μ and A respectively. At present ATT has settled on a value of μ equal to 255 and CCITT specifications use a value of A equal to 87.6. Substituting these constants into the original equation above obtain:

μ-law $Y = 0.18 \ln(1 + \mu|X|) \quad \operatorname{sgn} X$ for $-1 \le X \le 1$

A-law $Y = 0.18 \ln(1 + \ln|X|) \quad \operatorname{sgn} X$ for $1/A \le |X| \le 1$

$Y = 0.18\, A|X| \quad \operatorname{sgn} X$ for $0 \le |X| = 1/A$

The wideband (unfiltered) signal-to-distortion ratio over the useable dynamic range of voice transmissions is shown in Figure 1. This plot does not represent actual system performance; it is instead, a measure of the distortion which would be caused by an ideal quantizer.

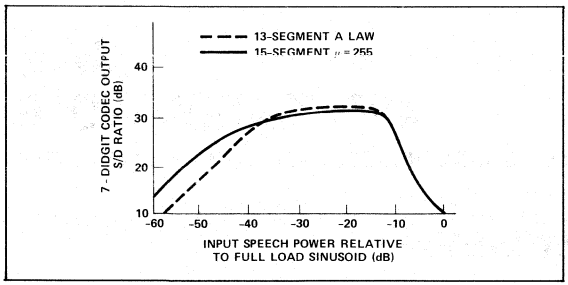

Figure 1. Input Speech Power Relative to Full Load Sinusoid (dB)

The practical implementation of the two transfer functions is accomplished by standardized piece-wise linear approximations. The transfer functions are implemented in chords or segments where the transfer function within any one chord is a linear staircase. Each chord has sixteen steps and the size of the step in each succeeding chord is double the size of the step in the preceeding chord. There are normally eight chords numbered zero through seven in both μ-law and A-law characteristics. For the A-law function the first two chords on either side of the origin have equal step sizes, whereas, for the μ-law function, the second chord after the origin has a step size which is double that of the first. For all remaining chords the steps double in size for each succeeding chord. This applies to both the μ-law and A-law functions. For the A-law function the four chords about the origin can be considered as a single segment so that the A-law characteristic is sometimes referred to as being a "13-segment" code. The A-law characteristic also differs from the μ-law characteristic in the manner in which the transfer function crosses the origin. The X-axis origin for the μ-law is at "mid-step" while the X-axis origin for the A-law is coincident with a "riser". This can be understood better from the "blow-ups" about the origin of Figures 2 and 3.

In order to obtain the best implementations of the transfer function, companded DACs are constructed such that encode and decode functions are offset by one-half step. With this technique the quantizing band for the encode DAC will be centered about the decode value. This can be seen in Figure 4, where the μ-law characteristics about the origin are shown. (The A-law characteristics would be identical except for the "mid-riser" phenomena at the origin.) As an example suppose that, for Figure 4, an analog input whose amplitude lies between levels 2 and 4 is being encoded. The best quantizing code to assign to this entire quantizing band is its mean value of 3. Thus the DAC used in the suc-

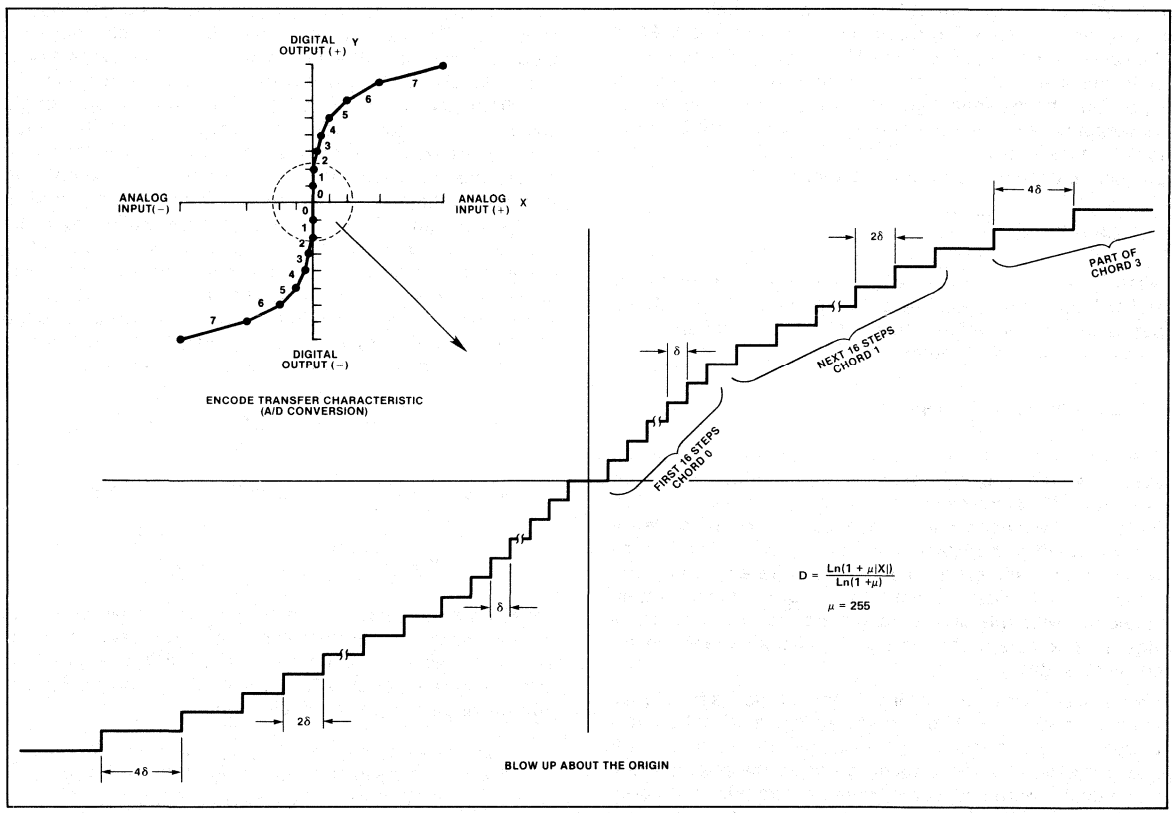

Figure 2. μ-Law Transfer Function

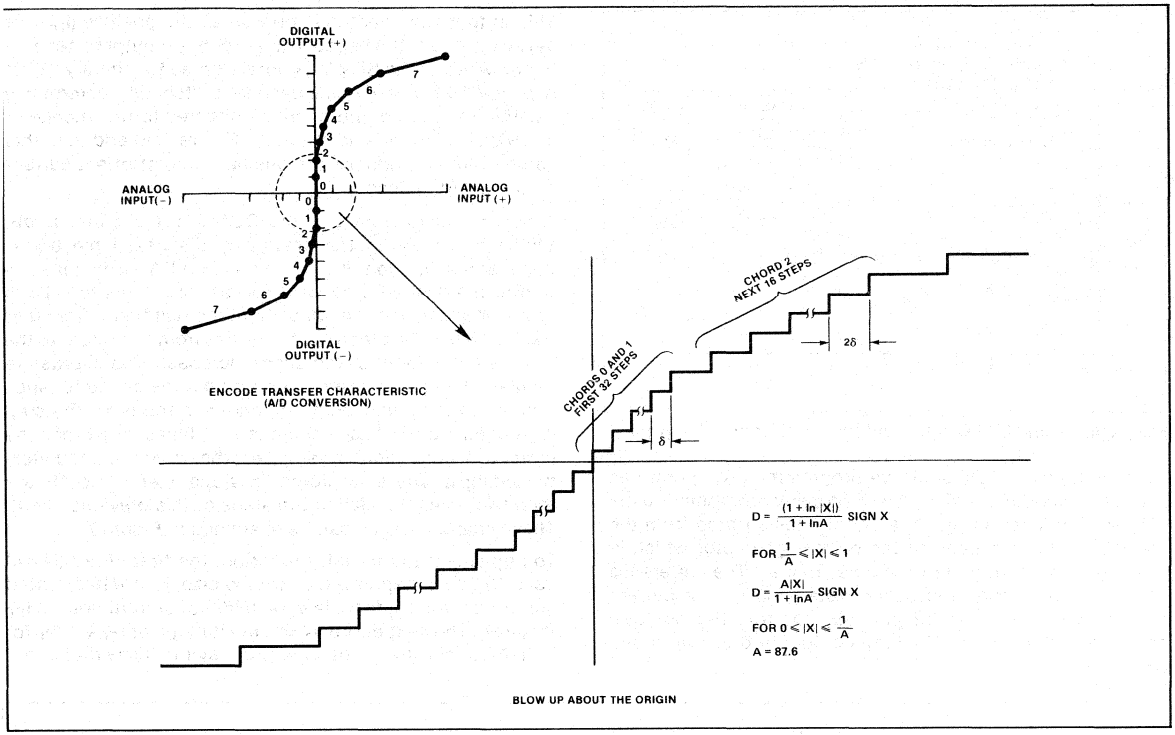

Figure 3. A-Law Transfer Function

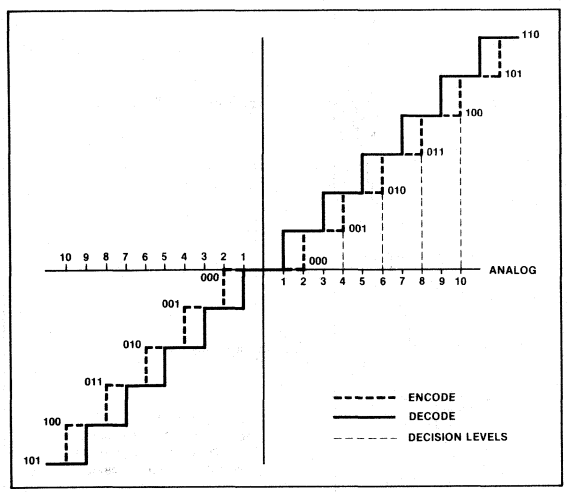

Figure 4. μ-Law Encode/Decode Characteristics About the Origin

cessive approximation feedback loop of the encode has output levels which represent the quantizing band edges. These can be referred to as decision levels. On the other hand the DAC for the decoder has output levels which represent the mean values of the quantizing bands which must, of necessity, be centered about the decoder output values. The end result is that a DAC used for decoding must be offset one-half step from the DAC used for encoding. This situation must exist over the entire range of the CODEC. A transmission system implemented with companding DACs is shown in Figure 5.

COMDAC® SYSTEM DESCRIPTION

A block diagram of ADI's companding DAC is shown in Figure 6. A single current output DAC is used to generate outputs for either the encode of decode mode of operation.

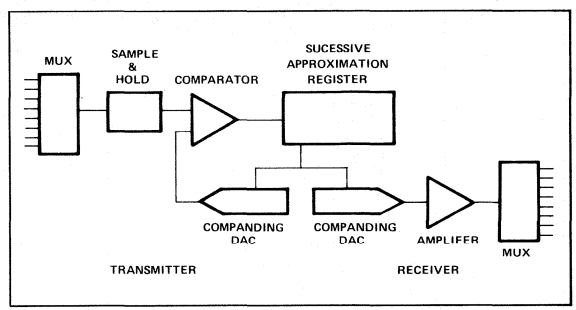

Figure 5. Transmission System Implemented with Companding DAC

COMDAC is a registered trademark of Analog Devices, Inc.

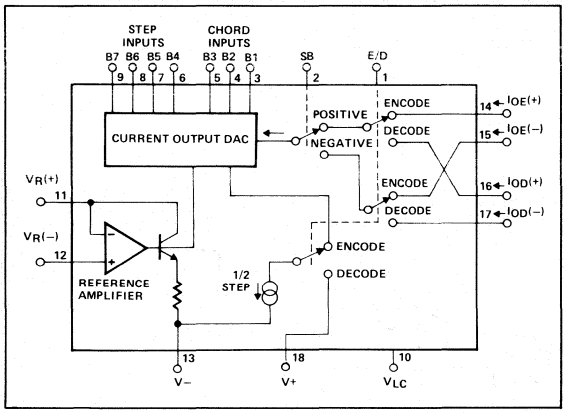

Figure 6. Equivalent Circuit and Pin Connection Diagram

Each companding DAC can be programmed to operate as either an encoder or a decoder by properly programming the E/D pin. The encode mode is offset one-half step from the decode mode by means of the current generator which is switched in during the encode mode. The reference amplifier establishes the current reference for the current output DAC. The sign bit pin (SB) controls the positive-negative switch which directs the output of the current output DAC.

This output will eventually end up at the positive (I_{OE+} or I_{OD+}) outputs or the negative (I_{OE-} or I_{OD-}) outputs depending on whether the SB pin is programmed to a binary "1" or a binary "0". The encode-decode switch E/D determines whether the DAC output shall be directed to the encode or decode terminations as shown in Figure 6. In addition, this same switch introduces the one-half step of offset current required during encode.

A better understanding of the COMDAC® circuitry is obtained by reviewing the previously discussed piece-wise linear approximation of the companding DAC transfer function to the desired μ-law or A-law transfer functions. Each chord or segment consists of 16 steps numbered from 0 to 15. The size of the steps double in size from one chord to the next as the number of the chord increases. The chords are numbered 0 to 7. In order to smooth out the characteristics during the transition from one chord to the next, the step current for step 0 of each chord is 1-1/2 times larger than the current of the highest step of the chord immediately preceding it. The succeeding 15 steps (steps 1 to 15) are then two times the size of the steps of this previous chord. These characteristics can be examined in Figure 7.

To implement the transfer function, the first chord (N = 0) uses 16 equal steps each of whose size, I_0 is 1/16 of chord current source I_{C0} for A-law, or 1/16.5 of current source I_{C0} for μ-law. The next chord, N = 1, must begin at $I_{C0} + 1.5I_0$ for both A-law or μ-law. Another way of saying this is that chord

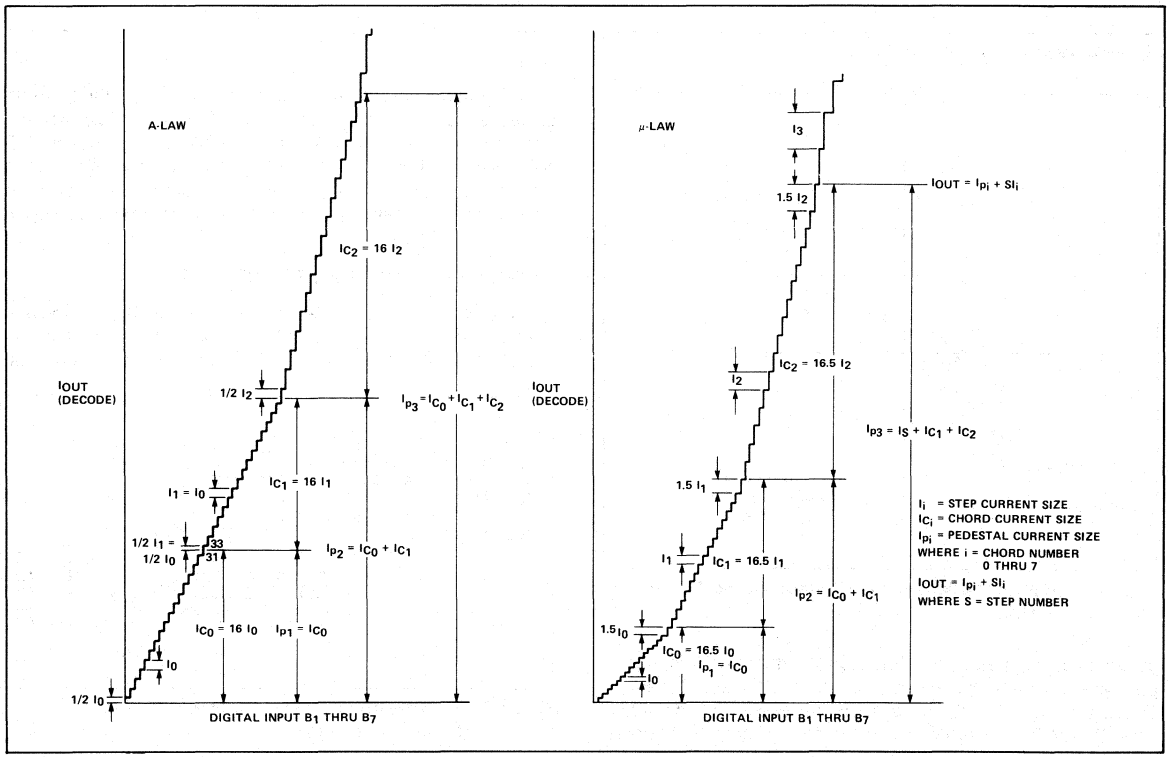

Figure 7. Construction of the Companding DAC Transfer Function

N = 1 begins 16.5 steps from the origin. In order to accomplish this a pedestal current must be directed toward the output whose magnitude is equal to $I_{C0} + 1.5I_0$. Chord C2 begins at $I_{C0} + 1.5I_0 + I_{C1} + 1.5I_1$ and ends at $I_{C0} + 1.5I_0 + I_{C1} + 1.5I_1 + I_{C2} + 1.5I_2$ and so forth. This process continues with pedestal currents for each chord number N described by the equation:

$$I_{PN} = \sum_{i=0}^{N-1} (I_{Ci} + 1.5I_i) = 16.5 \sum_{i=0}^{N-1} I_i$$

note that $I_{PO} = 0$.

A functional diagram of a companding DAC which implements the proper transfer function discussed above is shown in Figure 8, which operates in the following manner:

16.5 step currents (16.0 steps for A-law) where a step current is equal to the current step caused by changing the least significant bit in the chord of interest. Note that this satisfies the requirement of the equation for pedestal current I_{PN}. The step generator has the ability to sum current I_E into the output mode to provide the one-half step offset required when the system is operating in the encode mode. This one-half step offset current is controlled by the E/D pin. The system is in the encode mode when the E/D pin is biased to a binary "1".

DETAILED CIRCUIT DESCRIPTION

All of the single pole double throw switches in Figure 8 are constructed of bipolar emitter coupled transistors. One such switch is shown as an example in Figure 9. When the

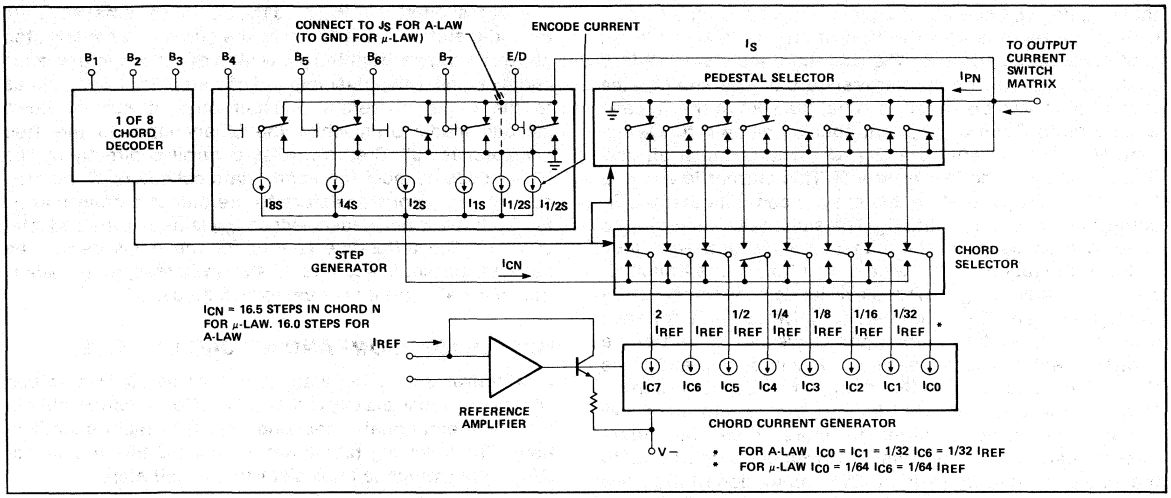

Figure 8. COMDAC® Companding DAC Functional Diagram

the reference amplifier sets the bias current for the chord generator by means of I_{C7} which is a current mirror whose output is equal to $2I_{REF}$. Next, due to the operation of an R−2R ladder which is described in a following paragraph, I_{C6} is made equal to one-half I_{C7} and is therefore equal to I_{REF}. I_{C5} is made equal to one-half I_{C6} and so forth. From I_{C3} down to I_{C0} a slave ladder is used rather than an R−2R ladder but the results are the same. The chord currents double in size progressing from I_{C0} to I_{C7} respectively (for A-law however $I_{C1} = I_{C0}$). The chord selector is programmed from the 1 of 8 decoder so that the chord identified by binary chord number N on leads B_1 to B_3 will switch I_{CN} to the step generator. All other chord currents are switched to the pedestal selector. The pedestal selector is programmed from the same 1 of 8 chord decoder such that chords I_0 to I_{N-1} are switched to the pedestal selector output in order to generate pedestal current I_{PN}. All other chord currents are switched to ground so that a pedestal current equal to the sum of the chord currents from I_{C0} to $I_{C(N-1)}$ will be directed to the output current switch matrix as I_{PN}. The I_{CN} flowing into the chord selector from the step generator is equal to

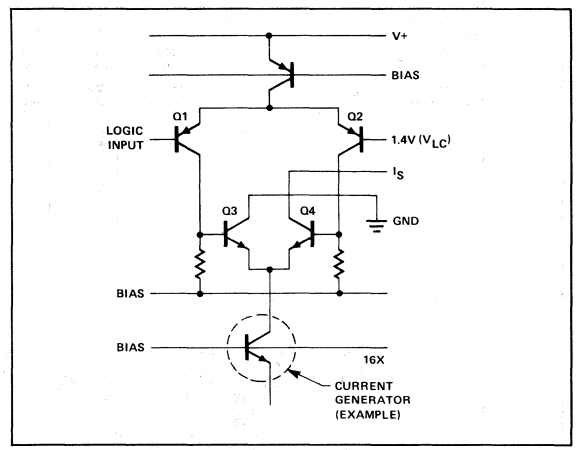

Figure 9. Double-Pole Double-Throw Switch Implemented with Emitter Coupled Transistors

logic input exceeds the logic level bias V_{LC} Q_1 is turned off and Q_2 is turned on. In turn Q_3 is turned off and Q_4 is turned on thus effectively switching the current generator, shown as an example, from the ground to I_S. Conversely, lowering the logic level input below V_{LC} will switch the current from I_S to ground. The V_{LC} Control permits the circuit to interface with a large range of logic levels.

The chord current generator circuit is shown in Figure 10. This circuit is the implementation of the chord current generator previously discussed. Q_0 is forced to operate at the reference input current I_{REF} and Q_1, with an emitter resistor one-half the size of the emitter resistor of Q_0, will then operate at $2I_{REF}$. Q_2 through Q_4 will operate at progressively smaller currents where each transistor operates at one-half the current of the transistor to its immediate left. To review this normal R–2R current-ladder function notice that Q_{4A} and Q_{4B} operate at equal currents and that the sum of their currents is equal to that of one transistor with an emitter resistor equal to R. When the series resistor R is added to the junction of the emitter resistors of Q_{4A} and Q_{4B} the current of Q_3 will be forced to equal the sum of the Q_{4A} and Q_{4B} currents. Thus Q_{4A} current equals one-half the Q_3 current. Now the current from Q_{4A}, Q_{4B} and Q_3 must all flow through the next series resistor R. This current is equal to twice that of Q_3; therefore it is easy to compute that the Q_2 current is twice that of the Q_3. The same reasoning may be used to proceed down the ladder to show that each transistor in the ladder sinks twice the current of the transistor on its immediate right. The slave ladder consisting of Q_5 through Q_{8A} and Q_{8B} continues to halve currents for each transistor proceeding to the right. However this part of the chord current generator uses scaled resistors instead of the R–2R ladder technique. Since Q_{4B} sinks constant current from the slave ladder, and since all the current must flow through the scaled emitter resistors, then the curent through each transistor must be inversely proportional to the size of its emitter resistor. By examination of the slave ladder it can be seen that each transistor proceeding to the right sinks one-half the current of the transistor to its immediate left. For the μ-law chord current generator Q_{8B} is simply diode connected such that the chord current for chord C_0 is roughly one-half the current of chord C_1. For the A-law chord current generator, however, the collectors of transistors Q_{8A} and Q_{8B} are tied together so that I_{C0} is exactly equal to I_{C1}. The currents flow to the chord current generator from an array of bipolar single pole double throw switches labeled "chord selector" in Figure 8. The actual switches are not shown in this paper.

The Step Current Generator is shown in Figure 11. Again the single pole double throw switches which connect the step generator to the output current matrix as shown in the companding DAC functional diagram are not represented. The step generator is connected to the chord selector which sinks I_{CN}. Ratioed emitters are used to divide the current. The largest emitter is 16 times the size of the smallest emitter and therefore sinks 16 times the current. The A-law step generator differs from the μ-law step generator in that each chord begins with a riser instead of a step. This also applies to the origin, therefore one-half step of current flows (decode mode) even when the binary input to the step generator is "0". Step switches controlled directly by the binary code connect the appropriate collectors of the step current generator transistors to the output current matrix. For both A-law and μ-law devices I_{CN} is one of the pedestal currents. The difference is that for the A-law device the pedestal current is equal to 16 steps whereas, for the μ-law, the pedestal current is equal to 16.5 steps.

NORMALIZED COMPANDING DAC OUTPUTS

It is convenient to generate tables of normalized values which correspond exactly to the CCITT (Consultative Committee for International Telephone and Telegraph) specifications. The following tables are normalized to the smallest DAC output which is equivalent to one-half step.

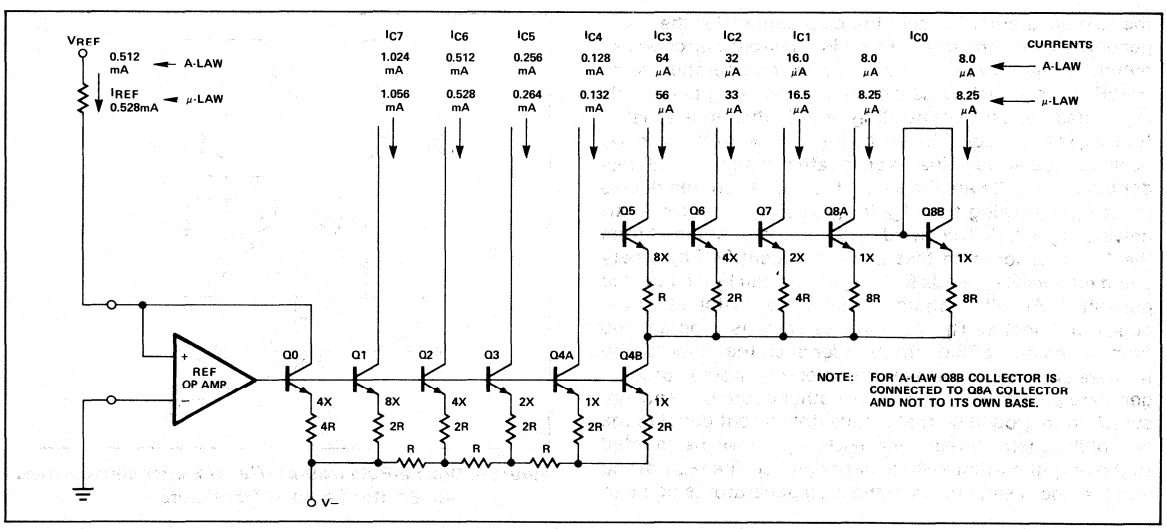

Figure 10. Chord Current Generator Diagram

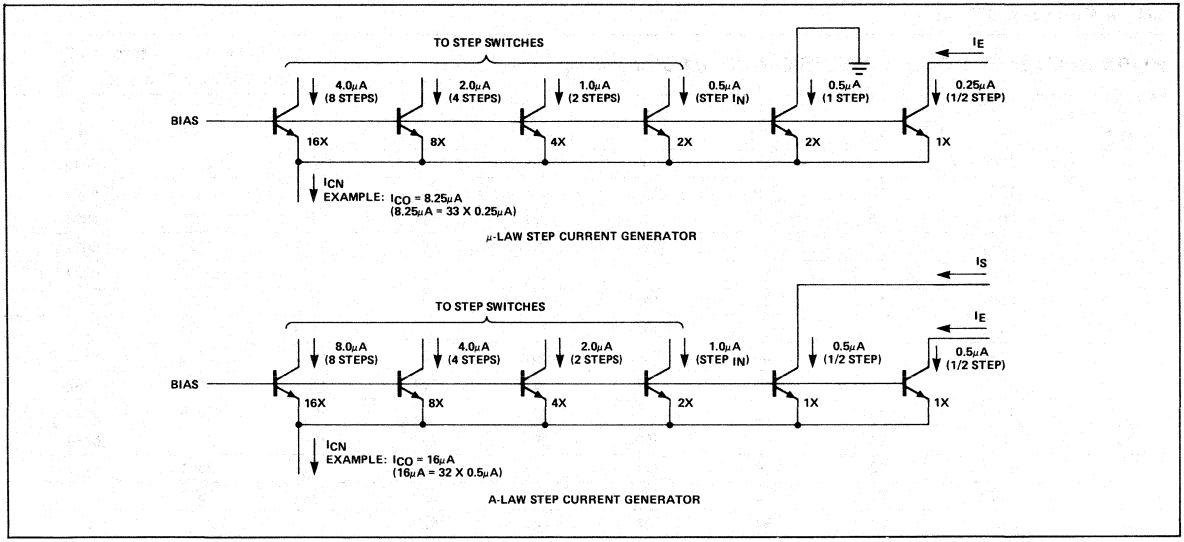

Figure 11. A-Law and μ-Law Step Current Generators

8

μ-Law Normalized Table

NORMALIZED DECODE OUTPUT (SIGN BIT EXCLUDED) $I_{C,S} = 2[2^C (S + 16.5) - 16.5]$

C = chord no. (0 through 7)
S = step no. (0 through 15)

STEP \ CHORD		0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	0000	0	33	99	231	495	1023	2079	4191
1	0001	2	37	107	247	527	1087	2207	4447
2	0010	4	41	115	263	559	1151	2335	4703
3	0011	6	45	123	279	591	1215	2463	4959
4	0100	8	49	131	295	623	1279	2591	5215
5	0101	10	53	139	311	655	1343	2719	5471
6	0110	12	57	147	327	687	1407	2847	5727
7	0111	14	61	155	343	719	1471	2975	5983
8	1000	16	65	163	359	751	1535	3103	6239
9	1001	18	69	171	375	783	1599	3231	6495
10	1010	20	73	179	391	815	1663	3359	6751
11	1011	22	77	187	407	847	1727	3487	7007
12	1100	24	81	195	423	879	1791	3615	7263
13	1101	26	85	203	439	911	1855	3743	7519
14	1110	28	89	211	455	943	1919	3871	7775
15	1111	30	93	219	471	975	1983	3999	8031
STEP SIZE		2	4	8	16	32	64	128	256

μ-Law Normalized Table

NORMALIZED ENCODE LEVEL (SIGN BIT EXCLUDED) $I_{C,S} = 2[2^C (S + 17) - 16.5]$

C = chord no. (0 through 7)
S = step no. (0 through 15)

STEP \ CHORD		0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	0000	1	35	103	239	511	1055	2143	4319
1	0001	3	39	111	255	543	1119	2271	4575
2	0010	5	43	119	271	575	1183	2399	4831
3	0011	7	47	127	287	607	1247	2527	5087
4	0100	9	51	135	303	639	1311	2655	5343
5	0101	11	55	143	319	671	1375	2783	5599
6	0110	13	59	151	335	703	1439	2911	5855
7	0111	15	63	159	351	735	1503	3039	6111
8	1000	17	67	167	367	767	1567	3167	6367
9	1001	19	71	175	383	799	1631	3295	6623
10	1010	21	75	183	399	831	1695	3423	6879
11	1011	23	79	191	415	863	1759	3551	7135
12	1100	25	83	199	431	895	1823	3679	7391
13	1101	27	87	207	447	927	1887	3807	7647
14	1110	29	91	215	463	959	1951	3935	7903
15	1111	31	95	223	479	991	2015	4063	8159
STEP SIZE		2	4	8	16	32	64	128	256

A-Law Normalized Table

NORMALIZED DECODE OUTPUT (SIGN BIT EXCLUDED) $I_{CS} = 2^{N-1}(33 + 2S)$ For $N > 0$; $I_{CS} = 2S + 1$ For $N = 0$

STEP \ CHORD		0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	0000	1	33	66	132	264	528	1056	2112
1	0001	3	35	70	140	280	560	1120	2240
2	0010	5	37	74	148	296	592	1184	2368
3	0011	7	39	78	156	312	624	1248	2496
4	0100	9	41	82	164	328	656	1312	2624
5	0101	11	43	86	172	344	688	1376	2752
6	0110	13	45	90	180	360	720	1440	2880
7	0111	15	47	94	188	376	752	1504	3008
8	1000	17	49	98	196	392	784	1568	3136
9	1001	19	51	102	204	408	816	1632	3264
10	1010	21	53	106	212	424	848	1696	3392
11	1011	23	55	110	220	440	880	1760	3520
12	1100	25	57	114	228	456	912	1824	3648
13	1101	27	59	118	236	472	944	1888	3776
14	1110	29	61	122	244	488	976	1952	3904
15	1111	31	63	126	252	504	1008	2016	4032
STEP SIZE		2	2	4	8	16	32	64	128

A-Law Normalized Table

NORMALIZED ENCODE DECISION LEVELS (SIGN BIT EXCLUDED)

$I_{CS} = 2^{N-1}(34 + 2S)$ For $N > 0$
$I_{CS} = 2S + 2$ For $S = 0$

STEP \ CHORD		0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	0000	2	34	68	136	272	544	1088	2176
1	0001	4	36	72	144	288	576	1152	2304
2	0010	6	38	76	152	304	608	1216	2432
3	0011	8	40	80	160	320	640	1280	2560
4	0100	10	42	84	168	336	672	1344	2688
5	0101	12	44	88	176	352	704	1408	2816
6	0110	14	46	92	184	368	736	1472	2944
7	0111	16	48	96	192	384	768	1536	3072
8	1000	18	50	100	200	400	800	1600	3200
9	1001	20	52	104	208	416	832	1664	3328
10	1010	22	54	108	216	432	864	1728	3456
11	1011	24	56	112	224	448	896	1792	3584
12	1100	26	58	116	232	464	928	1856	3712
13	1101	28	60	120	240	480	960	1920	3840
14	1110	30	62	124	248	496	992	1984	3968
15	1111	32	64	128	256	512	1024	2048	*4096
STEP SIZE		2	2	4	8	16	32	64	128

8

The numbers in these tables are directly proportional to the input reference current. However the exact relationship is somewhat complicated. A reference current of 528μA for the μ-law DAC will produce a step size of 0.5μA; thus, for the μ-law device driven by a reference current of 528μA, it is only necessary to multiply all the numbers in the normalized tables by one-half step or 0.25μA to obtain the output in μA. The table tabulated below corresponds to a 528μA reference.

μ-Law Current Output Table

IDEAL DECODE OUTPUT CURRENT IN MICROAMPS (SIGN BIT EXCLUDED)

STEP \ CHORD		0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	0000	0	8.25	24.75	57.75	123.75	255.75	519.75	1047.75
1	0001	0.5	9.25	26.75	61.75	131.75	271.75	551.75	1111.75
2	0010	1	10.25	28.75	65.75	139.75	287.75	583.75	1175.75
3	0011	1.5	11.25	30.75	69.75	147.75	303.75	615.75	1239.75
4	0100	2	12.25	32.75	73.75	155.75	319.75	647.75	1303.75
5	0101	2.5	13.25	34.75	77.75	163.75	335.75	679.75	1367.75
6	0110	3	14.25	36.75	81.75	171.75	351.75	711.75	1431.75
7	0111	3.5	15.25	38.75	85.75	179.75	367.75	743.75	1495.75
8	1000	4	16.25	40.75	89.75	187.75	383.75	775.75	1559.75
9	1001	4.5	17.25	42.75	93.75	195.75	399.75	807.75	1623.75
10	1010	5	18.25	44.75	97.75	203.75	415.75	839.75	1687.75
11	1011	5.5	19.25	46.75	101.75	211.75	431.75	871.75	1751.75
12	1100	6	20.25	48.75	105.75	219.75	447.75	903.75	1815.75
13	1101	6.5	21.25	50.75	109.75	227.75	463.75	935.75	1879.75
14	1110	7	22.25	52.75	113.75	235.75	479.75	967.75	1943.75
15	1111	7.5	23.25	54.75	117.75	243.75	495.75	999.75	2007.75
STEP SIZE		.50	1	2	4	8	16	32	64

*Virtual Decision Level

A similar exercise will yield a corresponding table for the A-law part. Multiplying all the numbers in the normalized A-law table, for instance, will produce a table of currents for a reference input of 512μA. A table based on 512μA reference current will have a step size of 1.0μA and is tabulated in the μ-law current output table.

A-Law Current Output Table

IDEAL DECODE OUTPUT CURRENT IN MICROAMPS (SIGN BIT EXCLUDED)

STEP \ CHORD		0	1	2	3	4	5	6	7
		000	001	010	011	100	101	110	111
0	0000	0.5	16.5	33	66	132	264	528	1056
1	0001	1.5	17.5	35	70	140	280	560	1120
2	0010	2.5	18.5	37	74	148	296	592	1184
3	0011	3.5	19.5	39	78	156	312	624	1248
4	0100	4.5	20.5	41	82	164	328	656	1312
5	0101	5.5	21.5	43	86	172	344	688	1376
6	0110	6.5	22.5	45	90	180	360	720	1440
7	0111	7.5	23.5	47	94	188	376	752	1504
8	1000	8.5	24.5	49	98	196	392	784	1568
9	1001	9.5	25.5	51	102	204	408	816	1632
10	1010	10.5	26.5	53	106	212	424	848	1696
11	1011	11.5	27.5	55	110	220	440	880	1760
12	1100	12.5	28.5	57	114	228	456	912	1824
13	1101	13.5	29.5	59	118	236	472	944	1888
14	1110	14.5	30.5	61	122	244	488	976	1952
15	1111	15.5	31.5	63	126	252	504	1008	2016
STEP SIZE		1	1	2	4	8	16	32	64

Reviewing the companding DAC functional diagram Figure 8 demonstrates the relationship between step size and I_{REF}. For a μ-law device I_{C0} equals 16.5 chord zero steps and for an A-law device I_{C0} equals 16 chord zero steps. I_{C6} is always equal to I_{REF} in either system. I_{C6} is then equal to 64 times I_{C0} for a μ-law system, and 32 times I_{C0} for an A-law system. The step size can then be related to I_{REF} by the following equations:

$$\text{step size} = I_{REF}/64 \times 16.5 = I_{REF}/1056 \text{ (μ-law)}$$
$$\text{step size} = I_{REF}/32 \times 16 = I_{REF}/512 \text{ (A-law)}$$

Now for a reference current of 528μA the step size for a μ-law system is 528/1056 or 0.5μA. For a reference current of 512μA the step size for an A-law system is 512/512 or 1.0μA. These values concur with those used to generate the tables.

In the design of the PMI DAC-89 the biasing resistors were not scaled to exactly integer values. This was done deliberately to standardize somewhat on 528μA input reference current for both A-law and μ-law parts. The performance of the device is not affected, however the actual scaling is somewhat complicated and will not be discussed in this paper.

Finally if encode output tables were desired for current output they could be obtained by scaling to proper step size the normalized encode tables or adding one-half step to each value in the decode table, where the step size depends on the chord number.

DAC ACCURACY

Companding DACs must be manufactured to satisfy a unique set of parameters. The performance of a companded DAC used for telephony must satisfy the requirements of a communication system on an end-to-end basis. A voice channel is first encoded by one CODEC then decoded by a second CODEC such that the system performance can be measured on an audio-in-audio-out basis. The CODEC performance will be almost completely dominated by the Gain Tracking requirement.

GAIN TRACKING

Gain Tracking refers to the ability of a system to track its input power level. The test is normally made with a system such as that shown in Figure 12.

Gain Tracking is measured by monitoring the input and output levels in decibels. At an input level of −10Bm0 the output is recorded as the output reference level. For ideal Gain Tracking, any change (in dB) of the input level must be matched exactly by the same change in the output level.

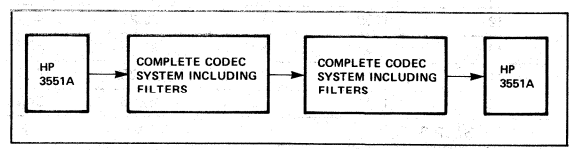

Figure 12. Gain Tracking or S/N Test

This condition is monitored over all input power levels of interest. The extent to which these power level changes differ again in dB) is a measure of Gain Tracking, also referred to as gain deviation. The ATT/D3 Gain Tracking specification is shown in Figure 13.

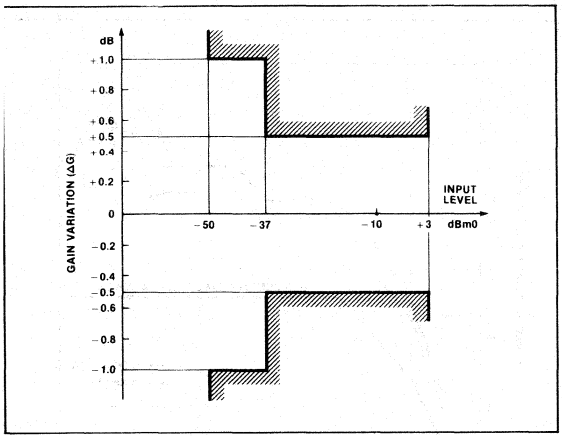

Figure 13. ATT/D3 Gain Tracking Specification

CCITT publishes two separate specifications for Gain Tracking. The apparatus used for making either of these tests is basically the same as that used in Figure 13 except that for the first part of the "method one" test the HP3551A would be replaced with a suitable white noise source at the input and an RMS reading voltmeter at the output. Gain Tracking masks equivalent to those found in CCITT publications are shown in Figure 14.

POWER LEVELS

For PCM channel performance measurements, power levels are characteristically expressed in dBm0. A reference level of 0dBm0 is established by referencing to a code in the **digital** transmission. The binary code pattern required to establish a reference level of 0dBm0 can be found in the CCITT publications. This pattern is reproduced in the PMI Telecommunications Handbook for the readers convenience. The constant repetition of these binary numbers at the normal sampling rate of 8kHz will produce a 1kHz sinusoid at a 0dBm0 reference level. Starting with this definition it can then be shown that a sinusoid whose peak value is just at the system saturation level (all "1s" PCM output) will have a power level of 3.14 and 3.17dBm0 for A-law and μ-law respectively.

SIGNAL-TO-DISTORTION MEASUREMENTS

Signal-to-Distortion is a measure of the total distortion a system will exhibit on an end-to-end basis. As with Gain Tracking this measurement is normally performed on an audio-to-audio basis. A typical setup for measuring Signal-to-Distortion is shown in Figure 15. A wideband (3kHz) filter may be substituted for the C-Message filter shown for some tests.

Figure 16 shows the ATT/D3 specification mask with the performance of a PMI demonstration COMDAC® based shared CODEC system superimposed. This method of measuring Signal-to-Distortion is applicable to either CCITT or ATT specifications.

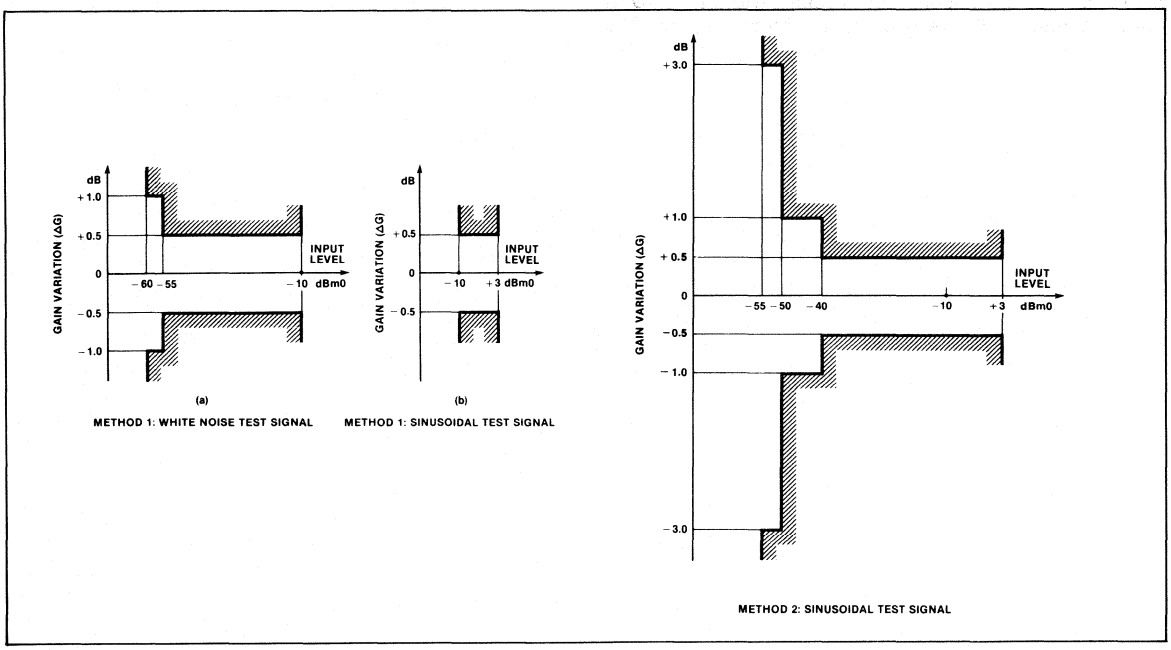

Figure 14. CCIT Gain Tracking Specification

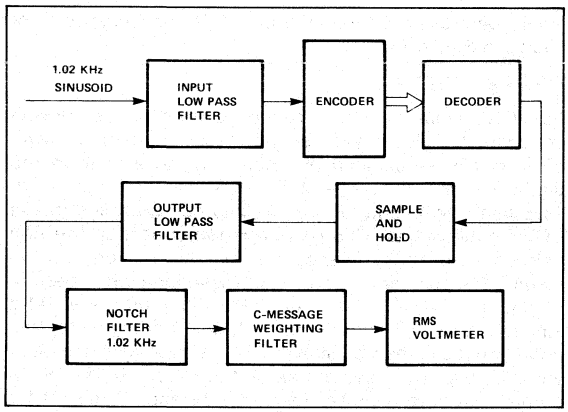

Figure 15. Signal-to-Distortion Test Setup

DAC ACCURACY VERSUS GAIN TRACKING AND SIGNAL-TO-DISTORTION RATIO

The analog portions of a PCM system usually make only a minor contribution to either Gain Tracking or Signal-to-Distortion errors. Thus, the major contribution to error is the inability of the companding DAC to accurately follow the encoding format. The process of quantizing and coding will cause some deviation from the ideal, however the errors made by the ideal CODEC system will be well within telephony specifications. To conform to the required Gain Tracking and Signal-to-Distortion specifications the DAC output currents must conform as closely as possible to the ideal transfer function as tabulated in the normalized tables. This corresponds to a specification of absolute error on the DAC output current with respect to its binary inputs. The DAC-86/89 companded DACs are guaranteed to plus or minus one-fourth step from ideal values in chord zero and to plus or minus one-half step elsewhere. This information can be transformed into tabular form by adding the allowable error to the DAC tables. Either the normalized tables or the current output tables can be used as a basis for this exercise.

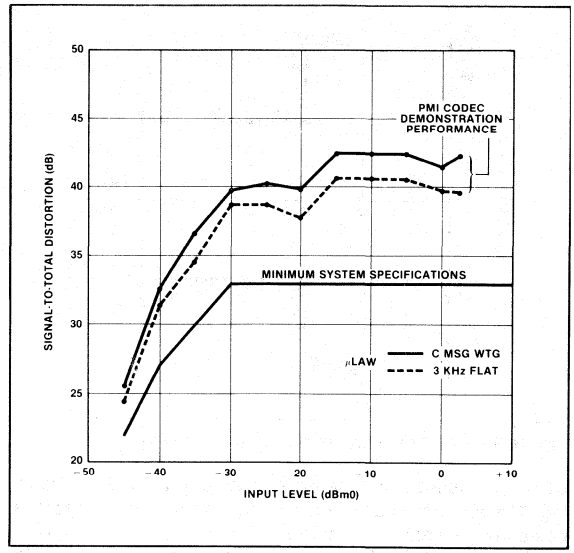

Figure 16. ATT/D3 Signal-to-Distortion Mask

ANALOG DEVICES

AN-41
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Improved Shared-Channel CODEC Design with Analog Devices' Companding DACs

DESIGNERS FACE CHOICE

Designers of telecommunications systems are faced with a fundamental design decision; they must base the design of digital voice transmission systems on either the use of a CODEC shared over several analog channels or on the use of one CODEC per channel.

Since 1976, users of ADI's shared-channel approach point to the economical advantages of a system that incorporates an encoder and a decoder capable of accommodating multiple voice channels. Proponents of single-CODEC-per-channel systems generally cite the straight-forward techniques involved in implementing this concept as the motivation for its selection.

The use of the shared-CODEC configuration is appealing because fewer integrated circuits per channel are required and less circuit board area per channel is needed. Because of its lower cost per channel and lower total system cost, the shared-channel CODEC concept looms as the logical choice for designers provided it can meet the performance needs of their systems.

Recently, new telecommunication components were introduced by Precision Monolithics Incorporated that improve the performance that can be obtained from a shared-CODEC system; the availability of these devices could influence future design decisions in favor of the shared CODEC concept over the single-CODEC approach. These newly developed devices, including the DAC-88 and the DAC-89 COMDAC® companding D/A converters, and the DMX-88 demultiplexer, not only provide performance which is superior to previously existing products, they are also easier to apply.

The degree of improvement offered by these devices and the present capabilities of a shared-CODEC system can be demonstrated using the circuit configuration shown in Figure 1. This configuration, an eight-channel digital transmission system, was designed and breadboarded by PMI as a vehicle for measuring the analog-input to analog-output transmission parameters commonly used to specify CODEC performance regardless of the particular system configuration.

Two of the components employed in the transmission system shown in Figure 1, the companding digital-to-analog converter and the analog multiplexer represent improved versions of previously existing products and are key contributors to the superior performance the system demonstrates over previous versions.

In the redesign of the companding DAC, it was felt that the best results would be achieved by improving the device's response within chord 0. Two design goals were set: to establish a reasonable settling time within chord 0 and to provide a guaranteed better than ±1/4 step linearity within that chord. Excellent results were obtained for both μ-law and A-law devices.

The new version of the companding DAC typically settles within 500ns, thus overcoming a restriction that had previously reduced the maximum number of channels for encoding. The IC's nominally settle to within ±1/8 step of the theoretical level in chord 0 and are 100% tested to be no worse than ±1/4 step.

Because of testing time restrictions, the settling time is given as a nominal specification. The guaranteed linearity specifi-

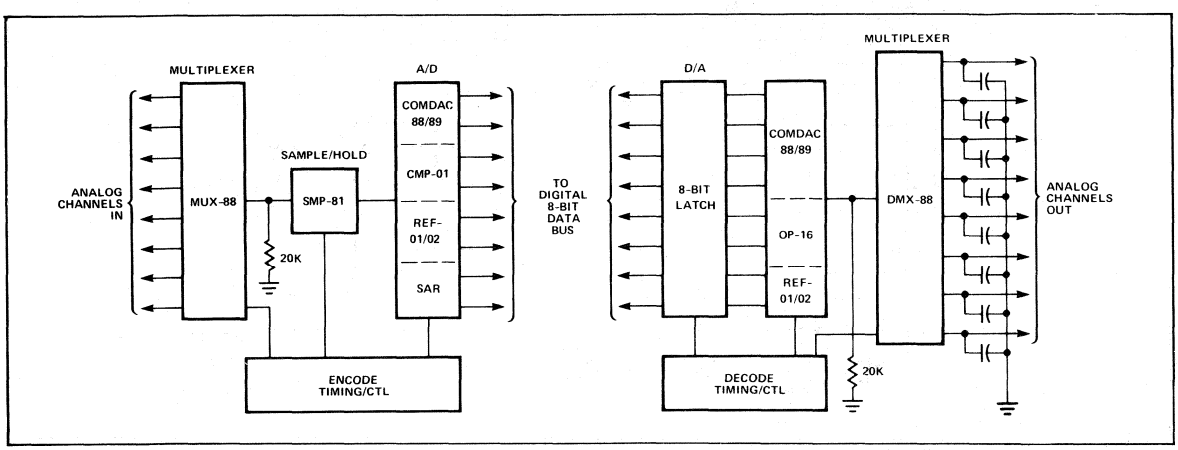

Figure 1. Eight-Channel Test Configuration

COMDAC is a registered trademark of Analog Devices, Inc.

cation, in conjunction with the nominal settling time data, can provide the designer with the data needed to determine if the performance needs of his system can be satisfied.

When the earlier version of the output multiplexer was used as a sample-and-hold and switch, it was found that certain characteristics of the device caused idle channel noise and transmission degradation. An analysis showed that reduction of the charge injected during the switch turn-off would enhance the performance of the device. The effect of the charge injection becomes important because of the capacitance added to the MUX output. This output drives a high-impedance load (the PCM filter) and without a discharge path, the charge adds to the analog output being switched through the multiplexer. Because of the use of an improved JFET switch structure, the new multiplexer exhibits a discharged only 1/4 of 1/5 of the value for the previous device. Tests reveal that the idle channel noise is reduced by several dB when the DMX-88 is used as the output switch and sample-and-hold. In addition, the reduced amount of charge permits the use of a smaller value capacitor and thus increases the number of output channels that can be decoded.

DEMONSTRATOR MODIFICATIONS

The most obvious system improvement provided by the eight-channel system depicted in Figure 1, compared to an earlier design developed by ADI (described in the ADI application note AN-37), is a simplified encoder clocking scheme. In the original circuit, additional settling time was required because of the slower changing bits.

The new version of the clocking circuitry still employs a programmable read-only memory (PROM) for flexibility in performing future modifications and experimentation. However, the clock pattern is markedly changed. The new SAR clock timing diagram is shown in Figure 2. As can be seen, bit clocking is accomplished with a set 772 kHz clock. This

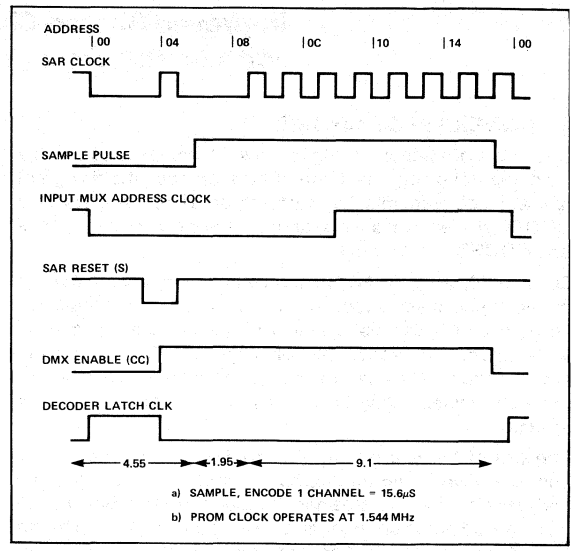

Figure 2. Encode Timing: DMX Control

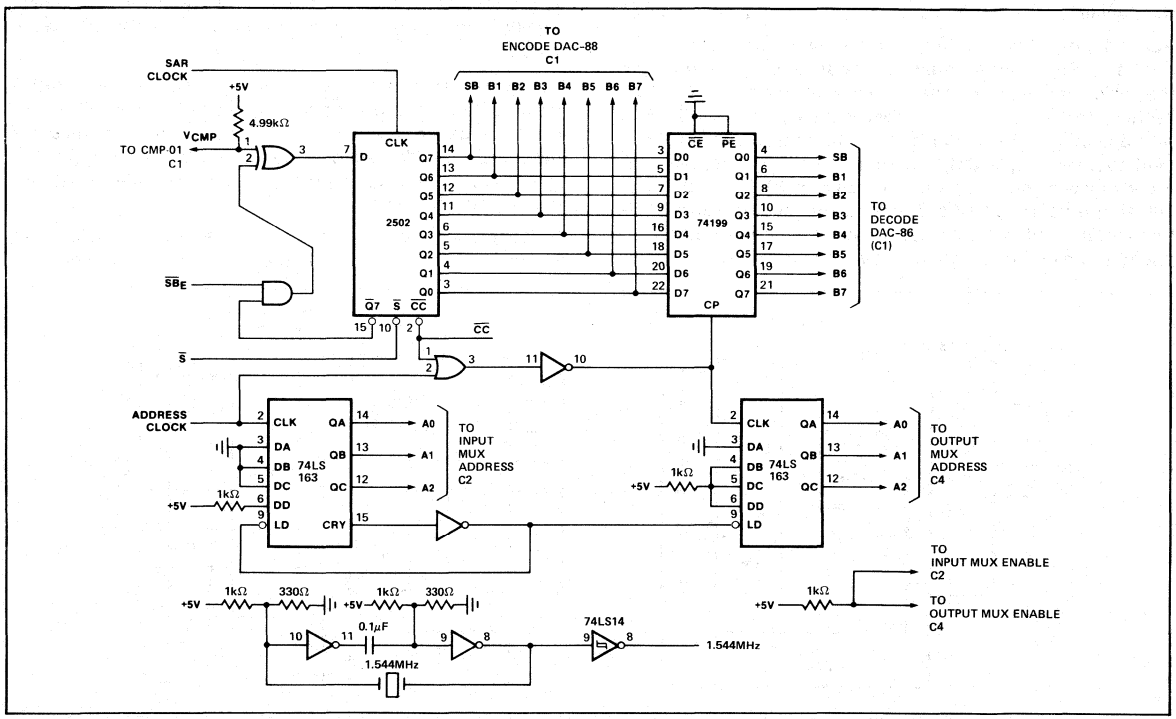

Figure 2A. Encode/Decode Controller

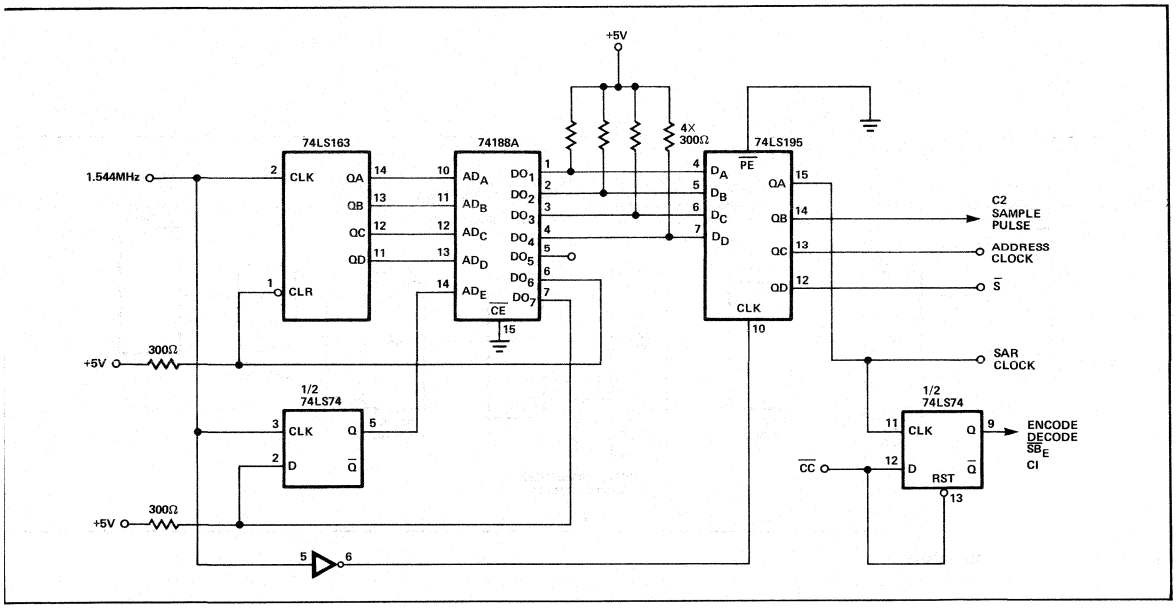

Figure 2B. Encode Clock

allows 9.1 μs (eight-channel rate) for encoding; the remaining cycle time (6.5 μs) is used for sampling the analog signal and holding the level to be encoded. The remaining cycle time is divided as follows: sample period, 4.55 μs; transition time between the hold signal and the sign bit acquisition, 1.95 μs.

The sampling time for the new clocking scheme is longer than it was for the old design (4.5 μs as compared to 3.2 μs) and the hold settling time has increased from 0.65 to 1.95 μs. As a result, the values measured for gain tracking differential (linearity) at low input levels (−55 dBm0) provide an indication of the improved response.

Since the clocking pattern has been simplified, the number of TTL gates needed is less than had been required by the original configuration. An even simpler clocking scheme can be designed by replacing the PROM, address counter and data latch with a "D" flip-flop and some additional gates. The use of a programmable clock generator, however, was advantageous in demonstrating the effects of various circuit components on the overall transmission performance. For example, by increasing the sample time and thereby reducing the hold settling time (keeping the SAR clock the same frequency), the system tends to show different characteristics. The gain tracking stays within spec, but signal-to-total distortion increases at levels below −40dBm0. The crosstalk performance (adjacent channel) also deteriorates. Apparently both of these effects are results of the output settling of the sample-and-hold device. The point is that by designing with a system consisting of individual devices the user can more precisely determine those components that have the greatest effect on system characteristics. The design can then be adapted to maximize certain performance attributes in lieu of other, less-important characteristics. The system as finalized in these notes is a compromise system, one aimed at providing adequate performance in various applicaitons. The final design modifications are left up to the individuals responsible for the specific systems.

The DAC-88 and DAC-89 have idle currents present on the selected output leads; these currents are equal on both the positive and negative outputs and are normally around 10μA. The DAC-89EX is specified at a lower reference current than the DAC-88. The new reference is 16μA less than the original value. The tests described here were performed with constant reference current for both the DAC-88 and the DAC-89. The effect on the A-law measurements means the full-scale output is 2079μA instead of 2016μA. Although all steps are slightly expanded, for the purpose of the data collected here the reference current difference is negligible.

The normal testing configuration used was to provide a test signal input in one channel and monitor the output of that channel (or adjacent channels for crosstalk) with a PCM receive filter and the prescribed receiver. The test diagram is shown in Figure 3. It becomes important to ground all unused encoder inputs to provide the proper termination. It also is very important when laying out system boards to generate sufficient ground planes and proper isolation between analog and digital ground areas. The common point of these ground areas should be as close to the power supply as possible. Also "daisy-chaining" of ground returns should be avoided. Careful consideration of grounding can help improve all system parameters.

The transmission test results collected with this system are shown in Figures 4 through 6. An example of results with a different clock pattern is also presented.

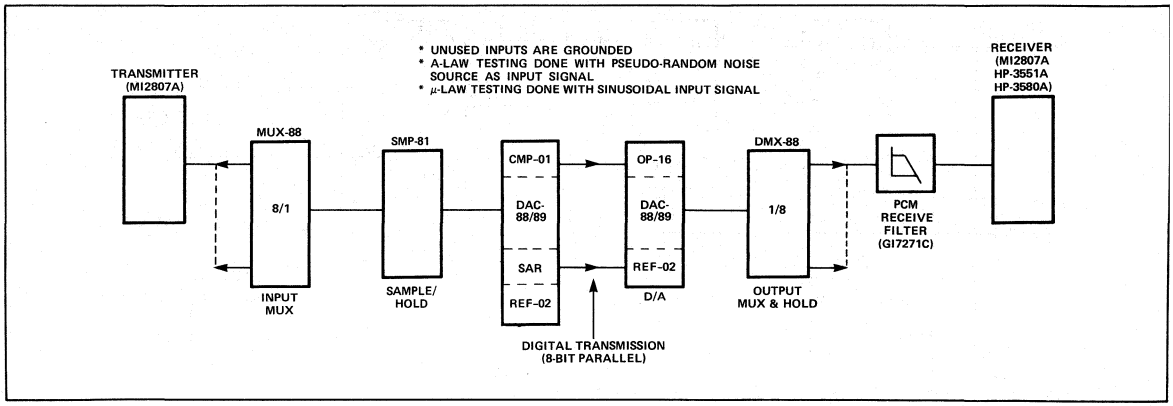

Figure 3. Test Configuration

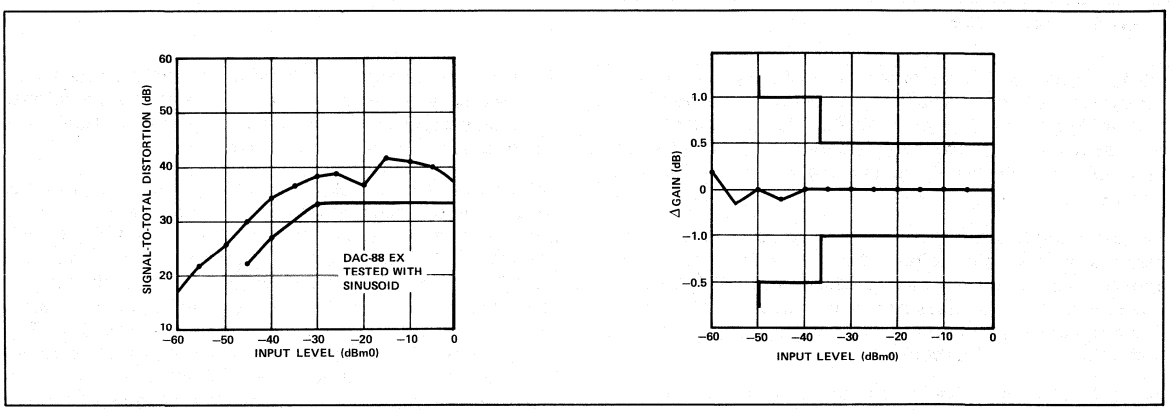

Figure 4. DAC-88EX Tested with Sinusoid

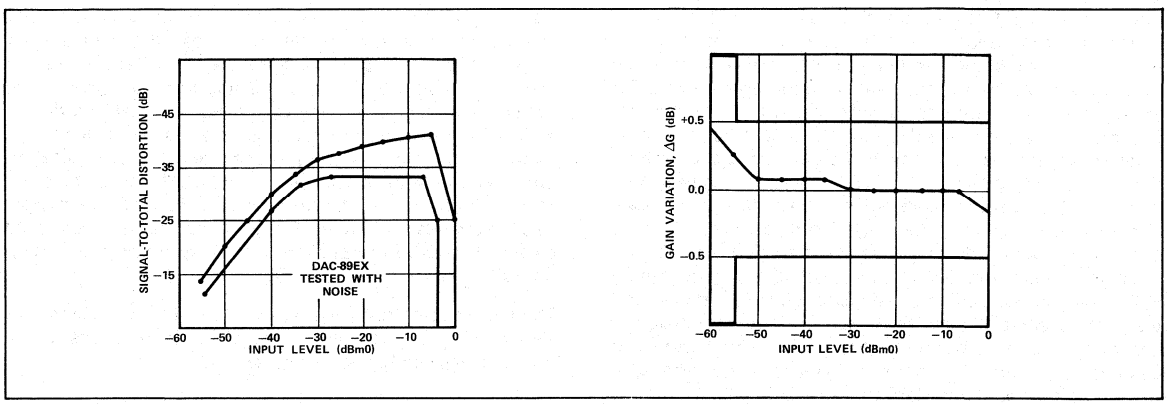

Figure 5. DAC-89EX Tested with Noise Source

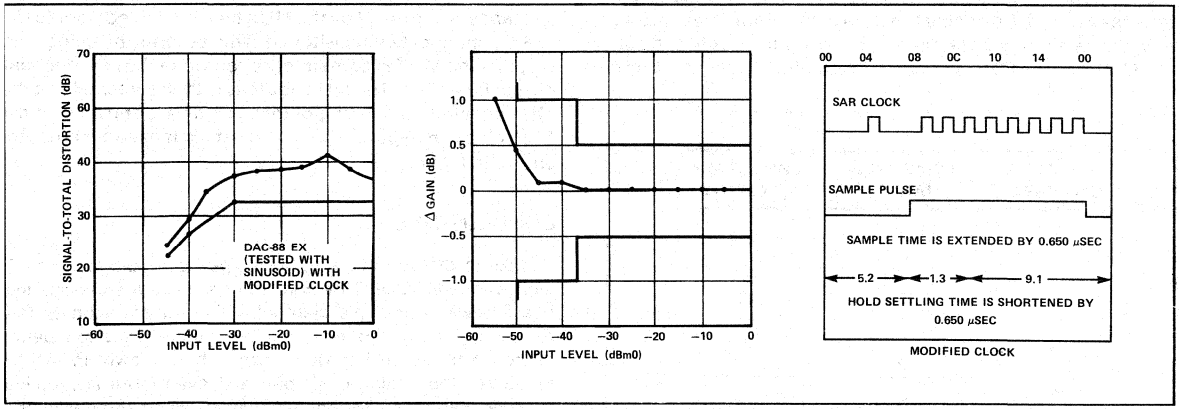

Figure 6. DAC-88EX with Altered Encode Clock (Reduced Hold Settling Time)

Comparisons of idle channel noise measurements using the MUX-88 and the DMX-88 are shown in Figures 7 and 8. It is important to notice that these measurements were made at the input to the PCM output filter. Collecting data at the filter output is not feasable because the noise values are too low for the equipment being used. To show the difference between the new and old components, a measurement was made that yielded some data. However, for both devices the idle channel noise is well within the normal system guidelines. The test data shows the difference due to reduced device charge injection of the DMX. Using a 10,000pF hold capacitor produces at least a 10dB improvement for the DMX and when the hold capacitor is reduced, the improvement is even more obvious. The smaller hold capacitor allows the D/A circuit to drive more channels. Since the current capability of the multiplexer switch is limited, a smaller capacitor means less charge-up time is required and a faster settling time is possible. The "demultiplexer" allows the user the option of reducing the hold capacitance without affecting the idle channel noise performance.

8

	CHANNEL MEASURED	IDLE CHANNEL NOISE	HOLD CAP
MUX-88	1	−60.8	10000pf
	2	−61.1	10000pf
	3	−60.3	10000pf
	4	−61.1	10000pf
	5	−60.9	10000pf
	6	−61.3	10000pf
	7	−61.2	10000pf
	8	−60.0	10000pf
DMX-88	1	All Values <−70dB	10000pf
	2		10000pf
	3		10000pf
	4		10000pf
	5		10000pf
	6		10000pf
	7		10000pf
	8		10000pf

NOTE: VALUES MEASURED BEFORE PCM OUTPUT FILTER WITH HP-3551A (−70dB IS LOWEST MEASUREMENT)

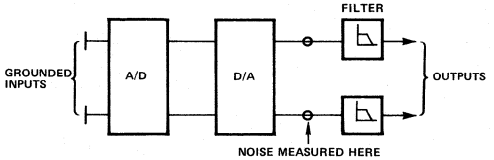

Figure 7. Idle Channel Noise

	CHANNEL MEASURED	IDLE CHANNEL NOISE	HOLD CAP
MUX-88	1	57.2	1000pf
	2	−57.2	1000pf
	3	−57.2	1000pf
	4	−57.2	1000pf
	5	−57.2	1000pf
	6	−57.2	4300pf
	7	−57.2	4300pf
	8	−57.2	4300pf
DMX-88	1	−68.6	1000pf
	2	−68.6	1000pf
	3	−68.6	1000pf
	4	−68.6	1000pf
	5	−68.6	1000pf
	6	<−70	4300pf
	7	<−70	4300pf
	8	−69.9	4300pf

NOTE: MEASURED AS SHOWN IN PREVIOUS FIGURE.

Figure 8. Idle Channel Noise

Another demonstrator design change was added to show how additional reduction of crosstalk is possible through proper control of the output multiplexer. The first design did not make use of the enable function of the output switch. As a new digital word was latched to the decode circuit, the output MUX address was switched. The timing involved in these two sequences is such that some signal feedthrough is seen due to the data latch and D/A circuit (DAC-88/89 and OP-16) settling more quickly than the MUX switch can open. Performance is improved by the MUX being disabled prior to the analog channel being switched. This assures, by using lead CC, that the MUX is completely open while the new decoder output is settling. The MUX address is then switched and the MUX enabled. The timing of these signals is shown with the SAR clock in Figure 2. The crosstalk measurements are shown in Figure 9.

INPUT CHANNEL	MEASURED CHANNELS	CROSSTALK (1.02KHz)
1	4	-72
	3	-80
	2	-80
	8	-85
	7	-85
	6	-85
	5	-78
2	8	-72
	7	-80
	6	-85
	5	-85
	1	-85
	4	-84
	3	-85
3	2	-70
	8	-80
	7	-85
	6	-85
	1	-85
	4	-74
4	3	-72
	2	-78
	8	-82
	7	-85
	6	-85
	5	-80
	1	-75

- CHANNEL NUMBERS AS SHOWN ON ENCODER & DECODER BOARDS.
- INPUT LEVEL = 0dBm0
- MEASURED WITH HP3580A SPECTRUM ANALYSER.

Figure 9. Crosstalk

CONCLUSIONS

Thanks to the use of newly developed components, the AD eight-channel CODEC demonstrator reveals the performanc possible with a multiple channel digital transmission system As illustrated by the test results, the transmission perfor mance has been improved from that shown in AN-37 Moreover, the design is simpler and even more economical Working with an eight-channel system allows the designer t investigate additional improvements in and advantages o the multiple-channel approach. Because of this, ADI make a set of boards available (encoder, decoder, controller) t any customer interested in investigating the advantages to shared-channel design.

The system is still the basic design presented in AN-37 an the work described in that note has aided in continuin improvement of ADI components. A complete demonstra tion system schematic is shown in Figure 10.

ADI is committed to providing telecommunication compo nents capable of meeting and exceeding all requirements fo digital switching and transmission systems. We feel th shared-channel approach provides economic and spac advantages over the use of single-channel CODEC's in man designs. It has also been seen by several of our customers that since our designs use individual components, by prop erly specifying these parts, the overall system performanc can be guaranteed. This can provide savings in componen and board-level testing costs. The system designer need only to evaluate the ADI approach to become aware of th possibilities it holds in terms of digital transmission syste design.

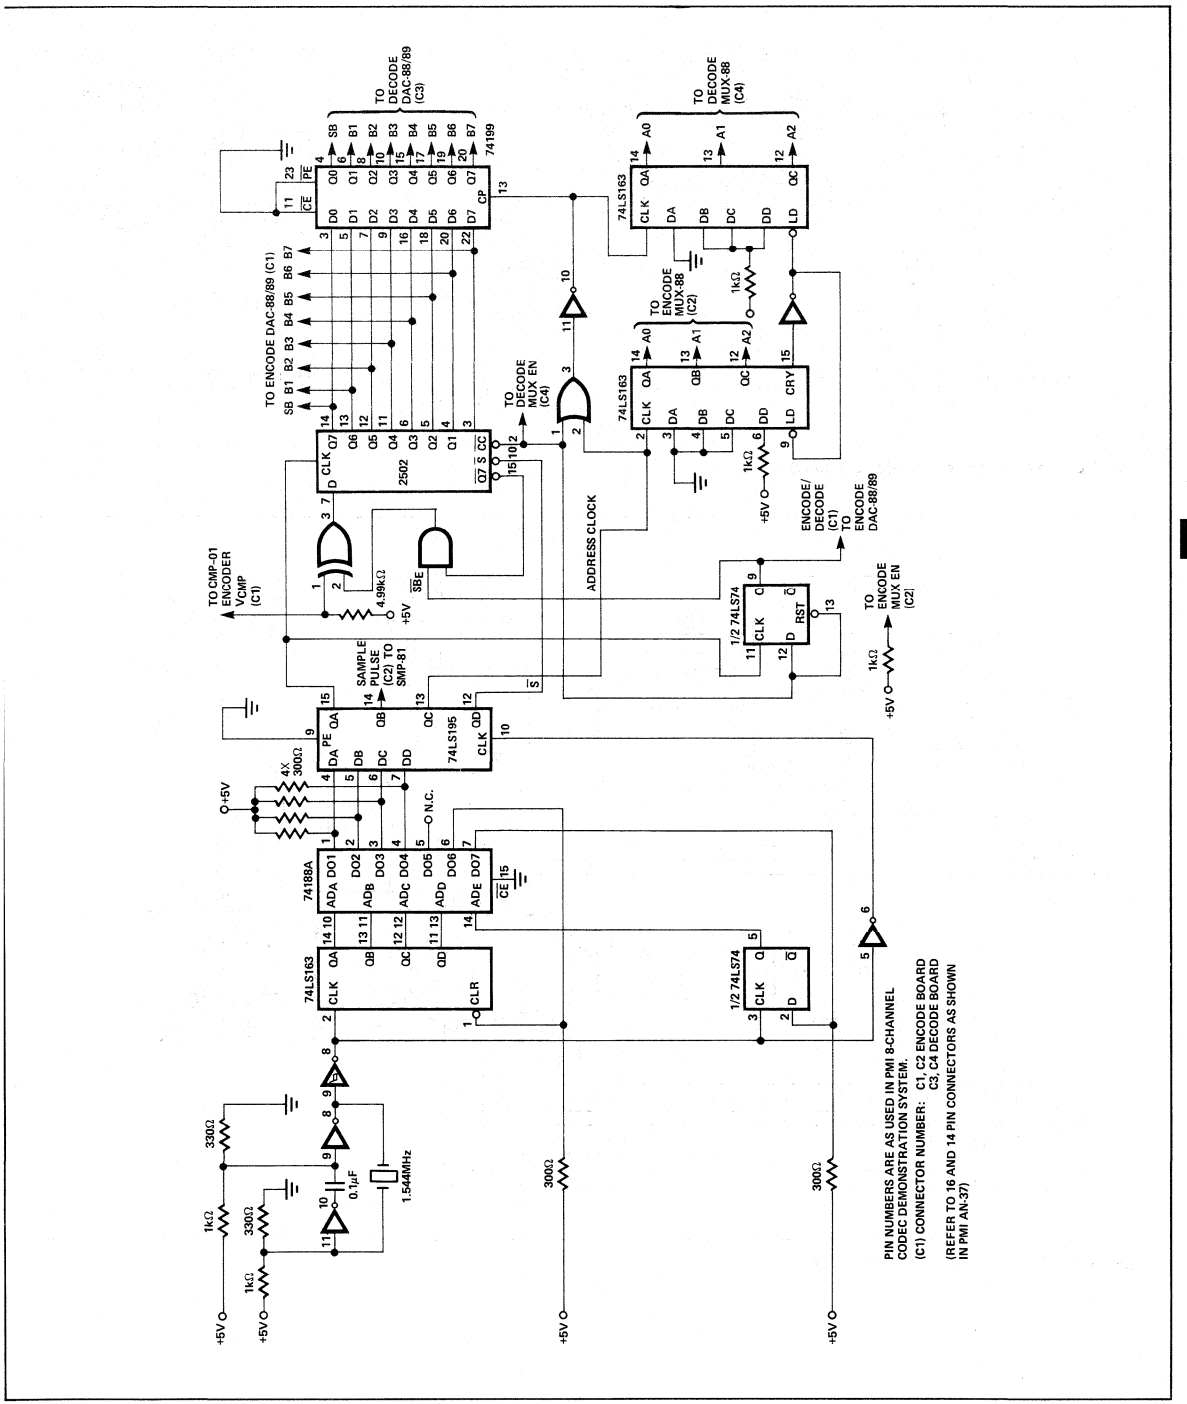

Figure 10A. CODEC Demonstration System Digital Control Board

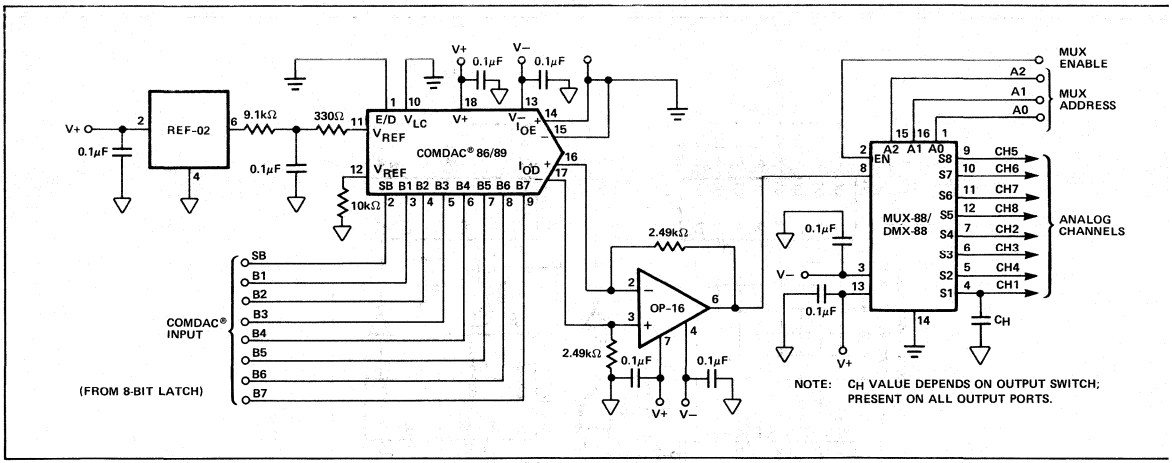

Figure 10B. Demo System Decode Board

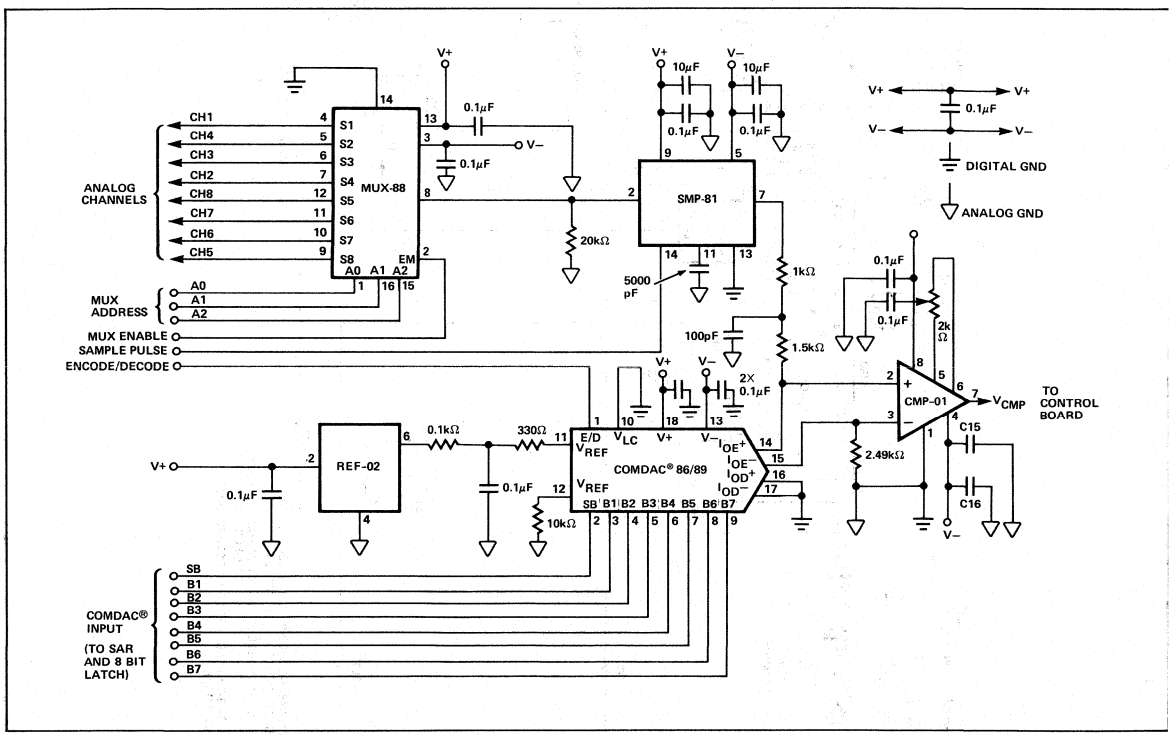

Figure 10C. Demo System Encode Board

ANALOG DEVICES

AN-101
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Cross-Plot Generator Allows Quick A/D Converter Evaluation

INTRODUCTION

The testing of analog-to-digital converters (ADCs) can be a difficult and time consuming process. The cross-plot generator described here provides a quick, low-cost way to evaluate ADCs.

ADCs AND THEIR TESTING

Analog-to-digital converters generate a digital code for each discrete value of analog input voltage. The number of codes is a function of resolution. For example, an 8-bit device has 256 codes. The resolution is equal to 1/256 of the full-scale analog voltage. 12-bit converters have 4096 codes and can resolve to 1/4096 of the full-scale voltage. The number of codes can be expressed as 2^n where n = number of bits.

Testing ADCs can be a difficult task. Modern high-resolution devices tax the speed and accuracy of measurement equipment. Linearity tests are highly data-intensive and dynamic measurements require complex and costly test apparatus.

CROSS-PLOT GENERATOR

The cross-plot generator is a simple approach to ADC evaluation. The system uses basic test equipment, simple logic and timing, and is compatible with converters of any resolution.

The test circuit (see Figure 1) consists of two D/A converters (DACs) which reconstruct the digitized information and display the data in a dot-matrix format (see Matrix Photos). This format allows the user to view the entire transfer function in real time.

Each dot in the matrix represents one digital code. Missing or narrow codes are identified as missing or dim dots in the displayed matrix.

This test method yields four types of data:

1. Missing code identification
2. Dynamic response (conversion speed)
3. Converter code-width
4. System noise

FIGURE 1: Block Diagram

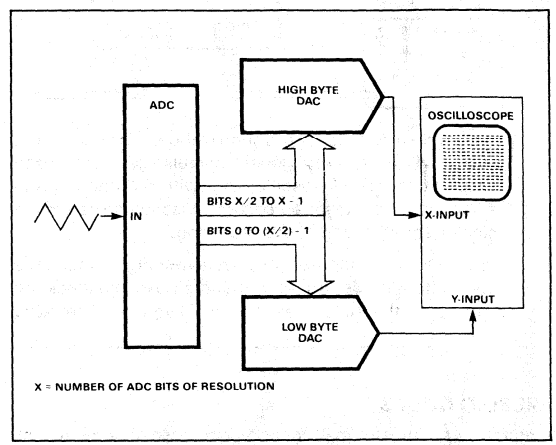

8

CROSS-PLOT MATRIX

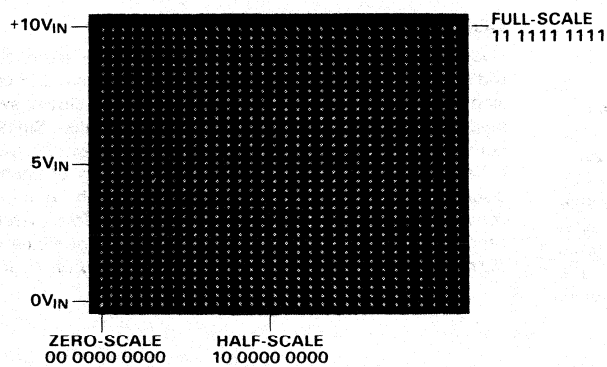

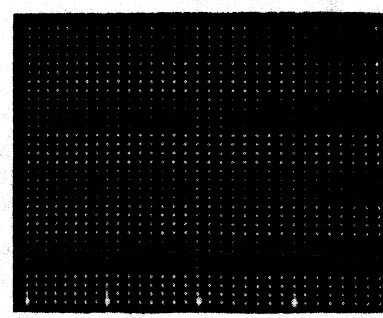

The dot-matrix field for a 10-bit converter is pictured above. The photo illustrates a 32 × 32 dot-matrix or 1024 dots.

The DUT is an ADC-910 operating with a nominal 6μs conversion time. The inputs are a 10Hz, 0 to 10V triangle wave and a 1MHz clock pulse train. These inputs result in approximately three conversions per code.

As illustrated in the photo above, missing and narrow codes are clearly visible as weak or absent dots in the array.

In this case, as conversion speed is increased beyond the ADC's capability, missing codes are observed.

FIGURE 2: Block Diagram

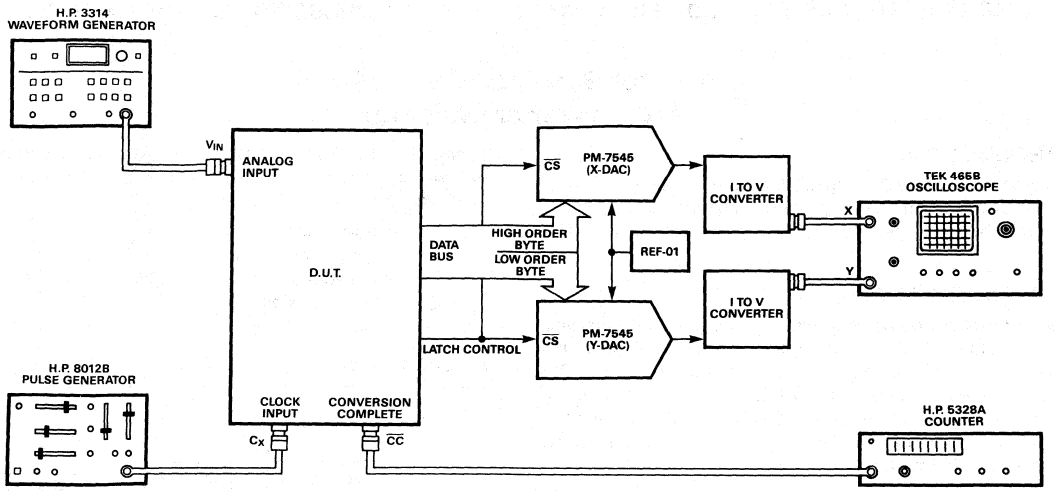

The cross-plot test set-up is shown above. The equipment required is: a waveform generator, pulse generator (used as an external clock), oscilloscope with X-Y function, and a frequency counter or second oscilloscope (to determine clock frequency and conversion time).

The ADC is configured in a free running (continuous conversion) mode. An input triangle wave and clock are applied to the ADC (see Timing Diagram). The ADC executes a conversion and presents the data to the Data Bus. The data is then divided into higher and lower order bytes and latched into the DACs by the latch control line. Following current to voltage conversion, the DAC outputs drive the oscilloscope (higher order driving the X-input, lower order driving the Y-input), creating the dot-matrix field.

MISSING CODES

One form of converter malfunction occurs when some digital codes are absent or missing. Missing codes can cause serious problems in systems level applications. These missing codes can be identified as holes in the dot-matrix field.

DYNAMIC TESTING

Conversion speeds can be measured with the cross-plot generator. A slow triangle wave is applied to the ADC's input. The waveform's voltage is adjusted to sweep from just below zero-scale to just above full-scale. The triangle-wave test frequency should be chosen so that the analog input remains within 1/2 LSB during the converter's conversion time. The conversion clock frequency is then increased until conversion codes (dots in the array) first begin dropping out. A frequency counter is used to determine the clock frequency and conversion time. Alternately, a second oscilloscope may be used, monitoring the ADC's Conversion Complete line, to determine directly the minimum conversion time.

CODE WIDTH

Code-width, or differential nonlinearity, is a measure of the width of the converter's quantizing bands. Each band or code-step has a finite value. As the input voltage is slowly swept, several conversions take place for each code value. Since the number of conversions is dependent upon code-width, shorter codes produce fewer conversions. The intensity of each dot displayed in the matrix is a direct function of the number of conversions and hence the code-width. Dimmer dots, therefore, denote narrower codes and higher differential nonlinearities. Narrow code-widths are depicted in the dot-matrix photograph.

SYSTEM NOISE

The cross-plot generator can be used to determine the repeatability of ADCs. By slowing down the input triangle wave frequency to 0.1Hz, the oscilloscope should ideally display one dot at a time when the input voltage is between code transitions. As the input voltage approaches a code transition, two dots will appear on the oscilloscope screen. Once the input voltage moves to the next code, a single dot will again appear on the oscilloscope screen. If the ADC system has noise problems, either due to layout or the ADC itself, the oscilloscope display will show two, three, or more dots simultaneously. This is an indication of undesirable noise in the system causing multiple output codes for a single analog input voltage.

Checking prototype circuits with this test can help quickly identify the noise level and source. Proper location of bypass capacitors, central analog ground, and digital signal lines can be easily determined, speeding the system debugging process.

FIGURE 3: Schematic

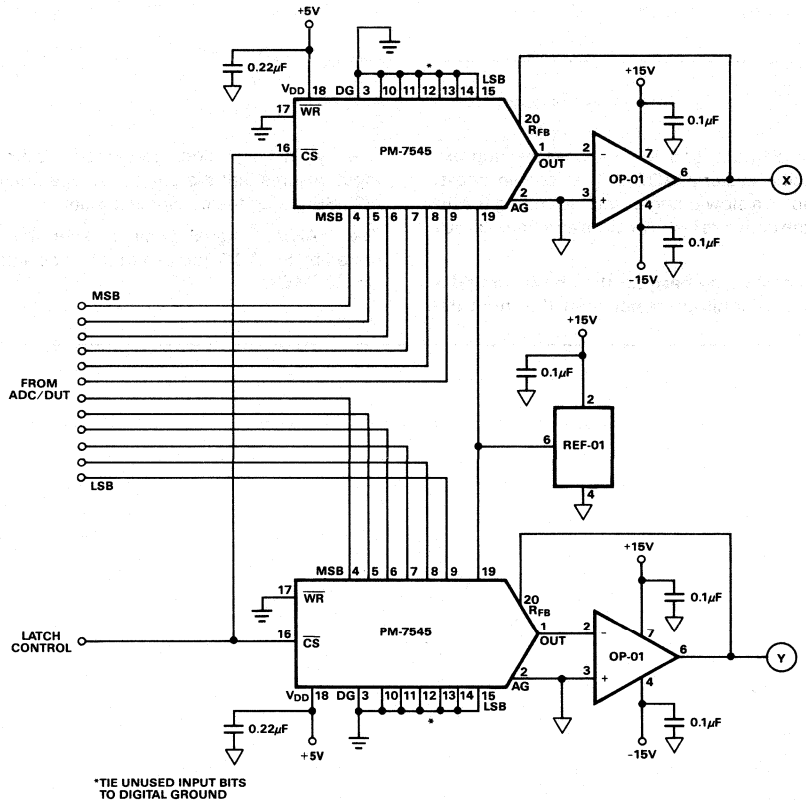

The cross-plot schematic is shown above. A pair of PM-7545s are driven directly by the ADC's outputs. The DACs drive the X-Y inputs of the oscilloscope after current-to-voltage conversion by the OP-01s. Reference voltage for the converters is provided by a REF-01.

FIGURE 4: Timing Diagram

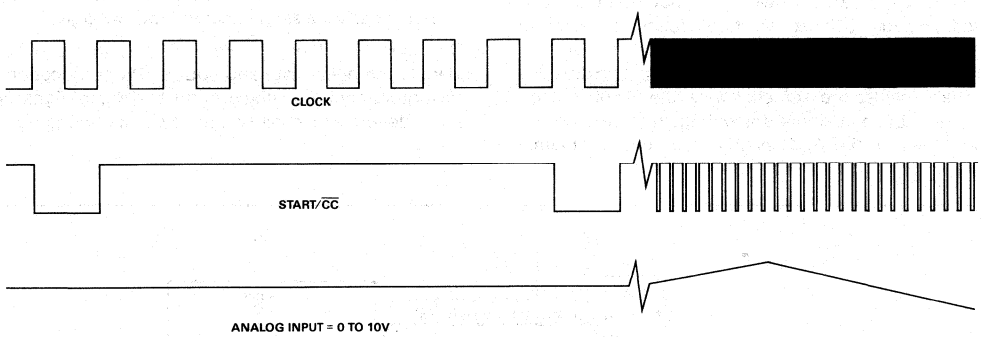

A typical timing diagram for a 10-bit successive approximation conversion system (ADC-910) is shown above. The analog input is a slow triangle wave sweeping the full input range of the converter (shown here as a unipolar 10V range).

The clock provides a time base for the converter and is several orders of magnitude faster than the input frequency. This high clock frequency, relative to the analog input, insures that the analog voltage changes less than 1/2 LSB during the conversion time.

The START/$\overline{CC}$ signal (sometimes termed $\overline{BUSY}$) is generated by the ADC and can function as a data-load pulse for the DACs.

ANALOG DEVICES

AN-205
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Video Formats and Required Load Terminations

by Bill Slattery

A number of international standards exist for specifying video levels used in television and video monitors. This application note describes some of the more common standards and compares their similarities. It also details the required load terminations for Analog Devices video RAM-DACs and explains how alternate video standards can be implemented by altering the load termination.

VIDEO STANDARDS

The NTSC standard is the one most commonly used in North America and Japan, while Europe uses PAL and SECAM video standards. Figure 1 shows an RGB video waveform, and Table I shows the associated current, voltage and IRE relationships for the various video standards.

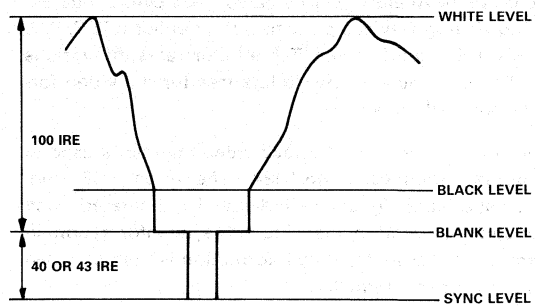

Figure 1. RGB Video Waveform

Table I. Levels Associated with Various Video Formats

	Video Output Levels	IRE Units	Volts	Singly Terminated Line mA (typ), 75Ω Monitor	Doubly Terminated Line mA (typ), 75Ω Monitor
NTSC RS-343A	Blank to White	100	0.714 ±0.1	9.52	19.04
	Blank to Black	7.5 ±5	0.054 (typ)	0.714	1.43
	Blank Level		0	0	0
	Blank to Sync	40 (typ)	−0.286 ±0.05	−3.81	−7.62
NTSC RS-170	Blank to White	100	1.0 ±0.05	13.33	26.67
	Blank to Black	7.5 ±2.5	0.075 (typ)	1	2
	Blank Level		0	0	0
	Blank to Sync	40 ±5	−0.4 (typ)	−5.33	−10.67
PAL	Blank to White	100	0.714 (typ)	9.52	19.04
	Blank to Black	0	0	0	0
	Blank Level		0	0	0
	Blank to Sync	43 (typ)	−0.307 (typ)	−4.09	−8.19
SECAM	Blank to White	100	0.714 (typ)	9.52	19.04
	Blank to Black	0 to 7	0 to 0.049	0	0
	Blank Level		0	0	0
	Blank to Sync	43 (typ)	−0.307 (typ)	−4.09	−8.19

NOTE
This table indicates the Blank Level as being the zero reference level while the Sync Level is given a negative value. In the case where composite sync is asserted using the DAC, Analog Devices video RAM-DACs have the Sync Level as the zero reference level; the Blank, Black and White Levels are all offset positively by the value of Sync Level. This will have no effect on the implementation of a particular standard as this is determined by the relative magnitude of the Blank Level relative to the White Level. The Blank to White Level remains unchanged whether or not sync is being asserted by the DAC.

The composite video waveform illustrates the relationship between the white level and blank level (gray scale or video portion) as well as the black and sync levels. The amplitude level between the blank level and white level is defined to be 100 IRE units. This corresponds to a voltage level of either 1V or 0.714V. The newer international standards specify the lower voltage level of 0.714V. The RS-343A, PAL and SECAM standards all specify a blank to white level of 0.714V, while RS-170 specifies a level of 1V. The blank to black level, also known as the setup or pedestal, is used to ensure a blacker than black beam level during retrace. The amplitude of the blank to black level varies between 0 and approximately 7.5 IRE units, depending on the video standard used. An additional 40 to 43 IRE units are required to drive the beam to the sync level. The sync levels of 40 IRE units for NTSC and 43 IRE units for both PAL and SECAM are close enough in tolerance to the 40 IRE levels of Analog Devices video RAM-DACs, thus enabling Analog Devices parts to output either NTSC, PAL or SECAM video formats. Table I illustrates the various amplitude levels and their tolerances for the video formats outlined above.

The most common of the four video standards used in computer graphics is RS-343A. The three RGB (red, green and blue) signals are individually generated, each one containing video, blanking and sync information. In many cases however, sync information is only encoded onto the green channel.

LOAD TERMINATIONS

Analog Devices video RAM-DACs are capable of driving 75Ω monitors using either doubly terminated or singly terminated loads. Table I shows the currents associated with the various video standards, for both 75Ω load termination and 37.5Ω (doubly terminated 75Ω) load termination. Figures 2 and 3 show the electrical connections between the video RAM-DAC and monitor. Any of the video standards listed in Table I can be implemented using either of these two terminations. Note that for singly terminated loads, I_{OUT} will have to be changed by adjusting I_{REF}.

IMPLEMENTATION OF RS-343A AND RS-170

Analog Devices video RAM-DACs can implement either RS-343A or RS-170. This is achieved, in the case of a doubly terminated configuration, by varying the value of the source termination resistance Z_S. Figures 4a to 4d show the required terminations as well as the associated RGB video waveforms for both RS-343A and RS-170 implementation when using the ADV478/ADV471. The assertion or nonassertion of sync is also distinguished in these diagrams. The advantage of using this technique to implement the different video standards lies in the fact that the output current level of the DAC need not be altered. The relationship between DAC output current, load termination resistance and voltage across the monitor is given by

$$V_L = \frac{I_{OUT} \cdot Z_S \cdot Z_L}{Z_S + Z_L}$$

V_L = voltage developed across monitor
I_{OUT} = DAC output current
Z_S = source termination resistance
Z_L = cable/monitor impedance

Since the relevant video standard is determined by the voltage developed across Z_L, altering the value of Z_S is a simple method of selecting any of the video standards.

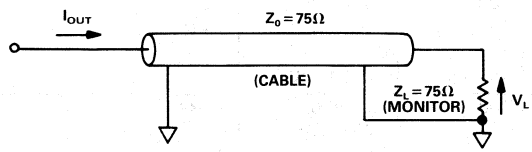

Figure 2. Singly Terminated 75Ω Load

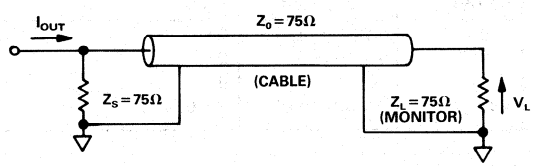

Figure 3. Doubly Terminated 75Ω Load

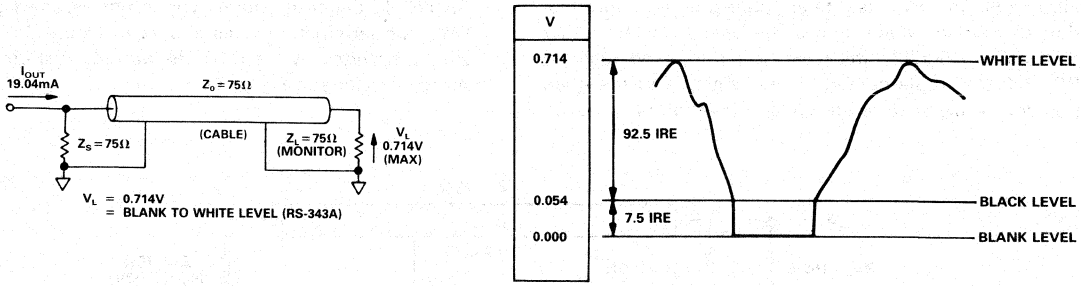

Figure 4a. RS-343A Load Termination & RGB Video Waveform (SYNC Not Asserted)

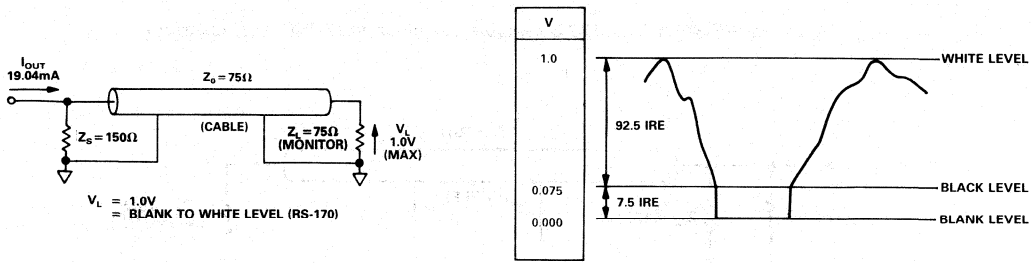

Figure 4b. RS-170 Load Termination & RGB Video Waveform (SYNC Not Asserted)

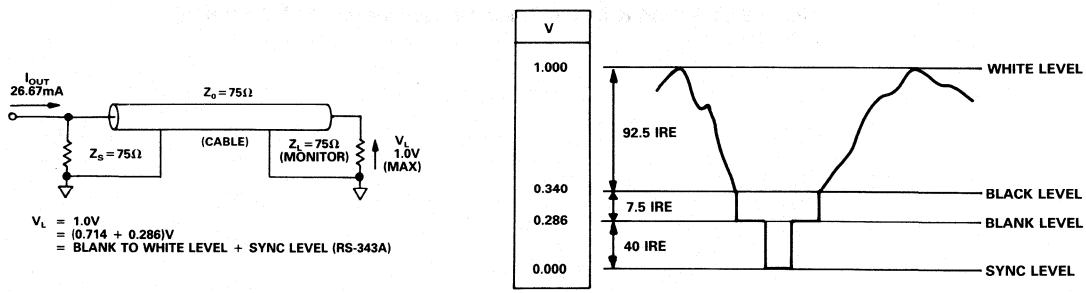

Figure 4c. RS-343A Load Termination & RGB Video Waveform (SYNC Asserted)

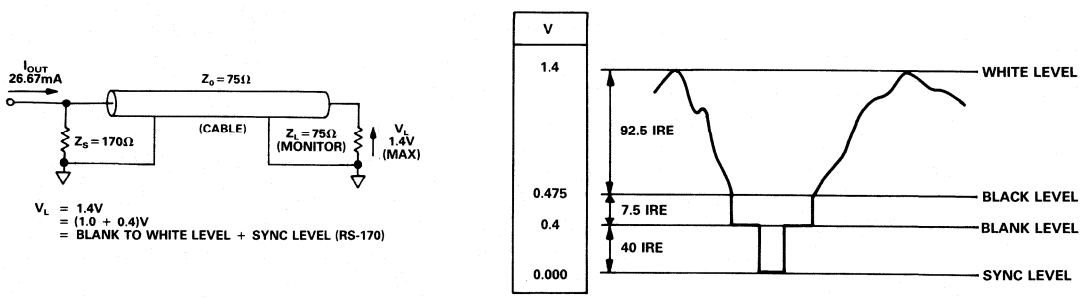

Figure 4d. RS-170 Load Termination & RGB Video Waveform (SYNC Asserted)

SELECTABLE TERMINATION

Figures 5a and 5b illustrate an interesting load termination technique which allows the user to select either RS-343A or RS-170. If the switch is in the closed position, RS-343A is implemented. If the switch is in the open position, a blank to white voltage level of 1V is developed across the 75Ω monitor load, corresponding to RS-170. In the case where sync is not asserted by the DAC, the termination is as shown in Figure 5a. When sync is asserted by the DAC, the output must be terminated according to Figure 5b.

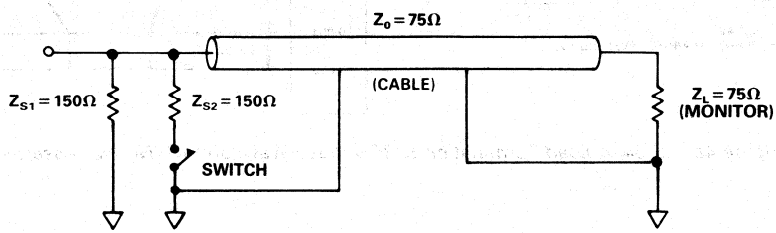

Figure 5a. RS-343A & RS170 Selectable Termination (SYNC Not Asserted)

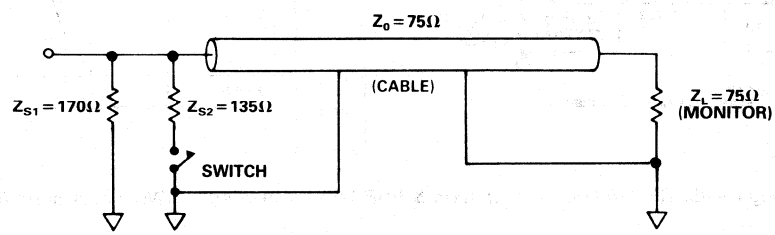

Figure 5b. RS-343A & RS170 Selectable Termination (SYNC Asserted)

ANALOG DEVICES

AN-330
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Animation Using the Pixel Read Mask Register of the ADV47x Series of Video RAM-DACs

by Bill Slattery & Eamonn Gormley

INTRODUCTION

The Pixel Read Mask Register, which is an integral part of IBM's VGA* graphics system, can be used as a hardware-level Pixel Processing Unit. This allows real time motion or animation to be implemented with minimal software overhead. This application note examines the operation and structure of such a pixel processing unit with the pixel read mask register as the central controller. A practical application which uses the pixel read mask register to animate a picture scene is described. A complete listing of the Turbo-C source code is given in the appendix.

No additional hardware is required for existing VGA graphics systems to implement this application.

VIDEO RAM-DAC

Analog Devices produces a range of video RAM-DACs, which are specifically designed for IBM's Personal System/2* VGA. The range includes the ADV478, ADV471 and ADV476, all of which are monolithic +5 V CMOS video RAM-DACs. These parts are specified over a number of speed grades; 35 MHz, 50 MHz, 66 MHz and 80 MHz. The RAM-DACs are packaged as 44-pin PLCC and 28-pin plastic DIP devices.

The ADV471 and ADV476 each contain a triple 6-bit digital-to-analog converter and a 256 location by 18 bits deep color look-up table. The devices also include an asynchronous pixel input port and bidirectional microprocessor (MPU) port. These devices and the associated control circuitry allow for flexible interface to many graphics systems configurations. The ADV478 differs from the ADV471 only in terms of its color resolution. The ADV478 has a triple 6-bit/8-bit D/A converter with a 256 × 24/18 color look-up table. The color resolution of the ADV478 is user selectable between 6 bits and 8 bits. The higher 8-bit performance can be used with IBM's 8514/A* graphics standard (upgrade on standard VGA). More detailed information on these and other video RAM-DACs can be obtained in the relevant product data sheets.

Built into all three devices is an 8-bit register known as the Pixel Read Mask Register. Figures 1 and 2 are block diagrams of the ADV478/ADV471 and ADV476 which show the Pixel Read Mask Register.

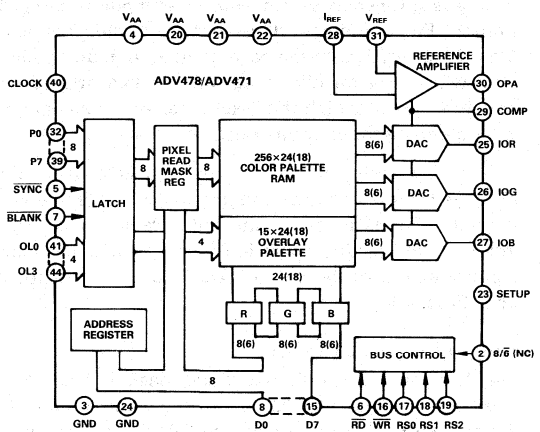

Figure 1. ADV478/ADV471 Functional Block Diagram

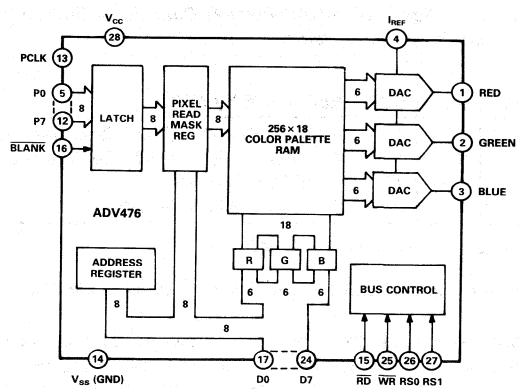

Figure 2. ADV476 Functional Block Diagram

*IBM, VGA, Personal System/2 and 8514/A are trademarks of International Business Machines Corp.

Some of the uses to which the Pixel Read Mask Register can be put include on-screen special effects such as real time animation, flashing objects and overlays.

PIXEL READ MASK REGISTER

The Pixel Read Mask Register is placed in the path of the pixel input stream of data as shown in Figure 3.

The input pixel data stream (P0–P7) is gated with the contents of the Pixel Read Mask Register. The operation is a bitwise logical ANDing of the pixel data. The contents of the Pixel Read Mask Register can be accessed and altered at any time by the MPU (D0–D7). Table I shows the relevant control signals. Under normal operating conditions, this register is loaded with all 1s, i.e., transparent mode.

In a VGA graphics system, the Pixel Read Mask Register is memory mapped and is accessible (read/write) by addressing memory location 36CH.

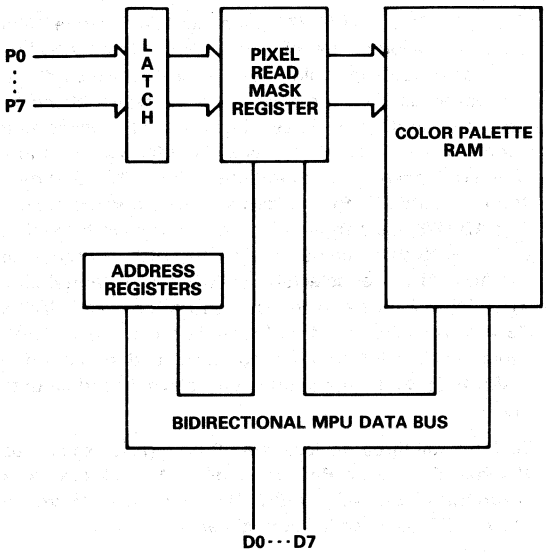

Figure 3. Video RAM-DAC Pixel & Data Ports Showing the Pixel Read Mask Register

RS2*	RS1	RS0	Addressed by MPU
0	0	0	Address Register (RAM Write Mode)
0	1	1	Address Register (RAM Read Mode)
0	0	1	Color Palette RAM
0	1	0	Pixel Read Mask Register

*RS2 is only present on ADV478/ADV471

Table I. Control Input Truth Table for Video RAM-DAC

Figure 4 shows the internal architecture of the pixel input port. The input word Pi, which corresponds to an on-screen pixel location, is ANDed with the contents of the Pixel Read Mask Register, Pm.

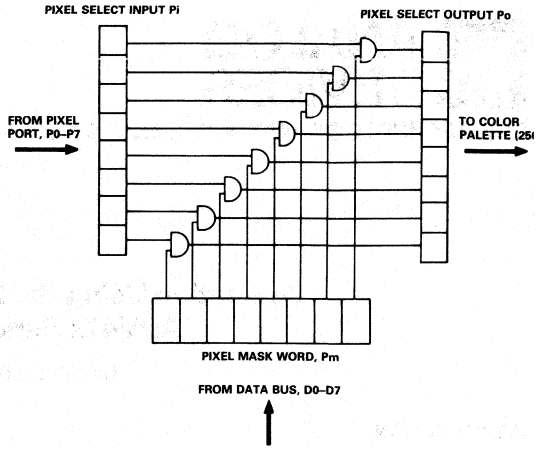

Figure 4. Internal Architecture of Pixel Input Port

The resulting output word, Po, determines which location in the color palette will be assigned to a particular on-screen pixel. Figure 5 shows the logical diagram for this masking operation.

Figure 5. Equivalent Logical Representation of Masking Operation

$$Po = Pi \cdot Pm \quad (1$$

If Pm = 1, transparent mode, then

$$Po = Pi \quad (2$$

The pixel stream of data, Po, which arrives at the color palette is a function of both the pixel input stream, Pi, from the Frame Buffer and the contents of the Pixel Read Mask Register, Pm. In the case of an animation application, which will be discussed later in this application note, the rate at which the pixel mask word is changed will determine the motion speed of the scene.

This pixel masking operation can be used to alter the displayed colors without changing the contents of either the video Frame Buffer or the Color Palette RAM. One interpretation of this operation is to consider the pixel input structure of the video RAM-DAC as an on board Pixel Processing Unit.

PIXEL PROCESSING UNIT

The Pixel Input Port (Pi), Pixel Read Mask Register (Pm) and Data Input Port (MPU) within the video RAM-DAC, are the hardware components of this Pixel Processing Unit (PPU). An associated software routine to control the operation is the final element in the complete PPU system. This interpretation enables the color palette to be configured as a multidimensional, paged memory address space, see Figure 6.

In the case of the ADV471 and ADV476 the color palette can be perceived as being broken into an even number of 18-bit color planes instead of just one 18-bit deep color plane. (In the case of the ADV478, each plane is 24 bits deep.)

The palette can therefore be partitioned to produce up to a total of 256 discrete contiguous color memory planes, some of which are shown in Figure 6. A tradeoff, however, must be considered when dividing the color palette into multiple color planes. The number of simultaneously displayable screen colors is inversely proportional to the number of color planes within the color palette. Table II illustrates this relationship. This contiguous configuration is not, however, the sole way of segmenting the memory within the palette. Other noncontiguous configurations including interleaving can be implemented. The choice of memory configuration will be determined by the particular application so as to make most efficient use of the available video memory (Image Frame Buffer and Color Palette RAM).

Number of Color Planes	Number of Simultaneously Displayable Colors
1	256
2	128
4	64
.	.
.	.
256	1

Table II. Simultaneously Displayable Screen Colors versus Number of Color Planes

To operate the PPU, two principal steps must be taken:

1. Load the image data in the correct paged configuration to both the frame buffer and color look-up table, e.g., four frame composite images to the frame buffer corresponding to four discrete planes of color to the palette RAM.
2. Generate the corresponding pixel mask words. These words are individually written to the Pixel Read Mask Register, Pm (at VGA memory location 36CH) and select which of the color planes is to be assigned to the incoming pixel data stream. In the case where four planes are implemented (see Figure 6), four pixel mask words are required.

The overall VGA System Block Diagram showing the PPU and an associated 8-page memory configuration of the color palette is shown in Figure 7.

8

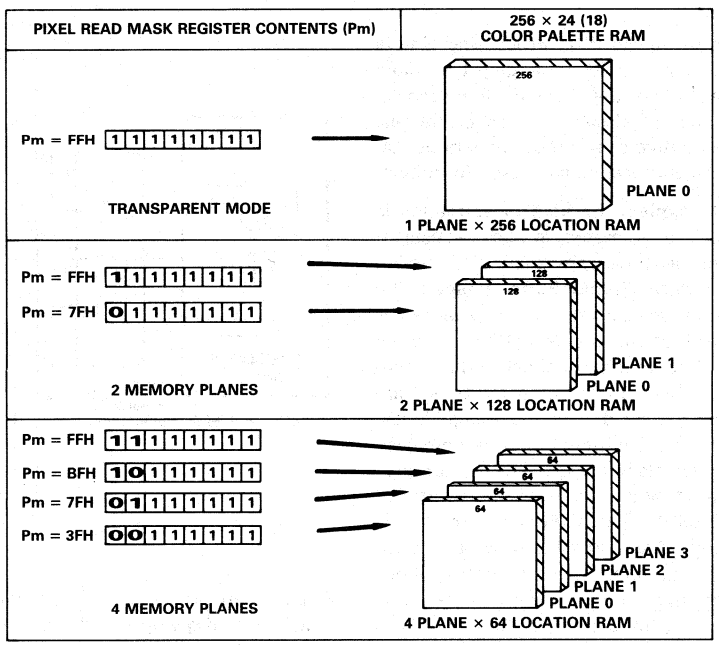

Figure 6. Some Color Palette RAM Configurations Showing Paged Memory Planes and Associated Pixel Read Mask Word (Pm)

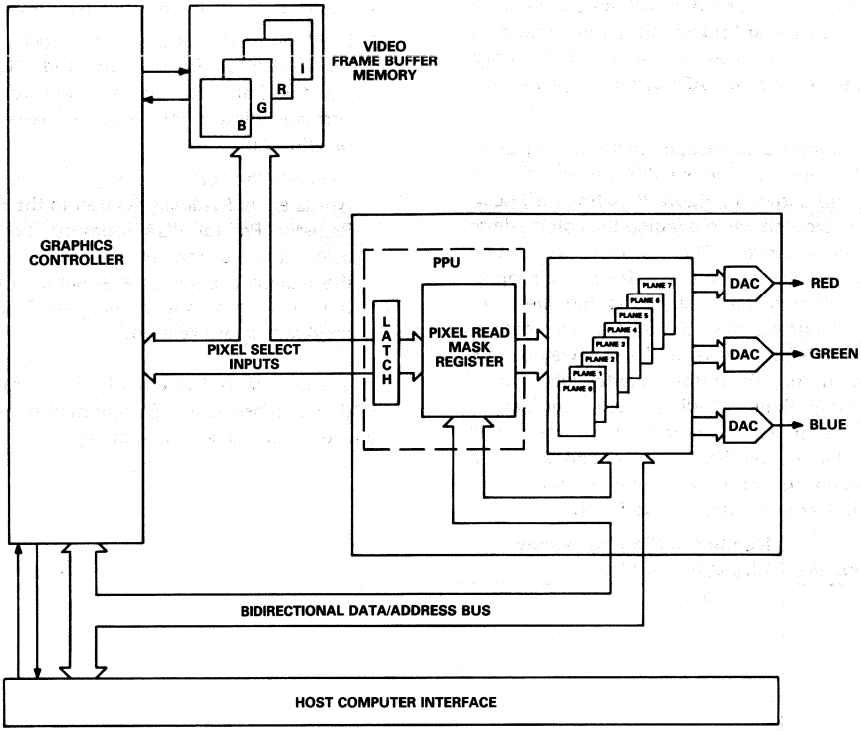

Figure 7. VGA System Block Diagram Showing Color Palette Broken into a Number of Color Planes

ANIMATION

Real time animation using the Pixel Processing Unit is based on the principle that rapidly changing the colors of a stationary object gives the illusion of motion. In other words, a number of similar images or frames, differing only by the relative position of the various colors, displayed in quick succession, can result in motion.

A simple example to explain the idea of animation is illustrated opposite. The animated image consists of three frames; each of the three frames is initially drawn as one composite picture (Frame Buffer image). The color palette contains three discrete, memory blocks or planes of color information, corresponding to three stages of animation. The animation effect in this example is "arm waving" of the cartoon character. By assigning the color planes one by one to the image in the Frame Buffer, the effect of animation can be perceived on the screen. The color plane assigned to the composite image is determined by the PPU which is controlled by the word in the Pixel Read Mask Register. Frame 1 is assigned Color Plane 0, this colors the down arm position in black, while the up and horizontal arm positions take on the background color. Frame number 2 is assigned Color Plane 1, this colors the horizontal arm position in black while the other two arm positions are assigned the back ground color.

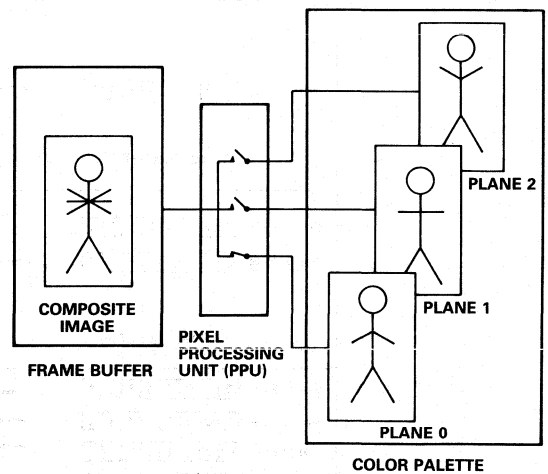

Finally, Frame 3 takes on Color Plane 2. This process is then repeated giving the illusion of motion. The rate at which each frame is selected determines the rapidity of the arm waving.

ANIMATION USING THE PPU

This section describes a particular animation example. The scene used in this example consists of traveling space ships and rotating planets. The program which draws the scene and implements the animation is described in the flow diagram of Figure 8. The associated source code, written in Borland's Turbo-C, is given in the Appendix. This application implements 8-stage animation.

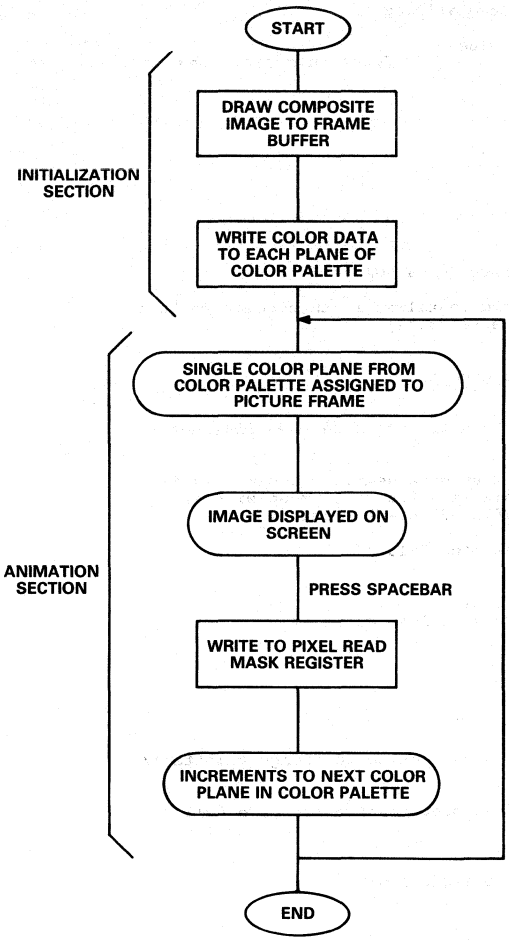

Figure 8. Flow Diagram Representation of Animation Using the PPU

The complete image is drawn to the Frame Buffer. This composite picture contains eight frames of information. The corresponding color planes for each of the eight frames in this composite image, are drawn to the color palette. The color information is arranged in a paged memory format corresponding to that shown in Figure 7. Each of these eight color planes has similar color data; they differ from each other only in terms of the relevant position of the particular colors. For example, Plane 0 could have blue in its first location and color yellow in its second location, while Plane 1 could have the opposite, yellow in Location 1 and blue in Location 2. During the display period, the color palette will only allocate colors to one of the eight frames (i.e., one color plane) at a particular instant. Each plane of color information is mapped to a particular frame within the Frame Buffer. The user-defined value of the Pixel Read Mask Register determines which of the color planes within the palette will be chosen for display at any particular instant. The hex codes, written to the Pixel Read Mask Register Pm, which correspond to each of the color planes (Plane 0 to Plane 7) are listed in Table III.

3C6F	: Address of Pixel Read Mask Register (Pm)	
8FH → Pm	: Plane 0 Selected	Pm = 1000 1111
9FH → Pm	: Plane 1 Selected	Pm = 1001 1111
AFH → Pm	: Plane 2 Selected	Pm = 1010 1111
BFH → Pm	: Plane 3 Selected	Pm = 1011 1111
CFH → Pm	: Plane 4 Selected	Pm = 1100 1111
DFH → Pm	: Plane 5 Selected	Pm = 1101 1111
EFH → Pm	: Plane 6 Selected	Pm = 1110 1111
FFH → Pm	: Plane 7 Selected	Pm = 1111 1111

Table III. Value Written to Pixel Read Mask Register and Associated Color Plane

Pressing the "Spacebar" increments the pixel read mask register corresponding to a jump of 16 locations in the color palette. As there are 16 colors in each color plane, a jump of 16 locations will select the corresponding color in the next highest plane. Continuously pressing the "Spacebar" cycles the incoming pixel stream of data through each of the eight color planes within the palette. This results in the apparent motion or animation of the image.

APPENDIX

C Program for ANIMATION EXAMPLE

Pixel Processing Using Video RAM-DACs
"Rotating Planets & Spaceships"

```
#include <stdlib.h>  /* Turbo C include files */
#include <math.h>    /* these are available under most versions */
#include <dos.h>     /* of C for the IBM & compatibles */
#include <graphics.h>
void palette(int col,int red,int green,int blue);
void plot13(int x,int y,int col);                    /* function definitions */
void mode13();
void circle13(int x,int y,int r,double tilt);
void planets();
void stars();
void line13(int x1,int y1,int x2,int y2,int col);
void triangle();

main()
{
 int gd=0,gm=0,opt;
 union REGS reg;
 detectgraph(&gd,&gm);                  /* check for a VGA card */
 if (gd != 9){
  printf("This program cannot find a VGA card installed in this computer.\n");
  printf("A VGA card is necessary to run the tests.");
  exit(1); }

 planets();                             /* do demo */
 reg.h.ah = 0x00;
 reg.h.al = 0x03;
 int86(0x10,&reg,&reg);                 /* return to text mode when finished */
}

void mode13()                           /* set up mode hex 13 = decimal 19 */
{                                       /* this is a 256 color mode with */
 union REGS reg;                        /* 320 x 200 pixel resolution */
 reg.x.ax = 0x0013;
 int86(0x10,&reg,&reg);                 /* set mode 0x13 */
}

void palette(int col,int red,int green,int blue)
             /* assigns a physical color to a logical color */
{
  union REGS reg;
     reg.x.ax = 0x1010;
     reg.x.bx = col;
     reg.h.dh = red;
     reg.h.ch = green;
     reg.h.cl = blue;
     int86(0x10,&reg,&reg);         /* call bios routine to change palette */
}

void plot13(int x,int y,int col)    /* special plot routine for mode 0x13 */
{
     union REGS reg;
     if(x>=0 && y>=0 && x <320 && y<200){
        reg.x.dx = y;               /* set up registers */
        reg.x.cx = x;
        reg.h.ah = 0x0c;
        reg.h.al = col;
        int86(0x10,&reg,&reg);      /* call bios plot routine */
     }}

void circle13(int x,int y,int r,double tilt)
{                                    /* routine draws a single planet */
```

```
int la,yy;
double ang,oldx,oldy,newx,newy,sintl,costl,rcos,rsin;
for(la = -r; la < r; la++)
{                                       /* routine uses a fairly simple */
 yy = sqrt(r*r - la*la) + 1;            /* algorithm to draw a solid circle */
 line13(la+x,y-yy,la+x,y+yy,14); }

costl = cos(tilt);                      /* set up some variables */
sintl = sin(tilt);
yy = 240;                               /* To draw lines of longitude: */
for(la = r; la >= -r; la-=r/15)         /* draw portions of ellipses */
  {                                     /* and rotate them by tilt radians */
   oldx = x-r*sintl;
   oldy = y+r*costl;
   for(ang = -1.57;ang <1.57;ang+=.195) {
    newx = x+la*cos(ang)*costl+r*sin(ang)*sintl;
    newy = y-r*sin(ang)*costl+la*cos(ang)*sintl;
    line13(newx,newy,oldx,oldy,yy); /* line segment of ellipse */
    oldx = newx;                        /* store endpoints */
    oldy = newy;  }
   yy = (yy==247) ? 240 : ++yy;         /* incrememt color used */
  }
for(ang=-1.57;ang<1.57;ang+=.39)        /* draw lines of latitude */
   {                                    /* ie sloped lines */
    rcos = r*cos(ang);
    rsin = r*sin(ang);
    line13(x+rcos*costl-rsin*sintl,y+rsin*costl+rcos*sintl
          ,x-rcos*costl-rsin*sintl,y+rsin*costl-rcos*sintl,15);
   }}

void planets()              /* routine to draw and animate the planets */
{
 int la,lb;

 mode13();
 palette(7,255,255,255);
 printf("       Pixel Read Mask Demo\n");
 printf("     ========================\n");
 printf("  This program contains an animated     picture scene which ");
 printf("is initially drawn  on the screen and then ANIMATED ");
 printf("using    the Pixel Read Mask Register.\n");
 printf("\n  Press the spacebar to draw scene and  hold it down ");
 printf("when scene is ready for    animation. When finished, ");
 printf("press any     other key.....");
 while(getch() != ' ');                 /* wait for keypress */

 mode13();
 for (la=8;la<16;la++)                  /* set up the palette for animation */
   {
    for (lb=0;lb<8;lb++)
    palette(la*16+lb,0,10,63);          /* set planet lines to blue */
    for (lb=8;lb<16;lb++)
    palette(la*16+lb,0,0,0);            /* stars are initially black */
   }
 for (la=128;la<256;la+=17)
    palette(la,63,63,63),               /* define one line on planet to white */
    palette(la+8,63,63,0);              /* and one star to yellow, per frame */

  palette(15,255,255,255);              /* set color 15 to pure white */
  palette(7,20,255,0);                  /* color 7 to green */
  palette(14,0,10,63);                  /* color 14 used for planet background */

  stars();                              /* draw stars in background */
  circle13(30,30,30,0.9);               /* draw the actual planets */
  circle13(280,35,35,4.0);
  circle13(130,100,70,-0.8);
```

```
  circle13(40,240,125,0.5);
  triangle();                          /* draw the spaceship thingy */
  gotoxy(30,21);printf("Space to");
  gotoxy(30,22);printf("animate.");    /* on screen instructions */
  gotoxy(30,24);printf("Other key");
  gotoxy(30,25);printf("to stop.");
   la=143;                             /* 143 = %10001111 */
   do
    outportb(0x3c6,la),                /* this part does the actual animation */
    la = (la<255) ? la+16 : 143;       /* loop through the palette */
   while((lb = getch()) == ' ');       /* while the spacebar is being pressed */
}

void stars()                           /* routine to plot in the stars */
{
 int la,lb,lc,ld,le,col = 248;
 long q;
 srand(time(&q) % 37);                 /* set up random background */
 for (la=0;la<200;la+=5)  {
    lc = (rand()&0x7)-0x4;
    ld = la;
    le = (rand()&7)+3;
    for (lb=1;lb<320;lb+=le,ld=la+lc*lb/64)
       plot13(lb,ld,col),     /* plot the star */
       col = (col == 255) ? 248 : ++col;
  }}

void line13(int x1,int y1,int x2,int y2,int col)
{                                          /* this routine draws a line in */
 int la,lb,lc;                             /* graphics mode 13H */
 if (abs(x1-x2) > abs(y1-y2))  {      /* line longer in x or y direction ? */
   lc = (x2-x1);lb = (x2 - x1 >=0) ? 1 : -1;
   for (la=x1;la!=x2;la+=lb)               /* loop works out the points on */
   plot13(la,y1+(la-x1)*(y2-y1)/lc,col);   /* the line and plots them */
  }
  else  {
   lc = (y2-y1);lb = (y2 - y1 >=0) ? 1 : -1;
   for (la=y1;la!=y2;la+=lb)
   plot13(x1+(la-y1)*(x2-x1)/lc,la,col);
  }}

void triangle()          /* This routine draws a simple spacecraft-type */
{                        /* object for animation. */
 int la,lb=19,col=248;   /* starting size = 19, color = 248 */
 double tilt=0.5236;     /* starting tilt */

 for (la=200;lb>0;la-=lb,lb--,tilt += .3)
 {                       /* loop to draw 19 objects */
  line13(200+la/2+lb*cos(tilt),la+lb*sin(tilt),
         200+la/2+lb*cos(tilt+2.0944),la+lb*sin(tilt+2.0944),col);
  line13(200+la/2+lb*cos(tilt+2.0944),la+lb*sin(tilt+2.0944),
         200+la/2+lb*cos(tilt+4.1888),la+lb*sin(tilt+4.1888),col);
  line13(200+la/2+lb*cos(tilt+4.1888),la+lb*sin(tilt+4.1888),200+la/2,la,col);
  line13(200+la/2,la,200+la/2+lb*cos(tilt),la+lb*sin(tilt),col);
  col = (col==255) ? 248 : ++col;     /* col = col + 1 until col = 255, when */
 }}                                   /* col returns to zero */
```

ANALOG DEVICES

AN-331
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Changing Your VGA Design from a 171/176 to an ADV471

by Bill Slattery

An ADV471 can be used in VGA* type graphic systems to replace the 171 or the 176. If the designer does not want to use the additional functions of the ADV471, it is quite simple to have these functions disabled. The diagrams below illustrate just this. Figure 1 shows the circuit connection diagram for a 171/176. Figure 2 is the connection diagram for an ADV471 in an application that was designed around a 171/176. It is an ADV471 functioning as a 171/176.

The additional pins of the ADV471 inputs (OL0–OL3) are connected as shown in Figure 2. The overlay inputs (OL0–OL3) are tied to ground as well as control inputs (RS2, SETUP and $\overline{SYNC}$). OPA is left floating while the pin V_{REF} is connected through a 0.1μF ceramic capacitor to V_{AA}. The additional V_{AA} pins are connected to the analog power supply while all GND pins are connected to the ground plane. Pins 1 and 2 (NC) are no connects.

All other external components and connections remain unchanged, including the RGB source terminations. If the 171/176 has 75Ω termination resistors, likewise the ADV471 should have 75Ω resistors. If the 171/176 uses 150Ω termination resistors (as is the case in the IBM PS/2*), the ADV471 should also be terminated using 150Ω resistors.

*IBM PS/2 and VGA are registered trademarks of International Business Machines Corp.

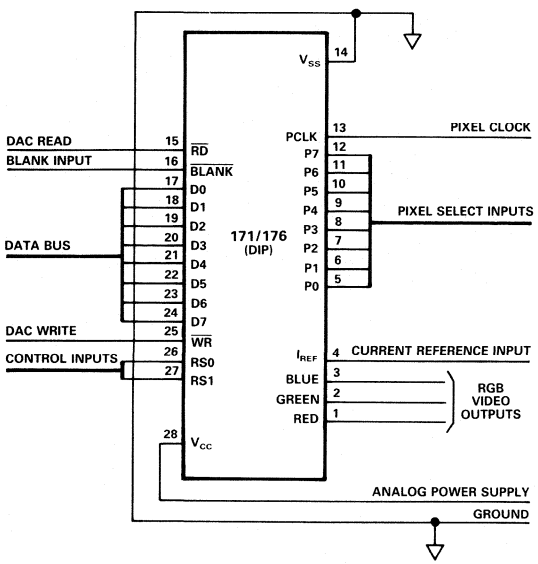

Figure 1. VGA Connection Diagram for 171/176

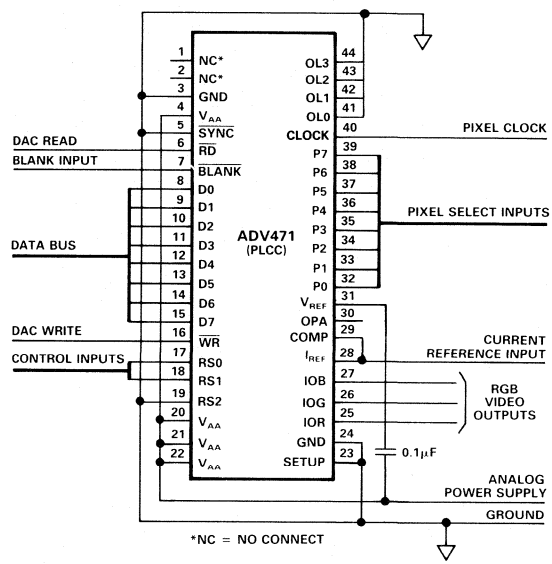

Figure 2. VGA Connection Diagram with the ADV471 Replacing the 171/176

ANALOG DEVICES

AN-332
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Improved PCB Layouts for Video RAM-DACs Can Use Either PLCC or DIP Package Types

by Bill Slattery

This application note describes printed circuit board layout schemes for the video RAM-DAC portion of a VGA compatible graphics card. The layouts allow either of two package types to be accommodated on one PCB.The user has the option of choosing either the 28-pin DIP 171/176 video DAC or the superior ADV471 which comes in a small 44-pin PLCC package. PCB patterns for two different types of layout are illustrated in this application note. The difference between the two layouts lies in the way in which the ADV471 PLCC is mounted on the board. The board designer has the option of either surface mounting or socket mounting the ADV471.

PLCC SURFACE MOUNT

Figure 1 illustrates how the PLCC can be directly surface mounted on the 28-pin DIP footprint. This is done by choosing a pin which can be used as a common point to both the DIP and PLCC pinouts. In this PCB design, Pin 1 (No Connect "NC") of the PLCC and Pin 24 (D7) of the DIP are the chosen coincident pins. The choice of these two pins as the common point optimizes board space as well as avoiding any manufacturing or solderability problems. This is ensured by selecting Pin 1 of the PLCC as the common pin. Pin 1 of the ADV471 is not internally connected within the package thus avoiding the need for it to be soldered onto the circular pad corresponding to Pin 24 of the DIP. Figure 1 is an X-ray view of the pattern showing both the component and solder (dotted-line traces) sides. It also indicates the DIP pin functions as well as some of the PLCC pin numbers. Figures 3 and 4 show the actual layout patterns for both the component side and solder side of the PCB.

PLCC SOCKET MOUNT

The designer can use a standard socket mount configuration for the PLCC, as shown in Figure 2. The PLCC is positioned with pin "D2" of both the DIP and PLCC sockets adjacent to each other. The PLCC socket recommended for use with this layout is available from AMP, part number 641747. Figures 5 and 6 give the PCB layout patterns for both the component side and solder side.

Note: PCB layouts on page 2 of this application note are at a scale of 2:1.

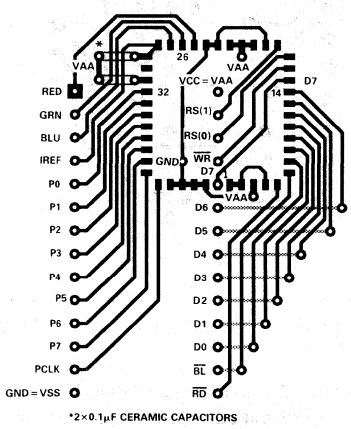

Figure 1. X-Ray View of Layout Using Surface Mounted PLCC

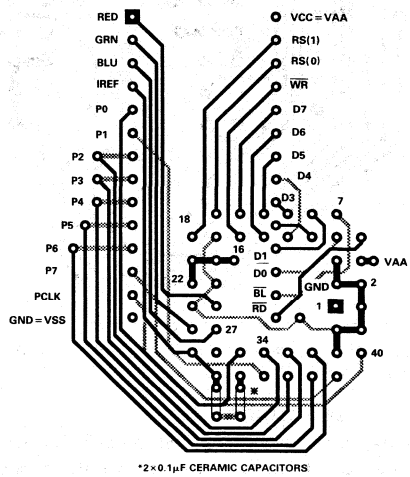

Figure 2. X-Ray View of Layout Using Socket Mounted PLCC

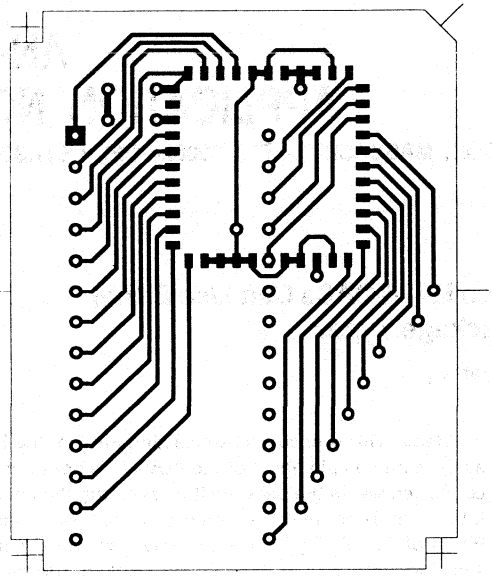

Figure 3. PCB Component Side Layout Pattern Using Surface Mounted PLCC

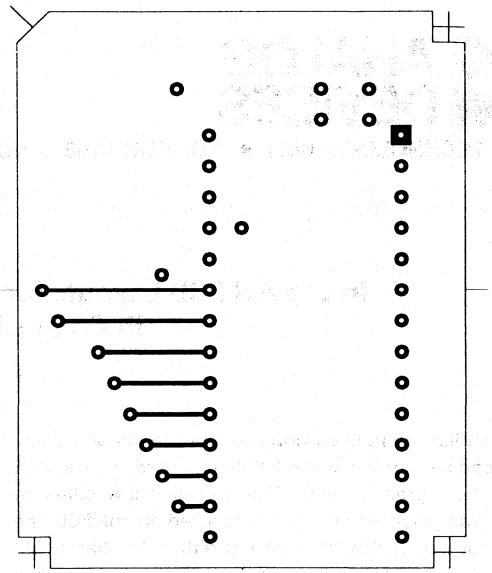

Figure 4. PCB Solder Side Layout Pattern Using Surface Mounted PLCC

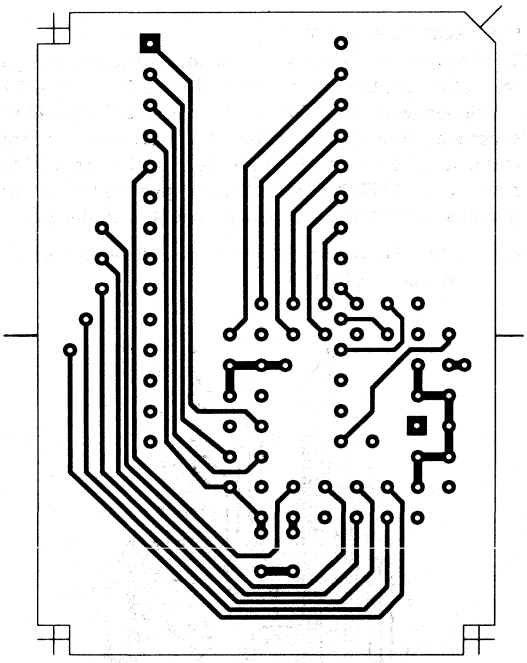

Figure 5. PCB Component Side Layout Pattern Using Socket Mounted PLCC

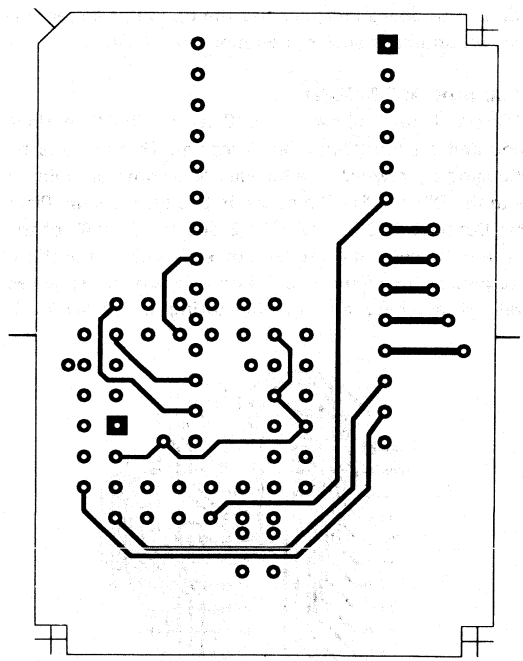

Figure 6. PCB Solder Side Layout Pattern Using Socket Mounted PLCC

LAYOUT CONSIDERATIONS

For optimum performance, the majority of PCB traces on both layouts vary between 12 and 15 mils, except in some cases where ground and power planes vary between 20 and 25 mils. The circular pads used are 42 mils in diameter on 28 mil diameter plated-through holes. Rectangular 55×36 mil surface mount pads are used for the PLCC surface mount option. Spacing between traces and pads for the analog circuitry is optimized to that required for a 75Ω transmission line.

ANALOG DEVICES

AN-333
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Design and Layout of a Video Graphics System for Reduced EMI

by Bill Slattery and John Wynne

The availability of low cost, high performance video RAM-DACs is a key element in the spread of personal computers into application areas previously considered the preserve of expensive, high-end computer systems. Applications such as Computer Aided Engineering (CAE), Computer Aided Design (CAD), Solids Modelling, Desktop Publishing, etc., are becoming more widespread as the cost of the necessary hardware drops. Successfully incorporating a video RAM-DAC into a personal computer or onto a video graphics plug-in board is not difficult once some basic concepts are grasped and some simple guidelines followed.

This application note is intended as a guide to the design of a video graphics system in terms of Electromagnetic Compatibility (EMC). EMC design will be considered as the technique of reducing radiated emissions and Electromagnetic Interference (EMI) from a high speed video graphics system. EMC implies that the system should not electrically or magnetically interfere with its surroundings, and conversely, the surroundings should not interfere with the operation of the system. In order to provide control of EMI in the radio spectrum, government agencies and other international organizations have established limits relating to EMI, most notably, the U.S. government's FCC Part 15.

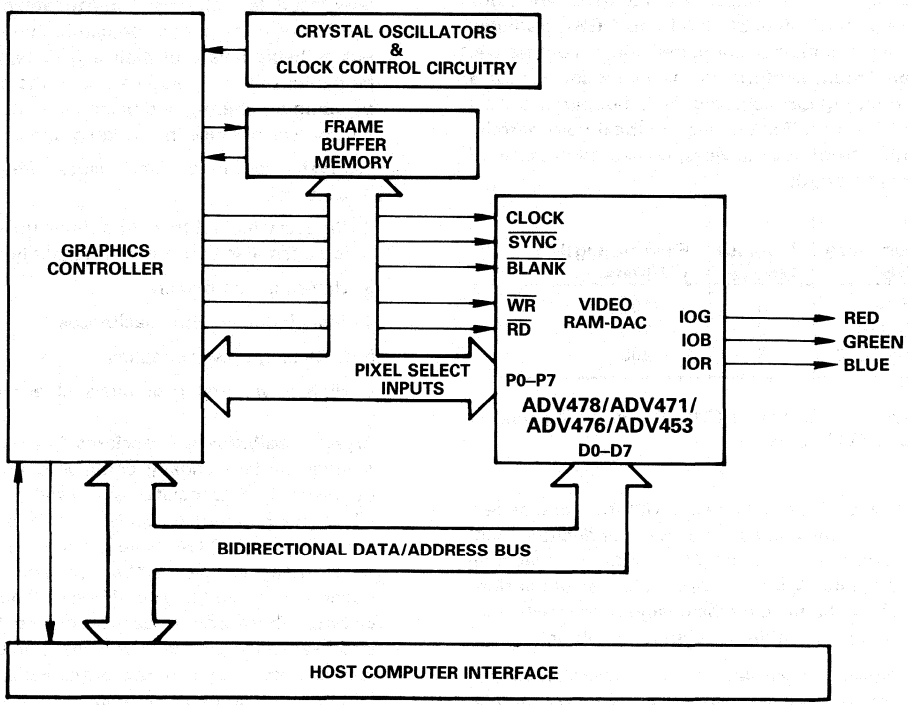

Simplified Block Diagram of a Typical Graphics System Using a Video RAM-DAC

OVERVIEW

This application note is divided into a number of sections as outlined below:

1. International EMI Regulatory Bodies – guidelines, testing and radiation limits.
2. System Noise Identification – identifying various sources of noise in a system.
3. PCB Layout & Design – component placement, multilayer boards, grounding, shielding and filtering components
4. Practical example of a VGA board design and associated FCC Testing.

REGULATIONS CONTROLLING EMI

The ultimate goal which must be achieved if EMC design is to be considered successful is the attainment of "Agency Certification." A number of international government agencies impose strict criteria on the allowable electromagnetic interference that electronic apparatus can emit. Electronic apparatus is required by law to conform to these agency limits, or else face severe government penalties.

In the United States, the Federal Communications Commission (FCC) is the national regulatory body which sets down strict controls on interference from computing devices. The FCC has divided computer interference into two principal types. The first type, and by far the most demanding of the two, deals with radiated emissions over the frequency range of 30 MHz to 1 GHz. Radiated emissions from personal computers, in a commercial environment, must conform to the limits set out for a Class B computing device pursuant to Subpart J of Part 15 of the FCC Rules. Table I lists the maximum permissible radiation from such devices, measured in terms of Electric Field Strength.

Frequency MHz	Distance Meters	Field Strength* μV/Meter
30–88	3	100
88–216	3	150
216–1000	3	200

Table I. Radiation Limits for Class B Computing Devices According to FCC Rules

The second type of emission deals with interference fed back onto the power lines. The FCC conduction limit on this interference is 250 μV* maximum over the frequency range 450 kHz to 30 MHz. This type of interference is heavily influenced by the design of the switched-mode power supply within the computer cabinet.

FCC Certification is awarded on the submission to the FCC of a complete report which consists of acceptable test results as well as a detailed description of the test and measurement procedure. Testing has to be carried out by an FCC-accredited test laboratory.

Class C Certification, which has less stringent limits, is allowable in certain commercial and industrial applications.

A list of the various international agencies is given in the Reference section. All agencies have very similar requirements to those of the FCC.

NOISE SOURCES

Identification of noise sources or potential noise sources in a system is the first and probably the most valuable step that has to be taken for successful EMI design.

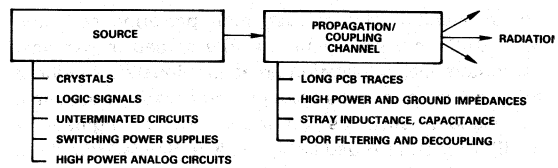

Figure 1. Noise Model of a Video Graphics System

It may even be possible to eliminate a particular noisy circuit completely from the design thus avoiding the later need for filtering. Unfortunately, many possible sources of noise cannot be eliminated from the design, but by being aware of their existence, their effects can be minimized at the source and in the coupling channel by optimum filtering and decoupling. Some of the inherent sources of noise in a video graphics system include:

1. Crystal oscillator and clock frequency division circuits.
2. Circuits with fast transition times (rise/fall times), e.g., logic families that are unnecessarily fast.
3. Unterminated circuits.
4. Stray inductances/capacitances.
5. Switching power supplies.
6. High power analog circuits such as video RAM-DACs.

Crystal Oscillators & Associated Circuitry

A video system usually contains a number of crystal oscillators and associated clock and pixel data circuits, which are required to achieve various on-screen pixel resolutions. In a VGA system, for example, there could be as many as five crystal oscillators varying in frequency from 25 MHz to 65 MHz, and maybe up to 80 MHz. These crystal oscillators and their associated circuitry tend to be rich in unwanted noise and harmonic components. They can be a prime source in the generation of EMI if some basic guidelines are not followed.

***Measured pursuant to §15.840 of the FCC Rules.**

Some of the important actions are:

1. Place crystal oscillator circuits as far as possible from analog circuitry and video output connectors.
2. Isolate power supply to crystals through the use of ferrite beads.
3. Avoid the mixing of clock buffers and other logic in the same IC package.
4. Use several low power data drivers or buffers for clock and pixel data lines distributed throughout the board, in preference to using a single high power driver.

With regard to the crystal oscillators themselves, the critical aspects which must be considered include the wave shape and the transition time (rise/fall time). Figure 2 is a plot of the output frequency spectrum of a typical 28.5 MHz crystal oscillator. It shows the amplitude of the harmonic components relative to the fundamental. It can be clearly seen that although the crystal's fundamental frequency lies outside the FCC's lower limit of 30 MHz, for radiated EMI, higher-order harmonic components exist throughout the FCC-controlled band. The frequency of the crystal oscillators should be kept to the required minimum, and transition times should be kept as slow as possible. This reduces the amplitude of unwanted harmonics, while still satisfying system functional performance needs.

The clock circuitry, which includes crystal oscillators and pixel data lines, is the primary source of most of a system's noise. Keeping this circuitry as far and as isolated as possible from other circuitry, especially analog circuitry, is all-important. On the other hand, it could be argued that by running long pixel data and clock lines from such circuitry to the Video RAM-DAC in itself is not desirable. Long lines increase noise coupling to other parts of the system. A tradeoff between length of pixel lines and the placement of high speed clock circuitry must be considered. The designer must attempt to optimize these two competing goals. As was mentioned earlier, the use of multiple low power buffers will help to ease such a conflict. A distance of less that three inches between buffers would be desirable in such circumstances.

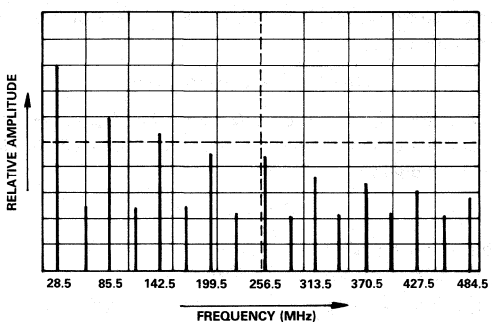

Figure 2. Magnitude of Harmonic Components Relative to the Fundamental for a 28.5 MHz Crystal Oscillator

Other Noise Sources

A number of other noise generators can be easily identified. Circuits with fast transition times including memory chips, logic circuitry and the graphics controller all contribute to the overall noise of the system. The faster the signal transition time, the greater will be the amplitude of the resulting harmonics, as was seen in the previous section relating to crystal oscillator circuits.

As a general rule, devices with the slowest possible rise/fall times that will achieve the system's tasks should be used. Table II lists some typical rise/fall times for various logic families.

Technology	IC Family	Transition Time
TTL	74	10 ns
	74LS	12 ns
	74ALS	3→20 ns
	74S	6 ns
	74AS	2→9 ns
	74F	1.2→8 ns
CMOS	74HC	20→150 ns

Table II. Comparison of Transition Times for IC Logic Families

Stray inductances and capacitances can cause signals to ring, to overshoot or undershoot the steady state voltage levels. This ringing is a source of EMI which can be minimized by keeping wires or traces short and adding series, damping resistances at the source or termination of long signal paths.

Unterminated circuits with floating signal lines should be avoided. Unwanted oscillations can result.

Power to all devices of a system is usually derived from switch-mode power supplies. While the design of the power supply is critical to the reduction of conducted noise in the FCC's band of 450 kHz up to 30 MHz, harmonics generated by the switching power supplies can extend well into the radiation frequency band and thus add to EMI.

Modern high resolution color graphics monitors are driven by analog signal levels direct from the DACs. The relatively high power, analog output levels from the Red, Green and Blue current sources of the video DAC require careful attention. As will be discussed in the PCB layout section, the power to the Video RAM-DAC should be isolated from the remainder of the PCB power plane. If noise on the high speed pixel and clock input section to the Video RAM-DAC has not been minimized, noise will be coupled through to the analog output section and onto the connecting cable to the monitor, causing this cable to act as an antenna. Filtering at the source termination of each of the three DAC outputs can be used if required to minimize the noise further. (See Appendix 1, *Three-Terminal Capacitor.*)

PRINTED CIRCUIT BOARD DESIGN

The extent of radiated emissions from a printed circuit board (PCB), will be determined by the effectiveness of the PCB to act as a propagation channel for unavoidable noise sources, its ability to couple this noise onto other circuitry, and the radiation into free space of this undesired noise. Apart altogether from a PCB's ability to radiate EMI, noise coupled from digital circuits on the board to the video RAM-DAC can adversely affect the system's functional performance.

Causes of EMI

The main sources which conduct or radiate EMI from a printed circuit board are as follows:

1. Common impedance coupling via power and ground traces.
2. Antenna loops formed by ICs and their bypass capacitors. Note that these loops also include the power and ground lead frame members within the IC packages.
3. Printed circuit board traces carrying signal currents. Note again that signal lead frame members within the IC packages are also included.
4. Crosstalk between adjoining signal traces.

Common Impedance Coupling

An example of common impedance coupling via power and ground traces is shown in Figure 3a where a number of logic gates are supplied with power over common printed circuit board traces. A typical V_{IN} input signal to one of these gates is shown in Figure 3b with the resulting transient and signal currents due to the gate switching also shown.

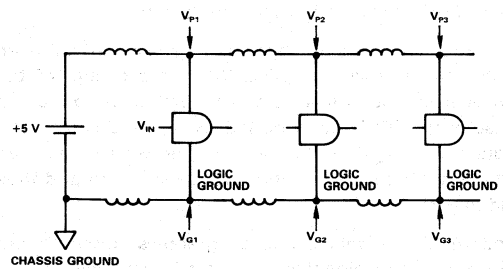

Figure 3a. Common Impedance Coupling via Power and Ground Traces

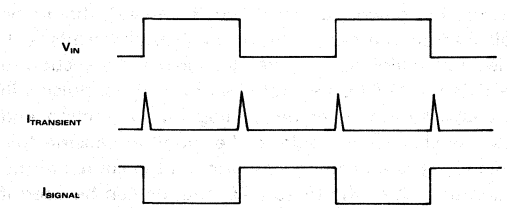

Figure 3b. Signal and Transient Currents Due to Gate Switching

The distributed trace inductances act as impedances to these switching currents spreading the resulting high frequency noise to all nodes common to the culprit. To get an idea of the magnitude of the generated high frequency noise, let's take a look at the effect of a single gate, as shown in Figure 4a. The associated worst signal current for a standard TTL gate is shown in Figure 4b.

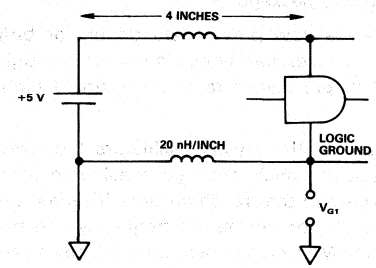

Figure 4a. Common Impedance Coupling Due to One Gate

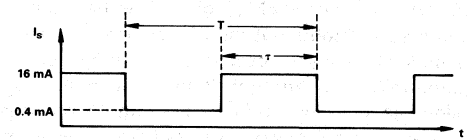

Figure 4b. Signal Current Due to TTL Gate Switching

Consider the harmonic components present in the switching signal current of Figure 4b. This assumes the gate is driving 10 standard TTL loads with a maximum sink current of 16 mA and a maximum source current of 0.4 mA.

With a mark/space ratio of τ/T and using a Fourier Series expansion, the formula for the amplitude of the nth harmonic is given by

$$I_n = \frac{2A\tau}{T}\left[\frac{Sin\left(\frac{n\pi\tau}{T}\right)}{\frac{n\pi\tau}{T}}\right]$$

where n = 1, 2, 3 . . . .

For a square wave $\tau/T = 0.5$, the amplitude of the third harmonic (n=3) is

I_3= 3.4 mA, zero to peak.

At 28.5 MHz, a standard VGA pixel clock frequency, the magnitude of the impedance of the ground trace in Figure 4a is given by

$$Z = [2\pi fL]$$

where f = 28.5 MHz
L = 20 nH/inch (typically)

hence Z = 3.58 Ω/inch.

At 85.5 MHz (3×28.5 MHz) the impedance is 10.74 Ω/inch. Thus the high frequency voltage at the logic

ground node of the switching gate due to the third harmonic alone is equal to

$$V_{G1} = (3.4 \times 10^{-3})(4)(10.74)\ V$$

$$= 146\ mV\ peak\ at\ 85.5\ MHz.$$

This high frequency component and other similar components will be circulated around the printed circuit board via the common ground traces. It will also appear on any cable shielding attached to this common ground trace and, depending on how efficient the cable shield is as an antenna, will be radiated into free space.

Antenna Loops

One of the most important principles of PCB layout and design for noise reduction can be described by the phrase:

"Minimize Signal Loop Areas."

In most circuit designs, we tend to think of the currents we're interested in as flowing "out" of one place, "through" some other place and "to" the target point. Unfortunately however, this often leads us to neglect to consider how these currents will eventually find their way back to their source. Ground and supply voltage points are considered "equivalent," and the fact that they are parts of a network of conductors through which currents flow and develop finite voltages is often not appreciated. These voltages can radiate to cause EMI, see Figure 5.

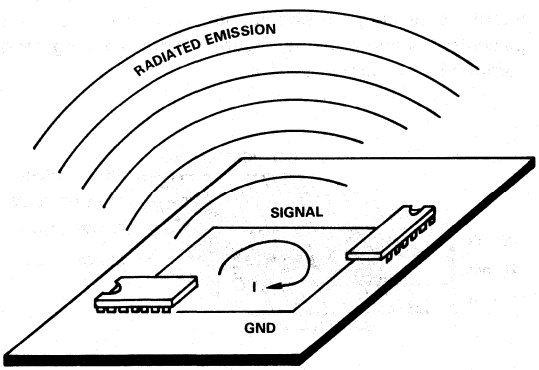

Figure 5. Currents Flowing in Large Loops Add to EMI

Voltages are generated because wires and traces do not have zero impedance due mainly to inherent inductances.

Many of the problems associated with power and ground loops can be avoided through the deployment of effective bypassing techniques.

The aim of effective bypassing is to maximize the charge stored in the bypass loop while simultaneously minimizing the inductance of this loop. Inductance in the loop acts as an impedance to high frequency transients and results in power supply spiking. Figure 6a shows a poor bypass arrangement and the associated inductances due to the large loop area are illustrated in Figure 6b.

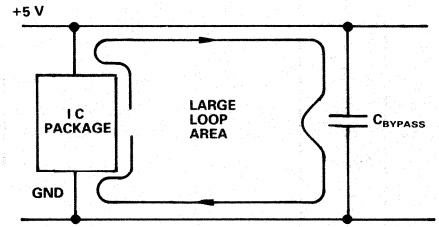

Figure 6a. Large Loop Associated with Poorly Placed Bypass Capacitor

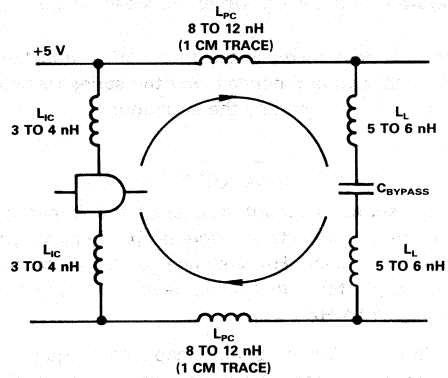

Figure 6b. Equivalent Circuit of Bypass Loop of Figure 6a

where L_L = inductance of the lead from the capacitor body to the PC board

L_{PC} = inductance of the trace between the lead arrival on the PC board and the IC pin

L_{IC} = inductance of the lead frame member carrying power within the IC package.

As well as loop inductances due to the above, the series inductance of the bypass capacitor itself must also be considered. It is well known that there is more inside a capacitor's body than a pure capacitance.

The simplified equivalent circuit of a 0.1 μF capacitor in Figure 7 shows an effective series resistance (ESR) and effective series inductance (ESL) in series with the ideal 0.1 μF capacitance.

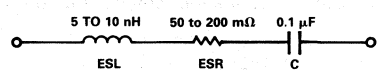

Figure 7. Equivalent Series Representation of a Bypass Capacitor

Figure 8 shows the complete inductive loop associated with the bypass circuit.

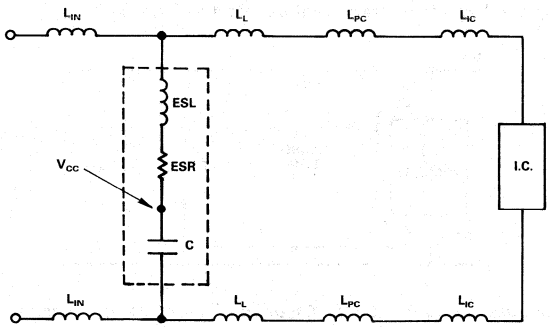

Figure 8. Inductive Loop of a Bypass Circuit

These inductances increase the total series inductance of the bypass loop and hence lower the series resonant frequency as determined by the equation:

$$f_O = \frac{1}{2\pi\sqrt{LC}}$$

Above the series resonant frequency the impedance becomes more inductive, increasing linearly with increasing frequency. For instance, a 0.1 μF ceramic radial lead capacitor with 1/4 inch leads generally resonates around 10 MHz.

The minimum value of bypass capacitor required is determined by the maximum amount of voltage drop allowable across the capacitor as a result of the transient current. An approximate value for a bypass capacitor is given as

$$C = \frac{I.\Delta t}{\Delta V} \; Farads$$

where I = Maximum Transient Current
Δt = Transient Duration
ΔV = Allowable Voltage Drop.

For example, a typical 74 HC I_{CC} transient is 20 mA high lasting 20 ns. If the voltage drop is to be kept below 100 mV, then the required bypass capacitor is

(20 mA) (20 ns) / (100 mV)

or 4 nF per output.

However, any series inductance in the bypass loop will cause additional voltage spiking. For any given magnitude of noise spike, an approximate expression for the maximum amount of series inductance is given by

$$L = \frac{V.\Delta t}{\Delta I} \quad Henrys$$

where V = Maximum Noise Spike
Δt = Transient Duration
ΔI = Transient Current.

The typical 74 HC I_{CC} transient of 20 mA has a rise/fall time of 4 ns. If we wish to restrict the inductive noise spike to, say, 100 mV peak, the maximum amount of series inductance is

(100 mV) (4 ns) / (20 mA)

or 20 nH.

Referring back to Figure 8, this means that the combined total of ESL, L_L, L_{PC} and L_{IC} must be kept below 20 nH. To a greater or lesser extent the first three terms are within the PC board designer's influence; the fourth term, the inductance of the IC lead frame member or L_{IC}, is invariable, being determined by the IC package. The use of PLCC packaged parts, such as the ADV478/ADV471, inherently reduces L_{IC} to 2–3 nH as against 10–12 nH for the more traditional DIP parts (see section on "Surface Mount Technology").

Appendix 1 examines in greater detail the characteristics and filtering capabilities of various bypass elements including two and three terminal capacitors.

Multilayer PC Boards

In the design of a high performance, high speed graphics system, it is recommended that a four-layer printed circuit board be used.

Figure 9 shows a cross-sectional view of a four-layer printed circuit board, with power and ground planes separating the signal-carrying traces of the component and solder sides of the PCB. As well as using multilayer boards, consideration should be given to the relevant placement of components. Figure 10 shows a suggested component placement scheme.

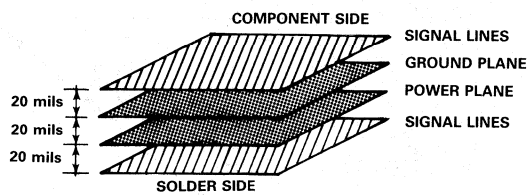

Figure 9. Four-Layer Printed Circuit Board Construction

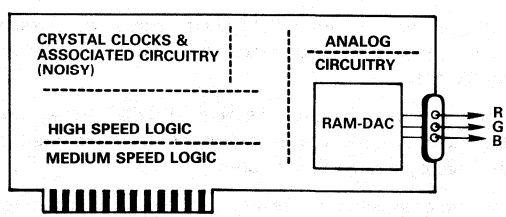

Figure 10. Printed Circuit Board Component Placement

Power and Ground Planes

Power supply decoupling attempts to contain the transient currents within the bypass loop. However it cannot be 100% successful, and some high frequency components will escape onto the power and ground traces. High frequency signal currents will also be flowing in the power and ground traces. In order to avoid common-impedance noise coupling due to these currents, it is necessary to reduce the impedance of the power and ground traces to an absolute minimum. The only satisfactory way to achieve this is not to use traces at all but to use power and ground planes. On a PC-card-sized, two-layer board with one side devoted to a ground plane, the impedance of the plane is in the tens of milliohms range.

A four-layer board allows another plane to be used as a power plane. Low impedance power and ground contacts are thus available over the full area of the board. Additionally, in a four-layer board with power and ground planes inside the board and signal traces on the top and bottom of the "sandwich," overlapping power and ground planes act as an inherent distributed capacitor, as shown in Figure 9. This provides some measure of high frequency decoupling. From the signal interconnect point of view, the major advantage of using a ground plane is the very substantial reduction in signal loop area it provides. In a typical PCB layout, signal current flows out through one trace and back through a ground trace. Such a path can include a large loop area; a large loop area, as has already been discussed, implies high inductance for the traces with follow-on consequences of signal ringing, EMI radiation and crosstalk. To reduce the inductance it is necessary to reduce the loop area through which the signal current flows. The use of power and ground planes minimizes loop areas, thereby reducing inductances and resulting EMI.

The electromagnetic fields associated with an idealized case of two parallel wires carrying equal and opposite currents are shown in Figure 11.

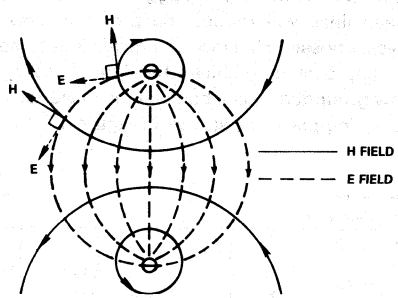

Figure 11. Electromagnetic Field about Two Parallel Conductors Carrying Equal and Opposite Currents

The two fields (electric, E, and magnetic, H) tend to be confined between or near the conductors. The electric field is strongest in the plane of the conductors. The magnetic field is nonzero at points close to the conductors, but farther away (relative to the wire spacing) the fields from both conductors tend to cancel out. Keeping the conductors together promotes field cancellation which can be viewed either as minimizing the loop area or minimizing the inductance; the results are the same.

Introducing a ground plane (sheet of copper) halfway between the wires, as shown in Figure 12, does not disturb the field pattern even when the lower wire is removed.

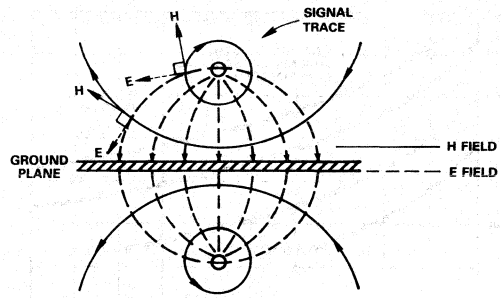

Figure 12. Electromagnetic Field about Two Parallel Conductors Separated by a Ground (Copper) Plane

8

A virtual image of the lower wire has been produced in the copper plane maintaining the original field configuration. This is the basis of microstrip. With a properly designed ground plane system, the return current will always flow under the signal trace, the path of lowest impedance.

Digital Signal Interconnections

The use of a ground plane allows the signal interconnects to be viewed as microstrip transmission lines whose characteristic impedances, propagation delays, etc., can be readily calculated. Microstrip is the name given to a transmission line which consists of a signal trace separated from a ground plane by a dielectric. Figure 13 shows the cross section of such a line.

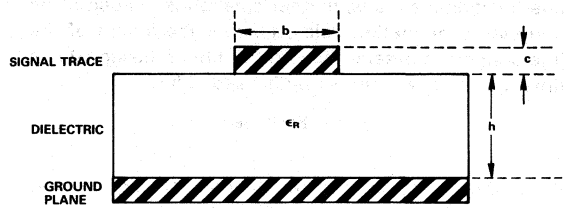

Figure 13. Cross Section of Microstrip Transmission Line

The characteristic impedance, Z_0, of this line is

$$Z_O = \frac{87}{\sqrt{\epsilon_R + 1.41}} Ln \left\{ \frac{5.98\,h}{0.89\,b + c} \right\} \Omega$$

where ϵ_R = Relative Dielectric Constant of Board

typically ϵ_R = 5 for glass/epoxy boards.

b, c, h = dimensions indicated in Figure 13

The propagation delay, t_{PD}, of a microstrip line is given by

$$t_{PD} = 1.017 \sqrt{0.475\,\epsilon_R + 0.67}\ ns/ft.$$

Note that this propagation delay is dependent only on the dielectric constant and not on the line geometry.

The graph below shows impedance values for various configurations of microstrip line.

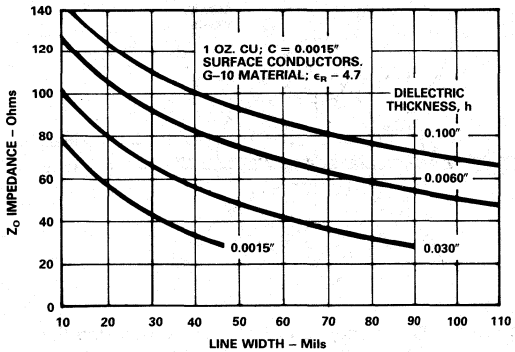

Figure 14. Impedance Versus Line Width & Dielectric Thickness for Microstrip Lines

Gross impedance mismatches between the transmission line's characteristic impedance and the source (driver output) or load (receiver input) impedances connected to the line reflect the signal back and forth on the line. These reflections will cause overshoot, EMI radiation and crosstalk. By properly terminating the line with either source or load impedances which match that of the transmission line, reflections can be eliminated or substantially reduced. However, not every signal interconnect demands line termination; the need is determined by the relationship between the rise (or fall) time of the signal and the time required for the signal to travel the length of the interconnect. As a general guideline for digital signals, line termination is needed if the one way propagation delay, t_P, over the length of the interconnect is greater than one eighth of the signal rise time, t_R, i.e., line termination is needed if

$$t_P \geq (1/8) \bullet t_R\ secs.$$

A number of dc and ac termination techniques exist which trade increased power dissipation against component count. The simplest termination technique which dissipates no extra power is a series termination one where a resistor is placed in series with the signal interconnect at the source end of the line, see Figure 15. The resistor should have a value equal to the characteristic impedance of the line minus the output impedance of the driver and should be of metal-film construction or some other low-inductance material. The load impedance is considered an open circuit. Series termination is most suitable for systems where only one receiver (e.g., ADV478/ADV471) is connected to the line. Note that if pull-up resistors are required on digital or clock signals, they should be connected to the PCB Power Plane (V_{CC}).

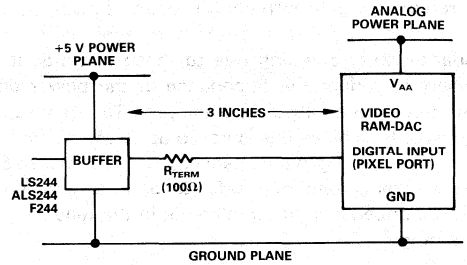

Figure 15. Series Termination of Signal Lines

Crosstalk

Crosstalk is any unwanted signal coupling between parallel PC-board traces due to mutual inductance (L_M) and capacitance (C_M), as illustrated in Figure 16.

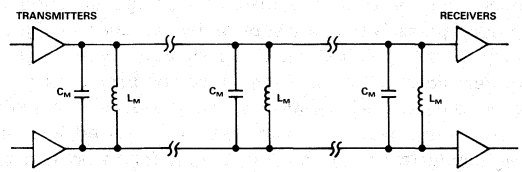

Figure 16. Capacitive and Inductive Coupling between Signal Traces

In general, crosstalk is directly proportional to line impedances, frequency and line lengths and inversely proportional to line spacing. Much of the induced crosstalk in a signal line is from immediately adjacent transmission lines which suggests that wider spacing between lines will reduce the problem. This may not always be possible in closely spaced circuits, so an alternative approach is to shield the signal lines by inserting narrow grounded traces between each signal line on the same wiring plane, as shown in Figure 17.

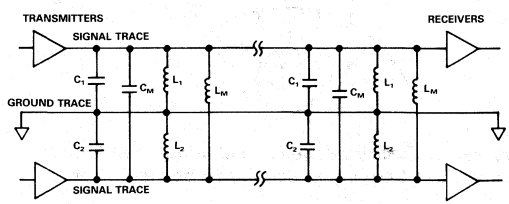

Figure 17. Ground Trace between Signal Lines Reduces Crosstalk

At high frequencies capacitive coupling dominates. The addition of a shield (or ground trace) between the signal lines changes the equivalent circuit. Crosstalk is now reduced since the inductance L_M is now much larger than either L_1 or L_2 and capacitance C_M is much smaller than either C_1 or C_2.

Separate Power Plane for Video RAM-DAC

To further isolate the Video RAM-DAC from the PCB's power supply, V_{CC}, it is recommended that a separate power plane, V_{AA}, be used for the video RAM-DAC and its associated circuitry. This analog power plane should be connected to the regular PCB power plane (V_{CC}) at a single point through a suitable filtering device such as a ferrite bead (see Appendix 1 – *Ferrite Bead Inductor*). This ferrite bead should be located no more than three inches away from the Video RAM-DAC. In the case of Analog Devices' ADV478 and ADV471, which have multiple power (V_{AA}) pins, it is important to connect all these V_{AA} pins to the analog power plane. This eliminates any possibility of latchup in the device.

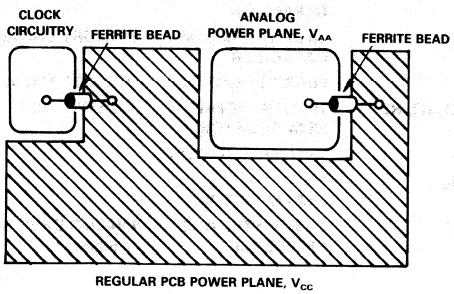

Figure 18. Power Plane Decoupling Using Ferrite Beads

Common Ground Plane

Due to the presence of RAM on board the ADV478/ADV471 and ADV476, it is not recommended to isolate the device's ground circuitry from the main PCB ground. Corruption of data could occur. These Video RAM-DACs should have all GND pins connected to the PCB's regular ground plane.

Analog Outputs

The analog outputs of Analog Devices' video RAM-DACs are driven by switched current sources. These parts are designed to drive either a singly or doubly terminated 75 Ω load. The doubly terminated configuration shown in Figure 19 is the preferred choice.

DAC output traces on a PCB should be treated as transmission lines. It is recommended that the Video RAM-DAC be placed as close as possible to the output connector, with the analog output traces being as short as possible (less than 3 inches). The 75 Ω termination resistors should be placed as close as possible to the DAC outputs and should overlay the PCB's ground plane.

As well as minimizing reflections, short analog output traces will reduce noise pick up due to neighboring digital circuitry.

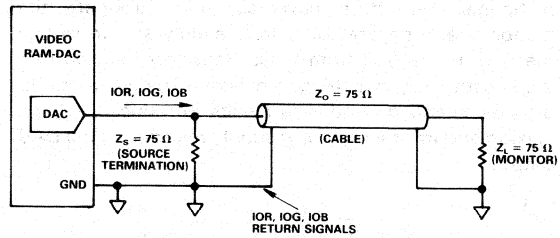

Figure 19. Recommended Analog Output Termination for Video RAM-DACs

Surface Mount Technology (SMT)

Surface mount technology (SMT) offers many EMI advantages over traditional through-hole designs. Because SMT allows the designer to place components on both sides of the printed circuit board as well as featuring smaller component sizes, it provides superior PCB integration. This means that loop lengths can be reduced and noisy signal traces can be shortened, all having a positive effect on EMI. The shorter lead lengths of SMT packages decrease inductance, thereby providing better high frequency performance.

Many, if not all components required for a video graphics system are available as surface mount devices. Memories, controller chips and logic are all available in small SMT packages. Resistors and capacitors can also be purchased in a small "chip" format.

As well as producing Video RAM-DACs in a dual-in-line package (DIP), e.g., the ADV476 (28-pin DIP), Analog Devices packages its ADV478/ADV471 in 44-pin plastic leaded chip carriers (PLCC). This package has substantial advantages over DIP packages in that the lead-frame inductance is small (2-3 nH) and constant for any pin around the package. A DIP package usually has power and ground on diagonally opposing corner pins which presents a much larger lead-frame inductance (10–12 nH). Additionally lead-frame inductance varies with pin position.

In addition to the PLCC package the ADV478/ADV471 video DACs have a number of design features intended to minimize EMI problems.

The pinout of the ADV478/ADV471 in Figure 20 shows four power pins and two ground pins. In operation, the four V_{AA} pins are tied together at the package and supplied with a single +5 V supply. Similarly the two ground pins are tied together at the package and connected to the PCB ground. Internally, however, power rail routing has been separated according to functionality. The various V_{AA} pins are used to drive different internal sections of the ADV478/ADV471. One V_{AA} pin provides a common power rail for "digital" logic;

another provides an "analog" power rail for the DAC and reference circuitry; while yet another provides power to the n-substrate of the device. Since all p-channel transistors have as their back-gate the n-substrate, a separate substrate supply acts to isolate the DAC power rail from noise transients injected into the substrate by switching transistors. The use of multiple power and ground pins results in reduced voltage spiking and improved power supply rejection at the DAC outputs.

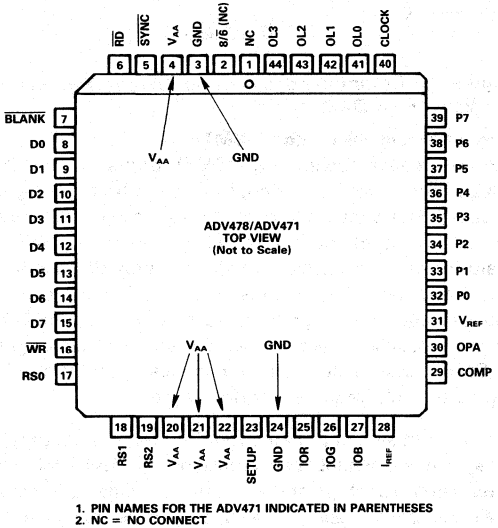

1. PIN NAMES FOR THE ADV471 INDICATED IN PARENTHESES
2. NC = NO CONNECT

Figure 20. ADV478/ADV471 PLCC Pin Assignment Showing Multiple Power and Ground Connection Points

CIRCUIT LAYOUT FOR THE ADV478/ADV471

A recommended layout and component listing for the ADV478/ADV471 is shown in Figure 21. More details regarding the characteristics of various decoupling and filtering components can be found in Appendix 1 of this application note.

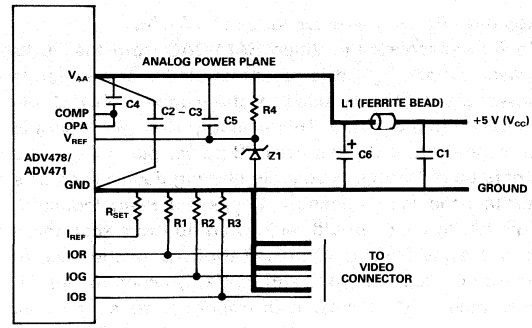

COMPONENT	DESCRIPTION	VENDOR PART NUMBER
C1 – C5	0.1 µF CERAMIC CAPACITOR	ERIE RPE112Z5U104M50V
C6	10 µF TANTALUM CAPACITOR	MALLORY CSR13G106KM
L1	FERRITE BEAD	FAIR-RITE 2743001111
R1, R2, R3	75 Ω 1% METAL FILM RESISTOR	DALE CMF-55C
R4	1 kΩ 5% RESISTOR	
R_{SET}	1% METAL FILM RESISTOR	
Z1	1.235 V VOLTAGE REFERENCE	ANALOG DEVICES AD589KH

Figure 21. Connection Diagram and Component Listing for the ADV478/ADV471

GETTING FCC CERTIFICATION

The final chapter in EMI design is the actual test. With good design practice and even a partial adherence to some of the issues raised in this application note, no difficulty should be encountered in achieving certification. The testing itself however must be carried out with great care in order to do justice to your design. It should be remembered that in the case of a peripheral device such as a VGA board, FCC testing is applied to a complete operating computer system. This complete system must pass with the VGA board present before the VGA board itself is considered to have passed.

A complete operating computer system is configured using the following:

1. a PC compatible computer with keyboard
2. a printer connected to the printer port
3. a mouse or modem connected to the serial port
4. a monitor.

The PC and all its peripherals must be operating when measurements are made.

It is paramount that only the best equipment is used. If for example a noisy monitor is used, the test results might not pass the agency limits, not because your board is at fault, but perhaps because of a poor quality, noisy monitor used in the test. An excellent choice of monitor is IBM's 851X series. Another principal culprit, causing EMI in such a system, is the parallel printer cable. A good quality, shielded cable must be used.

If you are using an outside test house, it is advisable to be present during testing, or at least have a representative from your company who understands the operation of the system and its various components.

AD/VGA

Analog Devices has designed its own high performance graphics board, AD/VGA for evaluation purposes. The board design is based on the ET3000AX* Video Graphics Controller from Tseng Labs and the high performance ADV478/ADV471 Color Palette RAM-DAC from Analog Devices. The board is fully compatible with all IBM PC† video standards as well as IBM PS/2† Video Graphics Array (VGA). It has additional modes including 800 × 600 resolution with 256 colors as well as 1024 × 768 in 16 colors. These modes require pixel data rates of up to 45 MHz or 66 MHz. The silkscreen showing component placement and type for the AD/VGA board is shown in the appendix.

The board has been certified to comply with the limits for a Class B computing device pursuant to Subpart J of Part 15 of FCC Rules.

FCC ID: HRF55L8826VGA

One actual set of measurements for radiated emissions from the AD/VGA board is shown in Appendix 2. Results can be compared to the FCC limits as listed in Table I.

The AD/VGA board was installed in a DELL SYS 200‡ computer using an IBM 8514 monitor. A modem, printer and mouse were attached. These measurements were made in 132 × 44 text mode.

A complete set of test results is available for inspection.

REFERENCES

1. International EMI Emission Regulations

USA:	FCC–15 Part J
West Germany:	VDE 0871/VDE 0875
Canada:	CSA C108.8–M1983
Japan:	CISPR(VCCI)/PUB 22

2. EMI Countermeasures – Application Guidance – Murata Mfg. Co. Ltd.
3. Henry. W. Ott, "Noise Reduction Techniques in Electronic Systems." John Wiley, N.Y., 1986.
4. Paul A. Brokaw, "An I.C. Amplifier Users' Guide to Decoupling, Grounding and Making Things Go Right for a Change." Analog Devices Data-Acquisition Databook 1984, Volume 1, Pages 20–13 to 20–20.
5. Motorola Inc., "MECL Design Handbook" 2nd Edition, 1972.

8

†IBM PC, IBM PS/2, IBM8514/A and VGA are trademarks of International Business Machines Corp.
‡DELL SYS 200 is a trademark of Dell Computer Corporation.
*ET3000AX is a trademark of Tseng Laboratories, Inc.

APPENDIX 1

EMI FILTERING COMPONENTS

No matter how well a PCB is laid out, there will always be a need for some kind of filtering. This section of the application note examines, in some detail, a number of filtering devices which are suitable for a high speed graphics system. The frequency characteristics of such devices as well as their inherent limitations will be discussed.

The effect of an EMI filter is generally expressed in terms of insertion loss. Noise suppression is described as a logarithm of the ratio of the output voltage without a filter to that with a filter in the circuit and is normally expressed in units of dBs. The simplest example is a first order device, a parallel capacitor or series inductor. A first order filter has an insertion loss slope of 20 dB per decade.

Table 1-1 shows a list of suitable filtering devices and their useful frequency bandwidth.

Filtering Device	Effective Bandwidth
Two-Terminal Capacitor	100 kHz–50 MHz
Ferrite-Bead Inductor	10 MHz–500 MHz
Three-Terminal Capacitor	1 MHz–800 MHz
Four-Terminal LC Filter	100 kHz–1 GHz

Table 1-1. Effective Bandwidth of Various Filtering Devices

The favorite and most widely used filtering device in electronic apparatus today is undoubtedly the "Bypass-Capacitor." It is simple to use, very effective and cheap. Unfortunately though, its effect in a high frequency graphics systems is sometimes the opposite to what is desired. A little thought regarding choice of capacitor, in terms of construction and value can lead to a dramatic improvement in noise performance.

Two-Terminal Capacitor

Let us first take a more detailed look at the structure of a real capacitor.

The impedance of an ideal capacitor connected between line (V_{CC}) and ground is given by

$$Z_C = \frac{1}{2\pi f C.}$$

From Figure 1-1, we can see that the insertion loss of an ideal capacitor increases with frequency at a rate equal to 20 dB/decade. In other words, the capacitor's filtering effect is greatest for higher value frequency components. Increasing the value of capacitance has the effect of filtering out lower frequency components. In the real world, however, a two-terminal capacitor has inductance in series with the capacitance, due to the inherent inductance of the lead wires as shown in Figure 1-2. The reactance characteristic of such a capacitor is shown in Figure 1-3 with the composite insertion loss characteristic shown in Figure 1-4. Clearly, the inductance, ESL, limits the insertion loss. The residual inductance of a capacitor is a function of both the electrode construction as well as lead length. This inductance can vary between 5 nH and 150 nH.

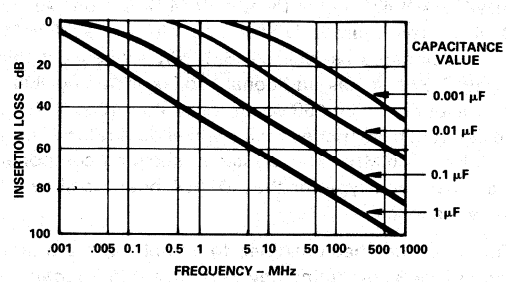

Figure 1-1. Insertion Loss Versus Frequency for Ideal Capacitors

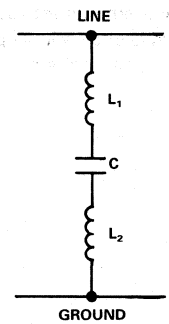

Figure 1-2. Real Capacitor Showing ESL Due to L_1 and L_2

Table 1-2 shows typical values of series inductance for various capacitor types.

Capacitor Type	Capacitance µF	Equivalent Series Inductance (ESL) nH
Lead Type Monolithic Ceramic	0.01	5
	0.1	5
	1.0	6
Disc/Lead Type Ceramic	0.0002	4.5
PolyethyleneTelephthalate	0.03	9
Mica	0.01	52
Polystyrene Film	0.001	12
	0.1	100
Tantalum Electrolytic (with Solid Electrolyte)	16	5
Aluminum Electrolytic		
RF Specific	470	13
Standard	470	130

Table 1-2. ESL for Various Capacitor Constructions and Values

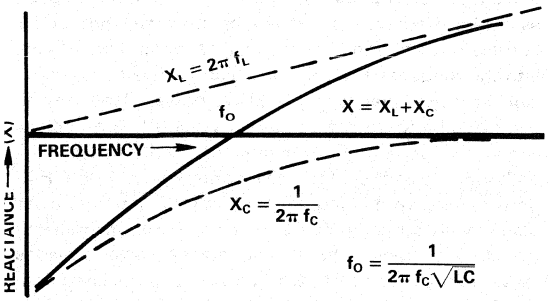

Figure 1-3. Reactance Characteristic of a Capacitor with Finite ESL

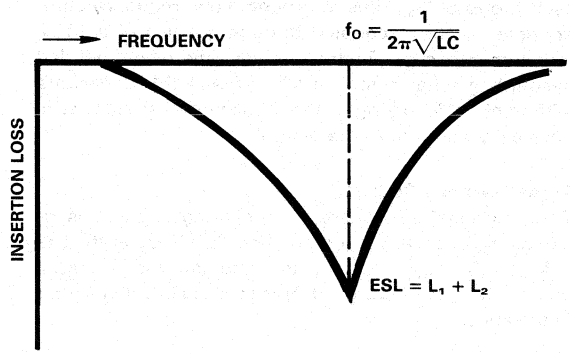

Figure 1-4. Insertion Loss of a Capacitor is Limited by f_O

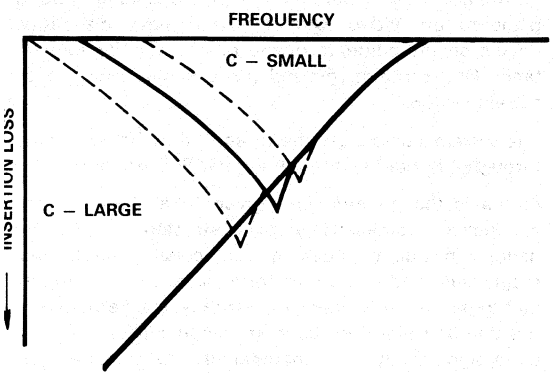

Figure 1-5. A Large Value Capacitor Has a Decreased f_O

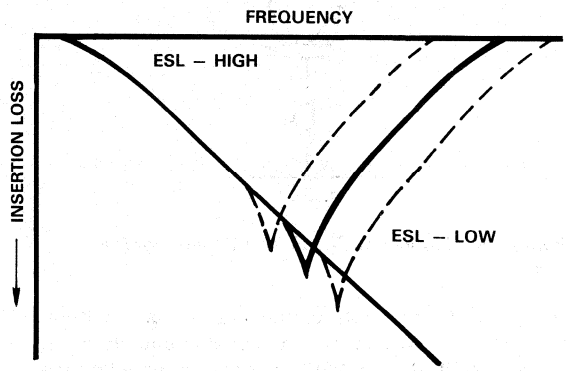

Figure 1-6. A Capacitor with Low ESL Has an Increased f_O

The resulting effect of this equivalent series inductance (ESL) is a reduction in the effectiveness of the capacitor at filtering frequency components beyond a certain point, known as the resonant frequency, f_O. It can be seen from Figures 1-4 to 1-6 that increasing the capacitance value can have the effect of reducing the resonant frequency thereby reducing the ability of this device to filter out higher frequency components. It is therefore imperative that a capacitor with low residual inductance be used. A good choice of high frequency decoupling capacitor would be a 0.1 μF lead type monolithic ceramic capacitor (MCC). This has a series inductance of approximately 5 nH. The resonant frequency, f_O, is thus kept high thereby maximizing high frequency rejection.

As well as series inductance, a capacitor will also contain resistance, referred to as equivalent series resistance (ESR). See Figure 1-7.

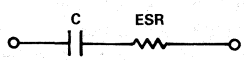

Figure 1-7. Real Capacitor showing ESR

This resistance imposes a finite limit on the ability of a capacitor to bypass high frequencies to ground. Total impedance of the device can never be lower than that imposed by the equivalent series resistance. A capacitor having a high ESR will exhibit a flattened insertion loss curve, compared with the sharp resonant point typically observed in capacitors with lower ESR values.

The overall equivalent circuit of a two-terminal capacitor with ESL and ESR is shown in Figure 1-8.

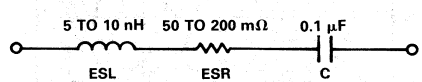

Figure 1-8. Equivalent Circuit of a Real Capacitor

Maximum noise reduction will be achieved through the selection of capacitors with lowest ESL and ESR values.

At least one decoupling capacitor should be used for each IC in the system. In the case of the Video RAM-DAC, it is recommended that for high frequency suppression, a 0.1 μF ceramic capacitor be used to decouple

each group of V_{AA} pins to ground. Low frequency components can be decoupled through the use of a 10 μF tantalum capacitor. Figure 21 shows the recommended decoupling scheme for the ADV478/ADV471. Capacitors with minimal lead length, should be placed as close as is physically possible to the device.

Three-Terminal Capacitor

Three-terminal capacitors have filtering characteristics extending to several hundred MHz. Two leads at the line side of the capacitor provide line input and output respectively, see Figure 1-9. This reduces effective series inductance.

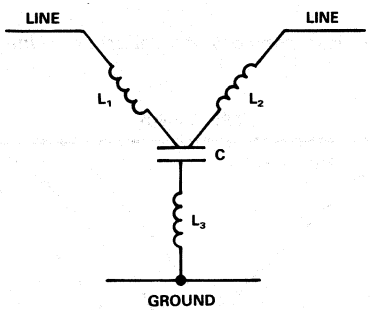

Figure 1-9. Equivalent Circuit of a Three-Terminal Capacitor

One particular example of the use of such a three-terminal device would be at the analog outputs of the Video RAM-DAC. The NFV510 series of three-terminal capacitors from Murata have a sharp insertion loss characteristic. This means that high frequency signals can be filtered without affecting the integrity of the signal itself. The NFV510 series has an excellent shape factor, due to an inherent roll-off of 100 dB/decade, see Figure 1-10. It therefore suppresses noise without reducing resolution. In the case of a graphics system with video output rates of 80 MHz, the NFV510-655 T2A 107 could be employed at the outputs of each of the red, green and blue DACs.

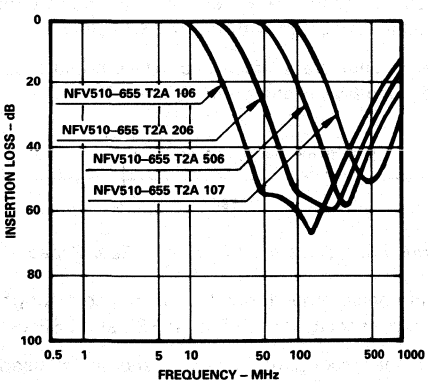

Figure 1-10. Frequency Response of the NVF510 Series of Three-Terminal Capacitors

Ferrite Bead Inductor

A ferrite is another very useful component in the effective suppression of EMI. One can consider a ferrite bead as a high frequency resistor with no resistance at dc. Ferrites have filtering characteristics which extend well into the megahertz range. They can be used to provide isolated power planes, e.g., a PCB's power plane may be subdivided into an analog power plane (V_{AA}) and clock circuitry power plane, see Figure 18. Ferrite beads work both ways; one prevents noise components on the PCB power plane from being coupled onto the analog power plane, while the other filters out noise components from the clock circuitry. The analog power plane provides power to the Video RAM-DAC and all its associated analog circuit, while the clock circuitry power plane provides power to the crystal oscillators and clock divide circuitry.

It is not, however, recommended to use a ferrite bead in separating ground planes; a separate analog ground plane to the Video RAM-DAC may have disastrous effects on the contents of the on-board color look-up table. One common ground plane should encompass the entire PCB.

The system's shield ground or earth may be connected through a ferrite bead to the regular PCB ground.

Ferrite is the generic term given to a class of non-conductive ceramics. They are constructed using combinations including oxides of iron, cobalt, nickel, zinc, magnesium and some rare earth materials. It is important to be careful in choosing ferrite beads. Ferrite composition and therefore filtering characteristics can vary quite significantly from manufacturer to manufacturer. Analog Devices recommends the use of the BL01/02/03 series of ferrite bead inductors from Murata as well as the Fair-Rite 2743001111. The frequency characteristic of a radial single bead ferrite bead is shown in Figure 1-11.

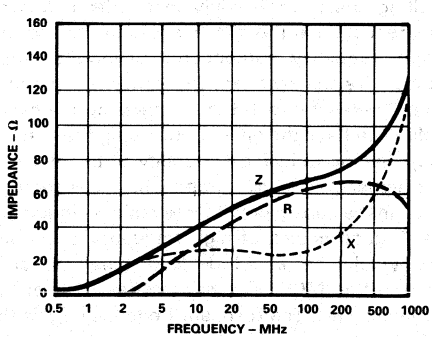

Figure 1-11. Frequency Characteristic of the BL01RN Ferrite Bead Inductor

DC Power Filter

In some cases it might be necessary to filter the power source's high frequency components. This is often the case when switched mode power supplies are used. If it is found that the system's power supply is excessively noisy, one could consider the use of a dc power filter

such as a BNX002-01 from Murata or equivalent. This filter which consists of a large value monolithic 4-terminal capacitor, a feed-through capacitor and beads, as shown in Figure 1-12, produces an effective filtering effect, 40 dB min over the frequency range 1 MHz to 1 GHz. Figure 1-13 shows the filtering characteristic of the BNX series of power filters from Murata.

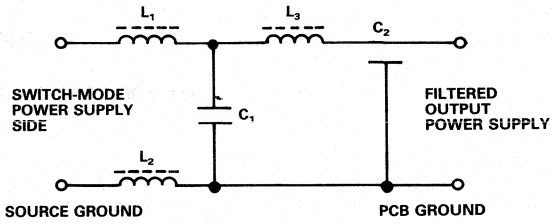

Figure 1-12. Equivalent Circuit of a dc Power Filter

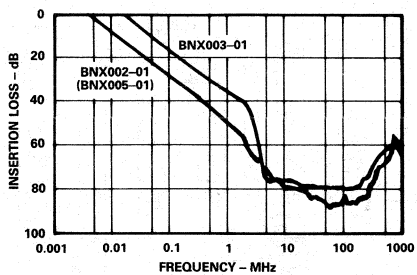

Figure 1-13. Frequency Response of the BNX Series of dc Power Filters

8

APPENDIX 2

FCC TEST RESULTS OF AD/VGA

Frequency	EUT Orientation	Ambient Radiation	Antenna		EUT Radiation			Field Strength
			Polarization	Height	Factor	Maximum	Corrected	
MHz	Degrees	dB μV	H/V	Meters	dB/M	dB μV	dB μV/M	μV/m
37.2	60	6.7	V	1	12.9	18.8	31.7	38.5
48.1	20	3.6	V	1	12.9	19.0	31.9	39.5
60.6	90	7.8	V	1	11.4	20.3	31.7	38.5
62.8	90	9.8	V	1	10.6	17.3	27.9	24.8
69.5	50	8.6	V	1	8.9	19.9	28.8	27.5
70.9	30	10.8	V	1.5	8.9	24.0	32.9	44.2
72.7	90	10.9	V	1	8.6	20.1	28.7	27.2
75.8	30	14.2	V	1	8.4	22.4	30.8	34.7
78.9	30	13.1	V	1	8.4	24.5	32.9	44.2
80.7	0	10.9	V	1	8.5	19.1	27.6	24.0
88.0	0	12.0	V	1.5	10.0	22.1	32.1	40.3
119.7	330	4.4	H	1.5	15.8	16.9	32.7	43.2

Sample of FCC Test Results for AD/VGA Board

APPENDIX 3

AD/VGA COMPONENT PLACEMENT (SILKSCREEN)

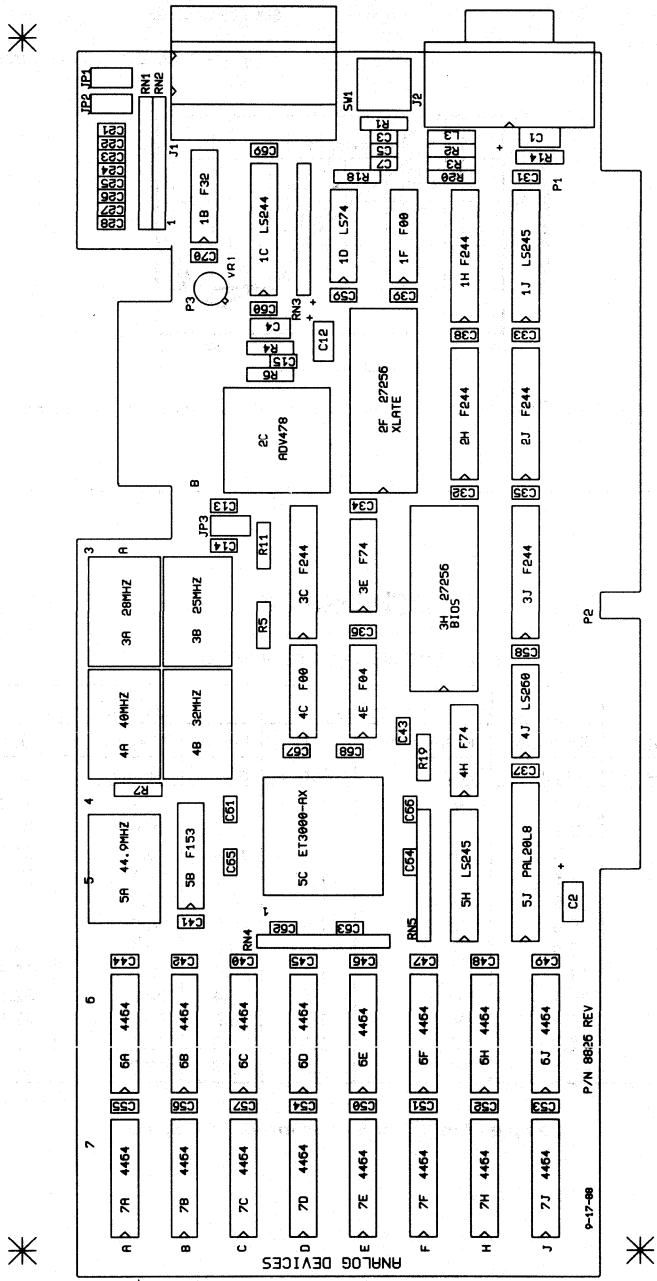

ANALOG DEVICES

AN-237
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Choosing DACs for Direct Digital Synthesis

by David Buchanan

INTRODUCTION

Direct Digital Synthesis (DDS) is a technique for deriving, under digital control, an analog frequency source from a single reference clock frequency. This technique provides high frequency accuracy; temperature and time stability; wideband tuning; and very fast, phase continuous frequency tuning. The performance of synthesizers using this technique is often limited by the performance of available digital-to-analog converters (DACs). This application note explains the basic architecture of DDS and some of its advantages in system design, and outlines some of the performance characteristics designers should look for when choosing DACs for DDS applications. Performance tradeoffs are also explored, and recommendations are made for characterizing the DACs.

DDS BACKGROUND

A simplified block diagram of a direct digital synthesizer is shown in Figure 1. The synthesizer has two digital inputs: a frequency control word (Δ Phase) and a reference clock signal (f_c). The output of the synthesizer is an analog sine wave with frequency f_a. The relationship between f_c and f_a is determined as

$$f_a = \frac{\Delta\ PHASE}{2^N} f_c$$

Where N is the resolution of the frequency control word, Phase.

As illustrated in Figure 1, the circuit is easily divided into three blocks: a phase accumulator, a phase-to-sine converter, and a DAC. The first two blocks are both digital circuits. The phase accumulator is simply an adder with a programmable step size, Δ Phase, representing the phase step taken by the output waveform during each clock cycle. On each clock cycle, the phase accumulator output represents the phase of the output sine wave, with all zeros representing 0 radians, and all ones representing 2π. This signal is a digital ramp with a frequency equal to the output sine wave.

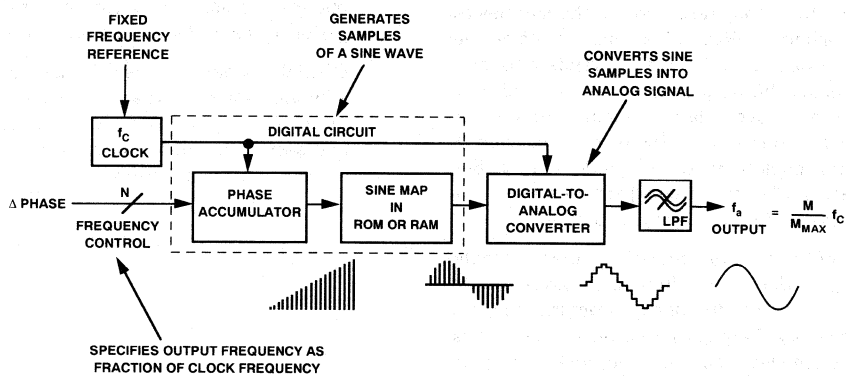

Figure 1. Block Diagram of DDS Generator

The phase-to-sine converter circuit takes the M most significant bits of the phase accumulator's output and provides an M-2 bit sine amplitude output; M-2 determines the resolution of the circuit to follow, usually a DAC. Truncation of the N-M least significant bits is necessary to reduce the complexity of the phase-to-sine conversion. This function may be performed by a look-up table stored in memory, or the sine value may be calculated from a digital algorithm to make a faster or smaller circuit.

The phase accumulator and phase-to-sine converter together form a DDS system with a digital output. The digital output is useful in many applications as a frequency reference (digital demodulation as an example), but most applications require a transformation of the digital sine wave into an analog frequency reference. This makes the digital to analog converter extremely important.

DDS PERFORMANCE CHARACTERISTICS

DDS has both advantages and disadvantages over other frequency synthesis techniques such as phased locked loops. While it is not the intent of this paper to explore fully the tradeoffs of choosing DDS architecture for a synthesizer, it will point out some of the more obvious ones and discuss the important performance characteristics.

A DDS system effectively provides a frequency reference that is a fraction of the clock input frequency. The DDS is digitally tuned by the Δ Phase input, usually controlled by a microcontroller or digital signal processor. Once the digital data are registered on board the phase accumulator, the controlling circuits are free to perform other functions in the system. The digital nature of the DDS eliminates the inconvenience associated with the "tweaking" of synthesizer designs that rely on analog component values to determine frequencies.

The frequency resolution is determined by N, the resolution of Δ Phase, as

$$\frac{1}{2^N} f_c$$

As an example, if the DDS clock reference, f_c, is 20 MHz, and N = 32, the frequency resolution of the synthesizer will be 4.66 millihertz (mHz). This is an advantage over phased locked loops, in which the reference frequency directly determines the frequency resolution and must be large enough to avoid large multiplication ratios. Most integrated solutions for DDS provide at least 24 bits of frequency resolution; some provide up to 48 bits. Many phase accumulators are designed so they can be cascaded to increase the frequency resolution.

DDS provides not only very accurate frequency resolution, but also provides a wide frequency range. As described above, the output frequency is determined as a fraction of the clock frequency reference. The lowest frequency is the smallest fraction of the clock that can be used, or the resolution of the phase accumulator (see above). Nyquist's theory holds that the DAC can reproduce signals up to one half of this clock frequency. Thus the upper limit on the frequency range of the DDS system is determined by the maximum clock rate the synthesizer will support $f_{MAX} = (f_c/2)$. Technology advances over time have increased the available clock rates of both the digital and DAC portions of the DDS circuit.

DDS TECHNOLOGY TODAY

As in other digital circuits, the phase accumulator and phase-to-sine conversion circuit designs must be optimized for cost and power. The DAC design must also address power and cost concerns, but dynamic performance of the converter is of premium consideration.

There are CMOS digital devices available that provide DDS solutions for clock rates up to 100 MHz. A few bipolar devices cover clock frequencies up to 300 MHz, and there are also GaAs devices that provide digital solutions up to 1.4 GHz clock rates.

Most DACs used in DDS are bipolar, although a few GaAs DAC designs see service in the higher frequency applications. Designers prefer monolithic DACs to keep the cost of the converter reasonable. Some 12-bit monolithic DACs can clock up to 100 MSPS (AD9713B), while higher speed applications can take advantage of monolithic 10-bit designs that clock up to 400 MSPS (AD9720). Above 400 MSPS, there are a few 8-bit devices. As a later section of this paper will point out, resolution and speed of a DAC do not always determine the DAC's suitability for DDS applications.

ADVANTAGES AND DISADVANTAGES

While the available clock frequencies described above indicate that DDS circuits can generate output frequencies well into the UHF band, in practice the frequency range of a DDS system is limited by the real world characteristics of the DACs. This is because the resolution and performance (in terms of power, cost, ease of use) that can be practically realized in the DACs at higher clock frequencies are limited. Designers may find that a particular system design could benefit from a DDS which would generate a frequency reference directly, but the required clock rate dictates power, cost, and performance tradeoffs that are unacceptable.

Often the benefits of the DDS can be preserved by using a hybrid design approach. An example of this would be to use DDS at an intermediate frequency with a mixer and second frequency reference, or a PLL to translate the signal to the RF band.

Furthermore, as the output frequency increases, the number of amplitude samples for each sine cycle decreases, making it more difficult for the DAC to reproduce accurately the output sine wave. The accuracy of the analog sine wave is often described in terms of its spectral purity (see below). Each system application of DDS will have a limit on spectral purity. In many applications, the DDS system will meet the spectral purity specifications over only a portion of the available $f_c/2$ bandwidth, effectively limiting the frequency range of the synthesizer. As a rule of thumb, the output frequency should be limited to one fourth of the clock frequency to preserve a reasonable level of spectral purity.

Figure 2 illustrates another feature of DDS, its frequency tuning speed. If the phase input is changed, the phase accumulator instantaneously changes to the new frequency at the next clock cycle. Pipeline delays used in the digital circuits are usually the only limit to switching speed. Frequency transitions in DDS are also phase continuous. Figure 2 illustrates how this instantaneous, phase continuous frequency switching might compare to that of a phase locked loop. The PLL transition includes a frequency transition period ($t_{SETTLING}$, usually a few microseconds) and a frequency overshoot. Coherent analog frequency synthesizers (fixed oscillators, mixers, and filters) also have fast switching speed, but are not phase continuous and cannot be implemented easily.

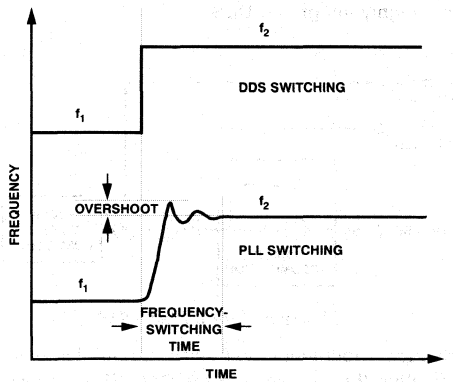

Figure 2. Comparison of DDS and PLL Frequency Switching Characteristics

DDS also has inherently low phase noise and drift. These characteristics are essentially those of the reference clock, f_c. In most DDS applications, a fixed crystal oscillator provides the reference frequency and therefore the phase noise and drift characteristics are excellent.

As mentioned earlier, the spectral purity of a DDS is often a limiting factor in its performance, and is determined by the dynamic performance of the DAC. Ideally, a synthesizer's output spectrum consists of a single frequency.

Since a DDS's output is based on a digital approximation of the sine wave, the theoretical output spectrum contains only the frequency of interest and a constant quantization noise of q/(12)1/2 (where q is the weight of one LSB) spread evenly from dc to $f_c/2$. This ideal spectrum is illustrated in Figure 3 (which ignores phase noise). As with all sampled systems, the amplitude response of the output frequency is weighted as

$$A = \frac{sin\ (\pi\ f_a/f_c)}{(\pi\ f_a/f_c)}$$

where A is the normalized output amplitude.

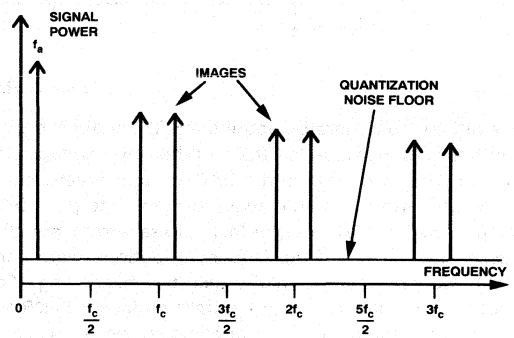

Figure 3. Ideal DDS Output Spectrum

This effect can be corrected by a digital, inverse (sine x)/x filter. Images of the Nyquist bandwidth (dc to $f_c/2$) will appear around each multiple of f_c, again weighted by the (sine x)/x function. These images are commonly removed with a low pass filter at the output of the DAC.

A more realistic illustration of a DDS output spectrum is shown in Figure 4. Here it becomes evident that transfer functions are not ideal. The additional signal content is generated by the digital-to-analog conversion process. The noise contribution is no longer uniform, and harmonics of the fundamental frequency and its images are created, along with other frequencies that have no obvious harmonic relationship.

In many applications, a synthesizer must meet a specified spurious free dynamic range (SFDR) over its output bandwidth. This specification defines the difference in power of the signal of interest and the worst case (highest) signal power of any other signal in the band of interest. This concept is illustrated in Figure 4 for the Nyquist bandwidth. Some other definitions of SFDR do not include direct harmonics, despite the fact the frequency band of interest may include them.

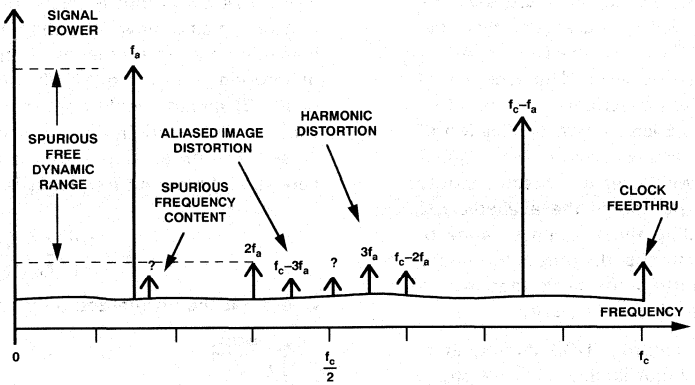

Figure 4. Real DDS Output Spectrum

It is difficult to generalize about the SFDR of direct digital synthesizers because of the number of variables involved. The resolution of the DAC, its ac characteristics, clock rate, and the circuit layout all come into play. DAC characteristics that contribute to these errors are discussed later. Even if the DAC contributions were constant, generalizations might still be misleading. For instance, the illustration in Figure 4 defines a SFDR for the Nyquist bandwidth; an application whose band of interest was limited from $f_c/8$ to $f_c/4$ would see an improvement in this specification. Applications also vary widely in their requirements for spectral purity. System designs which take advantage of the frequency accuracy and stability for channel selection often require a very high SFDR, and are thus restricted to smaller bandwidths. Systems that take advantage of the digital interface to implement IF modulators may be more forgiving in their spectral requirements, and therefore can use a wider bandwidth.

As the discussion of DACs below indicates, the only safe generalization about spectral purity is "each DDS application must be proven out and characterized in the lab." It is unfair, however, to prepare this application note without giving the reader some idea of the performance that can be achieved with available technology.

Ten- to twelve-bit DACs have been shown to provide up to 70 dB of SFDR in applications with clock rates less than 80 MHz and analog bandwidths less than a few MHz. Designs using higher clock rates have been limited to 8 bit DACs, and a SFDR of up to 45 dB is often quoted as a limit. Caution is advised in using these figures because they may not be applicable over wide analog bandwidths.

The modulation capabilities of a synthesizer are also important, and DDS is versatile in this respect. Frequency modulation is possible directly through the Δ phase data port. In fact, circuits ahead of the DDS could use the Δ phase input to specify both channel selection and FM modulation simultaneously.

Figure 5 shows a DDS block diagram modified to include digital control of the phase modulation (PM) and amplitude modulation (AM) of the waveform. Thus, a user could implement all three forms of modulation digitally with a highly integrated DDS.

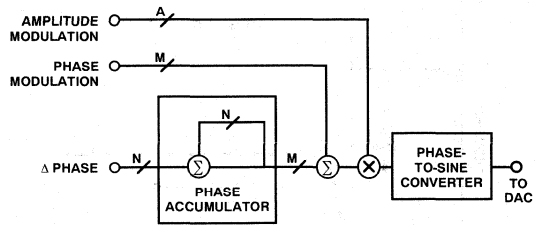

Figure 5. Modulated DDS

Amplitude modulation can also be implemented by modulating the reference to the DAC. Figure 6 shows a circuit using a second DAC to modulate the DDS DAC's reference under digital control. In such an application, the amplitude modulation data would have to be skewed in time to match any pipeline delays in the digital part of the DDS system. Many DAC designs integrate a reference and reference amplifier, but only a few provide enough bandwidth to support multiplying applications.

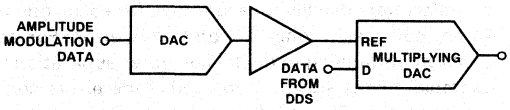

Figure 6. Amplitude Modulation Using a Multiplying DAC

To this point, DDS has been described as a wideband, frequency agile synthesizer with digital tuning and modulation capabilities, phase continuous frequency transitions, high dynamic range, and good phase noise and stability. This set of characteristics is unique, and not easily realized with any other synthesizer architecture. In terms of price and power dissipation, components that accommodate real world applications have been available for only a few years. The availability of complete digital solutions and higher performance DACs, optimized for price and performance, is sure to drive this frequency synthesis option into many applications.

There are, of course, tradeoffs and limitations that designers must make to take advantage of this versatile synthesizer architecture. Applications requiring a high SFDR will find their options limited to the higher resolution DACs with lower clock rates and analog bandwidths. Applications that can take advantage of the higher speed DDS solutions may find the cost and power prohibitive. Designers must evaluate their system's architecture to determine if DDS offers an advantage over traditional frequency synthesis techniques, or if a modified architecture could take advantage of DDS and improve system performance. Applications which can usually take advantage of DDS include (but are not limited to) advanced military radar, high performance instrumentation, digital communications links, and commercial cellular communications.

CHOOSING DACS FOR DDS

Since DDS provides a frequency reference, it makes sense that its key specifications are in the frequency domain. While the phase noise and stability requirements may be directly related to specifications of the frequency reference, f_c, it is difficult to relate frequency range and SFDR specifications directly to high speed DACs that are suitable for DDS applications. High speed DACs are traditionally specified in the time domain. This is reasonable, since the traditional applications (video, analog-to-digital converter building blocks, fast tuning voltage references) are concerned with the time domain characteristics of the DAC. These specifications are constant without regard to the end application.

Earlier sections of this paper described the high performance characteristics of the DDS architecture and how its use is likely to grow in communication, instrumentation, and military applications. It would seem natural, with such an opportunity, that DAC manufacturers would design and specify devices especially for DDS applications and solve the problem of choosing DACs for DDS. To some degree, this is already happening.

Highly integrated DACs (including on-board registers, reference circuits, and reference amplifiers) optimized for waveform synthesis are available today, and manufacturers are providing applications support for their use in direct digital synthesizers. The problem of specifying the DAC is not solved quite so easily.

There are two reasons this is unlikely to change. The first has been explained, and is simply because various system applications of DDS have a wide range of dynamic requirements. The second reason involves the inconsistency of DAC performance over a wide frequency range. It is the combination of these two characteristics that makes it so difficult to specify the devices. If the applications requirements were narrowly focused, DAC manufacturers would simply specify the devices accordingly, and ignore the inconsistency in performance.

Alternatively, if the DAC's performance were to degrade in some predictable fashion (as an amplifier might) the manufacturer could supply data or a set of performance curves that would allow designers to predetermine system performance. As it stands, the manufacturer does well to indicate simply that the part is intended for use in DDS applications, and to supply some sample data to back up this claim.

Since high speed DAC manufacturers are often unable to specify the devices adequately, the task of deciding whether a part is suitable for a particular DDS system application is left to the system designer. It is the goal of the remainder of this paper to explain the relevant dc and time domain specifications of the DAC, and how they can be related to frequency domain performance. It will become clear that these specifications will not provide all the necessary information to predict the frequency domain performance; consequently, some recommendations on how to characterize the DACs for DDS applications are also presented.

DAC DC SPECIFICATIONS

Figure 7 illustrates both an ideal and real world transfer function for a 3-bit DAC. Manufacturers typically specify offset, gain error (full scale in relation to reference),

differential nonlinearity (DNL), and integral nonlinearity (INL) as an approximation of this transfer function. Output offset is usually defined as a constant dc offset in the transfer curve, and therefore has no effect on the frequency domain characteristics of the output. Gain defines the full-scale output of the converter in relation to its reference circuit. The gain error is usually specified tightly enough so there is little concern about its effect on the frequency domain performance.

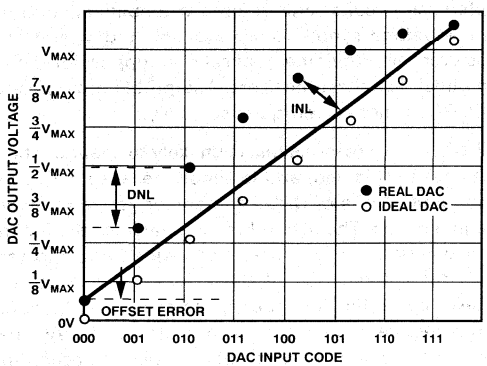

Figure 7. DAC Transfer Function

The linearity specifications, differential nonlinearity (DNL) and integral nonlinearity (INL), provide useful information about frequency domain performance. DNL is typically measured in LSBs as the worst case deviation from an ideal LSB step between adjacent code transitions, and can be a negative (less than 1 LSB step) or positive (greater than 1 LSB step) error. DACs with a DNL specification of < -1 LSB are not guaranteed to be monotonic. Figure 7 illustrates both positive and negative DNL errors, although the part is monotonic.

Armed with the DNL error of each code transition, a diligent individual could predict the frequency domain performance of an individual DAC's transfer function. This process would consist of modeling the digital part of a DDS; using the modeled DDS data as input to the DAC transfer function; collecting samples of the simulated DAC output levels, and performing a Fourier transform to observe frequency domain performance. This process would have to be repeated for every ratio of clock to analog output frequency. While this might be an interesting exercise, it is not a realistic way of selecting a DAC.

Manufacturers usually measure DNL for each code transition, but specify only a worst case error. For some applications (analog to digital converter applications, for example) this is sufficient to predict the errors contributed by the DAC's transfer function. For frequency synthesis, however, it raises some interesting questions. For example, all other things equal, will a DAC with a 1/2 LSB DNL error outperform a DAC with a 1 LSB error in the frequency domain?

The answer is not obvious. Theory predicts that a perfect DAC generating a full-scale sine wave will have an rms signal to rms noise ratio of 6.02N – 1.76 dB[1] over the full Nyquist bandwidth [ignoring (sine x)/x rolloff[2] where N is the resolution of DAC]. A perfect 12-bit DAC could therefore provide 70.48 dB of SNR. A measurement of the DNL error would predict a maximum reduction in this performance; for example a 12-bit DAC with a 1 LSB DNL error (every other quantization level missing) would have a minimum SNR of 64.46 [6.02 × (11) – 1.76] dB (if DNL were its only nonlinearity).

This reduction in SNR, however, does not predict how the distortion is spread over the Nyquist bandwidth. If the DNL errors caused the additional noise to be spread evenly over the Nyquist band, their effect would be negligible. However, if the DNL errors concentrate portions of the noise into a single frequency (usually a harmonic) this could limit the SFDR.

Getting back to the example, if the 1/2 LSB device has its maximum DNL error on a high percentage of its code transitions, and the 1 LSB device were perfect except for a single code transition, then the 1 LSB part would have better frequency domain characteristics. It is obvious from this example that DNL specifications can be misleading if not interpreted properly.

INL is measured as the worst deviation from a straight line approximation to the DAC's transfer function. The straight line approximation eliminates the dc errors (gain and offset) which have already been discounted from any frequency domain effect. In Figure 7, this straight line is drawn between the two end points. The INL characteristics of some devices are measured against a "best fit" straight line through the converter's code transitions. Like the DNL specification, the INL measurement is a worst case deviation. It does not indicate how many DAC codes reach this deviation, nor which direction away from the best straight line the deviation occurred.

Figure 8 illustrates how this specification might be misinterpreted. Each of the illustrated curves represents a transfer function that would have the same INL measurement, but three distinct effects in the frequency domain. For example, the transfer function of the "bow" INL curve will introduce a prominent 2nd harmonic distortion, while the symmetrical "S-curve" will tend to introduce 3rd harmonic distortion.

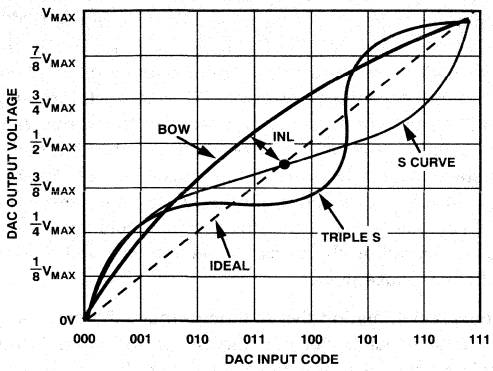

Figure 8. INL Curves—All 1 LSB

Most IC DAC designs recommended for use in frequency synthesis exhibit some predictable linearity pattern, with DNL and INL errors occurring due to architectural tradeoffs and matching of process parameters. This pattern determines how the DNL and INL will contribute to the overall frequency content of the signal. Because the linearity patterns of any two specific DAC designs will be different, it is impossible to compare their frequency domain performance by comparing their linearity specifications.

Since deviations in DNL and INL will probably not affect the ac characteristics (glitch, slew rate, etc.) of the DAC, it is possible to characterize the effects of linearity on frequency domain performance for a specific DAC. This characterization would consist of selecting units that exhibit good DNL and INL characteristics, and comparing their frequency domain performance with parts that exhibit poor linearity. This allows designers to determine if tighter linearity grades of a specific DAC will improve a synthesizer's performance.

AC TIME DOMAIN DAC SPECIFICATIONS

The ac time domain specifications of a DAC relate to its code-to-code transitions as illustrated in Figure 9. It is difficult to predict the frequency domain effects of any of these specifications, and practically impossible to predict the frequency domain effects of their combinations. It is, however, possible to consider what contributes to each one of these nonlinearities, and how its effect can be minimized in a final application.

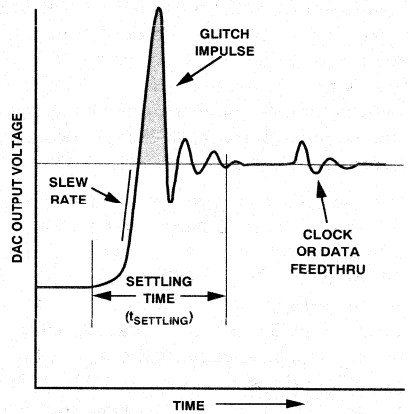

Figure 9. DAC Transition

The effect of DAC output slew rate on its frequency domain performance is not as straightforward as it is in linear applications due to the presence of glitch impulse and settling effects (discussed below). In general, a DAC with a high slew rate will produce an output transition that is closer to ideal than one which is slower. Designers should be careful to consider also the difference in rising and falling slew rate, as this will tend to concentrate energy in the harmonics of the fundamental output frequency.

Glitch impulse, often considered a key figure of merit in DDS applications, is simply a measure of the initial transient response (overshoot) of the DAC between two output levels. It is usually measured as the area of the transient, and ranges from 15 to 100 pV-s for devices commonly used in DDS applications. This transient is commonly associated with the time skew between the data bit transitions or by unequal propagation delays through the internal logic. In either case, the time skew would cause the DAC's output to approach an intermediate state, and probably add unwanted energy to the output frequency spectrum.

DAC architecture has a strong effect on the magnitude of the glitch impulse. Architectures employing an on-board register for data deskew and propagation delay matching, along with a segmentation of the major bits, have the lowest glitch impulse. The glitch illustrated in Figure 9 is a peak glitch interpretation. Designers are warned that not all manufacturers define glitch impulse in the same fashion, and comparison of data specifications is often misleading.

Before the advent of monolithic devices which minimize glitch impulse and other DAC nonlinearities, it was common practice to "deglitch" the outputs of hybrid converters using a track-and-hold amplifier (T/H). The idea is to place the T/H in "hold" mode during the DAC's transition to a new value. After the converter's output is settled to its final value, the T/H is allowed to return to "track" mode; and the nonlinearities of the T/H's transition to the newly acquired DAC output become the limit to ac performance.

T/H technology has been predominantly hybrid technology, making any possible improvement in performance a costly venture. Although the transition characteristics of the T/H may imply that deglitching could improve the DAC, overall performance may actually degrade because of limits in other performance characteristics such as slew rate and absolute accuracy. More recent monolithic T/H product offerings such as the AD9101 125 MSPS sampling amplifier may offer marginal improvement in overall ac performance at a reasonable cost.

After the initial transient, the DAC's output will settle to its final value. The settling time is usually measured as the time from the digital inputs' transitions to the time the DAC's output settles to within a certain error band (usually 1/2 or 1 LSB) of the final output value. Many manufacturers understandably argue that the digital propagation of the converter should not be included in this specification. Settling time should instead be measured as the interval from the time when the DAC's output leaves the error band around its initial value, and when it settles within the error band around its final value. This definition more accurately describes the non-ideal characteristics of the code transition.

Slew rate, glitch impulse, and settling time characteristics of a DAC's output are all dependent on the output loading circuit. Stray capacitive loading generally increases all of them. Most high speed DACs are current output devices and require an external load resistor to produce an output voltage reference. Design of the final DDS circuit should pay close attention to the DAC load, with signal traces kept short as possible, and/or with impedance matching techniques being used.

Feedthrough of the clock or data transitions to the DAC's output also adds to the frequency content of the DDS's output spectrum. These effects are often related to the test circuit layout and can be minimized with good circuit layout techniques. Many DAC manufacturers suggest the use of a series resistor in the input data connections to minimize data feedthrough. This series resistance works with the input capacitance of the DAC to form a low pass filter, and may alter the setup and hold characteristics of the device. This technique is not recommended on the clock connection as it will tend to add jitter (phase noise).

Theoretically relating the ac time domain characteristics of a DAC to its frequency domain performance is not feasible, nor recommended. One simple observation about the ac nonlinearities is possible: they become a more significant portion of each clock cycle as the clock or analog output frequency of a DDS system increases. Consequently, one can expect the spurious frequency content of the DAC to degrade at higher frequencies. Many DDS designs will operate well below the maximum clock frequency of the DAC to get optimum performance, and translate the output to a higher frequency using analog techniques (mixers).

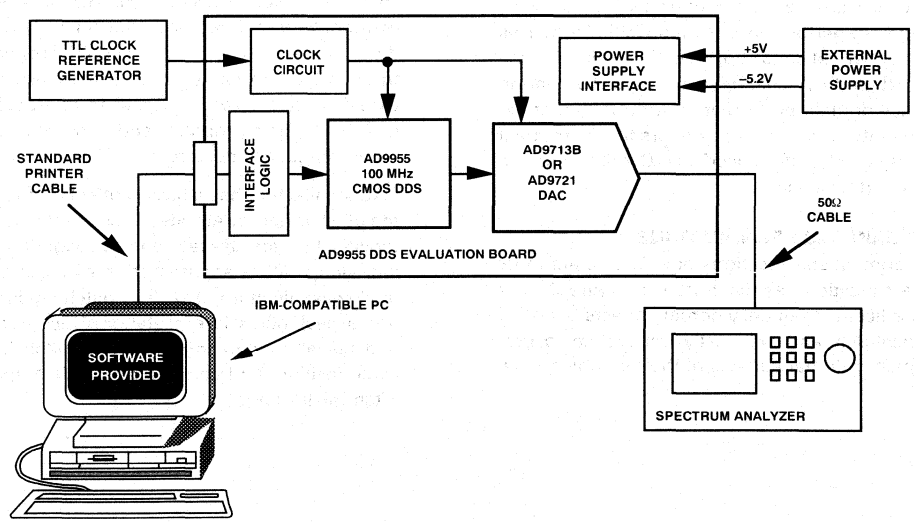

Figure 10. AD9955/PCB DDS Evaluation Board Setup

This degradation at higher clock rates, however, may not be the dominant effect in output spectral purity. Instead, for any given bandwidth of interest, the frequency domain performance may vary as the ratio of the clock to analog frequency changes. Spurious distortion will tend to appear in the output spectrum at alias frequencies, $Af_a \pm Bf_c$, where A and B are integers. The effects of these aliased components are concentrated near the fundamental frequency when the fundamental output frequency is nearly an integer fraction of the clock frequency. Designs that cover wide analog output bandwidths may find it hard to avoid this effect even at low clock rates, while narrow bandwidth applications may be able to find a particularly "clean" segment of output spectrum that allows taking advantage of the higher clock rates. This would require careful characterization of the DAC for the bandwidth of interest.

While the ac specifications of the DAC will not allow accurate predictions of the output spectral purity, these specifications are recommended as guidelines for determining if the DAC is worthy of being characterized for use in a DDS application. Devices that exhibit high slew rate, low glitch and feedthrough, and fast settling times should be considered first. Once the decision has been made to characterize, these time domain specifications are also useful as a guideline to determine if the DAC has been optimized in the test fixture. A block diagram of how this test is performed is illustrated in Figure 10.

The power supplies should be the linear (nonswitching) type if possible, and appropriate filtering needs to be included in the circuit to reduce noise to a minimum. For the purposes of evaluation, separate digital and analog supplies should be used as recommended by the DAC manufacturer. Once the characterization is complete, the designer can explore how combining digital and analog supplies would affect the spectral purity of the DDS.

The external controlling circuits should allow complete access to all digital circuitry. This may not seem appropriate if all of the features are not required in the final application, but may simplify the characterization process and allow the test set to be used in future design phases. The clock distribution is usually generated by an instrument during characterization, allowing a wide range of reference frequencies. The clock distribution circuit serves to buffer the incoming reference and sets up the appropriate timing between the digital part of the DDS and the DAC. The important specifications here are the DAC's input data setup and hold times. Even a slight violation of these specifications can cause an increase in glitch impulse if the latch is transparent, or capture of erroneous data if the latch is edge triggered.

Once the test set is constructed, functionality must be verified. This is most easily accomplished with a relatively slow clock rate. Checking functionality includes investigating the stability and level of the power supplies (at many points in the circuit), operation of the digital section of the DDS, stability of the DAC reference circuit, and the timing relationships at the DAC interface. The output of the DAC should be analyzed in the time domain to verify that the DAC is meeting the ac characteristics specified in the data sheet. If they are significantly different, the differences must be resolved.

After the circuit is functioning properly, the DAC can be characterized for its frequency domain performance. There are two approaches that can be taken: to characterize the part in general, or to characterize for a particular application.

To characterize the part for general use, a matrix of conditions is suggested. A maximum usable clock rate can be identified from the data sheet. The converter should be tested at several clock frequencies over the entire range. For example, a DAC billed as a 50 MSPS converter might be tested for clock rates of 2, 5, 10, 20, 35, and 50 MSPS. At each of these clock rates, the SFDR of the output should be characterized using a spectrum analyzer for several clock to analog output frequency ratios.

Determination of these ratios is a subject of debate, but the following selections would give a decent understanding of performance:

1. A low ratio of $<1/10\ f_c$, to determine the performance for the given clock frequency, and
2. Ratios of $1/3\ f_c \pm \Delta f$ and $1/4\ f_c \pm \Delta f$ to analyze the level of spurious frequency energy at alias frequencies.

For each of the three analog selections, the characterization of the spectrum will consist of two parts. The first considers performance across the entire Nyquist band, or a reasonable section of it. The second considers a narrow band around the fundamental to look for unfilterable spurious distortion. This process can become quite time consuming because of the time required for broadband, high dynamic range sweeps of the analyzer. A sample of data characterizing a prototype of the AD9721, a 10-bit, 100 MSPS TTL converter is shown in Figure 11.

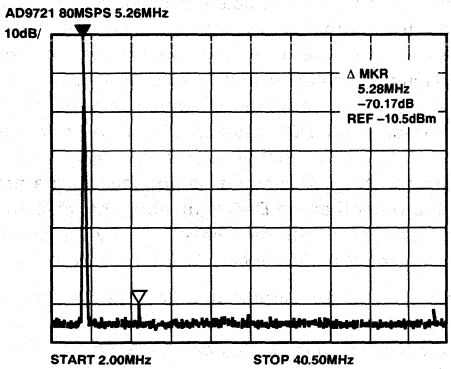

a.

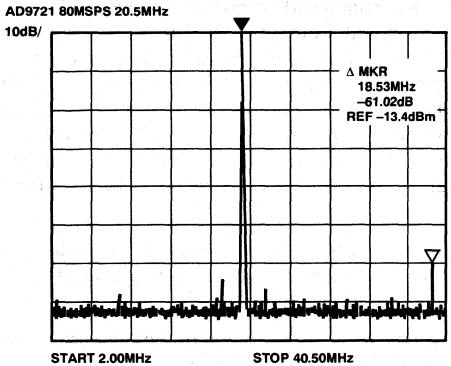

b.

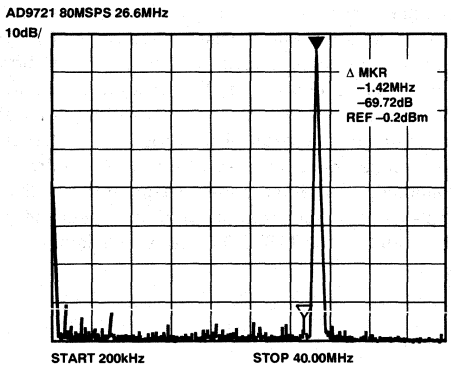

c.

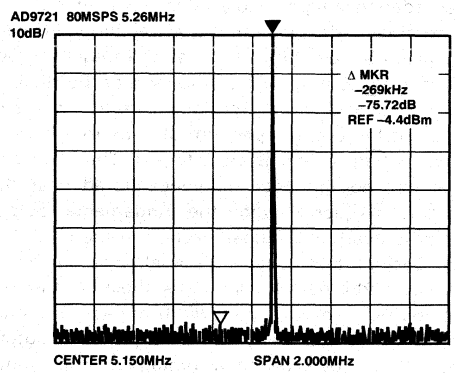

d.

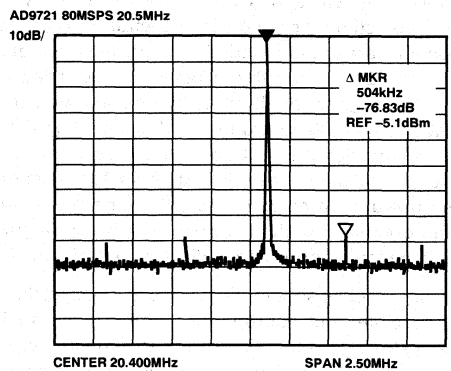

e.

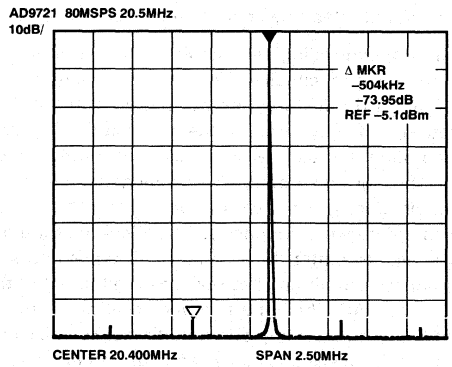

f.

Figure 11. DAC Characterization Plots

Once this general characterization is complete, some limited conclusions about the DAC can be made. For instance, in the worst case, the part may exhibit poor performance across the test matrix; and the device could be disqualified from any further consideration. Encouraging results might indicate that the device provides enough SFDR for the application of interest.

If so, the application specific characterization process begins. This would consist of developing a frequency plan to use the DAC to generate an IF band to be used in the system. Characterization of the DAC would focus on determining the clock rate which could give the best performance in this band.

CONCLUSION

This application note has explored the architecture and advantages of DDS, and pointed out that the performance-limiting block is the DAC. It should be clear from the discussion that frequency domain performance cannot be accurately predicted from device specifications such as linearity, glitch impulse, slew rate, and settling time. Instead, testing is required to characterize the spurious free dynamic range of a converter for a range of clock and analog frequencies. A suggested method of characterizing the DAC was also discussed.

REFERENCES

1. Seminar Notes: "Fundamentals of Frequency Synthesizer Design," Instructor: Fred Studenberg. Nov. 6–8, 1990, University of Maryland Center for Professional Development.
2. W. R. Bennett, "Spectra of Quantized Signals." BSTJ 27, pp 446–472, July, 1948.

Digital Signal Processing Products
Contents

		Page
AN-334	Digital Signal Processing Techniques	9–3
AN-344	Digital FIR Filters Without Tears	9–19
AN-218	DSP Multirate Filters	9–27
AN-217	Audio Applications of the ADSP Family (IIR Filters)	9–37
AN-336A	Image Compression: Spelling Out the Options	9–41
AN-336B	JPEG Compression	9–45
AN-227	Digital Control System Design with the ADSP-2100 Family	9–49
AN-338	Loading an ADSP-2101 Program via the Serial Port	9–61
AN-240	Considerations for Selecting a DSP Processor (ADSP-2101 vs. WE DSP16A)	9–65
AN-233	Considerations for Selecting a DSP Processor (ADSP-2101 vs. TMS320C50)	9–77
AN-235	Considerations for Selecting a DSP Processor (ADSP-21020/ADSP-21010 vs. TMS320C30)	9–93

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Digital Signal Processing Techniques

Digital Filtering

Real-time digital filtering is one of the most powerful tools of DSP. Apart from the obvious advantages of virtually eliminating errors in the filter associated with passive component fluctuations over time and temperature, op amp drift (active filters), etc., digital filters are capable of performance specifications that would, at best, be extremely difficult, if not impossible, to achieve with an analog implementation. In addition, the characteristics of a digital filter can be easily changed under software control. Therefore, they are widely used in adaptive-filtering applications such as modems, digital audio, digital mobile radio, and speech processing.

The actual procedure for designing digital filters has the same fundamental elements as that for analog filters. First, the desired filter responses are characterized and the filter parameters are then calculated. Characteristics such as transfer function and phase response are used in the same way. The key difference between analog and digital filters is that instead of calculating resistor, capacitor, and inductor values for an analog filter, coefficient values are calculated for a digital filter. So for the digital filter, numbers replace the physical resistor and capacitor components of the analog filter. These numbers reside in a memory as filter coefficients and are used along with data values from the ADC in performing the filtering calculations.

The digital filter, because it is a discrete function, works with digitized data as opposed to a continuous waveform, and a data point is acquired each sampling period. Because of this discrete nature, we can reference data samples by numbers such as sample 1, sample 2, sample 3, etc. Figure 7.1, illustrating the basic filtering function, shows a low frequency signal containing higher frequency noise which must be filtered out. This waveform must be digitized with an ADC to produce samples x(n). These data values are fed to the digital filter, which in this case is a lowpass filter. The output data samples, y(n), are used to reconstruct an analog waveform using a DAC.

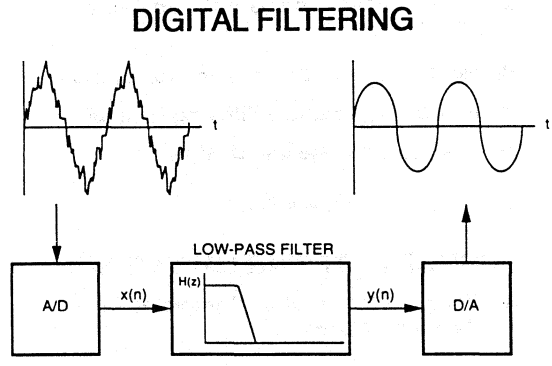

Figure 7.1

Digital filters, however, are not the answer to all signal processing filtering requirements. In order to maintain real-time operation, the DSP processor must be able to execute all the steps in the filter routine within one sampling clock period, $1/f_s$. This currently limits their use to primarily voice and audio bandwidth applications. However, it is possible to sacrifice software control and flexibility, and design special hardware digital filters which will operate at video-speed sampling rates. In other cases, the

speed limitations can be overcome by first storing the high speed ADC data in a buffer memory. The buffer memory is then read at a rate which is compatible with the speed of the DSP-based digital filter. In this manner, pseudo real-time operation can be maintained as in a radar system, where signal processing is typically done on bursts of data collected after each transmitted pulse. Even in highly oversampled sampled data systems, a simple analog antialiasing filter is usually required ahead of the ADC and after the DAC. Finally, as signal frequencies increase sufficiently, they surpass the capabilities of available ADCs, and digital filtering then becomes impossible, since we no longer have a sampled data system because we have no ADC. Active analog filtering is not even possible at extremely high frequencies because of op amp bandwidth and distortion limitations, and filtering requirements must then be met using purely passive components. The primary focus of the following discussions will be on filters which can run in realtime under DSP program control.

DIGITAL FILTERING ADVANTAGES

- **High Accuracy**
- **High Performance**
- **Linear Phase, Constant Group Delay (FIR Filters)**
- **No Drift Due to Component Variations**
- **Flexibility, Adaptive Filtering Possible**
- **Easy to Simulate and Design**

Figure 7.2

DIGITAL FILTERING LIMITATIONS

- **Computation Must be Completed in Sampling Period**
- **Limited to Voice and Audio Bandwidth Signals if Real-Time Operation is to be Maintained**
- **Hardwired Digital Filters Required for Video Frequencies**
- **Analog Filters Still Needed: Antialiasing and High Frequencies**
- **Lack of High Speed ADCs for Sampling**

Figure 7.3

FINITE IMPULSE RESPONSE (FIR) DIGITAL FILTERS

The simplest form of a digital filter is the finite impulse response filter (FIR), and the most elementary form of an FIR filter is a *moving average* filter as shown in Figure 7.4,

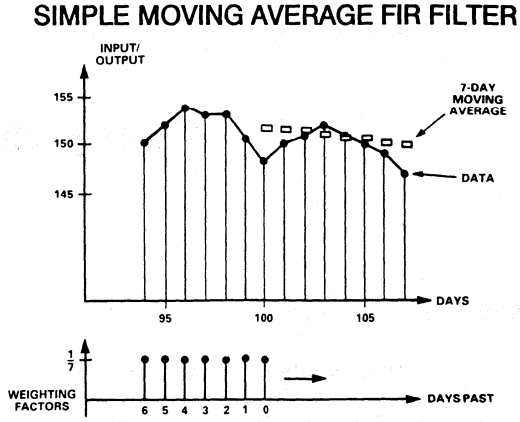

Figure 7.4

where we show a 7-day moving average of a dieter's weight plotted along with the daily weights. After 7 days worth of data samples are obtained, the first point on the moving average is computed by adding the 7 data samples together and dividing by 7. Another way to view the process is to *weight* each data sample by a factor of 1/7 and perform a summation. To obtain the second point on the moving average, the first weighted data sample is subtracted from the summation, and the 8th weighted data sample is added to the summation. This process continues, and can be viewed as a very crude lowpass filtering of the daily readings. The digital implementation of the process is shown in Figure 7.5 which shows the various multiplications, delays, and the summation. The Finite Impulse Response (FIR) filter gets its

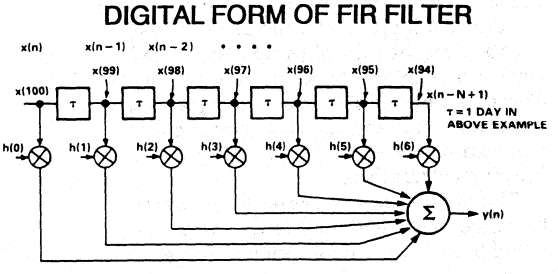

Figure 7.5

name because the impulse response is of finite duration; i.e., after seven zero-valued input samples, the filter output goes to zero. When processing an actual electrical signal, a moving average might look like Figure 7.6.

MOVING AVERAGE FIR FILTER APPLIED TO ANALOG SIGNAL

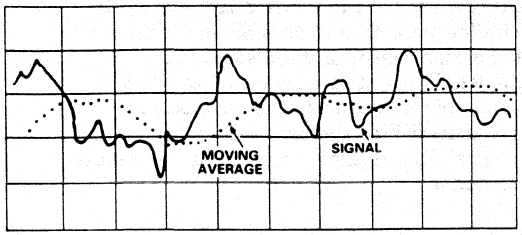

Figure 7.6

It is useful from a mathematical standpoint to view the moving average filter as a *convolution* of the filter impulse response h(t) with the sampled data points x(t) to obtain the output y(t) as shown in Figure 7.7. For a linear convolution, the operation involves multiplying x(t) by a reversed an linearly shifted version of h(t), and then summing the values in the product.

MOVING AVERAGE COEFFICIENTS CONVOLVED WITH SAMPLED WAVEFORM

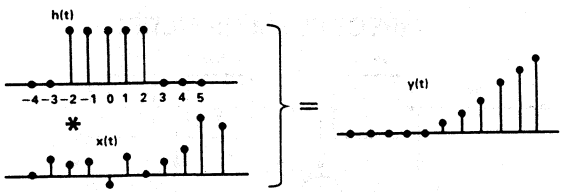

Figure 7.7

The sin(x)/x frequency response of the moving average filter is shown in Figure 7.8 for various numbers of taps, N. (Note: in this section N refers to the number of sample points and not the number of bits of resolution of an ADC or DAC!). Note that increasing the number of taps sharpens the rolloff characteristic of the moving average filter but does nothing to improve the undesirable sidelobes.

It is possible to dramatically improve the performance of the simple FIR moving average filter by properly selecting the individual

FREQUENCY RESPONSE OF MOVING AVERAGE FILTER FOR VARIOUS NUMBER OF TAPS

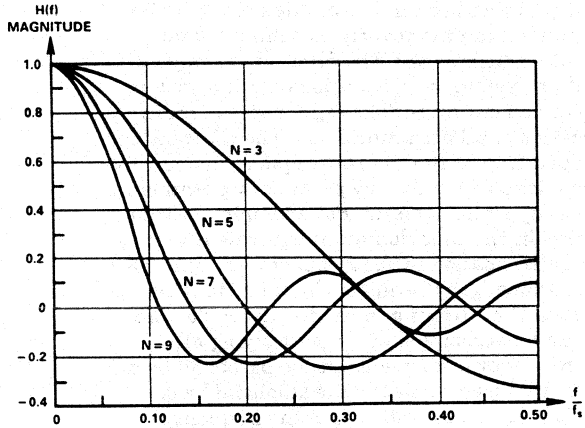

Figure 7.8

weights or coefficients rather than giving them equal weight. The sharpness of the rolloff can be improved by adding more stages (taps), and the stopband attenuation characteristics can be improved by properly selecting the filter coefficients. The essence of FIR filter design is the appropriate selection of the filter coefficients and the number of taps to realize the desired transfer function H(f). Various algorithms are available to translate the frequency response H(f) into a set of FIR coefficients. Most of this software is commercially available and can be run on PCs. *The key theorem of FIR filter design is that the coefficients h(n) of the FIR filter are simply the quantized values of the impulse response of the frequency transfer function H(f).* Conversely, the impulse response is the Fourier Transform of H(f).

FACTORS DETERMINING FIR FILTER TRANSFER FUNCTION H(f)

- **Number of Taps**
- **Proper Selection of Weighted Filter Coefficients**

Figure 7.9

The Duality of the Time and Frequency Domains

It is useful to digress for a moment and examine the relationship between the time domain and the frequency domain to better understand the principles behind digital filters such as the FIR filter. In a sampled data system, a convolution operation can be carried out by performing a series of multiplications and accumulations. The convolution operation in the time or frequency domain is equivalent to point by point multiplication in the opposite domain. For example, convolution in the time domain is equivalent to multiplication in the frequency domain. This is shown graphically in Figure 7.10. It can be seen that filtering in the frequency domain can be accomplished by multiplying all frequency components in the passband by a 1 and all frequencies in the stopband by 0. Conversely, convolution in the frequency domain is equivalent to point by point multiplication in the time domain.

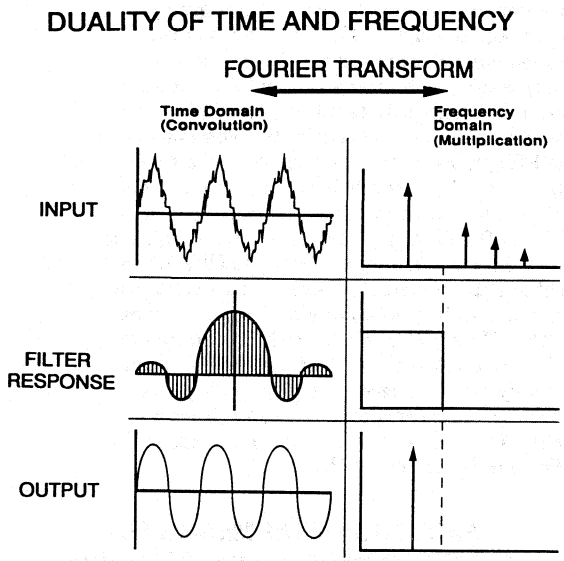

Figure 7.10

The transfer function in the frequency domain (either a 1 or a 0) can be translated to the time domain by the Fourier transform. This transformation produces an impulse response in the time domain. Since the multiplication in the frequency domain (signal spectrum times the transfer function) is equivalent to convolution in the time domain (signal convolved with impulse response), the signal can be filtered by convolving it with the impulse response. The FIR filter is exactly this process. Since it is a sampled data system, the signal and the impulse response are quantized in time and amplitude yielding discrete samples. The discrete samples comprising the impulse response are the FIR filter coefficients.

The mathematics involved in filter design (analog or digital) most always make use of transforms. In continuous-time systems, the Laplace transform can be considered to be a generalization of the Fourier Transform. In a similar manner, it is possible to generalize the Fourier transform for discrete-time sampled data systems, resulting in what is commonly referred to as the z-transform. Details describing the use of the z-transform in digital filter design are given in References 1, 2, and 3.

FIR Filter Implementation in DSP Hardware Using Circular Buffering

As has been discussed, an FIR filter (shown in Figure 7.11 must perform the following convolution equation:

$$y(n) = h(n)*x(n) = \sum_{i=0}^{N-1} h(i)x(n-i) \quad ,$$

where h(i) is the filter coefficient array and x(n-i) is the input data array to the filter. The number N, in the equation, represents the number of taps of the filter and relates to the filter performance as has been discussed above.

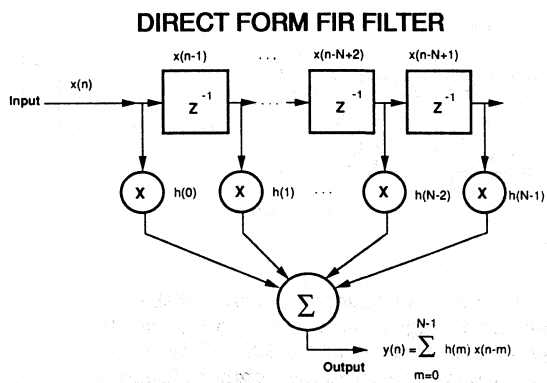

Figure 7.11

In the series of FIR filter equations, the N coefficient locations are always accessed sequentially from h(0) to h(N-1). The associated data points circulate through the memory; new samples are added replacing the oldest each time a filter output is computed. A fixed boundary RAM can be used to achieve this circulating buffer effect as shown in Figure 7.12 for a 4 tap FIR filter.

DATA MEMORY ADDRESSING FOR 4 TAP FIR FILTER

$$y(n) = h(n) * x(n) = \sum_{i=0}^{N-1} h(i)\, x(n-i)$$

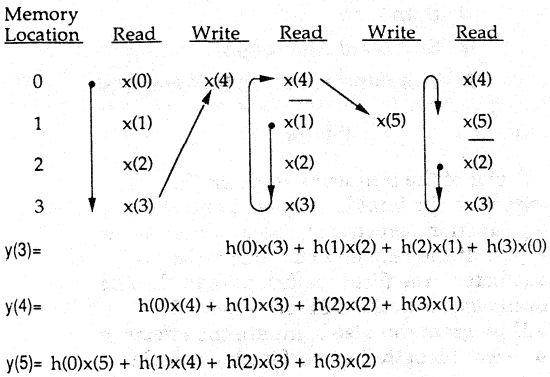

Figure 7.12

The oldest data sample is replaced by the newest after each convolution. A "time history" of the four most recent data samples is kept in RAM.

This delay line can be implemented in fixed boundary RAM in a DSP chip if new data values are written into memory, overwriting the oldest value. To facilitate memory addressing, old data values are read from memory starting with the value one location after the value that was just written. For example, x(4) is written into memory location 0, and data values are then read from locations 1,2,3,and 0. This example can be expanded to accommodate any number of taps. By addressing data memory locations in this manner, the address generator need only supply sequential addresses regardless of whether the operation is a memory read or write. This data memory buffer is called *circular* because when the last location is reached, the memory pointer must be reset to the beginning of the buffer.

The coefficients are fetched simultaneously with the data. Due to the addressing scheme chosen, the oldest data sample is fetched first. Therefore, the last coefficient must be fetched first. The coefficients can be stored backwards in memory: h(N-1) is the first location, and h(0) is the last, with the address generator providing incremental addresses. Alternatively, coefficients can be stored in a normal manner with the accessing of coefficients starting at the end of the buffer, and the address generator being decremented. In the example shown in Figure 7.12, the coefficients are stored in a reverse manner.

FIR Filter Design Techniques

FIR filter design calls for specifying a finite set of N coefficients, h(n), to approximate an idealized filter form. *The filter coefficients, h(n), in the time domain correspond to the impulse response of the filter transfer function H(f).*

KEY FIR FILTER DESIGN THEOREM

- **The Coefficients h(n) of an FIR Filter are Simply the Quantized Values of the Impulse Response of the Frequency Transfer Function H(f)**
- **The Impulse Response is Calculated by Taking the Fourier Transform of H(f)**

Figure 7.13

In Figure 7.14, 2nd, 4th, and 6th-order ideal Chebyshev lowpass filter transfer functions, optimized for 1dB in-band ripple, are compared with a 91-tap (i.e., 91 coefficients and 91 sequential circular buffer memory locations) digital FIR filter optimized for 0.002dB passband ripple. There is no practical analog equivalent; this is higher order than is realistic with analog hardware (greater than 70 poles using rule-or-thumb approximation). Since the response is flatter within the passband, the signal is reproduced more faithfully, and phase distortion in the passband is negligible, since all frequencies are delayed equally by the filter. This is another important characteristic of FIR filters (linear phase response and constant group delay) which makes them extremely attractive to digital audio applications.

91 TAP FIR FILTER RESPONSE COMPARED TO CHEBYSHEV ANALOG FILTER RESPONSE

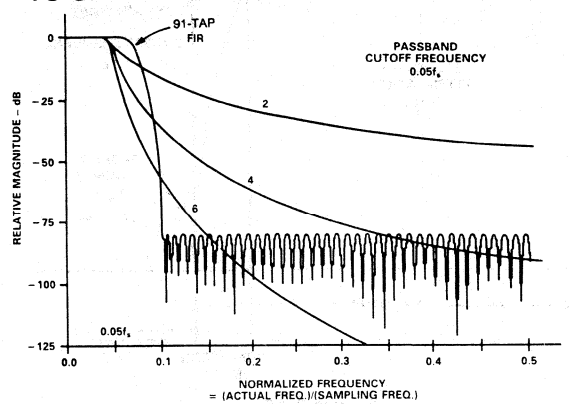

Figure 7.14

If the 91-tap FIR filter shown in Figure 7.14 is implemented in the ADSP-2101 microcomputer, each tap requires one processor cycle (80ns). The total processing time is therefore 7.3μs. This implies that sampling rates of up to about 136kHz can be achieved and still maintain real-time operation.

91 TAP FIR FILTER PERFORMANCE CHARACTERISTICS

- **0.002dB Passband Ripple**
- **Linear Phase**
- **80dB Stopband Attenuation**
- **136kHz Sampling Rate Possible with ADSP-2101 Processor (80ns Cycle Time per Filter Tap)**
- **No Analog Equivalent! (70 poles Required!)**

Figure 7.15

FIR Filter Design Using CAD Techniques

In actual practice, the concepts presented in the above discussions have been implemented in easy to use CAD programs which can be run on most PCs. It is only necessary to specify the desired FIR filter characteristics (sampling frequency, passband frequency, stopband frequency, passband ripple, and stopband attenuation) as shown in Figure 7.16. The CAD program calculates the number of filter taps required (N), the impulse response, and the filter coefficients.

KEY FILTER DESIGN PARAMETERS

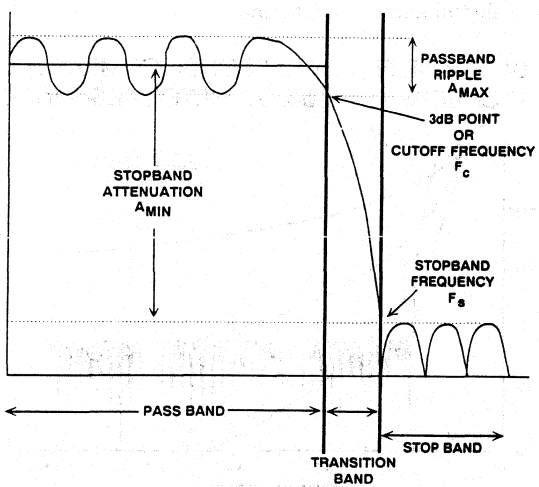

Figure 7.16

FIR FILTER DESIGN CAD PROGRAM INPUTS

- **Passband**
- **Passband Ripple**
- **Stopband**
- **Stopband Attenuation**
- **Wordlength, i.e., 16 Bit Fixed-Point**

Figure 7.17

A plot of the frequency response, H(f), along with the impulse response and the step-function response is also available as an output. If the response characteristics are satisfactory, the filter coefficients can then be downloaded into the DSP processor. The CAD program can also simulate the effects of finite word-length (i.e., performing calculations in 16 bit fixed point arithmetic) on the transfer function.

FIR FILTER DESIGN CAD PROGRAM OUTPUTS

- **Frequency Response Plot Showing Effects of Finite Wordlength Arithmetic**
- **Impulse Response Plot**
- **Step Function Response**
- **Number of Taps Required**
- **Filter Coefficients**

Figure 7.18

Other algorithms have been developed for CAD filter designs which optimize the filter performance for various characteristics. An example is the Parks and McClellan program (see Reference 1) which minimizes the maximum errors between the desired characteristic and the actual characteristic by using the Remez exchange algorithm from approximation theory.

Design Example for an FIR Digital Audio Filter Using CAD Program

For this example, we will design an audio lowpass filter that is designed to operate at a sampling rate of 44.1kHz (standard for CD players). The program is available from Momentum Data Systems, Incorporated (Reference 5). The program is menu-driven and IBM PC compatible. The filter will be implemented as a Direct Form FIR as shown in Figure 7.19.

DIRECT FORM FIR FILTER

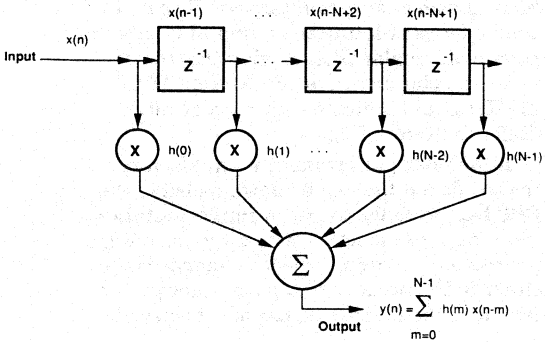

Figure 7.19

First, we select the type of filter to be designed from among the Main Menu shown in Figure 7.20. We choose the Equiripple FIR Design (Parks-McClellan)

FILTER DESIGN AND ANALYSIS SYSTEM MAIN MENU (Screen 1)

- **IIR Filter Design**
- **FIR Filter Design With Windows**
- ***Equiripple FIR Design (Parks-McClellan)***
- **Read Filter Specification File**
- **System Analysis (Z Domain Input)**
- **System Analysis (s Domain Input)**
- **Read System Analysis Input File**
- **Set System Defaults**
- **Exit to DOS**

Figure 7.20

The second screen then appears as shown in Figure 7.21. This screen is used to select the type of FIR filter (lowpass, highpass, bandpass, etc.) as well as specify the mode for frequency, gain, and whether sin(x)/x compensation is to be used.

FINITE IMPULSE RESPONSE FILTER DESIGN MENU (Screen 2)

Filter Type:	**1 - *Lowpass***
	2 - Highpass
	3 - Bandpass
	4 - Bandstop
	5 - Differentiator
	6 - Multiband
Frequency Mode:	**H - *Hertz***
	R - Radians/Second
Gain Specification Mode:	**1 - *Maximum Gain 1.0***
	2 - Nominal Gain 1.0
Filter Compensation:	**Enter X to Select**

Figure 7.21

The next screen appears as shown in Figure 7.22 where we enter the sampling rate, the band edges, and specifications for the passband ripple and stopband attenuation. The example we have chosen is a lowpass filter with a cutoff frequency of 18kHz.

FIR FILTER DESIGN LOWPASS FILTER (Screen 3)

Sampling Frequency:	***44100.0***
Passband Frequency:	***18000.0***
Stopband Frequency:	***21000.0***
Passband Ripple:	***1.00000E-02***
Stopband Ripple: (Attenuation)	***96dB***

Figure 7.22

The program will then calculate the required filter coefficients. When this calculation is complete, the screen shown in Figure 7.23 appears which lets us know the number of coefficients (taps) required to implement the filter. If the number of taps is compatible with the throughput of the DSP processor and the sampling rate, the user allows the program to proceed.

FIR DESIGN EXAMPLE (Screen 4)

Estimated Number of Taps of FIR Filter: 69

Enter Number of Taps Desired: *69*

Figure 7.23

If the 69-tap FIR filter is implemented in the ADSP-2101 microcomputer, each tap requires one processor cycle (80ns). The total processing time is therefore 5.5μs. This implies that sampling rates up to about 182kHz can be achieved and still maintain real-time operation.

ADSP-2101 PROCESSOR TIME FOR 69 TAP FIR FILTER

- **80ns (One Processor Cycle) per Tap**
- **69 Taps**
- **5.5μs Processor Time (80ns x 69)**
- **182kHz Sampling Rate for Real-Time Operation**

Figure 7.24

The next step shown in Figure 7.25 is to quantize the coefficients to the correct number of bits so that the coefficients are compatible with the DSP processor being used. In this example, the ADSP-2101 is to be used. It is a 16 bit fixed point machine, so the coefficients are quantized to 16 bits.

FIR DESIGN EXAMPLE (Screen 5)

Select the Desired Number of Bits for Quantization

Number of Bits (8 to 32): *16*

Figure 7.25

Now that the coefficients are calculated and properly quantized, we must see what effects on filter performance have been introduced by the quantization process. It should be noted that the filter design program initially calculates the coefficients with very high resolution. When these very accurate coefficients are quantized to a lower resolution, i.e. 16 bits, some accuracy is lost. This loss in accuracy may adversely affect the performance of the filter. To verify the proper performance, the filter is simulated. In this example the simulation is performed with 16 bit math. Figure 7.26 shows the simulated filter response so that the filter performance can be analyzed. Also available as outputs are the impulse response (shown in Figure 7.27) and the step response (shown in Figure 7.28). It should be clear that this filter has no analog counterpart. The rule of thumb for calculating the required number of poles of an analog filter having this transition band characteristic (85dB from 18 to 21kHz) would indicate a 65th order filter! (Refer to Section III).

If the filter performance is satisfactory, the coefficient file can be downloaded to the DSP hardware for the filter implementation. If the response is not satisfactory, the design process may be iterated with changes made either to the number of taps or other parameters until the desired response is achieved.

FIR FILTER DESIGN EXAMPLE
FREQUENCY RESPONSE

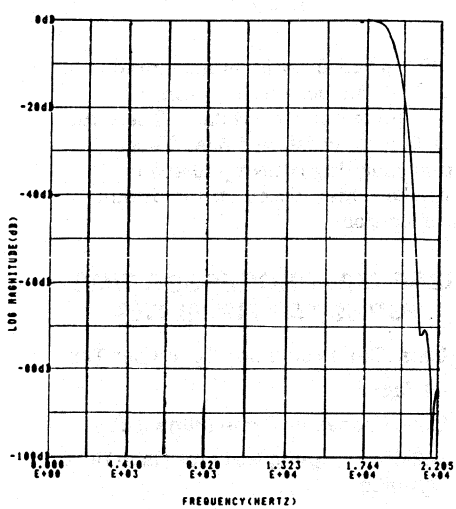

Figure 7.26

FIR FILTER DESIGN EXAMPLE
IMPULSE RESPONSE

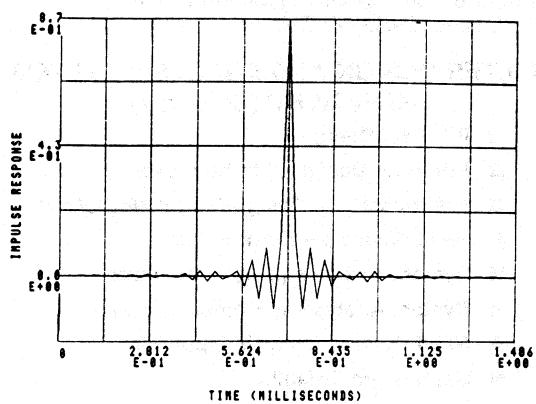

Figure 7.27

FIR FILTER DESIGN EXAMPLE
STEP RESPONSE

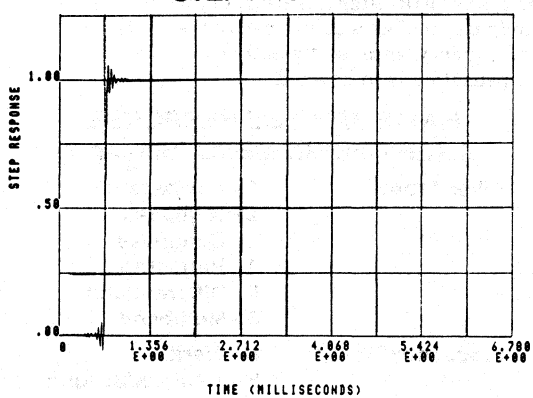

Figure 7.28

Insuring Linear Phase In FIR Filters

An advantage of FIR filters is they can always be made to have linear phase response which is a characteristic that makes them extremely attractive in audio and sonar applications. Linear phase means that all input frequencies are delayed by the same amount through the filter. In an FIR filter, this is the time required for the signal to propagate through the N taps. This delay is often referred to as *group delay* when applied to a band of frequencies. The group delay is constant for a linear phase FIR filter.

In order to insure phase linearity in an FIR filter, it is required that the filter coefficients are symmetric as in the case of a simple lowpass filter (Figure 7.29) or as in the case of a simple highpass filter (Figure 7.30). In addition, using an odd number of taps is also a requirement for linear phase.

SYMMETRICAL FILTER COEFFICIENTS PRODUCE LINEAR PHASE RESPONSE - LOWPASS FILTER

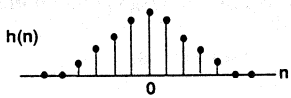

Figure 7.29

SYMMETRICAL FILTER COEFFICIENTS PRODUCE LINEAR PHASE RESPONSE - HIGHPASS FILTER

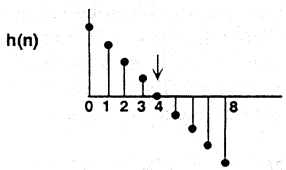

Figure 7.30

Decimation Using FIR Filters

FIR filters lend themselves to applications where data rate decimation is required, such as in oversampled sigma-delta ADCs. If we want to decimate the output data rate of an FIR filter by a factor of 2, for instance, we would take only every other sample point out of the filter. This also implies that the filter output computations need only be done every other sampling clock period. In other words, the DSP processor now has two sampling clock intervals to complete the convolution calculation. This implies that either more filter taps can be used, or perhaps a slower processor.

FIR FILTER PROPERTIES SUMMARY

- **Always Stable**
- **Have Linear Phase, Constant Group Delay**
- **Can be Adaptive**
- **Low Round-Off Noise**
- **Computational Advantages When Decimating Output**
- **Easy to Understand and Implement**

Figure 7.31

Infinite Impulse Response (IIR) Digital Filters

As was mentioned previously, digital FIR filters have no real analog counterparts, the closest analogy being the weighted moving average. In addition, FIR filters have only zeros and no poles. On the other hand, IIR filters have traditional analog counterparts (Butterworth, Chebyshev, and Elliptic) and can be analyzed and synthesized using more familiar traditional filter design techniques.

Figure 7.32 shows a second-order lowpass active filter, and its IIR digital filter equivalent is shown in Figure 7.33. This second-order IIR filter is referred to as the *biquad* (because it is described with a biquadratic

SECOND-ORDER ANALOG FILTER IMPLEMENTATION

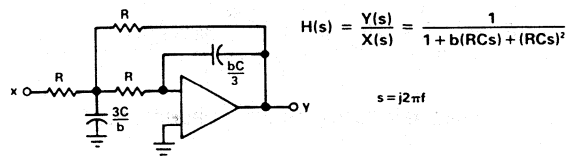

Figure 7.32

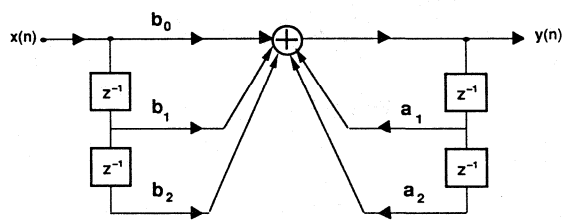

Figure 7.33

equation in the z-domain) and forms the basic building block for most higher order IIR designs. The difference equation which describes the characteristics of the filter with 5 coefficients is also shown in the figure.

The general digital filter equation is shown in Figure 7.34 which gives rise to the general transfer function H(z) which contains polynomials in both the numerator and the denominator. The roots of the denominator determine the pole locations of the filter, and the roots of the numerator determine the zero locations. Although it is possible to construct a high order IIR filter directly from this equation (called the *direct form* implementation), accumulation errors due to quantization errors (finite wordlength arithmetic) may give rise to instability and large errors. For this reason, it is common to cascade several biquad sections with appropriate coefficients rather than use the direct form implementation. The biquads can be scaled separately and then cascaded in order to minimize the coefficient quantization and the recursive accumulation errors. Cascaded biquads execute more slowly than their direct form counterparts, but are more stable and minimize the effects of errors due to finite arithmetic errors. In calculating the throughput time of a particular DSP IIR filter, one should examine the benchmark performance specification for a biquad filter section. For the ADSP-2101, the execution time for a single biquad section is 560ns, corresponding to seven instruction cycles.

GENERAL FILTER EQUATION

$$y(n) = \underbrace{\sum_{k=0}^{M} b_k x(n-k)}_{\text{FEEDFORWARD}} + \underbrace{\sum_{k=1}^{N} a_k y(n-k)}_{\text{FEEDBACK}}$$

GIVES RISE TO THE TRANSFER FUNCTION

$$H(z) = \frac{\sum_{k=0}^{M} b_k z^{-k} \quad (\text{ZEROS})}{1 - \sum_{k=1}^{N} a_k z^{-k} \quad (\text{POLES})}$$

Figure 7.34

IIR FILTER PROPERTIES SUMMARY

- **Feedback (Recursion)**
- **Potentially Unstable**
- **Usually Implemented as Cascaded Biquads Rather than Direct Form**
- **Non-Linear Phase**
- **More Efficient Than FIR Filters**
- **No Computational Advantage when Decimating Output**
- **Analogous to Analog Filters**

Figure 7.35

THROUGHPUT CONSIDERATION FOR IIR FILTERS

- **Determine How Many Biquad Sections are Required to Realize the Desired Filter Function**
- **Multiply by the Execution Time per Biquad (560ns for the ADSP-2101)**
- **The Result is the Minimum Sampling Period ($1/f_S$) Allowable for Real-Time Operation**

Figure 7.36

SUMMARY: FIR VERSUS IIR FILTERS

Choosing between FIR and IIR filter designs can be somewhat of a challenge, but a few basic guidelines can be given. Typically, IIR filters are more efficient than FIR filters because they require less memory and fewer multiplications are needed. IIR filters can be designed based upon previous experience with analog filter designs. IIR filters may exhibit instability problems, but this is much less likely to occur if higher order filters are designed by cascading second-order systems.

On the other hand, FIR filters require more taps and computations for a given cutoff frequency response, but do exhibit linear phase characteristics. Since FIR filters operate on a finite history of data, if some data is corrupted (ADC sparkle codes, for example) the FIR filter will ring for only N-1 samples. Because of the feedback, however, an IIR filter will ring for a considerably longer period of time.

If sharp cutoff filters are needed and processing time is at a premium, IIR elliptic filters are in order. If the number of multiplies is not prohibitive, and linear phase is a requirement, then the FIR should be chosen.

IIR VERSUS FIR FILTERS

IIR FILTERS	FIR FILTERS
More Efficient	Less Efficient
Analog Equivalent	No Analog Equivalent
May be Unstable	Always Stable
Non-Linear Phase Response	Linear Phase Response
More Ringing on Glitches	Less Ringing on Glitches
CAD Design Packages Available	CAD Design Packages Available
No Efficiency Gained by Decimation	Decimation Increases Efficiency

Figure 7.37

Fast Fourier Transforms

In many applications it is desired to process or analyze a signal in the frequency domain. In the analog world, this is easily accomplished using an analog spectrum analyzer. Mathematically, this process can be duplicated by taking the Fourier transform of the continuous-time analog signal. The Fourier transform yields the spectral content of the analog signal. In sampled data systems, however, this process must be accomplished by DSP processing of the ADC output data. Furthermore, there are two distinct differences between an analog and a digital spectral analysis. First, the output of the ADC is discrete quantized samples of the continuous input, x(t). In sampled data systems, the Discrete Fourier Transform (DFT) performs the transformation of the time domain samples into the frequency domain. In addition, the DFT must operate on a finite number of sampled data points, while the Fourier transform operates on a continuous waveform.

CONTINUOUS AND DISCRETE TIME-TO-FREQUENCY TRANSFORMATIONS

- **Fourier Transform Operates on Continuous-Time Waveforms**
- **Discrete Fourier Transform Operates on a Finite Number of Discrete Time Samples of a Waveform**

Figure 7.38

If x(n) is the sequence of N input data samples, then the DFT produces a sequence of N samples X(k) spaced equally in frequency. The DFT consists of a series of multiplications and additions where a data word is multiplied by a sinusoid value, and a number of these products are added together as shown in Figure 7.39.

THE DISCRETE FOURIER TRANSFORM (DFT) EQUATION

$$X(k) = \sum_{n=0}^{N-1} x(n)e^{-j2\pi nk/N}, \text{ where}$$

$$e^{-j2\pi nk/N} = \cos(2\pi nk/N) - j\sin(2\pi nk/N)$$

Figure 7.39

The DFT can be viewed as a correlation or comparison of the input signal to many sinusoids, evaluating the frequency content of the input signal. For example, a 1024 point DFT would require 1024 samples of the input signal and 1024 points from a sinusoid. Sinusoids of 1024 different frequencies equally spaced from $-f_s/2$ to $+f_s/2$ are used. Each pass of the DFT checks the sinusoid against the input signal to see how much of that frequency is present in the input signal. This is repeated for each of the 1024 frequencies. The result is shown in Figure 7.40 where N/2 discrete frequency components appear in the output spectrum. If the sampling frequency is f_s, then the spacing between the spectral lines is f_s/N, or $1/Nt_s$, where t_s is the sampling period, $1/f_s$.

Spectral analysis is most often performed with complex signals (having both real and imaginary components) so that phase information as well as amplitude and frequency information is obtained. In the above example, 1024 complex data values are multiply/accumulated with 1024 complex sinusoid values. This requires 1024 complex multiplies. This process is repeated for each of the 1024 frequencies for a total of 1024^2 multiplies, or in general terms, N^2 complex multiplies. Even for a powerful DSP device, this number of computations can be cumbersome and time consuming. This amount of compu-

TYPICAL FFT OUTPUTS FOR DIFFERENT RECORD LENGTHS

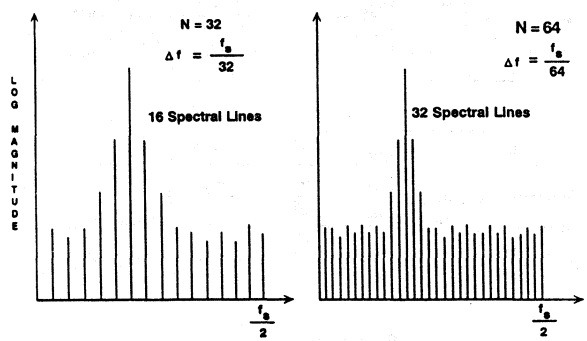

Figure 7.40

tation is only required when all output frequencies are to be calculated. If the value of frequency content for only one or a few frequencies is to be determined, the computational load is not as heavy.

In most spectral analysis situations, however, the entire frequency spectrum up to $f_s/2$ must be computed, so we must find a faster method! The FFT is simply an algorithm to speed up the DFT calculation by reducing the number of multiplications and accumulations required. It was popularized by J. W. Cooley in the 1960s and was actually a rediscovery of an idea of Runge (1903) and Danielson and Lanczos (1942), first occurring prior to the availability of computers and calculators-when numerical calculation could take many manhours.

The FFT is based on taking advantage of certain algebraic and trigonometric symmetries in the DFT computational process. For example, if a 1024 point DFT is performed, 1024^2 (1,048,567) complex multiplications are required. It is possible to break up the 1024 point DFT into two 512 point DFTs and end up with the same results. This is called *decimation*. Each 512 point DFT requires 512^2 (262,144) complex multiplications for a total of 524,288 complex multiplications. This is a significant reduction compared to the original 1,048,567. Figure 7.41 shows an N-point DFT broken up into two N/2-point DFTs. The presence of a phase factor W (sometimes called a *twiddle factor*) on a horizontal line indicates a multiplication by W. The points where the arrows intersect the horizontal lines indicates a summation. The presence of a -1 on the line indicates a sign reversal.

FIRST DECIMATION IN TIME OF 8-POINT DFT

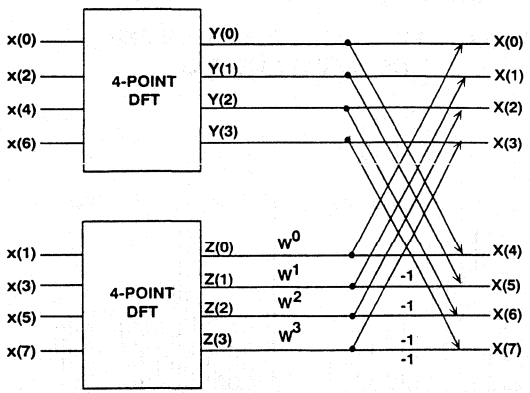

Figure 7.41

If it's possible to break up the 1024 point DFT into two 512 point DFTs and still get the same result, why can't each 512 point DFT be broken up into two 256 point DFTs for an even greater reduction in computations? Well, they can. This decimation process can continue until the original DFT is broken up into 2 point DFTs (the smallest DFT possible).

EIGHT-POINT DECIMATION-IN-TIME FFT NORMAL-ORDER INPUTS, BIT-REVERSED OUTPUTS

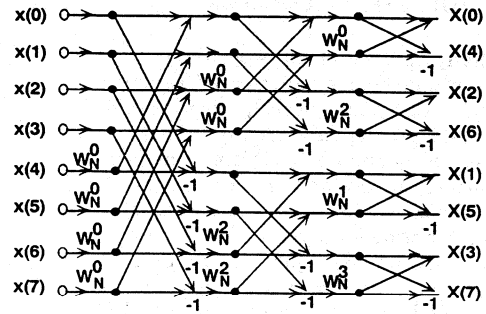

Figure 7.42

The final series of computations, after the decimation process is complete, is the FFT. This is shown for the 8 point DFT in Figure 7.42. Since the FFT was first decimated by a factor of 2, the FFT is known as a Radix-2 FFT. If the *initial* DFT was decimated by a factor of 4, it would be referred to as a Radix-4 FFT. Note that the input data points are taken in normal order, but the outputs are in bit-reversed order. Bit-reversing hardware is therefore common in DSP processors such as the ADSP-2101. The basic calculation, essentially the 2 point DFT, is commonly referred to as a *butterfly* calculation. The FFT is made up of many butterfly calculations. Figure 7.43 shows the basic butterfly for the Radix-2 decimation-in-time FFT which requires one complex multiply operation per butterfly.

RADIX-2 DECIMATION-IN-TIME FFT BUTTERFLY

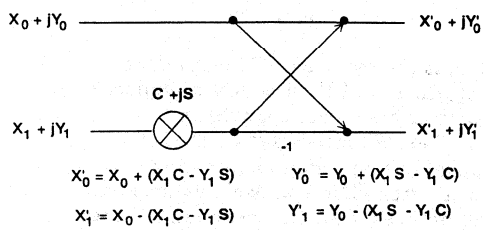

Figure 7.43

The significance of the FFT on the reduction in computations required to do the DFT is shown in Figure 7.44.

COMPUTATIONAL EFFICIENCY OF AN N-POINT FFT

DFT	FFT
N^2 Multiplications	$(N/2 \log_2(N)$
For N = 1024	For N = 1024
1,048,576 Multiplications	5,120 Multiplications

200 : 1

Figure 7.44

Note that the FFT results in the computation of *all* N/2 spectral outputs (all or nothing!). If only a few spectral points need to be calculated, the DFT is more efficient. Calculating a single spectral output using the DFT requires only N complex multiplications.

FFT Hardware Implementation

In general terms, the memory requirements for an N point FFT are N locations for real data, N locations for imaginary data, and N locations for the sinusoid data (sometimes referred to as the FFT coefficients or twiddle factors). As long as the memory requirements are met, the DSP processor must perform the necessary calculations in the required time. Many DSP vendors will either give a performance benchmark for a specified FFT size or a calculation time for a butterfly. When comparing FFT specifications, it is important to make sure that the same type of FFT is used in all cases. For example, a 1024 point FFT benchmark could have been derived from a Radix-2 or Radix-4 FFT and would not be compatible benchmarks since the number of computations required is different.

Once the basic hardware requirements are met, it is the job of the software to make the system realizable. With the same hardware, different software routines make possible a Radix-2, Radix-4, decimation-in-time or decimation-in-frequency algorithm just by manipulating the data in a different manner. An optimized Radix-4 FFT algorithm is given in Reference 6.

DSP FFT HARDWARE BENCHMARK COMPARISONS

- Radix-2, Radix-4 FFT?
- Butterfly Execution Time?
- Total FFT Execution Time?

Figure 7.45

FFT Design Considerations

The first step in designing an FFT is to determine the number of points required, N, or the *record* length. There are several ways to approach this problem. The sampling rate, f_s, must be at least twice the maximum input signal frequency of interest. Once the sampling rate is known, the spectral resolution of the FFT is then given by f_s/N. The more points in the FFT, the better the spectral resolution. This is a prime consideration in spectral analysis applications.

In real-time analysis of speech, for example, the signal bandwidth is approximately 4kHz, implying a sampling rate of 8kHz. The spectrum of speech is not stationary. The signal must be divided up into windows, T_w, short enough to ensure that individual features are not averaged out in the FFT; all meaning is lost in the *long-term* FFT of speech, for example. But T_w must be long enough to give adequate spectral resolution. It has been determined that for human speech phonemes, 20ms is adequate, hence T_w = 20ms.

REAL-TIME SPEECH ANALYSIS FFT EXAMPLE

- BW = 4kHz, Sampling Rate = 8kHz
- Window = 20ms
- N > 8kHz x 20ms = 160, Therefore use N = 256
- Can Processor Keep Up?
- ADSP-2101 Benchmark for N = 256 is 0.59ms
- Yes! With 19.41ms for Other Computations

Figure 7.46

Now, what determines if the FFT can keep up? The number of sample points in the window T_w is equal to $T_w f_s$, or 20ms x 8kHz = 160 points. This will be rounded up to the nearest power of 2, or 256 points. This says that the DSP processor must complete the 256 point FFT in less than the data acquisition time per window, T_w. Otherwise real-time processing is not possible, and the computation would have to be done off line. The ADSP-2101 completes a 256-point FFT in 0.59ms leaving 19.41ms for other computations.

Benchmark FFT processing times for most DSP processors are given by the manufacturer. Figure 7.48 shows Radix-4 benchmark times for the ADSP-2101. The 512-point benchmark time is for a Radix-2 FFT. In

evaluating various DSP processors, make sure to compare them under the same conditions. For instance, a Radix-4 FFT is somewhat faster than a Radix-2 FFT.

Figure 7.47 also shows the maximum sampling rates for real-time operation associated with the FFT execution times. These sampling rates indicate that modern DSP microcomputers such as the ADSP-2101 are capable of real-time FFT analysis of signals having bandwidths as great as 100 to 200kHz.

ADSP-2101 BENCHMARK FFT PERFORMANCE AND ASSOCIATED SAMPLING RATES FOR REALTIME OPERATION

FFT SIZE	EXECUTION TIME	MAXIMUM SAMPLING RATE
256	0.59ms	434kHz
512	1.3ms	394kHz
1024	2.9ms	353kHz
2048	6.5ms	315kHz
4096	14.2ms	288kHz

Figure 7.47

FFT OF SINEWAVE HAVING INTEGRAL NUMBER OF CYCLES IN WINDOW

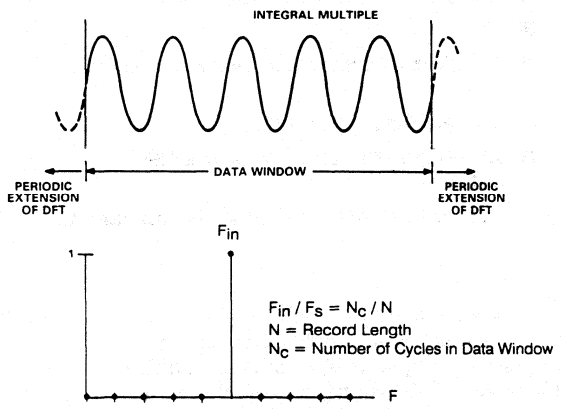

Figure 7.48

Spectral Leakage and Windowing

Spectral leakage in FFT processing can best be understood by considering the case of performing an FFT on a pure sinewave input. Two conditions will be considered. In Figure 7.48, the ratio between the sampling rate and the input sinewave frequency is such that precisely an integral number of cycles are contained within the data window (or record length). This results in a single tone FFT spectral response at the sinewave frequency as shown in the figure. Figure 7.49 shows the condition when the sinewave does not contain an integral number of cycles within the data window. The discontinuities at the endpoints are equivalent to multiplying the sinewave by a rectangular windowing pulse which has a sin(x)/x frequency domain response. The discontinuities in the time domain result in leakage in the frequency domain, because many spectral terms are needed to fit the discontinuity. Because of the endpoint discontinuity, the FFT spectral response shows the main lobe of the sinewave being smeared, an a large number of associated sidelobes which have the basic characteristics of the rectangular time pulse.

FFT OF SINEWAVE HAVING NON-INTEGRAL NUMBER OF CYCLES IN WINDOW

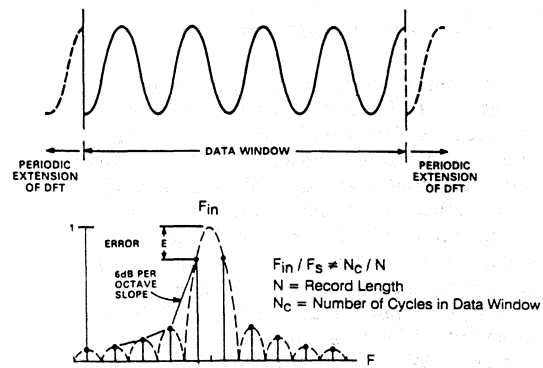

Figure 7.49

Since in practical FFT spectral analysis applications the exact frequencies are unknown, something must be done to minimize these sidelobes. This is done by choosing a windowing function other than the rectangular window. The input time samples are multiplied by an appropriate windowing function which brings the signal to zero at the edges of the window. The selection of an appropriate windowing function is primarily a tradeoff between main-lobe spreading and sidelobe rolloff. Leakage can also be reduced by padding the data with zeros and performing a correspondingly longer FFT. Reference 4 is highly recommended for an in-depth look at windows.

The time-domain and frequency-domain characteristics of a simple windowing function (the Hanning Window) are shown in Figure 7.50. A comparison of the frequency response of the Hanning window and the more sophisticated Minimum 4-Term Blackman-Harris window is given in Figure 7.51.

TIME AND FREQUENCY REPRESENTATION OF HANNING WINDOW

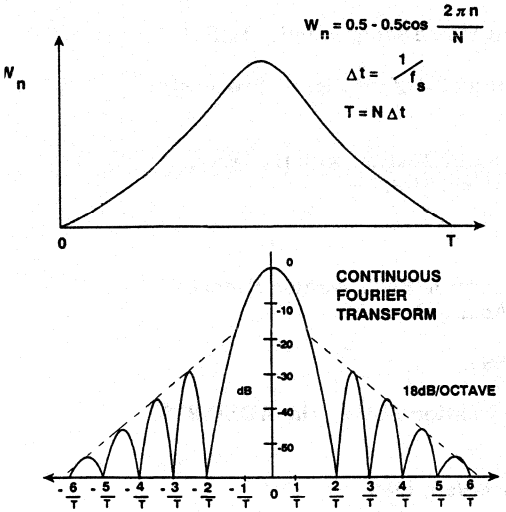

Figure 7.50

Data Scaling and Block Floating Point

The results of the butterfly calculation can be larger than the inputs to the butterfly. This data growth can pose a potential problem in a DSP with a fixed number of bits. To prevent data overflow, the data needs to be scaled before hand, leaving enough extra bits for growth. Alternatively, the data can be scaled after each pass of the FFT. The technique of scaling data after each pass of the FFT is known as *block floating point*. It is called this because the full array of data is scaled as a block regardless of whether or not each element in the block needs to be scaled. The complete block is scaled so that the relative relationship of each data word remains the same. For example, if each data word is shifted right one bit (divided by 2), the absolute values have been changed but relative to each other, the data stays the same.

COMPARISON OF WEIGHTING FUNCTIONS

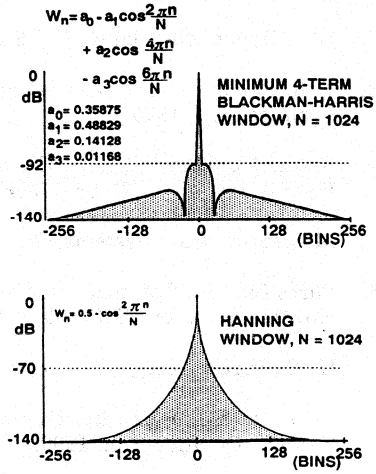

Figure 7.51

9

FFT SUMMARY

- The FFT is an Algorithm, not an Approximation
- Computational Speed is not Achieved at the Expense of Accuracy
- The FFT is a Fast Implementation of the DFT
- Resolution of the FFT in Frequency if f_S/N, N = Record Length
- Endpoint Discontinuities in Time Usually Require Smoothing Using Windowing Functions
- Real-Time FFT Processing Possible at Sampling Rates in Excess of 100kHz Using DSP Microcomputers

Figure 7.52

REFERENCES

1. Richard J. Higgins, **Digital Signal Processing in VLSI,** Prentice-Hall, 1990.

2. A. V. Oppenheim and R. W. Schafer, **Digital Signal Processing**, Prentice-Hall, 1975.

3. L. R. Rabiner and B. Gold, **Theory and Application of Digital Signal Processing**, Prentice-Hall, 1975.

4. Fredrick J. Harris, On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform, **Proc. IEEE**, Vol. 66, No. 1, 1978 pp. 51-83.

5. Momentum Data Systems, Inc. Costa Mesa, CA.

6. Fares Eidi, An Optimized Radix-4 Fast Fourier Transform (FFT), **Analog Devices Application Note E1329-5-9/89**. Available from Analog Devices.

7. **High Speed Design Seminar**, Analog Devices, 1990.

8. Amy Mar, Editor, **Digital Signal Processing Applications Using the ADSP-2100 Family**, Prentice-Hall, 1990.

9. C. S. Williams, **Designing Digital Filters,** Prentice-Hall, 1986.

10. R. W. Ramirez, **The FFT: Fundamentals and Concepts**, Prentice-Hall, 1985.

ANALOG DEVICES

AN-344
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Digital FIR Filters Without Tears

by Bill Windsor and Paul Toldalagi

Digital filters once required specialized design techniques, high-performance costly hardware, and complicated software to implement; for these reasons, they found only restricted application. Today, in sharp contrast, the availability of low-cost high-speed digital signal-processing ICs, such as multipliers and multiplier/accumulators, combined with easy-to-use standardized design procedures, has dramatically simplified filter implementation. Consequently, if your applications require filters with rolloffs in excess of 24 dB/octave, you should include digital filters in your design repertoire.

In these pages, we will compare digital and analog filters, discuss the various digital filter architectures, and—as an example—show you a step-by-step method of designing FIR (non-recursive) filters. A set of references will show you where to find information on topics only touched on lightly here.

COMPARE DIGITAL AND ANALOG FILTERS

Digital filters increasingly find their way into modems, radars, spectrum analyzers, and speech- and image-processing equipment, and for good reasons: Compared to analog filters, digital designs offer sharper rolloffs, require no calibration, and have greater stability with time, temperature, and power-supply variations. Simple software changes can alter a digital filter's response in real time, creating so-called "adaptive filters," whereas analog filters usually require hardware changes.

But digital filters do not satisfy every application. Analog techniques are usually most cost-effective in designs calling for rolloffs of up to about 24 dB/octave. As rolloff requirements exceed 24 to 36 dB/octave, however, digital filters increasingly make more sense. In fact, in applications calling for such steep rolloffs, many designers find digital filters significantly easier to develop. Prototypes can be easily altered through software changes. Also, software simulations of digital filter designs reflect the exact filter performance, whereas computer simulations of analog filters can only approximate true filter performance, since the parameters of analog filters are sensitive to component values that are initially inexact and can vary substantially.

DIGITAL FILTER BASICS

Common digital filter designs fall into two basic categories—non-recursive (finite impulse-response, FIR) and recursive (infinite impulse-response, IIR). Besides straightforward IIR designs, there is a growing interest in types embodying what is known as a lattice topology. But before examining these digital filter types, let us review some digital filter basics.

Digital filters are not as difficult to understand as you might at first think. A previous *Analog Dialogue* article introduced the subject to our readers (Vol. 17, Number 1, 1983, page 3); the references listed at the end of both that article and this one can provide greater detail.

Although filtering is often required for smoothing signals in the time domain, most designers understand the operation of a filter best in the frequency domain. The spectrum of the input signal is multiplied by the frequency response of the filter to produce an output signal with an altered spectrum. This multiplication in the frequency domain is equivalent to convolution of the waveform and a response function in the time domain. What then is convolution?

To understand the process, first consider a transfer function, H(f), with an ideal magnitude graph in the frequency domain as shown in Figure 1a. The function H(f) responds with unity gain to signals having frequency components from 0 Hz to f_1 Hz, where each frequency component is simply a cosine wave at a particular frequency. For instance, the signal, $\cos(2\pi 3t)$, represents a unity-amplitude frequency component at f = 3 Hz.

Figure 1b illustrates the spectrum of a signal, X(f), whose time value is $\cos(2\pi f_2 t) + \cos(2\pi f_3 t)$. X(f) therefore represents the sum of two equal components at f_2 and f_3. If you want to extract the f_2 component, leaving behind the f_3 component, you could simply pass the X(f) signal through a low pass filter. In fact, H(f) depicts just such a filter, with a cutoff frequency of f_1. Since H(f) equals 1 at f_2 and 0 at f_3, multiplying H(f) by X(f) gives you $1 \times \cos(2\pi f_2 t) + 0 \times \cos(2\pi f_3 t)$, or simply $\cos(2\pi f_2 t)$.

So far, we have been discussing continuous functions of time. However, in digital filters, we are dealing with sampled data, where a function of time consists of a finite number, k, of discrete values, x(n), per second, where k is the sampling rate and n/k is the discrete variable corresponding to time. Thus, a cosine waveform, in discrete time, is expressed as $\cos(2\pi fn/k)$.

Reprinted from Analog Dialogue 17-2 1983

The continuous Fourier transform provides a means for mapping continuous functions of time into the continuous complex frequency domain, and the inverse Fourier transform maps functions of frequency into the time domain. Similarly, the discrete Fourier transformation maps discrete functions of time into the discrete frequency domain, and its inverse transforms discrete functions of frequency into the discrete time domain.

If the function of frequency is the product of two functions—for example, the frequency content of a signal and the transfer function (i.e., the frequency response)—the corresponding time function is the same as the *convolution* of two functions in the time domain—i.e., the signal's time waveform and a time-response function, determined by the transfer function.

Thus, Fourier's theorem, which equates multiplication in the frequency domain to convolution in the time domain, provides a means of calculating the time response directly. As a consequence, the discrete-time convolution:

$$y(n) = [h * x](n) \quad (1)$$

is equal to the sum of the products of the signal and the frequency response, i.e.,

$$y(n) = \sum_{m=1}^{N} h(m) \cdot x(n - m) \quad (2)$$

for all values of n.

Equation 2 represents a series of multiplications and additions, which, if performed in a particular order, will automatically treat the input signal x(n) as if it were put through a low-pass filter. The equation assumes that h(n) is zero for $m < 1$ and for $m > N$, which happens to always be true for FIR filters, where N is the number of samples in h(n).

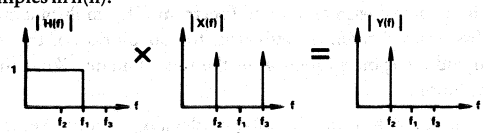

Figure 1. Only one frequency component remains (c) after the filter function (a) multiplies the signal at (b).

To perform the calculation of equation 1, using the Fourier theorem, all you need are the functions h(n) and x(n). These are the inverse discrete Fourier transforms of H(f) and X(f) of Figure 1. The transform of X(f) is a simple cosine wave, $x(n) = \cos(2\pi f_2 n/k) + \cos(2\pi f_3 n/k)$. As discussed in a later section, you can readily calculate the values of x(n) if you know f_2, f_3, and the sample rate, k – the rate at which your analog-to-digital converter is sampling the incoming time-domain signal x(t). You may have a little more difficulty computing the values of h(n), which are called the filter coefficients. But several good computer programs are available to help out, including one from Analog Devices.

To illustrate a practical example of equation 2, consider a 27th order filter, with N = 27. Then, the filter output value y(30), which depends on the 27 preceding values of x, will be:

$$y(30) = h(1)\cdot x(29) + h(2)\cdot x(28) + h(3)\cdot x(27) + \ldots + h(26)\cdot x(4) + h(27)\cdot x(3).$$

The physical meaning of this summation is that the filter's step response is synthesized by summing 27 successively delayed versions of the input step, each multiplied by its own coefficient, in effect building an arbitrary step response. For example, if each h(m) a gain of 1/27, the filter's response to a step will be a 27-step stai case (approximating an analog ramp), followed by constant ou put; with any input sequence, it performs a 27-interval runnir average.

The following sections discuss means for calculating the coefi cients.

DIGITAL FILTER TYPES

Figure 2 illustrates FIR and two of the most-prominent IIR digit filter topologies, the former straightforward and the latter in th form of a lattice. FIR, or finite impulse response, filters (Figure 2 have no feedback terms. Their outputs are a function only of a f nite number of previous input values (x(n)), and they are by defin tion nonrecursive. The filter in the example above is an FIR filte The IIR filters of 2b and 2c will be seen to have recursive term in which a value of the output is affected by previous values of th

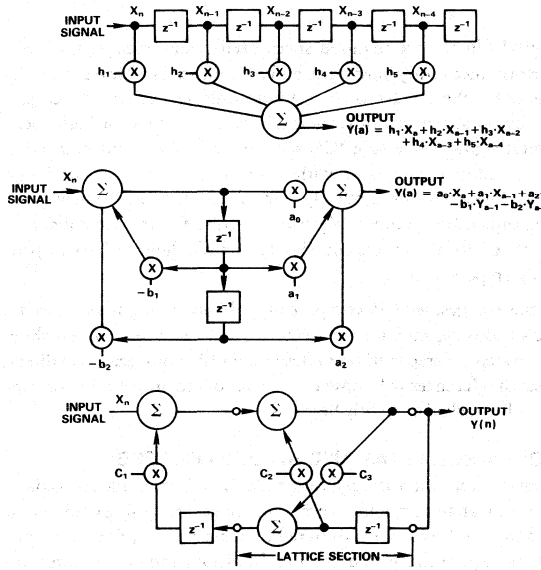

Figure 2. Three common digital filter topologies include FIF (a), IIR (b), and lattice type (c).

output (y(n)), as well as by input values. In comparison to the othe types, FIR filters offer:

Stability. FIR filters have no poles in their Z-plane transfer func tion. Thus, their output is always finite and stable. IIR filters ir contrast, require careful design to insure stability. Because FIR de signs are based on discrete time delays and have no poles, they car be used to construct filters for which there are no continuous ana log equivalents.

Linear Phase Response. You can design FIR filters with linea phase response – the phase delay of the output signal increases linearly with the frequency of the input signal. Linear phase response becomes particularly important in applications such as speech processing, sonar and radar. IIR filters, on the other hand, have nonlinear phase response. Linear phase response is difficult to achieve with continuous analog filters.

Ease of Design. Designers find the FIR filter the easiest of the three forms to understand, design, and implement, especially for indicial response in the time domain.

Low Sensitivity to Coefficient Errors. This permits FIR filters to e implemented with small word sizes – 12-to-16 bits for instance. ypical IIR filters need 16-to-24 bits per coefficient.

Accomodates Adaptive Designs. Adaptive FIR filters are comaratively easy to implement, by changes to the filter coefficients ı real time, to adapt the filter's characteristics to external condiions. Adaptive equalization filters in modems, for instance, are rogrammed to change their characteristics in response to changes ı the impedance of the transmission line.

IR* (infinite impulse response) filter outputs (Figure 2b) combine ıput values with previous output values, which have been fed ack into the circuit. IIR filters are therefore recursive. As in any eedback circuit, to avoid instabilities, IIR filter designs must avoid ositive feedback with gains equal to or greater than 1. IIR designs inear phase shift, and they need large coefficient word sizes to eep rounding errors small and insure stability. Nevertheless, IIR ilters have major advantages, including:

Highest Efficiency. IIR designs require fewer filter coefficients, hereby minimizing the number of multiplications and maximizing he throughput.

Least Memory Storage. Because the IIR filter has the least ıumber of coefficients, it requires the least amount of read-only nemory (ROM). For example, a typical highpass design requires ɔnly four coefficients in an IIR implementation, versus 19 for an FIR equivalent.

Lattice-type digital filters promise greater stability than IIR forms, with less hardware than FIR types. The newest form of digital filter, lattice filters presently have rapidly developing design theory. Although earlier lattice designs were highly sensitive to coefficient accuracy, recent designs have shown less sensitivity to filter parameters than the corresponding IIR filter (by 2 to 3 bits!). A big advantage of lattice filters is that the parameters used in each of the steps can be used for efficient encoding methods, as in linear predictive coding (speech).

DESIGNING FIR FILTERS

Specifications and Tradeoffs

Designers specify non-recursive (i.e., FIR) digital filters similarly to analog filters – a maximum amount of ripple in the passband, a maximum amount of attenuation in the stopband, etc. (See the adjacent definitions of digital filter terminology.) You will need to specify the following design parameters:

N, the number of taps in the filter, which equals the number of filter coefficients

f_p, the passband cutoff frequency

f_s, the stopband cutoff frequency,

$K = (\delta_1/\delta_2)$, the ratio of the ripple in the passband to the ripple in the stopband.

Figure 3 illustrates these parameters for lowpass, highpass and bandpass filters. Designers usually define the units of passband ripple in dB as $20 \log_{10} (1 + \delta_1)$, and the units of stopband ripple, also in dB, as $-20 \log_{10} (\delta_2)$. Passband ripple typically ranges from 0.001 to 1 dB, and stopband ripple from -10 dB to -90 dB. Frequencies f_p and f_s are normalized frequencies, which equal the ratio of the actual signal frequency to the sampling frequency. Consider, for example, a filter designed for a sampling frequency of 100 kHz,

*See Johnson, Matt, "Implement Stable IIR Filters Using Minimal Hardware." EDN, April 14, 1983, pp. 153-166.

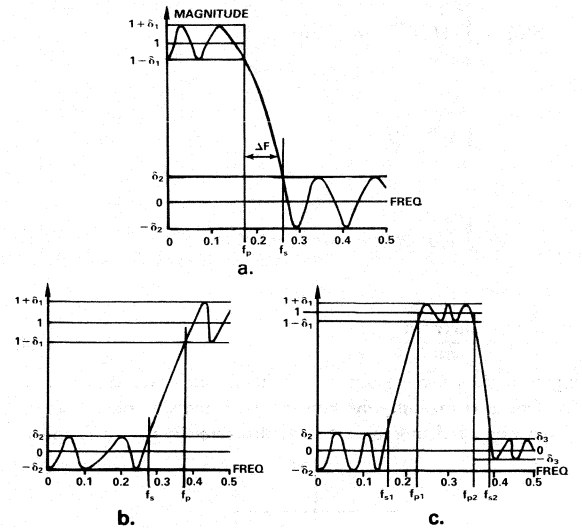

Figure 3. Design parameters defined for lowpass (a), highpass (b), and bandpass (c) filters.

a passband cutoff frequency (f_p) of 10 kHz, and a stopband cutoff frequency (f_s) of 20 kHz. Then,

$$f_p \text{ (normalized)} = 10\,\text{kHz}/100\text{kHz} = 0.1$$

$$f_s \text{ (normalized)} = 20\,\text{kHz}/100\text{kHz} = 0.2$$

Note that the normalized frequency axis extends from 0 to only 0.5, since a design in accordance with the Nyquist sampling theorem requires that a signal be sampled at more than twice its highest frequency in order to eliminate the possibility of aliasing.

As always, specifying these design parameters requires some tradeoffs. With a fixed number of filter taps, steeper rolloffs result in greater ripple. For both steep rolloffs and small ripple, you will have to increase the number of filter taps, and therefore the filter's complexity.

Designing FIR Filters Through Windowing

To design a digital filter, you must first calculate the filter's coefficients, h(m), in order to implement equation 2. The two most common design methods include "windowing" and the Remez Exchange algorithm. For almost 95% of design examples, Remez Exchange results in a significantly more efficient filter. The Remez Exchange algorithm has also been coded in Fortran, and is available from Analog Devices, as noted below.

Windowing methods are useful, however, because of their simplicity, and because they also aid in understanding filtering methods, so they are well worth examining. Keep in mind, though, that FIR designs developed through windowing do not perform as well as those obtained through other methods (see Ref. 7, for instance).

Consider the case of an FIR lowpass filter with stopband attenuation greater than 50 dB, normalized passband cutoff frequency (f_p) of 0.2, and normalized stopband cutoff frequency of 0.3. Figure 4 plots the filter's ideal transfer function, H(f). You can obtain the Fourier series coefficients by solving the inverse Fourier transform by equations 3.

$$h(n) = \int_{-0.5}^{0.5} H(\epsilon^{j2\pi f}) \cdot \epsilon^{j2\pi fn}\, df$$

$$= \int_{-f_1}^{f_1} \epsilon^{j2\pi fn}\, df$$

$$= \int_{-f_1}^{f_1} (\cos 2\pi fn + j \sin 2\pi fn)\, df$$

$$= \frac{\sin(\pi f_1 n)}{\pi n} \qquad (3)$$

Figure 5a shows the resulting set of h(n), which extend to ± infinity. You next multiply the Fourier coefficients by one of several window or weighting functions, as illustrated by Figure 5b and 5c.

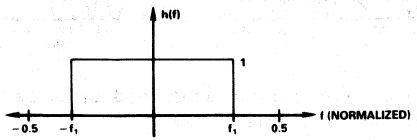

Figure 4. Idealized low-pass filter transfer function

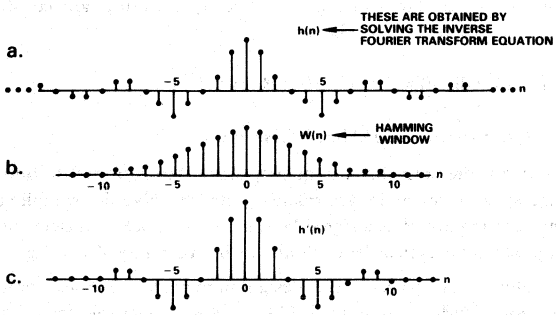

Figure 5. In the window method, a filter's Fourier coefficients (a) multiply a weighting function (b) resulting in (c).

The weighting function is equal to zero above and below some value, ν, which depends on the number of filter taps, N. Multiplying the Fourier coefficients by the weighting function generates a finite impulse response approximation to the desired transfer function, H(f). This guarantees that the Fourier series will converge.

Although several weighting functions will work, Figure 5b plots the widely used Hamming window. Others include the Kaiser, Blackman, and Hanning windows (see Ref. 1). After you choose your window, you can determine the number of coefficients or filter taps, N, from the desired rolloff band, $\Delta f = f_s - f_p$. For the Hamming window, this rolloff bandwidth relates to filter taps, N, by the conservative approximation:

$$\Delta f \approx 4/N \qquad (4)$$

For this design example,

$$\Delta f = f_s - f_p = (0.3 - 0.2) = 0.1.$$

Thus, N = 4/0.1 = 40. This approximation usually yields 2 to 5 more taps than needed, so specify N as 36.

Next, you obtain the filter coefficients, h'(n), of Figure 5c by multiplying each h(n) by the corresponding weighting function w(n) c Figure 5b. Since the coefficients are symmetrical about 0, you nee only compute their absolute values (i.e., half the coefficients). Th coefficients describe the windowed function, of the form (sin x)/x which is the Fourier transform of the low pass filter of Figure 4.

Remez Exchange Design

For most FIR applications, the Remez Exchange algorithm offer a more-powerful design technique than the windowing method. The Remez Exchange algorithm designs an optimal FIR filter a defined by the minimax error criterion (Ref. 7). The minima: criterion specifies a filter that, for a given number of coefficients *minimizes* the *maximum* ripple in the passband.

In general, for a given set of filter specifications, the Reme: Exchange algorithm quickly generates an FIR design with th smallest possible number of filter coefficients, particularly in com parison to the results of the windowing method. Also, passban ripples all have equal amplitude, as do all stopband ripples. Yo designate the ratio, K, of stopband to passband ripple, as note above.

The Fortran-coded Remez Exchange algorithm is easy to use. Con sider, for instance, the case of an FIR low-pass filter with the fol lowing specifications:

Sample rate = 50 kHz
f_p (actual) = 10 kHz, f_p (normalized) = 0.2
f_s (actual) = 14 kHz, f_s (normalized) = 0.28
minimum stopband attenuation = 40 dB
maximum passband ripple = 0.2 dB
ripple ratio K = 1 (equal ripple in pass- and stop-bands)

The Fortran program then prompts the user for inputs using fiv consecutive lines:

Line 1:

FILT = number of filter taps or coefficients. Set this equa to zero if the filter order (number of taps) is no known, as in this design example.

JTYPE = type of filter (set to 1 for low-pass, high-pass, o band-pass filters).

NBANDS = number of pass-bands plus stop-bands in th filter. For a low-pass filter or a high-pass filter a in this example, NBANDS = 2. A band-pass fil ter has NBANDS = 3.

JPUNCH = (normally set to zero).

LGRID = number of frequency points used in the Reme: Exchange algorithm. For most applications, suc as in this example, LGRID = 16 suffices. Fo high-performance filters with more than 50 taps set LGRID to 32.

Line 2:

Line 2 contains the normalized frequencies of the pass-band an stop-band edges. The number of values here equals twice th number of bands. For the case of the low-pass filter specifie above, the pass-band ranges from 0.0 to 0.2 in normalized fre quency, and the stop-band edge spans 0.28 to 0.5 in normalize frequency. Line 2 therefore carries these four numbers for this de sign example.

Line 3:

Line 3 specifies the magnitude of the desired transfer function,

V_{out}/V_{in}, in each band. In this example, the low-pass filter has unity gain in the pass-band, and zero gain in the stop-band, so Line 3 contains the numbers 1, 0.

Line 4:

Line 4 specifies the desired relative weights of the two bands. For this example, specify stop-band ripple equal to pass-band ripple, denoted by a 1, 1 on Line 4.

Line 5:

You need Line 5 only if the number of filter taps needed is unknown – as in this example (NFILT = 0). This line specifies the desired pass-band and stop-band ripple in dB. The program then estimates the number of required filter taps NFILT. Assume in this case that passband ripple does not exceed 0.2 dB, and that stop-band attenuation is 40.0 dB. Line 5 therefore includes the numbers 0.2 and 40.0.

With these inputs, the Fortran program estimates the filter order (number of taps) by approximating design relationships between the filter parameters (Refs. 5 and 6). The result usually falls within four taps of the correct number needed.

Figure 6 illustrates a typical computer result. The "filter length," determined by approximation as noted above, equals 24 taps. The "impulse response" gives the filter coefficients, and the next few lines of Figure 6 simply repeat the program input values for band 1, the pass-band, and band 2, the stop-band. The "desired value" indicates the desired filter transfer functions in the pass- and stop-bands. The "weighting" of the ripples is 1.00 in the passband and 1.00 in the stopband. The "deviation" is the ripple in each band, which equals 0.011 in the passband and 0.011 in the stopband. The "deviation in dB" represents the decibel value of the "deviation" numbers. "Extremal frequencies" denotes frequencies at which maximum passband and stopband ripple occurs.

```
**********************************************************

                FINITE IMPULSE RESPONSE (FIR)
                LINEAR PHASE DIGITAL FILTER DESIGN
                REMEZ EXCHANGE ALGORITHM

                BANDPASS FILTER

          FILTER LENGTH =  24

          FILTER LENGTH DETERMINED BY APPROXIMATION

          ***** IMPULSE RESPONSE *****
               H(  1) = -0.10748326E-01 = H(  24)
               H(  2) = -0.18704087E-02 = H(  23)
               H(  3) =  0.15714122E-01 = H(  22)
               H(  4) =  0.47213142E-02 = H(  21)
               H(  5) = -0.25240039E-01 = H(  20)
               H(  6) = -0.13135824E-01 = H(  19)
               H(  7) =  0.41533310E-01 = H(  18)
               H(  8) =  0.28864330E-01 = H(  17)
               H(  9) = -0.69702514E-01 = H(  16)
               H( 10) = -0.71426094E-01 = H(  15)
               H( 11) =  0.16081354E+00 = H(  14)
               H( 12) =  0.43491969E+00 = H(  13)

                    BAND  1          BAND  2          BAND
LOWER BAND EDGE     0.000000000      0.280000001
UPPER BAND EDGE     0.200000003      0.500000000
DESIRED VALUE       1.000000000      0.000000000
WEIGHTING           1.000000000      1.000000000
DEVIATION           0.011113827      0.011113827
DEVIATION IN DB     0.193075284    -39.082729340

EXTREMAL FREQUENCIES
   0.0000000   0.0416667   0.0807292   0.1197916   0.1562500
   0.1875001   0.2000000   0.2800000   0.2930208   0.3216665
   0.3607289   0.3997912   0.4414577

**********************************************************
```

Figure 6. Remez Exchange program output with NFILT initially = 0.

This initial computer run, with N = 24, results in passband ripple of 0.19 dB, and stopband attenuation of 39.08 dB, which do not meet the design specifications. Repeating the computer run, with successively higher values for N, leads to the acceptable results of Figure 7, for N equal to 27.

Hardware Design

Once fully defined, your filter can be readily implemented in hardware. Figure 8 is a functional diagram of a system implementing the 27-tap filter defined above, assuming 16-bit words. The tradeoffs in selecting word size will be discussed later.

```
**********************************************************

                FINITE IMPULSE RESPONSE (FIR)
                LINEAR PHASE DIGITAL FILTER DESIGN
                REMEZ EXCHANGE ALGORITHM

                BANDPASS FILTER

          FILTER LENGTH =  27

          ***** IMPULSE RESPONSE *****
               H(  1) =  0.37293066E-02 = H(  27)
               H(  2) = -0.72368127E-02 = H(  26)
               H(  3) = -0.90835225E-02 = H(  25)
               H(  4) =  0.77674030E-02 = H(  24)
               H(  5) =  0.15654538E-01 = H(  23)
               H(  6) = -0.11497792E-01 = H(  22)
               H(  7) = -0.28403830E-01 = H(  21)
               H(  8) =  0.14617096E-01 = H(  20)
               H(  9) =  0.50319906E-01 = H(  19)
               H( 10) = -0.17289791E-01 = H(  18)
               H( 11) = -0.97608425E-01 = H(  17)
               H( 12) =  0.19059505E-01 = H(  16)
               H( 13) =  0.31539205E+00 = H(  15)
               H( 14) =  0.48032677E+00 = H(  14)

                    BAND  1          BAND  2          BAND
LOWER BAND EDGE     0.000000000      0.280000001
UPPER BAND EDGE     0.200000003      0.500000000
DESIRED VALUE       1.000000000      0.000000000
WEIGHTING           1.000000000      1.000000000
DEVIATION           0.008834021      0.008834021
DEVIATION IN DB     0.153466403    -41.076831818

EXTREMAL FREQUENCIES
   0.0000000   0.0446429   0.0848214   0.1250000   0.1607143
   0.1875001   0.2000000   0.2800000   0.2911607   0.3179463
   0.3514283   0.3871424   0.4250887   0.4630350   0.5000000

**********************************************************
```

Figure 7. Remez Exchange program output, N = 27 taps.

The anti-aliasing filter of Figure 8 minimizes high-frequency signal and noise components reaching the a/d converter. In many cases, anti-aliasing filters require rolloffs no greater than 6-24 dB/octave. The a/d converter samples the incoming analog signal at a rate equal to about three times the highest input frequency. Although the Nyquist criterion specifies a sampling rate of at least two times the highest frequency, conservative design practice dictates a factor of three. The RAM stores the a/d converter's output. With a 27-tap filter, you will need 27 RAM locations, each 16-bits wide.

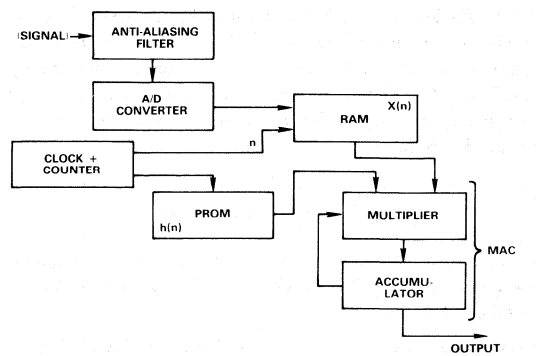

Figure 8. Digital FIR filter system functional block diagram.

A PROM stores the filter coefficients determined earlier. You may need RAM, instead of a PROM, particularly if you wish to implement an adaptive filter. The number of PROM locations equals the number of *different* filter coefficients. Because of symmetry, an FIR filter has N/2 different coefficients for N even (N is the number of taps), and (1 + N)/2 coefficients for N odd. A 27-tap filter therefore requires a PROM with 14 16-bit locations.

The clock and counter step through the RAM and PROM, presenting coefficients and input values to the multiplier. The multiplier/accumulator combination performs the multiplication and addition as specified by equation 2, and thus forms the heart of the digital filter.

Analog Devices offers a variety of multiplier/accumulator ICs which considerably simplify such digital filter implementations. The ADSP-1010*, for instance, multiplies two 16-bit numbers and accumulates the products in a 35-bit accumulator, which includes 3 bits of extended precision to accomodate overflows resulting from the addition of two or more 32-bit products.

Hardware Details

As a first step in implementing a detailed design, convert the filter coefficients to 16-bit fixed-point or block-floating point numbers. In fixed-point arithmetic, for instance, simply multiply the coefficients by 2^{15}.

Next, round off the coefficients to the nearest least-significant bit. Do not simply truncate the coefficients, since truncation destroys the accuracy of the filter coefficients, whereas rounding achieves performance close to the theoretical limits imposed by your word length. Store the rounded 16-bit coefficients in PROM.

You also must determine whether a standard μP can implement the filter, or whether you will need a dedicated high-speed multiplier IC. To determine the required computational speed, multiply the sampling rate by the number of filter coefficients.

In the above example, a sampling rate of 50 kHz and a filter with 27 taps requires (50kHz × 27) = 1.35 million 16-bit multiply-and-accumulate operations per second, or 740 nanoseconds per combined operation. Few microprocessors can handle such requirements; for instance, the 12.5-MHz version of the Motorola 68000 performs a 16-bit multiplication in 5.6 μs. The Analog Devices ADSP-1010 multiplier/accumulator (MAC), however, readily performs a multiply-and-accumulate operation in only 165 nanoseconds, at low cost and with low power consumption.

To ensure proper multiplications, the memory-control circuitry (RAM, PROM, counter) must retrieve the correct combination of words from memory. The stack and pointers of Figure 9 illustrate one method. Pointer 2 directs the storage of each new data point into RAM. For each new sample, the system computes the transformation by incrementing down from pointer 1 and up from pointer 3 as follows:

$$h(4)\cdot x(n-3) + h(3)\cdot x(n-2) + h(2)\cdot x(n-1) + h(1)\cdot x(n) + h(6)\cdot x(n-5) + h(5)\cdot x(n-4).$$

After computing each sample, pointers 2 and 3 increment; the pointers reset when they reach the stack boundary. Figure 9b details the operation.

Next, decide how to handle accumulator overflow. When a filter performs its multiply-and-accumulate operations, the number of bits in the accumulator will certainly exceed the 32-bit resolution of a single 16 × 16-bit multiplication. To handle overflow, first calculate a reasonable upper bound for the amount of overflow your filter could experience. By summing the squares of the filter coefficients, you can estimate a reasonable level of overflow. Compare this number to the maximum the accumulator can handle.

You can handle accumulator overflow in one of several ways. The ADSP-1010 MAC provides three additional bits of accumulator

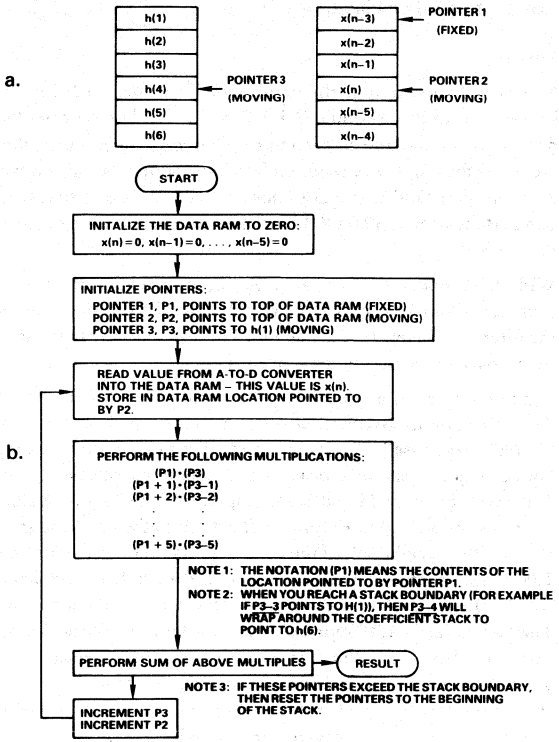

Figure 9. Filter uses pointers (a) to compute convolution as outlined in (b).

precision in addition to the 32 bits needed to handle a single 16 × 16-bit multiplication. This suffices for most applications.

Alternatively, you can scale down the coefficients, from 1 to 5 bits, at the sacrifice of some accuracy. To scale them down, divide them by 2 and apply the overflow test described above. Continue the process until the scaled coefficients pass the overflow test.

Finally, for some applications, you may not want to accomodate the full dynamic range of the input signal. Therefore, just let the accumulator saturate at its maximum value.

Occasionally, required multiply/accumulate speeds exceed the capabilities of even the fastest MAC ICs. In that case, you can combine two or more processors in parallel to increase the throughput. The circuit of Figure 10 combines two ADSP-1010 MACs operating in parallel, thus cutting the multiply/accumulate time per computed point to 75 nanoseconds, which is one-half the normal 150 nanoseconds for a single such component.

Avoid Rounding and Roundoff Errors

Most digital filter hardware errors result from two sources – rounding and roundoff. *Rounding* errors result from the rounding of filter coefficients, such as those generated in Figures 6 and 7,

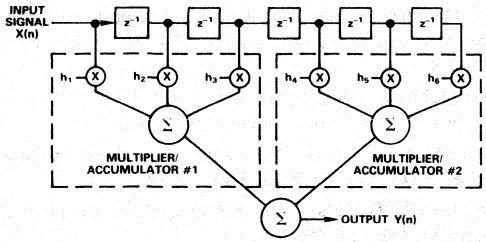

Figure 10. Paralleling two multiplier/accumulators doubles throughput.

by a high precision mainframe computer, to the 16-bits of typical digital filter hardware implementations. As noted earlier, rounding produces less error than truncating, but error nonetheless.

Roundoff errors result from consecutive finite precision multiply-and-accumulates. Roundoff errors are more significant than rounding errors, particularly in high-order filters.

Estimating the required word size to avoid such errors can prove tricky. Generally, if your design calls for more than 67 dB of stop-band attenuation or less than 0.05 dB of passband ripple, 16-bit words may lead to excessive errors. Such cases may require 24, and sometimes even 32-bit, word lengths. Software simulation, discussed in a following section, can help you determine word-length requirements before you commit your design to hardware.

To illustrate the significance of these errors, Figure 11 compares simulated performance of 16-bit fixed-point and 32-bit floating-point 27-tap low pass filters. Although the errors appear slight in this case, a similar comparison in Figure 12 for a 90-tap filter shows dramatic differences. For more than 80 dB of stop-band attenuation with 90 taps, more than 16 bits of precision are needed.

Software Simulations

The flow chart of Figure 13 shows a typical software program for simulating the performance of your digital filter with a high-resolution computer. You can obtain from Analog Devices a Fortran-coded version of this program for simulating 16-bit FIR designs, employing the ADSP-1010 16 × 16-bit multiplier/accumulator. It is available from the DSP Marketing Group, under the name, "FIR 16-bit simulation program."

The simulation repeats the steps in the hardware design process. It begins by obtaining the filter coefficients h(n) from the Remez-Exchange computer program, checks for overflow and scales the coefficients, and obtains the 16-bit fixed-point or floating-point coefficients, normally stored in PROM.

The program next simulates a digitized input signal array, x(n), which corresponds to the output of the A/D converter, normally stored in RAM. The number of values in the array equals the number of filter taps. Normally, you should begin the simulation with a cosine wave of frequency 0 Hz, and work you way up to higher frequencies.

The arithmetic operations of the multiplier/accumulator combination are readily simulated. The simulation program restricts the computer's word size to correspond to the limited precision (16 bits) of your filter's hardware implementation. In Fortran, for example, the INTEGER*2 or INTEGER*4 variable type declarations handle this for you.

The simulation program also includes an accumulator overflow check to verify the effectiveness of the initial coefficient scaling operation. If the computer flags an accumulator overflow, you'll have

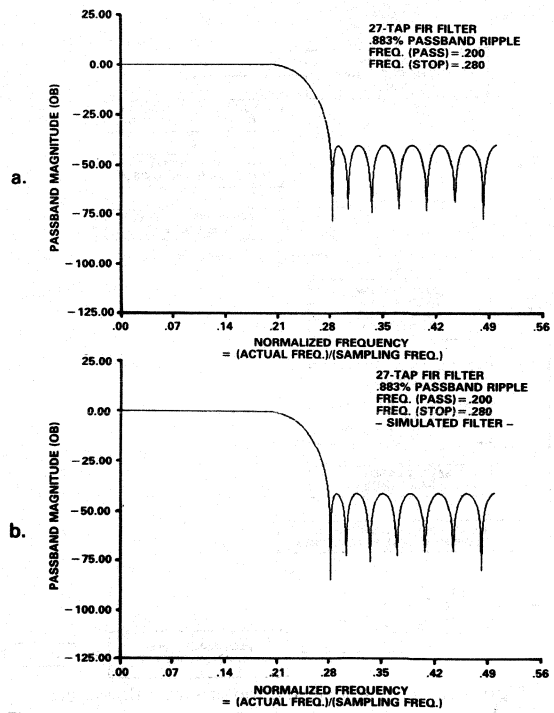

Figure 11. Computer-simulated response of a 27-tap low pass FIR filter using 32-bit arithmetic (a), and 16-bit arithmetic (b).

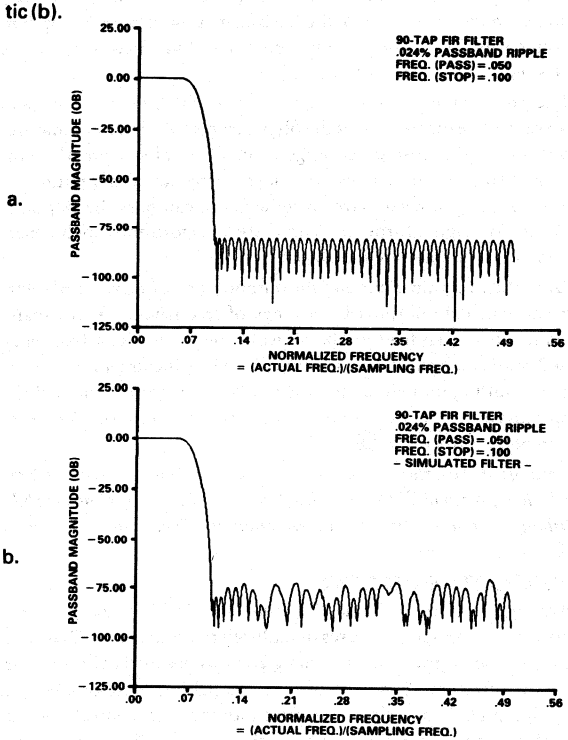

Figure 12. Computer simulated response of a 90-tap low pass FIR filter using 32-bit arithmetic (a), and 16-bit arithmetic (b).

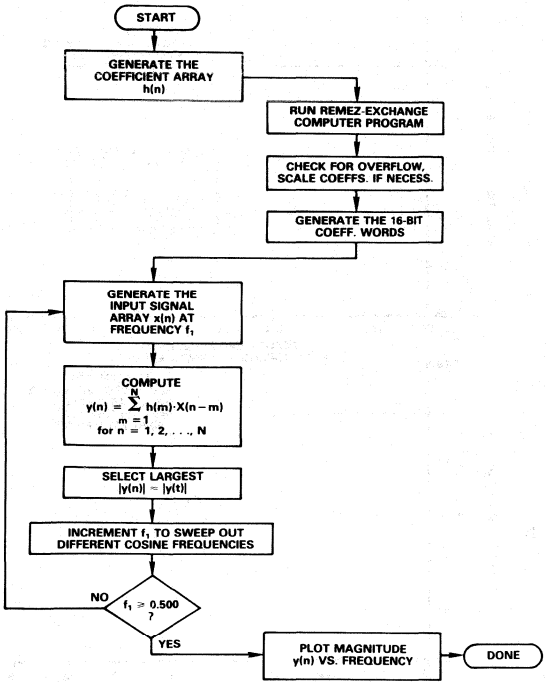

Figure 13. FIR filter simulation program flowchart.

to scale down the coefficients again and re-run the simulation.

The program next computes the filter output values, y(n), by setting up a loop to calculate the transformation of equation 2. For each cosine input, the program computes N values of y(n), where N is the number of filter taps.

The program next finds the magnitude of the filter's output (and therefore the magnitude of the filter's transfer function at that frequency), by choosing the largest absolute value from the y(n) array. This usually comes very close to the actual magnitude of y(n). But if you need better accuracy, you can pass the N points of the y(n) signal through a curve-fitting algorithm which generates a continuous-time signal, y(t).

Once the program computes the output for a cosine wave of 0 Hz, it can calculate outputs for a range of frequencies. You usually want it to sweep from 0 Hz to just below the Nyquist frequency of 0.5 (normalized), in normalized-frequency increments of 0.001. The resulting plot simulates your filter's transfer function. If the plot meets your expectations, you can proceed with the construction of the hardware. ◘

Portions of this article are adapted from an article by the same authors, which appeared in EDN, March 3, 1983, copyrighted by Cahners Publishing Company, Division of Reed Holdings, Inc, with permission of the copyright owner.

FOR FURTHER READING

Analog Dialogue, Volume 17, Number 1 (1983) carried a review article summarizing the uses of multiplier/accumulator ICs in a variety of DSP applications. Analog Devices has also collected a recently published series of 5 articles comprising cookbook designs applying DSP to various common filtering, and control applications. If you would like reprints of this set of articles, use the reply card; ask for "EDN series."

Ted Dintersmith, and Paul Toldalagi, "Apply Modern Contro Theory to Optimize Digital Systems," *EDN*, April 28, 1983, pp 165-179.

Matt Johnson, "Implement Stable IIR Filters Using Minima Hardware," *EDN*, April 14, 1983, pp. 153-166.

John Oxaal, "Temporal Averaging Techniques Reduce Imag Noise," *EDN*, March 17, 1983, pp. 211-215.

John Oxaal, "DSP Hardware Improves Multiband Filters," *EDN* March 31, 1983, pp. 193-197.

Bill Windsor, and Paul Toldalagi, "Simplify FIR-Filter Desigr With a Cookbook Approach," *EDN*, March 3, 1983, pp. 119 128.

REFERENCES

The following publications should prove useful to the reader seeking even more information on digital signal processing. These publications are *not* available from Analog Devices.

1. Frederick J. Harris, "On the Use of Windows for Harmonic Analysis With the Discrete Fourier Transform," *Proceedings of the IEEE*, Vol. 66, No. 1, January, 1978.

2. J. H. McClellan, T. W. Parks, and L. R. Rabiner, "A Computer Program for Designing Optimum FIR Linear Phase Digital Filters," *IEEE Transactions on Audio and Electroacoustics*, Vol AU-21, No. 6, December, 1973.

3. A. V. Oppenheim, and R. W. Schafer, *Digital Signal Processing* (Englewood Cliffs, New Jersey: Prentice-Hall, 1975), chapter 9.

4. A. Peled, and B. Liu, *Digital Signal Processing*, (New York, New York: John Wiley and Sons, Inc., 1976), chapter 2.

5. L. R. Rabiner, "Practical Design Rules for Optimum Finite Impulse Response Low-Pass Digital Filters," *The Bell System Technical Journal*, Vol. 52, No. 6, July-August, 1973.

6. L. R. Rabiner, "Approximate Design Relationships for Low-Pass FIR Digital Filters," *IEEE Transactions on Audio and Electroacoustics*, Vol. AU-21, No. 5, October, 1973.

7. L. R. Rabiner, and B. Gold, *Theory and Application of Digital Signal Processing*, (Englewood Cliffs, New Jersey: Prentice-Hall, 1975), chapter 3.

FILTER TERMINOLOGY

Attenuation – A decrease in output signal magnitude relative to input signal magnitude.

Cutoff frequency – The frequency at which the filter's response drops below the specified pass-band ripple ($1-\delta_1$).

Pass-band – The filter frequency range through which signals pass without more than a specified amount of attenuation.

Stop-band – The filter frequency range through which signals experience a specified degree of attenuation.

Stop-band attenuation – The minimum amount of attenuation in the stop-band.

Pass-band ripple – The maximum deviation from the desired output magnitude in the pass-band.

Sampling rate – The rate at which the system samples the input signal.

Filter coefficients – Numbers representing the inverse Fourier transform of the filter's transfer function. Coefficients define the filter's characteristics and form the basis of digital filter implementations.

Taps – Taps equal the number of sampled input values processed by the filter for each output point. Taps also equals the number of filter coefficients, and can represent a measure of the filter delay.

ANALOG DEVICES

AN-218
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

DSP Multirate Filters

MULTIRATE FILTERS

Multirate filters are digital filters that change the sampling rate of a digitally represented signal. These filters convert a set of input samples to another set of data that represents the same analog signal sampled at a different rate. A system incorporating multirate filters (a multirate system) can process data sampled at various rates.

Some examples of applications for multirate filters are:

- Sample-rate conversion between digital audio systems
- Narrow-band low-pass and band-pass filters
- Sub-band coding for speech processing in vocoders
- Transmultiplexers for TDM (time-division multiplexing) to FDM (frequency-division multiplexing) translation
- Quadrature modulation
- Digital reconstruction filters and antialias filters for digital audio, and
- Narrow-band spectra calculation for sonar and vibration analysis.

The two types of multirate filtering processes are decimation and interpolation. Decimation reduces the sample rate of a signal. It eliminates redundant or unnecessary information and compacts the data, allowing more information to be stored, processed, or transmitted in the same amount of data. Interpolation increases the sample rate of a signal. Through calculations on existing data, interpolation fills in missing information between the samples of a signal. Decimation reduces a sample rate by an integer factor M, and interpolation increases a sample rate by an integer factor L. Non-integer rational (ratio of integers) changes in sample rate can be achieved by combining the interpolation and decimation processes.

The ADSP-2100 programs in this application note demonstrate decimation and interpolation as well as efficient rational changes in sample rate. Cascaded stages of decimation and interpolation, which are required for large rate changes (large values of L and M) and are useful for implementing narrow-band low-pass and band-pass filters, are also demonstrated.

Decimation

Decimation is equivalent to sampling a discrete-time signal. Continuous-time (analog) signal sampling and discrete-time (digital) signal sampling are analogous.

Decimation Filter Structure

The decimation algorithm can be implemented in an FIR (Finite Impulse Response) filter structure. The FIR filter has many advantages for multirate filtering including: linear phase, unconditional stability, simple structure, and easy coefficient design. Additionally, the FIR structure in multirate filters provides for an increase in computational efficiency over IIR structures. The major difference between the IIR and the FIR filter is that the IIR filter must calculate all outputs for all inputs. The FIR multirate filter calculates an output for every Mth input. For a more detailed description of the FIR and IIR filters, refer to Crochiere and Rabiner, 1983.

The impulse response of the anti-imaging low-pass filter is h(n). A time-series equation filtering x(n) is the convolution

$$w(n) = \sum_{k=0}^{N-1} h(k)\, x(n-k)$$

where N is the number of coefficients in h(n). N is the order, or number of taps, in the filter. The application of this equation to implement the filter response $H(e^{jw})$ results in an FIR filter structure.

Figure 1a, shows the signal flowgraph of an FIR decimation filter. The N most recent input samples are stored in a delay line; z−1 is a unit sample delay. N samples from the delay line are multiplied by N coefficients and the resulting products are summed to form a single output sample w(n). Then w(n) is down-sampled by M using the rate compressor.

It is not necessary to calculate the samples of w(n) that are discarded by the rate compressor. Accordingly, the rate compressor can be moved in front of the multiply/accumulate paths, as shown in Figure 1b. This change reduces the required computation by a factor of M. This

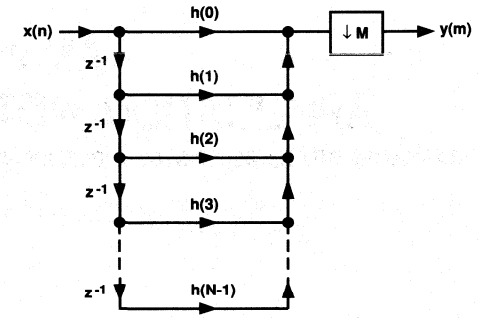

a.

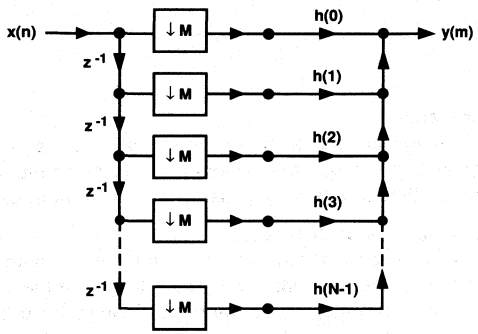

b.

Figure 1. FIR Form Decimation Filter

filter structure can be implemented by updating the delay line with M inputs before each output sample is calculated.

Substitution of the relationship between w(n) and y(m) into the convolution results in the decimation filtering equation

$$y(m) = \sum_{k=0}^{N-1} h(k)\, x(Mm-k)$$

Some of the implementations shown in textbooks on digital filters take advantage of the symmetry in transposed forms of the FIR structure to reduce the number of multiplications required. However, such a reduction of multiplications results in an increased number of additions. In this application, because the ADSP-2100 is capable of both multiplying and accumulating in one cycle, trading off multiplication for addition is a useless technique.

ADSP-2100 Decimation Algorithm

Figure 2 shows a flowchart of the decimation filter algorithm used for the ADSP-2100 routine. The decimator calculates one output for every M inputs to the delay line.

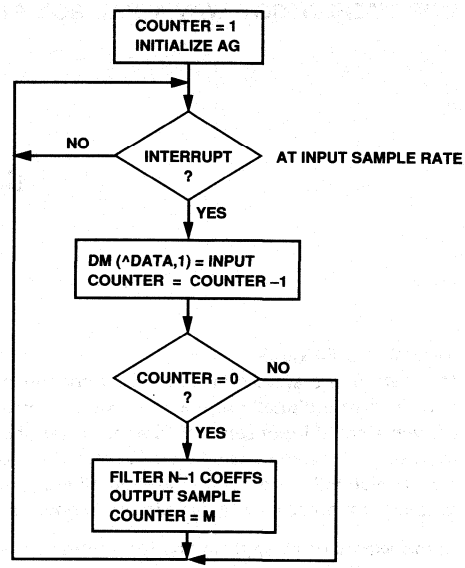

Figure 2. Decimator Flowchart

External hardware causes an interrupt at the input sample rate F_S, which triggers the program to fetch an input data sample and store it in the *data* circular buffer. The index register that points into this data buffer is then incremented by one, so that the next consecutive input sample is written to the next address in the buffer. The counter is then decremented by one and compared to zero. If the counter is not yet zero, the algorithm waits for another input sample. If the counter has decremented to zero, the algorithm calculates an output sample, then resets the counter to M so that the next output will be calculated after the next M inputs.

The output is the sum of products of N data buffer samples in a circular buffer and N coefficients in another circular buffer. Note that M input samples are written into the data buffer before an output sample is calculated. Therefore, the resulting output sample rate is equal to the input rate divided by the decimation factor: $F_S' = F_S/M$.

For additional information on the use of the ADSP-2100's address generators for circular buffers, see the *ADSP-2100 User's Manual,* Chapter 2.

The ADSP-2100 program for the decimation filter is shown in Listing 1. Inputs to this filter routine come from the memory-mapped port *adc*, and outputs go to the memory-mapped port *dac*. The filter's I/O interfacing hardware is described in more detail later in this application note.

```
{DECIMATE.dsp
Real time Direct Form FIR Filter, N taps, decimates by M for a decrease of 1/M times the input sample rate.

  INPUT: adc
  OUTPUT: dac

}
.MODULE/RAM/ABS=0    decimate;
.CONST               N=300;
.CONST               M=4;                    {decimate by factor of M}
.VAR/PM/RAM/CIRC     coef [N];
.VAR/DM/RAM/CIRC     data [N];
.VAR/DM/RAM          counter;
.PORT                adc;
.PORT                dac;
.INIT                coef:<coef.dat>;
                     RTI;                    {interrupt 0}
                     RTI;                    {interrupt 1}
                     RTI;                    {interrupt 2}
                     JUMP sample;            {interrupt 3= input sample rate}

initialize:          IMASK=b#0000;           {disable all interrupts}
                     ICNTL=b#01111;          {edge sensitive interrupts}
                     SI=M;                   {set decimation counter to M}
                     DM(counter)= SI;        {for first input data sample}
                     I4=^coef;               {setup a circular buffer in PM}
                     L4=%coef;
                     M4=1;                   {modifier for coef is 1}
                     I0=^data;               {setup a circular buffer in DM}
                     L0=%data;
                     M0=1;                   {modifier for data is 1}
                     IMASK=b#1000;           {enable interrupt 3}
wait_interrupt:      JUMP wait_interrupt;    {infinite wait loop}
{_________Decimator, code executed at the sample rate_________}
sample:              AY0=DM(adc);
                     DM(I0,M0)=AY0;          {update delay line with newest}
                     AY0=DM(counter);
                     AR=AY0-1;               {decrement and update counter}
                     DM(counter)=AR;
                     IF NE RTI;              {test and return if not M times}
{_____  code below executed at 1/M times the sample rate_____}
do_fir:              AR=M;                   {reset the counter to M}
                     DM(counter)=AR;
                     CNTR=N - 1;
                     MR=0, MX0=DM(I0,M0), MY0=PM(I4,M4);
                     DO taploop UNTIL CE;    {N-1 taps of FIR}
taploop:               MR=MR+MX0*MY0(SS), MX0=DM(I0,M0), MY0=PM(I4,M4);
                     MR=MR+MX0*MY0(RND);     {last tap with round}
                     IF MV SAT MR;           {saturate result if overflow}
                     DM(dac)=MR1;            {output data sample}
                     RTI;
.ENDMOD;
```

Listing 1. Decimation Filter

The routine uses two circular buffers, one for data samples and one for coefficients, that are each N locations long. The *coef* buffer is located in program memory and stores the filter coefficients. Each time an output is calculated, the decimator accesses all these coefficients in sequence, starting with the first location in *coef*. The I4 index register, which points to the coefficient buffer, is modified by one (from modify register M0) each time it is accessed. Therefore, I4 is always modified back to the beginning of the coefficient buffer after the calculation is complete.

The FIR filter equation starts the convolution with the most recent data sample and accesses the oldest data sample last. Delay lines implemented with circular buffers, however, access data in the opposite order. The oldest data sample is fetched first from the buffer and the newest data sample is fetched last. Therefore, to keep the data/coefficient pairs together, the coefficients must be stored in memory in reverse order.

The relationship between the address and the contents of the two circular buffers (after N inputs have occurred) is shown in the table below. The *data* buffer is located in data memory and contains the last N data samples input to the filter. Each pass of the filter accesses the locations of both buffers sequentially (the pointer is modified by one), but the first address accessed is not always the first location in the buffer, because the decimation filter inputs M samples into the delay line before starting each filter pass. For each pass, the first fetch from the data buffer is from an address M greater than for the previous pass. The data delay line moves forward M samples for every output calculated.

Data		*Coefficient*	
DM(0)	= x(n−(N−1)) oldest	PM(0)	= h(N−1)
DM(1)	= x(n−(N−2))	PM(1)	= h(N−2)
DM(2)	= x(n−(N−3))	PM(2)	= h(N−3)
	•		•
	•		•
	•		•
DM(N−3)	= x(n−2)	PM(N−3)	= h(2)
DM(N−2)	= x(n−1)	PM(N−2)	= h(1)
DM(N−1)	= x(n−0) newest	PM(N−1)	= h(0)

A variable in data memory is used to store the decimation counter. One of the processor's registers could have been used for this counter, but using a memory location allows for expansion to multiple stages of decimation.

The number of cycles required for the decimation filter routine is shown below. The ADSP-2100 takes one cycle to calculate each tap (multiply and accumulate), so only 18+N cycles are necessary to calculate one output sample of an N-tap decimator. The 18 cycles of overhead for each pass is just six cycles greater than the overhead of a non-multirate FIR filter.

Interrupt Response	2 Cycles
Fetch Input	1 Cycle
Write Input to Data Buffer	1 Cycle
Decrement and Test Counter	4 Cycles
Reload Counter with M	2 Cycles
FIR Filter Pass	7+N Cycles
Return from Interrupt	1 Cycle
Maximum Total	18+N Cycles/Output

A More Efficient Decimator

The routine in Listing 1 requires that the 18+N cycles needed to calculate an output occur during the first of the M input sample intervals. No calculations are done in the remaining M−1 intervals. This limits the number of filter taps that can be calculated in real time to:

$$N = \frac{1}{F_S\, t_{CLK}} - 18$$

where t_{CLK} is the instruction cycle time of the processor.

An increase in this limit by a factor of M occurs if the program is modified so that the M data inputs overlap the filter calculations. This more efficient version of the program is shown in Listing 2.

In this example, a circular buffer *input_buf* stores the M input samples. The code for loading *input_buf* is placed in an interrupt routine to allow the input of data and the FIR filter calculations to occur simultaneously.

A loop waits until the input buffer is filled with M samples before the filter output is calculated. Instead of counting input samples, this program determines that M samples have been input when the input buffer's index register I0 is modified back to the buffer's starting address. This strategy saves a few cycles in the interrupt routine.

After M samples have been input, a second loop transfers the data from *input_buf* to the data buffer. An output sample is calculated. Then the program checks that at least one sample has been placed in *input_buf*. This check prevents a false output if the output calculation occurs in less than one sample interval. Then the program jumps back to wait until the next M samples have been input.

This more efficient decimation filter spreads the calculations over the output sample interval $1/F_S'$ instead of the input interval $1/F_S$. The number of taps that can be calculated in real time is:

$$N = \frac{M}{F_S\, t_{CLK}} - 20 - 2M - 6(M-1)$$

which is approximately M times greater than for the first routine.

)EC_EFF.dsp
eal time Direct Form FIR Filter, N taps, decimates by M for a decrease of 1/M times the input sample rate. This version
ses an input buffer to allow the filter computations to occur in parallel with inputs. This allows larger order filter for a
iven input sample rate. To save time, an index register is used for the input buffer as well as for a decimation counter.

```
  INPUT: adc
  OUTPUT: dac
}
.MODULE/RAM/ABS=0          eff_decimate;
.CONST                     N=300;
.CONST                     M=4;                    {decimate by factor of M}
.VAR/PM/RAM/CIRC           coef[N];
.VAR/DM/RAM/CIRC           data[N];
.VAR/DM/RAM/CIRC           input_buf[M];
.PORT                      adc;
.PORT                      dac;
.INIT                      coef:<coef.dat>;
      RTI;                       {interrupt 0}
      RTI;                       {interrupt 1}
      RTI;                       {interrupt 2}
      JUMP sample;               {interrupt 3= input sample rate}
initialize:                IMASK=b#0000;           {disable all interrupts}
                           ICNTL=b#01111;          {edge sensitive interrupts}
                           I4=^coef;               {setup a circular buffer in PM}
                           L4=%coef;
                           M4=1;                   {modifier for coef is 1}
                           I0=^data;               {setup a circular buffer in DM}
                           L0=%data;
                           M0=1;                   {modifier for data is 1}
                           I1=^input_buf;          {setup a circular buffer in DM}
                           L1=%input_buf;
                           IMASK=b#1000;           {enable interrupt 3}
wait_M:                    AX0=I1;                 {wait for M inputs}
                           AY0=^input_buf;
                           AR=AX0-AY0;             {test if pointer is at start}
                           IF NE JUMP wait_M;
{_____code below executed at 1/M times the sample rate_____}
                           CNTR=M;
                           DO load_data UNTIL CE;
                               AR=DM(I1,M0);
load_data:                     DM(I0,M0)=AR;
fir:                       CNTR=N - 1;
                           MR=0, MX0=DM(I0,M0), MY0=PM(I4,M4);
                           DO taploop UNTIL CE;    {N-1 taps of FIR}
taploop:                       MR=MR+MX0*MY0(SS), MX0=DM(I0,M0), MY0=PM(I4,M4);
                           MR=MR+MX0*MY0(RND);     {last tap with round}
                           IF MV SAT MR;           {saturate result if overflow}
                           DM(dac)=MR1;            {output data sample}
wait_again:                AX0=I1;
                           AY0=^input_buf;
                           AR=AX0-AY0;             {test and wait if i1 still}
                           IF EQ JUMP wait_again;  {points to start of input_buf}
                           JUMP wait_M;
{_____sample input, code executed at the sample rate_____}
sample:                    ENA SEC_REG;            {so no registers will get lost}
                           AY0=DM(adc);            {get input sample}
                           DM(I1,M0)=AY0;          {load in M long buffer}
                           RTI;
.ENDMOD;
```

Listing 2. Efficient Decimation Filter

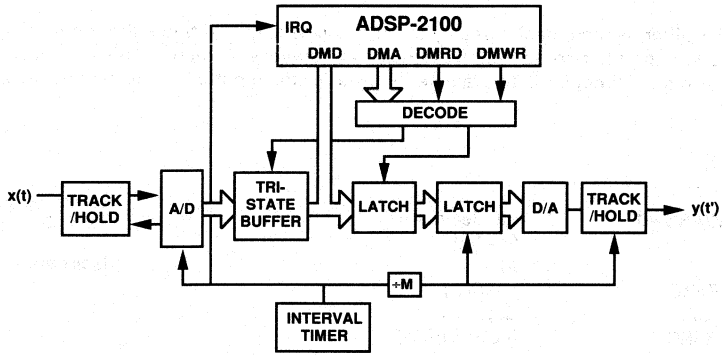

Figure 3. Decimator Hardware

Decimator Hardware Configuration

Both decimation filter programs assume an ADSP-2100 system with the I/O hardware configuration shown in Figure 3. The processor is interrupted by an interval timer at a frequency equal to the input sample rate F_S and responds by inputting a data value from the A/D converter. The track/hold (sampler) and the A/D converter (quantizer) are also clocked at this frequency. The D/A converter on the filter output is clocked at a rate of F_S/M, which is generated by dividing the interval timer frequency by M.

To keep the output signal jitter-free, it is important to derive the D/A converter's clock from the interval timer and not from the ADSP-2100. The sample period of the analog output should be disassociated from writes to the D/A converter. If an instruction-derived clock is used, any conditional instructions in the program could branch to different length program paths, causing the output samples to be spaced unequally in time. The D/A converter must be double-buffered to accommodate the interval-time-derived clock. The ADSP-2100 outputs data to one latch. Data from this latch is fed to a second latch that is controlled by an interval-timer-derived clock.

Interpolation

The process of recreating a continuous-time signal from its discrete-time representation is called reconstruction. Interpolation can be thought of as the reconstruction of a discrete-time signal from another discrete-time signal, just as decimation is equivalent to sampling the samples of a signal. Continuous-time (analog) signal reconstruction and discrete-time (digital) signal reconstruction are analogous.

Interpolation Filter Structure

Figure 4a shows a block diagram of an interpolation filter. The two major differences from the decimation filter are that the interpolator uses a sample rate expander instead of the sample rate compressor and that the interpolator's low-pass filter is placed after the rate expander instead of before the rate compressor. The rate expander, which is the block labeled with an up-arrow and L, inserts L−1 zero-valued samples after each input sample. The resulting w(m) is low-pass filtered to

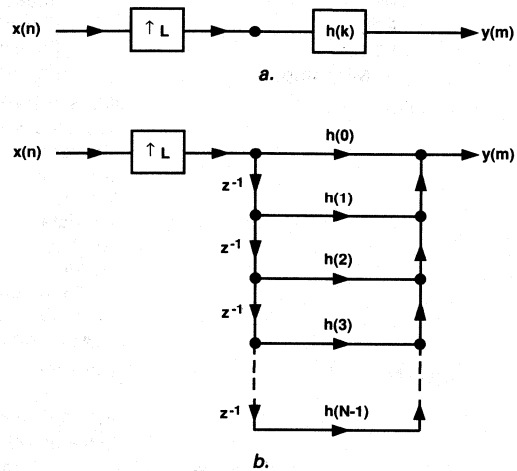

Figure 4. Interpolation Filter Block Diagram

produce y(m), a smoothed, anti-imaged version of w(m). The transfer function of the interpolator H(k) incorporates a gain of 1/L because the L−1 zeros inserted by the rate expander cause the energy of each input to be spread over L output samples.

The low-pass filter of the interpolator uses an FIR filter structure for the same reasons that an FIR filter is used in the decimator, notably computational efficiency. The convolution equation for this filter is

$$y(m) = \sum_{k=0}^{N-1} h(k)\, w(m-k)$$

N−1 is the number of filter coefficients (taps) in h(k), w(m−k) is the rate expanded version of the input x(n), and w(m−k) is related to x(n) by

$$w(m-k) = \begin{cases} x((m-k)/L) & \text{for } m-k = 0, \pm L, \pm 2L, \ldots \\ 0 & \text{otherwise} \end{cases}$$

The signal flowgraph that represents the interpolation filter is shown in Figure 4b. A delay line of length N is loaded with an input sample followed by L−1 zeros,

hen the next input sample and L−1 zeros, and so on. he output is the sum of the N products of each sample rom the delay line and its corresponding filter coeffi- ient. The filter calculates an output for every sample, ero or data, loaded into the delay line.

An example of the interpolator operation is shown in the signal flowgraph in Figure 5. The contents of the delay ine for three consecutive passes of the filter are high- ighted. In this example, the interpolation factor L is 3. The delay line is N locations long, where N is the num- ber of coefficients of the filter; N=9 in this example. There are N/L or 3 data samples in the delay line during each pass. The data samples x(1), x(2), and x(3) in the first pass are separated by L−1 or 2 zeros inserted by the rate expander. The zero-valued samples contribute L−1)N/L or 6 zero-valued products to the output result. These (L−1)N/L multiplications are unnecessary and waste processor capacity and execution time.

A more efficient interpolation method is to access the coefficients and the data in a way that eliminates wasted calculations. This method is accomplished by removing the rate expander to eliminate the storage of the zero- valued samples and shortening the data delay line from N to N/L locations. In this implementation, the data delay line is updated only after L outputs are calculated. The same N/L (three) data samples are accessed for each set of L output calculations. Each output calculation ac- cesses every Lth (third) coefficient, skipping the coeffi- cients that correspond to zero-valued data samples.

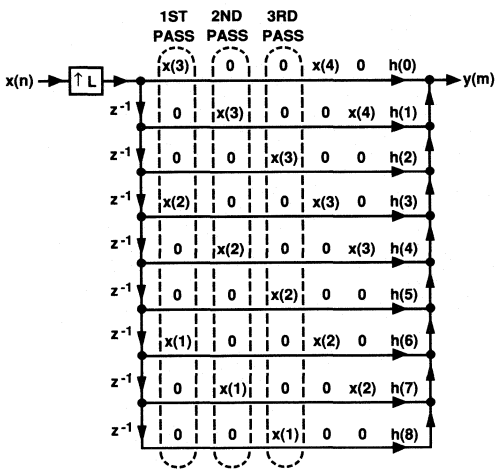

Figure 5. Example Interpolator Flowgraph

Crochiere and Rabiner refer to this efficient interpolation filtering method as *polyphase filtering*, because a differ- ent phase of the filter function h(k) (equivalent to a set of interleaved coefficients) is used to calculate each output sample.

ADSP-2100 Interpolation Algorithm

A circular buffer of length N/L located in data memory forms the data delay line. Although the convolution equation accesses the newest data sample first and the oldest data sample last, the ADSP-2100 fetches data samples from the circular buffer in the opposite order: oldest data first, newest data last. To keep the data/coef- ficient pairs together, the coefficients are stored in pro- gram memory in reverse order, e.g., h(N−1) in PM(0) and h(0) in PM(N−1).

Figure 6 shows a flowchart of the interpolation algo- rithm. The processor waits in a loop and is interrupted at the output sample rate (L times the input sample rate). In the interrupt routine, the coefficient address pointer is decremented by one location so that a new set of inter- leaved coefficients will be accessed in the next filter pass. A counter tracks the occurrence of every Lth out- put; on the Lth output, an input sample is taken and the coefficient address pointer is set forward L locations, back to the first set of interleaved coefficients. The out- put is then calculated with the coefficient address pointer incremented by L locations to fetch every Lth coefficient. One restriction in this algorithm is that the number of filter taps must be an integral multiple of the interpolation factor; N/L must be an integer.

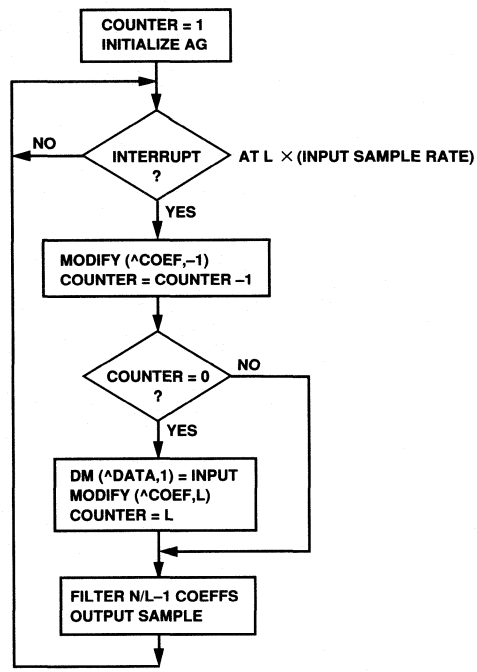

Figure 6. Interpolation Flowchart

Listing 3 is an ADSP-2100 program that implements this interpolation algorithm. The ADSP-2100 is capable of calculating each filter pass in ((N/L)+17) processor instruction cycles. Each pass must be calculated within the period between output samples, equal to $1/F_S L$. Thus the maximum number of taps that can be calculated in real time is:

$$N = \frac{1}{F_S\, t_{CLK}} - 17L$$

where t_{CLK} is the processor cycle time and F_S is the input sampling rate.

The interpolation filter has a gain of 1/L in the passband. One method to attain unity gain is to premultiply (offline) all the filter coefficients by L. This method requires the maximum coefficient amplitude to be less than 1/L, otherwise the multiplication overflows the 16-bit coefficient word length. If the maximum coefficient amplitude is not less than 1/L, then you must multiply each output result by 1/L instead. The code in Listing 4 performs the 16-by-32 bit multiplication needed for this gain correction. The MY1 register should be initialized to L at the start of the routine, and the last multiply/accumulate of the filter should be performed with (SS) format, not the rounding option. This code multiplies a filter output sample in 1.31 format by the gain L, in 16.0 format, and produces in a 1.15 format corrected output in the SR0 register.

```
{INTERPOLATE.dsp
Real time Direct Form FIR Filter, N taps, uses an efficient algorithm to interpolate
by L for an increase of L times the input sample rate. A restriction on the number
of taps is that N/L be an integer.

    INPUT: adc
    OUTPUT: dac
}
.MODULE/RAM/ABS=0 interpolate;
.CONST          N=30;
.CONST          L=4;                        {interpolate by factor of L}
.CONST          NoverL=75;
.VAR/PM/RAM/CIRC coef[N];
.VAR/DM/RAM/CIRC data[NoverL];
.VAR/DM/RAM     counter;
.PORT           adc;
.PORT           dac;
.INIT           coef:<coef.dat>;
                RTI;                        {interrupt 0}
                RTI;                        {interrupt 1}
                RTI;                        {interrupt 1}
    JUMP sample;                            {interrupt 3 at (L*input rate)}
initialize:     IMASK=b#0000;               {disable all interrupts}
                ICNTL=b#01111;              {edge sensitive interrupts}
                SI=1;                       {set interpolate counter to 1}
                DM(counter)=SI;             {for first data sample}
                I4=^coef;                   {setup a circular buffer in PM}
                L4=%coef;
                M4=L;                       {modifier for coef is L}
                M5=-1;                      {modifier to shift coef back -1}
                I0=^data;                   {setup a circular buffer in DM}
                L0=%data;
                M0=1;
                IMASK=b#1000;               {enable interrupt 3}
wait_interrupt: JUMP wait_interrupt;        {infinite wait loop}
{_____Interpolate_____}
sample:         MODIFY(I4,M5);              {shifts coef pointer back by -1}
                AY0=DM(counter);
                AR=AY0-1;                   {decrement and update counter}
                DM(counter)=AR;
                IF NE JUMP do_fir;          {test and input if L times}
{_____input data sample, code executed at the sample rate_____}
do_input:       AY0=DM(adc);                {input data sample}
                DM(I0,M0)=AY0;              {update delay line with newest}
                MODIFY(I4,M4);              {shifts coef pointer up by L}
                DM(counter)=M4;             {reset counter to L}
{_____filter pass, occurs at L times the input sample rate_____}
do_fir:         CNTR=NoverL - 1;            {N/L-1 since round on last tap}
                MR=0, MX0=DM(I0,M0), MY0=PM(I4,M4);
                DO taploop UNTIL CE;        {N/L-1 taps of FIR}
taploop:          MR=MR+MX0*MY0(SS), MX0=DM(I0,M0), MY0=PM(I4,M4);
                MR=MR+MX0*MY0(RND);         {last tap with round}
                IF MV SAT MR;               {saturate result if overflowed}
                DM(dac)=MR1;                {output sample}
                RTI;
.ENDMOD;
```

Listing 3. Efficient Interpolation Filter

```
MX1= MR1;
MR= MR0*MY1 (UU);
MR0= MR1;
MR1= MR2;
MR= MR+MX1*MY1 (SU);
SR= LSHIFT MR0 BY -1 (LO);
SR= SR OR ASHIFT MR1 BY -1 (HI);
```

Listing 4. Extended Precision Multiply

Interpolator Hardware Configuration

The I/O hardware required for the interpolation filter is the same as that for the decimation filter with the exception that the interval timer clocks the output D/A converter, and the input A/D converter is clocked by the interval counter signal divided by L. The interval timer interrupts the ADSP-2100 at the output sample rate. This configuration is shown in Figure 7.

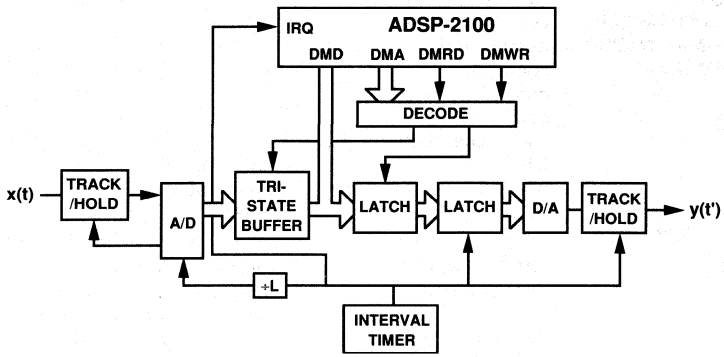

Figure 7. Interpolation Filter Hardware

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Audio Applications of the ADSP Family (IIR Filters)

INFINITE IMPULSE RESPONSE (IIR) FILTERS

Compared to the FIR filter, an IIR filter can often be much more efficient in terms of attaining certain performance characteristics with a given filter order. This is because the IIR filter incorporates feedback and is capable of realizing both poles and zeros of a system transfer function, whereas the FIR filter is only capable of realizing the zeros (although the FIR filter is still more desirable in many applications, because of features such as stability and the ability to realize exactly linear phase responses).

Direct Form IIR Filter

The IIR filter can realize both the poles and zeros of a system because it has a rational transfer function, described by polynomials in z in both the numerator and the denominator:

$$H(z) = \frac{\sum_{k=0}^{M} b_k z^{-k}}{1 - \sum_{k=1}^{N} a_k z^{-k}}$$

The difference equation for such a system is described by the following:

$$y(n) = \sum_{k=0}^{M} b_k x(n-k) + \sum_{k=1}^{N} a_k y(n-k)$$

In most applications, the order of the two polynomials M and N are the same.

The roots of the denominator determine the pole locations of the filter, and the roots of the numerator determine the zero locations. There are, of course, several means of implementing the above transfer function with an IIR filter structure. The "direct form" structure presented in Listing 1 implements the difference equation above.

Note that there is a single delay line buffer for the recursive and nonrecursive portions of the filter (Oppenheim and Schafer's Direct Form II). The sum-of-products of the a values and the delay line values are first computed, followed by the sum-of-products of the b values and the delay line values.

```
.MODULE  diriir_sub;

{
  Direct Form II IIR Filter Subroutine

  Calling Parameters
          MR1 = Input sample (x[n])
          MR0 = 0
          I0 → Delay line buffer current location (x[n–1])
          L0 = Filter length
          I5 → Feedback coefficients (a[1], a[2], . . . a[N])
          L5 = Filter length – 1
          I6 → Feedforward coefficients (b[0], b[1], . . . b[N])
          L6 = Filter length
          M0 = 0
          M1,M4 = 1
          CNTR = Filter length – 2
          AX0 = Filter length – 1

  Return Values
          MR1 = output sample (y[n])
          I0 → delay line current location (x[n–1])
          I5 → feedback coefficients
          I6 → feedforward coefficients

  Altered Registers
          MX0,MY0,MR

  Computation Time
          ((N – 2) + (N – 1)) + 10 + 4 cycles (N = M = Filter order)

  All coefficients and data values are assumed to be in 1.15 format.

}

.ENTRY diriir;

diriir:     MX0=DM(I0,M1), MY0=PM(I5,M4);
            DO poleloop UNTIL CE;
poleloop:   MR=MR+MX0*MY0(SS), MX0=DM(I0,M1), MY0=PM(I5,M4);
            MR=MR+MX0*MY0(RND);
            CNTR=AX0;
            DM(I0,M0)=MR1;
            MR=0, MX0=DM(I0,M1), MY0=PM(I6,M4);
            DO zeroloop UNTIL CE;

zeroloop:   MR=MR+MX0*MY0(SS), MX0=DM(I0,M1), MY0=PM(I6,M4);
            MR=MR+MX0*MY0(RND);
            MODIFY (I0,M2);
            RTS;

.ENDMOD;
```

Listing 1. Direct Form IIR Filter

9

Cascaded Biquad IIR Filter

A second-order biquad IIR filter section is shown on Figure 1. Its transfer function in the z-domain is:

$$H(z) = Y(z)/X(z) = (B_0 + B_1 z^{-1} + B_2 z^{-2})/(1 + A_1 z^{-1} + A_2 z^{-2})$$

where A_1, A_2, B_0, B_1 and B_2 are coefficients that determine the desired impulse response of the system H(z). Furthermore, the corresponding difference equation for a biquad section is:

$$Y(n) = B_0X(n) + B_1X(n-1) + B_2X(n-2) - A_1Y(n-1) - A_2Y(n-2)$$

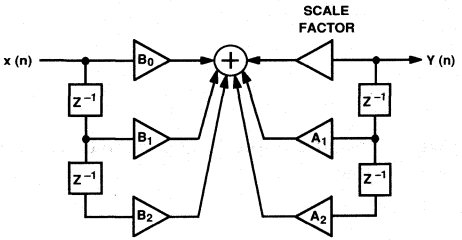

Figure 1. Second-Order Biquad IIR Filter Section

Higher-order filters can be obtained by cascading several biquad sections with appropriate coefficients. Another way to design higher-order filters is to use only one complicated single section. This approach is called the direct form implementation. The biquad implementation executes slower but generates smaller numerical errors than the direct form implementation. The biquads can be scaled separately and then cascaded to minimize the coefficient quantization and the recursive accumulation errors. The coefficients and data in the direct form implementation must be scaled all at once, which gives rise to larger errors. Another disadvantage of the direct form implementation is that the poles of such single-stage high-order polynomials get increasingly sensitive to quantization errors. The second-order polynomial sections (i.e., biquads) are less sensitive to quantization effects.

An ADSP-2100 subroutine that implements a high-order filter is shown in Listing 2. The subroutine is arranged as a module and is labeled *biquad_sub*. There are a number of registers that need to be initialized in order to execute this subroutine. It may be sufficient to do this initialization only once (e.g., at power-up) if other executed algorithms do not need these registers. In most typical cases, however, some of these registers may need to be set every time the *biquad_sub* routine is called. It may sometimes be beneficial, from a modular software point of view, always to initialize all the setup registers as a part of this subroutine.

The *biquad_sub* routine takes its input from the SR1 register. This register must contain the 16-bit input X(n). X(n) is assumed to be already computed before this subroutine is called. The output of the filter is also made available in the SR1 register.

After the initial design of a high-order filter, all coefficients must be scaled down separately in each biquad stage. This is necessary in order to conform to the 16-bit fixed-point fractional number format as well as to ensure that overflows will not occur in the final multiply-accumulate operations in each stage. The scaled-down coefficients are the ones that get stored in the processor's memory. The operations in each biquad are performed with scaled data and coefficients and are eventually scaled up before being output to the next one. The choice of a proper scaling factor depends greatly on the design objectives, and in some cases it may even be unnecessary. The filter coefficients are usually designed with a commercial software package in higher than 16-bit precision arithmetic. System performance deviates from ideal when such high precision coefficients are quantized to 16 bits and further scaled down. In systems that require stringent specifications, careful simulations of quantization and scaling effects must be performed.

During the initialization of the *biquad_sub* routine, the index register I0 points to the data memory buffer that contains the previous error inputs and the previous biquad section outputs. This buffer must be initialized to zero at powerup unless some nonzero initial condition is desired. The index register I1 points to another buffer in data memory that contains the individual scale factors for each biquad. The buffer length register L1 is set to zero if the filter has only one biquad section. In the case of multiple biquads, L1 is initialized with the number of biquad sections. The index register I4, on the other hand, points to the circular program memory buffer that contains the scaled biquad coefficients. These coefficients are stored in the order: B_2, B_1, B_0, A_2 and A_1 for each biquad. All of the individual biquad coefficient groups must be stored in the same order in which the biquads are cascaded, such as: B_2, B_1, B_0, A_2, A_1, B_2*, B_1*, B_0*, A_2*, A_1*, B_2**, etc. The buffer length register L4 must be set to the value of (5 × number of biquad sections). Finally, the loop counter register CNTR must be set to the number of biquad sections, since the filter code will be executed as a loop.

The core of the *biquad_sub* routine starts its execution at the *biquad* label. The routine is organized in a looped fashion where the end of the loop is the instruction labeled *sections*. Each iteration of the loop executes the computations for one biquad. The number of loops to be executed is determined by the CNTR register contents. The SE register is loaded with the appropriate scaling factor for the particular biquad at the beginning of each loop iteration. After this operation, the coefficients and the data values are fetched from memory in the sequence in which they have been stored. These numbers are multiplied and accumulated until all of the values for a particular biquad have been accessed. The result of the last multiply/accumulate is rounded to 16 bits and upshifted by the scaling value. At this point, the *biquad* loop is executed again, or the filter computations are completed by doing the final update to the delay line.

The delay lines for data values are always being updated within the *biquad* loop as well as outside of it.

The filter coefficients must be scaled appropriately so that no overflows occur after the upshifting operation between the biquads. If this is not ensured by design, it may be necessary to include some overflow checking between the biquads.

The execution time for an Nth order *biquad_sub* routine can be calculated as follows (assuming that the appropriate registers have been initialized and N is a power of 2):

ADSP-2101/2102: $(8 \times N/2) + 4$ processor cycles
ADSP-2100/2100A: $(8 \times N/2) + 4 + 5$ processor cycles

It may take up to a maximum of 12 cycles to initialize the appropriate registers every time the filter is called, but typically this number will be lower.

```
.MODULE biquad_sub;

{ Nth order cascaded biquad filter subroutine

  Calling Parameters:

        SR1 = input X(n)
        I0 → delay line buffer for X(n-2), X(n–1),
             Y(n-2), Y(n-1)
        L0 = 0
        I1 → scaling factors for each biquad section
        L1 = 0 (in the case of a single biquad)
        L1 = number of biquad sections
             (for multiple biquads)
        I4 → scaled biquad coefficients
        L4 = 5 × [number of biquads]
        M0, M4 = 1
        M1 = –3
        M2 = 1 (in the case of multiple biquads)
        M2 = 0 (in the case of a single biquad)
        M3 = (1 – length of delay line buffer)

  Return Value:
        SR1 = output sample Y(n)

  Altered Registers:
        SE, MX0, MX1, MY0, MR, SR

  Computation Time (with N even):
        ADSP-2101/2102: (8 × N/2) + 5 cycles
        ADSP-2100/2100A: (8 × N/2) + 5 + 5 cycles

  All coefficients and data values are assumed to be in 1.15 format

}

.ENTRY biquad;

biquad: CNTR = number_of_biquads
        DO sections UNTIL CE;
          SE=DM(I1,M2);
          MX0=DM(I0,M0), MY0=PM(I4,M4);
          MR=MX0*MY0(SS), MX1=DM(I0,M0), MY0=PM(I4,M4);
          MR=MR+MX1*MY0(SS), MY0=PM(I4,M4);
          MR=MR+SR1*MY0(SS), MX0=DM(I0,M0), MY0=PM(I4,M4);
          MR=MR+MX0*MY0(SS), MX0=DM(I0,M1), MY0=PM(I4,M4);
          DM(I0,M0)=MX1, MR=MR+MX0*MY0(RND);
sections: DM(I0,M0)=SR1, SR=ASHIFT MR1 (HI);
        DM(I0,M0)=MX0;
        DM(I0,M3)=SR1;
        RTS;

.ENDMOD;
```

Listing 2. Cascaded Biquad IIR Filter

ANALOG DEVICES

AN-336A
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Image Compression: Spelling Out the Options

by Chris Cavigioli

Been confused about the approaches being developed for this vital technology? This could help.

Image data compression can be used to reduce channel bandwidth requirements when transmitting still or moving images over a band-limited channel, and enables images to be stored in a smaller memory space. These are among the most important requirements for furthering imaging in the immediate future. Compression is a prime subject of discussion in electronic publishing, videoteleconferencing, medical image sharing, real-time missle guidance by visual control and space exploration—or in looking at the future of color copiers, still video cameras, ISDN networks (voice, video, text documents and images), video channels on commercial broadcast satellites, image databases, digital office FAX and PC FAX boards.

There exist many approaches to reducing image data; the end application dictates which technique is best suited for compressing data. For example, some techniques are *lossy* techniques while others are *lossless*. Medical images, for example, often require completely lossless compression because the exact value of each pixel could be significantly important in a mathematical computation or in a legal dispute. Browsing an ID card database, or guiding a missile in the battlefield, on the other hand, could be carried out without having to know the exact value of each pixel. The human visual system is adaptive enough to be "fooled," seeing a lossy-compressed image and confusing it—accepting it—for the original. In fact, most pictures can be processed to remove 97% of the original data, and the human eye would have a very hard time at noticing any degradation.

The compression techniques

Compression techniques include: transform-based coding, sub-band coding, entropy coding, run-length coding, vector quantization, predictive coding, block-truncation coding, and coding by modeling a scene. The latter has recently become quite exciting, as ways of applying completely new approaches involving fractals are being discovered. Since there are so many different compression techniques available, and since certain compression applications actually employ several of those techniques in combination, I will try to navigate through several techniques, discussing each along the way.

Transform coding applies a mathematical transform to the pixel data to put the resulting data values into another domain. A transform is chosen which consolidates pixel information into a more compact form. Once in this form, other techniques are employed to compress the information even further. For example, in a lossy system, thresholding and quantization can be used to select the more significant values from the less significant ones, and then transmit only the most important ones. For lossless systems, thresholding and quantization wouldn't be allowed. In either case, entropy coding (such as Huffman coding) and run-length encoding can follow to reduce the bit rate even further. The most popular transform today would be the discrete cosine transform (DCT). In the past, before high performance, programmable DSPs were available, the Hadamard transform enjoyed some usage. The DCT is relatively compute-intensive, requiring fast multiplies and sophisticated data addressing. The Hadamard transform was initially popular since it can be computed with simple ±1 operators, thereby eliminating the need for hardware multiplication.

Run-Length coding capitalizes on the high pixel-to-pixel correlation in the image, and then recodes the data to eliminate the redundancies. To achieve a high degree of compression, the run-length data is recoded into digital codes of varying length, with shorter codes assigned to the more likely original code values (see Figure 1). The common facsimile machine uses the CCITT Group 3 standard to first determine the length of a continuous run of white (or black) image pixels, and then assigns (via a modified Huffman code table) a corresponding code between 2 and 13 bits long.

Sub-band coding passes an image through a bank of special filters called quadrature mirror filters (QMFs). A QMF filter has the special property of reducing the input sampling rate by two and splitting the input bandwidth into two output channels, each having half of the original bandwidth. For example, say

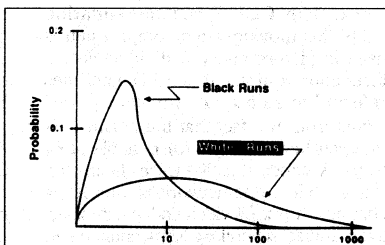

Image-Pixel Run-length	Corresponding Modified Huffman Codes: white runs	black runs
0	00110101	0000110111
1	000111	010
2	0111	11
3	1000	10
4	1011	011
5	1100	0011
6	1110	0010
7	1111	00011
8	10011	000101
9	10100	000100
10	00111	0000100
11	01000	0000101
12	001000	0000111
13	000011	00000100

Fig. 1. Run length coding.

Reprinted from Advanced Imaging – October 1990

we have a signal which is sampled at 100 KHz: the input bandwidth spans 0 Hz to 50 KHz. By passing this signal through a quadrature mirror filter, two output channels result—one has all the 0 Hz to 25 KHz information while the other has the 25 KHz to 50 KHz information. Both output channels are each sampled at 50 KHz (half of the original sampling rate). After several stages of QMF filtering, an image is separated into several smaller images, each image holding data within a known spatial frequency band. The lower frequency imformation is the most important information, while the high frequencies hold the crispy details of a picture. By selectively transmitting only some of the lower frequency portions, data compression is achieved.

Progressive coding and JPEG

Sub-band coding is ideally suited for a technique called *progressive coding*. Progressive coding is especially useful for archival database browsing. When browsing through a database, only the lowest frequency image portion are first retrieved: the viewer can get a good idea of the image. The longer the viewer observes the image before going on to the next, the more detail is added to the picture. This is done by progressively adding the next higher frequency portions until all the original data is recreated.

Progressive coding is also possible with transform-based techniques such as the discrete cosine transform. For example, the much-discussed JPEG still-image compression standard developed by the International Standards Organization (ISO) specifies a sequential DCT technique as its absolute minimum (baseline): this means that the algorithm performs its compression tasks on the whole input image in a sequential fashion. First an image is subdivided into smaller blocks; then the blocks are transformed using the DCT; then the DCT coefficients are thresholded and quantized; and finally they are entropy coded. After all the tasks are done, the data can then be stored or transmitted. On the receiving side, the compressed data is received and sequentially decompressed. When that is all done, the decompressed image data is finally displayed on the monitor.

JPEG permits the offering of more advanced features above the baseline requirement—one of these features is a progressively-coded DCT implementation. The user sees a rough image immediately, which is then progressively improved in quality. Progressive DCT coding in JPEG can be achieved in one of two ways: *spectral selection*, where selection of DCT coefficients to transmit is based on their associated frequency; and *successive approximation*, a method whereby an averaged approximation of the DCT coefficients is sent first, followed by successively better correction values sent later to improve the image.

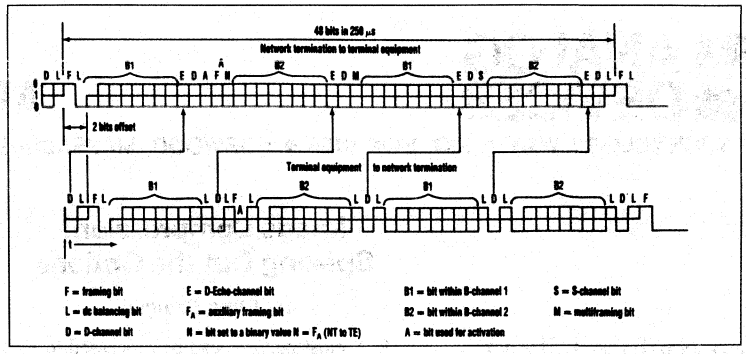

Fig. 2. Basic Access ISDN.

While progressive coding usually means that a larger bandwidth requirement is sacrificed for a longer transmit duration on a narrow bandwidth, *flexible bandwidth coding* can be used to describe a scheme in which the channel bandwidth itself can be altered in order to accomodate various picture quality requirements in real-time.

Phone transmission

A good example of this would be videotelephony on the telephone network. The CCITT recommendation H.261 for moving video compression is used in picture phone and videoteleconferencing equipment and is sometimes referred to as *p x 64*. The name p x 64 stems from the fact that the channel bit rate can be an integer (p) multiple of 64 Kbps. A standard digital voice channel on today's telephone network is a 64 Kbps channel. A p x 64 video codec can selectively utilize several 64 Kbps channels to achieve the desired picture quality. H.261 was originally drafted for video telecommunications over digital telephone lines having multiple channels multiplexed in time. It should also be mentioned that one of the 64 Kbps channels is reserved for the audio portion of the telephone call.

The most common *digital carriers* for H.261 would be the standard T1 carrier (North America) or the CEPT carrier (Europe): both of these carriers are refered to as the DS1 (digital service 1) layer. The phone line which comes to your house carries one voice channel at 64 Kbps—this is the DS0 layer. At the central office, several DS0 channels are time-multiplexed (TDMA) onto the DS1 layer. In North America, 24 channels are combined into a frame on the DS1 layer; this protocol yields a total bit rate of 24 x 64 Kbps or 1.536 Mbps and is called T1. In Europe, 32 DS0 channels are framed onto Europe's DS1 layer; this protocol is called CEPT and yields 32 x 64 Kbps or 2.048 Mbps.

Four *enabling factors* have emerged over the last few years to make videotelephony a reality: plummetting cost of DS1 services and associated tariffs, standardization of the H.261 recommendation by the CCITT (June 1990), DSP chips enabling affordable implementation of videotelephony products, and live demonstrations of working products by at least three firms in the U.S.A.

To make videotelephony even more attractive, the telephone company now also sells new digital services to the subscriber including: fractional T1, switched 56, and ISDN. *Fractional T1* implies that instead of purchasing all 24 channels of T1 service, it is now possible to purchase, for example, only 12 of the channels at a reduced cost. *Switched 56* is the common name for dial-up lines which carry a single channel of digital data at 56 Kbps. In this case, you would have to purchase one line for video, albeit at a reduced picture quality, and one line for the audio portion. The price of switched 56 is the factor that is making it immensely popular.

ISDN is a bit slower in being widely accepted in the U.S.A., but its time is definitely on the way. *Basic Access ISDN* consists of two independent B channels together with a D channel (see Figure 2). The B channels are 64 Kbps and the D channel is only 16 Kbps for basic access; *Primary Access ISDN* is at the DS1 level. In this case, there are 23 B channels (North America) in conjunction with a 64 Kbps D channel. In Europe, there are 30 B channels with a 64 Kbps D channel.

Images into the stream

With that said, let's look at how a picture is taken from a camera and compressed into a low data-rate bit stream. To get a better idea of the image compression task, consider the magnitude of the data reduction required. In the H.261 recommendation, the input picture frame size can be in one of two spatial resolutions: there is the *CIF* (Common Intermediate Format) which has a resolution of 360x288 pixels, and there is the *QCIF* (Quarter Common Intermediate Format)

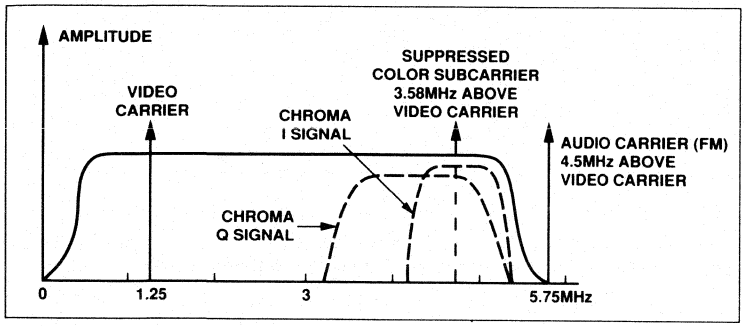

Fig. 3. Broadcast television signal.

which has a resolution of 180x144 pixels. Two formats enable basic units to support QCIF for a lower price, while higher-priced units support both formats. For a moving picture, the frame is updated many times every second: full motion is regarded as 30 frames per second. Since this represents a very large amount of data, the input data is often temporally subsampled—subsampling at a rate of 2, 3, or 4 yields a frame rate of 15, 10 or 7.5 frames per second respectively.

The number of bits per pixel depends upon whether each frame is color or not. The data format closely resembles that of standard television. In other words, instead of one frame each of red, green and blue as found in your computer CRT, television signals are based on luminance (the gray scale or black-and-white portion) information and chrominance (the color) information. (This was developed historically so that black-and-white television and color television could coexist.) Either receiver demodulates the luminance information, and color TVs have additional circuitry to demodulate the chrominance information. Chrominance comes in two parts which are quadrature modulated; television signals are sometimes refered to as being in YIQ format. Y stands for luminance, and I and Q are the two chrominance components (see Figure 3). Since the human eye is more responsive to luminance degradation versus chrominance degradation, the chrominance is often spatially subsampled. The CIF and QCIF formats subsample chrominance at a 2:1 rate.

Video signals are sampled at 8 bits resolution. Consequently, a full-motion, color, CIF bitstream runs at 37.3 Mbps! All this data must somehow be shoehorned into p x 64 Kbps—this formidable task can be achieved with the CCITT H.261 compression algorithm. To help reduce the data to be compressed, provisions for motion compensation have been included in the algorithm.

Motion compensation is a technique whereby a frame is compared to the previous frame (in time) and the movement of objects in the scene is estimated. Stationary objects are not required to be recoded and retransmitted since the receiver already has these objects in its frame memory from previous frames. Using this technique, it becomes possible to code a frame, then for several next frames, simply code and transmit the difference information to the receiver instead of completely new frames. Motion compensation aids tremendously in cases such videoteleconferencing where a subject's head and lip movements are practically the only moving objects in the frame sequence.

Motion compensation enables large data reduction performance at a very slight expense of increased control signal requirements. There are times when motion compensation will work and times when it cannot. Signalling bits are specified in H.261, informing the receiver whether it should be in *interframe* mode versus *intraframe* mode. Interframe mode indicates that data being sent is frame difference data which should be used to update the receiver current frame store. Intraframe mode means that the receiver should throw away its stored frame and start over from scratch with a completely new picture. Intraframe mode is applicable anytime that the interframe differences are too great and a fresh frame would be better, for instance, when the very first picture is being sent, a scene cut has happened, or too much time has elapsed since the last complete frame was sent.

We will look at working examples of practical methods of image compression in a future issue of this magazine— including a close look at discrete cosine transform-based compression employing DSP hardware.

ANALOG DEVICES

AN-336B
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

JPEG Compression

by Chris Cavigioli

Software-based? Or which hardware type? Lossy or lossless? More mysteries clarified

In part one of this image compression overview, ("Image Compression: Spelling Out the Options," *Advanced Imaging*, Oct. '90) we addressed the very general issue of compressing image data, looking at options for image transmission over limited bandwidth channels.

In this second part, we focus on the compression of continuous tone images using the much-discussed, sometimes understood ISO/JPEG compression algorithms. (JPEG, incidentally, stands for Joint Photographic Experts Group: this is a committee which is proposing a standard for compressing high-quality, still images. The standardization activity (e.g., JPEG) throughout the U.S. is coordinated by ANSI (American National Standards Institute), and ANSI, in turn, casts the U.S. vote at the International Standards Organization (ISO) meetings.)

Our focus is on two questions: Which JPEG variant should be used to compress a given image, and should compression be done with PC software, a JPEG VLSI chip set, or a programmable DSP?

Which JPEG?

JPEG is not a single, fixed algorithm, but rather a family of several techniques, each member best suited for certain applications. All JPEG coders must support the baseline technique, with the exception of lossless JPEG coders, which are not required to support baseline. Some JPEG coders will support extended systems which are above and beyond the fundamental baseline technique. Fig. 1 shows the relationship among all current JPEG algorithms.

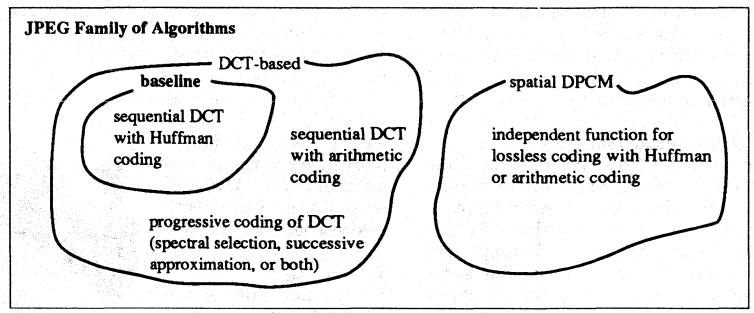

Fig. 1. The Family of JPEG algorithms.

JPEG compression algorithms have been tested on images of natural scenes. These images averaged 16 bits/pixel and were stored in the CCIR-601 4:2:2 (Y:Cb:Cr) format. Table 1 shows experimental compression ratios achieved and the resulting image quality.

Lossy, high compression techniques are DCT-based (Discrete Cosine Transform), while the lossless but low compression techniques are DPCM-based (Differential Pulse Code Modulation). Within the DCT-based realm, there is the baseline mode which all DCT-based JPEG coders must support as well as the extended systems. The full name for the baseline system is: *sequential DCT system with Huffman coding*. Let's see how this compares with other JPEG algorithms.

Table 1: JPEG Compression Ratios

Bits/Pixel	Image Quality	Compression Ratio
0.1	Recognizable image	160 : 1
0.25	Useful image	64 : 1
0.75	Excellent quality	22 : 1
1.5	Indistinguishable from original	11 : 1
8	Lossless JPEG method	2 : 1

Progressive vs. sequential coding

Sequential and progressive coding refer to the order in which the compressed information is sent from encoder to decoder and subsequently displayed. Sequential systems begin upper left-hand corner of the image, then proceed to completely encode the image from left to right across the rows, and top to bottom down the columns. When you observe the decoder reconstructing the image, you see the final image appearing in the same order, analogous to covering a picture by a piece of paper and slowly revealing the picture by sliding the paper away.

With the same decoder in progressive mode, a coarse approximation of the image instead fills the screen right away, giving you a good idea of what the final image looks like. The image progressively improves in quality until it is indistinguishable from the sequentially-generated image. The progressive technique would be better suited for applications such as image database browsing, while the sequential technique is practically mandatory in other applications, such as hardcopy printing.

Progressive coding choices

Two progressive coding techniques, spectral selection and successive approximation, are defined by JPEG. The DCT transforms spatial pixel data into frequency data. By picking only certain frequency bands of interest (by choosing only some of the DCT coefficients), it is possible to approximate an image by first using only low frequency components, then add detail by sending some missing high frequency components. When all the coefficients have been sent, the final picture results. This is the spectral selection technique.

Reprinted from Advanced Imaging – March 1991

The successive approximation technique sends the approximate value of the DCT coefficients by truncating their LSBs (least significant bits) in the first stage. The missing low order magnitude bits are sent in the later stages, one bit plane at a time. Hybrid algorithms employing a combination of both techniques are allowed.

All of the images in the examples *(at right)* are JPEG compressed at 15:1. There are three images in each compression mode, giving 10%, 80%, and 100% of complete image transmission. One set shows baseline mode, the second set, progressive mode using spectral selection only, and the final set, progressive mode using both spectral selection and successive approximation.

Entropy coding issues

Some of the JPEG extended systems replace the Huffman coding technique of the baseline system with adaptive binary arithmetic coding. Either method is a lossless way of entropy coding the quantized DCT coefficients. Huffman coding replaces the quantized coefficients with variable length codewords (VLCs). The coefficient is the look-up index into a table, and the VLC extracted from that location is the data sent. The codeword length is shortest for statistically more frequent coefficient values and longer for less frequent coefficients, so the total amount of data transmission is reduced. (Morse code is an example of Huffman-like coding for the characters in the English language.) Obviously, different images possess different statistical distributions of DCT coefficients.

For this reason, in addition to the baseline default tables, JPEG allows custom Huffman tables to be generated from optional signaling information. Custom tables can be used for certain classes of images or even specifically optimized for each individual image. Adaptive binary arithmetic coding provides adaptive entropy coding in a single pass, eliminating the need for custom tables used in the Huffman technique, by dynamically adapting its set of probability estimates to each image or even to regions within an image.

Independent function for lossless

A separate, independent JPEG function (not associated with DCT-based techniques) is specified by JPEG for sequential, lossless coding. This algorithm is a spatial algorithm in the pixel domain as opposed to the transform domain. It is based on the DPCM coding model originally developed for the DC coefficients of the DCT in the DCT-based systems. However, the model is extended by incorporating two-dimensional prediction.

Lossless compression cannot achieve high compression ratios like the lossy ones can—but lossless compression may be an end-application requirement. For example, medical images usually are not allowed to be lossy compressed. Other images where the pixel data is qualitatively important may be scientific data, remote sensed or computer-generated visualization images.

JPEG alternatives: 10%, 80% and 100% transmissions (left to right) using baseline (top), spectral selection only (middle), and both spectral selection and successive approximation (bottom). Note missing stripes on her clothes when transmission is only 80% complete.

Which way?

After deciding which JPEG algorithm is needed, the final issue is implementation—by what system should the JPEG compression algorithm be carried out?

CPUs and RISCs: Compressing images with software programs executing on standard computational platforms such as a PC, for example, has advantages. Software is available today, and such software can be modified, ported, or embedded into other products. Execution speed is the major drawback of this approach. Today's 80286 processors or even RISC chips cannot do the calculations fast enough for fast compression, let alone real-time compression. Using general purpose CPU software for JPEG compression only makes sense when compression requests are infrequent.

Building blocks and chipsets: ICs are available which perform computationally intensive tasks very quickly. These ICs—chip sets—can be connected together to implement JPEG compression on a circuit board in hardware. Examples of these building blocks would be 8x8 DCT chips, motion vector estimation chips, and now even some quantization and entropy coding accelerators. Both SGS-Thompson/INMOS and LSI Logic offer these types of chips. The advantage to using these chips is primarily their speed. Typically, you would select one vendor and use that chipset; mixing vendors may be difficult if the timing and other signals between chips are incompatible.

The disadvantages of chip sets are high cost of the completed implementation and relatively little flexibility. Flexibility is extremely important as users demand

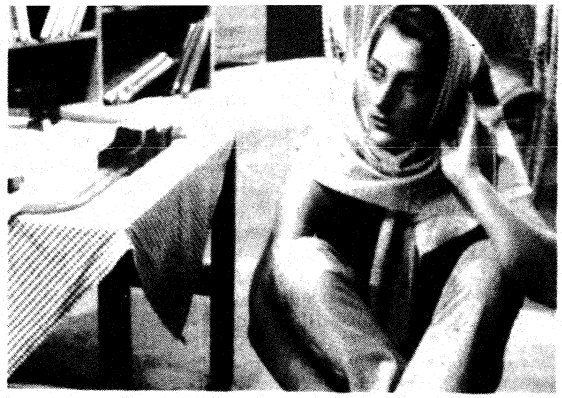

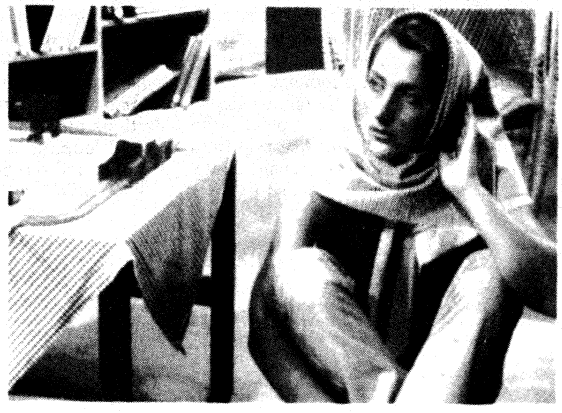
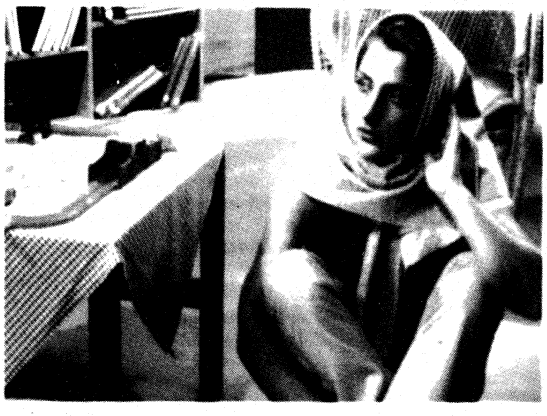

dynamic solutions and as systems migrate towards multimedia capabilities. For example, today's chipsets are fine for sequential DCT systems with Huffman coding (if the final standard is identical to the current version), but what happens if the end-application requires arithmetic coding, or lossless compression, or maybe even simply progressive coding?

JPEG VLSI chips: Single-chip VLSI JPEG coders, such as the CL550A from C-Cube, combine the separate building block ICs of chip sets into a single IC, reducing cost, power, and complexity. The drawback for the end-user is that this is the most rigid and inflexible solution. All the inflexibility of chip sets exist with single-chip solutions, but more severely, since the "chip set" has been already arranged for you in VLSI and there is no access to the internal architecture for modifications, no reprogrammability in a turnkey, baseline JPEG chip.

Case for programmable DSPs

For many applications, the best balance between cost, flexibility, and processing throughput is achieved by

reprogrammable digital signal processors. DSPs are optimized for signal processing computations, much more so than even the fastest RISC chips. Because DSPs are general purpose, cost is low due to the large volume of DSP chips manufactured. Development time is sharply reduced because of mature and powerful development aids. DSPs are available from Analog Devices, Texas Instruments, Motorola, and AT&T, for example.

Flexibility is their outstanding feature: the same DSP can perform JPEG compression at one instant, digital audio compression at another, even handle incoming telephone calls upon interrupt control, and interface to UARTs, MIDI controllers, as examples. However, DSPs are probably not the best choice in rigid, single-application, embedded systems where JPEG is the sole task to be performed at real-time, because those applications do not require the DSP flexibility and probably could not afford the slower DSP throughput.

DSPs today are available on bus plug-in boards. Future personal computers with multimedia options will use powerful DSPs on the motherboard, alongside the CPU, to perform various signal processing tasks as they are developed. Ironically, much of the progress in multimedia applications will be made by developing the algorithms right on the personal computer—by bringing this capability to the general engineering public.

By choosing a DSP in a related family, based on a powerful and flexible internal open architecture, code compatibility is assured. This allows an upgrade path and portable software without risking the investment already made in code development. An example of such a family is the ADSP-2100 family of DSPs from Analog Devices, Inc. with a variety of memory, I/O and packaging options, but with a unified, common register set and basic architecture.

For example, due to their algebraic syntax instruction set, assembly language programming on these DSPs is much easier than typical "assemblers." Gone are the cryptic mnemonics, operators, and operands associated with early (and many present) DSPs. Features of these DSPs which are needed for DCT computations include the nestable zero-overhead looping, the independent barrel shifter, and the data address generators (DAGs) with bit-reverse capability. For zigzag scanning of the DCT coefficients, the DAGs are indispensable, because of the ability to mix-and-match index registers with modify registers, allowing single-cycle access in a zigzag fashion. With each access, the index pointer is updated by one of four available modify registers, depending on which direction the zigzag scan proceeds. The independent barrel shifter (capable of N-bit left/right shifts in a single cycle) is especially powerful in variable length entropy coding (VLCs, Huffman, arithmetic).

With compression algorithms and associated ICs proving the enabling technology at affordable prices, a major revolution in imaging storage and transmission is imminent. Until now, the availability of moving color pictures with sound has been limited to non-computer electronics such as television. Providing this capability to the user at the computer interface will significantly reduce the barrier between computers and the general public.

Our appreciation to Jørgen Vaaben at autoGraph International A/S, Lyngby, Denmark, for the images.

ANALOG DEVICES

AN-227
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Digital Control System Design with the ADSP-2100 Family

by Kapriel Karagozyan

INTRODUCTION

The ADSP-2100 Family of Digital Signal Processors is well suited for implementing complex measurement and control algorithms in embedded control systems with high sampling rates. This is mainly due to their computing speed, which is much greater than that of conventional microcontrollers and microprocessors. Typical application areas include servo motor control, process control, robot arm control, disk drive head control, flight control and general servomechanisms.

This application note presents the implementation of several common control algorithms on the ADSP-2100 Family of DSP processors, and presents software and hardware design methods as well as guidelines for designing high speed digital control systems with the ADSP-2100 Family. A table of representative benchmarks for common digital control algorithms can be found at the end of this note.

DIGITAL CONTROL SYSTEMS OVERVIEW

A controller is a system used to control closed-loop feedback systems. It implements algebraic algorithms, such as filters and compensators, in order to regulate, correct, or change the behavior of a controlled system. Controllers can be implemented using analog or digital circuitry. A digital control system is comprised of a digital controller, the controlled plant (or system), and the necessary input/output devices. A general digital control system is shown in Figure 1. Note the analog-to-digital (A/D) and digital-to-analog (D/A) converters that are used to interface the digital controller with the plant (which is a continuous time system). There are several advantages to using a digital controller implementation instead of an analog one. In the case of digital controllers, complex control algorithms can be implemented in software or firmware rather than in special hardware. Digital controller designs and parameters can be changed without affecting the hardware. In digital control systems increased noise immunity is guaranteed and parameter drift is eliminated. Such systems are more reliable, maintainable, and testable. Finally, digital control systems feature reduced size, power, weight and costs.

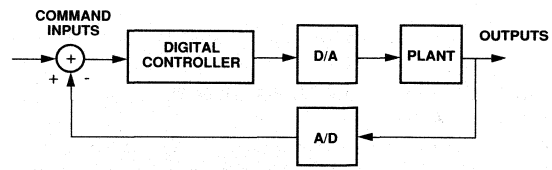

Figure 1. General Digital Control System

Analog Devices' ADSP-2100 Family of Digital Signal Processors has several beneficial features for implementing digital controllers. These features include the following:

- Single-cycle instruction execution
- Three arithmetic function units arranged in parallel
- Single-cycle 16x16-bit multiplications
- Single-cycle ([16x16]+40 bit) multiply-accumulate operations with 40-bit results
- Single-cycle 16-bit additions, subtractions and logical operations
- Single-cycle bit shifts (up to 32 bits at a time)
- Efficient execution of 32-bit (or higher) arithmetic operations
- Efficient modulo addressing for data and coefficient arrays in memory
- No cycle penalty for looped code execution
- Single-cycle access of internal and external memory
- Single or multi-cycle parallel accesses of external peripherals (i.e., A/D, D/A)
- Up to four levels of nested external interrupts
- On-chip interval timer and serial ports
- Low power consumption (CMOS) and power down "idle" mode
- Easy-to-read algebraic assembly language syntax
- Complete set of hardware and software development tools

DIGITAL CONTROL SYSTEM MODEL

Most practical control systems use feedback in their operation. Figure 2 shows a model for a typical closed-loop digital control system. R(z), E(z), U(z) and Y(z) are the z-transforms of the reference input, the error signal, the control signal, and the plant output respectively. G(z) is the transfer function corresponding to the digital controller, while P(z) is the transfer function describing the input-output behavior of the object to be controlled (e.g., plant). This does not imply that the object to be controlled (e.g., a plant) must be a discrete system, but rather that it must be modeled as one. P(z) is also assumed to contain the transfer characteristics of the A/D and the D/A converters that are needed to implement a real system.

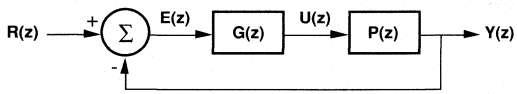

Figure 2. Digital Control System Model

DIGITAL CONTROL SYSTEM HARDWARE IMPLEMENTATION

A digital controller G(z) , as shown in Figure 2, must be able to observe and alter certain characteristics of the controlled system. For example, an ADSP-2100 Family-based digital controller can be used to control fast and accurate positioning of an actuator shaft upon an external command R(z). In this case, the output of the controller can be used to alter the amount of current U(z) that is fed into the actuator windings which in turn would move the actuator shaft. In a closed-loop system, the same controller would also need to observe the position of the actuator at all times. This can be achieved by recording the position Y(z) of the shaft at specific intervals and feeding it back to the controller. This would allow the controller to compare the desired shaft position to the actual measured position and make the necessary adjustments in the actuator current. This simple controller example can serve as a starting point for constructing an actual hardware implementation.

The block diagram for a digital control system based on the ADSP-2102 is shown in Figure 3. The ADSP-2102 performs the digital control algorithms by executing instructions from its on-chip program memory ROM. The ROM is also used to store fixed coefficients and scale factors. The processor uses its on-chip data memory RAM and program memory RAM to store incoming data values and other intermediate variables.

The ADSP-2102 accepts up to three external hardware interrupts. In a typical digital control system, the processor operation is interrupt-driven. In the system shown in Figure 3, an external clock (sample clock) drives one of the ADSP-2102 interrupts. The same clock is typically used to initiate A/D conversions at regular intervals. Other interrupts can be set by the host to notify the ADSP-2102 of new commands, expiration of a watchdog timer, etc.

The ADSP-2102 outputs its control current via a D/A converter, whose output is amplified before it is fed into the motor. The processor receives its feedback from a position encoder which can be an optical shaft encoder, a synchro-to-digital converter, a resolver-to-digital converter, or some other circuitry with an

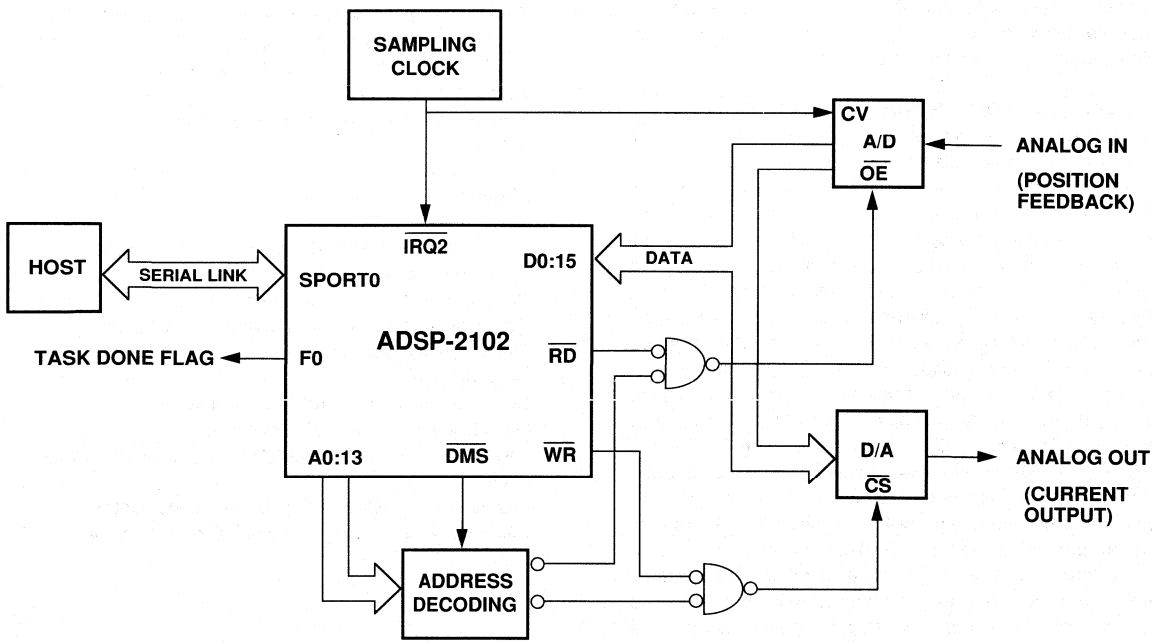

Figure 3. ADSP-2101 Based Actuator Controller

A/D converter. The A/D label is used in the figure since feedback essentially involves an analog-to-digital conversion. Data transfers between the processor and the converters are done over the 16-bit data bus. The converters are mapped into the ADSP-2102's external data memory space. This allows the processor to access them as memory locations. The address decoding circuitry shown in Figure 3 is used to map every converter to a separate data memory location. If the converters have slow data bus interfaces, the processor can extend the duration of the converter access cycles by inserting wait states. The data converters can also be tied to the ADSP-2102's serial ports if it is more convenient to do so. Several serial input and serial output converters are available from Analog Devices, such as the AD766, AD7772, AD7868, and AD7878.

The ADSP-2102 receives its reference position command R(z) from a host processor or an internal software routine that is running concurrently with the shaft positioning program. Figure 3 depicts the case where serial port 0 on the ADSP-2102 is used to exchange commands and results with a host processor. The flag out pin on the ADSP-2102 can be used to notify the host of the completion of a specific task.

The ADSP-2100 Family processors can interface with multiple A/D and D/A converters in order to monitor and control several motors, actuators, or processes. These converters can simply be added as memory-mapped peripherals like the ones shown in Figure 3. Several general purpose and special purpose data converters for digital control applications are available from Analog Devices.

DIGITAL CONTROL SYSTEM SOFTWARE IMPLEMENTATION

The software running in a digital controller system is responsible for executing the control algorithms which are represented by G(z) in the model on Figure 2. Typically, G(z) can be broken into smaller sub-tasks. For example it may be necessary to execute a state estimator, several notch filters and some PID (Proportional, Integral, Derivative) control as a whole function. Generally, a separate portion of the software must manage the input/output operations of the controller with the host and other peripherals. A diagnostic error checking and handling routine is also usually developed, to be run at powerup or at specified intervals during program execution. Finally there is a main manager routine that is responsible for the orchestration of these different subroutines.

The software must be organized in a modular manner in order for the main managing program to call every sub-task as a subroutine. The ADSP-2100 Family Development Software tools encourage modular programming. The subroutines can be written, assembled, and debugged as independent modules which can later be linked with the main manager program. Parameter passing and symbolic coding is supported on the assembler, linker and simulator. An example of a fully coded notch filter algorithm is shown in a later section of this application note. The ADSP-2100 Assembler and Simulator manuals describe the software tools.

Memory management is very straightforward in the ADSP-2100 Family processors. The Data memory (DM) space is typically used for variables and data storage. The incoming A/D samples can be stored in data memory buffers. A large number of variables and intermediate values can also be stored in DM space. The Program memory (PM) space is divided into two sections: the PM instruction space and the PM data space. The instruction space is used to store the programs to be executed.The PM space can also be freely used for additional data and variable storage. This data space is usually used to store filter coefficients and various other tables that may need to be present during program execution. The ADSP-2100 Family processors can read or write to both DM and PM locations in a single instruction cycle and execute an arithmetic operation at the same time. This not only allows classical control algorithms to execute at very high speeds but also allows very efficient implementation of adaptive control algorithms. This is due to the fact that in adaptive control, filter coefficients must be updated periodically with every new incoming sample. These coefficient values can be updated easily in the PM space and can be readily available on the next processing cycle.

The following sections discuss the implementation of first, second, and higher-order control algorithms with the ADSP-2100 Family processors.

DIGITAL PID CONTROLLER DESIGN

The controller G(z) shown in Figure 2 can be designed to vary its output U(z) in relation to the error feedback E(z). A PID (Proportional, Integral, Derivative) controller derives its name from the fact that its output U(z) is a weighted sum of the error signal, its integral, and its derivative. PID controllers are widely-used building blocks in a large variety of servo control applications.

Since analog PID controllers are well understood, it is often desirable to start a digital PID controller design in the continuous domain and then create discrete equivalents. In the continuous time case if E(t) is the error feedback, the PID output U(t) can be expressed as:

$$U(t) = K_p \cdot E(t) + K_d \cdot dE(t)/dt + K_i \cdot \int_0^t E(\tau)d\tau \quad (1)$$

where K_p, K_d, and K_i are the gains associated with the proportional, derivative, and integral terms, respectively. Equation (1) can be represented in the frequency domain by using Laplace transforms:

$$U(s) = K_p \cdot E(s) + K_d \cdot s \cdot E(s) + (K_i/s)\, E(s) \quad (2)$$

where it is assumed that the initial conditions are 0. The equations (1) and (2) are graphically represented in Figure 4.

The coefficients Kp, K_d and K_i must be determined during the design process. These coefficients will depend on the desired controller characteristics and will be varied in different systems.

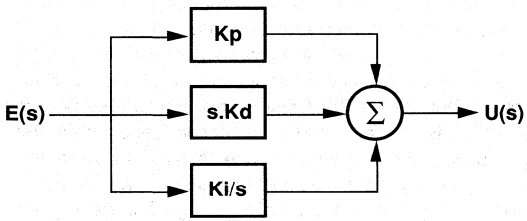

Figure 4. PID Block Diagram

Other types such as PD (Proportional, Derivative) and PI (Proportional, Integral) controllers can also be expressed in a similar manner to the PID relationships in (1) and (2). These controllers lack either the differential or the integral term that exists in the relationships above.

The next step is to derive the discrete equivalent for the controller described by equations (1) and (2). The backward difference is defined as the discrete-time equivalent for the continuous-time derivative of a function. It is obtained by:

$$\Delta f(t) = [f(t) - f(t-T)] / T \qquad (3)$$

where T is the sample period.

The definite sum is defined as the discrete time equivalent for the continuous time integral of a function. It is obtained by:

$$\int_a^b f(\tau) = T [f(a+T) + f(a+2T) + \cdots + f(b)] \qquad (4)$$

where T is the sample period.

By applying the relationships shown in (3) and (4) to the ones shown in (1) and (2) we obtain U(n) and U(z), which are the discrete equivalents of the PID output U(t):

$$U(n) = U(n-1) + A_1 \cdot E(n) + A_2 \cdot E(n-1) + A_3 \cdot E(n-2) \qquad (5)$$

and

$$U(z) / E(z) = [A_1 \cdot z^2 + A_2 \cdot z + A_3] / [z^2 - z] \qquad (6)$$

with

$$A_1 = K_p + K_i \cdot T + K_d / T$$
$$A_2 = -[K_p + 2K_d / T]$$
$$A_3 = K_d / T$$

Transfer functions and difference equations for PD and PI controllers can also be derived in a similar manner. Figure 5 shows the results for these as well as for the PID controllers.

PID CONTROLLER IMPLEMENTATION

An ADSP-2100 Family assembly language subroutine that implements the PID algorithm is shown in Listing 1. There are a number of registers that need to be initialized in order to execute this subroutine. It may be sufficient to do this initialization only once (e.g. on powerup) if other algorithms that are being executed do not need to use these registers. In most typical cases, however, some of these registers may need to be set every time the PID subroutine is called.

The PID routine in Listing 1 takes its input from the AR register. This register must contain the 16-bit error input E(n). E(n) is assumed to be already computed before the PID subroutine is called. The output of the PID algorithm, U(n), is made available in the SR1 register.

	PD Controller	PI Controller	PID Controller
Definitions	$A_1 = K_P + \frac{K_D}{T}$ $A_2 = -\frac{K_D}{T}$	$A_1 = K_P + K_I T$ $A_2 = -K_P$	$A_1 = K_P + K_I T + \frac{K_D}{T}$ $A_2 = -[K_P + 2\frac{K_D}{T}]$ $A_3 = \frac{K_D}{T}$
Transfer Function, $\frac{C(z)}{E(z)}$	$\frac{A_1 z + A_2}{z}$	$\frac{A_1 z + A_2}{z-1}$	$\frac{A_1 z^2 + A_2 z + A_3}{z^2 - z}$
Difference Equation	$U(n) = A_1 E(n) + A_2 E(n-1)$	$U(n) = U(n-1) + A_1 E(n) + A_2 E(n-1)$	$U(n) = U(n-1) + A_1 E(n) + A_2 E(n-1) + A_3 E(n-2)$

Figure 5. PD, PI, and PID Controllers

```
.MODULE        PID_CONTROLLER;

{   This is a PID controller subroutine that executes the following equation:

                                U(n) = B•U(n-1) + A0•E(n) + A1•E(n-1) + A2•E(n-2)

            Calling Parameters:

                AR=      error input E(n), [E(n) = Y(n) - R(n)]
                I0 ->    circular delay line buffer for E(n-2), E(n-1) and U(n-1)
                         this delay buffer must be initialized to zero at powerup
                I4 ->    circular buffer for the scaled coefficients A2, A1, A0, B
                M0,M4=   1
                L0 =     3
                L4 =     4
                SE=      scaling factor for the coefficients

            Return Value:
                SR1=  output sample U(n)

            Altered Registers:
                MX0, MX1, MY0, MR, SR

            Computation Time:
                ADSP2101/2  :    8 Instruction Cycles
                ADSP2105/6  :    8 Instruction Cycles
                ADSP2100/A  :    12 Instruction Cycles

            All coefficients and data values are assumed to be in 1.15 format
}

.ENTRY         PID;

PID:           MX0 = DM(I0,M0), MY0 = PM(I4,M4);
               MR = MX0*MY0 (SS), MX1 = DM(I0,M0), MY0 = PM(I4,M4);
               MR = MR+MX0*MY0 (SS), MY0 = PM(I4,M4);
               MR = MR+AR*MY0(SS), MX0 = DM(I0,M0), MY0 = PM(I4,M4);
               MR = MR+MX0*MY0 (RND), DM(I0,M0)= MX1;
               SR = ASHIFT MR1 (HI), DM(I0,M0) = AR;
               DM(I0,M0) = SR1;
               RTI;

.ENDMOD;
```

Listing 1. PID_CONTROLLER Routine

After the initial design of a digital PID controller, all coefficients must be scaled down by the same factor. This is necessary in order to conform to the 16-bit fixed-point fractional number format as well as to insure that overflows won't occur in the final stage of the multiply-accumulate operations. The scaled down coefficients are the ones that get stored in the processor's memory. The result of the multiply and accumulate operations is eventually scaled up before being output to the controlled system. The choice of a proper scaling factor depends greatly on the design objectives and in some cases it may even be unnecessary. The PID controller coefficients are usually designed with a commercial software package in higher precision arithmetic than 16 bits. System performance deviates from ideal when such high precision PID coefficients are quantized to 16 bits and further scaled down. In systems that require stringent PID specifications, careful simulations of quantization and scaling effects must be performed.

During the initialization for the PID routine, the scaling factor for the coefficients must be stored in the SE register. The index register I0 points to the circular data memory buffer that contains the previous error inputs and the previous PID output. This buffer must be initialized to zero at powerup unless some non-zero initial condition is desired. The index register I4, on the other hand, points to the circular program memory buffer that contains the scaled PID coefficients. These coefficients include a term "B" (for U(n-1)), which is equal to the value " 1/scaling factor ". This value is derived from the fact that the real coefficient for U(n-1) is " 1 " and that it must be scaled down along with the other coefficients. The order that these scaled coefficients are stored in program memory is: A2, A1, A0, B.

The PID core routine fetches the coefficients and data values from memory following the sequence that they have been stored. These values are multiplied and accumulated until all

of them are accessed. Note that both of the address generators are used in parallel with the multiply-accumulator throughout these operations. Finally, the data memory buffer is updated with the new samples and the output is obtained by scaling up the result of the multiplication and accumulation operations.

The PID routine executes in 8 instruction cycles on the ADSP-2101 and ADSP-2105. It executes in 12 instruction cycles on the ADSP-2100 and ADSP2100A. In the case that the initialization registers have been modified by other routines, it may be necessary to execute up to 7 overhead setup cycles before calling the core PID routine.

N'TH ORDER DIGITAL CONTROLLER DESIGN

There are several methods to design high order digital controllers. This section briefly outlines three approaches and cites some references on this topic. The three methodologies are "analog-controller-based digital design", "direct digital design" and "state-space design".

Analog-Controller-Based Digital Design

This is a very common way of designing digital controllers. In this method, an analog controller that satisfies the desired requirements is first created using well established analog design procedures. This controller is then transformed into the digital domain and implemented.

The analog controller design may be performed in the s-plane using common design techniques such as root-locus methods, Bode plots, the Routh-Hurwitz criterion, state variable techniques and other methods. The resulting analog transfer function is then transformed into a digital transfer function in the z-domain. Finally, the z-domain transfer function is inverse-z transformed into a difference equation that can be implemented on a digital processor.

The transformation from the s-domain to the z-domain can be accomplished using various techniques such as the matched pole-zero method, the bilinear (Tustin) transformation the method of mapping differentials, the impulse-invariance method, the step-invariance method, and the zero-order hold technique.

The most commonly used of these methods is the bilinear transformation. This transformation approximates the s-domain transfer function with a z-domain transfer function by use of the substitution:

$$s = (2/T)\,(z-1/z+1) \qquad (7)$$

Analog controllers that are in parallel or cascade maintain their respective structures after going through this transformation. The Tustin transformation maps the stable region of the s-plane exactly into the stable region of the z-plane although the entire jw-axis of the s-plane is stuffed into the 2π-length of the unit circle. Obviously a great deal of distortion takes place in the mapping in spite of the consistency of the stability regions. This distortion can be corrected by using a frequency pre-warping scheme. The pre-warping matches the single most important critical frequency in the analog domain and the digital domain. This method replaces each "s" in the analog transfer function with

$$(\omega 1 / \omega_2)\, s$$

where ω_1 is the frequency to be matched in the digital transfer function and with

$$\omega_2 = (2/T) \tan (\omega_1 T/2) \qquad (8)$$

Bilinear transformation with frequency pre-warping is one of the most commonly used analog based design techniques. The most significant drawback of this type of design is that the digital controller that results from it is only an approximation to the analog one. The analog controller puts an implicit upper bound on the digital controller's performance. More information on the other analog based controller design methods can be found in References 1 and 3.

Direct Digital Design

This method allows us to perform the control system design directly in the digital domain. Thus, the design can be carried out in the z-plane. The approximations and limitations that arise from starting in the s-domain and transforming into the z-domain are eliminated. Conventional design techniques can be used to place the closed-loop poles and zeros exactly where appropriate. Some of these z-domain techniques include the root-locus method, the pole-zero cancellation method and the w-transform. More detailed information on these methods can be found in References 1 and 3.

State-Space Design

The digital controller design methods discussed above are designated as classical design methods. Same design tasks can be accomplished by using a different set of techniques based on the state-space or modern control formulation. Modern control design methodology is especially advantageous when designing controllers for multi-input and multi-output systems. However, single-input and single-output systems that are discussed in this application note can also be efficiently designed using state-space methods. More detailed information on this topic can be found in References 1 and 3.

N'TH ORDER DIGITAL CONTROLLER STRUCTURES

Standard second order (N=2) digital controller implementation is directly analogous to IIR (Infinite Impulse Response) filter implementations. These second order controller blocks can be implemented as biquad second order IIR filter sections. A second order biquad section is shown on Figure 6 and its corresponding transfer function in the z-domain is given by:

$$G(z) = U(z)/E(z) = (B_0 + B_1 \cdot z^{-1} + B_2 \cdot z^{-2})/(1 + A_1 \cdot z^{-1} + A_2 \cdot z^{-2}) \qquad (9)$$

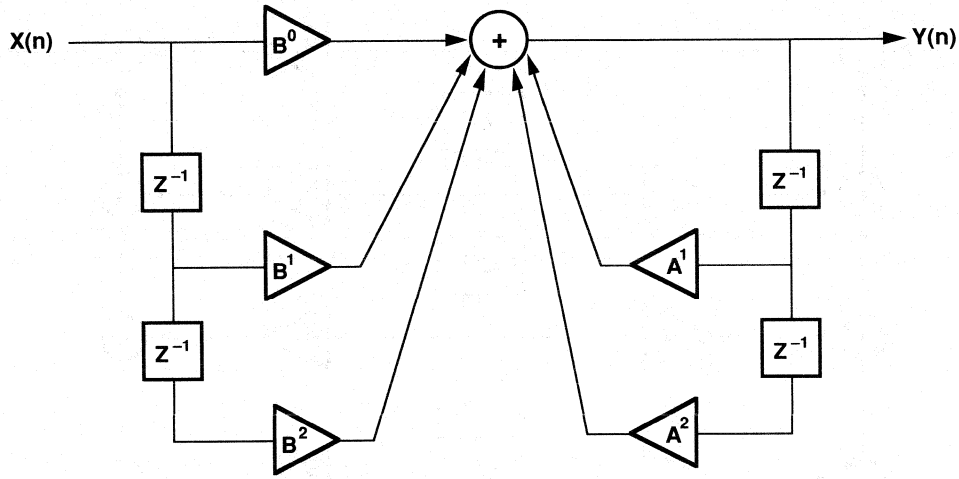

Figure 6. Second Order Biquad Structure

where A_1, A_2, B_0, B_1 and B_2 are coefficients that determine the desired impulse response of the system G(z). Furthermore, the corresponding difference equation for a biquad section is given by:

$$U(n) = B_0 \cdot E(n) + B_1 \cdot E(n-1) + B_2 \cdot E(n-2) - A_1 \cdot U(n-1) - A_2 \cdot U(n-2) \quad (10)$$

Higher order (N'th order) controllers can be obtained by cascading several biquad sections with appropriate coefficients. An example is shown on Figure 7 where three biquad sections are cascaded to construct the overall G(z) transfer function. Another way to design higher order controllers is to use only one complicated single section. This approach is also called the direct form implementation. The block diagram of a direct form fourth order controller is shown on Figure 8 as an example.

The direct form implementation executes faster but generates larger numerical errors than the biquad implementation. The biquads can be scaled separately and then cascaded in order to minimize the coefficient quantization and the recursive accumulation errors. The coefficients and data in the direct form implementation must be scaled all at once, which gives rise to larger errors. Another disadvantage of the direct form implementation is that the poles of such single stage high order polynomials get increasingly sensitive to quantization errors. The second order polynomial sections (i.e. biquads) are less sensitive to quantization effects.

9

N'TH ORDER CONTROLLER IMPLEMENTATION

An ADSP-2100 Family assembly language subroutine that implements a high order controller is shown in Listing 2. The subroutine is arranged as a module and is labeled BIQUAD_CONTROLLER. There are a number of registers that need to be initialized in order to execute this subroutine. It may be sufficient to do this initialization only once (e.g. on powerup) if other executed algorithms do not need these registers. In most typical cases, however, some of these registers may need to be set every time the BIQUAD_CONTROLLER subroutine is called. It may sometimes be beneficial, from a modular software point of view, to always initialize all the setup registers as a part of this subroutine.

The BIQUAD_CONTROLLER routine in Listing 2 takes its input from the SR1 register. This register must contain the 16

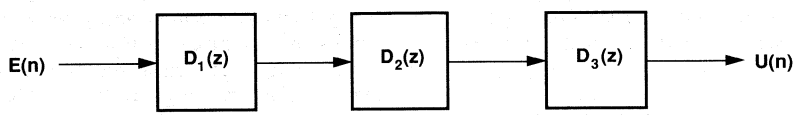

$$G(z) = \frac{U(z)}{E(z)} = D_1(z) \cdot D_2(z) \cdot D_3(z)$$

Figure 7. Cascased Biquad Sections

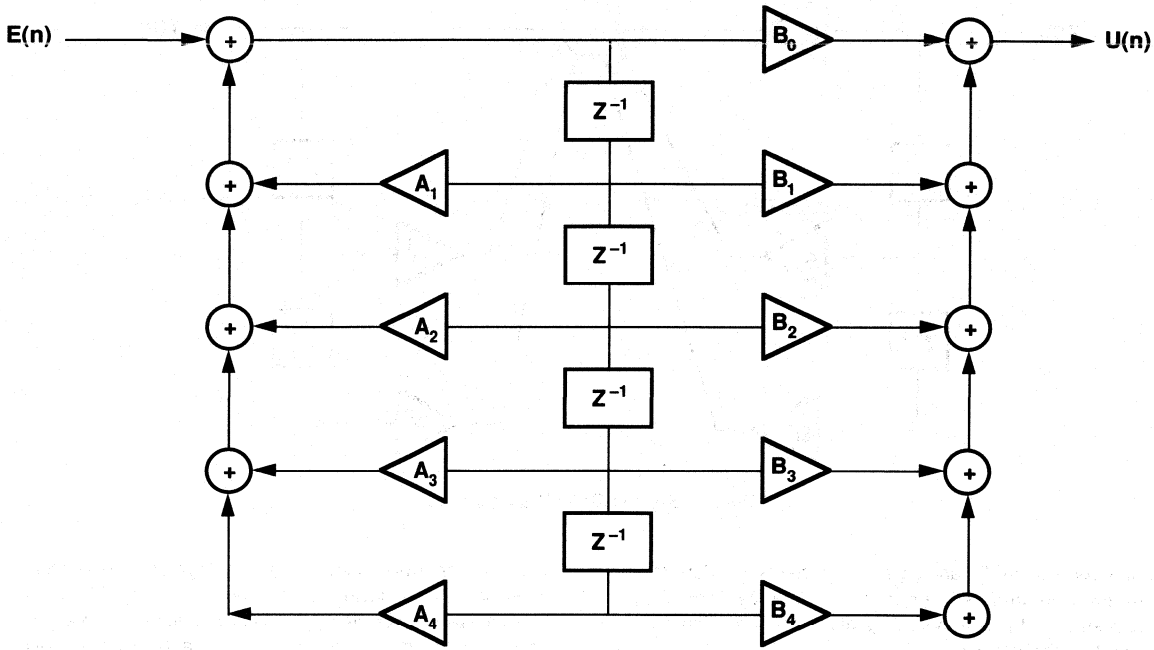

Figure 8. Fourth Order Direct Form Controller

bit error input E(n). E(n) is assumed to be already computed before this subroutine was called. The output of the controller is also made available in the SR1 register.

After the initial design of a high order controller, all coefficients must be scaled down in each biquad stage separately. This is necessary in order to conform to the 16 bit fixed point fractional number format as well as to insure that overflows won't occur in the final multiply-accumulate operations in each stage. The scaled down coefficients are the ones that get stored in the processor's memory. The operations in each biquad are performed with scaled data and coefficients and are eventually scaled up before being output to the next one. The choice of a proper scaling factor depends greatly on the design objectives and in some cases it may even be unnecessary. The controller coefficients are usually designed with a commercial software package in higher precision arithmetic than 16 bits. System performance deviates from ideal when such high precision coefficients are quantized to 16 bits and further scaled down. In systems that require stringent specifications, careful simulations of quantization and scaling effects must be performed.

During the initialization of the BIQUAD_CONTROLLER routine, the index register I0 points to the data memory buffer that contains the previous error inputs and the previous biquad section outputs. This buffer must be initialized to zero at powerup unless some non-zero initial condition is desired. The index register I1 points to another buffer in data memory that contains the individual scale factors for each biquad. The buffer length register L1 is set to zero if the controller has only one biquad section. L1 is initialized with the number of biquad sections in the case of multiple biquads. The index register I4, on the other hand, points to the circular program memory buffer that contains the scaled biquad coefficients. These coefficients are stored in the order : B2, B1, B0, A2, A1 for each biquad. All of the individual biquad coefficient groups must be stored in the same order that the biquads were cascaded in such as: B2, B_1, B_0, A_2, A_1, B_2*, B_1*, B_0*, A_2*, A_1*,B_2**...etc The buffer length register L4 must be set to the value given by: (5 x number of biquad sections). Finally, the loop counter register "CNTR" must be set to the number of biquad sections since the controller code will be executed as a loop.

The core of the BIQUAD_CONTROLLER routine starts its execution at the "biquad" label. The routine is organized in a looped fashion where the end of the loop is the instruction labeled "sections". Each iteration of the loop executes the computations for one biquad. The number of loops to be executed is determined by the CNTR register contents. The SE register is loaded with the appropriate scaling factor for the particular biquad at the beginning of each loop iteration. After this operation, the coefficients and the data values are fetched from memory in the sequence that they have been stored. These numbers are multiplied and accumulated until all of the values for a particular biquad have been accessed. The result of the last multiply accumulate is rounded to 16 bits and upshifted by the scaling value. At this point the "biquad" loop is executed again or the controller computations are completed by doing the final update to the delay line. The delay lines for data values are always being updated within the biquad loop as well as outside of it.

```
.MODULE          BIQUAD_CONTROLLER;

{   This is an Nth order cascaded biquad controller subroutine

                 Calling Parameters:

                      SR1= error input E(n),  [ E(n) = Y(n) - R(n) ]
                      I0 -> delay line buffer for E(n-2), E(n-1), Y(n-2), Y(n-1)
                      L0 = 0
                      I1 -> scaling factors for each biquad section
                      L1 = 0  (in the case of a single biquad)
                      L1 = number of biquad sections (for multiple biquads)
                      I4 -> scaled biquad coefficients
                      L4 = 5 x [number of biquads]
                      M0,M4 = 1
                      M1 = -3
                      M2 = 1 (in the case of multiple biquads)
                      M2 = 0 (in the case of a single biquad)
                      M3 = (1 - length of delay line buffer)

                 Return Value:
                      SR1 = output sample U(n)

                 Altered Registers:
                      SE, MX0, MX1, MY0, MR, SR

                 Computation Time (with N even):
                      ADSP-2101/2102:        (8 x N/2) + 5 instruction cycles
                      ADSP-2100/2100A:       (8 x N/2) + 5 + 5 instruction cycles

                 All coefficients and data values are assumed to be in 1.15 format
}

.ENTRY           CONTROLLER;

BIQUAD:          CNTR = number of biquads
                 DO SECTIONS UNTIL CE;
                      SE = DM(I1,M2);
                      MX0 = DM(I0,M0), MY0 = PM(I4,M4);
                      MR = MX0*MY0(SS), MX1 = DM(I0,M0), MY0 = PM(I4,M4);
                      MR = MR+MX1*MY0(SS), MY0=PM(I4,M4);
                      MR = MR+SR1*MY0(SS), MX0 = DM(I0,M0), MY0 = PM(I4,M4);
                      MR = MR+MX0*MY0(SS), MX0 = DM(I0,M1), MY0 = PM(I4,M4);
                      DM(I0,M0) = MX1, MR = MR+MX0*MY0(RND);
SECTIONS:             DM(I0,M0) = SR1, SR = ASHIFT MR1 (HI);
                 DM(I0,M0) = MX0;
                 DM(I0,M3) = SR1;
                 RTS;
.ENDMOD;
```

Listing 2. BIQUAD_CONTROLLER Routine

The controller coefficients must be scaled appropriately so that no overflows occur after the upshifting operation between the biquads. If this is not insured by design, it may be necessary to include some overflow checking between the biquads.

The execution time for an N'th order BIQUAD_CONTROLLER routine can be calculated as follows (assuming that the appropriate registers have been initialized and N is a power of 2):

ADSP2101/2105 : (8 x N/2) + 4 processor cycles
ADSP2100/2100A : (8 x N/2) + 4 + 5 processor cycles

It may take up to a maximum of 11 cycles to initialize the appropriate registers every time the controller is called. But typically this number will be lower in most applications.

NOTCH FILTER EXAMPLE FOR THE ADSP-2100A

A fully coded ADSP-2100A notch filter program is presented in this section. The program executes two cascaded biquad sections and is designed to run standalone on an ADSP-2100A. The processor reads the input samples from a data memory mapped A/D converter and sends the filtered output to a data memory mapped D/A converter. The operation of the ADSP-2100A is interrupt driven. The occurrence the IRQ0 interrupt notifies the processor that there is a new sample ready at the A/D converter output. The ADSP-2100A normally waits in a wait loop, processes the incoming samples upon an interrupt and returns to the wait loop again. This process starts with a reset or powerup and repeats until a powerdown or another reset occurs. The full code of the notch filter is shown in Listing 3.

The program starts with some variable, constant and port declarations. These declarations allow the program to refer to specific memory addresses symbolically. This greatly eases the software maintenance and debugging tasks at the assembly level. In order to run this program, the NOTCH_FILTER assembly module shown in Listing 3 must be first assembled using the ADSP-2100 assembler. It must then be linked with an architecture description file that was built using the ADSP-2100 system builder. This architecture file

```
{THIS IS AN ADSP2100A ASSEMBLY PROGRAM THAT EXECUTES A FOURTH}
{ORDER NOTCH FILTER IN A CASCADED BIQUAD IMPLEMENTATION}

.MODULE/RAM             NOTCH_FILTER;           {The name of the module}

.VAR/PM/CIRC            COEFFICIENTS[10];       {These are the declarations for}
.VAR/DM                 DATA_BUFFER[6];         {data and program memory buffers}
.VAR/DM/CIRC            SCALE_FACTORS[2];

.PORT                   AD_CONVERTER;           {There is one input port and one}
.PORT                   DA_CONVERTER;           {output port in the system}

.INIT                   COEFFICIENTS: <COEFF.DAT>;    {The memory buffers}
.INIT                   DATA_BUFFER: <INITIAL.DAT>;   {are initialized here}
.INIT                   SCALE_FACTORS: <SCALE.DAT>;

                JUMP BIQUAD;                    {Interrupt vector for IRQ0}
                RTI;
                RTI;
                RTI;

                I0 = ^DATA_BUFFER;              {These are the proper initializations}
                L0 = 0;                         {for the index, length and modify}
                I1 = ^SCALE_FACTORS;            {registers to be used}
                L1 = %SCALE_FACTORS;
                I4 = ^COEFFICIENTS;
                L4 = %COEFFICIENTS;
                M0 = 1;
                M1 = -3;
                M2 = 1;
                M3 = -5;
                M4 = 1;
                ICNTL = B#00001;                {Set IRQ0 to be edge sensitive}
                IMASK = B#0001;                 {Enable IRQ0 interrupt only}
WAIT:           JUMP WAIT;                      {Wait for IRQ0 interrupt to occur}
```

```
{The interrupt service routine below executes the two biquad sections of the filter}

BIQUAD:         SR1=DM(AD_CONVERTER);          {Read the A/D converter}
                CNTR = %SCALE_FACTORS;
                DO SECTIONS UNTIL CE;
                      SE = DM(I1,M2);
                      MX0 = DM(I0,M0), MY0 = PM(I4,M4);
                      MR = MX0*MY0(SS), MX1 = DM(I0,M0), MY0 = PM(I4,M4);
                      MR = MR+MX1*MY0(SS), MY0=PM(I4,M4);
                      MR = MR+SR1*MY0(SS), MX0 = DM(I0,M0), MY0 = PM(I4,M4);
                      MR = MR+MX0*MY0(SS), MX0 = DM(I0,M1), MY0 = PM(I4,M4);
                      DM(I0,M0) = MX1, MR = MR+MX0*MY0(RND);
SECTIONS:             DM(I0,M0) = SR1, SR = ASHIFT MR1 (HI);
                DM(I0,M0) = MX0;
                DM(I0,M3) = SR1;
                DM(DA_CONVERTER) = SR1;        {Send the filtered output to the D/A}
                RTI;                           {Return to the wait loop}
.ENDMOD;
```

Listing 3. NOTCH_FILTER Routine

should be specific to the particular hardware configuration that the ADSP-2100A is being built in. The assembly module also needs to be linked with three data files along with the architecture file. The first data file must contain the scale factors, the second one must contain the scaled coefficients and the third one must contain the initial values for the delay taps of the filter.

9

ADSP-2100 Family Benchmarks For Digital Control Applications		
	ADSP-2101 (60 ns)	ADSP-2105 (100 ns)
PID Loop	0.48 µs	0.8 µs
FIR Filter	60 ns/tap	100 ns/tap
IIR Biquad – 16 Bit (2nd Order)	0.72 µs	1.2µs
IIR Biquad – 16 Bit (nth Order)	4N + 4 cycles	4N + 9 cycles
IIR Biquad – 32 Bit (2nd Order)	2.64 µs	4.4 µs
IIR Biquad – 32 Bit (nth Order)	16N + 12 cycles	16N + 22 cycles
256-Point FFT (Complex)	0.405 ms	0.675 ms
Matrix Multiply (3x3 * 3x1)	3.12 µs	5.2 µs
Stochastic Gradient (LMS) N Tap Coefficient Update	2N + 9 cycles	2N + 9 cycles

REFERENCES

(1) Franklin, G.F., J.D. Powell and M.L. Workman 1990. "Digital Control of Dynamic Systems". Reading, MA: Addison-Wesley Publishing Company.

(2) Oppenheim, A. V. and A. Willsky 1983. "Signals and Systems". Englewood Cliffs, NJ: Prentice-Hall, Inc.

(3) Borrie, J. A. 1986. "Modern Control Systems". Englewood Cliffs, NJ: Prentice-Hall Inc.

(4) Kazanzides, P. 1985. "A Microprocessor Based Control System with Robotics Applications". Brown University, LEMS, Providence, RI.

ANALOG DEVICES

AN-338
APPLICATION NOTE

)NE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Loading an ADSP-2101 Program via the Serial Port

by Gerald McGuire

NTRODUCTION

For many DSP applications, it is desirable to have a DSP processor under the control of a host computer. In these situations, the host computer would download a program for he DSP to execute. The ADSP-2101 provides two serial ports suitable for program download from a host computer. This application note details the ADSP-2101 monitor program for downloading from a serial port. The monitor program itself would be booted from EPROM or other boot memory. While his example uses serial port zero, the principal could be extended to download via a memory-mapped parallel port.

A MONITOR

The task of the host computer is to download a series of nstructions to the ADSP-2101 for execution. The ADSP-2101 eceives the incoming instructions, loads them into program nemory and when all instructions have been received, executes them. Prior to and during the download from the host, he ADSP-2101 executes a monitor program. This monitor activates the serial port, receives the instructions and places hem in program memory for execution.

The ADSP-2101 instruction is twenty-four bits wide but many nosts, including eight-bit processors, more readily handle byte-wide data. Since the serial port can accommodate serial words from three to sixteen bits in length, byte-length data words are easily received.

Whenever a program memory write occurs, the sixteen most significant bits are supplied by the source register, explicitly named in the instruction, and the eight LSBs are supplied by he PX register. The basic tactic of the monitor program is to assemble the two most significant bytes in a data register using the Shifter) and load PX explicitly with the least significant byte. A program memory write then writes the correct twenty-four bit instruction.

n addition to the transfer of instructions through the serial port nto program memory, the monitor program must also know when the download is complete and execution can begin. A straight forward method is to count the number of instructions sent to the serial port. A count value is sent to the ADSP-2101 before the first instruction. This is the count of the instructions to follow. After each instruction is downloaded, the count can be decremented.

The downloaded program must avoid overwriting the monitor program while the monitor executes. The last instruction of the monitor program is identified by a global label which also identifies the beginning of the available space for downloaded code. The monitor program must be linked with the downloaded program so that the downloaded program makes the correct address references including the reference to this global label.

The indirect addressing capabilities of the Data Address Generators on the ADSP-2101 make it easy to cycle through the correct sequential locations starting with the label.

The final concern is the interrupt table. If the downloaded program is interrupt-driven, the interrupt table (program memory H#0000 to H#001C) must contain valid instructions for servicing expected interrupts.

There are several ways to do this. First, the monitor program itself could contain the valid interrupt table for the program to be downloaded. This assumes that the interrupt structure of the downloaded program is known when the monitor program is created. Second, the interrupt table may be downloaded through the serial port just as the rest of the program is. The DAG can loaded with the start address of the interrupt table and the instructions can be loaded, but you may not overwrite the interrupt being used to receive the data on the serial port until all instructions have been received.

The monitor program example does not load an interrupt table. The best approach is dependent on your application.

IMPLEMENTATION

The first task of the monitor program is to setup and enable the

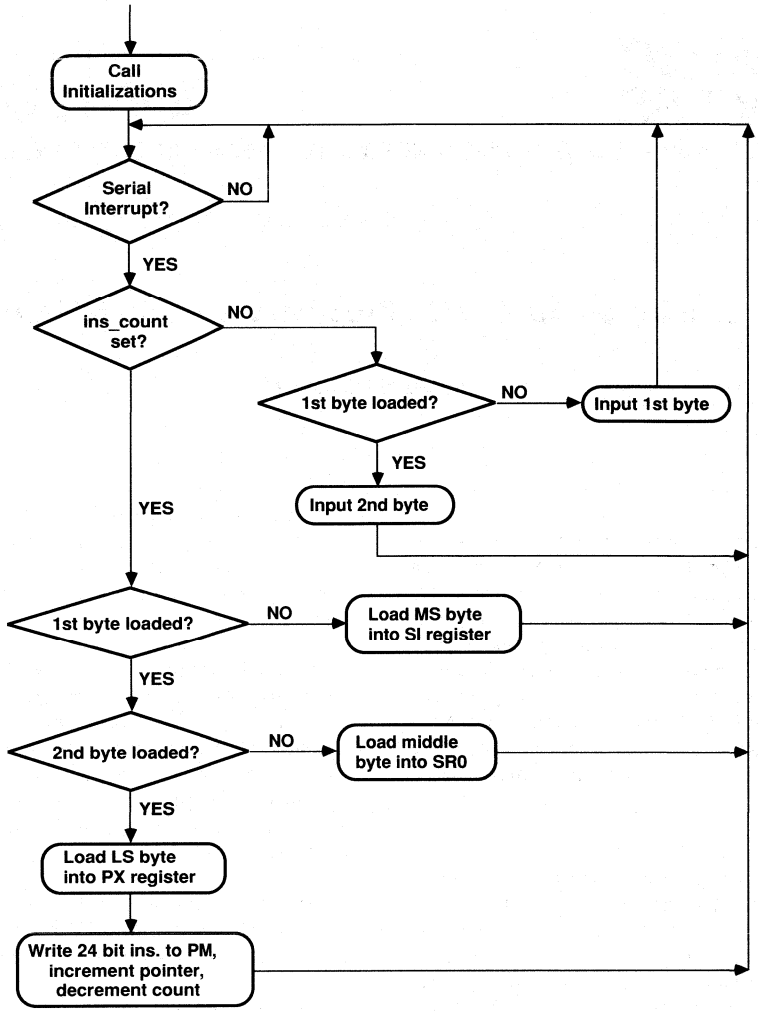

Figure 1. Boot Program Flow Diagram

serial port. Serial ports on the ADSP-2101 are extremely flexible in terms of framing options, word lengths and timing. The ADSP-2101 serial ports may receive the frame synch and serial clock from the host processor or generate them internally.

As the program is downloaded from a host computer, the ADSP-2101 looks to the host for serial port information. That is, the serial port frame synchronization and serial port clock are supplied by the host computer. For purposes of illustration, the code that appears at the end of this application note uses normal framing and external receive frame synchronization. For externally generated serial clocks the ADSP-2101 can support frequencies up to the processor instruction rate.

The flow for the monitor program is shown in Figure 1.

Once the serial ports are enabled, the monitor program waits for a serial port interrupt signifying that a serial word has been received. The first two serial words received are the instruction count. As the serial word is eight bits, two serial port words make up the instruction count. The separate bytes of the instruction count are combined in the shifter and loaded into data memory. This count represents the number of instructions to be downloaded from the host and does not include the interrupt table. The interrupts are handled automatically, as the interrupt table has a fixed length.

With the count downloaded, the ADSP-2101 is ready to accept instructions through the serial port. Instructions are downloaded a byte at a time just as the instruction count was. The most significant byte is first. It is loaded into the SI register and the byte count ("count") is decremented. The middle byte

If the instruction is loaded into the SR0 register. These two bytes are combined in the shifter with the results residing in the SR0 register. Once again the byte count is decremented. Finally, the least significant byte is loaded into the PX register. Now that all three bytes are loaded into registers on the ADSP-2101, the downloaded instruction can be written to program memory.

When all is downloaded, a jump to the new downloaded program is all that is necessary to begin execution.

The monitor program is listed at the end of this note.

SUMMARY

A monitor program initializes the serial port and receives instructions, writing them into program memory, then beginning execution. This method of booting is useful when the ADSP-2101 is under the control of a host computer or controller. Any size program may be downloaded (up to the full addressing capability of the ADSP-2101) with this particular method of implementation. Only the program memory used by the monitor program cannot be loaded. That space could, however, be used for program memory data storage.

```
.MODULE/RAM/BOOT=0                  serial_boot_monitor;
.VAR/DM                             count;              {counts bytes}
.VAR/DM                             ins_count;          {counts instructions}
.GLOBAL                             code_start;         {end of monitor space}


                    JUMP restarter; NOP; NOP; NOP;      {restart vector}
                    RTI; NOP; NOP; NOP;                 {IRQ2 not used}
                    RTI; NOP; NOP; NOP;                 {sport0 TX not used}
                    JUMP serial; NOP; NOP; NOP;         {sport0 RX }
                    RTI; NOP; NOP; NOP;                 {sport1 TX not used}
                    RTI; NOP; NOP; NOP;                 {sport1 RX not used}
                    RTI; NOP; NOP; NOP;                 {no timer used}

restarter:          CALL initializations;
wait_loop:          IDLE;
                    JUMP wait_loop;

initializations:    I4 = H#3ff3;                        {pointer to mem map reg}
                    I5 = ^code_start;                   {pointer to start of prog}
                    M4 = 0;
                    M5 = 1;
                    L4 = 0;
                    L5 = 0;
                    SR0 = 0;
                    SR1 = 0;

                    AX1 = 1;
                    DM(count) = AX1;                    {count val for # of bytes}
                    DM(I4,M5) = 0;                      {disable autobuffer}
                    DM(I4,M5) = 0;                      {no frame divide modulus}
                    DM(I4,M5) = 0;                      {no clk divide modulus}
                    DM(I4,M5) = H#2007;                 {extrnl RFS & SCLK, no compand}
                                                        {SLEN 8, no multichannel}
                    AX0 = H#1000;
                    DM(H#3fff) = AX0;                   {enable sport0}
                    AX0 = H#FFFF
                    DM(ins_count) = AX0;

                    IMASK = 8;                          {sport0 rec interrupt only}
                    RTS;

serial:             AY1 = DM(ins_count);
                    AR = PASS AY1;
                    IF GT JUMP next_instruction;        {get next instruction}
                    IF LT JUMP load_word_count;         {get number of instructions}
                    IF EQ JUMP code_start;              {start downloaded program}
```

```
{load the count, that is, the number of instructions to be downloaded}
{this happens in two bytes  The most significant byte first}

load_word_count:    AY0 = DM(count);                    {is this 1st or 2nd byte}
                    AR = PASS AY0;
                    IF NE JUMP first_byte;
                    IF EQ JUMP second_byte;

first_byte:         SI = RX0;                           {first byte decrem. count}
                    AR = AY0 - 1;
                    DM(count) = AR;
                    RTI;

second_byte:        SR0 = RX0;                          {second byte...}
                    SR = SR OR LSHIFT SI BY 8 (LO);     {put two bytes together}
                    DM(ins_count) = SR0;                {store in ins_count}
                    AX0 = 3;
                    DM(count) = AX0;                    {load count for ins.}
                    RTI;


{load the next instruction.  Instructions are 24 bits long and appear}
{at the serial port in 8 bit fragments.  The most significant byte 1st}

next_instruction:   AX0 = 2;                            {decide which byte is due}
                    AY0 = DM(count);
                    AR = AX0 - AY0;

                    IF LT JUMP most_sig_byte;
                    IF EQ JUMP middle_byte;
                    IF GT JUMP least_sig_byte;

most_sig_byte:      SI = RX0;                           {load MS byte into SI}
                    AR = AY0 - 1;                       {decrement count}
                    DM(count) = AR;
                    RTI;

middle_byte:        SR0 = RX0;                          {load Middle into SR}
                    SR = SR OR LSHIFT SI BY 8 (LO);     {put MS and middle together}
                    AR = AY0 - 1;                       {decrement count}
                    DM(count) = AR;
                    RTI;

least_sig_byte:     PX = RX0;                           {put LS byte into PX}
                    PM(I5,M5) = SR0;                    {write SR0 into PM}
                                                        {PX provides 8 LS bits}
                    AX0 = 3;
                    DM(count) = AX0;                    {reset byte count}
                    AR = AY1 - 1;                       {decrement ins count}
                    DM(ins_count) = AR;
                    RTI;

code_start:         NOP;

.ENDMOD;
```

ANALOG DEVICES

AN-240
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Considerations for Selecting a DSP Processor (ADSP-2101 vs. WE DSP16A)

by Bruce Wolfeld

INTRODUCTION

Digital signal processing systems contain high-performance numerical processing units that capable of performing rapid computations. The numerical performance of a DSP system is measured by the processing unit's capabilities in the following areas:

- Fast and flexible arithmetic with extended dynamic range
- Efficient operand access
- Circular buffering capabilities
- Overhead required for program loops

As with any microprocessor design, you, the system designer, must consider other factors including:

- Interface to external devices
- Quantity and speed of external memory
- Instruction set
- Development tools available for system debug

This application note first discusses the numerical processing capabilities of the Analog Devices ADSP-2101 and the AT&T DSP16A. The second part discuses other factors that you should consider.

ARITHMETIC CAPABILITIES

One indication of good arithmetic architecture is the ability to perform a wide range of arithmetic computations. These computations should be flexibly handled so that the algorithm implementation preserves the order of the arithmetic operations and operands. If the arithmetic architecture is too special purpose, this is impossible. Rearranging an algorithm to fit the architecture requires extra programming, increases the possibility of error and may increase product development time.

Arithmetic Architecture

Figure 1 shows a block diagram of the arithmetic section of the ADSP-2101 while Figure 2 shows that of the DSP16A. Both of these devices utilize a modified Harvard architecture which can fetch data operands from both program memory and data memory. By utilizing on-chip memory, the ADSP-2101 can fetch an instruction from program memory and fetch/store two data operands from memory in a single cycle. The DSP16A requires one cycle to fetch an instruction from program ROM. Two cycles are required if the instruction and immediate data both reside in the program ROM.

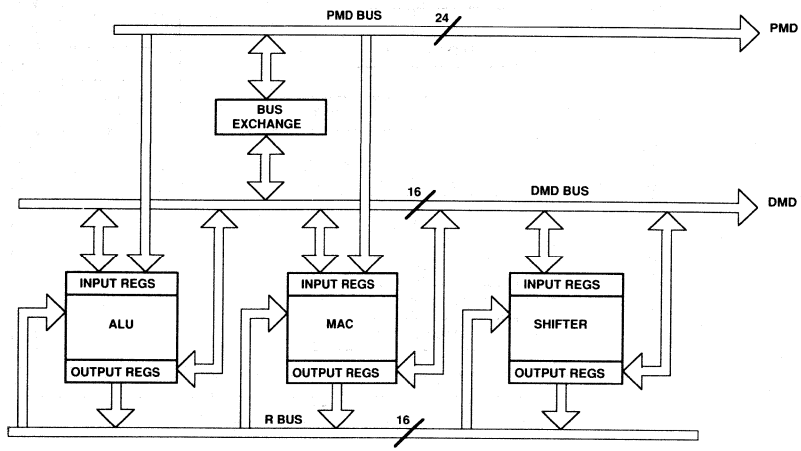

Figure 1. Block Diagram of Arithmetic Section of the ADSP-2101

The arithmetic section of the ADSP-2101 architecture was designed so DSP algorithms are easily coded and rapidly executed. Unlike many DSP processors, the ADSP-2101 uses an algebraic notation for a variety of multi-function instructions. These operations exploit the inherent parallelism of the ADSP-2101 architecture by providing combinations of register moves, memory transfers, and conditional computation. Each program statement assembles into a single 24-bit opcode that can execute in a single cycle. There are no multicycle instructions in the ADSP-2101 instruction set.

The ADSP-2101 has three independent computational units: an arithmetic/logic unit (ALU), a multiplier/accumulator (MAC), and a barrel shifter. They are connected by the result bus (R bus) so that the output of any computational unit can be used directly as the input for itself or any other unit on the next instruction cycle. In addition, the ALU and MAC are directly connected to both the program and data memory busses. Operands for ALU and MAC operations can come from program memory, data memory or any of the result registers. The arithmetic units have features that allow for 32-bit or 64-bit computations.

The arithmetic section of the DSP16A contains a multiplier unit with a scaling shifter and a ALU/shifter unit. The multiplier has three input registers and one output register. While all three input registers can be loaded from data memory, only the x register can be loaded from program memory. The ALU/shifter can operate on data in the y registers, the multiplier output register, or the accumulator registers. The y registers and the accumulator registers are accessible only through data memory. If arithmetic results are to be stored in external memory, they must first be loaded into one of the accumulator registers.

ADSP-2101 ALU

The ADSP-2101 ALU has two X and two Y input registers: AX0, AX1, and AY0, AY1. ALU operations are performed on any X, Y register combination. These registers may be loaded from program memory, data memory, the counter or a variety of other processor registers. ALU results can be stored in either the ALU result (AR) or the ALU feedback (AF) register. AR and AF are available as the X and Y operands (respectively) in subsequent ALU operations. In addition, the result registers of the MAC and barrel shifter can be used as the X input to the ALU.

The ALU performs mathematical operations on 16-bit data operands. An internal carry bit is always updated during a mathematical operation. Addition with carry and subtraction with borrow instructions are provided for implementing 32-bit arithmetic.

Two operands can be fetched or stored in parallel with an ALU operation. By fetching operands while executing ALU operations, the processor can perform one ALU operation per cycle without speed penalties due to operand fetching. All ALU operations can execute in a single cycle.

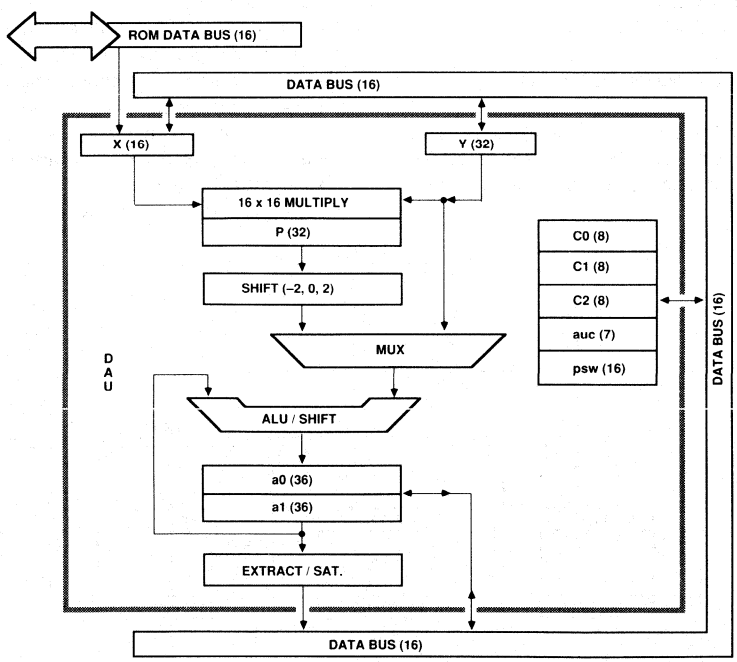

Figure 2. Block Diagram of Arithmetic Section of the DSP16A

Other features of the ADSP-2101 ALU include six status flags, division primitives and a complete set of background registers for fast context switching.

DSP16A ALU

DSP16A ALU operations require that one operand must come from an accumulator registers while the other must come from the multiplier output or from the data bus through yh or yl registers. Two-number addition is a two-step process. First, the accumulator must be loaded with the first data value. After the accumulator is loaded, a second number can then be added. Before the result can be used as an input value to the multiplier, it first must be stored back into data memory.

The ALU operates on 16- or 32-bit operands compared to 16-bit operands on the ADSP-2101. There is no carry bit or provisions for operating on larger operands. ALU result registers are 36 bits wide.

The number of cycles required to execute an ALU operation depends on the number of data transfer operations and if the instruction word is present in the cache. Generally, a single ALU operation executes in one cycle. An ALU operation with a data transfer may require two cycles.

ADSP-2101 MAC

As shown in Figure 1, the ADSP-2101 multiplier/accumulator (MAC) is separate from the ALU. All MAC operations occur in a single cycle. The unit performs both multiplications and MAC operations independent of the ALU. This is a key difference from the architecture of the DSP16A.

Like the ALU, the MAC has two X and two Y input registers, MX0, MX1 and MY0, MY1. Operations are performed on any X-Y pair of input registers. These registers can be loaded from program memory, data memory or other registers in the processor. The result of the operation appears in the multiplier result register (MR) or the multiplier feedback register (MF). As in the ALU, the feedback and result registers can also serve as the X and Y inputs to subsequent MAC operations. The barrel shifter result registers and the ALU result register can also be used directly as X inputs to the MAC.

The ADSP-2101 multiplier performs mathematical operations on 16-bit data. The data formats can be any combination of signed or unsigned values. This feature simplifies the implementation of 32-bit multiplication.

The multiplier result register (MR) is a 40-bit accumulate register. MR is divided into two 16-bit registers (MR0 and MR1) and one 8-bit extension register (MR2). DSP applications frequently deal with accumulator operations that require a large dynamic range. The extension register allows for 256 MAC overflows before a loss of data can occur.

All MAC operations can execute in a single cycle. Since two new operands can be loaded into the input registers in parallel with the computation, a new MAC operation can be executed every cycle.

Other features in the MAC include an overflow status bit, a 16-bit rounding option, and a complete bank of background registers for fast context switching.

DSP16A MAC

As shown in Figure 2, the DSP16A MAC consists of a 16x16 multiplier and the ALU described above. One multiply operand must come from the 16-bit x register. The second operand is chosen from one of the two 16-bit y registers. Because a multiply/accumulate operation requires both the ALU and the multiplier, two cycles are required to perform a single MAC operation.

There are only three registers available for MAC operands. The x register can be loaded from program or data memory. The y registers must be loaded from data memory. Two cycles are required to load the x register if the load instruction is not available in the cache. The 32-bit multiplier result is stored in the p register. To complete the MAC operation, the p value is added to an accumulator. A minimum of two cycles are required before a multiplier result can be used in another multiplication.

The DSP16A accumulate registers are 36 bits wide. A MAC operation can overflow 16 times before data is lost.

At least two clock cycles are required to execute a single MAC operation. Multiple MAC operations can be pipelined to reduce execution time. In addition, two new operands can be loaded in parallel with the MAC operation.

ADSP-2101 Barrel Shifter

The barrel shifter in the ADSP-2101 has an SI input register and can also accept inputs from any result register in the processor (e.g. MR1, AR) including its own result register, SR. Like the MAC result registers, the 32-bit SR is divided into two 16-bit registers, SR0 and SR1. The shifter also has an exponent register, SE, which is set automatically by the exponent adjust instruction and is used for normalization.

The shifter can place a 16-bit input value anywhere in a 32-bit field. The input can be shifted any number of bits from off-scale left to off-scale right. The shift can be sign-extended or zero-filled. Other functions such as exponent detection, normalization, denormalization, block floating-point exponent maintenance, and pattern merging can also be performed with the shifter. All shifter operations are performed in a single cycle regardless of the number of bits shifted.

DSP16A Shifters

The DSP16A contains two shifters. One shifter can shift the output of the multiplier two bits, left or right. This shifter is used for scaling multiplier results. The second shifter is contained in the ALU. Since the second shifter requires the same circuitry as the ALU, it has the same limitations as the ALU. The DSP16A can only shift operands located in the accumulators.

The ALU/shifter can perform a sign extended right-shift or a zero-fill left shift. The DSP16A has no provisions for block floating point, or pattern merging. Data can be shifted in a

single cycle only by 1, 4, 8 or 16 bits. A 6-bit shift requires three instruction cycles.

Arithmetic Summary

Table 1 summarizes the comparison of the arithmetic capabilities of the ADSP-2101 and the DSP16A.

DSP Requirement	ADSP-2101	DSP16A
Single cycle ALU operations	√	√[1]
Single cycle multiplication	√	√
Single cycle MAC operations	√	**No**[2]
Single cycle shifting	0-32 bits	1,4, 8, or 16 bits
Accumulator overflow protection	8 bits	4 bits
Signed, unsigned or mixed-mode multiplications	√	**No**

1. May require 2 cycles if data transfer is required.
2. Multiple MAC operations can be pipelined to reduce execution time.

Table 1. Arithmetic Capabilities

DATA ADDRESSING

A digital signal processor's ability to perform fast arithmetic is wasted if the required data cannot be fetched at a similar speed. DSP algorithms require that data operands and coefficients be available at the same time. Likewise, circular buffers are a natural method for accessing tables and coefficients. The Harvard architecture allows coefficients and data to be available in both program or data memory. Simultaneous fetches of two operands is necessary to make efficient use of this architecture.

Figure 3 shows the data address generators of the ADSP-2101 while Figure 4 shows the data address generators of the DSP16A.

ADSP-2101 Addressing

There are two independent data address generators (DAGs) in the ADSP-2101. One typically supplies addresses for program memory data fetches while the other supplies addresses for data memory. The ability to simultaneously fetch two operands makes the ADSP-2101's Harvard architecture very efficient. The address generators are completely separate from the program sequencer.

Each DAG has four I registers which store pointers (addresses), four M registers for address modifiers, and four L registers which store circular buffer lengths for modulo addressing.

DAG1 can bit reverse addresses as they are output to the address bus. This zero-overhead bit-reversing is useful for implementing FFTs. The 14-bit I, M and L registers can also be used for general purpose data storage.

ADSP-2101 Indirect Addressing

With indirect addressing, the address in an I register drives either the data or program memory address bus. As shown in

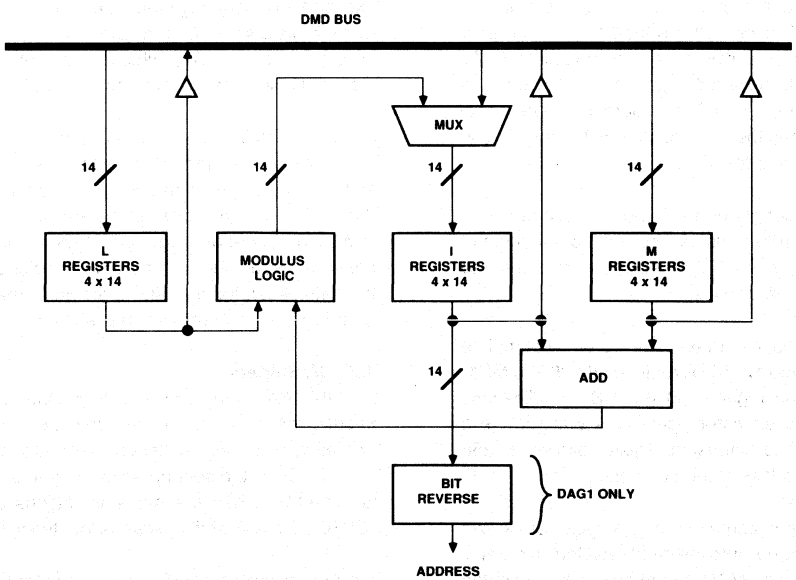

Figure 3. Block Diagram of ADSP-2101 Data Address Generators (DAGs)

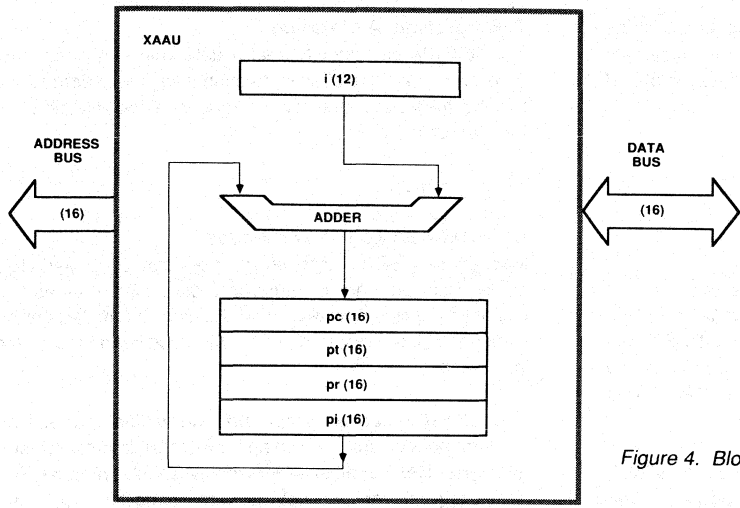

Figure 4. Block Diagrams of DSP16A Address Generators

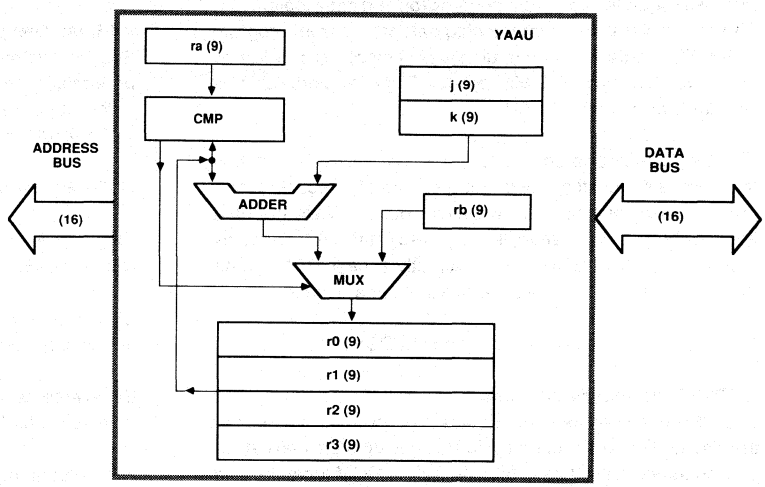

Figure 3, while memory is being accessed, The DAG updates the address simultaneously by adding to it the contents of any modify (M) register in the DAG. The specific pairing of the 14-bit I and M registers is up to the programmer. For example, I0 and M3 could be specified in the instruction as:

```
AX0=DM(I0,M3)  {Load AX0 from Data Memory}
```

The ability to mix I registers and M registers is especially useful for two-dimensional addressing or for supporting pointer increment and decrement operations without constantly loading different modify values. This instruction syntax shows explicitly what registers are used to generate the address and where the data is going; nothing is inferred.

The ADSP-2101 multifunction capability allows up to two data accesses and one arithmetic operation in a single instruction. This example instruction fetches two operands and performs an ALU operation:

```
AR=AX0+AY0(SS), AX0=DM(I2,M3), AY0=PM(M4,M5);
```

Loading the length of a circular buffer into the L register activates the modulus logic, guaranteeing that the address is kept inside the buffer in a modulo fashion. This structure is maintained automatically by the address generator hardware and does not have to be calculated explicitly by the programmer. Once initialized, circular buffers are used transparently and require no overhead instructions.

ADSP-2101 Direct Addressing

Due to the 24-bit width of the ADSP-2101 instruction, a full 14-bit address can be specified within a single-word instruction. This feature allows single cycle access to data located in data

memory. There is no paging or memory segmentation on the ADSP-2101. The programmer can specify an immediate address or a predefined variable. Below are some examples of direct addressing read instructions.

```
MX0=DM(beta);

AY1=DM(0x0FE3);
```

DSP16A Addressing

Like the ADSP-2101, the DSP16A contains two address generators - the ROM address arithmetic unit (XAAU) and the RAM address arithmetic unit (YAAU). Like the ADSP-2101, the ability to fetch two operands allows efficient use of the DSP16A's Harvard architecture. Only the YAAU can implement circular buffers.

The XAAU generates addresses for program memory ROM. It acts as an instruction sequencer for the CPU and as a data address generator. It has four 16-bit static pointer registers that are used for the program counter (pc), the subroutine return address (pr), interrupt return address (pi) and table lookup (pt). The pt register is used for fetching data words from program ROM. This register is automatically post modified by +1 or by the value stored in the 12-bit i register. There is no provision for circular buffers in the XAAU.

The YAAU is dedicated to addressing data RAM. It contains four pointer registers, two address modify registers and provisions for one modulo or circular buffer. While the data address bus is 16 bits wide, the registers in the YAAU are only 9 bits wide. The DSP16A can indirectly access 64K of data memory, but only 2K bytes are available at any time.

There is no bit-reverse capability on the DSP16A.

DSP16A Indirect Addressing

The DSP16A uses indirect addressing in both the XAAU and the YAAU. The XAAU contains four 16-bit registers that can be post-modified after data is fetched from a ROM location. The XAAU lacks circular buffering capabilities. If filter coefficients are stored in program ROM, the program must reset the pt register after every filter iteration.

The YAAU contains four 9-bit pointer registers used to indirectly address data RAM. The pointers can be post-modified by -1, 0, +1, +2 or by a value stored in one of two post-modify registers.

The DSP16A has provisions for one circular buffer, limited to 2K words. Only +1 is allowed as a post-modify value. The ra and rb registers contain the last and first address of the circular buffer. If an address register is post-modified to a value equal to that in the ra register, then the value in rb is written back into the address register.

DSP16A Direct Addressing

The DSP16A can directly load a data value from memory. Below is an example of an instruction using direct addressing to read from data memory (user-defined variable, ibuf) to register zero:

```
r0=ibuf
```

Other DSP16A Addressing Modes

Most of the DSP16A internal registers can be immediately loaded from program memory. In most cases, it takes two instruction cycles and two ROM locations to implement an immediate load. Registers in the YAAU can be loaded in one instruction.

The DSP16A has a compound addressing mode. The programmer can specify a swap between an internal register and a data RAM location addressed by a YAAU register. The YAAU register can be post-modified during this operation. Compound addressing requires two instruction cycles to execute.

Address Generation Summary

Sustained intensive arithmetic operations demand maximum performance from the data addressing section of a processor architecture. Table 2 summarizes the differences between the ADSP-2101 and the DSP16A.

DSP Requirement	ADSP-2101	DSP16A
Single cycle fetch of two operands	√	√
Generate new program memory and data memory addresses each cycle	√	√[1]
Bit reverse data memory addresses for FFT	√	**No**
Transparent, no overhead circular buffers	1-8	1
Maximum circular buffer size	16K words	2K words

1. The YAAU can directly access only 2K-bytes of data memory.

Table 2. Data Addressing Capabilities

PROGRAM SEQUENCING

Efficient architectures for signal processing require fast arithmetic and matching speed in data addressing and fetching capabilities. To fully deliver the performance required for real-world signal processing, a DSP machine must execute its program with little or no overhead spent on maintaining the proper program flow.

Efficiency in program sequencing has many different aspects. The comparison in this application note focuses on three areas:

- Execution of Loops
- Execution of branches and conditional instructions
- Processing of interrupts and subroutine calls

Loops are fundamental to the implementation of DSP algorithms. Many operations, such as the sum-of-products, are repetitive. If a program is efficiently expressed in looped form then coding is quite straightforward. In addition, modifying the program (for example, increasing the number of filter taps) is easy.

Branching on conditions is a natural way to code any program which must respond to its environment. Efficient response to interrupt requests is also important.

Figure 5 shows the architecture of the ADSP-2101 program sequencer. The program control section of the DSP16A is shown in Figure 6.

ADSP-2101 Program Sequencer

The program sequencer on the ADSP-2101 contains logic that selects a program memory address source and routes the address to the program memory address bus (PMA). This address selection occurs automatically in response to the

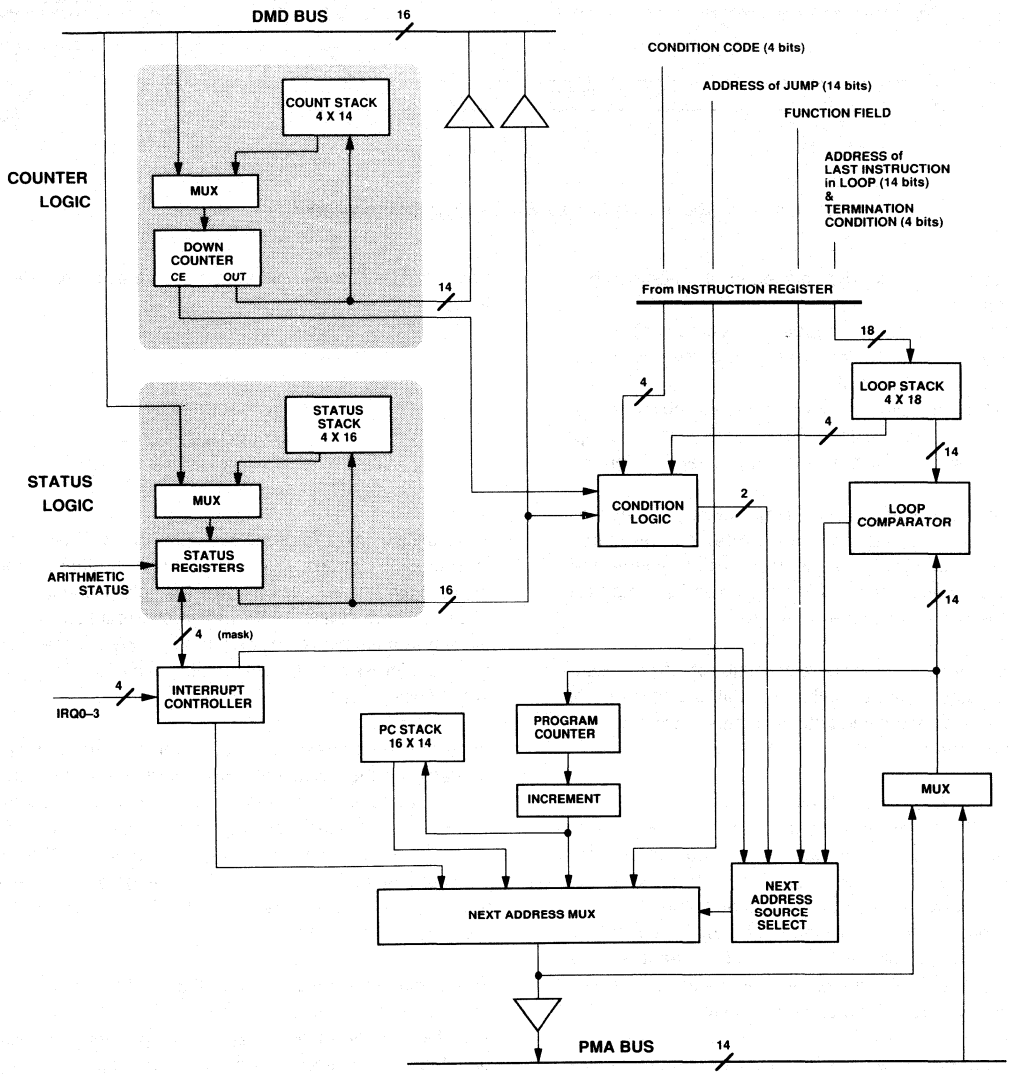

Figure 5. Block Diagram of ADSP-2101 Program Sequencer

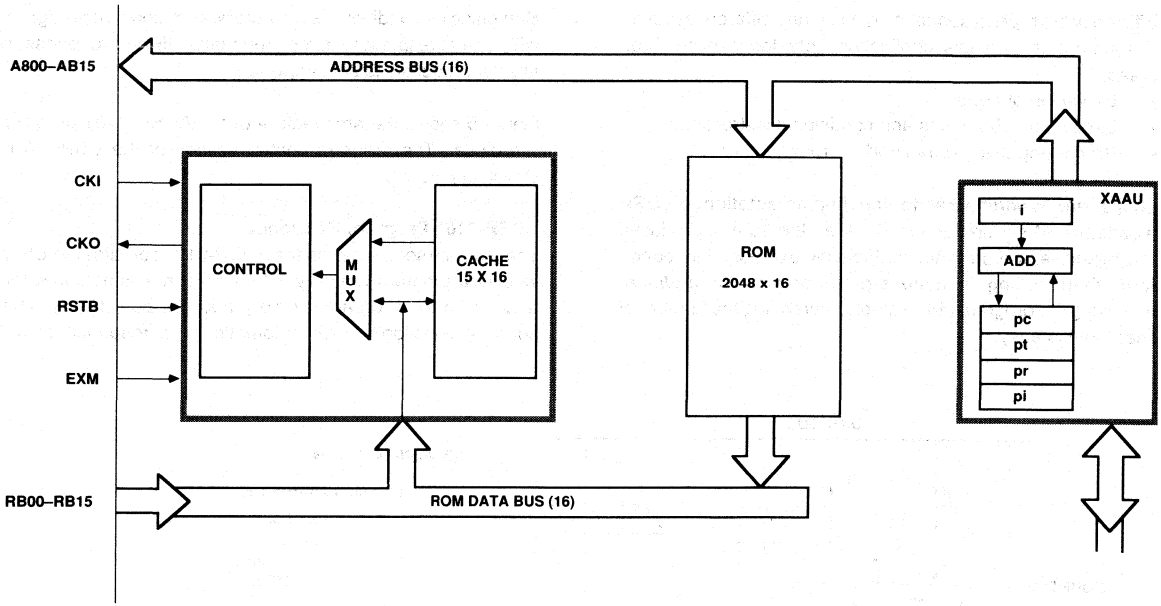

Figure 6. Block Diagram of DSP16A Program Sequencer

current instruction. The address placed on the address bus can come from

- The program counter (for sequential addressing),
- A 14-bit address in the instruction word itself, for direct jumps and subroutine calls,
- The PC stack, for returns from subroutines and interrupts,
- The interrupt logic, to automatically vector to the interrupt routine in response to assertion of any external interrupt.

All instructions can execute in a single cycle; this applies equally to jumps, calls, and interrupts. No pipelined instructions are required. Since the ADSP-2101 can fetch an instruction and access both on-chip program and on-chip data memory in one cycle, no cache is required.

When an interrupt occurs, the complete status of the processor (stack status, mode status, arithmetic status and interrupt mask) is automatically pushed onto the status stack as part of the interrupt vector process.

ADSP-2101 Looping Capabilities

The ADSP-2101 program sequencer can support zero-overhead loops. Using the loop stack and loop comparator, the processor can decide if the loop should terminate and determine the address of the next instruction (either the top of the loop or the instruction following the loop). These actions are performed in parallel with instruction execution, are transparent to the user and require no overhead cycles. Program loops can be nested up to four deep.

A DO UNTIL loop may be as large as program memory size permits. A loop may terminate when the 14-bit counter expires or when a specified arithmetic condition occurs. The example below shows a three instruction loop that is repeated 100 times:

```
CNTR = 100;
DO foo UNTIL CE;
        first loop instruction
        second loop instruction
foo:    third loop instruction
first instruction outside of loop
```

The first instruction loads the counter with 100, The DO UNTIL instruction contains the termination condition (counter expired) and the address *foo*, the last instruction in the loop. The execution of the DO UNTIL instruction causes the address of the first instruction of the loop to be pushed on the PC stack and the address of the last instruction of the loop to be pushed onto the loop stack.

As instruction addresses are output to the program memory address bus, the loop comparator determines if the instruction is the last instruction in the loop. If it is, the program sequencer checks the status and condition logic to see if the termination condition is satisfied. The program sequencer then takes either the address from the PC stack (to go back to the top of the loop) or simply increments the PC (to go to the first instruction outside the loop).

The looping mechanism of the ADSP-2101 is automatic and transparent to the user. Once the DO UNTIL instruction is specified, program flow, all stack updates, and counter updates are handled by the sequencer logic. No program overhead or extra cycles are required. After the last loop instruction is executed, the next instruction will be the first loop instruction or the first instruction following the loop.

he ADSP-2101 has up to six interrupts. Three of these terrupts are external to the processor. The interrupts are rioritized, maskable and can be edge or level sensitive. When n interrupt is asserted, the ADSP-2101 aborts execution of ie current instruction and vectors to the appropriate interrupt ervice routine. Because the interrupts are prioritized, oncurrent interrupts can be handled without external ardware.

he PC stack is 16 words deep. Up to seven nested interrupts re allowed. The size of the loop stack and the count stack llow four-deep nested loops.

onditional execution of instructions is available on the ADSP-101. All ALU, MAC and shifter instructions can be executed onditionally. The only exceptions are ALU division and nmediate shifts. All program jumps, subroutine calls and eturns can also be conditionally executed. All of these nstructions execute in a single cycle.

SP16A Program Sequencer

rogram sequencing in the DSP16A is controlled by the AAU, a 16 word cache and other control logic. The XAAU ontains the program counter, an interrupt return register and subroutine return register. The program counter can address p to 64K of program memory ROM.

he DSP16A can implement one level of program looping with ne internal cache. The loop is limited to 15 bytes of instruction nd cannot execute more than 127 times. An advantage of the nternal cache is that instructions that require a ROM memory ead execute more efficiently. However, interrupt requests are ot recognized if the DSP16A is executing instructions from ne cache.

he example below shows a two instruction loop that executes 00 times:

```
do 100 {
        first loop instruction
        second loop instruction
       }
```

here are two classes of instructions that can be executed onditionally: special-function instructions and branch nstructions. Special-function instructions consist of data hifts, accumulator increments and some data transfer perations. These instructions can be conditionally executed ased on the status of the multiplier/ALU status flags or a seudo-random sequence bit. Special-function instructions xecute in one cycle.

Branch instructions include goto, subroutine calls and return nstructions. These instructions require two cycles to execute. f the instruction is conditional, three cycles are required. Since here is only one subroutine return register, extra instructions are required to implement nested subroutine calls.

There are four internal and one external interrupts on the DSP16A. The internal interrupts are asserted by the parallel and serial interface ports. In addition, there is one software interrupt available. All interrupts are maskable. When an interrupt is recognized, the DSP16A completes the current instruction and vectors to the interrupt service routine. Note that the processor supports only one service routine. The service routine must determine which interrupt was processed. The DSP16A does not recognize interrupts when executing branch instructions or instructions in the cache. Interrupts are not prioritized. Providing for concurrent interrupts requires extra hardware and/or programming.

Program Sequencer Summary

Efficient looping capabilities are very import for DSP algorithms due to their repetitive nature. Table 3 summarizes the program sequencing capabilities of the ADSP-2101 and the DSP16A.

DSP Requirement	ADSP-2101	DSP16A
Zero-overhead looping	√[1]	√[2]
Conditional arithmetic instructions	√	√[3]
Single-cycle branching	√	**No**
Zero-overhead nested interrupts	Up to 7 deep	**No**

1. Four levels of nested loops.
2. One level of looping; up to 15 instructions long; can execute only 127 times.
3. Shifts and accumulator increments only.

Table 3. Program Sequencing Capabilities

PROGRAMMING EXAMPLE

A good illustration of the differences between the ADSP-2101 and the DSP16A is the implementation of the stochastic gradient algorithm for updating taps in an adaptive filter.

$$C_{n+1}=b*e*Y_n$$

A pseudo-code description of this algorithm is shown for a 256 tap filter.

```
MOD=beta*e;
Loop n=0 to 255;
   C(n+1)=Cn MOD*Yn;
EndLoop;
```

When implementing this algorithm you should consider the following:

How many cycles are required for loop overhead?

How many cycles are required for loading and storing operands?

How many extra cycles are required to move intermediate results?

Listing 1 shows an ADSP-2101 code segment for updating the taps in a 256-tap adaptive filter. The coefficients are stored in program data memory and the sample data is stored in data memory. Efficient operand manipulation is achieved by storing the MOD term in the multiplier feedback register. In the loop, data is fetched in parallel with instruction execution. No cycles are wasted fetching or storing data.

Listing 2 shows a DSP16A code segment for updating the taps in a 256-tap adaptive filter. Although data operands can be fetched from program memory ROM, is not possible to write updated values back to program ROM. Therefore, both coefficients and data must be stored in data memory. Since the multiplier output register is not available as an input for another multiplication, two extra instruction cycles are required to move the MOD term back into the x register. The program loop contains two instructions. Three clock cycles are required to execute this loop.

Interrupt requests are ignored when the DSP16A is executing from a cache loop. The code segment shown below periodically stops the loop so pending interrupt requests can be serviced.

The total number of clock cycles for each algorithm are:

	N Taps	*256 Taps*	*Time*
DSP16A	8N + 25	2073	69 µs (33 ns cycle)
ADSP-2101	3N + 15	783	63 µs (80 ns cycle)

CONNECTING THE PROCESSOR TO OTHER DEVICES

Although digital signal processors that have efficient architectures, fast arithmetic capabilities and efficient data addressing are well suited for high performance systems, a DSP system designer must also consider how the microprocessor communicates with other components in the system. These devices can include other DSP processor, a host processor, or peripherals such as CODECs, A/D converters and D/A converters. The configuration of off-chip memory and the size of on-chip memory can significantly affect the cost of a DSP system.

Connecting the ADSP-2101 to Other Devices

The ADSP-2101 has two bidirectional double-buffered serial ports. These ports can be used to communicate with serial devices such as CODECs, serial A/D converters and serial ports on other ADSP-2101s. The ADSP-2101 can communicate with parallel devices by memory-mapping them into external data memory space.

The serial ports (SPORTs) are synchronous and use framing signals to control data flow. Each SPORT can generate its own serial clock internally or use an external clock. The framing sync signals may be generated internally or by an external device. Word lengths may vary from three to sixteen bits. The serial ports can have a zero chip interface to other serial devices and can perform single cycle µ-law and A-law companding in hardware.

SPORT0 has a multichannel capability which allows the receiving or transmitting of arbitrary data words from a 24-word or 32-word bitstream. This is useful for implementing a T1 interface or a networking scheme for multiple processors.

SPORT1 may be optionally configured as two external interrupts, $\overline{IRQ0}$ and $\overline{IRQ1}$, and the Flag In and Flag Out signals instead of as a serial port.

The SPORTs also have an autobuffering capability. This feature allows a block of data to be loaded into memory while the ADSP-2101 is executing program code. When all data is loaded, the SPORT interrupts the processor.

The ADSP-2101 can interface to parallel devices such as parallel DACs and UARTs through external data memory. Software-controlled wait states in the ADSP-2101 allow for a simple interface to these devices. The only additional circuitry that may be required is address decode logic.

Connecting the DSP16A to Other Devices

The DSP16A has two ports for communicating with the other devices. The SIO is a serial communications port. The PIO is designed for parallel communications. Either the SIO or the PIO is required to load internal RAM from off-chip devices, such as external memory or a host processor.

The SIO is a synchronous port. The framing signals are ILD and OLD and must be synchronized to the input and output clocks respectively (ICK and OCK). Framing and clock signals can be generated internally or externally. The internal clock is programmable to one of four different clock speeds. Word lengths are limited to 8 or 16 bits. A time division multiplex (TDM) communications scheme can be used to communicate with up to seven other DSP16As.

The PIO interfaces to AT&T CODECs without any additional circuitry. A bit-reversing capability on the PIO allows for an easier software implementation of companding operations.

The DSP16A PIO is a parallel port that can be programmed to transmit or receive 16-bit data words or simultaneously transmit and receive 8-bit data words. The PIO can operate in either the active or the passive modes. In the passive mode, an external device provides the I/O strobe signals. In active mode, the strobe signals are provided by the DSP16A. The width of the strobe signals is programmable allowing connections to slow peripherals.

Memory Configuration on the ADSP-2101

The ADSP-2101 contains 2K words of on-chip program memory and 1K words of on-chip data memory. In a single cycle, the ADSP-2101 can access on-chip program RAM twice and on-chip data RAM once. All instructions execute in a single cycle when the processor is executing from on-chip memory.

The ADSP-2101 can address 16K words of program memory and up to 16K words of data memory. An additional 16K words of data memory are available with the PMDA pin. The processor can be configured with all memory off-chip or with a

```
        {                    DAG Initialization                    }
        {                                                          }
        {I2->Oldest input data value in delay line}
        {L2=Filter Length}
        {I6->Beginning of filter coefficient table}
        {L6=Filter Length}
        {M1,M5=1}
        {M6=2}
        {M3,M7=-1}

.ENTRY affir;
strt: CNTR=< Filter Length >
        MX0=< error >
        MY1=beta;
        MF=MX0*MY1(RND), MX0=DM(I2,M1);      {MF=error*beta}
        MR=MX0*MF(RND), AY0=PM(I6,M5);

        DO adapt UNTIL CE;
              AR=MR1+AY0, MX0=DM(I2,M1), AY0=PM(I6,M7);
adapt:  PM(I6,M6)=AR, MR=MX0*MF(RND);

        MODIFY(I2,M3);                       {Point to oldest data}
        MODIFY(I6,M7);                       {Point to start of table}
```

Listing 1. ADSP-2101 Code Segment for Stochastic Gradient Algorithm

```
        j=0
        k=1
        r0=beta                                        /*Load beta and e*/
        r1=e
                                          x=*r0
                                          y=*r1
                    p=x*y                              /*Calculate MOD*/
        a0=p
        x=a0                                           /*Store MOD in x*/
        r0=<address of coefficients C_n>               /*Init pointers*/
        r1=<address of data A_n>
                                          a0=*r0
                                          y=*r1++
                    p=x*y
 DO 2 {                                                /*Program loop */
        a0=a0+p                  y=*r1++
                    p=x*y        *r0jk:a0
        }
 REDO 127                                              /*Service Interrupts*/
```

Listing 2. DSP16A Code Segment for LMS Stochastic Gradient Algorithm

portion of memory on-chip. For off-chip memory, the program and data memory lines are multiplexed into a single off-chip address bus and a single off-chip data bus. Extra cycles may be required to execute instructions when more than one piece of data is located off-chip.

The boot capability of the ADSP-2101 allows loading internal program memory directly from an inexpensive boot EPROM. On reset, up to 2K words of instructions can automatically be loaded into on-chip program memory. Up to seven more 2K instruction blocks can be paged into the processor on demand. If desired, the boot code can contain a short program for downloading data and program code through the serial ports.

Software programmable wait states on the ADSP-2101 makes it compatible with slow external RAM or ROM. Using slower memories reduces overall system cost at the expense of reduced execution speeds. Regardless of the speed of external memory, the ADSP-2101 always executes at full speed when accessing on-chip program and data memories.

Memory Configuration on the DSP16A

The DSP16A contains 2K words of on-chip data memory and 4K words of on-chip program memory ROM. In one cycle, the DSP16A can access data RAM and instructions in the program memory ROM. A second cycle is required to fetch data from program ROM.

The DSP16A can address up to 64K of program ROM. Like the ADSP-2101, the part can be configured to use off-chip memory only or a combination of on-chip and off-chip program memory. Since there is no off-chip data bus, data memory cannot be expanded off-chip. The data memory address space can be expanded through the PIO port. However, this requires external circuitry and additional software.

Because the DSP16A is designed to execute ROM code, there are no read/write strobes associated with the program memory address and data lines. This makes it difficult to use RAM for program memory. Extra hardware is required to download code from a host processor.

The internal ROM on the DSP16A must be programmed by the factory. The extra costs associated with programming internal ROM make it impractical for low-volume applications. It is also expensive to fix bugs or upgrade code once the device is programmed. In applications with external program memories, the high clock speeds of currently available DSP16As require extremely fast external EPROMs. The 33 ns version requires EPROMs with 15 ns access times. A system that uses fast versions of the DSP16A will be very expensive to design and manufacture. There are no programmable wait states on the DSP16A.

SUMMARY

Digital signal processing is a specialized branching of processor design and application. Table 4 summarizes the fundamental requirements of DSP.

The DSP processors available today vary drastically in thei ability to meet the requirements described above. Some processors are optimized for very specific applications while others are flexible and allow for future growth. Analyzing the requirements of your DSP system and matching them to the capabilities of a DSP architecture assures efficient operation

Due to space limits, this application note does not cover many topics in detail. Consult the *ADSP-2101 User's Manual* and the *ADSP-2101 Cross-Software Manual* for a greater depth o information on this processor.

DSP Requirement	ADSP-2101	DSP16A
Fast arithmetic	√	√
Extended dynamic range on multiplication/accumulation	√[1]	√[2]
Hardware circular buffering (both on- and off-chip)	√	**No**
Zero overhead looping & branching	√	**No**

1. Accumulator may overflow 256 times.
2. Accumulator may overflow 16 times.

Table 4. Fundamental Requirements of DSP

ANALOG DEVICES

AN-233
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Considerations for Selecting a DSP Processor (ADSP-2101 vs. TMS320C50)

by Bob Fine and Gerald McGuire

INTRODUCTION

Digital signal processing systems demand high performance processors. But high performance cannot be measured by a processor's multiplication/accumulation speed or MIPS (Millions of instructions per second) rating alone. Many times a DSP processor is characterized mainly by its MIPS rate. Since the instruction of one DSP device is not necessarily equivalent to that of another DSP device, a MIPS rating can be misleading. Other architectural and performance requirements relating to a DSP processor's capabilities in areas such as arithmetic, addressing and program sequencing may be more important. What distinguishes DSPs from other types of microprocessor and microcontroller architectures is how well they perform in each of the following areas.

1. Fast and flexible arithmetic

   A DSP processor must provide single-cycle computation for multiplication, multiplication with accumulation, arbitrary amounts of shifting, and standard arithmetic and logical operations. In addition, the arithmetic units should allow for any sequence of computation so that a given DSP algorithm can be executed without being reformulated.

2. Extended dynamic range on multiplication/accumulation

   Extended sums-of-products are fundamental to DSP algorithms. Protection against overflow in successive accumulations ensures that no loss of data or range occurs.

3. Single-cycle fetch of two operands (from either on- or off-chip)

   Again, in extended sums-of-products calculations, two operands are always needed to feed the calculation. A processor must be able to sustain two operand data throughput. Also, flexible addressing capabilities for multiple data memories is important.

4. Hardware circular buffering (both on- and off-chip)

   A large class of DSP algorithms including most filters require circular buffers. Hardware to handle address pointer wraparound or modulo addressing reduces overhead (increasing performance) and simplifies implementation.

5. Zero overhead looping and branching

   DSP algorithms are naturally repetitive and can easily be expressed as loops. Program sequencing that supports looped code with zero overhead provides the best performance and the easiest programming implementation. Likewise, overhead penalties for conditional program flow are unacceptable in signal processing applications.

Not all processors currently used for DSP and DSP-like functions meet these architectural and performance requirements equally well. This article examines these considerations for selecting a DSP processor, comparing two 16-bit fixed-point processors, the ADSP-2101 from Analog Devices and the TMS320C50 from Texas Instruments.

The three sections that follow discuss the five points above. The arithmetic section discusses items one and two, the addressing capabilities sections discusses items three and four and the program sequencing section discusses item five.

Program examples and benchmarks can be found at the end of this article.

ARITHMETIC CAPABILITIES

The basis of a successful DSP implementation is the ability to perform fast math. Arithmetic capabilities are the foundation of DSP performance.

General Purpose Math

One indicator of a good arithmetic architecture is the ability to perform a wide range of arithmetic computations. These

computations should be handled in a flexible manner so that the algorithm can be implemented without rearranging the order of the arithmetic operations or operands. If the arithmetic architecture is fixed, too special-purpose or limited and the algorithm must be rearranged, this poses extra work for the DSP designer or programmer and delays getting a system running. Algorithm development frequently turns out to be much of the work of implementing a DSP system. If an algorithm can be used "as is" with no extra work, the design can be finished sooner and with less chance of error.

Arithmetic Architecture

Figure 1 shows a block diagram of the arithmetic section of the ADSP-2101 while Figure 2 shows that of the TMS320C50. Both of these devices utilize a modified Harvard architecture which can feed data operands from both program memory and data memory to the arithmetic section. Both of these devices work with 16-bit numbers.

ADSP-2101 Arithmetic Architecture Overview

The ADSP-2101 has three independent computational units: an ALU, a multiplier/accumulator (MAC), and a barrel shifter. They are connected (via the Result bus) so that the output register of any arithmetic unit may be operated on directly as an input by any other unit. In addition, the ALU and MAC are directly connected to both the program and data memory buses. Operands for ALU and MAC operations can come from both memories or any combination of off-chip memory and other data registers in the processor. All arithmetic operations are register based and a group of registers surrounds each arithmetic unit. A primary and secondary bank of registers is available to provide for fast context switching. All arithmetic registers can also be used as general purpose data registers.

TMS320C50 Arithmetic Architecture Overview

Figure 2 shows the block diagram of the arithmetic section of the TMS320C50. The TMS320C50 contains a multiplier, an ALU, a Parallel Logic Unit (PLU), a 16-bit scaling shifter and additional shifters at the outputs of both the accumulator and multiplier. The multiplier has an input register, TREG0, and an output register, PREG. The multiplier has direct input connections to both the program and data bus only for one operand or input. The ALU has direct access to only the data bus, not the program bus. Results are always sent to either the data bus or the accumulator registers. In some cases, the result must first be stored back in data memory before it can be used as an input for another calculation. Operations such as adding two data values from memory or multiply/accumulating with a data value can require multiple cycles.

As with the TMS320C25, there is no dedicated multiplier accumulator (MAC), which is required in many DSP algorithms. Instead the ALU must be used in conjunction with the multiplier for MAC operations. This may require some rearrangement of the algorithm or the temporary storage of intermediate results in data memory if the algorithm requires

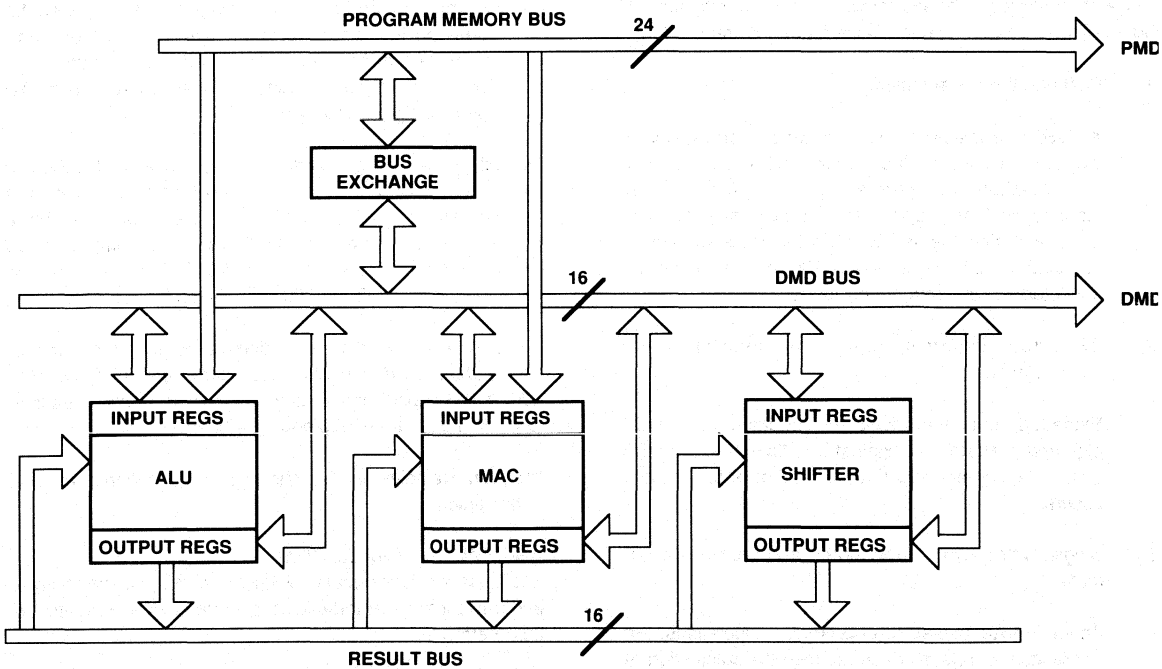

Figure 1. Block Diagram of Arithmetic Section of the ADSP-2101

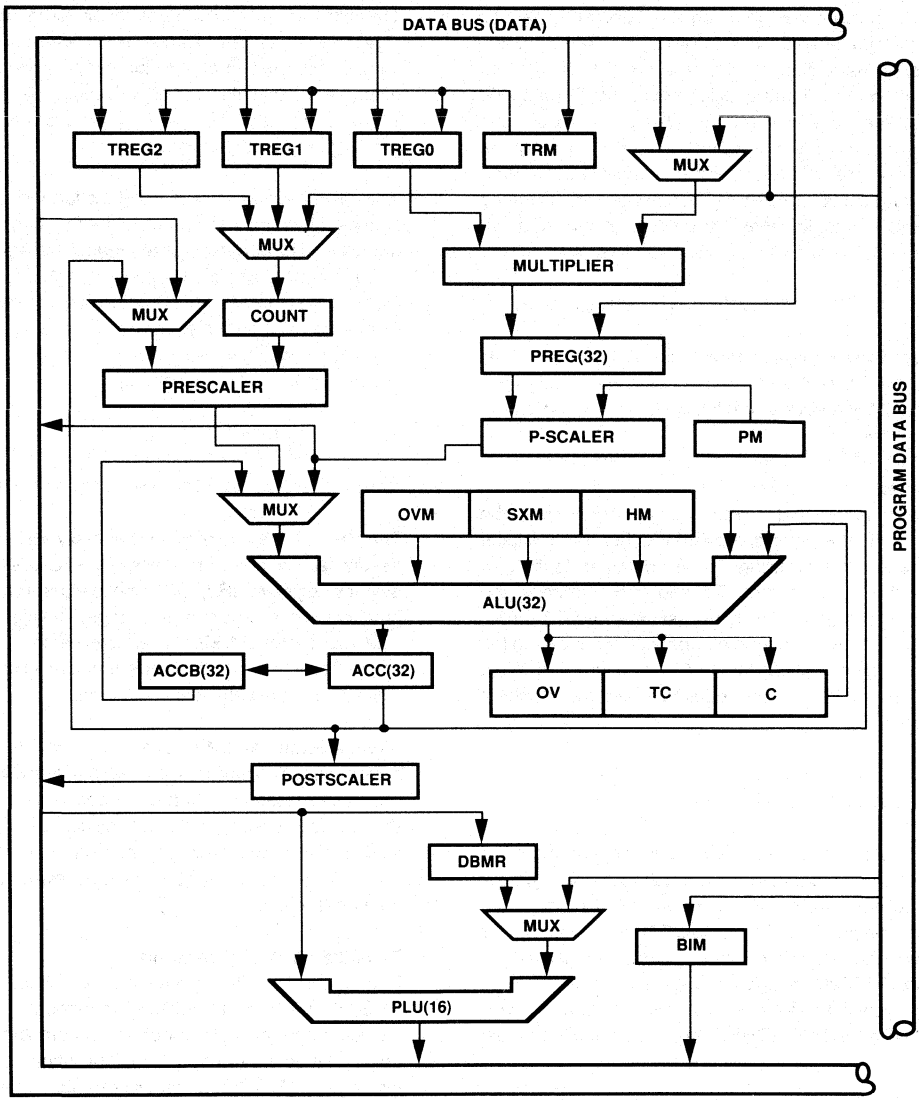

Figure 2. Block Diagram of Arithmetic Section of the TMS320C50

MAC operations interleaved with ALU operations. Also, there are arithmetic pipeline delays that are required to achieve sustained MAC operations. Basic multiply and ALU operations require multiple cycles as opposed to the single cycle operation of the arithmetic units in the ADSP-2101.

The availability of general purpose data registers and the flexibility of data movement in the TMS320C50 is limited. This may result in data bottlenecks and in extra cycles being required to move data into the right position prior to an arithmetic operation.

ADSP-2101 ALU

The ALU has two X and two Y input registers: AX0, AX1, and AY0, AY1. ALU operations are performed on any X-Y assortment of these input registers. They may be loaded from any combination of program and data memory or other data registers in the processor. The result of the operation appears in the ALU result (AR) or ALU feedback (AF) register. AR and AF can also be used as the X and Y operands (respectively) in any ALU calculation. In addition, the result registers of the MAC and barrel shifter can also be used directly as X inputs to the ALU (and vice versa).

ALU instructions are coded in a register transfer, algebraic syntax. An example of addition is shown below. This example is a multifunction instruction. The first "clause" of the instruction (up to the first comma) is the addition operation. The second clause loads the X input register from data memory ("DM") and the third clause loads the Y input from program memory. An addition (or any other ALU operation) can be executed on a sustained, single-cycle basis. (These operand fetching clauses of the instruction may be omitted, if they are not needed.)

```
AR = AX0 + AY1, AX0=DM(I0,M0), AY1=PM(I4,M4)
```

All ALU operations complete in a single 60 ns cycle. (All references to cycles for the ADSP-2101 assume a 16.67 MHz device.) The ADSP-2101 runs at full speed even with an off-chip memory access.

TMS320C50 ALU

ALU operations require that one operand must come from the accumulator while the other comes from either the multiplier output, the accumulator buffer, or from the data bus or accumulator through a shifter. To add two numbers, the accumulator must be loaded with the first data value. After the accumulator is loaded, a second number can be added to the accumulator. The instructions for the ALU are specified with a mnemonic. The two instructions required to add two numbers are shown below.

```
ZALR    <data memory address>
ADD     <data memory address>
```

For the result to be used as an input value for anything other than another ALU operation, the data must first be stored back into data memory from the accumulator. Not all ALU operations can be performed in a single 50 ns cycle; an add as shown above can be accomplished every two cycles. All references to TMS320C50 cycles assume a 20 MHz device with an 50 ns cycle time. Not all ALU instructions (i.e., ADD #k, SUB #k, ADD #lk, SUB #lk, ADRK) can be used with the repeat feature.

ADSP-2101 MAC

As shown in Figure 1, the ADSP-2101 multiplier/accumulator (MAC) sits next to the ALU. Like the ALU, it has two X and two Y input registers, MX0, MX1 and MY0, MY1. The unit performs both multiplications and MACs independent of the ALU. This is a key difference from the architecture of the TMS320C50.

MAC operations are performed on any X-Y assortment of input registers. They may be loaded from any combination of program and data memory or other data registers in the processor. The result of the operation appears in the MAC result register (MR) or the MAC feedback register (MF). Like the ALU, the feedback and result registers can also serve as the X and Y inputs for any multiplication or MAC operation. In addition, the result registers of the barrel shifter and ALU can also be used directly as X inputs to the MAC (and vice versa).

The instructions for the MAC are specified in a register transfer, algebraic syntax. An example is shown below. The first line shows multiplication of two signed operands and the second example shows multiplication with accumulation of one signed and one unsigned operand. (Signed and unsigned operands can be mixed in any combination.)

The second example is a multifunction instruction. The first "clause" of the instruction (up to the first comma) is the MAC operation. The second clause loads the X input register from data memory ("DM") and the third clause loads the Y input from program memory. Any MAC operation can be executed on a sustained, single-cycle basis. (These operand fetching clauses of the instruction may be omitted, if they are not needed, as in the first example.)

```
MR=MX0*MY0 (SS)

MR=MR+MX1*MY1 (SU), MX1=DM(I0,M0), MY1=PM(I4,M4)
```

The MR (MAC result) register is actually a 40-bit accumulator. It is divided into two 16-bit pieces (MR0 and MR1) and an 8-bit overflow register (MR2). DSP applications frequently deal with numbers over a large dynamic range. The eight "overflow" bits of MR2 allow for 256 MAC overflows before a loss of data can occur. The MAC also supports multiprecision operations as well as automatic unbiased rounding.

All multiplication and MAC operations execute in a single 60 ns cycle. (Please consult an *ADSP-2101 Data Sheet* for the most recent specifications.) Two new operands can be loaded into the input registers in parallel with the computation so that a new MAC operation with new operands can be started every cycle. The ADSP-2101 runs at full speed even with an off-chip memory access.

TMS320C50 MAC Operation

There is no dedicated multiplier/accumulator hardware in the TMS320C50. The TMS320C50 requires the use of both the multiplier and the ALU to perform a complete multiplication/accumulation operation. A multiplication is performed by loading the TREG0 register with the first operand. Once this data is loaded, a value from the data bus can be multiplied with the value in the TREG0 register. The instructions for the multiplier are specified with a mnemonic. The instructions for a multiplication are shown below.

```
LT      <data memory address>
MPY     <data memory address>
```

A product is obtained every two cycles.

A full multiplication/accumulation requires the use of the ALU as well as the multiplier. The instruction required to perform a MAC operation is shown below. This instruction requires two words of program memory storage.

```
MAC     <prog. mem. address> <data mem. address>
```

With both operands in on-chip memory, the MAC instruction takes three 50 ns cycles in non-repeat mode. In repeat mode, it will require *2 + n* cycles, where *n* is the number of repeats.

There are four different mnemonics used for the multiply/accumulate function: MAC, MACD, MADD, MADS. The specific use of each of these depends upon the source of the data. For a dual operand fetch, such as that needed for a digital filter, the MADD instruction should be used. The DMOV portion of the MADD instruction will not function with external memory. All data must reside on chip.

The TMS320C50 provides one bit of extension in the accumulator (a 31-bit accumulator with an overflow bit compared to the 40-bit accumulator of the ADSP-2101). After more than one overflow, the calculation of the TMS320C50 is corrupted. Automatic rounding is not supported in the multiplier. This is unlike the ADSP-2101, where up to 256 overflows can occur with no lost data and automatic rounding is performed in the same cycle as the multiply operation.

ADSP-2101 Shifter

The barrel shifter in the ADSP-2101 has an input register, SI, and accepts as inputs any result registers in the processor (e.g., MR1, AR) including its own result register, SR. Like the MAC result register set, the 32-bit SR is divided into two 16-bit registers, SR0 and SR1. The shifter also has an exponent register, SE, which is set automatically by the exponent adjust instructions and used for normalization instructions.

The shifter can place a 16-bit input value anywhere within a 32-bit field in a single cycle. The input can be shifted any number of bits from off-scale left to off-scale right with either an arithmetic or logical shift. Other functions such as exponent detection, normalization, denormalization, block floating-point exponent maintenance, and pattern merging can also be performed with this shifter. All shifter operations are performed in a single cycle. Numbers can be normalized, regardless of the number of bits to be shifted, in a single cycle.

TMS320C50 Shifter

The TMS320C50 has three scaling shifters. The P-scaler shifts the product 0, 1, or 4 bits to the left or 6 bits to the right. The prescaler at the input of the ALU shifts data to the left or right from 0 to 16 bits. The post-scaler at the output of the ALU can shift data coming from the accumulator left from 0 to 7 bits. These shifters add the advantage of being able to scale data during the data move instead of requiring an additional shifter operation.

Arithmetic Summary

Table 1 summarizes the comparison of arithmetic capabilities of these processors.

The side-by-side arithmetic architecture of the ADSP-2101 results in easier implementation of many DSP algorithms as compared to the fixed sequence, end-to-end architecture of the TMS320C50. Due to the dependency of the ALU on the multiplier for multiplication/accumulations in the TMS320C50,

DSP Requirement	ADSP-2101	TMS320C50
All ALU operations – single cycle	✔	No
Single-cycle multiplication	✔	No
Single-cycle MAC operations	✔	✔*
Single-cycle shifting	0-32 bits left or right	0-16 bits left or right 0-7 bits left 1 or 4 bits left 6 bits right
Accumulator overflow protection	8 bits	1 bit
Signed, unsigned or mixed-mode multiplications	✔	No mixed mode
Single-cycle normalization	✔	No

*Approaches single-cycle efficiency when using repeat mode

Table 1. Summary of Arithmetic Capabilities

MAC operations cannot be easily intermingled with ALU operations. This may require changing the order of calculations in an algorithm so that the interdependency of ALU and multiplier does not cause a problem. The local storage registers found in the ADSP-2101 make data movement for calculations easy. If data is to be used many times, it can reside in a register to eliminate the need of fetching it from memory each time. With local registers and the open architecture, it is easy to perform arithmetic operations in any order and to guarantee that input operands and results remain intact until explicitly overwritten or moved.

DATA ADDRESSING CAPABILITIES

A digital signal processor's ability to perform fast arithmetic is wasted if the required data cannot be fetched at sustained speed equal to the processing rate. Addressing hardware must support the dual operand fetches required to fully utilize the Harvard architecture found in most DSPs. A good DSP must have the ability to store two types of data operands, typically a coefficient and a data word. Maximum efficiency can be obtained if two different memory spaces are provided for the data operands so that two operands can be fetched in the same single cycle. Using both data memory and program memory to store data will allow maximum efficiency. Circular buffers are frequently useful in implementing DSP algorithms; hardware support of address pointer wraparound is another feature distinguishing a signal processor from other types of high-performance processors.

Figure 3 shows the address generation circuitry of the ADSP-2101 while Figure 4, on page 8, shows that of the TMS320C50. The addressing capabilities of the TMS320C50 are basically the same as those of the TMS320C25 with the addition of some circular buffering logic. Flexibility is still limited since there is only one modify register (AR0) and only two simultaneous circular buffers are supported compared to the eight modify registers and eight simultaneous circular buffers of the ADSP-2101. Also, due to instruction pipelining of the TMS320C50, the auxiliary registers cannot be used for as many as two cycles after certain register load instructions. These addressing limitations and restrictions of the TMS320C50 can present severe penalties in data addressing efficiency for DSP algorithms and can result in data bottlenecks and slower execution of DSP code.

ADSP-2101 Addressing

There are two independent address generators in the ADSP-2101. One typically supplies addresses for program memory data fetches while the other handles data memory, making efficient use of the modified Harvard architecture. Each

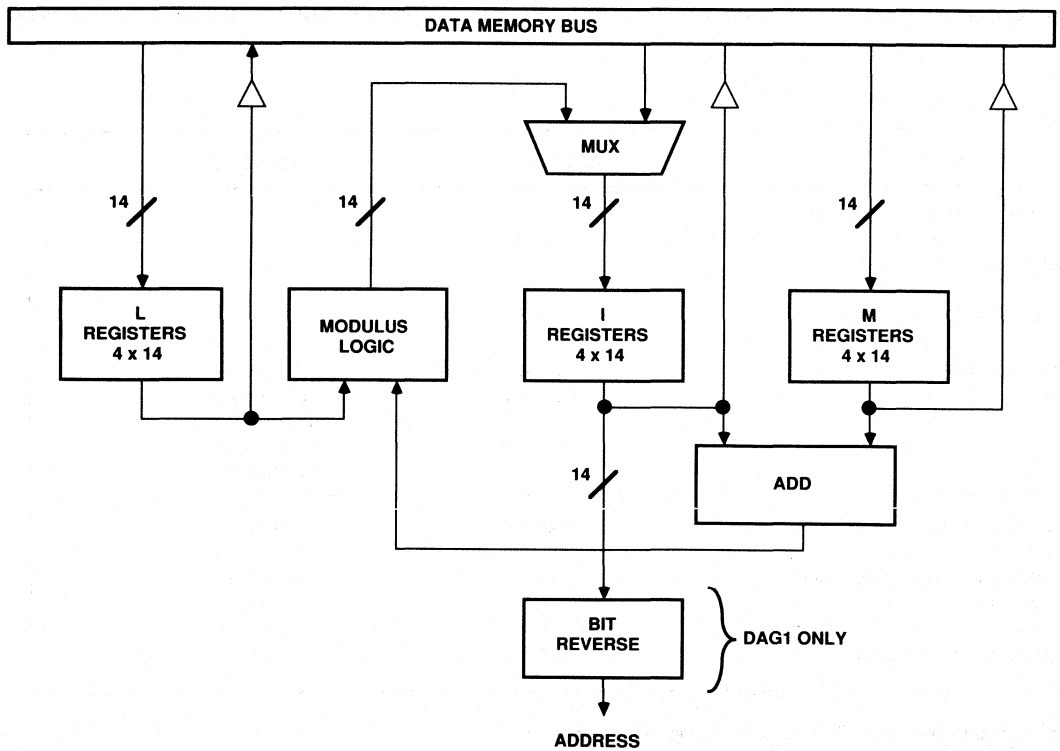

Figure 3. Block Diagram of ADSP-2101 Data Address Generators

address generator has four I (index) registers which store pointers (addresses), four M (modify) registers for address modifiers, and four L (length) registers storing buffer lengths for modulo addressing of circular buffers.

The address generator can bit-reverse an address as it is sent out to the address bus for zero-overhead bit-reversing for the FFT. The I, M, and L registers can be also used for general purpose data storage.

The address generators can also be used in conjunction with the serial ports to provide an automatic data buffering function. As data words come in or go out the serial port, data buffer addressing is automatically maintained and an interrupt is generated when the buffer is full or empty. This minimizes interrupt handling for serial port data transfers.

ADSP-2101 Indirect Addressing

With indirect addressing, the address in an I register drives either the data or program memory address bus. While the memory is being accessed, the address is simultaneously updated with the contents of any of the modify (M) registers, as shown in Figure 3. The specific pairing of I and M registers is up to the programmer. For example, I0 and M3 could be specified in the instruction as in

```
AX0 = DM(I0,M3);        {load AX0 from Data Memory
                         and modify I0 by M3}
```

The ability to mix I registers and M registers is especially useful for two-dimensional addressing or for supporting pointer increment and decrement without constantly reloading a new modify value. This instruction syntax shows explicitly what registers are used to generate the address and where the data is going; nothing has to be inferred.

Loading the length of a circular buffer into the L register activates the modulus logic, guaranteeing that the address is kept inside the buffer in a modulo fashion. This is maintained automatically by the address generator hardware and does not have to be calculated explicitly by the programmer. Circular buffers, such as for the delay lines of digital filters, are both transparent and require zero-overhead. Circular buffering is automatically maintained regardless of the modify value used.

ADSP-2101 Direct Addressing

Due to the 24-bit width of the ADSP-2101 instruction, a full 14-bit address can be specified within a (single-word) instruction for single-cycle access to any data. Figure 5 illustrates this. Below is an example of an instruction using direct addressing to read from data memory.

```
MX0 = DM(some_label);
```

ADSP-2101 Circular Buffering

Circular buffering is supported in hardware by the address generators of the ADSP-2101. Each address generator can maintain four simultaneous circular buffers for a total of eight. Circular buffers can be placed in either data or program memory. The length register (L registers) is simply loaded with the length of the circular buffer. The modulus logic detects when the pointer (updated index register value) has reached or exceeded the end of the buffer boundary. Operation is supported for going forwards or backwards through the buffer. The step size can be of any value that is less than the full buffer length. For applications such as interpolation filters, where the step size is equal to the interpolation factor, zero-overhead circular buffer operation is maintained.

TMS320C50 Addressing

The auxiliary register file of the TMS320C50 is used for storage of addresses and a single modifier. Only one address can be supplied at a time with the auxiliary register file so that two general purpose, indirect addressed data fetches cannot be achieved in a single cycle.

TMS320C50 Indirect Addressing

The auxiliary register file is connected to an arithmetic unit which will auto-index the contents of the auxiliary register or modify a register by the contents of auxiliary register number 0. The TMS320C50 has a single modify register. This limits the addressing capabilities for indirect addressing. Limited support is provided for circular modulo addressing; this diminishes the performance of DSP algorithms using circular buffers. Automatic circular buffering is only supported for increment and decrement address modifications. Modify values greater than 1 will not work.

TMS320C50 Direct Addressing

The TMS320C50 can directly access data within a 128-word block (compared to a 16K word block with the ADSP-2101). A 9-bit data page register is used in conjunction with the direct address to access a larger data space. To access data within a different block requires software overhead to update the 9-bit data page register. The update of the page register poses the requirement on the programmer to detect when the page boundary has been exceeded and when it is necessary to update the page register.

TMS320C50 Circular Buffering

Two circular buffers can be maintained by hardware in the address generation circuitry. A register (CBSR) is used to hold the start address of the circular buffer and a register (CBER) is used to hold the end address of the circular buffer. Since the auxiliary registers are used for pointers into the circular buffer, circular buffers in program memory (coefficients) are not possible. The circular buffer logic in the TMS320C50 checks only for a pointer equal to the end address, it does not check for a pointer that has skipped over the end address (i.e., using a step size greater than 1). For applications which require a step size greater than 1, such as interpolation filters, additional code (APL and OPL instructions) is needed to monitor the value of the pointer. This requires several cycles of overhead for each data word addressed. Also, the maximum circular buffer length supported by the TMS320C50 is 256, thus limiting the size of digital filters that can be used.

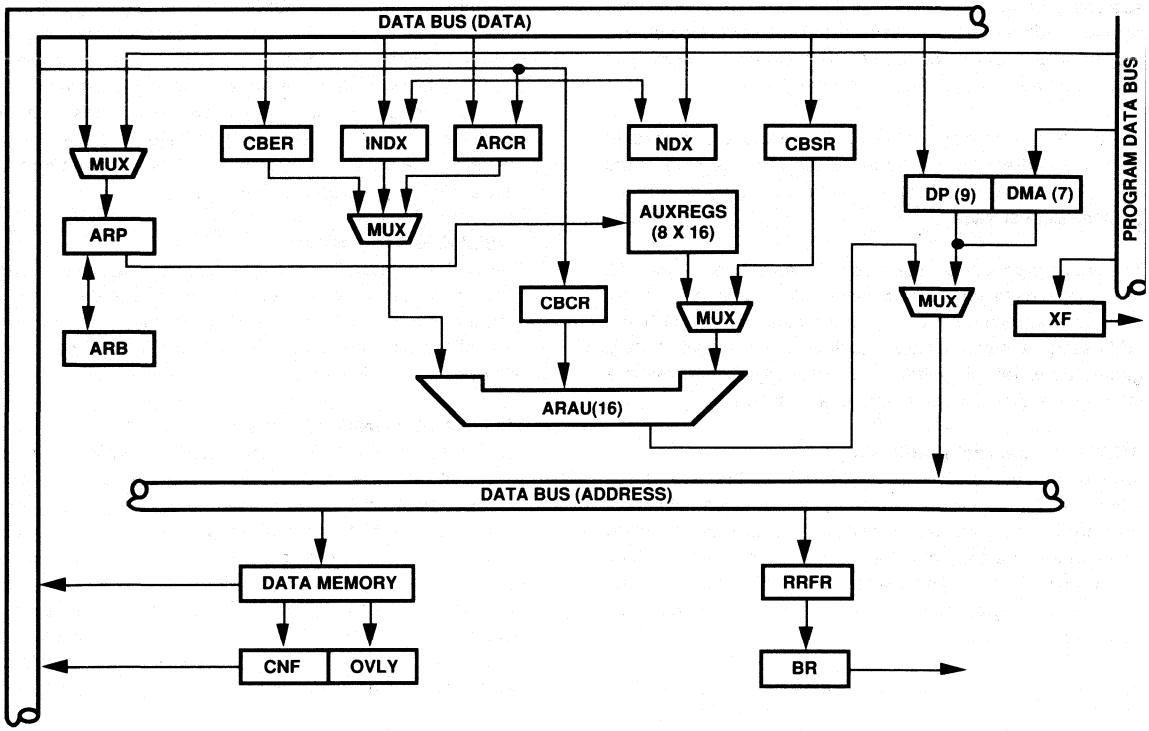

Figure 4. Block Diagram of TMS320C50 Address Generation Circuit

TMS320C50 Addressing Instructions

The instruction mnemonics of the TMS320C50 involve several addressing modes. Indirect and direct addressing is specified within arithmetic instructions and, depending upon the memory configuration, can impose several overhead cycles (overhead can be as high as eight cycles with external memory). Some general syntax examples are shown below.

```
ADD  {*|*+|*-|*0+|*0-|*BRO+|*BRO-} [,<next ARP>]
MPY  {*|*+|*-|*0+|*0-|*BRO+|*BRO-} [,<next ARP>]
```

Specific examples of these are shown below.

```
ADD     *
MPY     *0+
```

The first example uses the contents of an auxiliary register as the address and the second uses the contents of an auxiliary register as the address and adds the contents of auxiliary register 0 as a modifier. This instruction syntax can be hard to decipher because it does not directly name which auxiliary register is being used. That information is stored in the auxiliary register pointer (ARP).

The address generator can bit-reverse an address as it is sent out to the address bus for zero-overhead bit-reversing for the FFT. Auxiliary registers can also be used for general purpose data storage and the auxiliary ALU can be used for limited math.

ADDRESS GENERATION SUMMARY

Sustaining high rates of arithmetic operations demands maximum performance from the data addressing part of a processor's architecture. Table 2 summarizes the differences between the two processors in terms of their data addressing capabilities.

PROGRAM SEQUENCING CAPABILITIES

Efficient architectures for signal processing require fast arithmetic capabilities and matching speed in data addressing and fetching capabilities. To fully deliver the performance required for real-world signal processing, a DSP machine must execute its program with little or no overhead spent on maintaining the proper flow of control.

DSP Requirement	ADSP-2101	TMS320C50
Single-cycle fetch of two operands from on-chip	✔	No
Single-cycle MAC operations	✔	✔*
Modify two addresses by two different modify values on every cycle	✔	No
Bit-reverse data memory addresses for FFT	✔	✔
Automatic pointer wraparound for circular buffers	✔	✔**
Automatic circular modulo addressing	✔	No

* MAC, MACD, MADD & MADS instructions only
** For step size of 1 only, and cannot be used for program memory

Table 2. Summary of Data Addressing Capabilities

Efficiency in program sequencing has many different aspects; they cannot all be covered in this article. The comparison focuses primarily on two features

- the execution of loops and
- how branching and branching on conditions are handled.

Loops are fundamental to the way DSP algorithms are expressed in their natural mathematical form. Operations such as sums-of-products are repetitive. If the program can be efficiently expressed in a looped form, then coding is quite straight forward and changing the program (for example, to increase the number of taps in a filter) requires very little work.

Branching is fundamental to program structure. Branching on conditions (and executing arithmetic on conditions) is a natural way to construct any program which must respond to its environment.

Program Sequencer Architecture
Figure 5, which can be found on the following page, shows the architecture of the program sequencer of the ADSP-2101 and Figure 6, which can be found on page 12, shows that of the TMS320C50.

ADSP-2101 Program Sequencer
The program sequencer of the ADSP-2101 contains logic that selects a program memory address source and routes the address to the program memory address bus (PMA). This address selection occurs automatically in response to the current instruction. The address placed on the address bus can come from

- the program counter (for sequential addressing),
- a 14-bit address in the instruction word itself, for direct jumps and subroutine calls,
- the PC stack, for returns from subroutines and interrupts, and
- the interrupt logic, to automatically vector to the interrupt routine upon assertion of any external interrupt.

All instructions execute in a single cycle; this applies equally to jumps, calls and interrupts. No instruction pipelining is required in the ADSP-2101 so that program flow is simple to understand.

When an interrupt occurs, the complete status of the processor (stack status, mode status, arithmetic status and interrupt mask) is automatically pushed onto the status stack as part of the interrupt vector process.

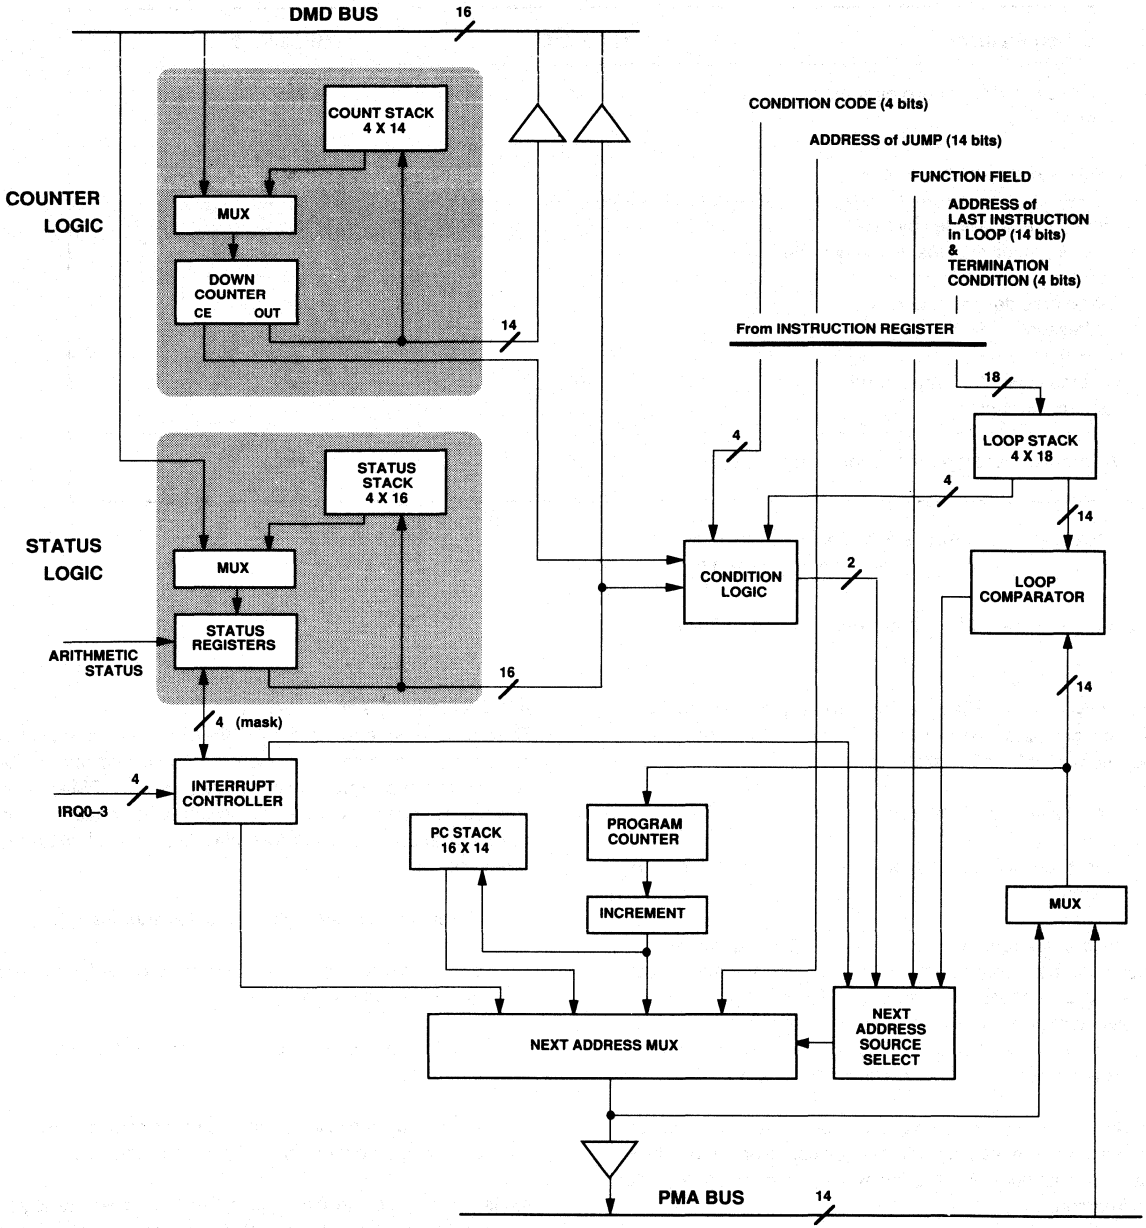

Figure 5. Block Diagram of the Program Sequencer of the ADSP-2101

ADSP-2101 Looping Capabilities

The ADSP-2101 program sequencer supports zero-overhead "DO UNTIL" loops. Using the count stack, loop stack and loop comparator, the processor can determine whether a loop should terminate and address the next instruction (either the top of the loop or the instruction after the loop) with no overhead cycle.

A DO UNTIL loop may be as large as program memory size permits, or as small as one instruction. A loop may terminate when a 14-bit counter expires or when any arithmetic condition occurs. The example below shows a three instruction loop that is to be repeated 100 times.

```
        CNTR = 100;
        DO Label UNTIL CE;
            First instruction of loop;
            Second instruction of loop;
Label:  Last instruction of loop;
        First instruction outside loop;
```

The first instruction loads the counter with 100. The DO UNTIL instruction contains the address of the last instruction in the loop (in this case the address represented by the identifier, *Label*) and also contains the termination condition (in this case the count expiring, CE). The execution of the DO UNTIL instruction causes the address of the first instruction of the loop to be pushed on the PC stack and the address of the last instruction of the loop to be pushed on the loop stack. (See Figure 5.)

As instruction addresses are output to the program memory address bus and the instruction is fetched, the loop comparator checks to see if the instruction is the last instruction of the loop. If it is, the program sequencer checks the status and condition logic to see if the termination condition is satisfied. The program sequencer then either takes the address from the PC stack (to go back to the top of the loop) or simply increments the PC (to go to the first instruction outside the loop).

The looping mechanism of the ADSP-2101 is automatic and transparent to the user. As long as the DO UNTIL instruction is specified, all stack and counter maintenance and program flow is handled by the sequencer logic with no overhead. This means that in one cycle the last instruction of the loop is being executed and in the very next cycle, the first instruction of the loop is executed or the first instruction outside the loop is executed, depending upon whether the loop terminated or not.

The ADSP-2101 can support four levels of nesting for loops. DSP routines such as matrix operations and two-dimensional processing, as well as more common algorithms such as the FFT, benefit from nested looping capabilities.

ADSP-2101 Program Sequencer Instructions

There are many conditional instructions for the ADSP-2101. Most arithmetic instructions as well as jumps, subroutine calls, returns from interrupts and returns from subroutines may all be conditional. The program sequencer decides on the fly whether the condition is true and what action to take, requiring zero overhead cycles. The coding of conditional jumps, subroutine calls and returns is straightforward. Some examples of the syntax are shown below.

```
IF condition JUMP label;
IF condition JUMP I4;
IF condition CALL label;
IF condition CALL I4;
IF condition RTS;
```

In the above examples, I4 references an address generator register for indirect branching. *Condition* refers to any of a set of 16 arithmetic conditions in the processor and *label* refers to any address or label in the program memory space.

TMS320C50 Program Sequencer

The program sequencer logic of the TMS320C50 controls instruction execution and consists of a program counter, stack and related hardware. Figure 6, on the following page, illustrates the logic used for program sequencing.

The TMS320C50 supports a block repeat function. The block repeat feature is controlled by three registers (PASR, PAER, and BRCR) which hold the top of loop address, the end of loop address, and the repeat count. Due to the limitations of instruction pipelining, the minimum size of a loop used with a block repeat is three instructions. A two instruction zero-overhead loop is not possible.

A loop is maintained automatically but since there are no local stacks or storage for loop count, top of loop address and bottom of loop address, there is no easy way to have nested loops. Logic is also included to repeat a single instruction as many as 256 times.

Instruction execution for the TMS320C50 utilizes a four-level pipeline consisting of a prefetch, decode, operand fetch, and execution stage. The four level pipeline imposes certain restrictions and extra cycles of overhead with operations such as loading data into registers, looping, branching, and executing certain instructions after other instructions. The ADSP-2101 has no such restrictions because it does not need the extra instruction pipelining to achieve its fast speed.

Anytime the flow of the program deviates from sequential instruction fetches, the instruction pipeline must be emptied and then refilled based on the destination address of the branch, call or interrupt vector. These types of operations require at least three cycles to execute when fetching the instruction from external memory or from internal program ROM. This type of instruction pipelining is not found in the ADSP-2101 (the fast instruction execution speed is achieved by other design techniques) and no extra overhead is encountered in the ADSP-2101 for jumps, subroutines or interrupts regardless of whether they are conditional or not.

A prefetch counter (PFC) contains the address of the next instruction to be prefetched. The prefetched instruction is loaded into the instruction register (IR), unless the instruction register still contains an instruction currently executing. In this case, the prefetched instruction is temporarily stored in the queue instruction register (QIR). The instruction pipeline, in conjunction with multi-cycle instruction execution, can make program flow complex and difficult to understand. Calculating a benchmark for a particular algorithm can also become difficult for the same reason. The following code examples illustrate the counter-intuitive sequence of events due to pipeline delays and the varying number of execution cycles for different instructions.

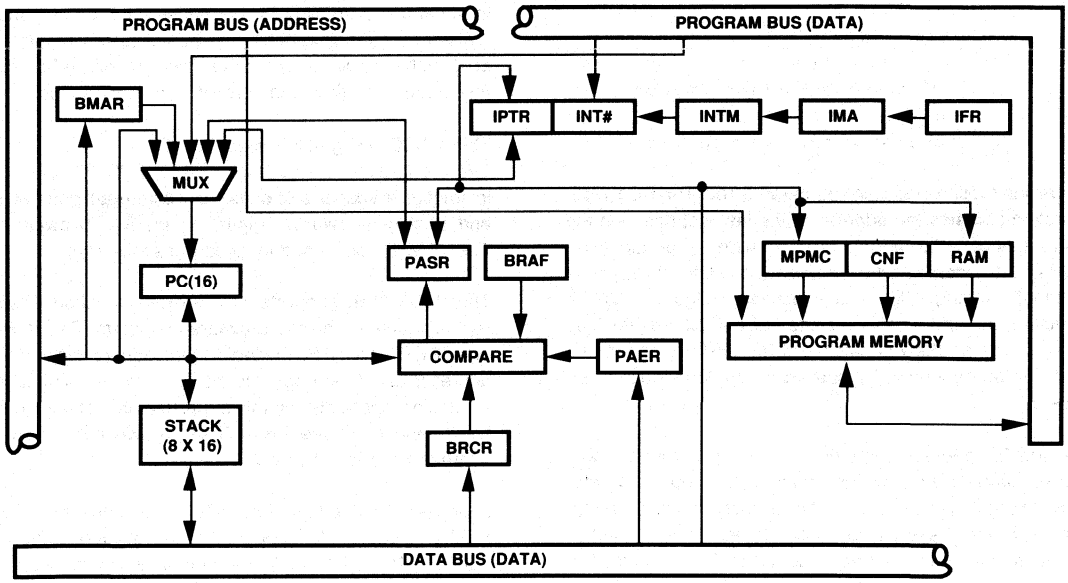

Figure 6. Block Diagram of the TMS320C50 Program Sequencer Circuit

```
PROB1  LAR    AR2,#067h  ;AR2 = 67.
       LACC   #064h      ;ACC = 00000064.
       SAMM   AR2        ;This update is overridden by *- updates
                         ;on the next two instructions
       LACC   *-         ;AR2 = 66.
       ADD    *-         ;AR2 = 65.
```

```
PROB2  LAR    AR2,#067h  ;AR2 = 67.
       LACC   #064h      ;ACC = 00000064.
       SAMM   AR2        ;LACC *- update happens before SAMM write
       LACC   *-         ;AR2 = 66.
       NOP               ;AR2 = 64  SAMM write to AR2 happens
                         ;  between instructions.
       ADD    *-         ;AR2 = 63.
```

The ADSP-2101 uses a single level of instruction pipeline where all instructions can execute in a single cycle. Therefore, none of these problems exist with the ADSP-2101.

The program counter can supply an address for sequential addressing. The single 8-deep PC stack is used for storage of return addresses as well as for providing the ability to push and pop data for the accumulator. An interrupt flag register (IFR) is used for the vectoring to an interrupt routine. Unlike the ADSP-2101, status is not automatically saved on the TMS320C50 for interrupts so that the programmer must perform any save and restore functions explicitly.

Branch instructions which contain a direct address require multiple program memory locations because both the instruction bits and the address cannot fit in the 16-bit instruction width. Delayed branches are required to minimize the overhead introduced by the instruction pipelining. Even with the use of delayed branches, as many as two cycles of overhead are required with the TMS320C50, where no overhead cycles are required with the ADSP-2101. Also, with the TMS320C50, the number of overhead cycles for a conditional branch will vary depending upon whether the condition is met or not.

TMS320C50 Program Sequencer Instructions

Arithmetic instructions cannot be conditional. Only branch instructions are conditional. Branch instructions with direct addresses require two program memory words due to the 16-bit instruction word.

```
BACC
BANZ <address>
```

There are many multiword instructions for the TMS320C50 because of the 16-bit size of the instruction word. This means that two or more fetches are required, which takes extra time. The ADSP-2101 has a 24-bit wide instruction and no multiword instructions are necessary.

PROGRAM SEQUENCER SUMMARY

Efficient looping capabilities are very important for DSP algorithms due to their repetitive nature. Also, zero-overhead jump and conditional branching is important where many decisions have to be made such as in speech processing. Table 3 summarizes the program sequencer capabilities of the ADSP-2101 and TMS320C50.

I/O HANDLING CAPABILITIES

A final area of efficiency is that of I/O handling. Memories, A/D and D/A converters, as well as EPROM for program booting will need to efficiently interface to the DSP processor to minimize extra logic and software overhead to drive external peripherals. The ADSP-2101 has several features relating to I/O handling which simplify DSP system design and which are not found on the TMS320C50.

Automatic Boot Loading From External Byte-Wide Memory

The ADSP-2101 directly interfaces to a single byte-wide EPROM for efficient program boot loading. No extra components are needed since the EPROM can directly connect to the address and data lines of the ADSP-2101. A boot memory select pin (BMS) on the ADSP-2101 is tied directly to the chip select pin of the EPROM and the read line (RD) is directly connected to the output enable pin of the EPROM. The boot memory space consists of an external 64K x 8 space divided into eight separate 8K x 8 pages. At reset, boot page 0 is automatically transferred in to the internal RAM of the ADSP-2101. Under program control, any of the eight pages can be boot loaded into the internal RAM of the ADSP-2101 with access time being programmable.

Flexible Serial Ports

Both devices have two serial ports. The serial ports of the ADSP-2101 can operate at the full speed of the processor where the serial port of the TMS320C50 can only operate at 1/4 the instruction cycle rate. The serial ports of the ADSP-2101 also have some additional features which makes their operation more flexible. The word width of the data to be transmitted and received is programmable and can be set for any size from 3 bits to 16 bits. On the TMS320C50, the word width is limited to 8 or 16 bits.

The address generators of the ADSP-2101 can be used in conjunction with the serial ports to provide an automatic data buffering capability. Normally, an interrupt is generated after each word is transferred through the serial port. If many words are to be transferred (i.e., data buffers filled for a speech application), there can be an excess of interrupt overhead associated with the serial ports. The ADSP-2101 allows autobuffering where a length is specified along with a buffer start address and a modify value (any integer value which is used to update the address). As each word is transferred through the serial port, the data is automatically read from or written to data memory, transparent to the user, with no interrupt being generated. An interrupt is generated only when the buffer is full or empty. One of the serial ports of the ADSP-2101 also supports a multichannel word stream for easy interface to a T1 or CEPT data stream.

9

DSP Requirement	ADSP-2101	TMS320C50
PC stack depth	16	8
Nested looping	4 levels	No
Conditional arithmetic instructions	✔	No
Zero-overhead branching	✔	No
Speed achieved with pipelining	Not required	✔ 4-level pipeline
Automatic status saving during interrupt vector	✔	No

Table 3. Summary of Program Sequencing Capabilities

SUMMARY

The DSP processors available on the market today vary drastically in their ability to meet the five key requirements of DSP processing. In fact, some DSP-oriented processors, like the TMS320C50, are better high-speed microcontrollers than they are DSP processors. Analyzing the requirements of your DSP system and matching them to the capabilities of a DSP architecture will assure efficient operation.

Due to space limits, this article does not cover many topics in detail. Consult the *ADSP-2101 User's Manual* and the *ADSP-2101 Cross-Software Manual* for a greater depth of information on this processor.

APPENDIX: PROGRAM EXAMPLE

To illustrate some of the issues discussed above, a code example is shown below for the ADSP-2101 and the TMS320C50. To avoid long listings and confusion, a short program which performs the LMS adaption of FIR filter coefficients is shown. Both processors perform identical tasks so that no interpretation of the type of algorithm is required. Both code examples do not show any initialization of pointers or the set up of any modes. For simplicity, the examples only focus on the core operation.

Because these examples are short, the performance advantages of the ADSP-2101 is not as apparent as in a more sophisticated example. Nevertheless, the ease of coding and the benefits of the instruction syntax and the architecture can be seen.

ADSP-2101 Code Example Description

The example shown implements an adaptive update of FIR filter coefficients. The formula used is expressed as

`Ck+1=Ck+Beta*Error*A(n).`

The program segment shown was taken from the book *Digital Signal Processing Applications Using The ADSP-2100 Family* published by Prentice Hall.

The code shown uses the looping capabilities of the ADSP-2101 and can be easily expanded for a larger number of coefficients by simply changing the number of loops (the value loaded into the counter). Indirect addressing is used to address the coefficient buffer Ck and the input data buffer A(n). The address registers I0 and I4 are used for addressing of these two buffers.

The first advantage of the ADSP-2101 is its algebraic syntax for assembly language code. The routine starts with a fetch of the error term from data memory. This value is loaded into the register AR. AR is the ALU result register, but it is used as a general purpose data register in this example. The next line of code loads an immediate value, the beta value, into register MY1. MY1 is one of the input registers of the multiplier for the Y operand.

With the error value in register AR and the beta value in register MY1, a multiplication of these two values is specified. The multiplication is performed with the result rounded to the most significant 16-bits with an unbiased rounding scheme. This multifunction instruction also specifies the fetch of the coefficient , Ck, from program memory and the data value A(n) from data memory. Note that the I register specifies which address register is used as a pointer and the M register specifies how the address is modified. This addressing capability is a key advantage to that of the TMS320C50. The multiplication, the program memory fetch and the data memory fetch all occur in a single cycle. The result of the multiplication is loaded into MF, the multiplier feedback register. This value is used immediately in the next cycle where a multiplication is performed using the MX0 register (holding the A(n) term) and the MF register (holding the product beta*error). Rounding is again specified.

The counter is next loaded with the number of coefficients to be updated and a DO UNTIL instruction is specified to set up the loop logic of the ADSP-2101. The core instructions of the loop calculate the result Ck+1 and also set up the calculations for the next update. Results are written into program memory in the last instruction of the loop.

Finally, a return from subroutine instruction is specified to return control back to the calling program.

```
AR=DM(Error);                                          {Get Err Value From Mem }
        MY1=Beta;                                      {Load Beta Value        }
        MF=AR*MY1(RND), AY0=PM(I4,M4), MX0=DM(I0,M0);  {MF=Beta*Err, Get Ck, A }
        MR=MX0*MF(RND);                                {MR=Beta*Error*A(n)     }

        CNTR=A;                                        {Set Loop Counter       }
        DO uloop UNTIL CE;                             {Tap Update Loop        }
          AR=MR1+AY0, AY0=PM(I4,M6), MX0=DM(I0,M1);    {AR=Ck+Beta*Error*A(n)  }
uloop:    PM(I4,M7)=AR, MR=MX0*MF(RND);                {Store CK+1, Do Next    }
        RTS;                                           {Return}
```

Performance Benchmark

The code section shown uses the looping capabilities of the ADSP-2101 and can be easily modified for any number of coefficients by simply changing the counter value. A total of nine instructions are used in the LMS adaption of FIR filter coefficients where each instruction executes in a single processor cycle. The two instructions in the core of the loop are repeated for each coefficient update. Therefore, the benchmark for the number of cycles required for this routine can be generally expressed as

*7+n*2*, where *n* is the number of coefficients to be updated.

For a 127 TAP filter (which requires 127 coefficients), an update can be performed in 7+127*2 = 261 cycles.

TMS320C50 Code Example Description

The example shown implements an adaptive update of FIR filter coefficients. The formula used is expressed as

```
a0(i+1)=a0(i)+Beta*err*X(i).
```

This is the same LMS adaptive update as shown for the ADSP-2101, the equation has just been stated with different terms. The program segment shown is described in the book *TMS320C5x User's Guide* published by Texas Instruments.

This is an example of looped code based on the RTPB (repeat block) instruction. Indirect addressing is used to address the coefficient buffer a(i) and the input data buffer x(i). The auxiliary registers AR2 and AR3 are used to address these two buffers.

The LMS adaption routine starts by loading the error stored in the memory location "ERR" into the TREG0 register for multiplication. The LT instruction is used to load the T register. Once the error is loaded, then the error is multiplied by the Beta value stored in the memory location "BETA". The results (error*beta) resides in the P register of the multiplier. Because of the inflexibility of the TMS320C50 architecture, the multiplier result must be moved explicitly into the accumulator. PAC is used to place the product into the accumulator for further computation. The error*beta term can then be rounded to 16-bit precision with the instruction "ADD ONE,14" and stored into a memory location with the SACH instruction. The rounding takes an extra instruction. On the ADSP-2101, this function can be performed as part of the multiply.

At this point in the program, the loop to calculate all of the new coefficients can be set up. There are 127 coefficients in this example, so the loop counter BRCR can be initialized with the constant 126. On the ADSP-2101, the programmer loads the loop counter directly with the number of loop iterations. Two instructions are required to load the loop counter, LACC and SAMM.

Indirect accesses using the auxiliary registers AR2 and AR3 are used within the loop. These registers can be initialized with the LAR instruction prior to entering the loop.

The error*beta term can then be reloaded back into the T register for multiplication with the tapped delay line values. The "LT ERRF" instruction loads this value and the "MPY *-,AR2" performs the first multiply outside of the loop. This reloading of

```
        LT    ERR               ; T=Err
        MPY   BETA              ; P=Beta*Err(i)
        PAC                     ; errf(i)=Beta*Err(i)
        ADD   ONE, 14           ; Round The Results
        SACH  ERRF,1            ; Save errf(i)

        LACC  #126
        SAMM  BRCR              ; 127 Coeffs To Update In The Loop
        LAR   AR2, #COEFFD      ; Point To The Coefficients
        LAR   AR3, #LASTAP      ; Point To The Data Samples
        LT    ERRF
        MPY   *-,AR2            ; P=Beta*Err(i)*x(i-255)

        RPTB  LOOP-1            ; For I=0, I<=126, I++
ADAPT   ZALR  *, AR3            ; Load ACCH With Ak(i)
        MPYA  *-,AR2            ; P=Beta*Err(i)*X(i-k-1),
                                  ACC=ak(i)+Beta*err(i)*x(i-k)
        SACH  *+                ; Store ak(i+1)

LOOP    ZALR  *,AR3             ; Final Update Last Coefficient a0(i)
        RETD                    ; Delayed Return
        APAC                    ; ACC=a0(i)+Beta*Err(i)*x(i)
        SACH  *+                ; Save a0(i+1)
```

partial results is required because of the inflexibility of the TMS320C50 architecture. The algorithm needs to be rearranged due to limitations of the hardware. The ADSP-2101 makes use of a more flexible bus structure where data can be fed back immediately without the need for temporary storage of intermediate results in memory.

The RTPB instruction performs the block repeat. As an argument, this instruction needs the end of loop address minus one. The three instructions ZALR, MPYA and SACH are executed in the loop 127 times. Notice that the programmer must label the instruction after the last instruction in the loop. The ADSP-2101 uses a much more understandable looping format where the last instruction in the loop is labeled. Any size loop is possible on the ADSP-2101. The TMS320C50 block repeat is useful only for loops of three instructions or larger.

After completion of the loop, the last tap is updated and a delayed return is executed. A delayed return is necessary because of the instruction pipelining found in the TMS320C50 This type of instruction pipelining is not found in the ADSP 2101 and delayed instructions are, therefore, not necessary

Performance Benchmark

A total of 19 instructions are used in the LMS adaption of FIF filter coefficients. All instructions, however, will not execute ir a single processor cycle. The three instructions in the core o the loop are repeated for each coefficient update. Therefore the benchmark for the number of cycles required for this routine can be generally expressed as

$17+n*3$, where n is the number of coefficients to be updated

For a 127 TAP filter (which requires 127 coefficients), ar update can be performed in 17+127*3 = 398 cycles.

ANALOG DEVICES

AN-235
APPLICATION NOTE

ONE TECHNOLOGY WAY • P. O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Considerations for Selecting a DSP Processor (ADSP-21020/ADSP-21010 vs. TMS320C30)

by Dan Ash

INTRODUCTION

Digital signal processing algorithms consist of multiplications and additions. In signal processing algorithms, additions are necessary to accumulate multiplier products. Most general purpose digital signal processors sustain single instruction cycle *multiply/accumulate* throughput. A digital signal processor must also meet the following requirements to efficiently process sampled signals in real time.

- Efficient access of dual data operands to sustain single-cycle multiply/accumulate throughput
- Circular/modulo data addressing to restrict index registers to a range of data addresses
- Hardware zero overhead program looping to eliminate the overhead of branching and of decrementing loop counters from time critical loops
- Extended arithmetic dynamic range to avoid overflows when accumulating fixed-point multiplier products

Many of today's digital signal processors satisfy these requirements, some more efficiently than others. This application note compares the core architecture of the Analog Devices ADSP-21020 (Figure 1) and the ADSP-21010 to that of the TMS320C30 (Figure 2) from Texas Instruments. The **ADSP-21010** is a pin-compatible, **32-bit only, low-cost** version of the ADSP-21020. To avoid confusion, most references in this application note are made to the ADSP-21020.

The three sections that follow discuss the four requirements mentioned emphasizing the differences between the Analog Devices ADSP-21020 and Texas Instruments TMS320C30 digital signal processors. The arithmetic section discusses fixed- and floating-point data formats, parallel operation of the multiplier and ALU, and the general purpose arithmetic capabilities of each processor. The data addressing section discusses the various addressing modes of each processor

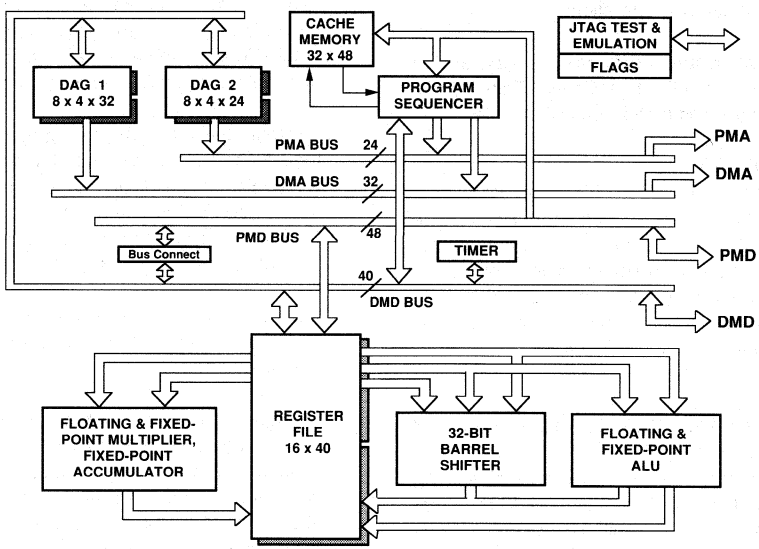

Figure 1. Analog Devices ADSP-21020

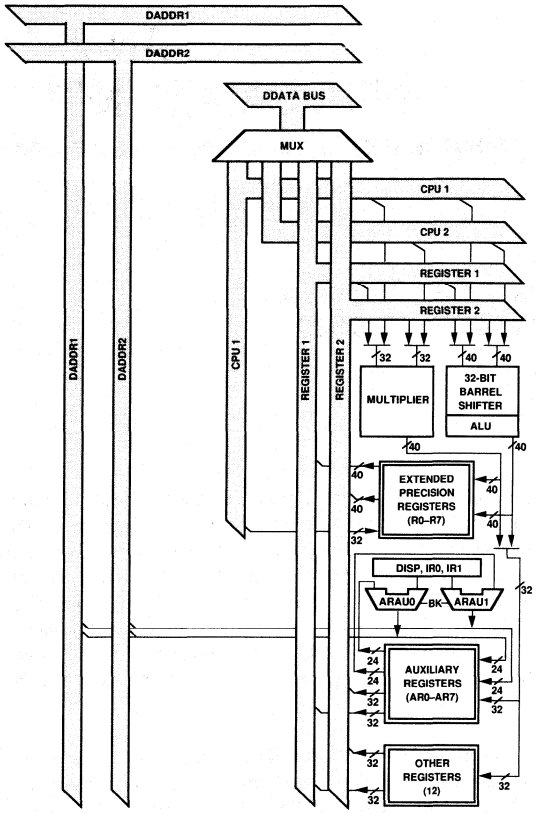

Figure 2. Texas Instruments TMS320C30 CPU

and the parallel operation of the arithmetic units with data moves. The third section discusses program sequencing including zero overhead program loops, conditional branching, and other conditional operations.

A comparison of assembly language syntax and a programming example (using the Fast Fourier Transform or FFT) are presented at the end of this application note.

ARITHMETIC CAPABILITIES

Digital signal processors must be able to multiply data operands and accumulate the products with single-cycle throughput. Some calculations require fixed-point arithmetic, others require the dynamic range of floating-point. A processor must also efficiently perform general purpose arithmetic and logical operations. Independent operation of the ALU and multiplier allowsefficient software to be written.

Data Formats

Both the ADSP-21020 and the TMS320C30 support fixed- and floating-point data formats. The ADSP-21020 supports a 32-bit and a 40-bit floating-point data format. The 32-bit floating point conforms to the IEEE 754/854 standard. This format specifies a 24-bit mantissa and an 8-bit exponent. The 40-bit format includes a 32-bit mantissa and an 8-bit exponent.

The TMS320C30 supports a 32-bit floating-point data format that specifies a 24-bit mantissa and 8 bit exponent. This format does not conform with the IEEE standard.

The TMS320C30 converts its generic floating-point format to IEEE format in 12 instruction cycles and from IEEE to TMS320C30 format in 23 instruction cycles worst case. The TMS320C30 ALU operates on 40-bit floating-point data, while the multiplier is limited to 32-bit floating-point inputs.

The ADSP-21020 multiplier and ALU operate on 32-bit fixed-point inputs as well as floating-point. The TMS320C30 multiplier only supports 24-bit fixed-point inputs while the TMS320C30 ALU operates on 32-bit fixed-point data.

The ADSP-21020 interfaces to either 32-bit or 40-bit data memory. The TMS320C30 interfaces to 32-bit data memory.

The ADSP-21020 supports two IEEE rounding modes; round-to-nearest and round-toward-zero. One of the two modes is always enabled so that both ALU and multiplier results are automatically rounded to 32-bit or 40-bit floating-point numbers. TMS320C30 results can be rounded using an instruction, RND. This instruction requires one instruction cycle to execute.

The ADSP-21020 can generate interrupts upon ALU/multiplier fixed-point overflow, floating-point overflow, floating-point underflow, or floating-point invalid operation. The TMS320C30 requires conditional branch instructions to trap on these arithmetic conditions.

Arithmetic Architecture

Figure 3 shows a block diagram of the ADSP-21020 register file and computational units. The register file is local storage for arithmetic input data and results. It consists of sixteen 40-bit registers and of eight data ports. In one instruction cycle, it can perform two memory transfers, source two inputs to the ALU, source two inputs to the multiplier, and store ALU and multiplier results. The two memory transfers can either both be reads, both writes, or one read and one write.

The register file consists of two groups of eight registers. Each register group can be swapped with a set of secondary registers in one instruction cycle. The secondary registers are typically used during interrupts or subroutines to eliminate time consuming memory transfers of register file data. Use of the 16 secondary registers provides a total of 32 register file registers.

Figure 4 shows the TMS320C30 register file and computational units. The register file consists of eight 40-bit registers.

The register file can source two inputs to the multiplier while the two ALU inputs are fetched directly from memory. Conversely, the inputs to the ALU can be sourced from the register file while the multiplier inputs are fetched directly from memory. The ALU and multiplier results are stored into the register file. The TMS320C30 register file does not have secondary registers.

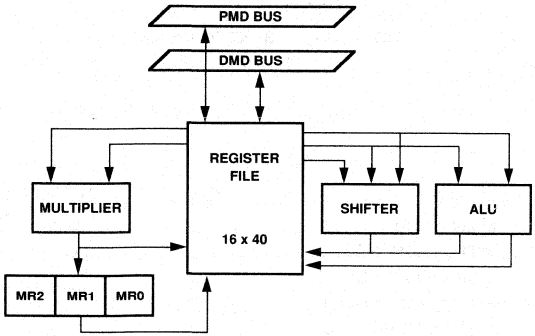

Figure 3. ADSP-21020 Register File/Computational Units

General Purpose Math

Signal processing and data compression algorithms require subtractions, division, square root, minimum/maximum comparisons, shifting, and logical operations, as well as multiply/accumulate operations. A powerful instruction set implements these functions directly. Direct implementation of such functions reduces programming errors. Single-cycle execution of such functions will result in faster execution of algorithms.

ALU

The ADSP-21020 ALU provides a complete set of logical and general arithmetic functions. The ALU operates on 32-bit fixed-point, 32-bit IEEE floating-point, or 40-bit floating-point data. In addition to performing fixed- and floating-point adds and subtracts, the ADSP-21020 ALU supports many non-traditional processor functions. All functions execute in one instruction cycle. The table below lists those functions.

INSTRUCTION	*21020*	*C30*
ABS X	x	x
–X	x	x
(X+Y)/2	x	
MIN(X,Y)	x	
MAX(X,Y)	x	
COMP(X,Y)	x	x
CLIP X BY Y	x	
X+Y,X–Y	x	
FLOAT Y	x	x
SCALE X BY Y **	x	
MANT X *	x	
LOGB X *	x	
FIX X *	x	x
FIX X BY Y **	x	
COPY X SIGN TO Y *	x	
ABS (X+Y) *	x	
ABS (X–Y) *	x	
SEED 1/X	x	
SEED 1/SQRT(X)	x	

* x and y are floating point

** x is fixed point, y is floating point

The ADSP-21020 ALU supports multiprecision fixed-point additions and subtractions preserving carry and borrow respectively. The ADSP-21020 has the ability to generate a seed for fast floating-point calculation of reciprocal and reciprocal square root. The TMS320C30 requires more instruction cycles to evaluate those expressions.

OPERATION	*21020*	*C30*
1/X	6 cycles	35 cycles
Y/X	6 cycles	36 cycles
1/SQRT(X)	9 cycles	74 cycles
SQRT(X)	10 cycles	39 cycles

The TMS320C30 ALU satisfies the basic requirements of addition, subtraction, absolute value, negate, and logic functions. The TMS320C30 however does not directly support many of the instructions available on the ADSP-21020. The ADSP-21020 maximum instruction R7=MAX(R5,R6) is coded on a TMS320C30 as:

```
LDF   R6,R7
CMPF  R5,R7
LDFLT R5,R7
```

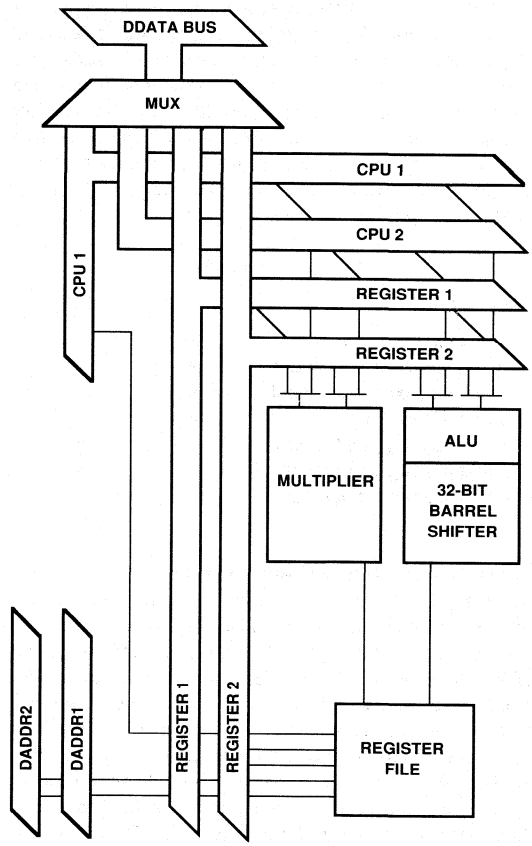

Figure 4. TMS320C30 Register File/Computational Units

ALU Summary		
FEATURE	*21020*	*C30*
40-bit floating point	x	x
32-bit fixed point	x	x
Multiprecision add/subtract	x	x
Min/max/avg/clip/scale	x	
Abs(x+y),abs(x–y)	x	
Simultaneous x+y,x–y	x	
8-bits accumulated cmp status	x	
Seed 1/x, seed 1/sqrt(x)	x	

Multiplier

The ADSP-21020 multiplier operates on 32-bit fixed-point, 32-bit floating-point, or 40-bit floating-point inputs. For 32-bit fixed-point multiplies, the upper 32 bits of the 64-bit product can be sent to a data register, or the 64-bit product can be added to or subtracted from one of two 80-bit fixed-point accumulators. Fixed-point inputs can be individually chosen as signed or unsigned, fractional or integer. Products of 32-bit or 40-bit floating-point multiplies must be sent to one of the sixteen 40-bit data registers.

The TMS320C30 multiplier operates on 24-bit fixed-point or 32-bit floating-point inputs. In the case of 24-bit fixed-point multiplication, a 48-bit product is generated. Inputs are assumed to be signed integers. Fractional fixed-point multiplication is not supported. The 32-bit ALU is used for accumulation of fixed-point products. A fixed-point product's magnitude is limited to 32 bits because only the lower 32 bits of the product are written to the destination register. Products of 32-bit floating-point multiplies are 40 bit.

Multiplier Summary		
FEATURE	*21020*	*C30*
32-bit fixed-point input	x	
Unsigned fixed-point input	x	
Fractional fixed-point input	x	
80-bit accumulators	x	
32-bit floating-point inputs	x	x
40-bit floating-point results	x	x
40-bit floating-point inputs	x	

Parallel Operation Of ALU And Multiplier

Since many DSP algorithms are modeled around a sum-of-products calculation (Figure 5), a DSP microprocessor must be able to perform multiply/addition operations efficiently. The ADSP-21020 and the TMS320C30 have the ability to multiply two operands and add (accumulate) the product from the previous multiply to a running sum simultaneously. There are other cases where code can be optimized by simultaneous operation of the ALU and multiplier. The ALU and multiplier may be performing completely independent tasks. The ADSP-21020 instruction set allows for such flexibility. The table below summarizes possible combinations of ALU and multiply operations on both the ADSP-21020 and TMS320C30.

$$y(n) = \sum_{K=0}^{M-1} a(k)\, x(n-k)$$

Figure 5. Example of Sum-of-Products Equation

OPERATION	*21020*	*C30*
Multiply/add	x	x
Multiply/subtract	x	x
Multiply/add/subtract	x	
Multiply/fixed-->float*	x	
Multiply/float-->fixed*	x	
Multiply/avg*	x	
Multiply/abs*	x	
Multiply/max*	x	
Multiply/min*	x	

* floating point only

Barrel Shifter

DSP microprocessors allow multiple bit shifting within a single instruction cycle. Bit and bit field manipulation is important for general purpose I/O functions and in graphics processing.

The ADSP-21020 instruction set includes many bit, bit field, and shift operations. The ADSP-21020 supports either left or right arithmetic, logical, and rotational shifts. The shift amount can originate from a register or can be specified in the instruction word. Shift results can be logically OR'd with other data registers. The ADSP-21020 supports bit test, bit set, bit toggle, and bit clear. The ADSP-21020 can extract arbitrary bit fields from one word and deposit the result anywhere inside another word. The ADSP-21020 can extract the exponent from a fixed-point input. It can count either leading ones or leading zeros in a fixed-point input.

The TMS320C30 supports only a subset of the ADSP-21020 shifter operations. The table below compares the instruction sets.

OPERATION	*21020*	*C30*
logical/arithmetic shift	x	x
rotate	x	x
shift/or with register	x	
bit test/clr/toggle/set	x	x*
field extract/deposit	x	
exponent extract	x	
count leading ones/zeros	x	

* requires masks stored in memory to accomplish

DATA ADDRESSING

Data addressing pertains to how data is read from memory or how data is written to memory. In order to perform multiplies each instruction cycle, new input data must also be made available at the same rate. Most DSP microprocessors satisfy this requirement of two data accesses in a single cycle. The basic addressing modes for most microprocessors are direct and register indirect addressing.

Direct Addressing

Direct addressing implies that the data address is specified directly in the instruction. The ADSP-21020 instruction, R2=DM(0x100000), loads data register R2 with the contents of address 1024K in data memory.

A 48-bit instruction word allows the ADSP-21020 to directly address the full 24-bit program memory or 32-bit data memory in one instruction cycle with one instruction opcode. Any general purpose or data register can be loaded from memory or stored into memory. The ADSP-21020 can also perform a single-cycle ALU/multiplier operation and direct memory read or write within a 64 location offset conditionally. Any of the 16 index registers can be used as base addresses for this offset.

The TMS320C30 only has a 32-bit instruction word. The TMS320C30 can only directly address data within 64K pages. The page base address is preloaded into the data page pointer (DP) register. There is only one data page pointer. The TMS320C30 can perform either an ALU or multiply instruction (but not both) with a direct memory read within a 64K page unconditionally. A direct memory write combined with a compute is not possible.

Indirect Addressing

Indirect addressing uses address generator registers as address sources. The ADSP-21020 instruction, PM(I9,M13)=R1, writes data from data register R1 into program memory. The program memory address is provided by index register I9. I9 is modified at the end of the cycle by the contents of modify register M13.

The ADSP-21020 has two data address generators (DAGs). One DAG is used to access data from data memory, the other allows data access of program memory. Each DAG (Figure 6) contains eight index registers and eight modify registers. Each index register can be used to address a different area in memory. When using indirect addressing, the programmer can pre-update or post-update the designated index register with any of the eight modify registers in the same DAG. The index register can also be modified by a 6-bit immediate offset.

The ADSP-21020 can perform a compute (multiplier + ALU), and two memory transfers (two reads, two writes, or one read and one write) in a single instruction cycle. The ADSP-21020 can perform a conditional compute (multiplier + ALU) and one memory transfer (read or write) in a single instruction cycle.

Each index register has a corresponding base address and length register. The base address and length registers allow one to confine the index register value within a range of data addresses. Each time the index register is modified, the resultant address is transparently tested by the DAG. If the resultant address is out of range, the index register is corrected with the modulo address. This feature is referred to as modulo or circular data addressing. The base and length registers allow placement of circular buffers within any range of addresses. Since the ADSP-21020 contains 16 index, base, and length registers, 16 circular buffers can be maintained simultaneously.

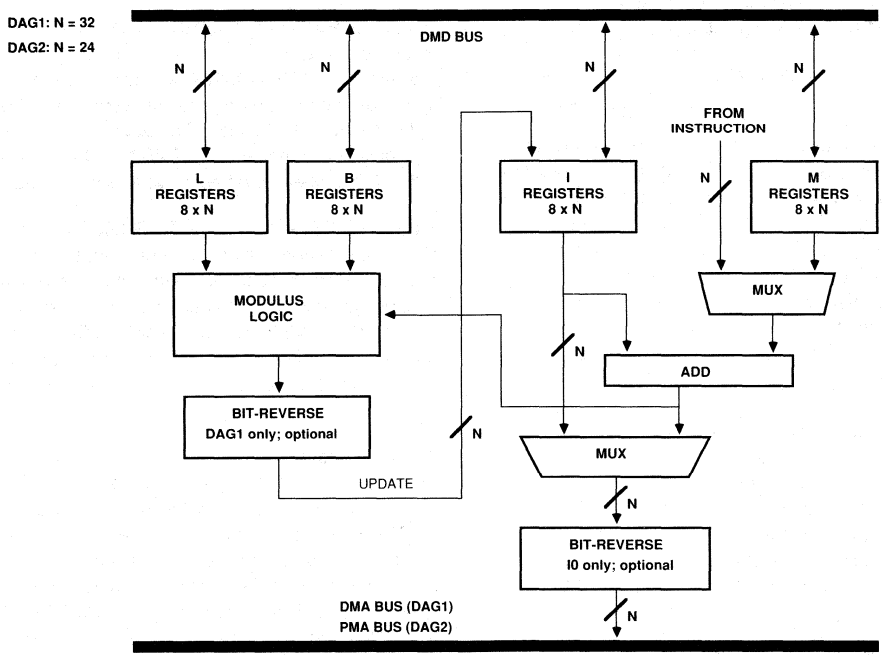

Figure 6. ADSP-21020 Data Address Generator

Each DAG's eight index, base, length, and modify registers has associated alternate registers configured in two groups of four. In the case of an interrupt or subroutine, one or both of the DAG's register groups can be swapped with secondary registers saving the overhead of register to memory transfers for a context save. The secondary DAG registers make available an additional 16 circular buffers.

The TMS320C30 also supports indexed addressing. It has eight 24-bit address registers AR7-AR0. These eight registers can be used to address different areas of memory. There are no alternate address registers. The TMS320C30 can fetch two operands from memory in a single cycle using two of the eight address registers. The address registers can be pre or post-modified by 0, 1, or one of two pre-programmed modify registers (IR1,IR0). Address registers can be modified by an 8-bit offset.

The TMS320C30 can perform a compute (multiplier + ALU) and two memory reads. It can perform a single compute (multiplier or ALU) with both one memory read and write.

The TMS320C30 supports circular buffers, but there is only one length register (BK). Therefore, if multiple circular buffers are used in one routine, they must all be the same length or the BK register must be modified for each circular data access. Circular buffers must be placed on the next power of two boundary greater than the buffer length.

Bit-Reversed Addressing

Bit- reversed addressing is necessary for proper addressing of data for the FFT. Both the ADSP-21020 and the TMS320C30 support bit-reversed addressing. Only the ADSP-21020 supports read back of bit-reversed addresses. Therefore, efficient, in-place bit-reversal is available.

Data Addressing Summary

FEATURE	*21020*	*C30*
Direct addressing (read or write)	32-bit	16-bit
Direct read, one compute	6-bit *	16-bit **
Direct read, two computes	6-bit *	n/s
Direct write, one compute	6-bit *	n/s
Direct write, two computes	6-bit *	n/s
Number of index registers	16	8
Number of modify registers	16	2
Number of length registers	16	1
Number of base address registers	16	0
Secondary index, modify, length, base registers	x	
Two indirect reads, two computes	x	x
Indirect read & write, one compute	x	x
Indirect read & write, two computes	x	
Two indirect writes, no compute	x	x
Two indirect writes, one compute	x	
Two indirect writes, two computes	x	
Compute, modify by immediate offset	6-bit	8-bit **
Bit-reversed addressing	x	x

* conditional or unconditional execution

** two operand compute only (reg1=reg1+src)

n/s not supported

PROGRAM SEQUENCING

An essential part of any microprocessor is the program sequencer. It is responsible for determining the program flow. In most cases, programs flow linearly from one address to the next sequential memory location. Branching is a form of program sequencing where the program jumps conditionally or unconditionally to another address. Subroutine calls are similar to branches, but store the pre-call address in order to return later. Interrupts also take advantage of the program sequencer. A DSP unconditionally branches to a dedicated program memory address upon interrupt. Each interrupt has a reserved memory location. Zero overhead program looping is another feature common to both the ADSP-21020 and TMS320C30.

Both the ADSP-21020 and TMS320C30 have these basic program sequencing capabilities. There are subtle differences that make programming easier and more efficient with the ADSP-21020.

Branching

Branching is necessary to change direction of program flow. Branch conditions are typically based on the arithmetic results of a bit operation, a subtraction of two inputs, or perhaps an overflow. Both the ADSP-21020 and the TMS320C30 can branch upon floating-point overflows or underflows.

The ADSP-21020 has the ability to branch on either downcounter status or on the level of any of four programmable input pins. It also has separate status bits for the ALU and multiplier. The ADSP-21020 also has complex conditional instructions where an arithmetic operation and a data move can all be executed on one condition.

Branch addresses can either be specified directly in the instruction word or can be indirectly specified through an index register. The ADSP-21020 allows a full 24-bit program address to be specified on any branch instruction. The TMS320C30 only allows full 24-bit addresses to be specified on unconditional branches. The TMS320C30 is limited to a +/– 32K (16-bit) offset from the conditional branch address.

Both the ADSP-21020 and TMS320C30 support indirect branch instructions where an index register is the address source. Using an index register allows a programmer to change a branch address during program execution. On the ADSP-21020, indirect jumps are pre-modified and execute in one instruction cycle. This is useful for implementing jump tables. The TMS320C30 does not have a pre-modified jump feature.

Pipeline Delays

When the ADSP-21020 fetches an instruction from memory, that instruction is not executed until two cycles later. The ADSP-21020 has a three-level (fetch, decode, execute) instruction pipeline. When a branch is taken, the two instructions fetched after the branch are in the pipeline waiting to be executed. When branching, the ADSP-21020 executes two NOPs instead of the two pipelined instructions. If the programmer specifies a delayed branch, the ADSP-21020 executes the two pipelined instructions. Eliminating the NOP

overhead allows the programmer to write more efficient code. The TMS320C30 has a delayed branch feature, but it has a four-level instruction pipeline. Non-delayed branches require four cycles to execute on the TMS320C30.

Subroutine Calls And Interrupts

When a subroutine call is made, the return address is stored on the program counter (PC) stack. The ADSP-21020 subroutine call is a single-cycle, single-operand instruction and is made using a direct or indirect address. Subroutine calls on the ADSP-21020 can be delayed allowing the two subsequent instructions to be executed. The TMS320C30 only supports non-delayed subroutine calls.

Interrupts

Interrupts are generated by internal or external events. Both the ADSP-21020 and the TMS320C30 have external interrupt input pins which generate interrupts upon high-low signal edges or in the presence of a low-level input. Both processors have internal timers to generate periodic software interrupts.

The ADSP-21020 responds (vectors) to an interrupt with no more than four cycles of latency. An interrupt mask register allows interrupts to be individually enabled by setting bits. There is one bit, IRPTEN, which allows masking of all interrupts. Each interrupt has eight reserved memory locations. Therefore, a delayed branch can be placed at the first interrupt vector address. There is also a five-level status stack. Anytime an interrupt occurs, arithmetic, mode, and interrupt mask status are automatically pushed onto the status stack. Status can be manually pushed onto this stack for subroutines. It is sometimes necessary that higher priority interrupts remain unmasked during lower priority interrupt routines. ADSP-21020 has a mode which allows such interrupt nesting. The ADSP-21020 has a special instruction, IDLE, which allows the processor to wait for interrupts in a low power mode. If an interrupt occurs while in IDLE, the processor stores the address of the instruction after IDLE as the interrupt return address and jumps to the interrupt vector.

The ADSP-21020 has four external interrupt input pins (IRQ3-0). It has several internally generated interrupts. An interrupt is generated any time the PC, status, or loop stacks overflow. The 32-bit internal timer can be made either a high priority or low priority interrupt. Interrupts can be generated anytime index registers I7 or I15 have modulo/circular buffer overflows. Interrupts can be generated upon fixed-point result overflow, floating-point result overflow, floating-point result underflow, or upon invalid floating-point input. The ADSP-21020 has eight user software interrupts which can be activated by setting their respective bits in the interrupt latch register (IRPTL). These are useful for creating multitasking software.

The TMS320C30 can vector to interrupts with no more than five-cycle latency. All interrupts can be masked individually or can all be masked by one master bit (GIE). Each interrupt only has one memory location. Therefore, a four-cycle, non-delayed branch to an interrupt subroutine must be placed at the interrupt vector address. Total interrupt latency is nine instruction cycles. The TMS320C30 does not have a status stack. Arithmetic and mode status must be manually saved when an interrupt occurs. An IDLE instruction is available while waiting for interrupts.

The TMS320C30 has four external interrupt input pins (INT3-0). Internal interrupts can be generated by both internal timers. The TMS320C30 can not generate interrupts upon arithmetic conditions or address register modulo overflows. The TMS320C30 has user software interrupt capability.

Subroutine And Interrupt Returns

At the end of a subroutine or an interrupt routine, program control has to be restored to the calling routine. Return instructions restore the pre-subroutine or pre-interrupt program address. The ADSP-21020 supports conditional subroutine or interrupt returns with an optional compute. Delayed branch returns allow the two instructions after the return to be executed. The TMS320C30 supports conditional subroutine or interrupt returns. It does not support delayed branch returns which incur three cycles of overhead after each return.

Zero Overhead Looping

Most digital signal processing algorithms are vector operations. Due to their repetitive nature, these vector operations are usually coded into short program loops. On most microprocessors, program loops require a branch instruction at the end of the loop for testing the loop condition. Both the ADSP-21020 and the TMS320C30 support zero overhead looping. This feature allows the end loop address and loop condition to be set up before entering the loop.

The ADSP-21020 can support six levels of nesting on zero overhead program loops. This is accomplished with a six-level loop stack. There is a six- level counter stack so that each loop can have its own 32-bit loop counter. The ADSP-21020 can ternimate a loop on loop counter status, arithmetic conditions, flag input pin states, bit test results, or unconditionally. The loop abort (LA) branch feature allows conditional branch out of a loop and automatically have the stacks, which maintain the loop, popped.

The TMS320C30 only supports one zero overhead loop. Zero overhead loop nesting is not supported. Loops can only be terminated on loop counter expired status. If the required loop condition is arithmetic, the loop requires a branch instruction and is not zero overhead. Loops are maintained by the RS (repeat start address register), RE (repeat end address register), and RC (repeat count register). If the main program is executing a zero overhead loop and an interrupt occurs, the interrupt routine must save the loop control registers (RS, RE, RC) before initializing its zero overhead loop if the interrupt routine uses zero overhead looping. The loop control registers must be restored before returning to the main program.

The two loop setup instructions, RPTS (repeat single) or RPTB (repeat block), flush the instruction pipeline and incur four instruction cycles of overhead on the first pass through the loop. RPTS is not interruptable. RPTB is interruptable. The TMS320C30 does not have a loop abort branch feature. If a loop must be aborted, the RC register must be loaded with zero after the branch is taken.

Cache Memory

Both the ADSP-21020 and TMS320C30 have on-chip instruction cache memories to minimize bottlenecks caused when both instructions and data must be accessed from the same external memory space.

The ADSP-21020 has a 32-word instruction cache. The cache is two-way, set associative. It only stores instructions which cause a bus conflict. Put another way, the ADSP-21020 cache memory only stores the last 32 instructions which require program memory data access. In the case of loops which contain more than 32 program memory data accesses, the cache can be frozen. The cache can be disabled allowing use of a logic analyzer to trace all instructions executed.

The TMS320C30 has a 64-word instruction cache. This cache stores all of the past 64 instructions fetched externally no matter if they cause a bus conflict or not. For loops of length greater than 64 instructions, the TMS320C30 cache can be frozen. In the case of the ADSP-21020, loops greater than 64 instructions can exist without bus conflict as long as no more than 32 of the instructions cause a bus conflict. The TMS320C30 can freeze, disable, or clear its cache memory.

Program Sequencing Summary

FEATURE	*21020*	*C30*
Single-cycle/operand branching with 24-bit address	x	x
Pre-modified indirect branching	x	
Conditional branch/subroutine cal with compute	x	
Instruction pipeline levels	3	4
Delayed branch	x	x
Delayed branch subroutine calls	x	
Memory locations at interrupt vector	8	1
Interrupts on index/address register modulo overflow	x	
Interrupts on arithmetic status	x	
Delayed branch returns	x	
Zero overhead loop nesting levels	6	1
Conditional branch with loop abort	x	
Status stack	x	
Instruction cache	x	x

EXTERNAL MEMORY

The ADSP-21020 efficiently interfaces to a large amount of external program and data memory. Any access, whether it be a read or a write, completes in a single cycle unless the programmer explicitly declares wait state values.

The access time requirements of external memory for the ADSP-21020 are not as stringent as those for the TMS320C30. An ADSP-21020 with a cycle time of 50 ns requires a memory access time of 35 ns.

The TMS320C30 is not as efficient with external memory accesses. A write to external memory always takes two cycles and in certain cases a read from external memory takes two cycles.

ASSEMBLY LANGUAGE

For more efficient operation, digital signal processors can be programmed in assembly language. Assembly language syntax usually takes the form of two, three, or four letter mnemonics followed by a register or immediate data field. The following is an example of a TMS320C30 assembly language instruction:

```
SUBI3 R7,R2,R0
```

Translated, this instruction performs an integer subtraction of R0 = R2 – R7 where R0, R2, and R7 are internal data registers. The ADSP-21020 has a much simpler, algebraic assembly language syntax. The same instruction programmed on the ADSP-21020 is:

```
R0 = R2 - R7;
```

If one programmer is trying to decipher another's code, the ADSP-21020 assembly language syntax is easier to understand.

The following comparison illustrates how easy it is to use the ADSP-21020 assembly language.

FUNCTION	*ADSP-21020*	*TMS320C30*
store reg direct	`DM(0x98A1)=R2;`	`STF R2,@98A1h`
load reg indirect	`R3=DM(I4,M1);`	`LDF *AR4++(IR1),R3`

parallel mult/add + data move	`R12=R2*R4,`
	`R8=R8+R12,`
	`R2=DM(I3,M1),`
	`R4=PM(I12,M8);`
		`MPYF3 R7,R4,R0 ||`
		`ADDF3 *-AR3,*AR5-(1),R3`

n example of zero overhead program loop initialization is nother case of the ADSP-21020 assembly language's igibility.

```
LCNTR=99, DO fir_lup UNTIL LCE;
```

lere is an example of zero overhead loop initialization on the MS320C30.

```
LDI  99, RC
RPTS fir_lup
```

ENCHMARKS

All digital signal processors have the ability to execute FIR ilters at one instruction cycle per filter tap. Algorithms which equire general purpose features and instructions intended for)SP algorithms can be used to demonstrate the differences)etween one processor and another. One example is the Fast Fourier Transform (FFT). The table below compares a 1024-)oint complex, radix-2 FFT on the ADSP-21020, ADSP-21010 and TMS320C30.

ADSP-21020 (40 ns)	*ADSP-21010 (80 ns)*	*TMS320C30 (50 ns)*
21314 cycles 0.852 ms	21314 cycles 1.705 ms	50660 cycles 2.533 ms

The ADSP-21020's FFT core loop (FFT butterfly) executes in four instruction cycles, while the TMS320C30's executes in nine. The ADSP-21020's assembly syntax is also more legible.

ADSP-21020 Radix-2 Complex FFT Butterfly (4 instructions)

```
R8=R1*R6,  R14=R11-R14,
             DM(I2,M0)=R10, R9=PM(I11,M8);
R11=R1*R7, R3=R9+R14, R9=R9-R14,
             DM(I2,M0)=R13, R7=PM(I8,M8);
R14=R0*R6, R12=R8+R12,
             R8=DM(I0,M0), PM(I10,M10)=R9;
R12=R0*R7, R13=R8+R12, R10=R8-R12,
             R6=DM(I0,M0), PM(I10,M10)=R3;
```

TMS320C30 Radix-2 Complex FFT Butterfly (9 instructions)

```
   subf  *ar2,*ar0,r2
   subf  *+ar2,*+ar0,r1
   mpy   r2,r6,r5
|| addf  *+ar2,*+ar0,r3
   mpyf  r1,*+ar4(IR1),R3
|| stf   r3,*+ar0
   subf  r0,r3,r4
   mpyf  r1,r6,r0
|| subf  *ar2,*ar0,r3
   mpyf  r2,*+ar4(IR1),r3
|| stf   r3,*ar0++(IR0)
   addf  r0,r3,r5
   stf   r5,*ar2++(IR0)
|| stf   r4,*+ar2
```

The key to the ADSP-21020's efficiency lies in its ability to simultaneously generate the sum and difference of the same two inputs, perform a multiply, store one number in memory, and fetch another from memory all in a single cycle. The TMS320C30 can not store data to memory while performing both a multiply and an add. The TMS320C30 can not simultaneously add and subtract the same two inputs.

SUMMARY

Digital signal processors must meet three basic requirements:

1. Perform multiply/accumulate operations at a single-cycle throughput rate
2. Maintain circular data buffers
3. Maintain loops with zero overhead

Floating-point digital signal processors such as the ADSP-21020 and TMS320C30 meet those requirements differently. As these products evolve, the core architectures will remain the same in order to retain upward compatibility. Therefore, the advantages and disadvantages mentioned above should be carefully analyzed.

This application note can not cover all topics in detail. Consult the *ADSP-21020 User's Manual*, *ADSP-21020 Development Software Manual*, ADSP-21020 data sheet, and ADSP-21010 data sheet for more detailed information.

Analog Devices' DSP bulletin board system (BBS) is available to the public for benchmarks, application notes and general product information.

Analog Devices Bulletin Board System
8 Data Bits, 1 Stop Bit, No Parity
300/1200/2400 Baud
(617)461-4258—On-Line 24 Hours

REFERENCES

Analog Devices, Inc., *ADSP-21020 User's Manual*, 1991.

Texas Instruments, Inc., *Third-Generation TMS320 User's Guide*, 1988.

Papamichalis, P., Simar, R., "The TMS320C30 Floating-Point Digital Signal Processor" in: *Digital Signal Processing Applications with the TMS320 Family, Volume 3*, Texas Instruments, Inc., 1990. Reprinted from IEEE Micro Magazine: Vol. 8, No. 6, December 1988.

Instrumentation Amplifiers
Contents

		Page
AN-244	A User's Guide to IC Instrumentation Amplifiers	10–3
AN-245	Instrumentation Amplifiers Solve Unusual Design Problems	10–11

10

ANALOG DEVICES

AN-244
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

A User's Guide to IC Instrumentation Amplifiers

by Jeffrey R. Riskin

INTRODUCTION

It is traditional to begin a discussion of instrumentation amplifiers by saying that an IA is not an operational amplifier. As obvious as this statement is to the informed user, and as awkward as a description by exclusion may be, such an approach is inevitable and perhaps necessary. When an engineer needs a signal conditioning gain block, the first thought that springs to mind is the nearly ultimate flexibility provided by the currently available assortment of low-cost IC op amps. It may well be that an op amp will suffice as an element in a given gain block, but in demanding applications, op amp circuitry will often require extensive and expensive additional circuit elements, specialized manufacturing and/or test instrumentation together with highly skilled personnel to make it all work. The purpose of this article is to explain when and where an instrumentation amplifier may best be employed and where its unique virtues give it an advantage over the more flexible op amp.

WHAT IS AN INSTRUMENTATION AMPLIFIER?

An instrumentation amplifier is a precision differential voltage gain device that is optimized for operation in an environment hostile to precision measurement. The real world is characterized by deviations from the ideal; temperature fluctuates, electrical noise exists, and voltage drops caused by current through the resistance of leads from remote locations are dictated by the laws of physics. Furthermore, real transducers rarely exhibit zero output impedance and nice neat zero-to-ten-volt ranges. Induced, leaked or coupled electrical interference (noise) is always present to some extent. In brief, even the best "cookbook" must be taken with a grain of salt.

Instrumentation amplifiers are intended to be used whenever acquisition of a useful signal is difficult. IA's must have extremely high input impedances because source impedances may be high and/or unbalanced. Bias and offset currents are low and relatively stable so that the source impedance need not be constant. Balanced differential inputs are provided so that the signal source may be referenced to any reasonable level independent of the IA output load reference. Common mode rejection, a measure of input balance, is very high so that noise pickup and ground drops, characteristic of remote sensor applications, are minimized.

Care is taken to provide high, well-characterized stability of critical parameters under varying conditions, such as changing temperatures and supply voltages. Finally, all components that are critical to the performance of the IA are internal to the device (with the exception of a single gain-determining resistor or resistor-pair). The manufacturer may then optimize, characterize and guarantee the specifications, while the user may in turn depend on a certain level of performance without having to provide his own precision application components or design expertise.

The precision of an IA is provided at the expense of flexibility. By committing to the one specific task of amplifying voltage, the IA manufacturer may optimize performance in this area. An IA is not intended to perform integration, differentiation, rectification, or any other non-voltage-gain function; although possible with an IA, these tasks are best left to operational amplifiers.

To put an instrumentation amplifier to work, the potential user does not require an intimate knowledge of its internal construction. Figure 1, a functional diagram of a basic IA, provides sufficient information for many applications.

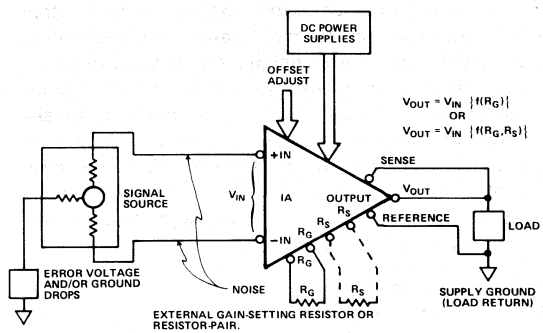

Figure 1. Basic Instrumentation Amplifier Functional Diagram

The two inputs shown permit direct interface to "floating" signal sources. The IA, being truly differential, detects only the difference in voltage between its inputs; any common-mode signals (signals present on both inputs), such as noise

and voltage drops in ground lines, are subtracted and cancelled at the inputs before amplification takes place.*

A single resistor or resistor-pair is used to program the IA for the desired gain. The manufacturer will provide a transfer function or gain equation that allows the user to calculate the required values of resistance for a given gain. Special requirements for that resistor or resistors, if any, are also spelled out by the manufacturer.

The output is single-ended and is designed to drive ground-referenced loads as normally found in measurement equipment. The load reference is common to the power supply return although careful consideration must be given to the overall grounding system (more on that later).

Of course, power must be supplied to the IA; as with op amps, this is normally a differential balanced voltage that may be varied over a specified range.

Most instrumentation amplifiers provide some means of adjusting offset voltage (that dc error voltage present at the output when both inputs are grounded). This adjustment is usually made by varying the setting of an external potentiometer. Sense and reference terminals allow remote sensing of output voltage so that effects of IR drops and ground drops may be minimized. For low current non-remote loads, the sense terminal may be tied directly to the output while the reference terminal may be tied to power supply common. There are other uses for sense and reference that will be discussed in the applications section of this article.

INSIDE AN INSTRUMENTATION AMPLIFIER

While there are many ways of designing an instrumentation amplifier, most such designs can be classified into one of two categories. The most common configuration consists of a number of interconnected operational amplifiers and a precision resistor network. This technique is popular in modular and hybrid instrumentation amplifiers where most practical designs utilize a minimum number of components.

In the other category are designs that, instead of employing op amps, use fundamental active-circuit elements, such as differential circuits and controlled current sources and reflectors; this eliminates all unnecessary or redundant features and tends to minimize active device (transistor) count and decrease the dependence upon accurate resistor matching. This technique is most often employed in the design of monolithic IA's where cost is inversely proportional to chip size. Some older modular IA's also use this technique because suitably precise IC op amps have only recently become readily available. Newer modular IA's may also use this technique because nonlinearity tends to be lower at high gains, although some sacrifice of linearity may exist at lower gains.

Op Amp Based IA's

The most simple (and crude) method of implementing a differential gain block with op amps is shown in Figure 2.

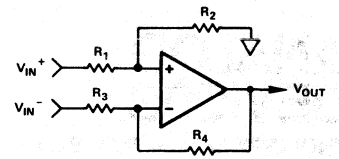

$$V_{OUT} = V_{IN}^{+}\left(\frac{R_2}{R_1+R_2}\right)\left(\frac{R_3+R_4}{R_3}\right) - V_{IN}^{-}\left(\frac{R_4}{R_3}\right)$$

Figure 2. Differential Input Voltage Gain Block (Simple Subtractor)

In this circuit, an expressions for V_{OUT} can be derived by superposition.

The output for V_{IN}^{+} (V_{IN}^{-} grounded) is:

$$V_{O_1} = V_{IN}^{+}\left(\frac{R_2}{R_1+R_2}\right)\left(\frac{R_3+R_4}{R_3}\right) \quad (1)$$

The output for V_{IN}^{-} (V_{IN}^{+} grounded) is:

$$V_{O_2} = -V_{IN}^{-}\left(\frac{R_4}{R_3}\right) \quad (2)$$

By superposition:

$$V_O = V_{O_1} + V_{O_2}$$
$$= V_{IN}^{+}\left(\frac{R_2}{R_1+R_2}\right)\left(\frac{R_3+R_4}{R_3}\right) - V_{IN}^{-}\left(\frac{R_4}{R_3}\right) \quad (3)$$

If $R_2 = R_4$, $R_1 = R_3$:

$$V_O = \left(V_{IN}^{+} - V_{IN}^{-}\right)\frac{R_4}{R_3} \quad (4)$$

Thus, we have created a simple differential voltage amplifier. The input impedances, however, are low and unequal. Furthermore, all 4 resistors have to be carefully ratio-matched to maintain good common mode rejection:

$$V_{OUT\,CM} = V_{OUT} \text{ for } V_{IN}^{+} = V_{IN}^{-}$$
$$= V_{IN}\left[\left(\frac{R_2}{R_1+R_2}\right)\left(\frac{R_3+R_4}{R_3}\right) - \left(\frac{R_4}{R_3}\right)\right] \quad (5)$$

If we are looking for a gain of 1, all resistors will be equal. For a 0.1% mismatch in just one of the resistors:

$$R_1 = R_3 = R_4 = R$$
$$R_2 = 0.999R$$
$$V_{O\,CM} = V_{IN}\left[\left(\frac{0.999R}{1.999R}\right)\left(\frac{2R}{R}\right) - \left(\frac{R}{R}\right)\right]$$
$$= 0.0005V_{IN} \quad (6)$$
$$CMR = 66dB$$

(Note that if the source resistance is not low and balanced, gain and CMR will be further degraded.)

*For applications involving extremely high common-mode voltages, or requiring complete galvanic isolation, *isolation amplifiers* should be used. Analog Devices manufactures a complete line of single and multi-channel isolators.

Considering what is available in the way of reasonably priced standard resistors, one can hardly expect to improve upon this mediocre level of performance. Considering the several serious drawbacks, it is not surprising that this configuration is not used in true instrumentation amplifier designs.

The two-amplifier approach shown in Figure 3 overcomes some of the weaknesses inherent in the simple subtractor of Figure 2.

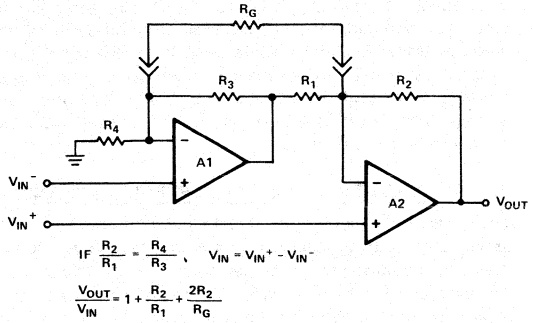

Figure 3. "Two-Amplifier" Instrumentation Amplifier

Input resistance is high, thus permitting the signal sources to have unbalanced, non-zero output impedance. Furthermore, gain may be changed by switching only one resistor thus allowing CMR to remain constant once initial trimming is accomplished. (CMR is still dependent upon the ratio-matching of four resistors.) The major disadvantage to this design is that the common mode voltage input range is a function of gain and can thus be very poor. By referring to Figure 3, it can be seen that A1 is called upon to amplify a common mode signal by the ratio $(R_3 + R_4)/R_4$; this could lead to saturation of A1 thus leaving no "headroom" to amplify the differential signal of interest. A few modules and hybrids use this configuration because of its simplicity, but it is not optimal.

The most popular configuration for op-amp based instrumentation amplifiers is shown in Figure 4:

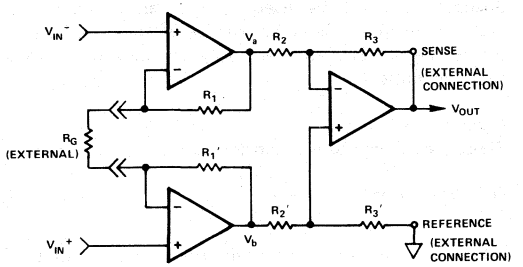

Figure 4. "Classic" 3 Op Amp Instrumentation Amplifier

The transfer function of this circuit can be calculated by superposition.

For $V_{IN}^+ = 0$

$$V_a = V_{IN}^- \left(\frac{R_1 + R_G}{R_G}\right) \quad (7)$$

$$V_b = V_{IN}^- \left(\frac{R_1'}{R_G}\right) \quad (8)$$

For $V_{IN}^- = 0$

$$V_a = V_{IN}^+ \left(\frac{R_1}{R_G}\right) \quad (9)$$

$$V_b = V_{IN}^+ \left(\frac{R_1' + R_G}{R_G}\right) \quad (10)$$

$$\therefore \quad V_a = V_{IN}^- \left(\frac{R_1 + R_G}{R_G}\right) - V_{IN}^+ \left(\frac{R_1}{R_2}\right) \quad (11)$$

$$\text{and } V_b = V_{IN}^+ \left(\frac{R_1' + R_G}{R_G}\right) - V_{IN}^- \left(\frac{R_1'}{R_G}\right) \quad (12)$$

$$V_{OUT} = -\left(\frac{R_3}{R_2}\right) V_a + V_b \left(\frac{R_3'}{R_2' + R_3'}\right)\left(\frac{R_3 + R_2}{R_2}\right) \quad (13)$$

If $R_3 = R_3'$, $R_2 = R_2'$, and $R_1 = R_1'$

$$V_{OUT} = (V_b - V_a)\left(\frac{R_3}{R_2}\right) \quad (14)$$

substituting for V_b and V_a and simplifying

$$V_{OUT} = (V_{IN}^+ - V_{IN}^-)\left(\frac{2R_1}{R_G} + 1\right)\left(\frac{R_3}{R_2}\right) \quad (15)$$

In this configuration, gain accuracy and CMR still depends upon the ratio-matching of R_2, R_2', R_3 and R_3'. It can be shown, however, that CMR does *not* depend on the matching of R_1 and R_1'.

10

$$V_{CM\ OUT} = (V_a - V_b) = V_{IN}^+ \left(\frac{R_1' + R_G}{R_G}\right) - V_{IN}^- \left(\frac{R_1'}{R_G}\right) - V_{IN}^- \left(\frac{R_1 + R_G}{R_G}\right) + V_{IN}^+ \left(\frac{R_1}{R_G}\right) \quad (16)$$

but $V_{CM\ IN} = V_{IN}^+ = V_{IN}^-$

$$V_{CM\ OUT} = V_{CM\ IN}\left[\frac{R_1' + R_G}{R_G} - \frac{R_1'}{R_G} - \frac{R_1 + R_G}{R_G} + \frac{R_1}{R_G}\right] \quad (17)$$

$$= V_{CM\ IN}\left[\frac{R_1'}{R_G} - \frac{R_1'}{R_G} + 1 - \frac{R_1}{R_G} + \frac{R_1}{R_G} - 1\right] \quad (18)$$

$$= V_{CM\ IN}\left[0\right]$$

$$= 0$$

Therefore, in theory at least, the user may take as much gain in the front end as he wishes (as determined by R_G) without increasing the common mode error signal. Thus, CMRR will theoretically increase in direct proportion to gain, a very useful property. Furthermore, common-mode signals are only amplified by a factor of 1 regardless of gain because no common-mode voltage will appear across R_G, hence, no common-mode current will flow in it (the input terminals of an op-amp operating normally will have no significant potential difference between them). This means that large common-mode signals may be handled independent of gain.

Finally, because of the symmetry of this configuration, first order common-mode error sources in the input amplifiers, if they track, tend to be cancelled out by the output stage subtractor. These features explain the popularity of this IA design technique.

IA's of this type may use either FET or Bipolar input operational amplifiers. FET input devices have very low bias currents and are well-suited for use with very high source impedances. FET input op amps, however, generally have poorer CMR than bipolar amplifiers due to non-geometry related mis-matches. (In other words, matching of FET's is largely a function of process control; matching bipolar transistors is less process dependent.) This will manifest itself in lower linearity and CMR for large input voltages. Furthermore, these mis-matches usually cause larger input offset voltage drifts. For these reasons, Analog Devices instrumentation amplifiers use bipolar input stages thus sacrificing low bias currents to achieve high linearity and CMR along with low input offset voltage drift. As technology develops, FET input IA's may become more viable.

Dedicated Design IA's

The second category of IA design is based on minimum active device count; a virtue for monolithic IC circuits. The basic schematic for such a design is shown in Figure 5.

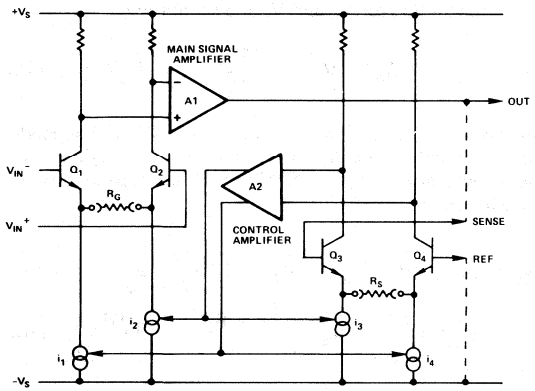

Figure 5. Typical IC IA Basic Schematic

Forward gain is provided by the input differential stage Q_1 and Q_2 whose current gain (transconductance) is $1/R_G$ (amps/volt) and the main signal amplifier A_1 which senses differences in input stage collector currents. When the output is connected back to sense (with reference grounded) differential stage Q_3 and Q_4 acts as a feedback error-sensing amplifier with a transconductance of $1/R_S$ (amps/volt). A_2 senses the collector current imbalance in that stage.

When a differential voltage is applied to the inputs, the collector currents of Q_1 and Q_2 tend to become unbalanced by $(V_{IN}^+ - V_{IN}^-)/R_G$. This is sensed by A_1 which develops an error voltage between the sense and reference points. This, in turn, tries to unbalance the collector currents in Q_3 and Q_4 by $(V_{SENSE} - V_{REF})/R_S$. That unbalance is sensed by A_2 which then adjusts I_3 and I_4 to equalize the collector currents in Q_3 and Q_4 ($I_4 - I_3 = (V_S - V_R)/R_S$). A_2 simultaneously adjusts I_1 and I_2 such that $I_1 - I_2 = I_4 - I_3$. Balance is reached when:

$$\frac{V_S - V_R}{(I_4 - I_3)\,R_S} = \frac{V_1 - V_2}{(I_1 - I_2)\,R_G} \qquad (19)$$

$$\text{if } \frac{V_S - V_R}{V_1 - V_2} = \frac{V_{OUT}}{V_{IN}} = \text{Gain} \qquad (20)$$

and $I_4 - I_3 = I_1 - I_2$

$$\text{Gain} = \frac{R_{SCALE}}{R_{GAIN}} \qquad (21)$$

It is apparent from this analysis that the requirement for carefully matched resistors changes to a requirement for carefully matched active devices. In IC technology, this is possible by utilization of precision photographic techniques along with careful design layout and well-controlled processing. The result is a good trade-off between high performance and low cost.

INSTRUMENTATION AMPLIFIER SPECIFICATIONS

To successfully apply any electronic component, a full understanding of its specifications is required. That is to say, the numbers contained in a spec sheet are of little value if the user doesn't have a clear picture of what each spec means. In this section, a typical instrumentation amplifier specification sheet will be reviewed. Each individual specification will be discussed in terms of how it is measured and what error it might contribute to the overall performance of the circuit. In some cases, a given specification may not affect a particular application; the more common situations of this type will be discussed.

At the top of the spec sheet is the statement that the listed specs are typical @ $V_S = \pm 15V$, $R_L = 2k\Omega$ and $T_A = +25°C$ unless otherwise specified. This tells the user that these are the normal operating conditions under which the device is tested. Deviations from these conditions might degrade (or improve) performance. When deviations from the "normal" conditions are likely (such as a change in temperature) the significant effects are usually indicated within the specs. This statement also tells us that all numbers are typical unless noted; "typical" means that the manufacturers characterization process has shown this number to be average, but individual devices may vary.

Specifications not discussed in detail are self-explanatory and require only a basic knowledge of electronic measurements. Those specs do not apply uniquely to instrumentation amplifiers.

Gain

These specs relate to the transfer function of the device.

$$\text{Gain Equation: } G = 1 + \frac{2(10^5)}{R_G} \qquad (22)$$

To select an R_G for a given gain, solve the equation for R_G

$$\text{(in ohms): } R_G = \frac{200{,}000}{G - 1} \qquad (23)$$

For example:

$G = 1 \quad : R_G = \infty$ (open circuit)

$G = 10 \quad : R_G = 22{,}222\Omega$

$G = 100 \quad : R_G = 2020.2\Omega$

$G = 1000: R_G = 200.20\Omega$

Of course the user must provide a very clean circuit board to realize an accurate gain of 1 since 200MΩ leakage resistance will cause a gain error or 0.1%.

Gain Range

Specified at 1 to 1000, this device may (and in fact will) work at higher gains, but the manufacturer will not promise any particular level of performance. In practice, noise and drift may make higher gains impractical for this device.

Equation Error

The number given by this specification describes maximum deviation from the gain equation. The user can trim the gain (above unity) or can compensate elsewhere in his design. If his data is eventually digitized and fed to an "intelligent system" (such as a microprocessor), he might be able to correct for gain errors by measuring a reference and multiplying by a constant.

Nonlinearity

Nonlinearity is defined as the deviation from a straight line on the plot of output versus input. Figure 6a shows the transfer function of a device with exaggerated nonlinearity. The magnitude of this error can be calculated thus:

$$\text{N. L.} = \left[\frac{\text{Actual Output - Calculated Output}}{\text{Rated Full-Scale Output Range}} \right]$$

To confuse matters, this deviation can be specified relative to *any* straight line or to a specific straight line. There are two commonly-used methods of specifying this ideal straight line relative to the performance of a precision measurement device.

The "Best Straight Line" method of nonlinearity specification consists of measuring the peak positive and negative deviations and adjusting the slope of the device transfer function (by adjusting the gain and offset) so that these maximum positive and negative errors are equal. This method yields the best specifications but is difficult to implement in that it requires that the user examine the entire output signal range to determine these maximum positive and negative deviations. The results of a best-straight-line calibration is shown by the transfer function of Figure 6b.

The "End-Point" method of specifying nonlinearity requires that the user perform his offset and/or gain calibrations at the extremes of the output range. This is much easier to implement but may result in nonlinearity errors of up to twice these attained with best-straight-line techniques. This worst case will occur when the transfer function is "bowed" in one direction only. Figure 8c shows the results of end-point calibration.

Most linear devices, such as instrumentation amplifiers, are specified for best-straight-line linearity. The user must take this into consideration when evaluating the error budget for his application.

Regardless of the method used to specify nonlinearity, the errors thus created are irreducible. That is to say that these errors are neither fixed nor proportional to input or output voltage and can not be reduced by adjustment.

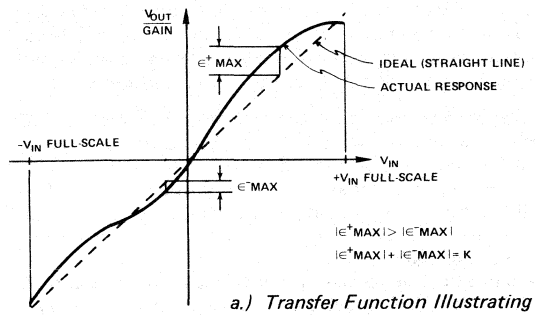

a.) Transfer Function Illustrating Exaggerated Nonlinearity

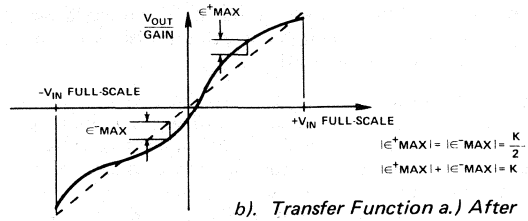

b). Transfer Function a.) After Calibration by Best-Straight-Line Method

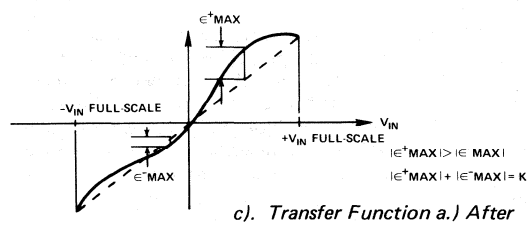

c). Transfer Function a.) After Calibration by End-Point Method

Figure 6. Nonlinear Transfer Function

Gain vs. Temperature

These numbers give both maximum and typical deviations from the gain equation as a function of temperature. An intelligent system can correct for this with an "auto-gain" cycle (measure a reference and re-normalize).

Settling Time

Settling time is defined as that length of time required for the output voltage to approach and remain within a certain tolerance of its final value. It is usually specified for a fast full scale input step and includes output slewing time. Since several factors contribute to the overall settling time, fast settling to 0.1% doesn't necessarily mean proportionally fast-settling to 0.01%. In addition, settling time is not necessarily a function of gain. Some of the contributing factors

include slew rate limiting, under-damping (ringing) and thermal gradients ("long tails").

Voltage Offset

Voltage offset specifications are often considered a figure of merit for instrumentation amplifiers. While initial offset may be adjusted to zero, shifts in offset voltage could cause errors. Intelligent systems can often correct for this factor with an auto-zero cycle, but there are many small-signal high-gain applications that don't have this capability.

Voltage offset and offset dritt comprise two components each; input and output offset and offset drift. Input offset is that component of offset that is directly proportional to gain, i.e., input offset as measured at the output at G = 100 is 100 times greater than at G = 1. Output offset is independent of gain. At low gains, output offset drift is dominant, while at high gains input offset drift dominates. Therefore, the output offset voltage drift is normally specified as drift at G = 1 (where input effects are insignificant), while input offset voltage drift is given by drift specification at a high gain (where output offset effects are negligible). All input-related numbers are referred to the input (RTI) which is to say that the effect on the output is "G" times larger. Voltage offset vs. power supply is also specified at one or more gain settings and is also RTI.

Input Bias Currents

Input bias currents are those currents necessary to bias the input transistors of a dc amplifier. FET input devices have lower bias currents, but those currents increase dramatically with temperature, doubling approximately every 11°C. Since bias currents can be considered as a source of voltage offset (when multiplied by source resistance), the change in bias currents is of more concern than the magnitude of the bias currents. Input offset current is the difference between the two input bias currents.

Although instrumentation amplifiers have differential inputs, there must be a return path for the bias currents. If this is not provided, those currents will charge stray capacitances, causing the output to drift uncontrollably or to saturate. Therefore, when amplifying "floating" input sources such as transformers and thermocouples, as well as ac-coupled sources, there must still be a dc path from each input to ground.

Common-Mode Rejection

Common-mode rejection is a measure of the change in output voltage when both inputs are changed equal amounts. These specifications are usually given for a full-range input voltage change and a specified source imbalance. "Common-mode rejection ratio" (CMRR) is a ratio expression while "common-mode rejection" (CMR) is the logarithm of that ratio. For example, a CMRR of 10,000 corresponds to a CMR of 80dB.

In most IA's the CMRR increases with gain. This is because most designs have a front-end configuration that does not amplify common-mode signals. Since the standard for CMRR specifications is referred to the output (RTO), a gain for differential signals in the total absence of gain for common-mode output signals will yield a 1-to-1 improvement of CMRR with gain. This means that the common-mode output error signal will not increase with gain, it does not mean that it decreases with gain! At higher gains, however, amplifier bandwidth decreases. Since differences in phase-shift through the differential input stage will show up as a common-mode error, CMRR becomes more frequency dependent at high gains.

Error Budget Analysis

To illustrate how instrumentation amplifier specifications are applied, we will now examine a typical case where an AD522 is required to amplify the output of an unbalanced transducer.

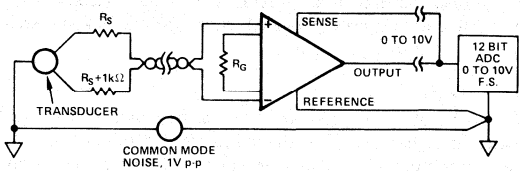

Figure 7. Typical IA Application

Figure 7 shows a differential transducer, unbalanced by 1kΩ, supplying a 0 to 1 volt signal to a remotely located IA. The output of the IA feeds a 12 bit A to D converter with a 0 to 10 volt input voltage range. There is 1 volt of peak-to-peak 0 to 10Hz noise on the ground return appearing as a common-mode signal at the inputs of the IA. The operating temperature range is −25°C to +85°C; calibration is performed at +25°C.

The input signal must be amplified by a factor of 10 in order to utilize the full resolution of the A to D converter. Solving the gain equation for G = 10 gives a value of 22.22kΩ for R_G.

Table 2 lists all applicable error sources and their corresponding effects on accuracy. Initial errors are defined as those errors that can be reduced to a negligible amount by performance of an initial calibration.

Reducible errors include these initial errors along with other errors that occur during normal operation that may be corrected by an adaptive or "intelligent" system. For example, changes in gain or offset may be measured during an auto-zero/auto-gain cycle by measuring two known voltages (a precision reference and ground, for example). This is a common practice in computer or processor-controlled equipment.

Irreducible errors are errors which can not be readily corrected either at initial calibration or in use. It could be argued that an array of precision references would permit a software linearity correction, but in most applications that would be unrealistically cumbersome.

The total error "as built" is approximately 5540ppm or 0.55%. If an initial calibration is performed, this number is reduced by 2210ppm to 3330ppm = 0.33%. Note that 3000ppm of this is gain drift.

In many applications, differential linearity and resolution are of prime importance. This would be so in cases where the absolute value of a variable is less important than changes in value. In these applications, only the irreducible errors (57.8ppm = 0.006%) are significant. Furthermore, if a system has an intelligent processor monitoring the A to D output, the addition of a auto-gain/auto-zero cycle will remove all reducible errors and may eliminate the requirement for initial calibration. This will also reduce errors to 0.006%.

In the above example, the system can justifiably make use of a 13 bit A to D converter for its differential linearity and

resolution. Dynamic range exceeds 84dB (14 bits). Absolute accuracy depends on calibration and system interaction capabilities; it might be as good as the resolution (0.006%) or as poor as the initial accuracy (0.55%).

INSTRUMENTATION AMPLIFIER APPLICATIONS

General Considerations

Whenever a precision high-gain device—such as an instrumentation amplifier—is used, certain precautions apply. Obviously, it is wise to have a clean layout, short wire runs where possible and a carefully considered grounding scheme. Figure 8 shows a well-thought out approach to IA interconnection.

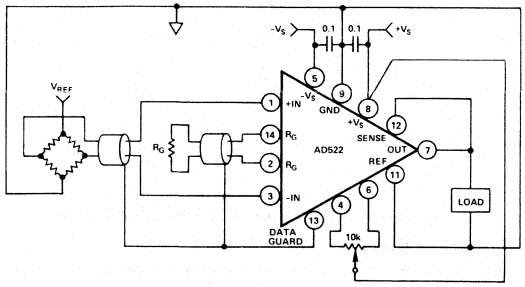

Figure 8. AD522 Interconnection

A properly designed instrumentation amplifier exhibits low sensitivity to power supply variations; the AD522, for example, shows an RTI offset variation of only 0.2μV per percent of power supply change at G = 1000. At increasing frequencies, however, this rejection factor will degrade as internal capacitances permit more power supply noise to find its way into the signal path. This effect can be minimized by bypassing the power supplies, as close to the IA as possible, with 0.1μF ceramic disc capacitors. Larger tantalum capacitors would be effective against lower frequency variations, but a competent IA is capable of rejecting most of these slower changes.

The offset adjustment pot usually affects the balance of the high gain differential input stage. Short wire runs to this pot will minimize injection of noise into a sensitive location.

The gain-determining resistor, R_G, is often remotely located for purposes of gain switching. A well-designed IA will tolerate this to a certain extent, but stray capacitances and wiring inductance may disturb the frequency compensation of the device. Sometimes it becomes necessary to install a series RC right at the R_G terminals of the IA to add a compensating zero to correct for LC resonances caused by stray inductances and capacitances. This lead compensation may improve stability at the cost of a peak in the frequency response curve at the high end. Unfortunately, this compensation, if required, depends on the individual application and is usually determined experimentally.

Most IA's are provided with "sense" and "reference" outputs. While there are several interesting uses for these features (to be discussed later), the most basic application is remote load sensing. This essentially puts the IR drops "inside the loop" of the IA and is most useful when driving remote and/or heavy loads or when the load ground is not firmly "anchored" to the power supply returns.

Grounding is a topic worthy of its own application note (see "An IC Amplifier User's Guide to Decoupling, Grounds, etc." by A. P. Brokaw). In the case of instrumentation amplifiers, the main thing to remember is that all signal and power returns must eventually have a direct or indirect common point. Direct coupling of IA inputs make it necessary to provide signal ground returns for input amplifier bias currents. Figure 8 shows a direct connection. If a "floating" source or ac coupling is used, indirect returns similar to those shown in Figure 9 must be provided.

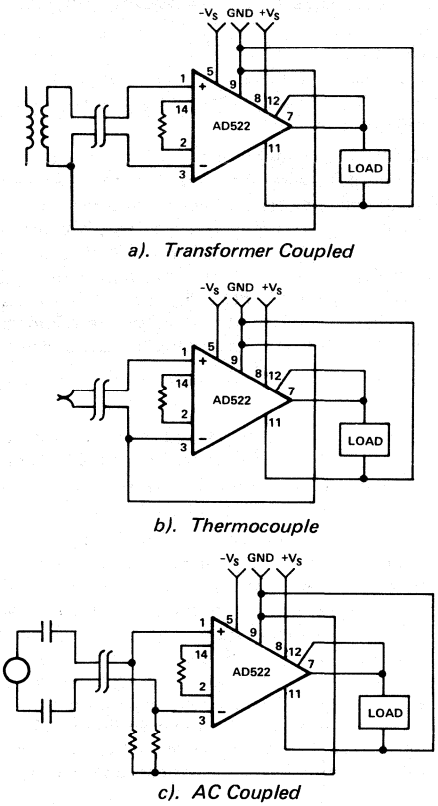

Figure 9. Indirect Ground Returns for "Floating" Transducers

Signals from remote transducers are often transmitted to the IA through shielded cables. While this may well serve to reduce noise pick-up, the distributed RC's in such cabling can cause differential phase shifts in those lines. When ac common-mode signals are present, these phase shifts will reduce common-mode rejection. The same effect will occur with remote R_G's located at the end of shielded cables. If the shields could be driven by the common-mode signal, the cable capacitance could be "boot-strapped" thus making the capacitance effectively zero for common-mode signals. The data guard output of the AD522 provides the common-mode component of the input signals and can be used to drive the shields of coaxial input cables and increase ac CMR. Figure 8 illustrates this connection; if not used, the data guard should be left unconnected.

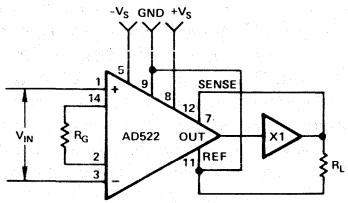

Figure 10. Current-Booster Output

Boosted Output

In the previous section, use of the sense terminal for remote load sensing was discussed. Another use of that terminal is illustrated in Figure 10.

Typically, IC instrumentation amplifiers are rated for a full ±10 volt output swing into 2kΩ. In some applications, however, the need exists to drive more current into heavier loads. Figure 10 shows how a high-current booster may be connected "inside the loop" of an instrumentation amplifier to provide the required current boost without significantly degrading overall performance. Nonlinearities, offset and gain inaccuracies of the buffer are minimized by the loop gain of the IA output amplifier. Offset drift of the buffer is similarly reduced.

Offset Load

The reference terminal may be used to offset the output by up to ±10V. This is useful when the load is "floating" or does not share a ground with the rest of the system. It also provides a direct means of injecting a precise offset.

Two caveats apply to the use of the reference pin. When the IA is of the three-amplifier configuration shown in Figure 4 (as is the AD522), it is necessary that nearly zero impedance be presented to the reference terminal. It can be shown that any significant resistance from the reference terminal to ground increases the gain of the non-inverting signal path thereby upsetting the common-mode rejection of the IA. An operational amplifier may be used to provide that low impedance reference point as shown in Figure 11. The input offset voltage characteristics of that amplifier will add directly to the offset voltage performance of the IA.

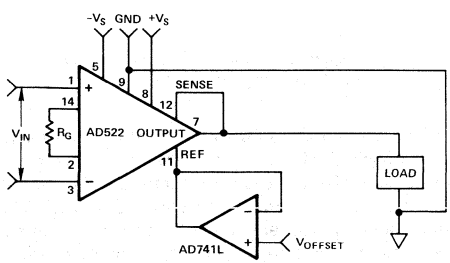

Figure 11. Use of Reference Terminal to Provide Output Offset

The other precaution is more obvious. The output voltage range of an IA is clearly specified; if that range is mostly used up by offset at the reference terminal not much range is left for the signal. In other words, the sum of the offset and signal may not exceed the specified output voltage range of the IA.

CMR Trim

The effect of resistance in the reference termination may be used to advantage. A short-term CMR improvement can be realized with the circuit shown in Figure 12.

While applying a low-frequency 20 volt peak-to-peak input signal to both inputs, the pot should be adjusted for an output null. In many cases this adjustment will not improve matters on a long-term basis since the common-mode rejection of the device is determined by the long-term stability of internal components (which will drift regardless of what happens externally).

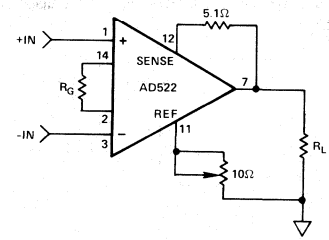

Figure 12. Common Mode Rejection Trim

Controlled Currents

An instrumentation amplifier can be turned into a voltage-to-current converter by taking advantage of the sense and reference terminals as shown in Figure 13.

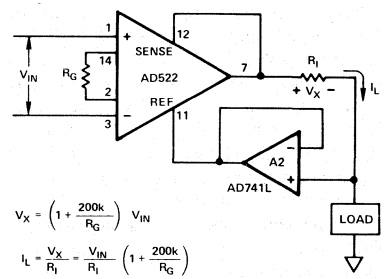

Figure 13. Voltage-To-Current Converter

By establishing a reference at the "low" side of a current setting resistor, an output current may be defined as a function of input voltage, gain and the value of that resistor. Since only a small current is demanded at the input of the buffer amplifier A_2, the forced current I_L will largely flow through the load. Offset and drift specifications of A_2 must be added to the output offset and drift specifications of the IA.

CONCLUSIONS

Thus characterized, the instrumentation amplifier stands ready to take its place in the Grand Order of Things. The preliminary contention that an IA is not a special sort of operational amplifier should now be obvious. Its versatility is limited in scope but its applications are limited only by the imagination of the potential user. As a precision linear device, an IA is qualified mainly by its specifications, a full understanding of which is necessary to successfully use it to advantage. Analog Devices, as a long time supplier of components for precision measurement applications, offers a full spectrum of instrumentation amplifiers in modular, hybrid and monolithic IC form, each ideally suited to particular applications. We hope that this article will help clarify the issues involved and will aid in the selection of a suitable device for a particular application.

ANALOG DEVICES

AN-245
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Instrumentation Amplifiers Solve Unusual Design Problems

by Scott Wurcer and Walt Jung

Traditionally considered only for transducer-conditioning applications, instrumentation amplifiers bring unique performance benefits to a range of other applications as well.

Instrumentation amplifiers traditionally find use in transducer-signal-conditioning applications that require gain blocks with true differential inputs and single-ended outputs. But the same instrumentation-amp performance characteristics that suit sensor-based circuits—such as strain-gauge and thermocouple amplifiers—also benefit many less traditional analog-signal-

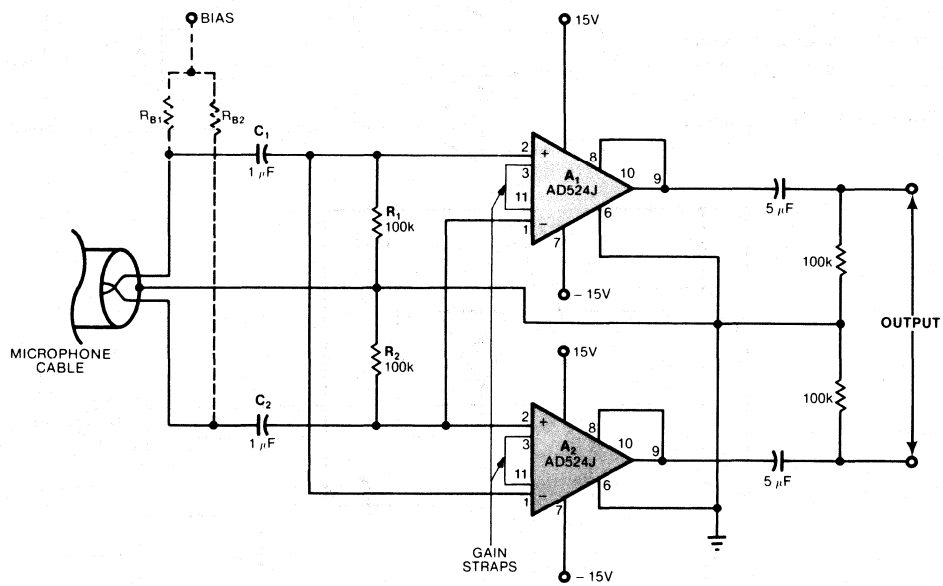

Fig 1—An active-transformer microphone preamp *uses two instrumentation amps for reduced noise. In this circuit, the instrumentation amps' output signals add linearly, while noise adds according to the square root of two. Thus, the circuit provides 3-dB better noise performance than do single-ended designs.*

Reprinted from EDN — August 4, 1983

Instrumentation amplifiers find diverse uses

processing applications. Indeed, as the circuits presented here illustrate, emerging low-cost monolithic instrumentation amplifiers often prove preferable to op amps as basic analog-circuit building blocks.

IC instrumentation amps' advantages stem from their use of multiple op amps to provide balanced inverting and noninverting inputs, high common-mode-rejection ratios (CMRRs) and programmable front-end gains (see **box**, "Three op amps are better than one"). High-performance devices combine low input noise voltages with wide bandwidths and stable dc characteristics—attributes that don't coexist in most monolithic op amps. Thus, instrumentation amps frequently outperform op amps in applications as simple as noninverting gain stages (see **box**, "The instrumentation amp as op amp").

Simplifying mike preamps

But instrumentation amps' most interesting applications involve circuits entailing performance problems that op amps typically can't solve. Microphone preamplifiers, for example, require active stages exhibiting

Three op amps are better than one

Instrumentation amplifiers typically consist of three op amps configured as a dual differential input stage and an output buffer (**figure,** part **(a)**). Internal feedback resistors connected to the input amplifiers allow one resistor (R_G) to set overall gain. Assuming a unity-gain output configuration, overall gain equals $((R_1+R_2)/R_G)+1$. Sense and reference inputs allow output gain scaling and offsetting.

The monolithic instrumentation amps used in this article's circuits, Analog Devices's AD524 and AD624 (**figure,** part **(b)**), possess the same gain characteristics as the classic 3-op-amp circuit but provide a few added features. They include transistor preamps (Q_1 and Q_2) that precede input amplifiers A_1 and A_2. As a result, input gain-bandwidth product depends on transistor transconductance, so the circuits provide greater high-gain bandwidth than do conventional instrumentation amps. The transistor preamps also increase overall instrumentation-amp gain and thus reduce closed-loop errors.

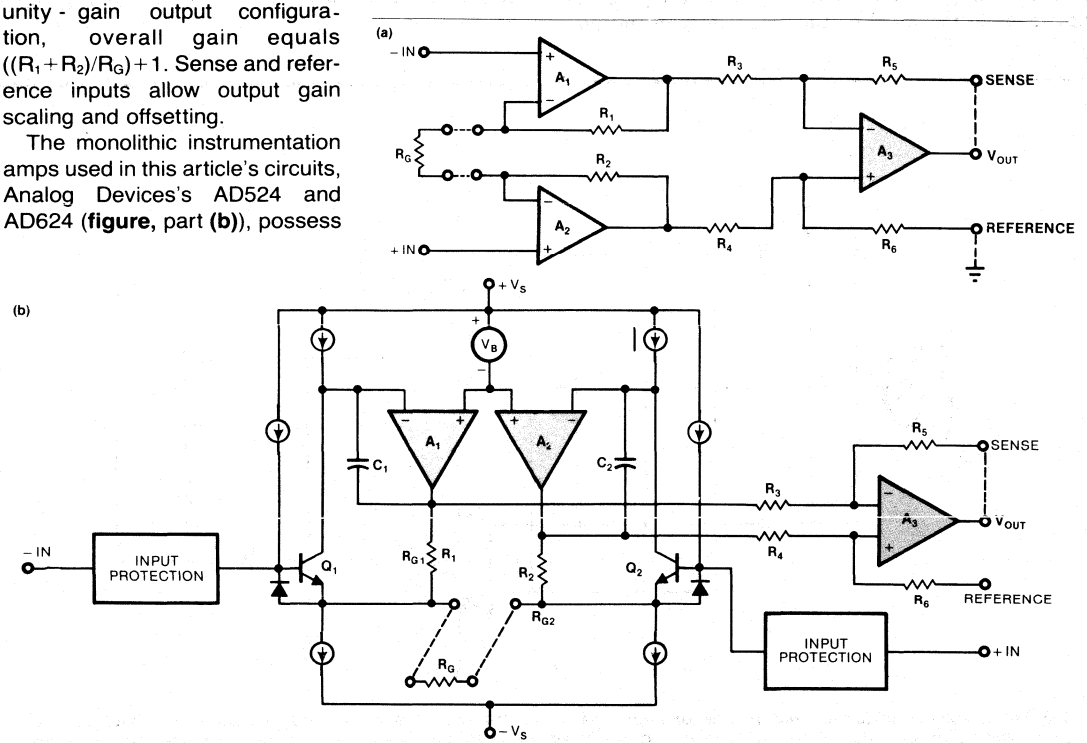

Three op amps *compose a typical instrumentation amp* **(a).** *Internal feedback resistors around the input amplifiers let one resistor (R_G) set overall circuit gain. The monolithic instrumentation amps used in this article's circuits* **(b)** *include transistor preamps that extend circuit bandwidth and reduce closed-loop distortion.*

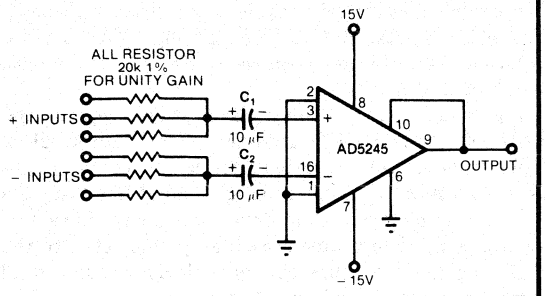

Fig 2—Multiple currents *are summed in this current-mode mixer. The circuit exploits the local feedback around inverting and noninverting instrumentation-amp inputs to allow signal addition and subtraction.*

both low noise and high CMRR and thus place excessive demands on conventional op amps. By providing dual differential inputs, low-noise input stages and easily settable gain, instrumentation amps greatly simplify mike-preamp design.

Fig 1 illustrates a preamplifier built from two Analog Devices AD524J monolithic instrumentation amps. Similar to the crossed-input stacked amplifier described in **Ref 1**, the circuit functions as a balanced-input/balanced-output active transformer. Its differential output configuration provides a 3-dB signal-to-noise-ratio increase compared with single-ended designs: The 2-amp configuration produces twice the signal voltage of a single-amp circuit while generating only $\sqrt{2}$ times as much noise. The preamp's differential output arrangement also suits non-audio-related applications. For instance, the circuit can function as a general-purpose active transformer that feeds instrumentation amps further down the signal paths.

For maximum application flexibility, the instrumentation amps' programmable front ends allow overall circuit gains of 20, 200 or 2000. When configured as shown for a gain of 20, the preamp provides 40V p-p output swings, 80-dB CMRR, 0.6-μV-rms noise over a 20-kHz bandwidth, and harmonic distortion less than 0.01%.

The preamp's inputs are ac coupled to a source microphone through capacitors C_1 and C_2; resistors R_1 and R_2 serve as return paths for the instrumentation amps' input bias currents. The coupling capacitors furnish ac and dc overload protection for levels to a few volts, but because input transients can become quite large (particularly when mike cables get connected or disconnected), you should ensure that the amplifiers remain adequately protected. The AD524 includes integral protection to 36V; other instrumentation amps might require 1-kΩ or larger resistors in series with the coupling capacitors. Note, though, that input resistors entail a noise penalty. A 1-kΩ resistor generates 4-nV/$\sqrt{\text{Hz}}$ equivalent input noise, which adds in rms fashion to instrumentation-amp input noise (4.7-nV/$\sqrt{\text{Hz}}$ for the AD524 in this circuit).

When constructing the mike preamp, be certain to shield the coupling capacitors and their associated input circuitry. In addition, for minimum hum pickup, use balanced, twisted-pair input cables.

Current summer combines audio signals

Fig 2 illustrates another instrumentation-amp application involving audio-frequency signals. Here, an instrumentation-amp-based current summer linearly combines an arbitrary number of input waveforms, allowing independent sign selection for each signal. The summer accommodates individual gain settings for each discrete input, and gain values can exceed or remain less than unity.

To provide dual-polarity summing, the circuit exploits the local feedback present in each of the instrumentation amp's input stages. Because the amplifier's input transistors incorporate local emitter feedback, when the transistor bases (pins 1 and 2) connect to ground, the total current fed into the emitters (pins 3 and 16) produces linearly dependent output voltages. The AD524 incorporates 20-kΩ feedback resistors, so unity gain results from the use of 20-kΩ input resistors.

In **Fig 2**'s circuit, pin 16 serves as a noninverting input and pin 3 acts as an inverting node. You can use the two inputs separately to sum single-ended currents, or together to algebraically combine differential inputs. Because of internal feedback connections, each input's dynamic impedance remains less than 1Ω at frequencies less than 1 kHz, so crosstalk between individual signal channels remains low. With 20-kΩ input resistors (unity gain), for example, interchannel crosstalk remains 86 dB below input levels.

Because the input transistors' grounded bases force emitter voltages to a level one V_{BE} drop less than ground, you must use input-coupling capacitors to ensure low offsets. Choose the values of these capacitors so that their reactances at frequencies of interest remain low compared with the input resistors' values; otherwise, gain errors result.

Thanks to the instrumentation amp's low input noise, the summer's signal-to-noise ratio over a 20-kHz bandwidth exceeds 90 dB relative to a 1V input. Moreover, total harmonic distortion remains less than 0.01% over the entire audio range. Gain accuracy, however, depends on matching between input resistors

and the instrumentation amp's internal feedback elements. The AD524's 20% feedback-resistor tolerance leads to uncompensated gain errors of ±1.6 dB in circuits set for unity gain; you can eliminate such errors by providing input-resistor trimming potentiometers. Note, though, that interchannel gain matching isn't affected by internal tolerances. Because all inputs feed a common node, gain matching depends only on input-resistor tolerances.

Note also that you can use the summer with current-mode analog switches, such as those found in DACs. In such applications, you can preserve dc coupling throughout the circuit by connecting the input transistors' bases to a positive voltage equal to the devices' V_{BE}. In that case, the emitters rest at ground.

Buffer a bucket brigade

Instrumentation amplifiers' true differential characteristics also prove beneficial in circuits incorporating analog delay lines. For minimum clock feedthrough and distortion in such circuits, dual CCD bucket-brigade registers—such as Reticon Corp's SAD-1024—provide two independently addressable delay lines with separate clock inputs. But fully exploiting CCD devices' differential capabilities requires buffer amplifiers with well-balanced inputs.

Fig 3 illustrates a delay circuit that employs the SAD-1024 512×2 CCD register in a differential-in/differential-out configuration. The circuit's input amplifiers, two AD524s, convert single-ended input signals to differential levels that drive the delay line's two inputs. A third instrumentation amp connected to the CCD register's outputs converts delayed differential signals to single-ended waveforms.

The delay line's clock inputs get driven by a pair of CMOS inverters that furnish buffered, complementary clock signals. To satisfy Nyquist's sampling criterion, the clock rate must remain at least twice the highest

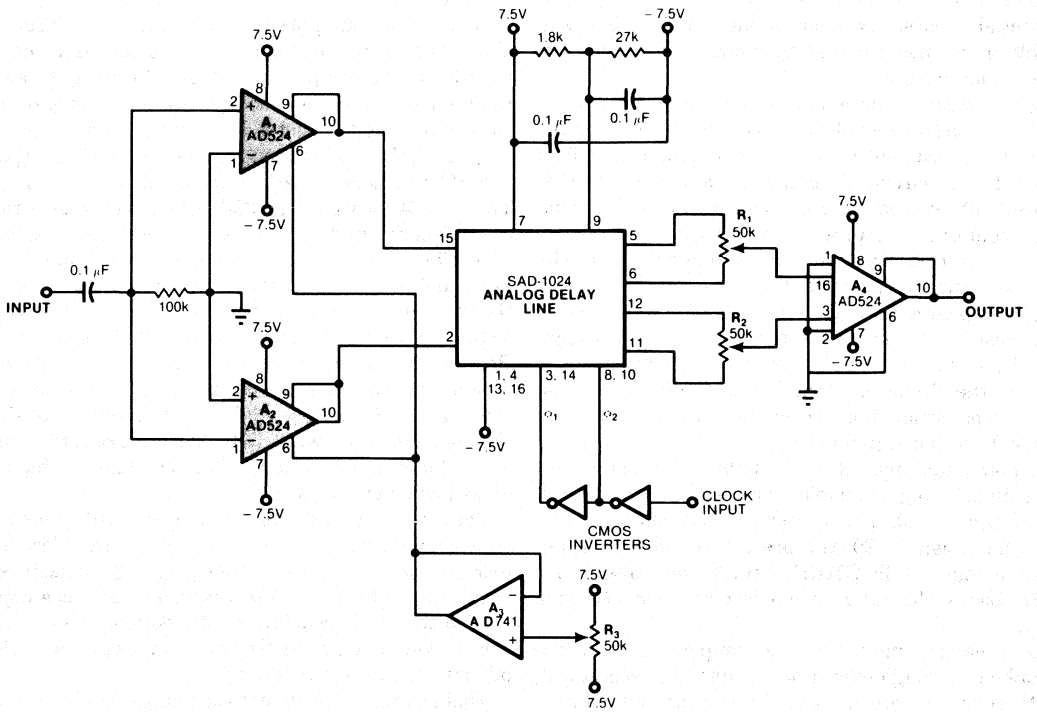

Fig 3—An analog delay circuit *exploits both halves of a CCD delay line. Because instrumentation amps maintain differential circuit operation, noise and clock feedthrough remain minimal.*

signal frequency. You can derive the clock from an external VCO and frequency divider.

To ensure minimal clock noise, each of the delay lines' outputs consists of two signals. Potentiometers R_1 and R_2 let you balance these signals' currents and thus maximize clock-noise rejection. The differential, current-summing output configuration further nulls clock feedthrough, resulting in an overall signal-to-noise ratio of 70 dB—a level unattainable using single-ended designs.

As shown, the circuit's input amplifiers exhibit unity gain, but you can configure them for gains to 1000. A_3 and its associated circuitry let you bias the input stages to a level that ensures minimum delay-line harmonic distortion. To trim the circuit, adjust R_3 for the bias level that results in minimum distortion with a 1V rms input signal. When properly adjusted and fed by a 150-kHz clock, the circuit typically operates with 0.2% total harmonic distortion.

Measuring distortion

You can further exploit instrumentation amps' differential properties to enhance the resolution of distortion-measuring equipment. Typical distortion meters null signal fundamentals by 80 to 100 dB and thus exhibit measurement limits of 0.01 to 0.001%. But by using an instrumentation amp to prenull a device-under-test's (DUT's) output (**Fig 4**), you can extend a distortion meter's measurement capabilities to 140 dB below fundamental amplitudes. Thus, you can measure distortion as low as 0.00001%.

The prenulling method entails several advantages in addition to greater measurement resolution. First, because an instrumentation amp subtracts signal-source inputs from DUT outputs, test-oscillator errors are cancelled during the measurement process. As a result, you can measure lower distortion levels than most oscillators allow.

In addition, the prenulling method relaxes the dynamic-range requirements normally imposed on distortion meters used for high-resolution tests. The distortion meter processes only signal errors—not fundamentals—so the analyzer need only possess sufficient dynamic range to monitor errors in the instrumentation amp's residual output. And because the test signal is cancelled through subtraction rather than through notch filtering, you can use complex waveforms for intermodulation testing and other harmonic analyses.

In **Fig 4**'s circuit, amplifiers under test drive a load resistor, R_L. A linear network with a gain of 1/A, where A equals the amplifier's gain, feeds the voltage across R_L to one input of the nulling amp. The other instrumentation-amp input serves as a reference port

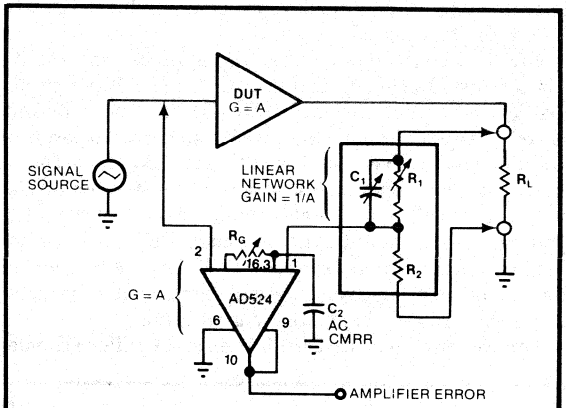

Fig 4—Enhance distortion-measurement resolution *with this prenulling circuit. A linear network between a device under test's output and an instrumentation amp's input compensates for DUT gain and frequency characteristics. The prenulling circuit relaxes the dynamic-range requirements of distortion analyzers, because those instruments must process only the error signals that emerge from the instrumentation amp's output.*

and connects to the DUT's input signal. With the instrumentation amp's gain set to A (via R_G), the nulling circuit's error-output level remains referred to the DUT's output-signal amplitude.

Adjusting the attenuator network proves the most difficult aspect of the nulling circuit's setup procedure. The first—and easiest—step involves adjusting R_1 for dc attenuation equal to $1 \div A$. Then you must trim compensation capacitor C_1 in the same manner that you compensate oscilloscope probes. Time constant R_1C_1 must correspond to the −3-dB bandwidth of the DUT; otherwise, the instrumentation amp won't fully null high-frequency signal components.

When constructing the attenuator, make certain that R_1 and R_2 are nonreactive and possess thermal- and voltage-nonlinearity coefficients small enough to prevent masking of DUT errors. In addition, you must perform several R_1/C_1 adjustment iterations to ensure optimum nulling.

If you intend to use a meter to display the instrumentation amp's output, you must ensure precise phase tracking between the DUT and the attenuator network; 1° of phase error limits rejection to only 35 dB. If you're feeding a spectrum analyzer, however, you can relax the phase-tracking requirement, so long as residual fundamentals don't overload the analyzer's front end. **Fig 5a** illustrates a spectrum-analyzer display that shows the results obtained from proper nulling. The fundamental is dispayed at a level 106 dB below a 3V rms test signal; the second harmonic,

Split equalization network improves phono preamp

therefore, is clearly displayed, even at a level 120 dB below the test signal.

Fig 5b shows the results of a prenull-based intermodulation test. In this measurement, 5- and 6-kHz signals are applied simultaneously to the DUT. The analyzer clearly displays the 1- and 11-kHz sum and difference frequencies as well as the 10- and 12-kHz second-harmonic components. Because the instrumentation amp cancels all but a small portion of the test signals, you can use a limited-dynamic-range FFT analyzer to recover output-waveform spectral data, even when harmonic and intermodulation components remain more than 100 dB below test-signal levels.

You can employ a similar difference-amplifier scheme

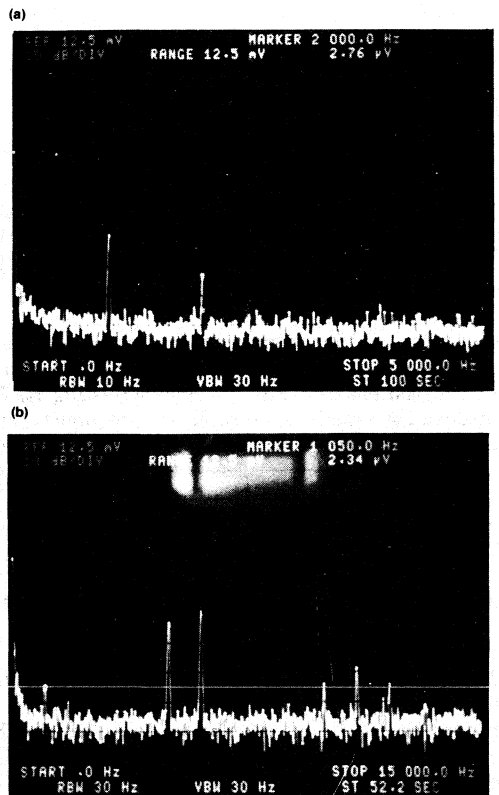

Fig 5—Harmonic spectra *illustrate the operation of* ***Fig 4****'s prenulling circuit. Trace* ***(a)*** *shows a harmonic-distortion test; the fundamental is prenulled to −106 dB, permitting clear indication of the second harmonic 120 dB below the DUT's output level. Trace* ***(b)*** *illustrates an intermodulation test in which a DUT processes 5- and 6-kHz test signals. The display shows 1- and 11-kHz intermodulation components as well as 10- and 12-kHz second harmonics.*

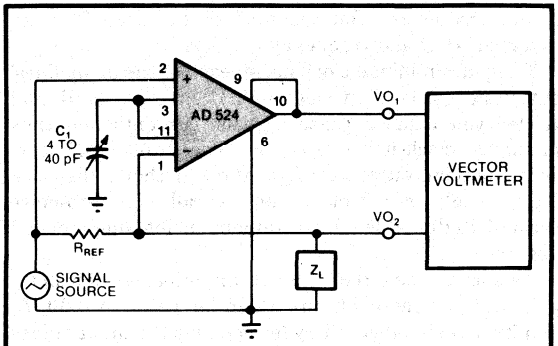

Fig 6—A current monitor *measures impedance. By amplifying the voltage drop across a reference resistor, this circuit lets a vector voltmeter detect* Z_L*'s magnitude and phase.*

to measure complex impedances. Some instrumentation amps feature very high high-frequency CMRR, trimmable via a single capacitor, so the circuits suit applications involving current sensing on lines with large time-varying components.

In **Fig 6**'s current monitor, for example, a low-impedance sine-wave source feeds a signal to a reference resistor (R_{REF}) and one input of an instrumentation amp. R_{REF} is a precision, nonreactive resistor exhibiting low impedance compared with the network under test (Z_L). The instrumentation amp's output equals the voltage across R_{REF} multiplied by the stage's gain and thus remains proportional to the current through Z_L. You can trim the circuit for minimum error (maximum CMRR) by shorting R_{REF} and adjusting C_1 for minimum circuit output.

You can use **Fig 6**'s circuit to measure complex impedances by connecting a vector voltmeter to the instrumentation amp's output and to Z_L. The vector monitor measures the magnitude and angle of $V_L \div I_L$ and computes $R_L=(V_L/I_L)\cos\theta$ and $X_L=(V_L/I_L)\sin\theta$.

Instrumentation amp improves phono preamp

High CMRR also brings advantages to phono-cartridge preamplifiers. Typically, preamps for moving-magnet phono pickups (the cartridges most commonly used in stereo systems) employ op amps with RIAA equalization networks (N) in their feedback loops (**Fig 7a**). These circuits are quite simple—they employ only one active device—but they impose stringent performance requirements on the op amps they employ. To process low-level signals between 20 Hz and 20 kHz while providing 1-kHz gains of as much as 40 dB, op amps used in conventional phono preamps must exhibit low noise and high speed at unity gain. And to drive complex feedback networks, they must furnish low

Prenulling circuit ups measurement resolution

output impedance at high frequencies. Most important, single-ended circuits provide poor low-frequency hum rejection.

An alternative phono-preamp approach (**Fig 7b**), however, exploits instrumentation amps' differential properties and low-input noise to provide excellent signal-to-noise characteristics and equalization-independent gain control. In this circuit, an AD624 instrumentation amp (A_1) converts a moving-magnet cartridge's balanced output to a single-ended amplified voltage suiting equalization and further processing.

Two termination resistors (R_{TA} and R_{TB}) provide dc input loading (47 kΩ typ), and a pair of capacitors (C_{TA} and C_{TB}) furnish cartridge damping (150 pF suffices for most cartridges). The differential network maintains signal balance for maximum hum rejection while ensuring proper cartridge termination. A 4- to 40-pF trimmer or 15-pF fixed capacitor connected to the instrumentation amp's gain terminals increases high-frequency common-mode rejection. You can further improve the circuit's high-frequency common-mode performance by adjusting the ratio of C_{TA} to C_{TB}.

Because most of A_1's input current flows through R_{TA} and R_{TB}, only the instrumentation amp's offset current flows through the cartridge. And because the AD624's offset current never exceeds 35 nA and typically equals only 10 nA, input-current-induced offset errors remain negligible compared with the instrumentation amp's 250-μV input-offset level.

To allow maximum component-choice freedom, **Fig 7b**'s design accomplishes RIAA equalization in two stages. A passive network (R_1 and C_1) between A_1 and output buffer A_2 furnishes 75-μsec de-emphasis; the responses' 318- and 3180-μsec components are generated in A_2's feedback loop. This split-network technique results in complete buffering of all equalization components. An alternative equalization technique involves lumping all three RIAA time constants in a single

The instrumentation amp as op amp

As monolithic instrumentation amps' performance improves, an interesting role reversal occurs: Instrumentation amps are composed of op amps, but they can also serve as op amps. Consider, for example, the nearby **figure**'s precision instrumentation-amp-based op amp. In this circuit, sense and reference resistors connect to the output terminal, thus configuring the instrumentation amp's output amplifier as a Howland current pump (EDN, January 20, pg 85). Because the current pump exhibits very high conductance, the output amplifier operates with considerable gain. Indeed, this gain is bounded only by CMRR restrictions and typically equals 80 dB.

However, the instrumentation-amp-based op amp's total gain exceeds 80 dB when the input amplifiers are connected to the output buffer. With input gain strapped for 100, for example, an op-amp-configured AD524 provides open-loop gain of 5×10^6 (134 dB) typ.

In the **figure**'s circuit, closed-loop gain equals 1000 and −3-dB bandwidth equals 56 kHz. Thus, the amplifier exhibits a gain-bandwidth product of 56 MHz. Such wide bandwidth, coupled with the instrumentation amp's low noise and high stability, results in performance typically unattainable in conventional IC op amps. As an added advantage over conventional op amps, the circuit's damping characteristics depend only on the setting of one resistor (R_G).

Such instrumentation-amp-based op amps prove useful in applications requiring closed-loop gains to 100,000 over wide bandwidths. An AD524, for example, provides gain-bandwidth products as high as 1 GHz when its input amplifiers are set for gains of 1000. Note, however, that instrumentation-amp-based op amps remain most effective at high gains; excess front-end phase results in instability at very low closed-loop gains.

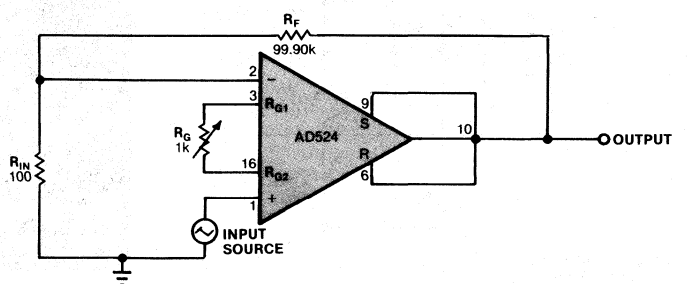

An instrumentation amp acts like an op amp *when sense and reference resistors connect to circuit outputs. The dual-feedback arrangement configures the instrumentation amp's output amplifier as a Howland current pump and thus results in high open-loop gain*

A unity-gain buffer overcomes drive limitations

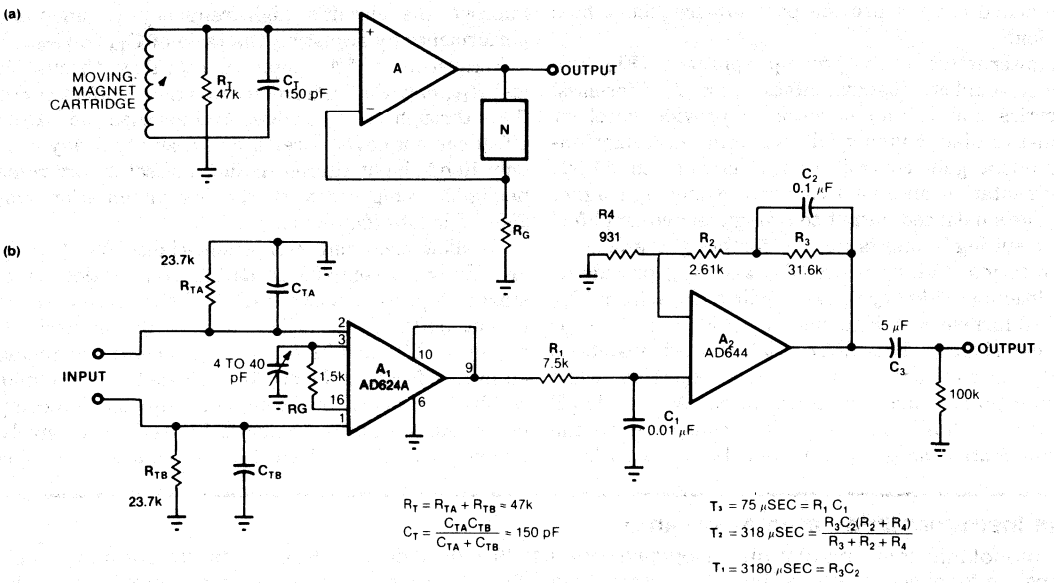

Fig 7—Conventional phono preamps (a) *employ a single RIAA-equalization network (N) in an op amp's feedback loop. Although such circuits are extremely simple, they impose undue demands on op-amp performance. A better circuit* ***(b)*** *uses an instrumentation amp to buffer cartridge outputs. The instrumentation amp's high CMRR results in hum rejection to 100 dB.*

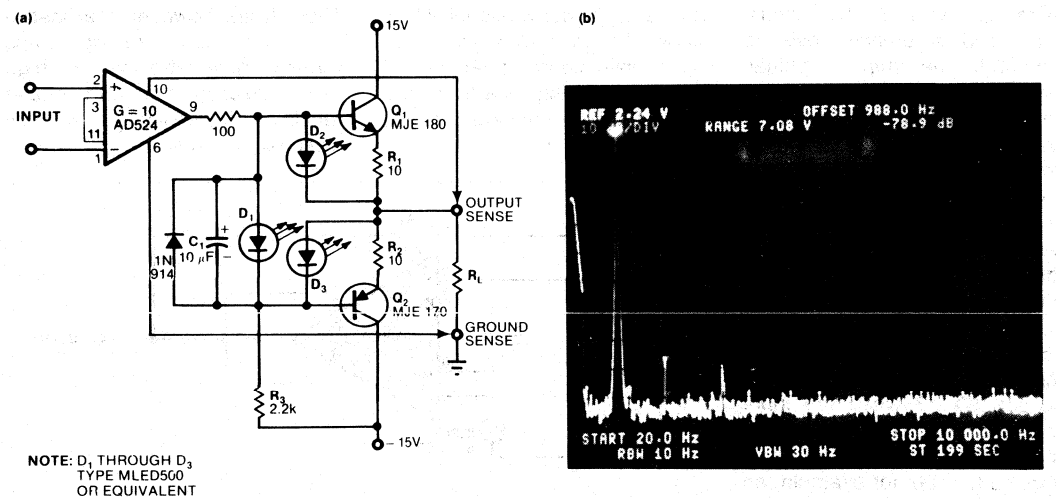

Fig 8—An output buffer *in an instrumentation amp's feedback loop provides high-current drive capability* ***(a)****. Photo* ***(b)*** *illustrates the buffer's low distortion; the spectrum shows the harmonic content of a 7V rms buffer output driving a 100Ω load.*

network between A_1 and A_2 (**Ref 2**), but such a scheme results in greater output noise than does the split-equalization method.

In **Fig 7b**'s circuit, lowest noise results when A_1 operates with high gain; greatest overload immunity results when A_1 exhibits relatively little gain. Therefore, you must select the input-stage gain that leads to the best compromise between two conflicting performance attributes. Typically, the best compromise results when A_1's gain equals approximately 20; A_2 then must provide the remainder of the network's gain. In **Fig 7b**'s design, A_1 operates with a gain of 28 and A_2 exhibits a dc gain of 38. Total dc gain thus equals 1064 or approximately 60 dB.

In general, the phono preamp's dc gain equals

$$G_{DC}=(1+40k/R_G)(1+(R_2+R_3)/R_4),$$

and RIAA-equalized 1-kHz gain equals

$$G_{1k}=0.101G_{DC}.$$

The constant 0.101 represents the ratio between dc and 1-kHz gains as set forth by the RIAA equalization standard. In **Fig 7b**'s circuit, 1-kHz gain equals approximately 40 dB. You can vary the circuit's gain without affecting its equalization characteristics by adjusting R_G's value. Similarly, you can alter the circuit's response—to adapt it for tape-head applications, for example—by selecting appropriate values for R_1 through R_3, C_1 and C_2.

As shown, with 40-dB midband gain, the phono preamp furnishes an 80-dB signal-to-noise ratio referred to a 1V output. Low-frequency CMRR equals 100 dB min. Maximum response accuracy and stability result when all equalization-network resistors are metal-film types and all capacitors employ polystyrene or polypropylene dielectrics.

Frequently, instrumentation amps must provide output currents exceeding 5 mA while maintaining full specified accuracy. Most IC instrumentation amps are specified to drive 5 mA into 2-kΩ loads, but their 100Ω typ open-loop output impedances restrict higher current operation. By placing a unity-gain buffer inside an instrumentation amp's feedback loop, however, you can overcome circuit drive limitations.

Fig 8a illustrates a buffered instrumentation-amp circuit that drives load impedances as low as 150Ω and delivers as much as 75 mA of output current. The circuit employs a complementary pair of 3A/40V power transistors (Q_1 and Q_2) that get forward-biased by LED D_1. Because of D_1's bias voltage, R_1 and R_2 drop approximately 140 mV each, resulting in an output quiescent current of approximately 15 mA with the component values shown. The 15-mA bias level proves sufficient to prevent output-stage crossover distortion. Two additional LEDs (D_2 and D_3) serve as voltage clamps, limiting the current through Q_1 and Q_2 to approximately 100 mA. These LEDs also act as overload indicators.

When constructing **Fig 8a**'s circuit, observe several precautions. First, for maximum linearity, select R_3 so that it forces A_1's output stage into Class A operation. You must also sense output voltages from the output terminal, not from internal circuit nodes. In addition, connect the load's common terminal to ground via its own power-supply connection. Otherwise, load currents might perturb the circuit's analog ground. Finally, to prevent high-frequency oscillations, mount bypass capacitors as close as possible to the output transistors' collectors.

References

1. Wurcer, S A, and Kitchen, C, "Stacked amplifiers lower noise," *EDN*, October 13, 1982, pg 184.
2. Jung, W, "Topology considerations for RIAA phono preamplifiers," 67th Audio Engineering Society Convention, November 1982, Reprint No 1719.
3. Wurcer, S A, and Counts, L W, "A programmable instrumentation amplifier for 12-bit-resolution systems," *IEEE Journal of Solid State Circuits*, Vol SC-17, No 6, December 1982, pgs 1102-1111.
4. Jung, W, *Audio IC Op Amp Applications (second edition)*, Howard W Sams and Co Inc, Indianapolis, IN, 1978.

Isolation Amplifiers
Contents

		Page
AN-246	Isolation Barrier Precisely Passes Transducer Signals	11–3

ANALOG DEVICES

AN-246
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Isolation Barrier Precisely Passes Transducer Signals

by Jeffrey A. Greenwald

Through the judicious use of isolation amplifiers that incorporate transformer coupling, you can design circuits that condition and amplify low-level transducer signals. Such circuits can precisely measure thermocouple, resistance-temperature-detector and strain-gauge outputs while providing common-mode isolation in harsh environments.

Isolation amplifiers protect data-acquisition components from potentially destructive voltages present at remote transducers. These amplifiers are also useful when you need to amplify low-level signals in multichannel applications. They can also eliminate measurement errors caused by ground loops. Amplifiers with internal transformers reduce circuit costs by eliminating the need for an additional isolated power supply. And in many cases it's possible to connect a transducer to an isolation amplifier with only a few external components.

The transducers that you'll most often encounter include thermocouples, resistance temperature detectors (RTDs), pressure sensors and strain gauges. These sensors' outputs usually require high amplification. The sensors may also require some sort of external excitation, and you may have to perform filtering or other signal conditioning to yield a useful signal. In some cases, you'll have to offset the sensor's output to place it into the proper signal range for the data-acquisition hardware. You can perform all of these functions with rather simple circuits and precise isolation amplifiers.

The isolation amplifiers you use in these circuits must have high enough breakdown potential to withstand ground faults and thus prevent system destruction. In many applications, you may be called upon to mount pressure transducers and strain gauges on metallic pipes and objects that are at different electric potentials than those of your data-acquisition system. High isolation comes in handy to break ground loops, to reduce noise, or to otherwise allow you to measure strain safely in the presence of several hundred volts.

The isolation amplifier needs to be able to respond to out of range transients that may overload pressure transducers. The AD208, for example, recovers from 5V input transients at a gain of 1000 in 5 msec without damage. Overload often happens when a valve closes, restricting the flow of fluid in a pipe. Since pressure is proportional to force and inversely proportional to cross-sectional area, the pressure rapidly rises to several times the full-scale pressure. This effect is called the water-hammer effect. The same condition is occurring when your pipes make a banging noise after you have shut off the water in the sink.

Possibly the most common transducer that you'll use isolation amplifiers with is a thermocouple. Thermocouples are low-cost sensors, and can measure temperatures over a wide range. **Table 1** lists some common types and their ranges. Thermocouples comprise two

Reprinted from EDN — June 21, 1990

Isolation amplifiers protect data-acquisition components from potentially destructive voltages present at remote transducers.

dissimilar metals. A phenomenon called the Seebeck effect generates a temperature-dependent microvolt signal when the two dissimilar metals join to form a closed-loop circuit.

Although thermocouples are inexpensive, they have one major disadvantage. In addition to the desirable thermocouple effect that you want to measure, undesirable thermocouple effects occur when you attach the thermocouple to the outside world, typically in the form of copper wire. If the thermocouple material is not also copper—it typically is not—additional thermocouple effects known as cold-junction errors occur.

Table 1—Common thermocouples and their characteristics

Thermocouple type	Maximum temperature range at 5V output (°C)	Maximum V_{IN} (mV)	Isolation amplifier's gain setting (V/V)	Suggested resistor values with reference to Fig 2		
				R_F (kΩ)	R_{TRIM} (kΩ)	R_Z (MΩ)
E	900	68.783	72.69	6.98	0.5	2
J	750	42.283	118.25	11.5	1	2
K	1250	50.633	98.75	9.53	1	2
R	1450	16.741	298.6	28.7	2	2
S	1450	14.973	333.9	32.4	2	2
T	350	17.816	280.6	27.4	2	2

For example, the two measurement terminals of a J-type thermocouple consist of iron (Fe) and constantan (Con) contacts. (Constantan is a nickel-copper alloy.) When you attach each side of the thermocouple to copper wire, two different metals join—iron and copper on one end and constantan and copper on the other—to create thermally induced voltages whose values are a function of ambient temperature. The induced voltages at the constantan-copper and iron-copper cold junctions are opposite in polarity to the Seebeck voltage created by the thermocouple.

You can model the two cold junctions as two voltage sources in series with the thermocouple (**Fig 1**). While the thermocouple is sensing temperature at a specific physical location, these cold junctions are functions of the ambient environment. The voltage at the measurement terminal, $V_M(T)$, is the difference between the voltage at the thermocouple and the sum of the induced voltages, as follows:

$$V_M(T) = V_{Fe\text{-}Con(sensor)} - V_{Cu\text{-}Con(ambient)} - V_{Fe\text{-}Cu(ambient)}.$$

Because the two cold junctions have a back-to-back, copper-to-copper connection, you can combine them in a single iron-constantan voltage source. You can then rewrite the voltage at the measurement terminal:

$$V_M(T) = V_{Fe\text{-}Con(sensor)} - V_{Fe\text{-}Con(ambient)}.$$

Canceling this cold-junction effect requires an opposing voltage in series with the thermocouple. This offsetting voltage should generate zero volts at zero degrees C and have a positive temperature slope equal to the thermocouple type you're using. In effect, you need a temperature sensor that measures the cold-junction ambient temperature and produces an offsetting voltage.

One method to measure and compensate for the cold-junction is to use a monolithic cold-junction compensator in series with the thermocouple (**Fig 2**). IC_1 generates a temperature-dependent output of 10 mV/°C. An internal laser-trimmed resistor network scales this signal to provide the correct slopes for most commonly used thermocouple types. You can access the scaled compensation voltage on one of four taps for various standard thermocouples: 60.9 μV/°C (E-type), 51.7 μV/°C (J-type), 40.6 μV/°C (K-type), and 6 μV/°C (R- and S-type).

The isolated ±8V dc supply at pins 5 and 6 of the isolation amp can supply as much as ±5 mA to the cold-junction compensator. However, because IC_1 only draws a maximum of 150 μA, there is sufficient power left over if you need to add more circuitry, such as an

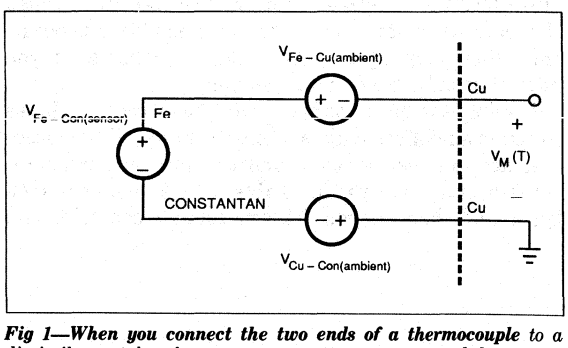

***Fig 1—When you connect the two ends of a thermocouple** to a dissimilar metal such as copper, you generate unwanted thermocouple effects called cold junctions. Voltage sources in series with the thermocouple's voltage model the effects of these junctions.*

input filter. In turn, the transformer-coupled isolation amplifier obtains its power from the power-clock driver, IC_3. The clock driver produces a 25-kHz, $15V_{p-p}$ square wave, which can drive up to 16 fully loaded isolation amplifiers.

You obtain the maximum sensitivity from the thermocouple circuit by amplifying the low-level input signal to the isolation amp's full-scale output range. The following equation yields the necessary gain from a few known values: the output voltage span of the thermocouple that corresponds to the thermocouple's temperature extremes ($V_{TC\text{-}HI} - V_{TC\text{-}LO}$); and the input-voltage requirements of your data-acquisition electronics, $V_{HI} - V_{LO}$.

$$\text{Gain} = (R_F/R_I + 1)$$

$$= (V_{HI} - V_{LO})/(V_{TC\text{-}HI} - V_{TC\text{-}LO}).$$

From the values in the **table**, you can see that a J-type thermocouple has a maximum temperature of 750°C, corresponding to a maximum output of 42.283 mV. Thus, the gain required to amplify the 0- to 42.283-mV signal to a 0 to 5V output span is (5 − 0V)/(42.283 mV − 0 mV), or 118.25.

You can determine the optimal feedback resistor value, R_F, by first setting R_I to 100Ω, and then solving for R_F using the equation $R_F = (\text{Gain} - 1)R_I$, which yields $R_F = 117.25 \times 100$. Note that in **Fig 2**, R_F is the sum of R_F' and R_{TRIM}. To obtain a bipolar adjustment range of ±5% for the gain, choose $R_F' = 0.95R_F$, and $R_{TRIM} = 0.1R_F$.

Similarly, you can use R_Z to set the offset adjustment range. The following equation relates this resistance to the gain, feedback-resistor value, isolation voltage, V_{ISO}, and two offset voltage values. One of these offset values, $V_{OS\text{-}RTI}$, is the referred-to-input offset; the other, $V_{OS\text{-}RTO}$, is the referred-to-output offset.

$$R_Z = V_{ISO}R_F/(V_{OS\text{-}RTO} + (\text{Gain} \times V_{OS\text{-}RTI}).$$

The AD208 isolation amplifier, for example, typically has a V_{ISO} of 8V, a $V_{OS\text{-}RTO}$ of 15 mV max, and a $V_{OS\text{-}RTI}$ of 250 μV max. Using these values and those for the gain and R_F, you can solve for $R_Z \approx 2$ MΩ.

Perform calibration first

For the circuit to achieve maximum accuracy, you'll need to calibrate the circuit in **Fig 2** to minimize offset and gain errors. You can adjust the offset by first shorting pin 1 of IC_2 to pin 2, the Input Common. (Or you can insert the thermocouple in an ice-point refer-

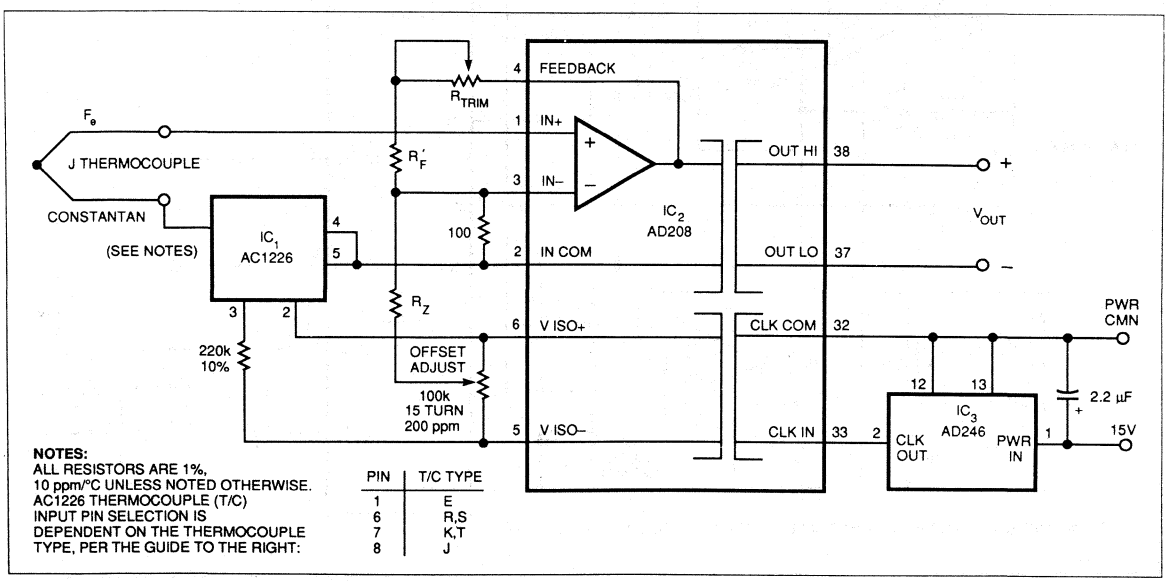

Fig 2—To compensate for cold-junction effects, *this circuit inserts temperature-dependent compensation, generated by IC_1, between the thermocouple and an isolation amplifier.*

A thermocouple relies on the Seebeck effect to generate a temperature-dependent microvolt signal when two dissimilar metals join to form a closed-loop circuit.

ence—a 0°C environment.) Then, adjust the offset potentiometer until the amplifier's output measures 0V. You can calibrate the gain by applying the full-scale thermocouple voltage (42.283 mV for a J-type thermocouple) to the input of IC_2. Then adjust the gain-trim potentiometer until the dc signal at the output of the isolator is 5V. Because the offset and gain may interact with one another, you may have to recalibrate for zero offset following the gain adjustment.

After removing the calibration source, the cold-junction compensator should cause no more than ±2°C of residual error. The temperature stability of the circuit in **Fig 2** (referred to output) is 212.2 μV/°C, yielding a sensitivity to ambient temperature of 0.032 μC/°C. Gain drift is 321 μV/°C, delivering a sensitivity to ambient temperature of 0.048 μC/°C.

Measure temperature more precisely

RTDs offer better stability and linearity than thermocouples, but are limited to temperatures below 850°C. The most common RTDs are made of platinum, because platinum is highly stable at elevated temperatures. Platinum RTDs typically have a 100Ω resistance and a temperature coefficient of +0.385%/°C at 0°C (European DIN Standard 43 760).

Because the RTD's resistance is the temperature-dependent parameter, the signal-conditioning circuit must supply a current to excite the sensor and then measure the voltage drop across the RTD. The circuit must apply relatively low excitation currents to avoid self-heating errors. As a result of the low excitation current and the low sensitivity of the RTDs, you'll need to employ a low-drift, high-gain amplifier to handle the RTD's millivolt output signal. Finally, to obtain the high-measurement accuracy that is possible with RTDs, you need to eliminate any initial offset voltage at the reference temperature, usually 0°C. You'll also need to reduce errors created by lead resistance in the wire that connects the RTD to the measurement electronics.

Fig 3's circuit is one way to achieve these goals. This measurement circuit combines excitation-current sources, lead-wire compensation and amplification for a 100Ω RTD. The LT1009, a 2.5V reference, amplifier

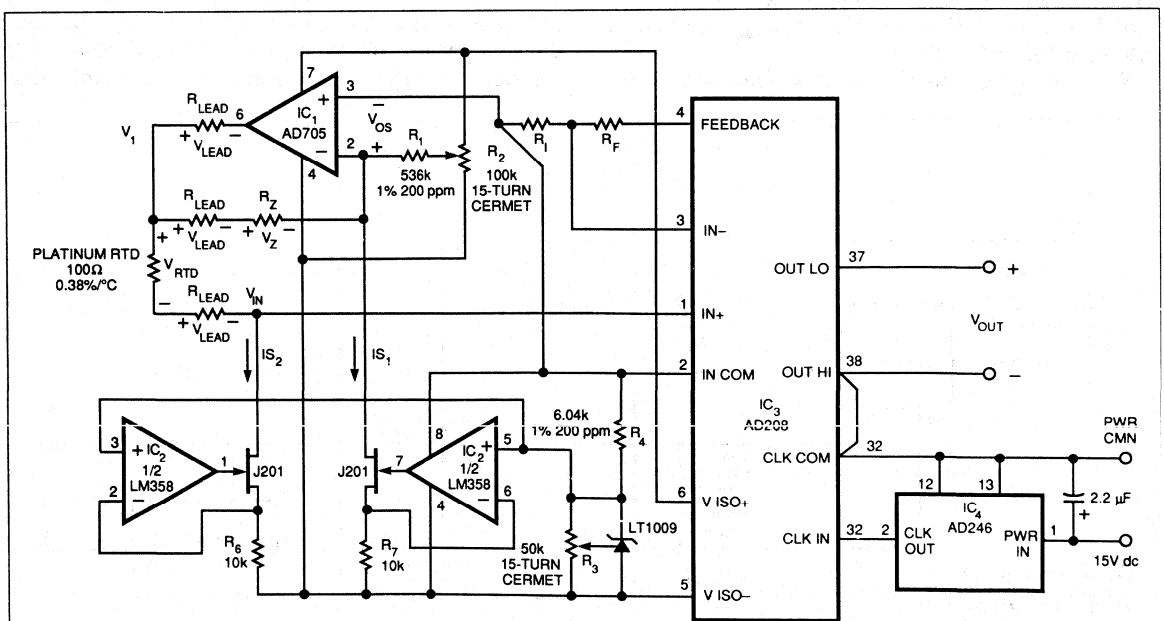

NOTE: ALL RESISTORS ARE 1%, 10 ppm/°C UNLESS NOTED OTHERWISE.

Fig 3—Dual currents and a ground-servo amplifier, IC_1, allow this circuit to excite and provide zero suppression to a 100Ω, platinum RTD over a 0 to 500°C range.

IC_1, dual LM358 amplifiers, IC_2, and a pair of low Vgs_{off} n-channel JFETs such as the J201 compose two low-power current sources. One current source, IS_2, excites the RTD. The other source, IS_1, supplies current to offset the RTD such that the voltage at the isolation amplifier's input is zero at 0°C. This zero-suppression technique saves you from having to reduce the gain sensitivity of the RTD circuit to keep the signal within the operating range of the isolation amplifier.

IS_2 excites the RTD with a low current, 250 μA, to minimize self-heating. By adjusting the trimming potentiometer, R_3, you can change the excitation current's range by as much as ±5%. The circuit implements zero suppression using a simple ground-servo amplifier comprising IC_1 and R_Z. (R_Z equals the RTD's 100Ω value at 0°C). The ground-servo amplifier provides an effective means of providing zero suppression without having to use an expensive differential amplifier. It also compensates for the finite lead resistance by remotely sensing the voltage at the RTD.

The current, IS_1, develops a voltage, V_1, equal to

$$V_1 = V_{OS} + V_Z + V_{Lead},$$

where V_{OS} is IC_1's offset voltage.

IS_2 in turn generates the voltage, V_{IN}, seen by IC_3's input amplifier. The value of V_{IN} equals

$$V_{IN} = V_1\text{-}V_{RTD} - V_{Lead}.$$

Substituting the previous equation for V_1 in this equation for V_{IN} yields

$$V_{IN} = V_{OS} + V_Z - V_{RTD}.$$

The trimming potentiometer, R_2, connected between $\pm V_{ISO}$ and the large injection resistor, R_1, nulls any residual offset. To minimize errors created by offset drift in IC_1, choose a low-power, low-offset, and low-drift amplifier for this function.

The sensitivity of a 100Ω platinum RTD with a 250-μA current excitation is approximately 95 μV/°C. Therefore, setting the isolation amplifier's gain to 105 will result in an output sensitivity of 10 mV/°C for a temperature range of 0 to 500°C. With this amplification, the amplifier's output will span a 0 to −5V output range. By connecting IC_3's OUT HI pin to ground (CLK COM) and using OUT LO as the output, the circuit inverts the output range to 0 to 5V.

The gain equation of the circuit in **Fig 3** is the same as the equation for **Fig 2** except for the additional excitation current term, IS_2, as follows:

$$\text{Gain} = (R_F/R_I + 1)$$

$$= (V_{HI} - V_{LO})/IS_2(V_{TC\text{-}HI} - V_{TC\text{-}LO}).$$

If you let R_I equal 100Ω, R_F should be 10.5 kΩ for a 0 to 500°C temperature range.

Fig 3's RTD circuit requires offset and gain calibration similar to that of the thermocouple-conditioning circuit. The easiest calibration method is to substitute a known resistance in place of the RTD. By replacing the RTD with a 100Ω resistor you can simulate a temperature input of 0°C. Then you can adjust the offset-trimming potentiometer, R_2, so that IC_3's output is 0V. Then exchange the 100Ω resistor with one that simulates the RTD at 500°C (280.9Ω, DIN 43 760 standard), and trim the reference voltage until the IC_3's output equals 5V.

A 100Ω platinum RTD that has a 250-μA excitation current and is amplified by a gain of 105 V/V will display a referred-to-output (RTO) sensitivity of 10.1 mV/°C. The circuit's RTO offset and gain drift is 433.2 μV/°C and 360 μV/°C, respectively. This offset and gain drift yields a total sensitivity to ambient temperature changes of 0.079°C/°C.

Condition strain-gauge signals

Another popular class of transducers used in fluid-flow monitoring (and in stress and strain measurements) are pressure transducers and strain gauges. Most of these transducer types consist of resistive elements whose resistance fluctuates in response to mechanical forces. These resistors, like RTDs, need an excitation signal and zero suppression. A common method of implementing these transducers is by placing them into one or more legs of a Wheatstone bridge. A differential amplifier then amplifies the small signal imposed on a large common-mode dc voltage.

Fig 4 shows a simple and accurate method of exciting and conditioning a strain gauge using the bridge approach—and without having to use a costly differential amplifier. This circuit connects a 100-mV strain gauge to the ±5V excitation signal derived from a precision voltage reference, IC_1. This reference contains a 10V buried-zener reference, half-bridge completion resistors (R_1 and R_2), and four amplifiers. Amplifier A_1

Thermocouples have cold-junction errors and voltage nonlinearity errors that complicate accurate temperature measurements.

performs as a ground-servo amplifier for zero suppression, and output amplifiers A_2 and A_3 add remote-sense capability. **Fig 4** configures IC_1 for ±5V operation with emitter-follower current boosters Q_1 and Q_2 to power the strain gauges. The current boosters can supply as much as 10 mA of load current into a 500Ω strain gauge. This current is limited by power limitations of the booster transistors, the power clock, and the saturation flux of the power transformer.

The ground-servo amplifier, A_1, and half-bridge completion resistors, R_1 and R_2, turn an otherwise differential measurement into a single-ended one. A_1 senses the reference voltage via pins 11 and 10 and forces its midpoint to zero. Thus, it forces the bridge to operate symmetrically around ground. If you need to make a remote measurement in cases where the lead resistance in the bridge can cause significant measurement errors, this circuit runs the sense leads of A_2 and A_3 to the remote strain gauge. A_2 and A_3 drive the bases of Q_1 and Q_2 such that the remote sense point will be V_{REF} ±5V.

Fig 4 amplifies a 100-mV, ground-referenced transducer signal by 50 to produce a ±5V output range. As in the previous circuits, you determine the isolation amp's gain by using the standard noninverting gain formula: $R_F/(R_I+1)$.

If you're using a 500Ω strain gauge that requires more than 10 mA, you'll need to add more power to the circuit. You can derive useful power from the power clock by using a simple half-wave rectifier, filter, and transformer comprising T_1, D_1, D_2, and 4.7-µF capacitors. The power clock's $15V_{p-p}$, 25-kHz square-wave output provides the necessary current to supply 20 mA dc from the ±8V output of the unregulated, isolated power supply. Because the AD588 voltage reference requires a maximum of 10 mA, 10 mA are available to the strain gauge.

Consider an alternative

Keep in mind that these voltage-output circuits require an external, unity-gain buffer amplifier, if you plan to drive a load that is less than 1 mΩ. As an option to all of the previous signal-conditioning circuits, you can convert the output voltage of the AD208 isolation

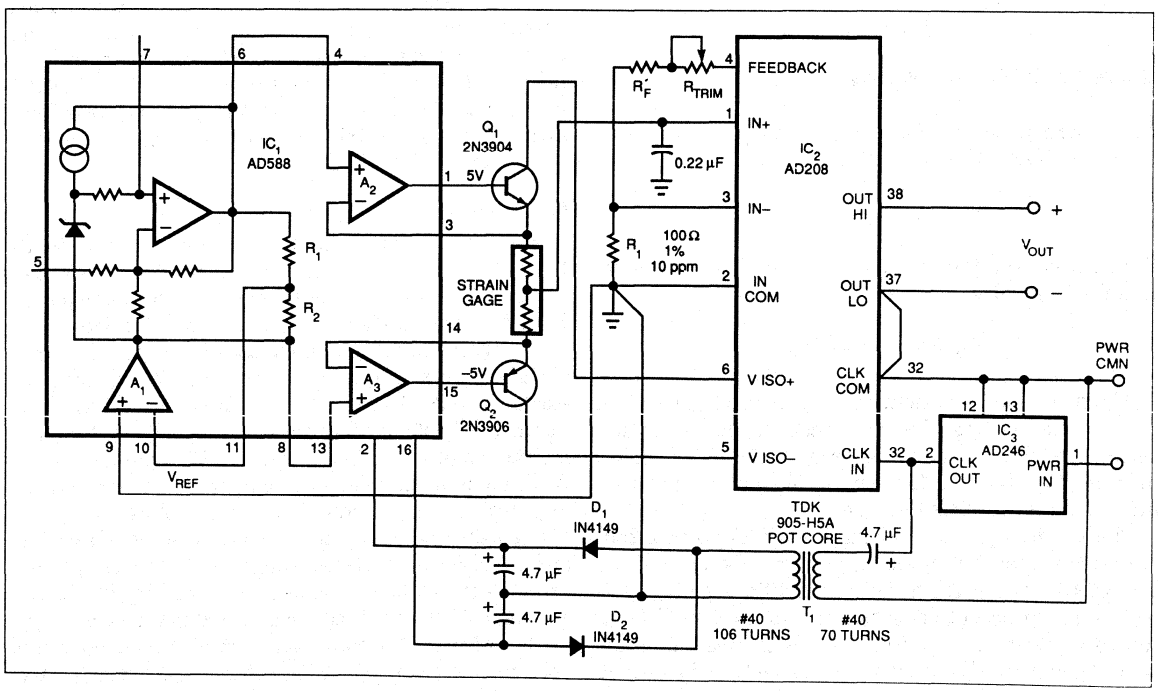

Fig 4—Amplifier A_1, part of a programmable voltage reference, *senses the reference midpoint and forces the bridge to operate around that value. T_1 and its associated components add power if needed.*

Resistance temperature detectors (RTDs) offer better stability and linearity than thermocouples, but have a narrower operating-temperature range.

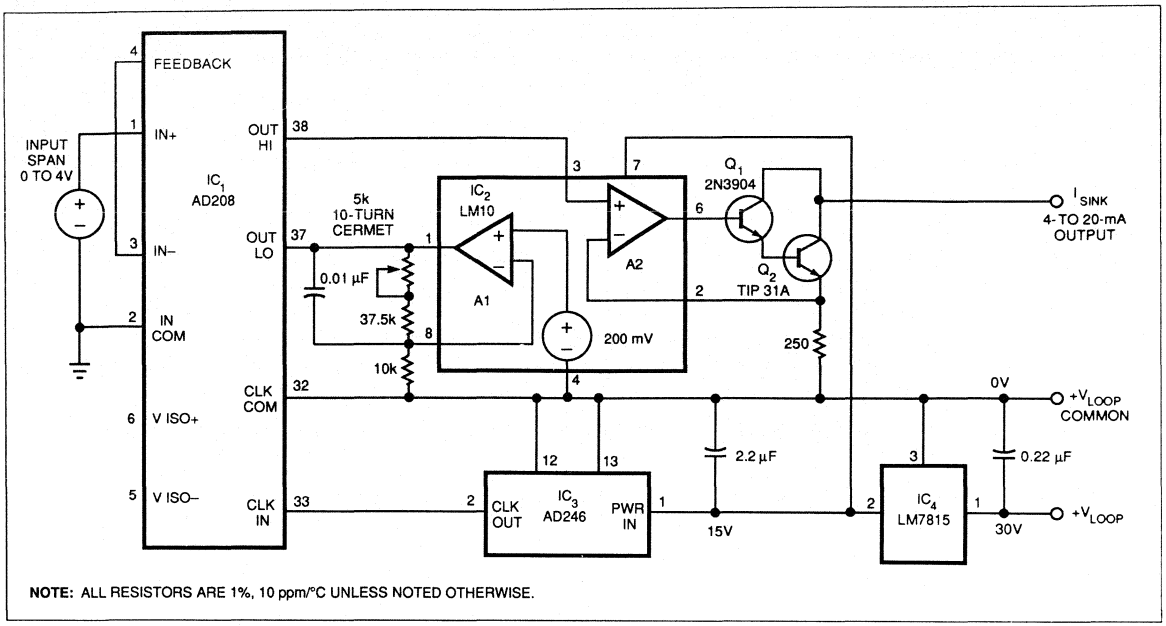

NOTE: ALL RESISTORS ARE 1%, 10 ppm/°C UNLESS NOTED OTHERWISE.

***Fig 5—This 4- to 20-mA current-loop circuit** isolates the offset adjustment at the OUT LO pin, from the gain adjustment at the input.*

amplifier to a 4- to 20-mA current (**Fig 5**). IC_1's common-mode output range, OUT LO – CLK COM, allows you to offset the OUT LO pin, thereby providing the offset necessary for the voltage-to-current conversion. IC_2's A_1 amplifies the 200-mV reference by 5, and drives OUT LO to 1V. You can then select resistor values to fix the gain of the AD208 such that the normal-mode output range is 0 to 4V. One requirement of this circuit is that the output common-mode voltage plus the normal-mode voltage must be less than or equal to 6V.

IC_1's OUT HI drives a voltage-to-current converter comprising A_2, Q_1, Q_2, and a 250Ω resistor. This converter generates the requisite 4- to 20-mA output with as much as 25V of compliance for a 30V supply. The circuit connects Q_1 and Q_2 in a Darlington configuration to keep base-current errors to a minimum. Q_2 must dissipate as much as 600 mW over your desired operating temperature range. Therefore, a medium-power transistor such as a TIP 31A works best. You can trim the 4-mA offset by adjusting the gain of A_1 and then independently adjusting the span at the input side of IC_1. This independent adjustment minimizes the interaction between offset and input span commonly encountered in other circuits.

IC_4, an LM7815 voltage regulator, preregulates a 30V input and supplies 15V to the power oscillator. If you are using more than a couple of isolation amplifiers in a multichannel application, you need to be aware that the AD208 isolation amplifier, in combination with the AD246 oscillator, represents a reactive load to your preregulator. Higher-than-average currents can flow into IC_3 and can force your preregulator to go into current limiting.

Multiplexers and Switches
Contents

		Page
AN-355	Behind the Switch Symbol: Use CMOS Analog Switches More Effectively When You Consider Them As Circuits	12–3
AN-248	ADG201A/202A and ADG221/222 Performance with Reduced Power Supplies	12–5
AN-32	Single-Supply Operation of JFET Multiplexers	12–9
AN-249	Overvoltage Protection for the ADG5xxA Multiplexer Series	12–11
AN-35	Understanding Crosstalk in Analog Multiplexers	12–13
AN-250	Bandwidth, OFF Isolation and Crosstalk Performance of the ADG5xxA Multiplexer Series	12–21
AN-251	R_{ON} Modulation in CMOS Switches and Multiplexers; What It Is and How to Predict Its Effect on Signal Distortion	12–27

ANALOG DEVICES

AN-355
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Behind the Switch Symbol

Use CMOS Analog Switches More Effectively When You Consider Them as Circuits

by Jerry Whitmore

CMOS analog switches are widely used to make or break circuits in such applications as multiplexing and function switching. Ideally, they have zero resistance when closed, infinite resistance when open, no leakage, instantaneous glitch-free response, and no parasitic capacitance. While these assumptions are reasonably valid for low-frequency applications at moderate impedance levels, the good designer will always challenge them, to establish what errors may be introduced and even to determine whether the circuit configuration is viable.

SWITCH CIRCUITS

Figure 1 is a reasonable approximation of the circuitry in a single-pole dielectrically isolated CMOS switch (e.g., AD7510DI or AD7590DI series*).The dielectric isolation makes possible protection against latchup and overvoltage to ±25V beyond the supplies. Note that, for one polarity, conduction is via an N-channel FET; for the other polarity, it is via a P-channel FET. The two types are not perfectly symmetrical.

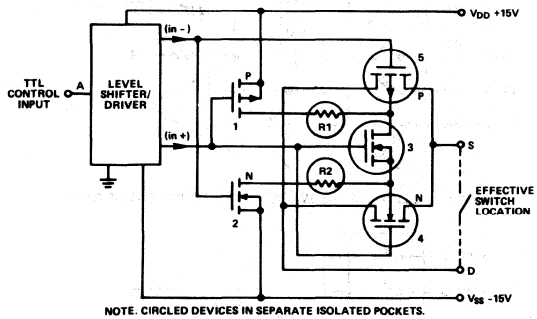

Figure 1. Typical output switch circuitry of the AD7590DI Series.

Figure 2 is an equivalent circuit of a pair of adjacent switches. The parameters are defined in Table 1. There are three principal categories of error one should be concerned about: low-frequency errors due to resistances and current leakage (switch open or closed), high-frequency and signal-transient errors due to stray capacitances (switch open or closed) and dynamic errors due to switching transients while the state of the switch is changing. Because of the present limitations of space, we shall for now consider just the first category, since it answers the most urgent question, "How well does the switch actually work for low-frequency signals?"

Although the leakage currents of the P- and N-channel transistors (devices 4 and 5 in Figure 1) might appear to tend to cancel, they don't, since the P channel is three times larger than the N channel.

Table 1. Nomenclature

C_{DS}:	Open-switch capacitance
C_S,C_D	Source, drain capacitance
R_{ON}	Series on resistance
S,D	Source, drain; electrically interchangeable
C_{SS},C_{DD}	Capacitance between any two corresponding switch terminals
I_{LKG}	Leakage current of back-gate diode

Because of the size mismatch of the reverse-biased source-or-drain-to-back-gate diodes, plus the differing lot-to-lot variations in breakdown voltage of the diodes, it is difficult to predict leakage or its tempco. However, maximum values at 25°C and over temperature are specified and 100% tested.

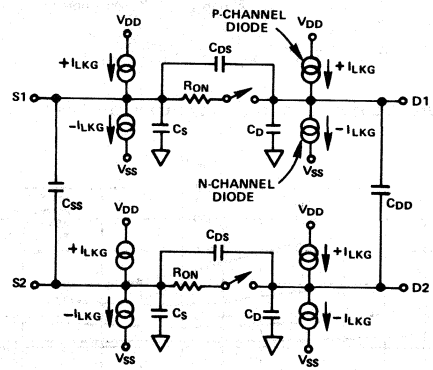

Figure 2. Equivalent circuit of a pair of adjacent switches.

Figure 3 shows the factors affecting dc performance for the *on* switch condition and how the various parameters affect the output voltage. Figure 4 shows typical curves of R_{ON}, as they appear on the product data sheet. They indicate how R_{ON}, is affected, as a function of input voltage, by supply voltages and by temperature, R_{IN} is lower and less signal-dependent at the higher supply voltages and lower temperatures.

How to minimize the influence of variable R_{ON} on circuit accuracy: Figure 5 shows a problem circuit–an inverting amplifier with four switched inputs. R_{ON}, in series with the 10-kilohm input resistor, affects the circuit gain. Even if it is compensated for at one level of suply voltage and analog input voltage, the input's variations will cause the gain to change and degrade the gain accuracy.

The most obvious solution–if the amplifier doesn't have to invert or act as a precision attenuator–is to use the amplifier in a nonin-

Reprinted from Analog Dialogue 15–2 1981

verting mode, as shown in Figure 6. Since there are no resistors in series, there is no effect on gain.

Another solution (Figure 7) is to connect the quad switch at the amplifier's summing point. Then the switch sees only millivolts–rather than volts–of signal variation, minimizing the variation of R_{ON} with signal. This solution can impair bandwidth, since capacitance C_S may require a capacitor in parallel with the feedback resistor for compensation. Also, I_{LKG}, flowing through the feedback resistor, may cause significant error, depending on the accuracy requirements. ($\Delta V_{OUT} = I_{LKG} \times R_F$).

Another possible solution is to use larger values of input and feedback resistance (Figure 8). Then the ΔR_{ON} variations will be small compared to the 1- megohm load. However, bandwidth will be affected by the larger R-C time constants.

Figures 7 and 8 do not compensate for the effects of variation of R_{ON} with temperature. A circuit that provides good compensation (Figure 9) uses one of the switches, wired *on*, in series with the feedback resistor. Its R_{ON} will tend to track that of the other switches on the same substrate with temperature; thus the feedback and input resistances will tend to track quite well, keeping the gain constant.

The principal dc effect in the switch *off* condition is that of I_{LKG} ($I_{D\ OFF}$ or $I_{S\ OFF}$), which will bias the output of a circuit by $I_{LKG} \times R_L$. Polarity of the error is determined by the dominant leakage polarity of a given switch.

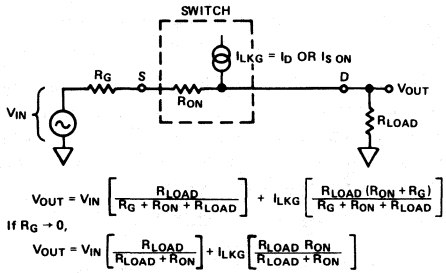

Figure 3. Effective circuit of switch in the ON condition.

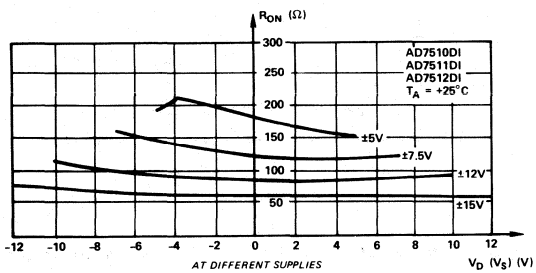

a. R_{ON} vs. V_D (V_S), as a function of $+V_{DD}$, ($-V_{SS}$.)

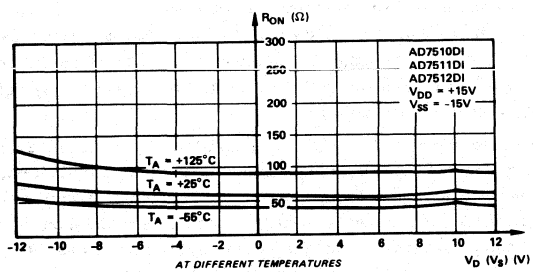

b. R_{ON} vs. V_D (V_S), as a function of temperature.

Figure 4. R_{ON} vs. input voltage as a function of supply voltage and temperature.

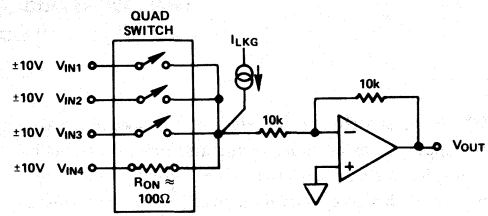

Figure 5. Unity-gain inverter with switched input.

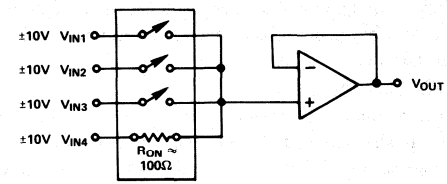

Figure 6. Noninverting solution.

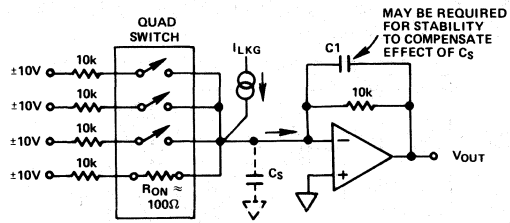

Figure 7. Connecting switch at the summing point.

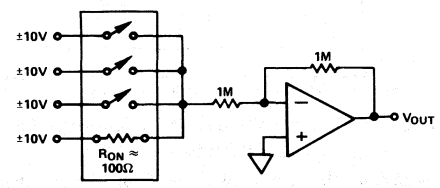

Figure 8. Using larger values of resistance.

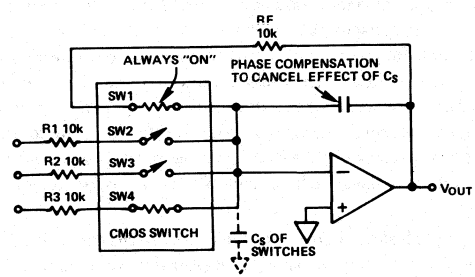

Figure 9. Switch in series with feedback resistor to compensate gain.

ANALOG DEVICES

AN-248
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

ADG201A/202A and ADG221/222 Performance with Reduced Power Supplies

by John Reidy

The ADG201A/202A and ADG221/222 are monolithic CMOS devices comprising four independently selectable switches. They are designed on an enhanced LC²MOS process and feature 44V supply maximum rating, a signal handling capability that extends to the supply rails, low R_{ON}, low leakage and low power dissipation. The ADG221/222 are latched versions of the ADG201A/202A.

The LC²MOS fabrication process enables these switches to operate over a wide range of supply voltages. They can comfortably operate anywhere in the 10V to 15V single or dual supply range, with only a slight degradation in performance. This application note discusses the variation in switch characteristics over this power supply range and contains graphs showing all the relevant performance curves. All test circuits and test conditions used to generate these curves are contained in the data sheets.

SWITCH ON-RESISTANCE, R_{ON}

The switch R_{ON} is defined by the parallel combination of the P and N channel devices of the output switch (See Figure 1). The gates of these devices are connected to the supply rails. Variations in the supply voltage will vary the gate drive and affect the switch R_{ON}. The R_{ON} sensitivity is about 5Ω/V for single supply and 3Ω/V for dual supply in the 10V to 15V supply range. See Figures 2 to 5.

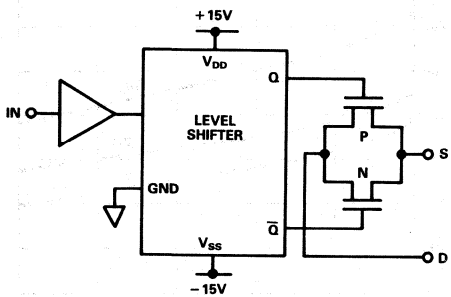

Figure 1. Output Switch Schematic

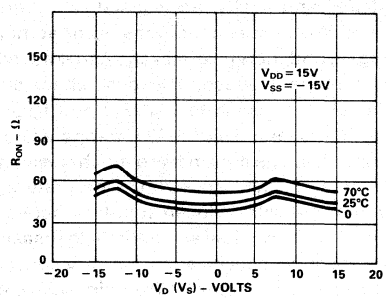

Figure 2. R_{ON} as a Function of V_D (V_S): Dual ± 15 Supplies

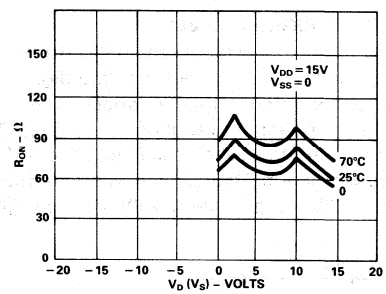

Figure 3. R_{ON} AS A Function of V_D (V_S): Single + 15V Supply

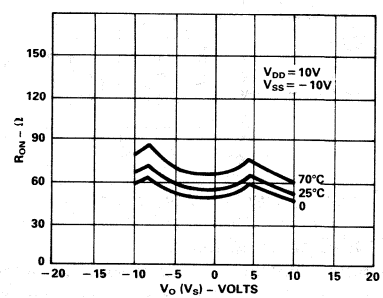

Figure 4. R_{ON} as a Function of V_D (V_S): Dual ± 10V Supplies

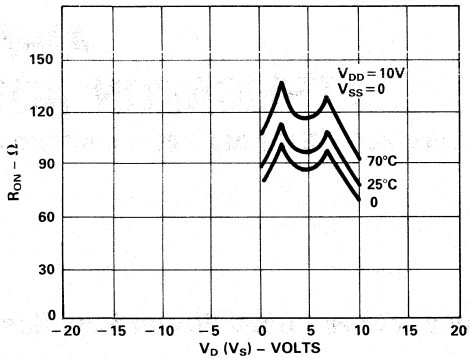

Figure 5. R_{ON} as a Function of V_D (V_S): Single +10V Supply

LEAKAGE CURRENTS I_S(OFF), I_D(OFF), I_D(ON)

I_S(OFF), I_D(OFF) and I_D(ON) are caused by reverse bias diode leakage. Worst case conditions occur at maximum stress voltages, and therefore, leakage currents reduce as the supply voltages reduce. The data sheet specified values at ±15V are 100nA for I_D(OFF) and I_S(OFF) and 200nA for I_D(ON) over the 0 to +70°C temperature range. In reality, these leakage currents are much lower, and are in the 1nA region in the 0 to +70°C temperature range. Since they improve with reduced supply voltages, actual leakage values are very small and difficult to measure precisely. The data sheet specification is a very conservative one and all these leakage currents remain well within this specification over the 0 to +70°C temperature range with reduced supplies.

DIGITAL INPUT TRIGGER LEVELS

Figure 6 shows the digital input schematic. An internal 5V generator provides a reference for all the digital inputs. This generator is immune to any variations in V_{SS}, and exhibits only very minor variations with V_{DD}. Typical changes in the input thresholds is 80mV, for 10V to 15V power supply variation. They remain well within standard TTL levels (<0.8V = Logic Low; >2.4V = Logic High). See Figure 7.

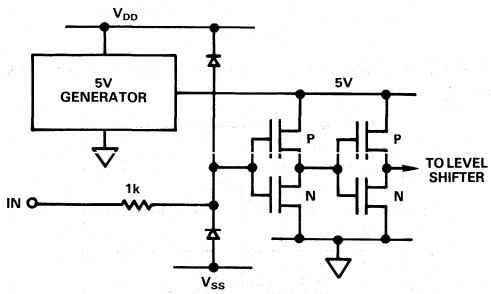

Figure 6. Digital Input Schematic

SWITCH TIMING (t_{ON}, t_{OFF}, t_{OPEN})

These are determined both by the input logic delays and the switch circuitry delays. The input logic timing versus power supply voltage remains fixed because of the internal 5V generator. Any variation in this timing is due only

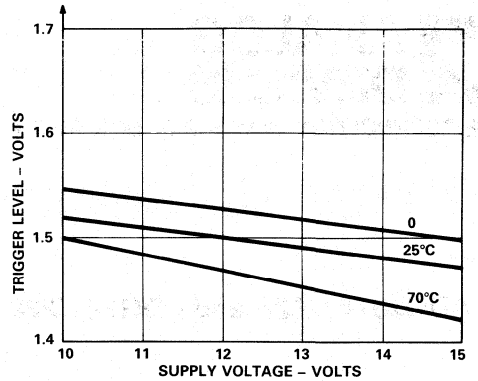

Figure 7. Trigger Levels vs. Power Supply Voltage, Dual or Single Supply Voltage

to the output switch circuitry and is confined to a 40ns change in the 10V to 15V range. See Figures 8 to 11.

LATCH TIMING (t_W, t_S, t_H, ADG221/222 ONLY)

These parameters are determined by the input circuitry only and remain constant over the 10V to 15V power supply variation.

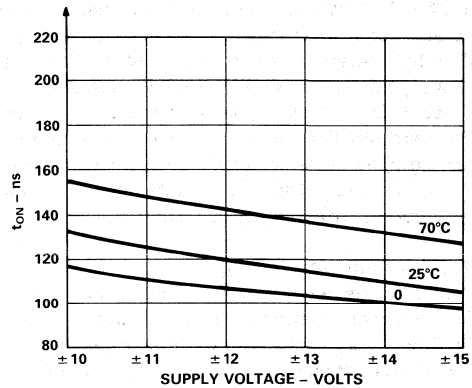

Figure 8. t_{ON} vs. Supply Voltage, (Dual Supply)

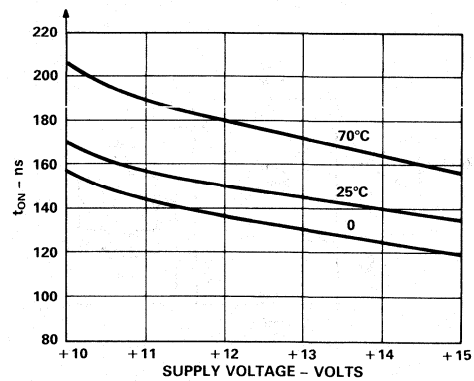

Figure 9. t_{ON} vs. Supply Voltage, (Single Supply)

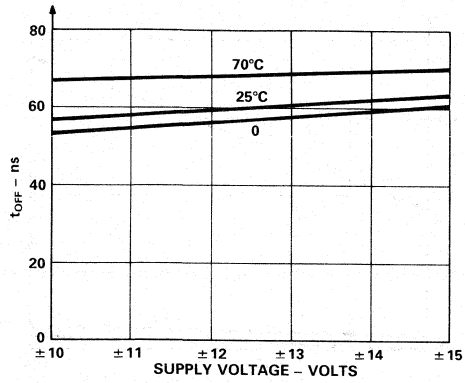

Figure 10. t_{OFF} vs. Supply Voltage, (Dual Supply)

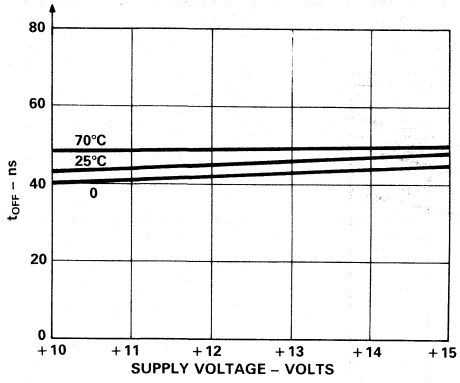

Figure 11. t_{OFF} vs. Supply Voltage, (Single Supply)

DYNAMIC CHARACTERISTICS

Off Isolation and channel-to-channel crosstalk are determined by stray capacitances. Some of these stray capacitances are junction diode capacitances and their values are modulated by the power supply voltage. However, the dominant influence is the package stray capacitance and these are immune to power supply variations. The net effect is that the switch AC performance does not alter with the power supply. See Figure 12.

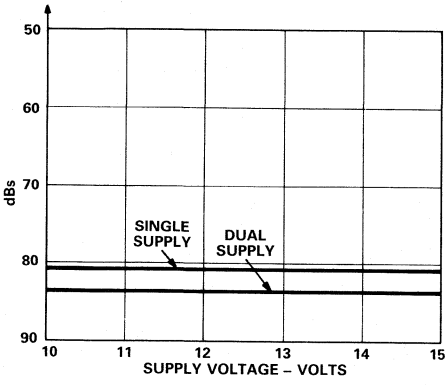

Figure 12. Off Isolation and Channel-to-Channel Crosstalk vs. Supply Voltage

CHARGE INJECTION

Charge Injection is caused by capacitive coupling between the digital control and the source and drain pins. When the digital inputs change, there is a corresponding level shifted change in the gates of the output switches (see Figure 6). This injects charge into the signal line. The amount of charge injected depends on the source voltage (V_S). The data sheet specification is 20 coulombs and is measured for $V_S = 0V$.

Figures 13 and 14 show how charge injection varies with Source Voltage for both 10V and 15V power supplies.

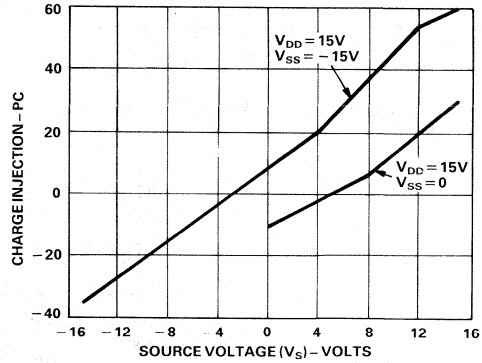

Figure 13. Charge Injection vs. Source Voltage (V_S) for Dual and Single 15V Supplies

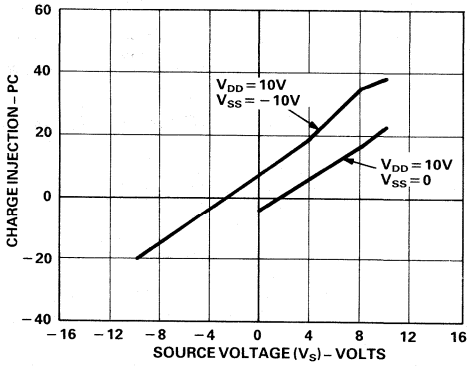

Figure 14. Charge Injection vs. Source Voltage for Dual and Single 10V Supplies

POWER SUPPLY CURRENT

Both I_{DD} and I_{SS} are directly proportional to the supply voltage. The values of these parameters decrease with reduced supply voltages, and thus they remain well within specification for a 15V to 10V power supply range. See Figures 15 to 17.

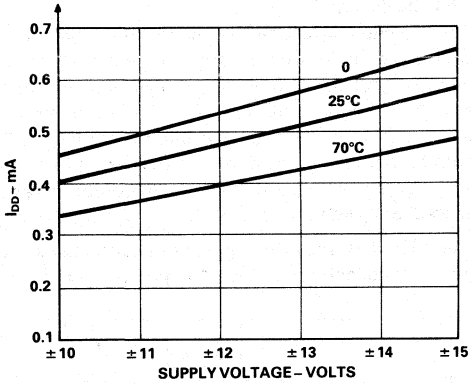

Figure 15. I_{DD} vs. Supply Voltage, (Dual Supply)

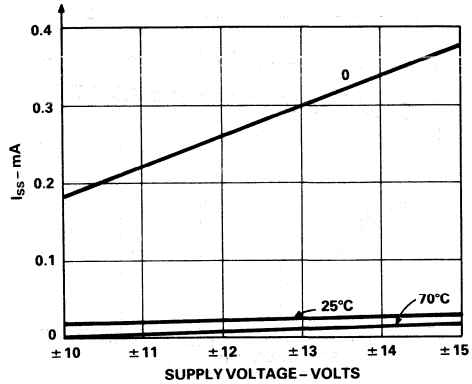

Figure 16. I_{SS} vs. Supply Voltage, (Dual Supply)

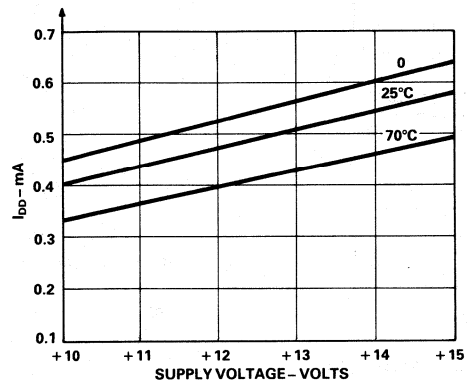

Figure 17. I_{DD} vs. Supply Voltage, (Single Supply)

ANALOG DEVICES

AN-32
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Single Supply Operation of JFET Multiplexers

INTRODUCTION

In addition to normal operation (+/–supplies), the ADI family of JFET multiplexers (MUX-08/88, MUX-24, MUX-16, and MUX-28) performs quite well in single supply systems. This Application Note explains single supply operation as it applies to JFET and CMOS multiplexers. Common requirements are in battery-operated systems and in microprocessor-based, single supply data acquisition systems. JFET and CMOS devices are compared for R_{ON} variation versus power supply voltage (V_S), then settling times.

CONNECTIONS FOR SINGLE SUPPLY OPERATION

Figure 1 shows single supply connections for the entire ADI JFET multiplexer family. Each multiplexer handles 0 to +10V signals with a +15V supply. The signal range is conservatively rated to be (V_S – 4V) as a maximum, and zero volt as a minimum.

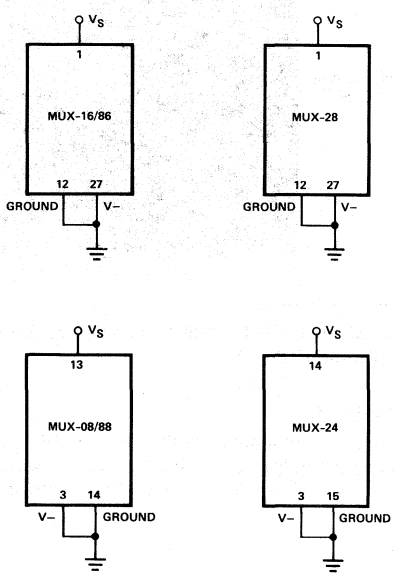

Figure 1. JFET Multiplexer Single Supply Connections

JFET VARIATION OF R_{ON} WITH V_S (MUX-08)

Figure 2 shows the test circuit and defines the test conditions (MUX-08). Figure 3 shows the performance of a MUX-08 driving a 1kΩ load. The positive voltage should be 1.10V and the negative voltage should be −0.4V. The reason for the output voltages being less (magnitude) than the above is due to the R_{ON} of the multiplexer switches. Curves 1 and 2 show that R_{ON} does not vary as V_S varies from +5V to +15V.

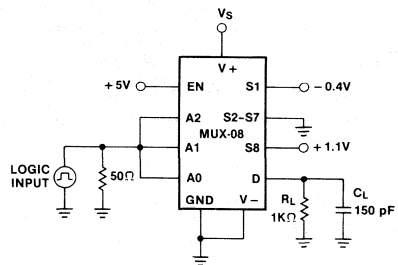

Figure 2. Test Circuit

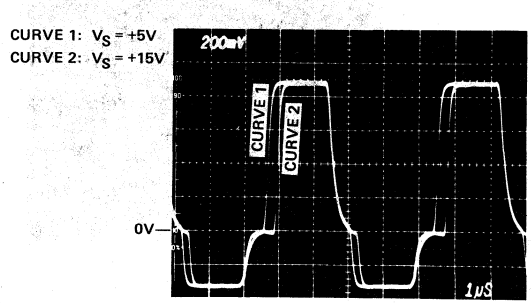

Figure 3. JFET Variation of R_{ON} with V_S

CMOS VARIATION OF R_{ON} WITH V_S (508 Pin-Compatible Device)

The CMOS multiplexer (connected as shown in Figure 4) **does** show a variation in R_{ON} as V_S is varied from +6V to +15V. This is evidenced by the curves shown in Figure 5. Note that while the positive peak voltages in Figure 3 are the same for both curves, the peaks differ in Figure 5.

One very important consideration when choosing a multiplexer is the nonlinearity (or distortion) introduced by the switch when it is ON. What is important is the **change** in R_{ON} which occurs because of external variations such as power supplies. In particular, the variation in R_{ON} shown in Figure 3 is 148 ohms. The R_{ON} at $V_S = +6V$ is 1000 ohms, while its value at $V_S = +15V$ is only 852 ohms. A change of 148 ohms represents a 1.48% error if the load resistor is 10,000 ohms. In battery-operated systems (which is what a lot of single supply applications are), distortion due to power supply variations is generally not acceptable.

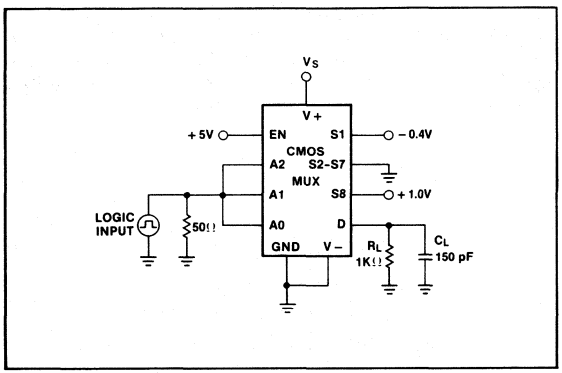

Figure 4. Test Circuit

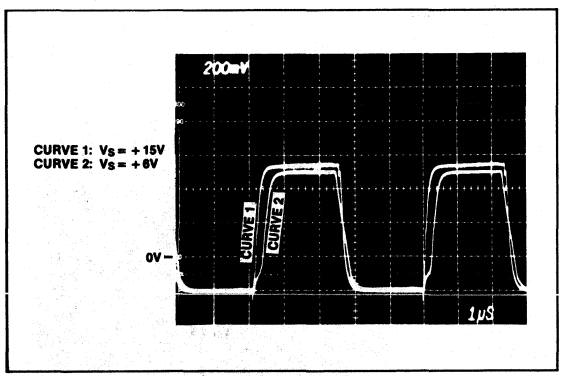

Figure 5. CMOS Variation of R_{ON} with V_S

CMOS vs JFET — EFFECT OF R_{ON} ON SETTLING TIME

Figure 6 defines the test conditions used for the JFET and CMOS multiplexer curves shown in Figure 7. In this case, R_L is large enough so that the output voltages will reach the input voltage levels. Note that MUX-08 does just that, while the CMOS multiplexer does not reach the final value.

The problem is settling time, and occurs because the R_{ON} of the CMOS device is considerably larger than the MUX-08 (852 ohms as opposed to 250 ohms). A final note concerns the fact that the multiplexers are switching signals at 400mV more negative than the negative supply voltage without appreciable distortion. In no circumstances should the input exceed one diode voltage below the negative supply voltage.

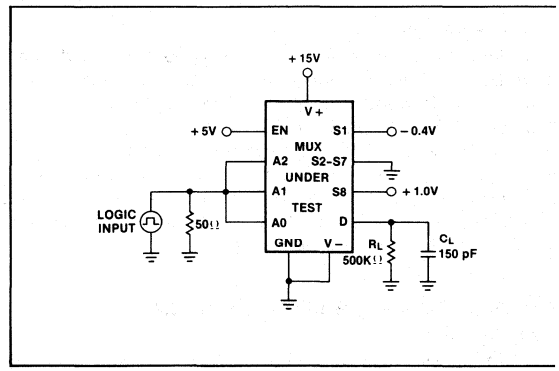

Figure 6. Test Circuit

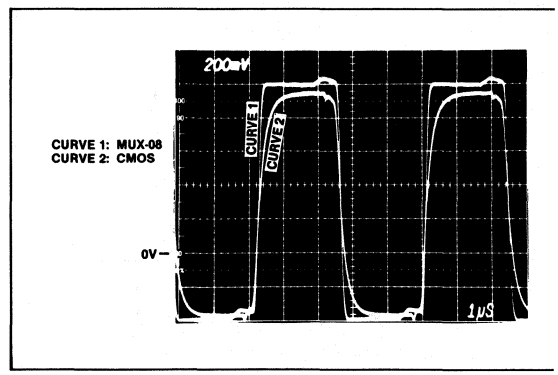

Figure 7. CMOS vs JFET Settling Time (Unloaded Output Voltage)

CONCLUSION

The information presented has shown how JFET multiplexer handle analog inputs in single supply systems, with R_{ON} independent of power supply variations, and with fast settling time.

)NE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Overvoltage Protection for the ADG5XXA Multiplexer Series

by Dan Sheehan

The ADG5XXA* multiplexer series is a family of single 8/16 channel and dual 4/8 channel parts. They are available in both latched (ADG52XA) and unlatched ADG50XA) versions. These are high performance multiplexers that offer the following notable features: specifications for both single and dual supply operation, 1nA max leakage current, 200ns max turn-on and turn-off times and TTL compatibility with reduced single or dual supplies down to 5V.

The multiplexers are specified to operate with an analog input signal range within the supply rails, i.e., $V_{SS} \leq V_S$ $V_D) \leq V_{DD}$. The parts are not internally overvoltage protected (i.e.,with resistors) to allow V_S (V_D) to exceed the power supply rails. Thus, the ADG5XXA series, without external overvoltage protection, is best suited for use in systems where the analog input signals come from sources within the system, such as from op amps powered from the same supplies as the multiplexers.

However, in many applications such as process control systems, the analog input signals can originate from sources external to the system which contains the multiplexer.This can be potentially destructive to the multiplexer for two principal reasons:

. The multiplexer power supplies may be turned off while the analog signals are still present.

2. The signal lines may receive induced voltage spikes which exceed the supplies to the multiplexer.

This application note addresses the above problems and outlines the protection circuitry required to allow the analog input signals to exceed the supply rails over the -40°C to +85°C temperature range.

OVERVOLTAGE PROTECTION: WHY?

The following is a brief and simplified analysis of what happens when a signal applied to the S(D) input of the multiplexer exceeds the power supply rails. The results outlined apply equally with the power supply rails at ±15V or GND.

*ADG506A/ADG507A, ADG508A/ADG509A, ADG526A/ADG527A, ADG528A/ADG529A.

The basic CMOS switch consists of an n-channel MOSFET in parallel with a p-channel MOSFET. This structure which yields an excellent analog switch contains p-n junctions between the signal path and power supplies. These p-n junctions or diodes are reverse biased under normal operating conditions, i.e., $V_{SS} \leq V_S$ $(V_D) \leq V_{DD}$. However, if V_S (V_D) exceeds either power supply rail by approximately 0.7V, the normally reverse biased junctions will become forward biased. This means that with an analog input overvoltage, the S(D) input of the multiplexer will appear as a diode connected to the relevant power supply voltage. Therefore, $V_S(V_D)$ is clamped to a maximum of 0.7V greater than either supply rail and large currents can flow that will destroy the parts unless restricted. The simplest form of protection uses resistors in series with the S (D) inputs to limit the input current to safe levels.

PROTECTING TYPICAL MULTIPLEXER APPLICATION CIRCUITS

This section shows two typical multiplexer application circuits and outlines the operating conditions and protection circuitry required for the safe operation of the parts with an overvoltage on the analog inputs. The ADG506A is shown in both circuits but the same conditions and results apply to any part in ADG5XXA series.

Generic Multiplexer Circuit

Figure 1 shows the general use of a multiplexer. R1–R16 (2.7kΩ) provide overvoltage protection.

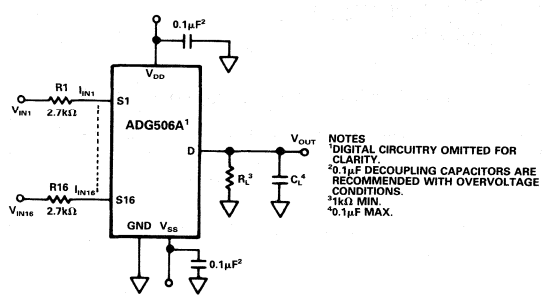

Figure 1. ADG506A Multiplexer Circuit

The series resistors are sufficient to protect the multiplexer over the −40°C to +85°C temperature range under the following conditions:

- V_{DD}/V_{SS} = GND with ±15V max applied continuously to all the analog inputs ($V_{IN1} - V_{IN16}$).

  With the supplies at GND, V_S and V_D are clamped to ±0.7V max. The total analog input current ($I_{IN1} - I_{IN16}$) is limited to 85mA max ([14.3V/2.7kΩ]×16).

  Power on/off is also allowed with $V_{IN1} - V_{IN16} = \pm 15V$, but V_{DD} and V_{SS} must power on/off within 0.5 seconds max of each other. The protection outlined is adequate with power supplies that remain low impedance in the off state. However, if the supplies go high impedance or open circuit in the off state, additional protection such as 15V TransZorbs* connected from both V_{DD} and V_{SS} to GND is recommended. The TransZorbs are recommended to suppress any transient voltages exceeding the Absolute Maximum Ratings of the multiplexer. Excessive transients commonly occur when powering-on supplies that are high impedance or open circuit in the off condition.

- V_{DD}/V_{SS} = +15V/−15V with ±35V max applied for a duration of 50ms max (voltage spike,10% duty cycle) to all analog inputs simultaneously.

  Under this condition, V_S and V_D are clamped to ±15.7V max.

- V_{DD}/V_{SS} = +15V/−15V with ±35V max applied continuously to any one analog input ($V_{IN1} - V_{IN16}$).

De-Multiplexer Circuit

Figure 2 shows a typical de-multiplexer circuit. In this application it is used to provide a 16-channel sample-and-hold circuit, allowing 16 separate digitally programmable voltages to be generated from the output of a single DAC.

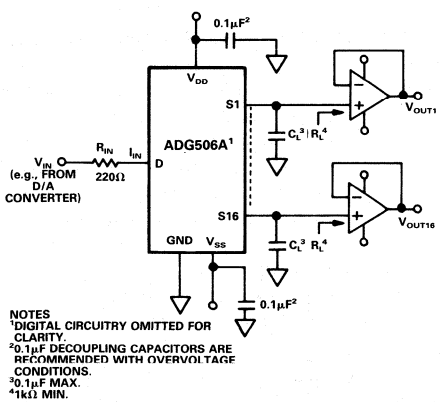

Figure 2. ADG506A De-Multiplexer Circuit

*TransZorb is a trademark of General Semiconductor Industries, Inc.

R_{IN} (220Ω) provides the necessary overvoltage protection.This protection is required if, for example, the supplies to the multiplexer are turned off but the supplies to the DAC are still present (e.g., separate supplies). However, the multiplexer is sufficiently robust when powered-off to allow the stored charge on C1 – C16 (0.1μF max) to discharge through it without causing failure. The protection outlined applies over the −40°C to +85°C temperature range for the following condition:

- V_{DD}/V_{SS} = GND with ±15V max applied continuously to V_{IN}.

  With this condition, V_S and V_D are clamped to ±0.7V max. I_{IN} is limited to 65mA max (14.3V/220Ω). In general, R_{IN} is not required because the output short circuit current of an op amp or a voltage output DAC is limited to less than 65mA.

  Power on/off is also allowed (0.5 second max timing skew) with $V_{IN} = \pm 15V$. If the supplies go high impedance or open circuit in the off condition, 15V TransZorbs should be connected from V_{DD} and V_{SS} to GND

ADVANTAGES OF EXTERNAL PROTECTION

The overvoltage protection outlined above uses external current limiting resistors in series with the S (D) inputs of the multiplexer. These resistors could also be incorporated on-board the die of the multiplexer, but this would lead to power dissipation problems and therefore restrict the overvoltage amplitude allowed. Overvoltage protected multiplexers from some manufacturers use on-chip resistors (with the associated power dissipation problem), but these are employed to protect the S inputs only. This means that in the de-multiplexer application shown in Figure 2, an external resistor (R_{IN}) is still required if the source of V_{IN} can supply greater than 65mA. In this instance, on-chip resistors actually cause problems by increasing the settling time of the sample-and-holds. In conclusion, external versus internal protection is preferable for the two reasons outlined above.

Note: The information furnished in this application note is believed to be accurate and reliable. It is based on experimental results from three fabrication lots. However, no responsibility is assumed by Analog Devices for its use.

ANALOG DEVICES

AN-35
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Understanding Crosstalk in Analog Multiplexers

NTRODUCTION

One of the most troublesome errors in analog multiplexers s crosstalk. Various schemes have been devised to reduce ts effects. One designer will terminate the multiplexer in a 10kΩ resistive impedance. Another will short the multiplexer node to ground between address changes with an analog switch. A third engineer will terminate the multiplexer node n 1MΩ because he doesn't want to live with the attenuation which comes about with any lower impedance. What is confounding about these three situations is that the solution is correct in each case. THE CORRECT SOLUTION IS DICTATED BY THE APPLICATION.

To understand why the solution is application dependent, it s necessary to dig rather deeply into what crosstalk really s. When this is done, crosstalk is found to have not one, but hree components in a multiplexer. To differentiate the components one from the other, it is convenient to give them names:

1. Static crosstalk (CT)
2. Dynamic crosstalk (DCT)
3. Adjacent Channel crosstalk (ACCT)

This application note explains the three crosstalk components qualitatively and quantitatively. The qualitative discussion tells what component(s) should be considered in various applications. The quantitative discussion uses both heoretical and empirical information to arrive at conclusions about what performance should be expected.

STATIC CROSSTALK (CT)

To introduce the concept of crosstalk, Figure 1 is helpful. A basic analog switch may be constructed with a FET (JFET or CMOS) and a suitable driver which switches it OFF and ON, as shown in Figure 1a. The equivalent circuit, as shown in Figure 1b, models the analog switch such that when the ideal switch SW) is closed, the switch has an ON resistance R_{ON}. When SW is open, the OFF impedance is determined by C_{EQ}. A wo-channel multiplexer circuit, made up of two analog switches connected as shown in Figure 1c, shows how signals from one channel can be coupled into the other channel. Theoretically, V_{OUT} consists of e_1 modified by the resistor divider formed by R_{ON1} and R_L (assumes reactance of C_L is $\gg R_L$). However, the capacitance of switch number two (C_{EQ2}) does couple some portion of e_2 into V_{OUT}. This is the simplest example of crosstalk.

The model which explains static crosstalk is relatively simple and may be derived from the OFF isolation model. Figure 2a shows the OFF isolation model as capacitive coupling from the input to the output of an OFF switch. This condition may be duplicated in Figure 1c by opening SW_1 and setting $e_2 = 0$. Coupling from input to output is accomplished through C_{EQ},

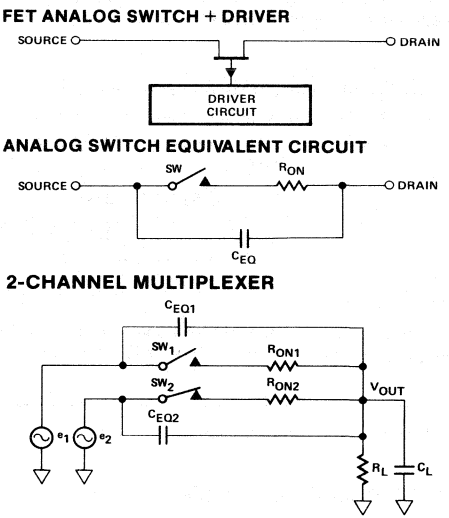

Figure 1. Essentials of an Analog Multiplexer

and this parameter may be computed from measurements of V_{IN}, V_{OUT}, and frequency. In the case of static crosstalk, C_{EQ} is shown coupling into a parallel combination of R_{ON} with R_L and C_L (Figure 2b). The two channel multiplexer shown in Figure 1c reduces to the circuit in Figure 2b, where $e_1 = 0$, $e_2 = V_{IN}$, and C_{EQ} is the coupling capacitance from e_2 to V_{OUT}.

Since R_L is generally 10kΩ or more, and typical analog switches are less than 1kΩ, static crosstalk is much smaller than OFF isolation. The crosstalk and OFF isolation numbers quoted on analog multiplexer data sheets are derived from the models shown in Figure 2. Unfortunately the one component of crosstalk specified is the least troublesome of the three. However the crosstalk figures on data sheets will alert the designer to those devices which absolutely will not satisfy his requirements.

There are applications where the static crosstalk specification given on data sheets is adequate. When the multiplexer is being used as a one-of-many switch, and is not being cycled through all channels on an automatic basis, then the static crosstalk component will give accurate prediction of the actual performance. Examples of such applications are:

1. Audio/Video Selector Switch
2. Programmable Gain Amplifier
3. Programmable Power Supply

(a) OFF ISOLATION EQUIVALENT CIRCUIT

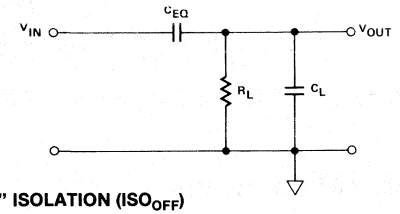

"OFF" ISOLATION (ISO_{OFF})

The proportionate amount of a high frequency analog input signal which is coupled through the channel of an "OFF" device. This feedthrough is transmitted through $C_{DS(OFF)}$ to a load comprised of $C_{D(OFF)}$ in parallel with an external load. Isolation generally decreases by 6dB/octave with increasing frequency.

(b) STATIC CROSSTALK EQUIVALENT CIRCUIT

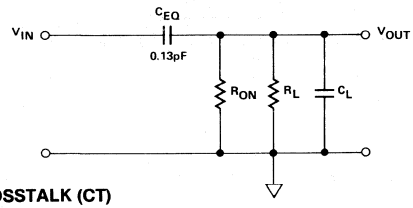

CROSSTALK (CT)

The proportionate amount of cross-coupling from an "OFF" analog input channel to the output of another "ON" output channel.

Figure 2. Model for Static Crosstalk

DYNAMIC CROSSTALK (DCT)

The dynamic crosstalk model can be derived from Figure 3. The switch SW_1 represents one condition on the multiplexer node (SW_1 is open). Actually SW_1 is continually switching between OFF and ON. This is represented in Figure 3b. In order to reduce crosstalk, multiplexers are designed to have break-before-make switching so that no two channels are addressed at the same time. The finite open time of SW_1 (shown in Figure 3b) represents the break-before-make action. There are two "open" conditions on the multiplexer node per cycle of the clock; thus the equivalent nodal resistance (R_{EQ}) may be computed as given in Figure 3b. Table I shows some typical values of static and dynamic crosstalk. Static crosstalk values are given in lines 1 and 12. There is a change in crosstalk as the clock frequency (f_{CLK}) is varied. Starting at line 4 notice the variation in crosstalk as R_L is varied from 10kΩ to 100kΩ while f_{CLK} remains constant at 100kHz. While Table I yields some theoretical values which give insight into the operation of dynamic crosstalk, a working multiplexer will have different values of f_{CLK} with respect to the maximum value of f_{SIG}. The real world situation will be analyzed in a later section of this paper.

Examples of multiplexer applications which are dynamic in nature are:

1. Industrial Process Control
2. Telephony
3. Data Acquisition Systems
4. Telemetry

Each one of the above applications are a form of Time Divi sion Multiplexing. In other words, these are sampled-dat

(a) DYNAMIC CROSSTALK EQUIVALENT CIRCUIT

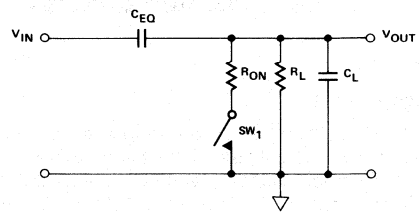

NOTE: SWI is a time dependent switch. Its characteristic is shown in Figure 3b.

(b)

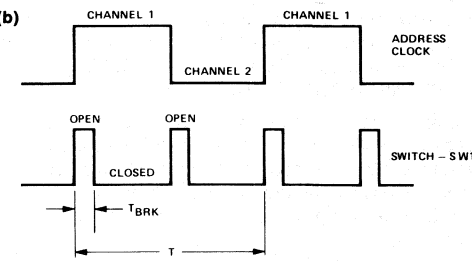

$T \equiv$ Period of address clock

$T_{BRK} \equiv$ Break-Before-Make Time

If $R_L \gg R_{ON}$, $R_{EQ} \cong \dfrac{R_{ON}(T - 2T_{BRK}) + 2R_L T_{BRK}}{T}$

Figure 3. Model for Dynamic Crosstalk

Table 1. Computed Values of Static and Dynamic Crosstalk

LINE NO.	f_{SIG} Hz	f_{CLK} Hz	T μsec	T_{BRK} μsec	R_{ON} OHMS	R_L OHMS	R_{EQ} OHMS	C_{EQ} pF	CROSS-TALK dB
1	10K	0	—	0.80	300	10K	291	0.30	105
2	10K	20K	50	0.80	300	10K	602	0.30	99
3	10K	40K	25	0.80	300	10K	913	0.30	95
4	10K	100K	10	0.80	300	10K	1845	0.30	89
5	10K	100K	10	0.80	300	20K	3448	0.30	84
6	10K	100K	10	0.80	300	40K	6650	0.30	78
7	10K	100K	10	0.80	300	100K	16.25K	0.30	70
8	20K	50K	20	0.80	300	10K	1068	0.30	88
9	20K	50K	20	0.80	300	20K	1872	0.30	83
10	20K	50K	20	0.80	300	40K	3474	0.30	78
11	20K	50K	20	0.80	300	100K	8275	0.30	70
12	20K	0	—	0.80	300	100K	291	0.30	99

stems where each channel is being continuously sampled
id the information for a given channel is contained in a
ven time slot. In these applications, the static crosstalk is
nost meaningless, since the wrong choice of R_L (or f_{CLK})
in be disastrous.

DJACENT CHANNEL CROSSTALK (ACCT)

djacent channel crosstalk is the most confusing compo-
ent of crosstalk. In addition to its confusing nature, in
ome cases, it is the most dominant component. While both
atic and dynamic crosstalk are capacitive in nature, i.e.,
ey vary with frequency at 6dB/octave, the adjacent chan-
el crosstalk is **invariant with frequency**. In other words, it is
ossible to have crosstalk **when multiplexing DC signals**
uch as the outputs of thermocouples, pressure trans-
ucers, etc. The parameters which must be dealt with are
L, C_L, R_{ON}, and f_{CLK}. In addition, the break-before-make
me ($= T_{BRK}$) of the multiplexer is of importance. Before div-
g into the details of this component of crosstalk, it will be
elpful to define what is meant by ACCT.

ne term "adjacent" refers to time only. In other words,
annel two is adjacent to channel one if channel two **im-
ediately** follows channel one in time slots. Since the chan-
el following is the "adjacent" channel, then channel one is
ot adjacent to channel two, but rather the other way
round. Figure 4 illustrates the concept of adjacent chan-
els. Assuming the multiplexer had, say, 1V on channel one,
V on channel two, etc., then the output would look like the
urve labeled "channel addressed." What is important
bout the waveforms in Figure 4 is the way the adjacent
hannel (in time) is shown. Note that while channel two is
djacent to channel one, channel one is itself adjacent to
hannel eight.

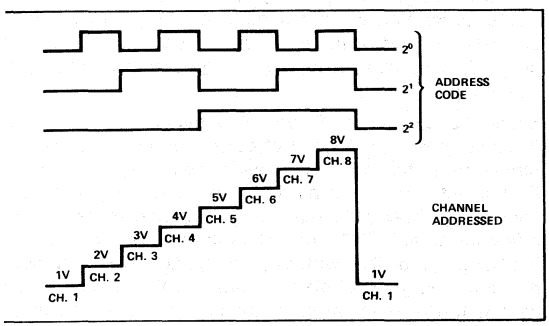

igure 4. Adjacent Channel Concept

he fact that information is "carried forward" from one
hannel to the next (in time) suggests a storage mechanism
s causing ACCT. Thus the multiplexer nodal capacitance
ecomes the prime suspect. Figure 5 illustrates how infor-
nation is carried forward from one channel to the next as
he addresses are changed. The address code is shown in
igure 5a, while Figure 5b shows the theoretical multiplexer
utput. Note that the even numbered channels have zero
olt on them, while the odd channels have their channel
umber in volts. This arrangement best illustrates how the

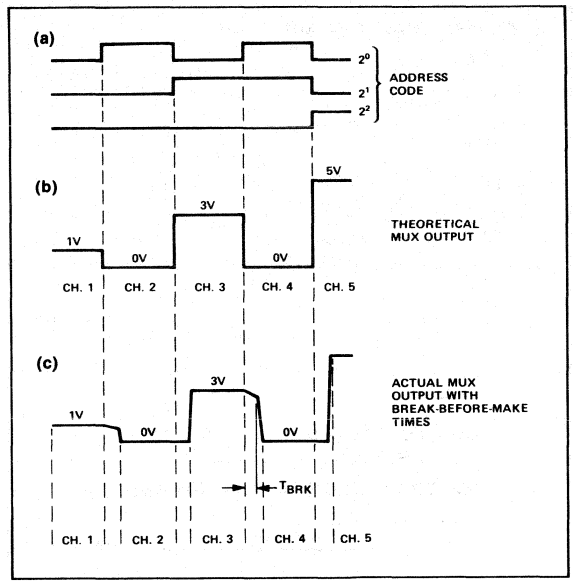

Figure 5. Adjacent Channel Crosstalk

information is transferred to the adjacent channel (as shown in Figure 5c). While the theoretical MUX output switches from channel three (3 volts) to channel four (0 volt) at the moment of the address change, note the delay in the actual MUX output caused by T_{BRK}. During this time the MUX node discharges along an RC curve determined by the load capacitance (C_L), and the load resistance (R_L). When the break-before-make time (T_{BRK}) is over, channel four is turned ON and the RC product is suddenly reduced to $R_{ON}C_L$. A curve which details how this all takes place is shown in Figure 6. Before leaving Figure 5, the arrangement suggests a method of avoiding adjacent channel crosstalk. In other words, the alternate grounding of channels prevents channel one signals from reaching channel three... channel three from reaching channel five, etc.

The curve in Figure 6a shows a typical nodal discharge for a set of real world conditions. The curve is normalized and T_{BRK} is chosen to be 900nsec. An accepted method of measuring T_{BRK} is from the 50% point of the channel which has been turned OFF to the 50% point of the channel which is being turned ON. This concept is illustrated in Figure 6b. In this case (Figure 6a) T_{BRK} is measured from the moment of the address change. While this is not totally correct, the agreement between theoretical and actual results is good enough to justify the simpler model which is derived. Since most designers are interested in crosstalk which is less than the resolution of the discharge curve, the ACCT vs. time graph gives crosstalk down to 90dB. In other words, the ACCT is down 90dB in less than 1.25μsec.

Adjacent channel crosstalk is a problem in every application where dynamic crosstalk must be considered; however there are techniques to minimize its effects. A popular way to diminish adjacent channel crosstalk is to short the

12

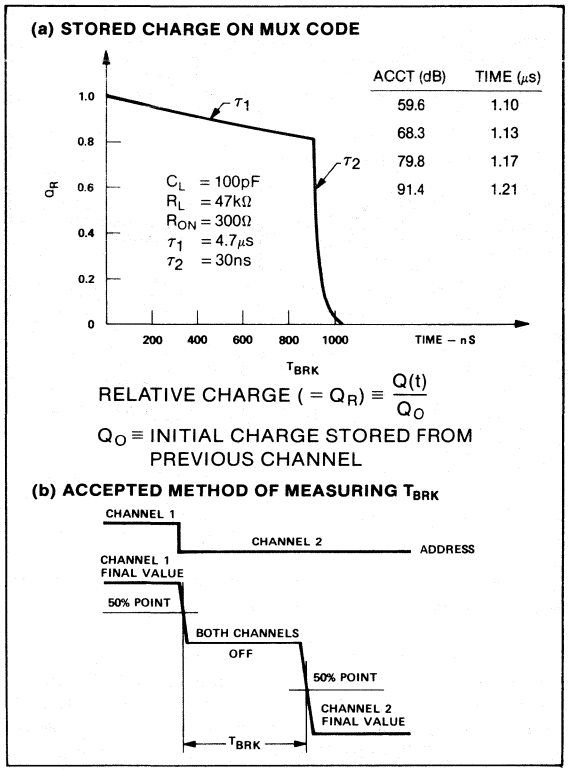

Figure 6. Stored Charge Decay and Definition of T_{BRK}

multiplexer node to ground between address changes. This requires an additional analog switch which should be fast and have low R_{ON}. An alternative approach to reducing adjacent channel crosstalk is to ground every other channel in a multiplexer. This technique was illustrated in Figure 5.

MEASUREMENT OF STATIC CROSSTALK

Figures 7 and 8 give the element values for a typical PMI JFET MUX-08 on channel three. In the case shown, the OFF isolation was first measured and found to be 75dB. With R_L and f_{SIG} known, then C_{EQ} was calculated. Once C_{EQ} is known, then R_{EQ} may be calculated from the static crosstalk measurement made in Figure 8. R_{EQ} is the parallel combination of R_L and R_{ON}; thus it is possible to compute R_{ON} and this value is also shown in Figure 8. The measurements thus far are relatively simple and only require a voltmeter which is capable of measuring signals which are 100dB below the reference signal. On the other hand, the measurement of dynamic crosstalk is a bit more involved, and requires a more complex system.

MEASUREMENT OF DYNAMIC CROSSTALK

The crosstalk measuring system shown in Figure 9 is to be used for measuring dynamic crosstalk. The signal from M_5 is fed into M_1 where it is multiplexed onto the OUT terminal.

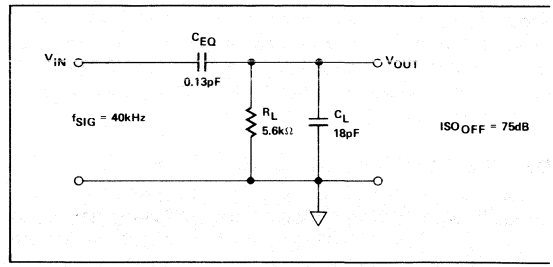

Figure 7. Typical OFF Isolation Element Values

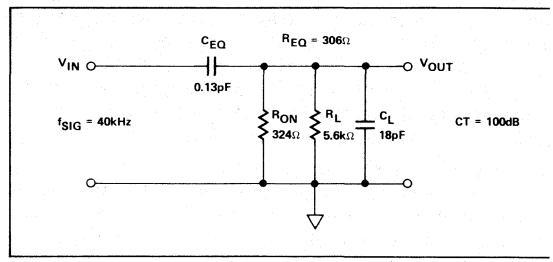

Figure 8. Typical Static Crosstalk Element Values

M_1 contains the multiplexer under test and a decoding ci cuit. The decoding circuit allows the selection of any tw channels to be used as a two channel multiplexer. M_2 is high-speed buffer used for driving the IN terminal of M_3. M contains a multiplexer operated in a demultiplexer mod along with decoding circuitry to allow several combination of two channel demultiplexing. The signal which appears o S_{3A} is fed through M_4 (high-speed buffer) to M_8 for spectrur analysis. In short, if no errors are introduced by the mult plexer-demultiplexer system, the output should be the sam as the input.

Since the system in Figure 9 is capable of measurin dynamic crosstalk, a good check of its performance is t repeat the static crosstalk measurements. M_3 is set to hav IN connected to S_{3A} at all times. M_1 is set to have S_3 cor nected to OUT, and the signal thus measured is taken as th reference signal. Static crosstalk is measured by connec ing S_1 (or S_8) to OUT, with V_{IN} still applied to S_3, and agai measuring V_{OUT}. The relative signal levels represent stati crosstalk. This measuring technique was used to verify th accuracy of the system.

The measurement of dynamic crosstalk leaves M_3 exactl as in the static case. With V_{IN} connected to S_3, M_1 is switche between S_1 and S_8. The signal frequency (f_{SIG}) was 40kH and f_{CLK} was 100kHz (see Figure 10). From the crosstal measured, the equivalent resistance (R_{EQ}) is computed to b 1150Ω (see Figure 10a). To verify the validity of this measure ment, R_{EQ} was calculated using the formula in Figure 10 (T_{BRK} was measured separately). Since there is very goo agreement between these two independently derive values, both the measurement technique and the dynami crosstalk model are valid.

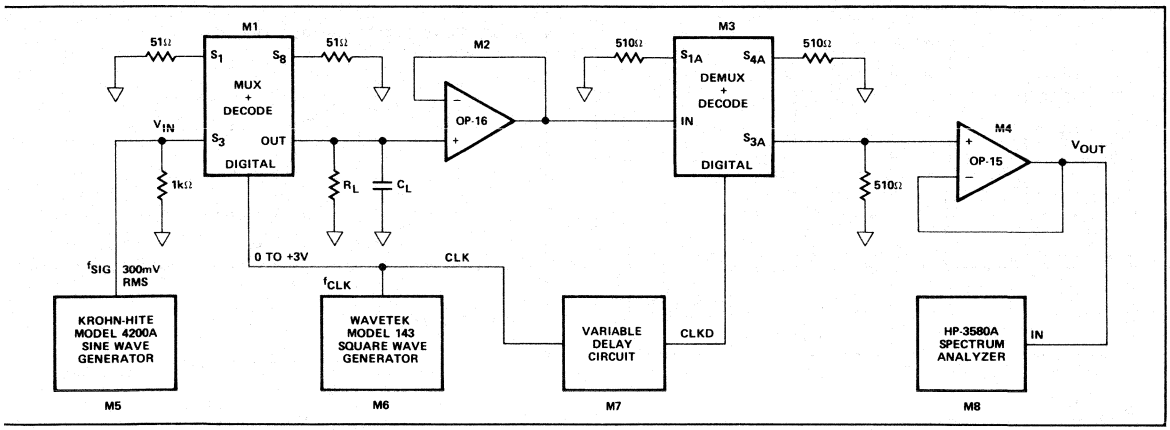

Figure 9. Dynamic Crosstalk Measuring System

(a) TYPICAL DYNAMIC CROSSTALK ELEMENT VALUES

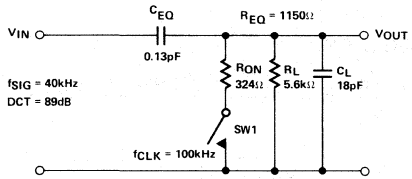

NOTE: SWI is a time dependent switch. Its characteristic is shown in Figure 10c.

(b) SYSTEM DYNAMIC CROSSTALK

f_{SIG}	f_{CLK}	R_L	R_{EQ}	DCT (C_{EQ} = 0.13pF)	DCT (C_{EQ} = 0.5pF)
Hz	HZ	Ω	Ω	dB	dB
10K	100K	10K	1718	97.1	85.4
10K	100K	22K	3463	91.0	79.3
10K	100K	33K	5059	87.7	76.0
10K	100K	47K	7090	84.7	73.0
10K	100K	100K	14.78K	78.4	66.7

(c)

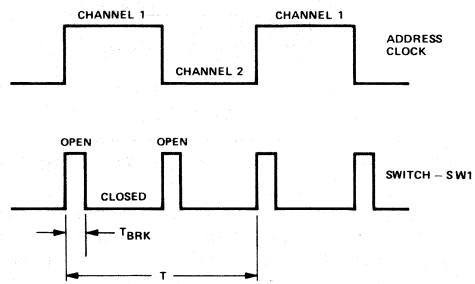

$T = 5\mu s$
$T_{BRK} = 725ns$
$R_{EQ} = 1089\Omega$
T ≡ Period of Address Clock
T_{BRK} ≡ Break-Before-Make Time

$$\text{If } R_L \gg R_{ON},\ R_{EQ} \cong \frac{R_{ON}\,(T - 2T_{BRK}\,(+\,2R_L\,T_{BRK})}{T}$$

Figure 10. Computed Dynamic Crosstalk for Actual Multiplexer

The numbers shown in Figure 10 apply to the measurement system, but are unlikely in a real multiplexer. To satisfy sampling theory limitations, f_{SIG} must be less than one-half the sampling frequency. Assuming f_{CLK} = 200kHz then each channel in a multiplexer is addressed for 5μsec. This means that it takes 40μsec to sample all channels of an eight channel multiplexer. In other words, **each channel** is sampled at a 25kHz rate. Thus the maximum value of f_{SIG} would be 12.5kHz. Figure 10b gives values of dynamic crosstalk (DCT) which would be experienced if the values of R_{ON} and T_{BRK} shown in Figures 10a and 10b were used. The first DCT column lists the values for a C_{EQ} of 0.13pF (measured value of channel three). The second DCT column shows the performance for C_{EQ} = 0.5pF. The purpose for the second column is to point out how critical minimizing stray capacitance is to good crosstalk performance.

MEASUREMENT OF ADJACENT CHANNEL CROSSTALK

The system shown in Figure 11 was used to measure adjacent channel crosstalk (ACCT). M_1 drives the address lines of the MUX system and the gating input of M_4. By setting the period of M_4 (T_2) to 10μsec, the pulse rate out of M_4 is controlled by the pulse rate of M_1 (40μsec) coming into the gate input of M_4. The output of M_4 is in the complement mode

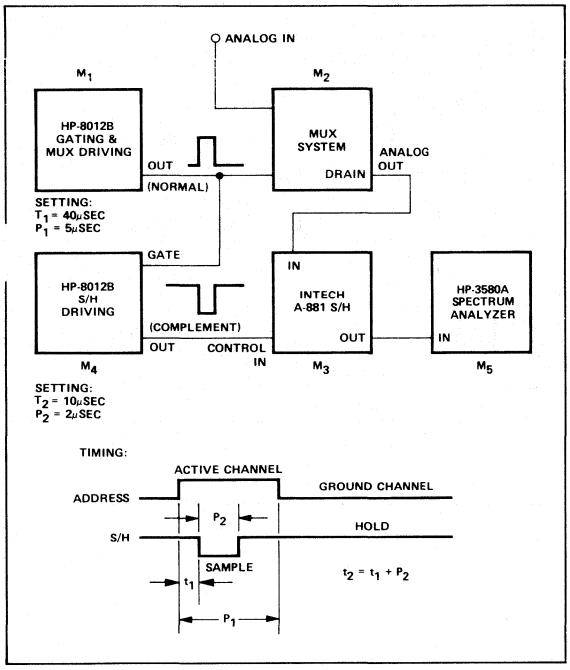

Figure 11. Adjacent Channel Crosstalk Measuring System

because the control input to M_3 causes the S/H to HOLD when the input is high (1). Thus the sample period occurs during the time P_2. M_4 also can delay its pulse relative to the pulse out of M_1, thereby allowing measurements of crosstalk versus t_1 (start of the sample time). This information is valuable because in many systems, a sample/hold is used with a successive approximation ADC to encode the analog output of the MUX. As will be shown, the ACCT can be made negligible if a sufficient time elapses before going to the HOLD mode for encoding the data. Since "time is money," the term "sufficient time" becomes important.

The nature of sample/holds and the nature of spectrum analyzers can cause some apparent discrepancies in the data observed by this measurement system. It is important to note the spectrum analyzer "sees" the average of **everything** that is presented to its input terminals. While it is true the sample/hold holds the last value it "saw," the spectrum analyzer also looks at the signal present during the sample/hold's sample time. Thus the equation which expresses the signal level present as a function of time must also account for the true averaging of the spectrum analyzer. Figure 12 shows the equations (12c) and the definitions of the terms used in the equations (12a and 12b). The term N_O is the **relative** signal level which the spectrum analyzer measures. If the model of the signal decay shown in Figure 12a is the correct one to explain the ACCT, then the computed value of N_O should correspond to the measured values. As will be shown in Figure 14, the agreement does in fact justify the model; however it was necessary to choose the measurement conditions **very carefully.**

(a) VOLTAGE DECAY ON MUX OUTPUT

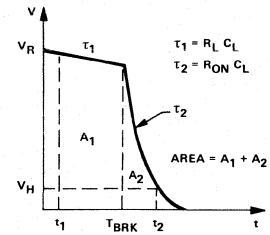

(b) SAMPLE/HOLD

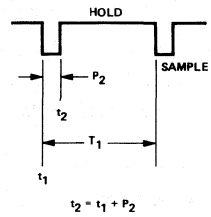

(c)

EQUATIONS:

1. $\frac{V_O}{V_R} \equiv N_O = \frac{N_H(T_1 - P_2) + S_1 + S_2}{T_1}$; Where

2. $\frac{V_H}{V_R} \equiv N_H = EXP\left[\frac{-t}{\tau_1}\right]$, $t \leq T_{BRK}$

$= EXP\left[\frac{-T_{BRK}}{\tau_1}\right] EXP\left[\frac{T_{BRK} - t}{\tau_2}\right]$, $t \geq T_{BRK}$

3. $\frac{A_1}{V_R} \equiv S_1 = \tau_1\left[EXP\left(\frac{-t_1}{\tau_1}\right) - EXP\left(\frac{-T_{BRK}}{\tau_1}\right)\right]$

4. $\frac{A_2}{V_R} \equiv S_2 = \tau_2\, EXP\left[\frac{-T_{BRK}}{\tau_1}\right]\left[1 - EXP\left(\frac{(T_{BRK} - t_2)}{\tau_2}\right)\right]$

Figure 12. Predicting the Measurement System Response

In order to get good correlation between lab data an theoretical predictions, it was necessary to use fairly lon time constants ($R_L = 22k\Omega$ and $C_L = 1000pF$). With $R_L = 22k$ and $C_L = 50pF$ ($R_{ON} = 300\Omega$), the theoretical plot of ACCT (a measured on the spectrum analyzer) vs. t_1 is shown i Figure 13. Note that the data is plotted between 900nse and 1025nsec. The curve shows that a **10nsec error in t_1 ca cause a 6dB error** in reading on the spectrum analyzer. Th results shown in Figure 14 confirm the necessity of usin large capacitances to obtain predictable results. The theo retical curve tracks the actual data well in both cases however the 1000pF curve is better than the 300pF curve Notice that there is good agreement both at DC and at 4kHz

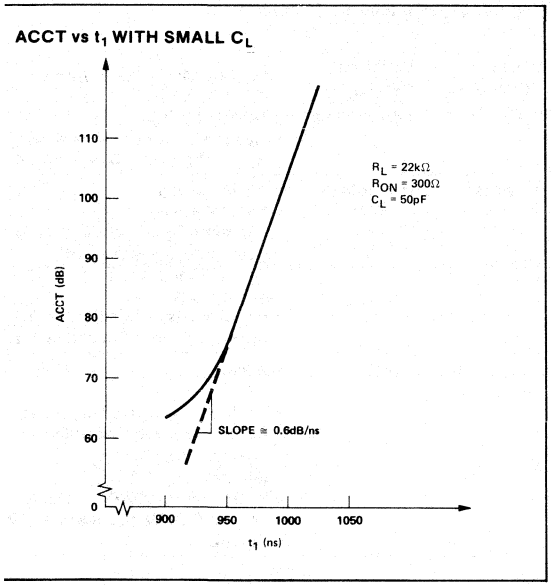

Figure 13. Measurement Errors Due To Small C_L

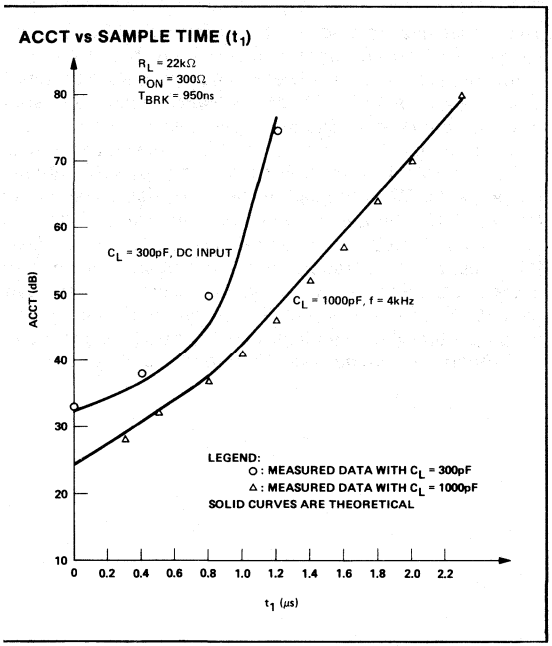

Figure 14. Agreement Between Measured and Computed ACCT

PREDICTING AND CONTROLLING ADJACENT CHANNEL CROSSTALK

The equations in Figure 12c can be used to predict how much adjacent channel crosstalk one might expect in an actual system. An all analog system will follow the MUX with a

A. Multiplexer-Demultiplexer System:

$N_H \equiv 0$ Therefore

1. $N_O = \dfrac{S_1 + S_2}{T_1}$, Where $T_1 = \dfrac{1}{f_{CLK}}$ x (No. of Channels)

2. $S_1 = \tau_1 \left[EXP\left(\dfrac{-t_1}{\tau_1}\right) - EXP\left(\dfrac{-T_{BRK}}{\tau_1}\right) \right]$

3. $S_2 = \tau_2 \ EXP\left[\dfrac{T_{BRK}}{\tau_1}\right]\left[1 - EXP\left(\dfrac{T_{BRK} - t_2}{\tau_2}\right)\right]$

Where $t_1 = T_D$ (Break-Before-Make Time of DEMUX)

$$t_2 = \frac{1}{f_{CLK}} - T_D$$

B. Multiplexer — Sample/Hold System

$S_1 = S_2 = P_2 \equiv 0$

4. $N_O = N_H = EXP\left[\dfrac{-t}{\tau_1}\right] \quad t \le T_{BRK}$

$$= EXP\left[\frac{-T_{BRK}}{\tau_1}\right] EXP\left[\frac{T_{BRK} - t}{\tau_2}\right], \quad t \ge T_{BRK}$$

Where: $t = t_H$ (Hold Command for Sample/Hold as measured from Address Change Time)

Figure 15. Predicting Adjacent Channel Crosstalk

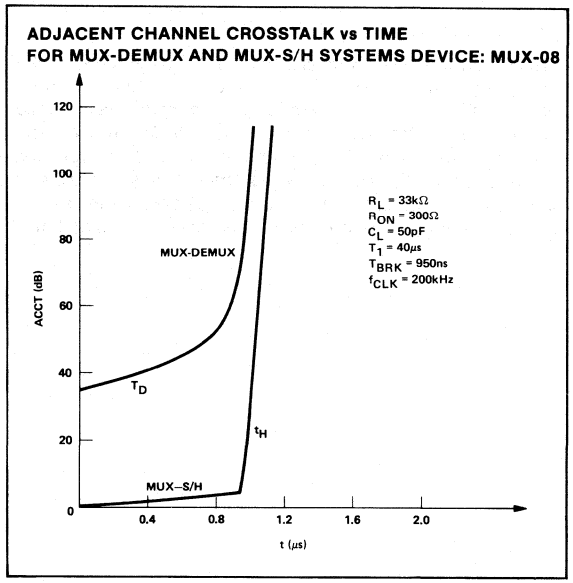

Figure 16. Computed ACCT vs Time for MUX-DEMUX and MUX-S/H Systems

demultiplexer, which will have its own break-before-make delay. An analog to digital system will have a sample/hold amplifier in front of the A/D converter. Since the equations which apply to these situations are different, they will be discussed separately. Figure 15 summarizes the conditions and the equations which apply to them.

Since there is no held voltage, then $N_H = 0$ in the multiplexer-demultiplexer system. This reduces N_O to the simple form shown in equation (1). S_1 and S_2 follow in equations (2) and (3). Since $t_1 = T_D$ (break-before-make time of the DEMUX), that time will have a significant effect on ACCT. The MUX-sample/hold system imposes the condition $S_1 = S_2 = P_2 = 0$; thus $N_O = N_H$. It will be instructive to compare the levels of ACCT in these two systems versus their appropriate times.

Figure 16 looks at a "typical" system which will give approximately one percent transmission error (33kΩ R_L and 300Ω R_{ON}), and has 50pF C_L. The value of C_L is somewhat on the high side (20pF being typical for MUX-08 connected to a buffer amp), but it does give a conservative value for analysis. What Figure 16 shows is rather startling. The adjacent channel crosstalk, while inherent in the multiplexer itself, can be eliminated in **both** systems by the proper timing. In the case of the sample/hold it is only necessary to delay the hold command for approximately 1.2μsec to have the ACCT vanish completely. This is no problem, since most sample/holds need at least 2μsec to accurately acquire the signal (this is particularly true of monolithic devices). The plot for the MUX-DEMUX system relates to T_D, which is **not adjustable** for a given DEMUX. What is possible is to add some delay to the address change for the DEMUX. In this way, the DEMUX will not "look" at the MUX output until the charge from the previous channel has had a chance to dissipate.

CONCLUSION

Table II summarizes the forms of crosstalk and lists ways of coping with them. Reduction of R_{ON} is helpful in all three cases. While T_{BRK} should be minimized as much as possible, it is important that no two channels are ON at the same time. In some cases, T_{BRK} is chosen such that even over temperature extremes, the break-before-make feature is maintained. Since all three components of crosstalk are present in a dynamic multiplexer, the "careful circuit board layout" is important even though it is not listed in the ACC section.

Table 2. How to Handle Crosstalk

Crosstalk Component	Variation with f_{SIG}	Ways to Minimize Effects
Static	**6dB/octave**	• **Minimize R_{ON}** • **Reduce stray capacitance (C_{EQ}) by careful circuit board layout.**
Dynamic	**6dB/octave**	• **Minimize R_{ON}** • **Minimize f_{CLK}** • **Minimize T_{BRK}, but $T_{BRK} > 0$ is needed to prevent shorting channels together.** • **Minimize R_L** • **Reduce stray capacitance (C_{EQ} by careful circuit board layout.**
Adjacent Channel	**NONE**	• **Minimize R_{ON}** • **Minimize f_{CLK}** • **Minimize T_{BRK}, but $T_{BRK} > 0$ is needed to prevent shorting channels together.** • **Minimize R_L and C_L** • **WAIT before allowing sample/hold or DEMUX to measure MU output.**

This paper has pointed out the fact that static crosstal (given on multiplexer data sheets) is only **one** of the **thre** components of crosstalk. The models for static and dynami crosstalk are relatively simple and were discussed to sho how they are related. The most troublesome component c crosstalk (adjacent channel crosstalk) was shown not to b quite so straight-forward. For one thing, adjacent chann crosstalk (ACCT) is **not signal frequency dependent** as ar CT and DCT. The mechanism which governs this form c crosstalk is stored charge on the MUX node. While CT an DCT must be minimized by careful layout and once presen in the multiplexer cannot be reduced, such is not the cas with ACCT. Even though ACCT is present in the multiplexe the proper timing of demultiplexer or sample/hold command can effectively eliminate ACCT from the total system.

ANALOG DEVICES

AN-250
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Bandwidth, OFF Isolation and Crosstalk Performance of the ADG5XXA Multiplexer Series

by Dan Sheehan and Matt Smith

The ADG5XXA* multiplexer series is a family of single 8/16 channel and dual 4/8 channel parts. They are available in both latched (ADG52XA) and unlatched ADG50XA) versions. These are high performance multiplexers that offer the following notable features: specifications for both single and dual supply operation, 1 nA max leakage current, 200 ns max turn-on and turn-off times and TTL compatibility with reduced single or dual supplies down to 5 V.

Multiplexers are widely used on the front end of data acquisition systems where there is an ever increasing requirement for higher accuracy. This requirement, coupled with the availability of increased resolution A/D converters that are capable of digitizing higher frequency signals, has dramatically heightened the need for a good understanding of the ac performance characteristics of multiplexers. Also, to be able to apply the multiplexers in applications such as RF, radar and video switching, the ac performance characteristics must be understood. The ac parameters of interest in a multiplexer are its varying impedance characteristics versus frequency and are principally bandwidth, off isolation and crosstalk.

Although the multiplexers of the ADG5XXA series are low cost, general purpose parts, their ability to switch high frequency signals is excellent. The parts are fabricated in Linear Compatible CMOS (LC^{2}MOS), an advanced process that features very low parasitic capactances and, consequently, improved bandwidth, off isolation and crosstalk performance.

This application note investigates the bandwidth, off isolation and crosstalk of the ADG5XXA series from a practical engineering viewpoint by presenting typical application circuits with results, over the dc to 20 MHz frequency range. The emphasis is very much on practicalities and so the results are stated rather than derived, though it is intended that the reader will gain some insight into the mechanisms which affect the high frequency performance of the multiplexers.

ADG506A/ADG507A, ADG508A/ADG509A, ADG526A/ADG527A, ADG528A/ADG529A.

MULTIPLEXER EQUIVALENT CIRCUIT

A simplified ac equivalent circuit for a pair of adjacent switches in a multiplexer is shown in Figure 1. For high frequency applications it is essential that the parasitic elements shown in the model are taken into account.

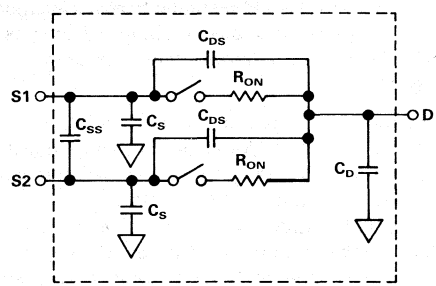

Figure 1. Equivalent Circuit for Adjacent Switches

Element	Title
C_{DS}	Drain-Source Capacitance
C_S	Source Capacitance to GND
C_D	Drain Capacitance to GND
C_{SS}	Capacitance Between Sources
R_{ON}	Channel ON Resistance

Table I. Parasitic Elements

The equivalent circuit simplifies further depending on which switch parameter (bandwidth, crosstalk or isolation) we are concerned with. These parameters will now be discussed separately.

BANDWIDTH

The 3 dB bandwidth frequency provides a measure of the high frequency usefulness of the ADG5XXA series in the ON state.

From an ac or bandwidth perspective, the ON state equivalent circuit of a multiplexer channel can be simplified to a series resistance and shunt capacitance model, as shown in Figure 2.

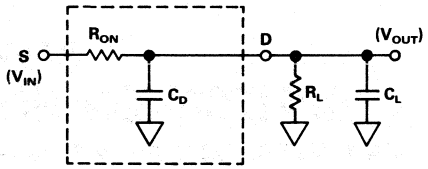

Figure 2. ON Channel Equivalent Circuit

The transfer function (V_{OUT}/V_{IN}) for the ON channel is:

$$\frac{V_{OUT}}{V_{IN}} = \frac{\frac{R_L}{R_L + R_{ON}}}{1 + j2\pi f\left[\frac{R_L \cdot R_{ON}}{R_L + R_{ON}}\right][C_D + C_L]} \quad (1)$$

The 3 dB break frequency f_o, is:

$$f_O = \frac{1}{2\pi\left[\frac{R_L \times R_{ON}}{R_L + R_{ON}}\right][C_D + C_L]} \quad (2)$$

Therefore, for dc or low signal frequencies (<f_o), the multiplexer channel functions as a simple resistive conduction path (R_{ON} = 180 Ω typ for all parts), and any ON insertion loss (attenuation of V_{IN}) is limited to the ratio of $R_L/(R_{ON}+R_L)$. It is with higher signal frequencies that the complex ac impedance shown in Figure 1 becomes significant, resulting in additional attenuation of the analog input signal.

From Equation 2 it is clear that the bandwidth is limited by R_{ON}, R_L, C_L and C_D . Normally $R_L >> R_{ON}$ in order to minimize the insertion loss, so the bandwidth is limited mainly by R_{ON}, C_D and C_L The ADG5XXA series of multiplexers exhibit extremely low parasitic drain capacitance C_D. To take advantage of this low C_D, the external load capacitance should be kept as low as possible. Stray capacitance at the MUX output contributes to C_L, so it is important that this be minimized by proper circuit board layout with signal line lengths as short as possible. Bandwidths in excess of 40 MHz may be achieved with reasonable care.The ON resistance R_{ON} may be minimized by operating with ±15 V power supplies and low input signal levels. (See ADG5XXA data sheets for further information on R_{ON} variation.) The frequency response roll-off rate is typically 20 dB/decade, i.e., single pole roll-off rate.

Figure 3 shows the test circuit used to evaluate the bandwidth performance of the ADG508A. Similar type test circuits apply for the other parts also.

Figure 4 shows a plot of bandwidth vs. load capacitance for the ADG508A and ADG528A. The same results apply for all the input channels and over all package types. From the plots it can be seen that C_L is critical to the bandwidth performance. Nevertheless, 12 pF yields excellent results, with f_O >40 MHz.

Similar results would apply for the ADG509A and ADG529A, but the break frequency (f_O) would be slightly higher due to the even lower C_D of these parts.

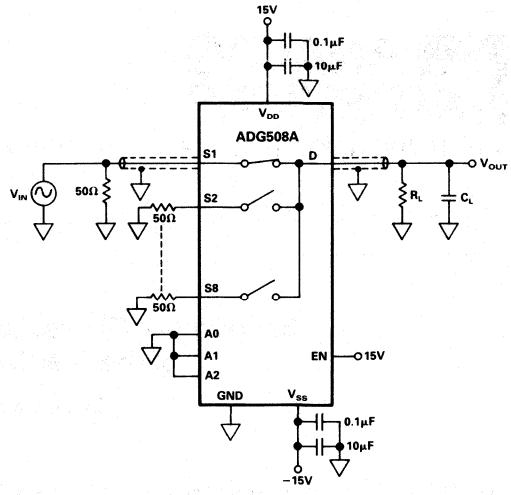

Figure 3. ADG508A Bandwidth Test Circuit

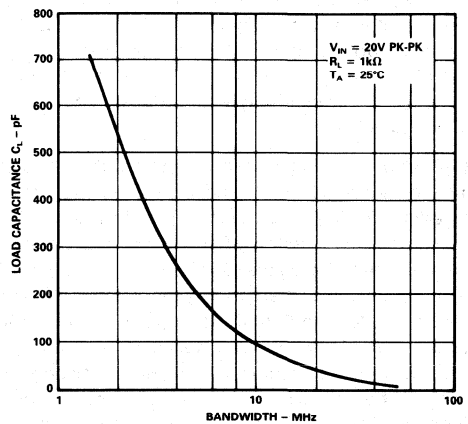

Figure 4. ADG508A and ADG528A Bandwidth Results

The 3 dB break frequency of the ADG506A and ADG526A, with an R_L and C_L of 1 kΩ and 12 pF respectively, is typically 30 MHz. Again, similar results apply for the ADG507A and ADG527A, but with a slightly higher f_O.

OFF ISOLATION

Although the bandwidth performance (ON state) is very important, most applications involving the switching of high frequency signals are frequency limited due to reduced OFF isolation rather than degraded ON performance. OFF isolation is a measure of the ability of the multiplexer to block ac signals in the OFF state and is defined by the following equation:

$$OFF\ Isolation^*(dB) = 20\ log\ (V_{IN}/V_{OUT}).$$

*Due to popular convention, OFF isolation and (crosstalk) values are positive number.

rom an OFF isolation point of view, the multiplexer hannel (or switch) can be modelled by the circuit hown in Figure 5. C_{EQ}, which allows signal coupling ·om the source (S) to the drain (D), determines the OFF npedance presented by the multiplexer to ac signals. 'herefore, as the signal frequency increases, C_{EQ} passes nore signal to the output. C_{EQ} is a combination of the)FF state parasitic capacitance of the relevant multi- lexer channel and external stray capacitances due to viring and circuit board layout.

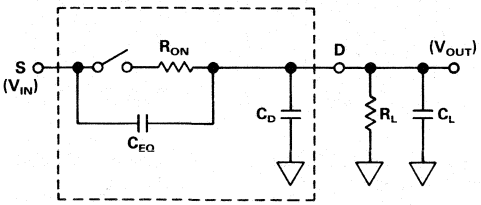

Figure 5. OFF Channel Equivalent Circuit

'he OFF state parasitic capacitance of the multiplexer hannel depends on a number of factors. These factors re principally: fabrication process (LC²MOS), pin con- iguration and package type (plastic DIP, Cerdip, PLCC). t has already been noted that the LC²MOS process fea- ures very low parasitic capacitances. The optimum pin onfiguration minimizes the parasitic capacitance by eeping the the analog input and output pins as far apart s possible. The effect of pin configuration and package ype is highlighted in the results.

Minimizing the stray capacitance, to achieve lower val- es of C_{EQ}, is critical in achieving the optimum OFF solation specifications. Coaxial cable should be used to ransmit the analog signals and a good circuit board ayout is vital. The circuit board should use short signal racks with guard traces between them, ground planes nd bypassed power supplies.

n addition to C_{EQ}, the OFF isolation is also determined o a large degree by C_D, C_L and R_L. The lower the value f R_L, the better the OFF isolation performance (C_{EQ}'s eactance becomes less effective), but note, a low value f R_L gives increased on-insertion loss (i.e., ratio of $R_L/[R_{ON} + R_L]$). Obviously, the choice of the R_L value nvolves tradeoffs to suit a particular application. Large alues of C_D and C_L improve the OFF isolation perfor- nance but again this leads to reduced bandwidth in the)N state.

'he results which follow apply for an optimized circuit oard layout and illustrate how some of the factors out- ined above influence the OFF isolation performance. A areless circuit board layout will give degraded results. igure 6 shows one of the test circuits used to measure he OFF Isolation of the ADG508A (similar test circuits or the other parts).

ADG508A Series* Results

'ypical OFF Isolation plots for the ADG508A and ADG528A are shown in Figures 7, 8, 9 and 10.

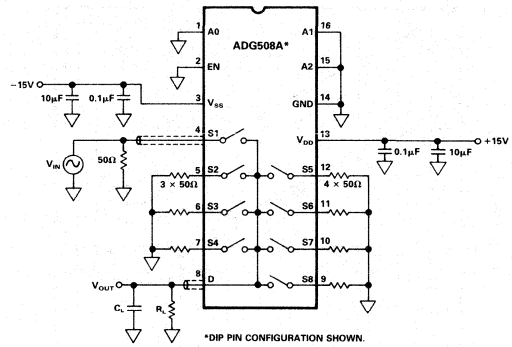

Figure 6. ADG508A OFF Isolation Test Circuit

The slope of most of the OFF isolation plots changes versus frequency, and this is very apparent from Figure 10 for different values of R_L. At low frequencies the fall off rate is typically 20 dB per decade. However, the OFF channel equivalent circuit has a pole frequency at $f_P = 1/2\pi\ (C_{EQ}+C_L+C_D)\ R_L$. Therefore, as the signal frequency approaches and then becomes greater than the pole frequency, the fall off rate decreases and flattens out.

The OFF isolation performance varies with package type due to the different lead frames used. Overall, the plastic leaded chip carrier (PLCC) gives the best results, followed by the plastic DIP and then the cerdip. The PLCC performs best because it has the smallest lead frame, giving the least parasitic capacitance between the input and output pins.

For both the plastic DIP and PLCC packages (No Connect pins are grounded), the input channel with the worst case OFF isolation is S4 (adjacent the output, D), followed by S3 with the best performance from S1 and S5 (same results). For the cerdip package, the worst case channel is again S4, but it is followed closely by S8 and then S3, with the best results from S1 and S5.The difference between this package and the plastic DIP is due to a different lead frame configuration. The plots also show the OFF isolation performance for all the input channels (S1 to S8) driven simultaneously.

From the plots for the ADG508AKN (see Figure 8), it can be seen that S4 couples more unwanted signal into the output than the remaining seven channels put together. Therefore, by sacrificing the use of Channel 4, the user can obtain excellent OFF isolation results from the remaining seven channels. Alternatively, to achieve even higher isolation from the seven channels, S4 may be utilized as a NORMALLY ON grounded channel. With this scheme, S4 is selected when all the other channels are off. This effectively grounds the output (via R_{ON}) giving improved OFF isolation results which are independent of R_L. Hence large values of R_L may be selected which minimize the ON insertion loss without degrading the isolation. The improvement offered with this scheme is included in Figure 10.

Similar results apply for the ADG509A and ADG529A, except at high frequencies where the OFF Isolation is marginally worse due to the lower C_D of these parts.

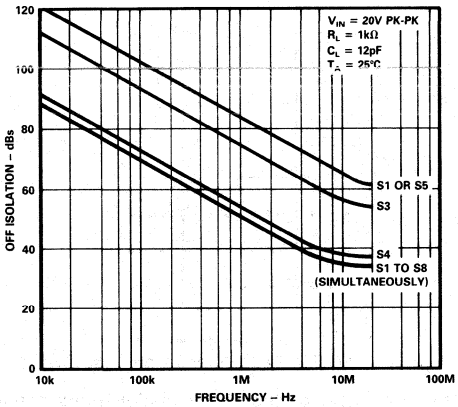

Figure 7. ADG508AKP and ADG528AKP (PLCCs) OFF Isolation vs. Frequency

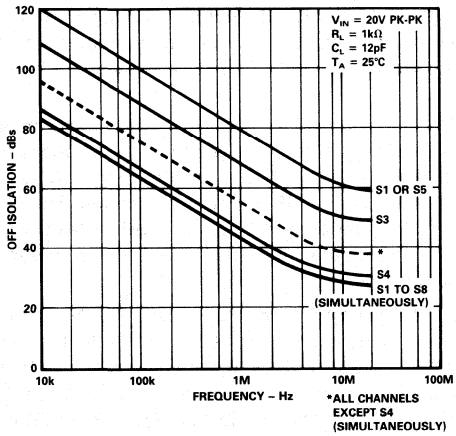

Figure 8. ADG508AKN and ADG528AKN (Plastic DIPs) OFF Isolation vs. Frequency

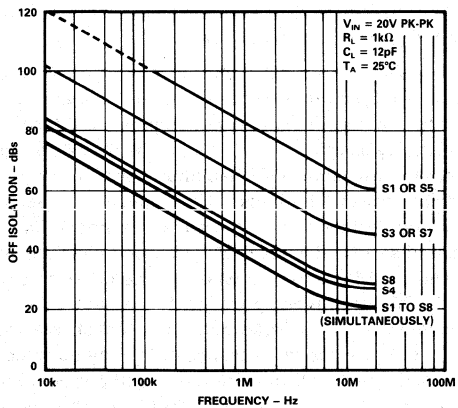

Figure 9. ADG508ABQ and ADG528ABQ (Cerdips) OFF Isolation vs. Frequency

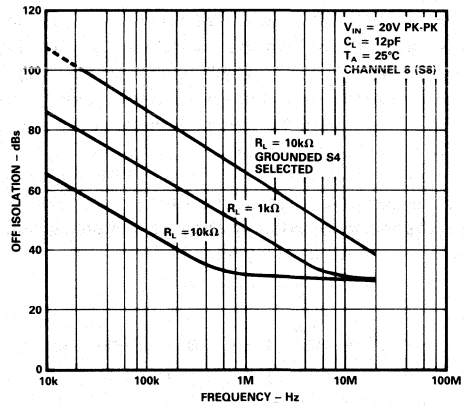

Figure 10. ADG508ABQ and ADG528ABQ (Cerdips) OFF Isolation vs. Frequency for Different R_L

ADG506A* Series Results

Typical OFF isolation plots for the ADG506A and ADG526A are shown in Figures 11 and 12. The ADG507A and ADG527A would yield similar results except at the high frequency end where the results would be marginally worse due to the lower C_D of these parts.

Since the ADG506A series has no analog input channel on a pin which is adjacent to the output, the results for this series surpass those for the ADG508A series. For both the DIP and PLCC packages (N.C. pins connected to GND), the input channel with the worst case OFF isolation is S8, followed by S7, with the best performance from S1 and S9.

In general, comments concerning the ADG508A series results apply for the ADG506A series also.

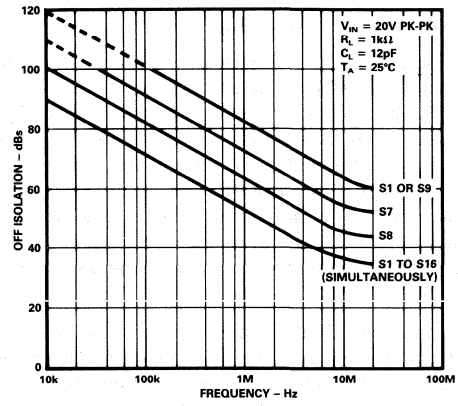

Figure 11. ADG506AKN and ADG526AKN (Plastic DIPs) OFF Isolation vs. Frequency

*ADG506A/507A, ADG526A/527A.

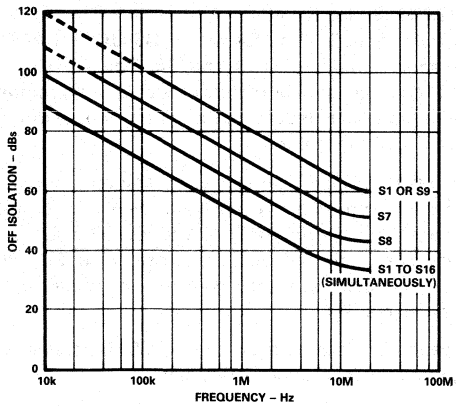

Figure 12. ADG506ABQ and ADG526ABQ (Cerdips) OFF Isolation vs. Frequency

CROSSTALK

Crosstalk, which is also called Channel-to-Channel Crosstalk, is another frequency-dependent factor. It is basically channel-to-channel isolation and occurs due to the OFF state parasitic capacitances of the multiplexer. Figure 13 shows a simple crosstalk model for a two-channel multiplexer. Assuming negligible ON insertion

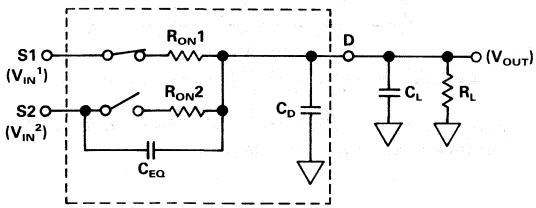

Figure 13. Crosstalk Equivalent Circuit

loss due to $R_{ON}1$ and R_L, V_{OUT} should equal $V_{IN}1$. However, C_{EQ} (see OFF isolation section for details) couples some portion of $V_{IN}2$ into V_{OUT}. Also, but to a lesser degree, additional signal coupling (not shown in Figure 13) occurs between the switches due to parasitic capacitances within the die and between the package leads. With $V_{IN}1$ connected to GND, crosstalk for the circuit shown in Figure 13 is defined by the following equation:

$$Crosstalk\ (dBs) = 20\ log\ (V_{IN}2/V_{OUT}).$$

The error signal couples into a parallel combination of R_L, $R_{ON}1$, C_D and C_L versus R_L, C_D and C_L only for OFF isolation. Because R_{ON} is typically 180 Ω, crosstalk results are generally 10 dB better than OFF isolation results.

While maintaining good bandwidth and OFF isolation performance, crosstalk can be minimized by choosing a multiplexer with a low R_{ON}. A good circuit board layout that shields the analog signals and keeps stray capacitances to a minimum is also vital.

Figures 14 and 15 show two of the test circuits used to measure the crosstalk of the ADG508A (similar test circuits used for the other parts).

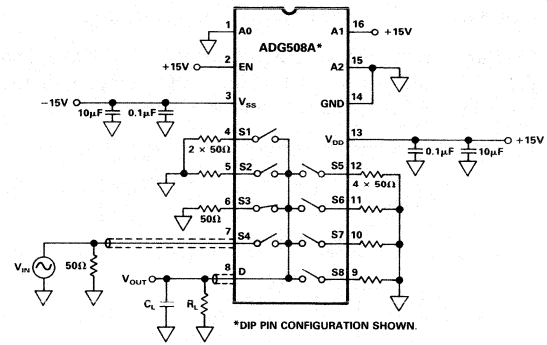

Figure 14. ADG508A Crosstalk Test Circuit (One OFF Channel Driven Only)

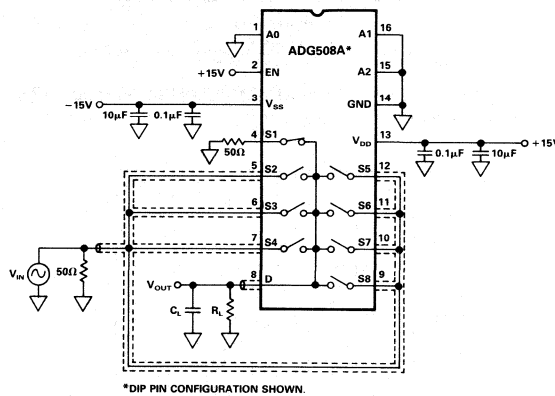

Figure 15. ADG508A Crosstalk Test Circuit (7 OFF Channels Driven)

ADG508A Series Results

Typical crosstalk plots for the ADG508A and ADG528A are shown in Figures 16, 17 and 18.

In Figure 14, the crosstalk is measured by driving S4 only, with S3 selected and the remaining OFF channels connected to GND via 50 Ω. Due to the proximity of S4 to the output pin, driving this channel gives the worst case crosstalk results (for one OFF channel driven only) – see Figure 16 Plot b. For example, driving S3 only with S4 selected would give far superior results – typically 15 dB better in the case of the ADG508AKN and ADG528AKN.

Crosstalk can also be measured by driving all of the remaining OFF channels simultaneously, with any one channel selected as in Figure 15. Under these conditions the worst case performance (see Figure 16 Plot a) occurs when S1 or S5 is selected and the best performance (see Plot c) results when S4 is selected.

The plastic DIP package gives better results than the cerdip package (similar to OFF isolation). The results for PLCC are not shown but would be approximately 5 dB better than for plastic DIP. Figure 18 shows that the crosstalk performance improves as R_L reduces.

Similar results would apply for the ADG509A and ADG529A. However, at the high frequency end, the results would be marginally worse due to their lower C_D.

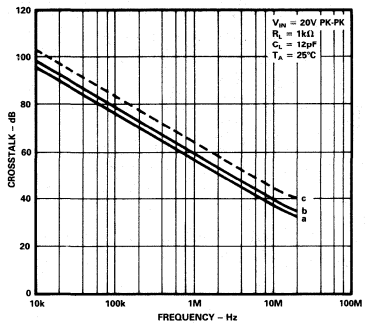

a. ALL CHANNELS DRIVEN EXCEPT CHANNEL 1, WHICH IS "ON."
b. S4 DRIVEN; S3 "ON"; REMAINING CHANNELS CONNECTED TO GND (SEE FIGURE 14).
c. ALL CHANNELS DRIVEN EXCEPT CHANNEL 4, WHICH IS "ON."

Figure 16. ADG508AKN and ADG528AKN (Plastic DIPs) Crosstalk Results vs. Frequency

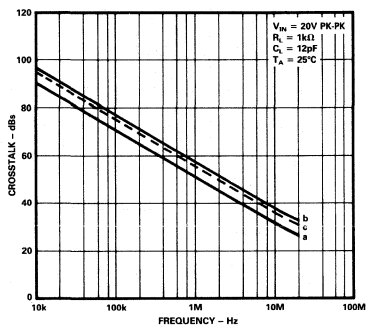

a. ALL CHANNELS DRIVEN EXCEPT CHANNEL 1, WHICH IS "ON."
b. S4 DRIVEN; S3 "ON"; REMAINING CHANNELS CONNECTED TO GND (SEE FIGURE 14).
c. ALL CHANNELS DRIVEN EXCEPT CHANNEL 4, WHICH IS "ON."

Figure 17. ADG508ABQ and ADG528ABQ (Cerdips) Crosstak Results vs. Frequency

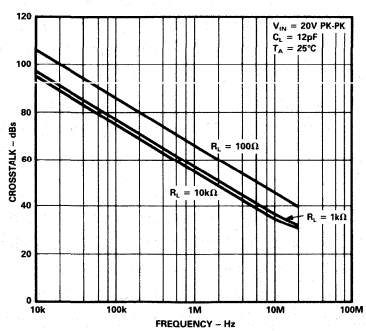

ALL CHANNELS DRIVEN EXCEPT CHANNEL 1, WHICH IS "ON."

Figure 18. ADG508AKN and ADG528AKN (Plastic DIPs) Crosstalk Results vs. Frequency for Different R_L

ADG506A Series Results

Typical crosstalk plots for the ADG506A and ADG526A are shown in Figures 19 and 20. The ADG507A and ADG527A yield similar results except at high frequencies where the results would be marginally worse due to the lower C_D of these parts.

Since the ADG506A series has no analog input channel and output on adjacent pins, the results for this series surpass those for the ADG508A series. S8 is the channel nearest to the output, and therefore driving S8 with any other channel selected contributes most to poor crosstalk results.

In general, comments on the ADG508A series results apply for the ADG506A series also.

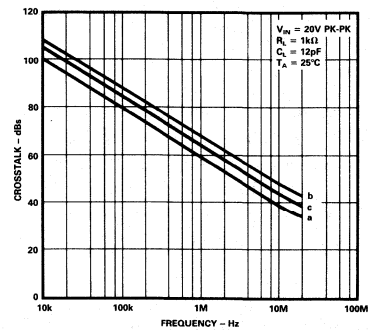

a. ALL CHANNELS DRIVEN EXCEPT CHANNEL 1, WHICH IS "ON."
b. S8 DRIVEN; S7 "ON"; REMAINING CHANNELS CONNECTED TO GND (SIMILAR TO FIGURE 14).
c. ALL CHANNELS DRIVEN EXCEPT CHANNEL 8, WHICH IS "ON."

Figure 19. ADG506AKN and ADG526AKN (Plastic DIPs Crosstalk Results vs. Frequency

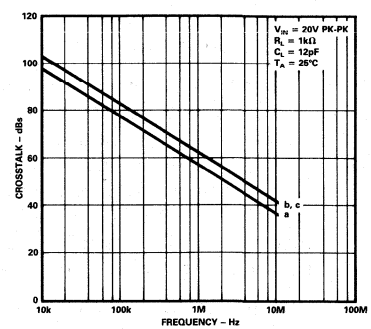

a. ALL CHANNELS DRIVEN EXCEPT CHANNEL 1, WHICH IS "ON."
b. S8 DRIVEN; S7 "ON"; REMAINING CHANNELS CONNECTED TO GND (SIMILAR TO FIGURE 14).
c. ALL CHANNELS DRIVEN EXCEPT CHANNEL 8, WHICH IS "ON."

Figure 20. ADG506ABQ and ADG526ABQ (Cerdips) Crosstalk Results vs. Frequency

COMMENTS

The results furnished in this application note are based on measurements from three fabrication lots. With the same conditions and device type, the results varied very little from lot to lot and from device to device. For example, the results for both OFF isolation and crosstalk varied by typically 1 dB only. For bandwidth, the variation in results was less than 0.2 MHz.

ANALOG DEVICES

AN-251
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

R_{ON} Modulation in CMOS Switches and Multiplexers; What It Is and How to Predict Its Effect on Signal Distortion

by John Wynne

A single CMOS switch or a single channel of a CMOS multiplexer essentially consists of an N-channel and a P-channel MOSFET transistor in parallel, see Figure 1a. The respective drains and sources of the two transistors are tied together to become the switch terminals while the gates of the two transistors are usually driven with the power supply voltages, V_{DD} and V_{SS}, to control the on-off action of the switch. Essentially the N-channel is ON for positive gate-to-source voltages and OFF for negative gate-to-source voltages (vice versa for the P-channel).

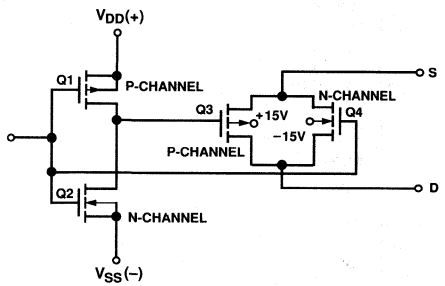

Figure 1a. Basic CMOS Switch

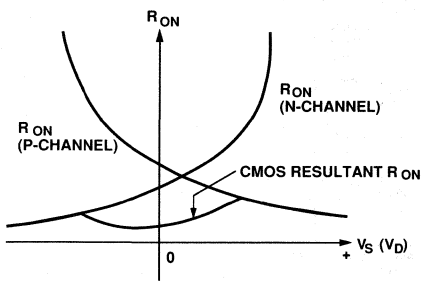

Figure 1b. Individual MOSFET R_{ON} Profiles vs. V_S (V_D)

With a fixed voltage on the gate, the effective drive voltage for either transistor varies in proportion to the polarity and magnitude of the signal passing through the switch. In Figure 1b where R_{ON} is plotted against applied switch voltage V_S (V_D), the resistance of the N-channel increases with positive voltage and the resistance of the P-channel increases with negative voltage. The resultant parallel combination (heavy line) exhibits the well-known "crown" or twin-peak characteristic. This variation in on-channel resistance with input signal is known as R_{ON} modulation.

Figures 2 and 3 show some typical R_{ON} and ΔR_{ON} profiles for Analog Devices' ADG5XXA multiplexer series. Figure 2 shows R_{ON} for three different power supply voltages. Figure 3 shows an expanded view of the change in resistance (ΔR_{ON}) under the same conditions. Note that both R_{ON} and ΔR_{ON} increase as the power supplies are reduced.

12

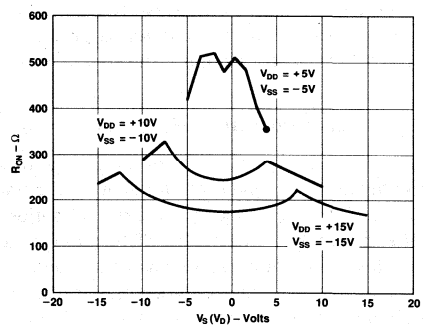

Figure 2. R_{ON} as a Function of V_S (V_D), T_A = +25°C

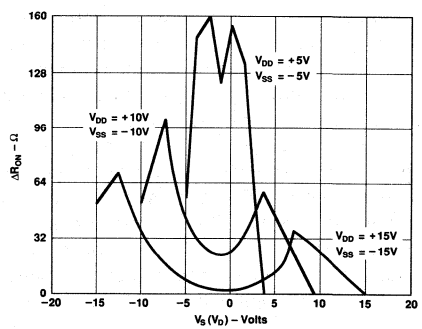

Figure 3. ΔR_{ON} as a Function of V_S (V_D), T_A = +25°C

R_{ON} modulation can be reduced by restricting the input signal range through the switch. For instance, for the ADG5XXA multiplexer series operating on ± 15 V power supplies, ΔR_{ON} is less than 4 Ω over an input signal range of ± 3 V, increasing to 12 Ω with a ± 5 V input signal range and increasing again to over 30 Ω with a ± 7 V input signal range. Configuring a switch or multiplexer to operate directly into the virtual earth of an op amp obviously ensures a very low voltage across the channel which in turn virtually eliminates R_{ON} modulation problems. However, many applications require high level signals to be passed through the channel. Figure 4 shows a typical situation where high level signals are multiplexed into a load resistance R_L.

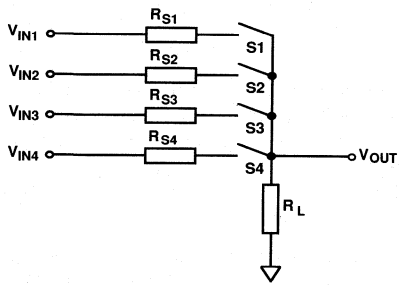

Figure 4. Multiplexing High Level Signals

In addition to the common load resistance, each multiplexer channel has a source resistance R_{SN} in series with it. Over the signal range of interest, the amount of distortion generated by R_{ON} modulation is proportional to the ratio of maximum to minimum channel resistance, i.e.,

$$Distortion \propto \Delta R_{ON}/(R_{EXT} + R_{ON}\ min)$$

The external resistance R_{EXT} is the sum of source and load resistance. For a given multiplexer (i.e., given R_{ON} minimum and ΔR_{ON}) , making the external series resistance large reduces the generated distortion through the channel. However, high resistor values result in high levels of circuit noise, so a compromise must be made when deciding on values for the external resistors.

PREDICT DISTORTION GRAPHICALLY

Figure 5 shows a nomograph which can be used to quickly predict, to a first order, the distortion generated by a single multiplexer channel in series with an external resistor. The nomograph scaling is based upon Figure 4. The left most scale represents the total channel resistance including the switch R_{ON} value at V_S (V_D) = 0 V. The middle scale represents distortion, and the right most scale represents ΔR_{ON}, the change in R_{ON} *over the signal range of interest*. The nomograph is used by drawing a straight line between the appropriate points on the outer scales; where the straight line intersects with the middle scale yields the predicted distortion. As an example of its use, consider a switch with R_{ON} = 400 Ω at V_S (V_D) = 0 V and ΔR_{ON} = 60 Ω over the input signal range. With R_S = 0 Ω and R_L = 4 kΩ, the total harmonic distortion (THD) through the channel is approximately 0.3% or −50 dB. With R_L = 40 kΩ, the THD improves to 0.048% or −66 dB.

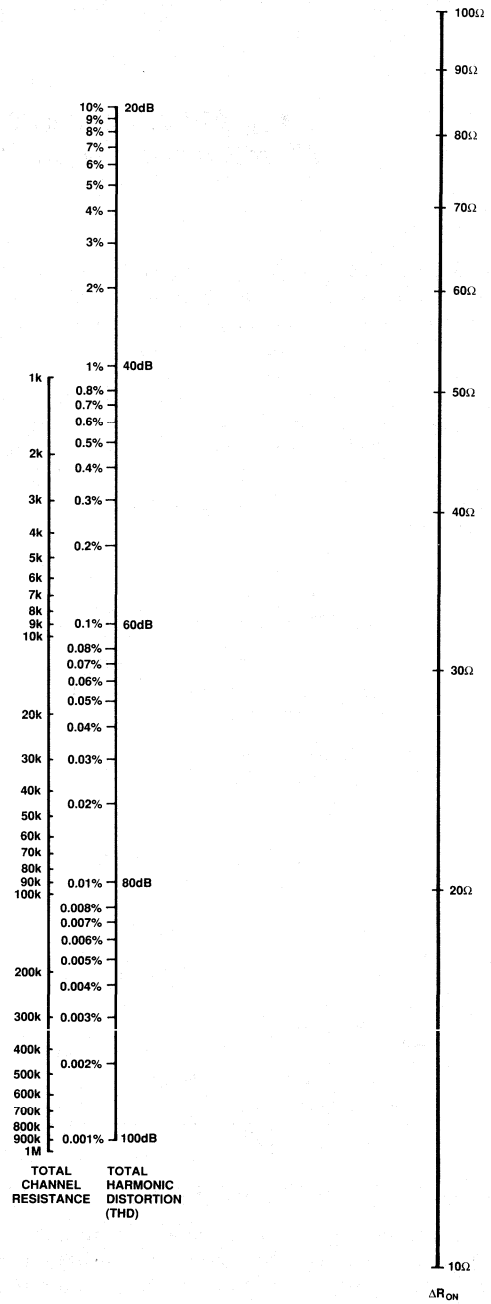

Figure 5. Nomograph to Determine THD Through a Single Switch or Multiplexer Channel

Operational Amplifiers
Contents

		Page
AN-356	User's Guide to Applying and Measuring Operational Amplifier Specifications	13–3
AN-357	Operational Integrators	13–11
AN-202	An I.C. Amplifier User's Guide to Decoupling, Grounding, and Making Things Go Right for a Change	13–17
AN-358	Noise and Operational Amplifier Circuits	13–25
AN-102	Very Low Noise Operational Amplifier	13–35
AN-253	Find Op Amp Noise with Spreadsheet	13–39
AN-107	Active Feedback Improves Amplifier Phase Accuracy	13–41
AN-106	A Collection of Amp Applications	13–49
AN-40	A Buffer Amplifier Applications Collection	13–67
AN-25	Applying the OP-06 Op Amp As a High Precision Comparator	13–71
AB-107	Two-Wire, 4–20 mA Current Transmitter	13–75
AB-108	A Micropower Single-Supply Precision Rectifier	13–77
AB-109	High Speed Precision Rectifier	13–79
AB-111	Single-Supply Wien Bridge Oscillator	13–81
AB-112	Single Resistor Controls Wien Bridge Oscillator Frequency	13–83
AB-113	Precision Ramp Generator	13–85
AB-114	Precision Current Regulator	13–87
AB-115	Micropower 1.23 V Bandgap Reference	13–89
AN-254	Add Programmable Gain, Attenuation	13–91
AN-255	Voltage-Controlled Amp Covers 55 dB Range	13–93
AN-359	Settling Time of Operational Amplifiers	13–95
AN-256	Accurately Testing Op Amp Settling Times	13–105
AN-108	JFET-Input Amps Are Unrivaled for Speed and Accuracy	13–107
AN-257	Careful Design Tames High Speed Op Amps	13–115
AN-110	OP-64 Advanced SPICE Macro-Model	13–125
AN-117	OP-42 Advanced SPICE Macro-Model	13–131
AN-120	OP-400 SPICE Macro-Model	13–137
AN-126	OP-260 Advanced SPICE Macro-Model	13–143
AN-132	OP-470 SPICE Macro-Model	13–149
AN-138	SPICE-Compatible Op Amp Macro-Models	13–155
AN-259	AD9617/AD9618 Current Feedback Amplifier Macro-Models	13–171
AN-360	Ask the Applications Engineer—6: Op Amp Issues	13–175
AN-281	Passive and Active Analog Filtering	13–177

ANALOG DEVICES

AN-356
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

User's Guide to Applying and Measuring Operational Amplifier Specifications

by Ray Stata

SINCE THERE ARE NO established standards for operational amplifier specifications we shall discuss here the terms used by Analog Devices to define operational amplifier characteristics as well as the limitations which must be observed in applying the published data to actual circuits. Wherever possible we show the test circuits used to measure these parameters. Although these test circuits are applicable to a wide range of operational amplifiers, special amplifiers such as FET, chopper stabilized or ultra fast response amplifiers may require changes in the recommended circuit values or in some cases different test methods to measure their specifications. As a general rule the power supply for these measurements should have line and load regulation of about 0.1% and ripple should be no more than a few millivolts.

Figure 1 gives a simplified equivalent circuit for an operational amplifier showing many of the sources of error which are discussed in the text The specifications should be referenced to this diagram to predict their effect in a closed loop circuit. For a single ended amplifier you would assume that the plus or non-inverting input is grounded.

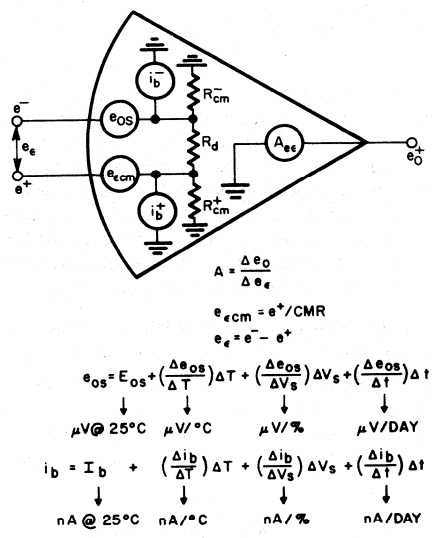

FIGURE 1. OPERATIONAL AMPLIFIER EQUIVALENT CIRCUIT

OPEN LOOP GAIN

Open loop gain, A, is defined as the ratio of output voltage to error voltage e_ε between inputs as shown in figure 1. Gain is usually specified only at DC (A_o), but in many applications such as AC amplifiers the frequency dependence of gain is also important. For this reason the typical open loop gain response is published for each amplifier. The open loop gain response of most amplifiers can be approximated by figure 2.

Reprinted from Analog Dialogue 1–3 1967

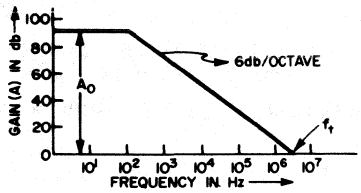

FIGURE 2. TYPICAL OPEN LOOP GAIN RESPONSE

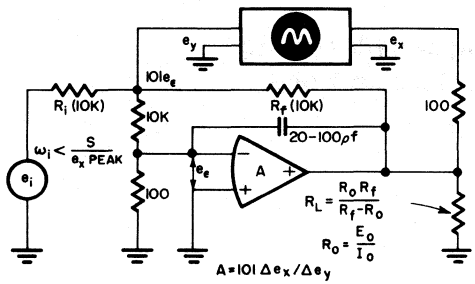

FIGURE 3. OPEN LOOP GAIN TEST CIRCUIT

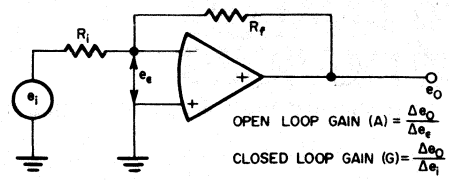

FIGURE 4. CLOSED LOOP CIRCUIT

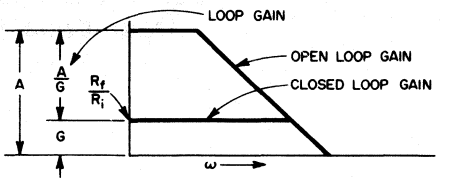

FIGURE 5. DETERMINATION OF LOOP GAIN

Open loop gain changes with load impedance (R_L), ambient temperature and supply voltage. As a rule, open loop gain will not change more than a factor of 10 between rated load and no load conditions. Most operational amplifiers have a positive gain temperature coefficient of about 0.5 to 1%/°C and gain changes with supply voltage at about 2%/%. Analog Devices specifies all open loop gains at rated load, 25°C and rated supply voltages.

A practical circuit for measuring open loop gain over a range of frequencies is shown in figure 3. The voltage divider on the negative input boosts the sensitivity of the error voltage by 100 times which makes it possible to measure gains up to one million. At low frequencies open loop gain is constant so that DC gain can be measured by a low frequency signal (about 5Hz). The voltage divider may not be necessary for low gain amplifiers (below 20,000) and it is not recommended for measuring gain at high frequencies where open loop gain is less.

At very best, noise pick up is a problem for measuring high gains and care must be taken to adequately shield the test circuit. At high frequencies the amplitude of the output voltage must be reduced to avoid exceeding the slewing rate of the amplifier. For this reason the output voltage should be adjusted, so that $e_x(\text{peak}) < \text{slew rate}/\omega_i$, where ω_i is the test frequency.

SIGNIFICANCE OF OPEN LOOP GAIN

Operational amplifiers are rarely used open loop. Instead negative feedback is used around the amplifier to improve the accuracy of the circuit. This introduces a second term, closed loop gain (G), which is defined as the gain of the circuit with feedback. The simple inverting amplifier in Figure 4 illustrates this point.

Linearity, gain stability, output impedance and gain accura are all improved by the amount of feedback. Figure 5 graph cally illustrates the relation between open loop gain and clos loop gain.

The excess of open loop gain over closed loop gain is calle loop gain. (Subtraction of dB is equivalent to arithmetic divi sion.) The improvement of open loop performance due t feedback is directly proportional to loop gain. As a genera rule for moderate accuracy, open loop gain should be 100 time greater than the closed loop gain at the frequency, or frequen cies, of interest (that is loop gain = 100). For higher accurac loop gain should be 1000 or more. To illustrate, we recall tha open loop gain stability for most operational amplifiers is abou 1%/°C. With loop gain of 100, closed loop gain stabilit would be 100 times better or 0.01%/°C. Likewise, closed loo output impedance would be 100 times less than open loo output impedance with a loop gain of 100.

RATED OUTPUT VOLTAGE AND CURRENT

Rated output voltage, E_o, is the maximum peak output voltag which can be obtained at rated output current before clippin or excessive non-linearity occurs. This measurement is made a rated power supply voltage; at other supply voltages the outpu will swing to within about 4 volts of the supply voltage. Als the output voltage swing will increase somewhat at lower loa current. Rated output current, I_o, is the minimum guarantee value of current at the rated output voltage. Load impedanc less than E_o/I_o can be used but E_o will decrease, distortion ma increase and open loop gain will be reduced. Driving larg capacitance loads at high frequencies will present a low loa impedance which may then exceed the rated output current Any convenient circuit such as figure 3 or figure 6 can be use to measure E_o and I_o.

UNITY GAIN SMALL SIGNAL RESPONSE

Unity gain small signal response, f_t, is the frequency at which the open loop gain becomes unity or zero dB (see figure 2). "Small signal" indicates that in general it is not possible to obtain large output voltage swing at high frequencies because of distortion due to slew rate limiting. Therefore in both measuring f_t and using the amplifier at high frequencies, the output voltage swing must be restricted to avoid slew rate limiting. This implies that the peak output voltage, e_o, for a sinusoidal signal at the unity gain frequency, f_t, must be less than $S/2\pi f_t$, where S is the slew rate.

For amplifiers with symmetrical response on each input, f_t may be measured by either the inverting circuit of figure 6 or the non-inverting circuit of the figure 7. Some units such as chopper stabilized amplifiers or wideband amplifiers with feed forward design have fast response only on the negative input which restricts testing and use to the inverting circuit. Remember that the closed loop unity gain response of figure 6 will be about one half the open loop unity gain response due to the loading of the feedback network. Moreover, large values of feedback resistance when coupled with stray capacitance may reduce the closed loop response and therefore the smallest possible value of R_f should be used, the limit being set by output current capability I_o.

Sometimes f_t is called unity gain-bandwidth product which implies that open loop gain at other frequencies can be predicted from this number. However, gain bandwidth product is constant only for amplifiers with 6dB/octave roll off! For fast roll off amplifiers, gain bandwidth product increases with gain and thus we publish the open loop response curve to give typical gain at each frequency.

FULL POWER RESPONSE

The large signal and small signal response characteristics of operational amplifiers differ substantially due to dynamic nonlinearities or transient saturation. An amplifier will not respond to large signal changes as fast as the small signal bandwidth characteristics would predict. The most prominent contributor to large signal response limitations is slew rate limiting in the output stages. Circuit and transistor capacitances can be charged and discharged only so fast due to the limited dynamic range of the driving circuits. Transient saturation can also occur in the input stages of the amplifier due to overloading the input stage or due to common mode voltage slew rate limiting, but this is rarely a problem as compared to saturation of the output stages.

Full power response, f_p, is the maximum frequency measured at unity closed loop gain, for which rated output voltage, $\pm E_o$, can be obtained for a sinusoidal signal at rated load without distortion due to slew rate limiting. Note that this specification does not relate to "response" in the sense of gain reduction with frequency. Instead it refers only to distortion in the output signal caused by slew rate limiting. For a sinusoidal signal, the maximum slope or rate of voltage change occurs at zero crossing and is proportional to the peak amplitude and the frequency. Thus we see that to a first approximation slew rate, S, and full power response, f_p, are related by equation 1.

$$\left.\frac{de_o}{dt}\right|_{max} = 2\pi f_p E_o = S \qquad \text{(equation 1)}$$

As the voltage swing is reduced below rated output, E_o, the operating frequency can be proportionally increased without exceeding the slew rate, S. In the limit the operating frequency approaches the unity gain bandwidth, f_t, and the corresponding voltage signal defines the maximum peak amplitude for "small signal" unity gain response. The circuits of figure 6 or figure 7 can be used to measure full power response depending on whether inverting or non-inverting parameters are measured. Where dynamic saturation of the output stages is the primary cause for slew rate limiting either test circuit will give equivalent results. For very fast response amplifiers, load capacitance and/or capacitance from the output to the negative input will cause apparent slew rate limiting and consequent degradation of full power response. This is due to saturation of amplifier output current in charging these capacitances and therefore such capacitances must be low.

Output distortion can be measured either by a distortion meter on the output or by observing a Lissajon pattern on an oscilloscope. There is no industry wide accepted value for the distortion level which determines the full power response limitation, but a number like 1% to 3% is a reasonable figure. One subtle point here is that closed loop output distortion depends on the amount of feedback or loop gain and therefore it depends on the closed loop gain of the measurement. Full power response is generally measured at unity gain where loop gain is the highest. At higher closed loop gains output distortion will increase for the same full power response frequency.

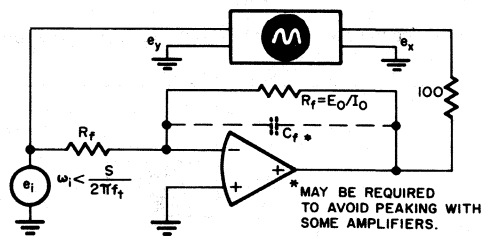

FIGURE 6. INVERTING CIRCUIT FOR MEASURING f_t, f_p, S, T

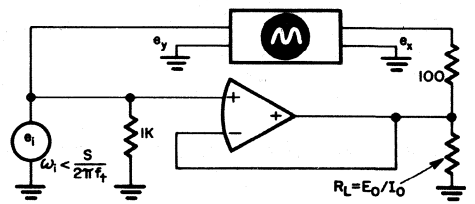

FIGURE 7. NONINVERTING CIRCUIT FOR MEASURING f_t, f_p, S, T, E_{cm}

In many applications the additional distortion which is caused by exceeding the full power response can be comfortably ignored. However, a far more serious effect, often overlooked, is that a DC offset voltage can be generated when the full power response is exceeded due to rectification of the unsymmetrical feedback waveform or due to overloading the input stage with large distortion signals at the summing junction.

These more subtle points in measuring full power response as well as the attendant side effects suggest the circuit of figure 8 as more satisfactory test circuit. By viewing the error voltage at the summing junction on an oscilloscope, distortion signals are more easily detected, signal generator distortion is eliminated from the measurement and frequency dependent DC offset can be readily observed.

SLEWING RATE

The origins of slewing rate limitations were discussed in the previous section. Slewing rate, S, usually expressed in volts/μsec defines the maximum rate of change of output voltage for a large step change.

Equation 1 suggests a convenient method to measure slewing rate by first measuring full power response, f_p, and then calculating S. Although this test method yields usable results for most amplifiers in most applications, the relationship of equation 1 does not apply under all conditions. First, slewing rate is a non-linear function of output voltage and equation 1 measures slewing rate only at zero volts output. This second order effect can usually be safely ignored in most applications. However, for certain amplifiers, particularly fast response types, the slew rate may be higher than that predictable from f_p. In these cases f_p is limited by factors other than slew rate such as DC offset errors which are generated by the rectification of large high frequency error voltages.

A more direct method to measure slewing rate is to apply low frequency square waves (about 100Hz) to the input of figu 6 or figure 7 which cause full voltage swing at the output an to observe the rise time from 10 to 90% on an oscilloscop (see figure 9). Small feedback resistors must be used to avoi degradation of slewing rate due to stray capacitance.

In applying operational amplifiers remember that repetitiv input waveforms whose rise time exceeds the amplifier's slewin rate will generate voltage spikes at the summing junctior These spikes are usually unsymmetrical and are also usuall clipped unsymmetrically by the input circuit of the amplifier – either or both of which effects will cause DC offsets at th output.

OVERLOAD RECOVERY

Overload recovery, τ, defines the time required for the outpu voltage to recover to the rated output voltage E_o from a satu rated condition. For this test the circuit of figure 6 or 7 is use with an input square wave adjusted to be 50% greater than th voltage required to saturate the amplifier output. The squar wave frequency should be adjusted to about 100Hz and th input-output signals should be compared on a dual trace oscillc scope as illustrated in figure 9.

In some amplifiers the overload recovery will increase for larg impedances (greater than 50KΩ) in the input circuit, either th summing impedance for figure 6 or the source input for figur 7. Published specifications apply for low impedances and as sume that overload recovery is not degraded by stray capacitanc in the feedback network.

Normally, overloaded recovery time runs about one milliseconc For the inverting configuration an external clamp circuit can b added to improve overload recovery as illustrated in figure 1C This circuit prevents the output from saturating and therefor circumvents any delays due to overload recovery. The only con straint for proper operation is that input current (e_i/R_i) sha be approximately less than the rated output current I_o minus th load current. The clamp circuit cannot be used with the nor inverting and differential configurations.

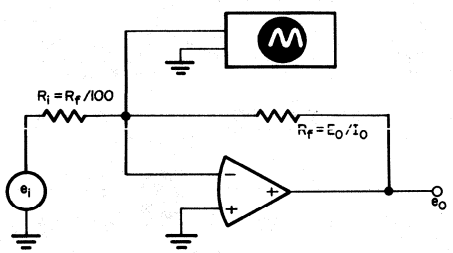

FIGURE 8. MEASURING FULL POWER RESPONSE

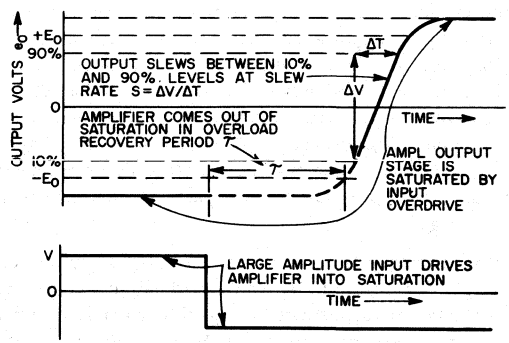

FIGURE 9. OVERLOAD RECOVERY AND SLEW RATE ILLUSTRATION

TIAL OFFSET VOLTAGE

et voltage, e_{os}, is defined as the voltage required at the it from a zero source impedance to zero the output, at any perature, supply voltage and time (see figure 1). Initial et voltage, E_{os}, defines the offset voltage at 25°C and rated ply voltages. In most amplifiers, provisions are made to ust initial offset to zero with an external trim potentiometer. ie amplifiers are internally trimmed to guarantee some imum limit on initial offset (usually ±1mV) which means in certain applications the external trim pot can be elimi-ed. On special order any amplifier from Analog Devices can internally trimmed to ±1mV initial offset or less. Initial et can be measured with the circuit of figure 11, where an ropriate fixed resistor is substituted for the external trim entiometer. There is a warm up drift of offset voltage owing the application of power supply voltage and it is ommended that you let the amplifier stabilize for at least 15 utes before making measurements.

TIAL BIAS CURRENT

s current, i_b, is defined as the current, at any temperature, ply voltage and time, required at either input from an infi-e source impedance to zero the output assuming zero com-n mode voltage. For differential amplifiers bias current is ignated by i_b^- for the negative input and by i_b^+ for the itive input. For single ended amplifiers, like chopper stabi-ed units, bias current refers to the current at the negative ut only.

tial bias current, I_b, is the bias current at either input mea-ed at 25°C, rated supply voltages and zero common mode tage. The designation (0,+) or (0,−) indicates that no ernal compensation is used to reduce initial bias current so it the polarity is always known. The sign tells to which wer supply voltage an external compensating resistor should connected to zero the initial bias current. The designation =) indicates that internal compensation has been used to luce initial bias current and that the residual bias current n be of either polarity. In general compensating initial bias rrent has little effect on the bias current temperature coeffici-t. The circuit of figure 11 is used to measure initial bias rrent.

IITIAL DIFFERENCE CURRENT

ifference current*, i_d, is defined as the difference between the as currents at each input from an infinite source required to ro the output assuming zero common mode voltage. The put circuitry of differential amplifiers is generally symmetri-l so that bias current at each tends to be equal and tends to ack with changes in temperature and supply voltage. Usually fference current is 3 to 5 times less than bias current at either put, assuming that initial bias current is not compensated. the impedance as seen from each input terminal to ground balanced then offset and drift errors are proportional to dif-rence current rather than to bias current. In most applica-ons, if the external impedances at each input are balanced en there is no particular advantage in using an amplifier here initial bias current is internally compensated. Initial dif-rence current, I_d, the difference current measured at 25°C and ted supply voltage, can be measured by the circuit of figure 11.

Previously called offset current.

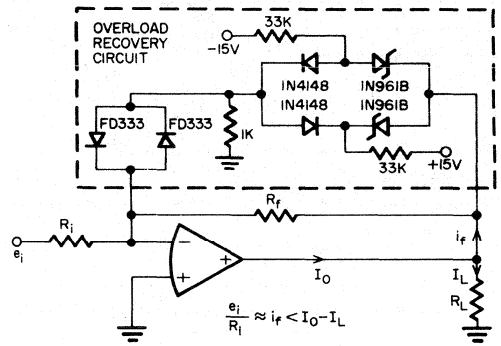

FIGURE 10. EXTERNAL CLAMP CIRCUIT FOR ±10 VOLTS

TEMPERATURE DRIFT

Offset voltage, bias current and difference current all change or "drift" from their initial values with temperature. By far this is the most important source of error in most applications. The temperature coefficients of these parameters, $\triangle e_{os}/\triangle T$, $\triangle i_b/\triangle T$ and $\triangle i_d/\triangle T$ are all defined as the average slope over a specified temperature range and are determined by subtracting the offset values at the end points of the temperature range and dividing by the temperature change. In general drift is a non-linear function of temperature and the slopes are greater at the extremes of temperature than around normal room ambient. The temperature drift coefficients are measured by the circuit of figure 11. The amplifier is used to boost its own low level input offset signal to a conveniently measurable voltage at the output. Gain is established by the ratio of R_f/R_i. The current sampling resistors, R_s, must be selected so that voltage drift is small compared to the drift due to difference current; that is $R_s \times \triangle i_d/\triangle T >> \triangle e_{os}/\triangle T$. Alternatively, voltage drift must be subtracted from the data for current drift.

One problem in using published voltage drift specifications is that this data applies only to static temperature conditions where the temperature of the module is assumed to be uniform. Voltage offset of most differential amplifiers is quite sensitive to thermal gradients, since drift performance depends on the cancellation of large offset in each transistor of the input differential pair. Therefore in environments where thermal gradients are present voltage offset may exceed that predictable from the drift coefficients. In this case where low drift over a narrow temperature range is critical, it is good practice to insulate or shield the amplifier to assure a uniform temperature. Bias current is not noticeably affected by thermal gradients and difference current, while affected, is far less sensitive to gradients than voltage offset.

Bias current and difference current for FET and varactor bridge amplifiers double each 10°C and therefore a linearized drift coefficient has little meaning except over a narrow operating temperature range.

13

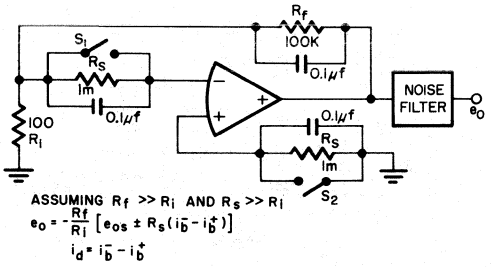

Parameter	Switch Position S1	S2	Output Voltage
e_{os}	Closed	Closed	$e_o = -(\frac{R_f}{R_i})e_{os}$
i_b^-	Open	Closed	$e_o = -(\frac{R_f}{R_i})R_s i_b^-$
i_b^+	Closed	Open	$e_o = -(\frac{R_f}{R_i})R_s i_b^+$
i_d	Open	Open	$e_o = -(\frac{R_f}{R_i})R_s i_d$

FIGURE 11. TEST CIRCUIT FOR OFFSET VOLTAGE, BIAS CURRENT AND DIFFERENCE CURRENT

SUPPLY VOLTAGE SENSITIVITY

Offset voltage, bias current and difference current will also change when supply voltage is varied. Usually errors due to this effect are negligible compared to temperature drift. Static or DC supply voltage coefficients, $\triangle e_{os}/\triangle V_s$, $\triangle i_b/\triangle V_s$, $\triangle i_d/\triangle V_s$ are measured with the circuit in figure 11 by varying supply voltages individually by ± 1 volt.

There is a common misconception that tracking power supplies whose plus and minus voltages change by the same amounts will improve supply voltage coupling. In general tracking supplies are of no benefit since the positive supply voltage coefficient is usually much larger than the negative supply voltage coefficient. Rejection of AC noise and ripple on the power supplies is not as good as static or DC rejection, but for almost all amplifiers AC rejection will be better than 1mV/V or 60dB over a wide range of frequencies.

DRIFT VS TIME

Offset voltage, bias current and difference current change with time as components age. Static data over long time periods is difficult to obtain because of the inherent time delays involved. But it is safe to say that the published time drift for amplifiers does not accumulate linearly. For example, voltage drift for a chopper stabilized amplifier (which by the way is by far the best amplifier type for long term stability) is usually quoted as $1\mu V$/day whereas cumulative drift over 30 days will usually not exceed $5\mu V$ nor $15\mu V$ in a year.

Long term voltage drift in differential input type amplifiers depends primarily on the aging of collector resistors in the input differential pair. The aging coefficient referred to the input is about $300\mu V/\%$ change of collector resistance. It not unlikely that carbon composition resistors will age by 1 2% over a year resulting in an offset voltage change of 300 $600\mu V$. The use of metal film resistors for the collector resist will greatly improve long term stability to the point wh base to emitter voltage aging is the determining factor. W metal film resistors, offset voltage for transistor amplifiers about $100\mu V$/year while FET amplifiers will age somewl more.

Long term bias current stability in differential input amplifi again depends on resistor stability when internal initial curr compensation is employed. In this case, multi-megohm carb composition resistors are used (since large value metal fi resistors are not available) to supply about 90% of the b bias current. If these resistors change by 1%, the specifi initial bias current will change by about 9% which can be substantial drift. Therefore one can conclude that amplifi without internal initial current compensation will exhibit mc stable bias current. Under these conditions long term bias c rent stability depends primarily on the stability of the transist or FET devices which may be better than 1%.

INPUT IMPEDANCE

Differential input impedance, R_d, is defined as the impedan between the two input terminals, measured at 25°C, assumi that the error voltage, e_ε, is nulled or very near zero volts (s figure 1). For a single ended amplifier, R_d, is the input impe ance since the plus input is grounded. To a first approximatio dynamic impedance can be represented by a capacitor, C_d, parallel with R_d.

Differential input impedance is among the most difficult pa meters to measure particularly for a high gain, high impedan type amplifier. In general this measurement can only be ma under laboratory conditions by an experienced engineer wi special fixtures to shield against noise pick up. For this reas most companies including Analog Devices rarely measure t parameter on a production line basis. Fortunately a preci knowledge of R_d is not required, since for most circuits, long as R_d is large compared to the external feedback impe ance, its value has little bearing on closed loop performance.

The circuits of figure 12 show in principle how R_d can be me sured with enough attention to reducing noise. These circu actually measure R_d in parallel with the negative input comm mode impedance. However, common mode impedance is usual 10 to 100 times greater than R_d so that the error is negligibl

Common mode impedance, R_{cm}, is defined as the impedan between each input and ground or power supply common a is specified at 25°C. (See figure 1.) For most circuits comm mode impedance on the negative input, R_{cm}^-, has little signi cance except for the capacitance which it adds to the summi junction. However, common mode impedance on the pl input, R_{cm}^+, sets the upper limit on closed loop input impe ance for the non-inverting configuration. Dynamic impedan can be represented by a capacitor, C_{cm}, in parallel with R which usually runs from 5 to 25 pf on the plus input.

The circuit of figure 13 can be used to measure R_{cm}^+ up about 500M ohms. Use an oscillator frequency of 1 to 5

d adjust R_s for 10% reduction at the output. Then $_m+ = 9R_s$. Above this impedance it is advisable to substitute picoameter for the resistor R_s and to measure DC bias cur-nt as a function of common mode voltage.

)mmon mode impedance is a non-linear function of both mperature and common mode voltage. For FET amplifiers mmon mode impedance is reduced by a factor of two for each)°C temperature rise.

s a function of common mode voltage, R_{cm} is defined as erage impedance for a common mode voltage change from ro to $\pm E_{cm}$, that is, maximum common mode voltage. Incre-ental R_{cm} about some large common mode voltage may be nsiderably less than the specified average R_{cm}, especially for ET input amplifiers.

IAXIMUM VOLTAGE BETWEEN INPUTS

nder most operating conditions, feedback maintains the error oltage, e_ε, between inputs very near to zero volts. However, some applications, such as voltage comparators, or where ie input voltage exceeds the level required to saturate the out-ut, the voltage between inputs can become large. E_d defines ie maximum voltage which can be applied between inputs ithout causing permanent damage to the amplifier. Placing arallel back to back diodes across the input terminals is one ay to provide added protection for the amplifier.

IAXIMUM COMMON MODE VOLTAGE

or differential input amplifiers, the voltage at both inputs can e raised above ground potential. Common mode voltage, e_{cm}, defined as the voltage above ground at each input when both nputs are at the same voltage. E_{cm} is defined as the maximum eak common mode voltage at the input before clipping or xcessive non-linearity is seen at the output. E_{cm} establishes the aximum input voltage for the voltage follower connection. See figure 7.)

E_{cm} is measured with the circuit of figure 7 by increasing the peak input voltage (sinusoidal waveform) until distortion is seen on the scope (about 1 to 3%). The input signal frequency must be well below the full power response frequency, f_p, for the non-inverting input.

COMMON MODE REJECTION

An ideal operational amplifier responds only to the difference voltage between inputs ($e+ - e-$) and produces no output for a common mode voltage — that is when both inputs are at the same potential. However, due to slightly different gains between the plus and minus inputs, common mode input voltages are not entirely substracted at the output. If we refer the output common mode error voltage to the input (dividing by gain) and call this the input common mode error voltage, $e_{\varepsilon cm}$, then common mode rejection (CMR) is defined as the ratio of common mode voltage to common mode error voltage. That is $CMR = e_{cm}/e_{\varepsilon cm}$. CMR is sometimes expressed in dB in which case you take 20 times the log (base 10) of the ratio. Errors due to common mode rejection can be represented in the equivalent circuit of figure 1 by a voltage generator, $e_{\varepsilon cm}$, in series with the input. Note that common mode error goes to zero when either input is grounded. Therefore the inverting configuration does not exhibit a common mode error since the plus input is grounded. Thus CMR is only a problem in the non-inverting and differential configurations where common mode voltage varies in direct proportion to the input signal. In this case $e_{\varepsilon cm}$ is a basic measuring error which affects the overall circuit accuracy.

For example, if a 10 volt signal, e_i, were applied to the input of the circuit in Figure 14 common mode voltage, e_{cm}, is equal to the input voltage, e_i. This would cause a common mode voltage, $e_{\varepsilon cm}$, of 2mV for an amplifier with 5,000 or 74dB CMR and thus a 0.02% measuring error.

13

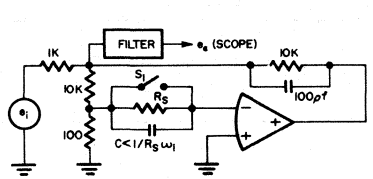

SET e_i FOR 1 TO 5 Hz, INCREASE R_s UNTIL OPENING S_1 INCREASE e_o BY FACTOR OF 2. THEN $R_s = R_d$

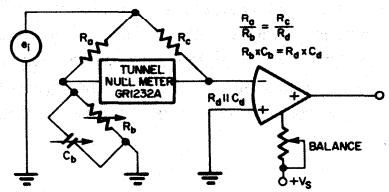

BALANCE DC OUTPUT REDUCE e_i TO SMALLEST POSSIBLE VOLTAGE AND FREQUENCY

FIGURE 12. DIFFERENTIAL IMPEDANCE TEST CIRCUITS

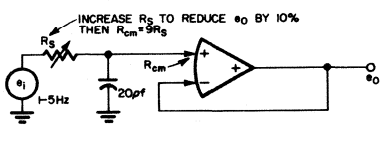

FIGURE 13. COMMON MODE IMPEDANCE TEST CIRCUIT

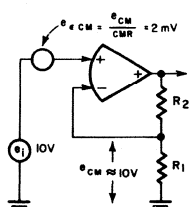

FIGURE 14. ILLUSTRATION OF COMMON MODE VOLTAGE

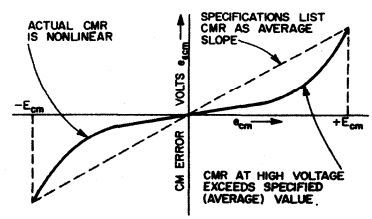

FIGURE 15. CM ERROR VS CM VOLTAGE

Precisely specifying CMR is complicated by the fact that common voltage error, $e_{\varepsilon_{cm}}$, can be a highly non-linear function of common mode voltage and it also varies with temperature. This is particularly true for FET input amplifiers. As illustrated by figure 15, CMR published by Analog Devices are average figures assuming an end point measurement at maximum common mode voltage, $\pm E_{cm}$. But the incremental CMR about some large common mode voltage may be less than the average CMR which is specified. In fact, if common mode error were a linear function of common mode voltage and if CMR were not strongly influenced by temperature then this source of error would be of little consequence. This follows since a linear CMR error can be viewed as a gain error which could be compensated for by adjusting the closed loop gain. Therefore linearity of common mode error, $e_{\varepsilon_{cm}}$, vs common mode voltage is actually more important for many applications than CMR itself.

The circuit in figure 16 provides a unique method to instrument CMR measurements as well as to measure the non-linearity of common mode errors. The oscilloscope display will duplicate the pattern of figure 15. A floated power supply allows a single ended oscilloscope to be used and almost any regulated power supply has floated outputs with sufficient isolation. Published CMR specifications apply only to DC input signals so that this measurement should be made with a signal frequency of 5Hz or less. CMR at higher frequencies, although not guaranteed by the specifications, can also be measured with this circuit. It is further assumed that the external circuit impedances of both the test circuit and the application are small compared to the common mode impedance to avoid additional common mode errors due to impedance unbalance.

INPUT NOISE

Input voltage and current noise characteristics can be in principle measured, specified and analyzed very much like offset voltage and bias current characteristics. In fact, offset voltage and bias current drift can be considered noise which occurs at very low frequencies. For this purpose in both the equivalent circuit of figure 1 and the test circuit of figure 11, replace e_{os} by e_n, an equivalent voltage noise generator, and replac i_b^- and i_b^+ by i_N^- and i_N^+, equivalent current noise generators The primary difference in measuring and specifying noise a opposed to DC drift is that bandwidth must be considered. A low frequencies, 100Hz or less, 1/f noise prevails which mean that the noise per root cycle increases inversely with frequency At the mid band frequencies noise per root cycle is constant o "white."

To measure noise a sharp cut off bandpass filter is added to th output of the circuit of figure 11. Furthermore the impedance gain and capacitors must be adjusted to assure that neither th amplifier nor the external feedback components limit the nois bandwidth of the measurement. For very low current noise, i becomes very difficult to make wideband measurements becaus of the interaction of stray capacitance and the large samplin resistor values needed to boost sensitivity.

Usually two noise measurements are taken. Low frequency noise in a bandpass of .01 to 1Hz is measured on a strip chart recorder and is specified as peak to peak with a 3σ uncertainty, meaning that 99% of the observed peak to peak exclusions will fall within the specified limits. Wideband noise in a bandpass of 5Hz to 50kHz is measured on a VTVM, preferably a true RMS type, and is specified as rms. Of course, shielding becomes very critical in these measurements to avoid power line frequency and radio frequency pick up.

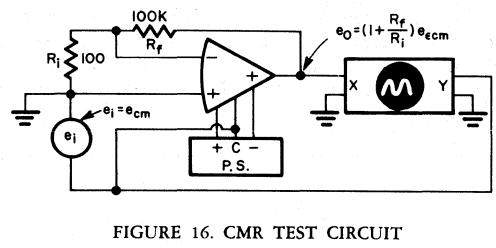

FIGURE 16. CMR TEST CIRCUIT

ANALOG DEVICES

AN-357
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Operational Integrators

by Ray Stata

Modern solid state operational amplifiers make remarkably good integrators. Almost any degree of accuracy can be achieved depending on the choice of the amplifier and the feedback capacitor. A great deal of literature exists† which discusses integrator error in analog computers and this subject will not be covered here. But we shall review the non-ideal characteristics of operational amplifiers (and to some extent capacitors) which limit the performance of integrators in instrumentation circuits. This we hope will help the reader make a better choice of amplifiers for his particular application.

†Korn and Korn, *Electronic Analog and Hybrid Computers* — McGraw Hill.

An ideal operational amplifier for integrator applications would have infinite open loop gain and input impedance and zero offset voltage and current (that is, $e_o = o$, when $e_s = o$). For this case, Figure 1 shows the characteristics of an ideal integrator.

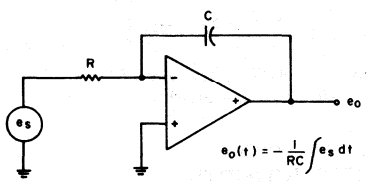

Figure 1. Ideal Operational Integrator

The gain (or characteristics time) of the circuit is given by 1/RC, which is to say that the output will change by (1/RC) volts/sec for each volt of input signal. The input impedance as viewed from the source voltage, e_s, is determined by the value for R.

OFFSET AND DRIFT ERRORS

By far the greatest source of error in integrators is due to offset and drift of the amplifier. An equivalent circuit is given in Figure 2 from which we can predict the errors due to offset. For the moment we shall assume that open loop gain, A, and open loop input impedance, R_d, are infinite.

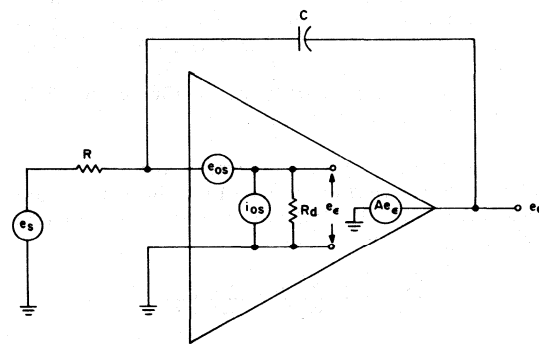

Figure 2. Equivalent Circuit for Integrator

$$e_{os} = \underset{\substack{\mu V \\ \text{at} \\ 25°C}}{E_{os}} + \underbrace{\frac{\Delta e_{os}}{\Delta T}}_{\mu V/°C}\Delta T + \underbrace{\frac{\Delta e_{os}}{\Delta V_s}}_{\mu V/\%}\Delta V_s + \underbrace{\frac{\Delta e_{os}}{\Delta t}}_{\mu V/day}\Delta t$$

$$i_{os} = \underset{\substack{\text{pa} \\ \text{at} \\ 25°C}}{I_{os}} + \underbrace{\frac{\Delta i_{os}}{\Delta T}}_{pa/°C}\Delta T + \underbrace{\frac{\Delta i_{os}}{\Delta V_s}}_{pa/\%}\Delta V_s + \underbrace{\frac{\Delta i_{os}}{\Delta t}}_{pa/day}\Delta t$$

As shown the offset voltage, e_{os}, and the offset current, i_{os}, can be calculated for any temperature, supply voltage and time period from the drift coefficients of the amplifier. It is usually possible to adjust the initial offset voltage and current, E_{os} and I_{os}, to zero by some biasing network.

Reprinted from Analog Dialogue 1–1 1967

The simplest way to analyze offset errors is to refer them to the source voltage as shown in Figure 3. In this case offset current is multiplied by R and becomes a voltage source. When viewed at the

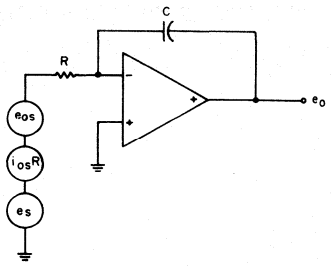

Figure 3. Offsets Referred to the Input

input, the offsets cannot be distinguished from the input signal and hence introduce a basic error in the integration of the signal. The percentage error would be, % error $= (e_{os} + i_{os}R)\ 100/\bar{e}_s$, where $\bar{e}_s$ is the time average of the input signal over the integration period. Notice that R should be as small as possible to minimize offset errors for a given amplifier. But remember that R also sets the input impedance for the integrator.

When referred to the input, the analysis of offset errors is not much different for an integrator than for an inverting DC amplifier. More detailed information is given on this subject in an Analog Devices' application note entitled "Part IV, Offset and Drift in Operational Amplifiers."

In some applications it is necessary to refer the offset error to the output in order to derive meaningful results. In this case the output error is a drift rate which is given by,

$$\frac{de_o}{dt} = \frac{e_{os} + Ri_{os}}{RC} = \frac{e_{os}}{RC} + \frac{i_{os}}{C}$$

Again, we see that output drift rate is minimized by using the smallest value for R and the largest value for C. This follows since the drift rate due to offset voltage is fixed by the gain of the circuit (1/RC) whereas the drift rate due to offset current is reduced by using a larger C.
The practical limits on the choice of R and C are as follows:

1. Source impedance sets a minimum value on input impedance which is equal to R.
2. The physical size, price and quality are all serious problems in using large value capacitors particularly when greater than 1 to 5μf.

For differential input amplifiers, the error due to offset current is generally reduced by balancing the impedance as seen from each input to ground. For the circuit in Figure 3 this would amount to inserting a resistor from the plus input to ground equal to R. Due to the input symmetry of a differential amplifier, offset current at each input tends to be equal and tends to track with temperature and thus the drift error is reduced by balancing impedances.

ERRORS DUE TO FINITE GAIN, INPUT IMPEDANCE AND BANDWIDTH

The open loop gain response for most operational amplifiers can be represented by the graph in Figure

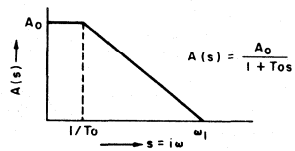

Figure 4. Typical Open Loop Gain Response

4. If we go back to Figure 2 and assume that the amplifier has to gain response of Figure 4 and an open loop input impedance, R_d, then the exact expression for integrator gain would be:

$$\frac{e_o(s)}{e_s(s)} = \underbrace{\left[\frac{-1}{RCs}\right]}_{\text{ideal}} \underbrace{\left[\frac{1}{1 + \left(\frac{1 + T_o s}{A_o}\right)\left(1 + \frac{1}{R_p C s}\right)}\right]}_{\text{error due to finite gain and bandwidth}} \quad (1)$$

where $R_p = R_d R / R_d + R$ (parallel sum)

Equation (1) can be simplified if we assume that $A_o >> 1$ (a very safe bet):

$$\frac{e_o(s)}{e_s(s)} = \left[\frac{-1}{RCs}\right]\left[\frac{1}{1 + \frac{s}{\omega_1} + \frac{1}{A_o R_p C s}}\right] \text{ for } A_o >> 1 \quad (2)$$

where $\omega_1 \approx A_o/T_o$ is the amplifier unity gain bandwidth.

HIGH FREQUENCY ERRORS DUE TO FINITE BANDWIDTH

Finite amplifier bandwidth imposes some limitation on the ability of the integrator to respond to instantaneous input changes. The transient behavior at t = o can be predicted by examining the behavior of equation (2) at high frequencies. In this case equation (2) becomes:

$$\frac{e_o(s)}{e_s(s)} = \frac{-1}{RCs}\left(\frac{1}{1 + s/\omega_1}\right) \text{ for } s >> \frac{1}{A_o R_p C} \quad (3)$$

This is the equation for an ideal integrator except for a time lag which is inversely proportional to the unity gain bandwidth, ω_1. To illustrate the error due to finite bandwidth, consider the response of (3) to a step function input as given by (4) and Figure 5.

$$e_o(t) \approx \frac{1}{RC}\left(t - 1/\omega_1\right) \text{ for } e_s(t) = -\mu_{-1}(t) \quad (4)$$

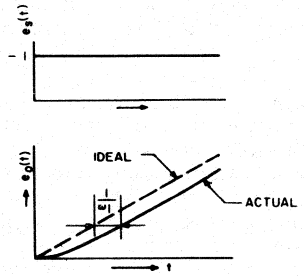

Figure 5. Step Response of Integrator at t = o

Note that the time lag depends only on amplifier open loop bandwidth, ω_1, and is independent of the values for R and C.

LOW FREQUENCY ERRORS DUE TO FINITE GAIN

The behavior of an integrator over long time periods can be predicted by the low frequency response of the circuit. In this case, where $s << \omega_1$, equation (2) becomes:

$$\frac{e_o(s)}{e_s(s)} = \left[\frac{-1}{RCs}\right]\left[\frac{1}{1+\frac{1}{A_oR_pCs}}\right] = -\frac{A_oR_p/R}{1+A_oR_pCs} \qquad (5)$$

Insight is gained into the operation of integrators at low frequencies by realizing that (5) is equivalent to the response of an ideal integrator with an infinite gain and input impedance amplifier, but with a feedback resistor A_oR_p in parallel with the feedback capacitor as shown in Figure 6.

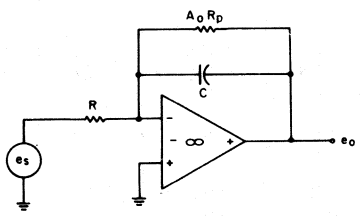

Figure 6. Integrator Low Frequency Equivalent Circuit

To illustrate more clearly the effect of low frequency errors, consider the response of (5) to a step function input as given by (6) and Figure 7.

$$e_o(t) = \frac{R_pA_o}{R}(1-e^{-t/A_oR_pC}), \text{ for } e_s(t) = -\mu_{-1}(t) \qquad (6)$$

Expanding (6) into a power series we have:

$$e_o(t) = \frac{t}{RC} - \frac{t^2}{2A_o(R_pC)(RC)} + - - -$$

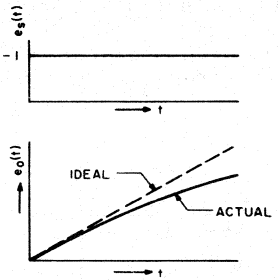

Figure 7. Step Response Showing Low Frequency Error Due to Finite Gain

The first term in this series is the response for an ideal integrator, while the second term is the principle error component which grows as the square of time.

In summary low frequency integrator errors are inversely proportional to finite open loop voltage gain. This follows from the fact that with finite gain the error voltage is not zero, as usually assumed, which tends to reduce input current as the output grows.

INTEGRATOR HOLD ERRORS

One important use of integrator circuits is to precisely remember or hold a voltage potential. Finite amplifier gain causes a fixed integrator output voltage to droop.

Intuitively it is apparent from Figure 6 that the effective leakage resistance, A_oR_p, due to finite voltage gain and input impedance tends to discharge any fixed voltage stored across the feedback capacitor.

To develop a quantitative expression for this error, assume that the circuit in Figure 6 has the initial condition $e(o) = E_o$ and that $e_s = o$. In this case, the output voltage is simply:

$$e_o(t) = E_oe^{-t/A_oR_pC}$$

By expansion this becomes:

$$e_o(t) = E_o - E_o\left[\frac{t}{A_oR_pC} - \frac{t^2}{2(A_oR_pC)^2} + - - -\right] \qquad (7)$$

The first item of (7) is the output of an ideal integrator, while the terms in the brackets represent the errors due to finite gain. Figure 8 shows integrator hold error.

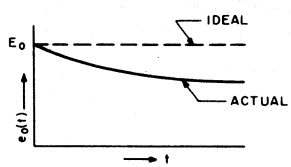

Figure 8. Integrator Hold Error

It is interesting to note that minimum error is obtained in hold operation when the input resistor is open circuited rather than short circuited. In this case the equivalent circuit is shown in Figure 9.

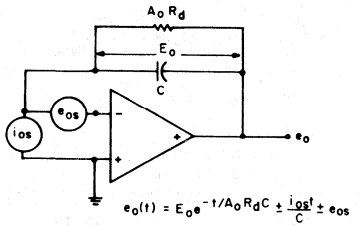

Figure 9. Hold Circuit with Input Open

Note that the equivalent feedback leakage resistor is determined only by the open loop input impedance and gain (A_oR_d). Further, the output drift rate is determined only by the offset current. Voltage offset appears at the output as a fixed offset.

FEEDBACK CAPACITOR

The performance of operational amplifiers have now reached the point where the quality of the feedback capacitor can limit the accuracy in the most precise applications. Therefore a brief discussion of capacitor limitations will be helpful in precise integrator design.

The chart in Figure 10 shows the salient characteristics for various types of high quality capacitors. This data was compiled with the assistance of Southern Electronics Corporation, Burbank, California, a capacitor manufacture who specializes in high quality capacitors for integrator applications. An application note from this company entitled "Capacitor Talk" gives more information on the interpretation of capacitor specifications.

In analog computers where scale factor accuracy is important it is common practice to enclose the feedback capacitor in a temperature controlled oven. In this case long term stability of capacitance value for polystyrene and mylar capacitors is about 0.1% per year.

Insulation Resistance — One important limitation for integrator capacitors is insulation or leakage resistance. The specification used to define this limitation usually is expressed in megohms - microfarads, which is equivalent to the time in seconds for a fixed voltage stored on the capacitor to discharge to 63% of its initial value. As a general rule the maximum insulation resistance is about two times the value for a one microfarad capacitor, which establishes the limit for insulation resistance of small capacitor values.

Dielectric	Mylar	Metalized Mylar	Poly-carbonate	Metalized Poly-carbonate	Polystyrene	Teflon	Metalized Teflon
Temperature Range							
Hi Temp (°C)	+125	+125	+125	+125	+85	+200	+200
Lo Temp (°C)	−65	−65	−65	−65	−65	−65	−65
Temperature Coefficient							
−65°C to 25°C (%)	−6	−6	−1.5	−1.5	+0.9	+1.9	+0.5
25°C to Hi Temp (%)	+12	+12	±0.5	±0.5	−0.6	−3.7	−1.0
Dielectric Absorption							
% @ 25°C	0.1	0.1	.05	.05	.02	.01	.02
Dissipation Factor							
@25°C (%)	0.3	0.5	0.1	0.2	0.02	0.01	0.1
@ Hi Temp (%)	1.2	1.7	0.07	0.6	0.04	0.02	0.2
Insulation Resistance							
@ 25°C (MΩ-μf)	$2x10^5$	$5x10^4$	$4x10^5$	$2x10^5$	$1x10^6$	$1x10^6$	$5x10^5$
@ Hi Temp (MΩ-μf)	$3x10^2$	$1x10^2$	$1.5x10^4$	$15x10^2$	$7x10^4$	$1x10^5$	$2.5x10^4$
Approximate Size for 50Vdc							
cubic inch/μf (uncased)	.12	0.06	.19	0.09	.44	1.1	0.39

Figure 10. Comparison of Capacitor Specifications

The effect of insulation resistance can be represented in Figures 6 and 9 as another resistance in parallel with A_oR_p or A_oR_d and the issuing equations are modified accordingly. The insulation resistance of the very best capacitors is about 10^{12} ohms. By comparison a chopper stabilized operational amplifier will have open loop input impedance, R_d, of 10^6 ohms and open loop gain, A_o, of 10^8 giving an equivalent resistance of 10^{14}. Even an inexpensive differential amplifier will have equivalent leakage resistance of 10^{10} to 10^{11} ohms. Thus we see that the capacitor and not the amplifier usually sets the limit on performance in this regard.

Dielectric Absorption — One of the single most important dynamic errors of integration is due to dielectric absorption. This error results from the fact that when a capacitor is charged or discharged not all of the dielectric polarization takes place immediately. Consequently there can be an appreciable residual voltage with a relatively long time constant. The specification given for this parameter is the residual voltage expressed as a percentage of the applied voltage measured approximately one second after the capacitor is discharged. Polystyrene and teflon are mostly used for precision integrators since these materials have small but measurable errors due to dielectric absorption. For additional information on analyzing this source of error you should refer to "An Analysis of Certain Errors in Electronic Differential Analyzers II—Capacitor Dielectric Absorption," P. C. Dow, IRE Trans. on Electronic Computers, Vol. EC-7, pp.17-22, March, 1958.

Dissipation Factor, which is related to dielectric absorption, can be termed the sum total of all the losses in the capacitor and is expressed as the percentage ratio of the effective series resistance to the reactive capacitance, or as the tangent of the loss angle. Dissipation factor is important in AC integrators or in analog computers where repetitive integration is performed.

LEAKAGE RESISTANCE

In the highest performance integrators, leakage resistance to the summing junction or across the feedback capacitor can play a large role in the attainable performance. It is extremely important to shield the summing junction and its leads from leakage paths to potentials other than ground. For example, offset current, which is one of the principle limitations to good integrator performance, in a good chopper amplifier is about 10^{-11} amps. The insulation resistance required to keep the leakage current from the 15VDC supply voltage less than 10^{-11} amp would have to be greater than one million megohms. The insulation resistance of most wire and connectors fall short of this requirement. However, by properly shielding the summing junction and its leads, leakage currents from active sources are shunted to ground, effectively creating extremely high insulation resistance from these potentials to the summing junction.

By the same token leakage resistance of the clamping circuits across the feedback capacitor used to reset the integrator should not be overlooked when calculating the effective leakage of a feedback capacitor. For example, the leakage resistance of a computer grade capacitor is typically 10^{12} ohms, which may be negligible compared to the leakage resistance of a relay or a solid state switch.

AC INTEGRATORS

In some applications it may be desirable to integrate AC signals over a reasonably long time and it may not be possible to reset the output to zero periodically. In this case the DC offset problem can be alleviated in part by bounding the DC closed loop gain as shown in Figure 11.

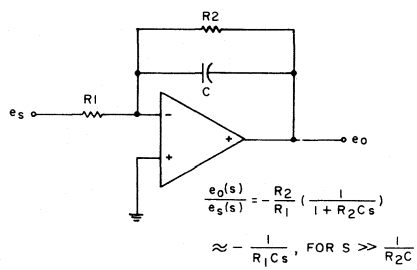

Figure 11. AC Integrator

The closed loop response for this circuit is shown in Figure 12.

For frequencies greater than $1/R_2C$ the response approaches that of an ideal integrator with gain of $1/R_1C$. For example, for signal frequencies, ω_s, a decade away from the corner frequency, $1/R_2C$, the gain error is only .5%.

The advantage of bounding the DC gain with R_2 is that the amplifier output will not drift into saturation. Instead the output will assume a DC value of $e_o = -R_2/R_1\,(e_{os} + R_1 i_{os})$. This output will limit the dynamic range for AC output signals; but, by choosing an amplifier with sufficiently low offsets, satisfactory operation can be obtained for many AC integrator applications.

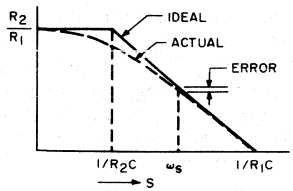

Figure 12. Gain Response for AC Integrator

For integration of very low frequency AC signals, the DC gain requirements of the previous circuit are so large as to cause saturation of the output. In this case the following circuit allows the DC gain to be reduced.

The lowest frequency which can be accurately integrated is limited by the size for C_2. The general expression for the corner frequencies ω_1 and ω_2 are rather complex and as a practical matter can only be determined by trial and error calculations. The lowest signal frequency, ω_s, should be at least ten times greater than ω_2.

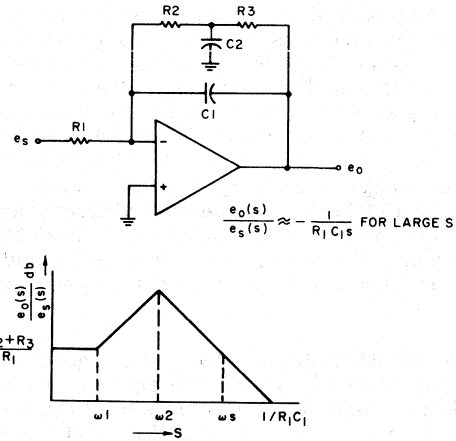

Figure 13. Low Frequency AC Integrator

ANALOG DEVICES

AN-202
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

An I.C. Amplifier User's Guide to Decoupling, Grounding, and Making Things Go Right for a Change

by Paul Brokaw

"There once was a breathy baboon
Who always breathed down a bassoon,
For he said "It appears
that in billions of years
I shall certainly hit on a tune"
(Sir Arthur Eddington)

This quotation seemed a proper note with which to begin on a subject which has made monkeys of most of us at one time or another. The struggle to find a suitable configuration for system power, ground, and signal returns too frequently degenerates into a frustrating glitch hunt. While a strictly experimental approach can be used to solve simple problems, a little forethought can often prevent serious problems and provide a plan of attack if some judicious tinkering is later required.

The subject is so fragmented that a completely general treatment is too difficult for me to tackle. Therefore, I'd like to state one general principle and then look a bit more narrowly at the subject of decoupling and grounding as it relates to integrated circuit amplifiers.

. . . Principle: Think—where the currents will flow.

I suppose this seems pretty obvious, but all of us tend to think of the currents we're interested in as flowing "out" of some place and "through" some other place but often neglect to worry how the current will find its way back to its source. One tends to act as if all "ground" or "supply voltage" points are equivalent and neglect (for as long as possible) the fact that they are parts of a network of conductors through which currents flow and develop finite voltages.

In order to do some advance planning it's important to consider where the currents originate and to where they will return and to determine the effects of the resulting voltage drops. This in turn requires some minimum amount of understanding of what goes on inside the circuits being decoupled and grounded. This information may be lacking or difficult to interpret when integrated circuits are part of the design.

Operational amplifiers are one of the most widely used linear I.C.'s, and fortunately most of them fall into a few classes, so far as the problems of power and grounding are concerned. Although the configuration of a system may pose formidable problems of decoupling and signal returns, some basic methods to handle many of these problems can be developed from a look at op-amps.

OP AMPS HAVE FOUR TERMINALS:

A casual look through almost any operational amplifier text might leave the reader with the impression that an ideal op-amp has three terminals: a pair of differential inputs and an output as shown in Figure 1. A quick review of fundamentals, however, shows that this can't be the case. If the amplifier has an output voltage it must be measured with respect to some point . . . a point to which the amplifier has a reference. Since the ideal op-amp has infinite common mode rejection, the inputs are ruled out as that reference so that there must be a fourth amplifier terminal. Another way of looking at it is that if the amplifier is to supply output current to a load, that current must get into the amplifier somewhere. Ideally, no input current flows, so again the conclusion is that a fourth terminal is required.

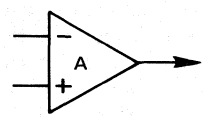

Figure 1. Conventional "Three Terminal" Op Amp

A common practice is to say, or indicate in a diagram, that this fourth terminal is "ground." Well, without getting into a discussion of what "ground" may be we can observe that most integrated circuit op-amps (and a lot of the modular ones as well) don't have a "ground" terminal. With these circuits the fourth terminal is one or both of the power supply terminals. There's a temptation here to lump together both supply voltages with the ubiquitous ground. And, to the extent that the supply lines really do present a low impedance at all frequencies within the amplifier bandwidth, this is probably reasonable. When the impedance requirement isn't satisfied, however, the door is left open to a variety of problems including noise, poor transient response, and oscillation.

DIFFERENTIAL TO SINGLE-ENDED CONVERSION:
One fundamental requirement of a simple op-amp is that an applied signal which is fully differential at the input must be converted to a single-ended output. Single ended, that is, with respect to the often neglected fourth terminal. To see how this can lead to difficulties, take a look at Figure 2.

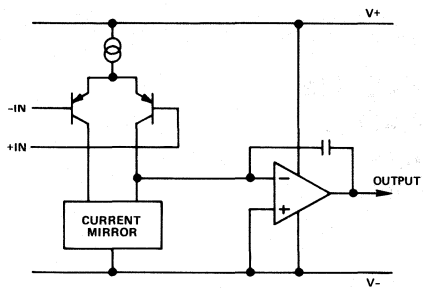

Figure 2. Simplified "Real" Op Amp

The signal flow illustrated by Figure 2 is used in several popular integrated circuit families. Details vary, but, the basic signal path is the same as the 101, 741, 748, 777, 4136, 503, 515, and other integrated circuit amplifiers. The circuit first transforms a differential input voltage into a differential current. This input stage function is represented by PNP transistors in Figure 2. The current is then converted from differential to single-ended form by a current mirror which is connected to the negative supply rail. The output from the current mirror drives a voltage amplifier and power output stage which is connected as an integrator. The integrator controls the open-loop frequency response, and its capacitor may be added externally, as in the 101, or may be self-contained, as in the 741. Most descriptions of this simplified model don't emphasize that the integrator has, of course, a differential input. It's biased positive by a couple of base emitter voltages, but, the non-inverting integrator input is referred to the negative supply.

It should be apparent that most of the voltage difference between the amplifier output and the negative supply appears across the compensation capacitor. If the negative supply voltage is changed abruptly the integrator amplifier will *force* the output to follow the change. When the entire amplifier is in a closed loop configuration the resulting error signal at its input will tend to restore the output, but, the recovery will be limited by the slew rate of the amplifier. As a result, an amplifier of this type may have outstanding low frequency power supply rejection, but, the negative supply rejection is fundamentally limited at high frequencies. Since it is the feedback signal to the input that causes the output to be restored, the negative supply rejection will approach zero for signals at frequencies above the *closed loop* bandwidth. This means that high-speed, high-level circuits can "talk to" low-level circuits through the common impedance of the negative supply line.

Note that the problem with these amplifiers is associated with the negative supply terminal. Positive supply rejection may also deteriorate with increasing frequency, but, the effect is less severe. Typically, small transients on the positive supply have only a minor effect on the signal output The difference between these sensitivities can result in ar apparent asymmetry in the amplifier transient response. I the amplifier is driven to produce a positive voltage swin across its rated load it will draw a current pulse from the positive supply. The pulse may result in a supply voltage transient, but, the positive supply rejection will minimize the effect on the amplifier output signal. In the opposite case, a negative output signal will extract a current from the negative supply. If this pulse results in a "glitch" on the bus, the poor negative supply rejection will result in a similar "glitch" at the amplifier output. While a positive pulse test may give the amplifier transient response, a negative pulse test may actually give you a pretty good look at your negative supply line transient response, instead of the amplifier response!

Remember that the impulse response of the power supply itself is not what is likely to appear at the amplifier. Thirty or forty centimeters of wire can act like a high Q inductor to add a high-frequency component to the normally overdamped supply response. A decoupling capacitor near the amplifier won't always cure the problem either, since the supply must be decoupled *to* somewhere. If the decoupled current flows through a long path, it can still produce an undesirable glitch.

Figure 3 illustrates three possible configurations for negative supply decoupling. In 3a the dotted line shows the negative signal current path through the decoupling and along the ground line. If the load "ground" and decoupled "ground" actually join at the power supply the "glitch" on the ground lines is similar to the "glitch" on the negative supply bus. Depending upon how the feedback and signal sources are "grounded" the effective disturbance *caused* by the decoupling capacitor may be larger than the disturbance which it was intended to prevent. Figure 3b shows how the decoupling capacitor can be used to minimize disturbance of V— and ground busses. The high-frequency component of the load current is confined to a loop which doesn't include any part of the ground path. If the capacitor is of sufficient size and quality, it will minimize the glitch on the negative supply without disturbing input or output signal paths. When the load situation is more complex, as in 3c, a little more thought is required. If the amplifier is driving a load that goes to a virtual ground, the actual load current does not return to ground. Rather, it must be supplied by the amplifier creating the virtual ground as shown in the figure. In this case, decoupling the negative supply of the first amplifier to the positive supply of the second amplifier closes the fast signal current loop without disturbing ground or signal paths. Of course, it's still important to provide a low impedance path from "ground" to V— for the second amplifier to avoid disturbing the input reference.

The key to understanding decoupling circuits is to note where the actual load and signal currents will flow. The key to optimizing the circuit is to bypass these currents around ground and other signal paths. Note, that as in figure 3a, "single point grounding" may be an oversimplified solution to a complex problem.

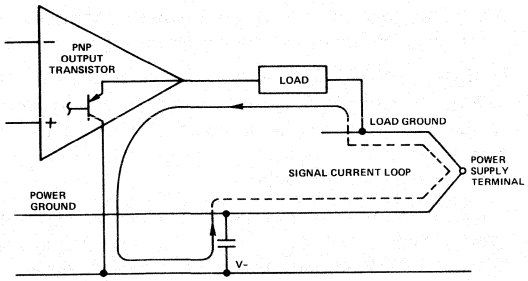

Figure 3a. Decoupling for Negative Supply Ineffective

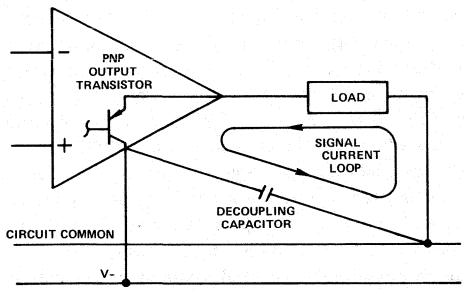

Figure 3b. Decoupling Negative Supply Optimized for "Grounded" Load

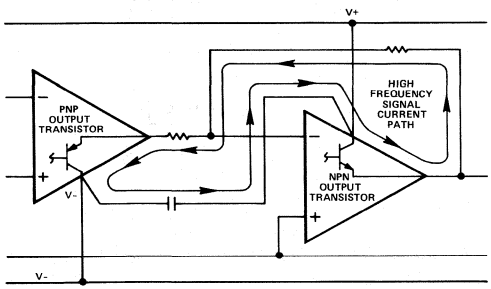

Figure 3c. Decoupling Negative Supply Optimized for "Virtual Ground" Load

Figure 3b and 3c have been simplified for illustrative purposes. When an entire circuit is considered conflicts frequently arise. For example, several amplifiers may be powered from the same supply, and an individual decoupling capacitor is required for each. In a gross sense the decoupling capacitors are all paralleled. In fact, however, the inductance of the interconnecting power and ground lines convert this harmless-looking arrangement into a complex L-C network that often rings like the "Avon Lady". In circuits handling fast signal wavefronts, decoupling networks paralleled by more than a few centimeters of wire generally mean trouble. Figure 4 shows how small resistors can be added to lower the Q of the undesired resonant circuits. The resistors can generally be tolerated since they convert a bad high-frequency jingle to a small damped signal at the op amp supply terminal. The residual has larger *low frequency* components, but, these can be handled by the op-amp supply rejection.

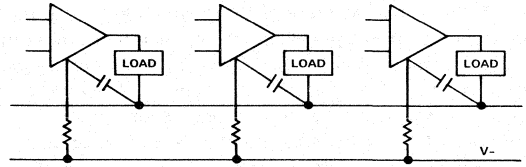

Figure 4. Damping Parallel Decoupling Resonances

FREQUENCY STABILITY:

There's a temptation to forget about decoupling the negative supply when the system is intended to handle only low-frequency signals. Granted that decoupling may not be required to handle low-frequency signals, but it may still be required for frequency stability of the op-amps.

Figure 5 is a more-detailed version of Figure 2 showing the output stage of the I.C. separated from the integrator (since this is the usual arrangement) and showing the negative power supply and wiring impedance lumped together as a single constant. The amplifier is connected as a unity gain follower. This makes a closed-loop path from the amplifier output through the differential input to the integrator input. There is a second feedback path from the collector of the output PNP transistor back to the other integrator input. The net input to the integrator is the difference of the signals through these two paths. At low frequencies this is a net, negative feedback. The high-frequency feedback depends upon both the load reactance and the reactance of the V– supply.

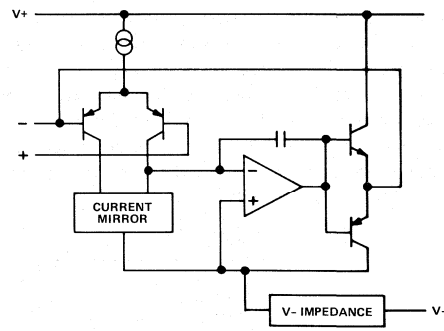

Figure 5. Instability Can Result from Neglecting Decoupling

When the supply lead reactance is inductive, it tends to destabilize the integrator. This situation is aggravated by a capacitive load on the amplifier. Although it's difficult to predict under exactly what circumstances the circuit will become unstable, it's generally wise to decouple the negative supply if there is any substantial lead inductance in the V– lead *or* in the common return to the load and amplifier input signal source. If the decoupling is to be effective, of course, it must be with respect to the *actual* signal returns, rather than to some vague "ground" connection.

POSITIVE SUPPLY DECOUPLING:

Up to this point we haven't considered decoupling the positive supply line, and with amplifiers typified by Figures 2

and 5 there may be no need to. On the other hand, there are a number of integrated circuit amplifiers which refer the compensating integrator to the positive supply. Among these are the 108, 504, and 510 families. When these circuits are used, it's the positive supply which requires most attention. The considerations and techniques described for the class of circuits shown in Figure 2 apply equally to this second class, but, should be applied to the positive supply rather than the negative.

FEED-FORWARD:

A technique which is most frequently used to improve bandwidth is called feed-forward. Generally, feed-forward is used to bypass an amplifier or level translator stage which has poor high frequency response. Figure 6 illustrates how this may be done. Each of the amplifiers shown is really a subcircuit, usually a single stage, in the overall amplifier. In the illustration, the input stage converts the differential input to a single-ended signal. The signal drives an intermediate stage (which in practice often includes level translator circuitry) which has low-frequency gain, but, limited bandwidth. The output of this stage drives an integrator-amplifier and output stage. The overall compensation capacitor feeds back to the input of the second stage and includes it in the integrator loop. The compromises necessary to obtain gain and level translation in the intermediate stage often limit its bandwidth and slow down the available integrator response. A feed-forward capacitor permits high-frequency signals to bypass this stage. As a result, the overall amplifier combines the low-frequency gain available from 3 stages with the improved frequency response available from a 2-stage amplifier. The feed-forward capacitor also feeds back to the non-inverting input of the intermediate stage. Note that the second stage is not an integrator, as it may appear at first glance, but actually has a positive feedback connection. Fed-forward amplifiers must be carefully designed to avoid internal oscillations resulting from this connection. Improper decoupling can upset this plan and permit this loop to oscillate.

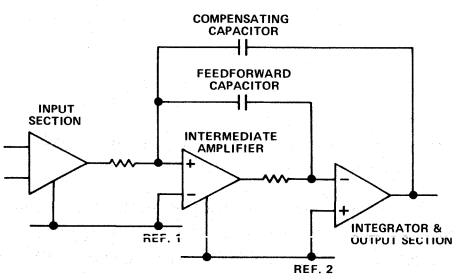

Figure 6. Fast Fed-Forward Amplifier

Note that the internal input stages are shown as being referred to separated reference points, Ideally, these will be the same reference so far as signals are concerned, although they may differ in bias level. In practice this may not be the case. Examples of fed-forward amplifiers are the AD518 and the AD707. In these amplifiers, signal Reference 1 is the positive supply, while signal Reference 2 is the negative supply. Signals appearing between the positive and negative supply terminals are effectively inserted inside the integrator loop!

Obviously, while feed-forward is a valuable tool for th high-speed amplifier designer, it poses special problems i application. A thoughtful approach to decoupling is re quired to maximize bandwidth and minimize noise, error and the likelihood of oscillation.

Some fed-forward amplifiers have other arrangements which include the "ground" terminal in inverting only am plifiers. Almost without exception, however, signals be tween some combination of the supply terminals get "in side" the amplifier. It is vital to proper operation that the involved supply terminals present a common low imped ance at high frequencies. Many high-speed modular ampli fiers include appropriate capacitive decoupling within the amplifier, but, with I.C. op amps this is impossible. The user must take care to provide a cleanly decoupled supply for fed-forward amplifiers. Figure 7 shows a decoupling method which may be applied to the AD518 as well as to other fast fed-forward amplifiers such as the 118. One capacitor is used to provide a low-impedance path between the supply terminals at high frequencies. The resistor in the V+ lead insures that noise on the supply lines will be rejected *and* prevents the establishment of resonances with other decoupling circuits. The second capacitor decouples the low side of the integrator to the load.

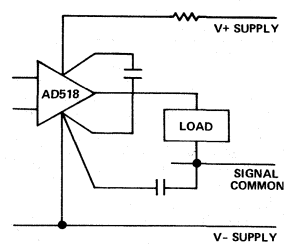

Figure 7. Decoupling for a Fed-Forward Amplifier

Alternatives include a resistor in both supply leads and/or decoupling from V+ to the load. In principle, the positive and negative supply should be tied in a "tight knot" with the signal return. To the extent that this cannot be done, there is a slight advantage to favoring the negative supply due to the high frequency limitations of PNP transistors used in junction-isolated I.C.'s.

OTHER COMPENSATION:

While most integrated circuit amplifiers use one of the three compensation schemes already described, a significant fraction use some other plan. The 725 type amplifiers combine a V– referred integrator with a network which the manufacturers recommend to be connected from signal ground to the integrator input. This makes the circuit extremely liable to pick up noise between V– and ground. In many circumstances it may be wiser to connect the external compensation to the negative supply, rather than to signal ground.

One more class of amplifiers is typified by the Analog Devices AD507 and AD509. In these circuits, a single capaci-

or may be used to induce a dominant pole of response without resorting to an integrator connection. The high-frequency response of the amplifier will appear with respect to the "ground" end of the compensation capacitor. In these amplifiers a small internal capacitance is connected between V+ and the compensation point. Unity gain compensation can be added in parallel and the pin-out is arranged to make this simple. The free end of the compensation capacitor can also be connected either to V– or signal common. It is extremely important that the signal common and the compensation connect directly or through a low-impedance decoupling.

Although the main signal path of these amplifiers can be compensated in a variety of ways, some care is required to insure the stability of internal structures. It's always wise to use extra care in decoupling wideband amplifiers to avoid problems with the output stage and other subcircuits which are similar to the main integrator problem illustrated by Figure 5. An effective compensation and decoupling circuit for the AD509 is shown in Figure 8. This arrangement is similar to Figure 7, and one of these two circuits is likely to be suitable for many types of wideband amplifier. Depending upon the power distribution, a small (10Ω to 50Ω) resistor may be appropriate in both of the supply leads to reduce power lead resonance and interference both to and from circuits sharing the power supply.

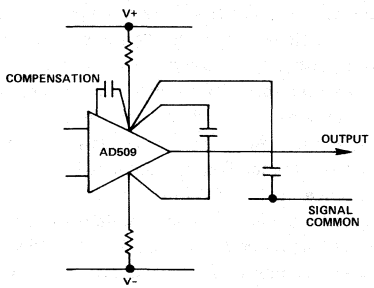

Figure 8. Decoupling a Wideband Amplifier

GROUNDING ERRORS:

Ground in most electronic equipment is not an actual connection to earth ground, but a common connection to which signals and power are referred. It is frequently immaterial to the function of the equipment whether or not the point actually connects to earth ground. I myself prefer some distinguishing name or names for these common points to emphasize that they must be *made* common. The term "ground" too often seems to be associated with a sort of cure-all concept, like snake oil, money or motherhood. If you're one of those who regards ground with the same sort of irrational reverence that you hold for your mother, remember that while you can always trust your mother, you should *never* trust your "ground." Examine and think about it.

It's important to have a look at the currents which flow in the ground circuit. Allowing these currents to share a path with a low-level signal may result in trouble. Figure 9 illustrates how careless grounding can degrade the performance of a simple amplifier. The amplifier drives a load which is represented by the load resistor. The load current comes from the power supply and is controlled by the amplifier as it amplifies the input signal. This current must return to the supply by some path; suppose that points A and B are alternative power supply "ground" connections. Assuming that the figure represents the proper topology or ordering of connections along the "ground" bus, connecting the supply at A will cause the load current to share a segment of wire with the input signal connection. Fifteen centimeters of number 22 wire in this path will present about 8 milliohms of resistance to the load current. With a 2k load, a 10-volt output signal will result in about 40 microvolts between the points marked "ΔV." This signal acts in series with the non-inverting input and can result in significant errors. For example, the typical gain of an AD510 amplifier is 8 million so that only 1¼μV of input signal is required to produce a 10 volt output. The 40μV ground error signal will result in a 32 times increase in the circuit gain error! This degradation could easily be the most serious error in a high-gain precision application. Moreover, the error represents positive feedback so that the circuit will latch up or oscillate for large closed-loop gains with R_f/R_i greater than about 250k.

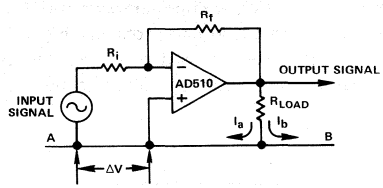

Figure 9. Proper Choice of Power Connections Minimizes Problems

13

Reconnecting the power supply to point B will correct the problem by eliminating the common impedance feedback connection. In a real system, the problem may be more complex. The input signal source, which is represented as floating in Figure 9, may also produce a current which must return to the power supply. With the supply at point B, any current which flows in additional loads (other than R_i) may interfere with the operation of the amplifier shown. Figure 10 illustrates how amplifiers can be cascaded and still drive auxiliary loads without common impedance coupling. The

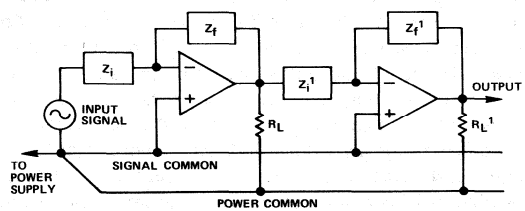

Figure 10. Minimizing Common Impedance Coupling

output currents flow through the auxiliary loads and back to the power supply through power common. The currents in the input and feedback resistors are supplied from

the power supply by way of the amplifiers as previously illustrated in Figure 3c. The only current flowing in signal common is the amplifier's input current, and its effect is generally negligibly small.

Having given an example of a simple "grounding error" and its solution, I will now get back on my soap box and say that grounding errors result from neglect based on the assumption that a ground, is a ground, is a ground. *Some* impedance will be present in any interconnection path, and its effect should be considered in the overall design of a system. Quantitative approaches are quite useful in specialized applications. In fast TTL and ECL logic circuitry the characteristic impedance of interconnections is controlled so that proper terminations can reduce problems. In RF circuitry the unavoidable impedances are taken into account and incorporated into the design of the circuit. With op-amp circuitry, however, impedance levels do not lend themselves to transmission line theory, and the power and ground impedances are difficult to control or analyze. The most expedient procedure, short of difficult and restrictive quantitative analysis, seems to be to arrange the unavoidable impedances so as to minimize their effects and arrange the circuitry to overcome the effects. Figures 9 and 10 illustrate the sort of simple considerations which can substantially reduce practical ground problems. Figure 11 illustrates how circuitry can be used to reduce the effect of ground problems which can't be corrected by topological tricks.

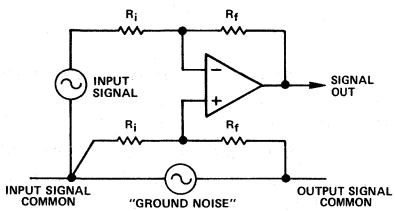

Figure 11. Subtractor Amplifier Rejects Common Mode Noise

GETTING AROUND THE PROBLEM:

In Figure 11 a subtractor circuit is used to amplify a normal mode input signal and reject a ground noise signal which is common to both sides of the input signal. This scheme uses the common-mode rejection of the amplifier to reduce the noise component while amplifying the desired signal. An important aspect of this arrangement, which is often overlooked, is that the amplifier should be powered with respect to the *output* signal common. If its power pins are exposed to the high-frequency noise of the input common, the compensation capacitor will direct the noise right to the output and defeat the purpose of the subtractor. It's just this kind of effect which makes it important to use care in grounding and decoupling. A subtractor or dynamic bridge, like Figure 11, will be ineffective in correcting a grounding problem if the amplifier itself is carelessly decoupled. In general, an op-amp should be decoupled to the point which is the reference for measuring or using its output signal. In "single-ended" systems it should also be decoupled to the input signal return as well. When it is impossible to satisfy both these requirements at once, there's a high probability of either a noise or oscillation problem or both. Frequently the difficulty can be resolved with a subtractor, like Figure 11, where a network like the single-ended feedback network (which needn't be all resistive) joins the input and output signal reference points and provides a "clean" reference point for the non-inverting input of the amplifier.

A problem with the subtractor is that it uses a balanced bridge to reject the common mode signal between the input and output reference points. The arms of the network must be carefully balanced, since to the extent they don't match, the unwanted signal will be amplified. Although even a poorly matched network will probably eliminate oscillation problems, noise rejection will suffer in direct proportion to any mismatches. An easier way to reject large "ground noise" signals is to use a true instrumentation amplifier.

INSTRUMENTATION AMPLIFIERS:

A true instrumentation amplifier has a very visible "fourth terminal." The output signal is developed with respect to a well defined reference point which is usually a "free" terminal that may be tied to the output signal common. The instrumentation amplifier also differs from an op amp in that the gain is fixed and well defined, but there is no feedback network coupling input and output circuits. Figure 12 shows how an instrumentation amplifier can be used to translate a signal from one "ground reference" to another. The normal mode input signal is developed with respect to one reference point which may be common to its generating circuits. The signal is to be used by a system which has an interfering signal between its own common and the signal source. The instrumentation amplifier has a high impedance differential input to which the desired signal is applied. Its high common mode rejection eliminates the unwanted signal and translates the desired signal to the output reference point. Unlike the dynamic bridge circuit, the gain and common mode rejection don't depend on a network connecting the input and output circuits. The gain is set, in Figure 12, by the ratio of a pair of resistors which are connected inside the amplifier. The amplifier has a very high input impedance, so that gain and common mode rejection are not greatly affected by variations or unbalance in source impedance.

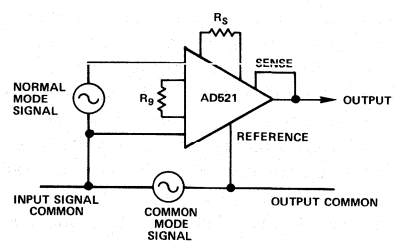

Figure 12. Applying an In-Amp

Since instrumentation amplifiers have a reference or "ground" terminal, they have the potential to be free of the power supply sensitivities of op amps. In practice, however, most instrumentation amplifiers have internal frequency

compensation which is referred to the power supply. In the case of the AD521, the compensation integrator is referred to the negative supply terminal. The decoupling of this terminal is particularly important, and it should be decoupled with respect to the output reference terminal, or actually to the point to which this terminal refers.

THE "OTHER" INPUT:

Most I.C. op-amps and in-amps include offset voltage nulling terminals. These terminals generally have a small voltage on them and by loading the terminals with a potentiometer the amplifier offset voltage can be adjusted. While their impedance level is much lower than the normal input, the null terminals can act as another differential input to the amplifier. Although the null terminals aren't generally looked at as inputs, most amplifiers are quite sensitive to signals applied here. For example, in 741 family amplifiers the output voltage gain from the null terminals is greater than the gain from the normal input!

An illustration of the type of problems that can arise with the "other" input is shown in Figure 13. The figure is an op-amp circuit with some of the offset null detail shown.

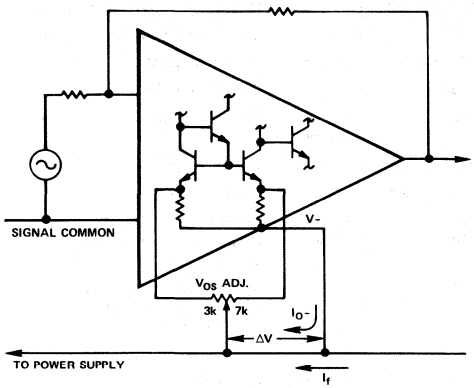

Figure 13. Details of V_{OS} Nulling – the "Other" Input

As it's drawn, the V_{OS} null pot wiper connects to a point along a V– "clothesline" which carries both the return current from the amplifier and currents from other circuits back to the power supply. These currents will develop a small voltage, ΔV, along the conductor between the amplifier V– terminal and the null pot wiper. If the null pot is set on center, the equal halves will form a balanced bridge with the resistors inside the amplifier. The effect of the voltage generated along the wire is balanced at the V_{OS} terminals and will have little effect on the amplifier output. On the other hand, if the null pot is unbalanced, to correct an amplifier offset, the bridge will no longer balance. In this case voltages developed along the "clothesline" will result in a difference voltage at the V_{OS} terminals. For instance, suppose that a 10k null pot balances out the op amp offset when it is set with 3k and 7k branches as shown in the figure. In a 741 the internal resistors are about 1k so that the difference signal at the V_{OS} terminals will be about 1/8 ΔV. The gain from these terminals is about twice the gain from the normal input, so that the disturbance acts as if it were an input signal of about 1/4 ΔV. Using the same assumptions as in the discussion of Figure 9. the current I_{O-} will result in a 10 microvolt input error signal. In this case, however, the error will appear *only* when the amplifier load current comes from the negative supply. When the load is driven positive the error will disappear. As a result, the V_{OS} input signal will result in distortion rather than a simple gain error!

An additional problem is created by I_f, a current returning to the power supply from other circuits. The current from other circuits is not generally related to the op amp signal, and the voltage developed by it will manifest itself as noise. This signal at the null terminals can easily be the dominant noise in the system. A few milliamps of V– current through a few centimeters of wire can result in interference which is orders of magnitude larger than the inherent input noise of the amplifier. The remedy is to make the connection from the null pot wiper direct to the V– pin of the amplifier, as shown in Figure 14. Some amplifiers such as the AD504 and AD510 refer to the null offset terminals to V+. Obviously, the pot wiper should go to the V+ terminal of this type of amplifier. It's important to connect the line directly to the op amp terminal so as to minimize the common impedance shared by the op amp current and the null pot connection.

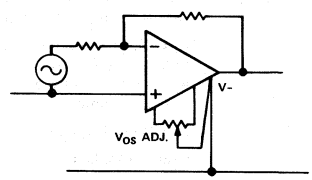

Figure 14. Connecting the Null Pot for Trouble Free Operation

The considerations for op-amp null pots also apply to the similar trimmers on almost all types of integrated circuits. For example, the AD521 In-Amp null terminals exhibit a gain of about 30 to the output. Although this is much less than in the case of most op-amps, it still warrants care in controlling the null pot wiper return. Table I lists the integrated circuits manufactured by Analog Devices, including some popular second-source families, and indicates how internal conversions from differential to single ended are referred. That is, the signals are made to appear with respect to the terminal(s) listed.

	Internal Integrator Referred to:	Comment
AD OP 07/ 27/37	V+, V−	Internal Feedforward Cap V+ to V− and Integrator V− to Output
AD380	V+	
AD390	V−	Output and Reference Amplifier
AD394/AD395	V−	Output Amplifiers
AD396	V−	Output Amplifiers
AD507	–	External Cap to Signal Common or V+
AD508	–	External Cap to Signal Common or V+
AD510	V+	
AD517	V+	
AD518	V+, V−	Internal Feedforward Cap V+ to V− and Integrator V− to Output
AD521	V−	Output Amplifier Integrator
AD524	V−	Output Amplifier Integrator
AD526	V−	Output Amplifier Integrator
AD532/AD533	V+	Multiplier Output Amplifier Integrator
AD534/AD535	V−	Output Amplifier
AD536A	V−, V+ Common	External Integrator to V+, Internal Feedforward V− to Common
AD538	V−	Internal Amplifiers
AD542/AD642	V−	
AD544/AD644	V−	
AD545A	V−	
AD546	V−	
AD547/AD647	V−	
AD548/AD648	V−	
AD549	V−	
AD557/AD558	Common	Output Amplifier and DAC Control Loop Integrator Referred to Common
AD561	V−, Common	DAC Control Loop Integrator and Ref. Amp Referred to Common and Ref. Bias Amplifier Referred to V−
AD565A/ AD566A	V−	DAC Control Loop Integrator Referred to V−. Reference Input Common to Control Loop Isolated from DAC Output Common
AD568	V+	Reference Amplifier
AD580	V−	Output Amplifier
AD581	V−	Output Amplifier
AD582	V−	Output Amplifier
AD584	V−	Output Amplifier
AD586/AD587	V−	Output Amplifier
AD588	V−	Output Amplifier
AD624/AD625	V−	Output Amplifier Integrator
AD636	V−, V+, Common	External Integrator to V+, Internal Feedforward V− to Common
AD637	V−, Common	Internal Feedforward V− to Common
AD645	V−	
AD650/AD652	V+	Internal Amplifier
AD662	Common	DAC Control Loop Integrator and Reference Amplifier Referred to Common
AD664	V−	Output Amplifiers
AD667	V−, Common	Output Amplifier Referred to V− and Reference Amplifier Referred to Common
AD668	V+	Reference Amplifier
AD688	V−	Output Amplifier
AD689	V−	Output Amplifier
AD704/AD705/ AD706	V+	
AD707/AD708	V+, V−	Internal Feedforward Cap V+ to V− and Integrator V− to Output
AD711/AD712/ AD713	V−	
AD736/ AD737	V−, Common	External Integrator to V− Internal Feedforward V− to Common
AD741	V−	
AD744/AD746	V−	
AD766	V−	Output and Reference Amplifier
AD767	V−, Common	Output Amplifier Referred to V− and Reference Amp Referred to Common
AD840/AD841/ AD842	V+, V−	
AD843	V+, V−	
AD844/AD846	V+, V−	
AD845	V+	
AD847/AD848/ AD849	V+, V−	
AD1856/AD1860	V−	Output and Reference Amplifier
AD1864	V−	Output and Reference Amplifier
AD2700/AD2710	Common	Output Amplifier
AD2701	V−	Output Amplifier
AD2702/ AD2712	V−, Common	Output Amplifiers
AD7224/AD7225	V−	Output Amplifiers
AD7226/AD7228	V−	Output Amplifiers
AD7237/ AD7247	V+, Common	Reference Amplifier to Common Output Amplifier to Both V+ and Common
AD7245/ AD7248	V+, Common	Reference Amplifier to V+ Output Amplifier to Both V+ and Common
AD7569/AD7669	V−	All Amplifiers
AD7769	Common	All Amplifiers
AD7770	Common	All Amplifiers
AD7837/AD7847	V+	All Amplifiers
AD7840	V+, Common	Output Amplifiers to V+ Reference Amplifier to Common
AD7845	V+	All Amplifiers
AD7846	V+	All Amplifiers
AD7848	V+, Common	Output Amplifier to V+ Reference Amplifier to Common

Table I.

This collection of examples won't solve all your potential grounding problems. I hope that it will give you some good ideas how to prevent some of them, and it should also give you some of the "inside story" on I.C.'s which you can put to work in very practical ways. There is no general grounding method which will prevent all possible problems. The only generally applicable rule is attention to detail, and remember that you can always trust your mother, but . . . .

ANALOG DEVICES

AN-358
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Noise and Operational Amplifier Circuits

by Lewis Smith and D.H. Sheingold

INTRODUCTION

An operational amplifier's output always contains signals that could not have been predicted from the inputs and an accurately-known closed-loop transfer function, taking into account all impedance loads, known or unknown. These unwanted signals are called *noise*. Noise can be introduced into a circuit by the amplifier itself, by the components used in its feedback loop, and by the power supply; or it can be coupled or induced into the input, the output, the ground return, or the measurement circuit, from nearby – or in some cases, quite distant – sources.

The degree of the designer's interest in noise depends, of course, on (1) the desired resolution of the circuit in the band of interest, and (2) the avoidance of transmitting noise in frequency bands that are *not* of immediate interest. As operational amplifiers are used more extensively as preamplifiers and high-accuracy signal processors, resolution in operational amplifier circuits becomes of increasingly broad interest.

Recognizing the folly and futility of attempting to deal too comprehensively or rigorously with the complete range of noise phenomena, the present essay will attempt to provide tools for practical understanding, characterization, computation, and reduction of noise in the immediate environs of the amplifier. Coupled and induced noise in ground leads, output leads, etc., will be dealt with somewhat cavalierly by suggesting simply that attention be paid to shielding, lead dress, and ground circuit design – all of which belong in the realms of system design and "good circuit practice."[1]

From the standpoint of noise, the operational amplifier is uniquely qualified to serve in low-level and high-accuracy circuitry because

1. The amplifier transfer function may be chosen to pass only those frequency bands of interest.

2. The specific amplifier chosen for the job may be picked from among a wide range of types having differing noise characteristics, in order to obtain the most nearly ideal characteristics in the band of interest.

3. If the sources of noise are known and properly evaluated, the noise behavior of a wide range of amplifier circuits can usually be predicted to sufficient accuracy to permit a particular design to be done "on paper", with good probability of successful verification.

Reprinted from Analog Dialogue 3–1 1969

THE BASIC MODEL – VOLTAGE NOISE AND CURRENT NOISE (Figure 1)

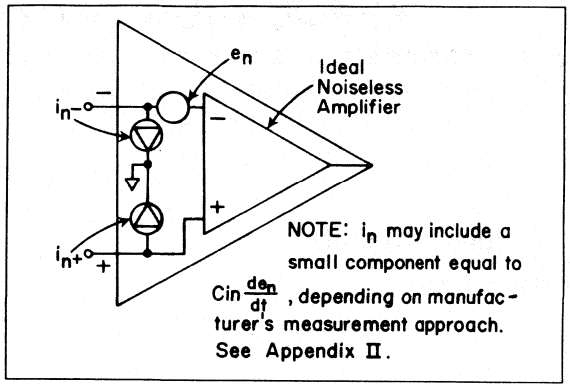

Figure 1. Voltage and Current Noise Model

A differential operational amplifier may be considered to be an ideal, noiseless amplifier, with noise current sources between each input terminal and common, and a noise voltage source effectively in series with one or the other input terminal. This model is quite similar to the E_{os}–I_{bias} model used for offsets[2], which should not be surprising, because E_{os} and I_{bias} can be considered as DC noise sources, modulatable by time, temperature, etc. For most practical purposes, noise voltage and noise current sources can be considered to be independent of one another. Ignoring the circuit and amplifier dynamics, for the moment, it is evident that, as with E_{os} and I_{bias}, the instantaneous voltage component of noise could be measured (Fig. 2) in a low-impedance circuit with high gain, and the instantaneous current component could be measured with a very large (ideally "noiseless") resistor. If there were no interaction between e_n and i_n, the output in the noise voltage measurement would be proportional to $(1 + R_2/R_1)$, and the output in the noise current measurement would be proportional to R_2 alone.

Note that the *instantaneous* sum of the two types of noise, appearing at the amplifier output, is

$$(1) \qquad e_0 = \left(1+\frac{R_2}{R_1}\right)e_n + R_2 i_n$$

and that the relative noise contributions of e_n and i_n are equal if

$$(2) \qquad \frac{e_n}{i_n} = \frac{R_2}{1+\frac{R_2}{R_1}}$$

that is, if the parallel combination of R_2 and R_1 is equal to the ratio of e_n to i_n. At impedance levels above e_n/i_n, current noise is dominant. The ratio of the RMS values of e_n and i_n is sometimes known as the "characteristic noise resistance" of the amplifier, in a given bandwidth, and it is a useful figure of merit for choosing an amplifier to match a given impedance level, or vice versa.

Noise coupled from *external sources* to the amplifier inputs can be considered as an additive voltage signal if the voltages and impedances are known, or as an additive current signal if its presence is determined by measurement at the amplifier site. This is shown compactly in Figure 3.

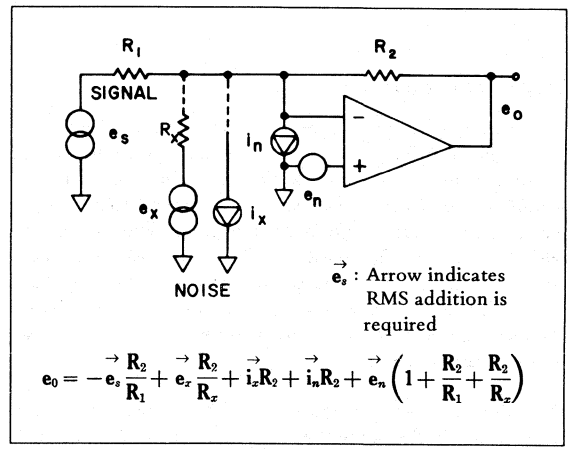

Figure 3. Contributions of Internal and External Noise Sources

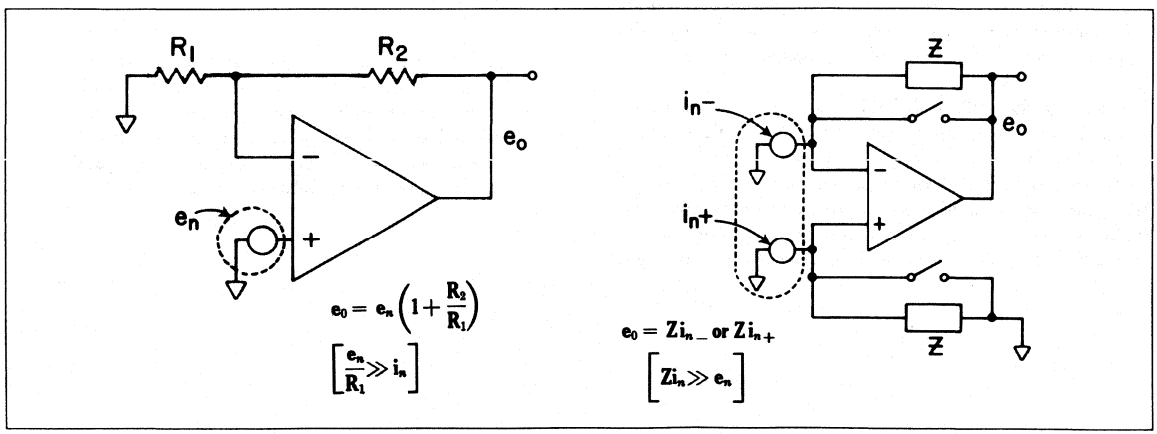

Figure 2. Rudimentary Measurement of e_n and i_n (Filters are used for narrow-band and spot noise measurements)

NOISE GAIN AND SIGNAL GAIN

Figure 4 shows the basic feedback model of an inverting amplifier with a number of resistive inputs. For large values of loop gain ($A\beta$), the noise gain for voltage noise is essentially $1/\beta$. If $A\beta$ is not much greater than unity, one should use the more exact expression

$$e_0 = \frac{1}{\beta}\left[\frac{1}{1+\frac{1}{A\beta}}\right] e_n \qquad (3)$$

The comparable current noise expression is

$$e_0 = i_n Z_f \left[\frac{1}{1+\frac{1}{A\beta}}\right] \qquad (4)$$

It should be noted that, for passive feedback elements, $1/\beta$ is never less than unity, and it also is larger than the closed-loop gain for any of the input signals. Therefore the total spectrum of e_n will appear at the output, with *at least* unity gain, despite less-than-unity signal gains, or narrow signal bandwidths. It should also be noted that, in the general case where A and β are both dynamic expressions, if the loop gain has substantially more than 90° phase shift, and the amplifier is in consequence lightly-damped in the vicinity of the frequency at which $A\beta = 1$, the noise gain at that frequency may peak at substantially greater than unity, *even if the signal gain rolls off smoothly at a very much lower frequency*. See Fig. 5 for a flagrant, easy-to-understand example.

COMBINING NOISE SIGNALS

This topic will be dealt with in some depth shortly. However, it is desirable to bear in mind the basic notion that RMS values of signal or noise voltage or current from essentially uncorrelated sources (for example, noise from different sources, or noise from different portions of the frequency spectrum of a given source) are combined by computing the square root of the sum of the squares. It should be evident that larger quantities will be emphasized and smaller quantities suppressed. For example, if X = 3Y the error incurred by simply ignoring Y is only about 5%. ($\sqrt{1^2 + 3^2} = \sqrt{10} \cong 3.16 \cong 3.0 \times 1.05$).

TO THE READER: The sections that follow are necessary in order to construct a coherent approach to a technique for predicting the "total RMS noise" of an operational amplifier circuit quickly and easily. The material is not the easiest of reading, and it covers a wide range of topics, the exact place of which in the final scheme may be somewhat difficult to second-guess. Nevertheless, it is a necessary part of the background, and, if read with patient attention, may prove interesting and informative. If you would like to get an idea of the resulting technique, look ahead to page 10.

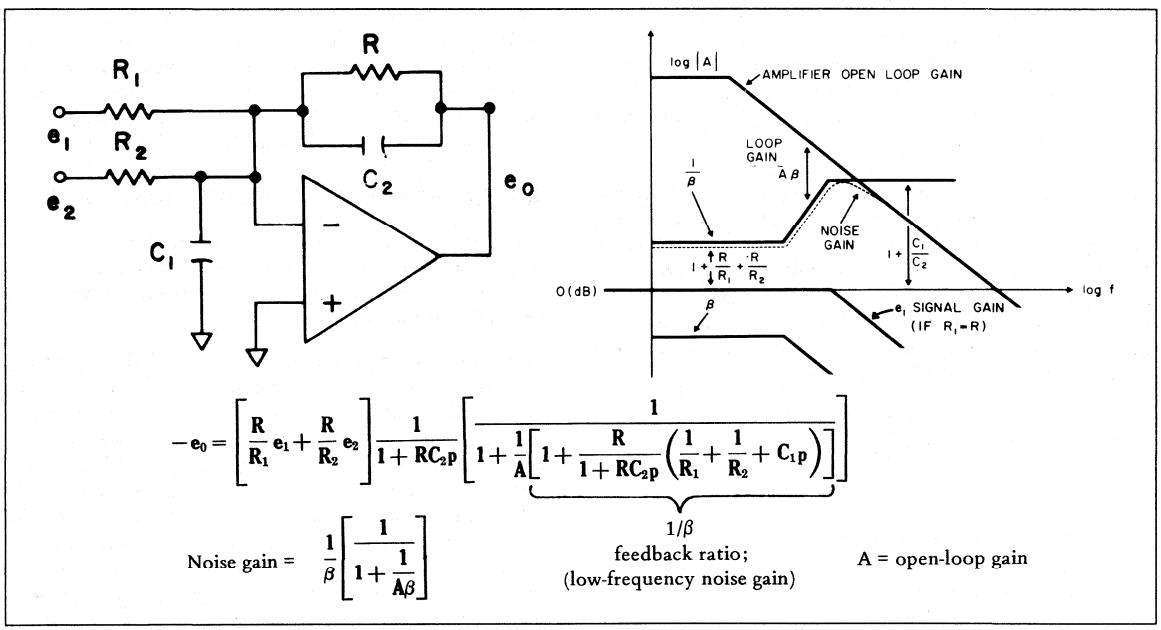

Figure 4. Feedback Circuit Relationships

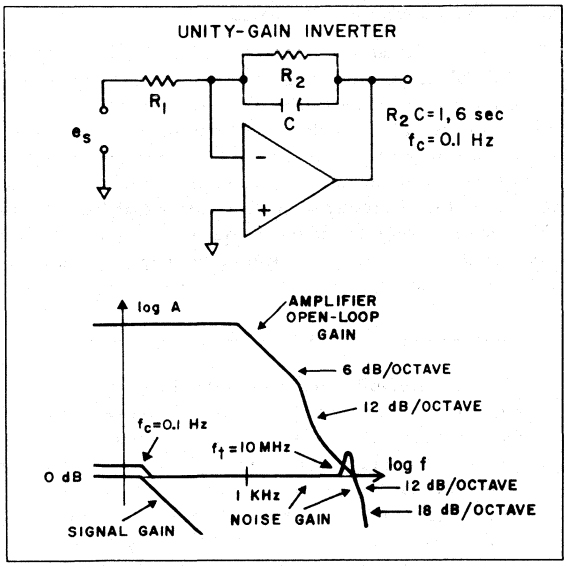

Figure 5. Noise Bandwidth vs. Signal Bandwidth

HOW NOISE IS CHARACTERIZED

Periodically-recurring noise may be described in terms of repetition rate, waveform, and magnitude (e.g., chopper noise.) Irregularly-recurring noise may be described in terms of its waveshape and magnitude only, since the repetition is not regular (to some extent, popcorn noise). Aperiodic noise not having a recurring waveshape (and this includes all the types of random noise to be described below) is usually dealt with in terms of some of its statistical properties: RMS value, peak value, and frequency content.

RMS Value. It is a characteristic of most kinds of random noise that if the averaging interval is made sufficiently long, the RMS value obtained is substantially repeatable. Therefore, RMS value in the bandwidth of interest, averaged over a sufficiently long interval, is a useful and meaningful way of characterizing such random noise. It is by far the most convenient for correlating noise numbers between vendors and users. The definition of RMS value of voltage is given by

$$E_{rms} = \sqrt{\frac{1}{T}\int_0^T e^2\,dt} \qquad (5)$$

where

E_{rms} = RMS voltage value
T = Interval of observation
e = Instantaneous noise voltage

Substituting the instantaneous value of current, *i*, gives I_{rms}, the RMS current value. When performing RMS measurement, either a "true RMS" meter must be used, or the reading of an AC average (sine wave "RMS"-calibrated meter) should be multiplied by a factor of 1.13.

Peak Values. Noise may also be characterized in terms of the difference between the largest positive and negative peak excursions observed during an arbitrary interval. For some applications, in which peak-to-peak noise may limit system performance, peak-to-peak measurement may be essential. Nevertheless, because noise is for all practical purposes Gaussian in amplitude distribution, and hence the highest noise amplitudes have the smallest (yet non-zero) probability, peak-to-peak noise is difficult to measure repeatably. Because RMS values are so easy to measure repeatably, and are the most usual form for presenting noise data non-controversially, the table below should be useful for estimating the probabilities of exceeding various peak values, given the RMS values.

Nominal "peak-to-peak"	Percentage of time that noise will exceed nominal "peak-to peak" value
2.0 x RMS	32%
3.0 x RMS	13%
4.0 x RMS	4.6%
5.0 x RMS	1.2%
6.0 x RMS	0.27%
6.6 x RMS	0.10%
7.0 x RMS	0.046%
8.0 x RMS	0.006%

Peak-to-peak vs. RMS (Gaussian distribution)

The casually-observed peak-to-peak noise varies between 3 x RMS and 8 x RMS, depending on the patience of, and amount of data presented to, the observer. Oscilloscope traces should be observed at high intensity, because considerable averaging occurs at low intensity, which will produce apparent results closer to the RMS value. Peak amplitude distribution analyzers are becoming available, for automating these measurements.

INTERFERENCE NOISE VS. INHERENT NOISE

There are two basic classifications of noise for a given circuit; they are *interference noise,* which is noise picked up from outside the circuit, and *inherent noise,* which is noise arising within the circuit itself.

Interference noise may be periodic, irregularly recurring, or essentially random, and it ordinarily may be reduced substantially (or forestalled) by taking precautions to minimize electrostatic and electromagnetic pickup from power sources at line frequencies and their harmonics, radio broadcast stations, arcing of mechanical switches, and current or voltage spikes resulting from switching in reactive circuits. Such precautions may include filtering, decoupling, electrostatic and electromagnetic shielding of leads and components, use of guarding potentials, elimination of ground loops, physical reorientation of leads and components, use of damping diodes across relay coils, choice of low circuit impedances where possible, and choice of power and reference supplies having low noise. Interference noise resulting from vibration may be reduced by proper mechanical design. A table outlining some of the sources of interference noise, their typical magnitudes, and ways of dealing with them is shown in Figure 6.

External Source	Typical Magnitude	Typical Cure
60Hz power	100pA	Shielding, attention to ground loops, isolated power supply
120 Hz supply ripple	3μV	Supply filtering
180Hz magnetic pick-up from saturated 60Hz transformers	0.5μV	Reorientation of components
Radio broadcast stations	1mV	Shielding
Switch-arcing	1mV	Filtering of 5 to 100MHz components, attention to ground loops and shielding
Vibration	10pA (10 to 100Hz)	Proper attention to mechanical coupling; elimination of leads carrying large voltages near input terminals (Note: $i = \frac{d(Cv)}{dt} = C\frac{dv}{dt} + v\frac{dC}{dt}$)
Cable vibration	100pA	Use low noise (carbon-coated dielectric?) cable
Circuit boards	0.01 to 10 pA/ $\sqrt{cps}$ below 10Hz	Clean boards thoroughly, use teflon insulation where needed, and guard well

Figure 6. Typical Sources of Interference Noise

But even if all interference noise is removed, *inherent noise* will still be present. Inherent noise is usually random in nature, most often arising in resistances and semiconductor elements, such as transistors and diodes. (An example of a *non*-random inherent noise component is chopper noise in chopper-stabilized amplifiers.) Random noise arising within resistances is known as *Johnson noise,* (or thermal noise). Random noise arising within semiconductor elements may be one of the following three types: *Schottky* (or shot) *noise, flicker noise* (1/f noise), and *popcorn noise.*

COMMON FORMS OF RANDOM NOISE

Johnson noise. Thermal agitation of electrons in the resistive portions of impedances results in the random movement of charge through those resistances, causing a voltage to appear corresponding to the instantaneous rate of change of charge (i.e., current) multiplied by the appropriate resistance. Ideally-pure reactances are free from Johnson noise.

The Johnson noise voltage within a bandwidth B generated by thermal agitation in a resistance is given by

$$(6) \qquad E_{rms} = \sqrt{4kTRB}$$

where

k = Boltzmann's constant = 1.374 x 10^{-23} Joules/° Kelvin
T = Absolute temperature (degrees Kelvin)
R = Resistance (ohms)
B = Bandwidth (cycles per second)

At room temperature, with more convenient units, this expression becomes

$$(7) \quad E_n = 0.13\sqrt{R \times B} \cong \frac{1}{8}\sqrt{R \times B} \text{ microvolts}$$

where

R = Resistance (megohms)
B = Bandwidth (cycles per second)

Johnson noise is quite often expressed as an equivalent current

$$(8) \quad I_n = \frac{E_n}{R} = 0.13\sqrt{\frac{B}{R}} \cong \frac{1}{8}\sqrt{\frac{B}{R}} \text{ picoamperes}$$

Ordinarily, Johnson noise is a less important source of noise *within the amplifier* than the noise produced within semiconductors; however it is the primary source of noise contributed by the signal source and the feedback circuitry, where resistance values are apt to be higher. In some situations, the Johnson noises of the components in the external circuit are completely dominant. For example, if the source resistance in a particular application is to be 10 megohms, there is little point to selecting an amplifier with *low voltage noise,* since the Johnson noise from the 10 megohm resistor will be the chief source of noise at the output, unless the amplifier's input *current noise* is high. In this example, the resistor noise in a 1kcps bandwidth would be 13 microvolts RMS, corresponding to current noise of 1.3 picoamperes.

Schottky noise. Shot noise arises whenever current is passed through a transistor junction. The noise is normally expressed as a current, which will, of course pro-

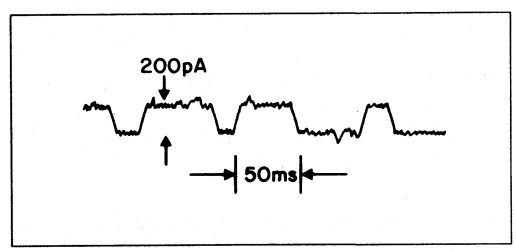

Figure 7. Typical "Popcorn" Noise

duce voltage drops in impedance, such as transistor emitter resistance. A convenient equation for shot noise is

$$(9) \qquad I_n = 5.7 \times 10^{-4} \sqrt{IB} \text{ picoamperes}$$

where

I = junction current in picoamperes
B = bandwidth of interest in cps

In a typical operational amplifier circuit using bipolar input transistors, the input transistor base current, flowing through the base-emitter junction, produces a Schottky noise component, which is a part of the amplifier's equivalent noise *current* source. Other P-N-junction-generated noise currents (internal to the operational amplifier), when divided by appropriate transconductances, will create an equivalent noise *voltage* at the input.

Flicker noise (1/f noise.) In the frequency range below 100Hz, most amplifiers exhibit another noise component that dominates over Johnson and Schottky components and becomes the chief source of error at these frequencies. Flicker noise is thought to be a result of imperfect surface conditions on transistors. Carbon composition resistors, if they carry much current, may contribute noise similar to flicker noise of transistors. (It is good engineering practice to use metal film or wirewound resistors wherever significant currents flow through the resistor and low noise at low frequencies is a definite requirement).

Flicker noise does not have an equal contribution at each frequency. The spectral noise density (to be defined below) of this type of noise typically exhibits a −3dB per octave slope.

Popcorn noise. Some transistors, especially those of integrated circuit monolithic construction (from some manufacturers), jitter erratically between two values of h_{fe} causing additional base current noise of the form shown in Fig. 7. The premium paid for most high-performance amplifier types includes tests to weed out units having transistors or IC's that exhibit this effect.

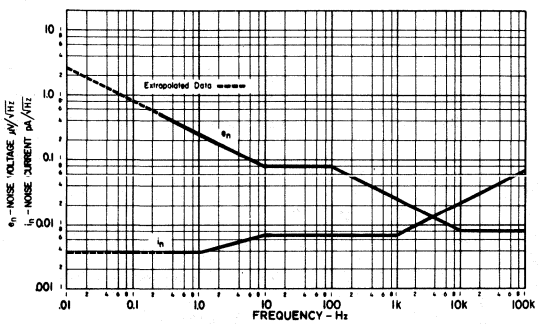

Figure 8. Noise Spectral Density. Analog Devices Model 144 (Economical FET Amplifier)

IN THE NEXT ISSUE

Settling Time. Based in part on a paper presented at NEREM 1968, this article will discuss the components of delay and error that contribute to the speed limitations of high-precision operational amplifier circuits, and their significance to the designer and tester of practical circuits. It will include a discussion of "glitch" phenomena in Digital-to-Analog Converter circuits. It will also present some surprising conclusions about the effects on settling time of deviation from 6 dB/octave slope in the open loop response of operational amplifiers employed in precision buffer circuits.

NOISE DENSITY SPECTRUM

Noise exists in all parts of the frequency spectrum, and the noise contribution of a resistor or amplifier varies with the range of frequency over which the observation is made. Probably the most useful manner of describing a noise characteristic is a spectral noise density plot, in which noise is shown as a function of frequency, usually on log-log axes.

Spectral noise density, e_n, at a given frequency f, is defined as the square root of the rate of change of mean-square voltage with frequency. Conversely, the RMS value in a given band is the square root of the definite integral of the square of spectral noise density over a given band of frequencies.

$$(10) \qquad e_n^2 \equiv \frac{d}{df}\left(E_n\right)^2 \qquad (11) \qquad E_n = \sqrt{\int_{f_1}^{f_2} e_n^2 \, df}$$

Inasmuch as average power is proportional to E^2, a simpler way of expressing this concept is: e_n^2 is proportional to the variation of noise power per cycle per second. Thus, e_n is expressed as the noise voltage per (square) root-cycle (per second.

From equation (11), we see that a lower and an upper frequency limit, f_1 and f_2, are required, as well as a knowledge of the manner in which e_n varies with frequency, in order to evaluate the integral and calculate the actual RMS noise voltage. A typical spectral density plot (both e_n and i_n) for Analog Devices types 144 is shown in Figure 8.

COMMON DENSITY SPECTRUM CONFIGURATIONS

White noise. In a white noise spectrum, e_n is constant as a function of frequency. Thus, the RMS noise value as measured in a given bandwidth via an ideal sharp-cutoff filter would be

$$(12) \qquad E_n = \sqrt{\int_{f_1}^{f_2} e_n^2 \, df} = e_n \sqrt{f_2 - f_1}$$

Because $f_2 - f_1$ defines a bandwidth (cf. B in equations 6,7,8 9), it is evident that both Johnson noise and Schottky noise are *white noise,* and that e_n, for Johnson noise, is $0.13\sqrt{R}$ microvolts per root cycle, or $0.13x\sqrt{1/R}$ picoamperes per root cycle (R given in megohms), and Schottky noise is $5.7 \times 10^{-4}\sqrt{I}$ picoamperes per root cycle (I also in picoamperes).

If f_1 is less than 10% of f_2, a simple expression for all the white noise from f_2 down to "DC", with less than 5% error, is

$$\textbf{(13)} \qquad E_n = e_n\sqrt{f_2}$$

Pink noise. A generic term applied to ideal 1/f noise, for which e_n is exactly proportional to $\sqrt{1/f}$, is *pink noise.* This may be expressed

$$\textbf{(14)} \qquad e_n = K\sqrt{\frac{1}{f}}$$

K is the value of e_n at f = 1 Hertz

The RMS noise in the band between f_2 and f_1 may be computed by substituting equation (14) in equation (11)

$$\textbf{(15)} \qquad E_n = K\sqrt{\int_{f_1}^{f_2}\frac{df}{f}} = K\sqrt{\ln\left(\frac{f_2}{f_1}\right)}$$

On a log-log plot of e_n vs. frequency, the slope will be −3dB per octave (e_n^2 vs. frequency would have a −6dB per octave slope). Because the RMS value of pink noise depends on the ratio of the frequencies defining the band of interest, every octave or decade of pink noise will have the same RMS noise content as every other octave or decade.

It is instructive to compute the RMS pink noise in the nine-decade realm below 1 Hz (10^{-9} Hz ≅ 1 cycle per 30 years). If the RMS value of pink noise in the decade from 0.1 to 1 Hz is 1 microvolt, then the total RMS noise over all nine decades is $\sqrt{9 \times 1^2} = 3\mu V$.

This tells us that, as the lower frequency f_1, goes to extremely small values ("DC"), RMS pink noise (and thus ideal flicker noise) will have less significance than drift caused by environmental factors, component aging, or perhaps even component life.

Spot noise. If we divide the spectrum of frequency into sufficiently narrow "spots" or intervals, Δf_1, Δf_2, Δf_3, etc., so that we may consider e_n (or its "average" value) essentially constant, then an incremental approximation may be used to evaluate the RMS noise, E_n, using equation (11). The RMS noise in the total interval to be computed is simply the root-sum-of-the-squares of the noise in the incremental intervals

$$\textbf{(16)} \qquad E_n = \sqrt{e_{n_1}^2\Delta f_1 + e_{n_2}^2\Delta f_2 + \cdots}$$

If the intervals are all equal = Δf, equation (16) becomes

$$\textbf{(17)} \qquad E_n = \sqrt{\Delta f}\sqrt{e_{n_1}^2 + e_{n_2}^2 + \cdots}$$

Filter skirt errors. Our discussion of frequency content has so far assumed that it is possible to combine noise source contributions, band by band, with perfect sharp-cutoff filters. In actual circuits, however, a quite common filter response is that of a single time constant (either lead or lag). Such a filter will have significant response in the band beyond its nominal cutoff frequency. For example, equation (13) for RMS white noise in the whole band below f_2 must be multiplied by 1.26 to account for the white noise passed at frequencies higher than f_2 by a first-order lag filter, having cutoff frequency f_2. In effect, the "noise bandwidth" is $f_2\pi/2$ and, ignoring the noise passed in frequencies under the filter skirt (i.e., beyond the cutoff frequency) will result in a 26% error. However, it should be noted that errors due to ignoring the skirts of higher-order filters are considerably less. The RMS contribution of the 6dB per octave skirt alone is $0.76e_n\sqrt{f_c}$.

NOTE TO THE READER: Thank you for your patience. You have now arrived at the point of this necessarily lengthy discussion: a graphical technique for easily predicting the noise spectrum and the total noise of an operational amplifier circuit, using the spectral density plots for voltage and current noise, the closed-loop gain and impedance level of the feedback amplifier circuit, and certain approximations made understandable by the background material.

BRIEF REVIEW OF RELEVANT POINTS

1. Pink noise contributes equal increments of RMS noise over each octave or each decade of the spectrum. Each increment will be 1.52K per decade, or 0.83K per octave, where K = e_n or i_n at 1 Hz.

2. Bandwidth for white noise is substantially equal to the higher frequency, if one is considering bandwidths greater than 1 decade.

3. Because of root-sum-of-the-squares addition, the greater of the RMS values of two uncorrelated noise signals will be substantially equal to their sum, if the greater is at least 3x the lesser (i.e., if 20 log-of-their-ratio is greater than 10dB).

4. Noise in different portions of a random spectrum typical of amplifier circuits is uncorrelated, and can hence be added by root-sum-of-the-squares.

THE "PINK NOISE TANGENT" PRINCIPLE

Consider Fig. 9, which shows an arbitrary plot of output noise spectrum on a log voltage vs. log frequency scale. It is a purely arbitrary choice, chosen simply as an example. Consider that it might have been arrived at by simple addition on a log scale (in effect obtaining the product) of e_n and an amplifier's closed loop gain (i.e., true noise gain).

First, let us compute the noise in each portion of the spectrum, using point (4)

Region	RMS Noise	Why
1	22μV	Pink noise, 2 decades, $\sqrt{2}$x1.52x10
2	100μV	White noise, 2 decades, 10x$\sqrt{100}$
3	152μV	Pink noise, 1 decade, 1.52x$\sqrt{100}$
4	72μV	6dB/octave skirt, 0.76x3x$\sqrt{1000}$
5	42μV	White noise, 0.3x$\sqrt{20,000}$
6	40μV	6dB/octave skirt, 0.76x0.3$\sqrt{30,000}$

The RMS total of these noises is

$$\sqrt{152^2 + 100^2 + 72^2 + 42^2 + 40^2 + 22^2} = 205 \text{ microvolts}$$

Just as a matter of interest, it is worth noting that the root-square sum of those portions nearest to the point of tangency of the K = 100 pink noise characteristic is

$$\sqrt{152^2 + 100^2 + 72^2} = 196 \text{ microvolts}$$

Before making our point, let us look at the example of Fig. 10, which shows another response. The noise bookkeeping is as follows:

Region	RMS Noise	Why
1	8μV	Pink noise, 3 decades, K=3.16$\sqrt{3}$x3.16 x 1.52
2	126μV	White noise and skirt (6dB/octave) 1.26 x 1.0 x $\sqrt{10^4}$
3	45μV	White noise, 0.1 x $\sqrt{200,000}$
4	42μV	6dB/octave skirt, 0.76 x 0.1$\sqrt{300,000}$

The RMS total of these noises is

$$\sqrt{126^2 + 45^2 + 42^2 + 8^2} = 140 \text{ microvolts}$$

Note that there is less than 10% error by assuming that all the noise is in Region 2, and negligible error in omitting Region 1 entirely. Furthermore, if one makes the purely arbitrary assumption that *all* the noise present can be approximated by one decade of pink noise in the region of tangency (10kHz), it would amount to 1.52 x 100 = 152 microvolts, which is less than 10% error in the *conservative* direction.

The point is this: If a characteristic −3dB per octave pink noise slope is lowered until it is tangent to the noise output characteristic of the amplifier, the only significant contribution to total noise output will come from those portions of the amplifier noise characteristic in the immediate vicinity of that pink noise slope. Any portions of the amplifier noise characteristic that are substantially (i.e., typically more than 10dB) below the pink noise slope will contribute insignificantly. The explanation for this is that the pink noise slope is the locus of equal per-octave (or per decade) contributions to total noise. In the region of tangency, it will be seen, lies the maximum noise contribution of the amplifier circuit. The noise contributions in any other comparable interval must be less, and in intervals more than typically 10dB below the −3dB per octave slope, they will be insignificant.

A TYPICAL EXAMPLE

Consider the circuit of Fig. 11. It is a summing amplifier with gains of 10 and 100, using a 1 megohm feedback resistor, paralleled by 160pF. The amplifier's open loop DC gain is 100dB, and f_t is 1MHz.

COMPUTE NOISE GAIN

Figure 11 shows the amplifier's assumed open loop gain-frequency plot, and the attenuation of the feedback network ("noise gain"). The loop gain (Aβ) is the difference between the two curves. The attenuation of the feedback network can be shown to be a constant attenuation of x111 up to 1kHz, and a 6dB/octave rolloff down to unity (i.e., 0 dB).

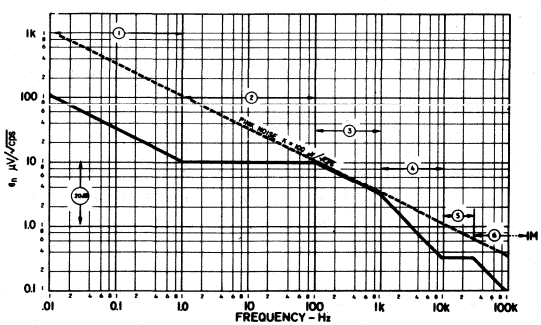

Figure 9.

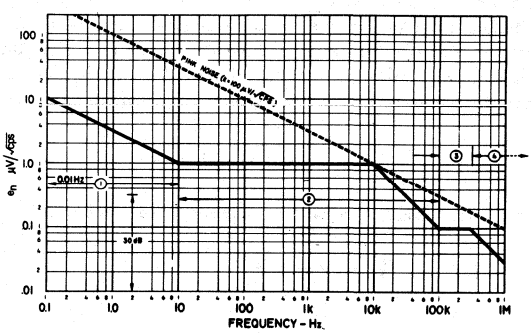

Figure 10.

First, the usual stability analysis is performed. In this case, it is pretty simple: The loop gain, $A\beta$, passes through a magnitude of unity with about 90° phase shift (90° from the amplifier, none from the feedback network) which implies not only stability, but absence of peaking as well. Thus the noise gain, after dropping to unity at about 100kHz, will resume its rolloff at about 1 MHz.

IDENTIFY AND COMPUTE CONTRIBUTIONS OF NOISE SOURCES

The sources of noise are: Amplifier e_n and i_n, and Johnson noise in the three resistors. Figure 12 shows (to begin with) a plot of i_n and e_n on a log-log scale. (Dashed and lower solid). We now plot the effective output noise contribution of all noise sources, in order to establish those that are most significant, the frequency bands in which they will appear, and to compute the total RMS "inherent" noise of the circuit, according to the method just described.

Amplifier voltage noise (e_n). At the output, the noise spectrum will consist of e_n, multiplied by the noise gain. Inasmuch as e_n and noise gain are both plotted on compatible log-log scales, we have simply to add the two spectra. (Upper solid line).

Amplifier current noise (i_n). At the output, the noise spectrum will consist of i_n, multiplied by the feedback impedance (1 megohm to 1 kHz, with a 6dB per octave rolloff thereafter.) This will consist simply of a curve paralleling the current curve at low frequencies, of magnitude $i_n R_f$, and departing from it at a 6dB per octave rate below the break frequency. (Coarse dotted line)

Resistor noise. It is illustrative to show the contribution of each resistor separately (since their noises, being uncorrelated, will be combined in RMS fashion). In the present example, it will be found convenient to treat each resistor as a noise *current* source. The current noise from each of the three resistors will be multiplied by the feedback impedance, in the same way as the amplifier's i_n. (dotted, starred, and light dashed lines)

Establish location of tangent pink noise line. Having plotted the spectral density of computed output noise from each of the five sources, a −3dB per octave slope is lowered until it touches one of the curves. This will be seen to be the e_n (x noise gain) curve, at 1 kHz. Qualitatively, we observe immediately the following:

The bulk of the random noise will occur in the vicinity of 1kHz, mostly contributed by the amplifier, but with some contribution from the 10k resistor.

There will be a minor additional contribution at 1 MHz.

If we (in a purely arbitrary manner) approximate all the noise by the decade of pink noise ($K = 160\mu V/\sqrt{cps}$) in the vicinity of 1kHz, the total RMS noise would be 1.52 x 160 = 244 microvolts. That this is quite conservative will be seen by the actual calculation, from the curves

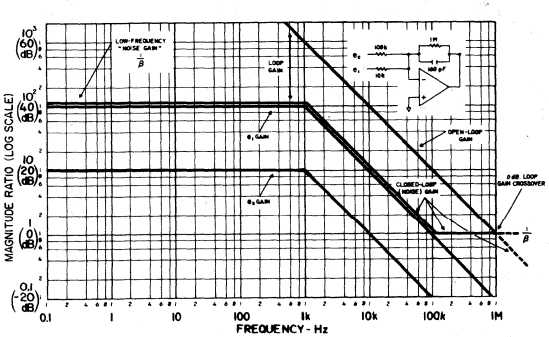

Figure 11. Closed-loop Gain Relationships

COMPUTE THE TOTAL NOISE

Region	E_{rms}	Method
100kHz up (amplifier)	76μV	White noise – 6dB/octave rolloff $1.26 \times 0.06 \times 10^3$
1kHz to 100kHz (amplifier)	120μV	White noise – 6dB/octave skirt $0.76 \times 5 \times \sqrt{1000}$
200Hz to 1kHz (amplifier)	116μV	Spot noise, summation of 100cps spots $10\sqrt{\Sigma e_n^2}$
All frequencies (resistor)	52μV	10k resistor, white noise, 1kHz rolloff 1.26 x 1.3 1000

The above constitute all the significant contributions, and add up, RMS fashion, to

$$\sqrt{76^2 + 120^2 + 116^2 + 52^2} = \textbf{190 microvolts RMS}$$

It may be a useful exercise for the reader to convince himself that all other sources and bands contribute in a minor or negligible degree.

13

FURTHER DISCUSSION OF THIS EXAMPLE

The total random noise of 190 microvolts may be referred back to either of the inputs to determine the signal-to-signal noise ratio. It will be seen that, referred to e_1, (gain of 100) the noise is 1.9 microvolts RMS, and, referred to e_2 (gain of 10), the noise is 19 microvolts RMS.

The signal "cutoff" frequency of this circuit is 1kHz, and it will be seen that, because of the absence of peaking, the noise contribution at 1MHz, though significant, is a relatively small contributor to the total. However, if an amplifier having less phase margin were used, and if the signal bandwidth were also less, (larger feedback capacitor used) it can be seen that the dominant noise could be in the region of 1MHz, *even though this is well beyond the passband of interest.* If this is the case in a given application, and an amplifier having narrower bandpass is not available, it may be desirable to follow the amplifier by a low-impedance low-pass R-C or L-C filter, to greatly reduce the high-frequency noise. If this is done, the post-filter transfer function is applied to the output noise spectrum, and the pink noise slope is lowered until it is once again tangent to one of the curves. An important suggestion to the circuit designer is: *Do not use an amplifier having greater bandwidth than is necessary for the application (if feasible).*

NOISE FIGURE

It will be noted that we have not yet discussed an amplifier's "noise figure." Noise figure is a measure of the additional noise contributed by the amplifier over-and-above that of the source resistance. It is often computed by the formula

$$(18) \qquad \text{N.F.} = 10_{\log} \frac{E_n^2 + I_n^2 R_s^2 + 4kTR_sB}{4kTR_sB}$$

where R_s is the source resistance. Noise Figure is expressed in dB, and is zero for an ideal noiseless amplifier. For a given amplifier, N.F. is minimum if $R_s = E_n/I_n$, a quantity known as the "optimum noise resistance." It should be noted that N.F. depends on bandwidth, and that $R_{optimum}$ may differ from frequency band to frequency band. A better way to express Noise Figure may be in terms of the actual closed-loop configuration, and the computed RMS noise. For the example given above, using this definition

$$\text{N.F.} = 10_{\log_{10}} \frac{\text{(Total output noise)}}{\text{(Source resistor noise)}}$$

$$= 10_{\log_{10}} \frac{(190)^2}{(1.26)^2(1.3)^2(1000)} = 11.3\text{dB}$$

Note that, for this amplifier, $R_{optimum}$ is

$$\frac{0.12\mu V/\sqrt{cps}}{10pA/\sqrt{cps}} = 12k\Omega$$

in the low frequency band, and rises to

$$\frac{0.05\mu V/\sqrt{cps}}{0.12pA/\sqrt{cps}} = 400k\,\Omega \text{ at } 1kHz.$$

Note also that the best noise figure does not always produce the lowest noise. This concept is perhaps of greatest usefulness when considering non-inverting amplifier configurations. When experimenting with various circuit impedance levels, it is good practice to add a small signal as a reference, in order to aid in maximizing signal-to-noise ratio, rather than merely minimizing output noise.

Footnotes:

[1]See Noise Bibliography, page 16

[2]See *Operational Amplifiers*, part IV: "Offset and Drift in Operational Amplifiers," Analog Devices, Inc. 1966

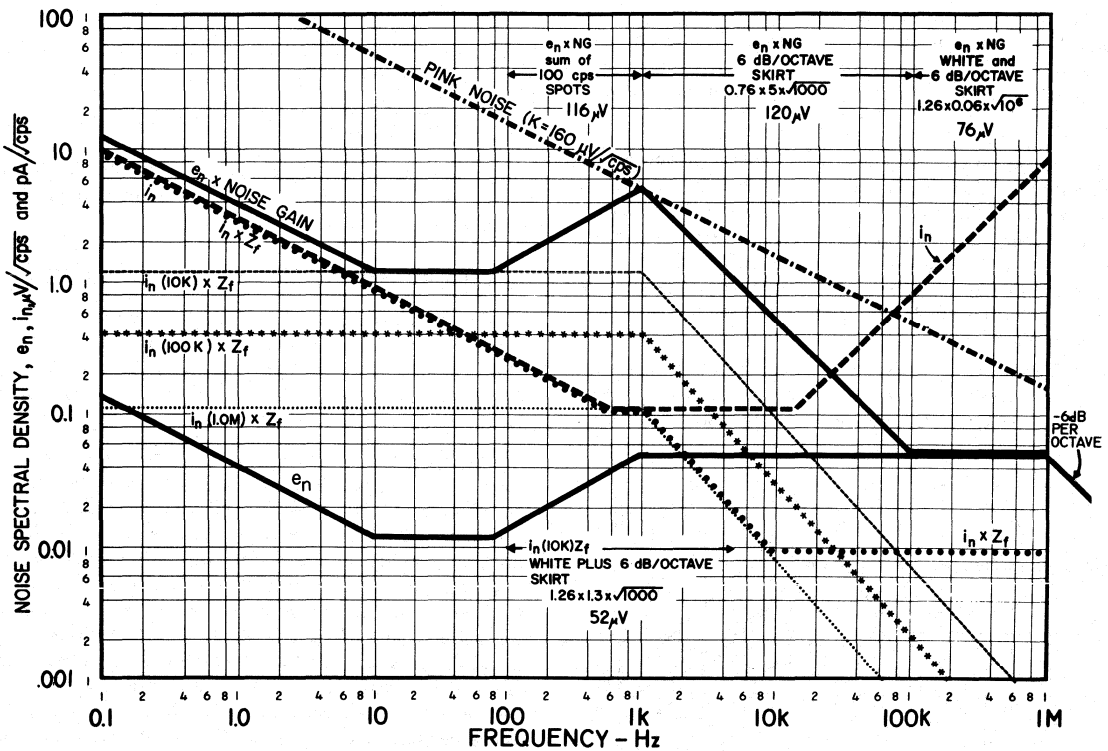

Figure 12. Closed-loop Noise Spectra

ANALOG DEVICES

AN-102
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Very Low Noise Operational Amplifier

by Derek F. Bowers

APPLICATIONS

- High Precision Instrumentation
- Microphone Preamplifier
- Tape-Head Preamplifier
- Strain-Gage Amplifier

FEATURES

- Very Low Voltage Noise ..... 500pV/$\sqrt{Hz}$
- High Gain-Bandwidth Product 150MHz
- High Open-Loop Gain ..... 3×10^7
- High CMRR 130dB
- Very Low Offset Voltage Drift ..... <0.1μV/°C

GENERAL DESCRIPTION

In situations where low output, low-impedance transducers are used, amplifiers must have very low voltage noise to maintain a good signal-to-noise ratio. The design presented in this application note is an operational amplifier with only 500pV/$\sqrt{Hz}$ of broadband noise. The front end uses SSM-2210 low-noise dual transistors to achieve this exceptional performance. The op amp has superb DC specifications compatible with high-precision transducer requirements, and AC specifications suitable for professional audio work.

PRINCIPLE OF OPERATION

The design configuration in Figure 1 uses an OP-27 op amp (already a low-noise design) preceded by an amplifier consisting of three parallel-connected SSM-2210 dual transistors. Base spreading resistance (R_{bb}) generates thermal noise which is reduced by a factor of $\sqrt{3}$ when the input transistors are parallel connected. Schottky noise, the other major noise-generating mechanism, is minimized by using a relatively high collector current (1mA per device). High current ensures a low dynamic emitter resistance, but does increase the base current and its associated current noise. Higher current noise is relatively unimportant when low-impedance transducers are used.

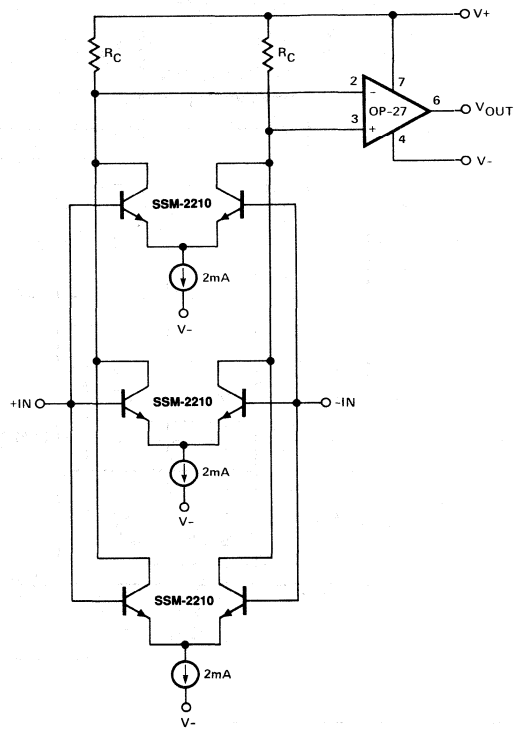

Simplified Schematic for Very Low Noise Operational Amplifier

Figure 1: Simplified Schematic

13

CIRCUIT DESCRIPTION

The detailed circuit is shown in Figure 2. A total input-stage emitter current of 6mA is provided by Q4. The transistor acts as a true current source to provide the highest possible common-mode rejection. R_1, R_2, and R_3 ensure that this current splits equally among the three input pairs. The constant current in Q4 is set by using the forward voltage of a GaAsP light-emitting diode as a reference. The difference between this voltage and the base-emitter voltage of a silicon transistor is predictable and constant (to within a few percent over the military temperature range. The voltage difference, approximately 1V, is impressed across the emitter resistor R_{12} which produces a temperature-stable emitter current.

R_6 and C_1 provide phase compensation for the amplifier and are sufficient to ensure stability at gains of ten and above.

R_7 is an input offset trim that provides approximately ±300μV trim range. The very low drift characteristics of the SSM-2210 make it possible to obtain drifts of less than 0.1μV/°C when the offset is nulled close to zero. If this trim is not required, the R_4, R_7, and R_8 network should be omitted and R_5/R_9 connected directly to V+.

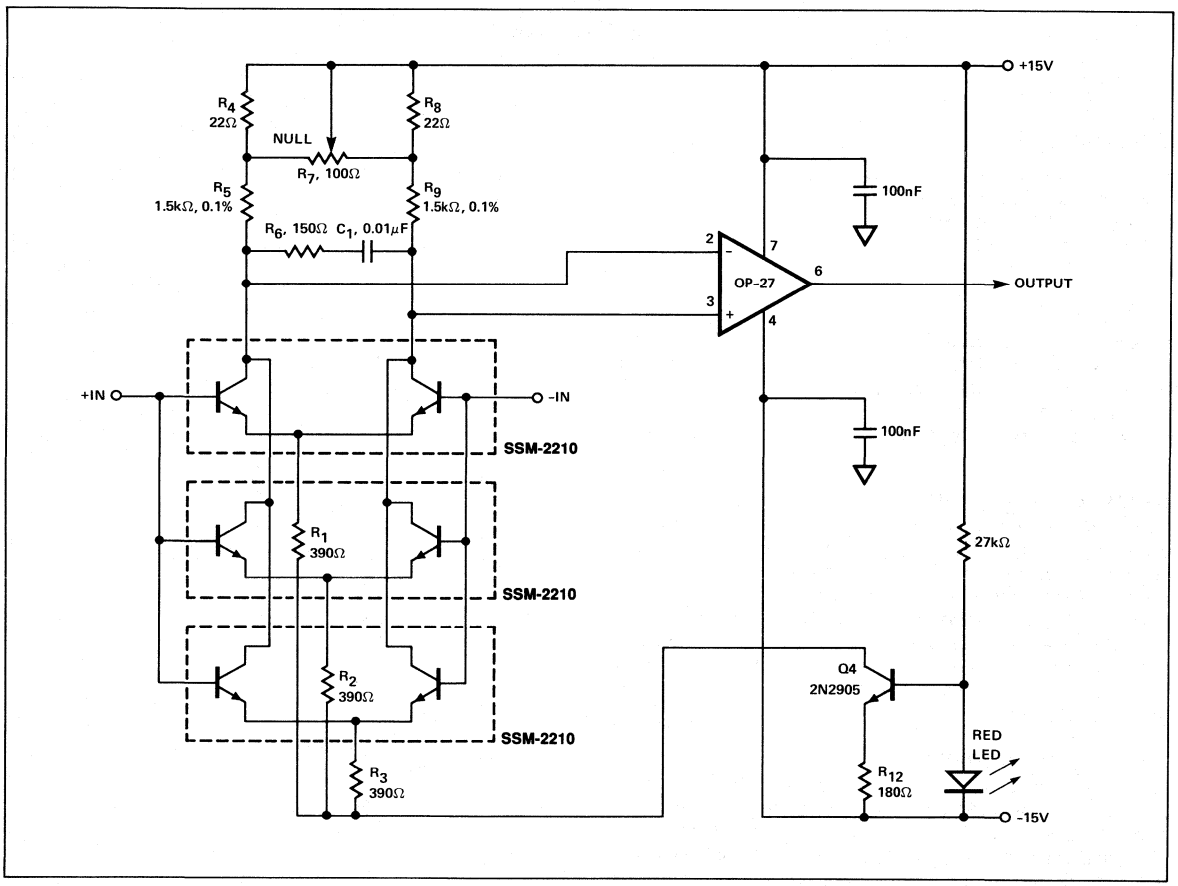

Figure 2: Complete Amplifier Schematic

AMPLIFIER PERFORMANCE

The measured performance of the op amp is summarized in Table 1. Figure 3 shows the broadband noise spectrum which is flat at about 500 pV/$\sqrt{Hz}$. Figure 4 shows the low-frequency spectrum which illustrates the low 1/f noise corner at 1.5Hz. The low-frequency characteristic in the time domain from 0.1Hz to 10Hz is shown in Figure 5; peak-to-peak amplitude is less than 40nV.

Table 1: Measured Performance of the Op Amp

Input Noise Voltage Density at 1kHz		500pV/$\sqrt{Hz}$
Input Noise Voltage from 0.1Hz to 10Hz		40nV$_{p-p}$
Input Noise Current at 1kHz		1.5pA/$\sqrt{Hz}$
Gain-Bandwidth	G = 10	3MHz
	G = 100	600kHz
	G = 1000	150kHz
Slew Rate		2V/μs
Open-Loop Gain		3×10^7
Common-Mode Rejection		130dB
Input Bias Current		3μA
Supply Current		10mA
Nulled TCV$_{OS}$		0.1μV/°C Max
T.H.D. at 1kHz	G = 1000	0.002%

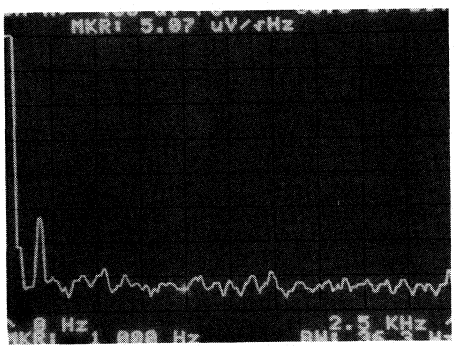

Spectrum analyzer display of broadband noise with a gain of 10,000.
Horizontal axis = 0 to 2.5kHz.
Normalized vertical axis = 830pV/$\sqrt{Hz}$ R.T.I.
e_n = 507pV/$\sqrt{Hz}$ at 1kHz.

Figure 3: Spectrum Analyzer Display — Broadband

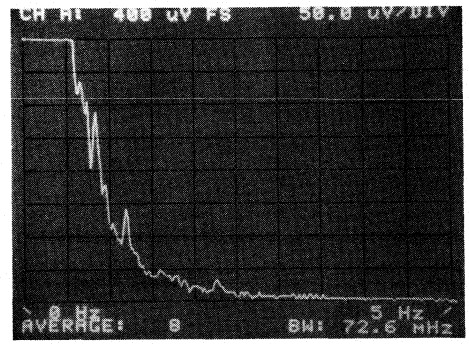

Low frequency noise spectrum at a gain of 10,000 showing a low 1.5Hz noise corner.

Figure 4: Spectrum Analyzer Display — Low Frequency

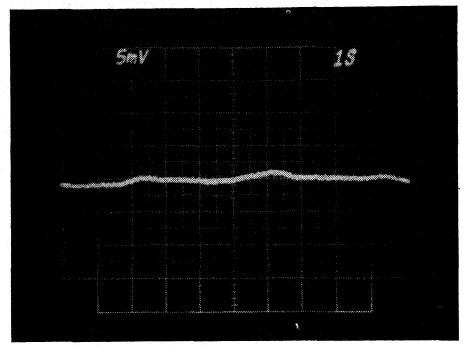

Peak-to-peak noise from 0.1 to 10Hz.
Overall gain is 100,000.

Figure 5: Oscilloscope Display

CONCLUSION

Using SSM-2210 matched transistor pairs operating at a high current level, it is possible to construct a high-performance, low-noise operational amplifier. The circuit uses a minimum of components and achieves performance levels impractical with monolithic amplifiers.

ANALOG DEVICES

AN-253
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Find Op Amp Noise with Spreadsheet

by Bob Clarke

By employing a spreadsheet's built-in graphics and programming capabilities, users can easily compare the noise performance of different op amps and plot their noise versus a variety of resistance and gain values. Using a noise model for the op amp *(Fig. 1)*, the expression for the effective integrated output noise (V_{on}) equals:

$$V_{on} = \{[I_{N-} R_{FB}]^2 + [I_{N+} R_P(1-G)]^2 + [V_N(1-G)]^2 + 4kT[R_{FB} + R_{FF}G^2 + R_P(1-G)^2]\}^{1/2} BW^{1/2}$$

and the expression for the effective integrated input noise (V_{in}) equals:

$$V_{in} = V_{on}/(1-G)$$

where

- V_{on} = the output noise voltage
- I_{N-} is the input noise current at the inverting input
- R_{FB} is the feedback resistance in ohms
- I_{N+} is the input noise current at the noninverting input
- R_P is the resistance at the noninverting input
- G is the circuit gain that equals $-R_{FB}/R_{FF}$
- V_N is the equivalent input noise voltage
- k is Boltzman's constant
- T is the absolute temperature in degrees Kelvin
- R_{FF} is the feedback resistance in ohms
- BW is the bandwidth in hertz.

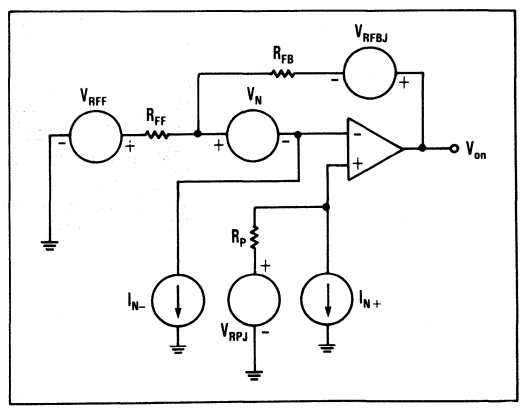

1. THIS OP AMP noise model accounts for noise current through the inverting (I_{N-}) and noninverting (I_{N+}) inputs and the input noise voltage (V_N). Each current induces a noise voltage in the resistors through which it flows.

Reprinted from Electronic Design – December 13, 1990

Programming these equations into a spreadsheet lets users compare different op amps as well as experiment with different component values in an interactive rather than a batch mode. This particular example was done using Microsoft's Excel spreadsheet program *(Fig. 2a)*. With component values entered in the cells across row 9, the formula for the effective integrated output noise in μV (entered in cell J9) is:

```
=((A9*E9)^2+(A9*F9*(1-I9))^2
+(C9*(1-I9))^2+H9*(E9+D9*I9^2
+F9*(1-I9)^2))^0.5*G9^0.5*1000000
```

This noise model for an op amp accounts for noise through the inverting and noninverting inputs as well as the input noise voltage. The noise versus circuit gain can be plotted by using the values obtained from the spreadsheet *(Fig. 2b)*.

OP AMP NOISE CALCULATIONS

Data Sheet Values			Circuit Values			Circuit Bandwidth	Thermal Contribution	Circuit Noise Gain	(μVrms)
I_{n+}	I_{n-}	V_n	R_{in}	R_f	R_p	(Hz)	4kT	$G=R_f/R_{in}$	
1E–11	1E–11	2E–09	1000	1000	0	10000000	1.645E–20	–1.00	38.59
1E–11	1E–11	2E–09	950	1000	0	10000000	1.645E–20	–1.05	38.81
1E–11	1E–11	2E–09	900	1000	0	10000000	1.645E–20	–1.11	39.06
1E–11	1E–11	2E–09	850	1000	0	10000000	1.645E–20	–1.18	39.34
1E–11	1E–11	2E–09	800	1000	0	10000000	1.645E–20	–1.25	39.66
1E–11	1E–11	2E–09	750	1000	0	10000000	1.645E–20	–1.33	40.02
1E–11	1E–11	2E–09	700	1000	0	10000000	1.645E–20	–1.43	40.44
1E–11	1E–11	2E–09	650	1000	0	10000000	1.645E–20	–1.54	40.93
1E–11	1E–11	2E–09	600	1000	0	10000000	1.645E–20	–1.67	41.51
1E–11	1E–11	2E–09	550	1000	0	10000000	1.645E–20	–1.82	42.21
1E–11	1E–11	2E–09	500	1000	0	10000000	1.645E–20	–2.00	43.05
1E–11	1E–11	2E–09	450	1000	0	10000000	1.645E–20	–2.22	44.11
1E–11	1E–11	2E–09	400	1000	0	10000000	1.645E–20	–2.50	45.45
1E–11	1E–11	2E–09	350	1000	0	10000000	1.645E–20	–2.86	47.22
1E–11	1E–11	2E–09	300	1000	0	10000000	1.645E–20	–3.33	49.64
1E–11	1E–11	2E–09	250	1000	0	10000000	1.645E–20	–4.00	53.13
1E–11	1E–11	2E–09	200	1000	0	10000000	1.645E–20	–5.00	58.54
1E–11	1E–11	2E–09	150	1000	0	10000000	1.645E–20	–6.67	67.91
1E–11	1E–11	2E–09	100	1000	0	10000000	1.645E–20	–10.0	87.46
1E–11	1E–11	2E–09	50	1000	0	10000000	1.645E–20	–20.0	148.64

(a)

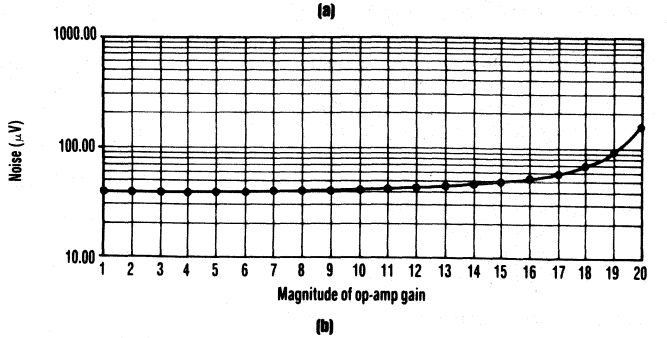

(b)

2. THE SPREADSHEET CALCULATIONS MAKE IT POSSIBLE for users to compare different op amps in the same circuit configuration or vary component values and look at the effects on noise (a). The spreadsheet's results can be plotted. Here, the noise is plotted versus circuit gain for an AD844 current-feedback op amp for a 10-MHz bandwidth (b).

ANALOG DEVICES

AN-107
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Active Feedback Improves Amplifier Phase Accuracy

by James Wong

Using matched op amps and active feedback, you can minimize the phase error and so extend the bandwidth of an amplifier by more than an order of magnitude. This technique is cheaper than using wideband amplifiers and less sensitive to the temperature-related drift specifications of passive components.

In applications such as sonar and image-processing systems, the phase relationship between two or more signals reveals essential information. These systems require accurate phase response in their amplifier circuitry to minimize measurement errors. In such cases, active feedback can often serve much better than other approaches. A typical op amp is insufficient in this situation because it introduces significant phase shift long before it reaches its −3dB frequency. The consequent phase error reduces the effective bandwidth of an op amp to something significantly less than the −3dB point.

You can use a wideband amplifier to overcome this phase-error problem. If the wideband amplifier operates with a −3dB bandwidth that is much higher than that of the signal that you intend to amplify, then the phase error at your signal's frequency decreases proportionately. The wideband amplifier's greater expense is the main drawback to this approach.

If you cascade two or three amplifiers, each of which has its gain reduced to share the overall gain of the composite amplifier, the gain reduction at each stage of the amplifier increases the −3dB bandwidth of each stage of the amplifier. Consequently, the overall bandwidth for a given phase accuracy is increased, but you pay for this improvement with increasingly higher costs and noise levels.

Another, less expensive way to solve the phase-error problem is to introduce extra circuitry in the amplifier's feedback loop, which provides frequency compensation. You can use an RC circuit to create a zero in the feedback loop that cancels the amplifier's pole. This cancellation improves the phase response markedly by lessening the amplifier's phase-response roll-off. The chief disadvantage of the RC technique is that it requires extensive tuning to match the zero with the pole. Furthermore, the different temperature coefficients of the RC components cause the zero to drift. And when the zero drifts, it no longer cancels the pole, and phase error becomes a problem once again.

Placing an op-amp circuit in the feedback stage of the amplifier creates the active feedback that can overcome the temperature drift of the RC networks. It is also a thriftier approach than using a wideband amplifier. You must make sure, however, that the op amps are very closely matched. Monolithically matched dual or quad op amps can provide the frequency-matching characteristics (to within 1 to 2%) necessary for the success of the active-feedback approach. This close matching is necessary across the full temperature range of your application. It's easier to achieve in an integrated dual or quad op amp than it is in discrete resistors and capacitors.

Figure 1 shows a basic active-feedback circuit. It requires an op amp and two external resistors to achieve phase-error cancellation. In the circuit, op amp A provides the forward gain of the composite amplifier. Resistors R_1 and R_1/K_1 determine the closed loop gain $A_V \gg 1 + K_1$. Amplifier A_2 provides active feedback to op amp A_1. The ratio of resistors R_2 and R_2/K_2 determines the amount of phase-error compensation and has no effect on the forward gain of the composite amplifier. You obtain optimum error cancellation when $K_1 = K_2$.

In terms of the complex frequency response, the error terms for the circuit are given by:

$$\text{magnitude error} \cong \left(\frac{\omega}{\beta\omega_T}\right)^2$$

$$\text{phase error} \cong -\left(\frac{\omega}{\beta\omega_T}\right)^3.$$

These equations let you compare the phase error associated with a single amplifier with that associated with the 2-op-amp, active-feedback approach. They arise from a complete analysis of the basic circuit and its derivatives (see "Analyzing compensation techniques").

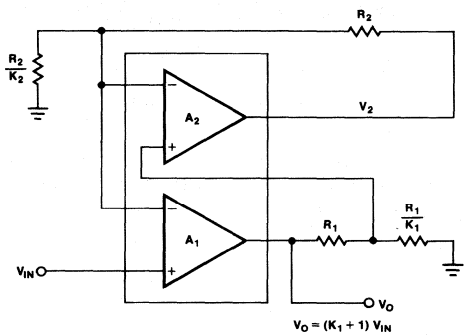

FIGURE 1: *Second-order compensation, as provided by the op amp in the feedback path of this circuit, extends the effective bandwidth of an amplifier substantially.*

Reprinted from EDN — September 17, 1987

TABLE 1: AC-Error Comparison

FREQUENCY	SINGLE-STAGE CONVENTIONAL PHASE (DEGREE)	CASCADED (TWO-STAGE) PHASE (DEGREE)	SECOND-ORDER COMPENSATION PHASE (DEGREE)
5kHz	–0.57	–0.36	0
10kHz	–1.15	–0.72	–0.0005
50kHz	–5.7	–3.62	–0.06
100kHz	–11.3	–7.21	–0.46
500kHz	–45.0	–45.0	–45.0

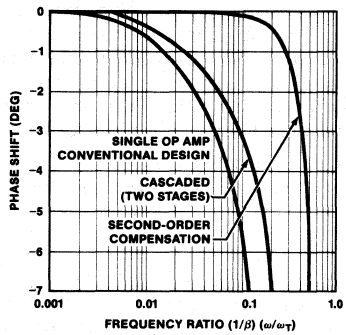

FIGURE 2: *Phase shift for single-op-amp, cascaded-op-amp, and second-order-compensation designs are charted here. Note the significant improvement brought about by second-order compensation.*

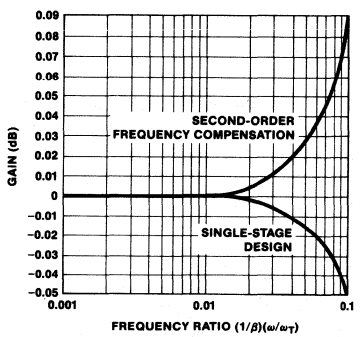

FIGURE 3: *Second-order compensation does carry the penalty of gain peaking, but this doesn't become significant until close to the corner frequency.*

DOMINANT POLES OCCUR AT 500kHz

Table 1 tabulates the phase error and magnitude error for an amplifier with a gain of 10. The comparison assumes that the op amp's unity-gain bandwidth is 5MHz. As you can see in the table, the dominant poles of all three amplifiers occur at 500kHz, where the phase shift for each technique is –45°. Clearly, the compensation techniques do not extend the bandwidth of the op amp itself. Rather, these second-order compensation techniques reduce the phase error at lower frequencies by adding an equal but opposite phase shift in the amplifier feedback loop.

If you must limit phase error to less than half a degree, the second-order compensation technique increases the effective bandwidth of your amplifier from about 5kHz to more than 100kHz. If you need to limit phase error entirely, the second-order compensation reduces phase error virtually to zero at frequencies to 50kHz; in contrast, a single-stage amplifier would be limited to well below 500Hz.

The phase and magnitude responses are plotted in Figures 2 and 3, respectively. The single op-amp response serves as a basis for comparison. The maximum scale of 1.0 represents the –45° phase-shift frequency. Figure 2 shows that the cascaded 2-stage amplifier offers only a slight improvement in bandwidth, and that the second-order compensation method offers a significant improvement in effective bandwidth.

THE TRADEOFFS ARE MINOR

The tradeoffs associated with the second-order feedback technique are minimal. Figure 3 illustrates the appreciable gain peaking incurred, but the circuit typically peaks about 3dB at the –45° phase frequency, well outside the useful range. Within the frequency range where phase shift is negligible, the gain error is also insignificant. In Figure 3, for example, at 1/10 of the corner frequency, gain error is only 0.1dB, about 1.2%.

In executing the second-order compensation design, it's extremely important to use op amps with frequency responses matched to within 1 to 2%. Op amps packaged separately can have mismatches as high as 10 to 20%, and high levels of mismatching cause either over- or undercompensation. Overcompensation creates excessive phase peaking, and undercompensation causes early phase roll-off.

Theoretically, second-order compensation works for any closed-loop gain. In practice though, at low gains, within the 1-to-5 range, the circuit may become unstable due to phase-margin degradation introduced by the active feedback. As a rule of thumb, you should work with a gain of 10 or greater. The match between theoretical and actual performance improves as the gain of the circuit is increased.

A circuit using the second-order compensation is illustrated in Figure 4. The op amp used in this example is an OP-470, which is unity-gain stable and provides a 6MHz unity-gain bandwidth. It's a quad op amp that offers a 1% match in ac characteristics between the four op amps on the chip. The circuit provides a gain of 10 for the amplifier.

The actual phase response of the circuit was measured using a network analyzer and is compared with a single-stage amplifier in Figure 5. The measurement confirms that the phase shifts of the second-order-compensated design and the single-stage design converge at −45° and −135°, respectively. The second-order response runs virtually flat with negligible phase shift to a much higher frequency before bandwidth limitation sets in. The roll-off is much steeper for the second-order system than for the single-stage system. Figure 6 illustrates more clearly that the second-order circuit's phase error remains nearly zero out to 100kHz, while the single-stage amplifier's low phase-shift bandwidth is limited to 2kHz.

Figure 6 also reveals a slight amount of phase peaking at about 100kHz, just before the response rolls off. The peaking stems from the op amps' second poles near the unity-gain frequency. With a closed-loop gain of 10, the dominant pole and the second pole are separated by only one decade in frequency. This proximity causes secondary effects on the phase response of the amplifier. However, if you increase the closed-loop gain of the circuit, you further separate the two poles, reducing the second-pole effect.

Fortunately, the second-order compensation's improvement in phase error does not exact much of a penalty in magnitude error. Figure 7 charts the magnitude response versus frequency. Some gain peaking is apparent, but as predicted, the peaking occurs well beyond the amplifier's useful bandwidth.

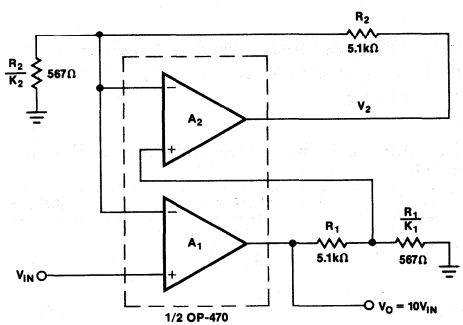

FIGURE 4: *This amplifier uses the second-order compensation technique with circuit values shown for a gain of 10.*

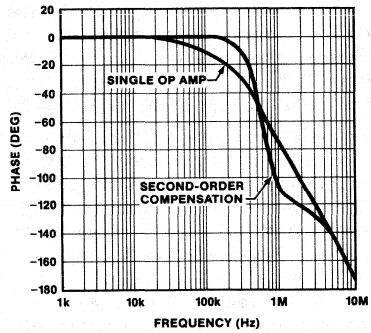

FIGURE 5: *The phase response of Figure 4's circuit illustrates the improved flatness of response that can be attributed to second-order compensation.*

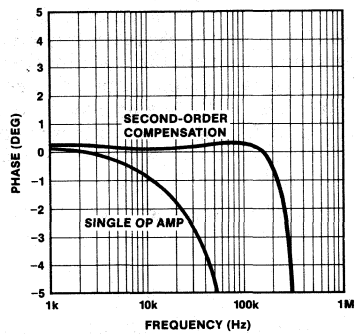

FIGURE 6: *Slight peaking in the phase response at 100kHz for the second-order-compensated circuit is illustrated in this close-up view of Figure 5.*

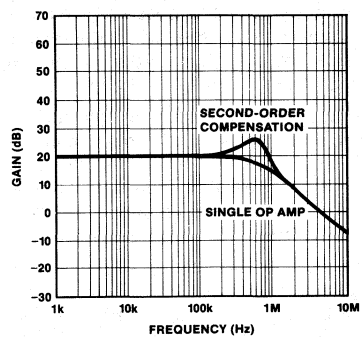

FIGURE 7: *The magnitude response of the second-order-compensated circuit exhibits peaking at a frequency that is well beyond the useful range of the circuit.*

(Figure 8 shows a close-up of the gain peaking for the second-order-compensated circuit.) In the frequency range below 100kHz, gain error is held to less than 0.25dB or 3%. In comparison to the single-stage approach, a second-order compensation offers both low phase error and low magnitude error over a significantly expanded bandwidth.

Either resistor tolerances or op-amp mismatches can degrade the compensation circuit's performance but to no great extent. What does make a big difference is variations in K_2 (Figure 9). If you assume that the op amps in Figure 1 match perfectly, the resistor ratio, K_2, is the only variable that affects the amount of phase compensation that the feedback circuit provides. If you decrease K_2, the gain of op amp A_2 is also decreased; in turn, the dominant pole of A_2 shifts to a higher frequency, thereby decreasing the compensation. In the extreme case, where op amp A_2's gain is reduced to unity, the circuit behaves as if it had no compensation at all. It then responds as a single-stage amplifier. On the other hand, if you increase K_2, the gain of A_2 is likewise increased, and A_2's bandwidth decreases; overcompensation results. The end effect is that phase and gain peak more.

Third-order compensation, which requires another op amp and two resistors in the feedback path, provides even greater phase compensation than afforded by second-order compensation. The schematic for the third-order compensation circuit is given in Figure 10. The additional compensation comes from op amp A_3 in the feedback path of op amp A_2. The basic circuit is the same as that of the second-order compensation circuit.

The sinusoidal transfer function when the phase difference is minimized may be written:

$$\frac{V_O}{V_{IN}} = (1 + K_1)\left[\frac{1 - j(\tau\omega)^5}{1 - 0.618034(\tau\omega)^2 - 0.618034(\tau\omega)^4 - (\tau\omega)^6}\right]$$

where: $\tau = \dfrac{1}{\beta\omega_T}$.

The numerator contains the phase information and the denominator contains the magnitude information. Table 2 contrasts the magnitude and phase errors for the second-order response with that of the third-order response. Magnitude error is 38% lower for third-order compensation. Phase error is reduced from a third-order to a fifth-order term.

The third-order frequency response is plotted in Figure 11 in comparison to that of the second-order compensation and of the single op-amp design. Given a prescribed error band — for example, one degree of phase error — third-order compensation doubles the effective bandwidth of an amplifier design over the second-order compensation technique.

As you might imagine, the extra op amp in the feedback loop causes the gain to peak more at the −45° corner frequency. However, in the frequency range where phase error is minimal, gain accuracy is also improved. At a frequency ratio below 0.1, magnitude errors do not exceed 0.1dB.

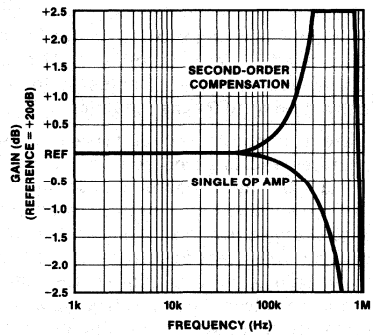

FIGURE 8: *Taking a closer look at the magnitude response, you see that below 100kHz, the frequency range where phase error is negligible, only a small degradation in gain accuracy occurs.*

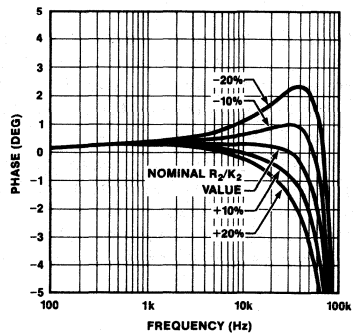

FIGURE 9: *Phase-response sensitivity to variations in the compensation resistor ratio results in slight over- and undercompensation.*

TABLE 2: Error Comparison of Second-Order vs Third-Order Compensation

	SECOND-ORDER	THIRD-ORDER
CLOSED-LOOP GAIN	$\frac{1}{\beta}$	$\frac{1}{\beta}$
MAGNITUDE ERROR	$\left(\frac{\omega}{\beta\omega_T}\right)^2$	$0.618\left(\frac{\omega}{\beta\omega_T}\right)^2$
PHASE ERROR	$-\left(\frac{\omega}{\beta\omega_T}\right)^3$	$-\left(\frac{\omega}{\beta\omega_T}\right)^5$

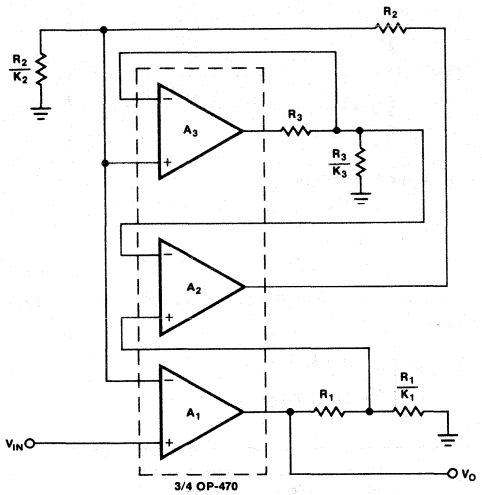

FIGURE 10: *Third-order compensation takes the concept of active feedback one step further, adding another level of compensation to the feedback path.*

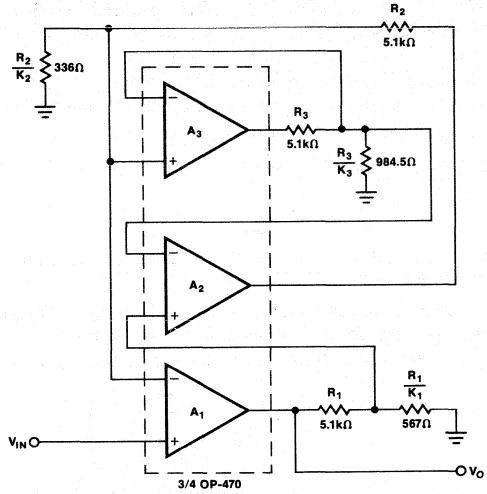

FIGURE 12: *A third-order compensation circuit with component values that set a gain of 10 for the amplifier is illustrated here.*

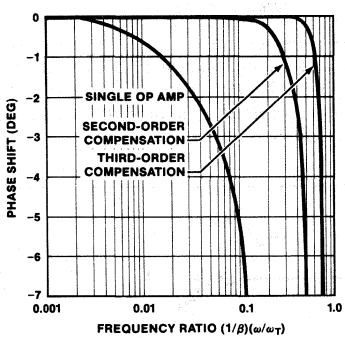

FIGURE 11: *For a given phase error, the third-order compensation technique offers twice the bandwidth that the second-order technique does.*

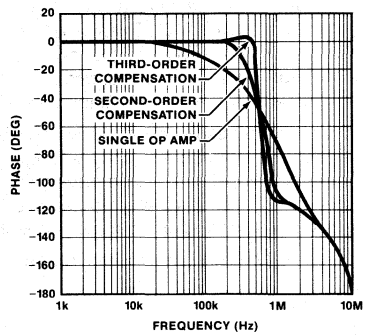

FIGURE 13: *The actual phase response of the third-order circuit remains flat to a higher frequency than the second-order circuit or the single-op-amp design.*

Figure 12 shows a test circuit that implements the third-order compensation. The phase response, measured by a network analyzer, is compared to that of the second-order compensation and of a single-op-amp design in Figure 13. The phase response of the third-order design remains flat beyond 400kHz, double the 200kHz provided by the second-order design. The magnitude response, on the other hand, has considerably higher peaking than the second-order design, but below 350kHz, the third-order design actually provides lower magnitude error than does the second-order design.

REFERENCES

1. Soliman, Ahmed M, "Design of High-Frequency Amplifiers," *IEEE Circuits and Systems*, June 1983.
2. Soliman, AM and Ismail, M, "Active Compensation of Op Amps," *IEEE Transactions on Circuits and Systems*, February 1979.

ACKNOWLEDGMENT

The author wishes to thank Tom Cate of Analog Devices, Inc., for his development of the mathematical model to minimize phase error.

ANALYZING COMPENSATION TECHNIQUES

A single-stage amplifier without any band-limiting circuit, as shown in Figure A, rolls off according to the closed-loop gain of the circuit and the amplifier's own natural frequency response. You can approximate the amplifier's open-loop frequency response in this way:

$$\text{open-loop gain } A(s) \cong \frac{A_O}{1+\frac{A_O}{\omega_T}s},$$

where A_O is the dc open-loop gain and, ω_T is the unity-gain radian frequency of the amplifier.

The fundamental assumption underlying this approximation is that the amplifier has a single-pole response. Thus, the closed-loop response as a function of frequency is given by:

$$\text{closed-loop gain } A_{CL}(s) = \frac{1}{\beta}\left(\frac{1}{1+\frac{1}{A_O\beta}+\frac{1}{\beta\omega_T}s}\right).$$

As long as the loop gain is much greater than unity – which is usually the case – you can rewrite the expression as:

$$\text{closed-loop gain } A_{CL}(s) \cong \frac{1}{\beta}\left(\frac{1}{1+\frac{1}{\beta\omega_T}s}\right).$$

Figure B illustrates the magnitude and phase response of the closed-loop amplifier. Notice that, at the -3dB corner frequency ω_c, the phase shift is -45°. Even at 1/10 of the corner frequency, the amplifier still has a 5.7° phase lag. The maximum frequency must be nearly two decades below the corner frequency to reduce phase error to less than one degree.

Figure 1 illustrates the second-order, active-feedback compensation circuit that you can model by writing the transfer equation for the two amplifier loops. In solving this equation, assume that each amplifier has a single-pole roll-off response and that the two amplifier responses match perfectly.

At the junction of R_1 and R_1/K_1, the resistor divider expression becomes:

$$V_O\left(\frac{1}{1+K_1}\right).$$

Similarly, the resistor divider expression for R_2 and R_2/K_2 is:

$$V_2\left(\frac{1}{1+K_2}\right).$$

For ac response, the loop equations can now be written:

$$A_1\text{loop: } V_O = \left[V_{IN} - \left(\frac{1}{1+K_2}\right)V_2\right]\frac{\omega_T}{s}$$

$$A_2\text{loop: } V_2 = \left[\left(\frac{1}{1+K_1}\right)V_O - \left(\frac{1}{1+K_2}\right)V_2\right]\frac{\omega_T}{s}.$$

Solving the simultaneous equations in terms of V_O and V_{IN}, the transfer function is:

$$\frac{V_O}{V_{IN}} = (1+K_1)\left[\frac{1+\frac{(1+K_2)}{\omega_T}s}{1+\frac{(1+K_1)}{\omega_T}s+\frac{(1+K_1)(1+K_2)}{\omega_T^2}s^2}\right].$$

Defining the time constants:

$$\tau_1 = \frac{1+K_1}{\omega_T} \text{ and } \tau_2 = \frac{1+K_2}{\omega_T},$$

and substituting those time constants into the previous equations yields:

$$\frac{V_O}{V_{IN}} = (1+K_1)\left[\frac{1+\tau_2 s}{1+\tau_1 s+\tau_1\tau_2 s^2}\right].$$

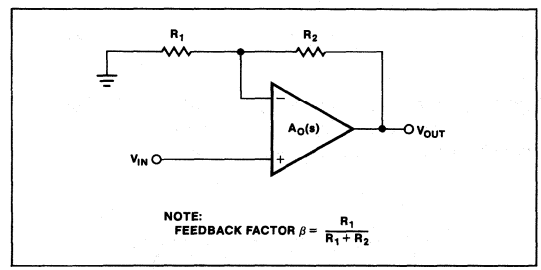

FIGURE A: *A basic single-stage amplifier without any band-limiting circuit rolls off according to the closed-loop gain of the circuit and the amplifier's own natural frequency response.*

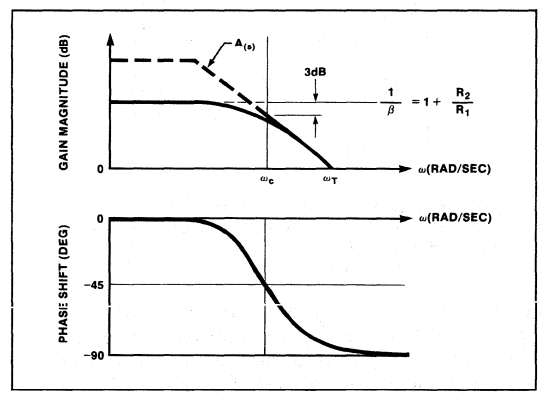

FIGURE B: *Magnitude and phase response are shown for the basic single-stage amplifier. Notice that, at the −3dB corner frequency ω_c the phase shift is −45°.*

Continued

ANALYZING COMPENSATION TECHNIQUES *Continued*

For optimum compensation, the time constants τ_1 and τ_2 are made equal:

$$\frac{V_O}{V_{IN}} = (1+K_1)\left[\frac{1+\tau s}{1+\tau s+\tau^2 s^2}\right].$$

For sinusoidal input, where phase information is important, the complex frequency domain is used:

$$\frac{V_O}{V_{IN}}(j\omega) = (1+K_1)\left[\frac{1+j\tau\omega}{1+j\tau\omega-\tau^2\omega^2}\right].$$

Now we calculate the complex conjugates to simplify:

$$\frac{V_O}{V_{IN}} = (1+K_1)\left[\frac{(1+j\tau\omega)(1+\tau^2\omega^2-j\tau\omega)}{(1-\tau^2\omega^2)^2+\tau^2\omega^2}\right]$$

$$= (1+K_1)\left[\frac{1-j\tau^3\omega^3}{1-\tau^2\omega^2+\tau^4\omega^4}\right].$$

From this equation, the magnitude error expression can be derived.

$$\text{Magnitude error} = \left|\frac{1-j\tau^3\omega^3}{1-\tau^2\omega^2+\tau^4\omega^4}\right| - 1$$

$$\text{for } \tau\omega \ll 1, \quad \cong \left|\frac{1}{1-\tau^2\omega^2}\right| - 1$$

$$\cong 1+\tau^2\omega^2 - 1$$

$$\cong \tau^2\omega^2.$$

Similarly, the phase error expression can be derived.

$$\text{Phase error} = \text{ARC}(1-j\tau^3\omega^3) \cong -\tau^3\omega^3.$$

$\tau\omega$ relates to closed-loop gain as $\frac{1}{\beta} = 1+K_1 = 1+K_2$,

therefore, $\tau\omega = \frac{\omega}{\beta\omega_T}$.

The transfer function can then be rewritten as:

$$\frac{V_O}{V_{IN}} = \frac{1}{\beta}\left[\frac{1-j\left(\frac{\omega}{\beta\omega_T}\right)^3}{1-\left(\frac{\omega}{\beta\omega_T}\right)^2+\left(\frac{\omega}{\beta\omega_T}\right)^4}\right]$$

Using this equation, the phase and magnitude behaviors as functions of frequency are tabulated in Table A. Note that, at 1/10 of the corner frequency, the compensated circuit produces a -0.057° phase shift, far superior to the -5.7° produced by the uncompensated circuit.

TABLE A: Second-Order Compensation Phase and Magnitude Behavior

$\frac{\omega}{\beta\omega_T}$	PHASE (DEGREE)	MAGNITUDE	20 LOG (MAG) IN dB
0.01	-5.7×10^{-5}	1.0000	0.0000
0.02	–0.0005	1.0004	0.0035
0.04	–0.00367	1.0016	0.0139
0.05	–0.00716	1.0025	0.0217
0.06	–0.01238	1.0036	0.0312
0.07	–0.01965	1.0049	0.0425
0.08	–0.0293	1.0064	0.0554
0.1	–0.0573	1.01	0.0864
0.2	–0.458	1.04	0.3404
0.3	–1.547	1.0896	0.745
0.4	–3.662	1.1576	1.271
0.5	–7.162	1.25	1.938
0.7	–14.428	1.3766	2.776
0.8	–27.1125	1.4598	3.2858
1.0	–45.0	1.4125	3.0

You can model the third-order compensation circuit of Figure 10 in a similar way. Assume that all three op amps in the circuit have a single-pole roll-off in the frequency response represented by:

$$A_O(s) \cong \left(\frac{\omega_T}{s}\right),$$

where ω_T is the amplifier unity-gain crossover frequency.

For ac sinusoidal response, the loop equations can be written:

$$A_1\text{loop: } V_O = \frac{\omega_T}{s}\left[V_{IN} - \frac{1}{(1+K_2)}V_2\right]$$

$$A_2\text{loop: } V_2 = \frac{\omega_T}{s}\left[\frac{1}{(1+K_1)}V_O - \frac{1}{(1+K_3)}V_3\right]$$

$$A_3\text{loop: } V_3 = \frac{\omega_T}{s}\left[\frac{1}{(1+K_2)}V_2 - \frac{1}{(1+K_3)}V_3\right].$$

Solving the three simultaneous equations in terms of V_O and V_{IN}, the transfer function is:

$$\frac{V_O}{V_{IN}} = \frac{ax^2+bx+1}{ax^3+bx^2+\left(1+\frac{1+K_3}{1+K_1}\right)x+\left(\frac{1}{1+K_1}\right)},$$

Continued

ANALYZING COMPENSATION TECHNIQUES *Continued*

where $a = (1 + K_2)(1 + K_3)$,

$b = 1 + K_2$

$x = \dfrac{s}{\omega_T}$.

Define the time constants $\tau_1 = \dfrac{1 + K_1}{\omega_T}$,

$\tau_2 = \dfrac{1 + K_2}{\omega_T}$, and $\tau_3 = \dfrac{1 + K_3}{\omega_T}$, then substitute,

$$\frac{V_O}{V_{IN}} = (1 + K_1)\left[\frac{1 + \tau_2 s + \tau_2 \tau_3 s^2}{1 + (\tau_1 + \tau_3)s + \tau_1 \tau_2 s^2 + \tau_1 \tau_2 \tau_3 s^3}\right].$$

The dc gain of the amplifier is $(1 + K_1)$. And the ac response has the general form:

$$\frac{V_O}{V_{IN}} = (1 + K_1)\varepsilon_p(s),$$

where $\varepsilon_p(s) = \dfrac{1 + \tau_2 s + \tau_2 \tau_3 s^2}{1 + (\tau_1 + \tau_3)s + \tau_1 \tau_2 s^2 + \tau_1 \tau_2 \tau_3 s^3}$.

Solve for ac response in terms of τ_1, τ_2, and τ_3:

$$\varepsilon_p(j\omega) = \frac{1 - \tau_2 \tau_3 \omega^2 + j\tau_2 \omega}{1 - \tau_1 \tau_2 \omega^2 + j\left[(\tau_1 + \tau_3)\omega - \tau_1 \tau_2 \tau_3 \omega^3\right]}.$$

Defining:

$a = 1 - \tau_2 \tau_3 \omega^2$, $b = \tau_2 \omega$,

$c = 1 - \tau_1 \tau_2 \omega^2$, $d = (\tau_1 + \tau_3)\omega - \tau_1 \tau_2 \tau_3 \omega^3$

$$\varepsilon_p(j\omega) = \frac{a + jb}{c + jd} = \frac{ac + bd + j(bc - ad)}{c^2 + d^2}.$$

The numerator determines the phase shift of the amplifier. Solve the numerator of $\varepsilon_p(j\omega)$:

$$\varepsilon_p(j\omega) = 1 + j\left[(\tau_2 - \tau_1 - \tau_3)\omega - (\tau_1 \tau_2^2 - \tau_2 \tau_3^2 - 2\tau_1 \tau_2 \tau_3)\omega^3 - \tau_1 \tau_2^2 \tau_3^2 \omega^5\right].$$

To minimize phase shift, make $\tau_2 - \tau_1 - \tau_3 = 0$, and

numerator $\varepsilon_p(j\omega) =$

$$1 + j\left[(-\tau_2)(\tau_1 \tau_2 - 2\tau_1 \tau_3 - \tau_3^2)\omega^3 - \tau_1 \tau_2^2 \tau_3^2 \omega^5\right].$$

The objective is to eliminate the ω^3 term:

$$-\tau_2\left[\tau_1(\tau_2 - 2\tau_3) - \tau_3^2\right] \to 0.$$

First substitute $\tau_1 + \tau_3$ for τ_2, then let $\tau_3 = \alpha\tau_1$, set the equation equal to zero and solve for α:

$$-(\tau_1 + \alpha\tau_1)\left[\tau_1(\tau_1 - \alpha\tau_1) - (\alpha\tau_1)^2\right] = 0$$

$$\alpha^2 + \alpha - 1 = 0.$$

Solving α for the two roots, you obtain:

$\alpha = 0.618034$

$\alpha = -1.618034$

For minimum phase shift, the relationships $\tau_1 + \tau_3 = \tau_2$ and $\tau_3 = 0.618\tau_1$, are used.

$\tau_1 = \tau$

$\tau_2 = 1.618034\tau$

$\tau_3 = 0.618034\tau$,

recalling that $\tau_1 = \dfrac{1 + K_1}{\omega_T}$, $\tau_2 = \dfrac{1 + K_2}{\omega_T}$, and $\tau_3 = \dfrac{1 + K_3}{\omega_T}$.

Since K_1 relates to the dc gain $1/\beta$ of the amplifier as $(1 + K_1) = 1/\beta$, the design equations are:

$$K_1 = \frac{1}{\beta} - 1.$$

Substituting and solving for K_2 and K_3 in terms of K_1,

$K_2 = 1.618034\,K_1 + 0.618034$

$K_3 = 0.618034\,K_1 - 0.382$.

These three last design equations produce optimum phase cancellation for an amplifier with any gain. For example, for an amplifier gain of 10, $K_1 = 9$, $K_2 = 15.18$, and $K_3 = 5.18$.

In order to derive the complete transfer function, the denominator is similarly solved. Therefore, with:

denominator $\varepsilon_p(j\omega) = 1 - 0.618\tau^2\omega^2 - 0.618\tau^4\omega^4 + \tau^6\omega^6$,

the complete ac response is:

$$\varepsilon_p(j\omega) = \frac{1 - j(\tau\omega)^5}{1 - 0.618034(\tau\omega)^2 - 0.618034(\tau\omega)^4 - (\tau\omega)^6}.$$

Finally, the complete amplifier transfer function is:

$$\frac{V_O}{V_{IN}} = (1 + K_1)\left[\frac{1 - j(\tau\omega)^5}{1 - 0.618034(\tau\omega)^2 - 0.618034(\tau\omega)^4 - (\tau\omega)^6}\right].$$

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

A Collection of Amp Applications

by James Wong

This note examines some of the numerous and widely-used applications of the operational amplifier. While not attempting to list every possible application, it presents several basic circuit configurations that can be modified to suit applications other than those listed. In each case, significant op amp or circuit characteristics are discussed to aid the user in adapting the circuit to a particular need. Table 1 provides an index to the applications contained in this note.

TABLE 1: Op Amp Circuit Applications

Figure	Application
1	±200mA Servo Motor Amplifier
2	Precision Programmable Gain Amplifier
3	Bilateral Current Source
4	Precision Current Pump
5	High-Sensitivity Voltage Comparator
6	Micropower 5V Regulator
7	Micropower 1.23V Bandgap Reference
8	Versatile Triangular Wave Generator
9	Wide Range, Low Current Ammeter
10	Precision Threshold Detector/Amplifier
11	Wide-Dynamic-Range Light Detector
12	Isolation Amplifier
13	Dual Programmable Window Comparator
14	±36V Low Noise Operational Amplifier
15	High Q Notch Filter
16	Piezoelectric Transducer Amplifier
17	High Stability Voltage Reference
18	Precision Dual Tracking Voltage Reference
19	RIAA Phono Pre-Amplifier
20	Headphone Amplifier
21	NAB Tape Head Pre-Amplifier
22	Microphone Pre-Amplifier
23	Micropower Wien Bridge Oscillator
24	Micropower Instrumentation Amplifier
25	Piecewise-Linear Amplifier (Decreasing Gain)
26	Piecewise-Linear Amplifier (Increasing Gain)
27	Current Monitor Circuit
28	Free-Running Square-Wave Oscillator
29	Precision Analog Multiplier/Divider
30	Precision Absolute Value Circuit
31	Thermocouple Amplifier with Cold-Junction Compensation
32	Instrumentation Amp (2 Op Amp Design)
33	±200V Low Offset Operational Amplifier
34	Impedance Transforming Amplifier
35	Precision Current Sinks
36	Low Noise AGC Amplifier
37	Amplifier With Active Output Clipping
38	Low Power Amplifier With Squelch

FIGURE 1: ±200mA Servo Motor Amplifier

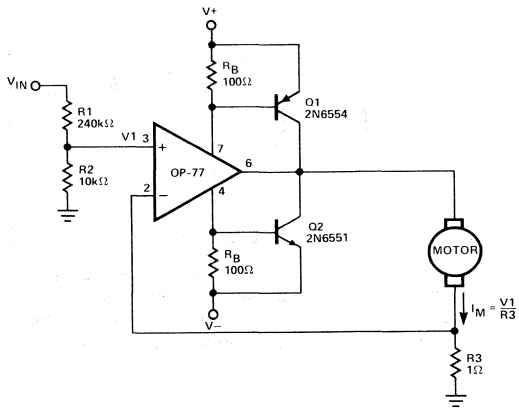

R_B sets the bias point for transistors Q1 and Q2. Because $V_{BE(ON)}$ varies greatly with temperature, a guardband is required to prevent Q1 and Q2 from conducting simultaneously. R_B should be selected such that the transistors do not conduct until I_M equals the op amp quiescent supply current, I_{SY}. The transistors will begin to conduct at about $V_{BE\,(ON)} = 0.5V$.

In this design,

$$R_B = \frac{V_{BE(ON)}}{I_{SY} + I_M} = \frac{0.5}{0.0025 + 0.0025} = 100\Omega$$

To maximize voltage swing across the motor, V1 must be minimized. If at full load V1 = 0.2V with V+ = 15V and V_{BE1} = 0.8V, the voltage across the motor will be:

$$V_M = (V^+ - 2) - V_{BE1} - V1 = (15 - 2) - 0.8 - 0.2$$
$$= 12.0V$$

V_{IN} may be scaled with a resistive divider as:

$$\frac{V_{IN}}{V1} = \frac{R1 + R2}{R2}$$

With R1 = 240kΩ and R2 = 10kΩ, V_{IN} = 5V will produce I_M = 200mA.

FIGURE 2: Precision Programmable Gain Amplifier

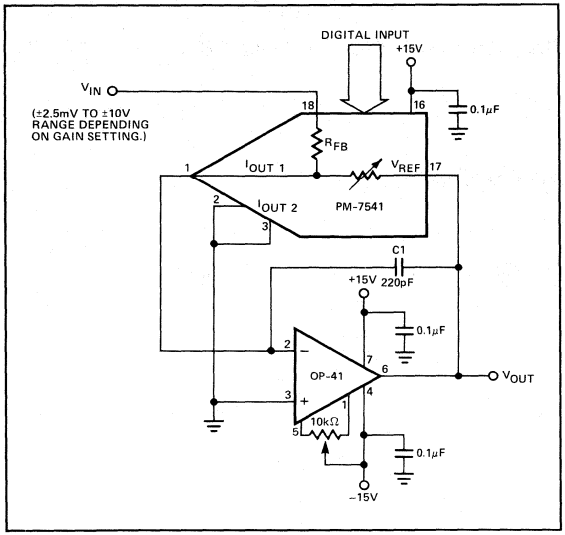

The digitally programmable gain has 12-bit accuracy over the range of –1 to –1024 and 10-bit accuracy to –4096. The low bias current of the OP-41's JFET input maintains this accuracy, while C1 limits the noise voltage bandwidth allowing accurate measurement down to microvolt levels.

DIGITAL IN	GAIN (A_V)
4095	–1.00024
2048	–2
1024	–4
512	–8
256	–16
128	–32
64	–64
32	–128
16	–256
8	–512
4	–1024
2	–2048
1	–4096
0	OPEN LOOP

FIGURE 3: Bilateral Current Source

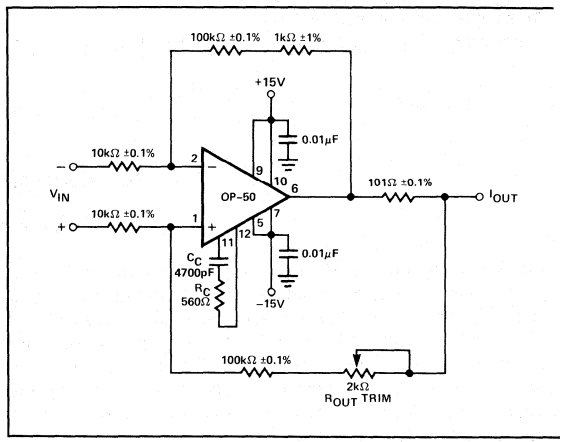

Compliance is better than ±11V at an output current of 20mA, and the trimmed output resistance is typically 2MΩ with $R_L \leq 500\Omega$. For the resistor values shown, the maximum V_{IN} is 200mV.

FIGURE 4: Precision Current Pump

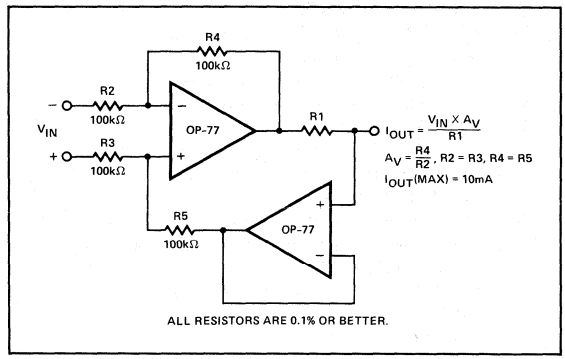

Accuracy of I_{OUT} is improved by using a noninverting voltage-follower in the feedback loop. To maximize voltage compliance of I_{OUT}, R1 should be minimized.

FIGURE 5: High-Sensitivity Voltage Comparator

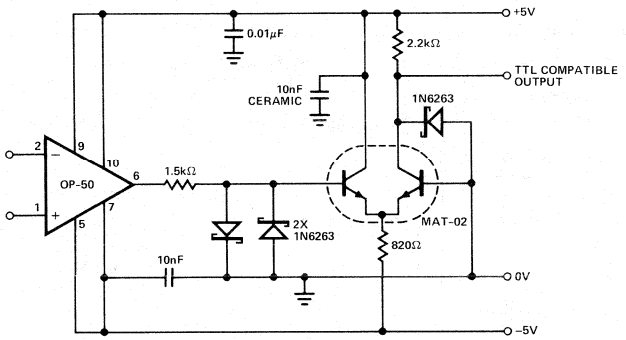

This comparator circuit is capable of resolving a submicrovolt difference signal. The OP-50, operating without feedback, drives a second gain stage which generates a TTL-compatible output signal. Schottky-clamp diodes prevent overdriving of the long-tailed transistor pair and stop saturation of the output transistor. Power supply voltage is set to ±5V to lower the quiescent power dissipation and minimize thermal feedback due to output stage dissipation. Operating from ±5V supplies also reduces the OP-50 rise and fall times as the output slews over a reduced voltage range. This, in turn, reduces the output response time.

It is common practice with voltage comparators to ground one input terminal and to use a single-ended input. The historic reason is poor common-mode rejection on the input stage. In contrast, the OP-50 has very high common-mode rejection and is capable of detecting microvolt level differences in the presence of large common-mode signals.

The comparator is not fast, but it is very sensitive and can detect signal differences as low as 0.3μV. With large input overdrives, the circuit responds in approximately 3μs. If sharp transitions are needed, the use of a TTL Schmitt-trigger input is recommended. A table of Response Time vs. Input Overdrive is shown below.

INPUT OVERDRIVE	100mV	10mV	1mV	100μV	10μV
Positive Output Delay	3.2μs	5μs	40μs	340μs	2.4ms
Negative Output Delay	1.8μs	5μs	50μs	380μs	4.5ms

13

FIGURE 6: Micropower 5V Regulator

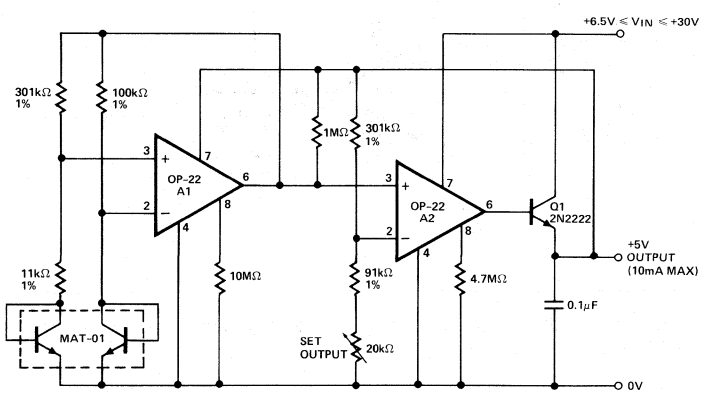

This 5V regulator is ideal for instrumentation requiring good power efficiency. Low-power 3-terminal IC regulators typically draw 2mA to 5mA quiescent current compared to only 50μA with this discrete implementation. Maximum load current is 10mA as shown, and can be increased by changing Q1 to a power transistor and proportionately increasing the set current of A2.

FIGURE 7: Micropower 1.23V Bandgap Reference

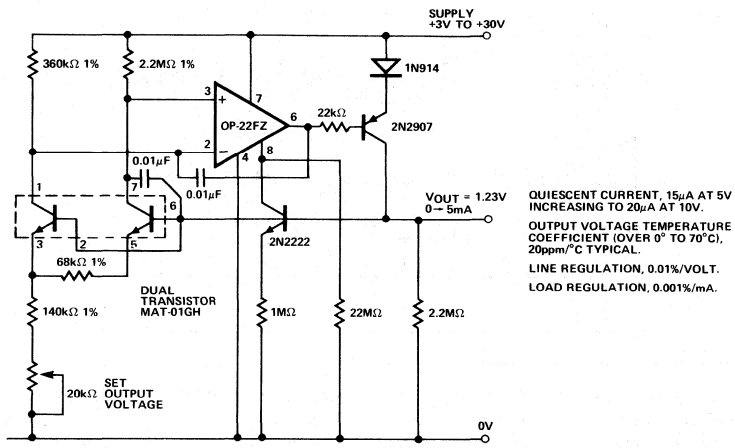

A micropower bandgap voltage reference operating at a quiescent current of 15μA may be constructed using an OP-22 and a MAT-01 dual transistor. The circuit provides a 1.23V reference with better performance than micropower IC shunt regulators and has the advantages of being a series regulator.

FIGURE 8: Versatile Triangular Wave Generator

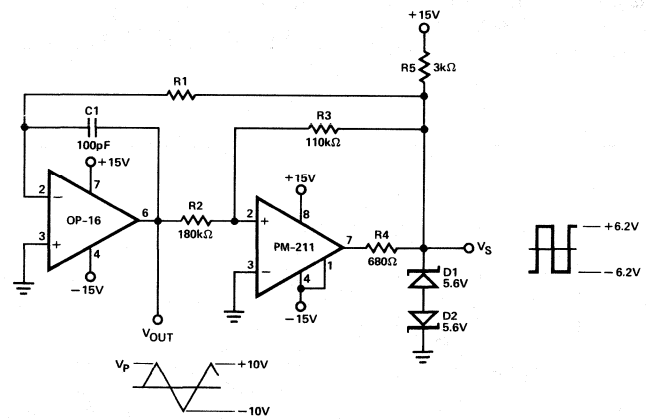

Triangular waveshapes of ±10V from 100Hz to 500kHz are obtained with values of R1 from 15MΩ to 3kΩ as given by:

$$R1 = \frac{V_S}{4V_P \cdot f_O \cdot C1}, \; V_S = 6.2V, \; V_P = 10V$$

The amplitude of the triangle wave may be adjusted via R2 as:

$$R2 = \frac{V_P \cdot 110k\Omega}{6.2V}$$

FIGURE 9: Wide Range, Low Current Ammeter

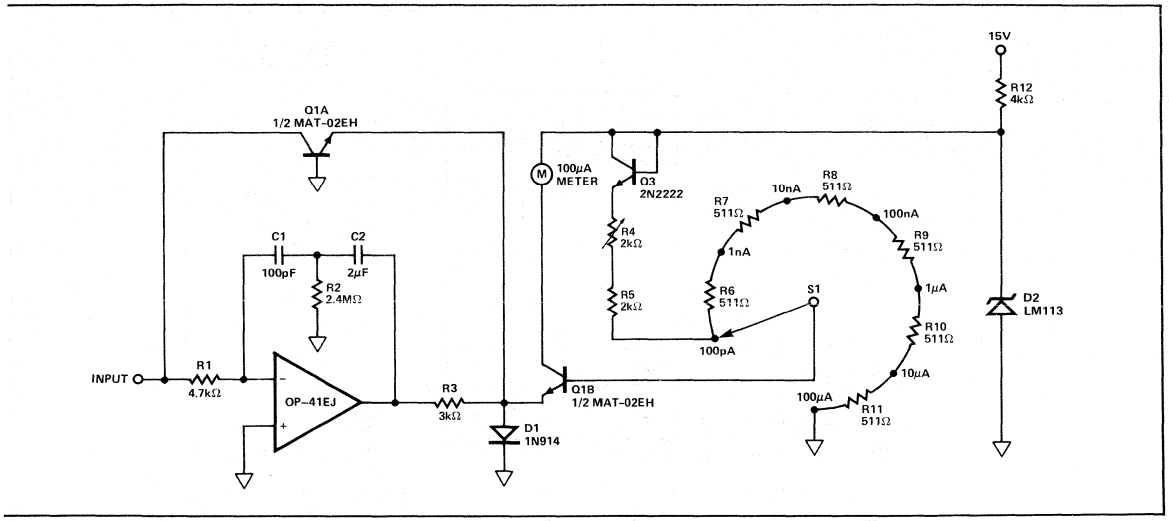

This ammeter can measure currents from 100pA to 100μA without the use of high value resistors. Accuracy is better than 1% over most of the range, depending upon the accuracy of the divider resistor and the input bias current of the op amp. Using the OP-41 as the input amplifier allows low end measurement down to a few pA due to the 5pA input bias current. Because the voltage across the inputs of the inverting amplifier is forced to virtually zero, the current meter's effective series voltage drop is less than 500μV at any current level.

Calibration is simple, requiring only one adjustment. R4 is used to adjust full scale deflection with a 1μA input current. This will give maximum accuracy over the operating range of currents.

The low V_{OS} and exceptionally good log conformance of the MAT-02 assure high accuracy over the full 6 decade operating range.

13

FIGURE 10: Precision Threshold Detector/Amplifier

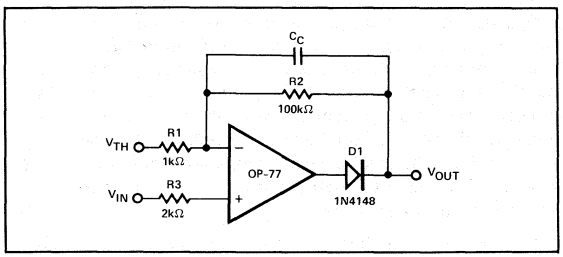

When $V_{IN} < V_{TH}$, the amplifier output swings negative, reverse biasing diode D1. Therefore $V_{OUT} = V_{TH}$.

When $V_{IN} \geq V_{TH}$, the loop closes, and

$$V_{OUT} = V_{TH} + (V_{IN} - V_{TH})\left(1 + \frac{R2}{R1}\right)$$

C_C is selected to smooth the loop response.

FIGURE 11: Wide-Dynamic-Range Light Detector

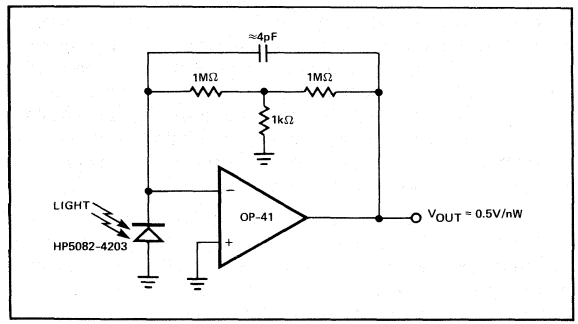

This circuit produces an output voltage proportional to light input over a 60dB range. The 5pA input bias current of the OP-41 assures a low output voltage offset.

FIGURE 12: Isolation Amplifier

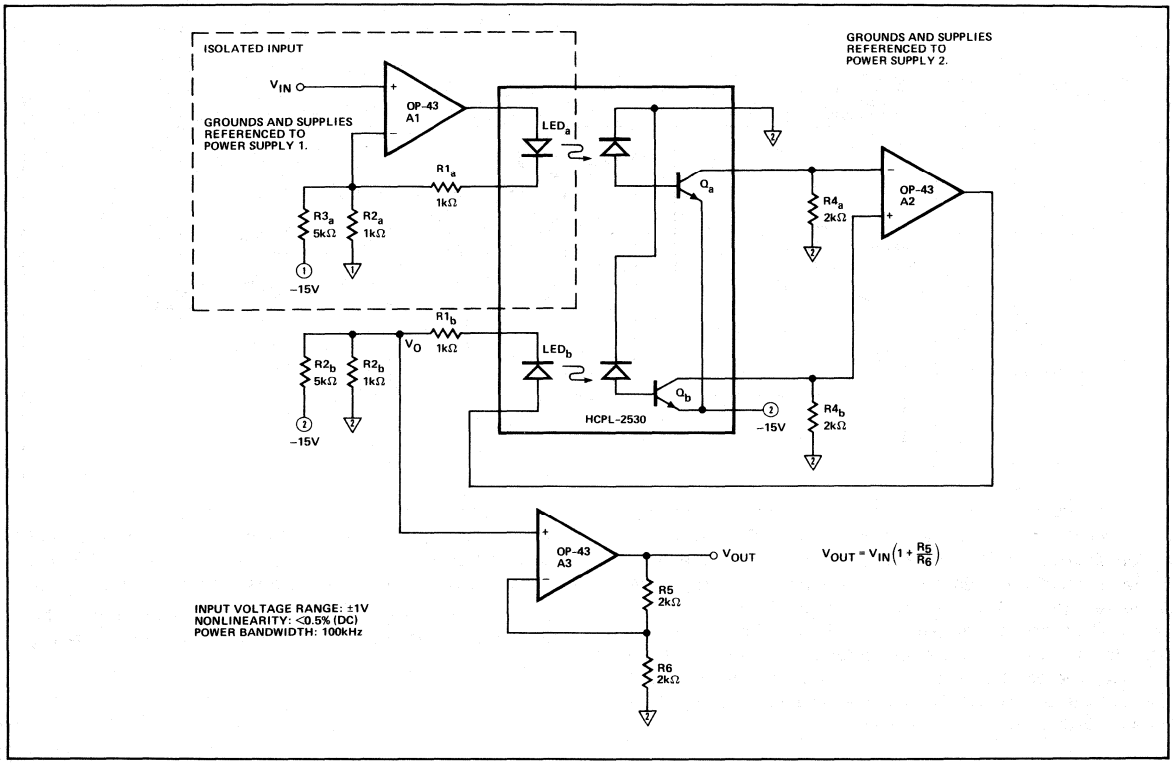

In conjunction with two optocouplers, or a dual optocoupler such as the Hewlett-Packard HCPL-2530, three OP-43's can be combined to create an isolation amplifier. In this sort of amplifier, the input and the output operate on separate power supplies, allowing extremely high common-mode voltages to be dealt with. In industrial applications, an isolation amplifier protects instrumentation from high voltages at the sensing site. When interfacing with a computing system, an isolation amplifier will protect the rest of the system from a sensor which accidentally becomes shorted to a high voltage.

The isolation amplifier operates on the principle that the nonlinearities of one optocoupler will be tracked by the nonlinearities of another, if they are well matched. By using an optocoupler in the feedback loop of A2, nonlinearities of the isolating optocoupler will be cancelled. V_O will equal V_{IN} plus an offset created by imperfect matching between a and b side resistors and optocouplers.

A3 is an output buffer for the isolation amplifier. The low bias current ensures that it does not affect the voltage it is amplifying. Gain is realized in this stage, and any offsets induced in the previous stages may be corrected by offsetting this op amp. Although shown in the circuit as a simple gain stage, this output amplifier may take any form desired. It may be configured as a filter or other waveshaping circuit as needed. The only requirement is that the buffer not disturb the currents in the optocoupler feedback circuit, thus noninverting amplifier configurations are preferred. For highest linearity, the currents in the two LEDs should track as closely as possible.

With stable supplies, the circuit has excellent reponse and displays less than 0.5% DC nonlinearity with a $2V_{p-p}$ signal. The high speed of the OP-43 gives the circuit a power bandwidth of 100kHz, while the majority of the power budget is consumed in biasing the LEDs. The dual optocoupler provides isolation against 600VDC common-mode voltages. Higher isolations may be achieved using two separate optocouplers, such as HP's 6N136.

FIGURE 13: Dual Programmable Window Comparator

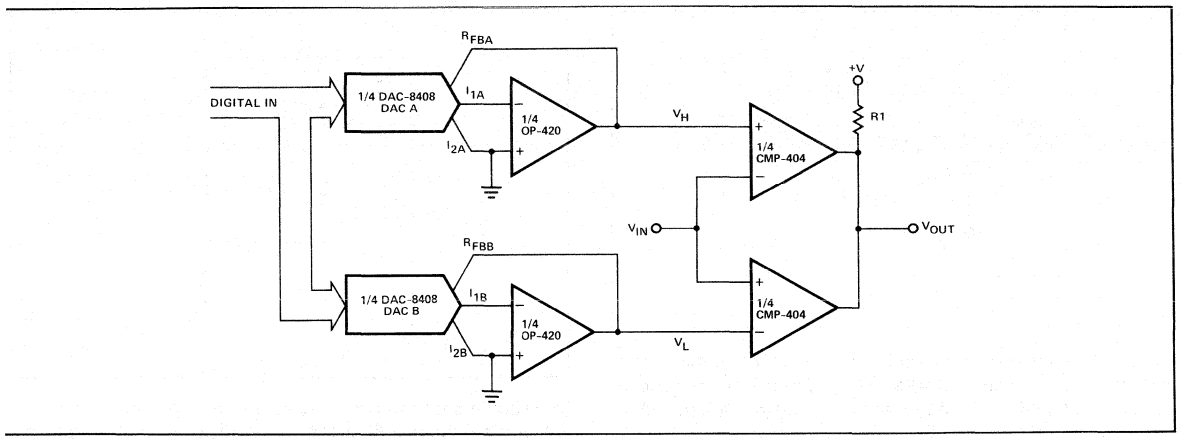

Only three ICs are required to fully implement two independent programmable window comparators. The quad, latched, 8-bit CMOS DAC-8408 together with the quad micropower OP-420 provide digitally-programmable HIGH and LOW thresholds to the CMP-404, quad low-power comparator. The outputs of the threshold comparators are wire-ORed with a common pull-up resistor producing $V_{OUT} = +V$ only when $V_L < V_{IN} < V_H$. Total supply current for the full circuit is less than 2mA.

FIGURE 14: ±36V Low Noise Operational Amplifier

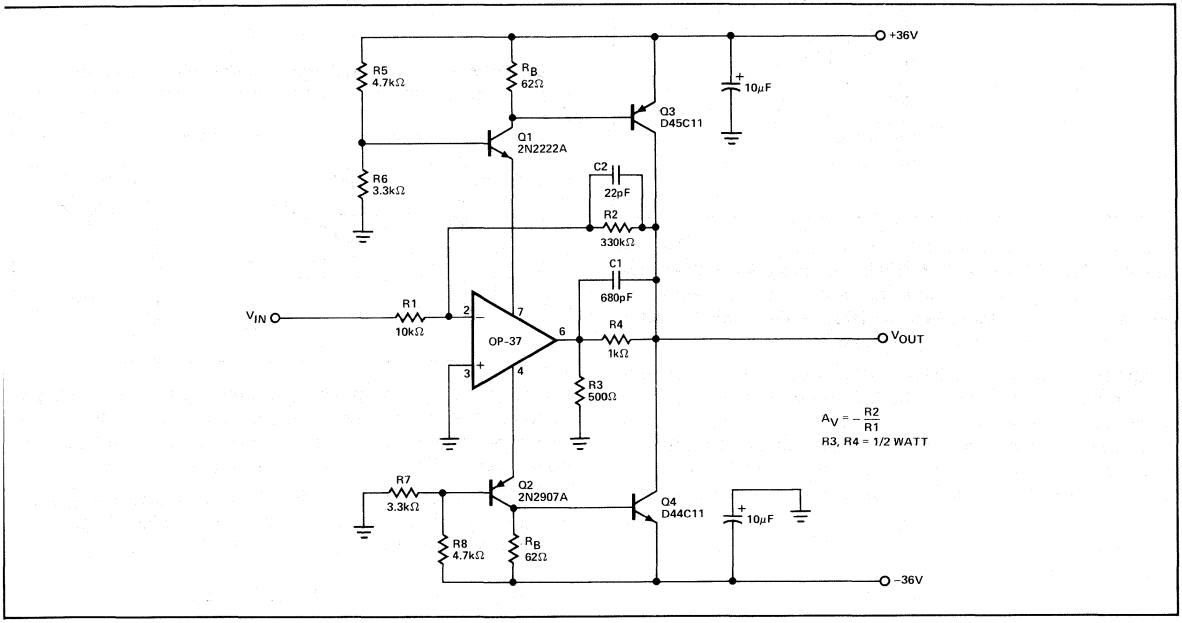

An OP-37 provides a low-noise front end for this amplifier which is capable of delivering over ±200mA to a load with a 70V peak-to-peak output swing. Transistors Q1 and Q2 are series regulators stepping down the supply voltage for the OP-37 to ±15V, while transistors Q3 and Q4 provide the high current output drive. R3 and R4 form an output voltage gain stage whose gain, $A_V = 3$, is reduced to unity at high frequencies by C1 to maintain stability.

FIGURE 15: High Q Notch Filter

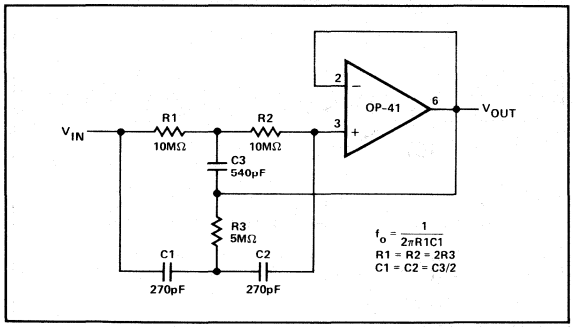

The low bias current and high input impedance of the OP-41 enable small value capacitors and large resistors to be used in this 60Hz notch filter. The 5pA bias current only develops 100μV across R1 and R2.

FIGURE 16: Piezoelectric Transducer Amplifier

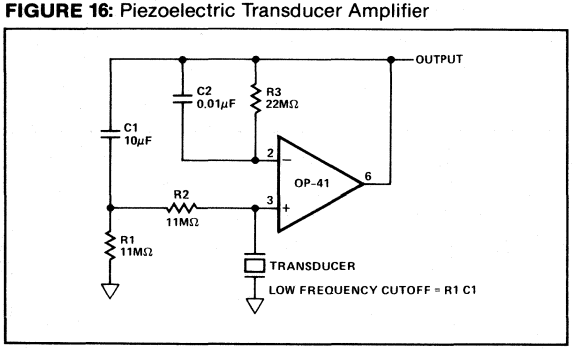

Piezoelectric transducers often require a high-input-resistance amplifier. The OP-41 can provide input resistance in the range of $10^{12}\Omega$, however, a DC return for bias current is needed. To maintain a high R_{IN}, large value resistors above 22MΩ are often required. These may not be practicable.

Using this circuit, input resistances that are orders of magnitude greater than the values of the DC return resistors can be obtained. This is accomplished by bootstrapping the resistors to the output. With this arrangement, the lower cutoff frequency is determined more by the RC product of R1 and C1 that it is by resistor values and the equivalent capacitance of the transducer.

FIGURE 17: High Stability Voltage Reference

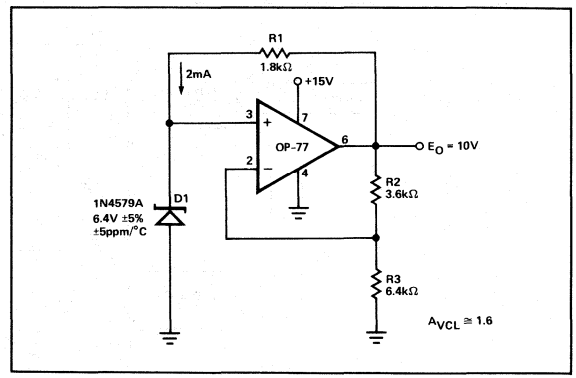

The simple bootstrapped voltage reference provides a precise 10V virtually independent of changes in power supply voltage, ambient temperature, and output loading. Correct zener operating current of exactly 2mA is maintained by R1, a selected 5ppm/°C resistor, connected to the regulated output. Accuracy is primarily determined by three factors: the 5ppm/°C temperature coefficient of D1, 1ppm/°C ratio tracking of R2 and R3, and operational amplifier V_{OS} errors. The OP-77, with TCV_{OS} of 0.3μV/°C, contributes only 0.05ppm/°C of output error, thus effectively eliminating TCV_{OS} as an error consideration.

FIGURE 18. Precision Dual Tracking Voltage References

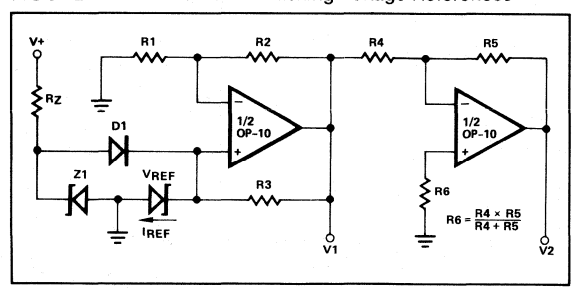

Precision dual tracking voltage references using a single reference source are easily constructed using OP-10. These references exhibit excellent stability vs. temperature/time, low noise, and excellent power supply rejection.

In the circuit shown, R3 should be adjusted to set I_{REF} to operate V_{REF} at its minimum temperature coefficient current. Proper circuit start-up is assured by R_Z, Z1, and D1.

$$V_{Z1} \leq V_{REF} \qquad V1 = V_{REF}\left(1 + \frac{R2}{R1}\right)$$

$$R3 = \frac{(V1 - V_{REF})}{I_{REF}} \qquad V2 = V1\left(\frac{-R5}{R4}\right)$$

Output Impedance (ΔI_L:1.0mA—5.0mA) ..... $0.25 \times 10^{-3}\Omega$

FIGURE 19: RIAA Phone Pre-Amplifier

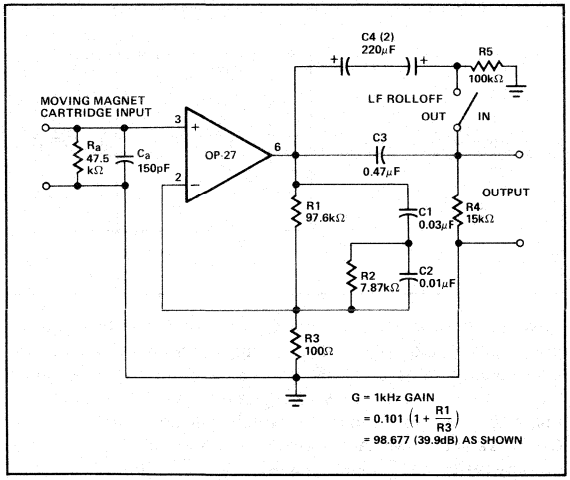

The OP-27 is used in this phono pre-amplifier circuit because of its low noise characteristics. It contributes only $3.2nV/\sqrt{Hz}$ voltage noise and $0.45pA/\sqrt{Hz}$ current noise to the circuit. To minimize noise from other sources, R3 is set to 100Ω. This generates additional voltage noise of only $1.3nV/\sqrt{Hz}$. With a 1kΩ source, the circuit noise measures 63dB below a 1mV reference level, unweighted, in a 20kHz bandwidth.

R1, R2, C1, and C2 form a very accurate RIAA network with standard component values providing the necessary time constants of 3180, 318, and 75µsec. For initial equalization accuracy and stability, precision metal-film resistors and film capacitors of polystyrene or polypropylene are recommended since they have low voltage coefficients, dissipation factors, and dielectric absorption. (High-K ceramic capacitors should be avoided here, though low-K ceramics—such as NPO types, which have excellent dissipation factors, and somewhat lower dielectric absorption—can be considered for small values.)

Capacitor C3 and resistor R4 form a simple −6dB-per-octave rumble filter, with a corner at 22Hz. As an option, the switch-selected shunt capacitor C4, a nonpolarized electrolytic, bypasses the low-frequency rolloff. Placing the rumble filter's high-pass action after the pre-amp has the desirable result of discriminating against the RIAA-amplified low-frequency noise components and pickup-produced low-frequency disturbances.

This circuit is capable of very low distortion over its entire range, generally below 0.01% at levels up to 7Vrms. At 3V output levels, it will produce less than 0.03% total harmonic distortion at frequencies up to 20kHz.

FIGURE 20: Headphone Amplifier

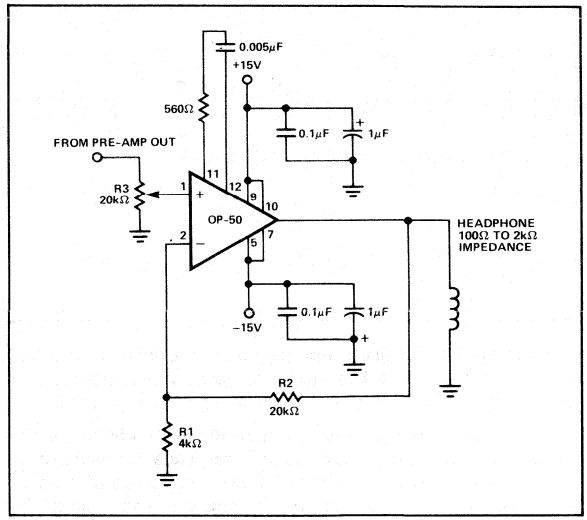

For low level Pre-Amp Out signals, the amplifier gain may be increased by reducing R1 according to:

$$R1 = \frac{20k\Omega}{A_V - 1}$$

Note that two amplifiers are required for stereo applications.

Performance: ($V_{OUT} = 6V_{RMS}$, R1 = 4kΩ)

T.H.D @ 100Hz = 0.0025%

@ 1kHz = 0.003%

@ 10kHz = 0.011%

Signal-to-Noise Ratio ≥ 80dB

Response Flatness = ±0.4dB from 10Hz to 20kHz

Bandwidth = −3dB @ 56kHz

FIGURE 21: NAB Tape Head Pre-Amplifier

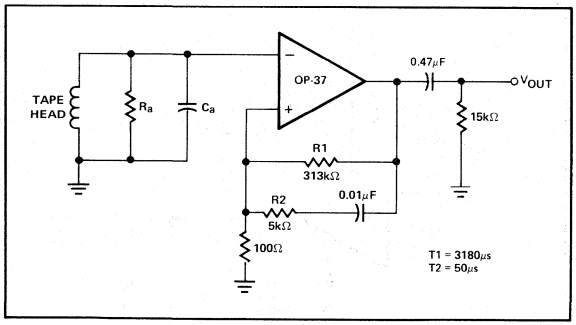

A pre-amplifier for NAB tape playback is similar to an RIAA phono pre-amp, though more gain is typically demanded, along with equalization requiring a heavy low-frequency boost.

The network values of this configuration yield a 50dB gain at 1kHz, and a DC gain greater than 70dB. Thus, the worst-case output offset is just over 500mV. The DC resistance of the tape head will add to this a bias-current-induced offset voltage. To minimize this contribution, the head's DC resistance should be low, preferably below 1kΩ. A single 0.47μF output capacitor can block the final output offset without affecting the dynamic range.

The tape head can be coupled directly to the amplifier input, since the worst-case bias current of 80nA with a 400mH, 100μin. head (such as the PRB2H7K) will not be troublesome.

One potential tape head problem is presented by amplifier bias current transients at power-up or power-down which can magnetize a head. Although the OP-37 is free of these bias current transients, it is always good design practice to control the speed of power supply rise and fall to eliminate transients.

FIGURE 22: Microphone Pre-Amplifier

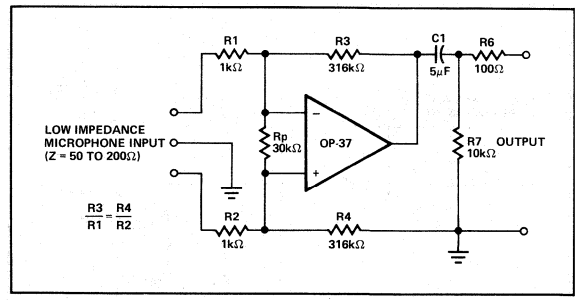

This simple, but effective, fixed-gain transformerless microphone pre-amp amplifies differential signals from low-impedance microphones by 50dB, and has an input impedance of 2kΩ. Because of the high working gain of the circuit, an OP-37 helps to preserve bandwidth, which will be 110kHz. As the OP-37 is a decompensated device (minimum stable gain of 5), a dummy resistor, R_p, may be necessary, if the microphone is to be unplugged. Otherwise the 100% feedback from the open input may cause the amplifier to oscillate.

Common-mode input-noise rejection will depend upon the match of the bridge-resistor ratios. Either close-tolerance (0.1%) types should be used, or R4 should be trimmed for best CMRR. All resistors should be metal-film types for best stability and low noise.

Noise performance of this circuit is limited more by the input resistors R1 and R2 than by the op amp, as R1 and R2 each generate a $4nV/\sqrt{Hz}$ noise, while the op amp generates a $3.2nV/\sqrt{Hz}$ noise. The rms sum of these predominant noise sources will be about $6nV/\sqrt{Hz}$, equivalent to 0.9μV in a 20kHz noise bandwidth, or nearly 61dB below a 1mV input signal. Measurements confirm this predicted performance.

FIGURE 23: Micropower Wien-Bridge Oscillator ($P_d < 500\mu W$)

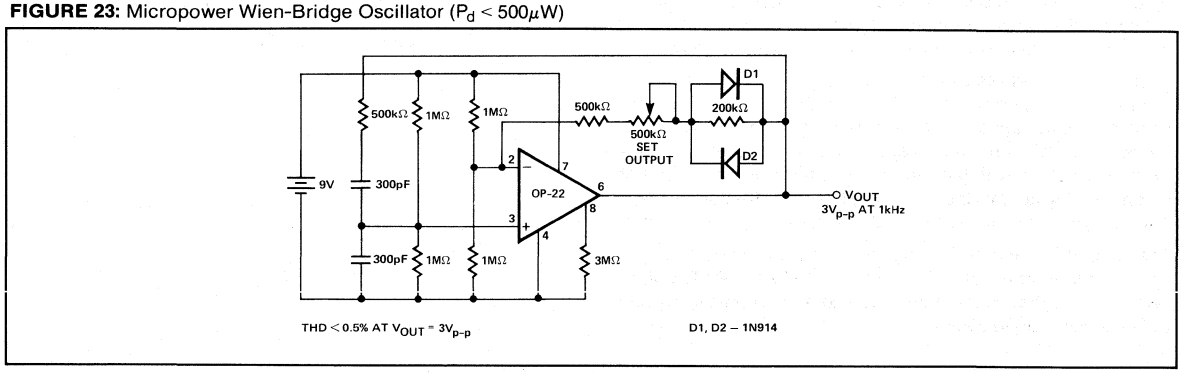

Requiring less than 60μA of supply current, this micropower Wien-bridge oscillator is ideal for battery-powered instrumentation. Output level is controlled by nonlinear elements D1 and D2. When adjusted for $3V_{p-p}$ output, the distortion level is below 0.5% at 1kHz.

FIGURE 24: Micropower Instrumentation Amplifier

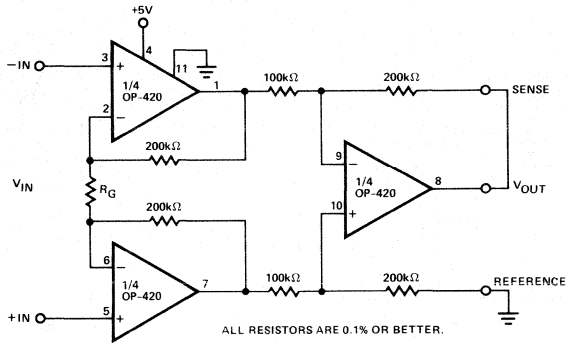

This instrumentation amplifier requires only 200μA total quiescent current (V_{CM} = 0V) and operates on single voltage supplies from +1.6V to +36V. Input and output voltage range is 0V to (V^+ − 1.5) volts with CMRR above 100dB. Differential gain, A_V, is adjusted with a single resistor, R_G, as given by:

$$R_G = \frac{800k\Omega}{A_V - 2}$$

FIGURE 25: Piecewise-Linear Amplifier (Decreasing Gain)

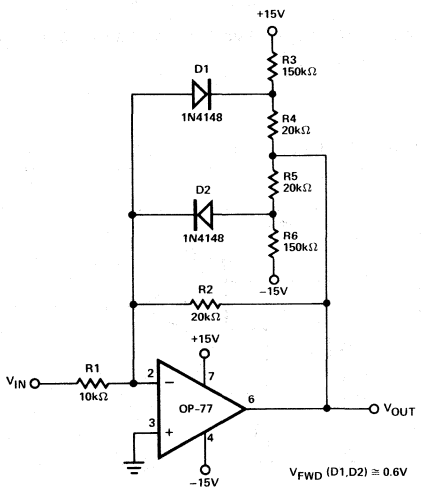

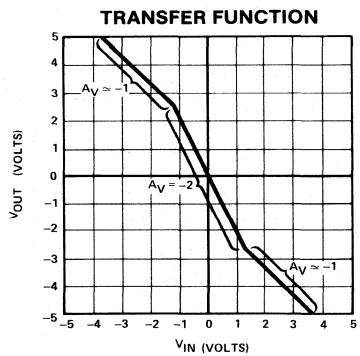

This circuit is useful in linearizing a nonlinear input signal or creating a nonlinear output function from a linear input. At V_{OUT} = 0V, both D1 and D2 are reverse biased, and V_{OUT}/V_{IN} = −R2/R1. As V_{IN} goes positive, V_{OUT} becomes negative according to that gain until a threshold is reached, $-(V_{OUT} + V_{FWD})/R4$ = 15V/R3, where D1 becomes forward biased. For more positive values of V_{IN}, the gain is:

$$\frac{V_{OUT}}{V_{IN}} = -\frac{(R2 \cdot R4)}{(R2 + R4)R1} - \frac{R2 \cdot V_{FWD}}{(R2 + R5)V_{IN}}$$

A similar action occurs as V_{IN} goes negative. Beyond the point where D2 becomes forward biased, $(V_{OUT} - V_{FWD})/R5$ = 15V/R6, the gain is:

$$\frac{V_{OUT}}{V_{IN}} = -\frac{(R2 \cdot R5)}{(R2 + R5)R1} - \frac{R2 \cdot V_{FWD}}{(R2 + R5)V_{IN}}$$

Additional diode/resistor combinations can be added to further contour the gain.

FIGURE 26: Piecewise-Linear Amplifier (Increasing Gain)

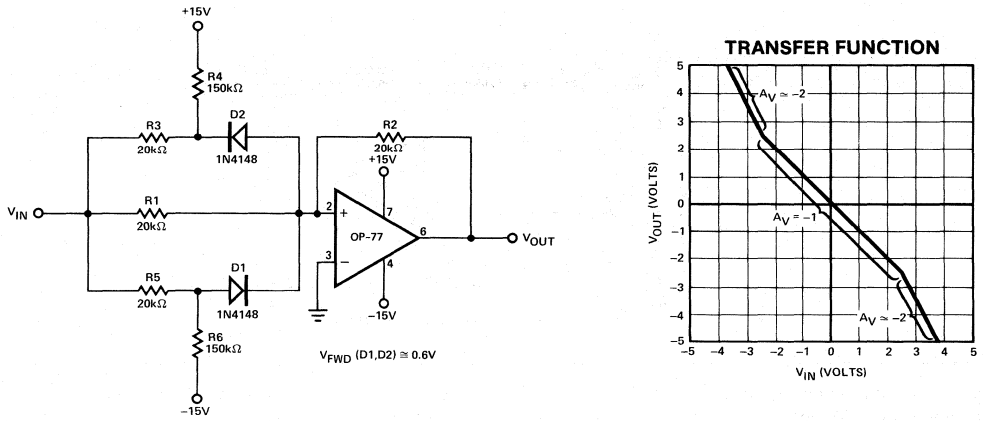

This circuit performs the linear-to-nonlinear or nonlinear-to-linear transformation by increasing gain beyond fixed thresholds. At $V_{IN} = 0V$, both D1 and D2 are reverse biased, and $V_{OUT}/V_{IN} = -R2/R1$. As V_{IN} goes positive beyond the threshold, $(V_{IN} - V_{FWD})/R5 = 15V/R6$, D1 conducts, and the gain becomes:

$$\frac{V_{OUT}}{V_{IN}} = -\frac{R2(R1 + R5)}{(R1 \cdot R5)} + \frac{V_{FWD} \cdot R2}{V_{IN} \cdot R5}$$

As V_{IN} goes negative beyond the threshold, $-(V_{IN} + V_{FWD})/R3 = 15V/R4$, D2 conducts, and the gain becomes:

$$\frac{V_{OUT}}{V_{IN}} = -\frac{R2(R1 + R3)}{(R1 \cdot R3)} + \frac{V_{FWD} \cdot R2}{V_{IN} \cdot R3}$$

Additional diode/resistor combinations can be added to further contour the gain.

FIGURE 27: Current Monitor Circuit

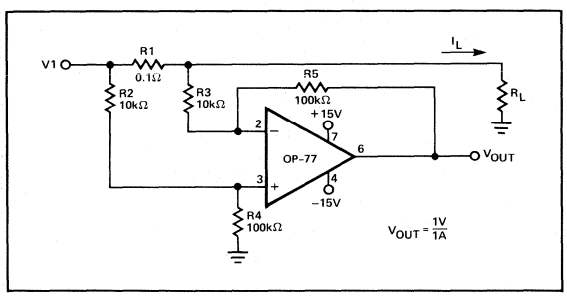

This versatile monitor circuit can typically sense current at any point between the ±15V supplies ($|V1| < 14.3V$ guaranteed). This makes it ideal for sensing current in applications such as full bridge drivers where bi-directional current is associated with large common-mode voltage changes. The 120db CMRR of the OP-77 makes the amplifier's contribution to common-mode error negligible, leaving only the error due to the resistor ratio inequality. Ideally, R2/R4 = R3/R5. This is best trimmed via R4.

FIGURE 28: Free-Running Square-Wave Oscillator

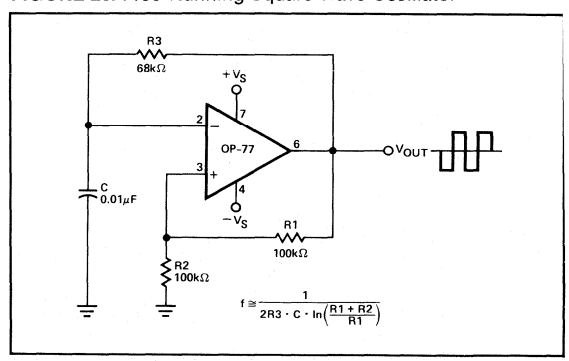

This simple oscillator creates a square-wave output of $\pm(V_S - 2V)$ at 1kHz for the values shown.

FIGURE 29: Precision Analog Multiplier/Divider

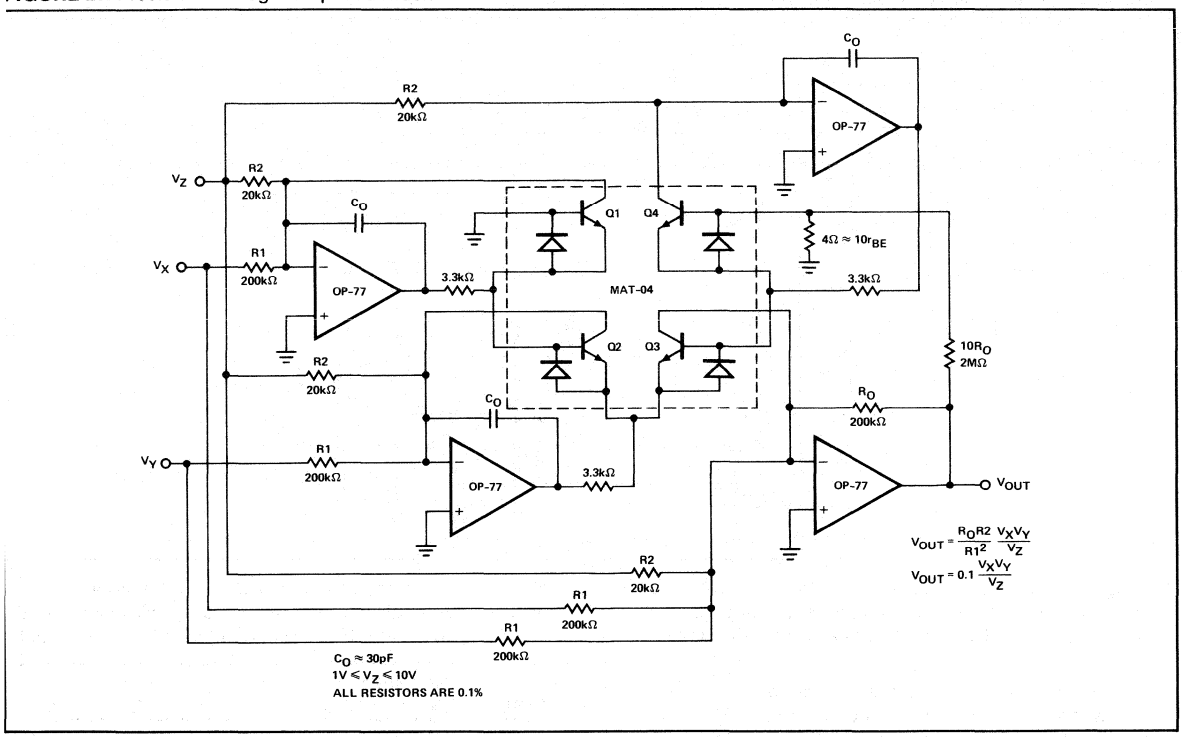

This multiplier/divider achieves its excellent performance through the low emitter resistance, r_{BE}, and superior V_{OS} matching of MAT-04 quad transistor array. In this circuit, linearity error due the transistors is less than ±0.1%. The OP-77 helps maintain accuracy with a V_{OS} less than 25μV. For even higher accuracy the offset voltages may be nulled.

13

FIGURE 30: Precision Absolute Value Amplifier

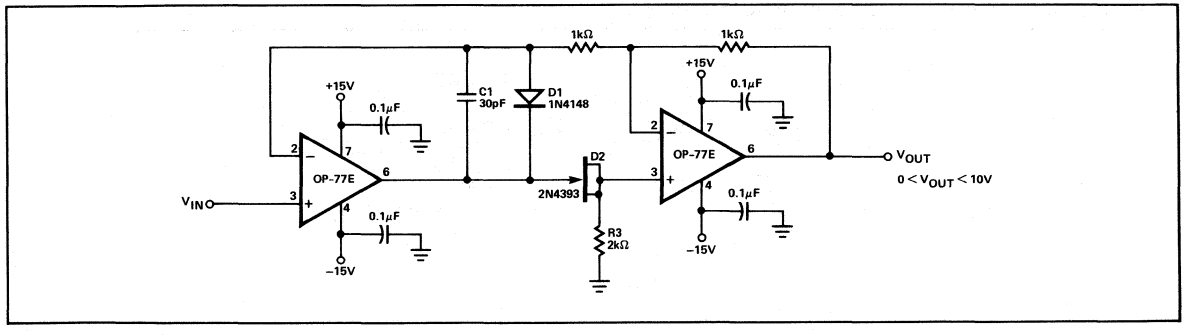

The high gain and low TCV_{OS} assure accurate operation with inputs from microvolts to volts. In this circuit, the signal always appears as a common-mode signal to the op amps. The OP-77E CMRR of 1μV/V assures errors of less than 2ppm.

FIGURE 31: Thermocouple Amplifier With Cold-Junction Compensation

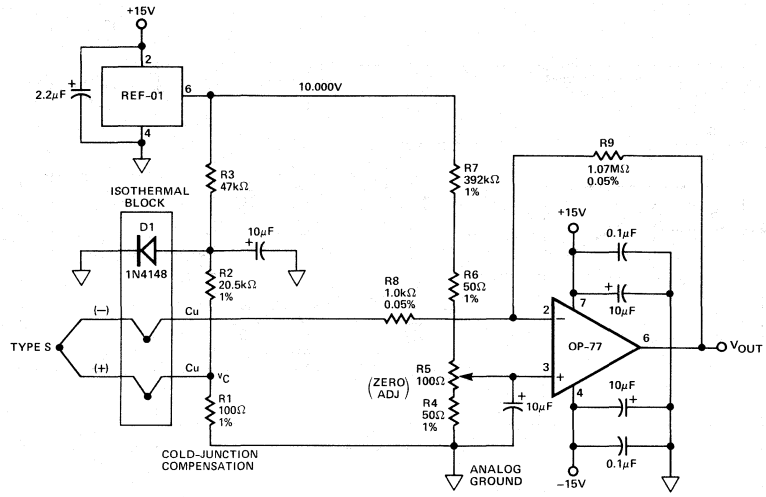

The high gain, low noise, and low offset drift of the OP-77 can be used to create a thermocouple amplifier with superb linearity. They combine to give a total accuracy typically better than ±0.5°C.

Cold-junction compensation is performed by R1, R2, and D1 which is mounted isothermally with the thermocouple terminating junctions. Calibration is done using R5 after the circuit has stabilized for about 15 minutes. A copper wire short is applied across the terminating junctions simulating a zero °C ice point. R5 is then adjusted for 0.000V output. The short is then removed, and the amplifier is ready for use.

This amplifier can be used for type S, J, and K thermocouples with the appropriate resistor values as shown below. R9 is chosen to give a +10.000V output for a 1,000°C measurement.

TYPE	SEEBECK COEFFICIENT, α	R1	R2	R7	R9
K	39.2µV/°C	110Ω	5.76kΩ	102kΩ	269kΩ
J	50.2µV/°C	100Ω	4.02kΩ	80.6kΩ	200kΩ
S	10.3µV/°C	100Ω	20.5kΩ	392kΩ	1.07MΩ

FIGURE 32: Instrumentation Amplifier (2 Op Amp Design)

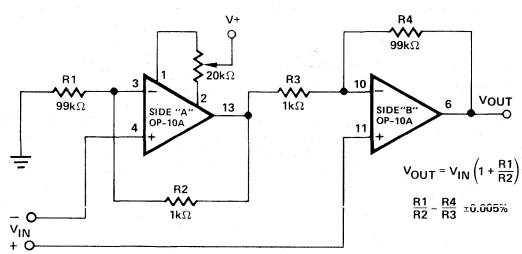

This differential amplifier offers high input impedance and a power supply rejection ratio greater than 100dB. Because the high circuit gain occurs in Side "B"; Side "A" offset should be adjusted with respect to V_{OUT}. This simultaneously corrects for Side "B" offset voltage. For the circuit values given, $A_V = 100$.

IGURE 33: ±200V Low Offset Operational Amplifier

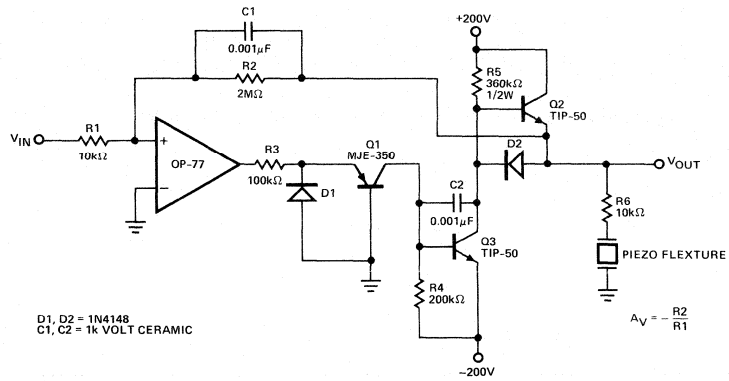

The OP-77 is the heart of this high voltage amplifier designed to drive piezo flexture elements in precise positioning applications. The high gain and low offset voltage of the OP-77 produce an accurate, stable drive voltage to the piezo device allowing predictable, repeatable sub-micron movements to be easily controlled.

The output of the OP-77 creates a proportional current drive through the common-base connected Q1 to the base of Q3. Q3 forms a Class A amplifier with bias resistor R5, and transistor Q2 is simply an emitter-follower used for output current boost. R6 is necessary to prevent oscillation caused by the capacitive loading of the piezo device on the output of the amplifier.

FIGURE 34: Impedance Transforming Amplifier

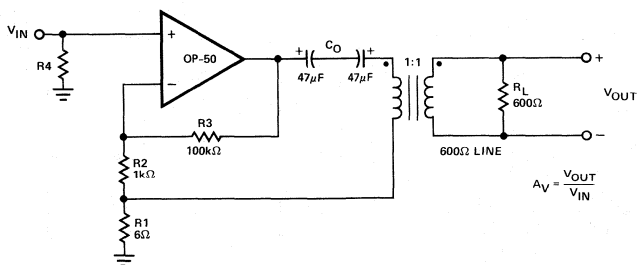

This is an efficient, flexible circuit simulating a source impedance equal to the load impedance. By definition, if $R_S = R_L$, then,

$$A_V = \frac{A_{V(UNLOADED)}}{2}$$

The unloaded gain is roughly R3/R2. When the output is loaded, a second feedback loop is closed whose gain, $A_V \simeq R_L/R1$, combines in parallel with $A_{V(UNLOADED)}$ to give:

$$A_{V(LOADED)} = \frac{A_{V(UNLOADED)} \cdot A_V}{A_{V(UNLOADED)} + A_V}$$

If R3/R2 = R_L/R1, then $A_{V(LOADED)} = \frac{A_{V(UNLOADED)}}{2}$

These approximations assume that R1 ≪ R2 and R2 ≪ R3. The OP-50 requires no compensation for the circuit values shown and can easily drive the 600Ω line.

FIGURE 35: Precision Current Sinks

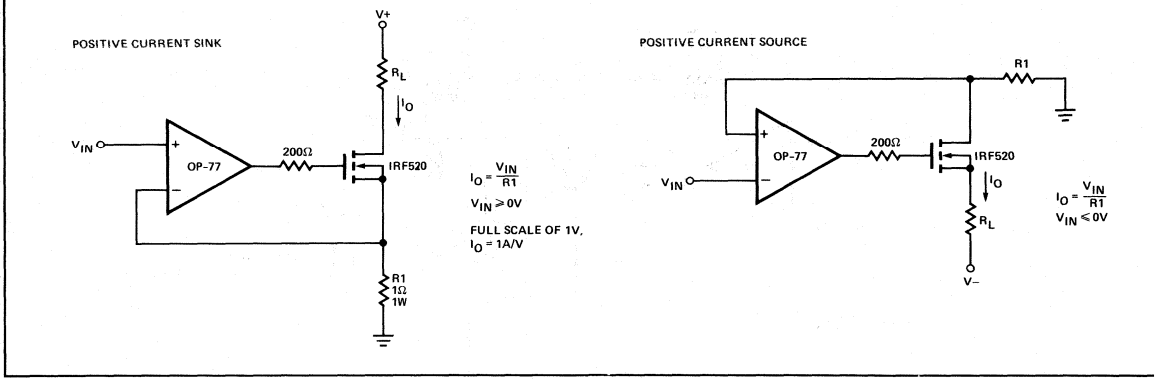

These simple high-current sinks require that the load float between the power supply and the sink. In these circuits, OP-77's high gain, high CMRR, and low TCV_{OS} assure high accuracy.

FIGURE 36: Low Noise AGC Amplifier

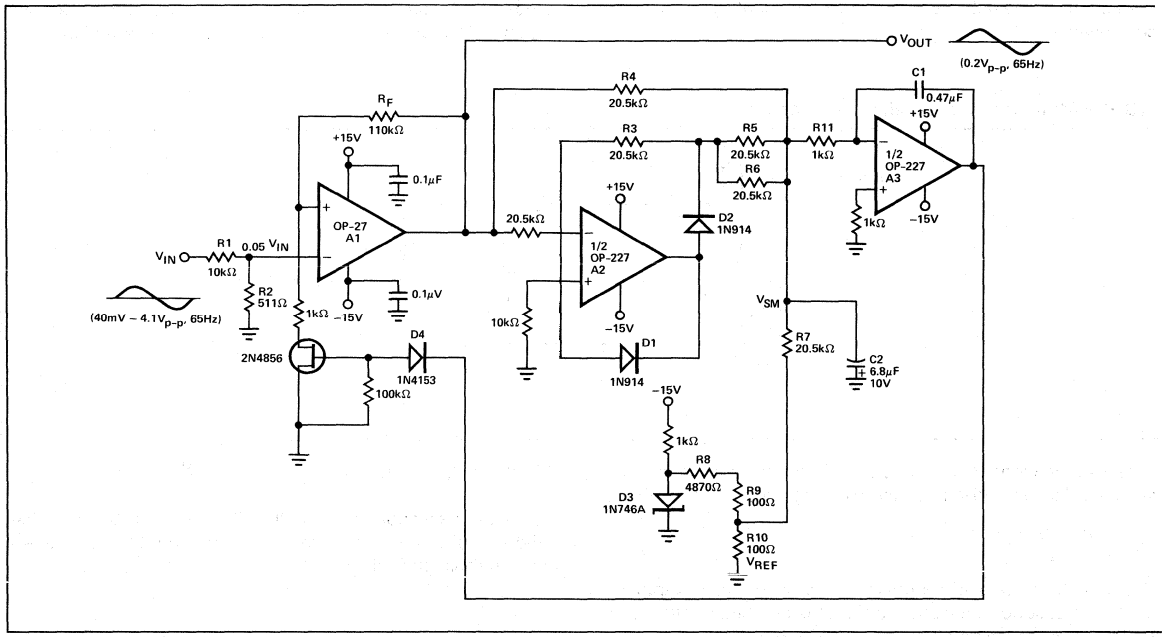

In this circuit, a JFET transistor is used to control the gain of the low-noise OP-27 amplifier over a two-decade input voltage range. For inputs from 40mV to 4.1V peak-to-peak, the AGC maintains a 0.2V peak-to-peak output.

Amplifier A2 performs an absolute-value operation on V_{OUT} and sums the result with a –0.2V reference on capaciator C2. The deviation of this sum, V_{SM}, from zero is amplified by A3 and controls the gate of the JFET. If the peak-to-peak amplitude of V_{OUT} exceeds 0.2V, V_{SM} becomes positive and drives the JFET gate negative. This increases the JFET's channel resistance lowering the gain of the A1. The reverse of this occurs if V_{OUT} falls below 0.2V peak-to-peak.

The values of C1 and C2 are chosen to optimize the circuits response time for a given input voltage frequency. This example was designed for a 65Hz signal. Higher frequencies would justify lower values for C1 and C2 to speed the AGC response.

GURE 37: Amplifier with Active Output Clipping

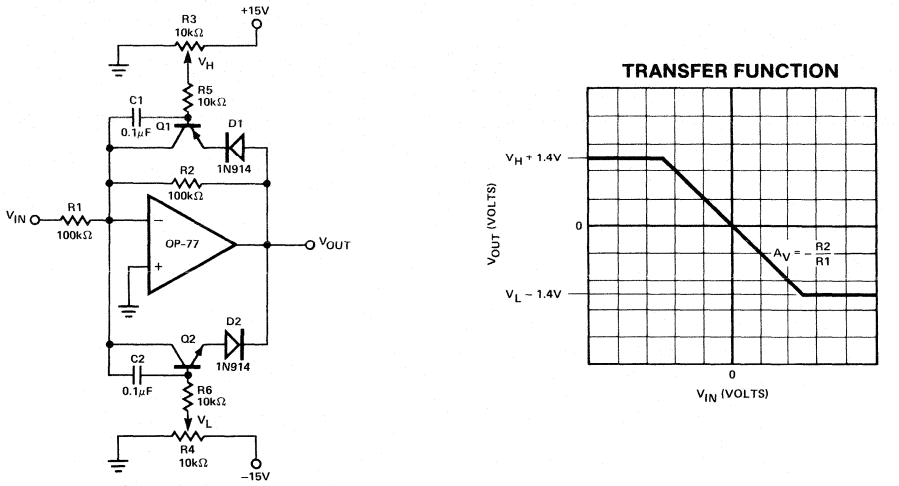

This configuration allows adjustment of the op amp's maximum output voltage. Below the clipping levels, circuit gain is $A_V = -R2/R1$. As V_{OUT} rises above ($V_H + 1.4V$), the base/emitter of Q1 becomes forward biased, allowing collector current to flow to the summing node, thus clamping V_{OUT}. A similar action occurs as V_{OUT} goes below ($V_L - 1.4V$) and is clamped by Q2.

FIGURE 38: Low Power Amplifier With Squelch

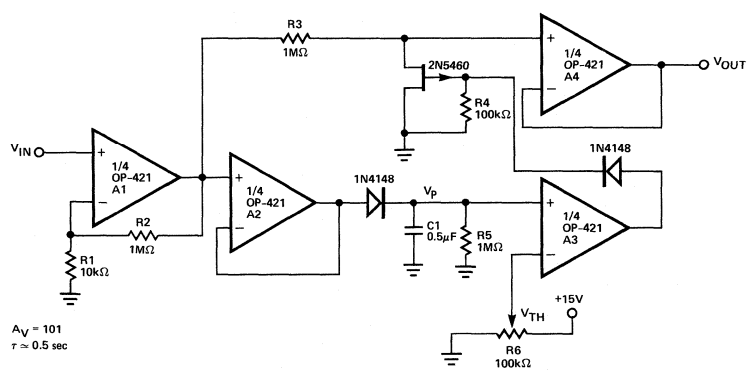

The OP-421 is the heart of this variable squelch amplifier which requires less than 2mA of supply current ($R_L = \infty$). A1 provides a high impedance input amplifier with $A_V = 1 + R2/R1$. Its output drives a unity gain output buffer, A4, and a peak detector with a time constant set by C1 and R5. When the output of the peak detector, V_P, exceeds the adjustable threshold, V_{TH}, the comparator, A3, drives the gate of the P-channel FET high, turning it OFF. At lower input signal levels, V_P falls below V_{TH}, and the FET turns ON, clamping the input of A4 to ground.

ANALOG DEVICES

AN-40
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

A Buffer Amplifier Applications Collection

INTRODUCTION

This Application Note consists of a collection of circuits which use buffer amplifiers in a variety of applications. As will be shown, buffers may be used to make filters, current sources, cable drivers, sample and holds, line drivers for multiplexers, current boosters, and high speed voltage output DACs.

INDUCTORS AND FILTERS

The active inductor in Figure 1 is realized with an eight-lead IC, two carbon resistors, and a small capacitor. A commercial inductor of 50 henries may occupy up to five cubic inches.

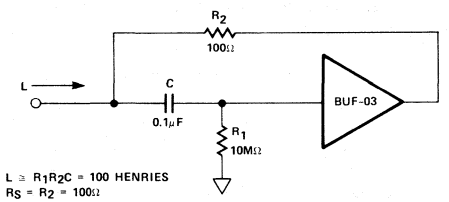

Figure 1. Active Inductor

The tuned circuit shown in Figure 2 uses the simulated inductor of Figure 1 (R_1, R_2, C_1) and C_2. Depending upon whether the circuit is driven at E_1 or E_2 the responses of Figures 3 or 4 result. The resonant response in both cases is

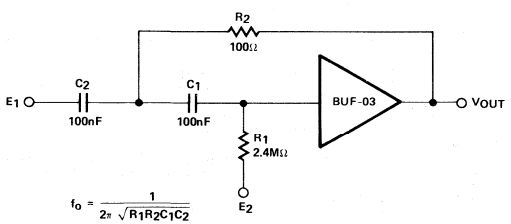

Figure 2. Tuned Circuit

+ 38dB at 103Hz. The Figure 3 response is + 2.5dB at 200Hz and − 10dB at 50Hz. On the other hand, the Figure 4 response is − 9dB at 200Hz and + 2.5dB at 50Hz.

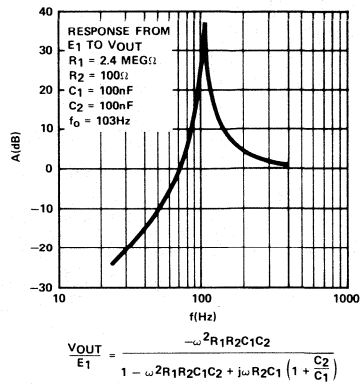

Figure 3. Response from E_1 to V_{OUT}

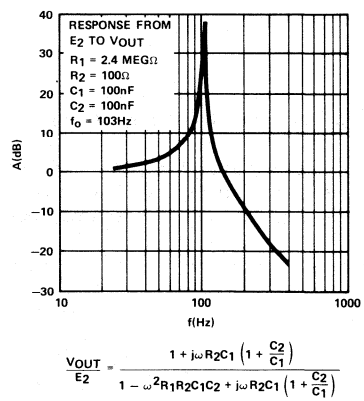

Figure 4. Response from E_2 to V_{OUT}

Figure 5 shows a low pass filter realized for f_o of 1MHz. What is remarkable about this filter is most ICs do not have the full power bandwidth to handle 1MHz signals in the 5 to 10 Volt range, while the BUF-03 has a greater than 4MHz full power bandwidth for a $20V_{p-p}$ sinewave. Similar comments apply to the filter in Figure 6. In other words, the outstanding bandwidth of the BUF-03 extends the bandwidth capability of certain classes of active filters.

13

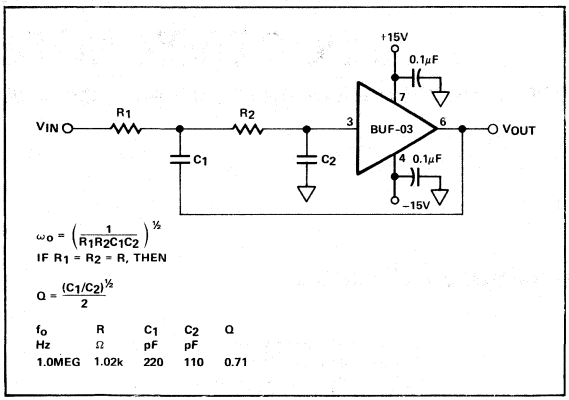

Figure 5. Low Pass Filter (High Frequency)

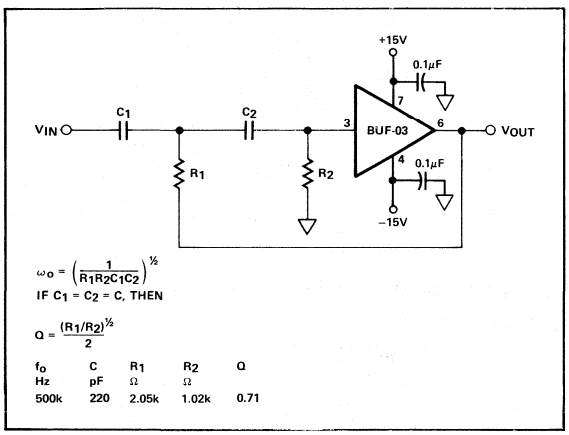

Figure 6. High Pass Filter (High Frequency)

The BUF-03 can be used to make a 4.5MHz trap for use in TV. This circuit is shown in Figure 7, and the elements are chosen such that no capacitor is less than 100pF.

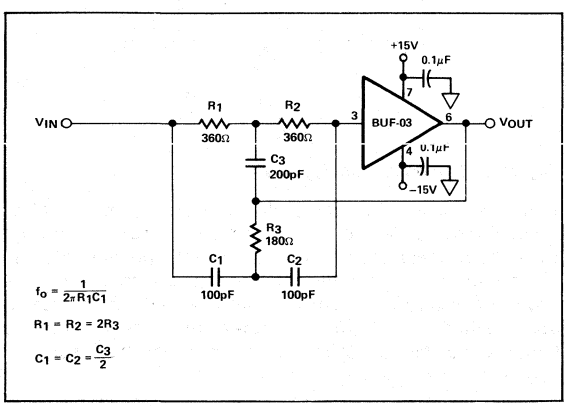

Figure 7. Notch Filter at 4.5MHz

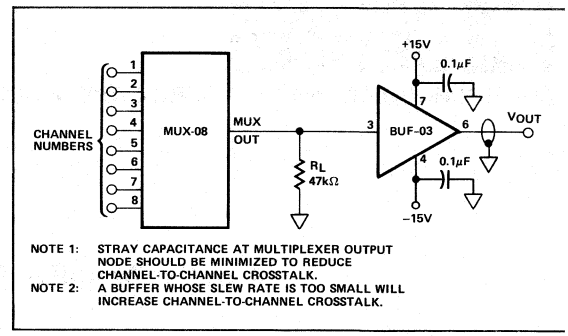

Figure 8. High Speed Line Driver for Multiplexers

The BUF-03 can be used as a data line driver because of its speed and current drive capabilities. The connection for this application is shown in Figure 8. The realization of a high speed sample and hold is possible using the BUF-03 and suitable analog switches. The circuit shown in Figure 9 provides the highest speed because there are no feedback loops to slow down the settling times. Typically the sample and hold is followed by a successive approximation analog-to-digital converter (ADC). The final application involves the BUF-03 and the DAC-08 (digital-to-analog converter). Figure 10 shows how it is possible to develop both V_{OUT} and $\overline{V}_{OUT}$. The output capacitance of the DAC-08 is approximately 15pF, thus as R_0 increases in value, so does the settling time for V_{OUT} (and $\overline{V}_{OUT}$).

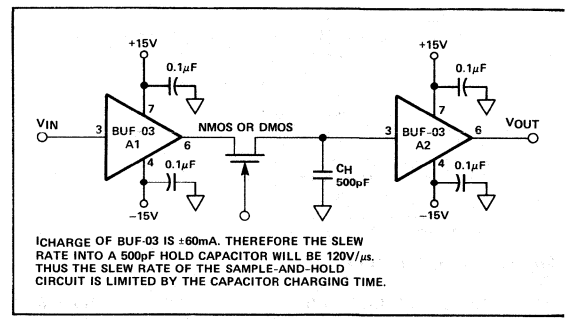

Figure 9. High Speed Sample and Hold

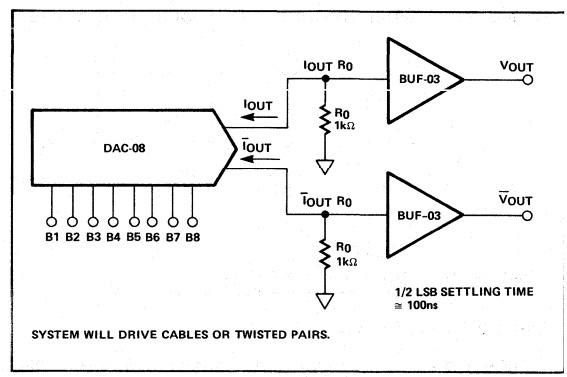

Figure 10. High Speed Voltage Output DAC

LINE DRIVER APPLICATIONS

If your FET "line driver" has the speed but not the stability or the current capability to drive coaxial cables, its output may be buffered with a BUF-03 as shown in Figure 11. Figure 12 shows an alternative connection when better accuracy and high current capability is needed. Note that the limitation on R_L being greater than 1kΩ does not apply in this case since the added error caused by lower impedances is imbedded inside the feedback loop of the op amp.

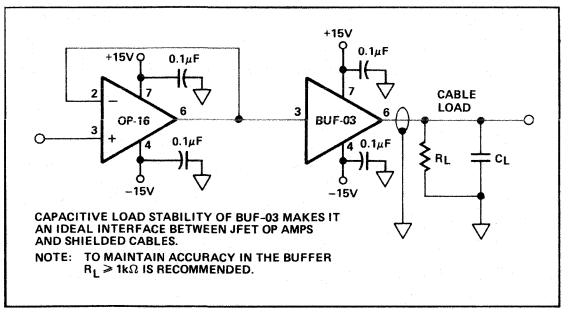

Figure 11. Convert FET Op Amp Into Cable Driver

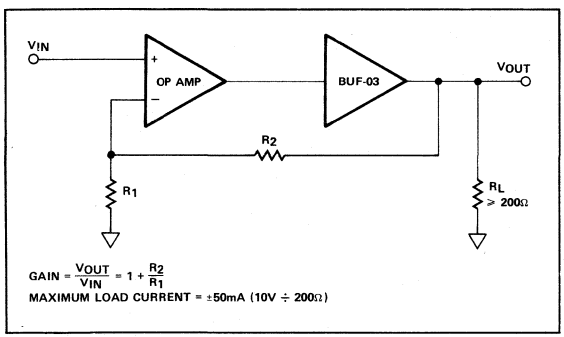

Figure 12. Current Booster

MISCELLANEOUS USES OF BUFFERS

Single supply applications can be realized using the BUF-03 as shown in Figure 13. The input is DC biased to +10V such that the BUF-03 operates in the linear region. Signal is AC coupled to the input and also AC coupled to the output load resistor.

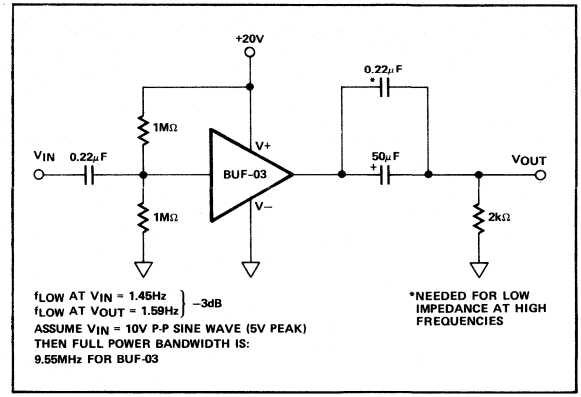

Figure 13. Single Supply AC Buffer (High Speed)

CONCLUSION

While the list is by no means all inclusive, this application note has attempted to point out some of the myriad of uses for the IC buffer. In particular, the BUF-03 makes possible a whole new class of high frequency filters and high speed current sources. Many problems in data acquisition systems can be solved by the use of buffers. In addition, the BUF-03 is useful in providing increased drive current, as well as the ability to drive long cables without instability. Finally, the versatility of the reference zener can be increased by using buffers, and for AC applications the buffer can be used on single power supplies.

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Applying the OP-06 Op Amp As a High Precision Comparator

INTRODUCTION

The Analog Devices OP-06 op amp makes an excellent comparator. In fact, for submillivolt signals, there is simply no comparator that performs as well. Using an external nulling potentiometer, the offset drift is typically 0.6μV/°C. With its high open loop gain to 1 million, only 30μV is required at the input to drive the output from one saturation level to the other. A 50°C change in temperature produces a 30μV change in V_{OS}; thus a *total* error band of 100μV *including temperature effects* is quite conservative. This performance is an order of magnitude better than other comparators. 100μV sensitivity is nice to have in 12-bit A/D converters, but it is essential in 14-bit converters. Where preamplifiers are typically needed with thermocouples and strain guages, the OP-06's sensitivity allows *direct* comparison of these low-level outputs. As a result system costs decrease, and reliability increases.

LOW-LEVEL PERFORMANCE MEASUREMENTS

The low-level capabilities of the OP06 comparator are graphically illustrated in Figures 1 and 2 using the test circuit below. Comparator voltage input, applied through a 100 to 1 attenuator, is $100\mu V_{p-p}$ in Figure 1 and $40\mu V_{p-p}$ in Figure 2. Note that the op amp output still reaches both positive and negative saturation.

TEST CIRCUIT

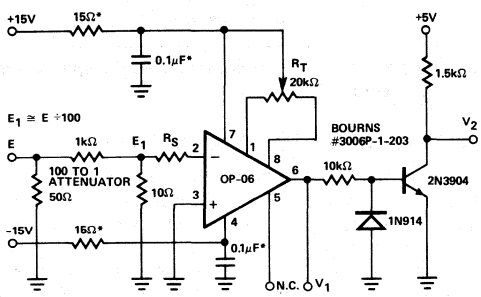

*DECOUPLING COMPONENTS NECESSARY TO PREVENT SPURIOUS OSCILLATION

PROCEDURE:
1. REDUCE E TO ZERO VOLTS
2. ADJUST R_T TO BRING V_1 INTO LINEAR RANGE (USE SCOPE)
3. APPLY E THAT IS SYMMETRICAL ABOUT GROUND, IE +E / -E
4. MEASURE V_1 AND V_2 WITH SCOPE

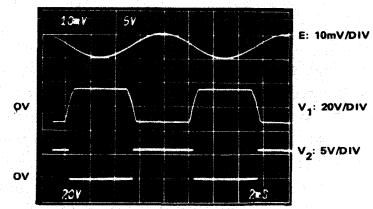

Figure 1. $100\mu V_{p-p}$ Sine Wave Response (R_S: 100Ω)

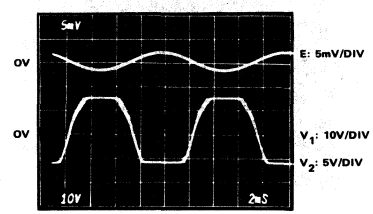

Figure 2. $40\mu V_{p-p}$ Sine Wave Response (R_S: 100Ω)

COMPARATOR RESPONSE TIME

While most comparators are specified for 2mV to 5mV overdrive, the OP06 operates very reliably with only 0.5mV overdrive. Figures 3 and 4 show the response times for both positive going and negative going inputs with 500μV and 5mV overdrives as measured at the logic output.

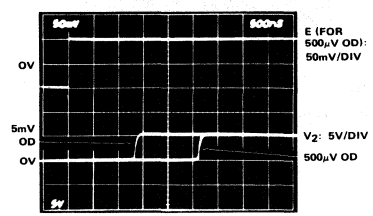

Figure 3. Positive Going Response Time (5mV and 500μV Overdrives)

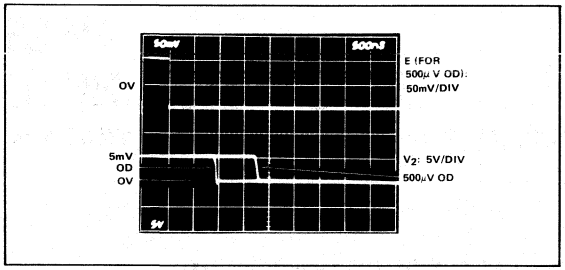

Figure 4. Negative Going Response Time (5mV and 500μV Overdrives)

OP AMP RESPONSE TIME

Primarily, the very high open loop gain (A_{VO}) of the OP06 — over one million — is responsible for its success as a comparator. The DC gain, as specified on op amp data sheets, is important for comparison sensitivity (V_{DET}). However, it is the shape of the gain curve with frequency that dictates how fast the op amp will switch as it passes through its linear region. When operated as a comparator the OP06 spends most of its time in either positive or negative saturation (see Figures 5 and 6). Saturation effects are discussed later; but for now notice the difference in slew rates for the overdrive levels.

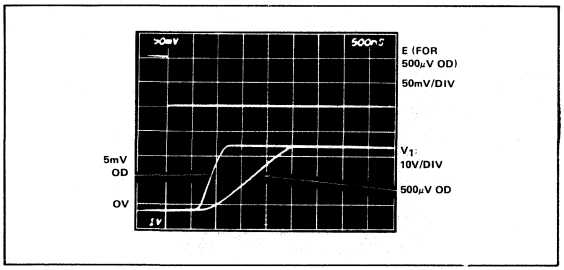

Figure 5. Positive Going Op Amp Response Time vs Overdrive

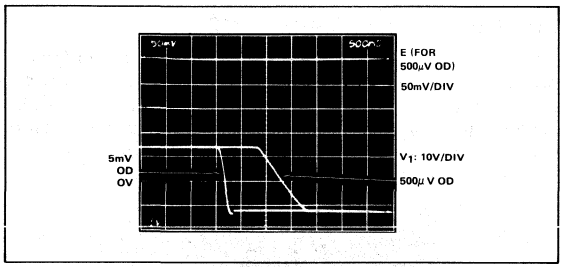

Figure 6. Negative Going Op Amp Response Time vs Overdrive

OP06 DYNAMIC PERFORMANCE CHARACTERISTICS

There is another factor which has a significant effect on slew rates — the way in which the AC gain is rolled off versus frequency. If we refer to Figure 7 we see gain compensation curves for closed loop gains from 1 to 10,000. Figure 8 relates the slew rate to these frequency response curves. These curves point out one of the tradeoffs between good op amp performance and good comparator performance. For example, if an OP06 has a gain compensation for $A_V = 10$, then its slew rate would be 0.08V/μs. This would result in a rise time (20 volt swing) of 250μs. The fact that a designer needs no compensation with the OP06 — when operated as a comparator — allows the rise times observed in Figures 5 and 6.

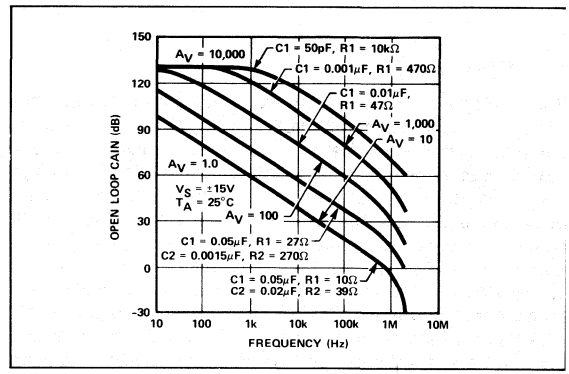

Figure 7. Open Loop Response for Values of Compensation

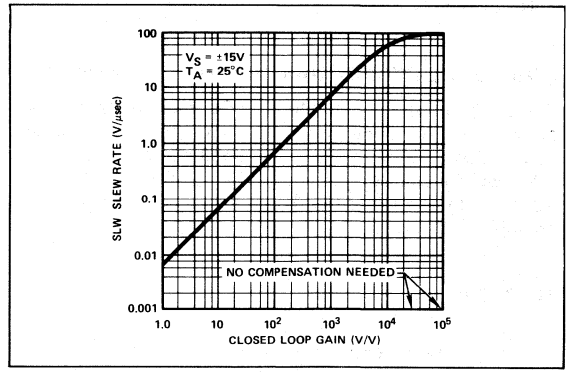

Figure 8. Slew Rate Using Recommended Compensation Networks

COMPARATOR, OP AMP SIMILARITIES

In an op amp, the offset voltage is that voltage which must be applied to the input to drive the output to zero volts. In a comparator this definition is modified to a *specified voltage range* at the output. In ths way the required voltage "window" includes the normal offset voltage of op amps and the signal voltage needed to move the output by some ΔV. Since most

op amps operate in ±15 volt systems, an output voltage range of ±15 volts (or a ΔV of 30 volts) has been chosen. Using this range assures saturation at both the positive and negative extremes (−14 volts and +12 volts for the OP06). Low offset voltage and high gain combine to produce the comparator "detector window."

OTHER FACTORS AFFECTING SPEED

To gain further insight into the relation between overdrive and the various switching times, the graph in Figure 9 was generated from measurements on the OP06 "comparator." To further characterize the OP06 performance, delay times were measured versus source resistance (R_S) with a fixed 5mV overdrive. This curve is shown in Figure 10. Since the rise and fall times were essentially constant with R_S variations, they were not plotted. The delay times are the main contributors to total comparator response time. Since the OP06 was not designed as a comparator, individual gain stages will go into saturation when the output voltage is driven to one of its limits. One of the differences between designing op amps and comparators is the addition of clamp diodes to prevent the above mentioned saturation.

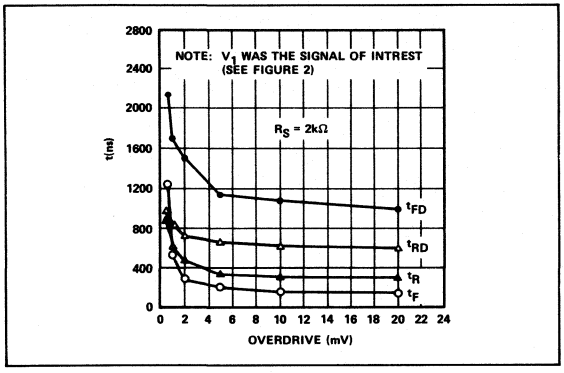

Figure 9. Rise, Fall and Delay Times vs Overdrive Signal

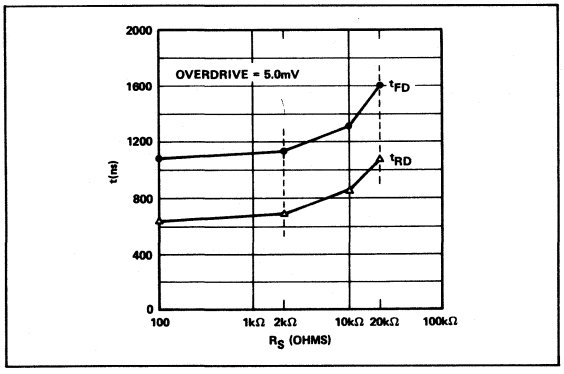

Figure 10. Delay Times vs Source Resistance (R_S)

NOISE AND POWER SUPPLY REJECTION

When dealing with sub-millivolt signals, noise referred to the comparator input becomes an important factor. Basically, the noise comes from two sources:

1) Normal input noise of an op amp;

2) Noise induced by power supply ripple.

Figure 11 shows the wideband noise-on an RMS basis-vs. system bandwidth. What is more important is the RMS to peak conversion factor. Table I shows the crest factors for gaussian noise. Note in particular that the crest factor is less than four 99.99% of the time, and less than five 99.9999% of the time. Thus the RMS noise is 1μV for 10kHz bandwidth and this yields a "worst case" of 5μV peak or 10μV peak-to-peak.

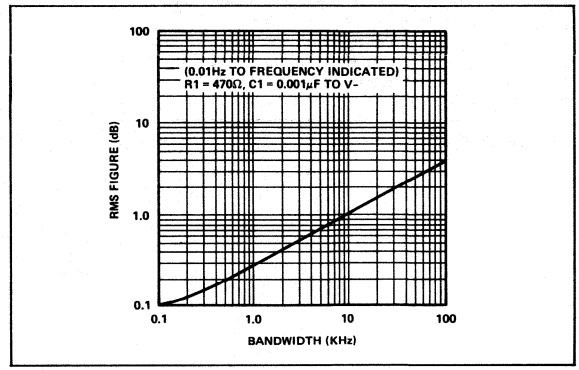

Figure 11. Input Wideband Noise vs Bandwidth

Table 1. Crest Factors for Gaussian Noise (1)

% OF TIME PEAK IS EXCEEDED	PEAK RMS	PEAK FACTOR IN dB = 20 $\log_{10}$ PEAK RMS
10.0	1.645	4.32
1.0	2.576	8.22
0.1	3.291	10.35
0.01	3.890	11.80
0.001	4.417	12.90
0.0001	4.892	13.79

The other source of noise comes from the power supplies. Looking at Figure 12 note that the power supply rejection ratio (PSRR) is 115dB (1.8μV/V) out to 300Hz. For example a power supply which had 1 volt (peak-to-peak) ripple would only produce 1.8μV peak-to-peak "noise." Thus it becomes obvious that the total noise performance of the OP06 is indeed outstanding.

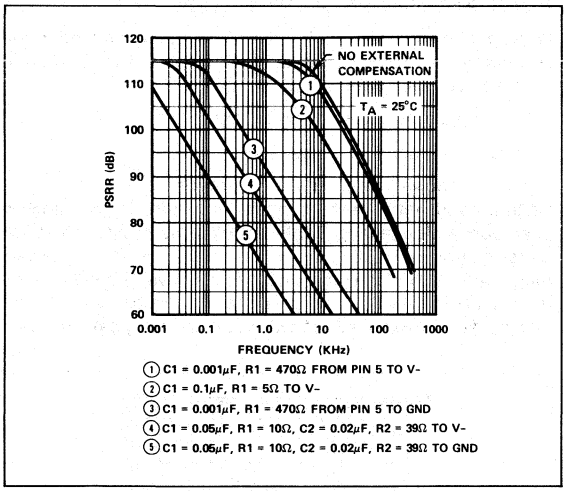

Figure 12. PSRR vs Frequency

CONCLUSION

The combination of ±30V input overvoltage protection, gain of 1 million, low noise, low drift and external compensation allow operation of the OP06 op amp as a low-level comparator. Low-level performance is unsurpassed by any presently available comparator.

BIBLIOGRAPHY

1. Bennett, W. R., "Electrical Noise," McGraw-Hill, New York, 1960, page 44.

ANALOG DEVICES

AB-107
APPLICATION NOTE

)NE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Two-Wire, 4–20 mA Current Transmitter

by James Wong

'his two-wire current transmitter provides an output of 4mA ɔ 20mA that is proportional to an input voltage V_{IN} plus an ffset. Current loops are particularly useful in process control ystems where remote analog signal conditioners must be ıterfaced to a central location. The loop can be powered by n inexpensive, unregulated DC voltage. The low supply cur- ent needs of the OP-22 programmable op amp and REF-02 andgap reference allow for "floating" operation. The trans- ıitter circuit uses less than 2mA and can therefore supply up ɔ 2mA at 5V as a transducer reference or bridge supply ithout exceeding the minimum loop current of 4mA. The)P-22 and REF-02 can be operated over a wide supply range. Vith a load resistor R_L of 50Ω and a sense resistor R_S of 100Ω, he maximum voltage from Ground to Signal Common is 150Ω < 20mA, or 3V. The REF-02 minimum limit is 7V, therefore V_S ıeeds to be above 10V.

'he OP-22 regulates the output I_O to satisfy the current sum- nation at the noninverting mode:

$$\frac{V_{IN}}{R_1} + \frac{5V}{R_2} - \frac{I_O R_S}{R_3} = 0$$

$$I_O = \frac{1}{R_S}\left(\frac{R_3}{R_1} V_{IN} + \frac{R_3}{R_2} 5V\right)$$

As a design example, consider a system need for:

$$I_O = \frac{16V_{IN}}{100\Omega} + 4mA$$

This would provide an output span of 4mA to 20mA for an input range of zero to 100mV. This requires a ratio of 16 for R_3/R_1, and a ratio of 0.08 for R_3/R_2. Choosing R_1 to be 5kΩ, then we need $R_3 = 80k\Omega$ and $R_2 = 1M\Omega$. Drift due to input bias current of the OP-22 can be minimized by making R_4 equal to the parallel combination of R_1, R_2, and R_3.

Designing for other input ranges or other values of R_S and R_L is straightforward. The sense resistor R_S does have an upper limit that is not obvious; the voltage drop across R_S at turn-on can pull the OP-22 noninverting input negative relative to its own negative supply rail. This can cause the OP-22 op amp output to go for the positive limit which drives Q1 into saturation and a possible latching condition. This is prevented by limiting the negative voltage at the noninverting input or by limiting the maximum drop across R_S.

This current transmitter has excellent linearity, operates well with very low supply currents, and is easily adaptable to a wide range of input signal levels.

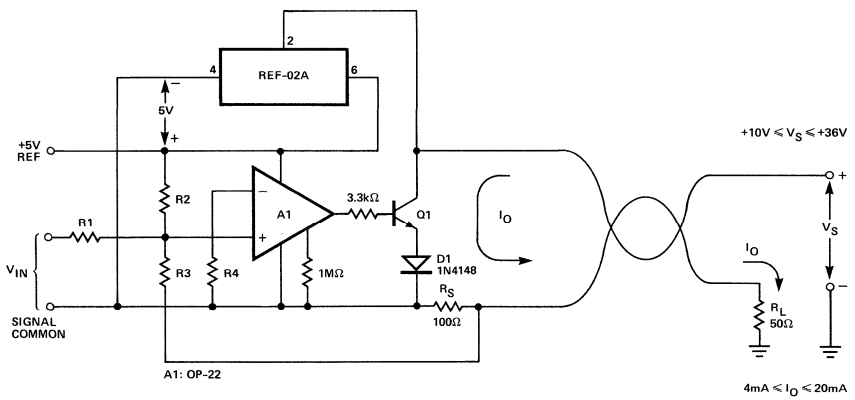

ANALOG DEVICES

AB-108
APPLICATION BRIEF

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

A Micropower Single-Supply Precision Rectifier

by James Wong

This precision full-wave rectifier circuit accepts AC inputs of up to ±3V, yet operates from a single +5V supply voltage. Quiescent supply drain is only 320μA. Rectifier gain is unity with the gain accuracy almost entirely dependent on the match between resistors $2R_1$ and $2R_1$. Frequency range is approximately DC to 2kHz. The single supply operation at very low, quiescent current drain makes this circuit particularly useful for battery-powered equipment.

For positive input voltage ($V_{IN} > 0$), A1 will drive Q1 and D2 to make the output voltage V_O equal to the input voltage. Output swing at V_O is approximately three diode drops below the supply voltage, thus the peak output voltage is near +3V. Amplifier A2 output goes to negative saturation, which is approximately +0.8V, and Q2 will therefore be back-biased and OFF.

For negative input voltage ($V_{IN} < 0$), A1 output goes into negative saturation and Q1 is thereby gated OFF. Amplifier A2 will serve as a unity-gain inverter. Since V_O will be equal to V_{IN} in magnitude, but opposite in polarity, V_O will be equal to the absolute value of V_{IN}.

Quiescent current drain is determined by the set current I_{SET}. With a 5V supply the set current will be $3.7V/R_{SET}$. Slew rate and bandwidth vary directly with the set current. Amplifier A1 essentially operates with unity-gain feedback, while A2 operates with a feedback gain of 0.5. The closed-loop gain-bandwidth is therefore made equal, and the frequency response symmetrical, by making the set current of A2 twice that of A1. Amplifier A2 has a set current of 3.7V/200k, which is 18.5μA, and amplifier A1 has a set current of 3.7V/390k, which is 9.5μA. These set currents will result in quiescent currents of 100μA for amplifier A1 and 220μA for amplifier A2.

The OP-22 input stage is a PNP Darlington, thus a negative input voltage can forward bias the collector-base junction of the input transistor. This potential problem is prevented by adding resistor R1 and diode D1 at the A1 input to limit the negative input voltage.

This simple circuit provides precise, unity-gain rectification of AC signals of up to ±3V in the frequency range of DC to 2kHz. It operates from a single +5V supply voltage with quiescent current drain of only 320μA.

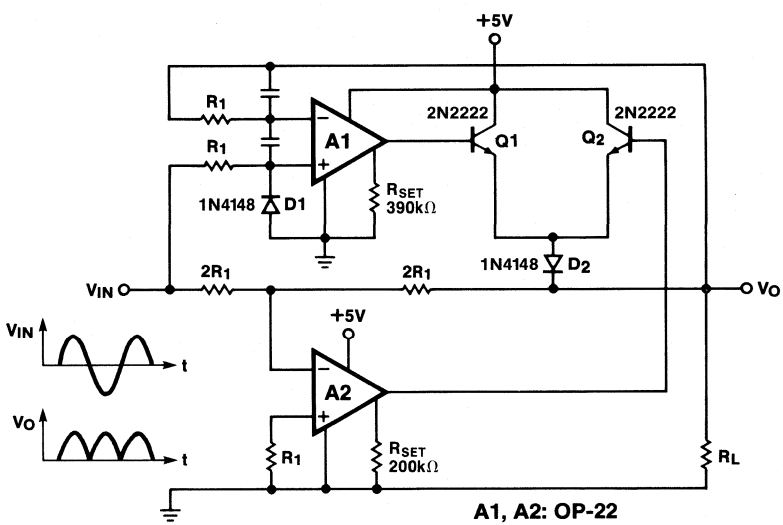

A1, A2: OP-22

ANALOG DEVICES

AB-109
APPLICATION BRIEF

INE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

High Speed Precision Rectifier

by James Wong

he low offsets and excellent load driving capability of the)P-27 are key advantages in this precision rectifier circuit. he summing impedances can be as low as 1KΩ which helps ɔ reduce the effects of stray capacitance.

or positive inputs, D2 conducts and D1 is biased OFF. Ampli-ers A1 and A2 act as a follower with output-to-input feedback nd the R1 resistors are not critical. For negative inputs, D1 onducts and D2 is biased OFF. A1 acts as a follower and A2 erves as a precision inverter. In this mode, matching of the wo R_1 resistors is critical to gain accuracy.

Typical component values are 30pF for C_1 and 2kΩ for R_3. The drop across D1 must be less than the drop across the FET diode D2. A 1N914 for D1 and a 2N4393 for the JFET were used successfully.

The circuit provides full-wave rectification for inputs of up to ±10V and up to 20kHz in frequency. To assure frequency stability, be sure to decouple the power supply inputs and minimize any capacitive loading. An OP-227, which is two OP-27 amplifiers in a single package, can be used to improve packaging density.

13

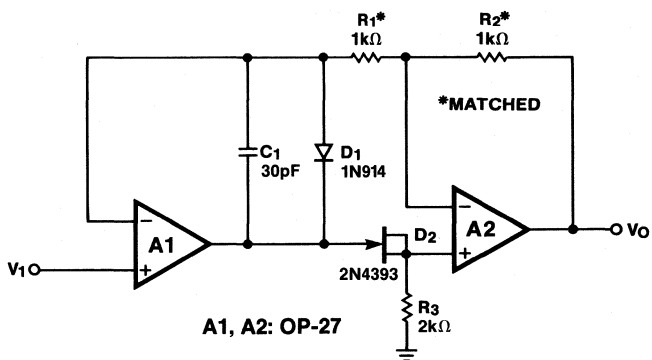

ANALOG DEVICES

AB-111
APPLICATION BRIEF

NE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Single-Supply Wien Bridge Oscillator

by James Wong

ien bridge oscillators have the advantage of requiring only ıe op amp, and this advantage is particularly important for ıttery-operated applications. This oscillator circuit operates ɔm a single 9V battery.

ıe conditions for Wien bridge oscillation are

$- R_1 R_2 C_1 C_2 \omega_O^2 = 0$ and

$$\frac{R_2 C_1}{{}_1 C_1 + R_2 C_2 + R_2 C_1} = \beta$$

here β is the ratio of output voltage feedback to the inverting put. If $R_1 = R_2$ and $C_1 = C_2$, then ω_O is 1/RC and β is 1/3.

his oscillator should be set to just diverge in amplitude. iodes are used to obtain a nonlinear feedback characteristic hich will limit the divergence without causing too much istortion. The condition for oscillation is

$$\frac{R_3}{{}_3 + 2\,(R_5 + R_4')} = \frac{1}{3}, \quad R_4' = \text{Parallel combination of } R_4 \text{ and diodes}$$

As a design example, consider

$C_1 = C_2 = 0.01\mu F$	$R_4 = 10k\Omega$
$R_1 = 15.8k\Omega$	$R_5 = 40k\Omega$ nominally
$2R_2 = 31.8$	Diodes = 1N914 or 1N4148
$R_3 = 50k\Omega$	$R_S = 1M\Omega$

Using these component values, f_O will be 1004Hz. Resistor R_5 must be adjusted for best amplitude stability. If R_5 is too low, the oscillation might converge; if too large, then the oscillation will diverge until the output clips. An oscillation output of 6V peak-to-peak when operating from a 9V battery is recommended. Resistor R_5 needs to be a nominal 40kΩ with a ±2.5kΩ adjustment range.

The OP-22 is operated with a 1MΩ set resistor for a set current of 7.8μA which corresponds to a supply current of approximately 100μA. Gain-bandwidth product and slew-rate vary directly with the set current, so R_S should be optimized for the specific oscillation frequency. Supply drain can be reduced for lower frequencies. The OP-22 works well for frequencies in the range of 100Hz to 1kHz; the OP-27 is recommended for higher frequencies.

13

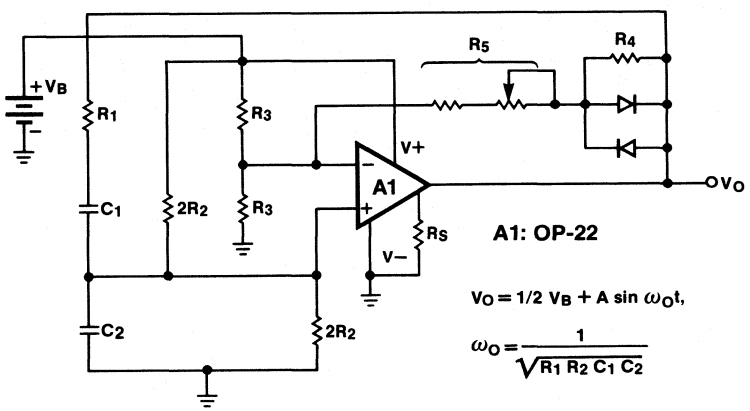

ANALOG DEVICES

AB-112
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Single Resistor Controls Wien Bridge Oscillator Frequency

By James Wong

Frequency control can be added to the conventional Wien bridge circuit by adding an op amp inverter (A1 in the diagram). The low power OP-221 dual works well in this circuit. Center frequency ω_O is $1/R_1C_1$ multiplied by a variable term $1/\sqrt{\alpha}$. The inverter gain is $1/\alpha$, where α is nominally unity.

The center frequency is given by

$$\omega_O = \frac{1}{R_1 C_1 \sqrt{\alpha}}$$

This circuit adds tuning capability to the Wien bridge oscillator circuit.

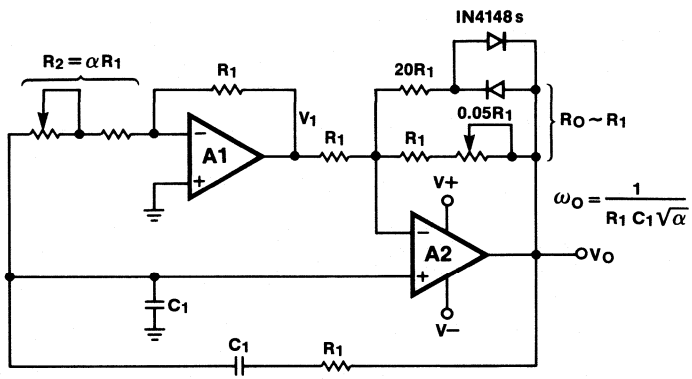

A1, A2: 1/2 OP-221

13

ANALOG DEVICES

AB-113
APPLICATION BRIEF

NE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Precision Ramp Generator

by James Wong

ecision ramps with well-controlled repetition rate and mplitude are generated by this circuit. Repetition rate is ntrolled by a DC input voltage (V_1). This circuit can also be ed as a simple voltage-to-frequency converter over a iited input range.

st after resetting, the A1 op amp generates a negative ramp th a slope proportional to I_1, which is $(V_1 - V_{D2})/R_1$. The ope is $-C_1 dV_R/dt$ and the A2 output is sitting at the positive mit. When the A1 output reaches −10V, the output of A2 flips the negative limit. This transition is given regenerative tion through capacitor C_2. The negative pulse from A2 scharges C_1 through diode D3 and drives it positive until ode D1 conducts. Since D2 sets the A1 inputs to −0.6V and has an equal drop when it conducts, the integrator will be set to zero volts. After the integrator reaches zero volts and has discharged, amplifier A2 flips back to positive saturaon and D3 is again back-biased. A key feature of this circuit the amplitude stability; the REF-01 output of +10V is very able and the reset zero is temperature compensated by the atching of D1 and D2. Thus the ramp amplitude of 10V is very accurate and stable over a wide range of operating conditions.

Exact circuit values and op amp choices depend on the desired operating range. For a range of 10Hz to 1kHz, the OP-215 can be used with the following values:

$R_1 = 1M\Omega$
$C_1 = 1200pF$
$R_2 = 10k\Omega$
$R_3 = 2.0k\Omega$
$C_2 = 200pF$

With these component values and using the OP-215 dual, we will have a reset time interval of approximately 5μs. The minimum ramp interval, assuming a maximum input voltage of 11.4V and 0.6V diode drop, is 1200pF × 10V/12μA = 1msec which corresponds to a 1kHz repetition rate. The ramp amplitude of zero to −10V is very accurate and stable over a range of 10Hz to 1kHz. The output of A2 is a 5μs pulse of approximately ±13V amplitude.

13

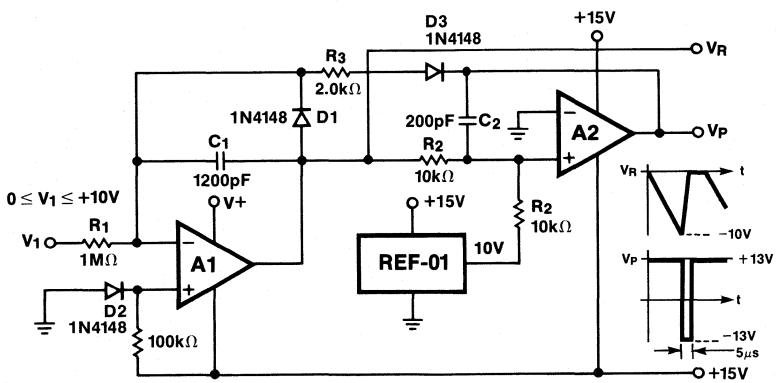

A1, A2: 1/2 OP-215 DUAL

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Precision Current Regulator

by James Wong

The wide input and output voltage range of the OP-22 programmable micropower op amp is very useful for driving power MOSFET devices. This simple circuit provides a very stable and accurate current source that operates over a range of 15V to 40V and up to 1A of current. The REF-02 and OP-22 both operate from the +15V to +40V supply voltage. The REF-02 output is a stable +5V that is divided down to drive the input voltage (V_R) to the current amplifier. The op amp drives the power MOSFET to make $I_O R_S = V_R$. A sense resistor of 5Ω provides a range of zero to 1A for an input range of zero to -5V. The high open-loop gain assures excellent current regulation.

While the nominal V_{GS} for the IRF533 is 2V to 4V and the OP-22 output will go to within 1.5V of the negative rail, there is a possible difficulty in swinging low enough at the OP-22 output to assure cut-off of the IRF533 over the full temperature range. Selecting the power MOSFET for minimum V_{GS} of 2V over the temperature range will provide a good safety margin. In any case, the IRF533 will require an adequate heat sink if operated over the full zero to 1A range.

The 200Ω resistor should be located near the gate; its purpose is to prevent any spurious parasitic oscillations in the power FET. The 100kΩ resistor is non-critical and just serves as a resistive load in parallel with the gate-to-source capacitance. If desired, a signal source can be used in place of the REF-02 to dynamically control I_O. Slew-rate is approximately 0.07V/μs at 20V supply voltage. Also, other power FET devices can be used in this circuit and other current ranges accommodated by simply changing the sense resistor R_S as needed.

13

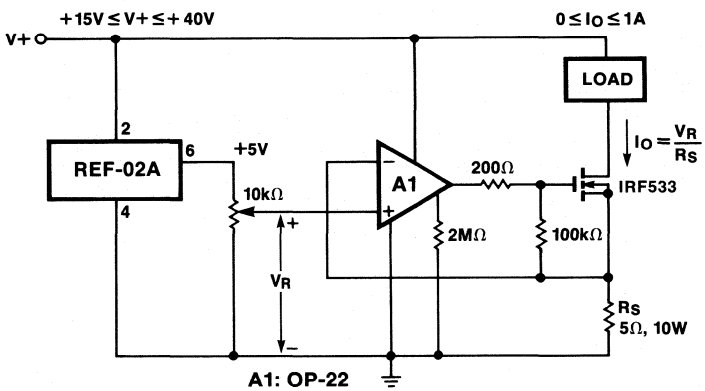

ANALOG DEVICES

AB-115
APPLICATION BRIEF

NE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Micropower 1.23 V Bandgap Reference

by James Wong

ıst bandgap references require at least 1mA of quiescent erating current.This circuit takes advantage of an ADI pro- ımmable micropower op amp, the OP-22, which provides ıh gain even at low quiescent currents. It has a further vantage of wide supply voltage range. Total quiescent cur- ıt for this 1.23V reference is only 20μA and the input vol- ıe range is +3V to +30V.

bandgap reference is generated by summing a V_{BE} drop ıt has a negative tempco with a fraction of a ΔV_{BE} that has a sitive tempco. The summation that gives zero tempco is ə energy bandgap of silicon (1.205V) plus kT/q (~25mV at °C) for a total of 1.230V. This varies slightly with process- ı, but the zero tempco point will be very close to 1.23V.

this circuit, a ΔV_{BE} is generated by the imbalance in collec- ' resistance.

$$V_{BE} = \frac{kT}{q} \ln \frac{I_{C1}}{I_{C2}}$$

$$= 25.7mV \times \ln \frac{R_2}{R_1}$$

$$= 25.7mV \times \ln \frac{2.2M}{360k\Omega}$$

$$= 46.5mV \text{ at } 25°C$$

The ΔV_{BE} is impressed across R_3, a 68kΩ resistor, and the resulting current of 0.684μA is summed with I_{C1}. This total current I_T is

$$I_T = I_{C1} + \frac{1}{R_3} \Delta V_{BE}$$

$$= \frac{\Delta V_{BE}}{R_3} \times \frac{R_1 + R_2}{R_1}$$

$$= 4.863\mu A$$

This current across a nominal 150kΩ for R_4 causes a voltage of 729mV. This 729mV adds to a V_{BE} of 500mV to generate the desired total of 1.230V at the base of Q1. Tempco of this 1.230V bandgap output will be low, typically 20ppm/°C over a 0°C to 70°C range.

The OP-22 op amp, followed by a 2N2907 PNP transistor, makes the collector voltages of Q1 and Q2 equal and regulates the output voltage. Line regulation is approximately 0.01%/V and load regulation is 0.001%/mA. Quiescent current for the OP-22 is set by the Q3 constant-current source. Set current for the OP-22 is 0.73V/1M which corresponds to a quiescent supply current of approximately 7μA to 12μA over the supply range of +3V to +30V.

This bandgap reference provides excellent stability and regulation at very low quiescent operating current, approximately 20μA at a 10V supply level.

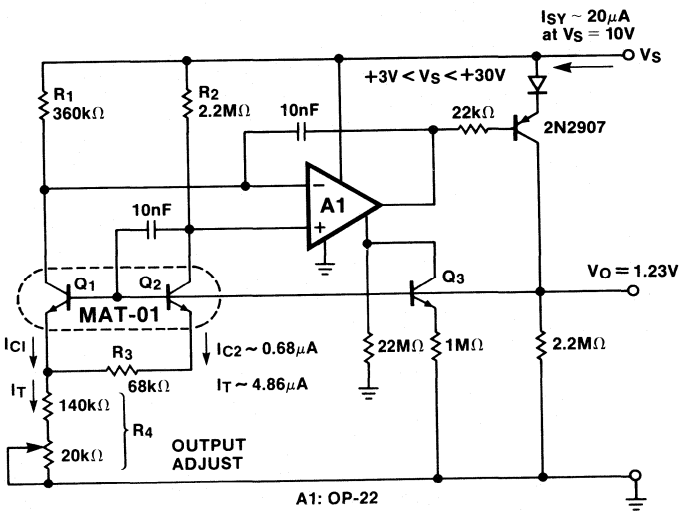

ANALOG DEVICES

AN-254
APPLICATION NOTE

NE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Add Programmable Gain, Attenuation

by James Wong

By adding two resistors to the output-amp feedback loop of a current-output digital-to-analog converter (DAC), both gain control and attenuation control can be achieved *(Fig. 1)*. This digitally programmable amplifier produces gain and attenuation in the range of 1/64 to 64. The circuit gets its range from a 12-bit CMOS DAC.

The design works because the transfer function from the DAC's input to its output is purely voltage attenuation. Connecting R_1 and R_2 in a "T" configuration inside the output amp's feedback loop produces a voltage gain from the resistor junction to the output. If R_1 is much less than R_{FB} (11 kΩ in this example), the gain produced nearly equals $1 + (R_2/R_1)$, or 64. The result is a programmable gain amp with a transfer function of

$$A_V = -(D/4096)(64),$$

where D represents the DAC's binary-weighted digital code. Of course, the added gain of the T-network increases the circuit's noise gain. Therefore, it's important to choose first a low-noise amplifier.

By using a low-noise, high-frequency op amp, such as the OP-61, the circuit will have a wide bandwidth performance even at high gain settings. The circuit's frequency response can be plotted at different gain settings *(Fig. 2)*. At high gains, the amp has a 1-MHz bandwidth.□

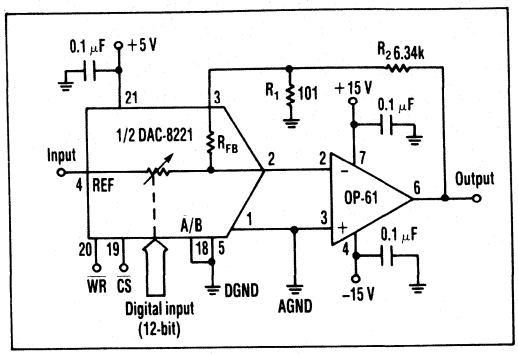

1. BY adding R_1 and R_2 in the feedback loop around a DAC, the circuit functions as a digitally-programmable amp. The gain or attenuation is variable over the range of 1/64 to 64. The resistors are connected in a T-configuration.

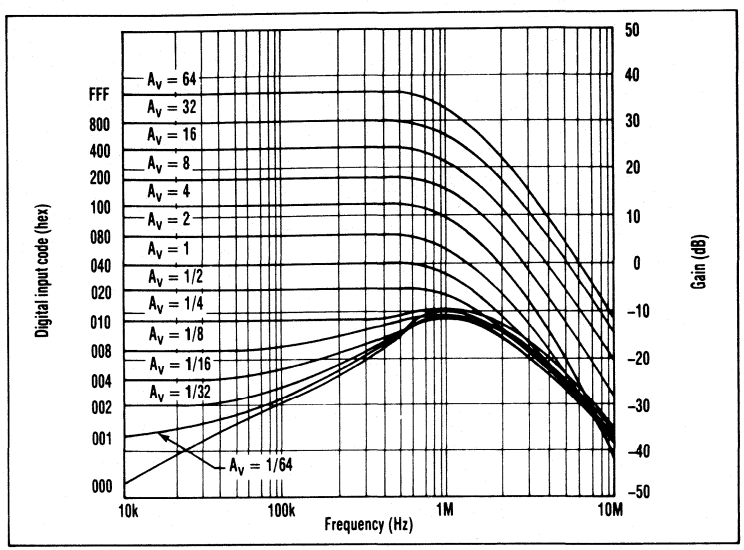

2. GAIN IS PLOTTED versus frequency for various digital inputs of the DAC. The amplifier has a 1-MHz bandwidth at high gains, but it drops for gains below 1/4.

Reprinted from Electronic Design – November 8, 1990

ANALOG DEVICES

AN-255
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Voltage-Controlled Amp Covers 55 dB Range

by Bob Clarke

INTRODUCTION

By using a dual multiplier chip and a dual high-speed op amp (Figure 1), you can build a 2-chip voltage-controlled amplifier with a dynamic range of 55 dB, a 3-dB bandwidth of 8 MHz, and exponential control. The amplifier's output ranges from 5 mV p-p at V_X = 0 V to 3 V p-p at V_X = 3 V for a 100 Ω load. The circuit's gain is unity at V_X= 2 V. You can also use Figure 1 to drive a reverse-terminated 50 Ω cable to 1.5 V p-p. Or you can use each multiplier-op amp combination separately to amplify two signals with control by a common voltage.

Figure 1's circuit connects the AD539's two voltage-in current-out multipliers in series. Each of the op amps acts as a current-to-voltage converter. V_X, a single 0 V to 3 V dc input, controls both multipliers. Because both multipliers are in series, the overall transfer function is

$$\frac{V_{OUT}}{V_{IN}} = \frac{V_X^2}{4\,V^2}$$

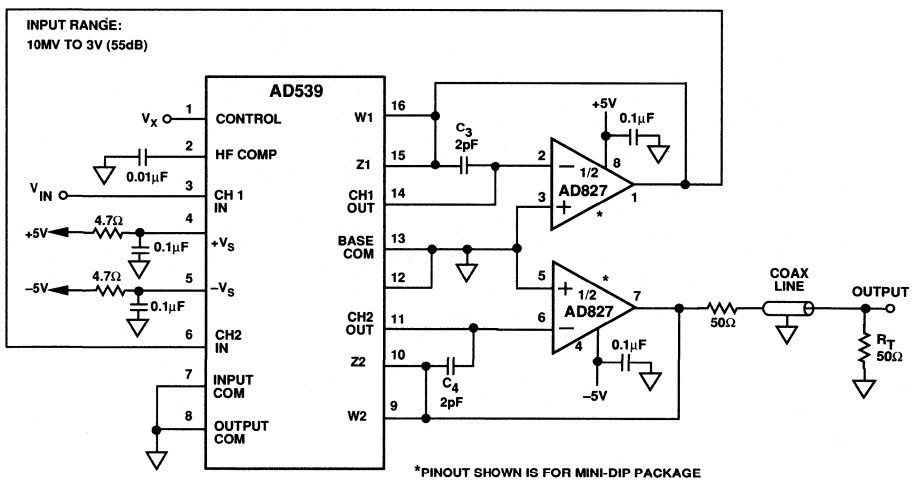

Figure 1. A Wide Range Voltage-Controlled Amplifier Circuit

Reprinted from EDN — July 19, 1990

The plot of V_X vs. the gain of this voltage-controlled amplifier on log-log axes is a straight line, which demonstrates the exponential gain response.

The square term in the denominator of the transfer function comes from connecting each of the multiplier's W and Z outputs. The W and Z pins of the AD539 are each connected to 6 kΩ resistors. Connecting the two pins sets the two resistors in parallel and thus halves the gain. The feedback resistor in each current-to-voltage converter halves from 6 kΩ to 3 kΩ and thereby reduces the amplifer's overall gain by a factor of four.

As an option to Figure 1's circuit, you can disconnect th Z outputs and use only the W outputs, thereby fixing th gain resistor at 6 kΩ. If so, the overall transfer function

$$\frac{V_{OUT}}{V_{IN}} = \frac{V_X^2}{1\ V^2}$$

The maximum gain is 9 when V_X = 3 V. You can trad decreased bandwidth for increased gain by adding a external scaling resistor, R_S, in series with the on-chi feedback resistors. In this case, the transfer function fc the dual-multiplier circuit becomes

$$\frac{V_{OUT}}{V_{IN}} = \frac{V_X^2}{1\ V^2}\left(\frac{R_S}{5R_S + 6.25}\right)^2$$

where the units of R_S is kΩs.

ANALOG DEVICES

AN-359
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Settling Time of Operational Amplifiers

Robert I. Demrow

FOREWORD

It is possible to measure voltages and currents to within small (even fractional) parts per million, traceable all the way to the Bureau of Standards. Typically, at the limits of precision, repeated measurements must be made over a period of time, after which an experienced practitioner can state with a great deal of certainty that a voltage or a current had, during that interval, an average value expressible to six or more significant figures.

That this can be accomplished is interesting, impressive, and of basic importance to all engineers, but it is of little direct relevance to a far more pervasive type of measurement. This is the measurement that calls for a more modest degree of accuracy – four significant figures at best (today), but the measurement must be completed, often within a microsecond or so, then converted, stored, and utilized by a computer. The time taken to complete the measurement to the desired degree of accuracy, be it ±1%, ±0.1%, ±0.01%, or ±½LSB, in terms of its final value (with the often-added provision that said final value is within a similarly small fraction of full scale of its nominal expected value) is usually called Settling Time. This interval may be quite short: at present ±0.01% within 0.5 to 1.0 microseconds is fashionable; next year engineers will consider an order of magnitude improvement in performance feasible (and many will attain it). Significant improvements will be reported in these pages.

The purpose of the present discussion is to provide an interpretation of Settling Time, and to set before the engineer the issues involved in terms of circuit design, hardware, specifications, applications, and measurements. We shall try, to the degree possible, to avoid the use of intimidating mathematical formulae, and to reveal practical circuits and mathematical shortcuts.

INTRODUCTION

Much has been written on the subject of small-signal frequency response characteristics of amplifiers; and most users of operational amplifiers are able to make use of the manufacturers' specifications to calculate closed-loop bandwidth, analyze stability, determine phase errors, etc.[1] Most manufacturers also provide data sufficient to characterize the full power response and slew rate of the amplifier.

The important difference of Settling Time investigations is that, although frequency response may be a tool, the results must be either implicitly or (preferably) explicitly measurable in the Time Domain. They are in general terribly nonlinear and subject to (so it would seem) every stray nastiness that Nature has evolved, because of the combined stress on both speed and accuracy.

Thus the recently-increased use of operational amplifiers in handling data in systems calling for high-speed switching rates, especially in connection with digital computers, has led to a requirement for *time-response* characterization of general- and special-purpose operational amplifiers. A typical problem is the application of rapidly-changing signals (e.g., step functions) to a buffer amplifier, which must faithfully reproduce the input to a high degree of accuracy within a period of the order of a microsecond. This requires that the amplifier be designed and optimized for *Settling Time*. Applications requiring fast settling time to high accuracies are typified by Sample-Hold circuits, Multiplexers, and amplifiers used with A/D and D/A converters. Settling time is important in these applications because it is the principal factor which determines the maximum data or information transfer rate for a given accuracy. A tradeoff is generally possible between data rate and accuracy: viz., higher data rates are possible at the sacrifice of accuracy. It should also be noted that the buffer amplifier, as but one step in a process, is now one of the major limitations on system speed.

13

DEFINITION

SETTLING TIME *is the time elapsed from the application of an ideal instantaneous step input to the time at which the closed-loop amplifier output has entered and remained within a specified error band, usually symmetrical about the final value. Settling time includes a very brief propagation delay, plus the time required for the output to slew to the vicinity of*

Reprinted from Analog Dialogue 4-1 1970

[1]Footnote references are to numbered items in the bibliography.

the final value, recover from the overload condition associated with slewing, and finally settle to within the specified error. Figure 1 illustrates this definition of Settling Time.

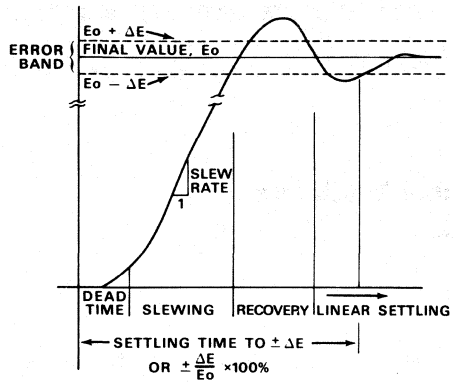

(a) Components of settling time (nonlinear scale)

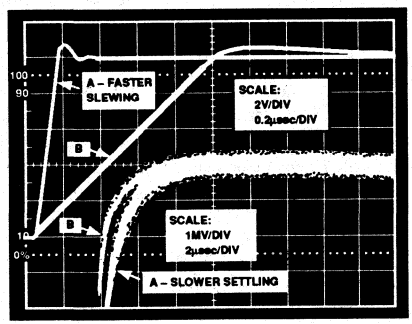

(b) Annotated wave forms illustrating comparative settling time of two typical op amps.

Figure 1. Pictorial Definition of Settling Time

CAUTION: The above definition has been used widely in the data acquisition field for at least the past 10 years. Nevertheless, some manufacturers new to this field (but not to the tricks of "specmanship") define "settling time" as the time required to settle within the linear small-signal region only. When considering the products and applications literature of these manufacturers, it is important to realize that their term usually corresponding to Settling Time is "Acquisition Time." (Until adopted for this usage, Acquisition Time had a specialized meaning of its own: minimum time required for a Sample-Hold switch to be closed to allow re-opening without loss of information, usually without the necessity that the amplifier finish settling.) We do not know at present what standards will eventually emerge, but we shall always identify small-signal settling time by attaching the "small signal" label wherever it would be misleading to do otherwise.

Settling time may also be defined in terms of the recovery time of the amplifier from an instantaneous error caused by a step or an impulse change in load. Because Settling Time is determined by a combination of amplifier characteristics, nonlinear as well as linear, and because it is a closed-loop parameter, it cannot readily be predicted from such open-loop specifications as slew rate, small-signal bandwidth, etc.

Two frequently-occurring applications requiring fast settling time for maximum transfer rate are indicated in Figures 2 and 3. Figure 2 depicts the typical Multiplexer problem (only two channels shown, for simplicity). One line is at 0V, the other is at -10V. When switching from one channel to the other, the amplifier's input may be subjected to a 10-volt step and must recover to the millivolt level in high-accuracy systems. Figure 3 indicates a typical waveform from a high-speed D/A Converter (DAC) showing large transient spikes ("glitches") caused by time skew between turn-on and turn-off of the DAC's switches. The settling time of the amplifier used (if the DAC uses an amplifier) limits the maximum bit rate for an A/D Converter (ADC) using the DAC.

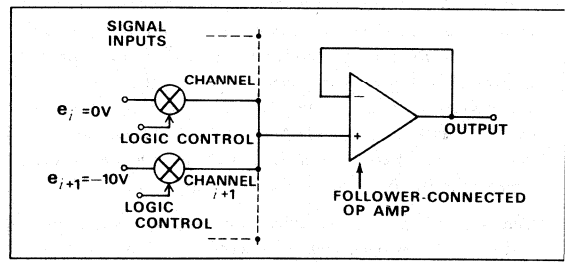

Figure 2. Simplified diagram of muliplexer circuit

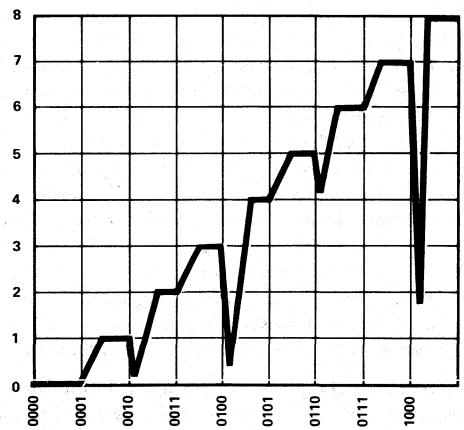

Figure 3. Output of counter-driven Fast D/A Converter showing "glitches"

AMPLIFIER DESIGN FACTORS AFFECTING SETTLING TIME

The design factors to be discussed here are those that the amplifier designer and the engineer who specifies the amplifier must consider. Other design factors concerning the circuit the amplifier is tested or used in will be discussed in a later section.

INEAR OPERATING RANGE. Although the essence of the roblem of settling time is nonlinear, the factors to be discussed ı this section pertain to the all-important linear-settling "tail", /hich applies to both small signals and large (though not lways identically). "Linear" means that the amplifier and all ssociated elements are operating in the linear range, i.e., the ıarametric relationship is independent of voltage or current ɛvel, or its previous history.

Noise. The amplifier's output can never settle within a given ıand of error if its output noise, however generated, is comıarable to the magnitude of the band defined for settling. In ıddition to inherent amplifier circuit noise, the designer must ılso consider the effects of interference noise, whether coupled rom the external environment, the power supply, the input ignal, or the logic signals in the associated circuitry. For a liscussion of noise in operational amplifier circuits, see *Analog Dialogue*, Vol. 3, No. 1.

DC Gain. To ensure the accuracy of the final value, the gain ɔf an amplifier which must settle to within ±0.01% should be at east 10,000 for unity-gain followers, 20,000 for unity gain in/erters, and more for higher-closed-loop-gain amplifiers. If the ımplifier's open-loop input-output characteristic is reasonably linear, the error caused by slightly-lower gains may be compensated for under a given set of conditions by trimming the feedback ratio. For follower applications, the CMRR should be commensurate with the desired gain accuracy.

Drift and Offset. For high-precision applications, offset must be low or adjustable, and drift over the temperature range should be within the error corresponding to the desired accuracy. For example, ±0.01% error in ±10V circuitry requires less than ±1 millivolt (±½mV for inverter applications). Bias current drift may not be disregarded: 1 millivolt in 10kΩ requires that bias current change less than 100 nanoamperes.

Dynamic Stability. Operational amplifiers designed for "optimum" response at high gains at low frequency often have transfer functions that provide only marginal stability when the loop is closed more tightly for near-unity gain and wide bandwidth. The amplifier used for fast settling to high accuracy should have a closed-loop response that is (at least theoretically) not much worse than critically-damped, as any oscillation or ringing may prolong settling time. Furthermore, in practical circuits, which have stray capacitance, the added lags caused by the external loop elements will cause an amplifier having insufficient phase margin to ring. For this reason, designers of fast-settling operational amplifiers strive to have the open-loop frequency characteristic be strongly dominated by a single time constant. This is stated in many ways, all having the same meaning: constant 90° phase shift, -6dB/octave (or -20dB/decade) rolloff, unit lag, exponential response, etc.

Amplifiers characterized by this form of response may be considered to be integrators with limited DC gain, characterized by the following equation:*

$$A_{OL} = \frac{1}{\tau_x p}\left[\frac{A_0\tau_x p}{1 + A_0\tau_x p}\right] \cong \frac{1}{\tau_x p} \cong -j\frac{f_x}{f}$$

*Note: p is the Heaviside differentiation operator, corresponding roughly to s or jω, A_{OL} is open-loop gain, A_0 is DC gain, f_x is unity gain bandwidth, $\tau_x = 1/2\pi f_x$

For many non-state-of-the-art applications, the single time constant response has the further advantage of being calculable virtually by inspection. Small-signal settling time of amplifiers with more complex transfer functions does not lend itself to ready calculation. Figure 4 shows the open- and closed-loop response of an amplifier having this popular characteristic, and provides a table indicating the relationship between small-signal settling time and amplifier unity-gain bandwidth (f_x) or characteristic time (τ_x).

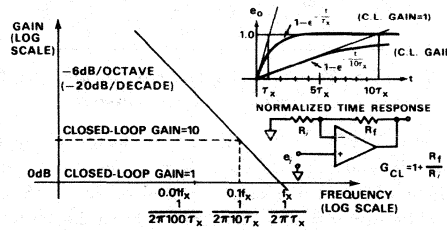

(a) Bode (amplitude vs. frequency) Plot, Feedback Circuit, Time Response

Percent error	$T_s = -2.303\tau_x \log_{10}\eta$ for τ_x (μs) or f_x(MHz)			
	Gain = 1		Gain = 10	
100η	T_s vs. τ_x (μs)	T_s vs. f_x	T_s vs. τ_x	T_s vs. f_x
10 %	$2.303\tau_x$	$0.367/f_x$	$23.03\tau_x$	$3.67/f_x$
1.0 %	$4.606\tau_x$	$0.734/f_x$	$46.06\tau_x$	$7.34/f_x$
0.1 %	$6.909\tau_x$	$1.101/f_x$	$69.09\tau_x$	$11.01/f_x$
0.01%	$9.212\tau_x$	$1.468/f_x$	$92.12\tau_x$	$14.68/f_x$

(b) Settling Time vs. Fractional Error η for First-order Step Response as a Function of Open-Loop Characteristic Time (τ_x) or Unity-gain Bandwidth (f_x)

Figure 4. Closed-loop Response of Amplifier having -6dB/Octave Open-Loop Amplitude Response

Well-Behaved Dynamic Response. Multistage amplifiers require the matching of poles and zeros in the linear transfer function to achieve a 6dB/octave characteristic.[2,3] That is, m = 1 in the general expression:

$$A_{OL} = \frac{1}{\tau_x p}\left[\frac{A_0\tau_x p}{1 + mA_0\tau_x p}\right]\left[\frac{1 + mA_1\tau_x p}{1 + A_1\tau_x p}\right]$$

If the mismatch term, m, is substantially different from 1.0, and if A_1 is substantially less than the open-loop gain required for the desired degree of accuracy, the output of the amplifier will settle rapidly to within a small fraction (about $1/A_0$) of the final value, undershooting if m>1, overshooting if m<1, then settle exponentially to the final value with a surprisingly long time constant. (Fig. 5) The details of this consideration are described, and illustrated by scope pictures of the response of a low-speed dynamic model in the Appendix on pages 10 and 11.

Propagation Delay. At frequencies substantially beyond f_x, the amplifier's open-loop response starts to roll off more steeply and with greater phase shifts. The time-domain analogy to this is a brief interval of "dead time" of the order of nanoseconds, small compared to τ_x if the circuit is to be stable closed-loop. Because τ_x itself is not usually a limiting factor on settling time, propagation delay is even less so, in linear operation.

13

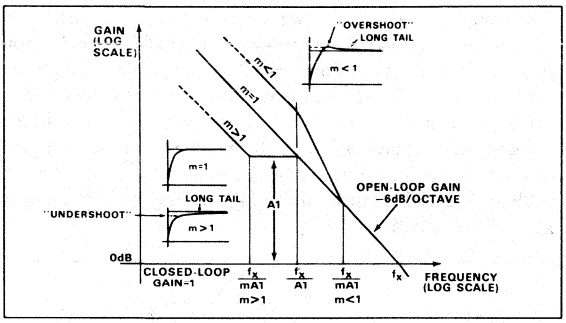

Figure 5. Pole-Zero Matching of open-loop gain and its effect on Closed-loop Linear Settling Time (see Appendix)

NONLINEAR EFFECTS. Amplifiers are available having f_x as high as 100 MHz and more. Using the formula given in Figure 4, this would imply a settling time to 0.01% of about 15 nanoseconds, if the response of such amplifier has a single unit-lag. Yet the specifications are more like 0.5 to 1.0 microseconds and more. Why? Here are some of the considerations:

Slew Rate. Within the amplifier, there are inherent semiconductor and circuit capacitances, as well as those added for stabilization; at the output, there is load capacitance. The rate of change of voltage at each point in the circuit is limited by the available current to charge the capacitance at that point ($dv/dt_{max}=I_{max}/C$). In the above example, if the amplifier were operating linearly over its whole output range, the initial rate of rise in response to a step would have to be 10V/1.6ns = 6,280 volts per microsecond! In order to drive a 10pF load without saturating, at least 62.8mA must be available.

ΔE	f_x (MHz)				τ_x (ns)				
	0.1	1.0	10	100	3.2	10	100	1000	10μs
10V	6.3	63	630	6,300	3,100	1000	100	10	1.0
5V	3.1	31	310	3,100	1,550	500	50	5	0.5
2V	1.3	13	130	1,300	620	200	20	2	0.2
1V	0.6	6	60	600	310	100	10	1	0.1

TABLE 1 — First-Order Response: Required Initial Rate of Change (Volts/μs) as a Function of Step Size, f_x, and τ_x.

$$\left[\left(\frac{de}{dt}\right)_{max} = \frac{\Delta E}{\tau_x} = \Delta E\, 2\pi f_x\right]$$

C ↓ / $\frac{de}{dt}$ →	1	10	20	50	100	200	500	1000
10pF	10μA	0.1	0.2	0.5	1.0	2.0	5.0	10.
20pF	20μA	0.2	0.4	1.0	2.0	4.0	10.	20.
50pF	50μA	0.5	1.0	2.5	5.0	10.	25.	50.
100pF	0.1	1.0	2.0	5.0	10.	20.	50.	100.
200pF	0.2	2.	4.0	10.	20.	40.	100.	200.
500pF	0.5	5.	10.	25.	50.	100.	250.	500.

TABLE II — Current (mA) to Maintain a Given Slew Rate (V/μs) As a Function of Capacitance. $\left[i = C\frac{de}{dt}\right]$

Tables I and II show the relationships between step size, slew rate required for exponential response, capacitor size, and required driving current, for a somewhat less exotic selection of parametric values. Table I gives examples of the slew rat requirement for linear behavior (i.e., exponential time response as a function of step size and either f_x or τ_x, and Table I indicates the current required to drive a given capacitive load a the appropriate value of slew rate. In practice, such slew rate (in relation to small-signal frequency response) cannot b achieved: the excessive currents required would degrade ac curacy and drift specifications; also the operational amplifie would be required to respond linearly to 100% error at its inpu (i.e., full-scale differential input). Typically, then, the amplifie slews, at a maximum rate 5 to 100 times less than that implie by its small signal bandwidth, up to the vicinity of full scale recovers, and settles exponentially (either aperiodically or witl oscillations).

Recovery. During nonlinear slewing, saturation imposes charg changes away from normal operating values on the circuit capa citances (including minority carrier storage in semiconductors) These must be discharged back to equilibrium values before th amplifier can operate normally. Thus, there is a period o recovery which is comparable to the period of slewing, but i may be substantially greater if many internal stages are involved Fast slew rate, therefore, is not by itself a good indicator of a fast-settling amplifier. Some amplifiers with extremely-larg slew rates have excessive recovery time and greater overal settling time than other amplifiers having more modest slev rates.

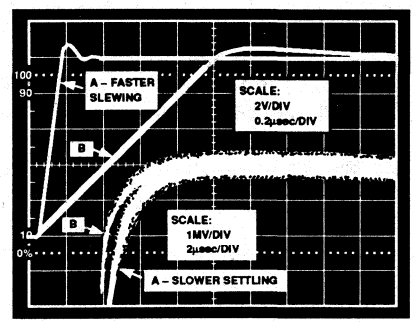

Figure 6. Comparison of Two Amplifiers Having Similar Settling Times, But Differing Slew Rates.

OTHER EFFECTS. Dielectric absorption (hysteresis), both internal to the amplifier and inherent in external bypass capacitors, may be an important factor in circuits requiring settling to 0.01%. Most capacitors used in electronic circuits have non-negligible dielectric hysteresis when used in precision applications; the popular ceramics, for example, may have typically 5%. When a capacitor is subjected to a sudden stress, it requires a period of "soaking" to return to internal charge equilibrium.[4] Thus the dC/dt term cannot be neglected in the definition of current, i = dQ/dt = d(CV)/dt = C dV/dt + V dC/dt. Dielectric absorption is responsible for the long "tails" (typically 10 to 100μs) often greatly extending the final-settling interval for otherwise fast amplifiers. With proper attention, an amplifier's settling time to 0.01% need not greatly be affected by dielectric absorption, and minimization of its effects in settling to 0.1% is even less difficult to achieve.

ıermal Transients ("self-heating" effects). During slewing, the ırmally equal distribution of dissipations in (particularly) the put stage may become grossly unbalanced. Amplifiers de- ŋned to be fast usually run at "rich" values of current (to ıable rapid handling of voltage changes across circuit capaci- nces and keep impedance levels down), but because a differ- ıtial temperature change of one degree in a bipolar ansistor stage can cause about 2 millivolts of drift (0.02% of)-volt full scale), the time required to recover thermal balance ter large input transients is a factor tending to limit the ıoice of currents in the input stage. In addition, a large change load current may significantly change dissipations in the ıtput stage (and hence temperatures inside the amplifier ıckage), and if these temperature changes are unequally con- ıcted to the two sides of the input stage, transient, as well as eady-state unbalances can result. Input circuit effects are less ırious in inverting amplifier configurations than they are in nity-gain followers, where common-mode swing (equal to ıll-scale input voltage swing) may cause substantial changes in put-stage dissipation. Self-heating effects will generally not fect the settling time to 0.1%, and – in well-designed ampli- ers – will not substantially increase the settling time to 0.01%. owever, where this design factor is neglected, they can cause tails" exceeding milliseconds in duration.

MINIMIZING SETTLING TIME SUGGESTIONS FOR THE CIRCUIT DESIGNER

f the designer of a typical fast-settling operational amplifier ıas done the job properly, taking into account all the factors ve have mentioned, he has made available a device that can be lemonstrated to have a specified accuracy of 0.01% and ettling time less than 1μs. However, an operational amplifier s a building block in a circuit that also has a feedback network, nput connections, power supply connections, output connec- ions, and a number of external components. What has been ɔainstakingly gained in amplifier design can be lost in short ɔrder through careless circuit design. Some of the elements of ŋood design are these:

CONNECTIONS. It is of utmost importance that the power ıupply leads be adequately bypassed directly at the amplifier's ɛrminals* and that especial care be taken in the signal and power ground circuits to avoid inducing or generating ex- traneous voltages in the ground signal paths.[5]

COMPONENTS. Resistors are preferably metal film types, because they have less capacitance and stray inductance than wirewound types, and are now available with excellent ac- ɔuracies and temperature coefficients.

Diodes are hot carrier types for the very fastest-settling applications, but 1N914 types are suitable for more routine uses.

Capacitors in critical locations are polystyrene, teflon, or poly- carbonate, to minimize dielectric absorption.

CIRCUIT. For the fastest settling times, keep leads short, orient components to minimize stray capacitance, keep circuit impedance levels as low as consistent with the output capa- bilities of the amplifier and the signal source, reduce all external load capacitances to the absolute minimum. Don't overlook sockets or printed circuit board mounting as possible sources of dielectric absorption. Avoid pole-zero mismatches in the feedback networks used with the amplifier. Minimize noise pickup.

*Most well-designed discrete op amps have bypass capacitors built in, but the designer should never take it for granted that they will be ade- quate for his application, until he can so demonstrate.

SPECIFYING SETTLING TIME

On page 2 is a statement relating to alternate definitions of settling time employed in some portions of the industry. When it comes to *specifying* settling time, there are as many formats as there are manufacturers. Until the format to be discussed and demonstrated below was developed, no known manufac- turer, including Analog Devices, had presented, as a settling time specification, much more than the stark statement: "settling time as an inverter to 0.01% of full scale is . . . .μs," plus a waveform or two, and a simplified schematic of a test setup.

Important unanswered questions were: "What are the trade- offs?" How much can settling time be reduced by reducing full-scale signal level (and thus the slewing and recovery periods)? Suppose one does not desire 0.01% accuracy, but wishes instead to determine the size of error band for a given step size and settling time? Is settling time symmetrical for positive and negative excursions?

V CURVES (Figure 7)

A graphical form of specification has been developed that helps answer many of the above questions. Because of the shape of the curves when plotted linearly, they have been dubbed "V Curves." For increased information content, however, a semi- log plot has been adopted. Each curve represents the measured settling time (microseconds) vs full-scale step magnitude (volts) for a given final error band. Three error bands were chosen: ±1mV, ±10mV, ±100mV, corresponding to ±0.01%, ±0.1%, ±1% of full scale, respectively. The percentage error is the ratio of error band to output step. For example, 1mV/5V = 0.02%.

Information is given for both positive and negative step polarities; the two standard configurations shown are as unity gain inverter and unity gain follower. It is of course possible to obtain data for curves for other conditions, and we welcome comments by our readers on the adequacy and relevance of this form of specification, along with suggestions for improve- ments to the format.

Figure 7 shows a set of theoretical V Curves for an ideal amplifier having no slew rate difficulties whose τ_x is 100ns (f_x = 1.6MHz).

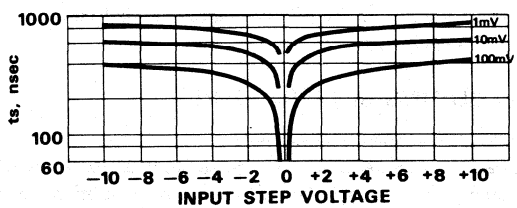

Figure 7. Example of V Curves. Theoretical Amplifier With First-Order Linear Response.

13

LIMITATIONS OF SPECIFICATIONS

CIRCUITS STUDIED. It is usually the case that for most operational amplifiers, the fastest-settling performance will be obtained when the amplifier is operated as either a unity-gain inverter or as a follower. (*Slewing* may be more rapid, however, for some types, at high closed-loop gains). For this reason, data has been taken at unity gain using both inverter and follower configurations. The amplifiers operate with no external resistive or capacitive load, other than the comparator used for the measurement, the leads to it, and the feedback circuit impedance in the inverting configuration. The resistance component has been confined to a pair of 2.5kΩ resistors, to enable amplifiers having 5mA output current to be measured.

EXTRAPOLATION OF SETTLING TIME MEASUREMENTS. Because of the many factors that influence settling time, it may be dangerous to assume that data may be interpolated (or worse, extrapolated), especially for those types having radical changes in shape between adjacent curves, or in different portions of the same curve.

AMPLIFIER STABILIZATION NETWORKS. The methods used for feedback stabilization can have quite drastic effects on the amplifier's settling time because the amplifier's settling characteristics are intimately related to the location of both the open loop and closed loop poles and zeros. Anything the user does that will reduce the amplifier's bandwidth will usually lengthen settling time. For example, multiple input summing resistors decrease loop gain; added summing point capacitance requires feedback capacitance to compensate, resulting in smoother (but often slower) settling; output cable capacitance can cause peaking, and may require a load isolation resistor, which will result in slower overall response; stray feedback capacitance. Even such a simple thing as placing a resistive load on the amplifier may reduce its open-loop gain (at all frequencies, including f_x), and thus its loop gain and settling time.

CAVEAT. Since settling time specifications cannot possibly include (or be extrapolated to) all possible situations, they should be used as a guide for preliminary selection of an amplifier to fit a given application. In all cases, and especially where the application conditions are significantly different from those used in specifying the amplifier, there can be no substitute for actual trial of the amplifier in the proposed circuit to determine whether it will perform as desired. Often, the manufacturer's experience with customers having a wide variety of applications can be helpful, and – especially where op amps are key elements in large important projects – designers are urged to make use of the manufacturer's applications engineers for helpful suggestions and judgements based on their experience.

MEASURING SETTLING TIME

INVERTING CONFIGURATION

There are a number of circuit approaches that may be used for observing the settling time of an inverting amplifier. Perhaps the simplest (or at least the most obvious) consists of measuring the voltage at the amplifier's own error point. Though superficially simple, this method has several deficiencies: the added capacitance of the measuring device affects the closed-loo dynamics, the dynamic input impedance of the amplifier ma cause errors in estimating the actual gain error, and feedbac circuit dynamics may introduce errors that will not be ob servable at the error point.

Instead of using the actual summing point of the amplifier, i is feasible to construct a quasi-summing point, as shown i Figure 10. Although measurement errors may be caused b capacitance at this quasi-summing point, it does not affec amplifier performance. Because the comparator which observe the error is comparing the proxy error signal with "ground,' this is by far the easiest kind of measurement to make. Th comparator's output is readily clamped to minimize overdriv of both the comparator and the monitoring oscilloscope, th comparator may have gain, and it may also settle more slowl than the amplifier under test (its settling time errors are second order, in this case.)

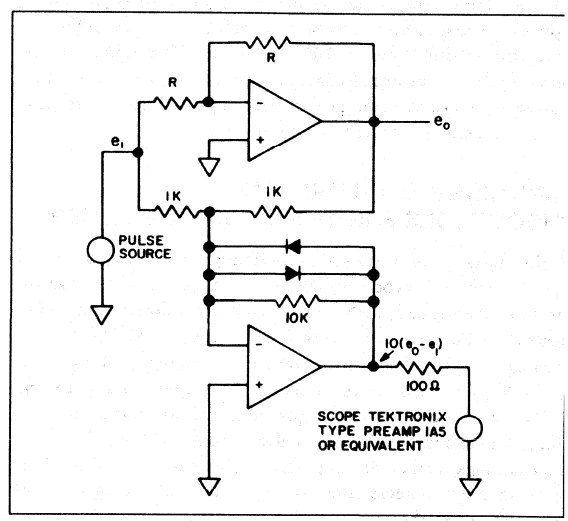

Figure 10. Settling Time Measurements at Quasi-Summing Point.

The above discussion pertains to an amplifier used with a voltage signal and input and feedback resistors. However, if the signal is a current step (amplifier used as a current-to-voltage transducer), the output step response must be compared with the expected final value. Since the comparator is taking the difference directly, rather than observing an error point, it will be subjected to wider swings, and it must settle substantially (at least 2 to 3 times) faster than the amplifier under test to avoid introducing excessive delay and time uncertainty. Needless to say (having already dwelt heavily on the problems of designing a fast-settling amplifier itself), the problem of designing and instrumenting a comparator to test settling time (and properly interpreting the results obtained with it) presents challenges that make the amplifier design job look easy.

NON-INVERTING CONFIGURATION. The most usual circuit to be tested is the unity-gain follower. It requires direct comparison between the input step and the output signal as they swing through the entire common-mode range. Thus the problem, and its instrumentation, may be similar to that of the inverter with current source.

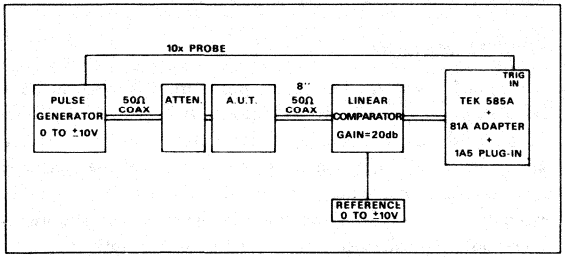

Figure 11. Instrumentation for Settling Time Measurement

EQUIPMENT. Figure 11 shows a typical measurement system for settling time. Commercially available comparators, such as the Tektronix Models W and 1A5, the Hewlett-Packard 1803A, and the Adage Ultranull NDI are useful, but their applicability is limited to settling times somewhat greater than 1μs for 0.01% measurements. As a rule, specially-designed and constructed equipment is required in order to assure oneself that the system requirements are met, including low noise, minimization of excess lead inductance, capacitance, and (naturally) adequate settling time.

The pulse generator used for the measurement should have rise times of the order of 10ns, with a flat top occurring within (say) a few hundred nanoseconds of the step transition. A square wave of about 100Hz should be applied initially to determine output amplitude and whether the amplifier exhibits slow recovery phenomena (due to thermal transients, dielectric absorption, etc.) Once the observer is satisfied of the absence of long tails, a much faster rate may be used to observe fine structure. Some low cost signal generators exhibit a degree of pulse "droop," that may cause what appear to be long tails in observations of settling time of an amplifier, especially in the follower configuration (because the amplifier output is being compared with a dc level). The dummy summing point approach of Figure 10 is less sensitive to pulse droop, because addition of input and output waveforms tends to cancel the droop, assuming that the amplifier reproduces the droop with fidelity.

Most of the measurements performed to obtain the V curves used the specially-constructed square wave source and comparator schematically illustrated in Figures 12 and 13. The square wave source is capable of supplying both positive and negative outputs variable from zero to ±10 volts, with a rise time of about 15 nanoseconds.

The comparator shown was used for all but the fastest amplifiers. It settled in about 0.75μs for measurements in a 1mV error band, and about 0.5μs for measurements in the 10mV error band. Its performance had been intentionally compromised by adjusting its input impedance so that it could be used for measuring outputs of G.P. amplifiers having only 5mA available. (Amplifiers designed for fast settling usually have outputs of 20mA or more).

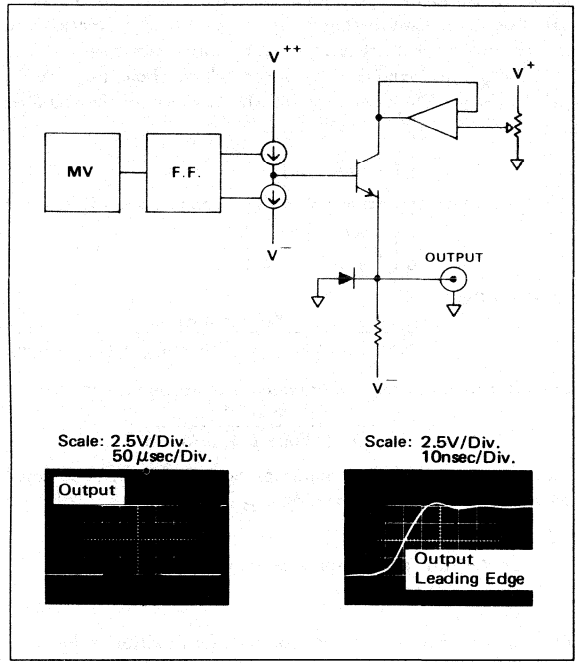

Figure 12. Square Wave Source, Simplified.

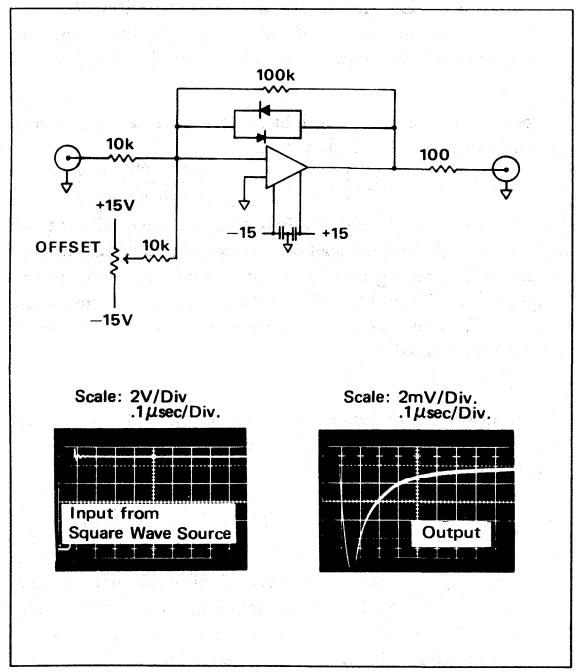

Figure 13. Comparator Simplified

APPENDIX

EFFECTS OF POLE-ZERO MISMATCH ON LINEAR-SETTLING INTERVAL

Consider an ideal* operational amplifier connected as a unity-gain follower. Assume that its open-loop transfer function is as shown in Figure 5, characterized by unity-gain bandwidth f_x ($= 1/2\pi\tau_x$), midband gain A_1, mismatch coefficient m, and DC gain $A_0(\rightarrow\infty)$. The open loop transfer function of the amplifier is

$$A = \left[\frac{A_0}{1 + mA_0\tau_x p}\right]\left[\frac{1 + mA_1\tau_x p}{1 + A_1\tau_x p}\right]$$

The gain of an operational amplifier connected as a follower is

$$G = \frac{1}{\beta}\left[\frac{A\beta}{1 + A\beta}\right] = \frac{A}{1 + A} \qquad (\beta = 1)$$

Substituting,

$$G = \frac{A_0(1 + mA_1\tau_x p)}{A_0(1 + mA_1\tau_x p) + (1 + mA_0\tau_x p)(1 + A_1\tau_x p)}$$

and after many steps, one of which is to assume $A_0 \rightarrow \infty$,

$$G = \frac{1 + mA_1\tau_x p}{1 + (A_1 + 1)m\tau_x p + mA_1\tau_x^2 p^2}$$

The quadratic in the denominator can, of course, be factored to an expression of the form $(1 + Bp)(1 + Cp)$

EXAMPLES

1. Consider now an example: (m = 1)
   $A_1 = 200$ $(A_0 \rightarrow \infty)$
   $m = 1$

One can compute the closed-loop transfer function to be:

$$G = \frac{1}{1 + \tau_x p}$$

As a first-order lag, the output will settle to within 0.01% in 9.2 time constants (see Fig.4). If $\tau_x = 0.1\mu s$, the output (assuming purely-linear settling) would settle within $0.92\mu s$.

2. $m > 1$

Suppose now that $A_1 = 200$, but m = 4. One can compute the closed-loop transfer function to be

$$G \cong \left[\frac{1}{1 + 0.996\tau_x p}\right]\left[\frac{1 + 800\tau_x p}{1 + 803\tau_x p}\right]$$

This can be seen to be equivalent to a unit-lag of about τ_x (as before) cascaded with a lead-lag having instantaneous response of 800/803, and lag time constant of $803\tau_x$! To go the final distance from 800/803 to 9999/10,000 requires settling to one part in 37, or about 3.6 time constants, or (in the case of $\tau_x = 0.1\mu s$) $290\mu s$!

3. $m < 1$

The situation is considerably better for the case where $A_1 = 200$ and m = ¼, but intolerable nevertheless. The closed-loop transfer function is

$$G \cong \left[\frac{1}{1 + 1.015\tau_x p}\right]\left[\frac{1 + 50\tau_x p}{1 + 49.2\tau_x p}\right]$$

Again, it is a unit lag of about τ_x, cascaded with a lead-lag. In this case, there is an initial value of 50/49.2, with a decay time constant of $49.2\tau_x$. The settling from 1.016 to 1.0001 requires 160:1, or about 5 time constants. In the case of $\tau_x = 0.1\mu s$, the final settling occurs after about $25\mu s$.

*"Ideal" in terms of input and output impedances, CMRR, and every other respect, except gain and frequency response.

In the above examples, neglecting the initial fast settling interval accounted for a conservative 7 time constants (i.e., for settling to well within 0.1%), or about $0.7\mu s$.

APPROXIMATION

A reasonable approximation for pole-zero mismatch response calculations, especially for m closer to unity than in the exaggerated example above, is to assume that the amplifier settles to within $1/A_1$ initially in exponential manner, with time constant τ_x. The number of time constants can be estimated quickly from the table in figure 4. The amount of overshoot or undershoot (positive for overshoot) is roughly $(m - 1)/A_1$. The final settling time constant is about $mA_1\tau_x$.

(continued on the next page)

BIBLIOGRAPHY

A. SPECIFIC FOOTNOTE REFERENCES

1. "Frequency Compensation Techniques For An Integrated Operational Amplifier"
   by James N. Giles
   Fairchild APP-117. 1965
2. "Latest Approach to Integrated Amplifier Design"
   by F. D. Waldhauer
   Electronics. May 31, 1963
3. "Operational Amplifier Frequency Response—It's the Shape that Counts" Parts I and II
   by Leonard Kedson and George Tempei
   Electronic Design. July 5 and 19, 1965
4. "An Analysis of Certain Errors in Electronic Differential Analyzers II—Capacitor Dielectric Absorption"
   by Paul C. Dow, Jr.
   IRE Trans on Electronics Computer. March, 1958
5. "Extending An Operational Amplifier's Bandwidth To 50mHz"
   by Richard D. Brugger
   Electronic Design. May 25, 1964

B. OTHER REFERENCES

"Specifying, Measuring, and Using Very High Speed Digital to Analog Converters"
by James J. Pastoriza
Analog Devices Technical Note

"High-Frequency Characteristics of Wide Band Inverter Op Amps"
by Heinrich Krabbe (now Analog Devices—West)
EEE. April, 1969

This paper is an outgrowth of a talk given at NEREM 1968 by Robert Demrow, entitled "How to Specify and Test Fast Settling Operational Amplifiers."

XPERIMENTAL RESULTS

he circuit, Bode plot, and results of a low-frequency simulation of an amplifier, in which $\tau_x = 2 \times 10^{-4}$, $A_1 = 20$, and m is qual to 1, ½, and 2, are shown in Figure 14. In each photograph, three traces are shown, corresponding to unity, x10, nd x100 magnification of the error band (i.e., the final settling nterval). In this example, the amplifier simulated was connected as a unity-gain inverter. The reason the simulation was erformed at low frequency was to retain as complete as ossible control over the idealized characteristic, eliminating lewing phenomena, dielectric absorption, etc.

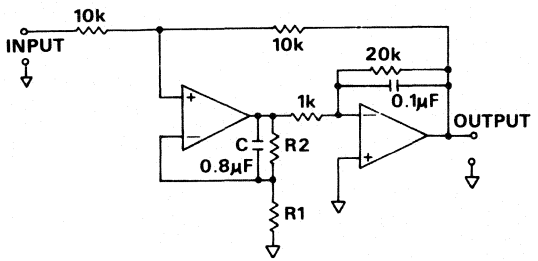

Circuit Schematic For Figures 14a, b, and c.

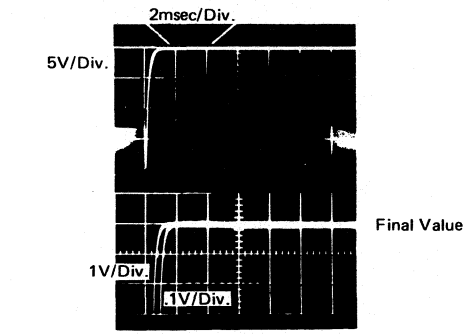

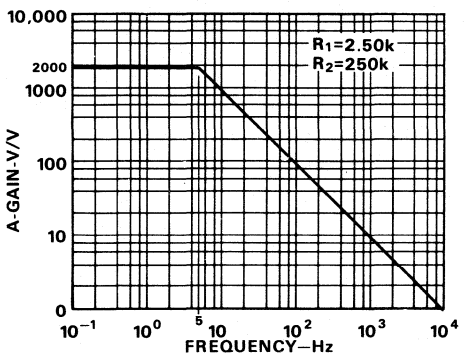

(a) Perfect Match (m = 1)

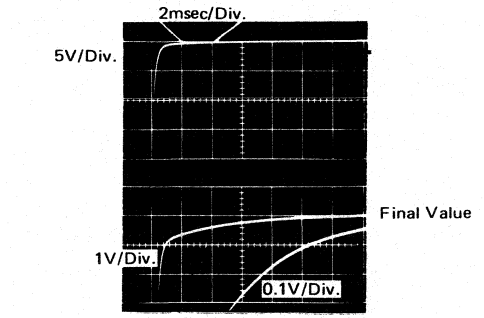

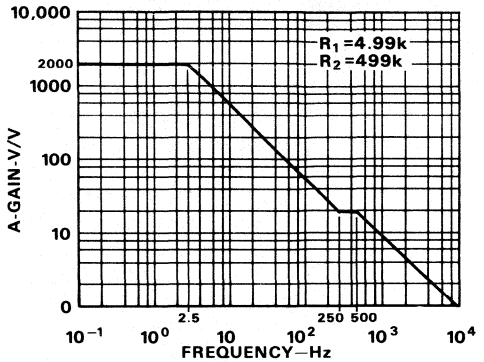

(b) Levelling-off of Response (m = 2)

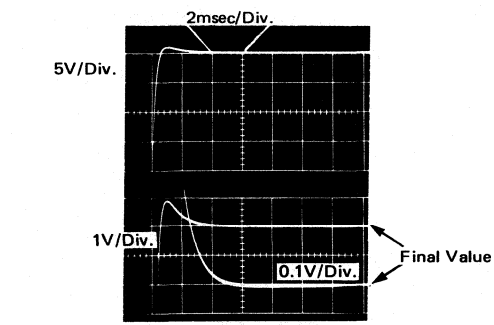

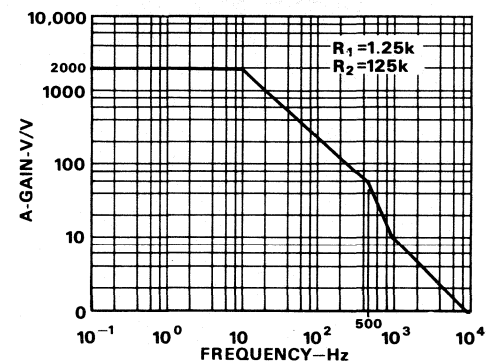

(c) Steepening Slope (m = ½)

Figure 14. Small-Signal Settling as a Function of Pole–Zero Match (Low Frequency Model)

13

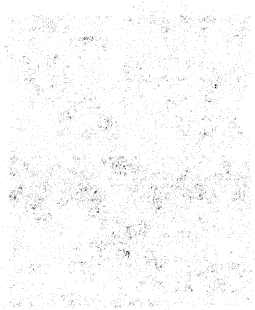

)NE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Accurately Testing Op Amp Settling Times

by Scott Wurcer and Charles Kitchin

ASIC EQUIPMENT FOR TESTING op-mp settling times consists of a igh-quality (and expensive) "flat-)p" pulse generator as input to the UT, a DUT test jig with power ıpply and bypass capacitors at op-mp-socket supply pins, and an os-lloscope to process and display ɜst output. Proper use of a pulse enerator containing a reed switch ın achieve the same result.

Settling-time tests require flat-'aveform inputs because, through-ut most of their frequency ranges, p-amps behave as integrators. herefore, their error-signal out-uts are derivatives of input sig-als. The derivative of an ideal-quare-wave input consists of a ulse at each edge transition and ero everywhere else. Any aberra-ions in the input square wave ap-ear in the error signal. Many ench-top square-wave generators xhibit thermal "tails" and other roblems that limit their suitability)r measuring settling times.

Figure 1 shows a practical circuit)r measuring the settling time of igh-speed op amps. In this case, a ommercial flat-top pulse generator

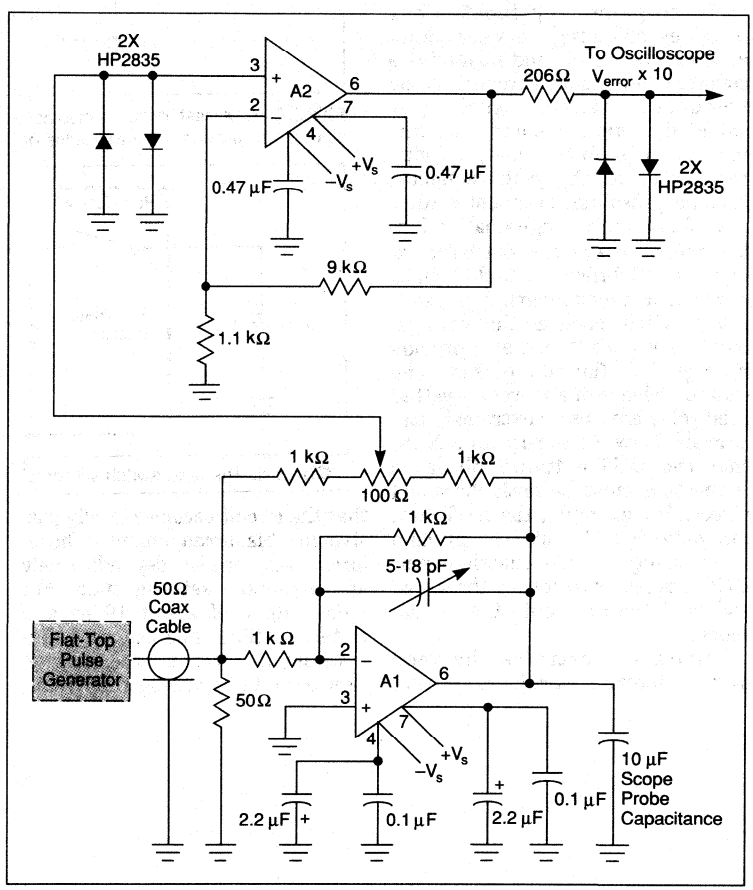

FIGURE 1. A practical circuit for measuring high-speed-op-amp settling time. The DUT-output-to-scope-input connection requires 50 Ω coax less than one foot long.

ɜprinted from Test & Measurement World — May 1990

drives the input. Figure 2 shows the scope's output waveform for a typical high-speed device. The error-signal output from DUT *A1* is clamped, then amplified by op-amp *A2* and clamped again. *A2*, a very high-speed device, provides a voltage gain of 10, amplifying the error-signal output by a factor of five. The clamps, each consisting of two high-speed Schottky diodes, prevent overloading the input of *A2* and the scope preamp. An unclamped high-intensity "tail" could exceed the desired error signal itself. We chose a Tektronix scope preamp type 7A26 because it can recover from the 0.4-V clamped signal quickly enough to accurately measure the DUT's 135-ns settling time.

This test setup uses coaxial cable to connect the pulse-generator output to the test-fixture input and to connect the fixture output to the oscilloscope-preamp input. Coax provides shielding between input and output signals and maintains a constant impedance throughout the frequency range. Audio-type shielded or unterminated coax wiring would produce widely varying responses with changing frequency, causing gross measurement errors.

Unfortunately, high-quality flat-top pulse generators are both expensive and difficult to find. A pulse generator incorporating a reed-relay switch, such as the arrangement in Figure 3, can also provide high-quality flat-top pulses. The contact release of a mercury-wetted reed relay creates an extremely fast current-flow interruption. Note that the DUT-output/oscilloscope connection must be both short and direct. For example, the DUT card can plug directly into the preamp input using a male chassis-mount BNC connector to convey the signal and hold the card onto the oscilloscope.

Figure 4 illustrates the reed switch's internal workings, showing

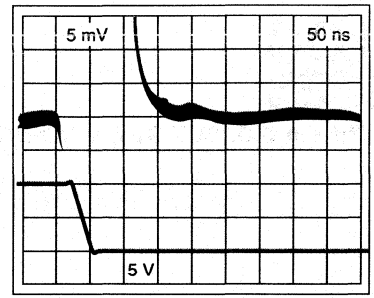

FIGURE 2. The oscilloscope's output waveform for a typical high-speed device, using the equipment in Figure 1.

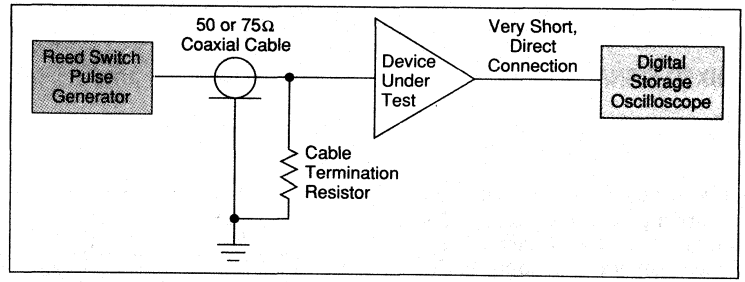

FIGURE 3. A test circuit including a pulse generator that contains a reed-relay switch can provide a flat-top pulse of the highest quality.

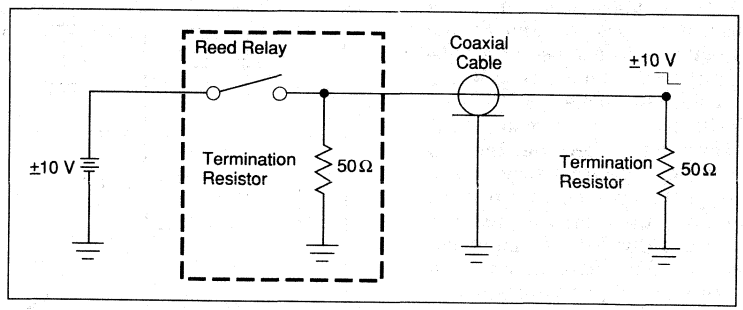

FIGURE 4. The reed switch's internal workings.

that the circuit becomes totally passive in this arrangement. Therefore, pulse quality depends solely on coax-cable settling time. For pulse edges of about 10 ns and faster, settling time can vary dramatically from case to case. This design permits bypassing the circuit, when necessary. However, reed switches only permit very low repetition rates, from about 100 Hz to about 2 kHz. A test can overcome this limitation by capturing and averaging repetitive waveforms on a digital-storage scope.

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

JFET-Input Amps Are Unrivaled for Speed and Accuracy

by Peter Henry

FET-input amplifiers provide an economical means of achieving high accuracy in applications that need wide bandwidths or large signals. They are ideal for pulse amplifiers, fast D/A converters, peak detectors, and logarithmic amplifiers.

FET-input operational amplifiers are an option for designers who require speeds greater than those provided by standard bipolar op amps such as the OP-07. The high slew rates of FET-input amplifiers make these devices attractive for pulse amplification and other applications that require wide bandwidths and handle large signals. Their low bias currents make them equally suitable for peak detectors and logarithmic amplifiers. Furthermore, their fast settling rates make them ideal for fast, high-precision DACs.

To obtain the full performance of which JFET-input amplifiers are capable, you'll need to take all the standard precautions in designing and laying out your pc boards, along with a few extra precautions that are specific to JFET-input devices (see "Caveats, warnings, and design reminders").

A JFET-input amplifier can help a CMOS D/A converter achieve a fast settling time while converting the DAC's current output into voltage levels and reducing output impedances. You'll obtain the fastest settling times from bipolar DACs, because they have lower output capacitance than their CMOS counterparts, but CMOS devices have the advantages of low price and the availability of a wide variety of interface options. The primary disadvantage of CMOS DACs is their large output capacitance, which can be 50 to 120 pF for an 8-bit DAC, and as much as 70 to 150 pF for a 12-bit DAC. This large capacitance increases the settling time. However, if you add a JFET-input amplifier (such as PMI's OP-42) to create a voltage output, you can compensate for the DAC's output capacitance. A CMOS DAC will then settle in approximately 3μs to within 0.01% of a 10V full-scale output step.

13

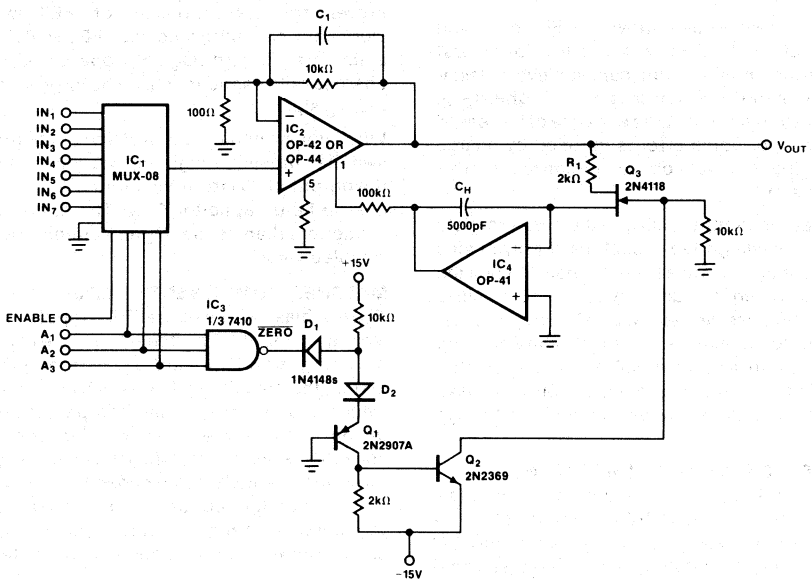

FIGURE 1: *This autozeroing amplifier multiplexes seven inputs to a common output. The eighth input grounds the input of C_2 to zero the amplifier's offset voltage. Before you select a signal, hold the three address lines high for 200μs to ensure proper zeroing.*

Reprinted from EDN – May 14, 1987

The offset voltage of many older JFET amplifiers suffered from large thermal drifts. However, newer state-of-the-art precision JFET-input amplifiers exhibit relatively little drift with temperature, and the resulting output error is generally insignificant unless you operate the amplifier at a high gain level. If your application requires the minimum possible offset error, however, you can use a servo loop that automatically corrects offset-voltage and drift errors. In Figure 1's circuit, for example, IC_1 multiplexes eight analog channels to the input of an OP-42 or OP-44 amplifier, which has a gain of 100. One of the analog channels grounds the amplifier input in order to correct for V_{OS} (offset) errors; the other channels are available for signals.

To correct V_{OS} errors during a conversion, you first drive the three multiplexer-address lines (A_1 — A_3) high, so that multiplexer IC_1 grounds the input of JFET-input amplifier IC_2. At the same time, AND gate IC_3 drives the $\overline{\text{Zero}}$ line low and thereby causes the switching circuitry consisting of transistors Q_1 and Q_2 to turn on JFET switch Q_3. This action connects JFET-input amplifier IC_4 into the feedback path of IC_2 via that IC's null pins (1 and 5) and thereby forces the output of IC_2 to assume the value of the offset voltage at the input of IC_4. The current in the feedback loop then develops a voltage across hold capacitor C_H.

KEEP LEAKS AWAY FROM CORRECTION CIRCUITRY

After a time period that's determined by the RC time constant of R_1C_H and the current through Q_3, you can change the multiplexer address so that the $\overline{\text{Zero}}$ line goes high and turns off Q_3. The voltage across C_H holds the output of IC_4 (which is also the offset-voltage compensation for IC_2). IC_1 has a relatively long switching time, so Q_3 turns off before IC_1 connects a new input to IC_2; consequently, the signal cannot leak into offset-correction circuits.

If you use the component values shown in Figure 1, you should make sure that the $\overline{\text{Zero}}$ line remains low for at least 200μs to ensure proper nulling. You can achieve a faster nulling time by using a JFET switch that has a higher I_{DSS}, such as a 2N4393, but you'll do so at the expense of increased leakage and faster droop. Some error is induced by charge injection through Q_3 into C_H, but you can minimize this error by using a large value of C_H.

To minimize droop at the output of IC_4, make sure that hold capacitor C_H is a low-leakage type (such as a polystyrene device). For higher precision, add a potentiometer to null the offset voltage of IC_4. When you use the component values shown, the droop of the offset voltage at 25°C is only 1.3μV/s. Near the center of the adjustment range, each 100mV of swing at the output of IC_4 will cause a shift of 150μV in the V_{OS} of IC_2. The circuit is capable of correcting as much as 10mV of offset, so that it can handle some system offsets in addition to IC_2's offset.

YOU CAN SUBSTITUTE DIGITAL CORRECTION

If your application requires digital correction of the offset voltage, you can substitute an ADC/DAC combination for IC_4. With this scheme, when you ground IC_2's input, the ADC digitizes the IC's output (offset) voltage and passes a correction factor to the DAC, which in turn applies an analog nullin voltage to pin 1 of IC_2. This modification is of value in applic tions that digitize the output of IC_2, and it has the advantag that digital correction circuits do not droop with time. How ever, the scheme is a needless complication in systems th do not include a μP for other purposes.

The level-shifting circuitry (D_1 and D_2, Q_1 and Q_2) converts TTL signal to the levels necessary to drive JFET switche and you can use the same circuit in a wide variety of applic tions. When the TTL input signal ($\overline{\text{Zero}}$) goes low, it turns o transistor Q_1. This action forces the base of transistor Q_2 t −15V, turning off Q_2 and holding the gate of Q_3 at groun level. While these conditions continue, Q_3 presents a lo impedance to the signal applied to it through R_1. Conse quently, the inverting input of IC_4 follows the output signal c IC_2, and IC_2 charges C_H.

When the TTL $\overline{\text{Zero}}$ signal goes high again, it turns on Q_1 an Q_2, which in turn pull the gate of Q_3 to −15V. This action put Q_3 into a high-impedance state, so that the switch discon nects the input of IC_4 from the output of IC_2, and IC_2 main tains the charge across C_H.

FAST S/H AMPLIFIER EXHIBITS 0.01% ACCURACY

The characteristics of JFET-input amplifiers make them nat ural choices for fast sample-and-hold (S/H) amplifiers. Fig ure 2a's circuit has an aperture time of 80ns and can acquire 10V step in less than 1μs with an accuracy of 0.1%, and in 2μ with an accuracy of 0.01%. The corresponding settling time are 100ns and 300ns, respectively.

The $\overline{\text{Sample}}$/Hold control-input circuit uses the same level shifting circuit used by the autozero circuit described above with the addition of a Schottky clamping diode (D_3) and pullup resistor R_1 to accelerate transitions. Diodes D_4 and D prevent the forward-biasing of JFET switches Q_3 and Q_4 During the sampling time, the JFET switches conduct, so tha amplifiers IC_1 and IC_2 both operate at unity gain and IC charges hold capacitor C_{H1}. The V_{OUT} signal precisely track the V_{IN} signal.

When you switch the control line to the hold mode, the JFE switches present a high impedance, so that Q_3 disconnect the output of IC_1 from the input of IC_2, and Q_4 allows IC_2 t charge hold capacitor C_{H2}. The charges on the two hol capacitors then cause IC_2 to maintain the V_{OUT} signal at th sampled value.

MATCHED COMPONENTS REDUCE ERRORS

The low bias current of the JFET-input amplifiers and the low leakage of the JFET switches combine to minimize leakage of the charges on C_{H1} and C_{H2}. Furthermore, you can closely match the remaining leakage by matching the capacitor values and by using a matched pair of JFET switches contained in the same housing. This configuration has two advantages. First, it matches the amounts of charge injection through the switches into capacitors C_{H1} and C_{H2} and there by considerably reduces the hold step. In fact, adjusting a trimmer capacitor connected in parallel with one of the hold capacitors will let you reduce the hold step below 1mV.

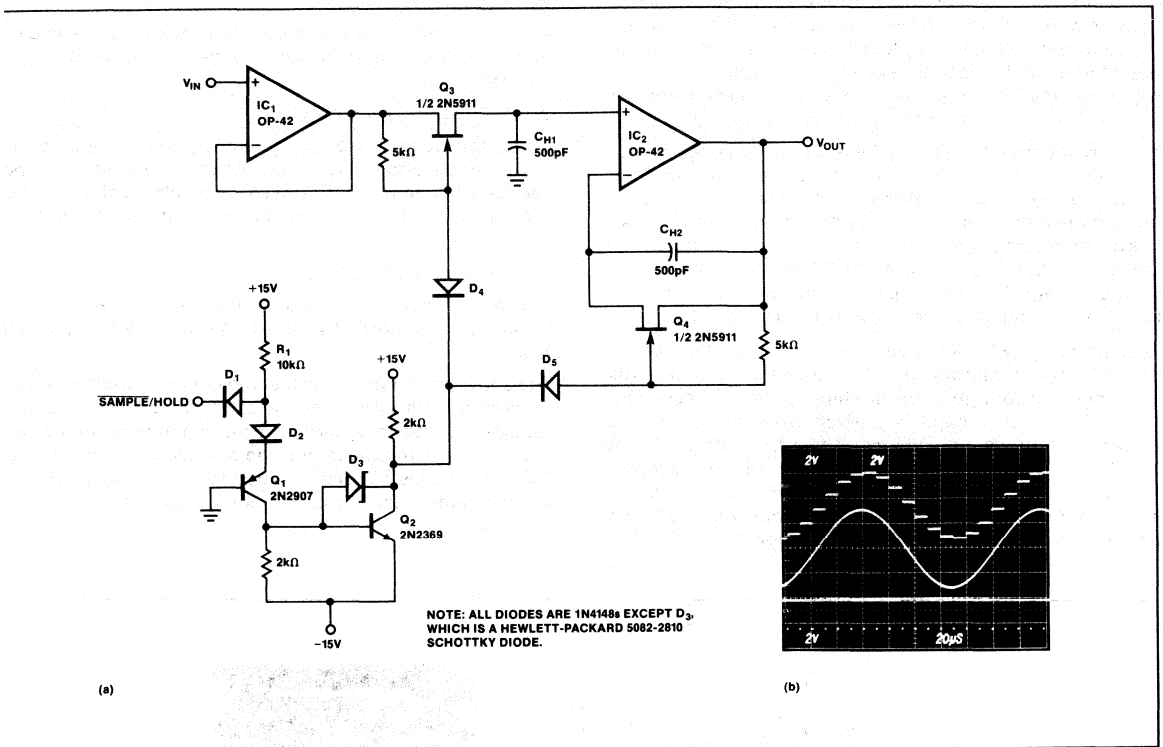

IGURE 2: *This fast sample/hold amplifier (a) uses two hold capacitors to minimize both the hold-step effect and the output roop. You can see the very small hold step and the absence of droop in the scope photo (b).*

13

econd, the scheme causes the I_{OS} of IC_2 and the leakage hrough the matched JFET switches (rather than absolute eakage levels) to control the voltage droop at the output erminal. Furthermore, the output voltage is controlled by the ifferential voltage between the two capacitors. The absolute oltages across the capacitors droop because of IC_2's bias urrent, but both droop at the same rate, making this voltage hange appear as a common-mode effect. The differential oltage is controlled by I_{OS}, which is usually much smaller han I_B. IC_2 can maintain a constant output voltage, even hough the actual voltage levels on the two capacitors may hange considerably during the hold period. In Figure 2a's circuit, this scheme reduces the droop rate to 7μV/ms. Figure 2b shows a sine wave applied to the S/H circuit, and the resulting samples at the output.

The control provided by these differential voltages breaks down when their absolute values fall below the minimum input voltage of IC_2. However, that condition will typically not occur until several seconds have elapsed. You can reduce both the hold step and the droop rate even more by using larger values for the capacitors, but you'll then sacrifice some speed. (For more on the examination of errors, see "Calculating error magnitudes.")

You can exploit the low bias currents inherent in a JFET-input amplifier to create a fast peak detector (Figure 3a) that can capture a 10V peak only 2.5μs wide. This circuit uses the level shifter employed by the two previous circuits, but it does not drive any FET switches. Instead, a Reset command causes Q_2 to discharge hold capacitor C_H directly.

After a reset operation, the voltage across C_H is negative, so the circuit will detect very small peaks even if the ac signal has a negative dc offset. Be aware, however, that the reset method prevents this circuit from detecting negative peaks. The droop rate is 3mV/ms and is mainly due to leakage of the charge on C_H through Q_2; diode D_5 contributes very little charge leakage, because it operates with only a small voltage across it. The action of the circuit is illustrated in Figure 3b.

You'll find that JFET-input amplifiers yield excellent performance in logarithmic amplifiers. Figure 4's circuit follows the usual log-amp configuration, in which the logarithm (to base 10) of the input voltage is derived from the differential between the threshold voltages of two transistors (Q_1 and Q_3). A precision reference source (IC_2) sets the collector current of Q_3, and the input voltage controls the collector current of Q_1. You can derive the threshold voltages (V_T) of the transistors using the following equation:

$$V_T = (kT/q)\ln(I_C/I_S),$$

where q is the electronic charge, T is the temperature, and k is Boltzmann's constant. The difference in V_T between the two transistors is then

$$\Delta V_T = (kT/q)\left[(\ln I_{C1}) - (\ln I_{C3})\right].$$

The output voltage is ΔV_T multiplied by a factor appropriate to your application. The circuit scales the output for 1V per decade with a zero-crossing at an input voltage of 100mV. For a V_{IN} greater than zero, this configuration yields a transfer function of

$$V_{OUT} = -\left[\log(V_{IN}/R_{IN}) + 4\right]$$

(for V_{OUT} in volts and the quantity V_{IN}/R_{IN} in amps). The circuit operates correctly for inputs from 1mV to more than 10V.

The above equations show that the temperatures of the log-conversion transistors have a direct effect on the output voltage. Conventional circuits provide temperature correction by using a thermistor instead of a resistor in place of R_7. This circuit uses a special method to provide an isothermal environment for the log-conversion transistors.

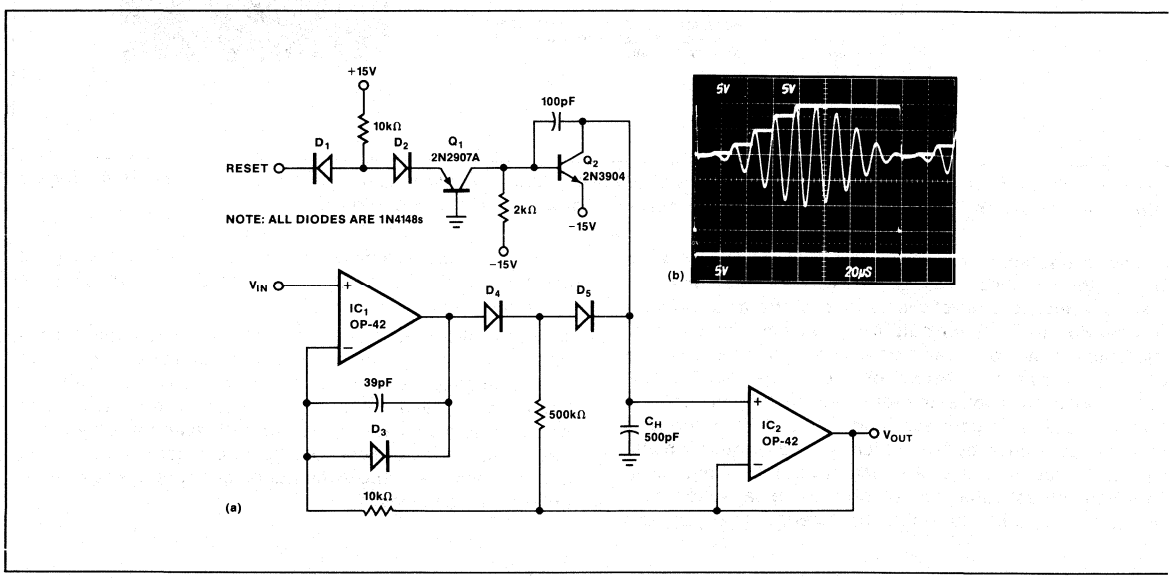

FIGURE 3: *This peak detector (a) can capture a 10V peak that's only 2.5μs wide. After a reset it can detect very small peaks even if the ac signal has a negative dc offset. The detector's JFET-input amplifiers have very low input bias currents and therefore minimize leakage from the hold capacitor. This low leakage is responsible for the absence of droop (b).*

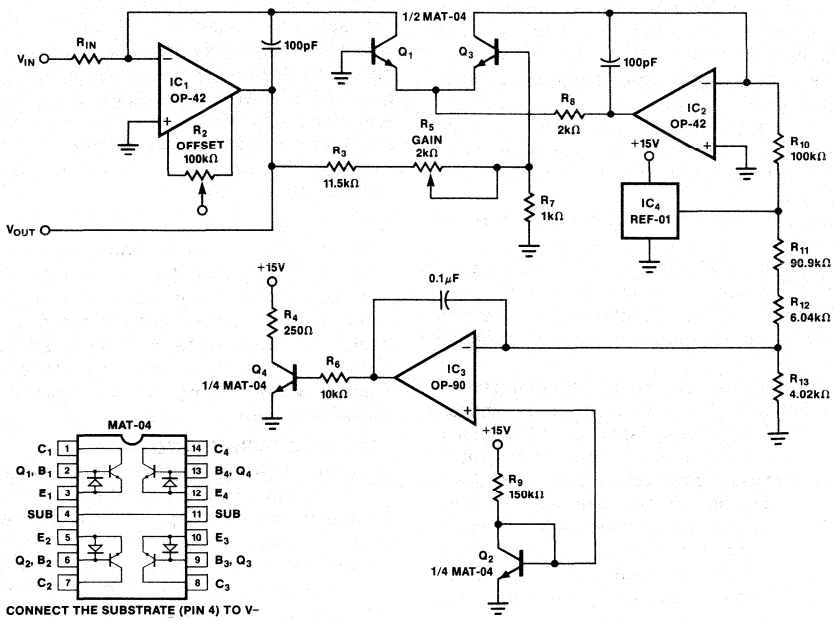

IGURE 4: *This logarithmic amplifier eliminates temperature effects by ensuring that the two log-conversion transistors are 'ways at the same temperature. The conversion transistors are on the same substrate as a heater transistor and a tempera- ire sensor, which are connected to each other in a feedback loop.*

13

he MAT-04 IC used in this circuit is a symmetric array of four ansistors placed at the corners of a square die. Two of these ansistors, located at diagonally opposite corners, act as the ıg-conversion elements. Of the remaining two transistors, ne (Q_4) acts as a heater and the other (Q_2) acts as a tempera- ire sensor. IC_3 forces the V_T of Q_2 to a specific value by ary ing the current through Q_4, and it maintains this value by ieans of the thermal feedback between Q_2 and Q_4. The ymmetrical layout of the IC ensures that the two log- onversion transistors are always at exactly the same tem- erature. The component values shown will maintain the IAT-04 die at approximately 60°C.

This operation may violate the rated specifications of the MAT-04 package and cause degradation of Q_4's β, but it does not hurt performance, because the characteristics of the heater are unimportant. You'll get the best results by encasing the MAT-04 package in thermally insulating foam, such as the urethane foam used for housing insulation.

To null the amplifier, you begin by setting the input voltage to 1mV. Adjust the offset voltage of IC_1 for an output of 3V. Next, raise the input to 10V and adjust the gain for an output of −1V. These two adjustments are interactive, so you may have to repeat them several times. You can modify the zero-crossing point by changing the value of resistor R_3.

CAVEATS, WARNINGS, AND DESIGN REMINDERS

The following design principles are critical for the performance of any high-speed amplifier. They are included here as a check list.

- Use separate supply traces and grounds for each amplifier.
- Bypass each supply, right at each amplifier, with a 1 to 10μF tantalum or electrolytic capacitor connected in parallel with a glass or ceramic capacitor that has a value of 0.01 to 1μF.
- Provide separate supply lines and ground lines for the digital and analog sections of your system.
- Switching power supplies can inject spikes of several hundred millivolts into the supply lines, and they can radiate EMI. Shield the analog sections and bypass all supply lines at the point where they enter the shielded enclosure.
- Be careful not to exceed an amplifier's maximum input-voltage specification. If a signal could be applied before the amplifier's supplies reach their full values, provide clamping diodes — but remember that these devices add leakage and capacitance to the circuit.
- Be careful not to exceed the maximum junction temperature or the maximum power-dissipation ratings of an amplifier. If you connect a capacitive load to the output of an op amp, be sure to include in your calculations the power dissipation caused by the rms ac currents delivered to the load.
- Internal power dissipation raises an IC's junction temperature. You can find the amount of the increase from the formula $\Delta T = P_D \Theta_{JA}$, where P_D is the power dissipation and Θ_{JA} is the thermal resistance of the package. For 8-pin packages lacking a heat sink, this parameter varies from 90°C/W (plastic DIP) to 200°C/W (TO-99 can). The junction temperature is the ambient temperature plus ΔT.

PRECAUTIONS FOR JFET AMPS

The following design principles also have general application, but the focus here is on their implications for designing with JFET-input op amps.

First, remember that a rise in junction temperature increases the bias current of a JFET-input op amp; a rise of 10°C may double the bias current. JFET-input op amps have a naturally low bias current, however; for PMI's OP-42, for example, the current is only 200pA at room temperature and less than 20nA over the full military temperature range. The errors produced by such small currents are usually insignificant; at a slew rate of 100V/μs, you would need a current of 1mA to drive stray capacitances amounting to only 10pF. Although some JFET-input op amps use cancellation methods to derease bias current, these techniques can create excessive phase shifts in high-speed amplifiers.

Remember also that the slew rate of an op amp varies according to the voltage difference between its two inputs. If you want to achieve the maximum slew rate specified in the data sheet, you must ensure a difference of about 2V between the inputs of a JFET-input op amp so that one side of the op amp's differential-input circuit turns completely off. At unity gain, such voltages are normal, but in circuits that have a higher gain, the input-voltage levels —and hence the slew rate — decrease. A JFET-input op amp that yields a slew rate of 60V/μs at unity gain might yield only 20V/μs if you operate it at a gain of 100 with a ±100mV input signal.

You should also keep in mind that an amplifier that has a high slew rate or a wide bandwidth doesn't necessarily settle fast. Many amplifiers with high slew rates obtain their speed at the cost of inducing excessive ringing in the output waveform; this ringing increases the settling time. Remember, too, that the ac characteristics of some amplifier types vary widely from part to part. Data sheets usually specify a typical settling time. Very few vendors guarantee a maximum value.

VARYING COMPENSATION

Most JFET-input op amps have input capacitances from 4 to 8pF. A small capacitor placed across the feedback resistor compensates for the pole created by the input capacitance. The amount of compensation needed depends on the performance you expect from the amplifier. Critical damping may give the fastest settling times to within very narrow error bands. In general, however, you'll improve the settling time, even to error bands as small as 0.01%, by providing compensation that yields slight underdamping. The optimum compensation is a function of the circuit and its layout, and you'll have to determine its value by experiment.

Proper compensation becomes critical when you use an op amp to convert the current output of a DAC to a voltage output. The output capacitance of the DAC, in parallel with stray capacitance and the input capacitance of the op amp, exacerbates any ringing and instability problems and you'll have to optimize the compensation for the combination of settling speed and accuracy that you want.

The gain-bandwidth product (GBW) is adequate to describe the ac response, at any frequency, of single-pole amplifiers such as the 741. The more complex design of a JFET-input op amp such as PMI's OP-42 yields higher slew rates with greater stability, but it distorts the meaning of the GBW. Nevertheless, you can derive an approximation of the cutoff frequency for any closed-loop gain (A_{VCL}) from the following formula:

$$f_C = (GBW/A_{VCL}).$$

This approximation is adequate for most purposes and is valid for most amplifiers, including PMI's OP-42 and OP-44.

The slew rate (SR) of an amplifier largely determines the maximum frequency at which it can operate with large

Continued

CAVEATS, WARNINGS, AND DESIGN REMINDER *Continued*

signals. You can calculate this frequency (known as the power bandwidth, or BW_P) from the equation $BW_P = SR/(\pi V_{p\text{-}p})$. An amplifier such as the OP-42, which has a 50V/μs slew rate, can operate at frequencies above 800kHz with only 1% distortion on a $20V_{p\text{-}p}$ signal. The OP-44 has a BW_P that's greater than 1.5MHz; it achieves a slew rate of 100V/μs in applications that have a closed-loop gain greater than three.

PMI's OP-42 guarantees settling times of 1μs or less to an accuracy of 0.01% — that is, to within ±1mV error bands — for a 10V input step. You can approximate settling times for input steps other than 10V by subtracting the slew time from the specification and adding the slew time for the desired output change. For example, to obtain the settling time for a 1V step, subtract the slew time for 10V (167ns at 60V/μs) from the 800ns typical settling time. To this result (800 − 167 = 633ns) add the slew time for 1V (16.7ns) to obtain a calculated settling time of approximately 650ns to 0.01%. The OP-42's measured settling time for a 1V step is somewhat better, being less than 600ns.

CALCULATING ERROR MAGNITUDES

For a concrete example on which to base an examination of errors, assume a circuit consisting of the autozeroing amplifier of Figure 1 as the first stage, and the S/H amplifier of Figure 2 as the second stage. Provide zeroing pulses to the amplifier in the first stage once every milli-second to eliminate thermal drifts and offset problems.

For the purposes of this discussion, assume that IC_2 is a PMI OP-44, with a GBW of 50MHz (think of the entire circuit of Figure 1 as this gain stage). Because IC_2 operates with a gain of 100, the system bandwidth is 500kHz before signals enter the S/H amplifier. The S/H circuit yields 2.5μs acquisition times and holds the output for 18μs; consequently, the sampling bandwidth is 50kHz. Assume that the power supplies are well regulated so that you can ignore power-supply rejection. Also, ignore phase errors.

ASSESSING GAIN-STAGE ERRORS

In the first stage, errors arise primarily from finite gain, common-mode rejection (CMR), and noise. Servo amplifier IC_4 nulls IC_2's offset voltage and drift; these values are no greater than 1mV over the full temperature range, without additional adjustment. The drift is negligible during the periods between nulling. IC_4 also nulls the offset voltage caused by bias current flowing through the multiplexer. If you assume that source impedances are no greater than 1kΩ, then the contribution traceable to bias current is less than $1.5\mu V \times A_{VCL}$, or an additional 1.5mV. (This includes leakage current from the switches in the MUX-08.) At dc, the common-mode error is only 0.004% (essentially nonexistent), but at 50kHz the CMR falls to 70dB and can contribute an error of 0.03%.

You can express open-loop gain errors as a percentage of the signal; at dc, these errors are approximately equal to the percentage calculated from A_{VCL}/A_{VOL}. Both the OP-42 and the OP-44 have an open-loop gain of more than 500,000. At dc, the gain error is less than 0.02%, but at 50kHz, gain errors can contribute amplitude errors as high as 0.5%. The amplitude error decreases rapidly as the operating frequency moves away from the 500kHz cutoff frequency (f_C).

You can obtain an approximate value of the error, for any frequency f, from the formula

$$\epsilon_A \approx 1/2(f/f_C)^2,$$

down to the frequency at which the formula yields a value that's less than the dc value. In Figure 1's circuit, this point occurs at 10kHz. Because the amplitude error is primarily a function of f_C, operating the amplifier at a lower closed-loop gain would result in a significantly smaller error at any given frequency.

You can calculate the noise by multiplying the square root of the bandwidth by the rms noise density. In a wideband amplifier, the noise is dominated by the high-frequency flatband noise, rather than by the higher-density low-frequency noise. For PMI's OP-42 amplifier, the flatband-noise density is typically $12nV/\sqrt{Hz}$. To get an idea of the worst-case performance, use $15nV/\sqrt{Hz}$, which yields a value of 10.6μV rms at the input and 1mV at the output. Always use the full bandwidth of the circuit for noise calculations, because noise frequencies higher than the maximum signal frequency can alias into the lower frequencies and affect the final value.

You'll see that, because error terms sum in an rms fashion, gain error is the dominant source of error. The OP-42 and OP-44 amplifiers minimize, but do not eliminate, these errors. Consequently, you'll find that the output of the autozeroing amplifier is accurate to within 0.5% at 50kHz, and to within 0.02% at frequencies below 10kHz, with the addition of a 7mV noise contribution and 2.5mV of offset error.

These gain-stage errors are passed on to the S/H amplifier, which has its own sources of error, arising from aperture time and aperture jitter, hold steps, droop, and finite CMR (think of the entire circuit of Figure 2 as the S/H stage). However, because the amplifiers of the S/H stage operate at unity gain, they eliminate gain error that's caused by bandwidth limitations or high closed-loop gain,

Continued

CALCULATING ERROR MAGNITUDES *Continued*

and they greatly reduce the other errors associated with the multiplexed stage.

You can simplify matters by eliminating the input amplifier (IC_1) and driving the S/H buffer (IC_2) directly from the output of the multiplexed stage. Then the primary source of error affecting the output will consist of the hold step (1mV) and acquisition error (0.01%). Errors arising from aperture uncertainty and aperture jitter will be minor and will add no more than 0.01%.

The CMR error can be more significant, depending on the frequency. It's difficult to estimate what frequency to use for calculating errors in the S/H stage; some designers use the dc specifications, on the assumption that the amplifier operates in dc mode after the hold settling time. This assumption would be true for long hold times, but in fact the S/H stage not only reproduces the primary waveform but also adds high-frequency components to form steps. For this reason you should always use actual operating frequencies in your error calculations. Consequently, at 50kHz, the 70dB CMR of the OP-42 amplifier can contribute 0.03% error. The error terms again sum in rms fashion to yield a total of slightly less than 0.04% error from the S/H stage. Noise is negligible in comparison with that delivered by the gain stage, and the dc offset is simply that of the OP-42, an additional 0.75mV.

IMPROVING PERFORMANCE

The total system error for both the gain stage and the S/H stage consequently becomes approximately 0.05% at frequencies less than 10kHz and increases to 0.5% at 50kHz. This is in addition to the 7mV noise contribution and 2.25mV dc offset. You could eliminate gain-stage errors by cascading two amplifiers, each with a gain of 10, instead of using a single amplifier with a gain of 100. The error terms contributed by the S/H stage would dominate in such a configuration.

You can calibrate and eliminate many of the error terms discussed above. For example, in a system that uses a fast Fourier transform to examine the spectral content of a waveform, you could apply a correction factor to correct for gain roll-off at high frequencies. You could also correct for CMR errors by applying an additional correction factor based upon the signal level. DC offsets are simply eliminated by subtracting a constant term from the signal. Noise can be eliminated only by averaging many repetitive signals.

ANALOG DEVICES

AN-257
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Careful Design Tames High Speed Op Amps

by Joe Buxton

Operational amplifiers continuously push the limits of speed and bandwidth. Today's high-speed IC op amps reach gain-bandwidth products in the hundreds-of-megahertz range—numbers unheard of just a few years ago. With such performance, designers must be extremely careful in preserving the op amp's stability without sacrificing bandwidth. Circuits and layouts previously used when designing with lower-frequency devices must be rethought in detail. Otherwise, circuit stability and ac performance can be impaired significantly.

Many factors cause instability in high-speed op amps. These include capacitive loading, inadequate power supply bypassing, input capacitance, and lead-lag compensation. Designers have dealt with both capacitive loading and supply bypass problems for some time. But at high frequencies, their effects are more critical and potentially harmful. Moreover, input capacitance and lead-lag compensation are usually ignored for lower-frequency circuits. Now their consideration is vital.

Driving a capacitive load is one of the most troublesome and difficult problems to overcome because it can easily cause circuit oscillation. When combined with an amplifier's output resistance (R_0), a capacitive load (C_L) creates a pole in the feedback loop that increases the closed-loop phase shift *(Fig. 1)*. Depending on its frequency, the phase shift can reduce phase margin, potentially causing the circuit to become unstable. This phase shift is easily calculated from the pole frequency (f_c):

$$f_c = 1/(2\pi R_0 C_L)$$

$$\text{Additional phase shift} = \text{TAN}^{-1}(f_U/f_C)$$

where f_u is the open-loop unity-gain frequency (the op amp's unity-gain bandwidth).

To verify these equations, compare the open-loop gain and phase responses of a 10-MHz, unity-gain-stable op amp, such as the OP-42 with and without a capacitive load *(Fig. 1, again)*. The network analyzer plots indicate the no-capacitive-load condition *(Fig. 1a)*. They also show a loading of 450 pF *(Fig. 1b)*. The 450-pF capacitive load combined with the OP-42's 45-Ω output impedance introduces a

Reprinted from Electronic Design – April 11, 1991

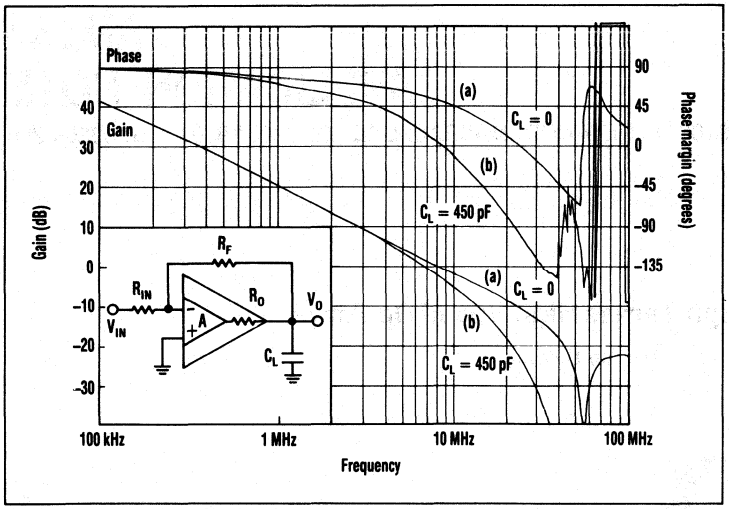

1. A CAPACITIVE LOAD (C_L) combines with an op amp's output impedance (R_O) to alter the gain and phase response, and thus the phase margin. The upper curves (a) are for an unloaded device, the lower curves (b) are for the same op amp driving a 450-pF load.

pole at f_c = 8 MHz. This results in an additional 45° of phase shift. Consequently, what used to be a stable circuit with 50° of phase margin degrades to a phase-margin of only 5°—potentially causing instability.

Classy Classics

The classical way to maintain stability is to compensate for load capacitance by adding a resistor (R_X) in series with the amplifier's output impedance and a shunt capacitor (C_F) in the feedback path *(Fig. 2, left)*.[1] The basic technique requires adding the proper shunt capacitance and series resistance so that the external feedback network adds a net 0° of phase shift to the loop. Amplifier stability depends on the phase shift of the signal that's fed back to the op amp's inverting input. The signal's phase shift must be less than 180° when the loop gain is greater than or equal to 1. If the feedback network contributes 0° of phase shift, the signal is only phase shifted by the op amp. Assuming the op amp has enough phase margin for the particular gain used, stability is ensured.

Now that the stability goal of the feedback network is established, how is it achieved? Redrawing the external feedback network helps clarify the analysis *(Fig. 2, right)*. Each capacitor contributes a pole and a zero to the feedback network. Intuitively, if the pole and zero contribution of one capacitor cancels the zero and pole contribution of the other capacitor, there will be 0° of phase shift. With this in mind, just derive the pole and zero locations for each capacitor, then set them equal to each other and solve for R_X and C_F. While it is straightforward in concept, the actual derivation is extremely involved and time consuming. But it can be approximated by taking an intuitive approach.

Because capacitive reactance changes with frequency, it can be assumed that a capacitor is an open circuit at 0 Hz and a short circuit at infinite Hz. To simplify network analysis, this principle is applied to one capacitor at a time. For the first case, assume that C_F is a short circuit, resulting in both a pole and a zero location as a function of C_L *(Fig. 3a)*. Next assume that C_L is an open circuit, again providing a pole and a zero location, but as a function of C_F *(Fig. 3b)*. Now there are two poles and two zeros. By equating the poles to the zeros, the necessary value for R_X and C_F can be found with the following two equations:

$$R_X = R_O R_{IN}/R_F$$

$$C_F = (1+1/|A_{CL}|)[(R_F + R_{IN})/R_F^2]\, C_L R_O$$

where A_{CL} is the closed-loop gain.

By experimenting, it was found that the $1/A_{CL}$ term needed to be added to the equation for C_F. Just these two equations enables virtually any op-amp circuit to be compensated for virtually any capacitive load. A complete derivation performed at the PMI division accurately predicts the previous two equations, including the $1/A_{CL}$ term.

Though this method of compensation yields a stable circuit for any capacitive load, it reduces circuit-bandwidth drastically. The bandwidth is no longer determined by the op amp, but rather by the external components. C_F and R_F dominate, creating a closed-loop bandwidth of:

$$f_{-3\,dB} = 1/2\pi C_F R_F$$

To show the limiting factors, substitute the equation for C_F into the previous equation and simplify:

$$f_{-3\,dB} = 1/2\pi C_L R_O (1 + 1/A_{CL})^2$$

This equation shows that the load capacitance (C_L) and the op amp's output impedance (R_O) need to be made as small as possible. Because R_O is internal to the op amp, the only way to minimize it is to choose an op amp with low R_O. However, there is much more flexibility in controlling C_L. Consider all of the possible sources for C_L in the circuit and try to minimize them. For example, driving an unterminated coaxial cable can add significant amounts of load capacitance. It's important to back-terminate the cable to remove this capacitive load. A pc-board trace with a surrounding ground plane can also add a small amount of capacitance. To reduce its capacitive loading, keep the trace short and keep the ground plane away from it. Even if capacitive loads don't cause the circuit to oscillate, they should be minimized so as not to limit the closed-loop bandwidth.

Power-supply bypassing is often

dealt with by connecting a 10-μF tantalum capacitor in parallel with a 0.1-μF ceramic capacitor between the supply line and ground. Though this works in many cases, an understanding of what's actually happening helps prevent oscillations where bypassing is critical.

To start, an equivalent circuit of a typical power supply line with the bypass capacitor in place can be represented by an inductor in series with the supply line and a capacitor connected from the IC to ground. The line from the power supply is far from a perfect zero-impedance source, especially after it's routed through a wiring harness and pc-board trace. In addition, every supply line has parasitic inductance that reacts with the decoupling capacitance at some frequency. This causes the impedance, as seen by the IC, to become infinite and experience a rapid phase change. The op amp's phase margin degrades and the circuit may oscillate. To see why it may oscillate, calculate the equivalent parallel impedance of L and C as seen by the IC, Z_{eq}.

$$Z_{eq} = sL(1/sC) / (sL + 1/sC)$$
$$= (1/C)s / (s^2 + 1/LC)$$

where: $s = \pm j(1/\sqrt{LC})$ the complex poles of the equation.

These complex poles cause the equivalent impedance to become infinite at a frequency:

$$f = (1/2\pi)(1/\sqrt{LC})$$

To show this, substitute the value for s into the original equivalent impedance equation:

$$s = j\omega = j(1/\sqrt{LC})$$

$$Z_{eq} = (1/C)s / \{[j(1/\sqrt{LC})]^2 + 1/LC\}$$

$$Z_{eq} = (1/C)s / (-1/LC + 1/LC)$$

The denominator then goes to zero, resulting in an expected infinite impedance. In reality, the impedance does not become infinite because of line losses—the small parasitic resistances that keep the denominator from going to zero. More importantly, the phase almost instantly changes by –180° at the pole frequency.

Like the supply line, the bypass capacitor also has a parasitic inductance, the equivalent series inductance (ESL). ESL creates another resonant frequency in combination with the bypass capacitor. But because they're in series, the capacitance and inductance form two zeros. These complex zeros cause the capacitor's impedance to go to zero and the phase to shift by +180°. A parasitic resistance, the equivalent series resistance (ESR) of the capacitor, dampens this response. Remember that ESR and ESL depend heavily upon the type of capacitor used, and thus should be considered.

A more complete equivalent bypass circuit along with a network analyzer gain-phase plot of an actual LC circuit illustrates the transfer function from the supply (V_P) to the IC's power supply pin (V_{IC}) *(Fig. 4)*. Notice the large gain peak, combined with –180° of phase shift, at the resonant frequency:

$$\text{Freq.} = (1/2\pi)(1/\sqrt{L_P C_C}) = 1 \text{ MHz}$$

when $L_P = 250$ nH and $C_C = 0.1$ μF. The gain dips and the phase shifts by +180° at 16 MHz for $C_C = 0.1$ μF and $L_C = 1$nH.

What does all of this mean to the op amp attached to this supply line? To start, a power supply with high impedance means that any current drawn by the op amp causes significant voltage noise on the power supply line. Secondly, any phase shift in the power supply can also feed into the op amp and cause additional phase shifts on its output. Remember that in a transistor-level analysis of an op amp, the power supplies are assumed to be an ac ground. This is true for most frequencies, but at the resonant frequency (with the bypass capacitor), this ac ground becomes a very high impedance with –180° of phase shift. These phase shifts can affect the output of the op amp, degrade the phase margin, and cause the op amp to oscillate. This effect is reduced somewhat by the power-supply rejection ratio (PSRR), which falls off at higher frequencies. Unfortunately, resonant peaks generally occur at high frequencies where most op amps are unable to reject it.

Wide-bandwidth op amps can easily have such a resonant frequency near their 0-dB (unity-gain) frequency, playing havoc with its gain and phase characteristics, and possibly again causing oscillation. For example, a network analyzer plot was done on the OP-42 with a supply line like that of figure 4. The plot clearly shows the effects on the op amp's output *(Fig. 5)*.

The first gain peak doesn't appear on the output but the gain dip does. This is a function of the drop in PSRR as the frequency rises. The OP-42's PSRR is high enough at 1 MHz, about 45 dB, to reject the first peak. But at 15 MHz, the PSRR has fallen to about 15 dB, allowing the gain dip to feed through to the output. There's a rapid change in both gain and phase at this frequency. In this

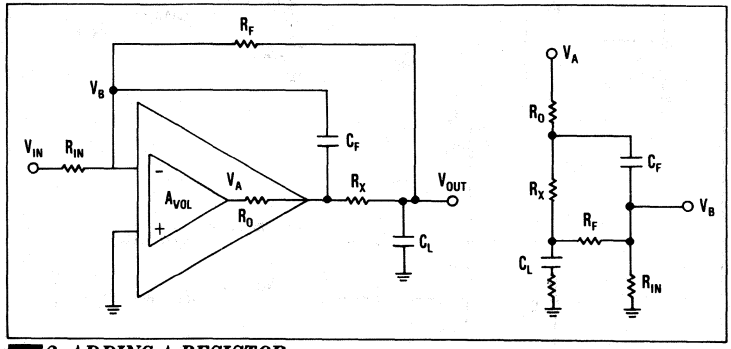

2. ADDING A RESISTOR (R_X) in series with the output and a capacitor (C_F) between the output and input of an op amp can reduce phase shift in the feedback loop to a value close to zero. This restores stability to an op amp which would oscillate when driving a capacitive load (C_L).

example, it happens that the gain dip is due to the complex zero and the phase jumps up, so the op amp maintains its phase margin.

However, if the complex pole had occurred at 15 MHz, the phase would have dipped drastically and could have resulted in –180° of phase shift, causing the op amp to oscillate. Clearly, what happens on its power supply line can severely degrade op-amp performance.

Because the supply line's inductance and the capacitor's ESL appear to be the main causes of circuit resonance, the inductance should be reduced. This is often easier said than done, and the inductance can never be completely eliminated. Furthermore, some inductance may actually be desirable to act as a filter. Consequently, liberal bypass capacitance is needed to move the resonance frequency lower to a point where the PSRR is high enough to reject the gain and phase changes. In addition, the bypass capacitor should be located as close to the IC as possible to minimize trace inductance between it and the IC.

Reducing the parasitic inductance within the capacitor is a matter of selecting the correct capacitor type for the job. Typically, a 10-μF tantalum in parallel with a 0.1-μF ceramic capacitor is specified for supply bypassing, and with good reason. The total bypass capacitance combined with the supply-line inductance sets the gain peak's position. To diminish its effect, the peak should be much lower in frequency than the amplifier's 0-dB frequency. As a result, a large capacitor is needed to move the peak lower in frequency to where the PSRR is high. A 10-μF tantalum fits this mold because large capacitance values are available in a reasonably sized component. In fact, in some applications where it's critical not only to maintain stability but also to ensure gain and phase-flatness out to the amplifier's 0-dB frequency, an even larger bypass capacitor is required to ensure rejection.

Unfortunately, tantalum capacitors aren't perfect because they have high ESR. On the other hand, ceramic capacitors have low ESR. Typical curves of impedance versus frequency for these two capacitors show that the ceramic capacitor has a much sharper dip in impedance (well below 1 Ω) at a much higher frequency due to its low ESR and ESL. However, the tantalum capacitor has much higher ESR, and thus has a shallow dip down to the range of 1-10 Ω. If just the tantalum is used, the impedance of the bypass won't come close to an ideal ac ground at high frequency. Consequently, a lower-value ceramic capacitor is used, in parallel, to further reduce the high-frequency impedance of the supply bypass circuit. The tantalum capacitor reaches its ESR limit around 1 MHz. Above this frequency, the ceramic continues to lower the bypass impedance until its own ESL dominates around 10 MHz.

Typical curves of ceramic capacitors will usually include a few tenths of an inch of lead length. Most of the ESL that causes the impedance to rise sharply above 10 MHz actually comes from the leads. Shortening the leads reduces ESL. This illustrates the importance of placing bypass capacitors close to the IC. Chip capacitors can be valuable for high-speed circuits for this reason. Because they're surface mounted, chip capacitors have almost no lead length except the pc-board trace and the amplifier leads. Minimizing every source of lead length ensures clean supply bypassing at high frequencies.

The discussion so far dealt with only one amplifier connected to the supply, but usually many amplifiers share the same power source. When this is true, the supply line should be as short and wide as possible, and each IC should be bypassed individually. This reduces noise on the supply line, which arises from the op amps' rapidly changing supply-current demands.

Stray capacitance on the input of lower bandwidth op amps, such as the 741, can often be overlooked without significantly impacting circuit performance. However, when dealing with high-speed circuits, this capacitance becomes critical. Not only will the input capacitance cause the closed-loop bandwidth to drop, it can also cause the op amp to oscillate. This capacitance, which comes from both the op amp's input circuitry and the pc-board (or breadboard) layout, can be looked at as a capacitor to ground on the inverting input *(Fig. 6a)*.

To see how this capacitance causes

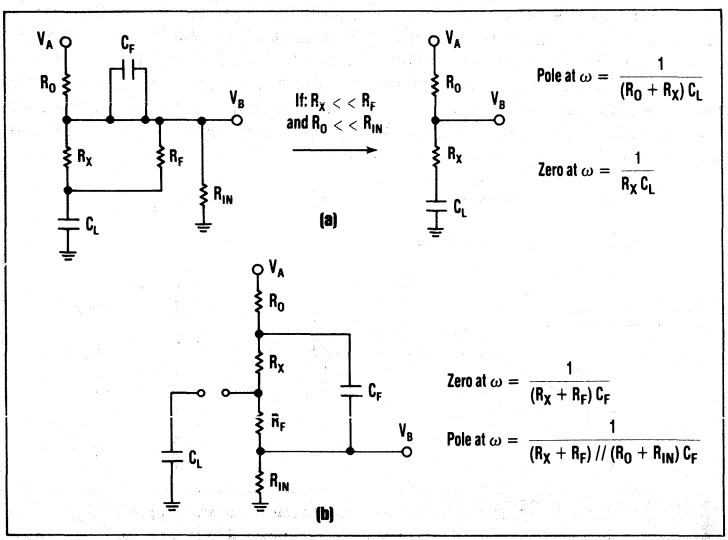

3. TO FIND THE POLE and zero locations caused by the load and feedback capacitors in an op-amp circuit, assume that C_F is a short circuit (a) and C_L is an open circuit (b).

instability, consider the feedback network that includes the amplifier's output impedance *(Fig. 6b)*. The capacitor's phase contribution to the feedback signal is determined by analyzing the circuit's transfer function from V_A to V_B. In other words, the signal is phase shifted by the amplifier's open-loop phase characteristic and then further shifted by the feedback network's phase response. The transfer function is easily determined to be:

$$V_B/V_A = R_1/[(R_1+R_2+R_0) + sC_1R_1(R_0+R_2)]$$

which produces a pole at:

$$fc = (R_1 + R_2 + R_0)/2\pi C_1R_1(R_0 + R_2)$$

$$\approx (R_1 + R_2)/2\pi C_1R_1R_2$$

(Assuming $R_0 << R_1$ and R_2)

This pole causes a phase shift at the unity-gain frequency, f_U, of:

$$\text{Phase Shift} = \text{TAN}^{-1}(f_U/f_C)$$

If this phase shift is large enough, the amplifier may oscillate. As an example, consider the OP-42 with 5 pF of input capacitance and $R_1 = R_2 = 10$ kΩ. This creates a pole at 6.4 MHz, creating a phase shift of 51° at the unity-gain crossover frequency of 8 MHz. Because the OP-42 has a phase margin of 48°, a phase shift of 51° can start the amplifier oscillating.

Fortunately, a feedback capacitor can be added in parallel with the feedback resistor to compensate for the input capacitance. The optimum value for the feedback capacitor is easily calculated by determining the pole and zero locations of the feedback network and setting them equal to each other:

$$\text{Pole} = 1/2\pi(C_1 + C_2)R_1//R_2$$

$$\text{Zero} = 1/2\pi C_2R_2$$

equating and solving these gives:

$$C_2 \geq C_1R_1/R_2$$

Using this value for C_2 provides for zero degrees of phase shift in the feedback network. A shift of 0° ensures stability (assuming the op amp itself is stable). An interesting note is that most resistors have about 1-2 pF of stray capacitance across them, which helps to stabilize the circuit. In the previous example, to compensate for the 5-pF input capacitance of the OP-42, C_2 also needs to be 5 pF.

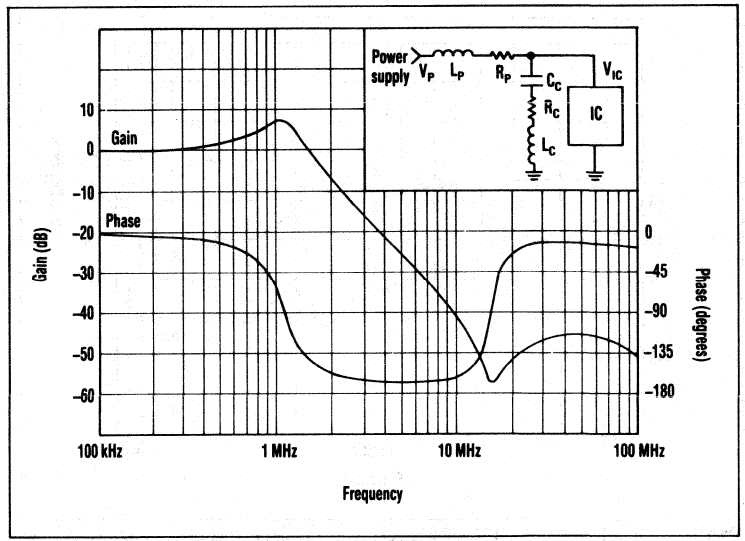

4. DUE TO RLC PARASITICS, the bypassed power supply line of an IC can resonate at several frequencies. In this instance, pronounced gain and phase changes occur at 1 and 16 MHz.

The board layout can be a major source of stray input capacitance. This capacitance arises from the input traces to the summing junction of the op amp.

As a reference point, 0.025 in. of pc-board trace with a ground plane surrounding it, on the opposite side of the board, represents about 10 pF of capacitance per inch. Of course, this value varies depending upon the board thickness and the material its made of. But there can easily be enough input capacitance to cause an amplifier to oscillate.

The effects of the input capacitance can be reduced by moving the pole further out in frequency to reduce its effects near the 0-dB frequency. Most op amps have 3 to 5 pF of input capacitance—a combination of the differential-and common-mode capacitances.

For inverting applications, the two capacitances add. However, for noninverting configurations, the differential capacitance is effectively zero, leaving about 1 to 3 pF of common-mode capacitance on each input. This capacitance is fixed for any given op amp and can't be reduced.

On the other hand, some control of pc-board stray capacitance does exist. There are two key ways to reduce this capacitance. First, keep the input traces as short as possible. Put the feedback resistor and the input source very close to the op amp's input, minimizing pc-board trace length. Keep the analog section close together to further reduce trace lengths. Second, don't place the ground plane near the op amp except where it's needed for the circuit. Be especially careful to keep the ground plane away from the op amp's inputs. The obvious exception to this occurs when the noninverting pin is grounded. When ground really is needed, bring it in with a wide trace to ensure a low resistance ground. Don't locate the ground plane on the opposite side of the board from the analog section. Combining all of these measures will go a long way toward keeping the stray input capacitance to a minimum.

The best possible pc-board layout for high-speed analog circuits would pack the analog parts close together, attenuating trace length. Surface-mounted devices and chip capacitors

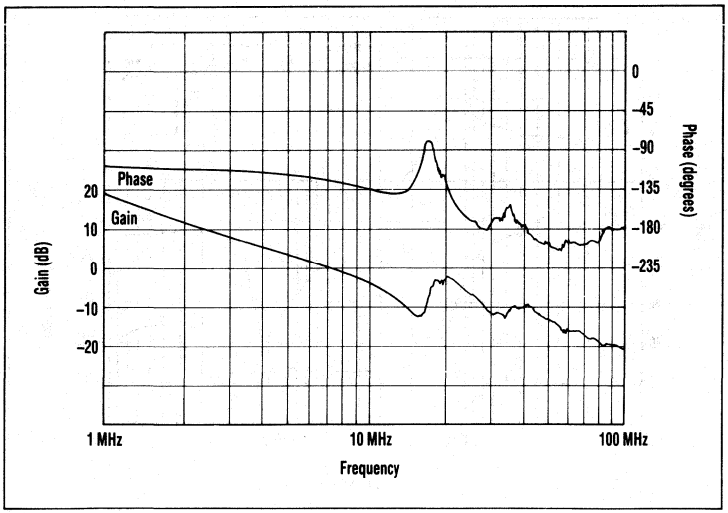

5. THE RESONANCES SEEN by the supply pins of an op amp translate into rapid changes in op-amp gain and phase with an attendant loss in phase margin and increased potential for oscillation. The effect is shown here at 16 MHz.

for power-supply bypassing can really help. The ground plane should be around the perimeter of the analog section and only come in through traces to make contact where ground is required. Try to avoid using sockets to mount the ICs on the board, because they can add another 1 to 3 pF of capacitance to the devices' input pins.

Referring back to the pole equation, it can be seen that minimizing the capacitance isn't the only way to reduce its effects. The parallel combination of the resistors, R_1 and R_2, should also be kept small. The op amp itself determines how small a feedback resistor can be used. Somewhere around 1-2 kΩ is best for most high-frequency amplifiers.

Lowering the resistors by a factor of 10 is the same as reducing the input capacitance by the same factor. However, the output circuit must be able to drive the feedback circuit and the load.

By compensating for the input capacitance with the feedback capacitor, C_2 can stabilize the circuit. However, it reduces bandwidth. Capacitor C_2 forms a pole with the feedback resistor, R_2, which limits the bandwidth to:

$$\text{Bandwidth} = 1/2\pi C_2 R_2$$

Clearly the best way to deal with input capacitance is to minimize both it and the feedback resistance. This reduces the possibility of the circuit oscillating and preserves maximum closed-loop bandwidth.

Trick Those Op Amps

Unfortunately, not all op amps are created stable—at least not at unity gain. Many high-speed op amps are stable only for gains greater than five or even ten. These broadband op amps sacrifice unity-gain stability to achieve a much higher gain-bandwidth product. What happens if a design calls for unity gain and you can't find a unity-gain-stable op amp that fits the application? It's not hopeless—by simply adding a capacitor and resistor across the inputs, almost any amplifier can be made stable at unity gain *(Fig. 7)*. This configuration reduces the feedback factor beta (β) at high frequencies, and the amplifier "thinks" it's running at a gain greater than unity.

To understand this compensation technique, first assume that the capacitor is a short at high frequencies, so all that remains is R_C and R_F. The calculation for the value of R_C is based on the amplifier's minimum stable gain If the amplifier needs a gain of at least 5, make $R_C = R_F/4$ to get $\beta = 1/5$. Because the feedback is the equivalent of the minimum gain needed for stability, the amplifier thinks it's at a closed-loop gain of 5 and is therefore stable. However, the signal sees a closed-loop gain of unity, which can be shown by deriving the expression for the gain:

$$V_O = -A_{OL}(R_C/R_C + R_1)V_O + A_{OL}[V_{IN} - (R_1/R_C + R_1)V_{IN}]$$

where A_{OL} is the open-loop gain of the amplifier.

This expression simplifies to:

$$V_O/V_{IN} = A_{OL}\beta / (1 + A_1\beta) = 1/(1/A_{OL}\beta + 1)$$

where

$$\beta = R_C/(R_C + R_1)$$

When the open-loop gain is large, $V_O/V_{IN} = 1$. Unity gain in the signal path is maintained even though the amplifier thinks it is at a closed-loop gain of five.

Next consider the value of the compensating capacitor, C_C. It must be large enough to ensure that the amplifier satisfies the β requirement at a low enough frequency to ensure stability. A minimum value for C_C should provide an impedance equal to that of R_C at a frequency at least a decade below the corner frequency for the amplifier's lowest stable gain:

$$C_C = 1/2\pi R_C(f_C/10)$$

For example, consider the OP-64, a high slew-rate op amp. Its lowest stable gain (5) yields a corner frequency at 16 MHz. A feedback resistor of 1000 Ω results in an R_C of 250 Ω and a C_C of 398 pF—the minimum value that should be considered for C_C.

This equation holds true for operation in a noninverting configuration. Now consider an inverting circuit, whose analysis is very similar. The closed-loop gain equation becomes:

$$V_O/V_{IN} = -1/(1 + 1/A_{OL}\beta)$$

where

$$\beta = (R_C//R_1) / (R_C//R_1 + R_2)$$

This expression is similar to the noninverting case except for the sign and the value of β. The expression for β indicates that the input resistor, R_1, is parallel with R_C at high frequencies. This parallel combination calculates the value of R_C for minimum stable gain. The capacitor's value is calculated the same as for the noninverting case.

WHAT'S THE CATCH?

An op amp's bandwidth and settling time can be affected by lead-lag compensation. Don't make the mistake of thinking that because the amplifier is in a unity-gain configuration that its signal bandwidth equals the amplifier's full gain-bandwidth product. Using lead-lag compensation doesn't increase the bandwidth above that at the minimum stable gain. For example, the OP-64 has a gain-bandwidth product of 80 MHz but it's stable only for closed-loop gains of five or more. Therefore, its bandwidth is 16 MHz for a gain of five. When compensated for unity-gain operation, its bandwidth is still 16 MHz, as is seen in the results of the circuit's Spice analysis *(Fig. 7, again)*. To understand this effect, look at the closed loop gain expression for $\beta = 1/5$:

$$V_O/V_{IN} = 1/(5/A_{OL} + 1)$$

and compare it to a typical unity gain expression:

$$V_O/V_{IN} = 1/(1/A_{OL} + 1)$$

Due to β, the compensated amplifier's gain bandwidth is one-fifth that of the uncompensated OP-64.

Lead-lag compensation can also affect an op amp's settling time. This can be illustrated by examining the transient response of the OP-64 *(Figs. 8a and 8b)*. These tests were performed with $C_C = 470$ pF and $R_C = 250\ \Omega$. The circuit's settling time to 0.1% increases to 600 ns in contrast to the $A_V = 5$ settling time of 390 ns. However, this effect is easily overcome by increasing the value of C_C. Another test was performed with $C_C = 1\ \mu F$ and the results prove to be much better *(Fig. 8c)*. The settling time in this instance drops to 310 ns—comparable to the $A_V = 5$ settling time. Using a smaller capacitor while ensuring stability still causes significant overshoot, resulting in a long settling time.

Changing the input voltage from –1 V to +1 V almost instantly creates a large differential voltage between the op amp's inverting and noninverting inputs. The smaller capacitor, combined with the compensating resistor, has a relatively small RC time constant. It charges quickly towards the full differential voltage. When the amplifier overshoots while trying to return to the final output voltage (the input voltage times the closed-loop gain, the capacitor is charged to enough differential voltage to continue forcing the output high. The long decay in the overshoot is caused by the slow discharge of the compensation capacitor. When the larger capacitor is used, the RC time constant becomes very large and the capacitor can only charge to a fraction of the differential voltage. Consequently, when the amplifier overshoots, the capacitor isn't charged to enough of the differential voltage to continue to force the output high. As a result, it quickly returns to the final voltage.

While overall settling time varies considerably, the slew rate (and thus the full power bandwidth) is constant. Regardless of the capacitor chosen, the slew rate is the same as that for a minimum-stable-gain condition. The amplifier's slew rate is a function of its internal structure—it's independent of the compensating resistor and capacitor.

This discussion may raise such questions as: "How big a capacitor

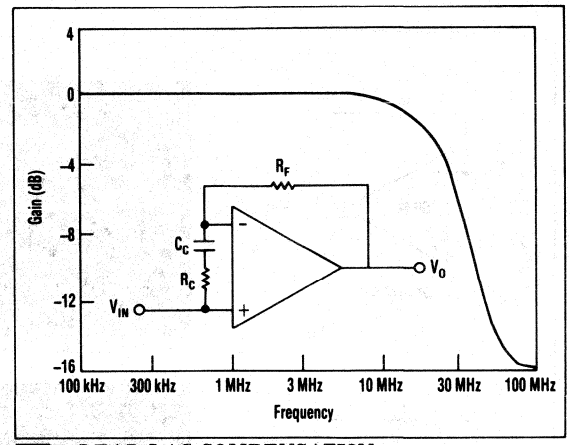

7. LEAD-LAG COMPENSATION formed by capacitor C_C and resistor R_C can stabilize at unity gain. Typically, this op amp is stable only at a higher gain.

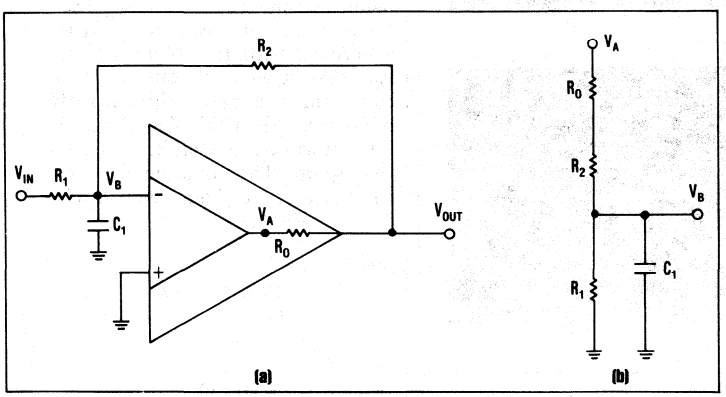

6. PARASITIC CAPACITANCE on an op amp's input (C_1) reduces phase margin and can cause oscillation (a). Its effect can be analyzed by considering C_1 as part of the op amp's feedback network (b).

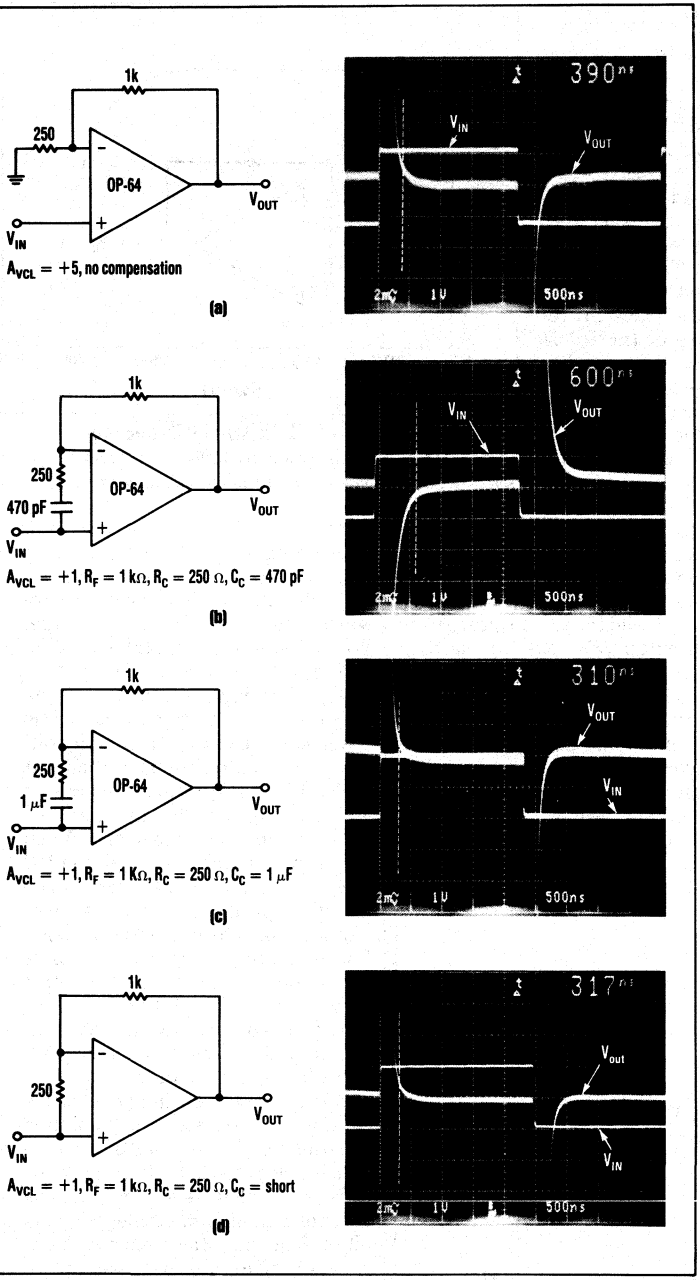

8. WHEN AN OP AMP with a minimum stable gain of 5 is lead-lag compensated to run at a gain of one, its original settling time of about 400 ns (a) increases to about 600 ns (b). However, increasing the size of the compensation capacitor from 470 pF to 1 μF drops the settling time to about 300 ns (c). Shorting the capacitor shows similar results (d).

should be used? And if a bigger capacitor is better, then why use one at all?" To answer the first question experimentally, a capacitor that's roughly 1000 times the calculated value usually works to keep the overshoot down and the settling time to a minimum. The settling time with just the resistor in place was also measured *(Fig. 8d)*. It is 317 ns, almost the same as when using C_C = 1 μF. However, dc errors will creep in if the capacitor isn't used. With the capacitor missing, the dc noise gain of this circuit is now also five. Therefore, any dc errors, such as voltage offset, are increased by a factor of five at the output. Because high-speed amplifiers typically have relatively large offsets, dc errors can be significant and certainly need to be considered.

More Noise Gain

High-frequency or ac noise gain must also be considered. Because the amplifier's feedback is now equivalent to a gain of 5, its ac noise is boosted by a factor of five at the output. Voltage noise is typically modelled as a noise generator on the amplifier's inverting input. A Spice analysis was performed to measure the noise gain by placing the input source on the inverting input and measuring the amplitude of the output for two different compensation capacitors, 470 pF and 1 μF *(Fig. 9)*. At unity gain, the noise gain would usually be one, but as the plot shows, the lead-lag compensation increases it to five (14 dB). By looking at these two plots, it seems that using a 1-μF capacitor generates significantly more noise than the 470-pF capacitor. However, the difference isn't that great. The equation for total rms noise over a given bandwidth is:

$$E_n = e_n \sqrt{(f_H - f_L)}$$

where:

e_n = the spectral voltage noise density

f_H = upper frequency limit

f_L = lower frequency limit

Examining the plot, both capacitors appear to have the same f_H—20

MHz—but the f_L is different. It's about 500 kHz for the 470-pF capacitor and about 500 Hz for the 1-μF capacitor. However, this only creates a difference of 1.3% for the total noise. Therefore, the noise trade-off with the 1-μF capacitor is minimal when compared to the difference in settling time resulting from the two capacitors. Lead-lag compensation is a valuable tool in dealing with high-frequency op amps. And, as the aforementioned discussion shows, it can offer stability for op amps that typically would be unstable at low gains. However, there are trade-offs, so care must be exercised when using this compensation scheme.□

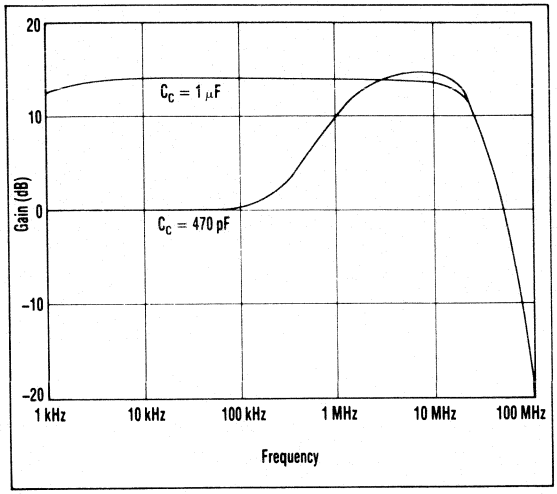

9. UPPING THE VALUE of a lead-lag compensation capacitor from 470 pF to 1 μF raises the low-frequency noise gain by about 10 dB. Over a given bandwidth, however, the total rms noise is only raised by a small percentage.

Reference:

1. Precision Monolithics Inc., 1988 Analog Applications Seminar.

For a complete discussion of the op-amp stability criterion:

Gary, P., and Meyer, G. *Analog Integrated Circuits.* (New York: Wiley, 1984), pp. 527-70.

ANALOG DEVICES

AN-110
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

OP-64 Advanced SPICE Macro-Model

by Joe Buxton

INTRODUCTION

This application note describes the SPICE macro-model for the OP-64 high speed, wideband operational amplifier. This model was tested with and is compatible with PSpice* and HSpice**. The schematic and net-list are included here so that the model can easily be used. This model can accommodate multiple frequency poles and multiple zeroes, which is an advanced concept that results in more accurate AC and transient responses necessary for simulating the behavior of today's high speed op amps. For example, the OP-64 has 9 poles and 2 zeroes, which this advanced model can easily accommodate.

The OP-64 macro-model uses almost entirely linear elements to model the op amp's behavior. With the exception of two NPN transistors for the input stage, the rest of the model uses voltage-controlled current sources. Simulation of the OP-64's AC behavior is achieved by using multiple poles and zeroes modelled simply by RL and RC circuits. Doing so not only saves computer time but also makes the development and calculations easy.

MODEL DESCRIPTION

The schematic (Figure 1) and net-list (Figure 2) describe the complete OP-64 model. This model breaks up the OP-64 into many distinct stages as described below.

INPUT STAGE

The input stage (Figure 1a) uses two NPN transistors, Q_1 and Q_2, to model the actual OP-64 input. Their forward current gain, i.e. BF, is chosen such that the base current matches the OP-64's input bias current specification. Because these transistors are modelled as an ideal and perfectly matched differential pair, they do not take into account non-ideal parameters such as V_{OS}, I_{OS}, and C_{IN}. Consequently, individual circuit elements are added to model these error sources. Lastly, C_2 is added to create one of the poles in the model.

GAIN STAGE

The open loop gain of the OP-64 is achieved entirely in the gain stage (Figure 1b), and all other stages have unity gain. The two voltage-controlled current sources, G_1 and G_2, have scaled transconductances that, when combined with R_7 and R_8, give an open loop gain of 15,000. This stage also uses C_3 and C_4 to create the dominant pole at 3.8kHz and to model the slew rate. Lastly, the diodes, D_1 and D_2, and voltage sources V_2 and V_3, are used to clamp the voltage of node 11 below the power supplies. Because the next stage (Figure 1c) has unity gain and its voltage-controlled current sources are controlled by node 11, it too is clamped, at node 12, below the power supplies. The same is true for the subsequent stages, including the output at node 38, such that their voltages are clamped below the power supplies.

POLE-ZERO STAGES

All of the pole-zero stages (Figure 1c-h, j) have unity gain. RC and RL networks are used to model the secondary poles and zeroes of the OP-64.

COMMON-MODE STAGE

The common-mode voltage that is used to create the CM error is created by the two input resistors, R_1 and R_2. This voltage is referenced by G_{17} and G_{18} in the common-mode stage (Figure 1i). The transconductances of these two sources are scaled such that, in combination with R_{29} and R_{30}, the V_{CM} is attenuated by the CMRR of 100dB. This error voltage is then inserted as an offset voltage in the input stage by E_{OS}. Also, the inductors, L_5 and L_6, model the pole in the CMR vs. frequency response of the OP-64.

OUTPUT STAGE

The output stage (Figure 1k) is modelled as an ideal output with an output resistance, R_{36} in parallel with R_{37}. An inductor, L_7, is added to model the rising output impedance of the emitter-follower output stage with frequency. The voltage sources V_6 and V_7 and diodes D_5 and D_6, limit the voltage across the resistors thus limiting the output current, I_{OUT} to 80mA.

PARAMETERS THAT ARE NOT MODELLED

To keep the OP-64 model as simple as possible and thus save computer time, not all features of the op amp were modelled as listed below:

- PSRR
- No Limits on Power Supply Voltages
- Maximum Input Voltage Beyond the Power Supplies
- Temperature Effects (i.e. Model Parameters Are Assumed at 25°C)
- Input Noise Voltage and Current Sources
- Parameter Variations for Monte Carlo Analysis (i.e. All Parameters Are Typical Only)

These parameters are not included because they aren't neces-

* PSpice is a registered trademark of MicroSim Corporation.
** HSpice is a tradename of Meta-Software, Inc.

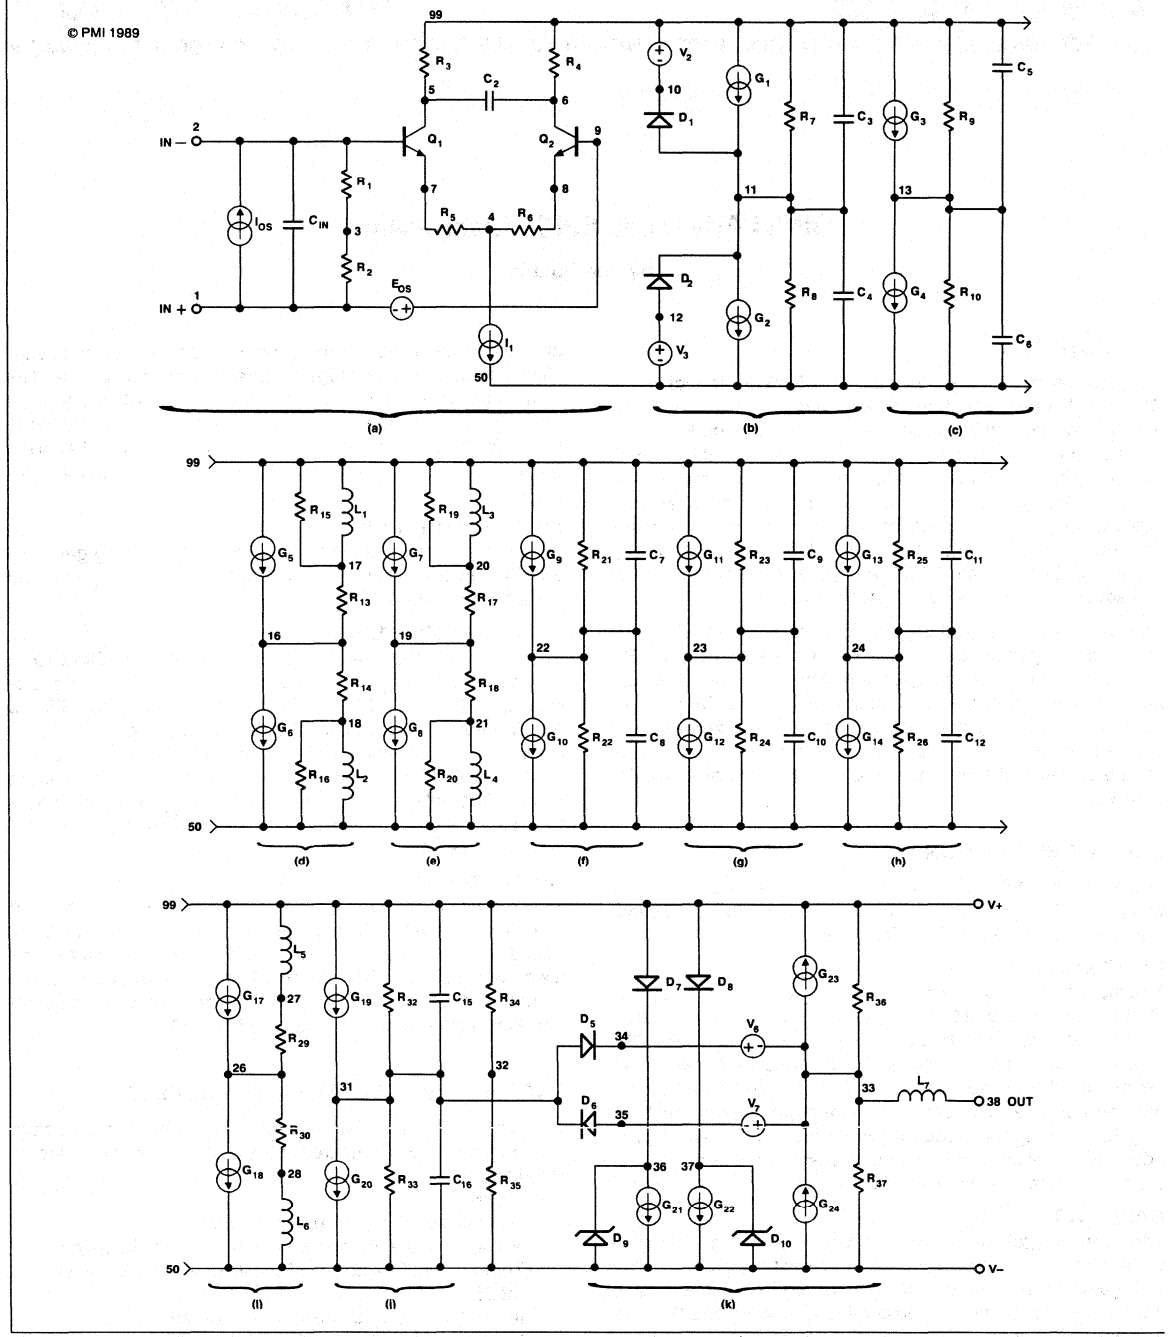

FIGURE 1: *OP-64 SPICE Macro-model Schematic and Node List*

```
OP-64 MACRO-MODEL © PMI 1989
*
• subckt OP-64 1  2 38  99  50
*
INPUT STAGE & POLE AT 39.8 MHz
*
r1   2   3        5E11
r2   1   3        5E11
r3   5   99       474.86
r4   6   99       474.86
r5   4   7        423.26
r6   4   8        423.26
cin  1   2        5E-12
c2   5   6        4.2106E-12
i1   4   50       1E-3
ios  1   2        1E-7
eos  9   1        poly(1)  26  32  4E-4  1
q1   5   2   7    qx
q2   6   9   8    qx
*
* SECOND STAGE & POLE AT 3.8 kHz
*
r7   11  99       7.1229E6
r8   11  50       7.1229E6
c3   11  99       5.88E-12
c4   11  50       5.88E-12
g1   99  11       poly(1)  5  6  4.31E-3  2.1059E-3
g2   11  50       poly(1)  6  5  4.31E-3  2.1059E-3
v2   99  10       2.25
v3   12  50       2.25
d1   11  10       dx
d2   12  11       dx
*
* POLE AT 39.8 MHz
*
r9   13  99       1E6
r10  13  50       1E6
c5   13  99       4E-15
c6   13  50       4E-15
g3   99  13       11  32  1E-6
g4   13  50       32  11  1E-6
*
* ZERO-POLE PAIR AT 26.5 MHz /159 MHz
*
r13  16  17       1E6
r14  16  18       1E6
r15  17  99       5E6
r16  18  50       5E6
l1   17  99       5.005E-3
l2   18  50       5.005E-3
g5   99  16       13  32  1E-6
g6   16  50       32  13  1E-6
*
* ZERO-POLE PAIR AT 31.8 MHz / 39.8 MHz
*
r17  19  20       1E6
r18  19  21       1E6
r19  20  99       2.5157E5
r20  21  50       2.5157E5
l3   20  99       1.006E-3
l4   21  50       1.006E-3
g7   99  19       16  32  1E-6
g8   19  50       32  16  1E-6
*
* POLE AT 100 MHz
*
r21  22  99       1E6
r22  22  50       1E6
c7   22  99       1.59E-15
c8   22  50       1.59E-15
g9   99  22       19  32  1E-6
g10  22  50       32  19  1E-6
*
*POLE AT 159 MHz
*
r23  23  99       1E6
r24  23  50       1E6
c9   23  99       1E-15
c10  23  50       1E-15
g11  99  23       22  32  1E-6
g12  23  50       32  22  1E-6
*
*POLE AT 159 MHz
*
r25  24  99       1E6
r26  24  50       1E6
c11  24  99       1E-15
c12  24  50       1E-15
g13  99  24       23  32  1E-6
g14  24  50       32  23  1E-6
*
*COMMON-MODE GAIN NETWORK WITH ZERO AT 20kHz
*
r29  26  27       1E6
r30  26  28       1E6
l5   27  99       7.9575
l6   28  50       7.9575
g17  99  26       3   32  1E-11
g18  26  50       32  3   1E-11
*
* POLE AT 159 MHz
*
r32  31  99       1E6
r33  31  50       1E6
c15  31  99       1E-15
c16  31  50       1E-15
g19  99  31       24  32  1E-6
g20  31  50       32  24  1E-6
*
* OUTPUT STAGE
*
r34  32  99       20.0E3
r35  32  50       20.0E3
r36  33  99       60
r37  33  50       60
l7   33  38       2.9E-7
g21  36  50       31  33   16.6666667E-3
g22  37  50       33  31   16.6666667E-3
g23  33  99       99  31   16.6666667E-3
g24  50  33       31  50   16.6666667E-3
v6   34  33       1.7
v7   33  35       1.7
d5   31  34       dx
d6   35  31       dx
d7   99  36       dx
d8   99  37       dx
d9   50  36       dy
d10  50  37       dy
*
* MODELS USED
*
•model qx NPN(BF=2500)
•model dx    D(IS=1E-15)
•model dy    D(IS=1E-15 BV=50)
•ends OP-64
```

FIGURE 2: *OP-64 SPICE Net-List*

13

sary to analyze the op amp performance under normal conditions. Furthermore, the chip's shutdown function, $\overline{\text{DISABLE}}$, and its associated currents, $I_{SY\ DIS}$ and I_{DIS}, are not modelled because they are not part of the normal operation of the OP-64 as an op amp.

MODEL PERFORMANCE

The graphs on these next two pages compare the actual responses of the OP-64 to the SPICE analysis. For example, the actual open-loop gain-phase response (Figure 3a) measured on a network analyzer shows very close agreement to the SPICE simulation (Figure 3b). This agreement holds true for both the gain and phase curves well above 10MHz. The same agreement holds true for a closed-loop gain of –10 (Figures 4a, b) and a closed-loop gain of –100 (Figures 5a, b) showing the accuracy of the model's frequency response.

Because of its ability to accommodate multiple poles and zeroes, the model's transient response (Figure 6b) shows remarkable resemblance to the actual response (Figure 6a), which was measured at a closed-loop gain of 5 with an 115pF capacitive load. Notice that the two figures show a slight difference in the damped settling characteristics occurring in the positive transition. This is due to the fact that the model's output stage is assumed to be symmetrical, unlike the OP-64's output stage which in reality has a slight asymmetrical behavior. Nevertheless, the model does predict, to a reasonable degree of accuracy, the damped oscillation. This accurate transient response, along with the accurate AC response, makes the OP-64 macro-model a powerful tool in analyzing circuit stability and AC behavior under various load conditions.

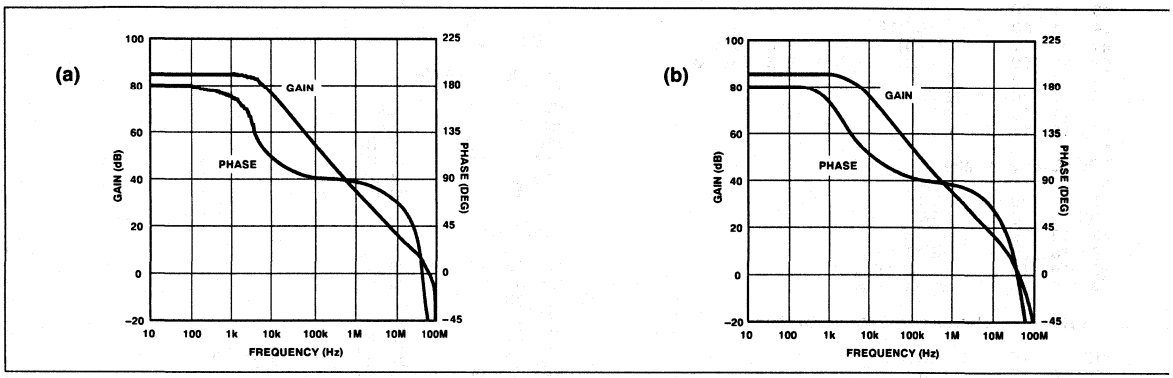

FIGURE 3: *Gain-Phase Plots, Open Loop Gain (a. Actual Response, b. Simulated Model Response)*

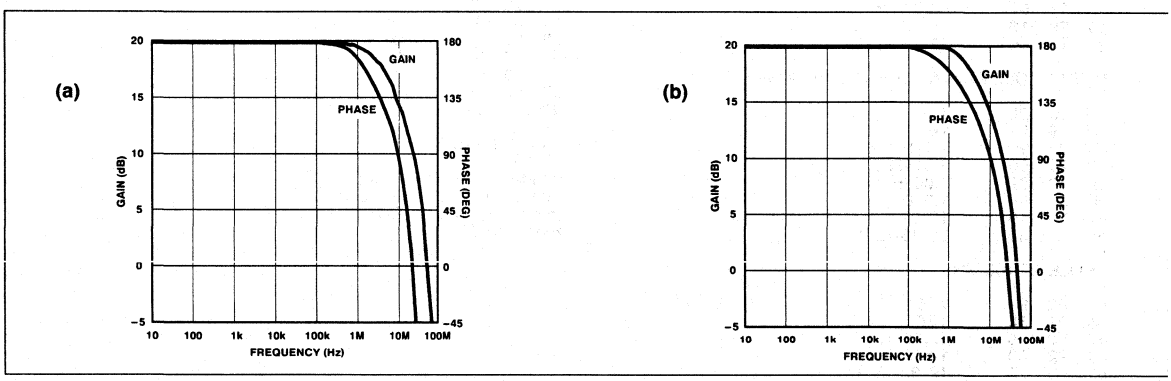

FIGURE 4: *Gain-Phase Plots, Closed Loop Gain of –10 (a. Actual Response, b. Simulated Model Response)*

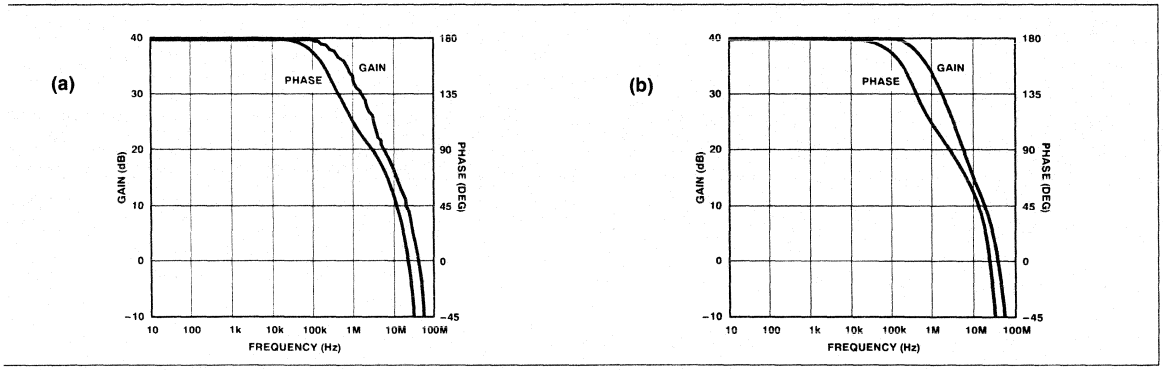

FIGURE 5: *Gain-Phase Plots, Closed Loop Gain of –100 (a. Actual Response, b. Simulated Model Response)*

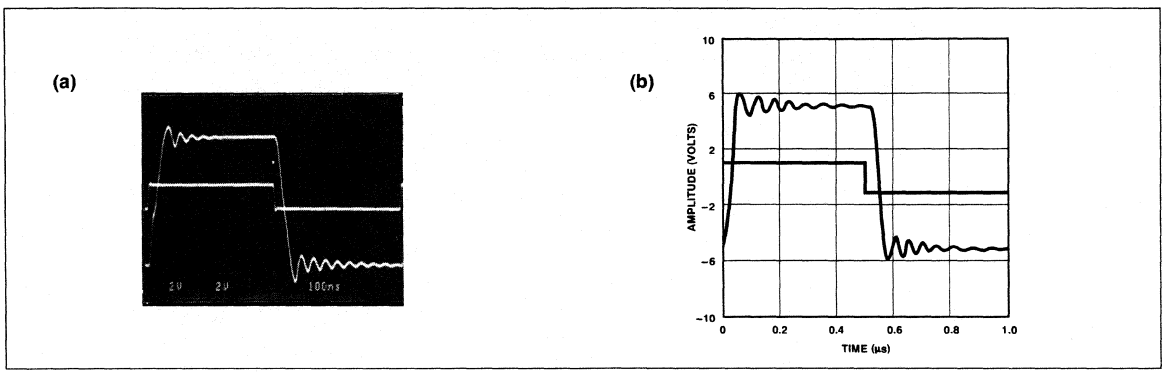

FIGURE 6: *Transient Response, 115pF Capacitive Load, Gain = 5 (a. Actual Response, b. Simulated Model Response)*

13

ANALOG DEVICES

AN-117
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

OP-42 Advanced SPICE Macro-Model

by Joe Buxton

INTRODUCTION

This application note describes the SPICE macro-model for the OP-42 high-speed, fast-settling precision operational amplifier. This model was tested with, and is compatible with PSpice* and HSPICE**. The schematic and net-list are included here so that the model can easily be used. This model can accommodate multiple frequency poles and multiple zeroes, which is an advanced concept that results in more accurate AC and transient responses necessary for simulating the behavior of today's high-speed op amps. For example, 8 poles and 2 zeroes are required to sufficiently simulate the OP-42, which this advanced model can easily accommodate.

Throughout the OP-42 macro-model, RC networks produce the multiple poles and zeroes which simulate the OP-42's AC behavior. Each stage contains a pole or a pole-zero pair. The stages are separated from each other by voltage-controlled current sources so that the pole and zero locations do not interact. The only nonlinear elements in the entire model are two p-channel JFETs which comprise the input stage. Limiting the model to almost entirely linear circuit elements significantly reduces simulation time and simplifies model development.

MODEL DESCRIPTION

The schematic (Figure 1) and net-list (Figure 2) describe the complete OP-42 model. This model breaks up the OP-42 into many distinct stages as described below:

INPUT STAGE

To correctly model the behavior of the OP-42, the model uses a differential pair of p-channel JFETs biased with a 1mA current source (Figure 1a). To keep this stage as simple as possible only those parameters necessary to the JFET model are specified, that is the threshold voltage VTO, the transconductance coefficient BETA, and the gate p-n saturation current IS, which is scaled to give the proper input bias current. All other JFET parameters are left at the model default values, most of which are zero.

As for non-ideal behaviors of the input stage, such as V_{OS}, I_{OS}, and C_{IN}, these are modelled with external circuit elements. For example, no gate capacitance is specified for the JFET model, therefore a capacitor, C_{IN}, is added across the inputs. Furthermore, since in this model the input JFETs are perfectly matched, V_{OS} and I_{OS} error sources are added using an external voltage source and current source, respectively. Lastly, the drain resistors R_3 and R_4 are chosen to be 1/gm of the JFETs to give a gain of unity in the input stage. C_2 is added to create one of the secondary poles in the model.

GAIN STAGE

The open-loop gain of the OP-42 is achieved entirely in the gain stage (Figure 1b), and all other stages have unity gain. The two voltage-controlled current sources, G_1 and G_2, have scaled transconductances that, when combined with R_5 and R_6, yield an open-loop gain of 250,000. This stage also uses C_3 and C_4 to create the dominant pole at 45Hz and to model the amplifier slew rate. Lastly, the diodes, D_1 and D_2, and voltage sources V_2 and V_3, are necessary to clamp the voltage of node 9 below the power supplies. Because the next stage (Figure 1c) has unity gain and its voltage-controlled current sources are controlled by node 9, it too is clamped, at node 11, below the power supplies. The same is true for the subsequent stages, including the output at node 32, such that their voltages are clamped below the power supplies.

POLE-ZERO STAGES

All of the pole-zero stages (Figures 1c-g, i) have unity gain, which is a result of the gm of the voltage-controlled current sources being the reciprocal of the resistors. The RC networks are used to model the secondary poles and zeroes of the OP-42.

COMMON-MODE STAGE

The common-mode voltage that is used to create the CM error is created by the two input resistors, R_1 and R_2. This voltage is referenced by G_{13} and G_{14} in the common-mode stage (Figure 1h). The transconductances of these two sources are scaled such that, in combination with R_{21} and R_{22}, the V_{CM} is attenuated by the CMRR of 96dB. This error voltage is then inserted as an offset voltage in the input stage by E_{OS}. The inductors, L_1 and L_2, mimic the pole in the CMR vs. frequency response of the OP-42.

OUTPUT STAGE

The output stage (Figure 1j) is modelled as an ideal output with an output resistance, R_{28} in parallel with R_{29}. An inductor, L_3, is added to model the rising output impedance of the emitter-follower output stage with frequency. The voltage sources V_6 and V_7, and diodes D_5 and D_6, combine to limit the voltage across the resistors, thus limiting the output current, I_{OUT} to 30mA.

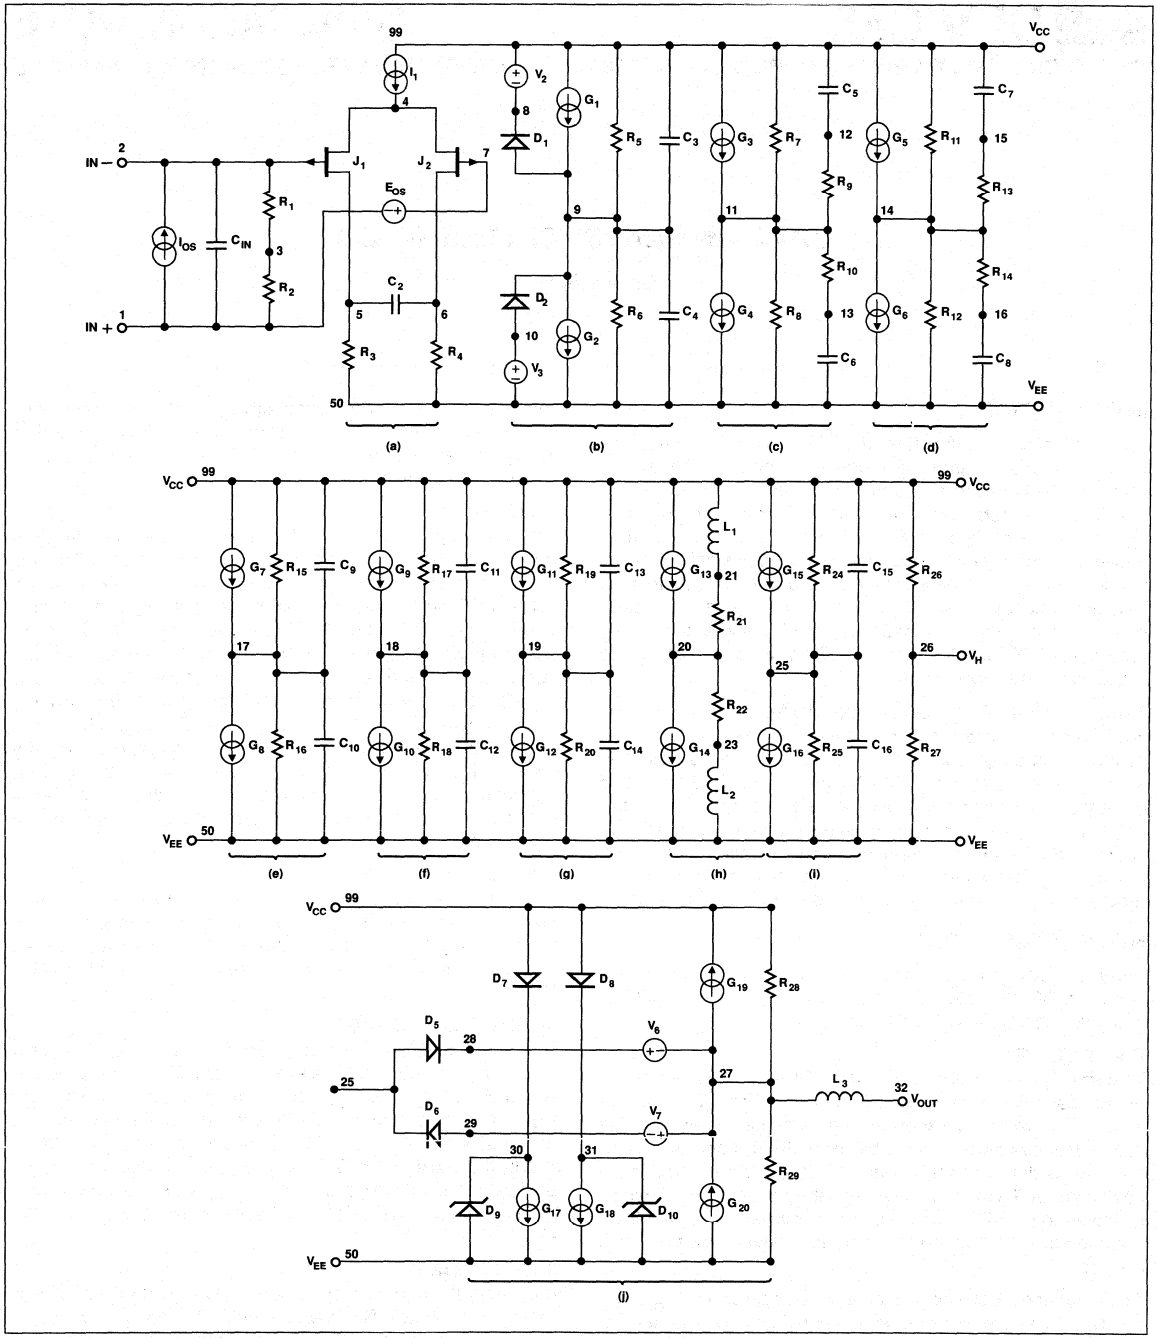

FIGURE 1: *OP-42 SPICE Macro-Model Schematic and Node List*

```
OP-42 MACRO-MODEL © PMI 1989
*
* SUBCKT OP-42 1  2 99 50 32
*
* INPUT STAGE & POLE AT 15.9 MHZ
*
R1    2   3        5E11
R2    1   3        5E11
R3    5   50       707.36
R4    6   50       707.36
CIN   1   2        5E-12
C2    5   6        7.08E-12
I1    99  4        1E-3
IOS   1   2        2E-12
EOS   7   1        POLY(1) 20 26 0.3E-3 1
J1    5   2   4    JX
J2    6   7   4    JX
*
* SECOND STAGE & POLE AT 45 HZ
*
R5    9   99       176.84E6
R6    9   50       176.84E6
C3    9   99       20E-12
C4    9   50       20E-12
G1    99  9        POLY(1) 5 6 3.96E-3 1.4137E-3
G2    9   50       POLY(1) 6 5 3.96E-3 1.4137E-3
V2    99  8        2.5
V3    10  50       3.1
D1    9   8        DX
D2    10  9        DX
*
* POLE-ZERO PAIR AT 1.80 MHZ/2.20 MHZ
*
R7    11  99       1E6
R8    11  50       1E6
R9    11  12       4.5E6
R10   11  13       4.5E6
C5    12  99       16.1E-15
C6    13  50       16.1E-15
G3    99  11       9  26  1E-6
G4    11  50       26  9  1E-6
*
* POLE-ZERO PAIR AT 1.80 MHZ/2.20 MHZ
*
R11   14  99       1E6
R12   14  50       1E6
R13   14  15       4.5E6
R14   14  16       4.5E6
C7    15  99       16.1E-15
C8    16  50       16.1E-15
G5    99  14       11 26 1E-6
G6    14  50       26 11 1E-6
*
* POLE AT 53 MHZ
*
R15   17  99       1E6
R16   17  50       1E6
C9    17  99       3E-15
C10   17  50       3E-15
G7    99  17       14 26 1E-6
G8    17  50       26 14 1E-6
*
* POLE  AT 53 MHZ
*
R17   18  99       1E6
R18   18  50       1E6
C11   18  99       3E-15
C12   18  50       3E-15
G9    99  18       17 26 1E-6
G10   18  50       26 17 1E-6
*
* POLE AT 53 MHZ
*
R19   19  99       1E6
R20   19  50       1E6
C13   19  99       3E-15
C14   19  50       3E-15
G11   99  19       18 26 1E-6
G12   19  50       26 18 1E-6
*
*COMMON-MODE GAIN NETWORK WITH ZERO AT 100 KHZ
*
R21   20  21       1E6
R22   20  23       1E6
L1    21  99       1.5915
L2    23  50       1.5915
G13   99  20       3  26 1.58E-11
G14   20  50       26  3 1.58E-11
*
*POLE AT 79.6 MHZ
*
R24   25  99       1E6
R25   25  50       1E6
C15   25  99       2E-15
C16   25  50       2E-15
G15   99  25       19 26 1E-6
G16   25  50       26 19 1E-6
*
*OUTPUT STAGE
*
R26   26  99       111.1E3
R27   26  50       111.1E3
R28   27  99       90
R29   27  50       90
L3    27  32       2.5E-7
G17   30  50       25 27 11.1111E-3
G18   31  50       27 25 11.1111E-3
G19   27  99       99 25 11.1111E-3
G20   50  27       25 50 11.1111E-3
V6    28  27       0.7
V7    27  29       0.7
D5    25  28       DX
D6    29  25       DX
D7    99  30       DX
D8    99  31       DX
D9    50  30       DY
D10   50  31       DY
*
* MODELS USED
*
.MODEL JX PJF(BETA=999.3E-6 VTO=-2.000 IS=8E-11)
.MODEL DX   D(IS=1E-15)
.MODEL DY   D(IS=1E-15 BV=50)
.ENDS OP-42
```

FIGURE 2: *OP-42 SPICE Net-List*

PARAMETERS THAT ARE NOT MODELLED

To keep the OP-42 model as simple as possible and thus save computer time, not all features of the op amp were modelled as listed below:

- PSRR
- No limits on power supply voltages
- Maximum input voltage beyond the power supplies
- Temperature Effects (i.e., model parameters are assumed at 25°C)
- Input Noise Voltage and Current Sources
- Parameter Variations for Monte Carlo Analysis (i.e., all parameters are typical only)

These parameters are considered second-order effects and are not considered necessary for circuit simulation under normal operating conditions. However, users can easily add these functions as needed.

MODEL PERFORMANCE

The graphs on these next two pages compare the actual responses of the OP-42 to the SPICE analysis under various load and gain conditions. To show the accuracy of the model's frequency response it was tested with different gains. Shown below are the gain-phase plots for an open-loop gain (Figure 3) and a closed-loop gain of –1 (Figure 4). The actual responses (part "a" of each figure), which were measured using a network analyzer, show very close agreement to the SPICE simulations (part "b" of each figure). This agreement holds true for both the gain and phase curves well above 10MHz, which demonstrates the accuracy of the model's frequency response.

The model can also accurately simulate the OP-42's transient response, which is stable for a closed-loop gain of –1. The comparison was done with a 430pF capacitive load to create the damped oscillations for both a small signal and a large signal response. Because of the model's complete complement of poles and zeroes, its large signal transient response (Figure 5b) shows close resemblance to the actual response (Figure 5a). Also notice that the slew rate limit of 58V/µs is accurately modelled. The small signal response (Figure 6) again shows the same accuracy. In both responses the actual device shows a slight difference in the damped oscillations between its negative and positive transition, which is due to the inherent asymmetrical behavior of the OP-42's output stage. On the other hand, the model assumes a perfectly symmetrical output stage so its responses do not exhibit this difference. Nevertheless, as can be seen from the figures, the model agrees closely with the actual device's performance. This accurate transient response, along with the accurate AC response, makes the OP-42 macro-model a powerful tool in analyzing circuit stability and AC behavior under various load conditions.

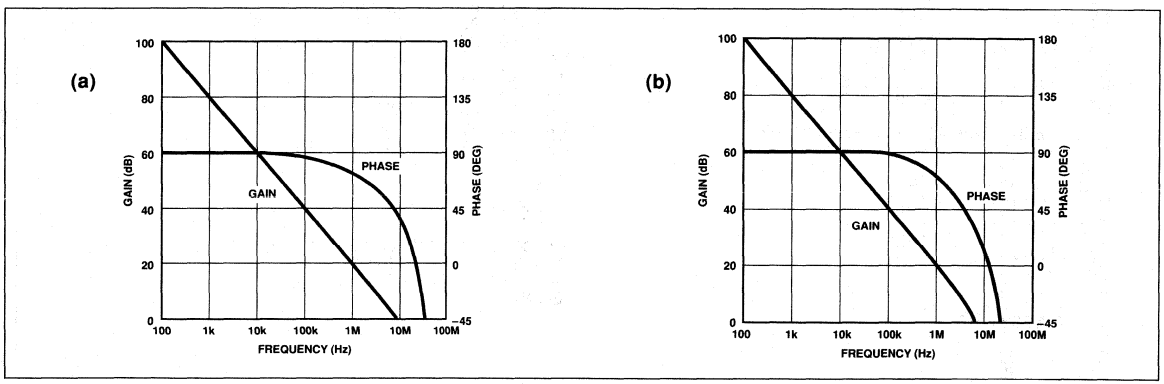

FIGURE 3: *Gain-Phase Plots, Open-Loop Gain (a. Actual Response, b. Simulated Model Response)*

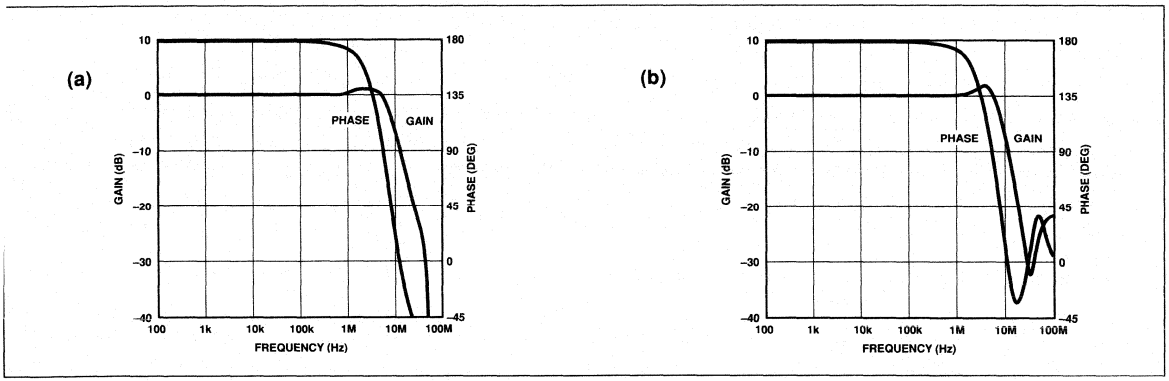

FIGURE 4: *Gain-Phase Plots, Closed-Loop Gain of –1 (a. Actual Response, b. Simulated Model Response)*

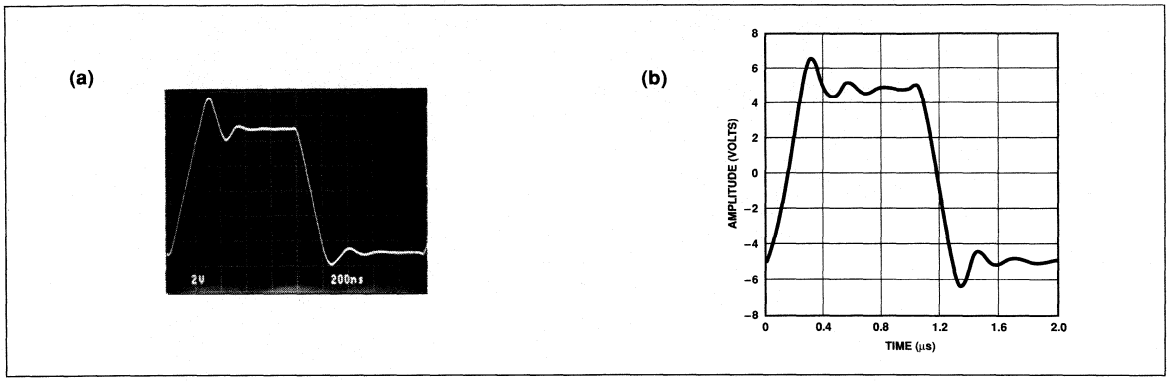

FIGURE 5: *Large-Signal Transient Response (a. Actual, b. Model)*

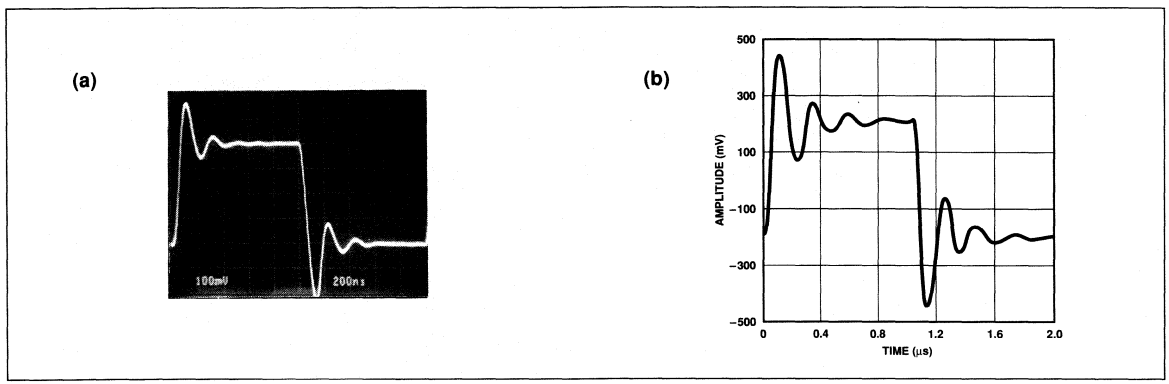

FIGURE 6: *Small-Signal Transient Response (a. Actual, b. Model)*

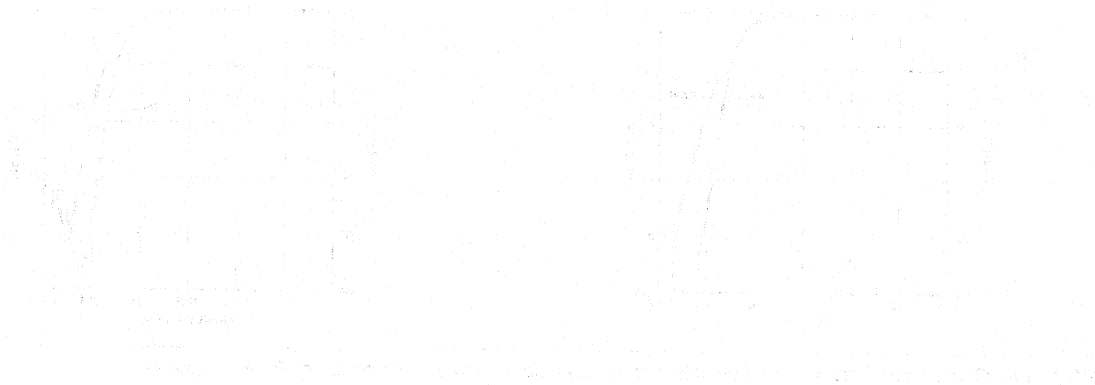

ANALOG DEVICES

AN-120
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

OP-400 SPICE Macro-Model

by Joe Buxton

INTRODUCTION

This application note describes the SPICE macro-model for the OP-400 quad low-offset, low-power operational amplifier. This model was tested with, and is compatible with PSpice* and HSPICE**. The schematic and net-list are included here so that the model can easily be used. This model can accommodate multiple frequency poles and multiple zeroes, which is an advanced concept that results in more accurate AC and transient responses. For example, 5 poles are required to sufficiently simulate the OP-400, which this advanced model can easily accommodate.

This macro-model represents one of the four amplifiers which are in the OP-400 package. To use the quad amplifiers, the simulation circuit needs only to call up the model four separate times and specify the same power supplies each time. Throughout the OP-400 macro-model, RC networks produce multiple poles which simulate the OP-400's AC behavior. The stages, which each contain a single pole, are separated from each other by voltage-controlled current sources so that the pole locations do not interact. Stages with a zero can also be added but are not necessary for the OP-400. The only nonlinear elements in the entire model are two NPN transistors which comprise the input stage, and diodes for voltage clamping. Limiting the model to almost entirely linear circuit elements significantly reduces simulation time and simplifies model development.

MODEL DESCRIPTION

The schematic (Figure 1) and net-list (Figure 2) describe the complete OP-400 model. This model breaks up the OP-400 into many distinct stages as described below.

INPUT STAGE

To correctly model the OP-400's input behavior, the model uses a differential pair of NPN transistors biased with a 100µA current source (Figure 1a). To keep this stage as simple as possible, only the forward beta, BF, is specified in the NPN-BJT model. This is chosen to give the correct input bias current, I_B, for the op amp. All other BJT parameters are left at the model default values, most of which are zero. As can be seen in the simplified schematic on the data sheet, the OP-400 has a voltage limiting network to protect the input transistor pair. This network is essentially two back-to-back diodes, which are included in the model. Furthermore, two 1kΩ resistors are placed in series with the inputs to limit the current if either diode turns on.

As for non-ideal behaviors of the input stage, such as V_{OS}, I_{OS}, and C_{IN}, these are modelled with external circuit elements. For example, no gate capacitance is specified for the NPN model, therefore a capacitor, C_{IN}, is added across the inputs. Furthermore, since in this model the input NPNs are perfectly matched, V_{OS} and I_{OS} error sources are added using an external voltage source and current source, respectively. Lastly, the drain resistors R_3 and R_4 are chosen to be 1/gm of the NPN transistors to give a gain of unity in the input stage. C_2 is added to create one of the secondary poles in the model.

GAIN STAGES

Because the open-loop gain of the OP-400 is so large, the model uses two stages (Figures 1b, c) to achieve it, and all other stages have unity gain. The voltage-controlled current sources for the two gain stages have scaled transconductances that, when combined with their respective resistors, yield a combined open-loop gain of 5,000,000. The first gain stage uses C_3 and C_4 to create the dominant pole at 0.125Hz and to model the amplifier slew rate. In addition, the second gain stage uses C_5 and C_6 in combination with R_7 and R_8 to add a secondary pole at 4.5MHz.

The diodes, D_1 and D_2, are needed to clamp the voltage at node 8 below the power supplies. More importantly, the sources V_2 and V_3, and diodes D_3 and D_4, are necessary to clamp the voltage of node 10 below the power supplies to get the proper output voltage swing. Because the next stage (Figure 1d) has unity gain and its voltage-controlled current sources are controlled by node 10, it too is clamped, at node 12, below the power supplies. The same is true for the subsequent stages, including the output at node 25, such that their voltages are clamped below the power supplies, which gives the correct output voltage swing.

POLE STAGES

In addition to the three poles in the input stage and the gain stages, additional poles are easily added with separate stages. The pole stages (Figures 1d, f) both have unity gain, which is a result of the gm of the voltage-controlled current sources being the reciprocal of the resistors. The pole locations are determined by the resistor and capacitor values for each of the stages.

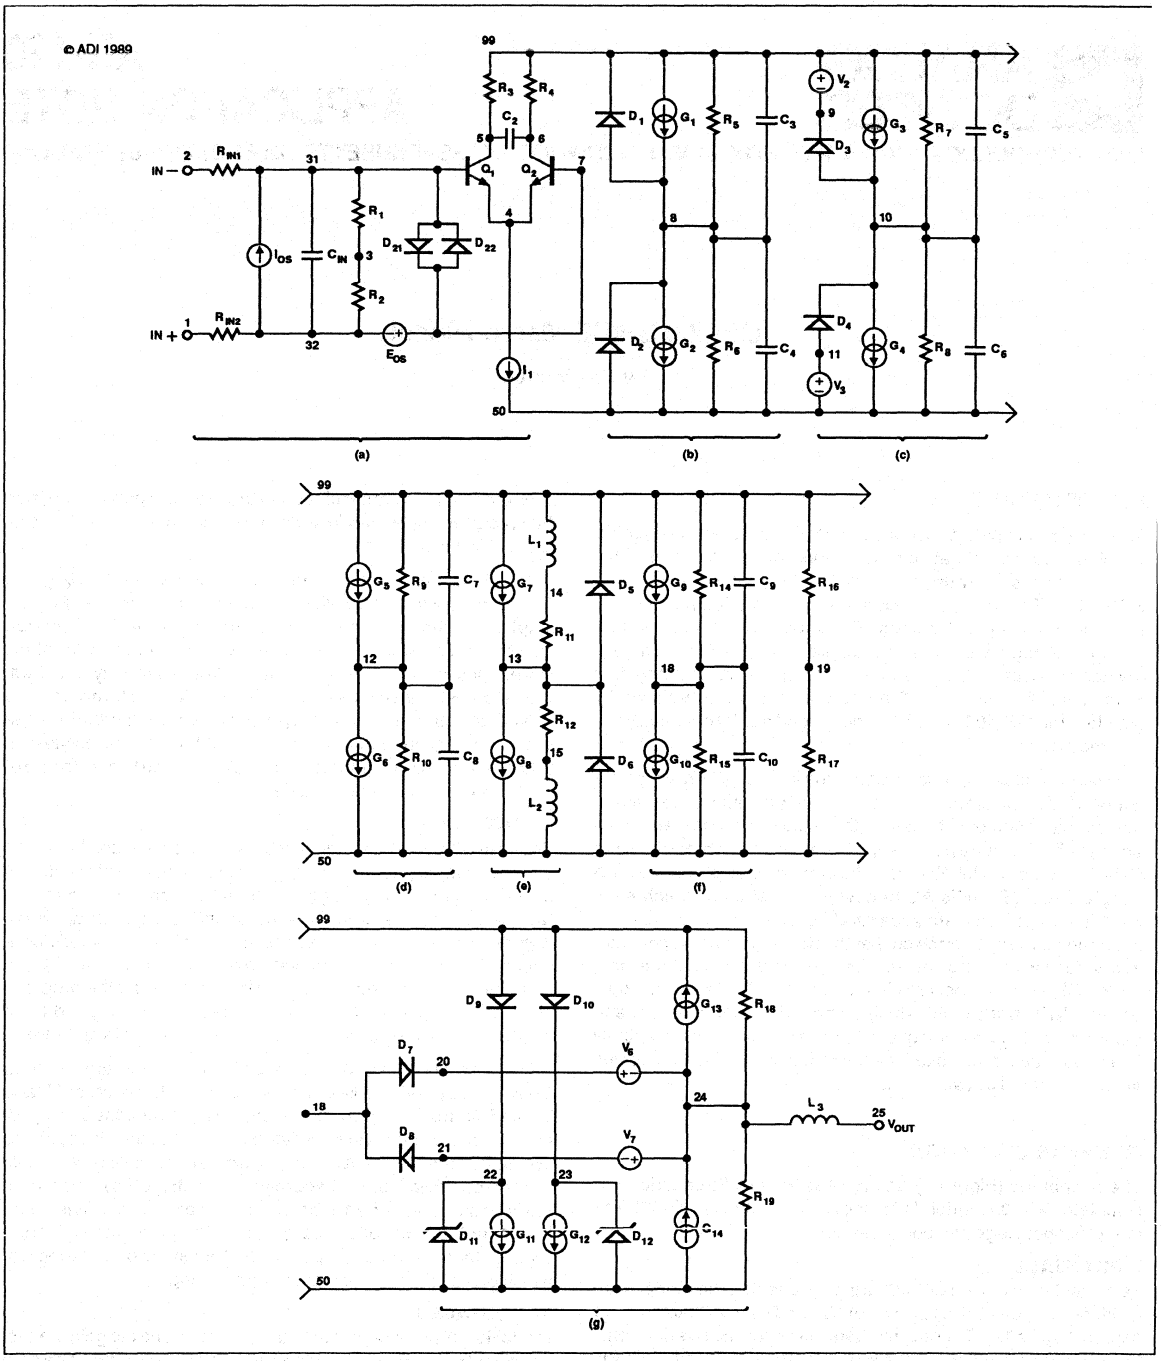

FIGURE 1: *OP-400 SPICE Macro-Model Schematic and Node List*

```
OP-400 MACRO-MODEL © ADI 1990
*
.subckt OP-400 1 2 25 99 50
*
* INPUT STAGE & POLE AT 4.5 MHz
*
rin1 2   31       1E3
rin2 1   32       1E3
r1   31  3        5E5
r2   32  3        5E5
r3   5   99       516
r4   6   99       516
cin  32  31       3.2E-12
c2   5   6        34.271E-12
i1   4   50       0.1E-3
ios  32  31       1E-10
eos  7   32       poly(1) 13 19 4E-5 1
q1   5   31  4    qx
q2   6   7   4    qx
d21  31  7        dx
d22  7   31       dx
*
* FIRST GAIN STAGE & DOMINANT POLE AT 0.125 Hz
*
r5   8   99       1.9098E9
r6   8   50       1.9098E9
c3   8   99       666.67E-12
c4   8   50       666.67E-12
g1   99  8        poly(1)  5   6   200E-6  1.938E-3
g2   8   50       poly(1)  6   5   200E-6  1.938E-3
d1   8   99       dx
d2   50  8        dx
*
* SECOND GAIN STAGE & POLE AT 4.5 MHz
*
r7   10  99       1E6
r8   10  50       1E6
c5   10  99       35.368E-15
c6   10  50       35.368E-15
g3   99  10       8 19  1.3509E-6
g4   10  50       19 8  1.3509E-6
v2   99  9        2.7
v3   11  50       2.7
d3   10  9        dx
d4   11  10       dx
*
* POLE AT 4.5 MHz
*
r9   12  99       1E6
r10  12  50       1E6
c7   12  99       35.368E-15
c8   12  50       35.368E-15
g5   99  12       10 19  1E-6
g6   12  50       19 10  1E-6
*
* COMMON-MODE GAIN NETWORK WITH ZERO AT 20 Hz
*
r11  13  14       1E6
r12  13  15       1E6
L1   14  99       7.9578E3
L2   15  50       7.9578E3
g7   99  13       3 19      1E-13
g8   13  50       19 3      1E-13
d5   13  99       dx
d6   50  13       dx
*
* POLE AT 4.5 MHz
*
r14  18  99       1E6
r15  18  50       1E6
c9   18  99       35.368E-15
c10  18  50       35.368E-15
g9   99  18       12 19  1E-6
g10  18  50       19 12  1E-6
*
* OUTPUT STAGE
*
r16  19  99       62.22E3
r17  19  50       62.22E3
r18  24  99       250
r19  24  50       250
L3   24  25       7E-7
g11  22  50       18  24  4E-3
g12  23  50       24  18  4E-3
g13  24  99       99  18  4E-3
g14  50  24       18  50  4E-3
v6   20  24       3
v7   24  21       3
d7   18  20       dx
d8   21  18       dx
d9   99  22       dx
d10  99  23       dx
d11  50  22       dy
d12  50  23       dy
*
* MODELS USED
*
.model   qx   NPN (BF=133333)
.model   dx   D (IS=1E-15)
.model   dy   D (IS=1E-15 BV=50)
.ends    OP-400
```

FIGURE 2: *OP-400 SPICE Net-List*

COMMON-MODE STAGE

The common-mode voltage that is used to create the CM error is created by the two input resistors, R_1 and R_2. This voltage is referenced by G_7 and G_8 in the common-mode stage (See Figure 1e). The transconductances of these two sources are scaled such that, in combination with R_{11} and R_{12}, the V_{CM} is attenuated by the CMRR of 140dB. This error voltage is then inserted as an offset voltage in the input stage by E_{os}. The inductors, L_1 and L_2, mimic the pole in the CMR vs. frequency response of the OP-400.

OUTPUT STAGE

The output stage (Figure 1g) is modelled as an ideal output with an output resistance, R_{18} in parallel with R_{19}. An inductor, L_3, is added to model the rising output impedance of the emitter-follower output stage with frequency. The voltage sources V_6 and V_7, and diodes D_7 and D_8, combine to limit the voltage across the resistors, thus limiting the output current, I_{OUT} to 30mA.

PARAMETERS THAT ARE NOT MODELLED

To keep the OP-400 model as simple as possible and thus save computer time, not all features of the op amp were modelled as listed below:

- Channel Separation
- PSRR
- No Limits on Power Supply Voltages
- Maximum Input Voltage Beyond the Power Supplies
- Temperature Effects (i.e., model parameters are assumed at 25°C)
- Input Noise Voltage and Current Sources
- Parameter Variations for Monte Carlo Analysis (i.e., all parameters are typical only)

These parameters are considered second-order effects and are not considered necessary for circuit simulation under normal operating conditions. However, users can easily add these functions as needed.

MODEL PERFORMANCE

The graphs on these next two pages compare the actual responses of the OP-400 to the SPICE analysis under various load and gain conditions. To show the accuracy of the model's frequency response, it was tested with different gains. Shown below are the gain-phase plots for an open-loop gain (Figure 3) and a closed-loop gain of –1 (Figure 4). The actual responses (part "a" of each figure), which were measured using a network analyzer, show close agreement to the SPICE simulations (part "b" of each figure). This agreement holds true for both the gain and phase curves, which demonstrates the accuracy of the model's frequency response.

The model can also accurately simulate the OP-400's transient response, which is stable for a closed-loop gain of –1 even with large capacitive loads. The comparison was done with a 10nF capacitive load, which shows the stability of the OP-400. Because of the model's complete complement of poles, its large-signal transient reponse (Figure 5b) shows close resemblance to the actual response (Figure 5a) including the slew rate limit of 0.15V/μs. The small-signal response (Figure 6) again shows the same accuracy, which is even more apparent because of the damped oscillations. In both reponses the actual device shows a slight difference between its negative and positive transitions, which is due to the inherent asymmetrical behavior of the OP-400's output stage. This is most apparent in the slew rate for the large-signal response, and the overshoot for the small-signal response. On the other hand, the model assumes a perfectly symmetrical output stage so its responses do not exhibit this difference. Nevertheless, as can be seen from the figures, the model agrees closely with the actual device's performance. This accurate transient response, along with the accurate AC response, makes the OP-400 macro-model a powerful tool in analyzing circuit stability and AC behavior under various load conditions.

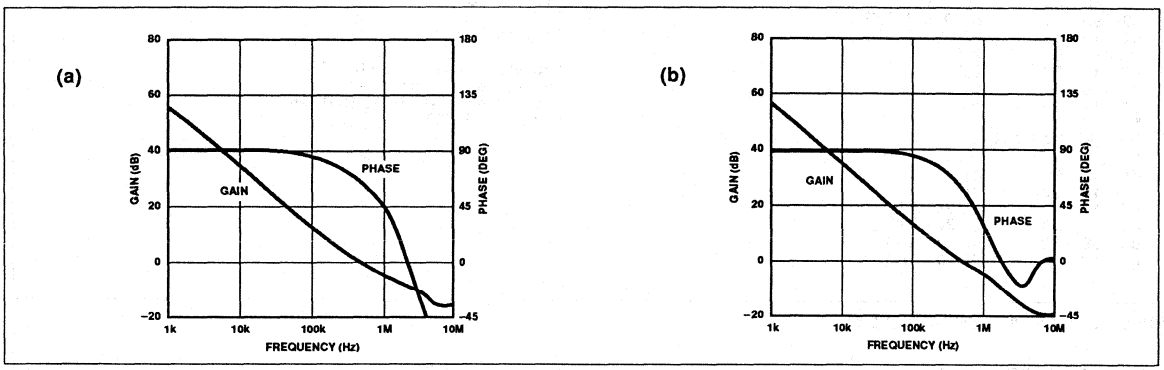

FIGURE 3: *Gain-Phase Plots, Open-Loop Gain (a. Actual Response, b. Simulated Model Response)*

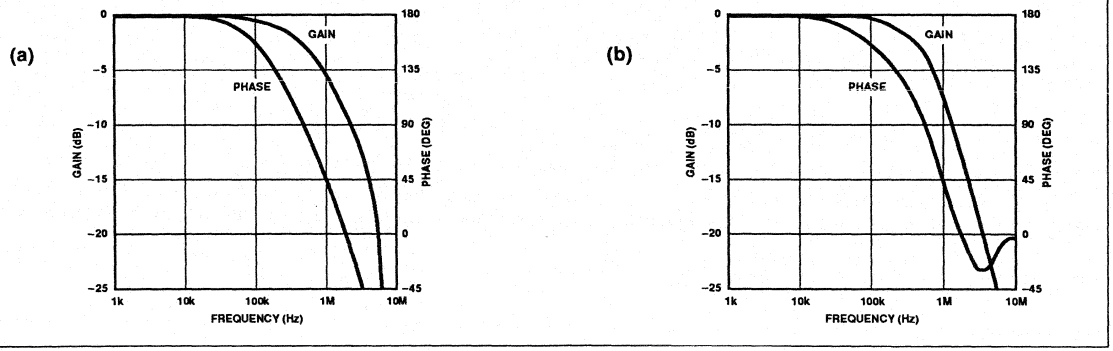

FIGURE 4: *Gain-Phase Plots, Closed-Loop Gain of –1 (a. Actual Response, b. Simulated Model Response)*

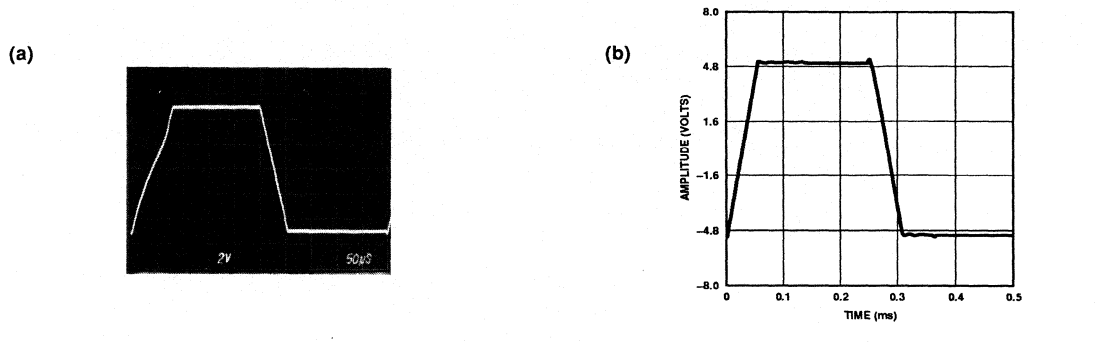

FIGURE 5: *Large-Signal Transient Response, Gain = – 1, 10nF Capacitive Load (a. Actual Response, b. Simulated Model Response)*

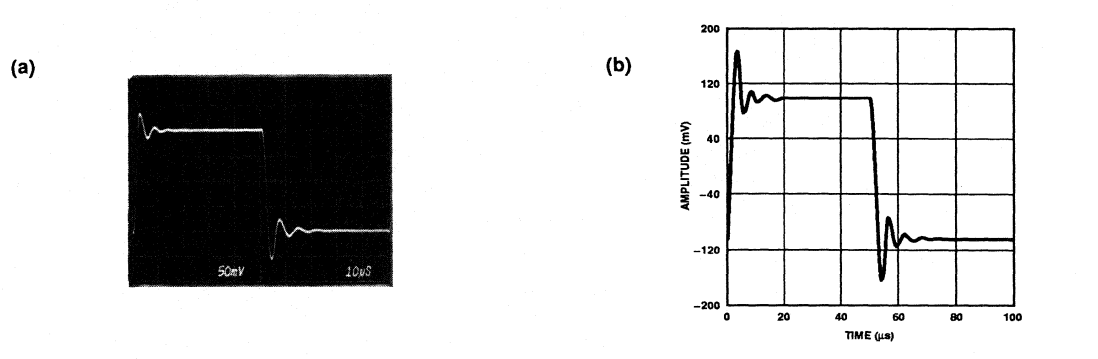

FIGURE 6: *Small-Signal Transient Response, Gain = –1, 10nF Capacitive Load (a. Actual Response, b. Simulated Model Response)*

ANALOG DEVICES

AN-126
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

OP-260 Advanced SPICE Macro-Model

by Joe Buxton

INTRODUCTION

This application note describes the SPICE macro-model for the OP-260 dual, high-speed, current feedback operational amplifier. This model was tested with, and is compatible with PSpice* and HSpice**. The schematic and net-list are included here so that the model can easily be used. This model uses a unique current feedback topology to accurately model both the AC and DC characteristics of the OP-260. In addition, this model can accommodate any number of poles and zeros to further shape the AC response.

This model consists of one of the op amps in the dual OP-260 package. To use this model as a dual, just call up the model twice and specify the same power supplies for each op amp. The OP-260 SPICE macro-model uses four BJT transistors to create the input buffer just as the actual device does. However, the rest of the model contains only ideal linear elements and ideal diodes to model the OP-260's behavior. Using only four transistors reduces simulation time and simplifies model development.

MODEL DESCRIPTION

The schematic (Figure 1) and net-list (Figure 2) describe the complete OP-260 model. This model breaks up the OP-260 into many distinct stages as described below:

INPUT STAGE

The topology of the model's input stage (Figure 1a) is actually quite similar to the actual device's input stage. The four transistors Q_1-Q_4 make up the input buffer and are biased with the 150μA current sources. The current is chosen so that the inverting input will have the correct buffer output of 100Ω. The resistors R_1 and R_2 are used by the following gain stage to sense the current flowing through Q_3 and Q_4 respectively. It is the difference in current between Q_3 and Q_4, i.e., the current flowing into or out of the inverting input, that provides the current feedback which is gained up by the gain stage.

Most of the input error sources are modelled with external circuit elements, such as V_{OS}, I_{B1}, and G_{B1}. Not only does G_{B1} model the constant 3μA of bias current in the inverting input, but it also models the inverting input common-mode rejection, $CMRRI_{B-}$. The two capacitors, C_{S1} and C_{S2}, model the parasitic stray capacitance on the inverting input, which greatly affects the amplifier response at high frequencies by causing peaking. The two RC networks model a pole in the input stage and also the input stage's slew rate, which typically exhibits 1000V/μs. Common-mode rejection is also modelled in the input stage by including the forward Early voltage in the transistor models, which varies the input offset voltage.

GAIN STAGE

The open-loop gain of the OP-260 is actually given as a transimpedance. The differential current flowing in the input stage is sensed as the voltage drop across R_1 and R_2 by the voltage controlled current sources, G_1 and G_2. These current sources in combination with the resistors, R_5 and R_6, produce the single-ended output voltage. This stage also limits the output voltage of the amplifier by clamping the voltage at node 12 to the output voltage range. Since all succeeding stages, which are controlled by this voltage, have unity gain, then the output will also be clamped.

The capacitors, C_3 and C_4, produce the dominant pole in combination with external feedback resistor. These capacitors also limit the slew rate of the amplifier. The slew rate is set by the size of these capacitors and the maximum amount of current that can flow in this stage. This current is from the current sources, which are controlled by the voltage drop across R_1 and R_2. Thus limiting the voltage drop across these resistors will limit the slew rate, which is done by the voltage sources and clamp diodes across these resistors. By changing the value of the voltage sources, the slew rate can be easily changed.

POLE STAGES

All of the pole stages (Figures 1c - e) have unity gain, which is a result of the g_m of the voltage-controlled current sources being the reciprocal of the resistors. The RC networks are used to model the secondary poles of the OP-260.

OUTPUT STAGE

The output stage (Figure 1f) is modelled as an ideal output with an output resistance, R_{15} in parallel with R_{16}. An inductor, L_1, is added to model the rising output impedance of the emitter-follower output stage with frequency. The voltage sources V_5 and V_6, and diodes D_5 and D_6, combine to limit the voltage across the resistors, thus limiting the output current, I_{OUT} to 30mA. A capacitor, C_{F1}, is added from the output to the inverting input to simulate the parasitic pin-to-pin capacitance of the actual device.

13

* PSpice is a registered trademark of MicroSim Corporation.
** HSpice is a tradename of Meta-Software, Inc.

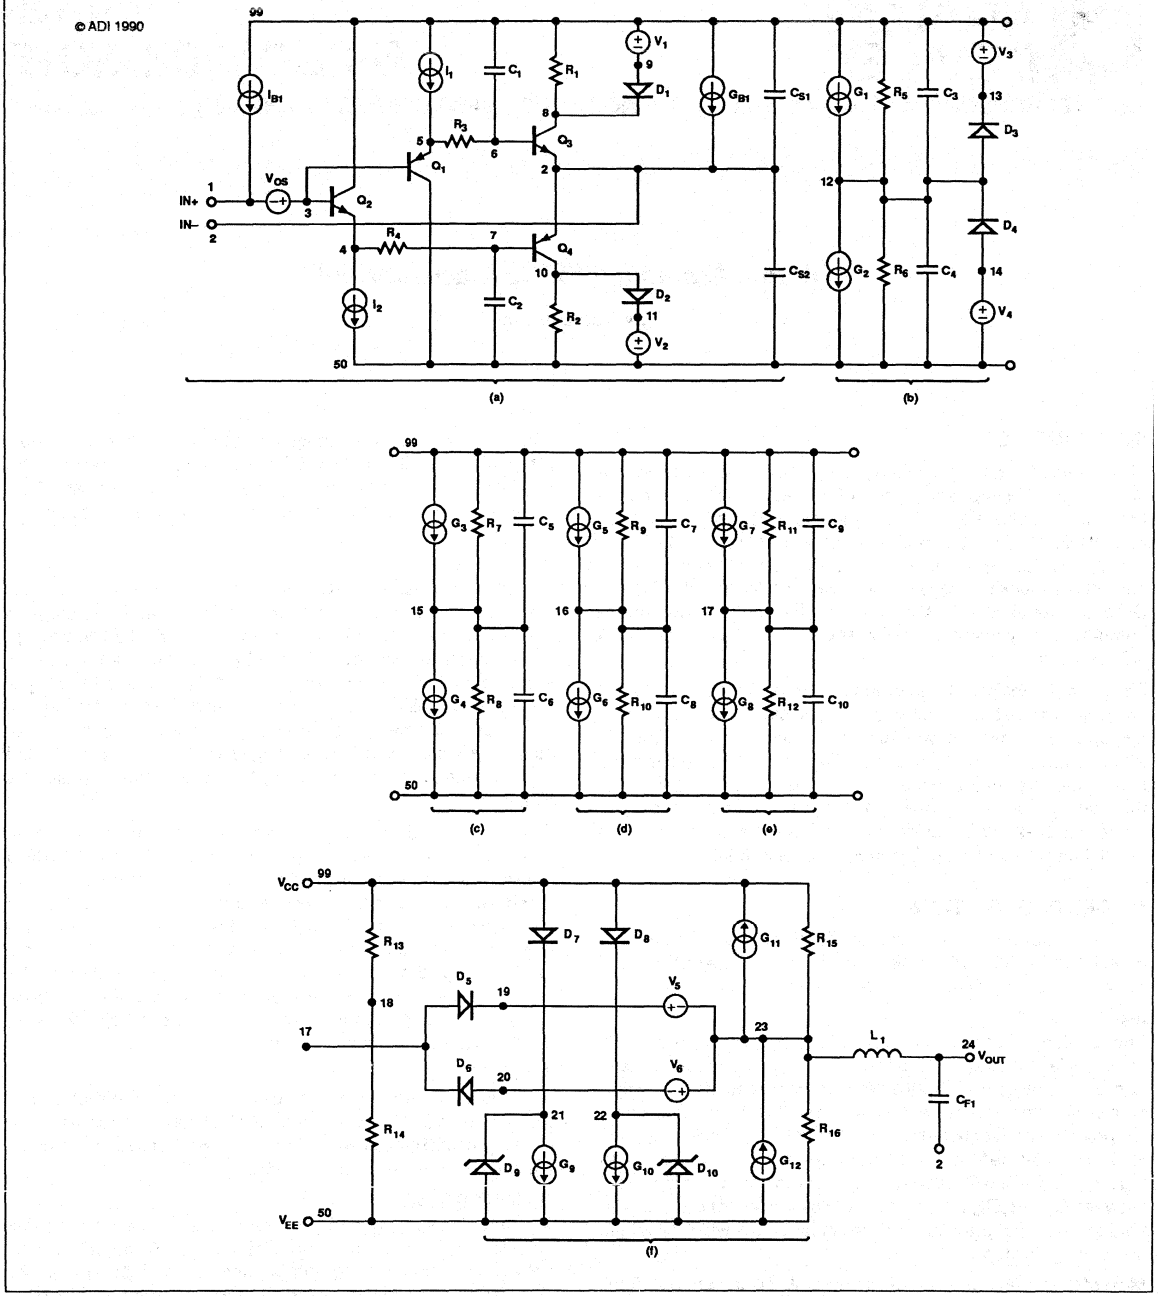

FIGURE 1: *OP-260 Macro-Model Schematic*

```
OP-260 MACRO-MODEL © ADI 1990
*
* SUBCKT OP-260 1  2 24 99 50
*
* INPUT STAGE
*
R1    99   8          4K
R2    10   50         4K
V1    99   9          1.1
D1    9    8          DX
V2    11   50         1.1
D2    10   11         DX
I1    99   5          15OU
I2    4    50         15OU
Q1    50   3    5     QP
Q2    99   3    4     QN
Q3    8    6    2     QN
Q4    10   7    2     QP
R3    5    6          14.3K
R4    4    7          14.3K
C1    99   6          0.133P
C2    50   7          0.133P
*
* INPUT ERROR SOURCES
*
GB1   99   2          POLY(1)   1  18  3E-6  4E-8
IB1   99   1          2E-7
VOS   3    1          1E-3
CS1   99   2          0.8E-12
CS2   50   2          0.8E-12
*
* GAIN STAGE & DOMINANT POLE
*
R5    12   99         10E6
R6    12   50         10E6
C3    12   99         0.6P
C4    12   50         0.6P
G1    99   12         POLY (1)  99   8    4E-3     0.25E-3
G2    12   50         POLY (1)  10   50   4E-3     0.25E-3
V3    99   13         2.2
V4    14   50         2.2
D3    12   13         DX
D4    14   12         DX
*
* POLE AT 64 MHz
*
R7    15   99         1E6
R8    15   50         1E6
C5    15   99         2.5E-15
C6    15   50         2.5E-15
G3    99   15         12   18    1E-6
G4    15   50         18   12    1E-6
*
* POLE AT 72 MHz
*
R9    16   99         1E6
R10   16   50         1E6
C7    16   99         2.2E-15
C8    16   50         2.2E-15
G5    99   16         15   18    1E-6
G6    16   50         18   15    1E-6
*
* POLE  AT 80 MHz
*
R11   17   99         1E6
R12   17   50         1E6
C9    17   99         2E-15
C10   17   50         2E-15
G7    99   17         16   18    1E-6
G8    17   50         18   16    1E-6
*
* OUTPUT STAGE
*
R13   18   99         3.333E3
R14   18   50         3.333E3
R15   23   99         150
R16   23   50         150
L1    23   24         1.5E-8
CF1   24   2          1.8P
G9    21   50         17   23    6.66667E-3
G10   22   50         23   17    6.66667E-3
G11   23   99         99   17    6.66667E-3
G12   50   23         17   50    6.66667E-3
V5    19   23         1.55
V6    23   20         1.55
D5    17   19         DX
D6    20   17         DX
D7    99   21         DX
D8    99   22         DX
D9    50   21         DY
D10   50   22         DY
*
*MODELS USED
*
.MODEL   QN  NPN (BF = 1E9   IS = 1E-15   VAF = 150)
.MODEL   QP  PNP (BF = 1E9   IS = 1E-15   VAF = 150)
.MODEL   DX  D (IS = 1E-15)
.MODEL   DY  D (IS = 1E-15   BV =50)
.ENDS OP-260
```

FIGURE 2: *OP-260 SPICE Net-List*

13

PARAMETERS THAT ARE NOT MODELLED

To keep the OP-260 model as simple as possible and thus save computer and development time, not all features of the op amp were modelled as listed below:

- PSRR
- Crosstalk
- Varying slew rate with closed-loop gain
- No limits on power supply voltages
- Maximum input voltage range
- Temperature effects (i.e., model parameters are assumed at 25°C)
- Input noise voltage and current sources
- Parameter variations for Monte Carlo analysis (i.e., all parameters are typical only)

These parameters are considered second-order effects and are not considered necessary for circuit simulation under normal operating conditions. However, users can easily add these functions as needed.

MODEL PERFORMANCE

The graphs on these next two pages compare the actual responses of the OP-260 to the SPICE analysis under various load and gain conditions. To show the accuracy of the model 's frequency response, it was tested in different gains with a varying load resistor as was done in the data sheet. Shown below are the gain and phase plots for an closed-loop gain of +2 (Figures 3, 4) and a closed-loop gain of +10 (Figures 5, 6). The actual responses (part "a" of each figure), which were measured using a networ analyzer, show very close agreement to the SPICE simulation (part "b" of each figure). This agreement holds true for both th gain and phase curves including the peaking in the gain respons for A_V = +2. Also notice that the model predicts the OP-260' behavior with varying resistive loads and by having a nearly cor stant 3dB point.

The model can also accurately simulate the OP-260 's transier response. The comparisons were done with the amplifier in closed-loop gain of +5 with a 20pF capacitive load and a 2.5kΩ feedback resistor. The 20pF load was used to create the dampe oscillations in the small-and-large signal responses. Because o the multiple poles and the accurate modelling of the parasiti elements, the model's small-signal transient response (Figure 8b shows close resemblance to the actual response (Figure 8a). The large-signal response (Figure 7) again shows good accuracy. Ir the large-signal response the actual device does show a differ ence in the damped oscillations, which is due to the inheren asymmetrical behavior of the OP-260's output stage. On the othe hand, the model assumes a perfectly symmetrical output stage so its response does not exhibit this difference. The actual device also shows a damping of the overshoot of the first oscillation, which is not modelled. Nevertheless, as can be seen from the figures, the model agrees closely with the actual device's performance. This acurate transient response, along with the accurate AC response makes the OP-260 macro-model a powerful tool in analyzing circuit stability and AC behavior under various gains and load conditions.

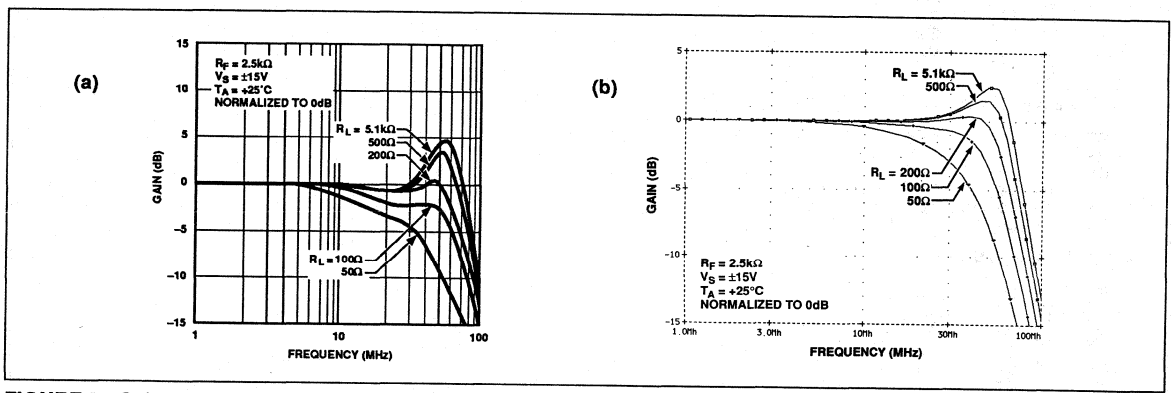

FIGURE 3: *Gain Plots, Closed-Loop Gain of +2 (a. Actual Response, b. Simulated Model Response)*

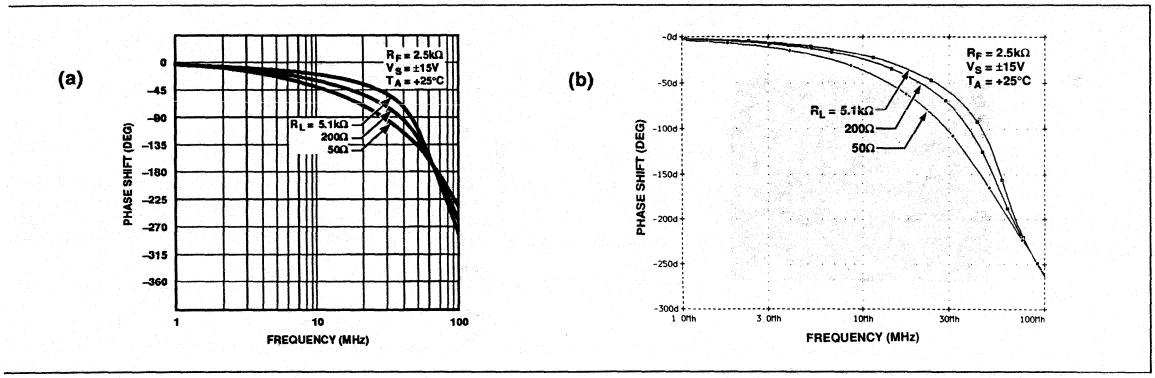

FIGURE 4: *Phase Plots, Closed-Loop Gain of +2 (a. Actual Response, b. Simulated Model Response)*

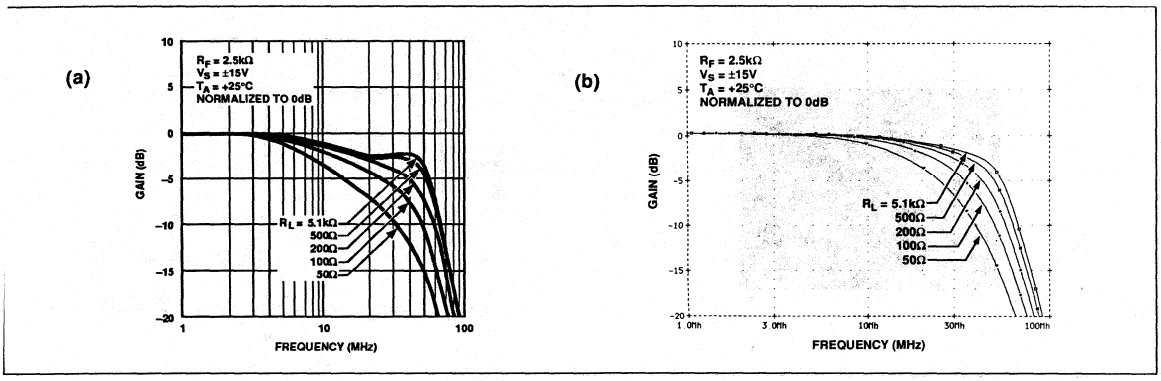

FIGURE 5: *Gain Plots, Closed-Loop Gain of +10 (a. Actual Response, b. Simulated Model Response)*

13

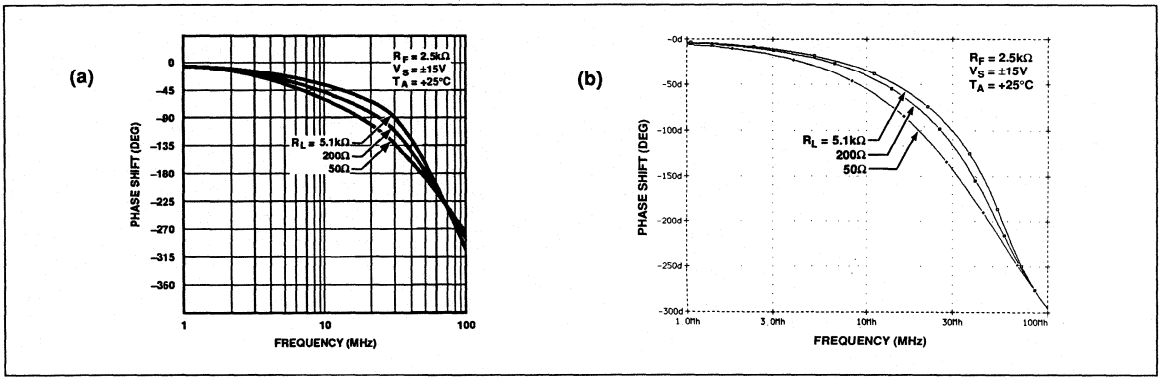

FIGURE 6: *Phase Plots, Closed-Loop Gain of +10 (a. Actual Response, b. Simulated Model Response)*

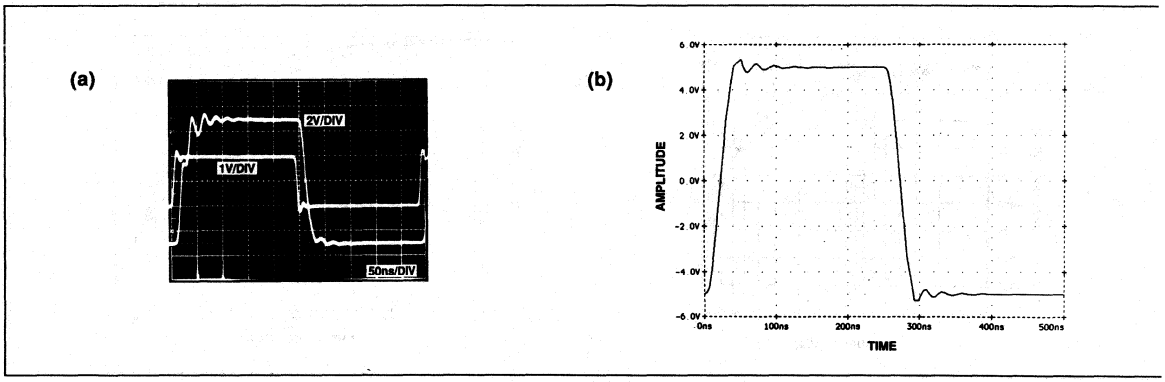

FIGURE 7: *Large-Signal Transient Response, $A_V = +5$ (a. Actual , b. Model)*

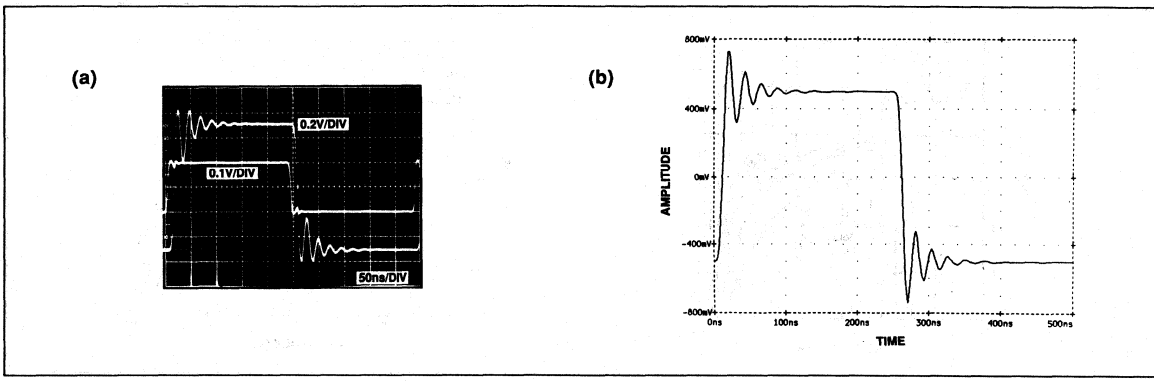

FIGURE 8: *Small-Signal Transient Response, $A_V = +5$ (a. Actual, b. Model)*

ANALOG DEVICES

AN-132
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

OP-470 SPICE Macro-Model

by Joe Buxton

INTRODUCTION

This application note describes the SPICE macro-model for the OP-470 very low noise quad operational amplifier. This model was tested with and is compatible with PSpice* and HSpice**. The schematic and net-list are included here so that the model can easily be used. This model can accommodate multiple frequency poles and zeros, which is an advanced concept that results in more accurate AC and transient responses. For example, 6 poles and 2 zeros are required to accurately model the OP-470, which this model can easily accommodate.

This macro-model represents one of the four amplifiers on one chip. To use the quad amplifiers, the simulation circuit needs only to call up the model four separate times and specify the same power supplies each time. Throughout the OP-470 macro-model RC networks produce multiple poles which simulate the amplifiers AC behavior. The stages, which each contain a single pole or a pole- zero pair, are separated from each other by voltage-controlled current sources so that it is easy to set individual pole locations. The only nonlinear elements in the entire model are two NPN transistors which comprise the input stage, and diodes for voltage clamping. Limiting the model to almost entirely ideal, linear circuit elements significantly reduces simulation time and simplifies model development.

MODEL DESCRIPTION

The schematic (Figure 1) and net-list (Figure 2) describe the complete OP-470 model. This model breaks up the OP-470 into many distinct stages as described below:

INPUT STAGE

To correctly model the OP-470's input behavior, the model uses a differential pair of NPN transistors biased with a 1mA current source (Figure 1a). To keep this stage as simple as possible, only the forward beta, B_F, is specified in the NPN-BJT model. This is chosen to give the correct input bias current, I_B. The voltage limiting network of the OP-470 input is added with the diodes in series with the voltage source across the inputs. A voltage source is used instead of another diode to save computer time.

As for nonideal behaviors of the input stage, such as V_{OS}, I_{OS}, and C_{IN}, these are modelled with external circuit elements. For example, no junction capacitance is specified for the NPN model, therefore a capacitor, C_{IN}, is added across the inputs. Furthermore, since the input NPNs in this model are perfectly matched, V_{OS} and I_{OS} error sources are added using an external voltage source and current source, respectively. Lastly, the collector resistors R_3 and R_4 are chosen to be $1/g_m$ of the NPN transistors to give a gain of unity in the input stage. C_2 is added to create one of the secondary poles in the model.

GAIN STAGES

The open-loop gain of the OP-470 is acheived entirely in the gain stage (Figure 1b), and all other stages have unity-gain. The gain is taken in two steps due to a limit for only one step based on the combination of the slew rate, dominant pole, and open-loop gain. The voltage-controlled current sources, G_1 and G_2, have scaled transconductances that, when combined with R_7 and R_8, give the gain in the first step at node 12. The gain in the second step is produced the same way, and together the two steps give the open-loop gain of 1,000,000. The current sources G_1 and G_2 are controlled by the voltage drops across R_3 and R_4 in the input stage, providing the differential to single ended voltage conversion. C_3 and C_4 create the dominant pole at 5.5Hz and model the amplifier's slew rate. Lastly, the diodes, D_5 and D_6 and voltage sources V_3 and V_4, are necessary to clamp the voltage of node 15 below the power supplies. Because the next stage (Figure 1c) has unity-gain and its voltage is controlled by the voltage at node 15, then it too is clamped, at node 18, below the power supplies. The same is true for subsequent stages, including the output, such that the output voltage is clamped to within the specified voltage range.

POLE STAGES

In addition to the poles in the input stage and the gain stage, additional poles are easily added with separate stages. The pole and pole-zero stages (Figure 1c-e, g) have unity-gain, which is a result of the g_m of the voltage-controlled current sources being the reciprocal of the resistors. The pole and zero locations are determined by the resistor and capacitor or inductor values for each stage.

13

* PSpice is a registered trademark of MicroSim Corporation.
** HSpice is a tradename of Meta-Software, Inc.

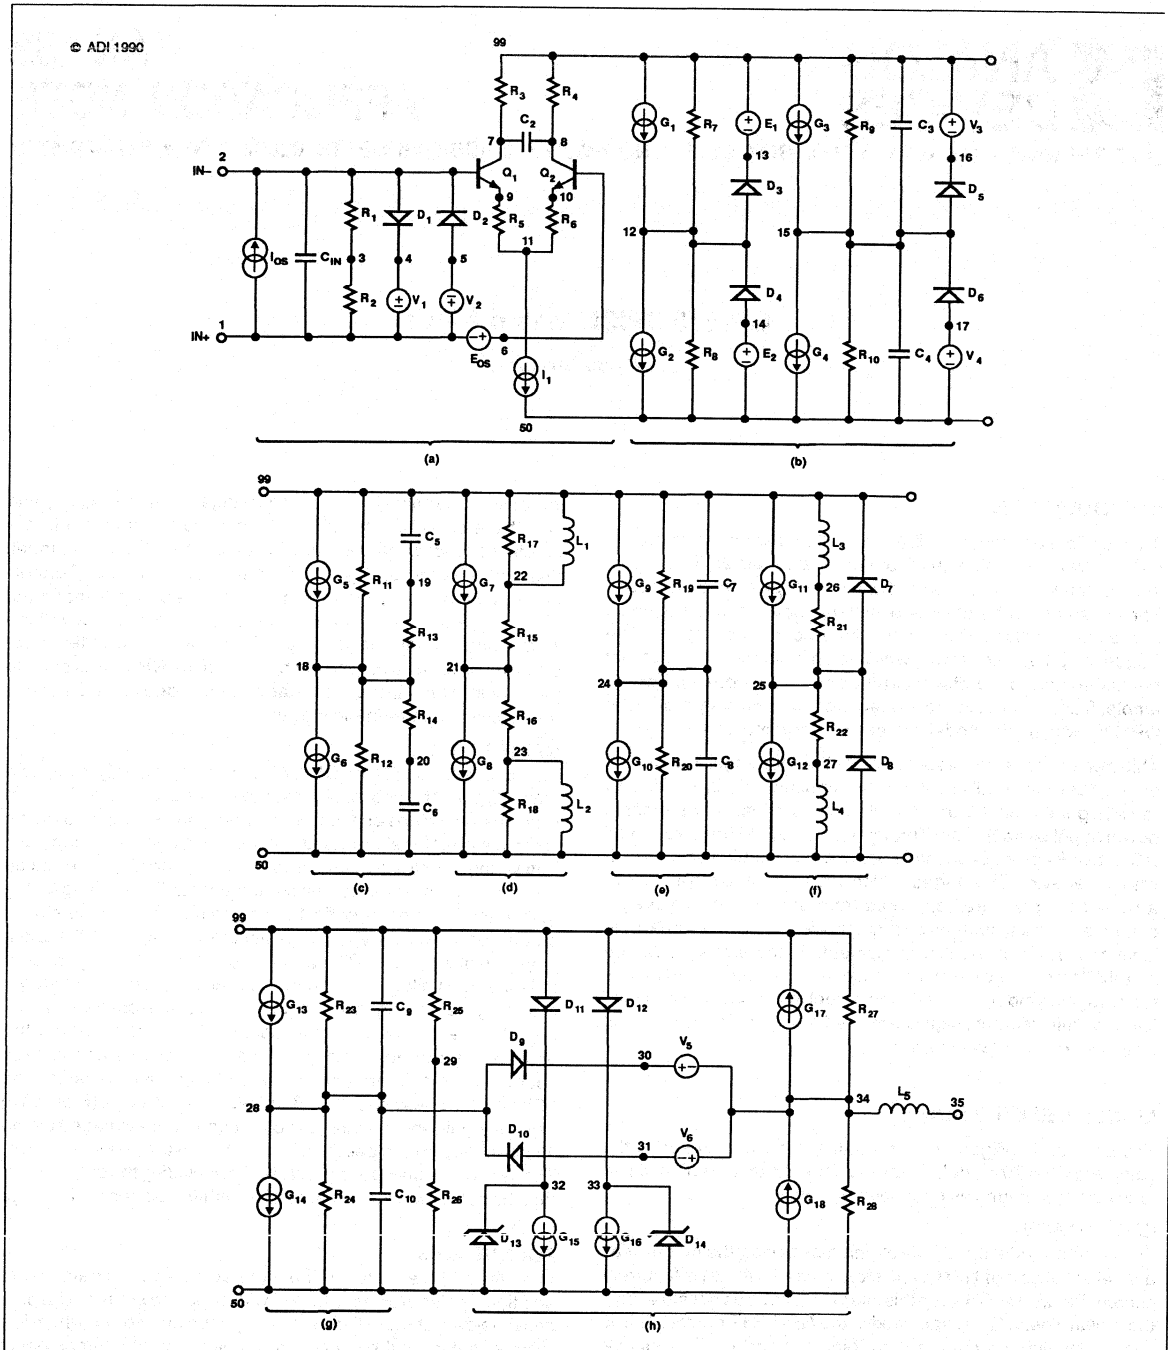

FIGURE 1: *OP-470 SPICE Macro-Model Schematic and Node List*

```
OP-470 MACRO-MODEL © ADI 1990
*
.SUBCKT OP-470 1 2 35 99 50
*
* INPUT STAGE & POLE AT 60 MHZ
*
R1   2   3        2E5
R2   1   3        2E5
R3   7   99       101.6
R4   8   99       101.6
CIN  1   2        2E-12
C2   7   8        13.05E-12
I1   11  50       1E-3
IOS  1   2        3E-9
EOS  6   1        POLY(1) 25 29 1E-4 1
Q1   7   2    9   QX
Q2   8   6    10  QX
R5   9   11       50
R6   10  11       50
D1   2   4        DX
V1   4   1        0.7
V2   1   5        0.7
D2   5   2        DX
*
* FIRST GAIN STAGE
*
R7   12  99       1E6
R8   12  50       1E6
G1   99  12       7   8   64.1E-6
G2   12  50       8   7   64.1E-6
D3   12  13       DX
D4   14  12       DX
E1   99  13       POLY(1) 99 29 -2.6 1
E2   14  50       POLY(1) 29 50 -2.6 1
*
* FIRST GAIN STAGE & DOMINANT POLE AT 6.8 HZ
*
R9   15  99       46.8E6
R10  15  50       46.8E6
C3   15  99       500E-12
C4   15  50       500E-12
G3   99  15       POLY(1) 12 29 380E-6 333.3E-6
G4   15  50       POLY(1) 29 12 380E-6 333.3E-6
V3   99  16       1.9
V4   17  50       1.9
D5   15  16       DX
D6   17  15       DX
*
* POLE AT ZERO AT 1.9 MHZ/3.5 MHZ
*
R11  18  99       1E6
R12  18  50       1E6
R13  18  19       1.19E6
R14  18  20       1.19E6
C5   19  99       38.21E-15
C6   20  50       38.21E-15
G5   99  18       15 29 1E-6
G6   18  50       29 15 1E-6
*
* ZERO-POLE AT 5 MHZ/10 MHZ
*
R15  21  22       1E6
R16  21  23       1E6
R17  99  22       1E6
R18  50  23       1E6
L1   99  22       15.92E-3
L2   50  23       15.92E-3
G7   99  21       18 29 1E-6
G8   21  50       29 18 1E-6
*
* POLE AT 10 MHZ
*
R19  24  99       1E6
R20  24  50       1E6
C7   24  99       15.9E-15
C8   24  50       15.9E-15
G9   99  24       21 29 1E-6
G10  24  50       29 21 1E-6
*
* COMMON-MODE GAIN NETWORK WITH ZERO AT 100 HZ
*
R21  25  26       1E6
R22  25  27       1E6
L3   26  99       1.592E3
L4   27  50       1.592E3
G11  99  25       3   29  1E-12
G12  25  50       29  3   1E-12
D7   25  99       DX
D8   50  25       DX
*
* POLE AT 50 MHZ
*
R23  28  99       1E6
R24  28  50       1E6
C9   28  99       3.18E-15
C10  28  50       3.18E-15
G13  99  28       24 29 1E-6
G14  28  50       29 24 1E-6
*
* OUTPUT STAGE
*
R25  29  99       20E3
R26  29  50       20E3
R27  34  99       300
R28  34  50       300
L5   34  35       1E-6
G15  32  50       28 34 3.33E-3
G16  33  50       34 28 3.33E-3
G17  34  99       99 28 3.33E-3
G18  50  34       28 50 3.33E-3
V5   30  34       1.3
V6   34  31       1.3
D9   28  30       DX
D10  31  28       DX
D11  99  32       DX
D12  99  33       DX
D13  50  32       DX
D14  50  33       DX
*
* MODEL USED
*
.MODEL QX NPN (BF = 166667)
.MODEL DX D (IS = 1E-15)
.MODEL DY D (IS = 1E-15  BV = 50)
.ENDS OP-470
```

FIGURE 2: *OP-470 SPICE Net-List*

COMMON-MODE STAGE

The common-mode error is based on the common-mode voltage of the input, which is created by the two input resistors, R_1 and R_2. This voltage is referenced by G_{11} and G_{12} in the common-mode stage (Figure 1f). The transconductances of these two sources are scaled such that, in combination with R_{21} and R_{22}, the V_{CM} is attenuated by the CMRR of 120dB. This error voltage is then inserted back into the input stage as part of the offset voltage, which is done by the voltage controlled voltage source E_{OS}. The inductors, L_3 and L_4, model the pole in the CMR versus frequency response of the OP-470.

OUTPUT STAGE

The output stage (Figure 1h) is modelled as an ideal output with an output resistance, R_{27} in parallel with R_{28}. An inductor, L_5, is added to model the rising output impedance of the emitter- follower output stage with frequency. The voltage sources V_5 and V_6, and diodes D_9 and D_{10}, combine to limit the voltage across the resistors, thus limiting the output current, I_{OUT} to 30mA.

PARAMETERS THAT ARE NOT MODELLED

To keep the OP-470 model as simple as possible and thus save computer and development time, not all features of the op amp were modelled as listed below:

- PSRR
- Crosstalk
- No Limits on Power Supply Voltages
- Maximum Input Voltage Range
- Temperature Effects (i.e., model parameters are assumed at 25°C)
- Input Noise Voltage and Current Sources
- Parameter Variations for Monte Carlo Analysis (i.e., all parameters are typical only)

These parameters are considered second-order effects and are not considered necessary for circuit simulation under normal operating conditions. However, users can easily add these functions as needed.

MODEL PERFORMANCE

The graphs on these next two pages compare the actua responses of the OP-470 to the SPICE analysis under various load and gain conditions. To show the accuracy of the model's frequency response, it was tested with different gains. Shown below are the gain-phase plots for an open-loop gain (Figures 3 and 4) and a closed-loop gain of –1 (Figures 5 and 6). The actua responses (part "a" of each figure), which were measured using a network analyzer, show very close agreement to the SPICE simulations (part "b" of each figure). This agreement holds true for both the gain and phase curves, which demonstrates the accuracy of the model's frequency response.

The model can also accurately simulate the OP-470's transient response. The comparisons to actual results were done with the amplifier in a a closed-loop gain of +1 with a 155pF capacitive load (260pF load for large signal) and a 2kΩ feedback resistor. The capacitive load was used to create the damped oscillations in the small and large-signal reponses. Because of the multiple poles and zeros, the model's small-signal transient response (Figure 7b) shows close resemblance to the actual reponse (Figure 7a). The large-signal reponse (Figure 8) again shows good accuracy. Notice in the large-signal reponse that the slew rate is accurately modelled, and that the "bump" before the amplifier starts slewing is also modelled. As can be seen from the figures, the model agrees closely with the actual device's performance. This accurate transient reponse, along with the accurate AC reponse, makes the OP-470 macro-model a powerful tool in analyzing circuit stability and AC behavior under various gains and load conditions.

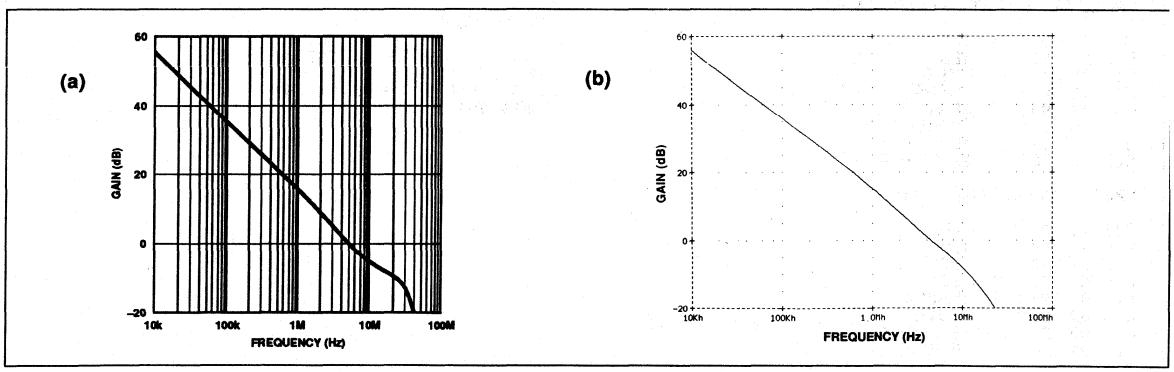

FIGURE 3: *Gain Plots, Open-Loop Gain (a. Actual, b. Simulated)*

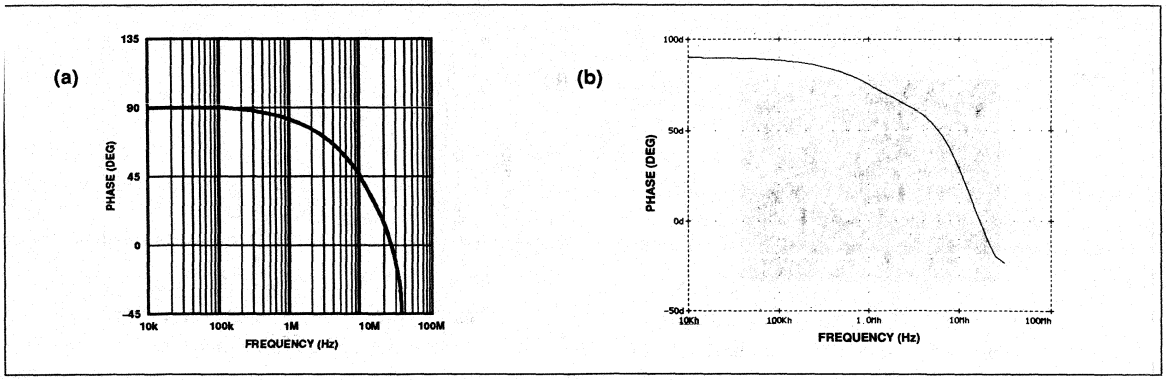

FIGURE 4: *Phase Plots, Open-Loop Gain (a. Actual, b. Simulated)*

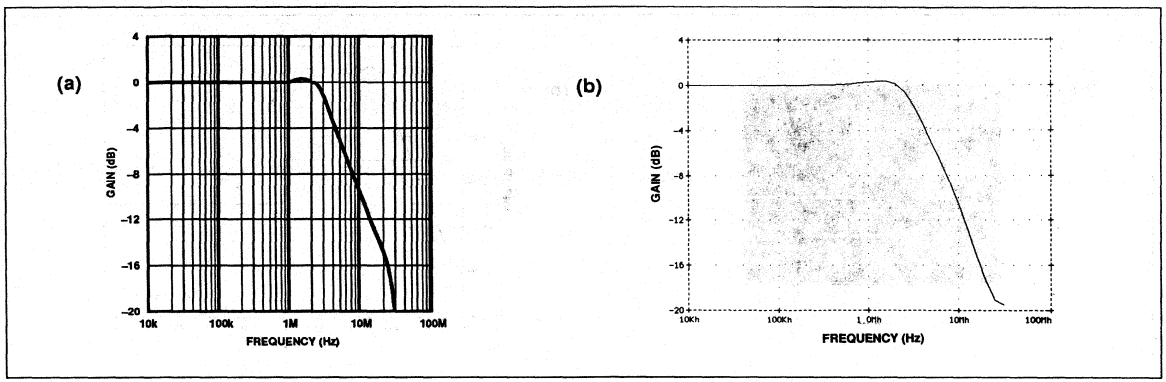

FIGURE 5: *Gain Plots, Closed-Loop Gain of –1 (a. Actual, b. Simulated)*

13

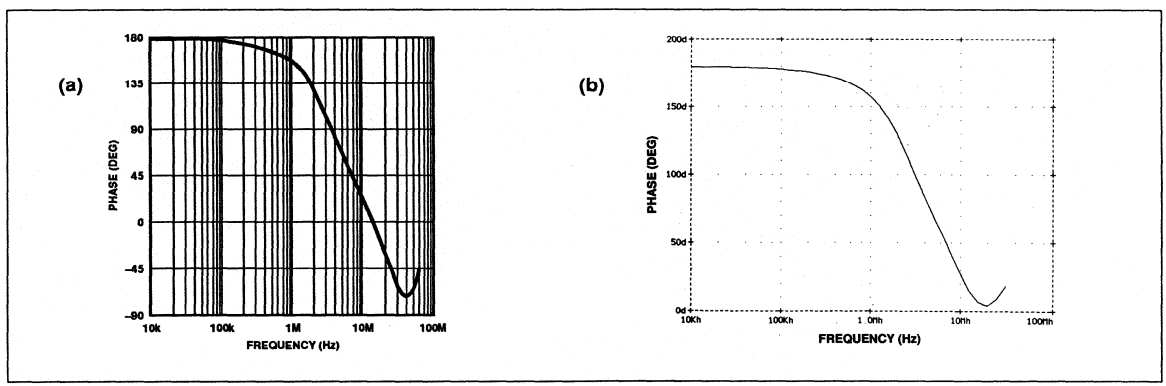

FIGURE 6: *Phase Plots, Closed-Loop Gain of –1 (a. Actual, b. Simulated)*

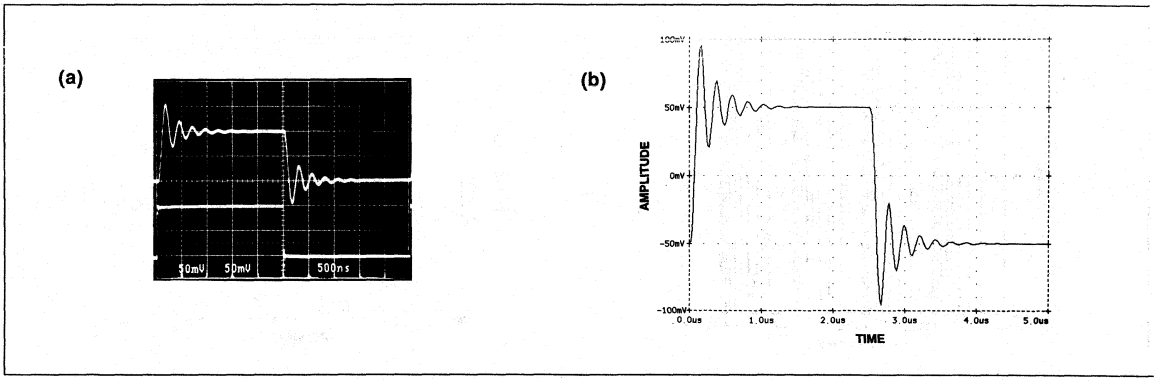

FIGURE 7: *Small-Signal Transient Response, Gain = +1, 155pF Capacitive Load (a. Actual, b. Simulated)*

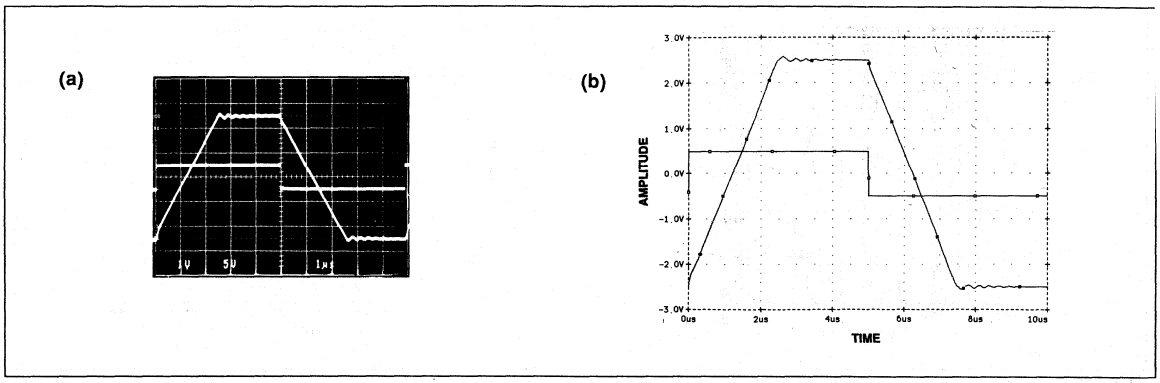

FIGURE 8: *Large-Signal Transient Response, Gain = +1, 260pF Capacitive Load (a. Actual, b. Simulated)*

ANALOG DEVICES

AN-138
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

SPICE-Compatible Op Amp Macro-Models

by Mark Alexander and Derek F. Bowers

There is a definite trend towards a comprehensive circuit-simulation approach. We believe that 75 percent of all installed circuit simulators are being used for system, rather than IC, design. Almost all of these simulators are variants of SPICE. With the growth of the electronics industry, system engineers have come to need increasingly accurate models for an ever larger number of integrated circuits, especially the ubiquitous operational amplifier. However, the increasing speed and complexity of these IC devices has caused problems that were never anticipated by the original developers of SPICE.

Because of the large number of active devices in a typical op amp, circuit simulations that use only transistor-level models can take an unacceptable amount of time, particularly when the circuit contains several op amps. Even simple models of semiconductor devices consume a large amount of computing time because of the multiplicity of nonlinear equations involved. In some cases, the time needed for a complete simulation might exceed the time necessary to build an engineering prototype. Obviously, such a situation would completely defeat the whole purpose of using SPICE.

Fortunately, you can reduce simulation time by using a macro-model that represents the op amp as accurately as possible without using large numbers of transistors or other nonlinear devices. However, it is quite a challenge to design a macro-model that, for all intents and purposes, exactly mimics the real device. For an op amp model to be of real use to the circuit designer, it must not only accommodate all important DC parameters, but also provide a reasonbly close approximation of the AC characteristics over a region that extends well beyond the unity-gain crossover frequency.

EXISTING MACRO-MODELS ARE INADEQUATE

Macro-models for many op amps already exist in the device libraries of several available software simulators. Most of these models are based on the original work done by Graeme Boyle and his colleagues (see Reference 1), who developed their macro-model during the mid-1970s to ease the CPU-time crunch on the already overloaded mainframe computers of the day. Boyle eliminated all but two transistors from his macro-model. The two remaining devices formed the differential-input stage of the op amp; all subsequent stages were implemented with linear controlled sources, passive components, and diodes. The transistors in the input stage were retained because they facilitated the simulation of real-world effects such as bias currents and variation of output dV/dt with the differential input voltage.

Because Boyle's method greatly reduces the number of overall nonlinear elements, the simulation time required per amplifier also decreases substantially. The Boyle structure is certainly an improvement over a full transistor-level simulation, but the structure still has several deficiencies, which prompted the development of the new macro-model. The deficiencies are as follows:

- The Boyle model provides only two poles (and no zeroes) for shaping the frequency response of the complete amplifier – a configuration that is barely adequate for slower op amps, and completely insufficient for today's faster devices.
- All internally generated node voltages are referenced to ground, even if the amplifier is "floated" with respect to ground. This configuration is not representative of the true operation of an op amp – almost none of the available devices provide a ground reference.
- The output-terminal current flows out of a controlled source connected to ground, instead of from the power-supply rails as it would in a real amplifier. This feature completely precludes the simulation of circuits that depend on the amp's output current splitting correctly between the supply rails.

IDEAL ELEMENTS CAN REDUCE COMPLEXITY

The circuit topology of the original Boyle model (Figure 1) was developed using two basic macro-modeling techniques (called simplification and build-up) that proved very useful in the development of the new macro-model as well.

The simplification technique successively reduces the complexity of major internal stages of the op amp by using simple ideal elements to replace real portions of the circuit. Therefore, you can expect a functional block that uses this approach to

Reprinted from EDN – February 15, 1990; and EDN – March 1, 1990

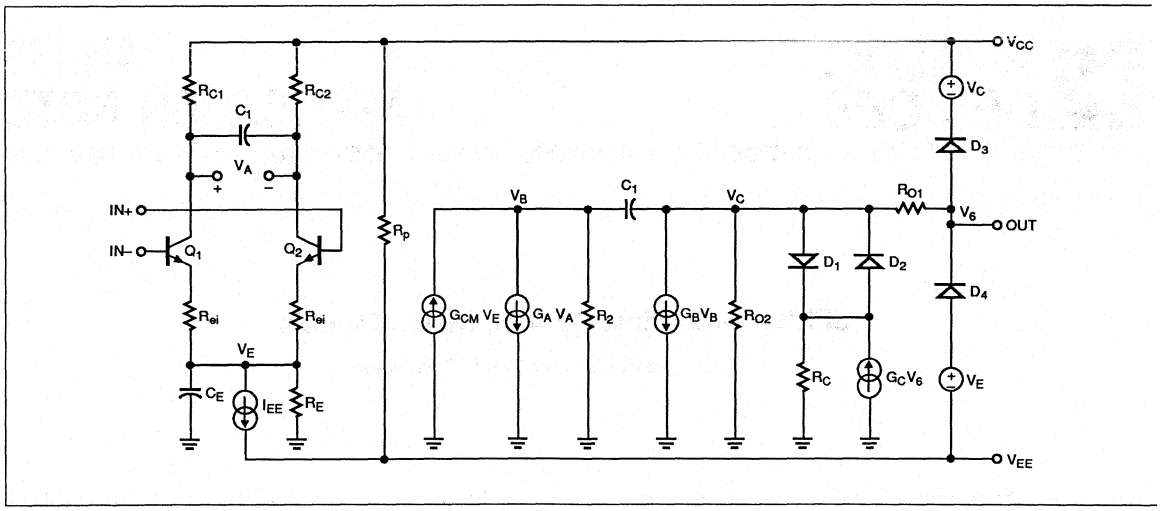

FIGURE 1: *A serious disadvantage to the Boyle op amp macro-model is that all voltages are referenced to ground.*

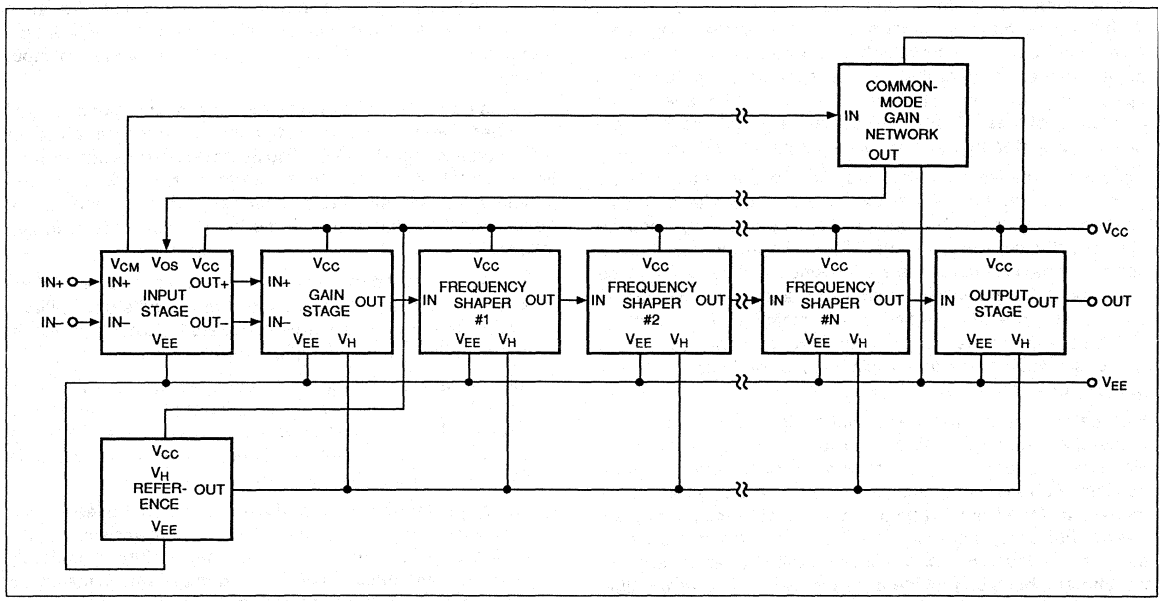

FIGURE 2: *The new op amp macro-model is inherently modular. You can cascade any of the building blocks to obtain any number of poles and zeroes in your op amp design.*

closely resemble the actual circuit. In Figure 1, the model of the input stage is a good example of simplification. The model retains the differential-input characteristics of an emitter-coupled pair, but eliminates any active loads; it replaces the tail-current source with an ideal element; and it assumes the task of generating the second amplifier pole. Adding a single capacitor (C_E) allows the model to provide a pole in this stage, and the reduction in overall component count makes the simulation run faster.

he build-up technique, on the other hand, lets you construct a ircuit block composed entirely of ideal elements, which very losely emulates the behavior of the real section of the device. nfortunately, the build-up technique often results in subsec- ons that bear little resemblance to their physical equivalents. igure 1's output stage is a good example: it provides the ecessary output voltage clipping, has the correct output resis- ance, and also provides short-circuit current limiting; but does not look like anything one would expect to find in the schematic of a real op amp.

DEVELOPMENT OF AN IMPROVED MACRO-MODEL

The impetus behind the new macro-model (Figure 2), arose out of the desire to create a model that operates like a real op amp. Yet it still had to be simple enough so that it would suffice as a generic model. Figures 3, 4 and 5 show that the model consists of several cascaded sections that process the input signal.

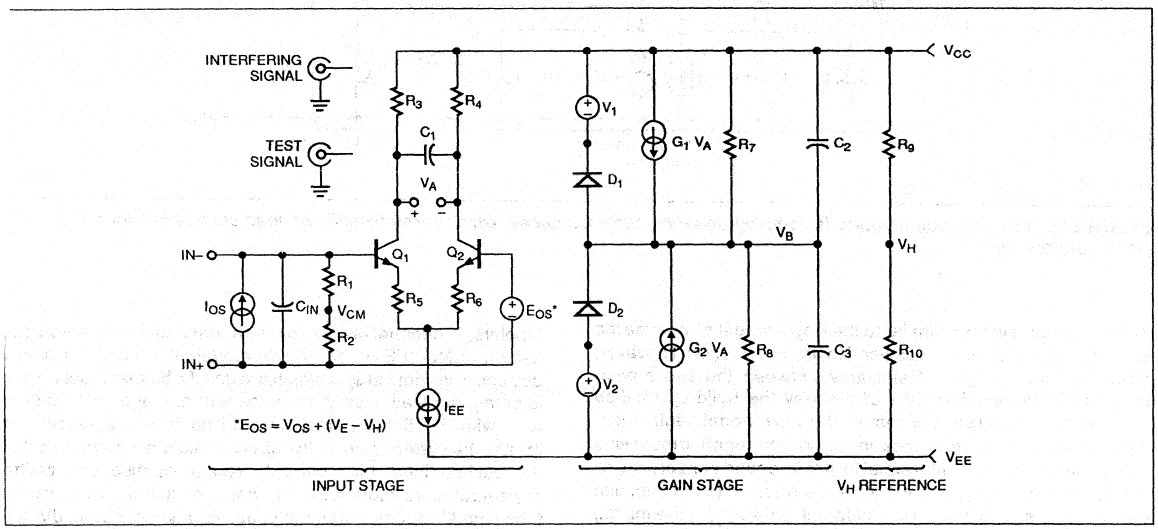

FIGURE 3: *The input stage of the new model resembles that of the Boyle, but all succeeding stages have a radically different structure.*

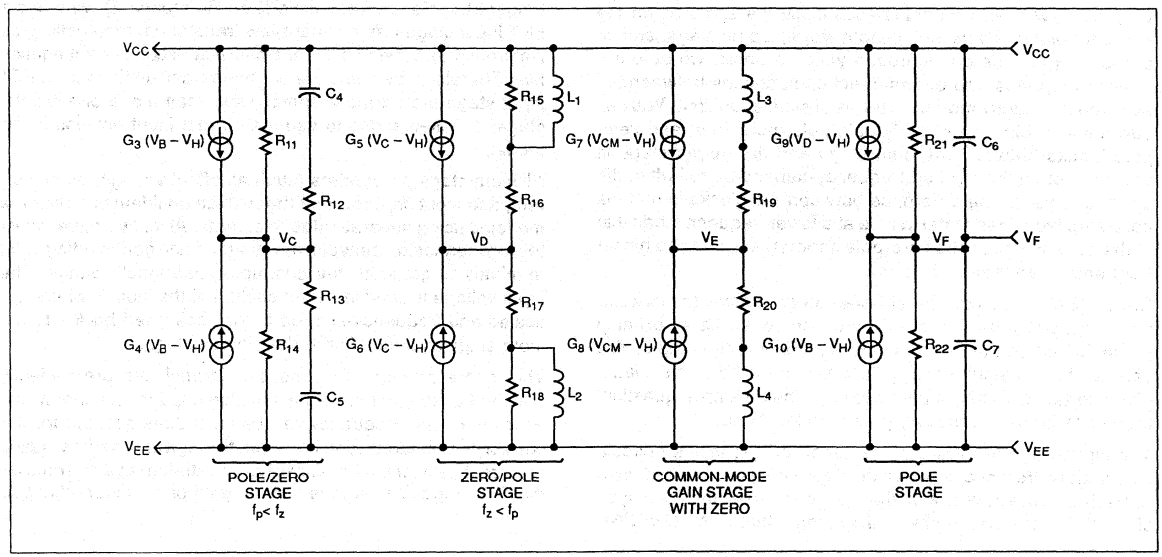

FIGURE 4: *Three types of frequency-shaping networks are available, in addition to a common-mode gain stage that provides a zero.*

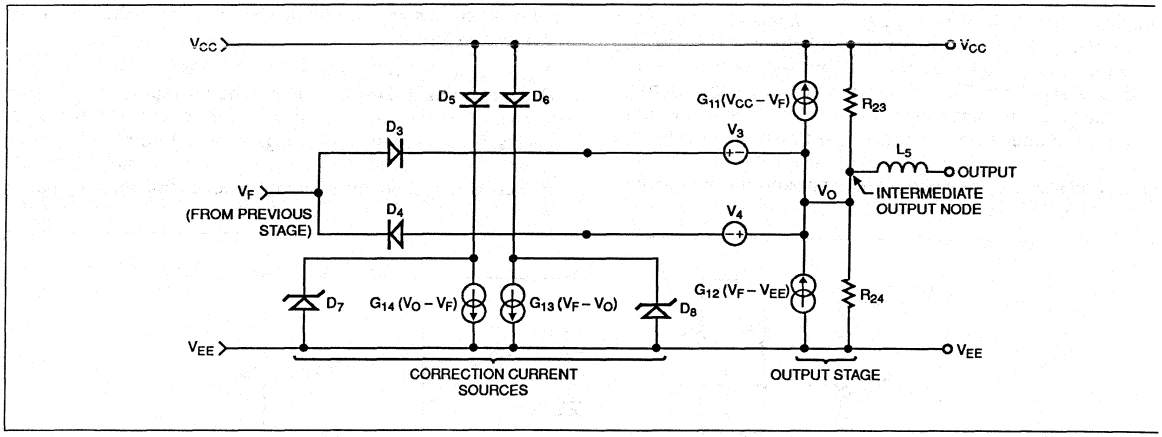

FIGURE 5: *The new output stage is complemented by current sources, which correctly split the load current between the power-supply rails.*

The input stage is very similar to the Boyle model's because the simplification technique was used in its construction. However, after the input stage, all similarity between the two macro-models disappears because of the way the build-up method was used to generate the rest of the new model. Notice that there is no ground reference in any of the signal-processing blocks. Instead, after differential to single-ended conversion, all internally generated node voltages are referred to the midpoint between the supply rails. This midpoint, called V_H in the model, is generated by two equal resistors connected between the supply rails.

The minimum requirement for modeling any particular op amp with the new macro-model is essentially the same as for the Boyle topology: a differential-input stage, a gain stage, and an output stage. This configuration yields a basic two-pole frequency response and allows direct comparisons between the two types of macro-models in terms of simulation time. You can add any combination of unity-gain pole, pole-zero, and zero-pole blocks between the gain stage and the output stage in order to obtain the desired frequency-dependent roll-off of the open-loop gain. The difference between the blocks is that the pole-zero block generates a pole at a lower frequency than that of the zero, whereas the zero-pole generates a pole at a higher frequency than that of the zero.

Box 2, "Calculation of Model Parameters," shows the calculations that you must perform in order to construct an op amp model based on the building blocks of Figures 3, 4 and 5. You can do these calculations quite easily on a hand calculator, given certain data sheet parameters of the op amp in question, together with the necessary pole-zero locations.

The input stage in Figure 3 is a simplified 2-transistor circuit. One major difference between the new model and its predecessor is that, in the new model, the input stage uses the same type of input devices as the physical op amp – that is, NPN or PNP bipolars, P-channel JFETs (or N-channel devices where appli cable), or MOSFETs. The Boyle model allows *only* for bipola devices in its input stage, which are fine if a bipolar input op amp is being modeled. However, when you model a FET input op amp with the Boyle technique, you have to dramatically in crease the current gain of the input transistors in order to obtain the desired input bias current. You must also use emitte degeneration to reduce their transconductance. The result o these modifications is that the usual variation of output dV/dt o a FET input amplifier over a fairly wide input-differential voltage range (typically 1 to 2V) will not be correctly simulated. A degenerated bipolar input stage has a linearized, hyperbolic tangent transfer characteristic (See Reference 2), whereas a FET input stage has a square-law transfer characteristic (See Reference 3). Obviously, these characteristics are not equivalent. Therefore, because the parameter calculations for a FET input stage are no more complicated than those of a bipolar stage, it makes sense to use the correct input devices in the model.

All input-stage parameters (such as offset voltage, offset current, and input capacitance) that exhibit nonideal behavior are modeled using separate ideal elements. Also, two equal resistors are connected between the inverting and noninverting input terminals to generate the common-mode input voltage. The input voltage is used in a later section of the model, where it is scaled and frequency-shaped before being fed back into the input stage as a modifier to the offset voltage.

The model assumes that the input transistors are perfectly matched and do not have any junction capacitance that would alter the overall frequency response. It does account for the correct input-bias current, however, through appropriate choice of current gain for the bipolar stage or gate-leakage currents for the FET stage. You set the voltage gain of the differential pair

THE EVOLUTION OF SPICE SIMULATORS

The electrical-circuit simulator SPICE, and the more powerful version called SPICE2, originally emanated from the University of California, Berkeley, during the 1970s (see Reference 4). Primarily written for the purpose of assisting design engineers with the analysis of integrated circuits at the transistor level (hence the acronym *S*imulation *P*rogram with *I*ntegrated *C*ircuit *E*mphasis), SPICE lets you use a computer to evaulate your designs more quickly and more thoroughly than is possible by means of laborious hand calculations. The popularity of SPICE soon spread to the system-level-design community for the same reasons that the IC designers embraced it.

The original version of SPICE was a public-domain program available at a purely nominal charge; however, many software vendors have recognized the need for a fully supported, adapted, and improved commercial circuit simulator. The first mainframe-based versions of such programs included HSpice from Meta-Software, I-Spice from NCSS timesharing and PRECISE from Electronic Engineering Software. Recently, most mainframe versions have been adapted for use on workstations, and some have been adapted for IBM PCs and compatibles.

The first PC-based version of SPICE was PSpice from MicroSim Corp. It was followed by others, such as IS-Spice from Intusoft. Other companies, including Analogy Inc (which offers a behavioral-simulation package known as Saber), have chosen to depart from the conventional SPICE format of using "boxed" circuit elements to construct models. Instead, Saber relies on rigorous defining equations written in a specific modeling language, called Mast, to control the behavior of any desired electrical circuit model.

› unity by making the load resistor value equal to the reciprocal f the transconductance of the transistors. This assumption implifies the calculations to determine the slew-rate-limiting omponents. Tail current for the input stage is nominally set to mA for convenience; however, it can be scaled down to 100μA, 0μA, or 1μA depending on the amp's total quiescent current.

;AIN STAGE FEATURES

'he model's open-loop gain is normally achieved in a single tage (See Figure 3), which consists of two voltage-controlled :urrent sources, two resistors, two capacitors, and a voltage-miting network. The conversion of signals from differential to ingle-ended form also takes place during this stage. The 'oltage-limiting network consists of a pair of diodes, each :onnected to its own voltage source. The network prevents the ain stage and the other internal nodes of the model from ;winging beyond the power-supply-rail voltages during an in-›ut-overdrive condition. Voltage limiting *must* take place in the ›pen-loop gain stage; otherwise, succeeding nodes could at-empt to simulate the generation of huge (hundreds of kilovolts) ;ignals.

Two capacitors, connected in parallel with the resistors, deter-nine both the dominant amplifier pole and the slew rate. At ›resent, the macro-model can handle only symmetrical positive ind negative slew rates, because symmetry is the easiest :ondition to simulate. However, future enhancements may allow for some variation between them. Finally, to each of the wo voltage-controlled current sources, the stage adds a DC :omponent that makes up the bulk of the amp's quiescent ;upply current.

nvestigation of op amp frequency response reveals that, in nost cases, accurate simulation of the gain and phase variation ›f real devices at high frequencies requires more than two ›oles. Further, different types of op amps have varying numbers of poles and zeroes. To allow each of these diverse types to be easily converted to a SPICE-compatible subcircuit - without having to start from scratch every time – a truly general model would have to be highly modular and permit arbitrarily large numbers of poles and zeroes. Therefore, the final structure uses a few basic building blocks that are common to all individual op amp models. These blocks are shown in Figure 4.

All of the frequency-shaping blocks have unity-gain at DC, because the g_m of each voltage-controlled current source (VCCS) is equal to the reciprocal of the resistance connected from each node of the VCCS (voltage-controlled current source) to the power-supply rails. This topology is advantageous because during model generation for a specific amplifier, you can comment out separate poles or pole-zero pairs. You can then see their effect, individually, on the net frequency response of the amplifier, so that pole-zero tweaking becomes rather easy. Because all the frequency-shaping blocks have unity-gain at DC, the procedure does not alter the model's DC open-loop gain.

The common-mode gain stage in Figure 4 consists of two VCCSs that drive two equal resistors, each resistor is connected in series with an inductor to one of the supply rails. The inductors simulate the typical fall-off of CMRR that most amplifiers exhibit as the input frequency increases. The input common-mode voltage, relative to the V_h node, controls the current sources. Each controlled source has a g_m equal to the reciprocal of the associated resistor value divided by the CMRR of the amplifier at DC.

Accordingly, the gain from the input common-mode network to the internal common-mode gain node is equal to the reciprocal of the amplifier's CMRR. (The term "gain" is a misnomer here, because the common-mode gain has a value that is much less than unity).

The inductors add a zero to the common-mode gain, which is the same as adding a pole to the CMRR. The common-mode voltage, after being scaled and appropriately frequency-shaped, is then added back to the input stage as dictated by theory. This step is done by making the offset-voltage source in the input stage a unity-gain voltage-controlled voltage source, which has a DC component equal to the amplifier's V_{OS}.

The operation of the output stage in Figure 5 is not entirely obvious. The internal op amp output signal, after receiving all the appropriate frequency shaping, appears as a voltage referenced to V_h at the last node prior to processing by the output stage. The two voltage-controlled current sources in the output block drive two equal resistors connected to the supply rails, as with the other blocks. Here, however, the g_m of the two controlled sources is arranged so that they act as active current generators. Consequently, each g_m source generates just enough current to provide the desired voltage drop across its paralleled resistor.

When there is no load on the output, the model draws no current from either power-supply rail. It thus behaves somewhat like an ideal, unity-gain, class-B output stage with no crossover distortion. Because the two resistors are each equal to twice the open-loop output resistance, the output stage appears to act as a voltage source referenced to V_h with the correct DC output resistance. Simulating the right output resistance means that the DC open-loop gain will be properly reduced as the amplifier is loaded.

BOX 2

CALCULATION OF MODEL PARAMETERS

The following equations allow you to create an improved macro-model for simulation of any op amp. The calculations are given separately for each of the available building blocks, and some power supply considerations are discussed.

INPUT- AND GAIN-STAGE CALCULATIONS

a. General Calculations

Refer to Figure 3 to identify components and signals mentioned here. First, Choose I_{EE} such that it is somewhat less than the amp's total quiescent current. For convenience, you may set I_{EE} to 1mA, 100µA, 10µA, or 1µA. Then,

$$C_2 = C_3 = \frac{I_{EE}}{\text{SLEW RATE}}$$

$$R_7 = R_8 = \frac{1}{2\pi f_{p1} C_2},$$

were f_{p1} = dominant amplifier pole.

$$G_1 = G_2 = \frac{A_{VOL}}{R_7},$$

where A_{VOL} = open-loop DC gain

$$R_3 = R_4 = \frac{1}{G_1}$$

$$C_1 = \frac{1}{4\pi f_{p2} R_3},$$

where f_{p2} = second amplifier pole

$V_1 = V_{CC} - (+V_{OUT\ MAX}) + V_T \ln(2I_{EE}/I_S)$

$V_2 = (-V_{OUT\ MAX}) - V_{EE} + V_T \ln(2I_{EE}/I_S)$

$V_T = 0.02585V$ at $T = 27°C$

$I_S = 1 \times 10^{-12}$ A (FOR BOTH DIODES).

You can substitute some data sheet parameters directly into the model. These parameters are:

E_{OS} = Input Offset Voltage (DC component only);

I_{OS} = Input Offset Current; C_{IN} = Input Capacitance.

b. Bipolar Input-Stage Calculations

First, you must evaluate the following equation to determine whether the op amp in question can be modeled using the new macro-model:

$$A_{VOL} \le \frac{\text{SLEW RATE}}{4\pi f_{p1} V_T},$$

Where $V_T = 0.02585V$ at 27°C.

If the equation holds true, then you may proceed with the rest of the calculations. If not, then you must modify the model to accommodate this particular op amp.

$$R_5 = R_6 = R_3 - \frac{2 V_T}{I_{EE}}$$

$$\beta_F = \frac{2 I_{BIAS}}{I_{EE}},$$

where β_F is the forward current gain of the input transistors, and I_{BIAS} is the input bias current.

$$R_1 = R_2 = \frac{1}{2\left(\frac{1}{R_{ID}} - \frac{1}{2 \beta_F R_3}\right)} \le 5 \times 10^{11}\ \Omega,$$

where R_{ID} is the differential input resistance. If R_{ID} is not a specified data sheet parameter, then set both R_1 and R_2 to the value $5 \times 10^{11}\Omega$.

c. JFET Input-Stage Calculations

If your design has a JFET input stage, use the default value of –2.000V for the gate-to-source cutoff voltage V_{TO}. Also, change the name of the first-stage current source to I_{SS}. The main calculation is to determine β, the JFET gain factor:

$$\beta = \frac{(G_{1/2})^2}{2\ I_{SS}},$$

where I_{SS} is the first-stage tail current.

For maximum output dV/dt, the tail current must originate from one side only of the differential pair; this condition requires a differential input voltage equal to:

$$V_{ID} = \frac{\sqrt{2}\ (\text{SLEW RATE})}{2\pi\ A_{VOL}\ f_{p1}}$$

Also, the input bias current is composed of the gate-drain and gate-source leakage currents. Thus,

$$I_S = \frac{I_{BIAS}}{2},$$

where I_{BIAS} is the input bias current at 27°C. Further,

$$R_1 = R_2 = \frac{R_{ID}}{2},$$

where R_{ID} is the differential input resistance (normally 1 x $10^{12}\Omega$.

Finally, you can set the values of R_5 and R_6 to zero, because degeneration is not normally needed with JFET input amplifiers.

Frequency-Shaping-Stage Calculations

To identify parameters of the frequency-shaping stages, refer to Figure 4. In all three types of the frequency-shaping stage, set G_3 and G_4 to I x 10^{-6} times A/V, for convenience. Further, f_z is the zero frequency and f_p is the pole frequency.

Then for the pole-zero stage,

$$R_{11} = R_{14} = 1 \times 10^6$$

$$R_{12} = R_{13} = \frac{R_{11}}{f_z/f_p - 1};$$

$$C_4 = C_5 = \frac{1}{2\pi f_z R_{12}}$$

for the zero-pole stage,

$$R_{16} = R_{17} = 1 \times 10^6$$

$$R_{15} = R_{18} = (f_p/f_z - 1)\, R_{16};$$

$$L_1 = L_2 = \frac{R_{15}/_{18}}{2\pi f_p}$$

for the pole stage,

$$R_{21} = R_{22} = 1 \times 10^6$$

$$C_6 = C_7 = \frac{1}{2\pi f_p R_{21}}.$$

Common-Mode-Gain-Stage Calculations

To identify the parameters of the common-mode gain stage, refer to Figure 4.

$$R_{19} = R_{20} = 1 \times 10^6$$

$$G_7 = G_8 = \frac{1}{R_{19} \times CMRR},$$

$$L_3 = L_4 = \frac{R_{19}}{2\pi f_{p(CM)}}$$

where f_{CM} is the common-mode pole.

Output-Stage Calculations

To identify the parameters of the output stage, refer to Figure 5. The breakdown voltage of diodes D_7 and D_8 is nominally set to 50V. The value of inductor L_5 is determined by experiment. R_{OUT} is that open-loop output resistance; V_T is 0.02585V at 27°C; and I_S is 1 x 10^{-12}A for all diodes. Then,

$$G_{11} = G_{12} = G_{13} = G_{14} = \frac{1}{2\ R_{OUT}}$$

$$R_{23} = R_{24} = 2 \times R_{OUT}$$

$$V_3 = I_{SC}(+VE)R_{OUT} - V_T \ln(20 \times 10^{-6}/I_S)$$

$$V_4 = |I_{SC}(+VE)R_{OUT}| - V_T \ln(20 \times 10^{-6}/I_S)$$

You can determine the values of resistors R_9 and R_{10} in Figure 3 by means of the following equation:

$$R_9 = R_{10} = \frac{1}{2(dI_{SY}/dV_{SY})},$$

where dI_{SY}/dV_{SY} represents the variation of supply current caused by a change in the supply voltage. The total quiescent current that flows between V_{CC} and V_{EE} in the model is thus:

$$I_{SY} = I_{EE} + (N+1)\left(\frac{V_{CC} - V_{EE}}{2R}\right) + I_{DC} + \left(\frac{V_{CC} - V_{EE}}{R_9 + R_{10}}\right),$$

where N is the total number of frequency-shaping and common-mode gain stages in the model; R is normally 1 x $10^6\Omega$; this value, along with the transconductance of the G-sources in the frequency shaping stages, can be appropriately scaled for low power op amps;

and I_{DC} is the DC offset added to the G_1 and G_2 sources to make up the difference between the current drawn by the rest of the model and the quiescent current of the

A subtle problem exists with this simple push-pull output stage, however. Regardless of whether that stage is sinking or sourcing current, the load current is always split equally between both rails – not exactly what a real output stage would do. So, with a sourced-load current, for example, the net positive supply current increases by only half the amount flowing in the load. The negative supply current decreases by the same amount. To compensate for this anomaly, you force a current from the positive rail to the negative rail exactly equal to half the load current. This correction current must always flow in the same direction – even if the output current reverses polarity.

The effect, therefore, of the two correction sources in Figure 5 is to force a unipolar compensating current between the power-supply rails that is equal to half the output current. In SPICE, because there is no easy way to implement an absolute value VCCS, there must be two linear correction sources – one for each half cycle of output current. The diodes in series with each

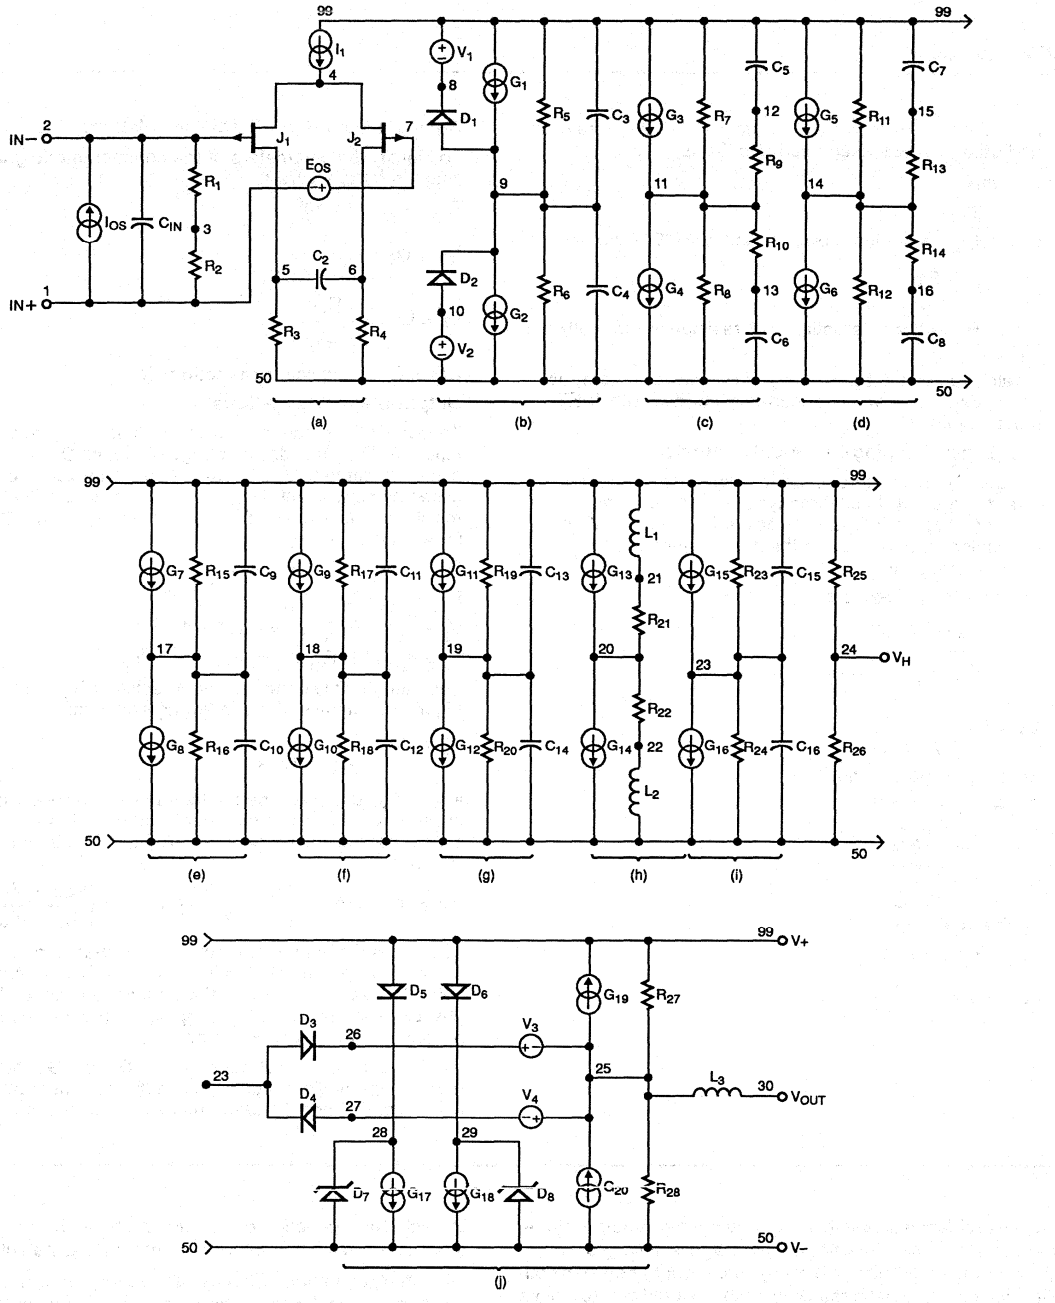

FIGURE 6: *The OP-42 macro-model is more complex than its Boyle counterpart, requires more time for simulation, but pays off in greatly improved accuracy.*

ISTING 1: OP-42 SPICE Macro-Model Net List

```
OP-42 MACRO-MODEL  ©PMI 1990
*
* SUBCKT OP-42 1  2 30  99  50
*
* INPUT STAGE & POLE AT 15.9 MHZ
*
R1    1   3        5E11
R2    2   3        5E11
R3    5   50       707.36
R4    6   50       707.36
CIN   1   2        5E-12
C2    5   6        7.08E-12
I1    99  4        1E-3
IOS   1   2        4E-12
EOS   7   1        POLY(1) 20 24 1E-3 1
J1    5   2   4    JX
J2    6   7   4    JX
*
* SECOND STAGE & POLE AT 45 HZ
*
R5    9   99       176.84E6
R6    9   50       176.84E6
C3    9   99       20E-12
C4    9   50       20E-12
G1    99  9        POLY(1) 5 6 3.96E-3 1.4137E-3
G2    9   50       POLY(1) 6 5 3.96E-3 1.4137E-3
V1    99  8        2.5
V2    10  50       3.1
D1    9   8        DX
D2    10  9        DX
*
* POLE-ZERO PAIR AT 1.80 MHZ/2.20 MHZ
*
R7    11  99       1E6
R8    11  50       1E6
R9    11  12       4.5E6
R10   11  13       4.5E6
C5    12  99       16.1E-15
C6    13  50       16.1E-15
G3    99  11       9  24  1E-6
G4    11  50       24  9  1E-6
*
* POLE-ZERO PAIR AT 1.80 MHZ/2.20 MHZ
*
R11   14  99       1E6
R12   14  50       1E6
R13   14  15       4.5E6
R14   14  16       4.5E6
C7    15  99       16.1E-15
C8    16  50       16.1E-15
G5    99  14       11 24 1E-6
G6    14  50       24 11 1E-6
*
* POLE AT 53 MHZ
*
R15   17  99       1E6
R16   17  50       1E6
C9    17  99       3E-15
C10   17  50       3E-15
G7    99  17       14 24 1E-6
G8    17  50       24 14 1E-6
*
* POLE  AT 53 MHZ
*
R17   18  99       1E6
R18   18  50       1E6
C11   18  99       3E-15
C12   18  50       3E-15
G9    99  18       17 24 1E-6
G10   18  50       24 17 1E-6
*
* POLE AT 53 MHZ
*
R19   19  99       1E6
R20   19  50       1E6
C13   19  99       3E-15
C14   19  50       3E-15
G11   99  19       18 24 1E-6
G12   19  50       24 18 1E-6
*
*COMMON-MODE GAIN NETWORK WITH ZERO AT 100 KHZ
*
R21   20  21       1E6
R22   20  22       1E6
L1    21  99       1.5915
L2    22  50       1.5915
G13   99  20       3  24 1.58E-11
G14   20  50       24  3 1.58E-11
*
*POLE AT 79.6 MHZ
*
R23   23  99       1E6
R24   23  50       1E6
C15   23  99       2E-15
C16   23  50       2E-15
G15   99  23       19 24 1E-6
G16   23  50       24 19 1E-6
*
*OUTPUT STAGE
*
R25   24  99       111.1E3
R26   24  50       111.1E3
R27   25  99       90
R28   25  50       90
L3    25  30       2.5E-7
G17   28  50       23 25 11.1111E-3
G18   29  50       25 23 11.1111E-3
G19   25  99       99 23 11.1111E-3
G20   50  25       23 50 11.1111E-3
V3    26  25       0.7
V4    25  26       0.7
D3    23  26       DX
D4    27  23       DX
D5    99  28       DX
D6    99  29       DX
D7    50  28       DY
D8    50  29       DY
*
* MODELS USED
*
.MODEL JX PJF(BETA=999.3E-6 VTO=-2.000 IS=8E-11)
.MODEL DX   D(IS=1E-15)
.MODEL DY   D(IS=1E-15 BV=50)
.ENDS OP-42
```

source perform half-wave rectification, and the zeners ensure that there is always a conductive path for each source when its current reverses direction. The net result of all these additional elements is an output stage model whose DC behavior very closely mimics that of the phycial circuit.

To account for the typical rise in emitter-follower output stage impedance with frequency, the macro-model includes an output inductor connected between the intermediate output node and the actual macro-model output node (see Figure 5). You determine the value of this inductor, on a trial-and-error basis, by using capacitive loads on the model until the amount of overshoot is very close to that which you observed with the real op amp and the same load.

Short-circuit current limiting is also a necessary feature of any good op amp macro-model. In Figure 5, limiting is accomplished by clamping the output voltage (V_F) from the previous frequency-shaping stage to the intermediate output node (V_0), using diodes D_3 and D_4, and voltage sources V_3 and V_4. Remember that the signal from the previous stage is always equal to the ideal output voltage with no load, and that the output stage behaves as a voltage source with a finite output resistance. The action of the diodes and voltage sources is then equivalent to clamping the voltage drop across the effective output resistance. You can obtain the required output current limiting by an appropriate choice of each voltage source.

Because the main goal of the new structure is to provide improved AC accuracy, the model must also correctly represent the common-mode behavior. So, the modeling team selected the PMI OP-42, a JFET-input op amp, as its first guinea pig, largely because the Boyle model cannot properly accommodate a JFET input stage. Although the team had to work out all the equations pertaining to the JFET input stage before they

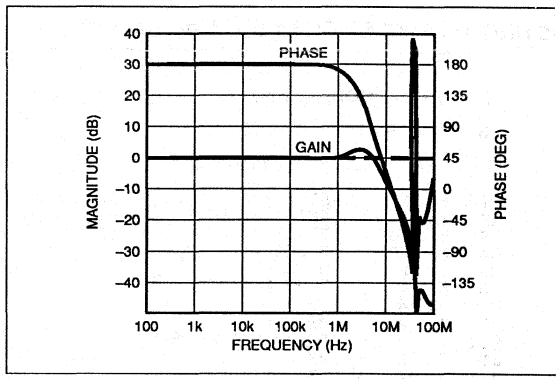

FIGURE 7: *When you connect the OP-42 in a unity-gain, inverting configuration, the gain response shows a slight peak at about 6MHz; there is a rapid increase in phase shift above 2MHz.*

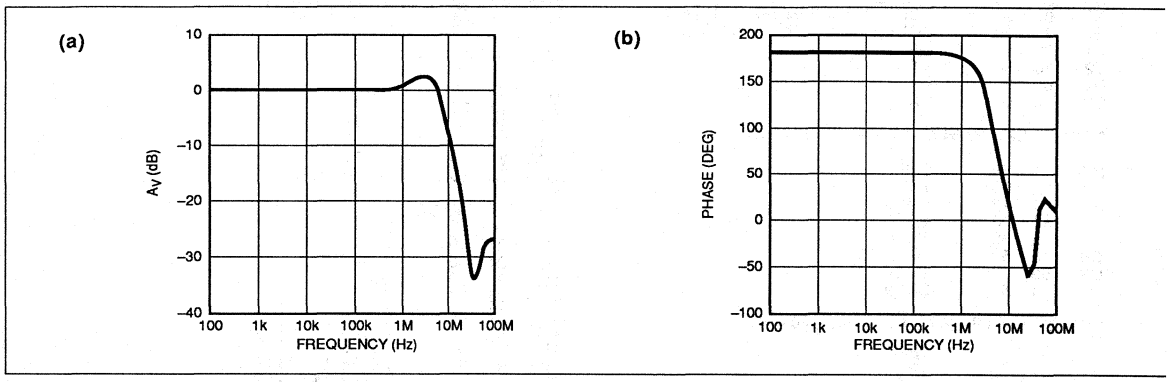

FIGURE 8: *Using the new macro-model, the simulated gain response (a) of the OP-42 is very like that of the real device, with a slight peak at 4MHz. The phase-response (b) is very good. This curve closely follows that of the real device.*

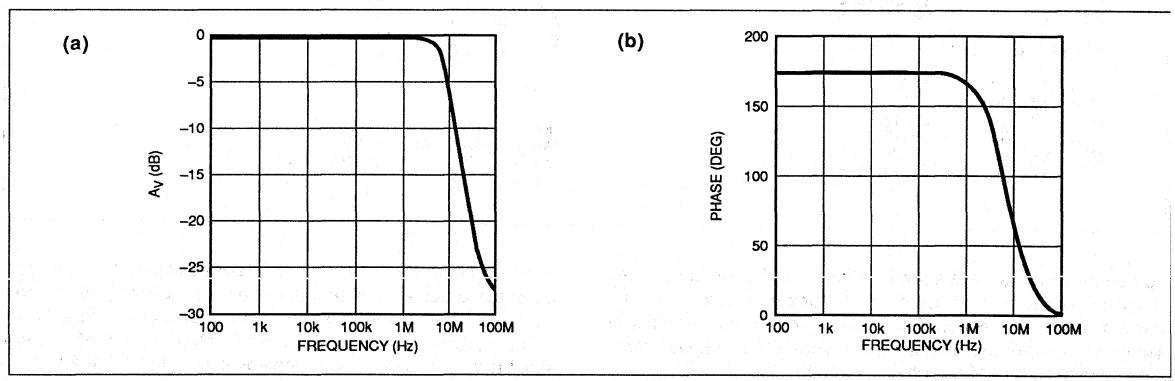

FIGURE 9: *The Boyle model of the OP-42 (a) does not show the amplitude peak at 4MHz that is characteristic of the real device. The phase response (b) also is not very accrurate, especially in the regoin beyong 10MHz.*

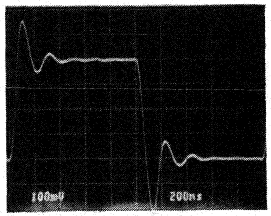

FIGURE 10: *An OP-42 with a 430pF capacitive load shows both overshoot and undershoot when driven with a 500kHz, 200mV peak square wave.*

could test the complete model, this stage turned out to be fairly easy to handle mathematically and did not hinder the development of the final macro-model structure.

Figure 6 shows the resulting implementation. The physical OP-42 has a gain-bandwidth product of approximately 10MHz and a symmetrical slew rate of 50V/µs. The CMRR-versus-frequency curve of this amplifier indicates that a zero at about 100kHz is necessary in the model's common-mode gain stage.

Listing 1 shows the net list for the OP-42 macro-model, which has 8 poles, 2 zeroes, plus a zero in the common-mode gain stage at 100kHz. The model of even a relatively stable amplifier needs that many poles and zeroes in order to accurately mimic the gain and phase behavior of the physical device at high frequencies.

Inspection of the output-stage section of the net list shows that the open-loop output resistance is 45Ω. A 250nH inductor, connected in series with the output terminal, compensates for the rise in effective open-loop output impedance at high frequencies. The current-limiting network formed by diodes D3 and D4 and voltage sources V_3 and V_4 clamps the maximum output current at approximately ±30mA.

SIMULATION-ACCURACY COMPARISONS

Figure 7 shows the gain and phase response of a physical OP-42 connected as an inverting, unity-gain amplifier that has 1kΩ input and feedback resistors and runs from ±15V supplies. You can see a small amount of peaking (about 2dB) in the closed-loop gain curve, and the phase shift increases rapidly above 2MHz. Figures 8a and 8b show the gain and phase response of the new OP-42 macro-model under the same conditions. The gain response shows the same amount of closed-loop peaking as that of the real circuit; the phase response almost exactly matches that of the real device to at least 10MHz.

Figures 9a and 9b, which show the corresponding output curves from the Boyle implementation, clearly demonstrate the deficiencies in the Boyle model's response accuracy. The gain response does not show the 2dB peak, indicates too steep a roll-off, and is quite inaccurate above 10MHz. The Boyle model's phase response does not even come close to the real circuit's response. The OP-42 macro-model, with its multiple pole-zero complement, emulates the AC response of the actual circuit more accurately.

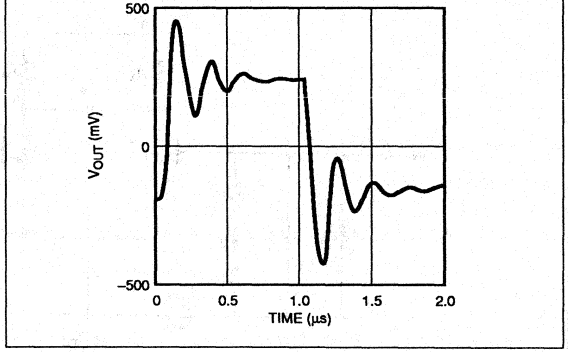

FIGURE 11: *The new macro-model's simulation of an OP-42 with a capacitive load of 430pF shows the symmetrical nature of the model's output stage.*

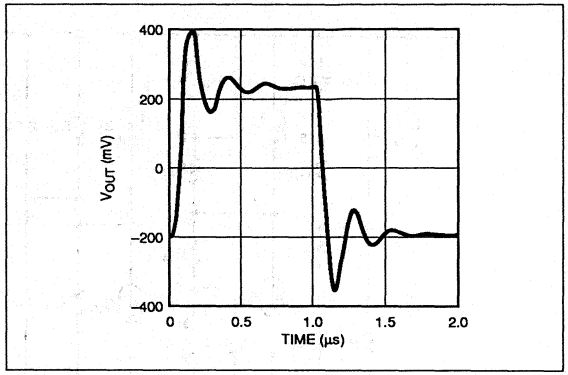

FIGURE 12: *The Boyle model of the OP-42 simulates approximately the right amount of overshoot, but its ringing frequency is too low.*

Figure 10 shows the measured transient response of the inverting, unity-gain OP-42 amplifier with a 430pF capacitive load. For a 400mV_{p-p} input signal, there is about 75 percent overshoot and 100 percent undershoot. The simulation results from the new macro-model (see Figure 11) show about 115 percent of both overshoot and undershoot. This simulated value is quite close to the actual value on the negative half of the waveform, but differs from the actual value on the positive half. The expanation for this anomaly is that although the new macro-model has a perfectly symmetrical output stage, the op amp being modeled may not. The OP-42, in fact, has an asymmetrical, all NPN-transistor output stage. As a result, the high-frequency, open-loop response is variable and depends on whether the output stage is sinking or sourcing current.

The Boyle configuration, too, models an op amp's output stage as a perfectly symmetrical voltage source and, as Figure 12 shows, it incorrectly simulates the undershoot on the negative half of the output waveform. It does come reasonably close on

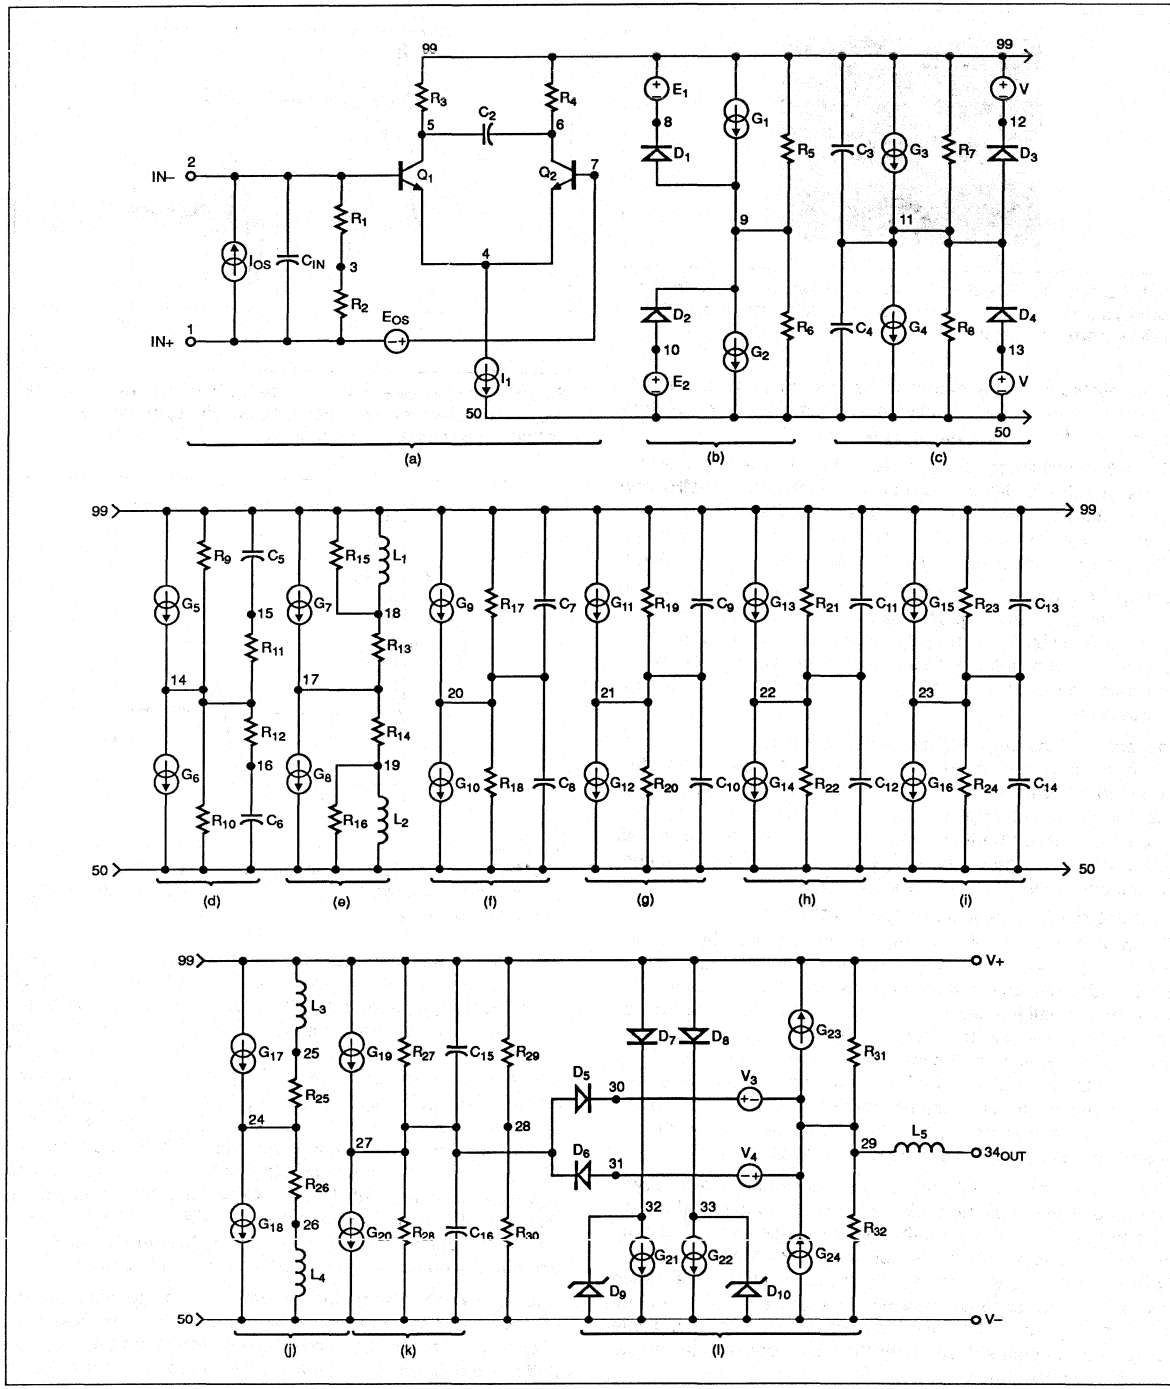

FIGURE 13: *The model schematic of the OP-61 looks similar to that of the OP-42, except that it has an additional gain stage.*

LISTING 2: OP-61 SPICE Macro-Model Net List

```
OP-61 MACROMODEL © PMI 1989
*
* SUBCKT OP-61  1  2 34 99 50
*
* INPUT STAGE & POLE AT 300 MHZ
*
R1    1    3        5E11
R2    2    3        5E11
R3    5    99       51.6
R4    6    99       51.6
CIN   1    2        5E-12
C2    5    6        5.141E-12
I1    4    50       1E-3
IOS   1    2        2E-7
EOS   7    1        POLY(1) 24 28 400E-6 1
Q1    5    2    4   QX
Q2    6    7    4   QX
*
* FIRST GAIN STAGE
*
R5    9    99       1E6
R6    9    50       1E6
G1    99   9        5 6 2E-4
G2    9    50       6 5 2E-4
E1    99   8        POLY(1) 99 28 -4.4 1
E2    10   50       POLY(1) 28 50 -4.4 1
D1    9    8        DX
D2    10   9        DX
*
* SECOND GAIN STAGE & POLE AT 2.5KHZ
*
R7    11   99       5.1598E6
R8    11   50       5.1598E6
C3    11   99       12.338E-12
C4    11   50       12.338E-12
G3    99   11       POLY(1)  9 28 4.24E-3 9.69E-5
G4    11   50       POLY(1) 28  9 4.24E-3 9.69E-5
V1    99   12       2.3
V2    13   50       2.3
D3    11   12       DX
D4    13   11       DX
*
* POLE-ZERO PAIR AT 4MHZ / 8MHZ
*
R9    14   99       1E6
R10   14   50       1E6
R11   14   15       1E6
R12   14   16       1E6
C5    15   99       19.89E-15
C6    16   50       19.89E-15
G5    99   14       11 28 1E-6
G6    14   50       28 11 1E-6
*
* ZERO-POLE PAIR AT 85MHZ / 300MHZ
*
R13   17   18       1E6
R14   17   19       1E6
R15   18   99       2.529E6
R16   19   50       2.529E6
L1    18   99       1.342E-3
L2    19   50       1.342E-3
G7    99   17       14 28 1E-6
G8    17   50       28 14 1E-6
*
* POLE AT 40MHZ
*
R17   20   99       1E6
R18   20   50       1E6
C7    20   99       3.979E-15
C8    20   50       3.979E-15
G9    99   20       17 28 1E-6
G10   20   50       28 17 1E-6
*
* POLE AT 200MHZ
*
R19   21   99       1E6
R20   21   50       1E6
C9    21   99       .796E-15
C10   21   50       .796E-15
G11   99   21       20 28 1E-6
G12   21   50       28 20 1E-6
*
* POLE AT 200MHZ
*
R21   22   99       1E6
R22   22   50       1E6
C11   22   99       .796E-15
C12   22   50       .796E-15
G13   99   22       21 28 1E-6
G14   22   50       28 21 1E-6
*
* POLE AT 200MHZ
*
R23   23   99       1E6
R24   23   50       1E6
C13   23   99       .796E-15
C14   23   50       .796E-15
G15   99   23       22 28  1E-6
G16   23   50       28 22 1E-6
*
* COMMON-MODE GAIN NETWORK WITH ZERO AT 40KHZ
*
R25   24   25       1E6
R26   24   26       1E6
L3    25   99       3.979
L4    26   50       3.979
G17   99   24       3  28 1E-6
G18   24   50       28  3 1E-6
*
* POLE AT 300MHZ
*
R27   27   99       1E6
R28   27   50       1E6
C15   27   99       .531E-15
C16   27   50       .531E-15
G19   99   27       23 28 1E-6
G20   27   50       28 23 1E-6
*
* OUTPUT STAGE
*
R29   28   99       20.0E3
R30   28   50       20.0E3
R31   29   99       30
R32   29   50       30
L5    29   34       1.65E-7
G21   32   50       27 29 33.3333E-3
G22   33   50       29 27 33.3333E-3
G23   29   99       99 27 33.3333E-3
G24   50   29       27 50 33.3333E-3
V3    30   29       0.2
V4    29   31       0.2
D5    27   30       DX
D6    31   27       DX
D7    99   32       DX
D8    99   33       DX
D9    50   32       DY
D10   50   33       DY
*
* MODELS USED
*
.MODELI QX NPN(BF=1250)
.MODEL DX    D(IS=1E-15)
.MODEL DY    D(IS=1E-15 BV = 50)
.ENDS OP-61
```

the positive half, but the ringing frequency is lower than that of the real circuit.

This inability to model nonsymmetrical output-stage behavior is inherent in the Boyle approach and is still, unfortunately, shared by the new macro-model. However, it is a drawback that you can work around. If, during the model-generation process, you find that the overshoot is different from the undershoot, you should use the larger of the two values in calculations pertaining to the output inductor. Then with capacitive loads, the inductor value will yield the worst-case overshoot and undershoot results.

EXECUTION-TIME COMPARISONS

Assuming that no convergence problems exist in the macro-model, the time taken for SPICE to produce an operating-point calculation or a DC-transfer curve is largely a function of the number of circuit elements specified in the net list. Consequently, the new OP-42 macro-model was almost exactly twice as slow as its Boyle counterpart and required 2.27 times as many iterations to reach the final solution. Similar remarks apply to the AC-analysis case, where the run-time overhead of the new macro-model was almost exactly twice that of the Boyle macro-model. However, the two models required about the same number of iterations for AC-response simulation.

Evaluating the computational overhead for a transient analysis is quite difficult, because of the large number of factors involved. In particular, the new macro-model will exhibit considerably more detail than the Boyle model. The simulator must therefore use a much finer time step and perform correspondingly more calculations. However, the large number of ideal elements in the model results in a very good probability on convergence. Therefore, you can sometimes speed up the analysis by allowing more iterations per time step, a procedure which often allows the simulator to maintain a coarser time step and reduces the number of backtracks.

Most SPICE simulators default the number of transient iterations to 10. You can override this default by setting ITL4 to a larger number (say 40) in the .OPTIONS section. Additionally, relaxing RELTOL to 0.01 (the default value is usually 0.001) will also speed up the run time by slightly reducing the accuracy. This reduction is quite permissible because the macro-model is only an approximation anyway. Note, however, that Figures 11 and 12 were generated with RELTOL set to 0.001 rather than 0.01, so that the curves would be more accurate. Another way of speeding up the transient analysis is to use GEAR rather than TRAPEZOIDAL integration; however, such integration can generate results that appear considerably less oscillatory than they actually should be.

Using 0.01 for RELTOL, 40 for ITL4 and trapezoidal integration, the OP-42 macro-model proved to be 3.64 times slower on transient runs than the Boyle model and required 2.15 times as many iterations. The reduction in simulation speed, though large, is acceptable, and is outweighed by the advantage of greatly improved accuracy.

THE OP-61 MACRO-MODEL

The OP-61 is a bipolar-input, wideband, precision op amp that typically has a gain-bandwidth product of 200MHz (at a test frequency of 1MHz) and a slew rate of 40V/μs. The model of this device, shown in Figure 13, is only slightly more complicated than that of the OP-42. The OP-61's common-mode rejection starts to roll off at a lower frequency than the CMRR of the OP-42, but at 1MHz, it is still a respectable 80dB. The net list (see Lisiting 2) indicates that the OP-61 model requires 9 poles and 2 zeroes to mimic the open-loop frequency response, and a common-mode gain of zero at 40kHz.

Notice that this model has an additional gain stage (stage b in Figure 13) between the differential input stage and the main gain stage (see Figure 13c), which generates the dominant amplifier pole. The extra gain stage is necessary in this particular model because the OP-61 does not satisfy the limiting equation, which relates the slew rate, open-loop gain, and the dominant pole frequency for the bipolar input stage (see Box 1). The OP-61 model requires an open-loop gain of 100dB and slew rate of 40V/μs, but the gain-bandwidth product (and hence the dominant pole frequency) is too high to allow a single stage to generate all of the open-loop voltage gain.

Therefore, this model uses two gain stages, which together give the requisite 100dB of gain. The first gain stage has a gain of 200; the second has a gain of 500. You have to provide clamping in the first gain stage, in order to limit the maximum drive voltage applied to the voltage-controlled current sources in the second gain stage. This clamping action then limits the amount of peak current delivered to the compensation capacitors C_5 and C_6, and thus limits the maximum dV/dt in the second gain stage.

The first gain stage must provide a fair amount of gain, because the maximum differential output voltage of the input stage is only 51.6mV. To facilitate clamping with voltage sources and diodes, you need a much larger voltage. A gain of 200 in the first gain stage would result in an unclamped voltage of ±10.32V relative to V_h during slewing, but the clamping circuit limits this to approximately ±5.0V regardless of the rail voltages. This configuration allows reliable clamping action even when the power supply voltages are as low as ±4.4V. It also results in the desired slew rate of 40V/μs.

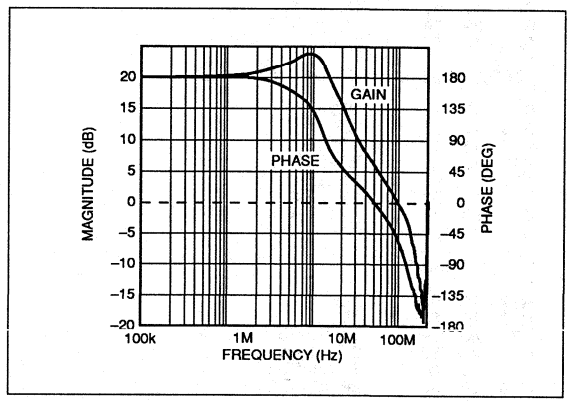

FIGURE 14: *When you connect a real OP-61 as an inverting amplifier with a gain of 10, the gain response shows a 3dB peak of 10MHz.*

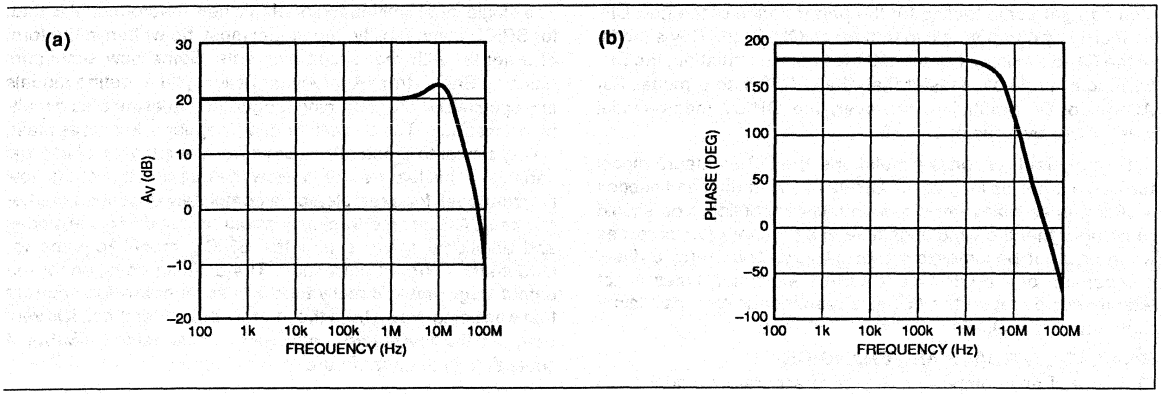

IGURE 15: *The simulated gain (a) of the OP-61 macro-model shows the correct amount of peaking at 10MHz. Further, its hase response (b) at 40MHz differs by only 10° from that of the real device.*

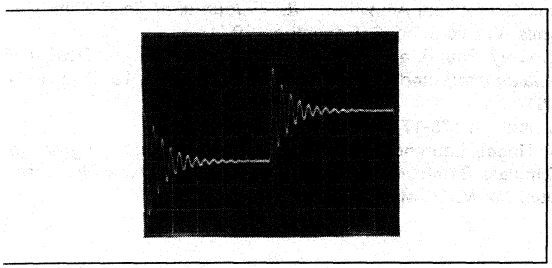

FIGURE 16: *The transient response of a real OP-61, when connected as an inverting amplifier with a gain of 10 and a capacitive load of 207pF, shows some asymmetry. The input signal is a 500kHz sqaure wave with a peak amplitude of 10mV. The vertical scale is 0.1V/div, and the horizontal scale s 0.2µs/div.*

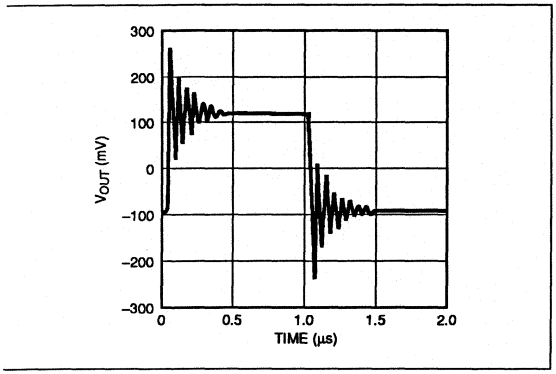

FIGURE 17: *The simulated transient response of the OP-61 macro-model quite closely matches the transient response of the real device.*

SIMULATION-ACCURACY COMPARISONS

Figure 14 shows the measured gain and phase responses of a physical OP-61 configured as an inverting amplifier with a gain of 10. Here, a 1kΩ feedback resistor, a 100Ω input resistor, and ±15V power supplies were used. The amplitude response exhibits a definite peak of about 3dB in the 10MHz region, and the phase shift also increases quite rapidly above 10MHz. The corresponding responses of the new macro-model (see Figures 15a and 15b) show excellent conformance to the measured gain response of the OP-61. The gain curve exhibits the requisite gain peak of slightly over 2dB just above 10MHz. The phase-response accuracy is also quite good; the error is only about 10° at 40MHz, and is probably within the range of variation one would expect to see on a breadboard because of parasitic capacitances and other physical effects. This new macro-model is therefore a useful tool in predicting the performance of the OP-61, even before you evaluate the breadboard.

Figure 16 shows the transient reponse of the OP-61, which might appear to be rather unstable until you notice that the device is driving a 207pF capacitive load. The waveform exhibits some asymmetry between the amounts of overshoot and undershoot (180% versus 220%), but the OP-61, like the OP-42, does not have a perfectly balanced output-stage structure. The choice of the output inductor (L_5 in the model) largely determines how closely the simulated transient response will mimic the real response. In fact, the simulation shown in Figure 17 yields symmetrical overshoot and undershoot of about 150%, which is a little low, and a ringing frequency which is a little high, compared to those of Figure 16. This discrepancy is unlikely to be of much importance to the user. If it is important, however, you could easily bring the simulated response closer to that of the real device by slightly increasing the value of the output inductor.

13

You can get some feeling for the performance of the new OP-61 model by comparing it to that of the OP-42 (no Boyle model of the OP-61 exists). For the DC bias point calculation, the OP-61 macro-model was faster than the OP-42 macro-model. For AC-response simulation, however, the OP-61 macro-model was slower by a factor of 1.18.

In the transient response simulations, the OP-61 macro-model took 1.76 times as long as the OP-42 macro-model and needed 1.56 times as many iterations. In this connection, you should remember that the simulation time of a transient run increases as the output becomes more oscillatory. Therefore, a direct comparison of the OP-42 and OP-61 execution times is not exactly fair because the OP-42's response is less oscillatory than that of the OP-61.

SIMULATION GOALS ARE CHANGING

The goal of any computer model is to accurately model some physical phenomenon; the more complex the phenomenon, the longer the time required for the computer to perform the necessary calculations. The goal of the Boyle op amp model was to reduce the number of nonlinear elements that required simulation, and hence to decrease the run time to an acceptable value. The Boyle model was not created with ultimate accuracy in mind, but it could correctly predict the low-frequency performance of an op amp, and was satisfactory for the relatively low-performance devices of its day.

Today, however, there is more and more demand for ever higher performance, and accurate prediction of a new device's performance can help to avoid design errors that would be expensive to correct at the manufacturing stage. Thus, accurate modeling of the high-frequency performance is essential, and in that region, the Boyle model is inadequate. The improved op amp macro-model described here not only models the high-frequency response and transient behavior of an op amp much more accurately than the Boyle model, but also does not need too much more CPU time to do its job. Today, with powerful desktop workstations available, the emphasis in modeling is on improving simulation accuracy rather than shaving every last bit from the execution times. The new macro-model is thus a good compromise.

The single most limiting factor of this new macro-model is that, for SPICE compatibility, the model must be written in the form of a net list with real circuit elements. Some new simulators (such as Saber, from Analogy Inc) allow you to define models in a specialized programming language that eliminates circuit-type constructs. The Saber modeling language, known as Mast, is very similar to C and allows powerful manipulation of internal variables. This feature would allow the output stage of the new macro-model, for example, to be completely described mathematically. A Saber model simply would not need all of the diodes and additional sources that the SPICE model requires for output-stage current correction. The defining equation for the output stage would directly take into account any load current that was being drawn from the model's output terminal. It is very likely that the new macro-model will be implemented in Saber at some time in the near future.

REFERENCES

1. Boyle, Graeme R., et al, "Macromodeling of Integrated Circuit Operational Amplifiers," *IEEE Journal of Solid State Circuits*, Vol. sc-9, no. 6, December, 1974.
2. Gray, Paul R. and Robert G. Meyer, *Analysis and Design of Analog Integrated Circuits*, John Wiley & Sons, 1977, pp 159-161.
3. ibid, pp 175-177.
4. Nagel, Laurence W., "SPICE2: A Computer Program to Simulate Semiconductor Circuits," *U.C. Berkeley Memorandum No. M520*, May, 1975.

ANALOG DEVICES

AN-259
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

AD9617/AD9618 Current Feedback Amplifier Macromodels

by William E. Tolley

Models of electronic circuits which are based on semiconductor process variables and device geometries may have limited usefulness because of their inability to accommodate multiple frequency poles and multiple zeroes. Without this ability, it is difficult to simulate the complexities and subtleties of analog components, although reasonably accurate models of digital components have been available for years.

Macromodels, by contrast, allow designers to simulate high speed op amps and other electronic components by using programs such as SPICE, or its many derivatives. Macromodels use a simplified equivalent circuit and describe the devices within this circuit in terms of their voltages, currents, parameters, and dependencies. By using a macromodel, a designer can investigate an extremely complex component containing dozens of active and passive devices. The macromodel makes this possible by reducing the component to a multiterminal "black box" behavioral model, represented by the netlist, and allows the user to interconnect the component into the final system schematic.

The macromodel listings for the AD9617 and AD9618 current feedback amplifiers included here are based on the equivalent circuits which are shown, and four basic transistor device models. The SPICE user of this macromodel calls up a five-terminal device (LGAMP—low gain amp—for the AD9617 and HGAMP—high gain amp—for the AD9618): model node 1 is IC pin 3 (+ INPUT); node 2 is pin 2 (−INPUT); node 15 is pin 6 (OUTPUT); and nodes 100 and 110 are the power rails of pins 7 and 4 ($+V_S$ and $-V_S$), respectively, paralleled with pins 8 and 5 (optional $\pm V_S$ connections). Some macromodels are based on overall specifications or simplified representations of the circuit and may be so oversimplified they ignore how the IC is actually built. The models shown here, however, attempt to reflect the realities of the analog world. As an example, they emulate diodes Q1 and Q2 in the AD9617/AD9618 input stage by using transistors because that is how these diodes are built. In this way, it becomes possible to do accurate simulations of changes in input circuitry and dc performance versus temperature. The graphs which are shown compare measured and modeled open-loop transimpedance magnitude and phase.

13

Although the macromodels allow accurate assessment of the AD9617 and AD9618 amplifiers in circuits, not all performance features of the amplifiers are included because many of them are considered second-order effects, not necessary for circuit simulation under normal operating conditions. In this way, the model is simplified and reduces computer run-time. Examples where the published models do not correspond to actual performance include:

1. Power supply voltage maximum limits
2. Maximum output voltage range
3. Temperature effects (model parameters assume +25°C)
4. Input voltage and current noise

AD9617 Netlist

```
.SUBCKT LGAMP 1 2 15 100 110
# 1=VIN +, 2=VIN −, 15=VOUT, 100=VCC, 110=VEE
C1 13 5 1.15P
C2 17 6 1.15P
C3 12 2 0.5P
C4 5 6 3.5P
C5 5 0 1.0P
C6 6 0 1.0P
C7 12 0 1.5P
CINN 1 0 1.5P
CINI 2 0 1.5P
F1 100 13 VM1 1.5
F2 17 110 VM2 1.5
GM1 12 9 POLY(2) 0 5 0 6 2.5M, −2.0M, −2.0M
GM2 10 12 POLY(2) 0 6 0 5 2.5M, 2.0M, 2.0M
I1 100 3 DC 1.98M
I2 100 5 DC 3.93M
I3 4 110 DC 2.0M
I4 6 110 DC 4.0M
I5 100 12 5.0M
I6 12 110 5.0M
Q1 3 3 1 110 QNA 1.05
Q2 5 3 2 110 QNA 1.05
Q3 4 4 1 100 QPA 1.05
Q4 6 4 2 100 QPA 1.05
Q5 110 12 13 100 QPA 1.25
Q6A 100 13 14 110 QNB 1.05
Q6B 100 13 14 110 QNB 1.05
Q7 100 12 17 110 QNA 1.05
Q8A 110 17 16 100 QPB
Q8B 110 17 16 100 QPB
R1 12 15 500
R2 14 15 5
R3 15 16 5
R4 100 5 190K
R5 5 0 25K
R6 6 110 395K
R7 6 0 12K
R8 7 20 350
VB1 5 7 DC 1.6
VB2 20 6 DC 1.6
VM1 100 10 DC 0
VM2 9 110 DC 0
*
.MODEL QNA NPN RB=75, IRB=0, RBM=7, RC=20, RE=0.7,
+IS=540E-18, XTB=2.4, BF=380, IKF=15M, VAF=30,
+ISE=22E-16, ISC=35E-21, TF=25P, CJE=16E-14, VJE=1.0,
+CJC=2.2E-13, XCJC=.2, CJS=4E-13, MJS=0.3
*
.MODEL QNB NPN RB=24, IRB=0, RBM=2, RC=6, RE=0.5,
+IS=18E-16, XTB=2, BF=380, IKF=49M, VAF=30,
+ISE=72E-16, ISC=115E-21, TF=25P, CJE=54E-14, VJE=1,
+CJC=6E-13, XCJC=.2, CJS=7E-13, MJS=0.3
*
.MODEL QPA PNP RB=83, IRB=0M, RBM=14, RC=18, RE=0.5,
+IS=26E-17, XTB=2, BF=190, IKF=24M, VAF=15, ISE=7E-15,
+ISC=30E-19, TF=35P, CJE=11E-14, VJE=1, CJC=3.1E-13,
+XCJC=0.2, CJS=9E-13, MJS=0.35,
*
.MODEL QPB PNP RB=25, IRB=0, RBM=4, RC=5, RE=0.2,
+IS=88E-17, XTB=2, BF=190, IKF=84E-3, VAF=15,
+ISE=22.3E-15, ISC=100E-19, TF=35P, CJE=36E-14, VJE=1,
+CJC=8.5E-13, XCJC=0.2, CJS=1.4E-12, MJS=0.35,
*
.ENDS
.END
```

AD9618 Netlist

```
*
VCC 5 0 DC 5 AC 0
VEE 6 0 DC -5 AC 0
V3 2 7 DC 0
RFB 4 2 1000
RFF 1 2 110
RL 4 0 100
*
X1 3 7 4 5 6 HGAMP
* 3=VIN +, 7=VIN −, 4=VOUT, 5=VCC, 6=VEE
*
V1 1 0 DC 0 AC 0
V2 3 0 DC 0 AC 1
*
*
*
.SUBCKT HGAMP 1 2 15 100 110
* 1=VIN +, 2=VIN −, 15=VOUT, 100=VCC, 110=VEE
C1 12 5 .45P
C2 12 6 .45P
C3 12 2 0.2P
C5 5 0 1.25P
C6 6 0 1.25P
C7 12 0 1.25P
CINN 1 0 1.5P
CINN 2 0 1.5P
F1 100 13 VM1 1.2
F2 17 110 VM2 1.2
GM1 12 9 POLY(3) 0 5 0 6 5 6 1.6M,-4.0M, -4.0M
GM2 10 12 POLY(3) 0 6 0 5 6 5 1.6M,4.0M,4.0M
I1 100 3 DC 2.00M
I2 100 5 DC 3.00M
I3 4 110 DC 2.00M
I4 6 110 DC 3.29M
15 100 12 5.0M
I6 12 110 5.0M
```

```
1 3 3 1 110 QNA 1.05
2 5 3 2 110 QNA 1.05
3 4 4 1 100 QPA 1.05
4 6 4 2 100 QPA 1.05
5 110 12 13 100 QPA .75
6A 100 13 14 110 QNB 1.05
6B 100 13 14 110 QNB 1.05
7 100 12 17 110 QNA .75
8A 110 17 16 100 QPB
8B 110 17 16 100 QPB
1 12 0 20000
2 14 15 6
3 15 16 6
4 100 5 120K
5 5 5 0 14K
6 6 110 450K
7 6 0 4.5K
8 7 20 280
/B1 5 7 DC 1.70
/B2 20 6 DC 1.70
/M1 100 10 DC 0
/M2 9 110 DC 0
*
.MODEL QNA NPN RB=75, IRB=0, RBM=7, RC=20, RE=0.7,
+IS=540E-18, XTB=2.4, BF=380, IKF=15M, VAF=30,
+ISE=22E-16, ISC=35E-21, TF=25P, CJE=16E-14,
+CJC=2.2E-13, XCJC=.2, +CJS=4E-13, MJS=0.3, VJE=1
*
.MODEL QNB NPN RB=24, IRB=0, RBM=2, RC=6, RE=0.5,
+IS=18E-16, XTB=2, BF=380, IKF=49M, VAF=30, ISE=72E-16,
+ISC=115E-21, TF=25P, CJE=54E-14, CJC=6E-13, XCJC=.2,
+CJS=7E-13, MJS=0.3, VJE=1
*
.MODEL QPA PNP RB=83, IRB=0, RBM=14, RC=18, RE=0.5,
+IS=26E-17, XTB=2, BF=190, IKF=24M, VAF=15, ISE=7E-15,
+ISC=30E-19, TF=35P, CJE=11E-14, CJC=3.1E-13, XCJC=.2,
+CJS=9E-13, MJS=0.35, VJE=1
*
.MODEL QPB PNP RB=25, IRB=0, RBM=4, RC=5, RE=0.2,
+IS=88E-17, XTB=2, BF=190, IKF=84E-3, VAF=15,
+ISE=22.3E-15, ISC=100E-19, TF=35P, CJE=36E-14,
+CJC=8.5E-13, XCJC=.2, +CJS=1.4E-12, MJS=0.35, VJE=1
*
.ENDS
.END
```

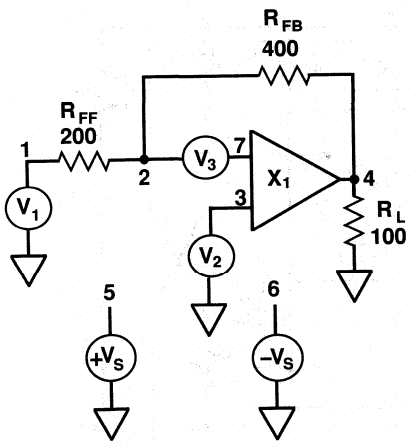

AD9617 and AD9618 Application Circuit

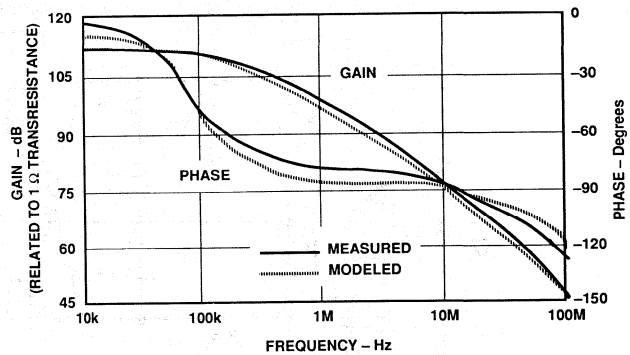

Open-Loop Transimpedance Gain [T (s)] for AD9617

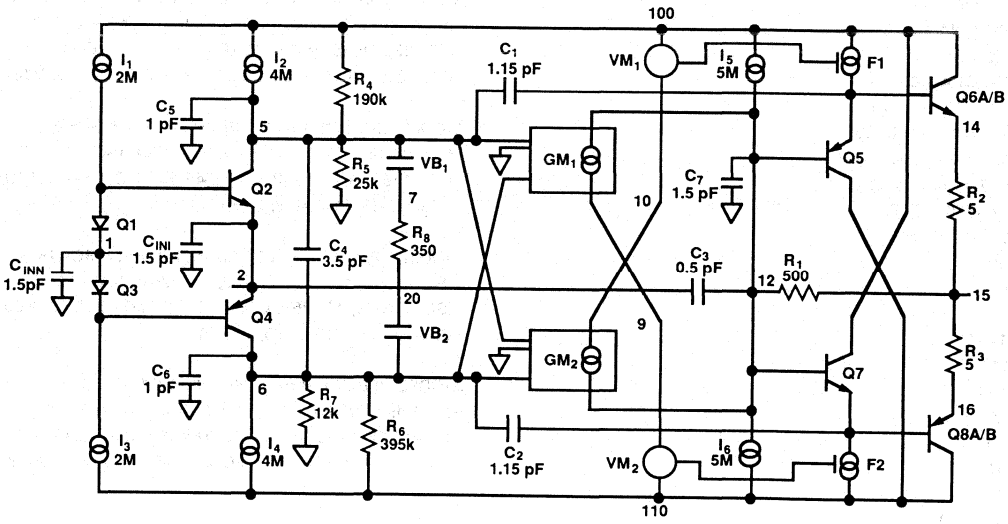

AD9617 Macromodel (LGAMP)

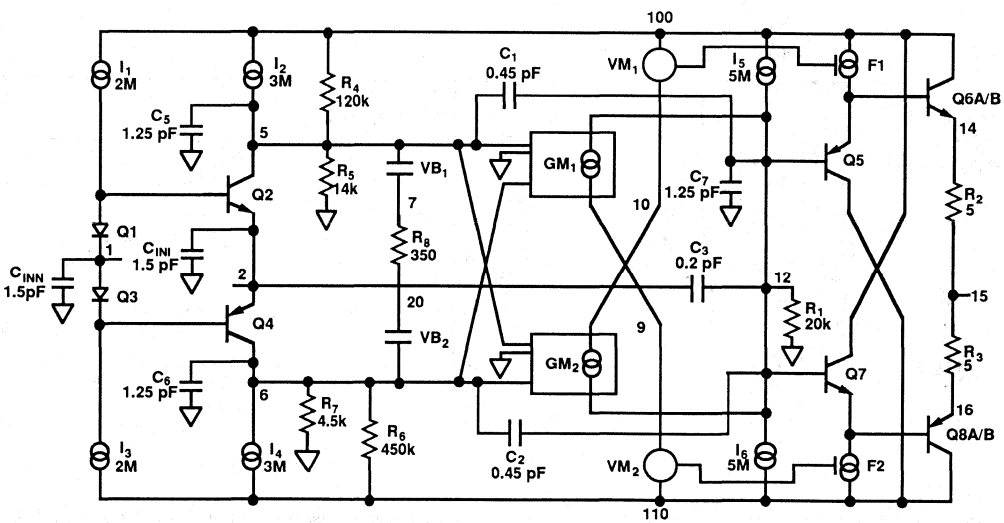

AD9618 Macromodel (HGAMP)

ANALOG DEVICES

AN-360
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Ask the Applications Engineer—6

by James Bryant

OP-AMP ISSUES

Q. Why are there so many different types of operational amplifier?

A. Because there are so many parameters that are important in different applications, and because it is impossible to optimize all of them at once. Op amps may be selected for speed, for noise (voltage, current or both), for input offset voltage and drift, for bias current and its drift, and for common-mode range. Other factors might include power: output, dissipation, or supply, ambient temperature ranges, and packaging. Different circuit architectures and manufacturing processes optimize different performance parameters.

Q. Is there any common factor in the design of op-amps?

A. Yes—most classical (voltage input) op-amps are three-stage devices, consisting of an input stage with differential input and differential output—with good common-mode rejection—followed by a differential-input, single-ended output stage having high voltage-gain and (generally) a single-pole frequency response; and, finally, an output stage, which usually has unity voltage gain.

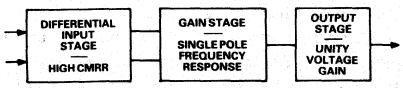

Q. So where are the differences?

A. There are many possible variations on this basic design. One of the most fundamental is the structure of the input stage. This stage is almost always a long-tailed pair—that is to say, a pair of amplifying devices connected as in the figure—but the choice of devices has a profound effect on the input parameters of the op amp. The figure was drawn with thermionic tubes to avoid any suggestion of partiality in favour of any particular semiconductor device. Since thermionic devices at present are not generally available in IC chip form, a monolithic op-amp will have an input stage built with bipolar or field-effect transistors.

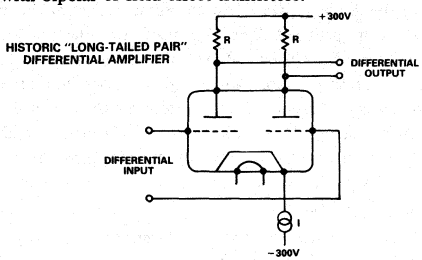

A long-tailed pair built with bipolar transistors is shown in the next figure. Its strong features are its low noise and, with suitable trimming, low voltage offset. Furthermore, if such a stage is trimmed for minimum offset voltage it will inherently have minimum offset drift. Its main disadvantage stems from the proportionality of the emitter and base currents of the transistors; if the emitter current is large enough for the stage to have a reasonable bandwidth, the base current—and hence the bias current—will be relatively large (50 to 1,000 nA in general-purpose op-amps, as much as 10 μA in high-speed ones).

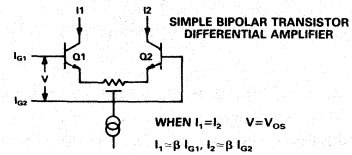

The bias currents in the inverting and non-inverting inputs are unipolar and well matched (their difference is called *offset current*), and they decrease in a minor way with increasing temperature. In many applications, the accurate matching may be used to compensate for their high absolute value. This figure shows a bias compensation circuit where the bias current in the non-inverting input flows in R_c (known as the bias compensation resistor); this compensates for the voltage drop as the bias current in the inverting input flows through R2. R_c is made nominally equal to the parallel combination of R_1 and R_2.—it can be trimmed to minimize error due to non-zero offset current).

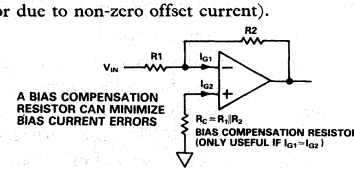

Such bias compensation is only useful when the bias currents are well-matched. If they are not well-matched, a bias compensation resistor may actually *introduce* error.

If a bipolar input stage is required without the drawback of such a high bias current, a different form of bias compensation may be used by the chip designer (next figure). The same long-tailed pair is used, but the major portion of the current required by each base is supplied by a current generator on the chip. This can reduce the external bias current to 10 nA or less without affecting the offset, temperature drift, bandwidth or voltage noise. Bias current variation with temperature is quite low.

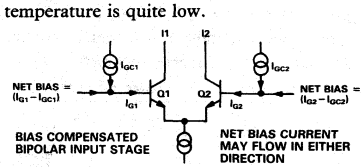

There are two disadvantages to such an architecture: the current noise is increased and the external bias currents are not well matched (indeed, they may actually flow in opposite directions, or change polarity as chip temperature changes). For many applications these features are no drawback; indeed, one of the most popular low-offset op-amp architectures, the OP-07, uses just such an architecture, as do the OP-27, OP-37 and the AD707, which has a guaranteed offset voltage of only 15 μV. Bias-compensated amplifiers of this type are often recognizable when their data sheets explicitly specify *bipolar* bias current, for example, ±4.0 nA.

Where bias currents of even a few nanoamps are intolerable, bipolar transistors are usually replaced by field-effect devices. In the past, MOSFETs have been somewhat noisy for op-amp input

stages, although modern processing techniques are overcoming this drawback. Since MOSFETS also tend to have relatively high offset voltages, *junction* FETs (JFETs) are used for high-performance low-bias-current op amps. A typical JFET op-amp input stage is shown in this figure.

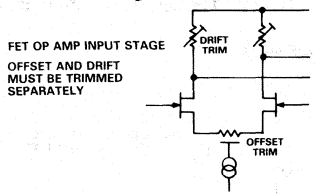

The bias current of a JFET bears no relationship to the current flowing in the device, so even a wideband JFET amplifier may have a very low bias current—values of a few tens of picoamperes are commonplace, and the AD549 has a guaranteed bias current of less than 60 fA (one electron per three microseconds!) at room temperature.

The qualification "at room temperature" is critical—the bias current of a JFET is the reverse leakage current of its gate diode, and the reverse leakage current of silicon diodes approximately doubles with every 10°C temperature rise. The bias current of a JFET op-amp is thus not stable with temperature. Indeed, between 25°C and 125°C, the bias current of a JFET op-amp increases by a factor of over 1,000. (The same law applies to MOSFET amplifiers, because the bias current of most MOSFET amplifiers is the leakage current of their gate-protection diodes.)

The offset voltage of a JFET amplifier may be trimmed during manufacture, but minimum offset does not necessarily correspond to minimum temperature drift. It has therefore been necessary to trim offset and drift separately in JFET op-amps, which results in somewhat larger values of voltage offset and drift than are available from the best bipolar amplifiers (values of 250 μV and 5 μV/°C are typical of the best JFET op-amps). Recent studies at Analog Devices, however, have resulted in a patented trimming method which is expected to yield much better values in the next generation of JFET op-amps.

We thus see that there are trade-offs between offset voltage, offset drift, bias current, bias current temperature variation, and noise in operational amplifiers—and that different architectures optimize different features. The table compares the features of the three commonest op-amp architectures. We should note one more category, typified by the new AD705 (introduced briefly on page 18), using bipolar *superbeta* input transistors; it *combines low offset voltage and drift with low bias current and drift.*

CHARACTERISTICS OF OP-AMP INPUT STAGES

	SIMPLE BIPOLAR	BIAS-COMPENSATED BIPOLAR	FET
OFFSET VOLTAGE	LOW	LOW	MEDIUM
OFFSETDRIFT	LOW	LOW	MEDIUM
BIAS CURRENT	HIGH	MEDIUM	LOW–VERY LOW
BIAS MATCH	EXCELLENT	POOR (CURRENT CAN BE IN OPPOSITE DIRECTIONS	FAIR
BIAS/TEMP VARIATION	LOW	LOW	BIAS DOUBLES FOR EVERY 10°C RISE
NOISE	LOW	LOW	FAIR

Q. *What other features of op amps should the user know about?*

A A common problem encountered with JFET op-amps is phase inversion. If the input common-mode voltage of a JFET op-amp approaches the negative supply too closely, the inverting and non-inverting input terminals reverse functions. Negative feedback becomes positive feedback and the circuit may latch up. This latchup is unlikely to be destructive, but power may have to be switched off to correct it. This figure shows the effect of such phase inversion in a circuit where latch-up does not occur. The problem may be avoided by using bipolar amplifiers, or by restricting the common-mode range of the signal in some way.

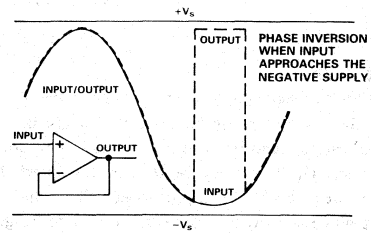

A more serious form of latchup can occur in both bipolar and JFET op-amps if the input signal becomes more positive or negative than the respective op-amp power supplies. If the input terminals go more positive than +Vs + 0.7 V or more negative than –Vs – 0.7 V, current may flow in diodes which are normally biased *off*. This in turn may turn on thyristors (SCRs) formed by some of the diffusions in the op-amp, short-circuiting the power supplies and destroying the device.

To avoid such destructive latch-up it is important to prevent the input terminals of op-amps from ever exceeding the power sup plies. This can have important implications during device turn-on: if a signal is applied to an op-amp before it is powered it may be destroyed at once when power is applied. Whenever there is a risk, either of signals exceeding the voltages on the supplies, or of signals being present prior to power-up of the op-amp, the terminals at risk should be clamped with diodes (preferably fast low-forward-voltage Schottky diodes) to prevent latchup from occurring. Current-limiting resistors may also be needed to prevent the diode current from becoming excessive (see the figure).

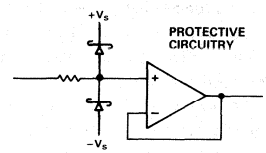

This protection circuitry can cause problems of its own. Leakage current in the diode(s) may affect the error budget of the circuit (and if glass-encapsulated diodes are used, their leakage current may be modulated at 100 or 120 Hz due to photoelectric effects if exposed to fluorescent ambient lighting, thus contributing *hum* as well as dc leakage current); Johnson noise in the current-limiting resistor may worsen the circuit's noise performance; and bias current flowing in the resistor may produce an apparent increase in offset voltage. All these effects must be considered when designing such protection.

The important subjects of noise, interference, bypassing, and grounding demand discussion—but we're out of space! We'll come back to them again in future chats; meanwhile you may want to take a look at some of the references in the footnotes on page 7 of **Analog Dialogue** 23-3.

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Passive and Active Analog Filtering

Filtering is an important part of analog signal processing. Filtering can be used to reduce unwanted signals, limit bandwidth, help recover wanted signals, minimize aliasing in sampled data systems, and smooth the output of DACs. There are five classes of filters. *Lowpass* filters pass all frequencies below the cutoff frequency and block all frequencies above the cutoff frequency. *Highpass* filters are the inverse of the lowpass filters. They block the low frequencies and pass those above the cutoff frequency. *Bandpass* filters pass those frequencies between the lower cutoff and upper cutoff frequencies and reject all others. *Bandstop* filters are the inverse of bandpass filters. They reject frequencies between the cutoff frequencies and pass all others. *Allpass* filters pass all frequencies equally but introduce a predictable phase delay to the signal.

CLASSES OF PASSIVE AND FILTERS

- Lowpass
- Highpass
- Bandpass
- Bandstop
- Allpass

Figure 2.56

Traditional filters were passive, that is designed with no active elements. Active components were too costly and had very poor performance characteristics. Inductors, capacitors, and resisters were used to synthesize the filter. This approach has several difficulties because inductors become physically large for low frequency filters and have poor characteristics at high frequencies. There is a great deal of interaction between the different sections of the filter. Impedance levels must be precisely controlled. Close component tolerances are difficult to manufacture and maintain. Despite these limitations passive filters are still dominant at high frequencies, primarily due to dynamic performance limitations of op amps.

PASSIVE FILTERS

- Designed with Inductors, Capacitors, Resistors
- Large Inductors Required for Low Frequency Filters
- Interaction Between Filter Stages
- Component Tolerances Difficult to Manufacture and Maintain
- Still the Only Solution at High Frequencies Due to Active Component Limitations

Figure 2.57

Active filters answer some of the limitations of the passive filter by offering isolation between stages and eliminating the need for inductors. Their use at high frequencies is limited by the dynamic performance of the active elements.

ACTIVE FILTERS

- Eliminate Need for Inductors
- Good Interstage Isolation
- High Frequency Use Limited by Op Amp Dynamic Performance

Figure 2.58

Reprinted from EDN – February 18, 1988

A filter can be specified in terms of five parameters as shown in Figure 2.59. The *cutoff frequency* F_c is the frequency at which the filter response leaves the error band (or the -3dB point for a Butterworth filter). The *stopband frequency* F_s is the frequency at which the minimum attenuation in the stopband is reached. The *passband ripple* A_{max} is the variation (error band) in the passband response. The *minimum passband attenuation* A_{min} defines the signal attenuation within the stopband. The *order* M of the filter is the number of poles in the transfer function.

KEY FILTER DESIGN PARAMETERS

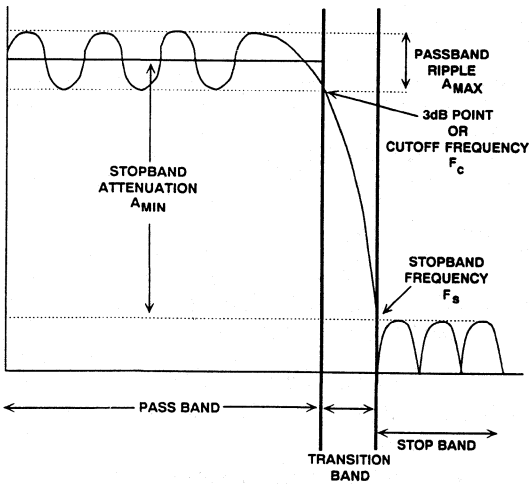

Figure 2.59

FILTER SPECIFICATIONS

- **Cutoff Frequency, F_C**
- **Stopband Frequency, F_S**
- **Passband Ripple, A_{MAX}**
- **Stopband Attenuation, A_{MIN}**
- **Filter Order, M**

Figure 2.60

Typically, one or more of the above parameters will be variable. For instance, if you were to design an antialiasing filter for an ADC you will know the cutoff frequency, the stopband frequency, and the minimum attenuation. You can then go to a chart or computer program to determine the other parameters.

There are many transfer functions that may satisfy the requirements of a particular filter. The *Butterworth* filter is the best compromise between attenuation and phase response. It has no ripples in the passband or the stopband and is called the *maximally flat filter* because of this. The Butterworth filter achieves its flatness at the expense of a relatively wide transition region from passband to stopband.

The *Chebyshev* filter has a smaller transition region than the same-order Butterworth filter, but it has ripples in either its passband or stopband. This filter gets its name because the Chebyshev filter minimizes the height of the maximum ripple—this is the Chebyshev criterion.

The Butterworth filter and the Chebyshev filter are all-pole designs. By this we mean that the zeros of the transfer function are at one of the two extremes of the frequency range (0 or ∞). For a lowpass filter the zeros are at $f = \infty$. We can add finite frequency transfer function zeros as well as poles to get an *Elliptical Filter*. This filter has a shorter transition region than the Chebyshev filter because it allows ripple in both the stopband and passband. The Elliptical filter also has degraded phase (time domain) response.

These are by no means all possible transfer functions, but they do represent the most common.

POPULAR FILTER DESIGNS

- **Butterworth: All Pole, No Ripples in Passband or Stopband, Maximally Flat Response**
- **Chebyshev: All Pole, Ripple in Passband, Shorter Transition Region than Butterworth for Given Number of Poles**
- **Elliptical: Ripple in Both Passband and Stopband, Shorter Transition Region than Chebyshev, Degraded Phase Response, Poles and Zeros**

Figure 2.61

Once the order of the filter and the specifications of filter have been determined, the design charts (see Reference 10) or computer programs are consulted, and the linear and quadratic factors of poles for the transfer function are determined. All filters, regardless of order, are made up of one- or two-pole sections. The single pole section is defined by its resonant frequency, which is the -3dB point. The pole pair in a two-pole filter section is defined by its resonant frequency (F_o) and Q, which indicates the peaking of the section. Sometimes alpha (α) is used instead of Q ($Q=1/\alpha$).

Armed with the various values F_o and Q, you then choose the configuration for the realization of the filter: Butterworth, Chebyshev, or Elliptical.

For passive filters, these values, along with the filter characteristic impedance determine the inductor, capacitor, and resistor values.

For active filters, you must decide which of the realizations you are going to use. The three most common are the *Sallen-Key* (voltage controlled voltage source), *multiple feedback*, and *state variable*. Each realization has its own advantages and disadvantages.

The Sallen-Key configuration shown in Figure 2.62 is the least dependent on the performance of the op amp, and the signal phase is maintained. For this filter the ratio of the largest resistor value to the smallest resistor value and the ratio of the largest capacitor value to the smallest capacitor value is low. The frequency term and Q terms are somewhat independent, but they are very sensitive to the gain parameter. The Sallen-Key is very Q-sensitive to element values for high Q sections. The design equations are also given in Figure 2.62.

VOLTAGE CONTROLLED VOLTAGE SOURCE (SALLEN-KEY) REALIZATION

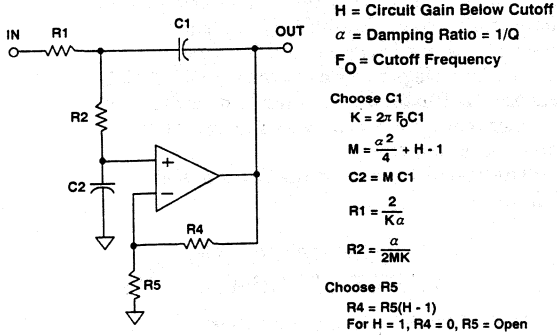

Figure 2.62

The multiple feedback realization shown in Figure 2.63 uses an op amp in the inverting configuration. The dependence on the op amp parameters are greater than in the Sallen-Key realization. It is hard to generate high Q sections due to the limitations of the open loop gain of the op amp. The maximum to minimum component value ratios are higher than in the Sallen-Key realization. The design equations are also given in Figure 2.63.

MULTIPLE FEEDBACK REALIZATION

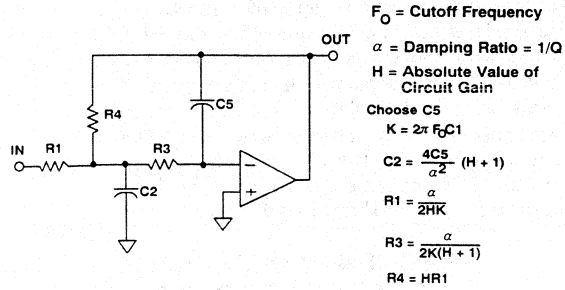

Figure 2.63

The state-variable realization shown in Figure 2.64 offers the most precise implementation, at the expense of many more circuit elements. All parameters can be adjusted independently, and lowpass, highpass, and bandpass outputs are all available simultaneously. The gain of the filter is also independently variable. Since all parameters of the state variable filter can be adjusted independently, component spread is minimized. Also variations due to temperature and component tolerances are minimized. The design equations for the state variable filter are given in Figure 2.64.

STATE VARIABLE REALIZATION

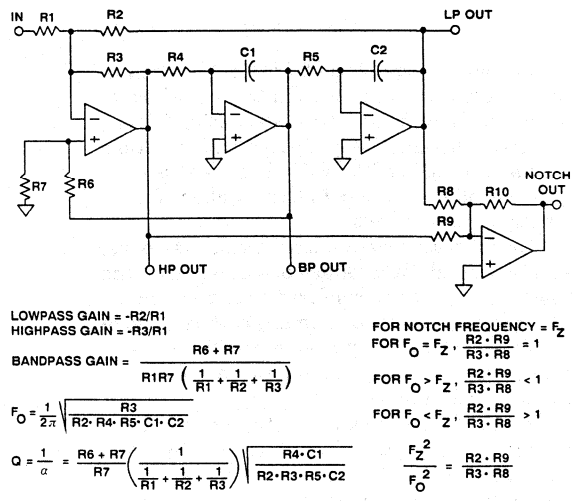

Figure 2.64

Another active filter realization that has recently become more popular is the *Frequency Dependent Negative Resistor* (FDNR), which is a subset of the *General Impedance Converter* (GIC). In the FDNR the passive realization goes through a transformation by

1/s. Therefore inductors, whose impedance is sL, transform into a resistor of value L. Similarly, a resistor of value R becomes a capacitor of value R/s. A capacitor of impedance 1/sC transforms into a frequency dependent variable resistor, which is given the designation D. Its impedance is $1/s^2C$. The transformations to the FDNR configuration and the GIC implementation of the D element are given in Figure 2.65.

FREQUENCY DEPENDENT NEGATIVE RESISTOR 1/S IMPEDANCE TRANSFORMATION

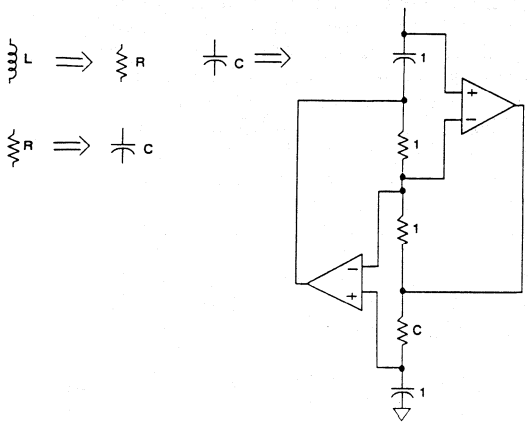

Figure 2.65

The advantage of the FDNR realization is that there are no op amps in the signal path which can add noise. This realization is also relatively insensitive to component variation. The advantages of the FDNR come at the expense of an increase in the number of components required.

For all of the realizations discussed above, the tabulated filter values are in terms of the lowpass function normalized to a frequency of 1 radian/second with an impedance level of 1. To realize the final design, the filter values are scaled by the appropriate frequency and impedance.

Similarly, the lowpass prototype is converted to a highpass filter by scaling by 1/s in the transfer function. In practice this amounts to capacitors becoming inductors with a value 1/C and inductors becoming capacitors with a value of 1/L for passive designs. For active designs resistors become capacitors with a value of 1/R, and capacitors become resistors with a value of 1/C.

Transformation to the bandpass response is a little more complicated. If the corner frequencies of the bandpass are widely separated (by more than 2 octaves) the filter is made up of separate lowpass and highpass sections. In the case of a narrowband bandpass filter the design is much more complicated and is usually done using a computer program or design tables.

SOME ACTIVE FILTER REALIZATIONS

- **Sallen-Key: Good Phase Response, Least Dependent on Op Amp Performance, Sensitive to Element Values for High Q Sections**
- **Multiple Feedback: Less Sensitive to Element Values, High Q Sections Difficult due to Op Amp Open Loop Gain Limitations**
- **State-Variable: Most Precise, More Components, All Parameters Independently Adjustable**
- **Frequency Dependent Negative Resistance (FDNR): Op Amps not in Signal Path, More Components, Relatively Insensitive to Component Variations**

Figure 2.66

Antialiasing Filter Design Example

We will now design a passive and active antialiasing filter based upon the same specifications. The active filter will be designed in four realizations: Sallen-Key, multiple feedback, state variable, and Frequency Dependent Negative Resistance (FDNR). We choose the Butterworth filter in order to give the best compromise between attenuation and phase response.

The specifications for the filter are as follows:

ANTIALIASING FILTER SPECIFICATIONS

- **Cutoff Frequency F_C = 8kHz**
- **Stopband Attenuation F_S at 50kHz = 70dB**
- **Best Balance Between Attenuation and Phase Response**
- **Choose Butterworth Design**
- **From Design Charts, for f = 6.25 (50kHz/8kHz), M = 5**

Figure 2.67

Consulting the design charts (Reference 10, p. 82), we see that for 70dB of attenuation at a frequency of 6.25 (50kHz/8kHz) a

fifth order filter is required.

We now consult the tuning tables (Reference 10, p. 341) and find:

ALPHA AND F_0 VALUES FROM TUNING TABLES

STAGE	ALPHA	F_0
1	---	1.000
2	1.618	1.000
3	0.618	1.000

Figure 2.68

The first stage is a real pole, thus the lack of an alpha value. It should be noted that this is not necessarily the order of implementation in hardware. In general you would typically put the real pole last and put the second order sections in order of decreasing alpha (increasing Q).

For the passive design we will choose the zero input impedance configuration. From the design table (Reference 10, p. 313) we find the following normalized values for the filter:

NORMALIZED PASSIVE FILTER VALUES FROM TABLES

L1 = 1.5451 C2 = 1.6944
L3 = 1.3820 C4 = 0.8944
L5 = 0.3090

Figure 2.69

These values are for a 1 rad/second filter with a 1 ohm termination. To scale the filter we divide all reactive elements by the desired cutoff frequency, 8kHz (50265 rad/sec). We also need to scale the impedance. For this example, we choose a value of 1000 ohms. To scale the impedance we multiply all resistor and inductor values and divide all capacitor values by the impedance scaling factor. After scaling, the circuit looks like Figure 2.70.

EXAMPLE FILTER PASSIVE IMPLEMENTATION

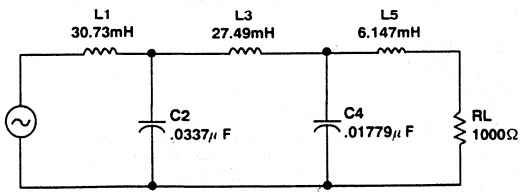

Figure 2.70

For the Sallen-Key active realization, we use the design table shown in Figure 2.62. The values for C1 in each section are chosen to give reasonable resistor values. The implementation is shown in Figure 2.71. For the Sallen-Key realization to work correctly, it is assumed to have a zero-impedance driver and a return path for dc. Both of these criteria are approximately met when you use an op amp to drive the filter.

EXAMPLE FILTER SALLEN-KEY IMPLEMENTATION

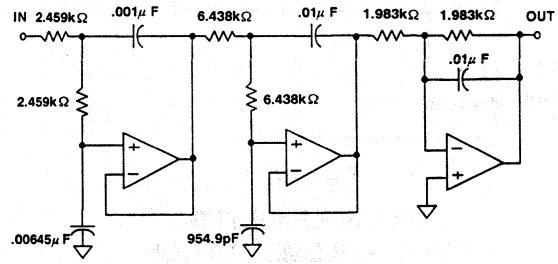

Figure 2.71

Figure 2.72 shows a multiple feedback realization of our filter. It was designed using the equations in Figure 2.63.

EXAMPLE FILTER MULTIPLE FEEDBACK IMPLEMENTATION

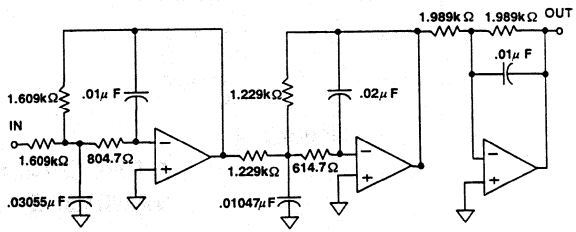

Figure 2.72

The state variable realization is shown in Figure 2.73, and the Frequency Dependent Negative Resistance (FDNR) realization is shown in Figure 2.74. In the conversion process from passive to FDNR, the D element is normalized for a capacitance of 1F. We then scale the filter to a more reasonable value (0.01μF in this case).

In all of the filters above the values shown are the exact calculated values. These exact values are rarely obtainable. We must therefore either substitute the nearest standard value or use series/parallel combinations. Any variation from the ideal values will cause a shift in the filter response char-

EXAMPLE FILTER STATE VARIABLE IMPLEMENTATION

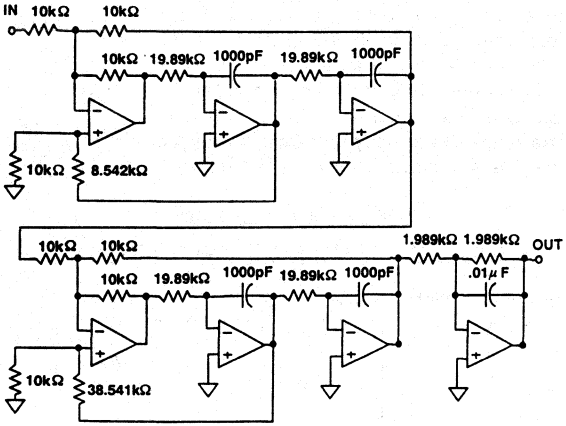

Figure 2.73

EXAMPLE FILTER FDNR IMPLEMENTATION

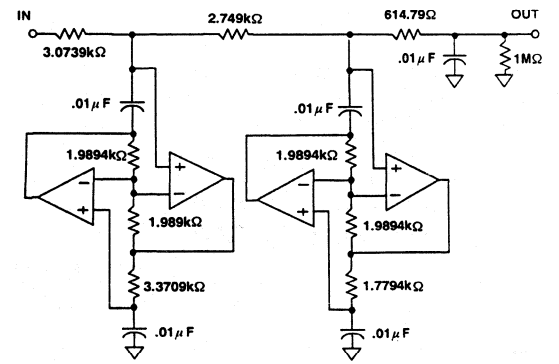

Figure 2.74

acteristic, but often the effects are minimal. The computer can be used to evaluate these variations on the overall performance and determine if they are acceptable.

In active filter applications using op amps, the dc accuracy of the amplifier is often critical to optimal filter performance. The amplifier's offset voltage will be passed by the filter and may be amplified to produce excessive output offset. For low frequency applications requiring large value resistors, bias currents flowing through these resistors will also generate an output offset voltage.

In addition, at higher frequencies, an op amp's dynamics must be carefully considered. Here, slewrate, bandwidth, and open loop gain play a major role in op amp selection. The slewrate must be fast as well as symmetrical to minimize distortion.

A Programmable State Variable Filter

A realization of a programmable state variable filter using DACs is shown in Figure 2.75. DACs A1 and B1 control the gain and Q of the filter characteristic, while DACs A2 and B2 must accurately track for the simple expression for f_c to be true. This is readily accomplished using two AD7528 DACs and one AD713 quad op amp. Capacitor C3 compensates for the effects of op amp and gain-bandwidth limitations.

This filter provides lowpass, highpass, and bandpass outputs and is ideally suited for applications where digital control of filter parameters is required. The programmable range for component values shown is f_c = 0 to 15kHz, and Q = 0.3 to 4.5.

A PROGRAMMABLE STATE VARIABLE FILTER CIRCUIT

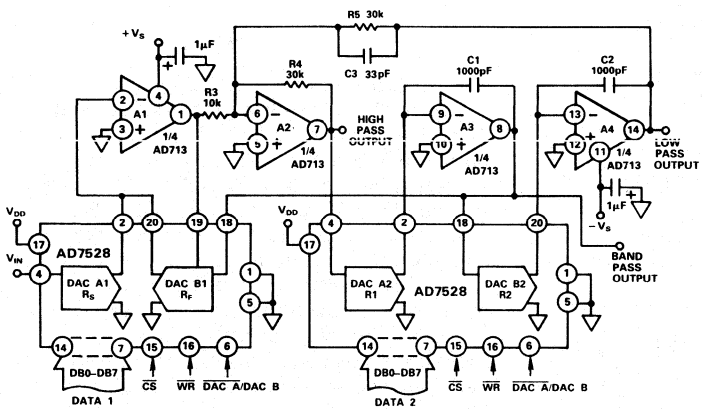

CIRCUIT EQUATIONS

$C_1 = C_2$, $R_1 = R_2$, $R_4 = R_5$

$$f_C = \frac{1}{2\pi R_1 C_1}$$

$$Q = \frac{R_3}{R_4} \cdot \frac{R_F}{R_{FBB1}}$$

$$A_0 = -\frac{R_F}{R_S}$$

NOTE:
DAC equivalent resistance equals
256 x (DAC Ladder resistance)
DAC Digital Code

Figure 2.75

Figure 2.76 shows a 7-pole antialiasing lter for a 2x oversampling (88.2kSPS) igital audio application. This filter has less han 0.05dB passband ripple and 19.8 ± 0.3μs delay, dc-20kHz. The filter will handle a 5V rms signal (V_s = ±15V) with no overload at any internal nodes. The frequency response of the filter is shown in Figure 2.77.

20kHz FDNR AUDIO ANTIALIASING FILTER

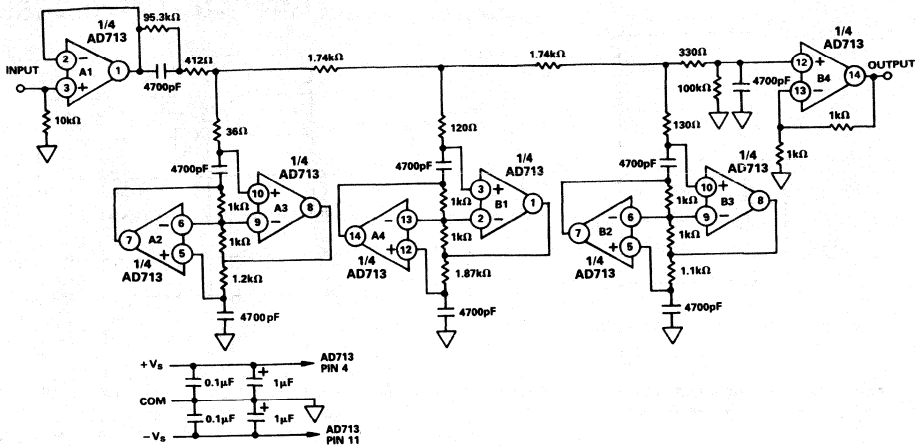

Figure 2.76

AUDIO ANTIALIASING FILTER RESPONSE

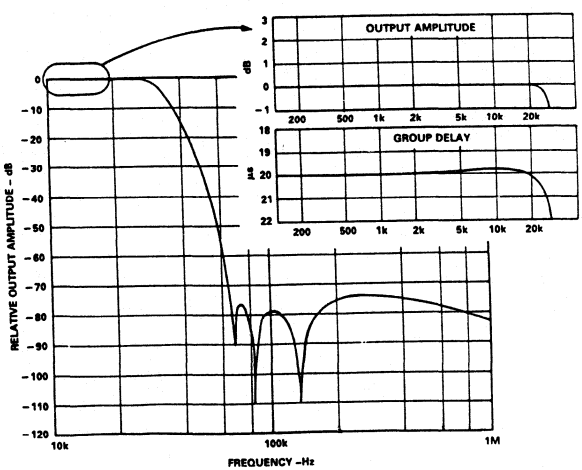

Figure 2.77

A Wideband Sallen-Key Filter

Figure 2.78 shows an AD843 FET input op amp used in a 1MHz Sallen-Key filter. This circuit also works well with the AD841, AD845, or AD847. The circuit is designed to be a maximum-flatness filter with a Q of 0.575 and a dc gain of 1.26. The frequency response of the filter to a 0dBm input signal is shown in Figure 2.79.

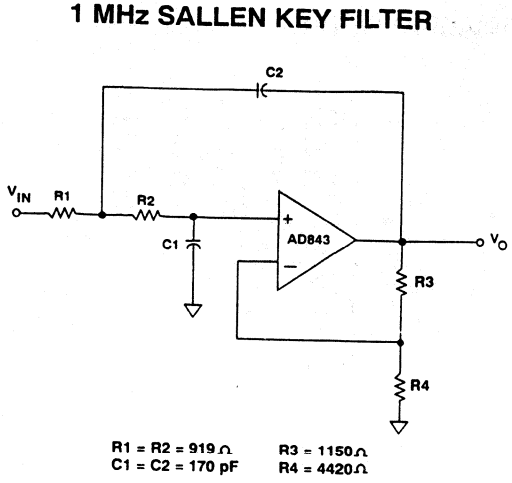

Figure 2.78

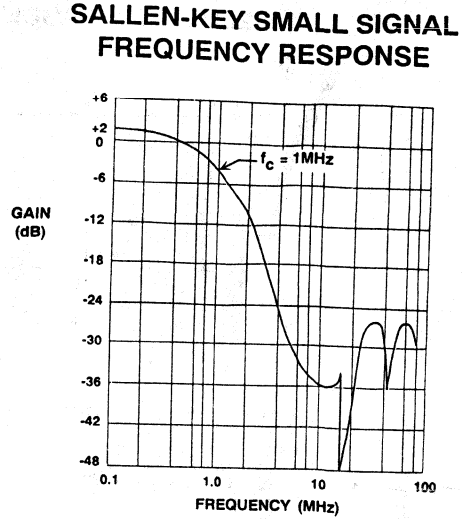

Figure 2.79

Programmable Delay Generators
Contents

		Page
AN-260	Using Digitally Programmable Delay Generators	14–3
AN-261	Programmable-Delay ICs Control System Timing	14–9

14

)NE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Using Digitally Programmable Delay Generators

by Allen Hill

'he AD9500 and AD9501 digitally programmable delay enerators are versatile parts, useful in numerous appliations. The parts are designed for use in automatic test quipment as a deskew element for digital data lines. 'he versatility of the AD9500 and AD9501 for generating rogrammable delays allows them to be used in appliations which range from ATE to communications, comuters, disk drives, lasers, and ultrasound systems. This ote describes how best to apply these parts in some of hese applications.

ENERAL DESCRIPTION

A digitally programmable delay generator delays a digial edge by a programmed amount of time. Figure 1 hows the basic function of a programmable delay genrator. The delay through the device is controlled by an N-bit digital word. This is the programmed delay. A rigger pulse is applied to the input, and after a fixed ropagation delay, (t_{PD}), the pulse edge appears a program delay later at the output.

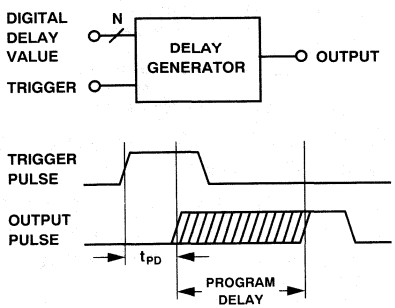

Figure 1. Programmable Delay Generator

The AD9500 (ECL) and AD9501 (TTL) use a ramp/comparator/DAC architecture as shown in Figure 2. One input of a high speed comparator is driven by a digital-to-analog converter (DAC). The DAC is used to set a reference voltage at this comparator input. The other input is connected to a ramp generator. The ramp generator is started by applying a pulse to the trigger input of the delay generator. When the ramp voltage crosses the comparator threshold set by the DAC, the output of the comparator switches.

This output is delayed from the trigger pulse by an amount of time which is proportional to the DAC digital input code and the slope of the ramp. Altering the digital delay value changes the DAC output voltage, which alters the delay through the circuit. The slope of the ramp is controlled with external components.

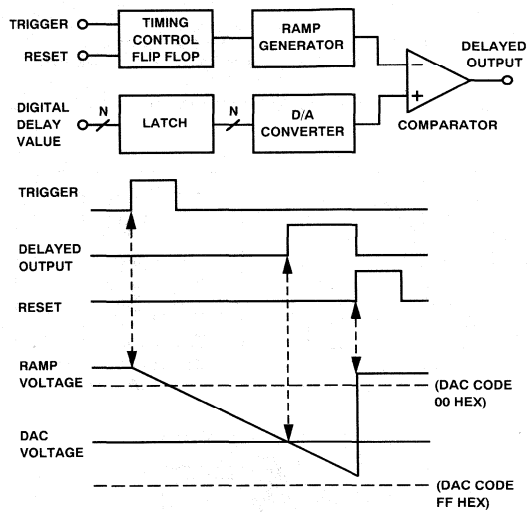

Figure 2. Delay Generator Block Diagram and Basic Timing

Once the comparator has switched, the ramp generator and comparator must be reset so that the device can be triggered again. One method of accomplishing the reset is to connect the output of the delay generator to the reset pin. This results in an output pulse width which is equal to the reset propagation delay of the device (7 ns to 15 ns). An alternate, and versatile method of resetting the device is to use an external signal which meets the timing requirements of the part. An external reset signal allows the pulse width to be controlled and makes system integration of the delay signal easier.

Full-Scale Range Setting

The full-scale range of the generator is the span over which the delay can be programmed. This range is divided into 256 equal delays by the 8-bit digital delay value. The full-scale range of the delay generator is configured by connecting R_{SET} and C_{EXT} as shown in Figure 3. Additional information is available on the AD9500 and AD9501 data sheets. The range can be adjusted from a minimum of 2.5 ns out to 10 μs and beyond.

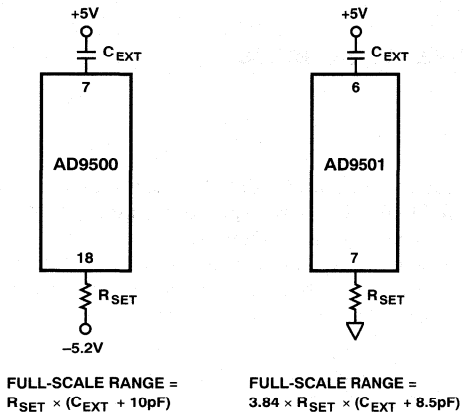

Figure 3. Setting Full-Scale Range on the AD9500 and AD9501

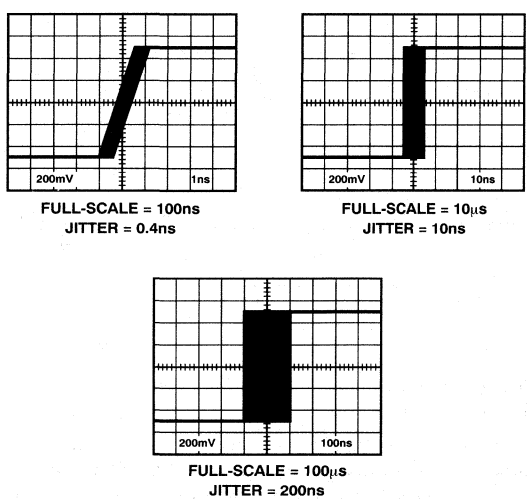

Figure 4. Output Jitter at Various Full-Scale Settings for the AD9500 and AD9501

The maximum full-scale range of these devices depends on the amount of jitter the application can tolerate. Jitter is the variation of delay through the device with subsequent trigger pulses. An increase in full-scale results in an increase in jitter of the output delay. As the full-scale is increased, the slope of the ramp is decreased, resulting in a longer period of time that the ramp is in the comparator transition region. Any noise on the ramp or DAC during this time can cause the comparator to switch. Figure 4 illustrates the output jitter with various full-scale delay settings.

The full-scale range can be adjusted by switching different values of resistance into the R_{SET} pin. A digitally controlled resistor can also be used to provide this function. A programmable full-scale circuit is generated by using a DAC to control the current flow into the R_{SET} pin as shown in Figure 5. The R_{SET} pin of the AD9500 is biased at approximately −4.3 V and requires a current source to set full scale. The bias on the R_{SET} pin of the AD9501 is about 0.5 V and can be used in a similar manner.

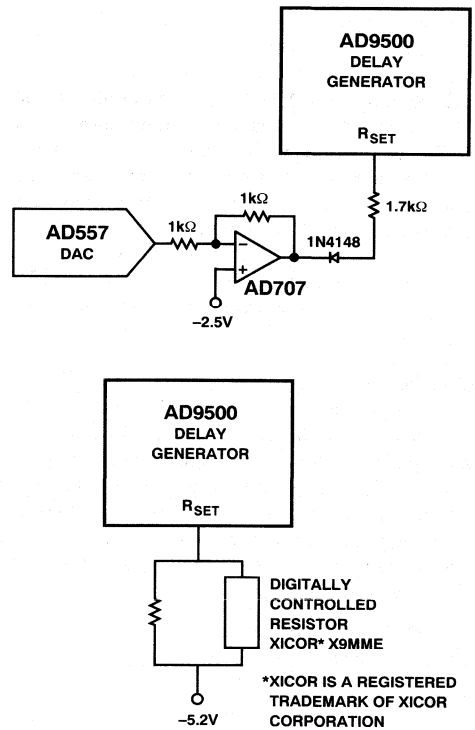

Figure 5. Programmable Full-Scale Range Using a DAC or a Digitally Controlled Resistor

C_{EXT} should not be switched on the AD9500 or AD9501 to provide a programmable full-scale range. The noise pickup on the interconnects can cause false triggering at the input of the device.

When C_{EXT} is increased to extend full-scale range, the reset propagation delay increases because the larger capacitance must be discharged to reset the ramp.

The propagation delay is the time required for the ramp to reach the first DAC threshold. The slope of the ramp is determined by the full-scale setting, which means the propagation delay will change with full-scale range. The propagation delay of the AD9500 and AD9501 can be calculated from the equations shown below.

AD9500 Prop Delay = 5 ns + 0.18 × (Full-Scale Range); Offset Pin Open

AD9501 Prop Delay = 8 ns + 0.18 × (Full-Scale Range); Offset Pin Grounded

Each device has an offset pin which can be used to adjust the prop delay. This adjustment should only be used to match the prop delays of more than one device in a system. There is not sufficient range to zero the prop delay, and if the offset is adjusted too low, false triggers will occur.

Timing for the AD9500 and AD9501

Figure 6 illustrates the timing requirements of the AD9500 and AD9501. The trigger delay through the generator is made up of four separate parts:

1. Trigger Circuit Delay
2. Ramp Generator Delay
3. Program Delay
4. Comparator Delay

Trigger circuit delay, ramp generator delay, and comparator delay make up the propagation delay of the device. Ramp generator delay is dependent on full-scale range, as discussed earlier. The program delay is the amount of delay programmed by the 8-bit digital value.

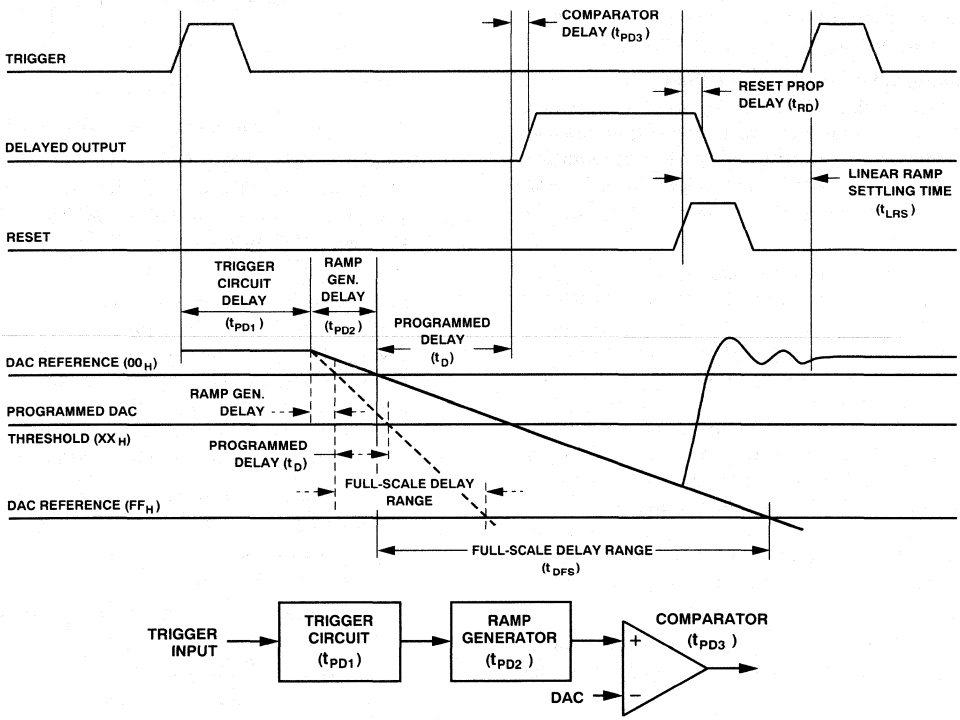

MINIMUM PROPAGATION DELAY = $(t_{PD}) = (t_{PD1}) + (t_{PD2}) + (t_{PD3})$

MAXIMUM PROPAGATION DELAY = MINIMUM PROPAGATION DELAY (t_{PD}) + FULL-SCALE DELAY RANGE (t_{DFS})

TOTAL DELAY = $(t_{PD}) + (t_D)$

Figure 6. Delay Generator Timing

The reset input also has delays associated with it:

1. Reset Propagation Delay
2. Linear Ramp Settling Time

The reset propagation delay is the time from the edge of the reset signal until the output of the comparator resets. The linear ramp settling time is the time from the edge of the reset signal until the ramp has settled to 8-bit accuracy. The reset delays are not cumulative because they are reset in parallel. This means that the linear ramp settling time determines the reset time when triggering at a random rate because it is the longer of the two delays.

However, if the device is triggered at a constant rate, the reset propagation delay determines the reset time. A constant update rate causes the ramp to settle to the same value, even though this is not the theoretical final value. This means the ramp will always start at the same point.

The trigger delay, combined with the reset delay, determines the maximum rate at which the device can be updated. When the update is constant and the output pulse is used for a reset signal, the following equation can be used to determine the maximum update rate.

Max Update = 1 / [Trigger Delay + Reset Delay + Reset to Trigger Holdoff]

where

Trigger Delay = Prop Delay + Full-Scale Delay

and

Reset Delay = Reset Prop Delay

In applications where the update is not constant, the only change in the above equation will be that the reset delay will be equal to the linear ramp settling time. When the digital code is being changed between trigger pulses, the settling time for the DAC must also be included to determine the maximum update rate.

Considerations When Using the AD9500

The $\overline{QR}$ output of the AD9500 is included to provide a separate output to reset the device. This output does not have the drive capability of the Q and $\overline{Q}$ outputs and should be terminated with higher values of resistance. Typically 4 kΩ connected to the minus supply will be acceptable.

Feedthrough of noise from the data bus can cause jitter on the delayed output of the AD9500. If the AD9500 is in the latched mode and the data bus is switching while a delay is being performed, jitter will be seen on the delayed signal. In applications where this situation occurs, it is recommended that an external latch be used. When using the latch of the AD9500, care must be taken to ensure that the voltage on the data bus does not exceed the $+V_S$ of the device. Data stored in the latch will be lost if this condition exists. The problem is prevalent when the data bus is tristated, or the driving logic has high V_{OH} (HCT), or high slew rate (ACT). In each of these conditions noise or switching can carry the voltage above $+V_S$.

Using an external latch with the device is the best way to alleviate the problem. The AD9501 does not exhibit this condition.

Using the AD9500 with TTL Signals

In applications where maximum update rate is required and TTL compatibility must be maintained, the AD9500 can be configured as shown in Figure 7. The $\overline{TRIGGER}$ pin is biased to a level which is at the switching threshold of a TTL signal. The TRIGGER input is connected to a TTL signal. The RESET and $\overline{RESET}$ inputs cannot tolerate TTL levels and must be resistively divided as shown.

The outputs of the AD9500 cannot drive TTL logic directly. Their swing is approximately 1 V peak-to-peak.

These output signals can be offset and the levels made TTL compatible with the addition of the output circuit shown in Figure 7.

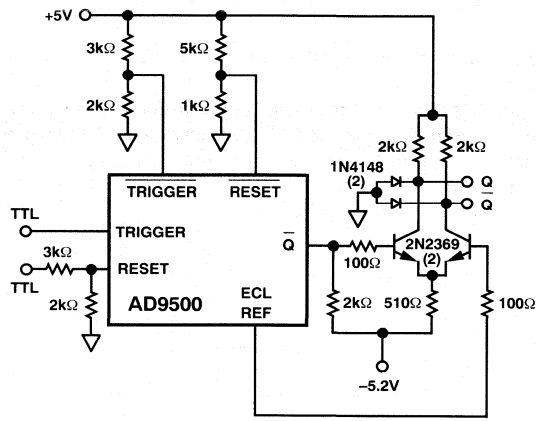

Figure 7. Interfacing the AD9500 with TTL

The smaller prop delay and the faster reset time of the AD9500 can be utilized for faster update rates. If the performance specifications of the AD9501 fit the application, it should be used in lieu of an AD9500 configured in the TTL mode.

Programmable Oscillator with Enable

An AD9501 configured as a programmable oscillator is shown in Figure 8. The frequency of oscillation is controlled by programming the AD9501 delay and can be determined by the equation shown below. The enable input allows the oscillator to be switched on and off, which is a useful feature in many applications.

Freq = 1 / [AD9501 Prop Delay + AD9501 Program Delay + AD9501 Reset Prop Delay + (3 × U1 Prop Delay)]

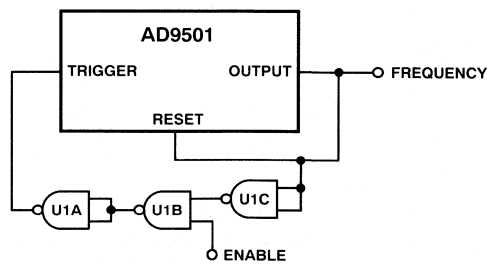

Figure 8. The AD9501 Configured as a Programmable Oscillator

Delaying A Pulse

The AD9500 and AD9501 delay only one edge of the input trigger pulse. The trigger pulse is not replicated at the output of the device. To preserve the pulse width of the input signal, two delay generators must be used as shown in Figure 9. The full-scale range and program delay are set to the same values on both devices. Delay Generator 1 delays the rising edge of the input pulse, while Delay Generator 2 delays the falling edge.

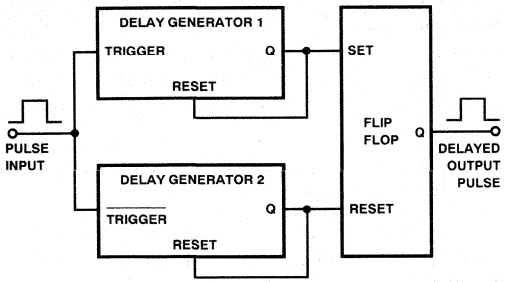

Figure 9. Two Delay Generators Used to Delay Both Input Pulse Edges to Main Pulse Width at the Output

Multichannel Deskewing

High speed systems with parallel signal paths require that close delay matching be maintained. Delay mismatch can cause error in data transfer. Much of this skew can be eliminated by running each signal through a delay generator and adjusting the program delays to minimize the timing skews. With the very fine timing adjustments possible from the AD9500 and AD9501, most high speed systems should be able to adjust automatically to extremely tight tolerances. Figure 10 illustrates a typical deskew application. A method of performing the calibration of multiple delay generators is shown in Figure 11. At the test head, the DUT socket is shorted from input to output. The Q output of the flip flop is monitored for the switching threshold of the delay generator as the digital delay code is varied. Each delay generator is adjusted separately with this offset delay stored in its latch. Calibration is complete when each delay generator is adjusted.

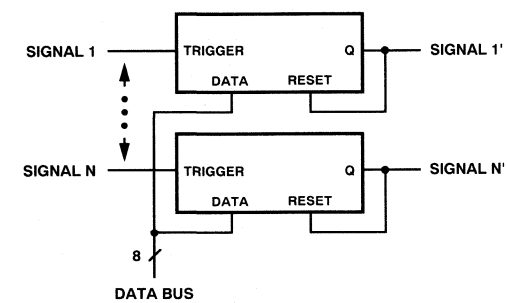

Figure 10. Multiple Signal Deskewing

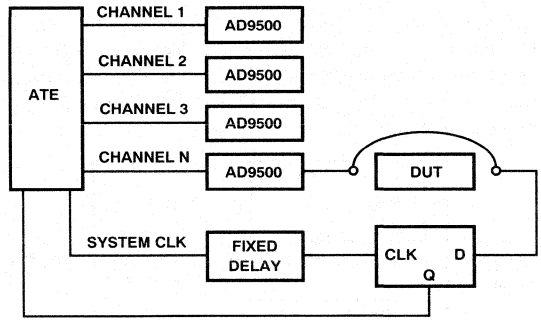

Figure 11. Automatic Calibration for ATE Deskew

Laser Applications

Laser light differs from other light sources in that it is coherent. Laser light is usually monochromatic and highly collimated. Because lasers are stable, emit sharp spectral lines, and are convenient to use, they are found in applications such as range measurement; atmospheric monitoring; industrial cutting, welding, and drilling; laser printers; computer laser disk drives; spectrometry; and communication.

In most of these applications, the laser is a pulsed beam of light which must be synchronized to an external event. The AD9500 and AD9501 provide an easy means for this synchronization, and allow calibration. A programmable delay generator can be used to control the timing and duration of a laser pulse with respect to an external event. In systems with two or more lasers, synchronization can also be controlled with a delay generator.

The duration of laser pulses is usually short, tens of nanoseconds, and the repetition rate is fairly low, tens of kilohertz. Full-scale range must be set small so that jitter induced by the delay generator does not corrupt the required synchronization of the system. This usually means the full-scale range is set to much less than a cycle of the laser pulse repetition rate, typically 100 ns–300 ns. A fixed coarse delay may have to be used in conjunction with a fine tune programmable delay generator to cover the range of interest to provide the required accuracy.

Pulse Width Modulation (PWM)

Another application for the AD9500/AD9501 is pulse width modulation (PWM) in laser printers. PWM allows a laser printer to print gray shades instead of just black and white images.

In normal operation a pulse is supplied to a laser diode driver which turns the laser diode completely on or off. When the diode is on, toner is deposited on the drum, and when the diode is off, no toner is deposited. This image is then transferred to paper. The image will be black and white.

By controlling the laser diode pulse width, varying shades of gray may be generated. Using two AD9501s as shown in Figure 12, the pulse width can be modulated digitally. Digital modulation allows tremendous flexibility and repeatability in creating gray shade images.

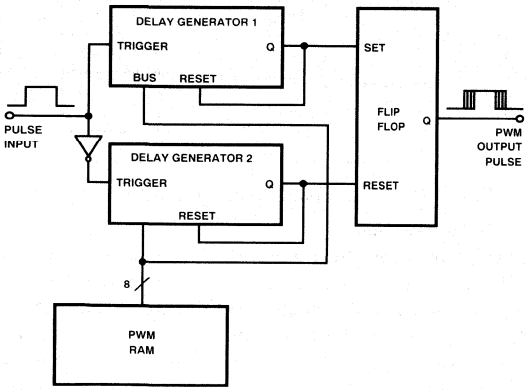

Figure 12. Laser Printer PWM

Layout Considerations

Although the AD9500 and AD9501 are designed to interface with digital signals, they are inherently analog circuits. It is therefore critical to use high speed analog circuit layout techniques. The ground connections to the device should be a low impedance connection to a solid ground plane. The plane should extend under the device to shield it from digital switching signals.

Most socket assemblies add significant interlead capacitance, and should be avoided whenever possible. If sockets must be used, individual pin sockets such as AMP Part Number 6-330808-0 should be used. Prototyping on wirewrap or vector board is strongly discouraged.

Power supply decoupling is also critical for high speed design; a 0.1 μF capacitor should be connected as close as possible to each supply pin.

Even in applications where the devices will be updated at slow rates and the programmed delays will be very long, the above considerations must be taken into account.

ANALOG DEVICES

AN-261
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Programmable-Delay ICs Control System Timing

by Craven Hilton and Jeff Barrow

Low cost, low power, and small package size extend the application of digital-to-time converters in system timing applications. By exploiting the programmability features of these devices, you can both simplify timing-system design and gain greater control of timing parameters than you can by using analog time-delaying methods.

Accurate control of pulse timing is extremely important in digital electronic systems in those applications where system requirements dictate digital control of delays. Until now, you've had to use an analog method, employing a high-speed comparator to detect the incremental delays on a linear ramp, and a D/A converter to set the threshold level of the comparator. This design uses as much as one watt of power; now, however, monolithic digital-to-time converters (DTCs), such as the AD9500, accomplish the same function while only dissipating 300 mW. You can use the AD9500 to control time delays having intervals as small as 10 psec in a full-scale span of 2.5 nsec min.

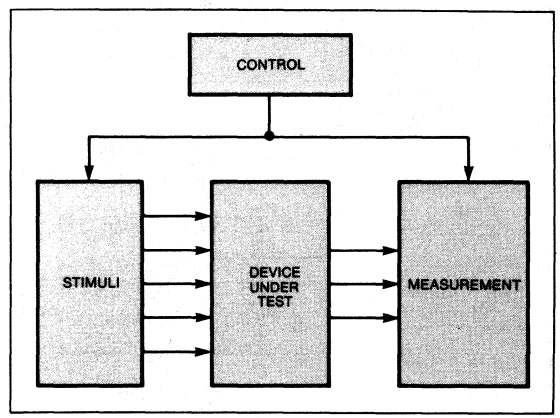

Fig 1—This test configuration, typical of virtually all electronic measurement systems, *comprises four blocks: stimuli, control, measurement, and the device under test.*

Some circuit examples illustrate the benefits of using a monolithic DTC in such applications as LSI and VLSI automatic test systems, which present significant challenges in pulse generation and distribution. For instance, although you can achieve repeatable delays of less than 100 psec by using an analog technique with an RC time reference, this method will not provide you with *variable* delays having such short intervals. The key to the flexibility of the monolithic DTC is the device's programmability.

Fig 1 is a generic block diagram of virtually all

Reprinted from EDN — February 18, 1988

electronic measuring systems. Such a system can evaluate any device (the device under test, or DUT) for virtually any performance criteria if you apply the proper stimuli and use the appropriate measurement circuits. This electronic measuring system will serve as a model for the timing circuits throughout the remainder of the text.

One way to exploit the programmability of the DTC is to use two DTCs triggered from the same clock to program both the leading and trailing edges of an output pulse. This application is illustrated in **Fig 2a.** The first DTC (IC_1), which produces the leading edge of the output pulse, drives the clock input of IC_3, a D-type flip-flop whose D input is tied to a logic one. After IC_1

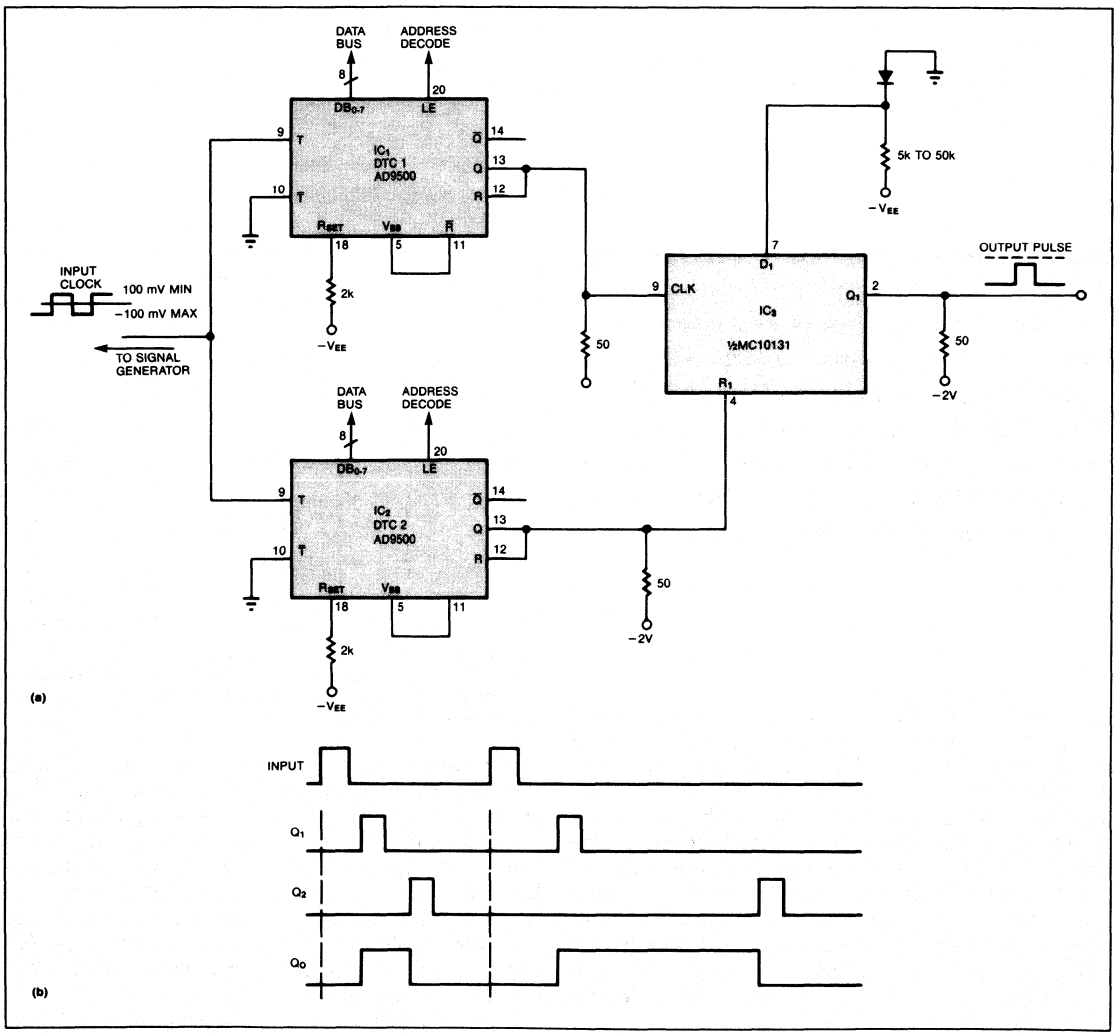

Fig 2—Two DTCs control the output's leading and trailing edges *in this digitally controlled pulse generator. The timing diagram in* ***b*** *brings out the fact that the leading edge of the output pulse, Q_0, occurs after an interval equal to the propagation delay plus the programmed delay.*

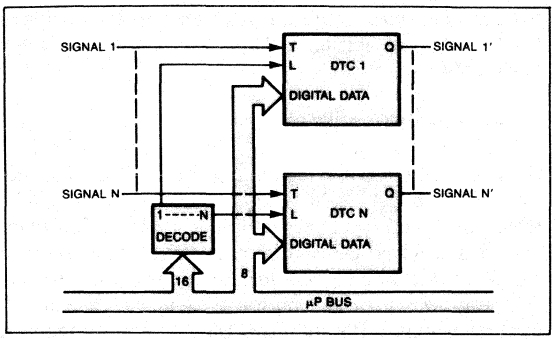

Fig 3—Providing precise delay matching in critical applications, *this circuit uses multiple DTCs to compensate for differences in the delays inherent in different signal paths. The closed-loop circuit provides a deskewing function.*

clocks the one through the flip-flop, the second DTC (IC_2) resets the flip-flop, thereby producing the falling edge.

At a time equal to the propagation delay plus the programmed delay of the first DTC (**Fig 2b**), the flip-flop produces the leading edge of the output pulse. Because the propagation delays of the two DTCs cancel each other, the width of the output pulse is exactly the difference between the programmed delays of the two DTCs. You can determine the programmed delay of each DTC from

$$t_D = t_{PD} + \frac{XX_{16}}{FF_{16}} (R_{SET}\ C_{SET}).$$

The circuit of **Fig 3** provides precise delay matching for those applications in which you need to distribute a

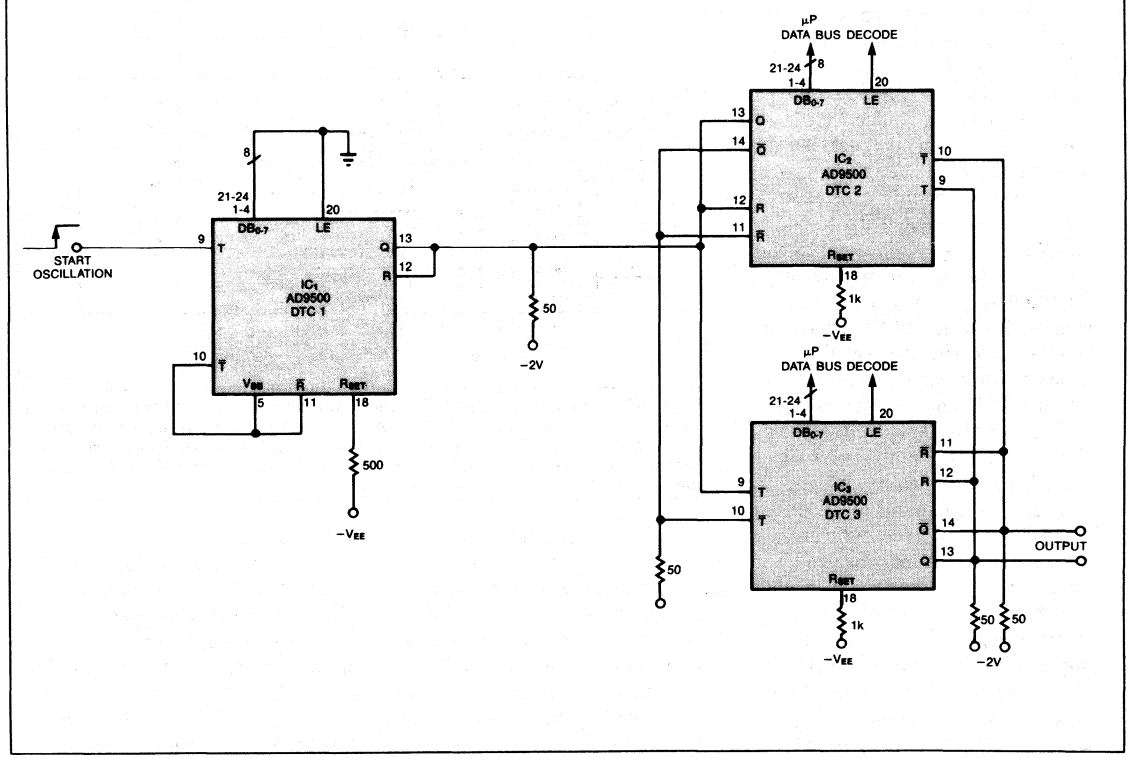

Fig 4—Using DTCs configured to start and stop the oscillation, *this digitally programmable oscillator gives you complete control over the start-up, shutdown, and frequency of oscillation.*

number of pulses and maintain good coherence between those pulses. Because individual test circuits may have extraneous delays in the signal paths to the DUTs, close matching in the initial tester delays will not be sufficient to guarantee close matching between the delivered pulses. The combination of programmable deskewing circuitry and the closed-loop calibration scheme of **Fig 3** allows you to compensate for the timing variations in the circuit paths during your test-system setup cycle.

During the setup cycle, the closed-loop system measures the delay to each input pin of the DUT. It then modifies the delay values stored in each of the DTCs until the input pulses arrive at the DUT's pins simultaneously. This method allows you to match the delays to the DUT to a resolution of 20 psec for a full-scale delay period of 5 nsec.

Programmable oscillator

Fig 4 shows that you can also create a digitally programmable oscillator by using three DTCs. IC_1 acts as a start-up pulser to trigger the other two DTCs, which are configured as astable oscillators interconnected in a wired-OR configuration. You can generate the start-oscillation pulse locally, or you can use the system power-up-reset signal to generate the start-oscillation pulse. Because the DTC is edge triggered and the oscillator is stable in either the oscillating or

DTC uses analog, digital internal circuitry

Fig A displays the AD9500 along with the external circuitry required to configure it as a DTC. An 8-bit word sets the output voltage of the IC's internal D/A converter. The DAC's output voltage in turn establishes a threshold for the high-speed voltage comparator. You can latch the input word to the AD9500 by applying a one to the latch-enable input of the device. Alternatively, if you want to change the value of the input on the fly, hold the latch-enable input at logic zero and the latches will remain transparent.

Because the DTC controls precision delays for high-speed signals, its delay-path inputs and outputs are designed to be ECL compatible. The IC's differential-I/O structure affords maximum timing-noise immunity when you interface the chip to either 10K or 100K logic. For less demanding applications that use 10K ECL circuitry, you can use the on-chip ECL reference and operate the chip in a single-ended mode.

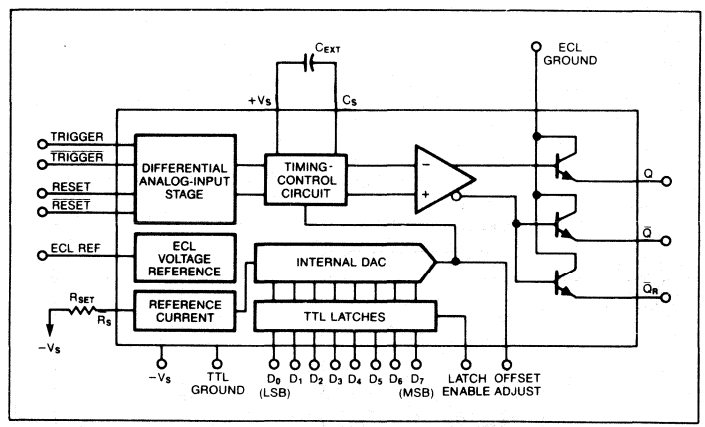

Fig A—Using both analog and digital internal circuitry, *the AD9500 requires only a resistor and a capacitor as external components in programmable-delay applications. R_{SET} and C_{SET} (C_{EXT}+C_{INT}) provide a reference for the internal timing-control circuit. The ramp generated by the R_{SET}/C_{SET} time reference remains linear despite the effects of time, temperature, and supply-voltage variations.*

The DTC's time reference is RC based; it serves as the ramp generator and as a timebase for the on-chip DAC. Because the DAC's gain is proportional to the time reference, any change in the ramp's slope is compensated by the DAC's gain. This compensation reduces the effects of environmental changes on full-scale time and timing linearity.

You determine the full-scale delay of the device through your selection of external, passive components. Although the recommended range of full-scale delays is from 2.5 nsec to 100 μsec, you can extend delays beyond 100 μsec if you can tolerate a degradation of the linearity and

nonoscillating mode, a single pulse from IC_1 will start the oscillation. By grounding the trigger input on either IC_2 or IC_3, you can stop the oscillation.

As **Fig 4** shows, each DTC resets itself as it triggers the alternate DTC. The programmed delay of each device is determined by the equation given earlier. This delay, in turn, determines the output frequency of the oscillator, which is simply the reciprocal of the sum of the two propagation delays plus the two programmed delays.

When you need to measure a time delay, you can use two DTCs in conjunction with two comparators, a D-type flip-flop, and a successive-approximation register (SAR) as illustrated in **Fig 5a.** Flip-flop IC_3 serves as a coincidence detector. The first DTC (IC_1) varies the delay of the pulse applied to the D input of the flip-flop. The coincidence detector serves as a time comparator, whose function is analogous to that of a voltage comparator in a successive-approximation A/D converter.

The clock input to the flip-flop is delayed by a period equal to the unknown ECL delay. The circuit compares the first cycle of the 1-MHz clock with the unknown delay and checks to see if the delay is greater than half-scale. Then it checks for one-quarter or three-quarters scale, one-eighth or seven-eighths scale, and so on. At the end of the test process, then, the output of the SAR provides an 8-bit representation of the delay through the DUT. To measure TTL-circuit delays, you

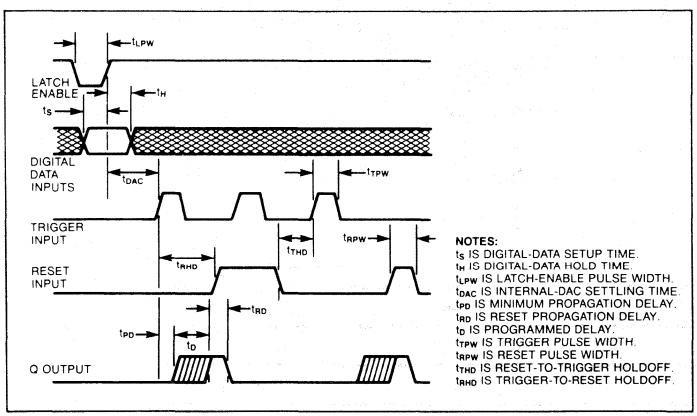

Fig B—Timing characteristics for the digital-to-time converter *are shown in this timing diagram. The illustration shows the relationship between the digital delay-coefficient data and the latch-enable strobe. The text delineates the limits you should observe for the various inputs, in order to obtain proper operation of the DTC.*

Fig B. Lines 1 and 2 show the timing relationship between digital-delay coefficient data and the latch-enable strobe. The minimum latch-enable pulse width is 2 nsec. The data setup time for the input latch is a maximum of 2.5 nsec, and the hold time is a minimum of 4.5 nsec. You must allow at least 25 nsec from the rising edge of the latch-enable pulse before you trigger an event, otherwise the internal DAC might not have time to settle.

Lines 3, 4, and 5 of **Fig B** show the relationship of the output to the reset and trigger events. The total delay through the DTC is the sum of the propagation delay and the programmed delay. The propagation delay equals the delay through the differential input stage, the comparator, and the delay attributable to ignoring the first, nonlinear portion of the ramp. The last of these components increases with full-scale delay.

repeatability of the delay. At the other end of the spectrum, if you choose a full-scale range of 2.5 nsec, then the smallest incremental delay available to you is 10 psec.

The maximum delay trigger rate is 100 MHz, but an offset adjustment in the device allows you to operate two DTCs in a ping-pong fashion to double the trigger rate. The IC's maximum differential nonlinearity is ±½ LSB at 25°C and ±1 LSB over the operating-temperature range. Maximum integral nonlinearity for the device is ±1.25 LSB for full-scale delays of 100 nsec or more over the operating-temperature range.

The timing characteristics of the AD9500 are illustrated in

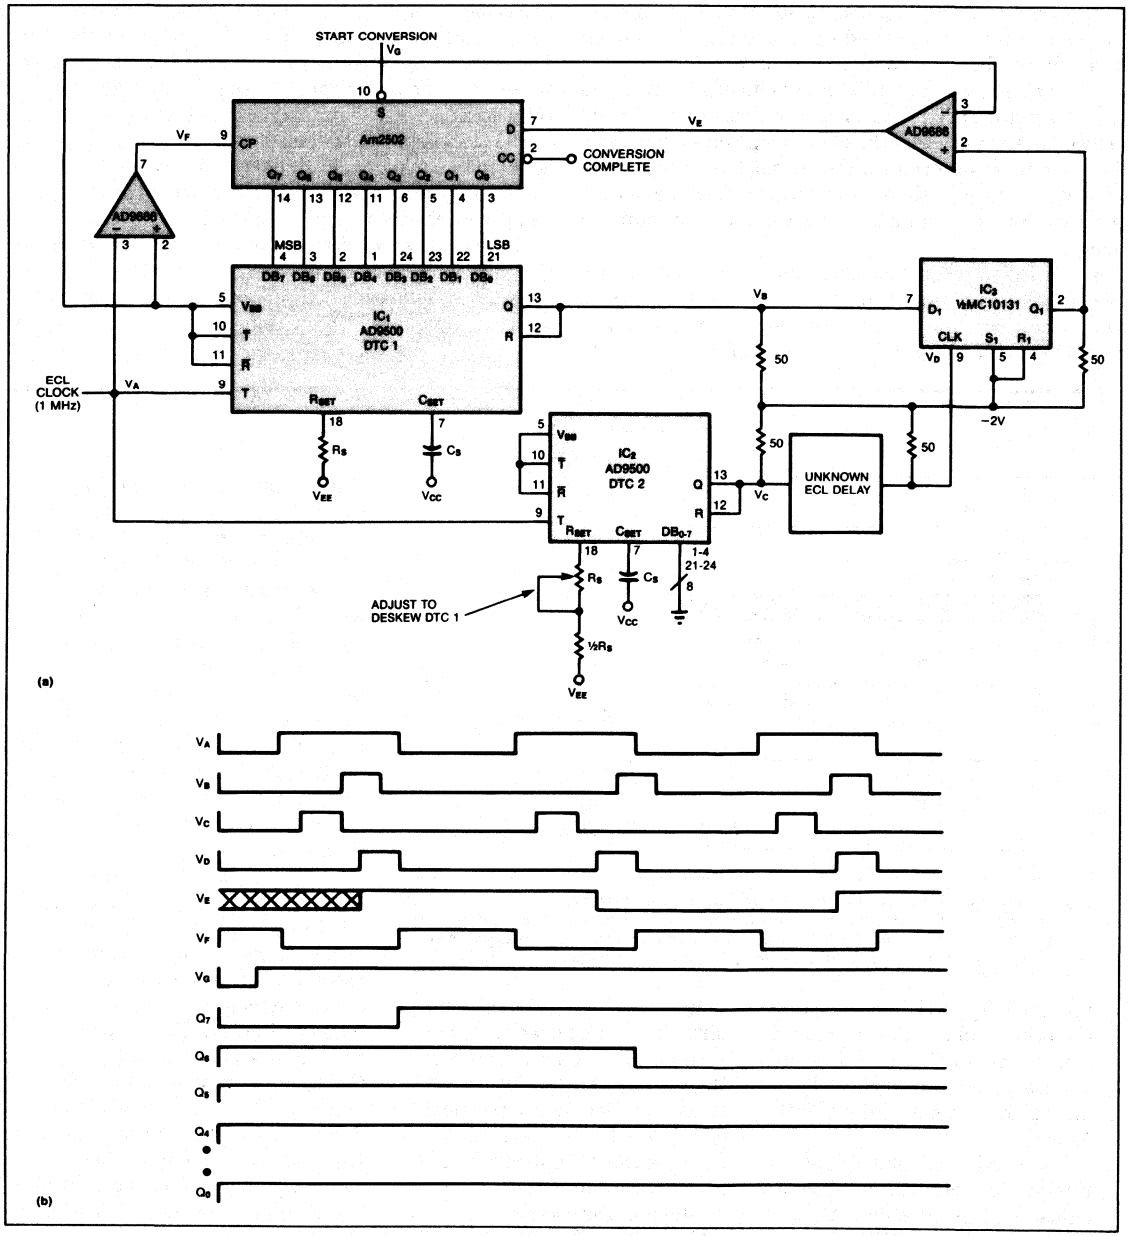

Fig 5—Operating in similar fashion to a successive-approximation A/D converter, *this circuit allows you to make precise delay measurements. The circuit in* ***a*** *handles ultrafast ECL circuitry; you can easily modify it to accommodate TTL circuits. The timing diagram in* ***b*** *illustrates the timing characteristics at various points in the circuit.*

add an AD9686 comparator between the Q output of the second DTC (IC_2) and the input to the unknown-delay circuit, and an AD96685 comparator between the output of the unknown-delay circuit and the flip-flop.

To calibrate the circuit of **Fig 5b,** you insert a shorting strap in place of the unknown-delay circuit to eliminate extraneous circuit delays and the flip-flop's setup time. To null the circuit, apply a digital code of 00_{16} to IC_2 and adjust potentiometer R_S. This adjustment varies the propagation delay through IC_2. The calibration is complete when the output of the SAR is also 00_{16}. You must apply start-conversion pulses during the calibration. This calibration procedure is equally valid for the modified, TTL-delay configuration. **Fig 5b** shows the timing for a typical conversion cycle.

You can use the circuit of **Fig 6a** to measure the settling time of analog signals—for example, the output of a D/A converter. The operation of this circuit is similar to the operation of the digital delay detector, but it uses a voltage-input window comparator in place

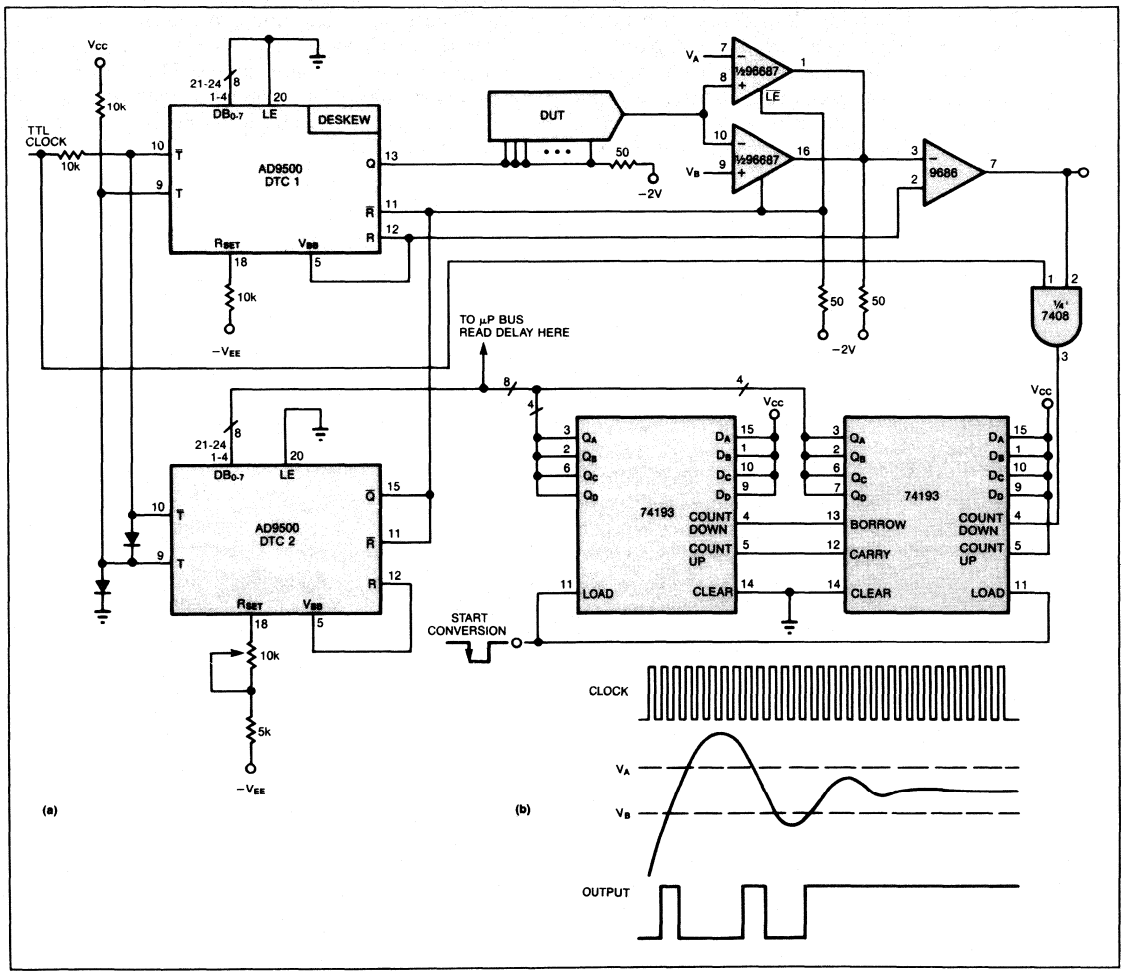

Fig 6—You can measure the settling time of analog signals *by using the circuit in* ***a.*** *The operation of this circuit is similar to that of the digital delay detector in* ***Fig 5a.*** *A typical analog voltage-settling waveform is illustrated in* ***b.***

of the coincidence detector, and down counters in place of the SAR.

Fig 6b shows a typical analog voltage-settling waveform, as well as the output of a window comparator that uses a constant high-frequency strobe. This continuous-clock method produces ambiguous results because the signal comes into the error-band window during three clock periods. The circuit of **Fig 6a,** however, produces a single strobe per cycle of the analog signal and homes in on the correct measurement in the following manner.

The first DTC (IC_1) controls the timing of the DUT switching. The second DTC (IC_2) delays the latch-enable strobe to the high-speed window comparators. The start-conversion pulse initiates the down counters to a full-scale setting. On each cycle of the clock, the counter decrements and the strobe to the window comparator moves back closer to the time when the DUT is switched. Because the circuit starts at full-scal time, the first strobe occurs well after the DUT ha settled. As successive clocks arrive, the circuit cause the strobe to back up until the DUT signal falls out o the range of the window comparator. As a result, th window comparator stops the down counter, whos output represents the settling time of the DUT.

To compensate for extraneous circuit delays, you ca adjust the R_{SET} potentiometer. Insert a shorting stra in place of the DUT and change the window referenc voltages V_A and V_B to $-1.28V$ and $-1.32V$, respective ly. Then adjust R_{SET} until you receive a zero outpu from the down counters. When you reset the voltage for the window around the DUT's output and th conversion takes place, the output of the down counte will represent the propagation delay plus the settlin time of the DUT.

References
Contents

		Page
AN-262	A Low Noise, Low Drop-Out Regulator for Portable Equipment	15–3

15

ANALOG DEVICES

AN-262
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

A Low-Noise Low Drop-Out Regulator for Portable Equipment

by James Wong

Designed for laptop computers and portable instruments, this low-noise, low drop-out regulator achieves all objectives of small size, low power and safe operation.

The primary goal of designing a battery power supply system for laptop computers or portable instruments is to maximize operating time before recharging becomes necessary. Translating the goal into design criteria, the system must have high efficiency and extraordinarily low standby power drain, not only during normal operation, but also during potential short-circuit occurrences. In addition, the regulator circuit must be capable of recovering rapidly and maintaining a stable output even when sudden surges of current occur in the load.

A linear series-pass regulator is particularly attractive when compared to a switchmode design. The linear design offers the following features: low noise, fast response, simplicity, and comparable efficiency. Low noise is inherent because the linear supply has no fast switching circuits. This feature is especially desirable for powering the linear circuits used in sensitive portable instruments where high frequency noise can easily enter high impedance circuits and destroy accuracy. The system can be designed for as wide a bandwidth as necessary, determined by the choice of op amp and its operating current. The linear circuit is quite simple compared to a switcher; fewer component parts helps ensure a reliable design. The full load efficiency can be designed to be comparable to a switcher because the input battery voltage can be chosen to be only slightly higher than the desired output and the input does not vary over wide limits. Furthermore, the zero load quiescent current can be made much lower than a switcher by the proper choice of components.

Figure 1 shows a complete 5V, 1A low drop-out linear voltage regulator that can operate at up to 90% conversion efficiency near drop-out voltage, while drawing less than 50μA of standby current. The design owes its low operating current to several contributing factors. The heart of the regulation loop is an ultra-low supply current op amp, the OP-90, which typically draws only 15μA of supply current. Another OP-90 is used to generate a precision bandgap reference that consumes less than 20μA current. Instead of using a bipolar power transistor for the pass element, a MOSFET power transistor is used. The reasoning behind selecting the MOSFET is that the MOSFET gate draws only picoamps rather than milliamps of steady-state current. In addition, a P-channel MOSFET allows the regulator to operate with an exceptionally low drop-out voltage and still maintain regulation. The details of the design are discussed in the following sections.

The Power Transistor

Using a P-channel MOSFET permits a low drop-out design to be easily implemented. The gate drive is simple because the FET is turned on when the gate is returned to ground. Several types of FETs are available that can handle amperes of current with a gate-source potential of only 5V.

To allow sufficient headroom for maintaining regulation the following inequality must be satisfied:

$$V_B \geq V_o + I_o\, r_{DS(on)}$$

where

V_B = Battery voltage at the limit selected for % discharge

V_o = Output voltage of regulator

I_o = Output current of regulator

$r_{DS(on)}$ = drain-source on-resistance of MOSFET

Reprinted from PCIM magazine – May 1990

The choice of MOSFET used is, as usual, a compromise between conflicting goals. A lower $r_{DS(on)}$ transistor allows V_B to more closely approach V_o. The penalty paid is higher part cost and higher capacitance, which lowers bandwidth. For the design example of *Figure 1*, the IRF 9530 MOSFET was chosen. It has a specified maximum $r_{DS(on)}$ of 0.5Ω. At operating temperature this provides enough headroom to operate satisfactorily when V_B is 6V and I_o is 1A.

The Reference Voltage

To achieve a stable, accurate regulation loop, a fairly precise reference voltage needs to be generated. It is difficult to find a reasonably accurate reference using a zener or a bandgap device that operates at microamps of current. In particular, low power zener diodes generally have poor dynamic resistance characteristics and poor initial device-to-device tolerance. Thus, as the battery voltage varies so will reference voltage accuracy.

However, a pair of matched transistors (MAT-01), closing the loop around an amplifier, easily forms a stable bandgap circuit, as shown in *Figure 1*. The circuit typically draws 17μA of quiescent current yet exhibits low temperature drift characteristics. The circuit requires one adjustment to set the bandgap voltage of 1.230V at the amplifier output. This calibration also establishes a low temperature drift coefficient point. Its typical drift is 5.5μV/°C over a −40°C to +85°C temperature range.

To achieve this high degree of precision, the proper passive components must be used. All resistors, except the one to the FET gate, must be 1% metal film types with a temperature coefficient of 50 ppm/°C or better supplied from the same manufacturer. The trimpot must have a similarly low temperature coefficient. It should also be a multiturn type in order to permit an accurate voltage adjustment.

In applications where low supply current is not of critical importance, the reference trim can be eliminated altogether using a monolithic precision reference device such as the REF-43 in *Figure 2*. The device's quiescent current consumption is on the order of 450μA. Besides offering a temperature stable (10ppm/°C) 2.5V output, its initial tolerance is 0.2% maximum over the full temperature range, which renders trimming unnecessary. In addition, it can tolerate an input drop to as low as 4.5V and still maintain regulation. Finally, the REF-43 has a temperature output that can be used to temperature compensate the circuit, saving components over a discrete circuit scheme.

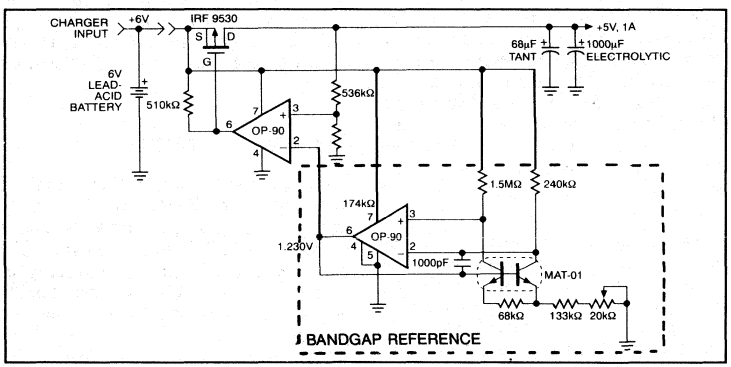

***Figure 1.** A Low Drop-Out Regulator That Supplied 5V, 1A Output. The Circuit Draws 50μA Standby Current.*

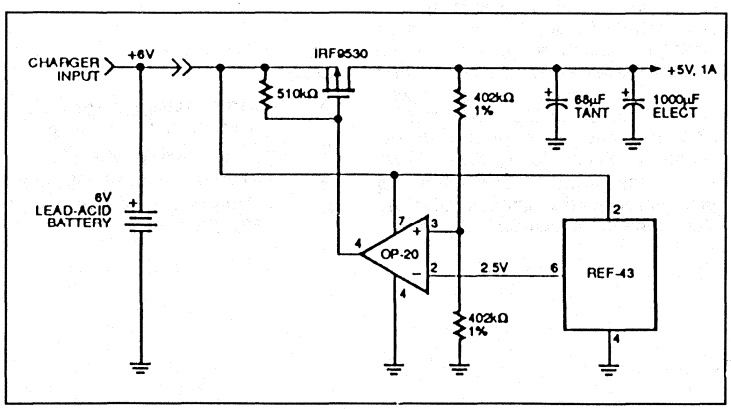

***Figure 2.** Less Critical Supply Current Applications Can Use an Adjustment-Free Design Using a Precision 2.500V Reference Supplies by the REF-43. Quiescent Current of the REF-43 Is 450μA.*

The Op Amp

For this design an ultra-low current, single supply op amp is used to regulate the loop. Of the many op amp types available, few can match the ultra-low 15μA supply current of the OP-90. In addition, the OP-90 has excellent output swing range, which is critical when operating on only 6V of supply voltage. For example, in the case where the regulator is not delivering any load current, the op amp's output must swing to within one volt of the positive supply rail in order to cut off the FET gate drive. Conversely, when a rated one amp current flows, the amplifier output must swing sufficiently low, to within one volt of ground, to ensure adequate enhancement drive to the FET.

The regulation loop is closed by feeding back a portion of the 5V regulated output to the noninverting input of the op amp. Negative feedback is accomplished because the P-channel MOSFET provides a phase inversion.

AC characteristics of the regulation loop are important whenever the load goes through a sudden disturbance. The lower the bandwidth the slower the recovery to a loop perturbation. *Figure 3* shows how the regulator of *Figure 1* reacts to such a load change. Starting from no load, a step change of 150mA of load current was applied to the regulator. Notice in *Figure 3a*, Trace 2, the output initially dipped by some 80mV. It took about 2msec before it recovered to a stable state. The amount of glitch noise and slow recovery may not be acceptable for certain equipment such as precision instruments where glitch noise in the supply can disrupt a sensitive measurement, resulting in errors. On the other hand, if the regulator is used to power logic circuits, the 80mV glitch is usually well within their noise margin.

Faster response can be obtained by substituting an OP-20 as the error amplifier. While the OP-20 consumes 50μA, versus 15μA for the OP-90, its bandwidth is considerably higher. The middle trace of *Figure 2b* shows that both glitch amplitude and recovery time improved significantly. The glitch amplitude is reduced from 80mV to about 35mV and its recovery time shortens from 2msec to 0.5msec.

The bottom trace of the photos of

Figure 2 shows how the two regulator circuits behave when a 150mA step current is applied on top of a normal 750mA quiescent load, which more realistically simulates the real-world condition. Notice in both cases, the peak glitch amplitude is on the order of 10mV, a much more tolerable behavior. The OP-20's performance shows a dramatic improvement in recovery time.

Battery Selection and Charging

The need for laptop computers or portable instruments to run for hours is a foremost design consideration. Advancements in CMOS circuits and battery technology today allow laptop computers to operate for as long as 4 to 6 hours without recharging.

Lead-acid batteries have several desirable characteristics that have caused sealed small cell units to be developed for portable products. One of their characteristics is that they will tolerate a relatively fast charging time, as fast as 2 hours, to reach 90% of capacity. Another is relatively high capacity for the weight and volume. Because of their extremely low internal impedance, they can deliver many amperes of current in a short period of time, although some care needs to be exercised to not deep-drain the battery. In addition, unlike Ni-Cad batteries, lead-acid cells do not have a "memory" problem, which after a while reduces the batteries' capability to recharge to full capacity, resulting in markedly reduced operating time.

Because of the regulator's low drop-out characteristic, it is well suited for use with a 6V lead-acid battery. The combination minimizes power loss and therefore maximizes conversion efficiency, resulting in considerable size and weight savings, particularly in the battery.

A unique characteristic of the lead-acid battery is that it requires a constant voltage source to charge as opposed to a current source charger for Ni-Cad batteries. Consequently, the same voltage regulator design as in *Figure 2* can be used. The battery requires a higher charging voltage of 7.3V so that it can be left float-charging indefinitely to maintain a full charge. A voltage-mode charging scheme also allows easy implementation of a temperature compensation circuit to achieve maximum charging efficiency.

For lead-acid batteries, the charging efficiency decreases as the battery temperature increases. Since charging and ambient environment can increase the battery temperature, a longer than expected charging time occurs. Applying a temperature-dependent charging voltage that has a coefficient of −2.5mV/°C/cell speeds up charging.

To implement a charging circuit that has this temperature coefficient, the previously mentioned temperature output of the REF-43 is utilized as shown in *Figure 4*. Its temperature output pin (pin 3), whose output impedance is quite high, is necessarily buffered with a unity gain amplifier. This pin normally outputs a DC voltage of 630mV at room temperature, with a temperature coefficient of +1.9mV/°C. This voltage is summed into the error amplifier that otherwise regulates at an output voltage of 7.3V. Because of phase inversion, the temperature coefficient of +1.9mV/°C is converted to a −2.5mV/°C at the regulated output by ratio of the 68.1kΩ and the two voltage divider resistors. To be sure the temperature is being measured accurately, it is necessary for the REF-43 to be in direct thermal contact with the battery body.

Current-Limiting Options

Because lead-acid batteries are capable of many amperes of current output for short periods of time, a short-circuit can potentially damage the cell by overheating. In addition, it is not a good idea to allow a deep discharge, as it reduces the charge capacity during subsequent recharges. A simple current-limiting circuit can provide the necessary protection. Two alternative arrangements, easily implemented into the circuits of *Figure 1* or *Figure 2*, are shown in *Figure 5*. To simplify the schematics the reference voltage circuit is omitted.

The first, and simplest, is a constant current limiter using a transistor acting as a threshold switch, plus a current sense resistor and a pull-up resistor. During normal operation, the regulator load current develops less than the 0.6V threshold voltage across the 0.6Ω sense resistor that is necessary to turn on the PNP transistor. Therefore normal voltage regulation ensues. However, as the load current exceeds the threshold of about 1.1A, which develops more than 0.6V across the sense resistor, the PNP transistor begins to draw current and thus pull up the gate voltage which prevents the MOSFET from passing more current. In this manner the PNP will maintain the current at the limit level, to the extreme point where the regulator output voltage begins to drop. Since only a low level of PNP collector current is required, the transistor need only be a general-purpose type, such as a 2N3905.

Although simple in implementation, there is a serious drawback with this design. During normal operation, the load current

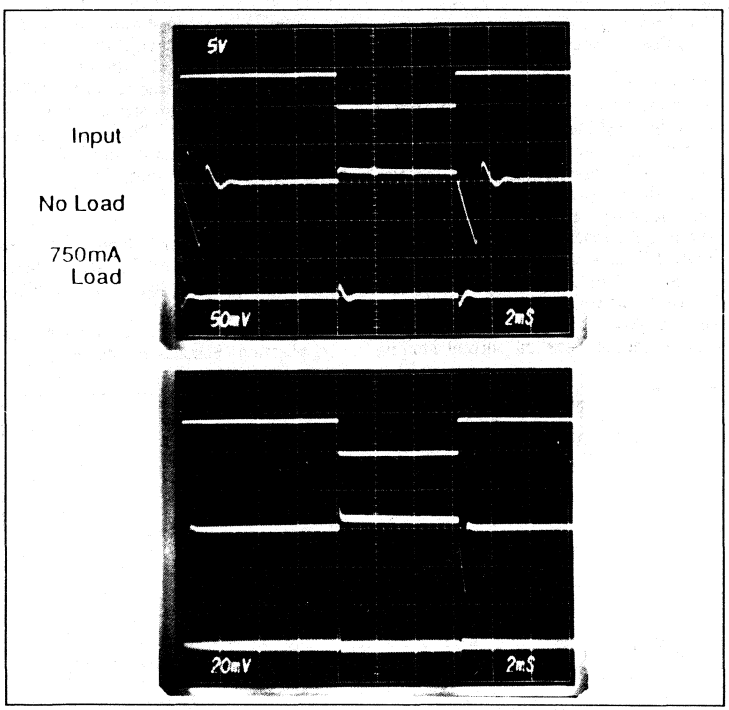

Figure 3. *Comparing a Regulator Load Recovery Between a Slower OP-90(a) and a Wider Bandwidth OP-20(b). The Top Trace Is the Control Signal to the Dynamic Load, Shown for Timing Reference. The Middle Trace Shows the Transient Recovery When a Step Current From No Load to 150mA Is Applied. The Bottom Trace Shows the Output When a Step Change From a DC 750mA to 900mA Occurs.*

drawn might be, say 0.75A. It would develop 0.45V across the sense resistor, which is in series with the pass element. In effect, this 0.45V reduces the drop-out "headroom" by the same amount, which is a serious drawback in a low drop-out design.

The second alternative is shown in *Figure 5b.* Although it costs a bit more circuitry, it is far better in terms of preserving the low drop-out feature. Instead of a 0.6Ω sense resistor, a lower 0.1Ω resistor is used. Thus at the current limit threshold of 1.1A, it only degrades the drop-out voltage by 0.11V, which is far more tolerable than the 0.66V in the other circuit.

The limit threshold voltage is set up by the voltage divider at the noninverting input of the current limit detector/amplifier. Sensing is done by the sense resistor plus the second voltage divider. If the threshold limit is not reached, the inverting input has a higher voltage than the noninverting input, causing the current limit amplifier output to go low, reverse-biasing the diode, thus opening the amplifier loop. Normal voltage regulation ensues. Once a 0.11V is developed across the sense resistor, which is the threshold of the current limit, the amplifier output rises to close the loop. Its low output resistance drives the gate of the MOSFET pass element high to maintain the current limit. This design preserves a low drop-out voltage that is critical to an efficient power system.

Even though constant current limiting saves the battery from self-destruction, sufficient heat sinking must be provided for the power MOSFET transistor in order to prevent it from being damaged under a sustained short-circuit condition. That is, although the output current is limited to 1.1A, a short circuit forces nearly 6V across the power MOSFET causing it to dissipate 6½W.

Foldback current limiting can be used to eliminate the need for a heat sink. Fortunately, such a feature comes free with the design. By moving the previously described current limit sensing circuit from the input side of the MOSFET to the output side as shown in *Figure 6,* foldback current limiting action results. The current limiting action initially works the same way as previously described. The foldback action kicks in as the current limit of 1.1A is reached, the output voltage begins to drop, pulling the threshold voltage on pin 6 down proportionally. As it drops, less and less voltage is needed across the sense resistor to trip the threshold. In effect, the regulated current "folds" back as the output voltage drops. For example, when the output drops to 1V, the output current drops to about 0.2A.

The significance of this design is that the power MOSFET is called on to dissipate less than 1W instead of more than 6W of power when in short circuit. The worst-case dissipation is indeed 2W when the output is three volts and the load current is 0.65A. This reduces heat sink size and improves operating reliability. Thus the design achieves all objectives of small size, low power and safe operation.

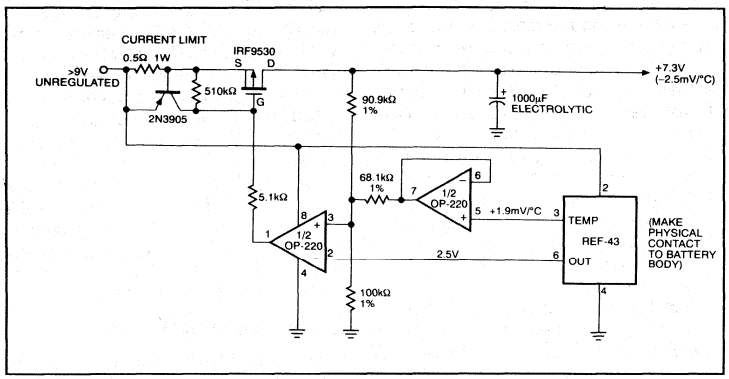

Figure 4. *Lead-Acid Battery Requires Simple Regulated Voltage Charging Circuit. Optimum Charging Efficiency Can Be Achieved by Providing a Temperature Sense and Compensation. Charging Voltage Has Temperature Coefficient of –2.5mV/°C/Cell.*

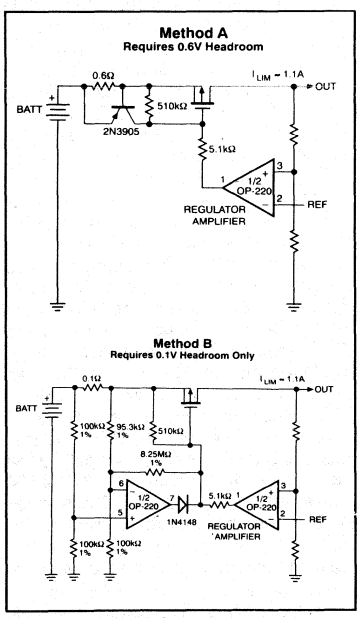

Figure 5. *Two Alternate Methods of Implementing Current Limiting. Method A(a) Is Simpler But Has Poor Drop-Out Characteristic. Method B(b) Is Preferred Because It Preserves Low Drop-Out Features.*

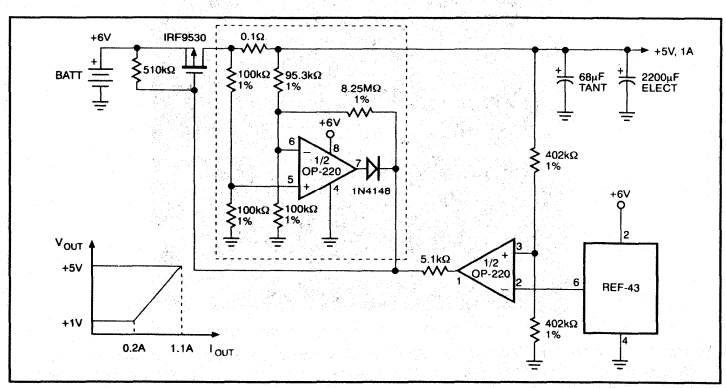

Figure 6. *By Moving the Current Limiting Circuit to the Output Side of the Regulation Loop, a Foldback Effect Is Achieved. As the Output Voltage Drops, the Output Current Also Drops, Reducing Power Dissipation in the Power MOSFET Device.*

References

1. Derek Bowers, Art Kappor, Andrew Jenkins, "Low Power Op Amp Delivers Precision at Low Signal Levels," EDN January 1987.
2. Precision Monolithics, Inc. 1989 Analog Applications Seminar Handbook.

Resolver-to-Digital Converters
Contents

		Page
AN-263	Resolver-to-Digital Conversion—A Simple and Cost Effective Alternative to Optical Shaft Encoders	16–3
AN-264	Dynamic Characteristics of Tracking Converters	16–7
AN-265	Circuit Applications of the AD2S81A and AD2S80A Resolver-to-Digital Converters	16–11
AN-252	Using the 2S80 Series Resolver-to-Digital Converters with Synchros: Solid-State Scott-T Circuit	16–17
AN-266	Passive Component Selection and Dynamic Modeling for the 2S80 Series Resolver-to-Digital Converters	16–21
AN-267	Vector Control Using a Single Vector Rotation Semiconductor for Induction and Permanent Magnet Motors	16–33
AN-230	Circuit Applications of the AD2S90 Resolver-to-Digital Converter	16–45
AN-234	Digital Resolver Integration	16–49
AN-236	Speed Measurement and Control with the AD2S83, a "Silicon Tachogenerator"	16–53

ANALOG DEVICES

AN-263
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Resolver-to-Digital Conversion — A Simple and Cost Effective Alternative to Optical Shaft Encoders

by John Gasking

INTRODUCTION

Ever since man invented the wheel, we have wanted to know, with varying degrees of accuracy the position of that wheel. This measurement of angular position is a craft that has been refined through the ages and accelerated as production and control methods have become more sophisticated. Most forms of motive power finish with a rotating shaft, the most popular example being the electric motor, and even linear motion is produced by mechanically converting that rotary motion with screw threads, cranks, etc. The basis, therefore, of most mechanical positioning systems is the measurement of the shaft angle. The Greeks taught us many centuries ago how to deduce position from angles, today's computing power just allows us to make the deduction more quickly.

Angular and linear measurement – encoders and resolvers . . . to a digital output.

There are fundamentally two alternatives open to the designer when it comes to choosing the type of angular transducer, each with its particular characteristics. I shall describe the alternative so that, what I consider are the advantages of a resolver, can be better understood.

Encoders fall into two groups, incremental and absolute. Incremental encoders simply count the passing of a division of the circle and give pulsed outputs that allow you to store that count and know its direction of rotation. This is known as the A Quad B output system illustrated in Figure 1, the direction is deduced from the occurrence of the edges of the A and B pulse trains. An "A" 0 to 1 transition occurs before a "B" transition with one direction of rotation and vice versa for the opposite rotation. A datum pulse is generated synchronously with A and B as the shaft rotates through zero. Incremental encoders usually produce these pulses from photo-electric devices, which precludes them from use in radiation hazardous areas such as may be required by military and aerospace applications.

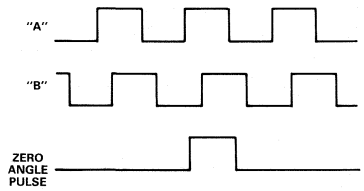

Figure 1.

Two disadvantages of incremental encoders are that the position is not known on "power up" and electrical interference can cause false counts.

Absolute encoders, as their name implies, give a parallel digital output that is generated from a pattern that is on a rotating disc attached to the shaft. The sensors used can be electrical contacts or a photo-electric system. Various codes are available, binary and Gray being the most popular. Very high resolutions and accuracies are available, 16 bits (20 arc seconds) and beyond. These are at very high cost, $1000's, and have the problem of parallel data transmission if the encoder is remote from the measuring electronics.

Probably one of the major reasons for the use of encoders is that it is obvious how they work and they give a digital output. Resolvers on the other hand, are a little less understood and the conversion appears complex. Recent advances in our conversion techniques and manufacturing technology makes application so simple that a complete understanding becomes less important.

16

The resolver is a rotating transformer whose output analog voltages are uniquely related to its input shaft angle. It is, therefore, an absolute position transducer with 0 to 360°of rotation.

The resolver as an angle measurement transducer has a number of advantages. Firstly, the resolver is a robust mechanical device that can withstand severe environments of dust, oil, temperature, shock and radiation. Secondly, being a transformer it provides signal isolation and a natural common-mode rejection of electrical interference. This feature coupled with the fact that only four wires are necessary for the angular data transmission, makes it unique in angle measurement and ideally suited

to the harsh world encountered in the heavy manufacturing and aerospace industry. Today, brushless resolvers are available that have no slip ring connections to the rotor, which greatly increases their life and reliability.

The resolver can be positioned where the angle needs to be measured and the electronics can be positioned where the digital output needs to be, in the card frame with the processor.

There are two methods of using a resolver to obtain output voltages related to the shaft angle.

In the first method the rotor winding (Figure 2) is excited by an alternating signal and the output is taken from the two stator windings. As the stator windings are mechanically disposed at right angles, the output signal amplitudes are related by the trigonometric sine and cosine of the shaft angle. Both the sine and cosine output signals have the same phase as the original excitation signal, only their amplitudes are modulated by sine and cosine as the shaft rotates.

The second method of obtaining an output signal uniquely related to shaft angle is to excite the two stator windings. When the two stator windings are excited by two alternating signals that are in phase quadrature to each other, then a voltage is induced in the rotor winding whose amplitude and frequency are fixed but whose phase varies with shaft angle. Consider Figure 2. When the shaft aligns the rotor winding R_1R_2 with the stator windings S_1S_3 the rotor output signal will be at a maximum of the 0° phase excitation signal. When the rotor windings R_1R_2 align with the stator windings S_2S_4 the output signal will be at a maximum of the 90° phase excitation signal. At angles between these two positions, the phase of the output signal will vary linearly with angle from 0° to 90° phase. By rotation of the shaft through the full 360°, the phase of the output rotor signal will itself vary from 0° through 90°, 180°, 270° to 360°, i.e., back to 0° phase.

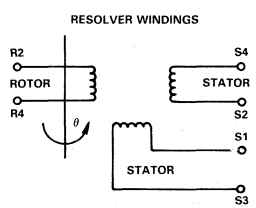

Figure 2.

DIGITAL CONVERSION

Most intelligent robots and numericaly controlled machine tools use digital computers as the calculation and memory element of the system. Therefore, in order that the computer understands the situation, all the angles being measured must be translated into digital words.

In order to fully appreciate the advantages of the tracking resolver-to-digital converter, it is necessary to understand the major alternatives.

Apart from the tracking converter, there are two other methods of conversion widely used in industry. Both methods have some severe weaknesses as the demands of industry for high accuracy increases.

The alternatives to the tracking converter have until now been cost effective, however, with the improvements in our techniques, the cost of the converter has reduced to the point that many users are making a reappraisal and switching to the tracking converter.

The most widely used method at the present time is the method known as the "Phase Analog Technique". This method adopts the mode of stator winding excitation by two signals in phase quadrature shown in Figure 3. In order to measure the shaft angle, accurate measurement of the phase of the rotor signal with respect to the excitation signal is necessary. A block diagram of such a scheme is shown in Figure 3.

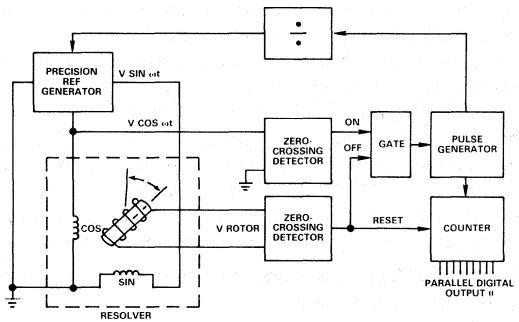

Figure 3.

In essence, the time interval between a zero crossing of the sine or 0° excitation signal and the zero crossing of the rotor signal is measured (see Figure 4). This is done by counting the numbers of pulses from the original reference clock oscillator that occur between the two zero crossings. By varying the division ratio between the reference clock oscillator and the two excitation generators, the resolution to which this measurement is made can be varied.

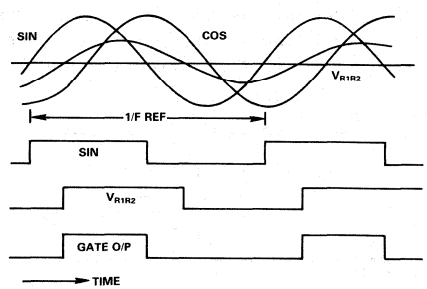

Figure 4.

*Footnote: The circle is divided into 360 degrees, it is then further subdivided into minutes and seconds. So that minutes and seconds of arc were not confused with minutes and seconds of time, the subdivisions are always referred to as arc minutes and arc seconds. Sixty arc seconds = 1 arc minute. Sixty arc minutes = 1 degree.

This division ratio has little effect on accuracy as that is determined by the accuracy to which the zero crossing intervals can be measured. The first and most obvious source of error is the electrical noise generated by the environment of the resolver causing noise on the rotor signal which makes the zero crossing point become indeterminate. The second and more subtle source of error comes from variations in the excitation. Any variation of the relative amplitudes or phases of the two excitation signals directly influences the phase of the output signal. As most systems use a square wave clock with integrated circuit dividers, very precise filtering of the signals is required to obtain the pure sinusoidal signal necessary. Such accurate filters are complex and it is difficult to obtain a good temperature versus phase coefficient. Any harmonic distortion of either or both the excitation signals will also affect the phase of the rotor signal. The summation of all, these errors limit the accuracy of the "Phase Analog" approach to about 10 bits or 21 arc mins.

The second most widely used method of resolver-to-digital conversion method is a sampling technique. There are many variations of this technique, however, in essence, they all make a sample of the sine and cosine output signals of a rotor excited resolver. The sample of the signal amplitudes is taken at a peak of the reference input amplitude and converted to digital signals by an A/D converter. The resulting digital words are used as a memory address to "look-up" the shaft angles in a processor.

The difficulties with such an apparently simple method comes from its total inability to deal with noise. If a noise disturbance occurs on the signal lines at the time of sampling, then a wrong answer results. If the noise only results in a single wrong reading, then the frequency pass band of the servomotor drive system acts as a filter and little error results. If, however, the noise is severe, then the errors can become unacceptable.

TRACKING RESOLVER-TO-DIGITAL CONVERISON

The tracking conversion techniques used in all our converters overcomes all the difficulties experienced by the other methods. The new 1S series of converters in particular compete with the other methods cost wise and yet provide superior accuracy, noise-immune operation. The basic method of operation and connection is shown in Figure 6.

A tracking converter operates ratiometrically, that means that it is only concerned with the ratio of the sine and cosine outputs of a rotor excited resolver. As the resolver acts as a transformer, any excitation waveform distortion or amplitude variation appears in the correct ratio on both sine and cosine and has little effect on accuracy. A tracking converter contains a phase demodulator and therefore frequency variation and incoherent noise do not affect accuracy. Tracking converters can operate with any reference excitation from square waves to triangular waves with only minor accuracy variation. Common-mode rejection is achieved by the natural isolation of the resolver, this combined with the integration of the tracking loop makes an angular-to-digital conversion system with excellent noise immunity.

As a tracking converter uses an internal servo system that contains integration in the control loop, there is little degradation of accuracy even at high input revolution rates up to more than 40,000 revs per minute.

APPLYING THE RESOLVER AND TRACKING CONVERTER

The most common misapprehension about resolvers concerns their working voltages and frequencies.

As the resolver is a linear transformer in which the primary winding rotates inside the secondary windings, it behaves as a normal transformer. There is a transformation ratio that is fixed within certain frequency limitations specified by the manufacturer. The rotor excitation or reference input can be varied to give the stator output voltages that are required to be directly connected to the converter inputs.

Variation of the reference frequency can cause changes in the phase angle between the stator rotor or signal outputs and the reference or rotor input. As the tracking converter is tolerant to both phase and amplitude, often no correction is needed. However, when using the 1S64 or 1S74 series of converters which provide a high quality velocity signal in addition to the digital position, it is essential that the phase shift is zero. This can be achieved by selecting the appropriate frequency, or using a resistor/capacitor phase shift network on the converter's reference input. (See also section "Maximum Lead Length Between Resolver and Converter".)

REFERENCE FREQUENCY

Converter tracking rate, that is the maximum speed up to which the input angle can be followed by the converter, is related directly to the reference frequency.

A rotor excited resolver when rotating gives an output that is a modulated carrier signal.

$V_{S153} = \text{Sin}\ \omega_2 t\, V_R \text{Sin}\ \omega_1 t$

where V_R = peak reference voltage
ω_1 = reference frequency in radians/sec.
ω_2 = shaft rotational frequency in radians/sec.

In order to convey the information ω_1 must always be greater ω_2.

REFERENCE OSCILLATOR

We currently offer a hybrid power oscillator, OSC1758, which is capable of supply 1.4VA.

The OSC1758 is in fact in two sections namely an oscillator and a power amplifier. They are separately accessible.

The frequency of the device is determined by two capacitors C1 and C2.

Remember that the tracking conversion technique is very tolerant on amplitude and frequency stability and, therefore, any drift in the capacitors will not unduly affect the

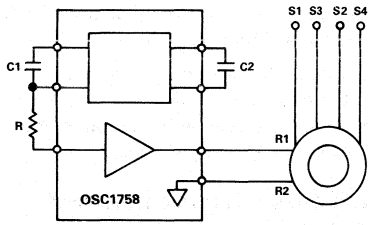

FREQUENCY	=	5kHz
RESOLVER ROTOR IMPEDANCE (Z_{SO})	=	$850 + j\,1250$
$\lvert Z_{SO}\rvert$	=	$\sqrt{850^2 + 1250^2}$
	=	1511Ω
TRANSFORMATION RATIO	=	0.5 = T.R.
RESOLVER PHASE SHIFT @ 5kHZ	=	−5°
REFERENCE VOLTAGE REQUIRED (V_R)	=	2/T.R. (NOTE 2V rms IS REQUIRED BY THE 1S CONVERTER)
V_R	=	2/0.5 = 4V rms
REFERENCE CURRENT	=	$\frac{V_R}{Z_{SO}} = \frac{4}{1511} = 2.6\text{mA}$
RESISTOR R	=	$\frac{37.5 \times 10^3}{V_{OUT}} - 5350\Omega$
	=	$\frac{37.5 \times 10^3}{4} - 5350$
	=	4025Ω
CAPACITORS C1 = C2	=	$\frac{1}{F_{OSC} \times 10^5}$ FARADS
	=	$\frac{1}{5.0 \times 10^3 \times 10^5} = 0.002\mu F$

Figure 5.

accuracy or performance of the overall system. However, it is suggested that high grade silver mica capacitors are used with a temperature coefficient of less than 50ppm/°C.

Connection to a typical brushless resolver is shown in Figure 5.

From the above it can be seen that the drive requirement is well within the capability of the OSC1758 and that many more resolvers could be powered from the same source when multichannel angular measurements are necessary.

MAXIMUM LEAD LENGTHS BETWEEN RESOLVER AND CONVERTER

The question is often asked as to the maximum length of lead between the resolver and the converter.

The answer is that this distance can be very long indeed providing the following simple precautions are taken.

1. It can be shown mathematically that as long as the phase shift between Sine/Cosine and Reference signals is zero, the additional errors due to capacitive effects of the interconnecting lead will be zero. The phase shift can be zero'd either by selecting the frequency at which zero phase shift occurs across the resolver or by introducing a phase lead or lag into the reference input to the converter.

   Simple phase lead and lag circuits are shown below.

PHASE LEAD

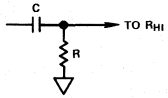

$$\text{Lead} = \text{ARC TAN}\frac{1}{2\pi fRC}$$

where f is the frequency.

PHASE LAG

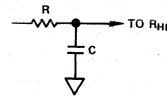

$$\text{Lag} = \text{ARC TAN}\,2\pi fRC$$

2. Ensure that the reference frequency is such that:

$$\frac{1}{2\pi fcL}$$

   is greater than X where:

   f = reference frequency

   c = Lead capacitance per metre of the cable.

   L = Distance between resolver and converter in metres.

   X = The reactive component of the resolver output impedance.

3. Ensure that the attenuation of the signal (sine/cosine) does not lower the voltage appearing at the inputs to the converter by less than 10% of the nominal voltage specified.

4. Use individually screened twisted pair leads for sine, cosine and reference signals. On four-wire output resolvers driving into converters which have only a "Sine", "Cosine", and "A GND" input, the common connection should be made at "A GND" on the converter.

Provided that the above precautions are followed, it is possible to transmit the angular data hundreds of metres without loss of accuracy and with the very considerable advantage of high-noise immunity.

CONCLUSION

The objective of this note was to indicate the simplicity and advantages of using resolvers in the cost effective measurement of the position. While such measuring devices are used within sophisticated control systems, recommendations beyond the digital interface are outside our brief as a component manufacturer, except that we recognize that they exist and have requirements of the transducer systems. To this end, we offer the best performance we can and publish with our data the Transfer Functions and dynamic specifications that allow system designers to calculate our converter's performance within their control circuits.

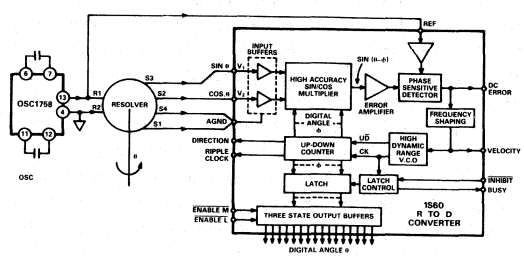

Figure 6.

ANALOG DEVICES

AN-264
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Dynamic Characteristics of Tracking Converters

by Mark Thomas

It is easy to get confused by the terms used to describe the dynamic performance of tracking resolver-to-digital converters:

"What is the *bandwidth* of this converter?"

"*Acceleration constant*! – what does that tell me about the performance?"

"I want to increase the *resolution*, how will that affect the *tracking rate*?"

The aim of this application note is to give a simple explanation of the meaning of the dynamic characteristic specifications, and to explain their particular relevance in the case of synchro- and resolver-to-digital converters.

All synchro/resolver-to-digital converters manufactured by Analog Devices utilize a Type 2 tracking loop in order to convert the analog input signal into a digital output signal.

Examination of the block diagram of a typical converter will reveal that it essentially consists of six circuit elements connected in a loop; these being the Control Transformer (Sine/Cosine Multipliers), Error Amp, Phase Sensitive Detector, Integrator, Voltage Controlled Oscillator and Up/Down Counter. The purpose of the converter is to produce a digital representation of the input (analog) signal. It is a closed-loop control system.

All closed-loop control systems can be classified according to the nature of their steady-state errors (i.e., the state where all transient signals have decayed), and in the particular case of the Type 2 system, the steady-state error is zero for both a stationary input signal and a constantly varying (constant velocity) input signal. The only time when the error signal is present is during periods of acceleration or deceleration.

The advantage of using a Type 2 loop is that the digital output will always represent the analog input (within the device accuracy specification) under the conditions of a stationary input and also a constantly varying (constant velocity) input, i.e., there is not velocity lag.

When considering the dynamic performance of the converter, the key specifications which will be of interest to the system designer will usually be found under the heading "Dynamic Characteristics"; the specifications being Bandwidth, Acceleration Constant, Reference Frequency, Tracking Rate, Velocity Scaling, Settling Time and Resolution. Let us look at each of these in turn.

Bandwidth

Being a Type 2 tracking loop, the converter is essentially a second order low pass filter. It has a predictable dynamic response, the critical frequency and damping characteristics being set by the values of the components which make up the loop.

Until very recently, the bandwidth of the converters has been set by the designer and could not be varied

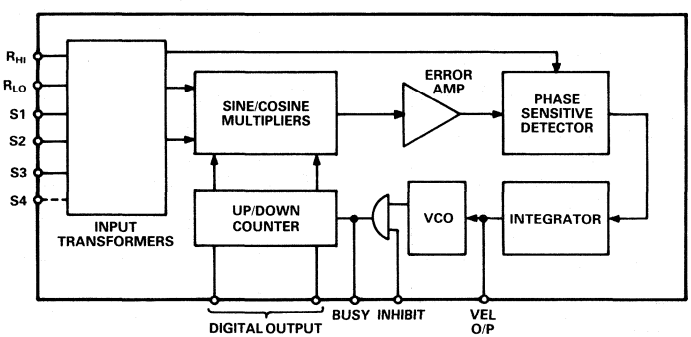

by the user. This meant that the user could not optimize the performance of the device in order to suit the system requirements. However, with the release by Analog Devices of the monolithic converter, the user can, for the first time, optimize the bandwidth and other dynamic characteristics.

The bandwidth of the converter indicates the small signal response of the loop, and is shown diagrammatically on the data sheets as a plot of gain and phase against frequency.

The plots show how the converter will respond if a small signal, sinusoidally varying about a fixed point, is applied to the input of the device. It is important to stress that the gain and phase plots only indicate the response of the device to small changes in the input signal; the sinusoidally varying signal must not cause saturation in any component part of the loop, and will in general have a maximum amplitude of about 5 degrees. The output from the converter will also be a small signal, sinusoidally varying about a fixed point; the gain and phase relationship between the input and output signals are shown on the plots. The frequency axis refers to the frequency of the sinusoidally varying input signal, and not, it should be stressed, to the frequency of the reference carrier. The bandwidth figure quoted for the converter is defined as the frequency of the sinusoidally varying input signal, at the point where the device has a gain of −3dB. For example, in the case of the 1S74 hybrid converter, the bandwidth quoted on the data sheet is 230Hz. If we apply a sinusoidally varying input signal with a maximum amplitude of 1 degree, oscillating at 230Hz about a fixed point, we should observe a sinusoidally varying output signal, with a maximum amplitude of about 0.7 degrees ($1/\sqrt{2} \times 1$ deg) and lagging by almost 90 degrees.

Both the bandwidth and the transfer function from which it is derived, are calculated specifications based on the individual transfer functions of the component parts of the converter loop. They allow the system designer to evaluate the frequency response of the converter and, in combination with the other components of the system, the frequency response of the overall system. The effect of increasing the bandwidth is to allow the converter to respond to faster (higher frequency) small changes in the input signal; hence, the response time is shorter. However, an increasing bandwidth also leads to an increase in the susceptibility to noise, as the higher frequency noise signals come into the range of frequencies "seen" by the converter (remember that the converter is effectively a low pass filter). One of the major detrimental effects of this increase in the susceptibility to noise, is the appearance of the noise on the analog velocity output signal.

Acceleration Constant (K_a)

This parameter specifies the ability of the converter to withstand and track input accelerations. It is specified in the data sheet with the dimension of "/sec^2" and is defined as:

$$K_a = \frac{\text{Input Acceleration}}{\text{Error in Output Angle}}$$

Once the K_a is known, the acceleration is defined in any units, e.g., with a K_a equal to 227,555 /sec^2, we can immediately see that there will be 1 arc minute of additional error for an input acceleration of 227,555 arc min/sec^2, or 1 degree of additional error for an acceleration of 227,555 degrees/sec^2.

However, the acceleration figure quoted in the specification is not the maximum acceleration which the converter can withstand. This maximum figure is governed by the maximum allowable output swing of the internal error amplifier. If this maximum acceleration is exceeded, then the digital output will irrecoverably lose track with the input. In general, the maximum acceleration will be reached when the error between the input and digital output is approximately 5 degrees. Therefore, a converter can withstand an acceleration in deg/sec^2 of about five times K_a value.

Reference Frequency

In general, the higher the bandwidth of the converter, the higher is the required reference frequency, which will also result in a faster possible tracking rate and shorter settling time. The only limitations on the maximum frequency the reference can have are the limits set by the components used in the converter loop. All converters have options available which will function with a 2.6kHz reference frequency, and many hybrids can operate up to 10kHz reference. The new monolithic converters will function with a reference frequency fo 20kHz. In order for the demod-

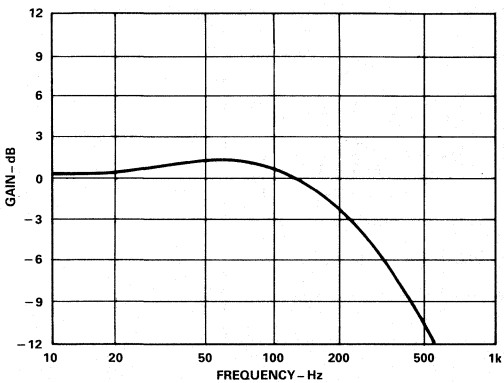

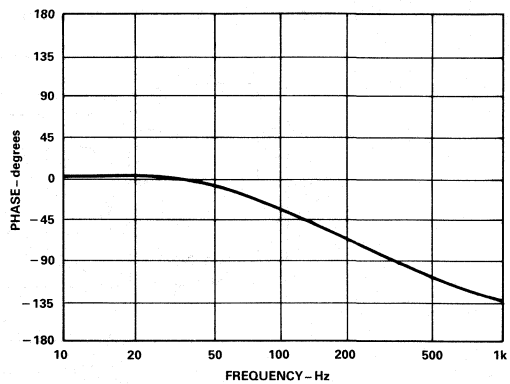

ulator to reject higher order multiples of the reference frequency, the reference frequency of the converter should be at least 2.5 times the closed-loop bandwidth; hence, possible bandwidth is limited by reference frequency used.

Tracking Rate

This is defined as the maximum angular speed for which the converter output will be able to keep track with the converter input. It should be noted that the tracking rate figure quoted in the data sheet actually indicates the guaranteed minimum limit to a range of tracking rates acceptable to the converter. This converter will, therefore, keep track in both directions with inputs varying from stationary up to the tracking rate figure quoted in the data sheet. However, due to component tolerances within the converter loop, the maximum limit of this tracking rate range may well exceed the figure quoted, and can vary from device to device. The higher the reference frequency at which the converter is intended to operate, then the higher will be the maximum possible tracking rate. However, the actual limit for the maximum tracking rate is set by the integrator and Voltage Controlled Oscillator (VCO).

The maximum recommended rate for the VCO sets the maximum possible tracking rate for the converter, e.g., if the maximum recommended VCO rate is 1MHz, then with a 12-bit converter the maximum possible tracking rate will be set a 1,000,000/4096 revs/sec.

The input to the VCO is obtained from the output of the integrator, and so the maximum voltage swing available at the integrator output will also limit the maximum tracking rate. In general, if the power supply rails are lowered, the maximum voltage swing available from the integrator will be lowered and the maximum possible tracking rate will also decrease. (N.B. this does not happen in all converters.) Signal voltage levels which are outside the specified limits for the converter will also degrade the dynamic performance, and the overall accuracy of the device.

Velocity Scaling

Closely linked to tracking rate is the scaling of the analog velocity signal output. The VCO rate is fixed for a given input current by the VCO scaling factor which relates the VCO output frequency to the input current. The input current is scaled by the velocity scaling resistor which is set by the manufacturer for the majority of converters. However, in the case of the monolithic converters, the value of this velocity scaling resistor may be selected by the user, allowing the velocity output to be scaled down to suit the user's requirements.

Settling Time

If a step input is applied to the converter, a finite time will be required for the internal loop to null and produce the required digital output. In all the data sheets, the settling time is given as the time required for the converter to settle within 1LSB of the quoted accuracy in response to a step input of 179 degrees (i.e., the worst possible case). Therefore, this specification gives an indication of the large signal response of the converter.

A typical settling time curve for steps of less than 179 degrees is shown below.

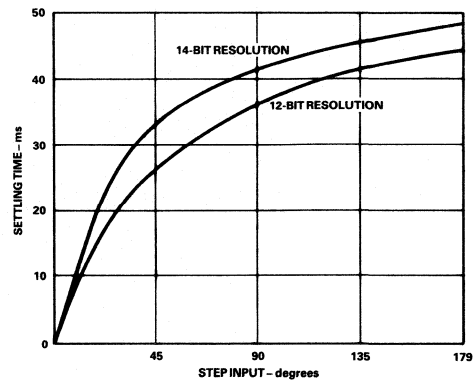

It can be seen that for step inputs within approximately 5-10 degrees, the response is linear, and can be derived from the small signal response of the converter (i.e., see bandwidth). Beyond this region, the internal error voltage will exceed the linear range of the converter.

Resolution

The VCO output clocks the up/down counter; therefore, for a given VCO output rate, increasing resolution will lead to a lengthening of the settling time of the output bits. The converters available have either a 10-, 12-, 14- or 16-bit output resolution, or in particular cases, the resolution may be varied to have any of these four output resolutions. Increasing the resolution from N to N +2 bits will lead to a factor of 4 decrease in the tracking rate and an increase in the settling time.

ANALOG DEVICES

AN-265
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Circuit Applications of the AD2S81A and AD2S80A Resolver-to-Digital Converters

by Hitesh Patel

This application note discusses some useful circuit applications of the AD2S80A and the AD2S81A monolithic resolver-to-digital converters.

It does not discuss the basic operation or specifications of either of the two converters; this information may be found in the data sheets. The following applications are discussed:

1. Interfacing to High Voltage Synchros and Resolvers
2. Connecting the Converter to a 15V Supply
3. Using the AD2S80A as a Control Transformer
4. Using the Converters with External Pitch or Revolution Counters
5. Simulating an Incremental Encoder Using the AD2S80A or AD2S81A
6. Connecting Inductosyns* to the Converters
7. Using the AD2S81A/AD2S80A to Build a Two Speed Coarse/Fine System.

The AD2S81A is a 12-bit tracking R/D converter packaged in a 28-pin ceramic DIP. The converter can track signals at rates up to 260 revolutions per second. Users can set the dynamic performance of the converter with external components, providing greater flexibility in tailoring the AD2S81A to suit system requirements.

The AD2S80A is a monolithic variable resolution, 10-, 12-, 14- or 16-bit tracking R/D converter packaged in a 40-pin DIP. The AD2S80A offers accuracy up to 2 arc mins ±1LSB and velocity linearity output of 1% at 25°C. The resolution and the dynamic performance of the AD2S80A can be set by the users to suit their own system requirements.

1. INTERFACING TO HIGH VOLTAGE SYNCHROS AND RESOLVERS

a. Interfacing with External Transformers

Transformers are often used to step down a higher voltage signal to a lower voltage signal. An additional advantage provided by using a transformer in this mode, is that it provides true isolation for the electronics from the resolver and other systems that may be attached to the resolver.

The 5S70 series of signal input transformers are ideal for use with the AD2S80A/AD2S81A converters. They accept all the standard synchro and resolver signal voltages and give the 2V rms signal output required by the monolithic converters.

Similarly the 5S72 transformer enables the reference input of the AD2S80A/AD2S81A to be isolated and accept higher voltages than the standard 2V rms. Figure 1 shows the circuit connections required when the 5S70 and the 5S72 transformers are used with the AD2S80A or AD2S81A.

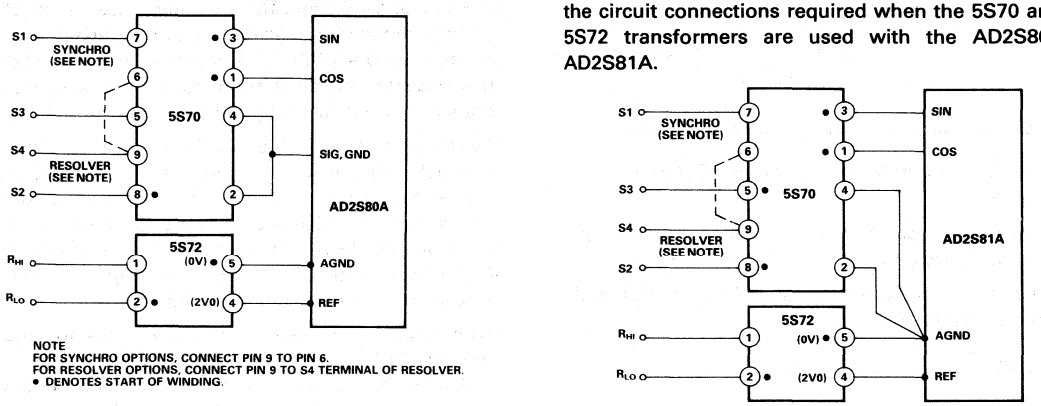

Figure 1. Connecting 5S70 and 5S72 to Either AD2S80A or AD2S81A

*Inductosyn is a registered trademark of Farrand Industries, Inc.

b. Using External Resistors for Scaling Inputs

Sometimes a requirement arises to interface a resolver with nonstandard output voltages to a resolver-to-digital converter. This requirement can either be met by designing a special transformer made to convert the nonstandard outputs of the resolver to the required input voltage on the converter or by using external resistors to scale the input voltages.

The reference and signal inputs to either the AD2S80A or AD2S81A may be scaled using external resistors to accept higher voltages than 2V rms. Let us assume that the signal and reference inputs are at A V rms, where A>2 V rms. Figure 2a shows how two external resistors can be used as
a potential divider to scale the reference input.

The value of R1 and R2 should be chosen such that:

$$\frac{R1}{R2} = \frac{(A-2)}{2}$$

Absolute resistor values are not critical and should be chosen in the kΩ range.

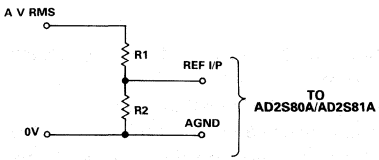

Figure 2a. Scaling Reference Input with External Resistors

It is not important that the absolute value resistors calculated for the reference input is used. The nearest value resistors available to the calculated values can be used.

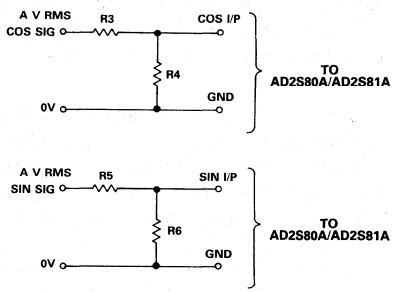

NOTE: GND FOR 2S81A IS "AGND" PIN AND GND FOR 2S80A "SIG GND" PIN

Figure 2b. Scaling Signal Inputs with External Resistors

Figure 2b shows how the signal inputs can be scaled using external resistors.

In the above external resistor network, R3 and R4 should be chosen such that:

$\frac{R3}{R4} = \frac{(A-2)}{2}$ where R3 and R4 are in the kΩ range.

$\frac{R3}{R4} = \frac{R5}{R6}$ It is important that this RATIO IS EQUAL. R5 and R6 should also be in the kΩ range.

The absolute value of the resistors R3, R4, R5 and R6 is not important but the ratio of R3/R4 = R5/R6 should be maintained otherwise the converter accuracy will be affected.

The additional error due to mismatch of the ratio R3/R4 = R5/R6 can be calculated using the formula given below:

$$\frac{\text{Gain Error of Resistors in ppm}}{10} = \text{Error in arc sec}$$

i.e., 1% gain error results in 16.7 arc min.

It is also advisable that the converter outputs be checked for input angle of 45 degrees and the resistor values adjusted accordingly if required.

2. USING THE CONVERTER WITH 15V SUPPLY

The AD2S80A/AD2S81A operates from a ±12V supply. In order to operate it from a ±15V supply requires stepping down from ±15V to ±12V. This may be done using the circuit shown in Figure 3a.

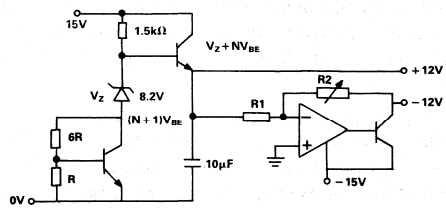

Figure 3a. Circuit for Connecting the Converter to ±15V Supply Requiring Velocity Reversion Error Trimming

The transistors and the Zener diode's combined temperature coefficient compensates for VCO rate temperature coefficient and will ensure a sufficiently accurate 12V supply to the converter over the full operating temperature range.

The −12V supply is derived using the +12V, an amplifier, transistor and resistors. The amplifier is used in negative feedback mode with resistor, R1, and a variable potentiometer, R2. One of the benefits of deriving the −12V in such a way is that it allows trimming of the reversion error on the velocity signal using the potentiometer.

The reversion error, or side-to-side nonlinearity, of the velocity output is a result of differences in the up and down rates of the VCO in the converter. The sensitivity of the VCO rate to −12V depends on the direction of the rotation; therefore, the reversion error can be reduced by varying the magnitude of −12V. By trimming, it is possible to achieve a reversion error of less than 1%.

An alternative to the above circuit which does not require the amplifier is shown in Figure 3b. This circuit may be used if the effect of reversion error on the velocity signal can be tolerated. The reversion error trim can still be included by replacing R_O with a 1kΩ potentiometer, but it will be more sensitive to supply and temperature variations.

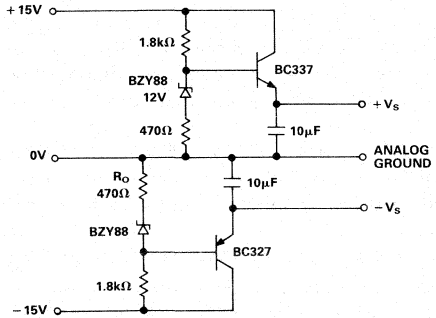

NOTE
IF R_O IS REPLACED BY A 1kΩ POTENTIOMETER, IT CAN BE USED TO TRIM THE REVERSION ERROR OF THE VELOCITY SIGNAL.

Figure 3b. Connecting the Converter to ± 15V Supply

3. USING THE AD2S80A AS A CONTROL TRANSFORMER

The ratio multiplier section of the AD2S80A can be used independently to the rest of the converter to perform the function of Control Transformer. In this mode the signal from the resolver inputs, θ, is compared to a digital angle, ϕ, loaded into the counters. Any difference between these two angles results in an analog voltage at the AC ERROR OUTPUT. To use the device in this manner the DATA LOAD pin is used.

Applying a logic "low" to the DATA LOAD pin will allow data to be loaded into the counters of the converter from the data lines. It is important that the data lines are placed in the high impedance state before pulling the DATA LOAD pin "low". A logic "high" on the enable input maintains the output data pins in the high impedance condition and a logic "low" presents the data in the latches to the output pins.

A switch or an open collector logic gate should be used to apply a logic "low" to the DATA LOAD pin because when it is in the "high" state the pin is internally pulled up to 12V. Figure 4 shows the circuit connections required to operate the AD2S80A as a control transformer.

To operate the AD2S80A as a tracking resolver-to-digital converter the DATA LOAD pin should be left unconnected as it is pulled high internally.

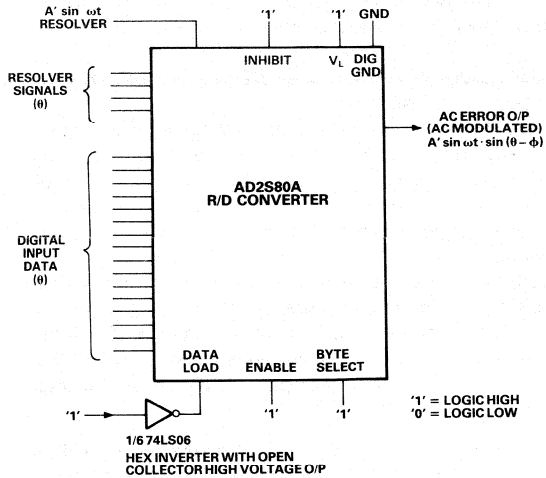

Figure 4. Using the AD2S80A as a Control Transformer

4. USING THE CONVERTERS WITH EXTERNAL PITCH OR REVOLUTION COUNTERS

Applications requiring pitch or revolution counts of the resolver can be easily done using external counters and gates. The input revolutions can be counted by using the RIPPLE CLOCK (RC) and DIRECTION (DIR) logic outputs provided on the AD2S80A/AD2S81A. Figure 5 shows the application circuit that can be used to perform this counting function.

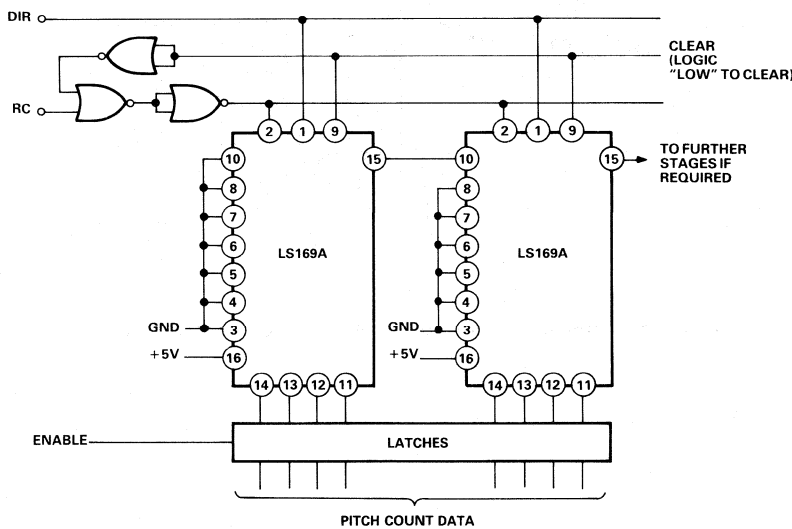

Figure 5. Circuit for Counting Pitches or Revolutions

To transfer the data into output latches on the converter, it is recommended in the data sheet that the INHIBIT input should be used. The data is valid 1μs after the application of a logic "low" to the INHIBIT. Hence the output latches used in Figure 5 should be enabled 1μs after the application of a logic "low" to the INHIBIT input. This will ensure that if an active BUSY pulse is present when the INHIBIT is applied it is given time to clear.

5. SIMULATING AN INCREMENTAL ENCODER USING AD2S80A OR AD2S81A

Incremental shaft encoders usually have two outputs similar to the A and B format, shown in Figure 6. These two outputs are normally used to either increment or decrement an up-down counter and hence show the current shaft angle position. They are often called A quad B outputs.

Where there is a requirement for interfacing to circuitry requiring incremental outputs, the monolithic converters can be easily configured to provide this type of output. Figure 6 shows a circuit which can be used to convert the resolver-to-digital converter parallel digital output into the A quad B outputs.

The "once per revolution" encoder marker pulse is supplied by the converter. An additional advantage of the resolver method is that the "Direction" signal is automatically provided. In the case of the encoder it would require extra circuitry to derive it from the A and B waveforms. N.B. Any of the digital outputs need to be buffered if they are to drive capacitive loads greater than 5pF.

6. CONNECTING INDUCTOSYN TO AS2S80A/AD2S81A

The AD2S80A and AD2S81A R/D converter with a preamplifier can be used as an Inductosyn-to-digital converter.

Inductosyns are similar in operation to resolvers. They consist of two magnetically coupled parts, a scale and a slider. The slider moves along the scale in association with the device to be positioned. There are two types of Inductosyns available: linear and rotary Inductosyns. Inductosyns are not inherently absolute, but it is often sufficient to start from a datum and keep track of pitches with an external counter.

The output signals from an Inductosyn slider are at a low level of the order of millivolts. These signals can be amplified with the hybrid Inductosyn preamplifier, IPA1764, to provide the necessary input signal amplitude required by the monolithic converters. The hybrid power oscillator, OSC1758, can be used to provide the drive for energization of the Inductosyn track.

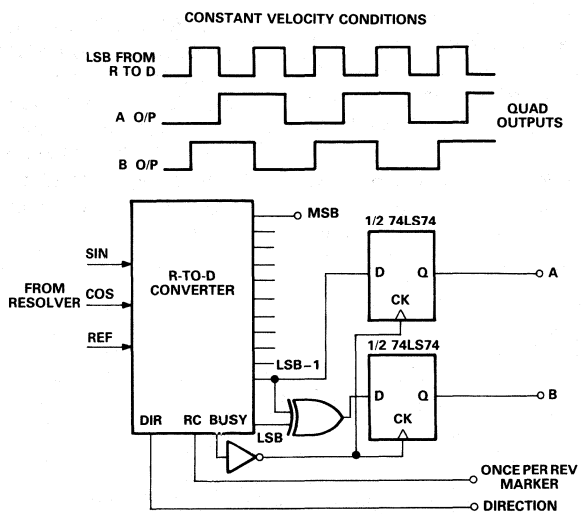

Figure 6. Simulating an Incremental Encoder Using a R/D Converter

7. USING AD2S80A/AD2S81A IN A TWO-SPEED COARSE/FINE SYSTEM

Often it is required to digitize the output of a coarse/fine mechanically geared synchro or resolver system or of an electrically geared or multiple resolver. Both requirements involve identical techniques.

Take for example a two-speed mechanical coarse/fine system with a gear ratio of 32:1. The AD2S80A and AD2S81A can be used in this system where the AD2S80A is used as a fine converter and the AD2S81A is used as the coarse converter. The combined digital output will be 17, 19 or 21 bits, depending on whether the AD2S80A is used as a 12-, 14- or 16-bit converter. The combinational logic required to combine the coarse and fine converter outputs is relatively easy to design. Figure 8 shows a circuit which may be used for this application. The data should be read using software if possible as this will remove the need for sophisticated timing electronics.

For detailed explanation on designing the combinational logic required for two-speed coarse/fine resolver systems, refer to pages 66 to 71 of the "Synchro and Resolver Conversion" book published by Memory Devices (Division of Analog Devices Inc.).

Alternatively, the TSL1612 (available from Analog Devices) which is a two-speed processor for coarse/fine synchro/resolver systems can be used to provide the combinational logic. The standard TSL1612 module is available for two-speed systems with gear ratio of 36:1, 18:1 or 9:1.

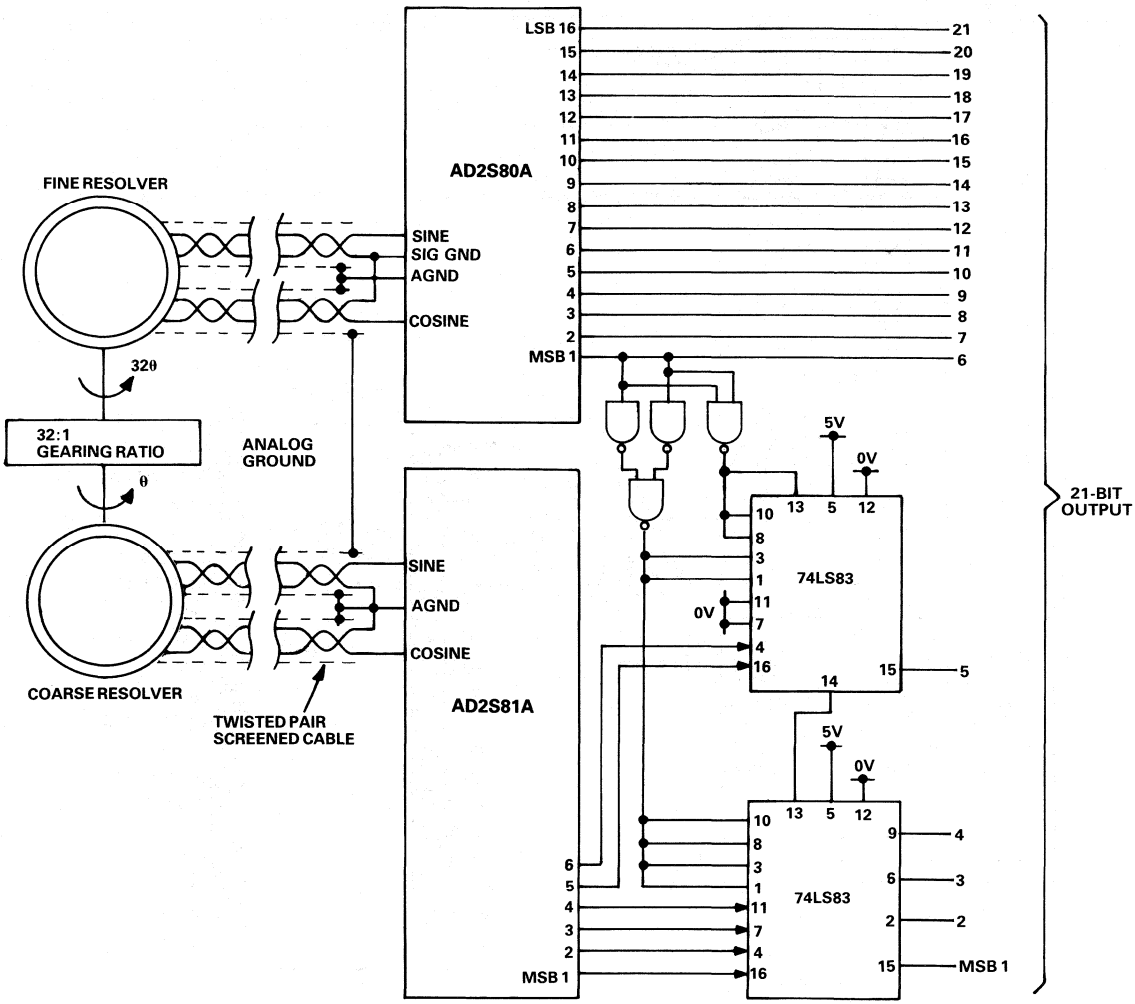

Figure 8. Two-Speed (32:1) System Using the AD2S80A and AD2S81A

Figure 7 shows the circuit configuration necessary to use the monolithic converters as Inductosyn-to-digital converters.

Precaution should be taken on the following points:

- Inductosyns have a 90 degrees phase shift built in between the reference and signal outputs. Using the OSC1758 can overcome this problem as it has two signal outputs, one 90 degrees in phase advance with respect to the other.
- The outputs of Inductosyns vary because of the mechanical tolerances involved with them. Inductosyns outputs must not be allowed to vary by more than ± 10%; otherwise any variation greater than this will be amplified proportionally by the IPA1764, and the inputs to the AD2S80A/AD2S81A will be greater than the maximum specified.
- Particular care should be taken in transferring the output signal from the slider to the amplifier because of the very low amplitude signals involved. Screened twisted pair cables should be used between the outputs of the Inductosyn slider and the amplifier and the outputs of the amplifier and the monolithic converter.

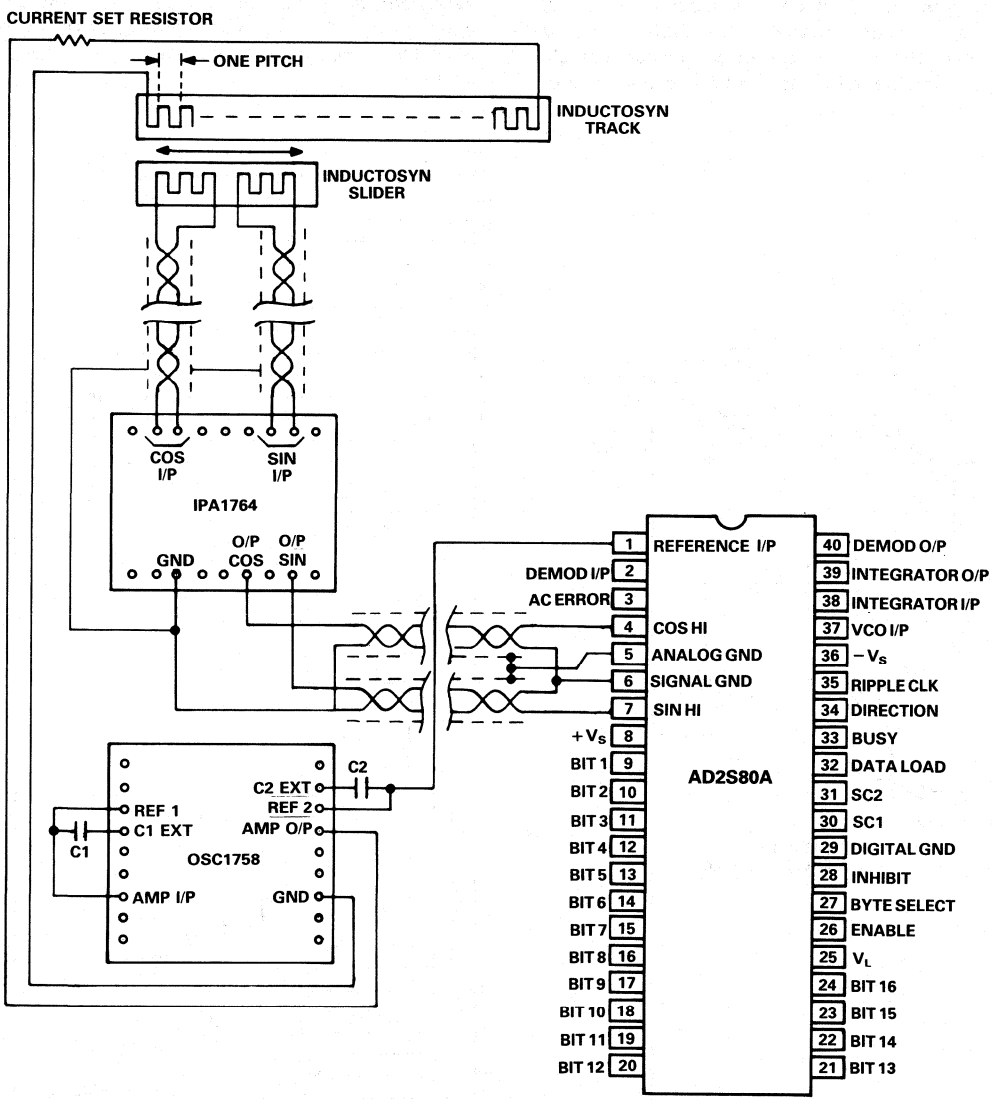

Figure 7. Use of AD2S80A/AD2S81A as an Inductosyn-to-Digital Converter

ANALOG DEVICES

AN-252
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Using the 2S80 Series Resolver-to-Digital Converters with Synchros: Solid-State Scott-T Circuit

by Mark Schirmer

INTRODUCTION

The 2S80, 2S81, and 2S82 are monolithic, tracking converters that are designed to interface to four-wire resolver format signals with nominal 2 V rms amplitudes. These devices can also be used with synchro format signals if an external circuit is employed which accomplishes the transformation from synchro to resolver format. Traditionally, this conversion has been accomplished with a Scott connected transformer, commonly referred to as the Scott-T transformer (Figure 1).

While a well designed Scott-T transformer makes for a very simple galvanically isolated synchro-to-resolver conversion system, the high cost and large size of transformers (particularly when operating at 60 Hz) may negate the benefits of a monolithic resolver-to-digital converter.

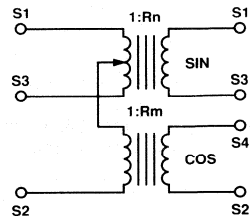

Figure 1. Synchro-to-Resolver Format Scott Connected Transformer

If the design does not require the galvanic isolation intrinsic to the Scott-T transformer, a simple solid-state circuit like that shown in Figure 2 can be used. By using precision resistors and/or trimming the circuit, performance can be achieved comparable with high quality Scott-T transformers, while realizing considerable savings in cost and size.

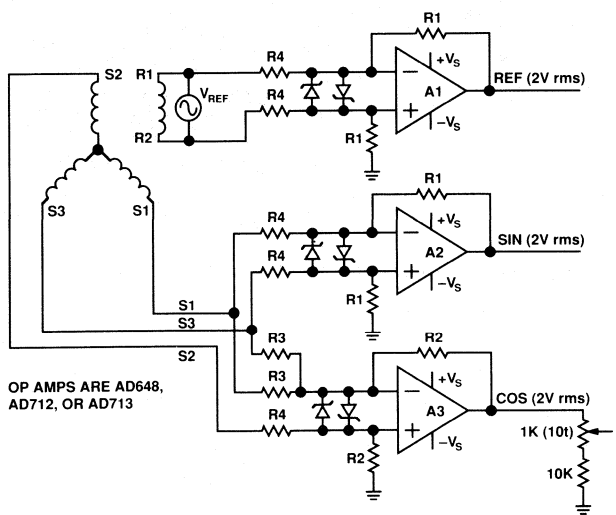

Figure 2. Solid-State Scott-T Circuit

RESOLVERS AND SYNCHROS

The resolver is an electromagnetic, rotational device that detects angular displacement. An equivalent electrical representation and diagram of typical output signal formats for a resolver are shown in Figure 3. An ac excitation signal applied to the primary (rotor) is inductively coupled to the secondary (stator). The transformation ratios are amplitude modulated by the sin and cosine of the angle of the rotor relative to the stator.

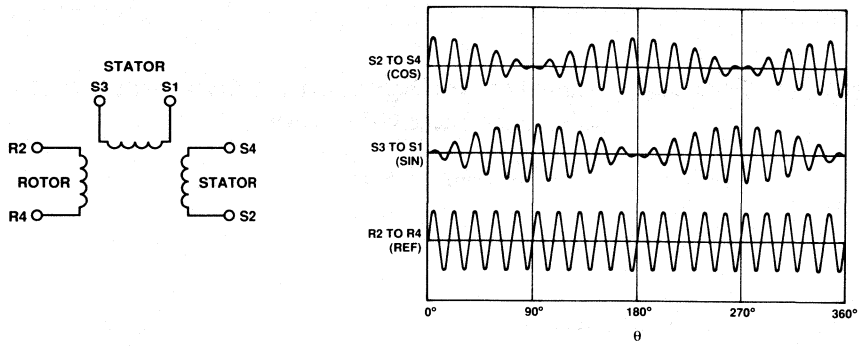

Figure 3. Electrical Representation and Typical Resolver Signals

The operation of the synchro, Figure 4, is very similar to that of the resolver. The fundamental difference is that the stator windings of the synchro are connected in a "Y" configuration, spaced 120 degrees apart, while the resolver has two isolated windings separated by 90 degrees.

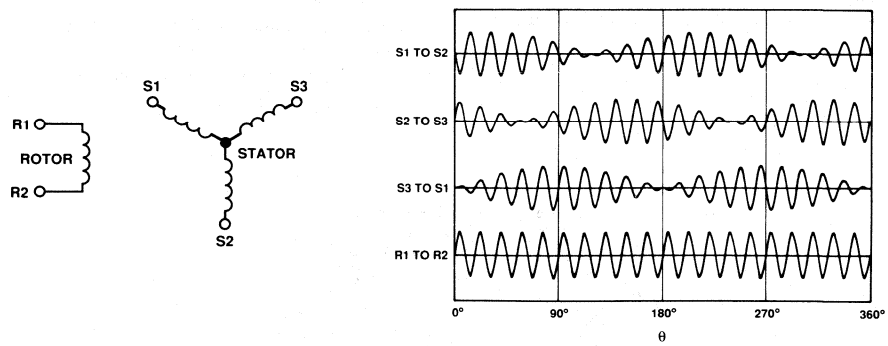

Figure 4. Electrical Representation and Typical Synchro Signals

THEORY OF OPERATION

The solid-state Scott-T circuit illustrated in Figure 2 uses two operational amplifiers to transform a synchro format signal into a resolver format. A third amplifier is used to provide a differential input for the reference signal.

The synchro format input voltages can be written in the form:

$$V_{S3-S1} = KV_{REF} \sin \theta \quad (1)$$

$$V_{S2-S3} = KV_{REF} \sin (\theta + 120°) \quad (2)$$

$$V_{S1-S2} = KV_{REF} \sin (\theta + 240°) \quad (3)$$

where K is the transformation ratio of the transducer and θ is the shaft angle.

In the above notation, V_{Si-Sj} refers to voltage between synchro stator terminals Si and Sj. The order of the indices indicates the polarity/phase of the signal, e.g., V_{S3-S1} represents the voltage at S3 measured with respect to S1. The above assume that the reference voltage applied across the rotor of the synchro is of the form:

$$V_{REF} = V_{R1} - V_{R2} = V_O \sin \omega t \quad (4)$$

The normal convention is for R1 to be taken as the high potential side of the reference excitation, thus making the quantity V_{R1-R2} positive, i.e., for shaft positions between 0° and 180°, V_{S3-S1} will be in time phase with V_{R1-R2}. These conventions are consistent with those of MIL-S-20708.

The objective of the circuit is to convert synchro signals to resolver format signals of the form:

$$V_{S3-S1} = KV_{REF} \sin \theta \quad \text{(SIN)} \quad (5)$$

$$V_{S2-S4} = KV_{REF} \cos \theta \quad \text{(COS)} \quad (6)$$

where for a resolver, the reference voltage is taken as:

$$V_{REF} = V_{R2} - V_{R4} = V_O \sin \omega t \quad (7)$$

Equations 5, 6, and 7 are based on the conventions used in MIL-R-21530 for the case where the reference excitation is applied across the R2–R4 rotor winding and R2 is taken as the high potential side.

MIL-R-21530 is essentially a specification for a brush type resolver. Brushless resolvers, typically with a single rotor winding, may use a different phasing convention. If it is desired to use this solid-state Scott-T circuit to emulate a brushless resolver, the resolver manufacturer's phasing equations should be consulted for consistency with Equations 5, 6, and 7. Different phasing conventions will result in positional offsets in multiples of 90° and/or sense of rotation reversals.

The signal represented by Equation 5 is directly satisfied by one of the synchro signals, Equation 1. Operational amplifier A2, acting as a differential inverting buffer, inverts V_{S1-S3} thus producing the resolver SIN signal. The Zener diodes on the inputs simply limit the applied voltages to the amplifier to guard against component damage. The series input resistors may be scaled to adjust the gain/attenuation of the amplifier.

The second resolver signal, Equation 6, can be shown to be a linear combination of the signals in Equations 2 and 3:

$$V_{S2-S3} - V_{S1-S2} = KV \sin \omega t \, (\sin [\theta + 120°] - \sin [\theta + 240°]) \quad (8)$$

Using the trigonometric identity:

$$\sin (A + B) = \sin A \cos B + \cos A \sin B \quad (9)$$

it is then possible to show that

$$\sin (\theta + 120°) - \sin (\theta + 240°) =$$

$$\sin \theta \cos (120) + \cos \theta \sin (120°)$$

$$-\sin \theta \cos (240°) - \cos \theta \sin (240°)$$

$$= -\frac{1}{2} \sin \theta + \frac{\sqrt{3}}{2} \cos \theta + \frac{1}{2} \sin \theta + \frac{\sqrt{3}}{2} \cos \theta$$

$$= \sqrt{3} \cos \theta \quad (10)$$

Operational amplifier A3 is configured as a differential summing amplifier and accomplishes the computation function:

$$V_{S3-S2} + V_{S1-S2} = -(V_{S2-S3} - V_{S1-S2}) \propto -\cos \theta \quad (11)$$

Since A3 also inverts the input signals, the output of A3 is thus the resolver COS signal.

The feedback resistor, R2, and series input resistors, R3, are chosen such that the gain of the amplifier is $1/\sqrt{3}$ for 2 V rms inputs and scaled proportionately for other signal levels, i.e.,

$$\frac{R2}{R3} = \frac{2}{\sqrt{3}\, V\, rms} \quad (12)$$

Since it is unlikely that the exact value of this resistor ratio can be obtained, it is suggested to use a slightly higher value for R2 and trim the output to the desired level using the potentiometer as shown in Figure 2.

16

RESISTIVE SCALING FOR HIGH VOLTAGE SYNCHROS

Since the 2S80, 2S81, and 2S82 all require nominal 2 V rms input signal amplitudes, it may be necessary to adjust the signal levels from the synchro. This can be easily accomplished by adjusting the values of the input resistors (R3, R4) to the op amps shown in the schematic. Note that because the amplifiers may be operating at less than unity gain, internally compensated amplifiers are recommended in order to reduce the susceptibility to oscillation. The AD712 (dual), AD713 (quad), and AD648 (dual) are suggested.

In addition to galvanic isolation, transformer coupling also significantly improves the common mode rejection on the inputs to the converter. The circuit shown in Figure 2 also enhances the common mode rejection of the system by virtue of the differential amplifier configuration of the circuitry on both the reference and signal inputs. As with any differential amplifier, the common

mode rejection ratio will be enhanced with closely matched resistance on the inverting and noninverting inputs of the operational amplifier. For example 0.01% tolerance resistors will yield a typical 80 dB CMRR with a worst case of 68 dB.

The table below summarizes suggested values of the resistors in the circuit for various standard synchro voltages. The values indicated are standard values for precision (1% or better) resistors. The last entry in the table gives generalized formulae for the resistor values as a function of an arbitrary signal voltage. In addition, please note that the value of R4 may differ in the signal and reference differential amplifier circuits (e.g., an 11.8 V signal synchro is often excited at 26 V rms).

Signal Voltage (V rms)	R1	R2	R3	R4
2.0	11K	12.7K	22.6K	11.3K
11.8	11K	12.7K	133K	66.5K
26	11K	12.7K	280K	140K
90	11K	12.7K	1.18M	590K
115 (REF only)	11K	–	–	620K
V	R	1.155*R	V*R	V*R/2

CIRCUIT ACCURACY

The accuracy of the solid-state Scott-T circuit is determined primarily by the accuracy of the resistors. While absolute values are not critical, the resistors should be matched in pairs on the inputs to the operational amplifiers in order to obtain maximum accuracy. Typically, 0.1% or better tolerance values are best.

Most critical is to maintain the ratios of R1:R2 = 1.1547 and R3:R4 = 2. Note that the number of different resistor values required can be reduced if R4 is constructed by using two resistors of value R3 in parallel (alternatively, R3 can be two R4 value resistors in series). In either case, the use of resistor networks may offer significant cost and size reductions.

If 0.1% tolerance resistors are used, the angular errors introduced by the solid-state Scott-T circuit will range from 2 arc minutes (typical) to 7 arc minutes (maximum). This is additional error which must be added to the accuracy specification of the converter in a worst case analysis. Substantially better accuracy is achievable through the use of matched pairs and/or networks, as the angular errors are proportional to resistor inaccuracies.

If the necessary precision resistors can not be obtained, it is suggested that the gain of the COS circuit in Figure 2 be increased by increasing the value of R2. The output of the circuit may then be attenuated using a potentiometer so that its gain is precisely matched to that of the SIN circuit.

ANALOG DEVICES

AN-266
APPLICATION SOFTWARE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Passive Component Selection and Dynamic Modeling for the 2S80 Series Resolver-to-Digital Converters

by Mark L. Schirmer

The 2S80 series of resolver-to-digital converters are unprecedented in their flexibility. This is a direct result of the use of external passive components to set the dynamic characteristics of the converter. This software package simplifies the computation of the values of these components. In addition, the closed loop transfer function and step responses (positional and velocity) can be computed and graphically displayed. The software also allows for component values to be adjusted and the resulting impact on the converter dynamics modeled.

The software can be used as a design tool for all options of the following devices:*

AD2S80A	Variable Resolution, Monolithic Resolver-to-Digital Converter (DIP)
AD2S81A	Low Cost, Monolithic 12-Bit Resolver-to-Digital Converter
AD2S82A	Variable Resolution, Monolithic Resolver-to-Digital Converter (PLCC)
AD2S83	Variable Resolution, Monolithic Resolver-to-Digital Converter with High Quality Velocity Output

16

The user needs only to specify the reference frequency and the desired digital resolution for the converter. Optionally, values for the tracking bandwidth and maximum tracking rate may be entered, otherwise default values may be accepted. Throughout the software, simple on-line help screens are offered minimizing the need to reference the component data sheet and/or program documentation.

*This software is also applicable to the earlier generation 2S80, 2S81, and 2S82. Please see item 3 under "STARTING THE PROGRAM" on page 2.

SYSTEM REQUIREMENTS

- IBM PC/XT/AT,† PS/2,† or "Compatible" Running MS-DOS‡ Version 2.0 or Higher
- 256K RAM Available Memory
- Graphics Display/Adapter (Monochrome, CGA, MCGA, EGA, VGA, or Hercules)
- 1 Floppy Disk Drive (5.25" 360KB Format)
- 8087/80287/80387 Math Coprocessor Recommended (Not Required)

STARTING THE PROGRAM

1. Boot up the PC with MS-DOS.
2. Insert the supplied floppy disk into a 360KB 5.25" drive.
3. From the DOS prompt, type "n:AD2S80" followed by <Enter> (<Return> on some keyboards) where "n" is the drive address. For example, if the PC was booted from drive C: and the software is installed on drive A:, the user would enter:

   C:> A:AD2S80 followed by <Enter>
4. The opening screen of the program should now be displayed:

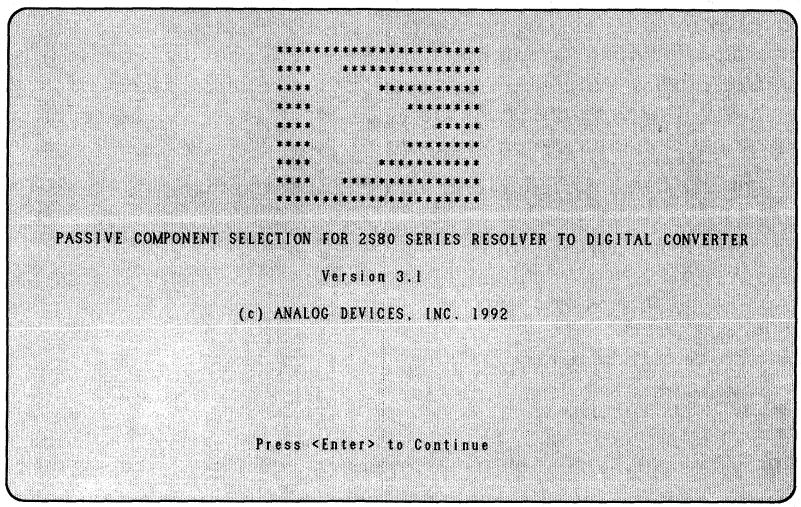

5. Press <Enter> to move to model # selection screen.

If the message "Sorry . . . No Graphics Card Recognized" is displayed, one of the graphics displays/adapters listed in SYSTEM REQUIREMENTS was not found. Your hardware will have to be reconfigured in order to run the software.

RUNNING THE PROGRAM

Press the <Enter> key to clear the opening screen and move to the passive component selection procedure.

NOTE: <Ctrl> + <Break> may be pressed at any time to exit the program.

†IBM PC/XT/AT and PS/2 are trademarks of International Business Machines Corp.
‡MS-DOS is a registered trademark of Microsoft Corporation.

If the software is being run on a machine with either a VGA, EGA, or High Resolution Monochrome (MCGA) graphics adapter, the prompt "Disable Color in Graphics Displays [Y/N/?]?" will be displayed. Entering a "Y" followed by <Enter> will prevent large colored (or gray for the MCGA) areas from being displayed so that useful screen dumps can be made using the PC's <PrintScrn> key. This is necessary because some printers map one or more colors to black giving unpredictable results. Entering a "N" or pressing <Enter> with no entry will give normal color displays. Entering "?" or any character other than "Y" or "N" will give a help screen. Please note that the GRAPHICS statement from DOS must be executed prior to running the software in order to enable graphics screen dumps.

PASSIVE COMPONENT SELECTION

In response to the prompts, the following values should be entered followed by <Enter>:

i. Reference Frequency

ii. Closed Loop Bandwidth

iii. Resolution – (Enter 12 bits for the 2S81)

iv. Maximum Tracking Rate

v. Value of Resistor R1

In the case of parameters i. and iii., pressing <Enter> with no entry will display a context sensitive "Help" screen. For parameters ii., iv. and v., pressing <Enter> with no entry will accept a default value; the "Help" screen can be accessed by entering a "0" followed by <Enter>. In all cases, entry of unacceptable values will generate an informative error message.

The correct values for the passive components given the input parameters specified will be displayed in a tabular format as shown below.

```
    PASSIVE COMPONENT VALUES FOR 2S80 SERIES RESOLVER-TO-DIGITAL CONVERTERS

Reference Frequency:  1000  Hz          Closed Loop Bandwidth:  200 Hz
Maximum Tracking Rate:   62.5 rev/sec  Digital Resolution:  14  bits
Small Signal Step Response Settling Time:  29.2 ms

R1 & C2 FITTED                          R1 & C2 OMITTED

R 1=  24.0   Kohm                         N/A
R 2=  24.0   Kohm                         R 2= 100.0   Kohm
R 3= 100.0   Kohm                         unchanged
R 4=  33.0   Kohm                         R 4=  99.0   Kohm
R 5=  75.0   Kohm                         unchanged
R 6=  56.0   Kohm                         unchanged
R 7=  68.0    ohm                         unchanged

C 1=   6.8   nF                          C 1>   9.1   nF
C 2=   6.8   nF                          N/A
C 3>   9.1   nF                          unchanged
C 4=   9.1   nF                          unchanged
C 5=  43.0   nF                          unchanged
C 6= 470.0   pF                          unchanged


 F1 = Repeat Selection  F2 = Transfer Function  F3 = Step Response  F4 = Quit
```

The values listed are standard 5% tolerance values that will give the closest possible performance to the requirements input by the user. These requirements are also reproduced in the tabular summary. All values are computed using the procedure and designations (R1, C2, etc.) defined in the relevant component data sheet and as detailed in the appendix. Alternative values are given for the mode where R1 and C2 are omitted.

NOTE: The values given for the maximum tracking rate and closed loop bandwidth are the same as those input by the user. These values will be adjusted if the transfer function is computed and found

to exhibit different behavior. These adjusted values will be given when the tabular componen summary is recalled after either the step response or transfer function is computed anc displayed.

Once the component values are displayed, four options, selected by function keys <F1> through <F4>, are available to the user:

<**F1**> **Repeat Passive Component Selection Program** —user is prompted to reenter input parame ters, starting with the reference frequency.

<**F2**> **Compute Transfer Function**—the closed loop transfer function of the converter is computec and plotted for the component values displayed.

NOTE: On PCs with slow clock speeds or without a math coprocessor, this computation may take several seconds.

<**F3**> **Compute Small Signal Step Response**—the normalized response of the digital outputs to a step change in position input is computed and plotted for the component values displayed.

NOTE: On PCs with slow clock speeds or without a math coprocessor, this computation may take several seconds.

<**F4**> **Quit**—terminates program execution and returns the user to the DOS prompt.

TRANSFER FUNCTION

The closed loop transfer function is plotted on a logarithmic grid and consists of two curves, the Gain (in dB) and the Phase (in degrees). The resulting display should look similar to the following:

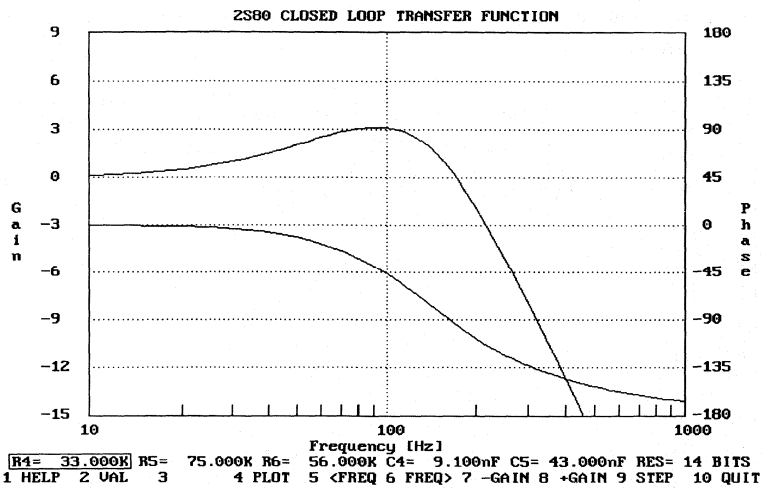

The closed loop transfer function is a measure of the response of the converter to dynamic inputs. The gain curve gives a measure of the amplitude response of the converter digital output to a small signal, sinusoidally varying about some fixed position, applied to the SIN and COS inputs. The phase curve is a measure of the temporal response of the converter, e.g., the phase delay from converter input to output. In both cases, the frequency axis of the transfer function refers to the frequency of the sinusoidally varying input signal, not to either the rotation speed (e.g., RPS) of the resolver or to the reference frequency.

As an example, consider a converter with a 200 Hz closed loop bandwidth as shown in the preceding transfer function plot. If a signal representative of a $\pm 1°$ sinusoidally varying position, oscillating at a 200 Hz rate about a fixed angle, is applied to the converter's SIN and COS inputs, we will observe a

sinusoidally varying output signal with an amplitude of approximately $\pm 0.7°$ ($1/\sqrt{2} \times \pm 1°$), down by 3 dB from the input amplitude. In addition, the output variation will lag the input by approximately 110° or 1.5 milliseconds at 200 Hz.

It should be stressed that the transfer function is valid only for small variations in the input signal; the variations must not drive any component of the loop into saturation. For the 2S80 series of converters, these variations should be less than about 5°.

Modifying Passive Component Values

This portion of the program will allow the user to modify the values of components suggested by the earlier selection portion. The program is controlled by the PC function and cursor control keys. The functions are as follows:

<**F1**> Displays "Help" screen summarizing the operation of the function and cursor keys.

<**F2**> Displays the current component values and operating parameters in a tabular summary identical to the first section of the program.

<**F3**> (Not Used)

<**F4**> Recomputes and displays the transfer function. This key is normally used to update the graphic display after one or more of the passive component values have been altered through use of the cursor control keys. If the present graphic display does not correspond to the displayed passive component values, the message "PLOT INVALID: Press <F4>" will be displayed in the upper left corner of screen.

<**F5**> Adjusts the horizontal scale (frequency) of the transfer function display by shifting 1 decade to the left (i.e., lower frequencies).

<**F6**> Adjusts the horizontal scale (frequency) of the transfer function display by shifting 1 decade to the right (i.e., higher frequencies).

<**F7**> Adjusts the vertical scale of the transfer function gain plot by shifting upper and lower limits down by 3 dB.

<**F8**> Adjusts the vertical scale of the transfer function gain plot by shifting upper and lower limits up by 3 dB.

<**F9**> Displays the small signal step response of the converter for the selected component values. Note that on PCs running at low clock speeds or without math coprocessors this computation may take several seconds.

<**F10**> Terminates program execution and returns the user to the DOS prompt.

< ↑ > Increases the selected component value box through its standard 5% tolerance values.

< ↓ > Decreases the selected component value box through its standard 5% tolerance values.

<← > Moves the box indicating a selected component value (i.e., subject to adjustment with the < ↑ > and < ↓ > keys) to the left.

<→> Moves the box indicating a selected component value (i.e., subject to adjustment with the < ↑ > and < ↓ > keys) to the right.

CAUTION: Using this software, it is possible to manipulate the component values used with the 2S80 series to the extent that converter will not function correctly. Particular care should be taken in adjusting the resolution and the value for R4, the series input resistor to the integrator. These two parameters determine the level of internal hysteresis inside the converter used to stabilize the least significant bit (LSB) for static and low speed applications. Ideally, the two parameters should meet the constraint:

$$2^R \cdot R4 \approx 5.5 \times 10^8 \text{ bit-}\Omega$$

where R4 is in ohms and R is the resolution of the converter (10, 12, 14, or 16 bits). Higher values of

this product will yield less hysteresis possibly leading to flicker of the LSB at low speeds. Lower values will give good stability but compromise the repeatability of the outputs.

In addition, the frequency at which the transfer function crosses −3 dB (the bandwidth) should be compared with the reference frequency. If the ratio of the reference frequency to the bandwidth is less than about 2.5:1, the converter may not be stable and excessive ripple may be observed on the analog velocity signal.

STEP RESPONSE

The step response of the converter is a measure of the response of the output of the converter to instantaneous changes in input position applied to SIN and COS. Such conditions may occur in multiplexed systems or applications with low inertia and high acceleration. The software has provisions for calculating responses of the both the position output and the analog velocity output to steps in input position. NOTE: The response of the analog velocity output of the converter to a velocity step is exactly the same as the response of the position output to a position step.

As is the case for the transfer function computation, the step response models generated by this program are valid only for small signal changes. For the 2S80 series, the largest step that can be applied without saturating any of the loop components is approximately 10 degrees, although the model gives good qualitative predictions for steps up to about 20 degrees. An example plot of the step response is given below:

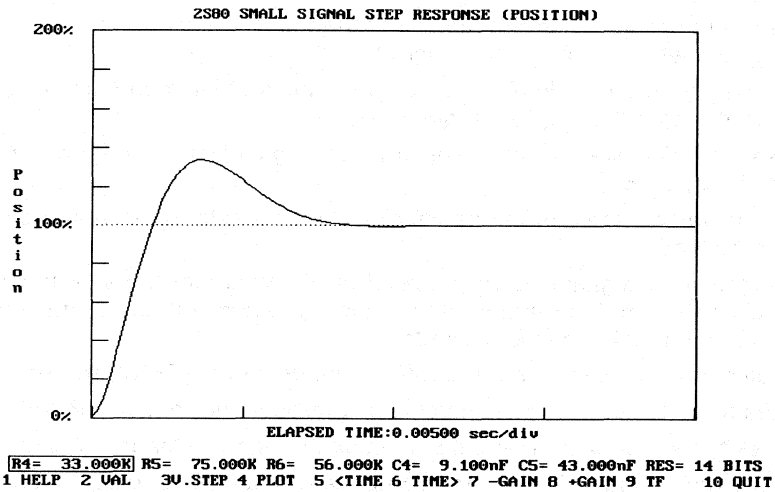

Both the position and velocity step responses are calculated using numerical Fourier transform convolution techniques on the transfer function. Because of both round-off errors in the PC and limitations of the frequency of sampled data in the computations, some artifacts may be introduced in the step response. The user should be aware of the following:

1. If the displayed positional step response does not start at 0% (0 V for the velocity step response) and at zero time, the total time displayed on the horizontal axis of the plot is too short to allow for accurate computations. The user should increase the total time interval by pressing <F6> [TIME>] until the response starts at 0% (0 V). This will rescale the limits internal to the calculation and result in a valid display. An example of an improperly scaled plot follows:

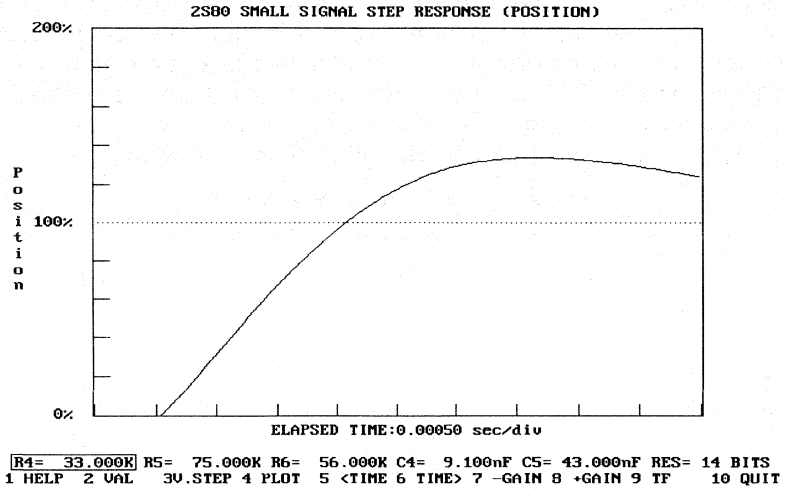

2. If apparent high frequency noise or "ringing" is observed on the displayed step response curve for long time scales, the total time displayed on the horizontal axis of the plot is too long for accurate computations. The user should decrease the total time interval by pressing <F5> [<TIME] until the response is stable over the entire time interval. An example of an improperly scaled plot exhibiting this problem is given below:

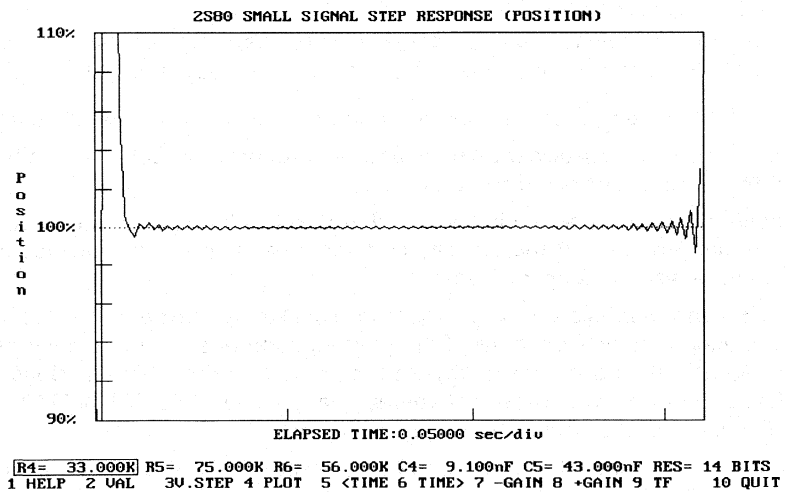

3. The positional step response is a relative calculation, hence the dimensionless vertical scale on the plot. The display shows the output position as a percentage of the magnitude of the input step. For example, if the input position steps from 15° to 20°, the 0% level corresponds to 15° and the 100% level to 20°, the response is normalized for the 5° input step. Please note that the calculation is only valid for steps less than approximately 10 degrees.

4. The velocity step response vertical scale is relative with the units [volts/radian] where 1 radian = 57.3°. In order to estimate the velocity response of the converter to an *x* radian position step, the readings should be multiplied by *x* to yield velocity values in volts. Please note that the actual velocity signal on the 2S80 series is constrained to less than approximately ±9 volts. As a result, if the model predicts a voltage peak greater than 9 volts for an *x* radian step, the step would in practice actually drive the converter into saturation, thus invalidating the calculation.

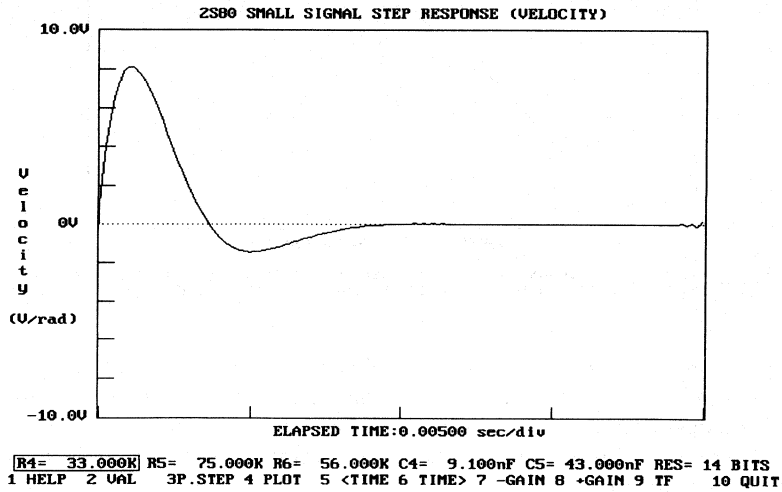

Modifying Passive Component Values

As in the case for the transfer function computation, the step response calculation is controlled by the PC's function keys. Their function closely parallels the transfer function section:

<**F1**> Displays "Help" screen summarizing the operation of the function and cursor keys.

<**F2**> Displays the current component values and operating parameters in a tabular summary identical to the first section of the program.

<**F3**> If the positional step response is displayed, this key will compute and display the velocity step response. Similarly, if the velocity step response is presently displayed, <F3> will generate the positional step response.

<**F4**> Recomputes and displays the step response. This key is normally used to update the graphic display after one or more of the passive component values have been altered through use of the cursor control keys. If the present graphic display does not correspond to the displayed passive component values, the message "PLOT INVALID: Press <F4>" will be displayed in the upper left corner of screen.

<**F5**> Adjusts the horizontal scale (time) of the step response display by reducing the total time displayed in a 5-2-1 sequence.

<**F6**> Adjusts the horizontal scale (time) of the step response display by increasing the total time displayed in a 1-2-5 sequence.

<**F7**> Adjusts the vertical scale of the step response plot:

Position: Increases vertical range by 10× symmetrically about 100% to a maximum of 0–200%.

Velocity: Increase vertical range by 10× to a maximum of ±100 V/radian.

<F8> Adjusts the vertical scale of the step response plot:

Position: Decreases vertical range by 10× symmetrically about 100% to a minimum of 99.1–100.1%.

Velocity: Increase vertical range by 10× to a maximum of ±0.1 V/radian.

<F9> Displays the closed loop transfer function of the converter for the selected component values. Note that on PCs running at low clock speeds or without math coprocessors this computation may take several seconds.

<F10> Terminates program execution and returns the user to the DOS prompt.

<↑> Increases the selected component value box through its standard 5% tolerance values.

<↓> Decreases the selected component value box through its standard 5% tolerance values.

<←> Moves the box indicating a selected component value (i.e., subject to adjustment with the <↑> and <↓> keys) to the left.

<→> Moves the box indicating a selected component value (i.e., subject to adjustment with the <↑> and <↓> keys) to the right.

APPENDIX

The following gives the relevant equations on which the calculations in this software are based.

Passive Component Selection

The program computes the values for R1 through R7 and C1 through C6 using the following equations. The closest 5% tolerance value to that suggested by the calculation is displayed. More detail on these relationships can be found in the relevant component data sheet. Note that all resistance values are specified in ohms and all capacitances are in farads.

$$R1 = R2 \leq 56\ k\Omega$$

$$C1 = C2 = \frac{1}{2\pi R1 f_{REF}}\ F$$

where f_{REF} = Reference Frequency [Hz].

$$R3 = 100\ k\Omega$$

$$C3 \geq \frac{1}{10^5 \times f_{REF}}\ F$$

The value of R4 depends on the configuration of the HF filter between the AC ERROR output and the DEMOD input. This filter may be either a bandpass configuration (R1, C2 fitted) or a simple high pass filter (R1, C2 omitted).

If R1 and C2 are fitted, then:

$$R4 = \frac{5.46 \times 10^8}{2^R}\ \Omega$$

otherwise with R1 and C2 omitted then:

$$R4 = \frac{1.64 \times 10^9}{2^R}\ \Omega$$

where R is the converter resolution (10, 12, 14, or 16 bits).

The tracking rate range of the converter is set by the resistor in series with the VCO input driven by the integrator output. This resistor is a function of the resolution of the converter as follows:

$$R6 = \frac{8 \times F_{VCO}}{T \times 2^R}\ \Omega$$

where T is the maximum tracking rate in revolutions/second [RPS], and F_{VCO} is the VCO rate in Hz/A.

The bandwidth selection components around the integrator are chosen to yield an optimally damped system ($\zeta = 0.707$). Note that indirectly (by virtue of their dependence on R6) the values of C4, C5, and R5 are function of the resolution and maximum tracking rate:

$$C4 = \frac{20.2}{R6 \times f_{BW}^2}\ F$$

$$C5 = 5 \times C4$$

$$R5 = \frac{4}{2 \times \pi \times f_{BW} \times C5}\ \Omega$$

For AD2S83, substitute

$$C4 = \frac{21}{R6 \times f_{BW}^2}$$

where f_{BW} is the closed loop bandwidth [Hz].

The phase compensation components for the VCO input are fixed with the following values:

$$C6 = 470\ pF$$
$$R7 = 68\ \Omega$$

For AD2S83, substitute

$$C6 = 390\ pF$$
$$C7 = 150\ pF$$
$$R7 = 68\ \Omega$$

The small signal step response settling time is approximated by the equation:

$$t_2 = \frac{5}{f_{BW}} \times \frac{R}{12}$$

Transfer Function

The closed loop positional transfer function of the 2S80 series is given by the equations:

$$f(s) = \frac{\theta_{OUT}}{\theta_{IN}}(s) = \frac{K_A m \omega_2 (1 + s/\omega_2)}{s^3 + m\omega_2 s^2 + K_A m s + K_A m \omega_2}$$

where

$$\omega_2 = \frac{1}{R5 \times C5}$$

$$m = \frac{C4 + C5}{C4}$$

$$s = j\omega$$

$$K_A = \frac{54.6 \times F_{VCO}}{2^R \times R4 \times R6 \times (C4 + C5)}\ sec^{-2}$$

Note that in the previous equations, the transfer function is specified in terms of angular frequency [radians/second] rather than rotational frequency [Hz].

The passive component selection procedure for the 2S80 series is structured so that the phase margin of the loop is maximized and the damping is near optimal. In this instance the following relations are true:

$$K_A = \sqrt{m}\,\omega_2^2$$

$$\zeta = \frac{\sqrt{m}-1}{2}$$

$$\omega_{BW} \approx 4\,\omega_2$$

where K_A is the acceleration constant, ζ is the loop damping factor, and ω_{BW} is the loop bandwidth. It should be observed that the above relationships may not remain valid if component values are substantially altered from those generated by the passive component selection algorithm.

Step Response

The step response of the 2S80 series is given by the inverse Laplace transform of the closed loop transfer function convolved with the Laplace transform of a unit step.

The Laplace transform and its inverse are defined as:

$$\pounds\{F(t)\} = f(s) = \int_0^\infty e^{-st} F(t)dt$$

$$\pounds^{-1}\{f(s)\} = F(t) = \frac{1}{2\pi j}\int_{a-j\infty}^{a+j\infty} e^{st} f(s)ds$$

The transform of the unit step is given by:

$$\pounds\{\mu(t)\} = \frac{1}{s}$$

This gives the following expression for the positional step response of the 2S80 series:

$$P(t) = \pounds^{-1}\left\{f(s)\cdot\frac{1}{s}\right\} = \int_{a-j\infty}^{a+j\infty} \frac{K_A m\omega_2 (1 + s/\omega_2)}{s^4 + m\omega_2 s^3 + K_A m s^2 + K_A m\omega_2 s}\,ds$$

This expression gives the response of the position output to a step change in input position. It should also be noted that the same relationship will hold for the response of the analog velocity signal to a step change in input velocity.

The velocity step response (velocity response to position changes) is given by a similar transformation, but using the velocity transfer function for the 2S80 series. The velocity transfer function can be derived from the positional transfer function by multiplying by the Laplace variable *s* and dividing by the VCO gain, K_{VCO}.

$$K_{VCO} = \frac{2 \times \pi \times F_{VCO}}{R6 \times 2^{R}} \left[\frac{\text{rad}}{\text{V} \cdot \text{sec}}\right]$$

$$V(t) = \mathcal{L}^{-1}\left\{\frac{s \cdot f(s)}{K_{VCO}} \cdot \frac{1}{s}\right\} = \int_{a-j\infty}^{a+j\infty} \frac{K_A m \omega_2 (1 + s/\omega_2)}{K_{VCO}(s^3 + m\omega_2 s^2 + K_A m s + K_A m \omega_2)} ds$$

Note that the factor of *s* introduced by the velocity transfer function is cancelled by the $1/s$ term for the step function transform. This expression gives the response of the analog velocity signal (in volts) to a step change in input position.

For computational efficiency and general case applicability, rather than perform the integrations detailed above, the program actually uses Fast Fourier Transform (FFT) techniques to perform the calculations. The results, however, are the same as those that would be generated using the equations outlined above.

ANALOG DEVICES

AN-267
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Vector Control Using a Single Vector Rotation Semiconductor for Induction and Permanent Magnet Motors

by F. P. Flett

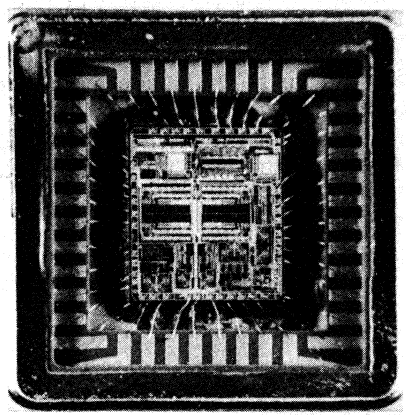

ABSTRACT

This paper describes the implementation and architecture of an integrated circuit which is dedicated to the role of vector rotation for AC induction and DC-Permanent Magnet synchronous, PMSM, brushless motors.

Sampling time and resolution are normally the factors which limit the performance of software control schemes in motor control. Now with the ever increasing tasks required by digital processing in motor control the gains in performance are often limited by the additional software computation time taken to achieve them.

At the other end of the spectrum many control schemes only require one, or, two vector transformations. In these cases the choice of a hardware solution that avoids the pitfalls of software intensity which either limits the bandwidth of control, or even worse, number truncation which reduces the signal-to-noise ratio, can be avoided.

A new semiconductor - AD2S100 - that has been developed and built in concert with an AC induction vector controller now removes the time intensive task of calculating vector coordinate transformations in software and realizes an extremely economic and efficient approach to vector control.

The AC Induction Motor

It is quoted that the implementation of variable speed drives could save between 22 and 27 percent of the total electricity used in the USA. This refers in particular to the use of AC induction motors with variable speed control.

Induction motors are often determined as asynchronous device when operating either under constant, or, adjustable frequency control. For control purposes however the induction motor can be thought of as a synchronous brushless motor where the rotor magnets rotate in the air gap between the stator winding and the squirrel cage rotor.

Additionally the induction motor likens itself to the DC brushless motor in that the rotor flux interacts with current in the stator to produce torque. Torque is maximized when the stator current is perpendicular, or in quadrature, with the rotor flux.

A fixed relationship exists between the rotor magnets and flux in a permanent DC permanent magnet motor. Thus, by applying stator current in time quadrature with the flux optimum torque is achieved.

The same relationship does not hold for the induction motor, however. There are two components from the modulus of stator current which must be generated and managed in this motor. The first realizes the in-phase flux component while the second concludes the quadrature torque component. Both currents hold no fixed natural phase relationship to the shaft position.

It is the rotating magnetic field generated by the electromagnets in the induction motor which constitutes the major control differentiation between the induction and DC permanent magnet motor.

The diagrams below aid this explanation:

DC PERMANENT MAGNET MOTOR

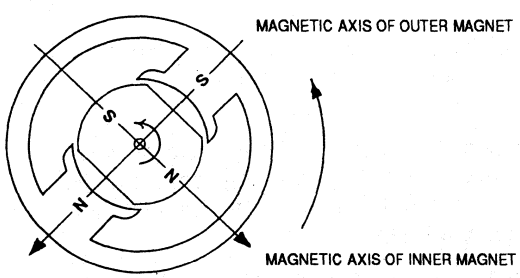

Figure 1.

AC INDUCTION MOTOR

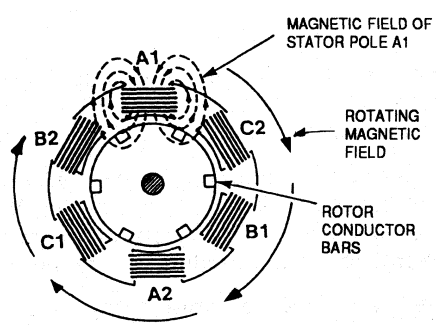

Figure 2.

Figure 1 identifies a DC motor where the armature winding is normally on the rotor and the field winding is on the stator. Permanent magnets mounted in the rotor interact with the field winding on the stator to produce maximum torque when both sets of magnets, those fixed on the rotor and the rotating magnetic field, are at ninety degrees to one another.

In contrast the armature winding of an AC induction motor is associated with the stator. The three phases A, B, and C, on the two pole motor, identified in figure 2, create a rotating magnetic field which follows the armature field as it rotates around the motor. This rotating magnetic field is a product of the electromagnets created by the stator induced field current in the rotor bars of the squirrel cage rotor.

AC induction motor control must therefore manage both the flux and torque current simultaneously while in the case of the DC permanent magnet motor these can be controlled independently. This is where the necessity for vector control arise.

Applied Vector Control

Three variables are attributed to the torque component of induction motors. These are stator current, slip frequency and flux. Vector control of induction motors calls for the independent control of two of these components, namely, current and slip.

At the simplest level of control a configuration which identifies the application of a vector control block is shown below. More complicated configurations will be discussed later.

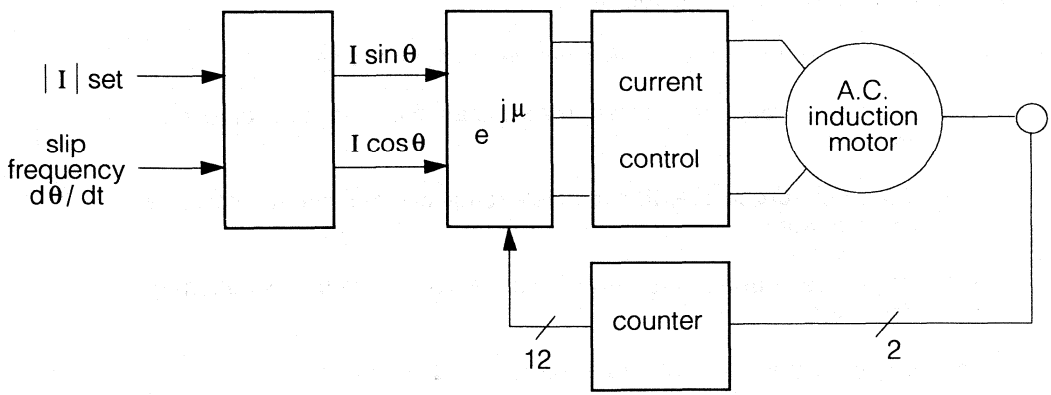

Constant slip frequency control using a vector rotation block

Figure 3.

Current and slip control was pioneered by Leonhard in 1973.

A control strategy such as the one depicted above results in good torque control over a wide speed range and high efficiency of AC induction motors when a constant slip frequency is applied [Jian, 1983].

The optimal slip frequency is dependent on the rotor resistance and saturation which varies with temperature.

If it is essential to maintain a constant torque at all times the rotor slip should increase with the rotor resistance and saturation [Bossche, 1991].

A position feedback device is also required and this can be either a resolver, or, an incremental encoder. Normally an incremental encoder featuring a line count of 1024 is adequate for the majority of controllers. In the case of a DC permanent magnet motor absolute position feedback is necessary and a resolver is the natural choice, here.

This form of control may be applied in applications where a slow dynamic response is adequate, pumps, fans and compressor for example, and where optimal motor efficiency is important. This is also the most opportunistic area for electricity cost reduction highlighted earlier.

Alternatively dynamic control applications demand the rotor be kept magnetized in order to react to sudden changes in load inertia by simultaneously exerting the corresponding torque current requirements. Normally for this purpose the high dynamic velocity and position loops contain an inner control loop. This inner loop is a major contributor to the overall performance of the system. Gain changes above 30% seriously affects the speed and position performance of dynamic systems and is one of the fundamental reasons for the adoption of vector control algorithms.

The key merits of vector control have often been stipulated as

(1) High motor efficiency due to magnetization current control related to RPM.

(2) Constant torque over the complete range of motor speed, inclusive of zero speed.

(3) An increase in the speed/torque envelope due to field weakening schemes.

(4) Good velocity control under varying load conditions.

Drawbacks to vector control schemes normally quote both economics and sampling performance as deterrents to adoption.

There is of course a direct relationship between these electives. The sampling interval is the most critical element in the application of digital signal processing systems and one

that has a fundamental bearing on the cost/performance of the controller. To preserve the information in the analog signal it must be sampled at a rate that is at least twice the rate of the highest frequency component.

In digital control systems the sampling rate is usually chosen to be between 8 to 10 times the system bandwidth.

Common sampling rates start at 1kHz and can reach up to 100kHz. However the hardware used in many motion controllers for factory automation limit the sampling frequency to around 2kHz.

This has dramatic effects on system stability and slows down the response time considerably.

The complexity of the vector calculations normally carried out by software is one overriding reason why the sampling time is reduced. This is not only due to the actual matrix multiplication, which normally carries out the Cartesian to polar coordinate conversion, but also to the "round off noise" caused by intermediate truncation when inadequate computing power is applied. A higher order DSP solution to solve these limitations can be implemented - but at a cost which many would be adopters have failed to justify to themselves.

With the advance of DSP technology many of these drawbacks have been overcome. However, there is now more demands being required from the processor to achieve superior controller performance.

Adaptive control to changes in rotor resistance and real time diagnostics with predictive maintenance routines that calculate fail-safe control strategies are just two tasks which many drive designers would like to see adopted as standard. These tasks sound futuristic but will become a reality as specific motion control computing devices become available.

A dedicated vector control IC - the AD2S100 - is now available to perform the time intensive complex vector rotation calculations. Control schemes that use this semiconductor will be described and the performance and economic benefits analyzed

Vector Math

As previously quoted, the main objective of vector control is to determine the torque and flux constituents from the applied magnitude of stator current and relate these currents to the moving rotor reference frame. In many articles denoting these currents the nomenclature used is "Iqs" for torque current and "Ids" for the flux current. These conventions will be carried through here.

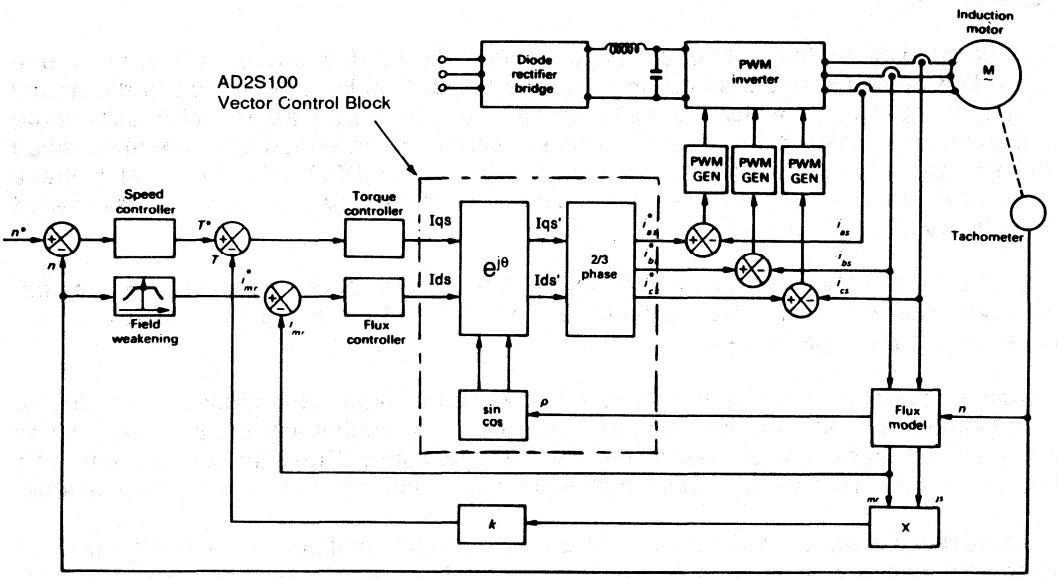

Induction motor speed control by direct field orientation with a current-controlled PWM inverter.

Figure 4.

Figure 4 relates a complete field oriented scheme used in adjustable-frequency induction motor drives. The vector block identified will be the topic of concern in this section. The scheme shown here is commonly known as indirect field control due to the fact that the rotor flux position 'ρ' is deduced from a shaft position transducer and the slip frequency of the rotor.

The rotor flux vector has an instantaneous value which is the sum of the shaft angular velocity, ωm, and the corresponding slip angular velocity, ωs.

If we are trying to determine the actual position of the rotor flux vector ωr, then the following holds:

Instantaneous value of rotor flux vector = ωr
Instantaneous value of shaft angular velocity = ωm
Instantaneous value of slip angular velocity = ωs

$$\omega r = \omega m + \omega s \quad \text{---------- 1}$$

also $\omega r = \omega m + \dfrac{Iqs}{\tau r\ imr}$ **where** $\omega s = \dfrac{Iqs}{\tau r\ imr}$

$$\omega s = \omega r - \omega m \quad \text{---------- 2}$$

and is also equal to sωr

$$\therefore\ \omega s = s\omega r \quad \text{---------- 3}$$

Where s is the frictional slip of the rotor with respect to rotor flux vector ωr.

The parameter ωr is normally represented by a digital number and is one of the input ports on the AD2S100 vector controller.

The kernel of the coordinate transformation is the vector operator:
ejθ or e-jθ, where ejθ = [cosθ + jsinθ].

It is this operator which relates the stator currents to the moving rotor reference frame.

We can describe the flux current Ids and torque current Iqs in complex numbers as:

$$Ids + jIqs$$

As previously described these current are the in-phase and quadrature currents respectively.

If we now want to relate these currents to the rotor frame reference the vector operator is used as follows:

$$Ids' + Iqs' = [\ Ids + jIqs\ ]\ [ej\theta]$$

if we expand this we get:

$$[\ Ids + jIqs\ ]\ [Cos\theta + jSin\theta]$$

which ends up as $Ids' = IdsCos\theta + IqsSin\theta$
and $Iqs' = -IdsSin\theta + IqsCos\theta$

The outputs Ids' and Iqs' are shown at the output of the vector control block. The 2/3 phase block is simply the conversion from sine and cosine signals to a three phase 120 degree representation.

The matrix representation that is used by digital processing is

$$\begin{bmatrix} Ids' \\ Iqs' \end{bmatrix} = \begin{bmatrix} Cos\theta & -Sin\theta \\ Sin\theta & Cos\theta \end{bmatrix} \begin{bmatrix} Ids \\ Iqs \end{bmatrix}$$

There are two other points of note in figure 4 which relate to field weakening and the multiplication of the rotor magnetizing current, imr, and iqs at the output of the flux model. Field weakening is a technique where the base speed of the induction motor can be extended at a constant torque by reducing the rotor magnetization current. It also lowers the problem of pulsating torque at low speeds but only at the expense of fundamental torque, however.

The multiplication of imr and iqs is quite simply to determine the actual motor torque T and relate this to the input speed controller for compensation.

Systems on Silicon in Motion Control - The AD2S100 & 2S110

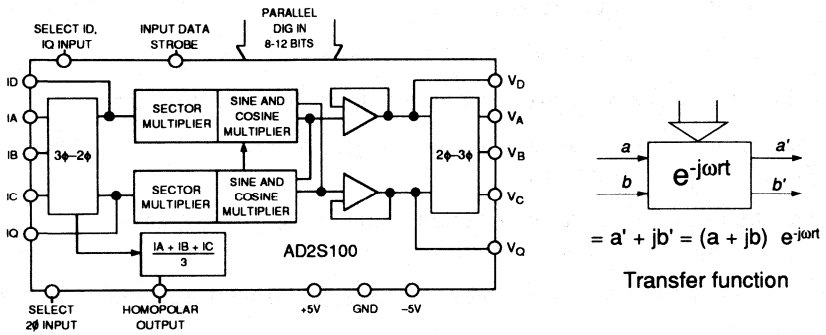

AD2S100

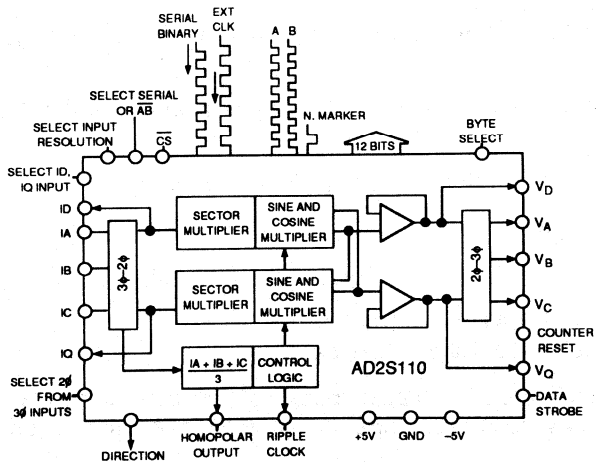

AD2S110

Figure 5.

The block diagrams above represent the key elements of the AD2S100 and AD2S110 vector controllers.

The prime difference between the two components relates to the position transducer selected. An absolute position device, a resolver, is used with the AD2S100 while the AD2S110 accepts both an incremental encoder A quad B input and serial absolute position input. This serial absolute position input is also generated by a resolver and can be converted by a low cost resolver to digital convertor such as the AD2S90.

General Product Description

These three-port device provide a digital input port that accommodates the most popular sensors used for rotor position and velocity measurement.

The parallel binary input on the AD2S100 accepts a twelve bit digital word which is derived from a resolver via a resolver to digital converter, (R/D). Another input form accepts incremental A quad B signals from an optical encoder, (AD2S110).

In the case of resolver based systems a natural by-product of the type-two tracking loop conversion used in monolithic R/D converters is the availability of a velocity signal which can be used to replace a tachogenerator.

An analog input port which senses three phase stator currents, normally commuted by hall effect devices, is shown at the left side of the vector controllers above. These real time signals may either be three phase 120 degree separated signal, or, quadrature sine and cosine signals.

These signals undergo the coordinate transformation according to the vector equations previously disclosed and are then presented as analog output in either three phase, or, quadrature representation. A complete vector transformation is completed in two microseconds.

The partitioning of the input and output sections allows these blocks to be used either for forward, or, reverse vector transformations. The implementation is shown below.

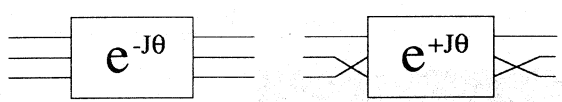

3 Phase in - 3 Phase out

Figure 6.

Additionally decode logic has been incorporated in the vector control chips to allow a user to select only two of three 120 degree phases at the input. This assumes that the sum of the three phases is zero and that the third phase can be deduced from the knowledge of the other two.

In this respect the economy of using only two hall effect devices is made at the expense of the homopoiar output. This output senses if there is an imbalance between any of the three phase inputs and therefore require that all three inputs are used in a three phase system. This output can be very useful in the detection, and prediction, of a dielectric breakdown resulting in an earth leakage condition.

AD2S110

The AD2S110 hosts differences and additional functions which reflect the use of an incremental encoder as the rotor position senor.

Inputs can in fact be either absolute serial binary or incremental encoder format. The parallel digital port is in fact an output in this case. Here the instantaneous incremental position input is represented as a parallel digital number which can be sent in that format to a peripheral device. Two other signals, direction sense, and a ripple clock which senses a once per revolution marker are also available as output logic states.

A Beta Site Evaluation

Vector Rotation Setup

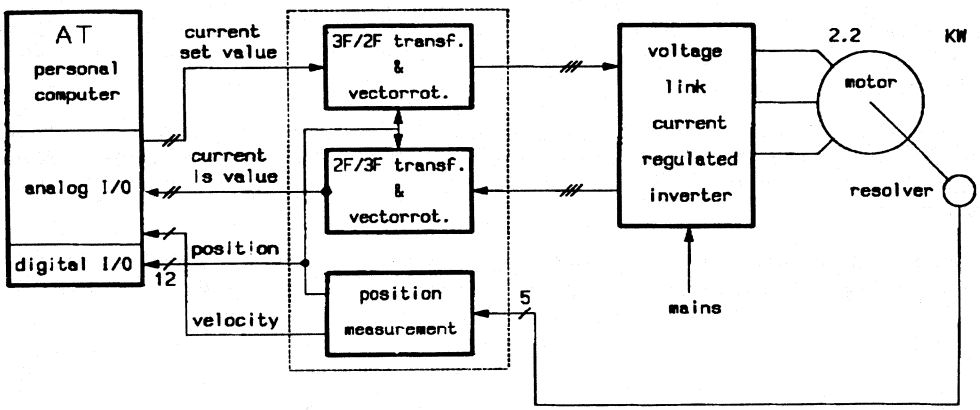

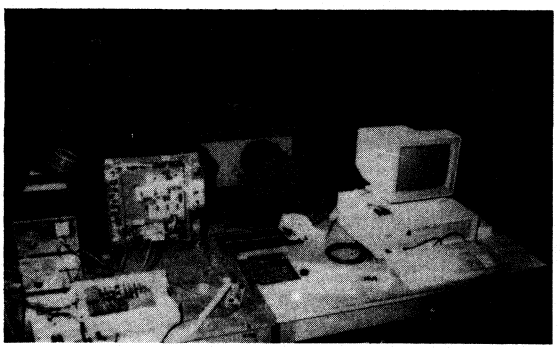

Figure 7.

The test system used in the evaluation of the vector control semiconductor is shown above. Two software routines were implemented. The first refers to a conventional slip regulation algorithm which was used for manual control of the motor. An additional on-line program was implemented to relate the performance of the vector control semiconductors when steering the control of the motor.

The benchmark performance of these semiconductor in the motor control system shown will be referred in the presentation of this paper.

References

[Leonhard, 1973] W. Leonhard, "Regelung in der elektrischen Antriebtechnik", Elektrotechnik, Teubner Studienbucher, 1973.

[Jian, 1983] T.W. Jian et al, "Characteristic induction motor slip values for variable voltage part load performance optimization", IEEE Transactions on Power Apparatus and Systems, Vol. Pas-102, No. 1, Jan. 1983, pp. 36 - 38.

[Bossche, 1991] A.P.M. Van den Bossche, P.J. Coussens "Single chip vector rotation blocks and induction motor field control", PCIM Conference, Nürnberg, Germany 1991.

ANALOG DEVICES

AN-230
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Circuit Applications of the AD2S90 Resolver-to-Digital Converter

By Dennis Shi-hsiung Fu

The AD2S90 is a 12-bit resolution monolithic resolver-to-digital converter. It provides a complete solution for digitizing resolver signals without using external components. The advantage of using the AD2S90 over other traditional R/D converters in the market is that this chip provides not only absolute position data output (in serial form), but also standard A Quad B format signals. When upgrading the incremental encoder system to a resolver based system, this chip plus a resolver will give out a complete replacement solution without revising the remainder of the system. The absolute position data output gives additional useful information in the upgraded system. The serial form of the digital position data is especially suitable for long distance data transmission.

This application note does not discuss the general features or the basic operations of the AD2S90; that information can be found in the data sheet. The following applications are discussed in this application note:

Absolute position data serial-to-parallel conversion
Absolute position data EIA-485 communication
4096 line incremental encoder emulation

ABSOLUTE POSITION DATA SERIAL-TO-PARALLEL CONVERSION

There are many different ways to perform AD2S90 absolute position data serial-to-parallel conversion. Proper digital timing can be generated using microprocessor or hardware circuits. The Figure 1 block diagram demonstrates the basic AD2S90 serial-to-parallel data conversion in a hardware mode. In this scheme we use two 74HC595s as shift registers. The timing is generated by using a counter/divider 74LS169, 74121 monostable and a delay component. Figure 2 shows the implementation circuit in schematic form.

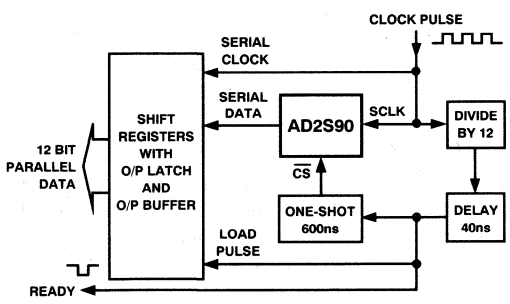

Figure 1. Serial-to-Parallel Conversion Block Diagram

Considering the timing diagram of the device as shown in the data sheet, the clock pulse waveform (SCLK) should have the following relationships with the chip select pulse ($\overline{CS}$) and load pulse:

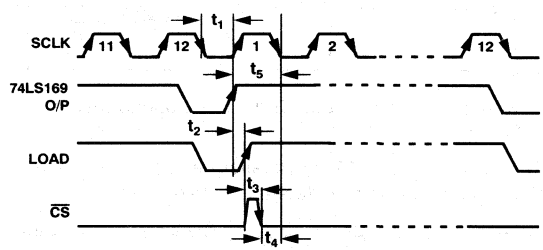

Here, t_1 = 250 ns minimum
t_2 = 40 ns (delay)
t_3 = 600 ns minimum
t_4 = 600 ns minimum

The value for t_5 can be easily calculated as:

$t_5 = t_2 + t_3 + t_4 = 1240$ ns minimum.

In this application circuit, there is a maximum possible delay of 40 ns between the upgoing edge of SCLK and the upgoing edge of the 74LS169 output signal pulse. This has to be taken into consideration when calculating the timing.

The one-shot used in the Figure 2 circuit (74121) has the pulse width:

$$t_W = \ln 2 \cdot C \cdot R$$

with C = 10 pF and R = 86.6 kΩ, t_W = 600 ns

So the t_5 value in this circuit will be:

$$t_5 = t_2 + t_3 + t_4 + 40 = 1280 \text{ ns minimum}$$

16

The calculation suggests the SCLK frequency range for this demonstrating circuitry, i.e.,

$$For\ 50\%\ duty\ cycle,\ SCLK \le \frac{1}{2 \times t_5} = 390\ kHz$$

$$For\ 84\%\ duty\ cycle,\ SCLK \le \frac{1}{t_5 + 250\ ns} = 653\ kHz$$

Notice here t_2 is the time period needed for 74HC595 to finish data transfer (from shift registers to the storage register). The load pulse can be used independently or gated with other signals to indicate the READY status of the data conversion. For further information on detailed timing, please refer to corresponding data sheets.

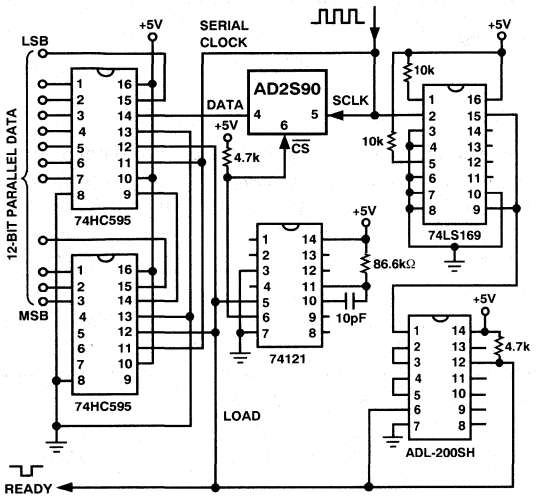

Figure 2. Synchronous Serial-to-Parallel Conversion Schematic

The read-out rate (12-bit parallel data) for this application circuit can be calculated as follows:

Assume that the SCLK = 653 kHz, (79% duty cycle) then:

$$Read\text{-}out\ rate = \frac{653 \times 10^3}{12} = 54416\ words/second$$

Here, 1 word = 12-bit digital data.

Notice that for the AD2S90, the maximum serial clock rate could be as high as 2 MHz (50% duty cycle). The corresponding conversion circuitry is not discussed here but can be readily developed based on the same principle. The maximum read-out rate corresponding to the 2 MHz serial clock is:

$$max\ read\text{-}out\ time = [600 + (12 \times 500) + 600 + 100] = 7.3\ \mu s$$

Then,

$$max\ read\text{-}out\ rate = \frac{1}{7.3\ \mu s} = 136{,}986\ words/second$$

(1 word = 12-bit digital data)

ABSOLUTE POSITION DATA RS-485 COMMUNICATION

High speed data transmission between a computer system and a remote transducer over long distances, under high noise conditions, usually proves to be very difficult if not impossible with single-ended drivers and receivers. In this case, EIA standard RS-422 or RS-485 balanced digital voltage interfacing is recommended. One benefit of the AD2S90 is its serial absolute position data form, which provides a basis for a simple long distance data transmission configuration. Figure 3 shows the block diagram of RS-485 communication scheme for AD2S90 long distance data transmission. Figure 4 is the circuit using a pair of driver/receiver chips to implement RS-485 communication. In this circuit, the serial position data, clock pulses and load pulse are generated on the driver side and sent to a differential transmitter (75172). On the receiver side, this data can be easily converted to a parallel format with a differential receiver (75173).

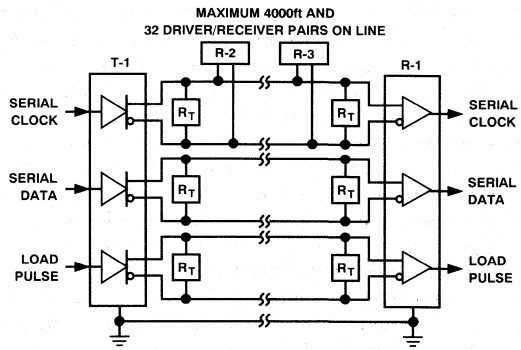

Figure 3. Block Diagram for AD2S90 RS-485 Communication

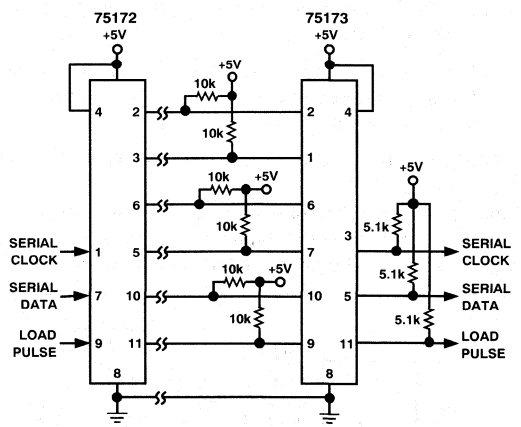

Figure 4. AD2S90 RS-485 Communication

The above 7-wire RS-485 data transmission configuration is basically designed for transmitting the AD2S90 serial data to a remote location where all the data is then converted to parallel. The block diagram of the whole system is shown in Figure 5.

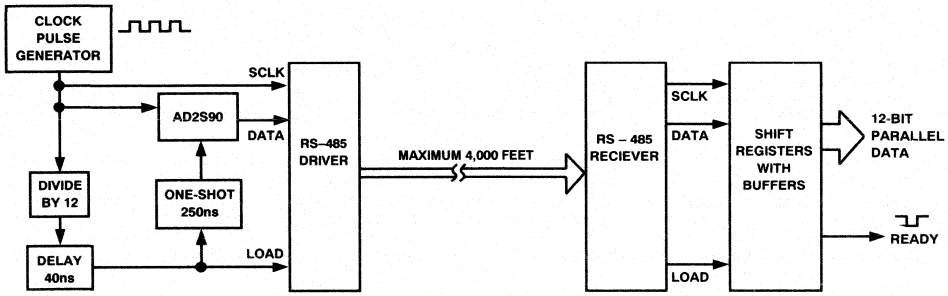

Figure 5. Serial-to-Parallel Data Transmission System Diagram

4096 LINE INCREMENTAL ENCODER EMULATION

One important feature that makes the AD2S90 unique among R/D converters is the quadrature signal output. The device gives out 1024 line A Quad B and north marker signals exactly emulating those from a 1024 line incremental optical encoder. The resolution of the A Quad B can be increased by multiplying the 1024 line count by 4. This produces a 4096 A Quad B output which can be decoded to 16384. An example of 4096 line incremental encoder emulation is shown in Figure 6. Cautions must be taken in selecting the inverter type shown in Figure 6. For example, 4000-Series CMOS will give you longer pulse durations than the AS logic family. In certain cases, inverters can be replaced by delay components to generate more precise timing.

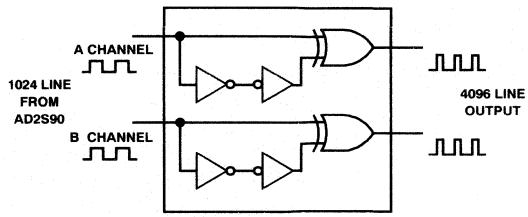

Figure 6. ×4 Circuitry Offering 4096 Incremental Encoder Emulation Using the AD2S90

16

ANALOG DEVICES

AN-234
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Digital Resolver Integration

By Dennis Shi-hsiung Fu

Resolver (Synchro) technology has been used in the industrial environment for many years. This technology still provides the best angular position transducer available in terms of ruggedness, reliability and resolution. The resolver operates on the same principle as a rotating transformer which has one primary winding (usually on the rotor for modern brushless resolvers) and two secondary windings (usually on the stator). Figure 1 shows the trigonometrical relationship between resolver rotors and stator windings:

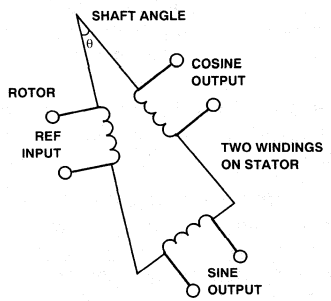

Figure 1. Trigonometrical Relationship Between Resolver Windings

If there is a sinusoidal signal applied across the rotor input, i.e.,

$$V = A\,Sin\omega t$$

then two stator windings will have a certain induced voltage output which has the form:

$$A\,Sin\theta Sin\omega t$$
$$A\,Cos\theta Sin\omega t$$

Here, θ is the resolver shaft angle, assuming that the primary/secondary ratio is 1:1.

Obviously the output signals are quite similar to those amplitude modulated signals with carrier angular frequency (ω). Figure 2 shows the waveform from rotor and stators.

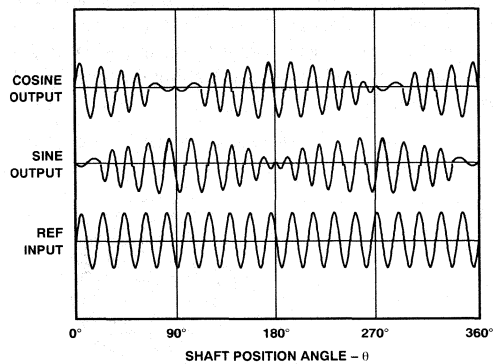

Figure 2. Resolver Signal Waveform

To get position angle θ in digital format, both sine and cosine channel signals must be sent to a resolver-to-digital converter for decoding. Because of the size, output digital format, cost and the complexity of the peripheral circuitry of the resolver-to-digital converter in the past, it is neither easy nor economical to integrate the converter into the resolver, (especially smaller ones like the size 11 or size 15). Now, with the advent of the AD2S90 (with a 20-pin, PLCC package), comes the real solution for industrial applications—a very rugged, low cost angular position transducer with digital position output. This digital resolver integration consists of a sine wave oscillator and a AD2S90 R/D converter. Figure 3 shows the block diagram. There are two sets of signals coming out from the digital resolver, serial absolute position data and A Quad B incremental encoder outputs. All the input/output signals are TTL logic compatible.

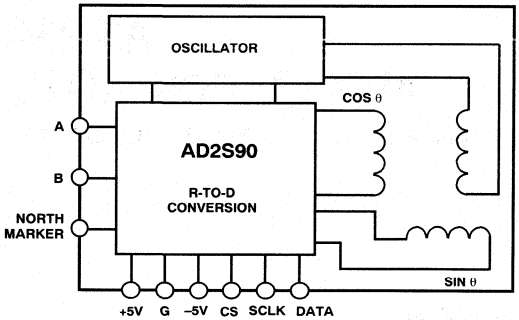

Figure 3. Block Diagram of Digital Resolver

The next section details the implementation of the concept outlined above.

Practical Implementation

The digital resolver designed, prototyped and tested is based on a size 11 brushless resolver and has the original profile of a size 11 resolver with an increase in the overall length. Figure 4 shows the outline dimensions of this prototyping model.

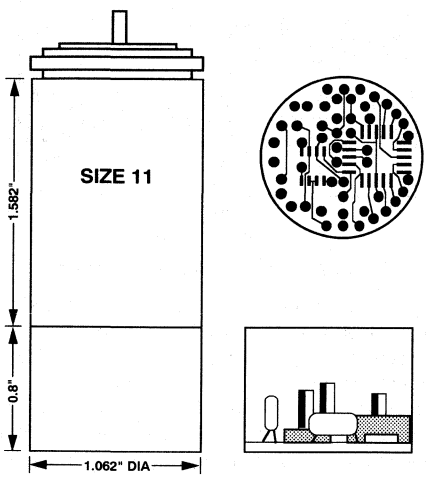

Figure 4. Outline Drawings of the Prototype Digital Resolver (Actual Size)

FEATURES

Size 11, Heavy Duty Resolver
Totally Digital I/O
Absolute Position Data Output
12-Bit Resolution
Very High Readout Rate
Low Cost
Incremental Encoder, A Quad B and NM O/P

There are five inputs (+5 V, −5 V, GND, Serial Clock, Chip Select) and four outputs (Absolute Serial Data, Incremental A Quad B, and North Marker). For high ambient noise and long distance data transmission applications, this digital resolver can be easily upgraded to a configuration which has differential line driver/receivers built in. For more information on this, please refer to application note "Circuit Applications of AD2S90 Resolver-to-Digital Converter."

SYSTEM SPECIFICATIONS

Resolution	12 Bit
Accuracy	15 arc min
Max Read Rate	137 Kbyte/s (1 Byte = 12 Bits Digital Word)
Max Rotating Speed	22,500 rpm
Settling Time	1° Step – 5 ms 179° Step – 20 ms
Logic I/O	Standard TTL
Power Supply	±5 V
Operating Temperature Range	−40°C to +85°C
Outline Dimensions	See the Left Figure (Actual Size)
Weight	≈ 19 oz

For more detailed information on digital timing, specifications and application circuits, please refer to the AD2S90 data sheet and other AD2S90 application notes.

The circuit diagram of this design is shown in Figure 5. A Wienbridge oscillator is used here as an excitation source for the resolver. This particular circuit provides an oscillating frequency at about 3000 Hz, which is determined by R and C.

$$f \approx \frac{1}{2\pi RC}$$

$R = R_1 = R_2$
$C = C_1 = C_2$.

The two resistors at the feedback loop determine waveform shape and distortion. All resistors and capacitors in this application circuit should be chosen with low temperature coefficients, e.g., better than 50 ppm/°C. High quality polycarbonate capacitors should be used to guarantee the temperature stability of the oscillating frequency and to minimize sine wave distortion.

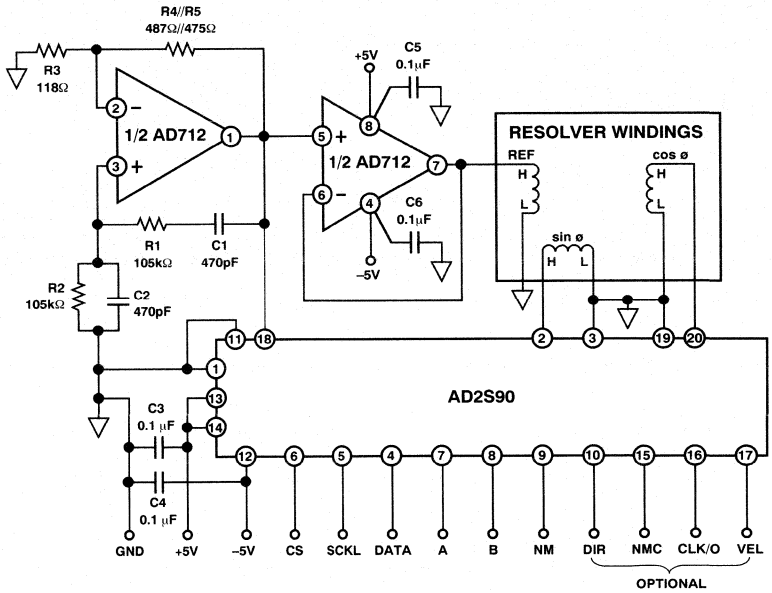

Figure 5. Digital Resolver Circuit Diagram

One application for this digital resolver is shown in Figure 6, which uses one microcontroller to select and read 16 remote location shaft position angles. When the chip select signal is logical HIGH, the corresponding serial data output bus must be in the high impedance state. This simplifies the addressing and reading process.

Another application is the direct replacement and upgrading of the incremental encoder system because of the unique A, B, NM and additional absolute position the digital resolver provides. The resolver can be used as a direct replacement for the encoder without the need to change or upgrade the interface.

COMPONENTS PART LIST

Size 11 Brushless Resolver (1)
AD2S90 Resolver-to-Digital Converter (1)
AD712 Dual Op Amp (1)
RN55C Metal Film Resistors (5)
Polycarbonate Caps (2)
High Quality Ceramic Caps (4)
PCB and 9 Pin D-Connector (1)

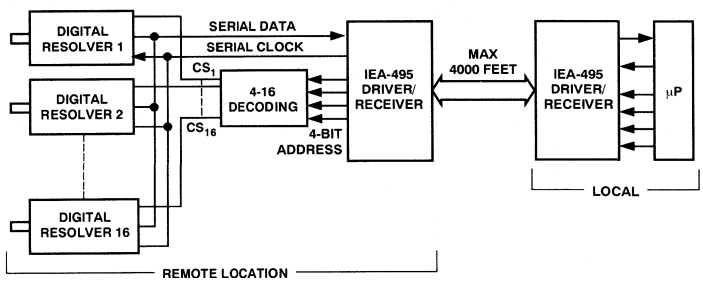

Figure 6. Remote Access to 16 Shaft Angle Position Information Using One Microcontroller

ANALOG DEVICES

AN-236
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Speed Measurement and Control with the AD2S83, a "Silicon Tachogenerator"

By Dennis Shi-hsiung Fu

INTRODUCTION

Continuous and precise control of speed with long-term stability and good transient performance is an important feature of machine control. At the present time, the most popular speed transducer used in the industry is still the mechanical dc tachogenerator, which has the advantages of small size, simple connections and good linearity. But its mechanical commutator is often undesirable because of the regular maintenance required. Mechanical tachogenerators normally use silver-graphite brushes for commutation. They usually have quoted life time of 10^9 revolutions in industrial conditions, which is equivalent to 347 days life with a motor running on average at 2000 rpm. When interruptions of the operation cannot be tolerated or when the transducer is used in inaccessible locations, constant maintenance is difficult; conversely, brushless resolvers have a robust rotor construction which permits reliable maintenance-free operation at a wide speed range. Because the resolver is an absolute position device, it intrinsically provides velocity information.

SILICON TACHOGENERATOR CONFIGURATION

The block diagram of a typical type II tracking R/D converter is shown in Figure 1. With the resolver shaft moving at a constant speed v, the VCO output will correspond exactly to the update rate which is equal to the rate of change in θ per unit of time where instantaneously $\theta = \phi$. This infers that $d\phi/dt$ always tracks the velocity of the input $d\theta/dt$ with zero velocity error (also, without a position error), i.e., $d\phi/dt = d\theta/dt = v$. The signal $d\phi/dt$ is used to represent accurately the resolver shaft rotating speed. Theoretically, the only error in this process is the momentary errors during acceleration or deceleration. How fast the transient error settles usually depends on the dynamics of the tracking converter, which are determined by external passive components (set by user) on some modern monolithic converters. Traditionally monolithic RDCs were optimized for position information outputs only. The AD2S83, based on the AD2S80A, can now provide a very high quality velocity signal with a linearity error as low as ±0.15% of its nominal value.

16

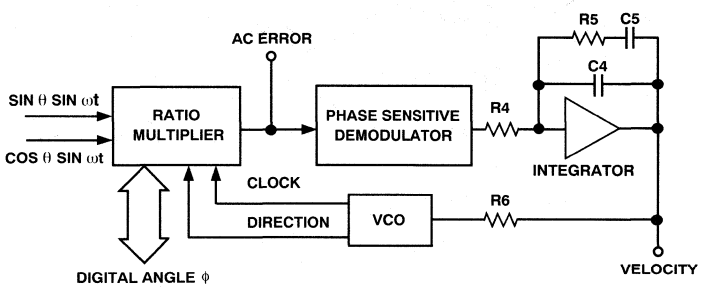

Figure 1. Tracking Resolver-to-Digital Converter Block Diagram

Practically, there are several error sources in the R/D converter velocity signals.

1. Static errors caused by resolver inaccuracies including resolver phase shift and transformer ratio mismatch;
2. An ac component of 2× reference frequency superimposed on the dc velocity signal output, which has the amplitude (rms value) ranging from several millivolts to a hundred millivolts depending on the converter loop bandwidth set by user;
3. LSB rate ripple caused by least significant bit update noise, which usually has a magnitude of approximately sub-millivolt peak if the VCO gain is correctly scaled.

These errors can be minimized with careful consideration in the sizing of resolvers and other passive components. The combination of resolver and R/D converter provides a suitably high quality velocity signal for high precision motion control.

Other advantages of using a resolver with an AD2S83 in speed control include the infinite velocity resolution compared to an encoder system, high speed ratings due to its brushless structure, and offering a cost competitive solution.

AD2S83 AND MECHANICAL TACHOGENERATOR PERFORMANCE TEST

In this performance test, we randomly picked a dc tachogenerator and a size 11 brushless resolver, with the following key specs which are typical of industrial applications:

Tachogenerator:	0.1% linearity, 0.25% reversion error and 3% ripple, over 0–3600 rpm
Brushless resolver:	size 11, 7 arc min accuracy and 0.015 V null voltage. Excitation frequency is 5000 Hz.

The resolver-to-digital converter used is the AD2S83. It has 8 arc min position accuracy and has been specifically designed for applications requiring a high quality velocity signal. The resolution of the position data is 10-, 12-, 14- and 16-bit user selectable, and the converter dynamics are set by external passive components. In the test, the resolution was set to 12 bits with a bandwidth of 520 Hz. Maximum velocity tracking rate in 12-bit mode is 15,600 rpm. To guarantee the accuracy of the test results, two velocity transducers are coaxially connected via a precision gearbox (with minimal backlash). Any velocity error from either the drive or gearbox will be common to both devices and, therefore, only the relative accuracy is compared. The hardware configuration of the testing system is shown in Figure 2. The velocity source is guaranteed to have overall 0.05% accuracy during the test. The test data is detailed in the curves shown in Figure 3 (percentage linearity error), Figure 4 (percentage reversion error) and Figure 5 (ac ripple content). From the comparison curves, it could be clearly seen that the performance of a resolver plus AD2S83 R/D converter is superior to that of a dc tachogenerator.

For resolver plus R/D, the test curve an be extended to its full range of 15,600 rpm, if the velocity source permits. The 0–3,000 rpm range was chosen for the test as tachogenerators usually have limited linear speed range, typically 0–3,000 rpm.

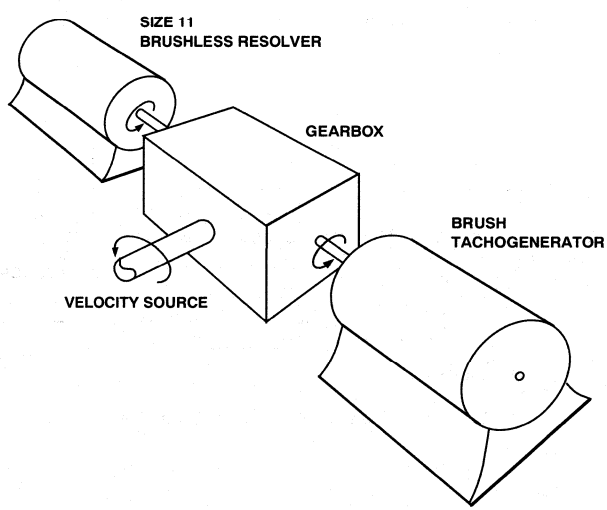

Figure 2. Hardware Configuration

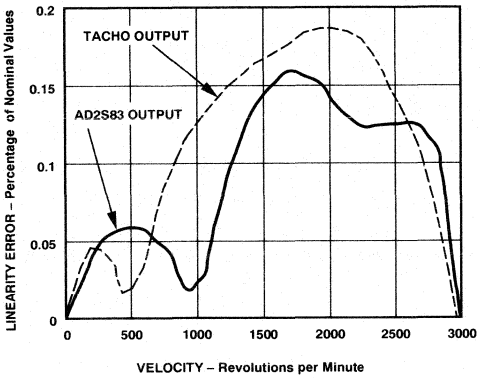

Figure 3. Linearity Error Comparison Curve

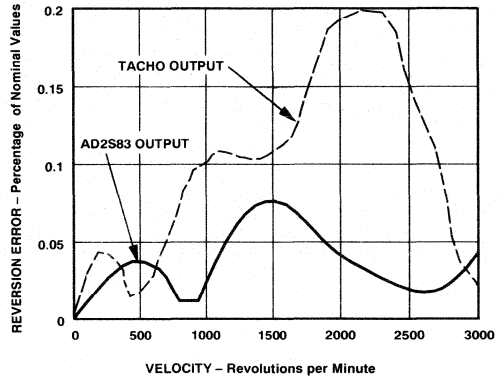

Figure 4. Reversion Error Comparison Curve

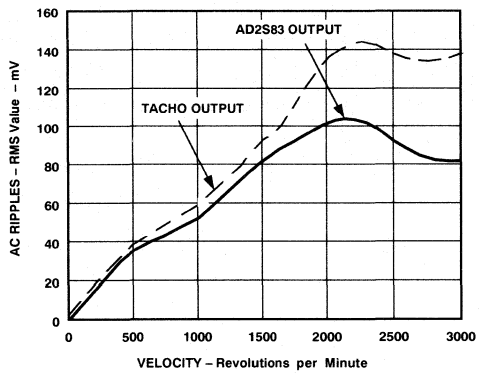

Figure 5. AC Ripple Content Comparison Curve

16

RMS-to-DC Converters
Contents

		Page
AN-268	RMS-to-DC Converters Ease Measurement Tasks	17–3
AN-269	Extend RMS-to-DC Converter's Range	17–15

ANALOG DEVICES

AN-268
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

RMS-to-DC Converters Ease Measurement Tasks

RMS-to-DC converters compute the true rms value of a signal without regard to waveform. You can use them as building blocks in a variety of measurement circuits.

by Bob Clarke, Mark Fazio, and Dave Scott

INTRODUCTION

RMS-to-DC converters provide a dc output equal to the rms value of an ac or fluctuating dc input. Analog Devices provides a selection of five such converters: the AD536A, AD636, AD637, AD736, and AD737. The theory and applications of the AD536A, AD636, and AD637 are covered in great detail in the *RMS-to-DC Conversion Application Guide, Second Edition.*

The newer AD736 and AD737 rms-to-dc converters, however, are not covered in the guide. This application note supplements the guide and discusses the theory and applications of the AD736 and AD737. It also contains information on how to enhance the accuracy and reduce the settling time of the AD637.

There are five main sections: How RMS-to-DC Converters Work, How to Select an RMS-to-DC Converter, Theory of the AD736 and AD737, Applications of the AD736 and AD737, and Applications of the AD637. Other sources of information on rms-to-dc converters include the *RMS-to-DC Conversion Application Guide, Second Edition,* the *Nonlinear Circuits Handbook, Second Edition,* and the data sheets for the AD536A, AD636, AD637, AD736, and AD737; all are available from Analog Devices.

WHY USE AN RMS-TO-DC CONVERTER?

Early multimeters used a simple rectifier and averaging circuit for ac measurements. These meters were then calibrated to read the rms value, but this was correct only for one waveform, invariably, a sine wave. In contrast to averaging circuits, true rms-to-dc converters measure the rms value of an input signal without regard to waveform.

Waveforms differ in their crest factor, which is defined as the ratio of the peak signal amplitude to the rms amplitude, that is, Crest Factor = V_{PEAK}/V_{RMS}. Many common waveforms, such as sine and triangle waves, have relatively low (≤2) crest factors. Such other waveforms as low-duty-cycle pulse trains and SCR waveforms have high crest factors.

To obtain accurate results using an averaging circuit, the user would have to know the waveform in advance and apply a correction factor. RMS-to-DC converters provide an accurate answer for a variety of crest factors. The AD637 handles crest factors as large as 10 with no more than 1% additional error; the AD736 and AD737 handle crest factors as large as 5. Table I contrasts true rms values and the measurement errors introduced by average-responding circuits for various waveforms.

HOW RMS-TO-DC CONVERTERS WORK

The rms-to-dc converters discussed here solve an implicit equation for the rms value of a voltage. The following discussion will show the transformations that lead from the definition of rms voltage to the implicit equation. It will then explain the implementation of the implicit equation in a monolithic rms-to-dc converter.

The definition of the rms value of a voltage is

$$V_{RMS} = \sqrt{\frac{1}{T}\int_0^T [V(t)^2]\,dt} \qquad (1)$$

where V_{RMS} is the rms value, T is the duration of the measurement, and V(t) is the instantaneous voltage, a function of time, but not necessarily periodic.

Squaring both sides of this equation yields

$$V_{RMS}^2 = \frac{1}{T}\int_0^T [V(t)^2]\,dt\,. \qquad (2)$$

The integral can be approximated as a running average:

$$Avg\,[V(t)]^2 = \frac{1}{T}\int_0^T [V(t)^2]\,dt \qquad (3)$$

then Equation 2 simplifies to

$$V_{RMS}^2 = Avg\,[V(t)^2]. \qquad (4)$$

17

Table I. Error Introduced by an Average Responding Circuit When Measuring Common Waveforms

Waveform Type 1 Volt Peak Amplitude	Crest Factor (V_{PEAK}/V_{RMS})	True RMS Value	Average Responding Circuit Calibrated to Read RMS Value of Sine Waves Will Read	% of Reading Error* Average Responding Circuit
Undistorted Sine Wave	1.414 ($\sqrt{2}$)	0.707 V	0.707 V	0%
Amplitude-Symmetrical Square Wave	1.00 (Exact)	1.00 V	1.11 V	+11.0%
Triangle Wave	1.732 ($\sqrt{3}$)	0.577 V	0.555 V	−4%
Gaussian Noise (98% of Peaks <1 V)	3	0.333	0.266	−20.2%
Examples of Unipolar Pulse Trains	2 10	0.5 V 0.1 V	0.25 V 0.01 V	−50% −90%
SCR Waveforms 50% Duty Cycle 25% Duty Cycle	2 4.7	0.495 V 0.212 V	0.354 V 0.150 V	−28% −30%

$$\%\ of\ Reading\ Error = \frac{Average\ Responding\ Value - True\ RMS\ Value}{True\ RMS\ Value} \times 100\%$$

Dividing both sides by V_{RMS} yields

$$V_{RMS} = \frac{Avg\,[V(t)^2]}{V_{RMS}}. \quad (5)$$

This expression provides the basis for the implicit solution for V_{RMS}, and is the technique used in Analog Devices' line of monolithic rms-to-dc converters.

Note that taking the square root of both sides of Equation 4 yields

$$V_{RMS} = \sqrt{Avg\,[V(t)^2]} \quad (6)$$

which is an alternate way of expressing the rms (root of mean of square) value of the function.

The implicit method of rms computation is preferable to the explicit method (successively squaring, averaging, and taking the square root of the input signal) for practical reasons that result in a superior dynamic range. Using the explicit method, the output of the squarer will vary over a 10,000:1 dynamic range (1 mV to 10 V) for a 100:1 (0.1 V to 10 V) instantaneous input. Since the input squarer used in the explicit method will have errors greater than 1 mV, the error will strongly depend on signal level, resulting in an overall dynamic range of less than 100:1.

Figure 1 shows the implicit method of rms-to-dc conversion. The circuit is essentially an analog computer that solves Equation 5. The Analog Devices AD536A, AD636, AD637, AD736, and AD737 all use variations on this theme.

The input stage is a unity-gain buffer, which is uncommitted in the AD536A, AD636, and AD637 and committed in the AD736 and AD737. "Uncommitted" in this context means that both inputs and the output connection are accessible; the user has the option of using this buffer as a high impedance input for the converter, using it to build an active filter to follow the rms-to-dc converter's own averaging filter, or simply leaving it unconnected.

An absolute-value circuit (that is, a precision full-wave rectifier) follows the input buffer. The output of the absolute-value circuit drives a squarer/divider. The squarer/divider squares the input signal and divides it by the output signal, which is the averaged output of the squaring circuit. By closing the loop around the divider, Equation 5 is solved continuously.

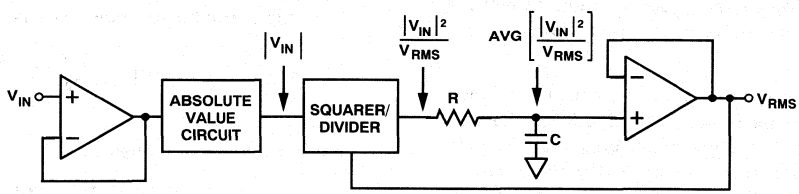

Figure 1. Implicit Method of RMS-to-DC Conversion Used in AD536A, AD636, AD637, AD736, and AD737

HOW TO SELECT AN RMS-TO-DC CONVERTER

Selecting an rms-to-dc converter means picking the product whose attributes best match the requirements of the application. Unfortunately, no one converter fits every situation, so trade-offs must be made between accuracy, bandwidth, power consumption, input signal level, crest factor, and settling time.

The AD637 accepts input voltages as high as 7 V rms and is Analog Devices' most accurate *and* widest-bandwidth rms-to-dc converter. Its −3 dB bandwidth is 8 MHz for a 1 V rms input. It has an auxiliary dB output that is proportional to the logarithm of the input signal over a 60 dB range and a power-down feature that reduces its quiescent current from 3 mA to 450 μA.

Both the AD736 and AD737 are optimized for use in portable instruments; they consume less than 200 μA of quiescent current and accept signal levels from 0 to 200 mV rms; as we'll see later, external attenuators can be added for other signal ranges. The AD737 also has a power-down input, which allows the user to reduce its quiescent current from 160 μA to 40 μA in portable applications. Table II summarizes these specifications for the AD637, AD736, and AD737. The AD637 is the best all-around performer, with its superior combination of accuracy, dynamic range, crest factor, and settling time. It also has the widest bandwidth, as shown in Table III.

The AD637 should also be chosen if the application requires high accuracy and a quick response for large, abrupt changes in signal level. The AD637's settling time is independent of signal level, while, for a given value of averaging capacitor, the settling time of the AD736 and AD737 depends on signal level, being longer for low level signals and shorter for high level signals.

Although they have less bandwidth, the AD736 and AD737 perform better than the AD637 for low level signals (<10 mV) and consume less power. (Later in this application note, we will see how to improve the AD637's performance for low level (<20 mV) signals by using an external preamplifier.) They can also be used as general purpose parts, replacing such op amp circuits as averaging converters and precision rectifiers. Both the AD737 and AD637 also have a power-down feature.

Table II. RMS-to-DC Converter Selection Guide

Model	Conversion Accuracy ±mV±% Reading	Maximum Power Consumption	Continuous Input (V_{RMS})	Crest Factor for ≤ 1% Additional Error	Relative Settling Time	Comments
AD637J	±1 mV ± 0.5%	3 mA @ ±15 V	7 @ V_S = ±15	≤10	Fast	Highest Accuracy Highest Bandwidth
AD637K	±0.5 mV ± 0.2%	3 mA @ ±15 V	7 @ V_S = ±15	≤10		Precision Applications
AD736A/J	±0.5 mV ± 0.5%	0.2 mA @ ±5 V	1 @ V_S = ±5	≤3	Slow	Low Cost Low Power
AD736B/K	±0.3 mV ± 0.3%	0.2 mA @ ±5 V	1 @ V_S = ±5	≤3		Output Buffer
AD737A/J	±0.4 mV ± 0.5%	0.16 mA @ ±5 V	1 @ V_S = ±5	≤3	Slow	Low Cost Lowest Power
AD737B/K	±0.2 mV ± 0.3%	0.16 mA @ ±5 V	1 @ V_S = ±5	≤3		No Output Buffer

17

Table III. RMS-to-DC Converter Bandwidth vs. Accuracy

Bandwidth (kHz) for 1% Additional Error for	AD637	AD736		AD737	
		Pin 1	Pin 2	Pin 1	Pin 2
V_{IN} = 1 mV	NA	1 kHz	1 kHz	1 kHz	1 kHz
V_{IN} = 10 mV	NA	6 kHz	6 kHz	6 kHz	6 kHz
V_{IN} = 20 mV	11 kHz	NA	NA	NA	NA
V_{IN} = 200 mV	66 kHz	90 kHz	33 kHz	90 kHz	33 kHz
3 dB Bandwidth (kHz) for					
V_{IN} = 1 mV	NA	5 kHz	5 kHz	5 kHz	5 kHz
V_{IN} = 10 mV	NA	55 kHz	55 kHz	55 kHz	55 kHz
V_{IN} = 20 mV	150 kHz	NA	NA	NA	NA
V_{IN} = 200 mV	1000 kHz	460 kHz	190 kHz	460 kHz	190 kHz

THEORY OF THE AD736 AND AD737

To better understand how the AD736 works, consider the simplified block diagram, first as drawn on the AD736 data sheet (Figure 2), and then redrawn with the averaging (C_{AV}) and filter (C_F) capacitors in Figure 3 to better show signal flow.

The input to the AD736 and AD737 is through a FET-input op amp connected as a unity-gain buffer. This amplifier allows both a high impedance, buffered input (Pin 2) or a low impedance input (Pin 1) that provides a wider dynamic range. The high impedance input, with its low input bias current, is well suited for use with high impedance input attenuators.

The output of the buffer drives a full-wave rectifier or absolute value circuit, which in turn drives a 2-quadrant squarer/divider. The output of the squarer/divider drives the summing node of an inverting op amp connected as a current-to-voltage converter. Pin 3 gives access to this node to connect a filter capacitor in parallel with the 8 kΩ feedback resistor to form a 1-pole low-pass filter.

The AD737 (Figures 4 and 5) is similar in design and function to the AD736 except that the AD737 lacks an output buffer, which is omitted in order to reduce power consumption, and has a power-down feature that further reduces power consumption. Its output stage is a simple open-collector NPN transistor with an 8 kΩ load resistor.

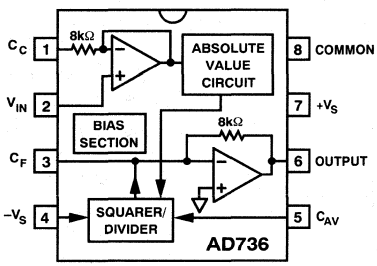

Figure 2. Simplified Block Diagram of the AD736

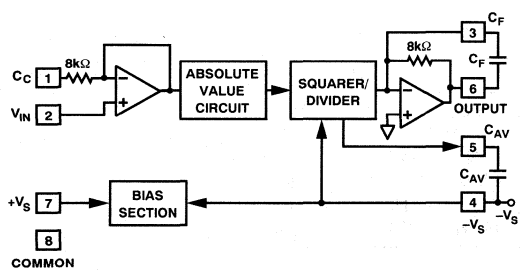

Figure 3. Redrawn Simplified Block Diagram of the AD736

It develops its output voltage by sinking current through this resistor. The external averaging capacitor (C_{AV}) for the AD736 and AD737 is connected between Pins 4 ($-V_S$) and 5 (C_{AV}), which places it across a transistor's base-emitter junction in the rms core. This means that the resistance in parallel with the averaging capacitor is that of a diode, and thus signal-level dependent. The resulting time constant is inversely proportional to the rms value.

Because the external averaging capacitor, C_{AV}, "holds" the rectified input signal during rms computation, its value directly affects the accuracy of the measurement —especially at low frequencies. (The larger the value of C_{AV}, the lower the error.) Also, because the averaging capacitor appears across a base-emitter junction in the squarer/divider whose resistance varies with signal level, the averaging time constant will increase linearly as the input signal is reduced.

Consequently, as the input level decreases, errors due to nonideal averaging will *decrease* while the time it takes for the circuit to settle to the new rms level will *increase*. Therefore, lower input levels allow the circuit to perform better (due to increased averaging) but increase the waiting time between measurements because the capacitor takes longer to discharge. Thus, a trade-off between computational accuracy and settling time is required. This topic is discussed in detail in the *RMS-to-DC-Conversion Application Guide, Second Edition.*

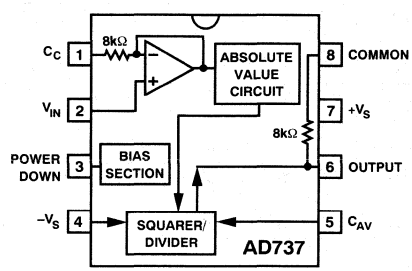

Figure 4. Simplified Block Diagram of the AD737

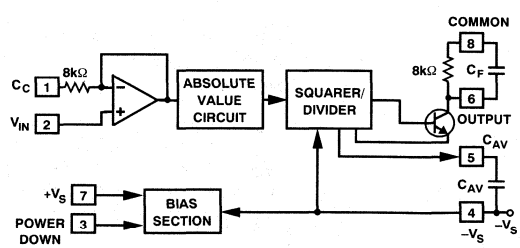

Figure 5. Redrawn Simplified Block Diagram of the AD737

)C Error, Output Ripple, and Averaging Error

'igure 6 shows the typical output waveform of the \D736 and AD737 with a sine wave input applied. The deal output of V_{OUT} = rms (V_{IN}) is never achieved; nstead, the output contains both a dc error and an ac ipple component.

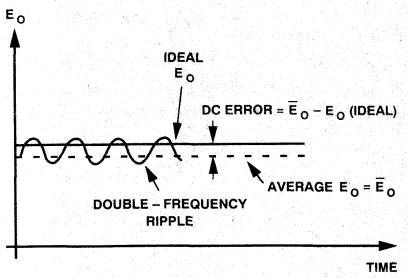

Figure 6. Output Waveform of the AD736 and AD737 for a Sine Wave Input Voltage

The dc error is the difference between the average of the output signal (when the ripple in the output has been removed by filtering) and the ideal dc output. The dc error component is therefore set solely by the value of averaging capacitor used—no amount of post filtering (i.e., using a very large C_F) will reduce this error, although the ripple may be removed by using a large value of C_F.

In most cases, the combined magnitudes of both the dc and ac error components need to be considered when selecting values of capacitors C_{AV} and C_F. This combined error, representing the maximum uncertainty of the measurement is termed the "averaging error" and is equal to the peak value of the output ripple plus the dc error.

As the input frequency increases, both the dc and ac error components decrease rapidly: if the input frequency doubles, the dc error and ripple reduce to 1/2 and 1/4 their original values, respectively, and rapidly become insignificant. Table IV provides practical values of C_{AV} and C_F for several common applications. Figure 7 shows the additional error versus crest factor of the AD736 and AD737 for various values of C_{AV}.

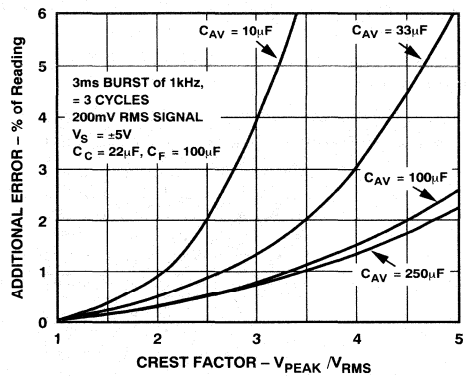

Figure 7. Additional Error vs. Crest Factor of the AD736 and AD737 for Various Values of C_{AV}

Table IV. Practical Values for C_{AV} and C_F for AD736 and AD737

Application	RMS Input Level	Low Frequency Cutoff (−3 dB)	Max Crest Factor	C_{AV}	C_F	Settling Time* to 1%
General Purpose RMS Computation	0 V–1 V	20 Hz	5	150 μF	10 μF	360 ms
		200 Hz	5	15 μF	1 μF	36 ms
	0–200 mV	20 Hz	5	33 μF	10 μF	360 ms
		200 Hz	5	3.3 μF	1 μF	36 ms
General Purpose Average Responding	0 V–1 V	20 Hz		NONE	33 μF	1.2 sec
		200 Hz		NONE	3.3 μF	120 ms
	0 mV–200 mV	20 Hz		NONE	33 μF	1.2 sec
		200 Hz		NONE	3.3 μF	120 ms
SCR Waveform Measurement	0 mV–200 mV	50 Hz	5	100 μF	33 μF	1.2 sec
		60 Hz	5	82 μF	27 μF	1.0 sec
	0 mV–100 mV	50 Hz	5	50 μF	33 μF	1.2 sec
		60 Hz	5	47 μF	27 μF	1.0 sec
Audio Applications						
Speech	0 mV–200 mV	300 Hz	3	1.5 μF	0.5 μF	18 ms
Music	0 mV–100 mV	20 Hz	10	100 μF	68 μF	2.4 sec

*Settling time is specified over the stated rms input level with the input signal increasing from zero. Settling times will be greater for decreasing amplitude input.

Calculating AD737 Settling Time

The graph of Figure 8 may be used to approximate the time required for the AD736 or AD737 to settle when its input level is reduced in amplitude. The total time required for the rms converter to settle will be the difference between two settling times extracted from the graph—the initial settling time minus the final settling time.

As an example, consider the following conditions: a 33 μF averaging capacitor, an initial rms input level of 100 mV, and a final (reduced) input level of 1 mV. From Figure 8, the initial settling time (where the 100 mV line intersects the 33 μF line) is around 80 ms. The settling time corresponding to the new or final input level of 1 mV is about 8 seconds. Therefore the net time for the circuit to settle to its new value will be dominated by the final settling time.

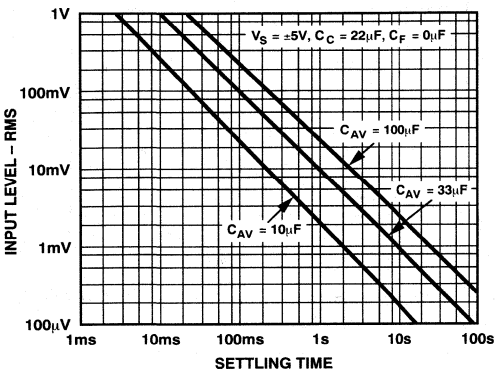

Figure 8. Settling Time vs. RMS Input Level of the AD736 and AD737 for Various Values of C_{AV}

APPLICATIONS OF THE AD736 AND AD737

AD736 as Precision Rectifier

Building a precision rectifier requires two op amps, two diodes, and a handful of matched resistors. An easy way to replace all these parts and save some board space is to use an rms-to-dc converter. Just omit the averaging capacitor and disconnect the feedback; this uses only the converter's internal precision rectifier (Figure 9), which, being monolithic, has inherently matched diodes.

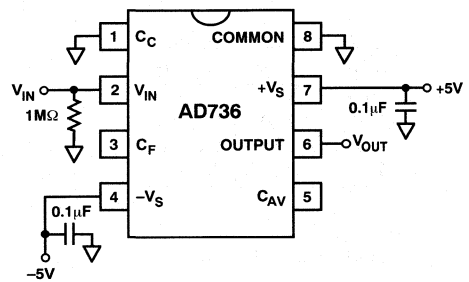

Figure 9. AD736 Connected as a Precision Rectifier

One note about precision rectifiers: as the input waveform crosses through zero, the op amp must instantaneously switch one diode on and the other off. For this reason, the bandwidth of precision rectifiers is always much less than one might otherwise expect based on

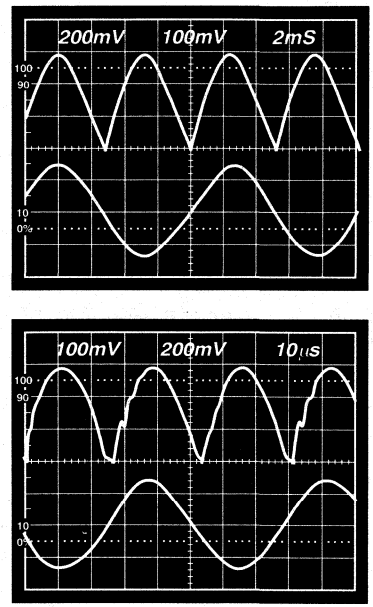

Figure 10. Performance of AD736 Precision Rectifier at 1 kHz (Top Photo) and 19 kHz (Bottom)

the gain-bandwidth product, open-loop gain, and slew rate of the op amp. The monolithic precision rectifiers used in rms-to-dc converters fare much better than discrete precision rectifiers in this regard, as the photo of the AD736's performance as a precision rectifier in Figure 10 shows.

Extending the AD736 and AD737 Full-Scale Input Ranges

The high impedance input (Pin 2) of the AD736 and AD737 allows simple resistive attenuators (Figure 11) to be used to extend their input range. Without input attenuation, both the AD736 and AD737 can accurately measure input signals as large as 200 mV rms with crest factors of 1 to 3.

The external attenuator simply reduces the full-scale input to the 200 mV rms input range of the AD736 or AD737. For a maximum 7 V rms input (10 V peak) input, for example, the attenuator should be a 35:1 (7/.2) voltage divider. The reading of the converter should be scaled by the factor of attenuation used. An external attenuator can also be used with the converter's low impedance input (Pin 1), as will be shown later in Figure 13.

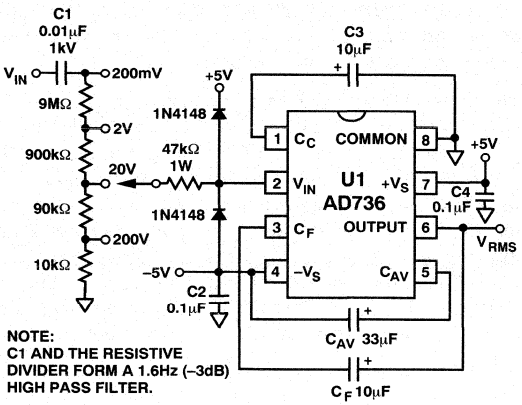

Figure 11. By using an external input attenuator, you can extend the measurement range of the AD736 and AD737. Although the AD736 is shown, this technique also works with the AD737.

Single-Supply Operation of the AD736

In dual-supply operation, the output (Pin 6) of the AD736 is at 0 V, halfway between the supply rails. But in single-supply operation, the output is at 1/2 V_{CC}. By adding a single-supply op amp as a differential amplifier, however, you can build a true "0 V out for 0 V" single-supply circuit with a ground-referenced output (Figure 12). For this circuit, V_{RMS} = 0 when V_{IN} = 0 and V_{RMS} = 200 mV dc when V_{IN} = 200 mV rms.

In this circuit, a single 9 V positive supply powers the AD736. Resistors R7 and R8 form a voltage divider across the 9 V battery that establishes a local "ground" rail at 1/2 V_{CC}, or 4.5 V. The AD736's "COMMON" pin, its 22 MΩ input bias resistor, and the inverting input of U2 (via R4 and R5) are all connected to this rail. The quiescent output voltage of the AD736, which is referenced to its COMMON pin, is 4.5 V.

A single-supply op amp, U2, is connected as a unity-gain, differential amplifier. Large-valued feedback resistors (R2 through R5) are used to minimize loading of the 4.5 V rail. U2 amplifies the difference between local ground at 4.5 V and the output of the AD736, which is also at 4.5 V for 0 V rms input. As the rms input to the AD736 increases from 0 mV to 200 mV, the AD736's output increases from 4.5 V to 4.7 V. U2's output is the difference between the AD736's output and 4.5 V, or 0 mV to 200 mV dc.

The remainder of the circuit works as follows. The AD736's input is ac coupled; R1 provides a path for the BiFET op amp's input bias current (typically 1 pA) to flow. The offset voltage due to bias current flowing through R1's 22 MΩ resistance is negligible. C3 (10 μF), connected between U1's Pins 1 and 8, provides a low frequency cutoff of 2 Hz; you can select other cutoff values using

$$f = \frac{1}{2\pi RC} \tag{7}$$

where f is the −3 dB frequency in Hz, C is in farads, π is 3.1416, and R is fixed at 8 kΩ ±20% by the AD736.

The averaging capacitor, C_{AV}, is 33 μF and is connected between pins 4 and 5 of U1. An optional 10 μF filter capacitor, C_F, in parallel with an 8 kΩ feedback resistor across the output buffer forms a 1-pole low-pass filter with a 2 Hz cutoff frequency.

You can calculate the value of C_F using the expression

$$f = \frac{1}{2\pi RC_F} \tag{8}$$

where f is the −3 dB frequency in Hz, C_F is in Farads, π is 3.1416, and R is fixed at 8 kΩ ±20% by the AD736. Or, since R is fixed,

$$f = \frac{20\ Hz}{C_F\ (\mu F)}. \tag{9}$$

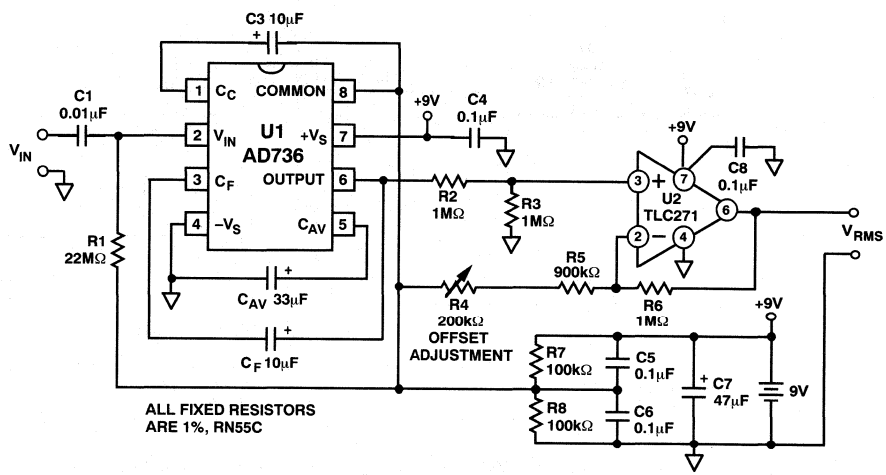

Figure 12. By using a single supply op amp to level shift the AD736's output, you can build a true single supply circuit that supplies 0 V output for 0 V input.

Single-Supply Operation of the AD737

You can also build a "true" single-supply circuit—0 V out for 0 V in—using the AD737. Note that the circuit in Figure 13 shows three design techniques: how to operate the AD737 from a single supply, how to use a resistive attenuator in series with the low impedance input of the AD737, and how to use a single-supply op amp to convert the AD737's output current to a voltage.

The combination of the input attenuator (R1 and R2) and the output op amp (U2) allows the circuit to provide 0 V dc to 2 V dc output for 0 V rms to 2 V_{RMS} full-scale input. You can also use this resistive attenuator with the AD736. The circuit consumes just 192 μA from a 9 V supply at 10 mV rms input and 240 μA from a 9 V supply at 2 V rms input.

At the input of the circuit, an additional resistance (formed here by the sum of R1 and R2) in series with Pin 1 serves as an attenuator. You can calculate the value of this resistance using the expression

$$R_{IN} = \frac{8\ k\Omega \times V_{FS}}{0.2\ V} - 8\ k\Omega \qquad (10)$$

where R_{IN} is the value of the series input resistance and V_{FS} is the desired full-scale input voltage. For the 2 V full-scale input voltage used here,

$$R_{IN} = \frac{8\ k\Omega \times 2\ V}{0.2\ V} - 8\ k\Omega \qquad (11)$$

which yields R_{IN} = 72 kΩ. For a 10 V full-scale input voltage, R_{IN} would equal 392 kΩ. Due to the thin film resistors used in the manufacturing process for the AD737 (and AD736), the tolerance of R_{IN} is 20%. Thus, the external resistance (R1 plus R2) must be 72 kΩ ±20% to compensate for the tolerance of the internal resistance.

As in the AD736 single supply circuit, a single 9 V battery powers the AD737, and two 100 kΩ resistors (R7 and R8) across the 9 V battery form a voltage divider that establishes a local "ground" rail at 1/2 V_{CC}, or 4.5 V. A single-supply op amp, U2, serves as an I/V converter that produces a 2 V full-scale output.

The output of the AD737 is an open-collector NPN transistor that normally sinks current through the 8 kΩ resistor connected through the COMMON pin to ground. In this circuit, the COMMON pin is left unconnected, and the AD737 sinks current from the node at the inverting input of U2. In order to maintain the same potential at its inverting and noninverting nodes, U2 increases its output voltage (V_{RMS}), which increases the current through feedback resistor R6. This circuit is balanced when the current through the R6 equals the current absorbed by the AD737.

The gain of the output stage can be increased by simply increasing the value of the feedback resistor, R6. The combination of C4 and R6 form the post filter in this circuit and, because R6 is an order of magnitude larger than the 8 kΩ internal resistor in the AD736 and AD737, this circuit allows a ×10 smaller value of C4 for the same amount of filtering (i.e., T = C4R6 = $C_F \times 8\ k\Omega$).

A 3-Chip Digital Panel Meter for Differential Current or Voltage Measurement

By using an AD22050 difference amplifier, an AD737 rms-to-dc converter, and an ICL7136 single-chip DMM, you can build a complete, 3-chip digital panel meter that measures the rms value of an alternating current (or voltage). The circuit uses very little power: the AD22050 draws less than 300 μA quiescent current and the AD737 draws less than 160 μA quiescent current.

Figure 14 shows the circuit. The transfer function of the circuit for a 200 mV full-scale reading is

$$20 \times I_{IN} R_{SENSE} = 200\ mV \qquad (12)$$

where 20 is U1's gain, I_{IN} is the input current in Amperes, R_{SENSE} is the value of the sense resistor in Ohms,

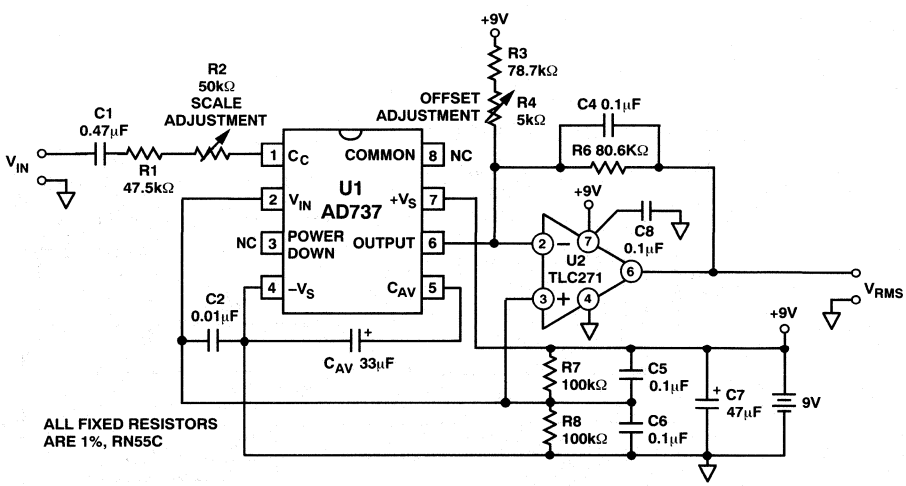

Figure 13. You can also build a true single-supply circuit using the AD737. Note that this circuit also shows how to use an input attenuator with the AD737's low impedance input.

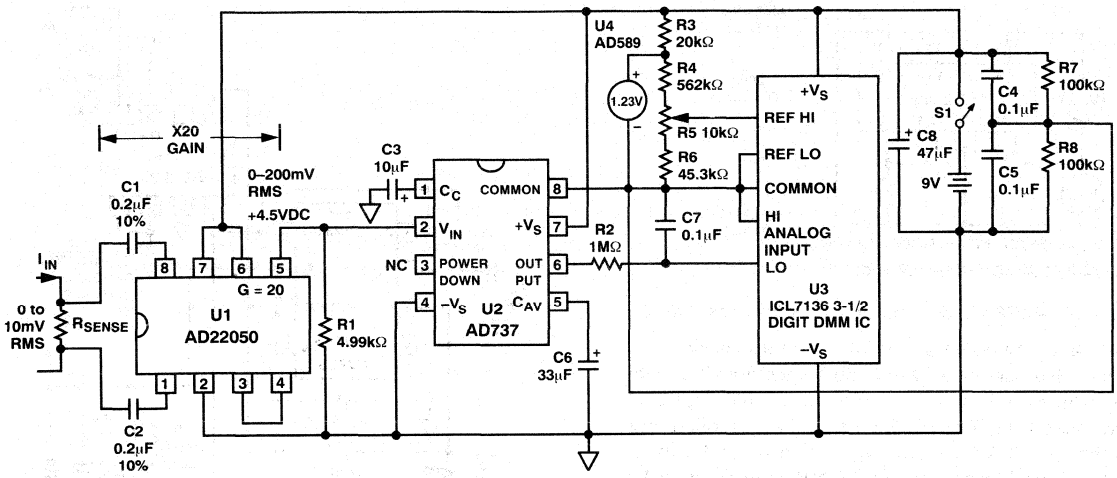

Figure 14. By using an AD737 rms-to-dc converter, an AD22050 difference amplifier, and an ICL7136 single-chip DMM, you can build a true rms low power digital panel meter.

and 200 mV is the full-scale reading of the ICL7136. A more convenient expression that provides a 100 mV full-scale reading for 10 mV across R_{SENSE} is

$$R_{SENSE} = \frac{5\ mV}{I_{IN}}. \qquad (13)$$

For example, to measure a full-scale current of 100 mA rms with a full-scale reading of 100 mV:

$$R_{SENSE} = \frac{5\ mV}{100\ mA} = 50\ m\Omega. \qquad (14)$$

Here's how the circuit works. The input current is converted to a voltage by R_{SENSE}. The input amplifier, U1, an AD22050, is a single-supply difference amplifier that amplifies the input signal by a factor of ×20; the gain of the AD22050 can be changed by adding external resistors (as shown on the next page). Its −3 dB bandwidth is 100 kHz and its slew rate is 0.1 V/μs. Note that a mismatch between C1 and C2 will degrade the CMRR of this circuit.

Connecting Pin 7 of the AD22050 to the positive supply pulls its zero-signal output at Pin 5 to 1/2 of the 9 V supply, or 4.5 V. Resistors R7 and R8 split the supplies for the AD737. The AD737 supplies a differential output between its Pins 6 and 8 to the COMMON and LO inputs of a 3-1/2-digit DMM IC, an ICL7136. (The 3-1/2-digit display and its connections are omitted for simplicity.) R2 and C7 form a simple RC filter. The small value of C7 is made possible by U3's high input impedance. U4, an AD589, provides an external 1.23 V reference to calibrate U3. To calibrate this circuit, adjust R5 to provide a 100 mV reference voltage between the REF HI and REF LO inputs of the ICL7136.

Changing the AD22050's Input Gain

The AD22050 consists of two stages: an input preamplifier and an output buffer. The gain of the AD22050 preamplifier, from the input Pins 1 and 8 to its output at Pin 3, is ×10, and that of the output buffer is ×2, making the overall gain ×20. Many applications will call for gains greater than or less than ×20. Both of these situations are readily accommodated by the addition of one external resistor. Note that this is possible because the output resistance of the input buffer (at Pin A1) is deliberately raised to 100 kΩ ±1%.

The gain may be raised by connecting a resistor from the output of the buffer amplifier (Pin 5) to its noninverting input (Pin 4) as shown in Figure 15. The gain is now multiplied by the factor R/(R−100), where R is in kΩ; for example, the gain is doubled for R = 200 kΩ. Overall gains as high as 160 are readily achievable in this way. Note that the gain becomes increasingly dependent on the accuracy of the resistor value at high gains.

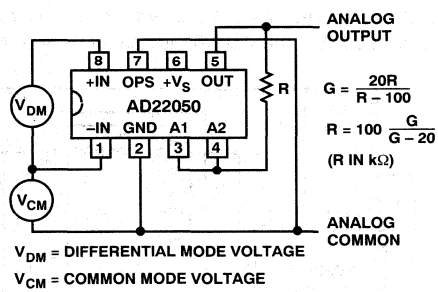

Figure 15. AD22050 Configured for Gains Greater than 20

Because the output of the AD22050 preamplifier has an output resistance of 100 kΩ (±1%), an external resistor connected from Pin 4 to ground (Figure 16) lowers the gain by a factor R/(100+R), where again R is in kΩ. When configuring the AD22050 for low gains, however, care should be taken not to exceed the output capabilities of the preamplifier, because the preamplifier with its gain of ×10 may saturate before the AD22050's output stage does.

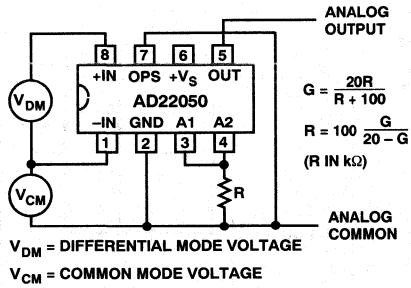

Figure 16. AD22050 Configured for Gains Less than 20

AD736 Single-Supply Circuit With Output Scaling

You can also use the AD22050 difference amplifier with the AD736 to build a 0 V in, 0 V out circuit. Figure 17 shows the circuit. The AD22050 amplifies the difference between the COMMON and V_{OUT} with a fixed gain of ×20, transforming the 0 mV to 200 mV input range to 0 V to 4 V. The output of the AD22050 can go within about 20 mV of ground, so the useful range of this circuit for 1% of reading accuracy is 10 mV to 200 mV ac rms input for 100 mV to 4 V dc output. The bandwidth of this circuit for less than 1 of reading error is 40 Hz to 6 kHz at 10 mV rms input, extending to 36 kHz at 200 mV rms input.

You can add an output low-pass filter by placing a capacitor from the junction of Pins 3 and 4 of the AD22050 to ground. The −3 dB cutoff frequency of this filter is

$$f = \frac{1}{2\pi C \times 100\ k\Omega} \qquad (15)$$

where C is in farads. Alternately, f = 1.59 Hz per μF.

Transmit Current Measurements Using 4-to-20 mA Transmitter

You can also measure alternating current and transmit the results on a 4-to-20 mA current loop. Figure 18 shows the AD22050, the AD736, and the AD694, a 4-to-

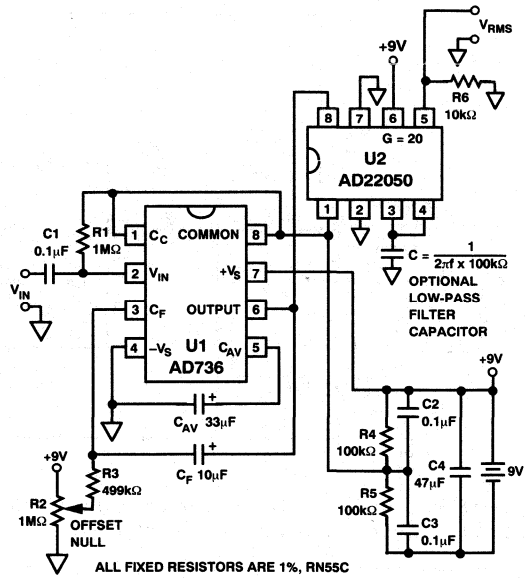

Figure 17. You can build a single supply circuit with X20 gain and optional filtering using the AD736 and AD22050.

20 mA transmitter, configured as a 0-to-10 mV ac rms in, 4-to-20 mA out current measuring subsystem for use in process-control loops.

This circuit builds on the techniques shown in the previous circuits. For example, the AD22050 provides a differential-in, single-ended out, current sensor. Here the AD22050 operates at a gain of ×20 as before, and drives the low impedance input (8 kΩ , Pin 1) of the AD736.

Because of their low power consumption, both the AD22050 and the AD736 can operate from 10 V supplied by the AD694's Pin 7 reference output. The AD694 operates from a +24 V single supply. Because this circuit

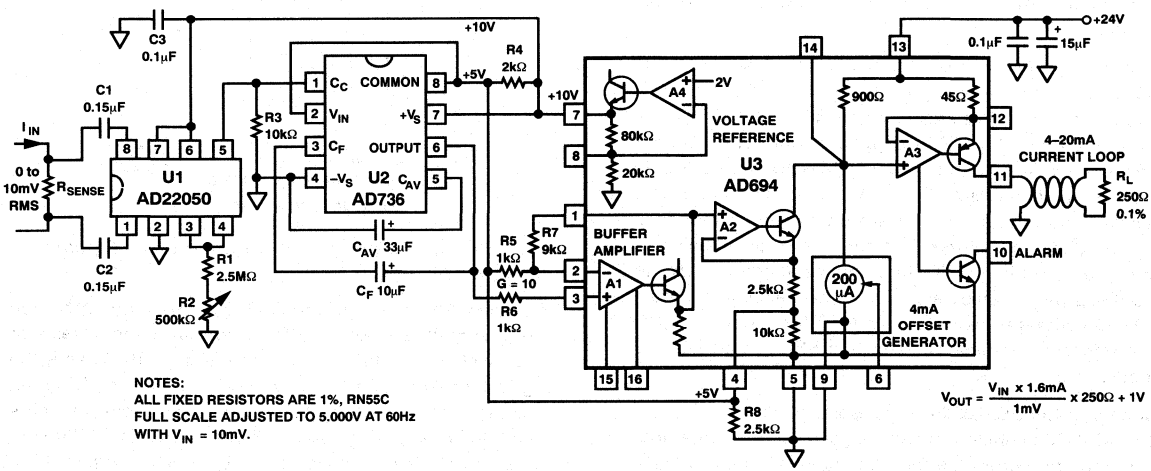

Figure 18. By using an AD694 4-to-20 mA current-loop interface IC, you can build a complete remote-monitoring system that measures true rms current or voltage with just three ICs. The entire circuit operates from a single +24 V supply.

operates from a single supply, you must bias the COMMON (Pin 8) input of the AD736 at 1/2 of the AD694's 10 V output, or 5 V. You can do this by creating a voltage divider at Pin 4 of the AD694 using R5 and R9, which is in parallel with the AD694's internal 10 kΩ resistor.

The AD694's buffer amplifier amplifies the *difference* between the AD736's output at Pin 6 and this 5 V rail—this *difference* ranges from 0 mV to 200 mV dc for a 0 mV to 10 mV rms input and produces a 4-to-20 mA current output from the AD694.

R2 serves as a gain adjustment. R5 and R7 set the gain of the AD694's amplifier A1 to ×10. R7 matches R5 to prevent offsets due to A1's input bias currents. This circuit's accuracy is 1.2% of reading from 20 Hz to 40 Hz and 1% of reading from 40 Hz to 1 kHz. Its −3 dB bandwidth is 33 kHz.

APPLICATIONS OF THE AD637

In the final section of this application note, we'll see how to enhance the accuracy of low level (<100 mV_{RMS}) measurements made with the AD637 rms-to-dc converter and reduce settling time.

Increased Accuracy for Low-Level Measurements

A problem in using all rms-to-dc converters occurs in making measurements of very low level signals. This problem arises due to slew rate limitations in the internal precision rectifier (or absolute value circuit) used to convert the bipolar input signal to a unipolar signal.

One way to circumvent this limitation is to shift the dynamic range of the input signal by amplifying it with a fixed gain stage. Figures 19's circuit uses an AD744 BiFET op amp configured as a fixed gain of ×10 amplifier. The AD744 was chosen because of its low cost, 13 MHz gain-bandwidth product (G = 2), and 75 V/μs slew rate. With a 75 V/μs slew rate, the AD744 can amplify a 10 V peak or 7 V rms signal at frequencies as high as 1.2 MHz without slew rate limiting.

The amplifier has a fixed gain of ×10. A 1 MΩ input resistor supplies the 100 pA maximum input bias current. Because the gain is ×10 and the circuit will be dc coupled to the AD637, an external offset trim is used. To minimize gain error, R4 and R5 were hand selected to set the gain of U1 at G = 1 + (R4 + R5)/R6 = 10.

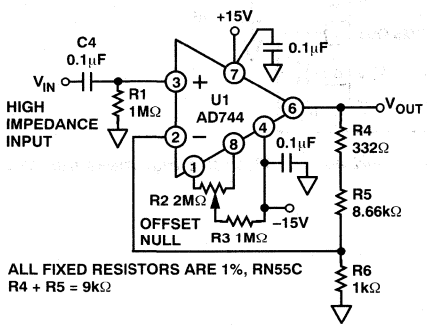

Figure 19. Adding an external preamplifier enhances the accuracy of the AD637 circuit for low level inputs. Here's an AD744 BiFET op amp configured as a gain of 10 amplifier.

A 3-Pole Ripple Filter for the AD637

The usual trade-off in designing any rms circuit is one of settling time versus accuracy and minimum output ripple. One way of reducing settling time is to use as small a value of C_{AV} as is practical while using a multipole output filter to reduce the residual ripple. Figure 20 shows how to construct a 5 Hz, 3-pole Bessel filter with constant phase using the AD637's uncommitted unity-gain buffer. Figures 21 and 22 show simulation results for the filter.

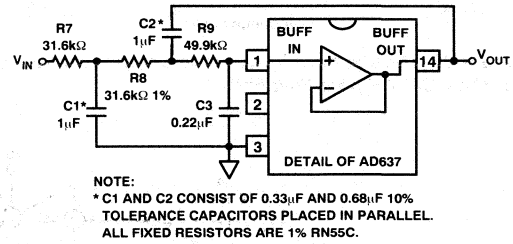

Figure 20. You can build a 3-pole Bessel filter using the AD637's on-chip buffer.

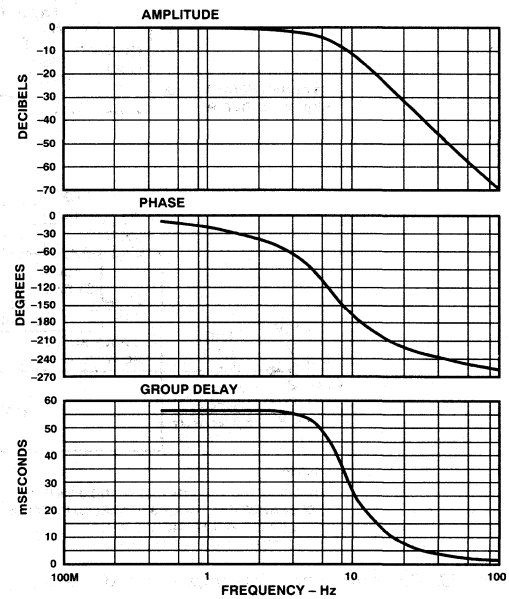

Figure 21. These simulation results show amplitude, phase, and delay for the AD637 ripple post filter.

17

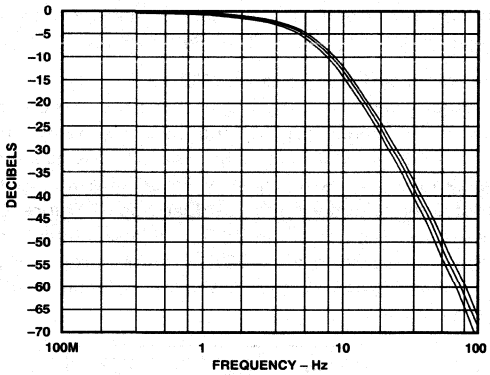

Figure 22. Results from Varying the Component Values for the AD637 Ripple Post Filter

High Accuracy AD637 Circuit for Low Level (<100 mV RMS) Measurements

Figure 23 shows how to use the preamplifier and 3-pole ripple post filter to enhance AD637 accuracy. With the preamplifier and the offsets nulled at V_{IN} = 10 mV at 1 kHz, the circuit's error is less than 0.5% of reading for inputs from 5 mV rms to 500 mV rms for frequencies from 40 Hz to 20 kHz. The 1% bandwidth of the AD744 preamplifier by itself (measured using a Fluke 931B RMS differential voltmeter) is 81 kHz at 10 mV input.

REFERENCES

Kitchin, Charles, and Counts, Lew, *RMS-to-DC Conversion Application Guide, Second Edition,* Analog Devices Inc., Norwood, MA, 1986.

Sheingold, Daniel H., Editor, *Nonlinear Circuits Handbook, Second Edition,* Analog Devices, Inc., Norwood MA, 1976.

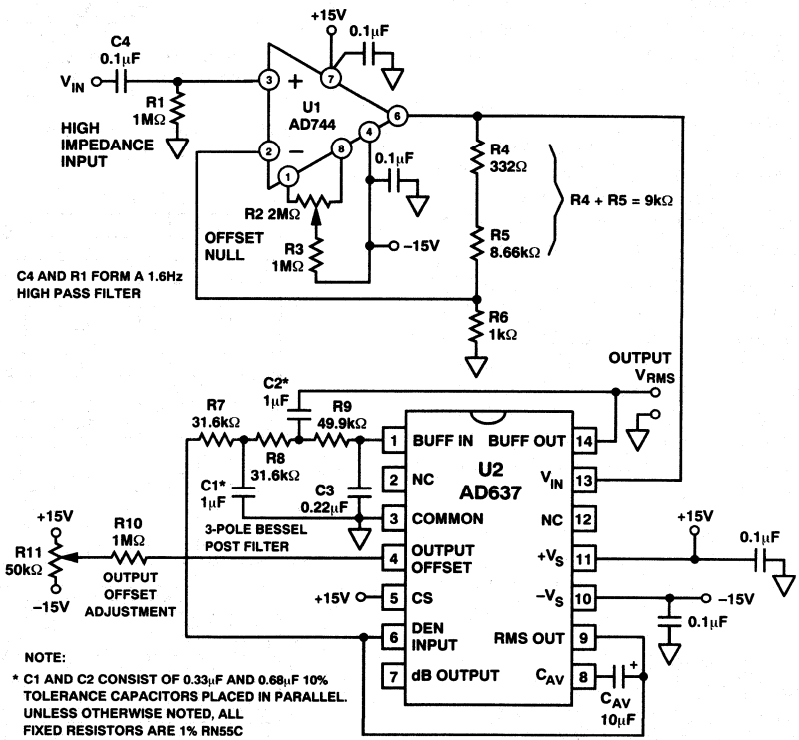

Figure 23. By combining a preamplifier and a 3-pole post filter, you can build an AD637 circuit that measures low level signals with precision and settles quickly for step inputs.

ANALOG DEVICES

AN-269
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Extend RMS-to-DC Converter's Range

by Mark Murphy and Jeff Smith

Comparator hysteresis can extend the dynamic range of ac signal measurement by autoranging an rms-to-dc converter. A circuit can be built with a dynamic range of 80 dB or more, and can resolve input signals from 1 mV to 10 V with less than 1% total error *(Fig. 1)*. The circuit also maintains high accuracy while measuring various input waveforms with different crest factors, including variable duty-cycle pulses. The autoranging method can be utilized in other signal-processing applications as well.

The AD736 rms-to-dc converter is a 1-V, full-scale device. Voltage comparator IC_3 can increase the full-scale input range to 10 V automatically. When the rms-to-dc operating range is exceeded, the comparator's hysteresis switches in an attenuator at the rms-to-dc converter's input. Concurrently, it adds a gain stage to the circuit's output-buffer amplifier, maintaining a steady output voltage.

The comparator has two stable overlapping output states. The lower hysteresis-transition point is set at 90 mV, while the upper point occurs at 1 V. This minimizes the situations in which a varying input waveform produces a confusing or unstable output state. The comparator's logic-high output is typically 4.8 V. Resistor-divider network R_5 and R_6 attenuate the output by a factor of 4.8, setting the noninverting input to 1 V. Similarly, the comparator's logic-low output is typically 0.44 V, setting the noninverting input to about 90 mV.

For input signals within the 100-μV to 1-V-rms operating range of IC_1, the state of the switches is such that the rms-to-dc converteris driven directly from the input signal and IC_2 operates as a unity-gain follower buffering the output of IC_1. When the inverting input to the comparator exceeds 1 V rms, its output goes low. The analog switches are active low, attenuator network R_1 and R_2, and the gain stage consisting of R_3, R_4, and R_g, are switched on.

R_1 and R_2 attenuate the input signal by a factor of 10. This means that inputs between 1 and 10 V now fall within the normal operating range of the rms-to-dc converter. The AD736's output is then amplified by 10 and the true rms value of the input signal is retained. When the output signal drops below 900 mV rms, the output of IC_1 falls below 90 mV, coming within the typical range of the

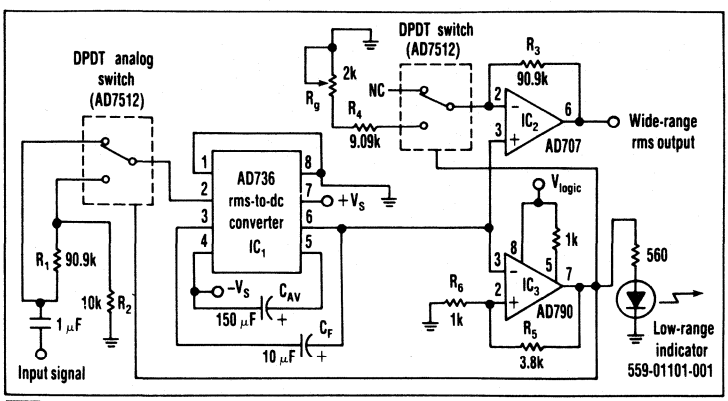

1. THE DYNAMIC RANGE of ac signal measurement can be extended by autoranging an AD736 rms-to-dc converter through comparator hysteresis.

17

Reprinted from Electronic Design – July 12, 1990

rms-to-dc converter again. Then, the comparator's output goes high, the input attenuator is bypassed, and the output-buffer amp functions as a voltage follower.

Due to the circuit's attenuation, the frequency response for 2 V acts similar to that of 200 mV, and 10 V acts like 1 V *(Fig. 2)*. This helps to improve the frequency response of large input signals. The low frequency cutoff (–3 dB) is 20 Hz. With the capacitor values shown, the rms converter's settling time is 360 ms for a 1-V symmetrical sine-wave input; the settling time is greater for input signals of decreasing amplitude. Settling time improves to 36 ms when $C_{av} = 15\ \mu F$ and $C_f = 1\ \mu F$. The low-frequency cutoff rises 200 Hz.

All measurements shown in the figure are for sine-wave inputs, and all voltage specifications are peak values. The circuit is calibrated using a sine-wave input greater than 1 V rms. The relationship between peak and rms value is given by:

The rms value of an untrimmed sine wave = $V_{PEAK}/\sqrt{2}$.

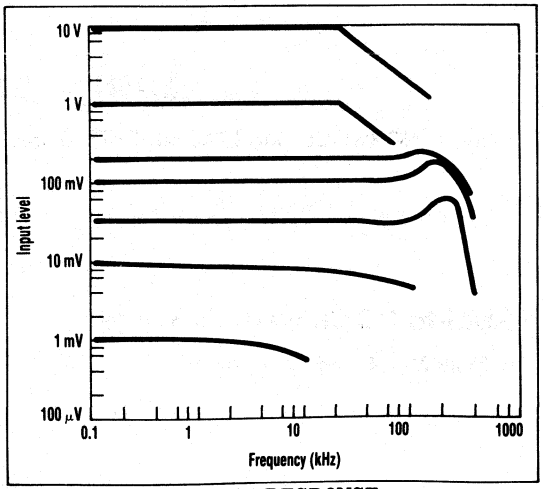

2. THE FREQUENCY RESPONSE of large input signals improves because the signals are attenuated by a factor of 10. For example, the frequency response for 10-V inputs is the same as for 1-V signals.

Sample/Track-Hold Amplifiers
Contents

		Page
AN-270	Applying IC Sample-Hold Amplifiers	18–3
AN-271	Build Precise S/H Amps for Fast 12-Bit ADCs	18–11
AN-27	Polarity Programmable Peak Detector	18–17
AN-204	Applications of the SMP-04 and the SMP-08/SMP-18, Quad and Octal Sample-and-Hold Amplifiers	18–19

ANALOG DEVICES

AN-270
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Applying IC Sample-Hold Amplifiers

by Walt Jung

The sample-and-hold (S/H) function is one which is basic to the data acquisition and A/D conversion processes. A S/H amplifier circuit has two basic and distinct operational states. In one state, an input signal is sampled and simultaneously transmitted to the output (SAMPLE). In the second, the last value sampled is held (HOLD) until the input is sampled again. In most applications, the S/H is used as a "front end" to an A/D converter in data acquisition systems. Used as such, it is the purpose of the S/H to maintain the analog input voltage at a constant level for the period of time required to perform an A/D conversion.

More specifically, the S/H is a system functional block necessary in data conversion systems, where the A/D converter in use requires a constant and accurate analog input during the period of its conversion process. An example of such a use is with a successive approximation type A/D. Ideally, the S/H "freezes" the last instantaneous input voltage prior to its HOLD command, and presents this voltage unaltered to the A/D converter. The A/D then converts the voltage to a corresponding digital word. In practice, there are many error factors which come into play in the process of implementing a S/H. Thus this application note will discuss these basic considerations, as well as representative device topologies and some representative applications.

BASIC OPERATION OF S/H AMPLIFIERS

Some S/H basics are illustrated by Figure 1, with the most basic form of S/H circuit operation diagrammed in 1a. Here, an analog input signal to be digitized is applied directly to the electronic switch, S1. Depending on the state of S1, the signal will be either transmitted to the holding capacitor, C_H, or it will be blocked. The state of switch S1 is controlled by the S/H control line, which is a digital input.

When S1 is closed, the input signal appearing across C_H is buffered by A1. It also appears at the S/H output (possible low-pass filter effects are disregarded for this discussion). If S1 is connected for a period of time as the input varies (as with an ac waveform), the operation may also be said to be *tracking*; that is, any input changes are also transmitted to the output.

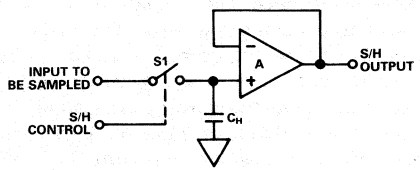

a. A Basic S/H Circuit, Consisting of Switch, Hold Capacitor, and Buffer Amplifier

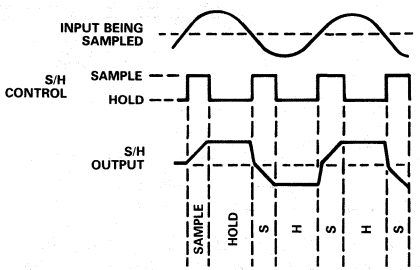

b. S/H Waveforms Showing the Input Being Sampled (Top), the S/H Control (Middle), and the S/H Output (Bottom)

Figure 1. S/H Basics

When S1 is opened, the last input voltage value is retained on C_H as a charge; that is, it is held. A1 continues to read this voltage until the next SAMPLE period. This action is illustrated by the input, output, and control waveforms in Figure 1b. S/H circuits are used for a wide variety of signal-processing functions; not only A/D interfaces, but also more general analog memory functions such as auto-zero amplifiers can be implemented.

The S/H operational waveforms of Figure 1b are nearly ideal and assume perfect switching, perfect tracking, ideal holding characteristics, and load/source immunity. In practice, S/H errors are found to exist for each of the four states of the device. These states are:

(1) HOLD-to-SAMPLE transition
(2) SAMPLE interval
(3) SAMPLE-to-HOLD transition
(4) HOLD interval

These errors are obviously potentially important to many applications, and for high accuracy applications (>10 bits, or accuracies of 0.1% or less), all are important. They are defined and illustrated below.

HOLD-TO-SAMPLE TRANSITION ERRORS

These errors are associated with the time interval when the device is switched from the HOLD state to the SAMPLE state. Since the input may have changed a large amount since the last sampled voltage (i.e., by as much as full scale), the S/H must reacquire the input signal and settle once again to within its rated accuracy band. This is shown in Figure 2.

Acquisition time is the time required for the S/H to acquire and then track the input signal after the SAMPLE command. It is usually specified for a full scale level change (−10V to +10V and vice versa), since this represents the worst case in terms of time necessary to acquire an arbitrary level signal. The output must assume the desired level to within a rated error band consistent with the level of accuracy required for the conversion or sample. For example, this may be 0.01% or 0.1%. An illustration of a HOLD-to-SAMPLE acquisition waveform is shown in Figure 2a.

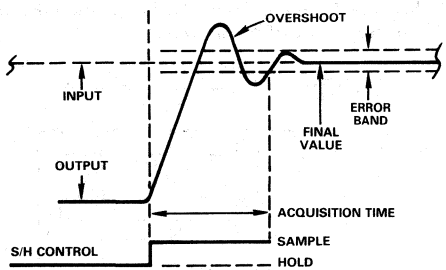

Figure 2a. HOLD-to-SAMPLE Node Acquisition Time Showing Acquisition of a New Signal (Top), and the S/H Control (Bottom)

A major portion of acquisition time for large HOLD-to-SAMPLE changes will be taken up by an initial slewing interval. After this high error interval, the output may overshoot, whereupon it will settle to within a rated accuracy band of ±2mV, which would be ±0.01% for a 20V scale for example. Note that acquisition time ends when the signal has settled and *remains* within the rated error band.

Acquisition time is the major HOLD-to-SAMPLE error component and is the primary determinant of how fast the S/H portion of a conversion system can be operated. Typical times are on the order of several microseconds to accuracies of 0.1% or 0.01% or better. Acquisition time is strongly dependent upon the value of the holding capacitor used as this capacitance (usually) influences slew rate.

Illustrated in Figure 2b is the **HOLD-to-SAMPLE transient**, a switching transient produced at the time of the transfer from HOLD into SAMPLE mode. Note that this transient will be present even in the case where there is no large difference between the previously held voltage and the new

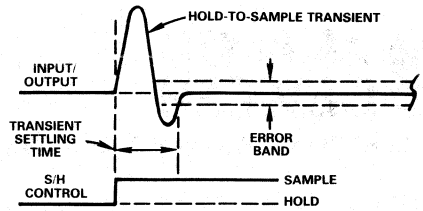

Figure 2b. HOLD-to-SAMPLE Mode Transient, and Settling Time

sample. Since the amplitude of this transient may be well in excess of the rated S/H accuracy (as much as several hundred millivolts), it must be allowed a sufficient time period to die out before the output voltage sample can be considered valid.

Since the settling time of this transient continues beyond the initiation of the HOLD-to-SAMPLE command, system timing must allow for this. However, in practice the settling time associated with any HOLD-to-SAMPLE transients will typically be much less than that of acquisition time. Thus, a time interval equal to the worst case (or acquisition time) will usually take into account the HOLD-to-SAMPLE transient error and its associated settling time automatically.

SAMPLE ERRORS

During the SAMPLE interval, the S/H device is tracking the input signal in a fashion which is very much like an op amp. In fact, most S/H devices are either specialized op amps or they are constructed using op amps with characteristics which are particularly suited to S/H use. Therefore, since most S/H amplifiers reduce to, or are equivalent to either an op amp voltage follower or an inverter, their SAMPLE mode errors can be calculated in a similar fashion.

Pure scaling errors within a S/H can generally be regarded as a benign error as they by and large can be easily nulled with a calibration adjustment. Usually, a convenient point for this to be done is at the A/D reference as this can remove all of the system scaling errors at once. Of course, this approach applies to the conventional usage, i.e., one S/H per A/D. If a number of S/H's are used in front of an A/D or if a S/H is just being used as part of another circuit, it may need to be gain trimmed locally via gain scaling resistors.

In any event, the worst case deviation from the ideal S/H scale factor must be known for use in the error budget. Typically, the scale factor will be 1 ± an error of 0.001% or less. That is, the type of gain error associated with a voltage follower hookup.

In the case where gain resistors are used, such as when the S/H is to be used for noninverting gains other than unity, resistor tolerances will increase this error substantially. For inverting mode operation S/H's, gain scaling resistors must be used, regardless. In either case, it is highly advantageous to have application resistors on the S/H chip, since they will have higher pretrimmed accuracies and a specification for their worst case drifts. The trend of

recent devices is towards sets of pretrimmed resistors to accommodate popular gains of −1, +2, etc.

Note that while a true gain (scaling) error can be adjusted by a system scale calibration, a gain *nonlinearity* is a non-trimmable error.

Gain nonlinearity is a critical S/H error and it appears as a deviation from the ideal transfer characteristic. This component is the dynamic deviation from the ideal numeric S/H gain (i.e., +1, +2, −1, −2, etc.) as the unit is exercised over its rated signal output range, usually ±10 volts. Its largest component is usually input stage common-mode error which will be typical of the case of a follower type hookup (by and large the most prevalent). In the case of an inverter type hookup, the common-mode errors disappear, but resistor matching errors can become an error source.

Typical S/H nonlinearity figures are 0.001% to 0.01% over the ±10V signal range. Obviously, the S/H nonlinearity must be better than the overall nonlinearity as established by the A/D in use to preserve system performance. A good rule of thumb for S/H nonlinearity is a figure of an order of magnitude better than the basic converter resolution. For example, a S/H nonlinearity of 0.01% or better would be used with a 10-bit converter. Note that the user may need to calculate nonlinearity from the S/H's CMRR, such as an 80dB CMRR, equivalent to 0.01% nonlinearity.

Offset is the dc shift which occurs between the input and the output with the S/H input grounded. It is usually adjustable to zero via an optional trimpot. A typical dc offset specification will be on the order of ±2mV or less. Pure offset by itself is not usually a major problem with S/H applications, as it can always be trimmed to zero as part of the overall system calibration. This can be part of the A/D trim calibration, either done manually or via software.

Offset temperature *drift* is another story, since it cannot readily be distinguished from real signal(s). Unless an auto-zero calibration cycle is included, S/H offset drift will be a nonreducible error component and will cause errors as temperature varies. Drifts for S/H's are typically on the order of 1 to 10μV/°C, so this error can be a serious one either for the higher accuracies or for wide temperature ranges.

S/H offset voltage will also vary with supply voltage to some degree, and this should also be specified. Typically, supply rejection is on the order of 80dB or 100μV/V. This parameter is usually not of major importance with well regulated supplies or with the use of auto-cal cycles.

Settling time is applicable to the SAMPLE mode for rapid input voltage changes. When tracking an input signal, a S/H is governed by dynamic limitations similar to other op amp configurations.

Settling time is governed by slew rate and small signal bandwidth, with slew rate being dominant for large scale step changes. Typical slew rates are 5 to 10V/μs, and settling times in the 5-10μs region. As noted under acquisition time, the exact specifications are strongly dependent upon the holding capacitor.

SAMPLE-TO-HOLD TRANSITION ERRORS

Aperture time or aperture delay is the time which elapses between the point when a HOLD command is issued and the actual opening of the S/H switch. With a rapidly changing input voltage, this time will introduce an error by determining the voltage which is actually held. The resulting voltage error will be equal to the change in input voltage which occurs during the effective aperture time interval.

Figure 3 generally illustrates aperture time related errors, of which Figure 3a shows how aperture delay creates an error in the voltage held. With a high rate-of-change input voltage, the S/H voltage will be changing by an amount which may be approaching 1/2LSB during the time the switch is passing from on to off.

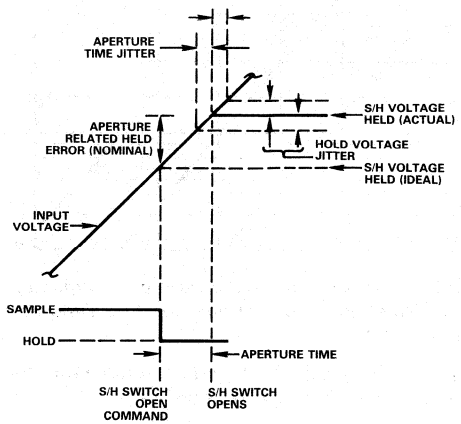

Figure 3a. Aperture Related Time/Voltage Errors. Analog Input/Output at Top, SAMPLE/HOLD Drive at Bottom.

As a general example of the effect of aperture time, consider an input signal with a rate of change (signal slope) of 1V/μs which is sampled with a 10ns aperture time. This will result in a 10mV sampled error due to the aperture time related dV/dt error.

This sort of error is usually significant. The effective aperture delay can be compensated by advancing the in system HOLD command timing the amount of nominal aperture delay, but this is not the entire picture.

If the nominal aperture delay is subtracted, the remaining error is known as the aperture *jitter* (or uncertainty) which is the true limit to S/H sampling errors with high signal slope inputs. Aperture jitter is defined as the net variation in the actual S/H switch timing from sample to sample. This jitter places the ultimate limit on aperture timing related error. For the 1V/μs slew rate example, a 1ns aperture jitter would result in a ±1mV voltage uncertainty.

A general relationship showing the limiting aperture time and the resulting allowable full scale level sine-wave input frequency can be drawn. This is shown in Figure 3b. This graph is based upon the maximum (full scale) sine-wave input frequency which will result in no more than 1/2LSB of error. This frequency, f_{max}, is calculated as:

$$f_{max} = 1/[2^{(n+1)} \pi t_a] \quad (1)$$

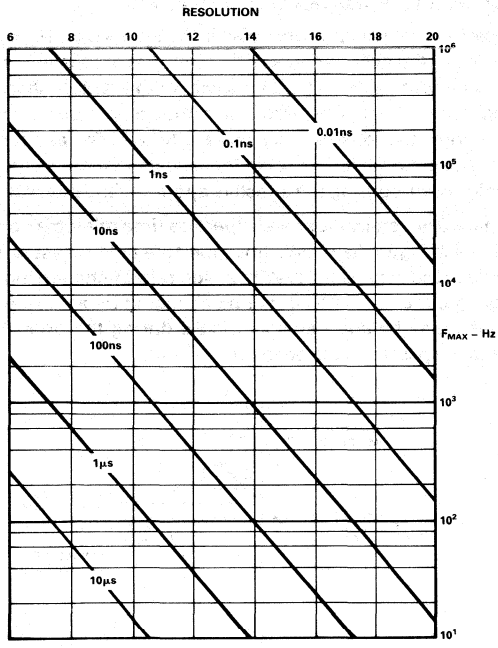

Figure 3b. Maximum Input Sine Wave Frequency (f_{max}) at Full Scale Level, for Various Aperture Times and Resolution(s)

where $\pi = 3.14$, t_a is the limiting aperture time, and "n" is the resolution of the converter in bits.

Since the data is for a limiting aperture case, it can be used for an A/D with a S/H or for an A/D operating alone. In the latter case, the A/D's conversion time will define the effective aperture time.

The data clearly illustrates the value of a S/H for maximizing the allowable input frequency. A 10µs 8-bit A/D without a S/H can only accept a maximum input frequency of roughly 60Hz. On the other hand, with the use of only a 100ns aperture S/H, the same A/D can accept a 6kHz maximum frequency. Note also a more general relationship; as the conversion resolution is increased, f_{max} is lowered for a given conversion time. Therefore, the need for a S/H becomes more critical as either resolution or frequency is increased.

Obviously, the lower the aperture time in a S/H, the better, as it will place less of a limit on the A/D used with it. While the times quoted for illustration are typical for medium speed I.C. S/H circuits, they are not the final limit to system timing since the maximum *throughput frequency* will often occur before the aperture time limited frequency.

S/H offset (also called S/H "pedestal", "jump" or "step") is the resulting analog error induced by a transient charge on the HOLD capacitor when the S/H switches from SAMPLE-to-HOLD. It is caused by the finite capacitance of the S/H switches used, and to a lesser extent, the layout and/or package capacitances. These capacitances feed through a portion of the digital control signal directly into the HOLD capacitor. In general, this error can be reduced with larger value HOLD capacitors since the parasitic coupling capacitance is fixed within a given device and layout. The effect is illustrated in Figure 4.

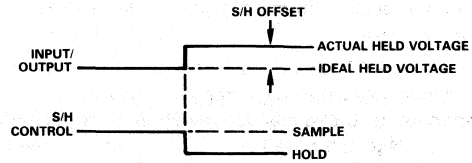

Figure 4. SAMPLE-to-HOLD Offset (Top), and S/H Control (Bottom)

This error, which may be on the order of several mV, can in some cases be compensated for by feeding an anti-phase control signal into the HOLD capacitance via a small external coupling capacitor. Generally, this error is reducible by reducing the p-p level of the digital control signal to an absolute minimum as well as by screening/guarding the coupling paths between this signal and the C_H node.

Note that SAMPLE-to-HOLD offset may not always be directly specified as such, particularly in IC devices which use *external* holding capacitors. In such cases, a specification may be given in terms of *charge*, in units of pC. In these cases, the S/H offset may be then calculated as:

$$\text{S/H offset}_{(V)} = \text{Charge (pC)}/C_H\text{ (pF)} \qquad (2)$$

A 10pC charge transfer with a 1000pF C_H would for example cause a S/H offset of 10mV. Obviously, the lower the charge transferred, the lower this error will be. In better units, charge transfer is as low as 1pC or less.

S/H offset can initially be considered as a trimmable error, just as can pure dc offset. If, however, it should change with time, input voltage, and/or temperature, these instabilities must also be considered. In some S/H units which use floating switches (see Figure 7b, type 2), the S/H offset will change with the value of the input signal. This is obviously not desirable, as it is difficult if not impossible to compensate. S/H type 3 (Figure 7c) does not show this problem, as the switch is maintained at a constant level (a virtual ground).

Sample-to-hold settling time is time for the S/H output to settle to within its rated accuracy following the HOLD command. It includes the time required for any switching transient to die out.

HOLD INTERVAL ERRORS

Droop, also called tilt, is the change in the HOLD voltage (ΔV) during the HOLD interval (Δt). It is caused by the net leakage current(s) into (or from) the HOLD capacitor. Since the currents which cause droop may be of either sign from one S/H to the next, droop current will cause the voltage on the capacitor to ramp either up or down, as shown in Figure 5. However, it is the magnitude of the error over the HOLD interval and not the sign that is important. Droop is governed by the simple charge/voltage relationship:

$$\Delta V/\Delta t = I_L/C_H \qquad (3)$$

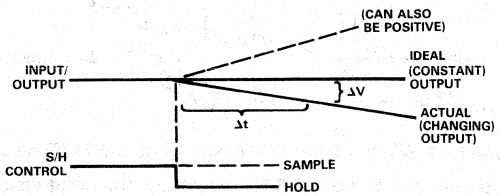

Figure 5. Droop Error Voltage (Top), and S/H Control (Bottom). Note that "Droop" Can Be Either Positive or Negative.

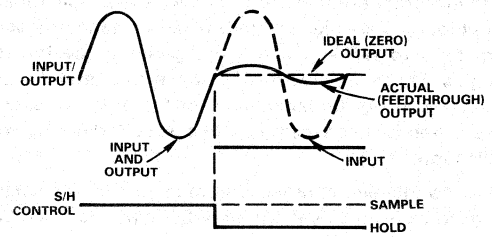

Figure 6. Feedthrough Error Voltage (Top), and S/H Control (Bottom)

where I_L is the leakage current and C_H is the HOLD capacitor value. As an example, a 1nA droop current with a C_H of 100pF will cause a droop rate of 10μV/μs. Over a HOLD period of 10μs, this will yield a max droop error of 100μV which is not likely to be a problem.

Within a given S/H device, the current I_L is given by a specification, but C_H is (usually) under user control. Therefore, droop can be adjusted by C_H and will become smaller as C_H is increased.

There is a tradeoff, however. As C_H is increased to lower droop, this typically results in longer acquisition times. This is because the current needed to charge C_H for acquisition slewing is fixed. As a result, it is usually desirable to hold the leakage current I_L as small as is practical in order to minimize droop. Typically, this parameter will be due to the S/H output amplifier's input bias current and switch leakage which usually increase with temperature. Droop is worst at high temperatures. Note that the value of I_L used to calculate droop should account for this temperature dependence of net leakage current. With FET input buffers, bias current doubles for each 10°C rise.

In practice, typical droop rates at medium temperatures in a S/H will be on the order of 1μV/μs, a fairly insignificant error. This may not necessarily remain so when maximum bias currents at the highest temperature are calculated. The leakage current which will be seen at the highest expected temperature can be calculated in order to determine the worst case droop rate.

Feedthrough is an analog error caused by leakage of ac signals through the S/H switch in the HOLD (off) state. Like S/H offset, it is basically caused by switching capacitance but it can also be influenced by layout dependent capacitive coupling. Since the switch capacitance and the HOLD capacitance form a voltage divider, feedthrough decreases for higher values of HOLD capacitance just like droop and S/H offset. The effect of feedthrough is shown in Figure 6.

Feedthrough is usually specified with a full scale peak-to-peak sine-wave input at a high frequency, such as 10kHz. Typically it will be on the order of 80dB (or more) in a good S/H. This is equivalent to a feedthrough error of 0.01% (or less). This parameter is important in cases where the S/H follows a multiplexer which can select high level inputs during the HOLD state and potentially cause errors due to the signal feedthrough.

DIELECTRIC ABSORPTION

With some common capacitor types, the dielectric does not completely release all of its energy after a charge/discharge cycle. The result of this phenomenon is an error in stored voltage after a period of time in the HOLD mode. This effect is minimized using certain dielectrics, in particular, the films. Dielectrics specified for low DA, such as Teflon, polystyrene, and poly propylene should be used. Note that consideration of DA will be most applicable in S/H's which either use an external C_H or where an external C_H is used to augment an internal one.

DRIFT AND NOISE

A S/H can possess different drift characteristics in the HOLD mode than those in the SAMPLE mode. In the HOLD mode, the output terminal sees only the drift of the *output* buffer amplifier. In the SAMPLE mode, it sees either the input amplifier alone or the composite drift of two series amplifiers.

Ordinarily, this different drift characteristic for HOLD versus SAMPLE is not detrimental. A conversion is usually made within a few microseconds, long before drift errors due to temperature can be a problem. For unusually long HOLD times, it may need to be considered.

Noise, however, can be a different story. Consider the S/H types 2 and 3 of Figure 7, for example, when in the HOLD mode. If the output amplifier has excessive noise, it will be seen during HOLD and digitized along with the desired signal. If this noise is high and the converter linearity is not well below 1/2LSB, certain codes will be in error due to noise modulation.

MOSFET input buffers are used in some S/H's for their very low input currents. Unfortunately, these types often have relatively high input voltage noise which can limit overall accuracy for some applications. JFET inputs will not have as low an input current as MOSFETs, but will have appreciably lower noise voltages.

S/H DESIGN TYPES

There are great numbers of S/H devices available today. Except for the very highest limits of performance, most of them tend to fall into a few categories of design topologies. These are illustrated below with a brief discussion on each. An understanding of the different circuit types is helpful in choosing a S/H best suited for a given application.

Similar to the elementary circuit of Figure 1a, the circuit of Figure 7a is an actual S/H amplifier type. Here A1, C_H, and the switch operate as described before, but an input stage buffer A 2 is also added.

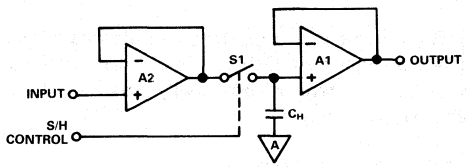

Figure 7a. Open Loop, Cascaded Follower S/H (Type 1)

The addition of the buffer provides increased charging current for C_H. This provides faster acquisition time without loading the source. In this S/H type, both amplifiers must have high slew rates, fast settling times, low offset voltages and low drift for best accuracy, since these errors are cumulative. A1 should be a FET input device to minimize loading of C_H, but this is not necessarily true for A2. This S/H type tends to be good for high-speed acquisition uses.

The S/H configuration of Figure 7b has the advantage that an overall feedback loop is returned around both amplifiers in the SAMPLE mode. As a result, the errors of A1 are minimized for the SAMPLE state, although they do appear in the HOLD mode. This circuit has potentially higher accuracy, but with some potential sacrifice in overall settling characteristic due to the multistage loop dynamics.

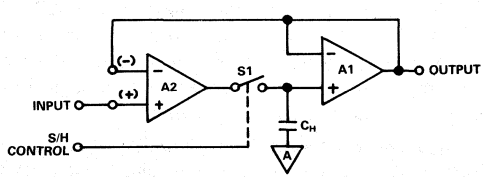

Figure 7b. Closed Loop, Follower Output S/H (Type 2)

Note that in this type of S/H connection, the input stage can be wired to have both the (+) and (−) input terminals available for external options, although this may not always be done in all types. With this flexibility, the S/H can be connected as either an (overall) inverting or noninverting type S/H. Application resistors may or may not be made available.

The third circuit in Figure 7c is similar in that it connects feedback around the two op amps and in this sense has advantages similar to 7b. In this case, the switch is operated at the virtual ground input of A1 and C_H is an integrator capacitor around the output amplifier stage.

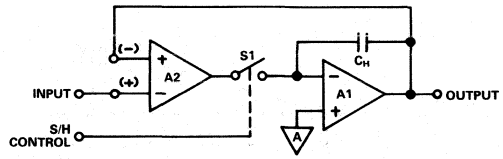

Figure 7c. Closed Loop, Integrator Output S/H (Type 3)

Note that in this type of S/H connection, the input stage can also have both the (+) and (−) input terminals available for use as previously noted. Therefore, this type can be connected as either an (overall) inverting or noninverting type S/H.

S/H APPLICATIONS

Implementing a conversion system with modern A/D and S/H devices, while never trivial, can be greatly aided by the applications versatility of the hardware employed. The 12-bit stand-alone A/D + S/H conversion system of Figure 8a is an example of a hookup which is easy to implement as shown but can also be easily modified for other scales, etc.

This circuit uses an AD585 S/H, connected in the noninverting unity gain mode, with $\overline{\text{HOLD}}$ active. The A/D is an AD574A connected for a unipolar 0 to +10V scale, with system GAIN and OFFSET calibration set by R_2 and R_1, respectively. The STATUS output of the AD574 drives the $\overline{\text{HOLD}}$ input of the AD585 A1, for the lowest possible S/H offset. This requires inversion in the TTL stage.

As shown, conversion is started by a CONVERT signal and begins on the falling edge, whereupon the STATUS line goes HIGH and the S/H goes to HOLD. For a 12-bit conversion, the AD574A will require 35 (max) μs to perform a conversion. The end of conversion is signaled by a STATUS LOW signal.

The A/D calibration trims shown provide for change of system gain and offset, sufficient to accommodate both S/H and A/D errors.

The maximum throughput time of this system is based on three factors, as it is represented here. These are: the A/D conversion time and the S/H aperture delay and acquisition time. The conversion time will be:

$$T = t_c + t_{ad} + t_{acq} \quad (4)$$

Respectively, these individual times will be 35μs + 35ns + 3μs, or 38.035μs, total. Maximum throughput frequency can be calculated, as the inverse of this time, as:

$$f_{throughput} = 1/T \quad (5)$$
$$= 26.3\text{kHz}$$

This frequency assumes a *single* sample per cycle; to meet the Nyquist criteria of 2 data points per cycle, the frequency would be halved. (Note also that any additional settling time period preceding the S/H also may need to be accounted for and can include an IA and/or a multiplexer, when used.)

A source of potential error in the application of high-speed successive approximation A/D's with S/H amplifiers is a *kickback error* from the A/D. A successive

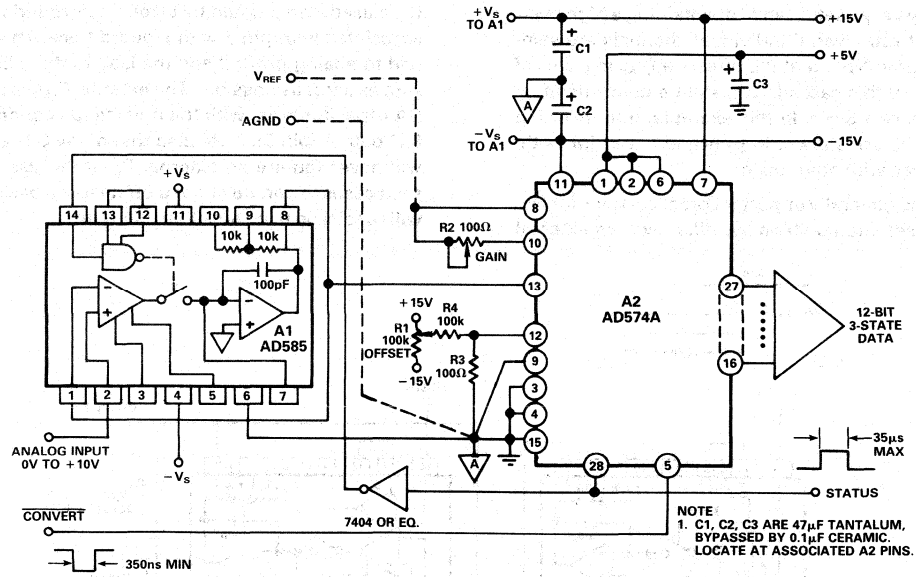

Figure 8a. Stand-Alone 12-Bit Unipolar A/D with S/H

approximation A/D represents a *dynamic* load to the S/H, and the MSB loading at the start of the conversion cycle can cause a transient at the A/D input voltage (S/H output) due to nonzero S/H output impedance. If the S/H impedance is not sufficiently low at high frequencies, the kickback error can exceed an LSB.

The solution to this particular problem lies in the use of a S/H with an intrinsically low impedance output stage or one with high feedback, for low kickback error. With the impedance sufficiently low to render the dynamic voltage error negligible, A/D loading will not cause a problem. The AD585 output stage has been designed for a low dynamic output impedance to minimize this error. As a demonstration of this, an expanded scale photo of an AD585 output driving an AD574A during conversion is shown in Figure 8b. Note that the negative going transient error at the MSB switch point is both low and short in its duration.

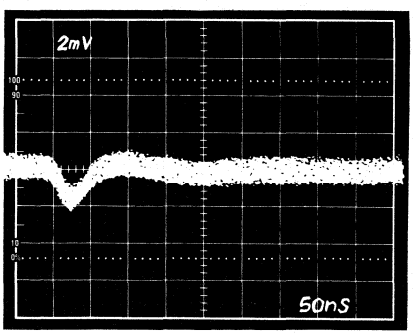

Figure 8b. Output of AD585 S/H Amplifier Driving AD574A, MSB Transition

For applications which demand both a fast acquisition time and a low droop rate, these conflicting performance parameters can be met together using a cascaded S/H. Figure 9 shows such a circuit with the $\overline{HOLD}_1$ and $\overline{HOLD}_2$ control lines driven as shown in the timing diagram. The basic idea is that the first S/H of the cascade acquires the input rapidly before the second has time to settle to its rated accuracy. The first S/H is then put into HOLD and the second one continues to acquire the "input" as it appears at the *output* of the first (fast) S/H. Since this constitutes a series path for the input signals, the errors of the two S/H's are additive.

This circuit uses two series connected AD585's with the first one configured normally for fast acquisition. Using the internal capacitor of the AD585, this stage will have a 1mV/ms (max) droop rate. This stage samples for a 5µs period (the width of the $\overline{HOLD}_1$ signal). The second S/H samples the output of the first for 500µs or the width of $\overline{HOLD}_2$. During this 500µs period, the output of the first stage will droop, up to 0.5mV. This is generally the $HOLD_2$ width (in ms), times 1mV/ms.

The second AD585 uses a 0.01µF external C_H, which will minimize droop of this stage by a factor of 100 (the ratio of the second to first stage C_H's). The effective droop of the entire circuit then becomes 0.5mV (first stage), plus 0.01mV/ms (second stage). For HOLD intervals in the second stage of up to tens of ms, the net droop of this cascaded S/H will therefore be close to that seen in the 500µs interval, or in this case, ≈0.5mV.

Of course, the tradeoff of this scheme is that the overall signal throughput is reduced. Practically speaking, there may be little or no penalty, as the application may be intended to be used with slower A/D's. Also, if a number of cascaded S/H's are being used before a multiplexer, this

scheme can be very useful, while maintaining a high overall throughput rate even if individual channels are sampled infrequently. Note that the effective f_{max} in terms of aperture time of this cascade will still be determined by the *first* of the two S/H's. In this example, non-inverting unity gain S/H's are shown. In principle either (or both) S/H may be used with other scales.

There are some generally important practical rules for layout that the user should abide by. Whenever an external C_H is used, care should be taken to surround the PC trace associated with pin 7 with a guard trace. This should be tied to analog ground and the lead to this pin should be kept as short as possible. The external C_H should be a low DA capacitor type, with the outside foil connected to the S/H output (pin 8). Note also that many S/H applications will never require an external C_H or the use of pin 7. In such cases, it can be clipped off near the package, which will cut pickup to an absolute minimum.

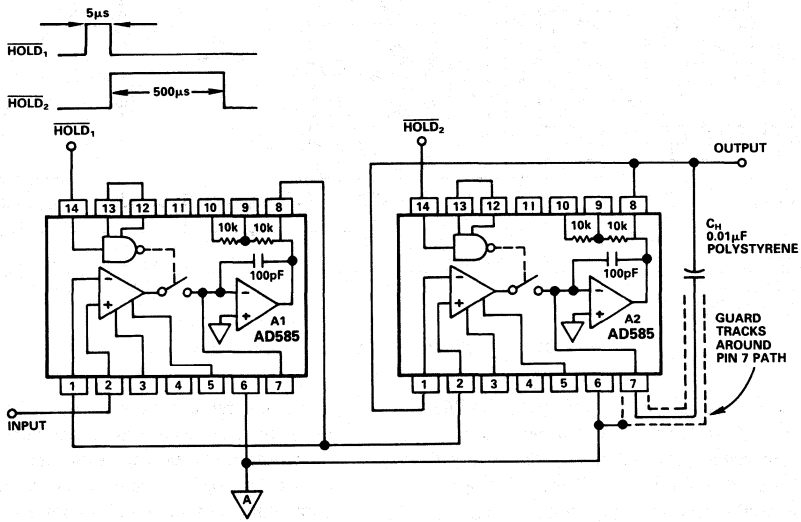

Figure 9. Loop Droop Cascade S/H

ANALOG DEVICES

AN-271
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Build Precise S/H Amps for Fast 12-Bit ADCs

by John Sylvan

TODAY'S FAST, PRECISION IC OP AMPS AND FET SWITCHES BUILD THE LOW-COST SAMPLE-AND-HOLD AMPLIFIERS NEEDED BY SUB-1 μS, 12-BIT A-D CONVERTERS.

Discrete sample-and-hold amplifiers place heavy demands on operational amplifiers. The op amps must have excellent ac performance and promise good dc specifications. While sampling, they must track the input signal without distortion—a demand for wide, full-power bandwidth. When switching from the hold to sampling mode they should acquire the signal quickly. Here, high slew rates and fast settling times are vital. To ensure 12-bit precision over a wide operating temperature range, dc specifications such as offset voltage and offset voltage drift must be low.

Low input-bias current is a mandatory requirement for an op amp to be used in a s/h amplifier. High bias current leads to excessive droop rates. Therefore, s/h amplifier circuits require FET-input amplifiers. Typical FET designs have delivered settling times of 1 μs to 0.01%. Newer complementary bipolar (CB) processes now used to manufacture FET op amps combine low picoampere bias currents with 35 MHz unity-gain bandwidths and 12-bit accurate settling times under 130 ns. Moreover, they work from standard ±15-V supplies, easily accommodating wide dynamic-range signals.

S/h amplifier circuits will use one of several standard architectures: closed-loop or open-loop, and inverting or noninverting configurations. No matter what circuit type designers employ, they need to incorporate a high-speed amplifier to buffer the input signal and a switch to connect and disconnect the input signal from the hold capacitor. In addition, a FET-input amplifier is needed to buffer the voltage on the hold capacitor and drive an a-d converter.

FAST, SIMPLE SAMPLING

Three low-cost CMOS switches, two CB op amps, and a few resistors and capacitors are all that's needed to build a fast, simple, low-cost s/h amplifier from discrete IC op amps

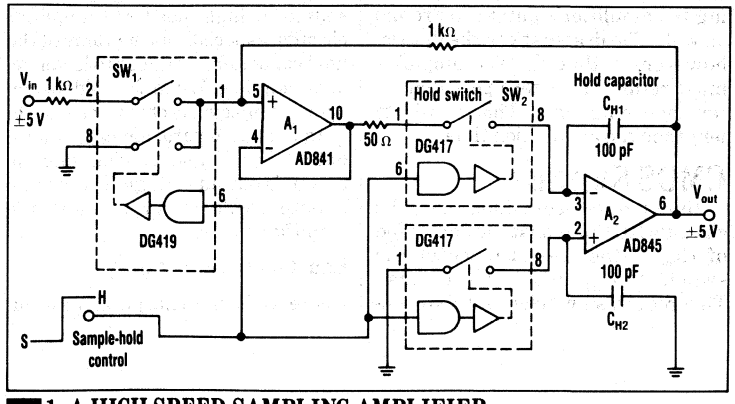

1. A HIGH-SPEED SAMPLING AMPLIFIER to be used with a 12-bit, 1-ms, a-d converter can be built with a pair of high-speed CB op amps from Analog Devices, three CMOS IC analog switches, and a handful of passive parts.

Reprinted from Electronic Design – January 25, 1990

(Fig. 1). The circuit's full-power bandwidth is in excess of 250 kHz, acquiring a ±5-V input signal in approximately 1 μs—and it can be built for about $15.

Operating as a unity-gain voltage follower, op amp A_1 buffers the input signal and supplies a minimum of 50 mA to charge the 100 pF hold capacitor, C_{H1}. Most older, fast, IC op amps don't have nearly this output current capacity—it's needed to charge a capacitor at high speed. In addition, many fast op amps would be unstable driving even 100 pF. This buffer, an AD841, is unity-gain stable and its full power bandwidth is close to 5 MHz.

Op amp A_2 does the sampling. It keeps the sampled charge on the hold capacitor and drives the next link in the system, for example the input of a high-speed, high-resolution a-d converter. A_2, an AD845 FET-input op amp, has an input bias current of just 500 pA, slews at 100 V/μs, and settles to within 0.01% of final value in under 250 ns for ±5-V output steps.

The circuit's overall design employs a noninverting, closed-loop configuration. Closed-loop designs typically trade off better accuracy and linearity for slower speed. In most closed-loop designs, the input amplifier usually saturates when it's disconnected from the second stage. This occurs because the feedback loop is broken.

But this circuit overcomes that limitation by having a separate feedback path around A_1 and a double-pole, single-throw switch at the amplifier's input. When the circuit is in the "hold" mode, the noninverting input of A_1 switches to ground, driving the amplifier's output to ground as well. The drawback to the design, however, is that the s/h amplifier must return from zero potential after each hold time. This serves to lengthen the acquisition time.

CMOS Switches

To minimize acquisition time, the circuit's RC time constant, composed of (R_{in} + the on-resistance of the switch) × C_{H1}, should be minimal. This requires switches with low on-resistance. New generations of fast CMOS switches, such as the DG417 and DG419 from Siliconix, Santa Clara, Calif., offer on-resistances of 35 Ω and turn off in just 200 ns. Moreover, these high-speed ICs operate from ±15-V op amp supplies readily accommodating a wide-range of input signals.

The choice of switch is a trade-off between circuit complexity and performance. Complete CMOS switches operate over a wide voltage range, but they typically have higher charge injection than discrete FETs, which leads to larger pedestal errors. They also turn off and turn on slower than discrete FETs. However, CMOS switches operate directly from TTL or CMOS logic signals, eliminating the need for extra level-shifting circuitry.

ALTHOUGH REQUIRING MORE DESIGN WORK, A DISCRETE APPROACH MAY BE A COST-EFFECTIVE ALTERNATIVE TO BUYING AN EXPENSIVE HYBRID.

Pedestal-error is a major error source in high-speed s/h amplifier circuits, especially as the value of the hold capacitor is reduced to lower acquisition time. When the hold-switch opens, it injects the charge held on its gate-to-source parasitic capacitance onto the hold capacitor changing the op amp's output voltage. The change in output voltage, ΔV_O, seen by the amplifier equals:

Equation 1: $\Delta V_O = \Delta Q / C_H$

where Q is the charge in picocoulombs and C_H is the value of the hold capacitor, typically in picofarads. The effects of charge injection are significant. For example, injecting just 5 pC of charge into a hold capacitance of 86 pF causes a change in output voltage of 58 mV—a 24 LSB error for a 12-bit a-d converter with a ±5-V full-scale input range.

To reduce pedestal error, a s/h amplifier can employ a differential mode of operation *(Fig. 1 again)*. When the circuit switches from sampling or tracking the input signal to the hold mode, a second hold switch (SW_2) injects an equal charge onto C_{H2}, balancing the effects of charge injection on C_{H1}. As a result, the pedestal error will be zero if the switches and capacitors exactly match and if the amplifier has high common-mode rejection. The pedestal error is seen as a common-mode error by the op amp and is rejected.

Driving a virtual ground with the hold-switch (that is with the storage capacitor C_{H1} in the feedback loop of an op amp), also lowers the s/h amplifier's pedestal errors. The charge injected into the op amp's summing node—a virtual ground—becomes constant regardless of the input voltage. The pedestal now also becomes a constant that can be trimmed to zero with a potentiometer. Alternatively, in a system autozero mode, the pedestal can be trimmed with a d-a converter. When used in this way, essentially as an integrator, the fast op amp must be stable at a noise-gain of unity.

Feedthrough Errors

Feedthrough of the input signal to the output when in the "hold" mode also haunts s/h amplifiers. Typically, feedthrough is a function of the input signal frequency and its amplitude. If the amplitude is less than 1/2 LSB, feedthrough effects won't show up in the spectrum of the digitized signal because 1/2 LSB is below the quantization noise floor. Feedthrough errors above 1/2 LSB cause erroneous codes, increased total harmonic distortion (THD), and lower signal-to-noise ratios (SNR).

Feedthrough error is determined by the equation:

Equation 2: $\Delta V_O = C_{DS}/C_H \Delta V_{in}$

where C_{DS} is the source-to-drain capacitance of the switch, and V_{in} is the peak-to-peak input voltage.

An obvious method to cut feedthrough is by reducing the peak-to-peak voltage seen by the hold switch. "Turning off" the input buffer amplifier by disconnecting the ac signal from the s/h amplifier's input achieves this with the help of SW_1, a single-pole, double-throw switch. It connects the noninverting terminal of op amp A_1 to ground during the hold mode and to the input signal during the sample mode. Therefore, SW_2 no longer sees a high-frequency ±5-V signal during the hold mode, but rather a "fixed" dc potential, which in this case is ground. As a result, feedthrough is virtually reduced to zero.

Switching the noninverting terminal to ground offers an added benefit. It drops the droop rate. Since the potential across the switch is zero volts (both sides of the switch see a virtual ground), there is virtually no leakage current through the open switch. The droop rate is dependent only on the input offset current of A_2, 20 nA maximum, and charge leaking through the capacitor.

Low-Cost, 1-MHz Sampler

Conversion rates of the latest generation of high-speed monolithic a-d converters have reached 2 MHz. Though hybrid s/h amplifiers can deliver the fast acquisition times demanded by these converters, they usually cost more than the a-d converter itself. Although requiring more design work, a discrete approach may be a cost-effective alternative to buying an expensive hybrid.

Another discrete approach employs two high-speed CB op amps, a DMOS quad switch, and a few passive components *(Fig. 2a)*. The circuit works with a high-speed a-d converter chip, such as the AD671, a monolithic 12-bit, 500 ns converter, and delivers a throughput rate of 1 Msample/s.

Even though a high input impedance would simplify signal conditioning, the s/h amplifier's circuit trades off a 2000-Ω input impedance for the lower distortion of an inverting configuration. Noninverting designs are subject to common-mode errors, pedestal errors that are a direct function of input signal level, and nonlinearities created by the switch. All three of these errors can lead to greater total harmonic distortion from the s/h amplifier.

High-speed acquisition is facilitated by the fast DMOS quad switch—

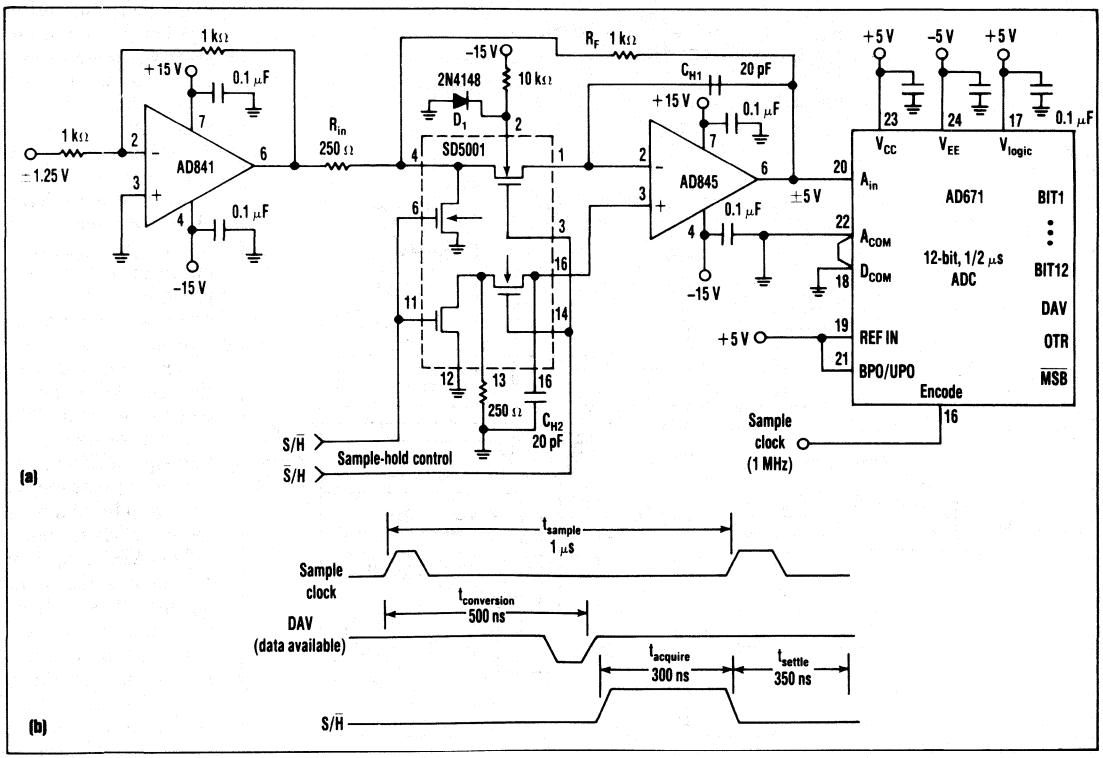

2. FOR THE ULTIMATE in speed and accuracy while sampling signals for a 12-bit, 1/2 μs a-d converter (for example, the AD671), it's often best to build a s/h amplifier with discrete CMOS transistor analog switches, such as Signetics SD5001 (a). Moreover, the s/h amplifier must be given time to acquire the signal and settle (b).

the Siliconix SD5001 turns on in just 1 ns. Biasing its body one diode-drop below ground with the 2N4148 diode produces a lower switch on-resistance for a given gate-to-source voltage. The switch also boasts a very small gate-to-drain capacitance of 1.6 pF, which further helps decrease charge injected onto the hold capacitor. The DMOS switches' control inputs require at least CMOS logic-level complementary voltages to turn them on. However, higher control levels for the gates of the switches will improve performance by lowering the on-resistance of the switches.

Because the AD845 works in a differential mode, droop rate is a function of input "offset" current rather than input "bias" current. The AD845 is specified to have a maximum 20 nA of input offset current. Using the equation:

Equation 3: $\Delta V/\Delta t = i_{offset}/C_{H1}$

where C_{H1} is the value of the hold capacitor, will designate the droop rate.

The voltage droop will be a maximum 1 mV/μs and is typically 1.25 μV/μs. Even if this calculation doesn't factor in the leakage current through the switch; the gate, source, and drain of the switch are all close to 0 V, creating a low differential voltage across it and negligible leakage current.

Minimizing Feedthrough

To minimize feedthrough, the fast s/h amplifier uses a T-switch configuration. When the circuit is in the hold mode, the input signal is shunted to ground through one of the SD5001's four switches. The switch in series with the 250-Ω input resistor acts as a voltage divider. With a typical on-resistance of 50 Ω and a ±1.25-V input signal, the maximum feedthrough voltage seen at the source terminal of the "hold" switch is 200 mVpk-pk. The high impedance of the open switch further reduces signal feedthrough more than 60 dB.

Because resistors R_F and R_{in} determine the gain of the s/h amplifier, designers can base their value on the maximum input signal and the full-scale input range of the a-d converter it's feeding. The values shown are chosen for a gain of −4, which increases the input amplitude from ±1.25 V to the ±5-V range of the AD671.

Even though the AD671 can perform a conversion in less than 1 μs, it's necessary to allow the s/h amplifier time to acquire the signal and then settle. However, the timing of the AD671 is not conventional. Its multistep conversion architecture doesn't require the input signal to settle to 12-bit accuracy until 150 ns of its conversion cycle has passed *(Fig. 2b)*.

A ping-pong s/h amplifier architecture improves throughput by employing two sampling circuits. While one is sampling or acquiring a signal, the other is in hold driving the a-d converter or some other load *(Fig. 3a)*. It takes just three high-speed IC op amps and a high-speed IC switch to build a ping-pong configuration.

The AD841 op amp A_1 follows the input signal. The DG542 switch, a dual wideband "T" circuit, alternately connects the input buffer A_1 to one of the two output op amps and its storage capacitor. At the same time, the switch also connects the amplifi-

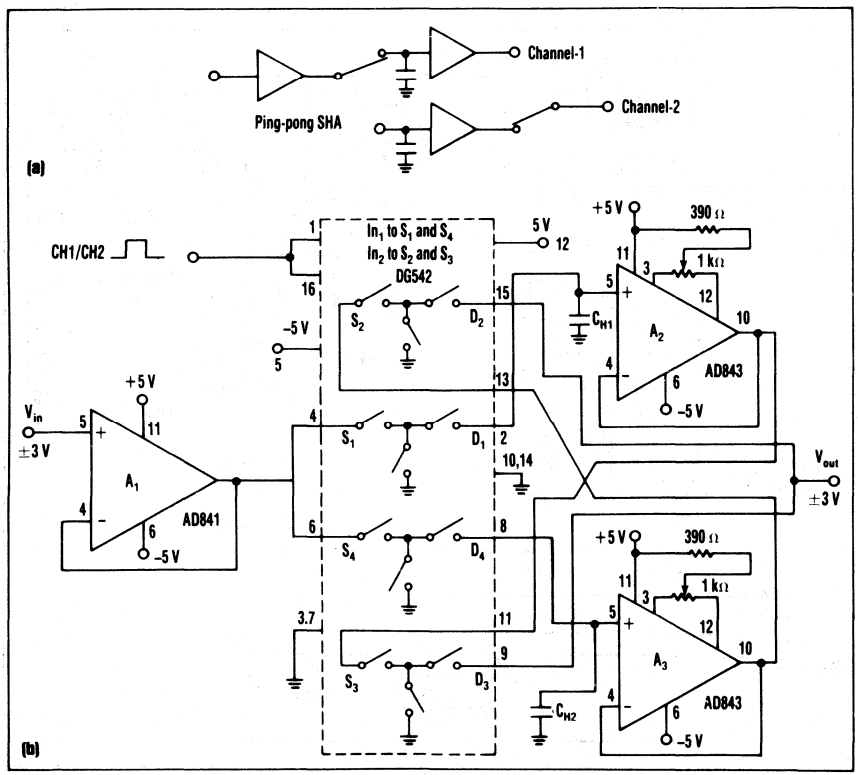

3. MAXIMUM THROUGHPUT RATE is achieved with a ping-pong s/h amplifier, where one channel samples the input while the other is in hold mode (a). Using the DG542 DMOS video switch, which employs T-type switches, minimizes feedthrough and crosstalk (b).

er's output in hold mode to the circuit's output V_{out}. For example, when pins 1 and 16 of the switch are at logic high, A_2 is tracking the input signal, and its output is disconnected from V_{out}. A_3 and its storage capacitor hold an analog voltage and it's output is connected to V_{out}. When the sample and hold command on pins 1 and 16 of the switch goes to logic low, the two output amplifiers switch jobs.

The ping-pong technique offers more than just greater throughput rates to high-speed s/h amplifier applications. For example, the technique makes it possible for designers to increase the value of the hold capacitors. Because the acquisition cycle occurs in parallel with the hold cycle, the RC time constant resulting from the finite on-resistance of the switch and the hold capacitance no longer limits throughput rate. If the acquisition time is less than the required hold-time for the conversion, then the slew rate and settling time of the output amplifier limit throughput. The two AD843 FET input op amps slew at 250 V/μs and settle to within $\pm$.01% of final value in 130 ns for $\pm$10-V steps.

The choice of the appropriate switch is critical in this type of design. The DG542 utilizes "T" switching techniques on each channel for exceptionally low crosstalk and high isolation.

The device further improves these specifications by locating ground pins between the signal pins. With an input frequency of 5 MHz, crosstalk and isolation are –85 dB and –75 dB, respectively.

The switch is limited to a maximum of –5 V on its negative supply input, making bipolar operation difficult. All three amplifiers should operate from the same –5-V rail as the switch to minimize potential latch-up. This limited supply range will restrict the amplitude of the inputs to video-level signals of $\pm$3 V. If a wider range is needed to maximize the full-scale range of an a-d converter, gain can be taken at the two output amplifiers by operating them as followers with gain. However, they'll have to be run from $\pm$12-V supply rails.

Because the input to the a-d converter consists of the alternating "held" values from the two amplifiers, the mismatch between their offset voltages will show up as nonlinearities and, therefore, distortion in the output signal. This demands that the amplifiers offer excellent ac performance and good dc precision. The AD843's maximum offset voltage is 1 mV, offset voltage drift is 10 μV/°C, and open-loop gain is 94 dB.

However, the AD843's balance pins accommodate offset trims. Users can ground the input and adjust the difference between their outputs through zero. The output is connected to a high gain oscilloscope and the S/H control input is toggled. The square wave created by the difference in offset voltages is adjusted to zero.

If offset voltage drift over temperature becomes a problem, especially with FET-input amplifiers, designers can employ an autocalibration circuit using two d-a converters. A "known" reference voltage, or ground, can be connected to the input op amp and the offset voltage adjusted to zero by manipulating the digital codes of two 8-bit d-a converters connected to the wipers of the op amps' trim-pins.

Choose The Right Caps

When selecting the appropriate types of capacitors for building a s/h amplifier, designers should employ devices with low dielectric absorption and low temperature coefficients (TCs). Silvered-mica capacitors exhibit low (0 to 100 ppm/°C) TCs and operate in excess of 200°C. Users should test the capacitors to ensure that their actual value matches their marked value. Not all manufacturers fully test all of their capacitors for absolute tolerance.

Aperture delay and aperture uncertainty, or jitter, represent two additional vital s/h amplifier specifications. Aperture delay produces an error that's determined by the equation:

Equation 4: $E_a = T_a\,(dV/dt)$

where E_a is the aperture error in volts, T_a is the aperture delay and dV/dt is the slew rate of the input signal. This error is a result of the finite time it takes a switch to open, and the effect of "averaging" the input signal over the closing-to-opening of the switch. However, as long as the switch opens in a repeatable fashion, the aperture time can be viewed as a phase shift in the input signal, not as an error source in a single-channel system. In fact, it can be calibrated out.

Circuit designers must guard against variable aperture delays or "aperture jitter." The maximum bandwidth of a s/h amplifier with a given aperture jitter (t_a) is determined by the equation:

Equation 5: $f_{MAX} = 1/(2\pi\, t_a)(2^{(n+1)})$

where n in the expression $2^{(n+1)}$ is the resolution of the a-d converter following the s/h amplifier. For a 12-bit converter and 10-ps aperture jitter, the maximum bandwidth is approximately 2 MHz. This equation illustrates that the s/h amplifier must maintain low jitter to be used with high-speed, high-resolution a-d converters.

Aperture jitter can arise from phase modulation in the sampling clock from either wideband random noise, power-line frequency noise, or digital noise due to poor grounding. Careful board layout can reduce these sources. The pc-board trace between the source of the s/h command and s/h amplifier's circuit should be as short as possible to minimize the effects of parasitic circuit inductance, resistance, and capacitance. Moreover, long runs of digital designers should make sure that control lines are not run parallel to any analog signal paths.□

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Polarity Programmable Peak Detector

Four ICs are used in this design to give positive or negative peak detection and reset levels over a ±10V range. A precision voltage comparator, a sample-and-hold amplifier, an open-collector exclusive-OR gate package, and a quad analog switch used as two SPDT switches are the principal components required.

BLOCK DIAGRAM

Figure 1 shows the basic circuit and the 4 modes of operation. A comparator continuously examines the difference between the present analog input voltage and a voltage peak held by the sample and hold (S/H) amplifier. If the present value exceeds the held value, the S/H is placed in the "Track" condition and acquires a new peak value returning to a "Hold" condition when coincidence is reached. The comparator's output is inverted or non-inverted depending on the polarity selected by POS/$\overline{\text{NEG}}$ SELECT. This same TTL control signal connects the appropriate voltage-programmed reset voltage to the S/H input during the RESET modes. In RESET the S/H is forced to a "Track" condition by the PEAK DETECT/$\overline{\text{RESET}}$ SELECT digital input. Figure 2 shows typical waveforms.

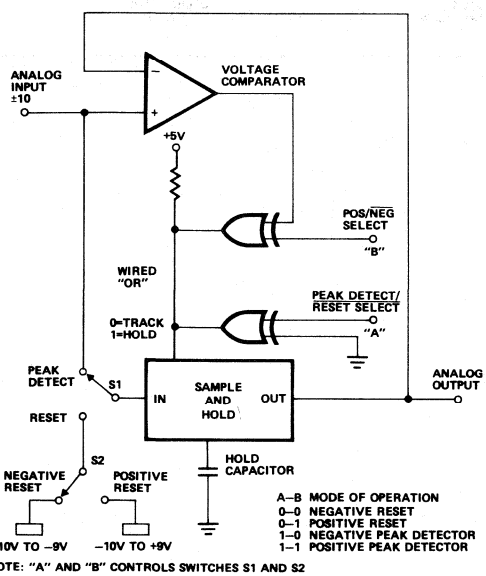

Figure 1. Programmable Peak Detector Block Diagram

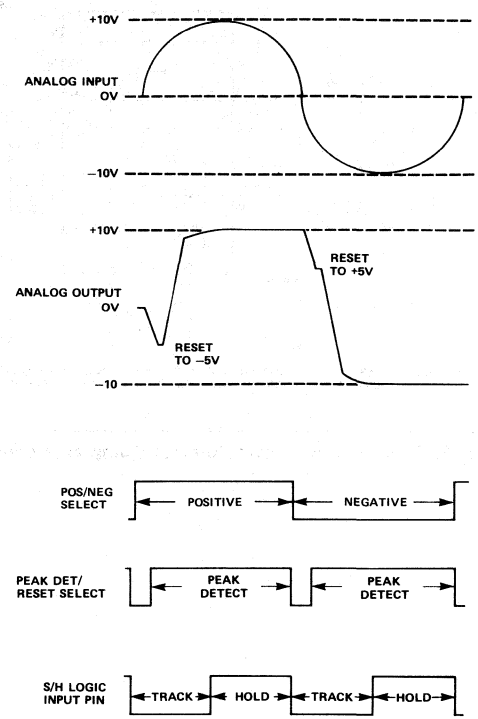

Figure 2. Programmable Peak Detector Waveforms

DETAILED CIRCUIT DESCRIPTION

In Figure 3, the SMP-11FY S/H is specified for three reasons: low cost, 2.5mV (Maximum) zero-scale error, and its 10V/μs slew rate. Together with the CMP-01CJ precision voltage comparator, system unadjusted DC accuracy is within 5mV at zero-scale and 10mV at full-scale. Comparator input over-voltage protection resistors and diodes are required as shown.

A SW-01 Quad Analog Switch connected at 2 SPDT switches connects the proper analog inputs to the S/H during the 4 modes of operation. 200 ohm current-limiting resistors are used as recommended by the switch manufacturer. Logic drive to the switch is provided by 1/2 of a SN74LS136 quad exclusive-OR gate used as 2 inverters.

Bypassing as shown is strongly recommended — 0.1μF ceramic dieletric capacitors for the comparator and the S/H, plus 10μF solid tantalum bypass capacitors physically close to the S/H. In addition the use of a ground plane is recommended to minimize ground path resistances. A 0.01μF hold capacitor is used to minimize overshoot when tracking high frequency analog inputs.

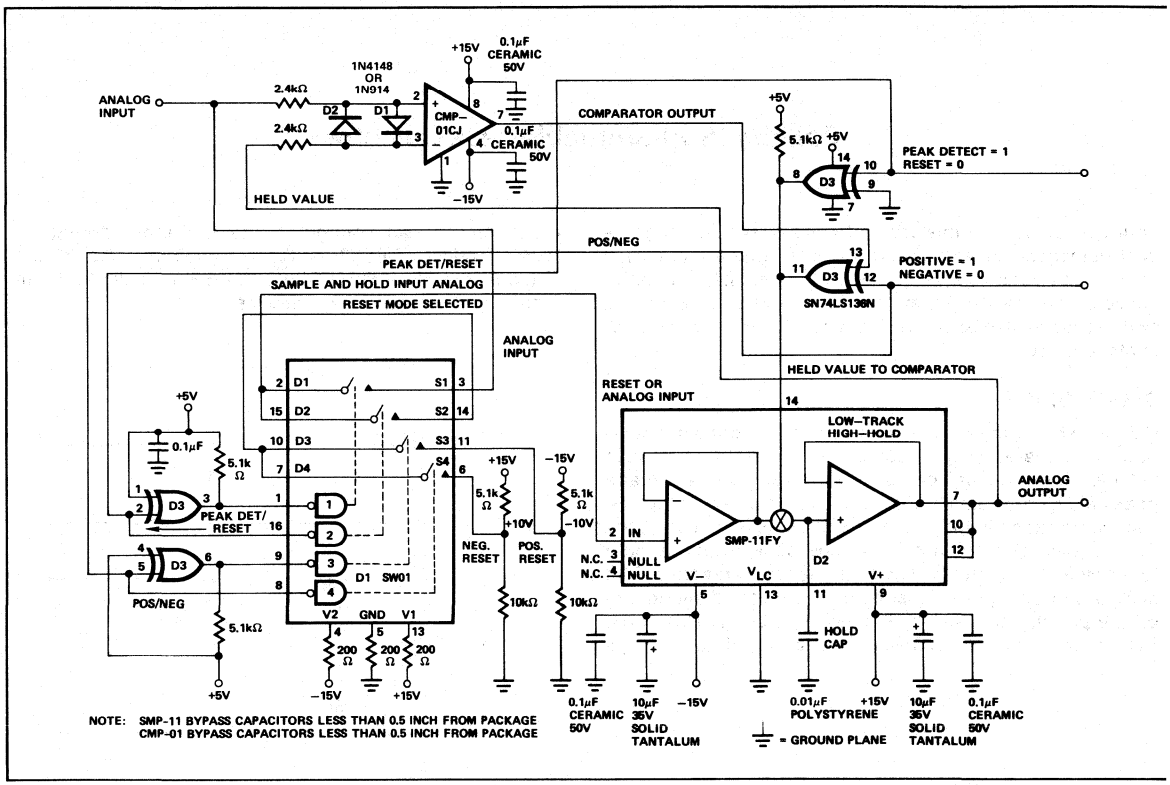

NOTE: SMP-11 BYPASS CAPACITORS LESS THAN 0.5 INCH FROM PACKAGE
CMP-01 BYPASS CAPACITORS LESS THAN 0.5 INCH FROM PACKAGE

Figure 3. Programmable Peak Detector Complete Schematic

ANALOG DEVICES

AN-204
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Applications of the SMP-04 and the SMP-08/SMP-18, Quad and Octal Sample-and-Hold Amplifiers

By Adolfo A. Garcia

The SMP-04 and the SMP-08 are monolithic quad and octal sample-and-hold amplifiers that greatly reduce the complexity and the cost of discrete approaches (op amps, capacitors, switches, and multiplexers). Built on an advanced analog CMOS process, the SMP-04 and SMP-08 exhibit very low droop rates of 2 mV/s and fast acquisition times of 7 μs. The SMP-18, a faster version of the SMP-08 (2.5 μs acquisition time) with similar specifications, is also available. The on-chip MOS capacitors save space and reduce cost by eliminating external, expensive low leakage capacitors. As a result, layout guard-ringing and details for board cleanliness are no longer necessary to maintain specified performance.

Architecturally, the SMP-04 contains four independent sample-and-hold amplifiers; whereas the SMP-08 and the SMP-18 contain eight sample-and-hold amplifiers multiplexing one common input to eight independent outputs. The SMP-04 will be used throughout this application note to describe all devices, unless otherwise noted. Simplified block diagrams of the SMP-04 and the SMP-08/SMP-18 are shown in Figure 1a and Figure 1b.

The versatility of these amplifiers opens up many useful applications which will be described below. This application note will point out special issues and techniques to derive the maximum performance from these devices.

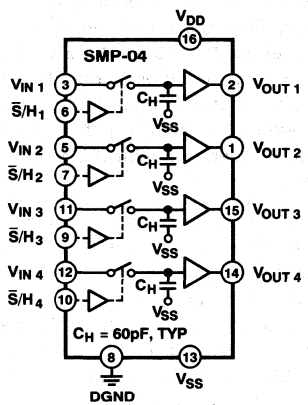

Figure 1a. Block Diagram of the SMP-04

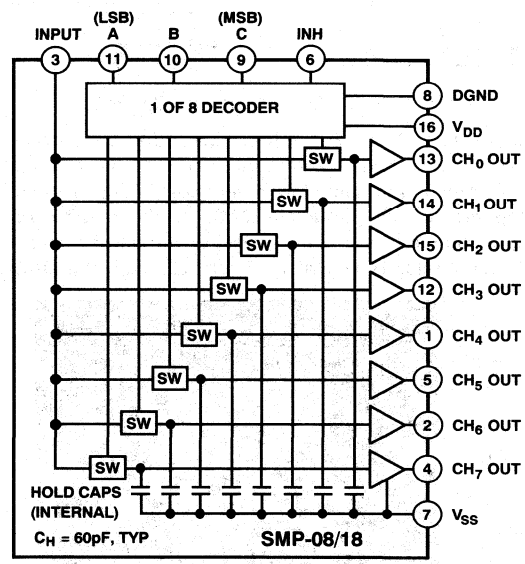

Figure 1b. Block Diagram of the SMP-08/SMP-18

APPLICATION CIRCUITS

An Eight-Channel Setpoint Controller

Figure 2 illustrates a low cost analog setpoint controller circuit by demultiplexing a single DAC output to eight independent channels using the SMP-08. With a maximum droop rate of 20 mV/s, the SMP-08 can maintain 8-bit accuracy (1/2 LSB at 5 V full scale) by refreshing

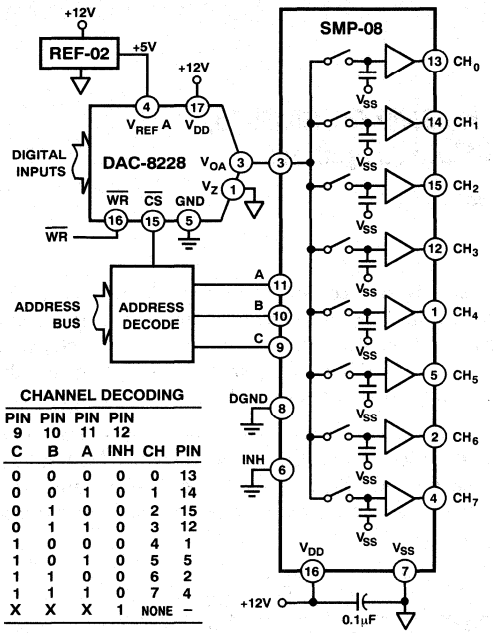

CHANNEL DECODING

PIN 9 C	PIN 10 B	PIN 11 A	PIN 12 INH	CH	PIN
0	0	0	0	0	13
0	0	1	0	1	14
0	1	0	0	2	15
0	1	1	0	3	12
1	0	0	0	4	1
1	0	1	0	5	5
1	1	0	0	6	2
1	1	1	0	7	4
X	X	X	1	NONE	–

Figure 2. An 8-Channel Setpoint Controller

once every 500 ms. For a 10-bit DAC, the refresh rate must be less than 122 ms and, for a 12-bit system, 31 ms. This implementation is very cost effective when compared to using multiple DACs as the number of output channels increases. Typical operation of the setpoint controller is shown in Figure 3a. With the DAC

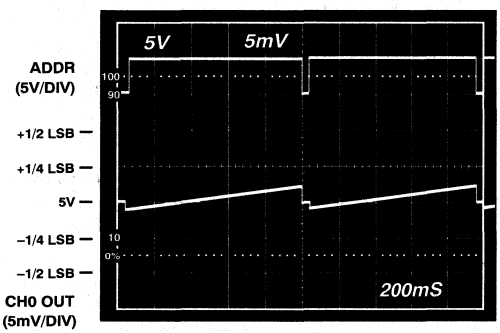

Figure 3a. Illustrates the Operation of the Setpoint Controller. The top trace is the address decoder selecting CH0 (LOW) or CH7 (HIGH). The bottom trace is CH0's output illustrating the SMP-08's typical droop and hold step.

output at 5 V, the address decoder drives all address lines concurrently, selecting either Channel 0 or Channel 7. As the photo shows, the SMP-08 maintains less than 1/4 LSB error (for 8-bit accuracy) for 1 second. An expanded view of the droop, the reset to 5 V, and a hold step of 1 mV is shown in Figure 3b.

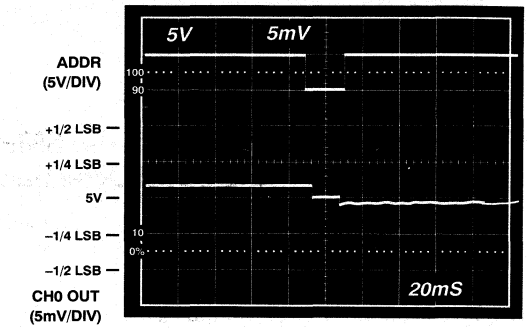

Figure 3b. Illustrates the Expanded View of the Setpoint Controller's Operation

A Single Event Sampler

The circuit in Figure 4 takes 16 equally timed consecutive samples to capture a single event on the input signal. Once captured, the sampled signals can be sequentially reconstructed repetitively on an oscilloscope, either at the original sampling clock rate or a slower clock rate.

By selecting the position of switch SW1, sampling can begin by triggering on the input's signal level, or on an external trigger (negative going) pulse that is no more than one clock period long. In the case where a 1 MHz clock is used, the pulse width must be 1 μs or less. Either signal resets the 74LS161 counter asynchronously. The sequential sampling begins at the rising edge of the second clock period. The reset pulse also clears both J-K flip-flops, removing the inhibit to the SMP-08s. The synchronous counter outputs ensure equal sampling periods to each sample-and-hold amplifier. The J-K flip-flops toggle and subsequently inhibit the SMP-08s from further sampling.

To reconstruct the sampled analog signal, the 16 sampled signals are sequentially addressed by the MUX-16, a 16-channel analog multiplexer. The 74LS163 binary counter cycles the multiplexer which generates a repetitive waveform as shown in Figure 5. The waveform will keep its shape for as long as ten seconds, but beyond that the SMP-08's 20 mV/s maximum droop rate causes it begin to lose its shape gradually.

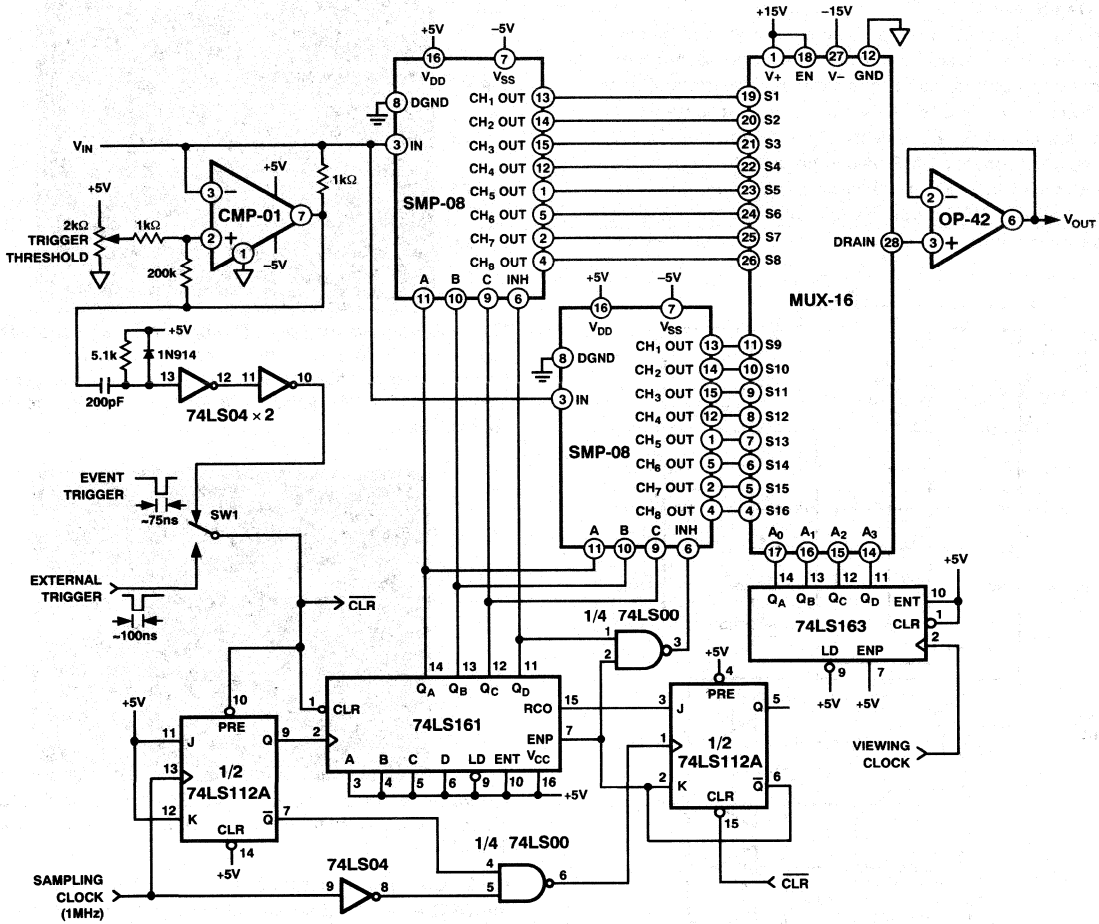

Figure 4. A 16-Channel Single Event Sampler

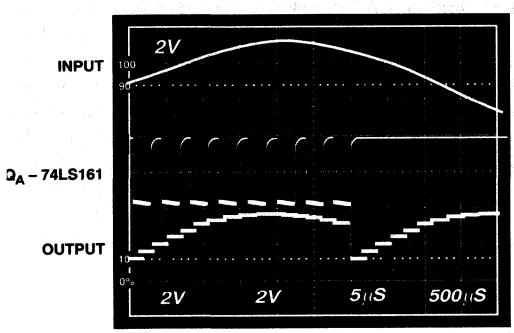

Figure 5. Illustrates the Operation of the 16-Channel Single Event Sampler

A variation of this concept has been demonstrated in other publications. One such design is described in Application Note AN-289 published by National Semiconductor Corp. That approach used eight separate sample-and-hold amplifiers plus eight capacitors. This design, however, illustrates how a considerably smaller circuit using less components, and therefore lower costs, can be realized using a single-chip octal sample-and-hold amplifier like the SMP-08. In fact, the beauty of this configuration is that the number of sample-and-hold channels can be easily expanded to increase the sampling resolution without a heavy penalty in increased circuitry and component costs.

DAC Deglitcher

Most DACs output an appreciable amount of "glitch" energy during a transition from one code to another. The glitch amplitude can range from several millivolts to hundreds of millivolts. This may be unacceptable in many applications. By selectively delaying the DAC's output transition, the SMP-04 can be used to smooth the output waveform. Figure 6 shows the schematic diagram of a quad DAC deglitcher circuit using only a write strobe command ($\overline{WR}$) generated by a microprocessor.

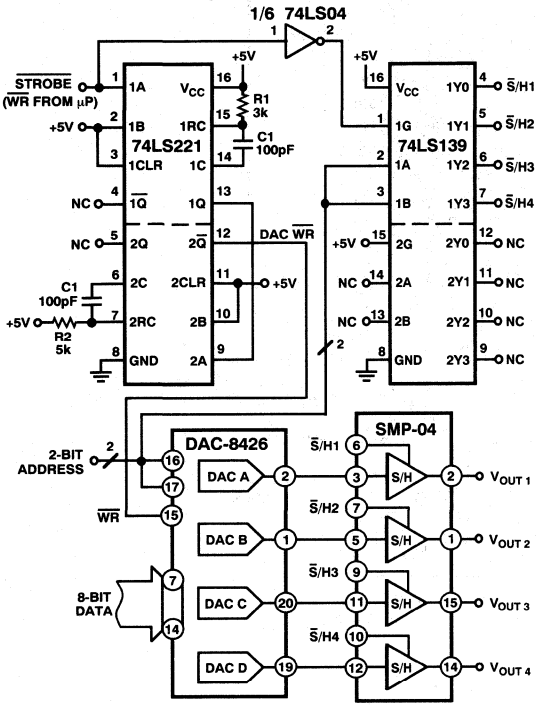

Figure 6. A Quad DAC Deglitcher

Dual one-shots (a 74LS221) and an inverter provide the proper delay timing for the DAC $\overline{WR}$ strobe and the $\overline{S}$/H control signal to the SMP-04. The 74LS139 demultiplexer is used to steer the $\overline{S}$/H control signal to the appropriate SMP-04 channel. Therefore, only one timing circuit is used to deglitch a quad DAC. In this example, a linear ramp signal is generated by the microprocessor whose $\overline{WR}$ strobe is 800 ns in duration. In Figure 7a, new data to

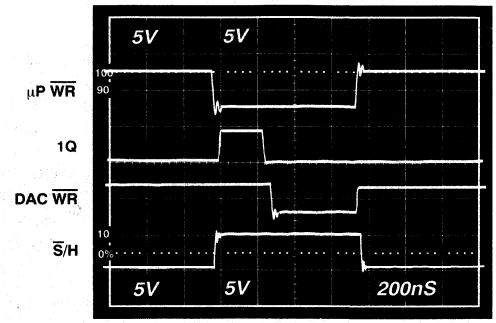

Figure 7a. Shows the Logic Timing of the Deglitcher. The top trace is the $\overline{WR}$ from the microprocessor. The second trace is used to strobe the second one-shot that generates the DAC's $\overline{WR}$ strobe. The bottom trace is the S/H control signal.

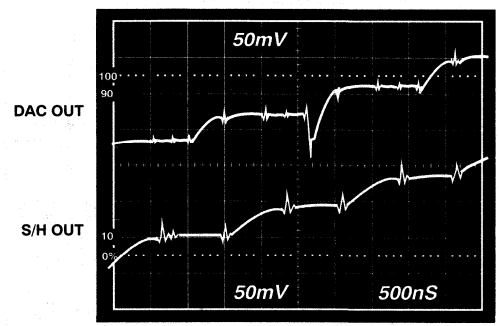

Figure 7b. Shows a Typical DAC Glitch Amplitude (Top Trace) and the Deglitched Output of the SMP-04 (Bottom Trace)

the DAC A input is set up at the S/H1's falling edge, but the DAC A output does not change until its $\overline{WR}$ strobe goes active. When that occurs, the S/H1 begins sampling and tracks DAC A's output. When the $\overline{S}$/H signal goes HIGH, DAC A's output voltage is held by S/H1. After a 300 ns settling, DAC A's $\overline{WR}$ strobe goes LOW to allow the DAC output to change. Any glitch that occurs at the DAC output is blocked by the SMP-04. As soon as the $\overline{WR}$ strobe goes HIGH, the digital data is latched; at the same time, the $\overline{S}$/H goes LOW, and the SMP-04 is allowed to track the new DAC output voltage. Figure 7b shows the deglitching operation. The top trace shows the DAC A's output during transition while the bottom trace shows the deglitched output of the SMP-04.

A Pipelined Data Acquisition System

If a sample-and-hold is used in the front of a successive approximation register (SAR) type ADC, the system throughput can be increased by adding a second sample-and-hold operating in a "ping-pong" fashion as shown in Figure 8. While the first channel is holding the input signal for the A/D conversion, the second sample-and-hold is acquiring another channel through the multiplexer. By the time the A/D conversion is complete on the first sample-and-hold's signal, the second sample-and-hold will have acquired its data and will be ready for the next conversion. This pipeline technique eliminates the acquisition time of the sample-and-hold in the signal path.

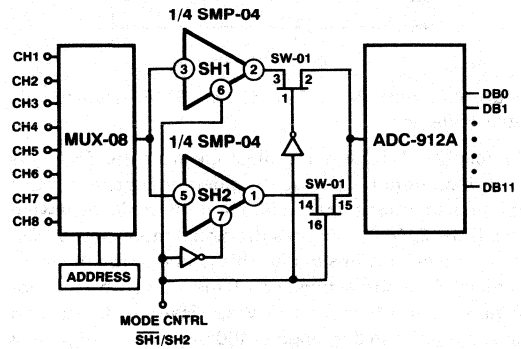

Figure 8. A Pipelined Data Acquisition System

A Low Cost Battery Tester

The condition of a new or old battery can be tested very quickly under a load with the circuit in Figure 9. When the momentary switch is pressed, the no-load battery voltage is sampled by the first sample-and-hold. R3 and R4 are chosen so that, according to the battery manufacturer's end-of-life specifications, the threshold of the PM-139 comparator is set to 87.5% of the no-load battery voltage. Since this is a very low frequency application, approximately 16 mV of hysteresis about the threshold voltage, set by R5 and R6, was chosen to prevent the comparator from oscillating. The gate of the power NMOS transistor is controlled by a differentiator whose time constant is set at 470 μs. When the power FET turns on, it forces a 500 mA load current from the battery. At the same time, the second differentiator controls the second sample-and-hold. Thus, the battery's voltage under load is sampled and applied to the inverting terminal of the comparator. If the battery's voltage is at or above the threshold, then the comparator's output goes low and triggers the GREEN LED through the D flip-flop, indicating a PASS condition; otherwise, the RED LED turns on. To clear the D flip-flop during the power-on sequence, a 100 ms R10-C3 combination was connected to the D flip-flop's RESET pin through an inverter.

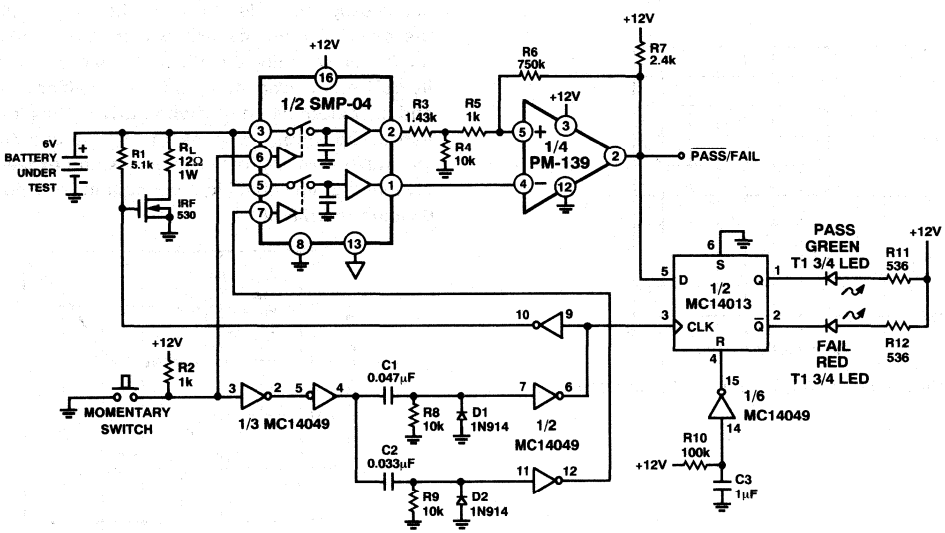

Figure 9. Low Cost Battery Tester

APPLICATIONS HINTS

Input Drive Considerations

For each channel of the SMP-04, the input looks directly into a CMOS transmission gate and an internal, low leakage MOS capacitor, as shown in Figure 10. The ON resistance of the transmission gate and the internal hold capacitor, which is typically 60 pF, largely determine the acquisition time of the SMP-04. Any drive circuit having a finite output resistance will increase the RC charge time and degrade the acquisition time. Consequently, good acquisition time can be preserved by driving the input of the SMP-04 from low output resistance sources.

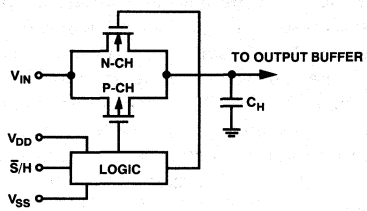

Figure 10. Input Circuit of the SMP-04 and the SMP-08/SMP-18

Since the transmission gate is constructed of enhancement-mode MOSFETs, its ON resistance is a function of both the supply voltages and the applied input signal. Figure 11a shows the transmission gate's ON resistance of the SMP-04 for a single +12 V supply as a function of the applied analog input voltage. Similarly, Figure 11b show the switch resistance for ±5 V supplies as a function of the applied analog input voltage.

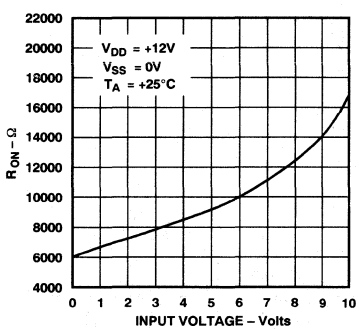

Figure 11a. SMP-04/SMP-08 Switch ON Resistance vs. Input Voltage

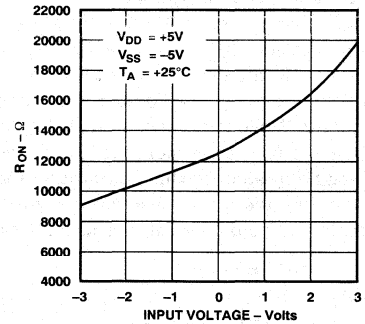

Figure 11b. SMP-04/SMP-08 Switch ON Resistance vs. Input Voltage

Figures 12a, 12b, and 12c illustrate how the SMP-04's acquisition time is affected by different supply voltages and applied input signals. Note that for positive-slewing signals the SMP-04 acquires the signal more slowly than for negative-slewing signals. This is due to the nonsymmetrical slew characteristics of the output buffer. For single supply +15 V or +12 V applications, the analog input should be in the range of 100 mV $\leq V_{IN} \leq V_{DD}-2$ V to avoid significant degradation of the sampled signal.

The input capacitance is a combination of the internal hold capacitor, the transmission gate's parasitic input capacitance, the output buffer's parasitic input capacitance, and the parasitic package pin capacitance. In the sample mode, the SMP-04's typical input capacitance is approximately 60 pF. In the hold mode, the input capacitance drops to under 10 pF. In the case of the SMP-08, the external input capacitance in the sample *or* hold modes is typically 80 pF, and typically 30 pF in the INHIBIT mode. Depending on the amplifier used, ringing and overshoot can occur when driving this input capacitance. Op amps, such as the OP-42 and the AD847, which have low output resistances and excellent capacitive load drive capability, are ideal for driving the SMP-04 and the SMP-08.

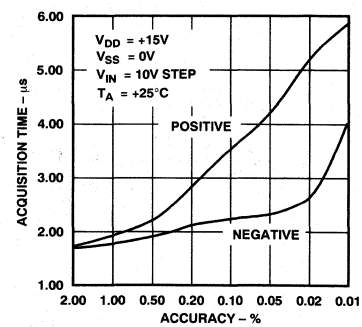

Figure 12a. SMP-04/SMP-08 Acquisition Time vs. Percent Accuracy

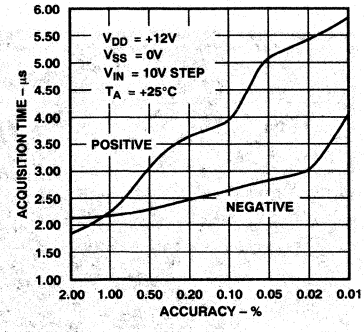

Figure 12b. SMP-04/SMP-08 Acquisition Time vs. Percent Accuracy

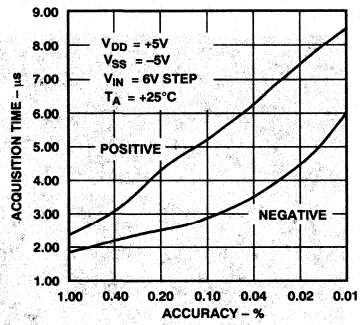

Figure 12c. SMP-04/SMP-08 Acquisition Time vs. Percent Accuracy

Using The SMP-04 As A Buffer

It is sometimes advantageous to use one or more sections of the SMP-04 as a buffer. This is done by setting the SMP-04 in a continuous SAMPLE mode by connecting the $\overline{S}$/H control pin to logic LOW. The SMP-04 has good distortion characteristics over the audio frequency range. For a single 12 V supply, Figure 13a illustrates the track-mode distortion characteristics of the SMP-04. Note that the SMP-04 maintains less than 0.04% total harmonic distortion over the entire audio frequency

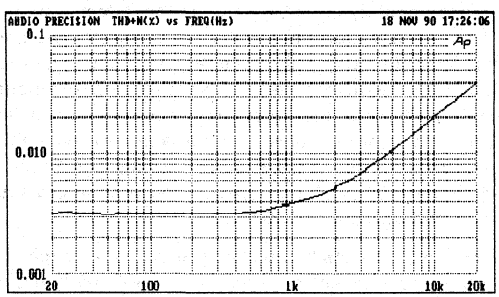

Figure 13a. SMP-04 Track Mode THD+N vs. Frequency. V_S = +12 V/GND, V_{OFFSET} = +6 V, R_L = 10 kΩ, and V_{IN} = 1 V p–p

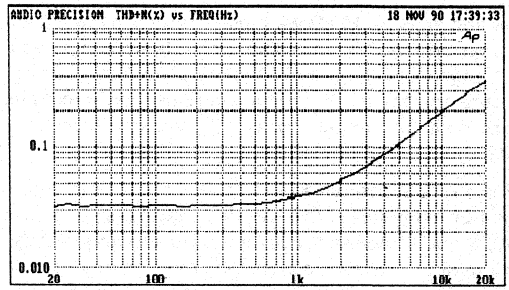

Figure 13b. SMP-04 Track Mode THD+N vs. Frequency. V_S = ±5 V, V_{OFFSET} = +0 V, R_L = 10 kΩ, and V_{IN} = 1 V p–p

band. For a ±5 V supply operation, the SMP-04 exhibits less than 0.4% total harmonic distortion over the same frequency band, as shown in Figure 13b. For audio distribution applications, one might consider using the SMP-08. Figure 14a illustrates its track-mode distortion characteristics for a single 12 V supply. The SMP-08 maintains less than 0.6% total harmonic distortion over the entire audio frequency band. Figure 14b illustrates the SMP-08's distortion characteristics for ±5 V supply operation. Distortion performance is only slightly worse (0.7%) than the single supply case over the entire audio frequency band. In all measurements, a standard 80 kHz

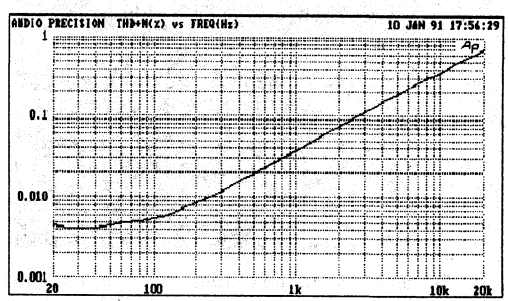

Figure 14a. SMP-08 Track Mode THD+N vs. Frequency. V_S = +12 V/GND, V_{OFFSET} = +6 V, R_L = 10 kΩ, and V_{IN} = 1 V p–p

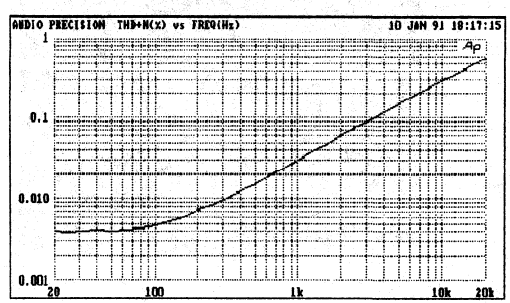

Figure 14b. SMP-08 Track Mode THD+N vs. Frequency. V_S = ±5 V, V_{OFFSET} = +0 V, R_L = 10 kΩ, and V_{IN} = 1 V p–p

filter and an R_L = 10 kΩ were used. If track-mode distortion performance is important, use a large supply voltage range (V_{DD}–V_{SS}) to yield lower distortion products. Be sure to remain well within the absolute maximum supply voltage ratings, that is, V_{DD}–V_{SS} less than 17 V, maximum.

Transient Load Response

When an amplifier is used to drive the input of a SAR ADC, the ADC generates abrupt current changes at the analog input. For accurate conversions, the drive amplifier must be capable of holding a constant output voltage under dynamically changing load conditions. Therefore, it is important to know how quickly the amplifier's output settles under these conditions. Amplifiers with low closed-loop output impedances recover from output load current changes quickly without long settling tails.

The SMP-04's low closed-loop output impedance (1 ohm, typical) allows the output to settle quickly during a SAR ADC's conversion cycle. Figure 15a illustrates a typical transient load response (to 12-bit accuracy) of the SMP-04 with a 1 mA transient load current. This load current is a worst case condition for driving a SAR ADC with a 2 mA reference current. Figure 15b illustrates the

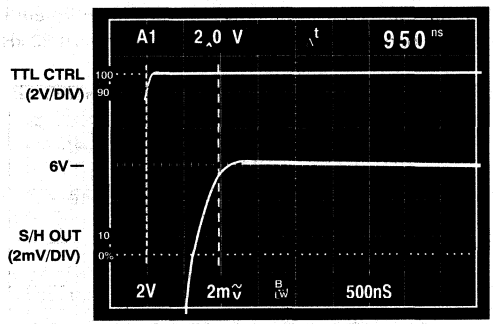

Figure 15a. SMP-04/SMP-08 Output Buffer Response to a 1 mA Load Current Transient. V_{DD} = +12 V, V_{SS} = GND

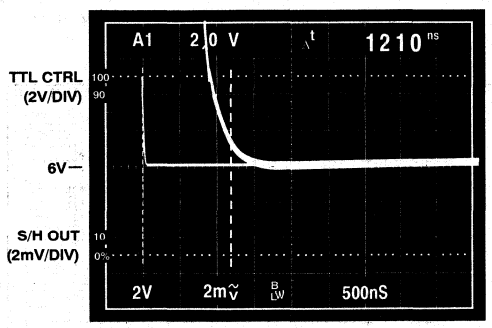

Figure 15b. SMP-04/SMP-08 Output Buffer Recovery from a 1 mA Load Current Transient. V_{DD} = +12 V, V_{SS} = GND

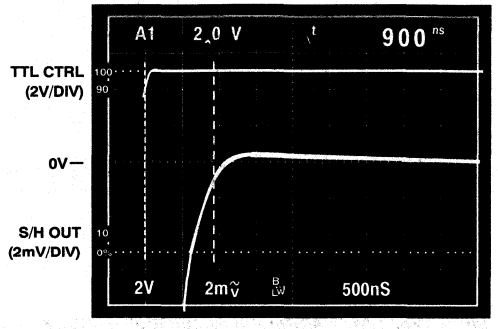

Figure 15c. SMP-04/SMP-08 Output Buffer Response to a 1 mA Load Current Transient. V_{DD} = +5 V, V_{SS} = −5 V

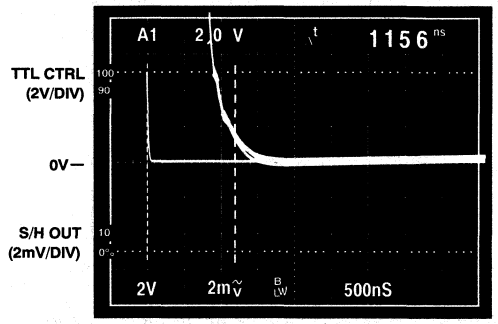

Figure 15d. SMP-04/SMP-08 Output Buffer Recovery from a 1 mA Load Current Transient. V_{DD} = +5 V, V_{SS} = −5 V

SMP-04's typical transient load response when the load current is turned off. The slower transient load recovery is due to asymmetric slew characteristics of the SMP-04's output buffer. For ±5 V supply applications, slightly faster response times to 12-bit accuracy are shown in Figures 15c and 15d. These results suggest that the minimum clock period for a 12-bit SAR ADC should be 1.3 µs.

Aperture Delay Time and Jitter

Two additional dynamic parameters that affect accuracy and total acquisition time are aperture delay time and jitter. Aperture delay is the time required by the internal switches to disconnect the hold capacitor from the input. Aperture jitter is the uncertainty in aperture delay caused by internal noise and variation of switching thresholds with the input signal. Typically, the SMP-04 and the SMP-08 exhibit an aperture delay time of 20 ns and an aperture jitter of 800 ps.

Interchannel Crosstalk and Hold-Mode Feedthrough

The four independent channels of the SMP-04 typically exhibit 100 dB of crosstalk rejection from dc to 100 kHz regardless of supply voltages used. In the hold-mode, its feedthrough measures −85 dB over the same frequency range.

SMP-08/SMP-18 Channel Decode Timing

For applications that require sequencing the SMP-08's (or the SMP-18's) output channels, Figure 16a illustrates the break-before-make timing relationships between the address lines and the analog input. When an address change occurs, the previously selected channel turns off in 45 ns. The next selected channel turns on in 90 ns. Therefore, to avoid potential crosstalk problems, the analog input must not change for 45 ns after an address change.

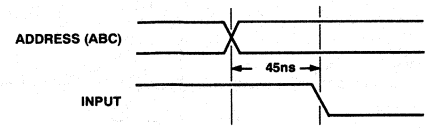

Figure 16a. SMP-08/SMP-18 Decode Timing Option 1

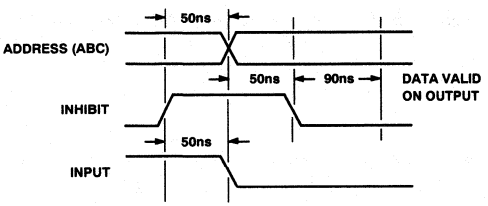

Figure 16b. SMP-08/SMP-18 Decode Timing Option 2

Another option is to use the INHIBIT function to disable all channels during address and analog input changes. The timing diagram in Figure 16b illustrates the relationships between the INHIBIT, the address lines, and the analog input. Once the INHIBIT goes LOW, the next selected channel's analog output is valid in 90 ns. It is therefore, recommended that the minimum INHIBIT pulse width be 100 ns to minimize feedthrough and glitching to the analog outputs. A simplified logic diagram of the 1:8 decoder is shown in Figure 17.

Single +5 V Operation

The SMP-04 and the SMP-08 can operate on a single +5 V supply but with reduced performance. Table I summarizes the performance characteristics of the SMP-04 and the SMP-08 with this supply voltage. For sampling applications, it is recommended that the analog input be in the range of 100 mV $\leq V_{IN} \leq$ 2.5 V to avoid degradation of the sampled signal. To maximize output voltage swing under this limited supply constraint, loads greater than 10 kΩ should be used.

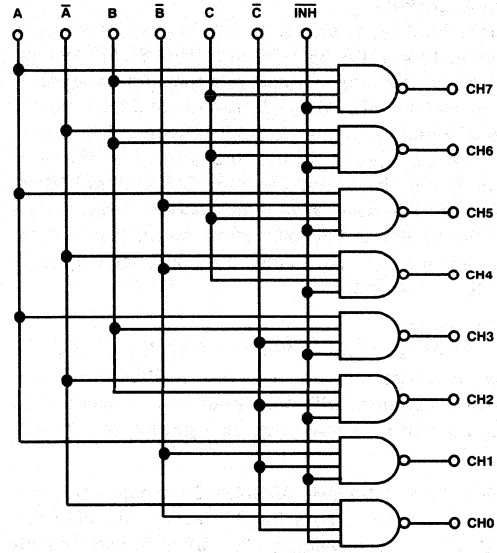

Figure 17. SMP-08/SMP-18 Decode Logic Diagram

Table I. Summary of SMP-04 and SMP-08 Single +5 V Characteristics

Buffer Offset Voltage V_{IN} = 2.5 V	3 mV
Input Voltage Range $R_L = \infty$	0.16 V to 4.1 V
Output Voltage Range $R_L = \infty$ R_L = 1 kΩ	0.15 V to 3.5 V 0.14 V to 1 V
Buffer Small-Signal Bandwidth	200 kHz
Buffer Phase Shift	53°
Acquisition Time 100 mV ≤ V_{IN} ≤ 2.5 V SMP-04 1% SMP-04 0.1% SMP-08 1% SMP-08 0.1%	 2.3 µs 8.3 µs 3.5 µs 14 µs
Total Harmonic Distortion 20 Hz–20 kHz, R_L = 10 kΩ, V_{IN} = 1 V p–p, Filter = 80 kHz LP SMP-04 SMP-08	 <0.3% <4%
Signal-to-Noise Ratio f_{SAMPLE} = 86 kHz, f_{IN} = 10 kHz f_{SAMPLE} = 14.4 kHz, f_{IN} = 1.8 kHz, t_{pw} = 10 µs	 18 dB 24 dB
SMP-04 Hold-Mode Feedthrough DC to 100 kHz	 85 dB
Supply Current SMP-04 SMP-08	 1.5 mA 3 mA

Use of IC sockets

Since the hold capacitors of the SMP-04 are internal and connected to V_{SS}, the SMP-04 and SMP-08/SMP-18 have no external capacitor connections that are exposed to leakage paths which degrade sample-and-hold performance. Therefore, unlike the LM398 that requires expensive, low leakage capacitors and special attention to leakage paths, the SMP-04 and the SMP-08/SMP-18 are tolerant of IC sockets. Wire-wrap sockets, however, are not recommended because of their susceptibility to noise pickup and larger parasitic capacitance compared to solder tail sockets.

ESD Protection

The SMP-04 and the SMP-08/SMP-18 have been designed to withstand over 900 volts of ESD on all pins. However, standard ESD prevention techniques should still be used to minimize potentially excessive high ESD voltages.

The SMP-04 and the SMP-08/SMP-18 are highly resistant to latchup. However, parasitic SCR action can still occur only if a current exceeding 200 mA is forced into the device's analog and digital inputs. This condition is not likely to ever occur because: (1) most amplifiers driving the analog inputs have built-in output short-circuit current limiting; and (2) ICs driving the digital inputs would be destroyed long before the 200 mA limit is reached.

Dual-to-Single Supply Interface

When the SMP-04, connected to a single supply, is driven by an amplifier operating from a dual supply, the circuit in Figure 18 should be used. The circuit configuration is a DAC current-to-voltage converter. The diode is used to clamp the SMP-04's input from negative voltages in the event that the amplifier's output goes to its negative output voltage limit during power on/off situations.

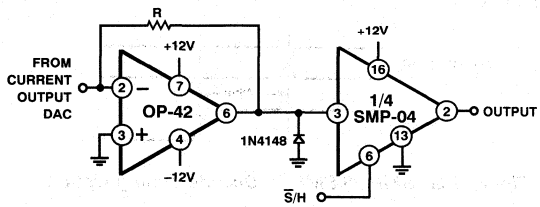

Figure 18. Dual-to-Single Supply Interface Circuit

Sensors
Contents

		Page
AN-272	Accuracies of the AD590	19–3
AN-273	Use of the AD590 Temperature Transducer in a Remote Sensing Application	19–7

19

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Accuracies of the AD590

The following tables contain maximum errors by grade for applications involving limited temperature spans. The tables reflect the worst-case nonlinearities of the AD590, which invariably occur at the ends of the specified temperature range. The "trims" in each table refer to the error correction circuits on pages 4 and 5 of the AD590 Data Sheet (Figures 4 and 7a). All accuracies given below are ±°C. For example, if ±1°C accuracy is required over the +25°C to +75°C range, then trimming a J grade device using the circuit of Figure 4 on the AD590 Data Sheet will result in a sensor of the required accuracy and range.

M GRADE

Number Of Trims	Temperature Span (°C)	Lowest Temperature In Span (°C)							
		-55	-25	0	+25	+50	+75	+100	+125
None	10	0.6	0.5	0.6	0.6	0.7	0.7	0.7	0.9
None	25	0.8	0.8	0.7	0.7	0.8	0.8	1.0	1.1
None	50	1.0	0.9	0.8	0.9	0.9	1.1	1.2	–
None	100	1.3	1.4	1.3	1.4	1.5	–	–	–
None	150	1.5	1.6	1.6	–	–	–	–	–
None	205	1.7	–	–	–	–	–	–	–
One	10	0.2	0.1	0.1	0.1	0.1	0.1	0.1	0.2
One	25	0.4	0.3	0.2	0.2	0.2	0.2	0.3	0.4
One	50	0.5	0.4	0.3	0.3	0.3	0.4	0.5	–
One	100	0.8	0.8	0.7	0.7	0.8	–	–	–
One	150	0.9	0.9	0.9	–	–	–	–	–
One	205	1.0	–	–	–	–	–	–	–
Two	10	0.1	*	*	*	*	*	*	0.1
Two	25	0.1	*	*	*	*	*	*	0.1
Two	50	0.2	*	*	*	*	*	0.2	–
Two	100	0.2	0.1	*	0.1	0.2	–	–	–
Two	150	0.3	0.2	0.3	–	–	–	–	–
Two	205	0.3	–	–	–	–	–	–	–

*Below ±0.05°C

19

L GRADE

Number Of Trims	Temperature Span (°C)	Lowest Temperature In Span (°C)							
		-55	-25	0	+25	+50	+75	+100	+125
None	10	1.0	1.0	1.1	1.1	1.2	1.3	1.4	1.6
None	25	1.3	1.3	1.3	1.4	1.5	1.6	1.7	1.9
None	50	1.9	1.8	1.7	1.8	1.9	2.1	2.4	–
None	100	2.4	2.4	2.4	2.4	2.7	–	–	–
None	150	2.7	2.6	2.8	–	–	–	–	–
None	205	3.0	–	–	–	–	–	–	–
One	10	0.2	0.1	0.1	0.1	0.1	0.1	0.1	0.2
One	25	0.5	0.4	0.3	0.3	0.3	0.3	0.4	0.5
One	50	1.0	0.8	0.6	0.6	0.6	0.8	1.0	–
One	100	1.3	1.2	1.1	1.1	1.3	–	–	–
One	150	1.4	1.3	1.4	–	–	–	–	–
One	205	1.6	–	–	–	–	–	–	–
Two	10	0.1	*	*	*	*	*	*	0.1
Two	25	0.1	*	*	*	*	*	*	0.1
Two	50	0.2	*	*	*	*	*	0.2	–
Two	100	0.3	0.2	0.1	0.2	0.3	–	–	–
Two	150	0.3	0.2	0.3	–	–	–	–	–
Two	205	0.4	–	–	–	–	–	–	–

*Below ±0.05°C

K GRADE

Number Of Trims	Temperature Span (°C)	Lowest Temperature In Span (°C)							
		-55	-25	0	+25	+50	+75	+100	+125
None	10	2.1	2.3	2.5	2.7	2.9	3.1	3.3	3.6
None	25	2.6	2.7	2.8	3.0	3.2	3.5	3.8	4.2
None	50	3.8	3.5	3.4	3.6	3.8	4.3	5.1	–
None	100	4.2	4.3	4.4	4.6	5.1	–	–	–
None	150	4.8	4.8	5.3	–	–	–	–	–
None	205	5.5	–	–	–	–	–	–	–
One	10	0.2	0.1	0.1	0.1	0.1	0.1	0.1	0.2
One	25	0.6	0.4	0.3	0.3	0.3	0.4	0.5	0.6
One	50	1.2	1.0	0.7	0.7	0.7	1.0	1.2	–
One	100	1.5	1.4	1.3	1.3	1.5	–	–	–
One	150	1.7	1.5	1.7	–	–	–	–	–
One	205	2.0	–	–	–	–	–	–	–
Two	10	0.1	*	*	*	*	*	*	0.1
Two	25	0.2	0.1	*	*	*	*	0.1	0.2
Two	50	0.3	0.1	*	*	*	0.1	0.2	–
Two	100	0.5	0.3	0.2	0.3	0.7	–	–	–
Two	150	0.6	0.5	0.7	–	–	–	–	–
Two	205	0.8	–	–	–	–	–	–	–

*Below ±0.05°C

J GRADE

Number Of Trims	Temperature Span (°C)	Lowest Temperature In Span (°C)							
		-55	-25	0	+25	+50	+75	+100	+125
None	10	4.2	4.6	5.0	5.4	5.8	6.2	6.6	7.2
None	25	5.0	5.2	5.5	5.9	6.0	6.9	7.5	8.0
None	50	6.5	6.5	6.4	6.9	7.3	8.2	9.0	–
None	100	7.7	8.0	8.3	8.7	9.4	–	–	–
None	150	9.2	9.5	9.6	–	–	–	–	–
None	205	10.0	–	–	–	–	–	–	–
One	10	0.3	0.2	0.2	0.2	0.2	0.2	0.2	0.3
One	25	0.9	0.6	0.5	0.5	0.5	0.6	0.8	0.9
One	50	1.9	1.5	1.0	1.0	1.0	1.5	1.9	–
One	100	2.3	2.2	2.0	2.0	2.3	–	–	–
One	150	2.5	2.4	2.5	–	–	–	–	–
One	205	3.0	–	–	–	–	–	–	–
Two	10	0.1	*	*	*	*	*	*	0.1
Two	25	0.2	0.1	*	*	*	*	0.1	0.2
Two	50	0.4	0.2	0.1	*	*	0.1	0.2	*
Two	100	0.7	0.5	0.3	0.7	1.0	–	–	–
Two	150	1.0	0.7	1.2	–	–	–	–	–
Two	205	1.5	–	–	–	–	–	–	–

*Below ±0.05°C

I GRADE

Number of Trims	Temperature Span (°C)	Lowest Temperature In Span (°C)							
		-55	-25	0	+25	+50	+75	+100	+125
None	10	8.4	9.2	10.0	10.8	11.6	12.4	13.2	14.4
None	25	10.0	10.4	11.0	11.8	12.0	13.8	15.0	16.0
None	50	13.0	13.0	12.8	13.8	14.6	16.4	18.0	–
None	100	15.2	16.0	16.6	17.4	18.8	–	–	–
None	150	18.4	19.0	19.2	–	–	–	–	–
None	205	20.0	–	–	–	–	–	–	–
One	10	0.6	0.4	0.4	0.4	0.4	0.4	0.4	0.6
One	25	1.8	1.2	1.0	1.0	1.0	1.2	1.6	1.8
One	50	3.8	3.0	2.0	2.0	2.0	3.0	3.8	–
One	100	4.8	4.5	4.2	4.2	5.0	–	–	–
One	150	5.5	4.8	5.5	–	–	–	–	–
One	205	5.8	–	–	–	–	–	–	–
Two	10	0.3	0.2	0.1	*	*	0.1	0.2	0.3
Two	25	0.5	0.3	0.2	*	0.1	0.2	0.3	0.5
Two	50	1.2	0.6	0.4	0.2	0.2	0.3	0.7	–
Two	100	1.8	1.4	1.0	2.0	2.5	–	–	–
Two	150	2.6	2.0	2.8	–	–	–	–	–
Two	205	3.0	–	–	–	–	–	–	–

*Below ±0.05°C

NOTES:

1. All accuracies excluding the 205°C span are guaranteed, not tested; the 205°C span accuracies are tested for by testing each device at -55°C, +25°C, +125°C and +150°C.

2. All one-trim accuracies excluding the 205°C span assume that the trim is made at the midpoint in the span; the 205°C span assumes a trim at +25°C.

3. All two-trim accuracies excluding the 205°C span assume that the trims are made at the endpoints of the span; the 205°C span assumes that trims are made at approximately 0°C and +140°C.

4. All accuracies exclude:
   a. Trim error in calibration technique used;
   b. Repeatability error;
   c. Long term drift errors.

In precision applications the actual errors encountered are usually dependent upon sources of error which are often overlooked in error budgets.

A) *Trim Error* is usually the largest error source. This error arises from such sources as:
   1. poor thermal coupling between the device to be calibrated and the reference sensor;
   2. reference sensor errors;
   3. device to be calibrated is not permitted to thermally settle;
   4. θ_{CA} is radically different for the trim and the application.

B) *Repeatability Errors* arise from a strain hysteresis of the package. The magnitude of this error is solely a function of the magnitude of the temperature span over which the device is used. For example, thermal shocks between 0°C and 100°C will result in an extremely low hysteresis, for a repeatability error of less than ±0.05°C. When the thermal shocks are widened to -55°C and +150°C, the device will typically exhibit a repeatability of ±0.05°C, with ±0.10°C maximum being guaranteed.

C) *Long term drift errors* are related to the average operating temperature and the magnitude of the thermal shocks experienced by the device. Extended use of the device at temperatures above 100°C typically results in a long term drift of ±0.03°C; the guaranteed maximum is ±0.10°C. Operating temperatures below 100°C induce no measurable drifts in the device. In addition to operating temperature the severity of the thermal shocks incurred will determine the absolute stability of the device. For thermal shock spans of less than 100°C, the drift is difficult to measure (<0.03°C). However, for 200°C spans the device may drift by as much as ±0.10°C after 20 such shocks. If severe, quick shocks are necessary in the application of the device, simulated life tests are recommended for a thorough evaluation of the error introduced by such shocks.

ANALOG DEVICES

AN-273
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Use of the AD590 Temperature Transducer in a Remote Sensing Application

by Paul Klonowski

INTRODUCTION

The AD590 is a two-terminal integrated circuit temperature transducer which produces an output proportional to absolute temperature. For supply voltages between +4V and +30V the device acts as a high impedance, constant current source supplying 1μA/K. Laser trimming of the chip's thin-film resistors is used to calibrate the device to an output of 298.2μA at 298.2K (+25°C).

A typical application for the AD590 is a remote temperature-to-current transducer. Figure 1 shows a thermometer circuit that measures temperature from −55°C to +100°C and whose output voltage is 100mV/°C. Since the AD590 measures absolute temperature (its nominal output is 1μA/K), the output must be offset by 273.2μA in order to read out in degrees Celsius. The output current of the AD590 flows through a 1kΩ resistance, developing a voltage of 1mV/K. The output of the AD580 2.5V reference is divided down by resistors to provide a 273.2mV offset, which is subtracted from the voltage across the 1kΩ resistor by an AD524 instrumentation amplifier. The amplifier provides a gain of 100, so that the output range corresponding to −55°C to +100°C is −5.5V to +10V (100mV/°C). An operational amplifier can substitute for the instrumentation amplifier, although care must be taken when designing with the op amp since the gain at the two input terminals will be different.

THE PROBLEM

A question often asked of Analog Devices Applications Engineers by customers using the AD590 in a remote temperature-to-current application is "How long can I make the cable and how can I eliminate any noise that the cable picks up?" Experiments were performed in an effort to provide some guidelines for answering this question using the circuit of Figure 1 with a 1000' shielded, though initially ungrounded, twisted pair cable (Belden 9461, style 2092). In order to duplicate actual conditions the experiments were performed in an industrial environment.

TYPES OF NOISE

There are three basic types of noise inherent in a data-acquisition system. The first type is *transmitted noise*: noise received with the original signal and indistinguishable from it. The second type is *intrinsic noise*: noise generated within the devices used in a circuit, e.g., resistors, op amps, etc. Included in this category are Johnson, shot, and popcorn noise. The third type is *induced noise*: noise picked up from the outside world and coupled into the circuit. This application note discusses methods of reducing induced noise, which is the only form of noise that can be influenced by choices of wiring and shielding.

19

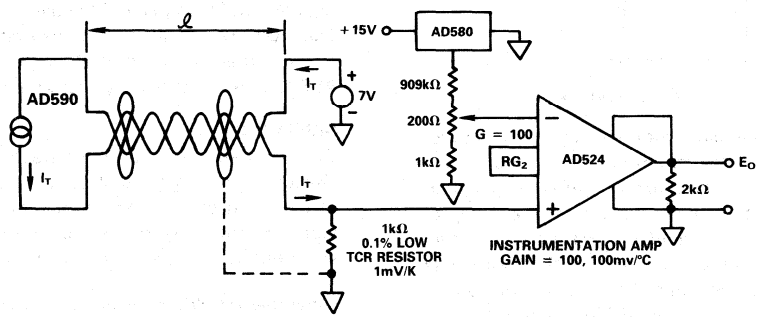

Figure 1. Thermometer Circuit

NOISE FACTORS

There are three elements involved in any noise problem. The first is the source of the noise. Possible noise sources include AM radio signals, logic signals, magnetic fields, and power line transients. The second element is the coupling medium. That is, how is the noise source entering the circuit? Possible coupling mediums include a common circuit impedance (Figure 2), stray capacitance (Figure 3) and mutual inductance (Figure 4). A brief description of each follows.

Common impedance noise is developed by an impedance common to several circuits. This might occur, as shown in Figure 2, when a pulse output source and an op amp's reference terminal are both connected to a "ground" point having tangible impedance to the power supply terminal. The noisy return current of Circuit 1 develops a voltage, V_{noise}, across the common impedance Z which will appear as a noise signal to Circuit 2. Possible solutions to this problem include proper circuits for distributing power and the use of isolation transformers and optical isolators.

Capacitively-coupled noise is produced by stray capacitance which couples the voltage changing noise source into high impedance circuits, as shown in Figure 3. The nature of the impedance Z determines the shape of the response. Methods of reducing capacitively-coupled noise include reducing the noise source, properly implementing shields, and reducing the stray capacitance.

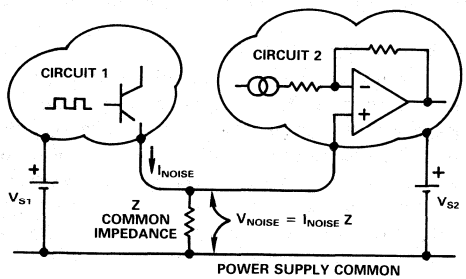

Figure 2. Common Impedance Noise

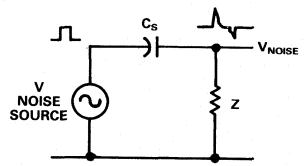

Figure 3. Stray Capacitance Noise

Magnetically coupled noise is produced by mutual inductance and can occur, for example, in an incorrectly shielded cable as shown in Figure 7. Figure 4 is a simple model of this incorrectly shielded cable, where L_s represents the inductance of the shield, L_c represents the inductance of one of the center conductors, and L_m represents the mutual inductance between the two. The noise current I(t) flows through L_s and establishes a magnetic flux; this time-varying flux also surrounds L_c and produces a

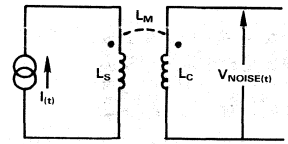

Figure 4. Mutual Inductance Noise

voltage $V_{noise(t)}$ proportional to the time rate of change of the current I (t) flowing through L_s. This voltage can be expressed as $V_{noise(t)} = L_m \frac{d\,I(t)}{dt}$

The third element involved in any noise problem is the receiver, or the circuit susceptible to the noise. It is important to understand the role of each of the three elements (the noise source, the coupling medium, and the receiver) in order to solve the noise problem.[1] In this experiment it was determined that the noise sources were 60Hz pick-up and AM radio signals, the coupling medium was stray capacitance, and the receiver was the AD524.

INITIAL NOISE EFFECT

The photograph in Figure 5 shows the output of the circuit in Figure 1 with the shield ungrounded and with the remote AD590 at 30°C. Ideally, the output should be a 3V (100mV/°C) dc signal. However, a 60Hz signal resulting from an electric field has been capacitively coupled into the circuit via the stray capacitance of the cable and then amplified by a gain of 100. Note, however, that the 60Hz signal is offset by a dc signal; when the output voltage of

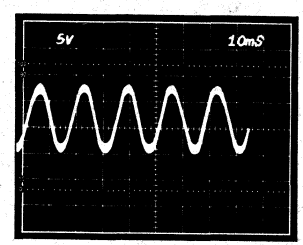

Figure 5. E_O (Figure 1) with Shield Ungrounded

the AD524 was measured with a dc voltmeter the value read was 3.0V. This is because the voltmeter read the value of the dc signal due to the AD590 along with the average value of the 60Hz sine wave noise signal. The average value of a sine wave is zero. In essence, notwithstanding the interferring signal the average value was correct. The accuracy of all the measurements were verified through the use of an RTD measurement system; the AD590 under test was physically attached to the RTD.

SHIELDING

One effective way to eliminate electrostatic noise is to use shielding. A charge resulting from an external potential cannot exist on the interior of a closed conducting surface. A shield is, in effect, a closed conducting surface that surrounds the twisted pair of wires within the cable.

[1]An excellent article dealing with this subject is "Understanding Interference-Type Noise" by Alan Rich, found in Section 20 of the Analog Devices Databook.

In order for a shield to be effective it should be connected to the reference potential of any circuitry contained within the shield. If the signal is connected to earth ground, then the shield should be connected to earth ground (see Figure 6). No voltage should exist between the reference potential and the shield conductor.

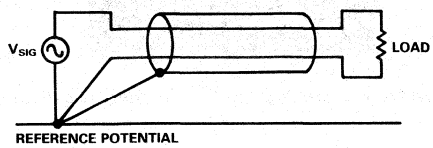

Figure 6. Correctly Shielded Cable

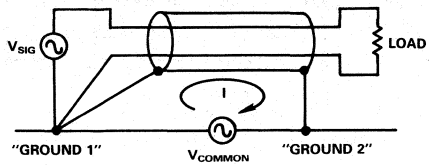

Figure 7. Incorrectly Shielded Cable

Only one end of the shield should be tied to "ground". Tying both ends of the shield to "ground" will produce a shield current equal to the difference in the potential of the two "grounds" divided by the series resistance of the shield (see Figure 7). As previously discussed, the mutual inductance between the shield and conductors will couple this noise current into the conductors in the form of a series voltage, V_{noise}.

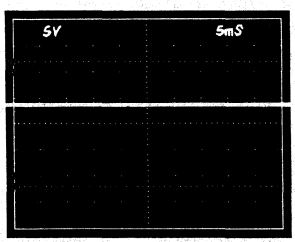

Figure 8. E_O with Shield Grounded

Although the 60Hz signal noise has been eliminated, voltage spikes are still visible in Figure 8. Figure 9, which is another view of the AD524 output with the scope ac coupled and the scale magnitude increased, shows the high frequency noise still being picked up by the cable. A look at Figure 10, which is a view of the signal being fed into the noninverting terminal of the AD524, shows that the noise being picked up is actually an AM radio signal.

RF NOISE

RF noise is a combination of an electric field and a magnetic field, or an electromagnetic field. This electromagnetic field extends into and beyond the space between conductors. That is, the electromagnetic field and hence the RF energy is "steered" by the conductors.

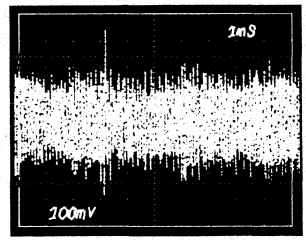

Figure 9. RF Noise at E_O

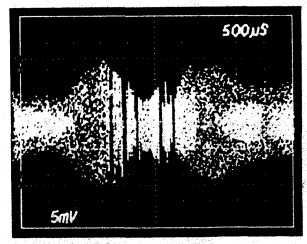

Figure 10. RF Noise at AD524 +Input

RF energy both enters and is reflected in a system wherever there is an impedance mismatch or discontinuity in the system. This includes discontinuities at the end of the signal run; at the AD590 end and the AD524 end of the cable. In order to eliminate these discontinuities it would be necessary to rf shield the entire circuit system, perhaps using conduit and metal boxes. In most cases however, this is rather impractical.

What is usually sufficient for most systems is to provide rf filtering using passive components at the critical point of interest. It is important to note that this filter may not eliminate the rf energy but it may just reflect the energy and redistribute the problem. A circuit topology which ensures that none of the external circuitry is affected by the rf noise is discussed later.

BYPASS CAPACITORS

A bypass capacitor can be used to divert some of the high frequency noise current to ground. Figure 11 is a simple schematic that shows the effect of the bypass capacitor. Recalling that the reactance of a capacitor $X_c = 1/2\pi fC$ and assuming that we have the minimum radio carrier frequency of 550kHz, using a 0.33µF capacitor the impedance seen by the noise current is:

$$X_C = 1/2\pi\,(550 \times 10^3)\,(0.33 \times 10^{-6}) = 0.88\Omega$$

Of course the direct current supplied by the AD590 will be unaffected by the bypass cap and will continue to flow into the 1kΩ resistor. Only the high frequency noise current will be affected. Figure 12, which is a view of the signal at the AD524 noninverting terminal, shows the result of tying a 0.33µF capacitor across the 1kΩ"load" resistor with the shield still grounded. Figure 13 is a look at the output signal of the AD524 with the 0.33µF capacitor. Compare Figures 10 and 12, and Figures 9 and 13. The bypass capacitor reduces the noise magnitude by a factor greater than 5.

In this instance the radio carrier frequency was 550kHz. It's important not to ignore rf noise, even if the bandwidth of the external circuitry is smaller than the bandwidth of the rf noise. A large rf component will overload the inputs of the external circuitry and be detected, causing an apparent dc shift in the component's output signal.

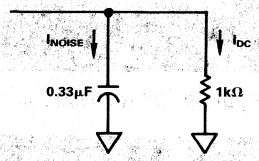

Figure 11. Bypass Caps Reduce RF Noise Effect

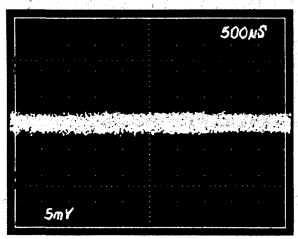

Figure 12. AD524 +Input with Bypass Cap

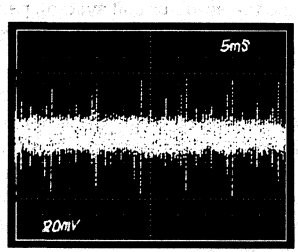

Figure 13. E_O with Bypass Cap

SERIES RESISTORS

Although an improvement, Figure 13 shows that noise is still entering the circuit. Another method of reducing the amount of noise seen at the noninverting terminal of the op amp is to limit the noise currents through the cable by adding 1kΩ resistors in series with the AD590.

This in effect forms a noise-voltage divider between the load impedance (the 1kΩ resistor and the 0.33μF capacitor) and the series resistors. Figure 14 shows the output of the AD524 with this final circuit configuration. The 10mV p-p amplitude of the output noise is actually 100 × the amplitude of the noise seen at the input of the AD524, which is 100μV p-p. Note also that the high frequency spikes in Figure 13 have been eliminated in Figure 14. Figure 15 is the schematic of the final circuit which has reduced the noise by a factor of 2000 over the circuit in Figure 1.

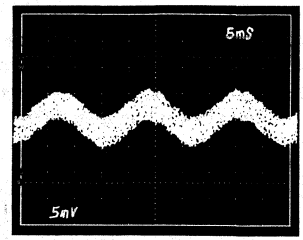

Figure 14. E_O with Bypass Cap and Series Resistors

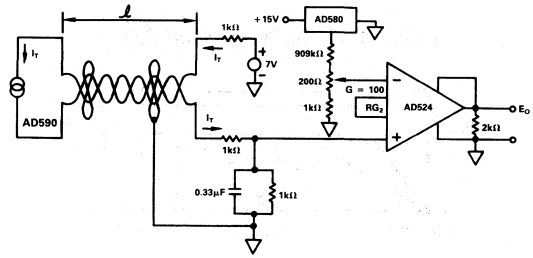

Figure 15.

CONCLUSION

In order to eliminate the effects of rf noise, the circuit of Figure 16 is recommended.

In Figure 16, the resistors and capacitors are placed at both ends of the cable. This ensures that the rf noise remains within the cable and does not affect the external circuitry. Further note that these resistors can be arbitrarily large as long as the voltage potential is large enough to supply the current (V = IR). The use of a bypass capacitor across the AD590 with the series resistors will filter rf signals. Theoretically the rf signal could be rectified by the AD590 and offset the accuracy of the device.

Using the techniques outlined above, noise and interference can be effectively eliminated through the use of shielded twisted pair cable and resistors and capacitors. Thus it is possible to drive the AD590 over 1000 feet of cable without a loss of accuracy.

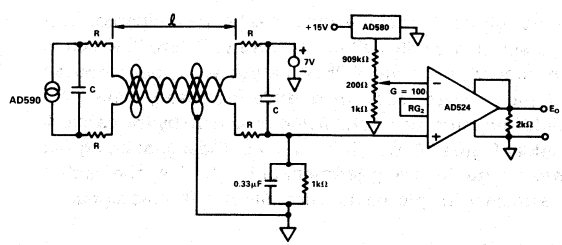

Figure 16.

Sigma-Delta Converters
Contents

		Page
AN-283	Sigma-Delta ADCs and DACs	20–3

20

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Sigma-Delta ADCs and DACs

SIGMA-DELTA OVERVIEW

Within the last several years, the sigma-delta architecture has become more and more popular for realizing high-resolution ADCs in mixed-signal VLSI processes. Until recently, however, the process technology needed to make these devices commercially viable has not been available. Now that 1 micron and smaller CMOS geometries are manufacturable, sigma-delta converters will become even more prolific in certain types of applications, especially mixed-signal ICs which combine the ADC, DAC, and DSP functions on a single chip.

Conceptually, the sigma-delta architecture is more digital than analog intensive. This does not, however, minimize the importance of the analog portion of the sigma-delta ADC. The design of a fifth-order sigma-delta modulator (as in the AD1879 dual 18 bit ADC) is certainly not a trivial matter, and neither is the digital filter. The sigma-delta converter is inherently an oversampling converter, although oversampling is just one of the techniques contributing to the overall performance. Basically, a sigma-delta converter digitizes an analog signal with a very low resolution (1 bit) ADC at a very high sampling rate. By using oversampling techniques in conjunction with noise shaping and digital filtering, the effective resolution is increased. Decimation is then used to reduce the effective sampling rate at the ADC output. The sigma-delta ADC exhibits excellent differential and integral linearity due to the linearity of the 1 bit quantizer and DAC, and no trimming is required as in other ADC architectures.

The key concepts involved in understanding the operation of sigma-delta converters are oversampling, noise shaping (using a sigma-delta modulator), digital filtering, and decimation.

SIGMA-DELTA CONCEPTS

- **Ideal Topology for Mixed Signal VLSI Chips**
- **Oversampling**
- **Noise-Spectrum Shaping Using Sigma-Delta Modulator**
- **Digital Filtering**
- **Decimation**
- **16-Bits and Higher Resolution Possible**

Figure 6.1

OVERSAMPLING

The concept of oversampling has been previously discussed in Section III, and is illustrated again in Figure 6.2 and 6.3. As was

NYQUIST SAMPLING WITH ANALOG LOWPASS FILTER

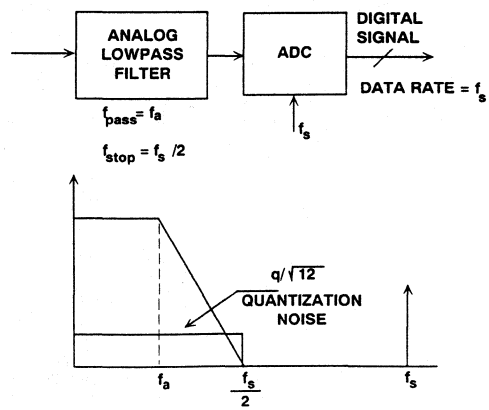

Figure 6.2

20

OVERSAMPLING WITH ANALOG AND DIGITAL FILTERING

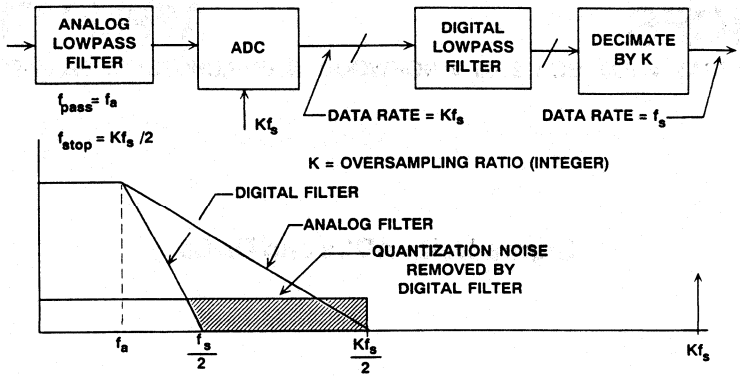

Figure 6.3

discussed, one significant benefit of oversampling is that the rolloff requirements on the analog antialiasing filter are relaxed. The quantization noise (rms value over Nyquist bandwidth is $q/\sqrt{12}$, where q is the weight of the LSB) which falls between $f_s/2$ and $kf_s/2$ is removed from the output by the digital filter (k is the oversampling ratio). This has the effect of increasing the overall signal-to-noise ratio by an amount equal to $10\log_{10}(k)$. Unfortunately this is a high price to pay for extra resolution, as an oversampling ratio of 4 is required just to increase the signal-to-noise ratio by a modest 6dB (1 bit). To keep the oversampling ratio within reasonable bounds, it is possible to shape the frequency spectrum of the quantization noise so that the majority of the noise lies between $f_s/2$ and $kf_s/2$, and only a small portion is left between dc and $f_s/2$. This is precisely what a sigma-delta modulator does in a sigma-delta ADC. After the noise spectrum is shaped by the modulator, the digital filter can then remove the bulk of the quantization noise energy, and the overall signal-to-noise ratio (hence the dynamic range) is dramatically increased.

SIGMA-DELTA MODULATORS AND QUANTIZATION NOISE SHAPING

A block diagram of a first-order sigma-delta ADC is shown in Figure 6.4. The first part of the converter is the sigma-delta modulator which converts the input signal into a continuous serial stream of 1's and 0's at a rate determined by the sampling clock

FIRST-ORDER SIGMA-DELTA ADC

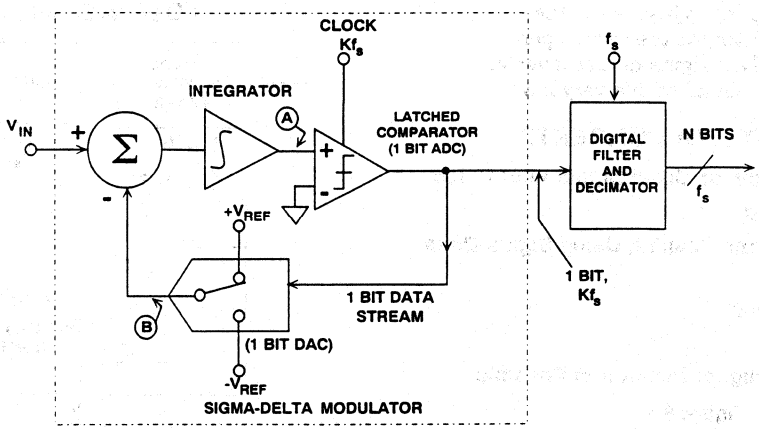

Figure 6.4

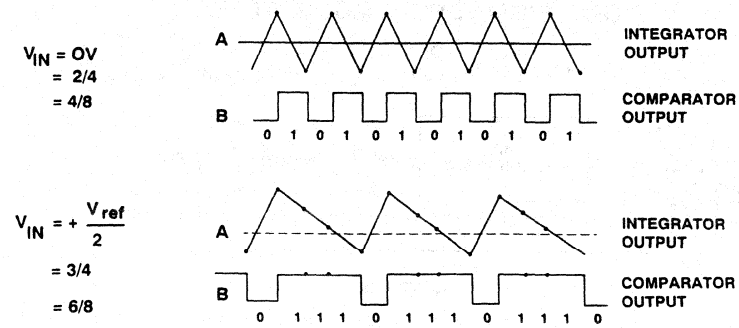

Figure 6.5

frequency, kf_s. The 1-bit DAC is driven by the serial output data stream, and the DAC output is subtracted from the input signal. Feedback control theory tells us that the average value of the DAC output (hence the serial bit stream) must approach that of the input signal if the loop has enough gain. The integrator can be represented in the frequency domain by a filter whose amplitude response is proportional to 1/f, where f is the input frequency. Since the chopper-like action of the clocked, latched comparator converts the input signal to a high-frequency ac signal, varying about the average value of the input, the effective quantization noise at low frequencies is greatly reduced (the integrator looks like a high-pass filter to quantization noise). The exact frequency spectrum of the resulting noise depends on the sampling rate, the integrator time constant, and the precise span of the voltage fed back.

For any given input value in a single sampling interval, the data from the 1-bit ADC is virtually meaningless. Only when a large number of samples are averaged, will a meaningful value result. The sigma-delta modulator is very difficult to analyze in the time domain because of this apparent randomness of the single-bit data output. If the input signal is near positive fullscale, it is clear that there will be more 1's than 0's in the bit stream. Likewise, for signals near negative fullscale, there will be more 0's than 1's in the bit stream. For signals near midscale, there will be approximately an equal number of 1's and 0's. Figure 6.5 shows the output of the integrator for two input conditions. The first is for an input of zero (midscale). To decode the output, pass the output samples through a simple digital lowpass filter that averages every four samples. The output of the filter is 2/4. This value represents bipolar zero. If more samples are averaged, more dynamic range is achieved. For example, averaging 4 samples gives 2 bits of resolution, while averaging 8 samples yields 4/8, or 3 bits of resolution. In the bottom waveform of Figure 6.5, the average obtained for 4 samples is 3/4, and the average for 8 samples is 6/8.

The sigma-delta ADC can also be viewed as a synchronous voltage-to-frequency converter followed by a counter. If the number of 1's in the output data stream is counted over a sufficient number of samples, the counter output will represent the digital value of the input. Obviously, this method of averaging will only work for dc or very slowly changing input signals. In addition, 2^N clock cycles must be counted in order to achieve N-bit effective resolution, thereby severely limiting the effective sampling rate.

Further analysis of the sigma-delta architecture is best done in the frequency domain using the linear model shown in Figure 6.6. Note that the integrator is represented as an analog filter with a given transfer function H(f). The transfer function has an amplitude response which is inversely proportional to the input frequency. The quantizer is modeled as a gain stage followed by the addition of quantization noise. One of the advantages of using frequency domain analysis is that algebra can be used to describe the signals. The output value y can be represented as the difference x - y from the summing node at the input multiplied by the transfer function of the analog filter (integrator), multiplied by the gain block, and then added with the quantization noise Q. If we set the gain to 1, and the transfer function is represented as 1/f, the following mathematical relationship results:

$$y = \frac{x - y}{f} + Q \text{ , or by rearranging,}$$

$$y = \frac{x}{f+1} + \frac{Qf}{f+1}$$

FREQUENCY DOMAIN LINEARIZED MODEL OF A SIGMA DELTA MODULATOR

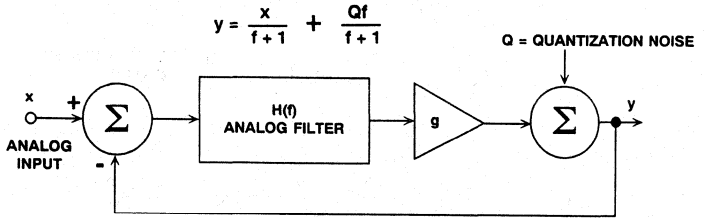

Figure 6.6

Note that as frequency f approaches 0, the output approaches x with no noise component. At higher frequencies, the value of x is reduced, and the value of the noise component is increased. For high frequency inputs, the output consists primarily of quantization noise. In essence, the analog filter has a low pass effect on the signal and a high pass effect on the noise component. For this reason, the analog filter of the modulator can be viewed as a noise shaping filter as shown in Figure 6.7.

SHAPED QUANTIZATION NOISE DISTRIBUTION

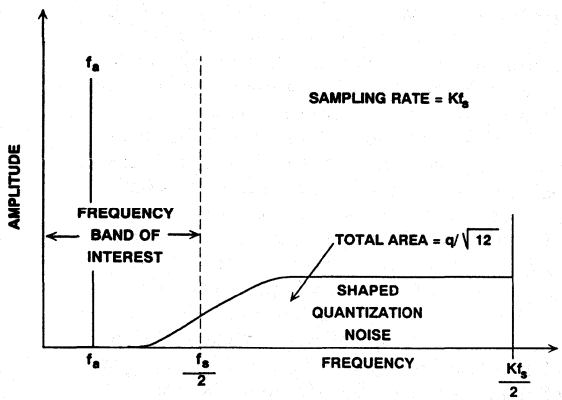

Figure 6.7

As with analog filters in general, higher order filters offer better performance. This is also true of the sigma-delta modulator, provided certain precautions are taken. A second order sigma-delta modulator is shown in Figure 6.8, and a comparison between the noise shaping functions is shown in Figure 6.9. Figure 6.10 shows a plot of the corresponding in-band signal-to-noise ratio (dynamic range) as a function of the oversampling ratio for a first and second order modulator. Note that the first order transfer function has a slope of 9dB per octave, while the second order transfer function slope is 15dB per octave. Higher order modulators (greater than second order) can realize even better performance, but the simple linear model must be used with great care, and sophisticated design techniques are required in order to insure stability. The curve shown in Figure 6.10 for the third-order loop represents an unrealizable condition and is shown for reference only.

The curves in Figure 6.10 can be used to determine the approximate ADC resolution achievable, given the modulator order and the oversampling rate. For instance, if the oversampling rate is 64x, an ideal second order system is capable of providing a signal to noise ratio of about 80dB. This implies an ADC resolution of approximately 13 bits. Although the filtering done by the digital filter can be done to any degree of precision

SECOND-ORDER SIGMA-DELTA ADC

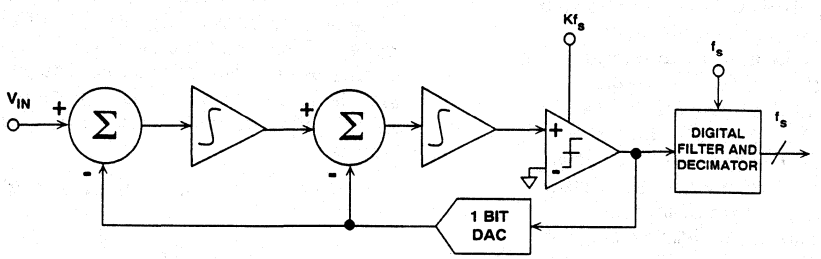

Figure 6.8

desirable, it would be pointless to carry more than 13 binary bits to the outside world. Additional bits would carry no useful signal information, and would be buried in noise.

FIRST AND SECOND-ORDER NOISE SHAPING FUNCTIONS

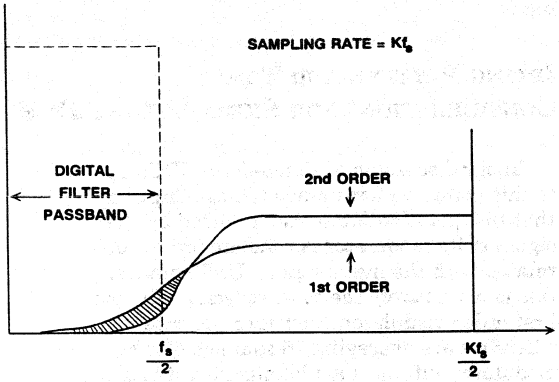

Figure 6.9

SNR VERSUS OVERSAMPLING RATIO FOR FIRST, SECOND, AND THIRD-ORDER LOOPS

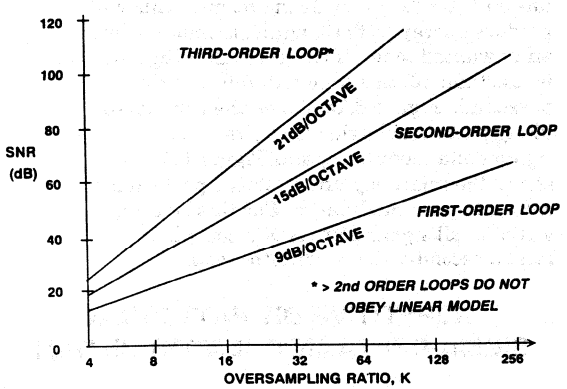

Figure 6.10

DIGITAL FILTERING AND DECIMATION

After the quantization noise has been shaped by the modulator and pushed into the frequencies above the band of interest, digital filtering techniques can be applied to this shaped quantization noise as shown in Figure 6.11. The purpose of the digital filter is twofold. First, it must act as an antialiasing filter with respect to the final sampling rate, f_s. Second, it must filter out the higher frequency noise produced by the noise-shaping process of the sigma-delta modulator.

EFFECTS OF DIGITAL FILTERING ON SHAPED QUANTIZATION NOISE

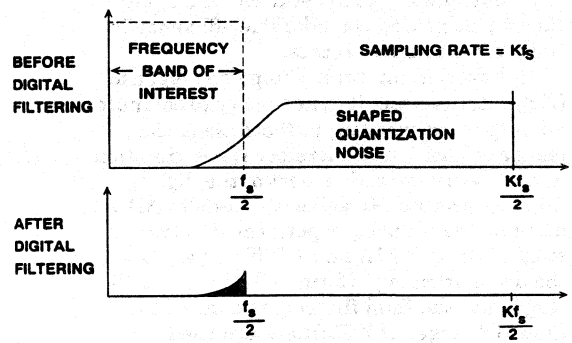

Figure 6.11

The final data rate reduction is performed by digitally resampling the filtered output using a process called decimation. The decimation of a discrete-time signal is shown in Figure 6.12, where the sampling rate of the input signal x(n) is at a rate which is to be reduced by a factor of 4. The signal is resampled at the lower rate (the decimation rate), s(n). Decimation can also be viewed as the method by which the redundant signal information introduced by the oversampling process is removed.

In sigma-delta ADCs it is quite common to combine the decimation function with the digital filtering function. This results in an increase in computational efficiency if done correctly.

Recall that a finite impulse filter (FIR) simply computes a moving weighted average (the weighting being determined by the individual filter coefficients) of the input samples. Normally, there is one filter output for every input sample. If, however, we wish to decimate the filter output by digitally

DECIMATION OF A DISCRETE-TIME SIGNAL

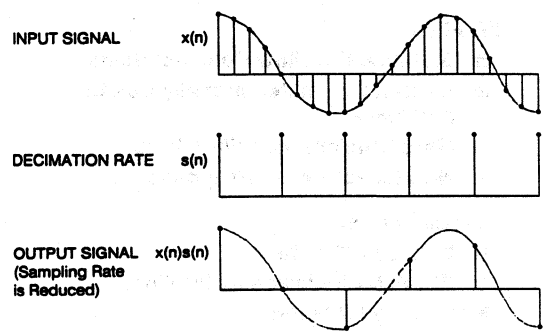

Figure 6.12

resampling at a lower rate, it is no longer necessary to compute a filter output for every input sample. Instead, we only compute filter outputs at the lower decimation rate, thereby achieving considerable efficiency in the computational process.

If, however, an infinite impulse response (IIR) filter is used, it is necessary to compute an output for every input (because of the feedback term), and therefore the decimation cannot be performed as part of the digital filtering process. In some sigma-delta ADC designs, the filtering is performed in two stages. If both FIR and IIR filters are used, the decimation is performed in the first FIR stage, and the final filtering is done in the final IIR stage. If FIR filters are used for both stages, it is usually more efficient to split the decimation between the two filter stages.

From the above discussion it should be clear that the design of a sigma-delta ADC digital filter involves many tradeoffs. FIR filters lend themselves to decimation, are always stable, and have linear phase characteristics (extremely important in audio and some telemetry applications). Although they are typically easier to design, they usually require more stages to realize a given transfer characteristic than a corresponding IIR filter. On the other hand, the IIR filter employs feedback which eliminates the possibility of decimation within the filter, but makes the filter more efficient (better filter performance with fewer calculations). The feedback used in IIR filters can lead to a potentially unstable filter implementation. Also, the IIR filter (which will closely emulate filter functions realized in the analog domain) exhibits non-linear phase characteristics. Because of the stability issues and the quantization effects in the feedback loop, IIR filters are more complicated to design correctly.

SIGMA-DELTA ADC DIGITAL FILTERING AND DECIMATION

FIR Filters:
- Easy to Design
- Easy to Incorporate Decimation
- Linear Phase Response
- Large Number of Coefficients May Be Required

IIR Filters:
- Stability, Overflow Considerations
- Cannot Decimate Internally Due to Feedback
- More Efficient than FIR Filters
- Non-Linear Phase Response

Combinations:
- 2-Stage FIR Filters
- FIR Filter Followed by IIR Filter
- 2-Stage IIR Filters

Figure 6.13

Idling Pattern and Tonal Considerations for Sigma-Delta ADCs

In our discussion of sigma-delta ADCs up to this point, we have made the assumption that the quantization noise produced by the sigma-delta modulator is random and uncorrelated with the input signal. Unfortunately, this is not entirely the case, especially for the first-order modulator. Consider the case where we are averaging 16 samples of the modulator output in a 4 bit sigma-delta ADC. Figure 6.14 shows the bit pattern for two input signal conditions: an input signal having the value 8/16, and an input signal having the value 9/16. In the case of the 9/16 signal, the modulator output bit pattern has an extra "1" every 16th output. This will produce energy at $f_s/16$, which translates into an unwanted tone. If the oversampling ratio is less than 16, this tone will fall into the passband. Figure 6.15 shows the correlated idling pattern behavior for a first order sigma-delta modulator, and Figure 6.16 shows the relatively uncorrelated pattern for a second-order modulator. For this reason, virtually all sigma-delta ADCs contain at least a second-order modulator loop.

REPETITIVE BIT PATTERN IN SIGMA-DELTA MODULATOR OUTPUT

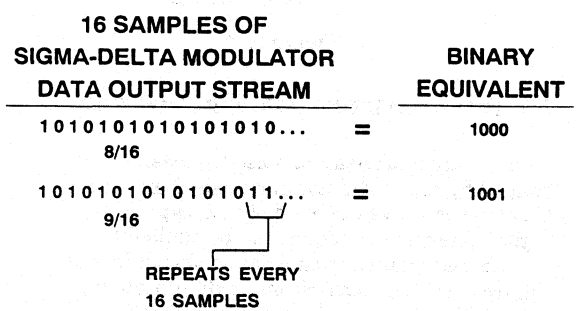

Figure 6.14

IDLING PATTERNS FOR FIRST-ORDER SIGMA-DELTA MODULATOR (INTEGRATOR OUTPUT)

)LE BEHAVIOR WITH 0 VOLTS INPUT

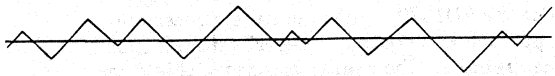

)LE BEHAVIOR WITH DC INPUT

Figure 6.15

DLING PATTERNS FOR SECOND-ORDER SIGMA-DELTA MODULATOR (SECOND INTEGRATOR OUTPUT)

DLE BEHAVIOR WITH 0 VOLTS INPUT

IDLE BEHAVIOR WITH DC INPUT SHOWING CORRELATED IDLING PATTERN

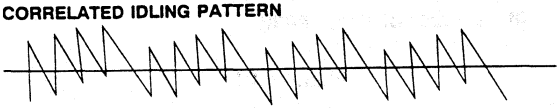

Figure 6.16

Higher Order Modulator Loops

In order to achieve wide dynamic range, sigma-delta modulator loops greater than second-order are necessary, but present real design challenges. First of all, the simple linear models previously discussed are no longer fully accurate. Loops of order greater than two are generally not guaranteed to be stable under all input conditions. The instability arises because the comparator is a nonlinear element whose effective "gain" varies inversely with the input level. This mechanism for instability causes the following behavior: if the loop is operating normally, and a large signal is applied to the input that overloads the loop, the average gain of the comparator is reduced. The reduction in comparator gain in the linear model causes loop instability. This causes instability even when the signal that caused it is removed. In actual practice, such a circuit would normally oscillate on power-up due to initial conditions caused by turn-on transients.

Instability in the AD1879 fifth-order modulator is sensed digitally by counting the number of consecutive ones or zeros in the modulator bit stream. A sufficiently long string of either ones or zeros indicates modulator instability. This triggers circuitry which resets the state in the integrators to put the modulator into a stable operating condition.

HIGHER ORDER LOOP CONSIDERATIONS (>2)

- **Increased Dynamic Range and Resolution is Achievable**
- **Higher Order Loops Minimize Idling Patterns and Tones**
- **Difficult to Analyze and Stabilize**
- **Non-Linear Stabilization Techniques Can Be Used Successfully: AD1879 18-Bit, 5th Order ADC**

Figure 6.17

Description of AD1879 18 Bit Sigma-Delta Audio ADC

The AD1879 is a state-of-the-art dual 18 bit sigma-delta ADC designed to meet the stringent requirements of professional digital audio. A block diagram of the device is shown in Figure 6.18, and performance specifications are given in Figure 6.19. The modulator is a fifth-order switched capacitor design which shapes the noise spectrum as shown in Figure 6.20. The oversampling ratio is 64x, which places the oversampling frequency at 3.072MHz for the standard audio sampling rate of 48kHz. Because of the high oversampling ratio, a single-pole analog antialiasing filter is sufficient at the input of the ADC.

AD1879 DUAL 18-BIT SIGMA-DELTA ADC

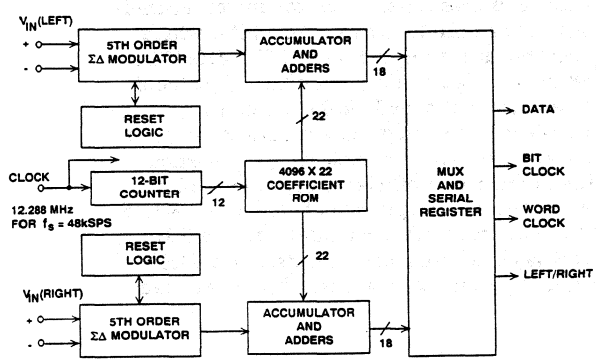

Figure 6.18

AD1879 18-BIT SIGMA-DELTA ADC KEY SPECIFICATIONS

- **Two 18-Bit Channels for Stereo Digital Audio**
- **Interchannel Crosstalk: –110dB at 1kHz**
- **SNR:104dB**
- **THD: 100dB**
- **Oversampling Ratio: 64x**
- **Output Word Rate: 55kHz Maximum**
- **Linear Phase Digital Filter**
- **Power: 900mW**
- **28-Pin, 600-mil Plastic Package**

Figure 6.19

AD1879 MODULATOR OUTPUT SPECTRUM

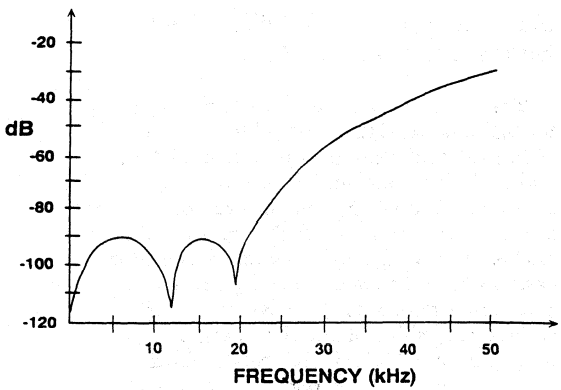

Figure 6.20

For audio ADCs such as the AD1879, the digital lowpass filter cannot be implemented using standard multiply-accumulate structures and present-day semiconductor technology. For example, we require a filter which operates at a sample rate of 3.072MHz (64 x 48kHz), is flat to 20kHz and has a stopband attenuation of over 115dB starting at 26.2kHz. If we plug these requirements into a standard FIR equiripple design program, the number of coefficients required is 4096. At an output sample rate of 48kHz, we would require a multiply-accumulate time of 5.1ns. This is clearly too fast for a standard FIR filter structure to implement because of semiconductor process limitations. For this reason, we must use either a parallel processing approach where more than one multiply-accumulate is being executed at any one time, or a multi-rate approach where the decimation is done in more than one step. For the AD1879, a novel parallel processing approach was chosen as described further in Reference 1. The characteristics of this filter are given in Figure 6.21, and the amplitude response in Figure 6.22.

AD1879 DIGITAL FILTER CHARACTERISTICS

- **Stopband Attenuation: 118dB**
- **Passband Ripple: ±0.0008dB**
- **Cutoff Frequency (48kHz output rate): 21.7kHz**
- **Stopband Frequency (48kHz output rate): 26.2kHz**
- **Number of Parallel Accumulators: 64 27-Bit Accumulators**
- **Coefficient Wordlength: 22-Bits**
- **Number of Taps: 4096**

Figure 6.21

The AD1879 ADC is a compound monolithic IC. One chip performs the sigma-delta modulation function, while the second chip performs the digital filtering.

AD1879 DIGITAL FILTER RESPONSE

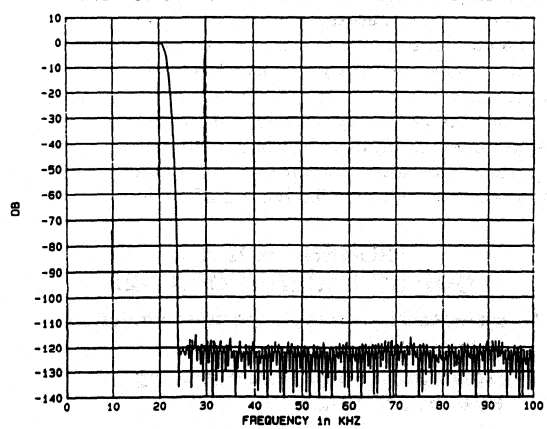

Figure 6.22

SIGMA-DELTA ADCS FOR LOW FREQUENCY MEASUREMENT APPLICATIONS

Applications such as industrial process control, weigh scales, temperature and pressure measurement instruments require ADCs which can digitize low frequency signals (usually less than 10Hz) to 16 bit or higher precision. In the past, this need was filled almost exclusively by integrating (or dual-slope) ADCs. Sigma-delta converters offer an attractive alternative. In addition to reduced cost and size, low frequency sigma-delta ADCs offer on-board digital filtering as well as system and self-calibration functions. Sampling rates allow processing of signals of up to 10Hz bandwidths, and power line frequency rejection of sigma-delta ADCs can be maintained over a much wider range of frequency variation than with traditional integrating ADCs. Power supply rejection of dual slope ADCs depends on the instantaneous line frequency variations since the sampling clock is synchronized to the line.

A functional block diagram of the AD7701 monolithic 16 bit sigma-delta ADC is shown in Figure 6.23 and key specifications in Figure 6.24.

AD7701 BLOCK DIAGRAM

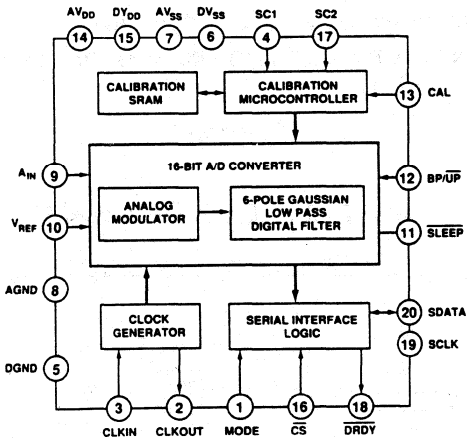

Figure 6.23

The AD7701 contains a second-order sigma-delta modulator which samples the analog input signal at a 16kHz rate when the external clock frequency is 4.096MHz. The quantization noise is therefore spread over the bandwidth 0 to 8kHz. The device contains a 6-pole gaussian lowpass digital filter which has a cutoff frequency of 10Hz at the maximum clock rate. The 16kHz sampling rate therefore represents an oversampling ratio of 800 with respect to the 10Hz cutoff frequency. The filter provides 55dB of 60Hz rejection under these conditions. If the clock frequency is halved to give a 5Hz cutoff, 60Hz rejection is better than 90dB. Power supply rejection is 70dB in the 0.1 to 10Hz bandwidth, and PSRR at 60Hz exceeds 120dB due to the digital filter. The frequency response of the digital filter at various clock rates is shown in Figure 6.25.

AD7701
LOW FREQUENCY MEASUREMENT ADC KEY SPECIFICATIONS

- **Monolithic 16-Bit ADC**
- **0.0015% Linearity Error**
- **4 kSPS Output Data Rate**
- **Programmable Low Pass Filter: 0.1Hz to 10Hz Corner Frequency**
- **On-Chip Self-Calibration Circuitry**
- **0 to +2.5V or ±2.5V Input Range**
- **40mW Power Dissipation**
- **20µW Standby Mode**
- **Flexible Serial Interface**

Figure 6.24

AD7701 DIGITAL FILTER RESPONSE

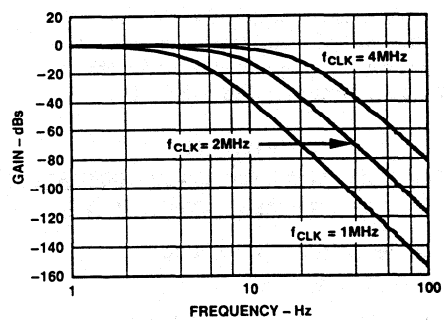

Figure 6.25

The long settling time of the internal digital filter (shown in Figure 6.26) in the AD7701 limits its use in multiplexed applications where channels are switched and converted sequentially at high rates. Switching between channels which may have different signal levels can cause a step change in the input. The AD7701is primarily intended for distributed converter systems using one ADC per channel. Multiplexing is possible, provided that sufficient settling time is allowed before data for the new channel is

AD7701 DIGITAL FILTER STEP RESPONSE

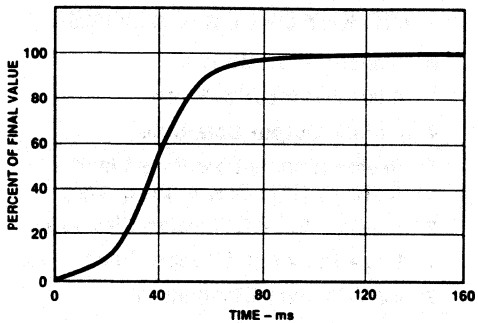

Figure 6.26

accessed. The gaussian response worst case settling time to ±0.0007% (±0.5LSB) is 125ms with a 4.096MHz master clock frequency.

The AD7701 offers two calibration modes using the on-chip calibration microcontroller and SRAM. In the self-calibration mode, zero-scale is calibrated against the analog ground pin (AGND), and fullscale is calibrated against V_{ref} pin. In the system-calibration mode, the AD7701 calibrates its zero and fullscale to voltages present on the analog input pin in two sequential steps, thereby allowing system offests and/or gain errors to be nulled out.

The AD7703 sigma-delta ADC has a similar architecture to the AD7701, but achieves 20 bits of resolution with 0.0003% linearity error. Key specifications for the AD7703 are summarized in Figure 6.27.

AD7703 LOW FREQUENCY MEASUREMENT ADC KEY SPECIFICATIONS

- **Monolithic 22-Bit ADC**
- **0.0003% Linearity Error**
- **4 kSPS Output Data Rate**
- **Programmable Low Pass Filter: 0.1Hz to 10Hz Corner Frequency**
- **On-Chip Self-Calibration Circuitry**
- **0 to +2.5V or ±2.5V Input Range**
- **40mW Power Dissipation**
- **20μW Standby Mode**
- **Flexible Serial Interface**

Figure 6.27

The AD7710, AD7711, and AD7712 ADCs constitute a family of 21 bit sigma-delta ADCs with on-chip signal conditioning for low frequency, low level measurement applications such as weigh scales, thermocouple temperature measurements, RTD (resistance temperature detector) temperature measurement, process controllers, and programmable loop controllers. Common features of the three devices are summarized in Figure 6.28, and a block diagram of the AD7110 is shown in Figure 6.29.

AD7710/7711/7712 MEASUREMENT ADC COMMON KEY FEATURES

- **21-Bit Sigma-Delta ADCs, ±0.0015% nonlinearity**
- **On-Board Differential Input PGA, Gains from 1 to 128 120dB CMR at 50, 60Hz**
- **First Filter Notch Frequency and Output Data Rate Programmable from 10Hz to 1kHz**
- **Sinc³ Filter Response with Cutoff Frequency of 0.262 times First Filter Notch Frequency**
- **Ability to Read /Write Calibration Coefficients**
- **Bidirectional Microcontroller Serial Interface**
- **Internal/External Reference Option**
- **Single or Dual Supply Operation**
- **Low Power (20mW) with Power-Down Mode (10μW)**

Figure 6.28

AD7710 BLOCK DIAGRAM

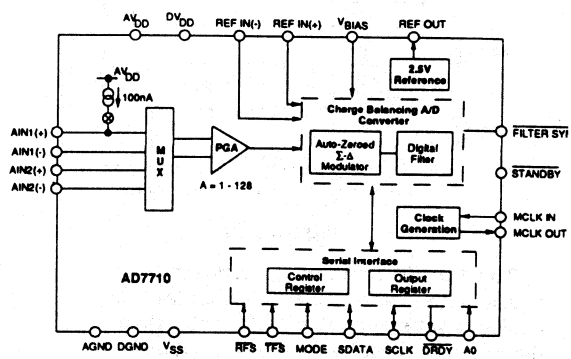

Figure 6.29

AD7710/7711/7712 DIGITAL FILTER RESPONSE

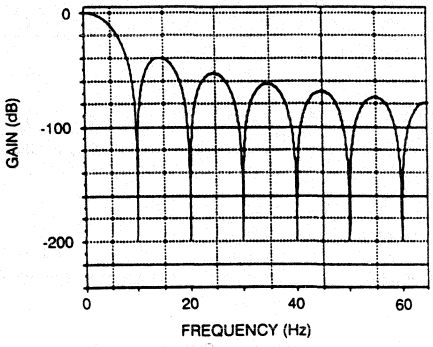

Figure 6.30

An on-board differential input PGA (gain = 1 to 128) enables the user to control fullscale voltage and voltage resolution. The effects of temperature drift are minimized by on-chip self-calibration which removes zero scale and fullscale errors. The internal digital filter has a $(\sin x/x)^3$ response, and 12 bits of data programmed into the control register determine the filter cutoff frequency, the position of the first notch of the filter, and the data rate. In association with the gain selection, it also determines the useful resolution of the device. The first notch frequency (which is also the output data rate) can be programmed from 10Hz to 1kHz. The corresponding -3dB frequency is equal to 0.262 times the first notch frequency. Figure 6.30 shows the filter frequency response for a cutoff frequency of 2.62Hz which corresponds to a first filter notch frequency of 10Hz. The filter response provides greater than 100dB of 50Hz and 60Hz common mode rejection.

Key device-specific features for the three devices are given in Figures 6.31.

AD7710/7711/7712 DEVICE-SPECIFIC FEATURES

AD7710:

- **Two-Channel Differential Low-Level PGA Input**

AD7711:

- **Single-Channel Differential Low-Level PGA Input**
- **RTD (Resistance Temperature Detector) Excitation Current Sources**

AD7712:

- **Single-Channel Differential Low-Level PGA Input**
- **High-Level Analog Input**

Figure 6.31

SIGMA-DELTA DACs

Sigma-delta D/A conversion can generally be thought of as the A/D conversion process in the reverse order, where all the basic functions of the digital filter and sigma-delta modulator previously discussed are the same. Sigma-delta DACs offer essentially the same advantages as sigma-delta ADCs. Because of the large oversampling ratio, the requirements on the antialiasing reconstruction filter are greatly relaxed. However, care must be taken to make sure the high frequency noise components contained in the one-bit DAC output are filtered sufficiently. If a higher order filter is required to reduce this noise, then some of the advantages of the sigma-delta DAC architecture are lost.

Accurate, low-cost, high resolution laser wafer trimmed DACs are readily available, and for this reason there has been less pressure to fully exploit sigma-delta DACs at the component level. The real incentive for developing the sigma-delta DAC technology is because it is the ideal architecture for mixed-signal ICs which require the chip-level integration of the ADC, DAC, and DSP functions.

SIGMA-DELTA DAC CONCEPTS

- **Basically a Sigma-Delta ADC in Reverse**
- **Low-Cost, High Resolution R/2R DACs Proliferate at the Component Level with Oversampling Capability**
- **Sigma-Delta DACs Ideal for Chip-Level Integration with ADC and DSP Functions**
- **Antialiasing Filter Must Remove High Frequency Noise**

Figure 6.32

The traditional approach to achieving high performance and wide dynamic range using R/2R-based DACs is shown in Figure 6.33. Due to the binary nature of the internal DAC switches, code-dependent transients, or glitches, typically produce some amount of harmonic distortion in the output spectrum. As discussed previously in the DAC section of this seminar, a technique called segmentation can greatly minimize these effects. For the ultimate in spectral purity, the remaining glitches can be removed with a sample-and-hold circuit which holds the DAC output voltage for the duration of the glitch. This technique can eliminate the code-dependent glitches (hence

CONVENTIONAL DAC DEGLITCHING

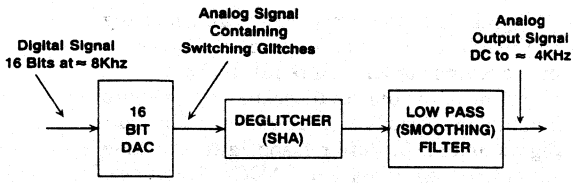

Figure 6.33

harmonic distortion) at the expense of introducing some additional energy at the sampling frequency. A lowpass, or smoothing filter is required at the output of the SHA to prevent aliasing as well as eliminate the energy at the sampling rate. The same basic considerations used to define the antialiasing filter used ahead of an ADC apply to the smoothing filter which follows the DAC. For this reason, oversampling relaxes the smoothing filter rolloff requirements in a similar manner. In fact, 2x, 4x, and 8x oversampling techniques are currently in widespread use in compact disk players which use conventional R/2R 16, 18, and 20 bit DACs.

The main elements used to implement a sigma-delta DAC are shown in Figure 6.34.

SIGMA-DELTA DAC

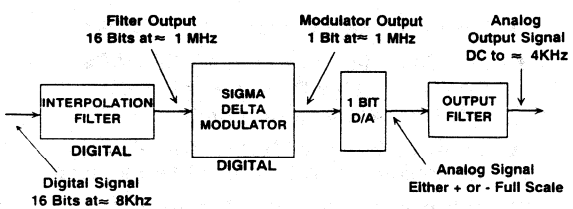

Figure 6.34

The example shown here is for a 16bit DAC which is updated at an 8kHz rate to produce a voiceband output signal having a bandwidth of 4kHz. The 16 bit digital word is fed to a digital interpolation filter where the sampling rate is increased to 1.024MHz, corresponding to an oversampling ratio of 128. This process can be viewed as the reconstruction of a new, higher rate digital signal from an older, lower rate digital signal. Figure 6.35 shows the interpolation of a discrete time signal by a factor of 4. The input signal x(m) is expanded by inserting three zero-valued samples between data samples. The resulting signal w(m) is lowpass filtered to produce y(m) whose sample rate is increased by a factor of 4.

The digital-input sigma-delta modulator noise-shapes the 16-bit 1.024MHz data

INTERPOLATION OF DISCRETE-TIME SIGNAL

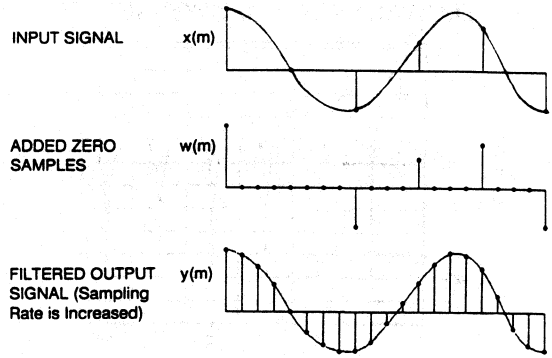

Figure 6.35

stream and reduces the sample width to one bit. Unlike the sigma-delta modulator in a sigma-delta ADC, this modulator is all digital. The transfer function is implemented in the digital domain with an IIR filter. This digital filter performs the same modulator function as in the ADC, where the input signal is effectively lowpass filtered, and the quantization noise is high-pass filtered.

As in the case of a sigma-delta ADC, the 1 bit DAC output is meaningless until it is averaged in some manner. Also, there is a need to remove the shaped quantization noise which resides in the upper frequency area. Finally, there is also a need to reject any images which result about the output Nyquist rate. The analog smoothing filter performs these functions, usually in several stages. It is important for the design of this filter that the filter characteristics match the requirements of the overall system. For example, an audio system would need to have its phase and amplitude response preserved while the output filter also provides the appropriate rejection of higher frequency components. If the smoothing filter is an active filter, care must be taken that the op amps used do not introduce distortion products in the final output due to slewrate limiting and noise.

THE ADSP-28MSP02 SIGMA-DELTA CODEC

The ADSP-28msp02 is a mixed-signal peripheral device available based on sigma-delta design. The device is a linear codec with a 16-bit sigma-delta ADC and DAC, thereby providing a complete analog front end and back end for high performance voiceband DSP applications. Key features of the IC are summarized in Figure 6.36 and a functional block diagram is shown in Figure 6.37.

KEY FEATURES OF THE ADSP-28msp02 SIGMA-DELTA CODEC

- 16-Bit Sigma-Delta ADC
- 16-Bit Sigma-Delta DAC
- On-Chip Antialiasing and Smoothing Filters
- 8kSPS Sampling Rate, 128x Oversampling Ratio
- On-Chip Voltage Reference
- 65dB SNR and THD
- Easy Interface to DSP Chips
- 24-Pin DIP/SOIC Package
- Single +5V Supply, 100mV Power Dissipation
- Ideal for Voiceband Applications

Figure 6.36

Compared to traditional m-law and A-law codecs, the ADSP-28msp02's linear coded ADC and DAC maintain wide dynamic range throughout the transfer function. An effective sampling rate of 8kSPS coupled with 65dB SNR and THD performance make the device attractive in many telecommunications applications such as digital cellular telephones. The part is packaged in a 24-pin DIP/SOIC package ensuring a highly integrated and compact solution to voiceband analog processing requirements. The ADSP-28msp02 easily interfaces to the ADSP-2101, ADSP-2105, ADSP-2111, MC56001 and TMS320C25 DSP processors via its serial I/O port; the serial port (SPORT) is used to transmit and receive data or control information to and from the device.

The encoder side of the ADSP-28msp02 consists of two selectable analog input amplifiers and a sigma-delta ADC. The gain of the input amplifiers can be adjusted with the use of external resistors from -12dB to +26dB. An optional 20dB preamplifier can be inserted before the modulator. The preamplifier and the multiplexer are configured by bits in the control register. The sigma-delta ADC consists of a sigma-delta modulator, an antialiasing decimation filter, and a digital high pass filter. The modulator noise-shapes the signal and produces 1-bit samples at a 1.024MHz rate. This bit stream, representing the analog input, is fed to an antialiasing decimation filter which consists of two low-pass filter stages. The first stage reduces the sampling rate to 40kHz and increases the sample width to 16 bits; the second further reduces the sampling rate to 8kSPS. Each resulting sample is then loaded into the SPORT for transmission.

The decoder consists of a sigma-delta DAC and a differential output amplifier. The DAC reads 16-bit samples at an 8kHz rate from the SPORT. The samples are low- and high-pass filtered by the digital anti-imaging and high pass filters. The anti-imaging filter interpolates the sampling rate in two stages, first to 40kHz, and then to 1.024MHz. The resulting 16-bit samples are processed by the digital sigma-delta modulator which reduces the sample width to 1 bit. This bit stream is fed to an analog smoothing filter which converts the data to an analog voltage. The gain of the smoothing filter can be adjusted via the control register from -15dB to +6dB is 3dB steps.

ADSP-28msp02 BLOCK DIAGRAM

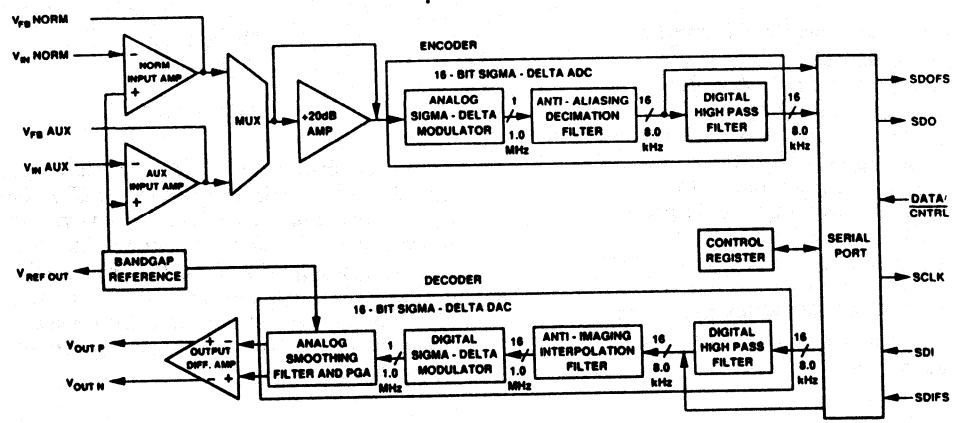

Figure 6.37

Multi-Stage Noise Shaping (MASH) Sigma-Delta Converters

As has been discussed, non-linear stabilization techniques have been successfully used to design a fifth-order sigma-delta loop in the AD1879 audio ADC. An alternative approach, called multistage noise shaping (MASH) utilizes cascaded stable first-order loops. Figure 6.38 shows a block diagram of a three-stage MASH ADC. The output of the first integrator is subtracted from the first DAC output to yield the first stage quantization noise, Q1. Q1 is then quantized by the second stage. The output of the second integrator is subtracted from the second DAC output to yield the second stage quantization noise which is in turn quantized by the third stage.

MULTI-STAGE NOISE SHAPING SIGMA-DELTA ADC (MASH)

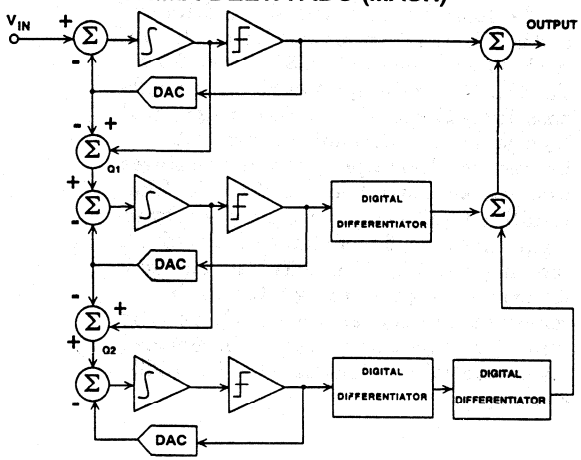

Figure 6.38

The output of the first stage is summed with a single digital differentiation of the second stage output and a double differentiation of the third stage output to yield the final output. The result is that the quantization noise Q1 is suppressed by the second stage, and the quantization noise Q2 is suppressed by the third stage yielding the same suppression as a third-order loop. Since this result is obtained using three first-order loops, stable operation is assured. A comparison of the MASH architecture with the higher-order single-loop architecture is given in Figure 6.39.

MASH TOPOLOGY VERSUS HIGHER-ORDER LOOP SIGMA-DELTA CONVERTERS

- MASH Cascades Single-Order Loops, therefore Easy to Stabilize
- Gain and Phase Matching Critical in MASH Converters for Errors to Cancel
- MASH Digital Differentiators Must Match Analog Integrators
- Single-Loop Higher Order Modulators Less Subject to Idling Patterns
- Single-Loop Higher Order Modulators More Difficult to Understand, Analyze, and Stabilize, But Can Be Done Using Non-Linear Techniques as in AD1879 (5th Order Modulator)

Figure 6.39

Multi-Bit Sigma-Delta Converters

So far we have considered only sigma-delta converters which contain a single-bit ADC (comparator) and a single-bit DAC (switch). The block diagram of Figure 6.40 shows a multi-bit sigma-delta ADC which uses an n-bit flash ADC and an n-bit DAC. Obviously, this architecture will give a higher dynamic range for a given oversampling ratio and order of loop filter. Stabilization is easier, since second-order and higher loops can be used. Idling patterns tend to be more random thereby minimizing tonal effects.

MULTI-BIT FIRST-ORDER SIGMA-DELTA ADC

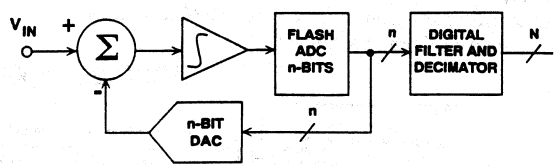

Figure 6.40

The real disadvantage of this technique is that the linearity depends on the DAC linearity, and thin film laser trimming is required to approach 16-bit performance levels. This makes the multi-bit architecture extremely impractical to implement on mixed-signal ICs. A comparison of the multi-bit versus single-bit sigma-delta converter is given in Figure 6.41.

MULTI-BIT VERSUS SINGLE-BIT SIGMA DELTA CONVERTERS

Multi-Bit:

- **Higher Dynamic Range for Given Oversampling Ratio and Loop Filter Order**
- **Higher Order Systems Easier to Stabilize**
- **Fewer Tonal Effects due to Idling Patterns**
- **Linearity Depends on DAC**
- **Thin Film Laser Trimming Required**

Single-Bit:

- **Perfect Linearity, no Strict Matching Requirements**
- **No Laser Trimming Required**
- **Perfect Topology for Mixed-Signal VLSI**
- **Non-Linear Techniques Required to Stabilize Higher Order Loops (AD1879)**

Figure 6.41

Sigma-Delta Summary

Although the concepts used in sigma-delta converters are not new by any means, their recent proliferation has been primarily driven by the need for converters which are compatible with mixed-signal VLSI chips. The sigma-delta architecture is ideal for converters for measurement, voiceband, and audio applications. Further exploration of various sigma-delta circuit topologies combined with the development of new processes is sure to push the maximum dynamic range and sampling rates even higher.

It is clear that the sigma-delta converter is not the answer to all data acquisition requirements at the present time. Upper sampling frequency is limited, thereby excluding video applications, multiplexing inputs is difficult due to the settling time of the internal digital filter, and out-of-range signals may cause saturation of the internal modulators.

On the other hand, the inherently good linearity performance without the need for laser trimming, the relaxation of antialiasing and anti-imaging filter requirements due to oversampling, and the basic sampling nature of the architecture without the need for a SHA will keep sigma-delta development moving at a rapid pace as mixed-signal ICs proliferate.

SIGMA-DELTA SUMMARY

- **Inherently Excellent Linearity**
- **Ideal for Mixed-Signal IC Processes, no Trimming**
- **No SHA Required**
- **Upper Sampling Rate Currently Limits Applications to Measurement, Voiceband, and Audio**
- **Out-of-Range Signals May Cause Modulator Saturation**
- **Analog Multiplexing Applications Limited by Internal Filter: Use one Sigma-Delta ADC per Channel!**

Figure 6.42

REFERENCES

1. J. Dattorro, A. Charpentier, D. Andreas, The Implementation of a One-Stage Multirate 64:1 FIR Decimator for use in One-Bit Sigma-Delta A/D Applications, **AES 7th International Conference**, May 1989.

2. W.L. Lee and C.G. Sodini, A Topology for Higher-Order Interpolative Coders, **ISCAS PROC. 1987.**

3. P.F. Ferguson, Jr., A. Ganesan and R. W. Adams, One Bit Higher Order Sigma-Delta A/D Converters, **ISCAS PROC. 1990**, Vol. 2, pp. 890-893.

4. R. Koch, B. Heise, F. Eckbauer, E. Engelhardt, J. Fisher, and F. Parzefall, A 12-bit Sigma-Delta Analog-to-Digital Converter with a 15MHz Clock Rate, **IEEE Journal of Solid-State Circuits**, Vol. SC-21, No. 6, December 1986.

5. Wai Laing Lee, A Novel Higher Order Interpolative Modulator Topology for High Resolution Oversampling A/D Converters, **MIT Masters Thesis**, June 1987.

6. D. R. Welland, B. P. Del Signore and E. J. Swanson, A Stereo 16-Bit Delta-Sigma A/D Converter for Digital Audio, **J. Audio Engineering Society**, Vol. 37, No. 6, June 1989, pp. 476-485.

7. R. W. Adams, Design and Implementation of an Audio 18-Bit Analog-to-Digital Converter Using Oversampling Techniques, **J. Audio Engineering Society**, Vol. 34, March 1986, pp. 153-166.

8 B. Boser and Bruce Wooley, The Design of Sigma-Delta Modulation Analog-to-Digital Converters, **IEEE Journal of Solid-State Circuits**, Vol. 23, No. 6, December 1988, pp. 1298-1308.

9. Y. Matsuya, et. al., A 16-Bit Oversampling A/D Conversion Technology Using Triple-Integration Noise Shaping, **IEEE Journal of Solid-State Circuits**, Vol. SC-22, No. 6, December 1987, pp. 921-929.

10. Y. Matsuya, et. al., A 17-Bit Oversampling D/A Conversion Technology Using Multistage Noise Shaping, **IEEE Journal of Solid-State Circuits**, Vol. 24, No. 4, August 1989, pp. 969-975.

11. P. Ferguson, Jr., A. Ganesan, R. Adams, et. al., An 18-Bit 20-kHz Dual Sigma-Delta A/D Converter, **ISSCC Digest of Technical Papers**, February 1991.

12. Steven Harris, The Effects of Sampling Clock Jitter on Nyquist Sampling Analog-to-Digital Converters and on Oversampling Delta Sigma ADCs, Audio Engineering Society Reprint 2844 (F-4), October, 1989.

Signal Conditioning Products
Contents

Page

AN-274 Taking the Uncertainty Out of Thermocouple Temperature Measurement (With the AD594/AD595) 21–3

21

ANALOG DEVICES

AN-274
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Taking the Uncertainty Out of Thermocouple Temperature Measurement (With the AD594/AD595)

by Bob LeFort and Bob Ries

Temperature is the most frequently measured physical parameter. However, the techniques of temperature measurement are grossly misunderstood, often resulting in serious inaccuracies or meaningless data. This application note is intended to clarify some of these common misunderstandings as well as present some interesting and useful circuit solutions.

TEMPERATURE TRANSDUCER TECHNOLOGIES

The most commonly-used electronic temperature measurement devices currently available include the thermocouple, the resistance temperature detector (RTD), the thermistor, and the integrated circuit temperature transducer. All have associated application benefits and limitations which are delineated in Table I.

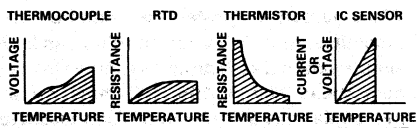

LINEARITY		*		**
SENSITIVITY			*	**
RUGGEDNESS	*			
COST	**		*	*
STABILITY		*		
ACCURACY		**		*
RESPONSE TIME	*			
NOISE IMMUNITY				*
POWER DISP.	**			*
MAX TEMP RANGE[1] (°C)	−270 TO +2980	−180 TO +630	−80 TO +150	−55 TO +150

*GOOD
**EXCELLENT
[1] TEMPERATURE RANGE INDICATED IS NOT NECESSARILY FOR A SINGLE VERSION OF THE TRANSDUCER TYPE.

Table I. Sensor Comparison

THERMOCOUPLE CHARACTERISTICS

The thermocouple is the most widely used temperature sensor for instrumentation purposes. Because of this, the National Bureau of Standards (NBS) has extensively characterized various metal combinations, i.e., type J (Iron – Constantan), type K (Chromel – Alumel), type E (Chromel – Constantan), and type T (Copper – Constantan). Thermocouple qualities include inherent accuracy, wide temperature range, fast thermal response, ruggedness, low cost, repeatability, and versatility of application. In addition to being widely used, the thermocouple is also the most misunderstood temperature sensor. Terms such as cold junction compensation, Seebeck coefficient, and isothermal connections or blocks have caused confusion and anxiety for many users. This application note explains those terms and provides information that allows the reader to measure temperature accurately and easily.

THE THERMOCOUPLE LOOP

Two wires of dissimilar metal, when joined together at both ends, constitute the basic thermocouple loop (see Figure 1a). This loop generates a voltage proportional to the difference in temperature between the two junctions. Since the thermocouple is basically a differential temperature measuring device, measuring a single temperature requires that the temperature of one of the junctions (a reference junction) be known. Users of thermocouples have relied on a variety of techniques to determine and compensate for the reference or "cold" junction temperature.

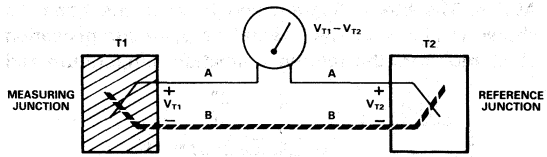

Figure 1a. Thermocouple Loop

21

ICE POINT REFERENCING

The voltage output of all thermocouples (as given in NBS tables) is referenced to 0°C. This means that the voltage across the thermocouple corresponds to the temperature of the measuring junction only if the reference junction is held at 0°C. This can be done with an ice point cell or an "ice bath" as shown in Figure 1b. Unfortunately these methods are awkward, expensive, and therefore only feasible in a laboratory setting. In a production environment it is impractical to maintain a reference junction at 0°C.

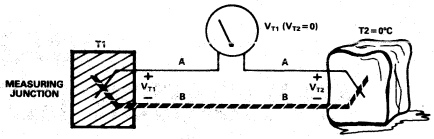

Figure 1b. Ice Point Reference

LAW OF INTERMEDIATE METALS

In practice, to eliminate the need for an explicit reference junction (as in Figure 1a) a direct connection equivalent to the basic thermocouple loop is made (see Figure 1c). The Law of Intermediate Metals states that a third metal (in most cases Copper) connected to the two dissimilar metals of a thermocouple will not have any effect on the output voltage, as long as the connections are at the same temperature.

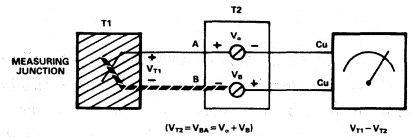

Figure 1c. "Indirect" Reference Junction

PRACTICAL THERMOCOUPLE MEASUREMENT

In a production environment an ice point reference can be eliminated by compensating for the voltage developed at the reference junction. This is done with a circuit which adds a voltage into the thermocouple loop, equal but opposite to that of the reference junction (see Figure 1d).

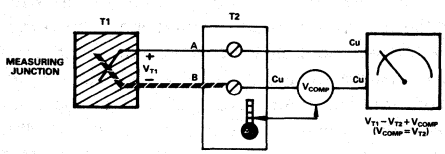

Figure 1d. Cold Junction Compensation

A device which accomplishes this and more is the AD594/AD595. The block diagram and basic connections are shown in Figure 2. The internal ice point compensation block monitors the reference junction temperature and adds the appropriate voltage into the thermocouple loop at the internal summing node. This net voltage is then amplified to a nominal output of 10mV/°C. The AD594 is factory calibrated for type J thermocouples, while the AD595 is set for type K.

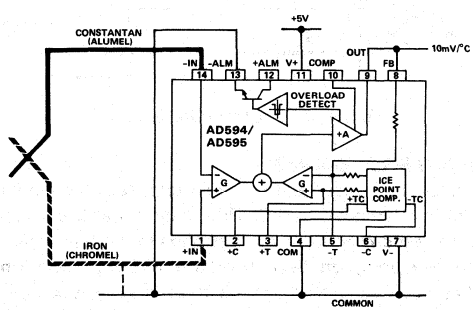

Figure 2. AD594/AD595 Block Diagram

SEEBECK COEFFICIENT

Seebeck coefficient of a thermocouple is defined as the rate of change of thermal voltage with respect to temperature at a given temperature and is usually expressed in μV/°C. Thermocouple nonlinearity is represented by the change in this coefficient over temperature. A graph of the Seebeck coefficient for various thermocouples is given in Figure 3.

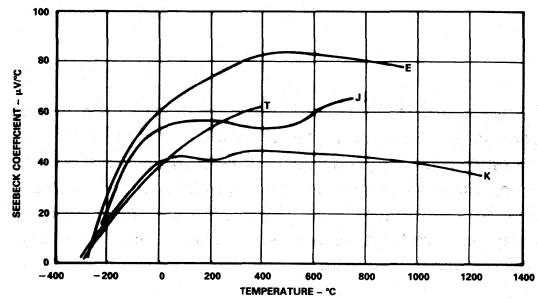

Figure 3. Seebeck Coefficient vs. Temperature

TYPES OF THERMOCOUPLES

The two characteristics generally used to differentiate thermocouple types are sensitivity and operating temperature range. The graph in Figure 4 portrays these characteristics for some popular combinations of metals.

While factory-calibrated to condition a J type thermocouple, the AD594 can condition an E type with just a simple external adjustment as shown in the AD594/AD595 data sheet. The AD595, calibrated for a K type thermocouple, may also be directly connected to a type T thermocouple with less than 0.2°C additional error.

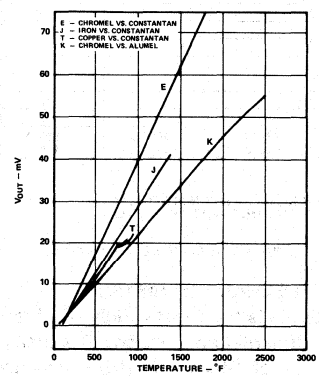

Figure 4. Thermocouple Output vs. Temperature

OPTIMIZING PERFORMANCE WITH THE AD594/AD595

Achieving the full rated accuracy from the AD594 or AD595 requires adherence to the following design guidelines:

1. **Cold Junction Errors**
   The AD594/AD595 has on-chip cold junction compensation. For this function to work correctly the device must be held at the same temperature as the thermocouple cold junction. Keep other components or heat sources from direct contact with the AD594/AD595 as their heat dissipation could cause cold junction compensation-related errors. (The AD594/AD595 draws only 160μA quiescent supply current; this minimizes self-heating related errors.)
2. **Circuit Board Layout**
   The printed circuit board connection layout (with the optional calibration resistors) illustrated in Figure 5 provides for thermal equilibrium between the cold junction and the AD594/AD595. Here the device and circuit board are thermally contacted in the copper printed circuit board tracks under pins 1 and 14. The reference junction is now composed of a copper-constantan (or copper-alumel) connection and copper-iron (or copper-chromel) connection, both of which are held at the same temperature as the AD594/AD595.

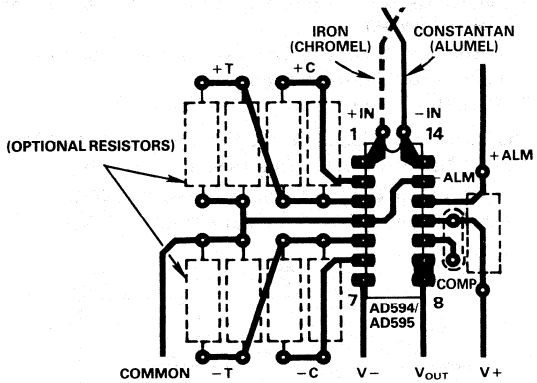

Figure 5. PCB Connections

3. **Soldering**
   To ensure secure bonding and to minimize I × R drops, clean the thermocouple wire to remove oxidation prior to soldering. Noncorrosive rosin flux is effective with iron, constantan, chromel, and alumel and the following solders: 95% tin – 5% antimony, 95% tin – 5% silver, or 90% tin – 10% lead.
4. **Grounding Considerations**
   The AD594/AD595 input stage has transistors which require bias currents to flow from the thermocouple inputs to ground. If this path is not provided these currents will drive the input stage into cutoff, causing the output to give a false reading. A direct connection to ground should be used to provide the return path.
5. **Minimizing Noise**
   Compensation capacitors between pins 9, 10 and 10, 11 will minimize the amplification of high frequency noise picked up by the thermocouple. The values shown in Figure 6 will provide a zero at 60Hz, but will increase the circuit's response time.

   To avoid I × R drops in the ground lines, all ground points should be directly connected to a central point. The 100Ω resistor combined with the 0.1μF capacitor will filter ripple and transient spikes on the supply lines.

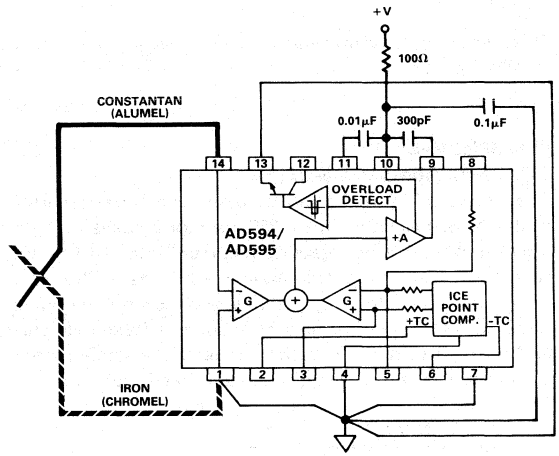

Figure 6. Reducing Errors through Filtering, Compensating and Grounding

EXTENDED AMBIENT TEMPERATURE ERROR CALCULATIONS

The ambient operating temperature range of the AD594/AD595 is specified from 0 to 50°C to minimize the errors associated with thermocouple nonlinearities (the varying Seebeck coefficient), and to optimize accuracy for a 25°C ambient. The AD594/AD595 ice point compensation voltage is linear and is matched to the best fit straight line of the thermocouple's output from 0 to 50°C. Outside of this range deviation between the thermocouple and the compensation voltage becomes more pronounced. This means that while the AD594/AD595 functions correctly outside of the rated temperature range, it may not remain within its specified temperature stability error limits. Table II provides a list of the calculated maximum errors associated with the commercial, industrial, and extended ambient operating temperature ranges. The ambient temperature refers to the device and reference junction. The measuring junction can be at any temperature within the thermocouple rated limits.

Ambient Temp. °C	AD594C Temp. Rej. Error °C	AD594C Total Error °C	AD594A Temp. Rej. Error °C	AD594A Total Error °C	AD595C Temp. Rej. Error °C	AD595C Total Error °C	AD595A Temp. Rej. Error °C	AD595A Total Error °C
−55	4.83	5.83	6.83	9.83	5.28	6.28	7.28	10.28
−25	1.98	2.98	3.23	6.23	2.04	3.04	3.29	6.29
0	0.62	1.62	1.25	4.25	0.62	1.62	1.25	4.25
+25	0.00	1.00	0.00	3.00	0.00	1.00	0.00	3.00
+50	0.62	1.62	1.25	4.25	0.62	1.62	1.25	4.25
+70	1.46	2.46	2.59	5.59	1.38	2.38	2.50	5.50
+85	2.25	3.25	3.75	6.75	1.99	2.99	3.49	6.49
+125	4.90	5.90	7.40	10.40	3.38	4.38	5.88	8.88

NOTE:
Temp. Rej. Error has two components. (a) Difference between actual reference junction and ice point compensation voltage times the gain (b) Offset and gain TCs extrapolated from 0 to 50°C limits. Total error is temp. rej. plus initial calibration error.

Table II. Maximum Calculated Errors at Various Ambient Temperatures

CIRCUIT IDEAS

OPTIONAL TRIMMING SCHEME

The circuit in Figure 7 nulls the residual calibration error of the AD594/AD595. The 15MΩ resistor insures that the offset is negative by injecting a current into −T (pin 5), corresponding to approximately a −3°C offset. The trimming potentiometer (R_{CAL}) allows the "forced" negative offset to be trimmed to zero by injecting a balancing current into +T (pin 3). This circuit makes it possible to null any calibration error with a single unidirectional trim.

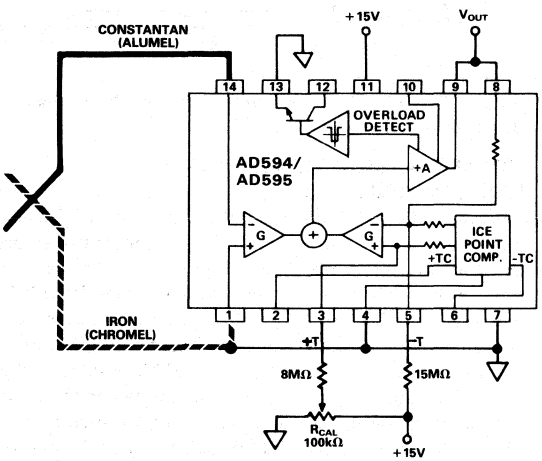

Figure 7. Calibration Error Adjustment

FAHRENHEIT OUTPUT

Figure 8 illustrates a circuit which allows the user to directly read a voltage output of 10mV/°F. The temperature scale conversion formula of;

Degrees Fahrenheit = (9/5) (Degrees Celsius) + 32

is entirely implemented in hardware. A current of 200nA/°C injected into pin 3 creates the 32°F offset while the resistor network on the output increases the gain by 9/5.

To calibrate the output:

1. Remove the thermocouple and input an ac signal to pins 1 and 14 of 10mV p-p, 100Hz. (By using an ac excitation the gain and offset adjustments are independent.)
2. Adjust R_{GAIN} for a p-p output of 3.481V (AD594) or 4.451V (AD595).
3. Reconnect a thermocouple which is in an ice bath or ice point cell at 0°C to pins 1 and 14.
4. Adjust R_{OFFSET} until the output reads 320mV.

The ideal transfer function for a Fahrenheit output AD594/AD595 when trimmed with thermocouple at 0°C is;

AD594 output = (Type J voltage + 919μV) × 348.12
AD595 output = (Type K voltage + 719μV) × 445.14.

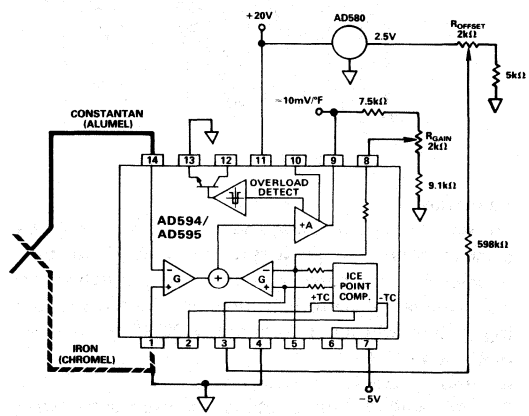

Figure 8. °C to °F Conversion

DIRECT MEAN TEMPERATURE

Average temperature can be measured directly with a single AD594/AD595 configured as shown in Figure 9. The output of this circuit equals $(T_1+T_2+T_3+\ldots T_N)/N$ (in °C) times a nominal 10mV/°C. With any number of thermocouple-resistor pairs in parallel the AD594/AD595 still provides for the correct cold junction compensation. The 300Ω series resistors minimize the currents circulating among the thermocouple branches. They will have positive/negative balancing voltage drops for thermocouples at temperatures lower/higher than the mean. This circuit also works well for generating a precise mean reading of an object which has significant thermal gradients.

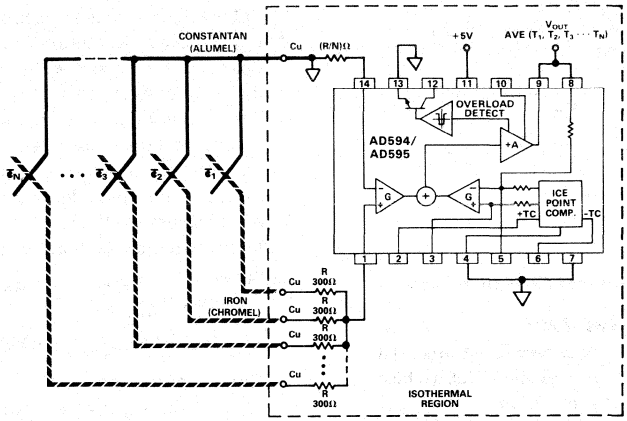

Figure 9. Measuring Average Temperature

TEMPERATURE MULTIPLEXING

Multiplexing thermocouple signals minimizes the number of AD594/AD595s required in a large temperature measuring data acquisition system (see Figure 10). This technique also serves another important function: it transforms the multiple reference junction connections from a terminal block assembly to a single junction at the AD594/AD595.

By placing a thermocouple beneath the AD594/AD595 (in thermal contact) and returning it to an isothermal connector, the reference junction voltages generated at the isothermal block are effectively cancelled. For a given multiplexer ON position, the Constantan (Alumel) – Copper junction in series with the Copper – Constantan (Alumel) junction contribute equal but opposite voltages. That is, since the block is isothermal, $V_1 = V_2$. Similar logic can be applied to the Iron (Chromel) – Copper junctions.

Because of these cancellations, the built-in cold junction compensator in the AD594/AD595 now compensates for the thermocouple directly underneath the IC. Thus, the terminal block can be at any remote location. The AD594/AD595 and the attached thermocouple, however, should remain in the 0 to 50°C range.

Using an AD7502, four temperatures can be monitored with a single AD594/AD595 (an AD7507 allows 8 temperatures). Approximately half of the thermocouple wire can be eliminated by using a single-ended multiplexer, however, the system will be more susceptible to common-mode noise pick-up.

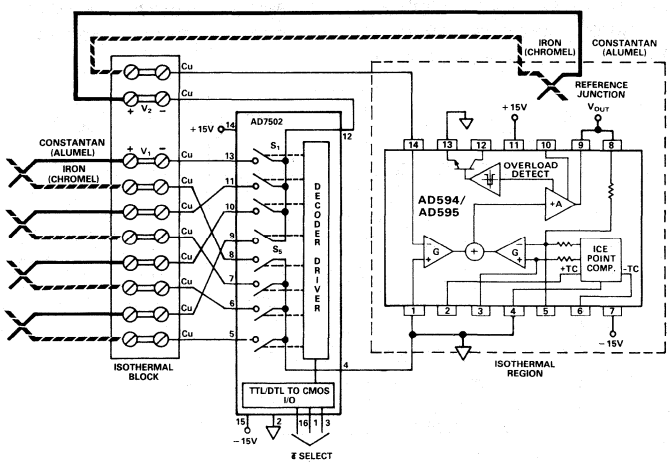

Figure 10. Multiplexing Thermocouples

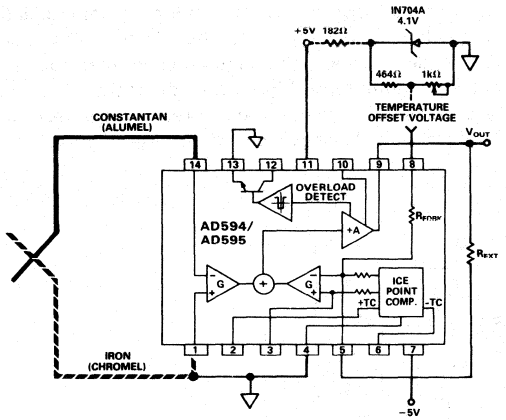

Figure 11. Zero Suppression and Sensitivity Change

OFFSETTING AND CHANGING THE GAIN

The circuit shown in Figure 11 serves two functions: 1) it allows the user to add an offset to the AD594/AD595 which shifts the output to read 0V at any temperature (i.e., suppress zero), and 2) it changes the gain of the AD594/AD595 enabling different output sensitivity (output voltage change per degree change of thermocouple).

Shifting the output to read 0V at a temperature other than 0°C is accomplished by applying an offset voltage to the feedback resistor (pin 8) and thus to the negative input of the right-hand amplifier connected to the summing node.

The offset voltage (obtained from the output column o Table I in the AD594/AD595 data sheet) transposes the zero voltage output to the temperature equivalent of the applied voltage. The sensitivity can be increased/decreased by replacing the internal feedback resistors with a larger/smaller external resistance. One method to calculate the value of the new feedback resistance is:

1. Determine the desired output sensitivity (in mV/°C).
2. Decide on a temperature range T1 to T2.
3. Calculate the average thermocouple sensitivity over that temperature range; $(V_{@1} - V_{@2})/(T_1 - T_2)$.
4. Divide the desired sensitivity by the average thermocouple sensitivity: result of (1) ÷ calculated value in (3). This value is the new gain (G_{NEW}) of the AD594/AD595. If the calculations are done correctly this result will be dimensionless.
5. Measure the actual feedback resistance (pin 8 to pin 5), R_{FDBK}.
6. $R_{INTERNAL} = \frac{R_{FDBK}}{193.4 - 1} = \frac{\text{Result of (5)}}{193.4 - 1}$

   NOTE: Use 247.3 for an AD595 instead of 193.4.
7. The new feedback resistance, $R_{EXTERNAL} = (G_{NEW} - 1)(R_{INTERNAL}) =$ (result from (4) − 1)(result from (6))

This technique makes it possible to measure a temperature range of 300 to 330°C with a 5V supply and an output of 100mV/°C starting with a 0V output at 300°C. A 4.1V zener diode and a resistor divider can be used to suppress zero.

CURRENT-MODE TRANSMISSION

When sending a signal through a noisy environment, it is preferred to transmit a current rather than a voltage. Figure 12 shows a method of transmitting the AD594/AD595 output signal as a current and then converting it back to a voltage at the control point.

In this circuit the feedback voltage at pin 9 forces the voltage across R_{SENSE} to equal the thermocouple voltage. Correctly choosing R_{SENSE}, (5.11Ω for AD594, 4.02Ω for AD595), generates a 10μA/°C current. Because the voltage across R_{SENSE} equals the thermocouple voltage, the reference junction voltage appears across the AD594/AD595 inputs. The amplifier +A drives the base of the 2N2222 transistor, converting the output voltage into a current.

Because the 160μA quiescent current flows through R_{SENSE}, it does not contribute any error. However, this means the minimum temperature that can be measured is 16°C. The accuracy of the circuit is determined by the initial AD594/AD595 calibration error and the match between R_{SENSE} and the 1kΩ current to voltage conversion resistor at the measurement point.

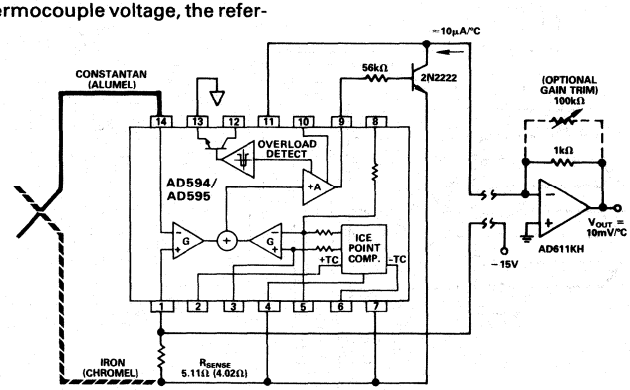

Figure 12. Remote Temperature Measurement

Voltage-to-Current Converters
Contents

		Page
AN-275	Versatile Transmitter Chip Links Strain Gages and RTDs to Current Loop	22–3

22

ANALOG DEVICES

AN-275
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Versatile Transmitter Chip Links Strain Gauges and RTDs to Current Loop

by Paul Brokaw

Replacing a board full of parts, a transmitter chip can be configured to match the 4-to-20-mA loop to a range of different process-control requirements.

The digital transmission of data may seem to be taking over from analog transmission. But in the process-control arena, the two-wire 4-to-20 mA current loop remains firmly entrenched as the traditional connection between transducer and control room. Moreover, that holds true even though converting the signal of a resistance temperature detector or resistive bridge into 4 to 20 mA demands about the same number of components as converting it into an 8- or 10-bit serial digital word.

In both cases, the same types of operations must be performed. The transducer must be excited, the signals conditioned, and the small floating voltages amplified. In addition, analog input voltages must be converted into bits or current, in the process creating an output signal that can drive a two-wire line several miles long. Also, that line must carry the circuit power.

The arrival of the AD693 process-control transmitter goes a long way toward guaranteeing the staying power of the 4-to-20-mA current loop. The chip tackles all the functions required to handle standard resistive transducers and to operate from and determine the current in a typical process-control loop. In addition, pin strapping enables the chip to work with other current ranges, and a few extra passive parts permit its use with a broad selection of sensors.

As important, it can replace the pc boards, modules, and hybrids now taking care of those functions. Consider that a pc board is a relatively large unit whose many ICs, transistors, and passive and adjustable components dictate careful design and expensive assembly, calibration, and testing measures. Modules eliminate those

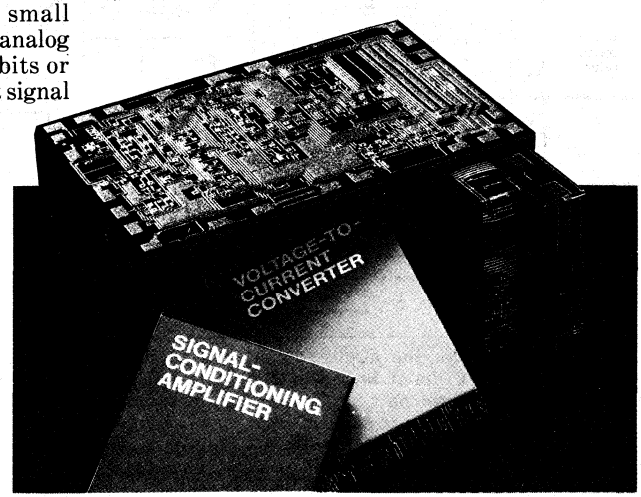

Reprinted from Electronic Design – April 18, 1985

Industrial Electronics: Versatile transmitter chip

concerns (apart from calibration) without necessarily reducing time and cost. Hybrids are small, but present problems of performance, functional completeness, and cost.

In sharp contrast, the 20-pin AD693 transmitter IC is an accurate, versatile, and tiny part that alleviates worries about functions or cost. Its precision reference sets the 4-mA zero point and also can excite a sensor. In fact, the chip's bag of signal-conditioning tricks covers virtually all common resistive sensors.

Loop-the-loop

The transmitter's operation is best demonstrated when its inputs are connected to a resistive bridge transducer and its output and power source to a 4-to-20-mA current loop (Fig. 1).

The loop is powered by a voltage source in series with a current-sensing, or current-monitoring, resistor, R_{sense}, located at the receiving end in the control room. The voltage across R_{sense} is the high-level signal to the controller. The AD693 connects to the two-wire loop at its remote end and couples it to the measurement represented by the output of the bridge transducer, or the sensor.

The reference on chip excites the bridge, the output of which connects directly to the chip's signal amplifier. That amplifier is an instrumentation-type circuit with a differential input and a single-ended output, relative not to ground but to the 6.2-V reference line. The output of the signal amplifier drives a voltage-to-current converter, which measures the total loop currrent with the voltage drop across the loop's sampling resistor, R_L. In turn it sets the loop's current, so that the voltage across R_L equals the input to the converter.

When a measurement unbalances the bridge, putting a small voltage at the input to the signal amplifier, the amplifier's output drives the converter's input positive, increasing the loop current. The gains of the signal amplifier and the resistance of the output resistor are laser-trimmed; in combination, they provide a precalibrated sensitivity relating the loop current directly to the bridge's output voltage. In addition, a current-limiting circuit monitors the current, so that if the input to the converter is overdriven, loop current is limited to 25 mA, protecting both loop and transmitter chip.

The zero input pin of the voltage-to-current converter is referred to a voltage that, because of the voltage drop across R_A, is slightly nega-

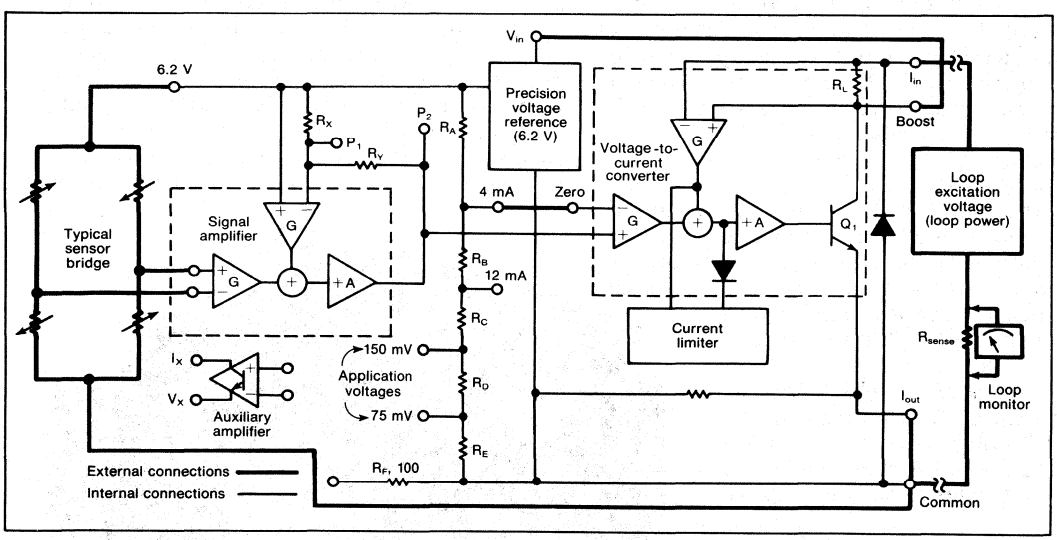

1. A one-chip transmitter, the AD693 converts the output voltage from resistive bridge sensors into a 4-to-20-mA process-control signal. Moreover, the chip excites the sensor, runs off loop power, and for many standard sensors, requires no span or zero trimming components.

tive with respect to the 6.2-V reference. So when the input signal from the bridge is zero, the converter's input is not zero and its output is 4 mA.

A voltage divider (R_A through R_E) from 6.2 V to common, or ground, is laser-trimmed so that connecting the converter's zero input to the divider's 4-mA tap results in 4 mA of loop current. Since that voltage is produced by a resistive divider, it can be easily adjusted to accommodate zero offsets in the bridge sensor. Moreover, the adjustment does not affect the signal amplifier's gain and thus is independent of the zero to full-scale range, or span.

An alternative connection for the converter's zero pin sets the loop current to 12 mA when the input is zero, meaning that input signals of both polarities can be handled. The zero pin also can be returned to the 6.2-V reference line, thereby reducing loop current to zero. That connection permits the chip to operate as a 0-to-20-mA converter, so long as the voltage reference and amplifiers are not powered from the loop. For most applications, the transmitter chip runs off loop power through the connection between the Boost pin and the V_{in} pin.

Connections between pins P_1, P_2, and the 6.2-V reference line provide the signal amplifier's negative feedback and take care of span adjustments. For instance, when the signal amplifier's gain is 2 (Fig. 1 again), the input span—that is, the signal that changes the loop current from 4 to 20 mA—is 30 mV. Connecting P_1 to P_2 changes the gain to 1 and the span to 60 mV; the span can be adjusted to a few millivolts by placing a resistor between P_1 and either P_2 or 6.2 V.

The way in which the input signal amplifier is designed produces a high differential input impedance and a wide common-mode range. The amplifier accommodates signals between 0 V and the 6.2-V reference and even, in some applications, signals greater than 6.2 V. The common-mode input range extends from −100 mV to within 4 V of the potential at V_{in}. The transmitter's ability to handle small negative voltages is valuable when ground-referenced signals have a small negative component. The positive range is useful when measurements are not powered directly by the reference.

The auxiliary amplifier solves a variety of excitation and measurement problems. For example, with the addition of two resistors, it operates at a noninverting gain of 10/6.2, which amplifies the reference, providing a low-impedance, 10-V source that can supply up to 3 mA to sensors.

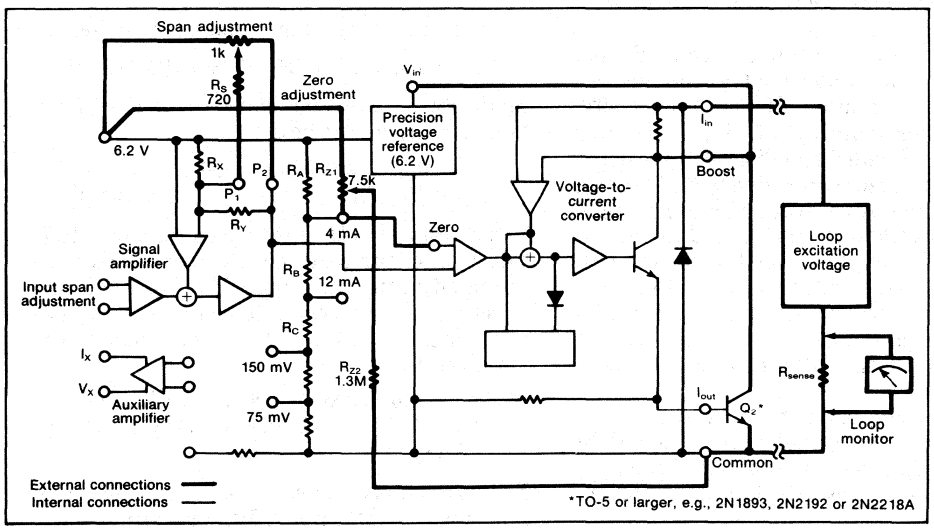

2. The transmitter chip adapts readily to the addition of zero and span adjustment potentiometers when used with nonstandard sensors. Incidentally, when an external transistor, Q_2, handles the loop current, the chip can operate at 50 mA full scale.

Industrial Electronics: Versatile transmitter chip

When the auxiliary amplifier is used as a voltage amplifier, the output stage must be powered by connecting the I_X terminal to a positive voltage, such as the 6.2-V line or V_{in}. In this case, the stage takes current from the boost terminal (see Fig. 1 again).

This amplifier can be used in most noninverting configurations and has the same common-mode range as the signal amplifier. The voltage divider, R_A through R_E, has 75- and 150-mV output taps that, when used with the auxiliary amplifier and the internal 100-Ω resistor, can form a signal-conditioning circuit with six precalibrated ranges for RTDs.

Exciting currents

The auxiliary amplifier also has a feature that proves useful for the current-mode, or constant-current, excitation of sensors. Since the I_X terminal supplies all of the auxiliary amplifier's output load current, it can be used as a high-impedance current source. When configured as a voltage follower, driving the 100-Ω resistor and connected to either the 75- or 150-mV pin, it makes a 750- or 1500-μA constant current available for sensor excitation.

The chip's I_{out} pin supplies the voltage-to-current converter with the excess current—beyond the operating and excitation currents—necessary to create the proper signal current for the loop. In many applications I_{out} is connected directly to common.

The pin also can be used to overcome the one disadvantage of a monolithic transmitter—its limited power dissipation. Though the AD693's power devices dissipate up to 7.5 W, most of that power can be offloaded to a discrete transistor. That setup minimizes errors due to self-heating and extends the temperature operating range.

If a resistor of about 5 Ω is placed in parallel with the loop's sampling resistor, R_L, the chip can be tweaked to operate in a 10-to-50-mA current loop. That configuration makes it desirable to use an external transistor, which enables the chip to work over the full industrial temperature range (Fig. 2).

In that circuit, I_{out} drives the base of a discrete npn transistor, Q_2, which pulls current out of the Boost pin and delivers it to common. The external transistor is in effect Darlington-connected to the normal pass device, Q_1. As a result I_{out} supplies only the base drive for Q_2, which handles most of the circuit's dissipation. The external transistor works with 4 to 20 and 0 to 20 mA and with any signal-conditioning circuits.

Essential to the AD693 and any other signal-conditioning circuit is the ability to adjust zero and the full-scale range. User-selected options afford three choices of loop current for an input signal of zero: 30 and 60 mV at 4 to 20 mA; ±15 and ±30 mV at 4 to 20 mA; and 37.5 and 75 mV at 0 to 20 mA. (In the last of those, the zero pin is tied to 6.2 V.) Although these values are laser-trimmed for high accuracy, they may require adjustment to accommodate variability between sensors or to add ranges. The circuit of Figure 2 indicates an optimum technique for adjusting zero and span for the AD693.

Voltages of 15 and 45 mV, negative with respect to the 6.2-V line, set the 4-mA and 12-mA operating points from a nominal source impedance of 450 Ω. The zero adjustment circuit is constructed by adding external resistors to modify the internal voltage divider to drive the zero pin.

To find the proper resistor values, the desired range of output current, I_A, is selected and substituted in a series of equations. For instance, to adjust the 4-mA tap, the equation to be used is:

$$\text{Potentiometer } R_{Z1} = (1.6\text{ V}/I_A) - 400\ \Omega$$

Then the series resistor would be evaluated as:

$$R_{Z2} = R_{Z1}\,(3.1\text{ V})/[15\text{ mV} + (3.75\ \Omega)\,(I_A)]$$

Similarly, for a divider connected to the 12-mA tap, the equations to be used are:

$$R_{Z1} = (4.8\text{ V}/I_A) - 400\ \Omega, \text{ and}$$
$$R_{Z2} = R_{Z1}\,(3.1\text{ V})/[45\text{ mV} + (3.75\ \Omega)(I_A)]$$

These equations take into account the ±10% tolerance of tap resistance and ensure a minimum adjustment range of I_A. For example, selecting an I_A of 200 μA delivers a zero adjustment range of ±1% over the 20-mA span. At the 4-mA tap, then:

$$R_{Z1} = 1.6\text{ V}/200\ \mu\text{A} - 400\ \Omega = 7.6\ \text{k}\Omega, \text{ and}$$
$$R_{Z2} = 7.6\ \text{k}\Omega\,(3.1\text{ V})/[15\text{ mV} + (200\ \mu\text{A})\,(3.75\ \Omega)]$$
$$= 1.49\ \text{M}\Omega$$

These resistance values can be rounded off to a convenient 7.5 kΩ and 1.3 MΩ, providing an ad-

Industrial Electronics: Versatile transmitter chip

justment range slightly in excess of ±200 μA.

The full-scale range is adjusted by controlling the feedback around the signal amplifier, achieved by paralleling internal resistor R_X or R_Y with an external resistor. (The internal resistors are laser-trimmed in ratio rather than in absolute value, but they are within ±10% of their nominal value of 800 Ω.)

Padding spans

To create a span, S_1, of less than 30 mV requires calculating the value of resistor R_{S1} located between P_1 and 6.2 V in parallel with R_X, according to the equation:

$$R_{S1} = 360\,\Omega/[(30\,\text{mV}/S_1)-1]$$

To set a span S_2 between 30 and 60 mV, a value must be calculated for resistor R_{S2} located between P_1 and P_2 in parallel with R_Y, according to the equation:

$$R_{S2} = 360\,\Omega/(1-60\,\text{mV}/S_2)/[(30\,\text{mV}/S_2)-1]$$

These are minimum values, and the resistors can be padded to obtain a particular gain; otherwise the circuit values in Figure 2 can be used with either R_{S1} or R_{S2} to get a range of span adjustments. For example, for continuous span adjustment between 20 and 40 mV, the 20-mV limit would be used to calculate RS_1 and the 40-mV limit to calculate RS_2, yielding the same value, 720 Ω. In this case the values happen to be equal, but in general, using the lower of the two values (R_S in Fig. 2) with a 1-kΩ potentiometer produces an adjustment range that encompasses at least the desired span.

Though that method of adjustment works

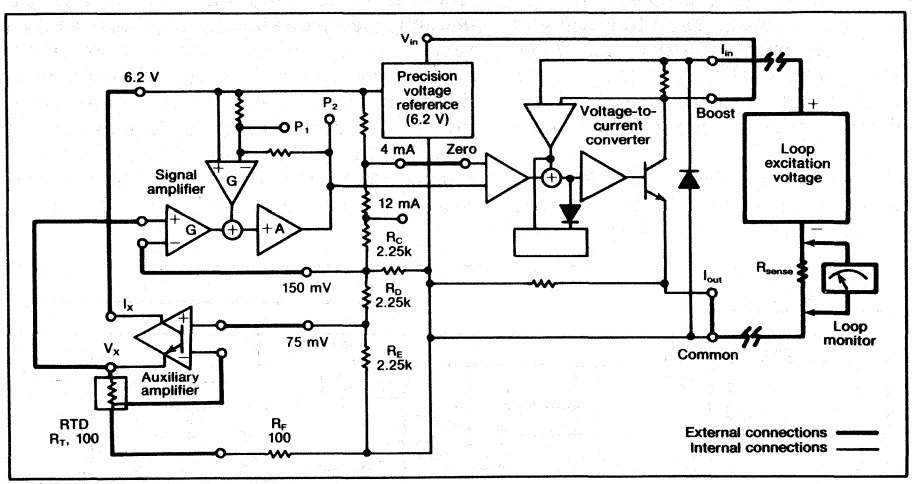

3. A complete loop-powered RTD measurement with six precalibrated ranges can be obtained with the AD693. Its auxiliary amplifier excites a dynamic bridge formed by the RTD, R_T, and the application resistors.

Temperature ranges (°C) for a 4-to-20-mA span with the AD693 and an RTD*

Input span	P_1 and P_2	Zero pin connected to 6.2-V pin	Zero pin connected to 4-mA pin	Zero pin connected to 12-mA pin
30 mV	Not connected	+25.9 to +130.5	0 to +103.9	−50.8 to +51.6
60 mV	Connected	+51.6 to +266.4	0 to +211.3	−100.6 to +103.9

*α = 0.00385 for a 100-Ω platinum RTD

22

Industrial Electronics: Versatile transmitter chip

well, the span of the AD693 can be adjusted with any scheme that is compatible with the built-in feedback network, as long as the amplifier is not overloaded. The network can also be used to inject currents, producing offsets or other additions to the signal applied to the voltage-to-current converter.

The excitation and signal-conditioning features of the chip easily adapt to measurements made with resistance temperature detectors, or RTDs. The application voltages at the bottom of the divider (between 6.2 V and common) set up precalibrated temperature ranges for standard 100-Ω platinum RTDs.

The auxiliary amplifier, application voltages, and 100-Ω resistor form a dynamic bridge circuit (Fig. 3). The auxiliary amplifier forces a constant 75 mV across the 100-Ω resistor. The current producing that voltage also passes through the RTD, putting a voltage across it too. The sum of these two voltages, appearing at V_X, is applied to the positive input of the signal amplifier, which compares, or sums, it with the 150-mV application voltage applied to the amplifier's negative input.

The resistance of the detector at 0°C is 100 Ω, making the voltage across it 75 mV and the output of the auxiliary amplifier 150 mV. The differential input to the signal amp is thus 0 V (150 mV − 150 mV). As the temperature rises, the resistance of the sensor increases and the differential input to the signal amplifier becomes:

$$V_X - 150\,mV = (R_T - 100\,\Omega)(75\,mV)/100\,\Omega$$

where R_T is the resistance of the sensor at any given temperature.

As R_T changes from 100 to 140 Ω (corresponding to a rise from 0° to +104°C), the signal amplifier's input changes from 0 to 30 mV, one of the precalibrated spans of the transmitter chip. Similarly, a temperature change from 0° to +211°C makes resistance rise from 100 to 180 Ω, resulting in a 60-mV signal—the other precalibrated span.

Other temperature ranges can be derived from various combinations of span and zero connections (see the table, p. 158). Also, still more ranges can be created by raising or lowering the application voltage—a change effected by modifying the divider, which has a nominal resistance of 1.5 kΩ at its 150-mV tap. For example, connecting a 1.5-kΩ resistor between the 150-mV tap and common doubles the resistance change required to put 30 or 60 mV into the signal am-

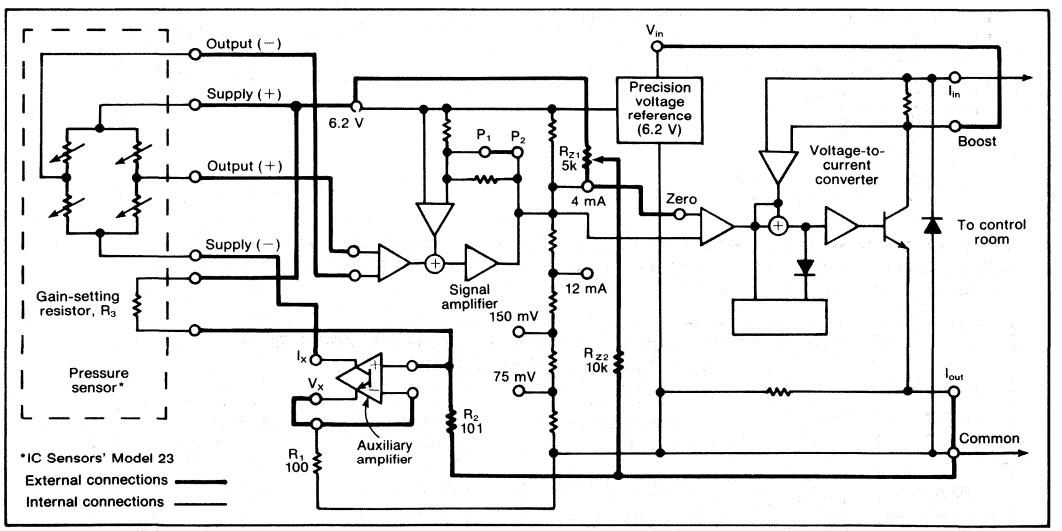

4. Teamed with the AD693's auxiliary amplifier, the calibration resistor (R_3) inside an IC pressure sensor can set the circuit's full-scale output span to 4 to 20 mA.

Industrial Electronics: Versatile transmitter chip

plifier and thus approximately doubles the temperature spans.

Furthermore, temperature span can be halved by doubling the gain of the signal amplifier, thereby reducing the span to 52°C. Alternatively, the application voltage can be doubled by placing a 62-kΩ resistor between the 6.2-V line and the 150-mV pin. Still other application voltages can be produced with a complete, external divider.

The AD693 easily lends itself to use with monolithic pressure sensors. For instance, it interfaces readily with the Model 23 from IC Sensors Inc. (Sunnyvale, Calif.). That sensor's internal calibration resistor makes the units interchangeable. When used with the AD693, the resistor forms part of a bias network that normalizes the full-scale output of the sensor to the 60-mV span of the transmitter (Fig. 4).

To get maximum stability over a wide temperature range, this semiconductor bridge sensor requires constant current excitation. The auxiliary amplifier in the transmitter chip is configured to deliver this current from its I_X terminals. All the current developed by voltage V_X across the chip's 100-Ω resistor is delivered through the I_X pin to the sensor. That voltage is equal to the output of the voltage divider, consisting of the sensor's calibration resistor, R_3, and a user-supplied 101-Ω resistor, R_2.

The circuit combines the preset input voltage span and the 6.2-V reference of the chip with the factory-calibrated pressure span of the transducer. Also it yields a calibrated 4-to-20-mA loop current requiring no span adjustments. Any residual zero offset in the sensor is eliminated by setting the output to 4 mA with R_{Z1}, the zero-adjustment potentiometer.

The circuit uses a Model 23-10 sensor rated at 10 psig; however, a selection of interchangeable sensors cover 5 to 250 psig. The effective span of a given sensor is halved by opening the jumper between P_1 and P_2.

If field interchangeability is not required, the same initial accuracy is available at lower cost with Series 22 sensors, which lack the factory-trimmed resistor. Instead, the user is supplied with the result of measurements of sensor output voltage at rated pressure and 1.5 mA of excitation. These sensors and similar devices from other suppliers are used in the same circuit by substituting a resistor for calibration resistor R_3. R_2 is rounded off to 100 Ω and the value of R_3 is calculated according to:

$$R_3 = [S_1\,(68\,\Omega)/\text{mV}] - 100\,\Omega$$

where S_1 is the full-scale sensor output in milli-

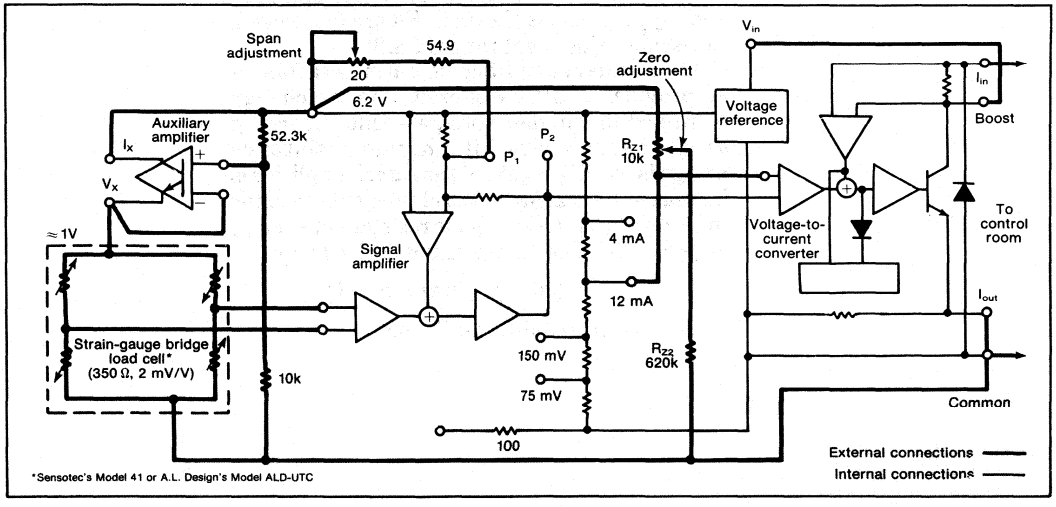

5. The chip's signal amplifier has enough gain to derive a 4-to-20-mA signal from the output of a foil strain gauge having 2-mV/V sensitivity. The auxiliary amplifier can be used to provide 1 V of bridge excitation, so that the bridge can be powered by the loop.

22

volts at an excitation of 1.5 mA. The calibration resistor normalizes the output so that the rated pressure span produces exactly 60 mV.

Load cells measure displacement of an elastic structure and the force that causes it. Metal-foil strain-gauge bridges, the most common sensing elements, are rugged, accurate, and reliable, but they are difficult to condition for 4-to-20-mA loops. Also, though they are more linear than the semiconductor units, they are much less sensitive; in addition, their usually low resistance limits the excitation voltage available from the 4-mA base-line current in the loop.

Load cells foil strain gauges

The AD693 transmitter can make the most of even that sensitivity. The internal circuitry of the chip leaves more than 3 mA of excitation current, yielding an excitation voltage of about 1 V on a typical 350-Ω bridge. The auxiliary amplifier and reference supply the excitation (Fig. 5), and the gain of the signal amplifier is raised to create a span of just 4 mV, enough to handle the typical ±2-mV/V sensitivity of the cell. (The 2-mV/V figure means 2 mV of output from the gauge for each 1 V of excitation.)

Load cells can be used in either a tension and compression mode, which produces a bipolar output, or compression only. For instance, the 12-mA tap matches bipolar tension-or-compression signals to the 4-to-20-mA span. The zero-adjustment circuit is composed of resistors R_{Z1} and R_{Z2}; the values shown in Figure 5 provide a full-scale adjustment range of ±2%.

The auxiliary amplifier, operating as a follower, provides a stiff 1 V of excitation, which may be raised to a higher voltage— thus affording greater sensitivity—if a higher resistance bridge is used. The gain of the signal amplifier is increased by placing a 54.9-Ω precision resistor in series with a 20-Ω variable resistor between P_1 and the 6.2-V line. These values derive from:

$$R_{S1} = (360\,\Omega)/[(30\,mV/S)-1]$$

where S is the required 4-mV span.

Voltage-to-Frequency Converters
Contents

		Page
AN-361	Ask the Applications Engineer—3: V/F Converters	23–3
AN-276	Analog-to-Digital Conversion Using Voltage-to-Frequency Converters	23–5
AN-277	Applications of the AD537 IC Voltage-to-Frequency Converter	23–13
AN-278	Operation and Applications of the AD654 IC V-to-F Converter	23–31
AN-279	Using the AD650 Voltage-to-Frequency Converter As a Frequency-to-Voltage Converter	23–59

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Ask the Applications Engineer — 3

by James Bryant

//F CONVERTERS

Q. *How do I send an analog signal a long distance without losing accuracy?*

A. An excellent solution to this common problem is to ship the signal as frequency using a *voltage-to-frequency converter* (VFC), a circuit whose output is a frequency proportional to its input. It is relatively easy to send a frequency signal over a long transmission path without interference via optical isolators, optical fibre links, twisted-pair or co-axial lines, or radio links.

If the data must be digital, the receiver will consist of a frequency counter, easily implemented in a single-chip microcomputer. Frequency is reconverted to analog voltage by a "frequency-to-voltage converter" (FVC)—generally a VFC configured to perform its inverse function, often using a phase-locked loop.

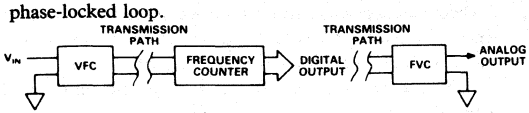

Q. *How does a VFC work?*

A. There are two common types: multivibrator- (AD537) and charge-balance (AD650) VFCs.*

In the *multivibrator* type, the input voltage is converted to a current which charges and discharges a capacitor. The switching thresholds are set by a stable reference, and the output, which has unity mark-space ratio, is a frequency proportional to the input.

The *charge-balance* VFC uses an integrator, a comparator and a precision charge source. The input is applied to the integrator, which charges. When the integrator output reaches the comparator threshold, the charge source is triggered and a fixed charge is removed from the integrator. The rate at which charge is removed must balance the rate at which it is being supplied, so the frequency at which the charge source is triggered will be proportional to the input to the integrator.

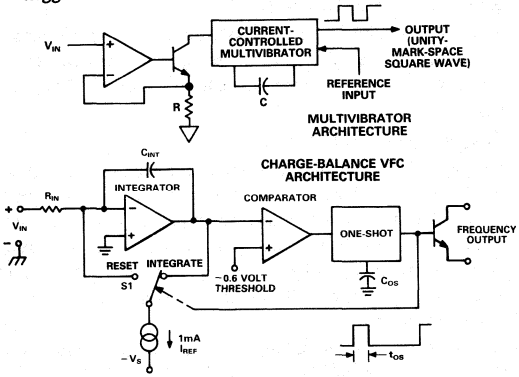

*Data sheets are available for any of the Analog Devices products mentioned here. An Application Note: "Operation and Applications of the AD654 V-to-F Converter," is also available without charge.

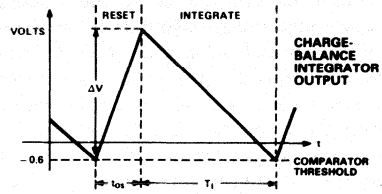

Q. *What are the advantages and disadvantages of the two types?*

A. The multivibrator is simple and cheap, demands little power, and has unity mark-space (M-S) output—very convenient with some transmission media. But it is less accurate than the charge-balance type and cannot integrate negative input transients.

The charge-balance type is more accurate, and negative input transients are integrated to contribute to the output. It has more-demanding supply requirements and a lower input impedance, and its output is a pulse train, not a unity M-S square wave.

Q. *What are the important types of error in a VFC?*

A. The same three as in most precision circuitry: offset errors, gain errors and linearity errors—and their variation with temperature. As with most precision circuitry, offset and gain can be trimmed by the user, but linearity cannot. However, the linearity of VFCs is normally very good (if the capacitors are properly chosen—see below).

Q. *How do you trim gain and offset in a VFC?*

A. The procedure suggested by theory is to trim offset first at zero frequency and then gain at full scale (FS). But this can give rise to problems in recognizing "zero frequency," which is the state when the VFC is just not oscillating. It is therefore better to trim offset with a small input (say 0–1% FS) and adjust for a nominal frequency, then trim gain at FS, and then repeat the procedure once or twice.

For example, suppose a VFC is being used with FS of 100 kHz at 10-volt input. Ideally, 10 V should give 100-kHz output and 10-mV input should give 100 Hz. Offset is, therefore, trimmed for 100 Hz with 10 mV applied; gain is then trimmed to give 100 kHz at 10 V. But gain error affects the 10-mV offset trim slightly, so the procedure may have to be repeated to reduce the residual error.

If a VFC is used with software calibration a deliberate offset is often introduced so that the VFC has a definite frequency for zero input voltage. The microcomputer measures the VFC outputs at 0 V and FS inputs and computes the offset and scale factor. It may also be necessary to reduce the gain so that the VFC cannot try to exceed its maximum rated frequency.

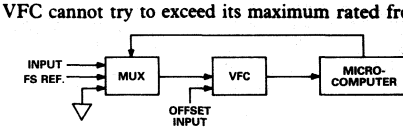

Reprinted from Analog Dialogue 23-2 1989

Q. *What circuit precautions are necessary when using a VFC?*

A. Apart from the usual precautions necessary with any precision analog circuitry (grounding, decoupling, current routing, isolation of noise, etc., a subject for a book, not a paragraph) the main precautions necessary when using a VFC are the choice of capacitor and separation of the input and output.

The critical capacitors in a precision VFC (the multivibrator's timing capacitor, and the monostable timing capacitor in a charge-balance type) must be stable with temperature variation. Furthermore, if they suffer from *dielectric absorption*, the VFC will be nonlinear and may have poor settling time.

If a capacitor is charged, discharged and then open-circuited it may recover some charge. This effect, known as dielectric absorption (DA), can reduce the precision of VFCs or sample-hold amplifiers using such capacitors. VFCs and SHAs should therefore use Teflon or polypropylene, or zero-temperature-coefficient (NP0, C0G) ceramic capacitors with low DA.

Coupling between output and input of a VFC can also affect its linearity. To prevent problems, decoupling practices and the usual layout precautions should be observed. This is critically important with opto couplers, which require high current drive (10-30 mA).

Q. *How do you make a frequency-to-voltage converter?*

A. There are two popular methods: the input frequency triggers the monostable of a charge-balance VFC that has a resistor in parallel with its integration capacitor; or the input frequency can be applied to the phase/frequency comparator of a phase-locked loop (PLL), which uses a VFC (of either type) as its oscillator. The basic principle of the first type is illustrated below.

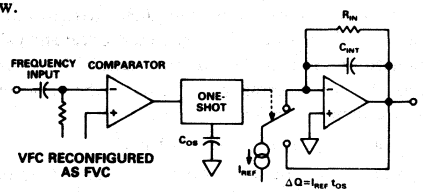

For each cycle of the input frequency, a charge, ΔQ, is delivered to the leaky integrator formed by *R and C. At equilibrium, an equal charge must leak away during each period, T (= 1/f)*, of the input, at an average rate, $I = V/R$. Thus, $V = \Delta Q \cdot f \cdot R$.

Though the mean voltage is independent of C, the output ripple is inversely dependent on C. The peak-to-peak ripple voltage, ΔV, is given by the equation, $\Delta V = \Delta Q/C$. This indicates that ripple is independent of frequency (assuming that the charge, Q, is delivered in a short time relative to the period of the input). The settling time of this type of FVC is determined by the exponential time constant, RC, from which the time to settle within a particular error band may be calculated.

From these equations, we see that the characteristics of this type of FVC are interdependent, and it is not possible to optimize ripple and settling time separately. To do this we must use a PLL.

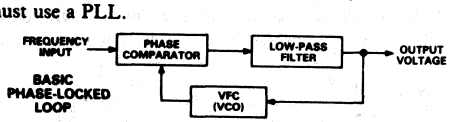

The phase-locked-loop FVC illustrated differs from any othe PLL in only one respect: the voltage-controlled oscillator o the normal PLL, which must be monotonic but not necessar ily linear, has been replaced by a VFC with a linear contro law. In the servo system, negative feedback keeps the VFC's output frequency equal to the input frequency. The outpu voltage, the VFC's input, is accurately proportional to the input frequency.

Designing PLL systems is beyond the scope of this discussion,[1] but if a 4000-series CMOS PLL, the 4046, is used *just as a phase detector* (its VCO's transfer characteristic is not sufficiently linear), we can build the FVC shown here, with an AD654 VFC.

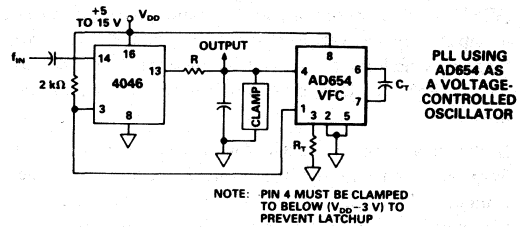

Q. *What is a synchronous VFC?*

A. A charge-balance VFC with improved linearity and stability, where the monostable is replaced by a bistable, driven by an external clock. The fixed time during which the precision current discharges the integrator is one clock period of the external clock.

A further advantage of the SVFC is that the discharge does not start when the integrator passes the comparator threshold (at a non-critical rate), but on the next clock cycle. The SVFC output is synchronous with a clock, so it is easier to interface with counters, µPs, etc.; it is especially useful in multi-channel systems: it eliminates problems of interference from multiple asynchronous frequency sources.

There are two disadvantages. Since the output pulses are synchronized to a clock they are not equally spaced but have substantial jitter. This need not affect the user of a SVFC for a/d conversion, but it does prevent its use as a precision oscillator. Also, capacitive coupling of the clock into the comparator causes injection-lock effects when the SVFC is at 2/3 or 1/2 FS, causing a small (4-6 bit at 18-bit resolution at 1-MHz clock) dead zone in its response. Poor layout or device design can worsen this effect.

Despite these difficulties the improvement in performance produced by the abolition of the timing monostable makes the SVFC ideal for the majority of high-resolution VFC applications.

Q. *Can you have a synchronized FVC?*

A. Yes, and with very good performance; it is best done with an FVC-connected SVFC and a clock that is common to both ends of the transmission path. If the input signal to a synchronized FVC is *not* phase related to the clock, severe timing problems can arise, which can only be solved by the use of additional logic (two D flip-flops) to establish the correct phase relationship. ◘

[1]See Gardner, F. M, *Phase-lock Techniques*, 2nd ed., New York: Wiley, 1979, for more detail; also Analog Devices' *Analog-Digital Conversion Handbook*.

ANALOG DEVICES

AN-276
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Analog-to-Digital Conversion Using Voltage-to-Frequency Converters

by Paul Klonowski

INTRODUCTION

A voltage-to-frequency converter (VFC) is a device which accepts at its input an analog voltage or current signal and provides at its output a train of pulses or square waves at a frequency which is proportional to the input value. A voltage-to-frequency converter can thus be used as a building block in an analog-to-digital (A/D) conversion system, by using the VFC to clock a counter for a certain period of time (the "count time" or "gate time" or "conversion time") and reading the output digital word. This digital word will be proportional to the analog input.

There are a number of advantages to using a voltage-to-frequency converter in an analog-to-digital conversion scheme. First, unlike converters based on binary-weighted networks, monotonicity is inherent under all supply and temperature conditions. Second, the fact that the signal is converted to an easily-transmitted serial bit stream allows the analog circuitry (the VFC and analog signal-conditioning circuits) to be located close to the signal source and the digital circuitry (the counter, timing gate and display circuitry) to be located elsewhere. This is especially advantageous when a large number of channels are required; the remote VFCs can be used to provide "converter-per-channel" data acquisition. Finally, since the digital number is accumulated over a large number of cycles, integration of and therefore reduction of unwanted signals is inherent.

The time required to convert an analog signal into a digital word is related to the maximum full-scale frequency of the VFC and the required resolution of the measurement. For example, the Analog Devices AD650 VFC has a full-scale frequency of 1MHz. If this device is used in an application where a resolution of 16 bits, or 1 part in 65,536 is desired, then the time required to convert the analog signal into a 16-bit digital word will be 65.536 ms. Resolution of 18 bits, or 1 part in 262,144 will require a count time slightly greater than 0.262 seconds. In general, the required count time for an analog-to-digital conversion using a VFC is:

$$T_{COUNT} = \frac{N}{FS_{out}}$$

where N is the number of codes for a given resolution and FS_{out} is the VFC full-scale output frequency.

Although VFC-based A/D converters are slower than successive-approximation and flash converters, they are comparable in speed to integrating A/D converters. VFC-based A/D converters are thus well suited for low frequency applications such as temperature and strain-gauge measurements. The resolution that the VFC can provide in these types of applications offsets the relatively long count time required to acquire the digital word corresponding to the analog input value.

The intent of this application note is to demonstrate different methods of using voltage-to-frequency converters as analog-to-digital building blocks. The focus, then, is on interfacing the output of a VFC. For detailed information on handling a variety of types of inputs (temperature, strain-gauge and photodiode signals), it is suggested that the reader consult the individual AD650, AD651 and AD654 data sheets and the AD654 application note, all available from the Analog Devices Literature Center (Phone # (617) 329-4700). Signal multiplexing schemes, isolation circuits, and individual device details are also included in the material listed above.

PULSE COUNTING

One way to perform an analog-to-digital conversion using a voltage-to-frequency converter is to have a single-chip microcomputer count the number of pulses that occur in a fixed time period. The total number of pulses counted during this period is then proportional to the input voltage of the VFC. For example, if a 1V full-scale input produces a 100kHz signal from the VFC and the count period is 100ms, then the total full-scale count will be 10,000. Scaling from this maximum is then used to determine the input voltage, i.e., a count of 5,000 corresponds to an input voltage of 0.5V.

Figure 1 shows the Analog Devices AD654 VFC output connected to the counter input, T1, of the Intel 8051 microcomputer. The AD654 is a low-cost, single-supply monolithic VFC with a full-scale frequency of up to 500kHz. The 8051 is a member of the Intel MCS-51 family of 8-bit microcomputers whose family members differ mainly in their internal memory capabilities. In the following text, the term "8051" is used to generically refer to all members of the MCS-51 family.

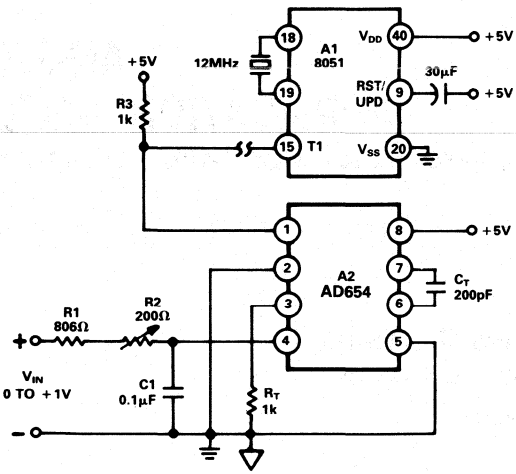

Figure 1. AD654 Pulse Counter

The analog input of the AD654 in Figure 1 is a 0 to +1V signal. The timing resistor and capacitor, R_T and C_T, are selected such that this 0 to +1V signal seen at Pin 4 results in a 0 to 500kHz output frequency. The pull-up resistor, R3, ensures that the AD654 output meets the logic levels required at T1 (Pin 15) of the 8051.

The 8051 has two 16-bit timer/event counters on-chip (the 8052 and the 8032 have three). These counters, Timer 0 and Timer 1, can be programmed independently to operate as 16-bit time-interval or event counters. The use of Timer 0 and Timer 1 is determined by two 8-bit registers, TMOD (timer mode) and TCON (timer control). The TMOD register is shown in Figure 2 below:

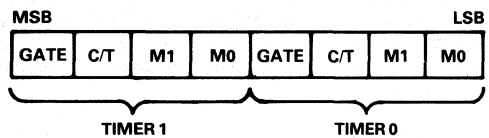

Figure 2. 8051 TMOD Register

"M1" and "M0" are used to select the mode of each timer. Mode 01 configures the timer as a 16-bit time-interval or event counter. "C/T" is the timer or counter selector and is cleared for timer operation. In this application, Timer 0 was configured as the timer (to provide the fixed time interval) and Timer 1 was configured as the counter (to count the pulses). Hereafter, the two timers will be referred to as Timer 0 and Counter 1. When running, Timer 0 increments at a rate equal to the external clock divided by twelve. Using a 12MHz crystal, this corresponds to once every microsecond. "GATE" is the gating control. When this bit is clear, timer/counter "X" is enabled whenever the "TRx" control bit found in the TCON register is set. The "TRx" bit is controlled via software. When the "GATE" bit is set, timer/counter "X" is enabled whenever the "TRx" bit is set and the signal level appearing at the $\overline{INTx}$ pin (Pin 12 or 13 for Timer 0 or 1, respectively) is high. Thus, when a GATE bit is clear that timer is controlled by software only, and when it is set that timer is controlled by a combination of software and hardware. In this application the GATE bits are clear; however in the next application they will be set.

Table I lists the software routine PLSECNT, which counts the number of falling edges that appear at T1 (the Counter 1 input) in a 50ms window. After Counter 1 is cleared, the value 15539 is loaded into Timer 0. Since Timer 0 is a 16-bit timer, the maximum possible count is 65535. With the Timer 0 interrupt enabled, a count of 65536 will cause a jump to the starting address (OBH) of the Timer 0 interrupt service program. With Timer 0 starting at 15539 and incrementing once every µs (based on a 12MHz clock), there will be 49,997 counts or 49.997ms before jumping to the service program. The 3µs difference from 50ms is made up in the speed of the interrupt response. The interrupt response latency ranges from 3µs to 7µs when using a 12MHz crystal. During this 50ms count period, control resides with the main program. Thus the 8051 is not tied up while Counter 1 is counting for 50ms. After the interrupt service is reached, Counter 1 and Timer 0 are stopped and the contents of Counter 1 are moved into RAM, where it may be accessed at the user's convenience. Control is then returned to the main program for which the subroutine was written. With a maximum frequency of 500kHz and a count window of 50ms, the maximum value of Counter 1 will be 25,000. This provides resolution greater than 14 bits. Appropriate scaling from this 1V full-scale reference point may then be performed in software.

```
        ORG     00H
        AJMP    MAIN
PLSECNT ORG     60H             ;PULSE COUNT SUBROUTINE
        MOV     TMOD, #51H      ;Put Time 0 and Count 1 in Mode 01
        MOV     TL1, #00H       ;Initialize Counter 1 Register
        MOV     TH1, #00H
        MOV     TL0, #0B3H      ;Load Time 0 With 15536 + 3.Will
        MOV     TH0, #3CH       ;Ovflw After 50ms + 3µs Delay
        SETB    PT0             ;Prioritize Time 0 Interrupt
        SETB    ET0             ;Enable Time 0 Interrupt
        SETB    EA              ;Enable Global Interrupt
        SETB    TR0             ;Start Timer
        SETB    TR1             ;Start Counter
        RET                     ;Return to Main Program
        ORG     0BH             ;TIME 0 INTERRUPT SUBROUTINE
        CLR     TR1             ;Stop Counter
        CLR     TR0             ;Stop Timer
        AJMP    COUNT
        ORG     40H
COUNT   MOV     50H,TL1         ;Move Counter Contents Into RAM
        MOV     51H,TH1
        RETI                    ;Return from Interrupt
        ORG     100H
MAIN    -       -               ;Main Program for Which PLSECNT
                                ;Subroutine Was Written
```

Table I. 8051 Pulse Count Routine

PERIOD TIMING

Another method of performing analog-to-digital conversion using a voltage-to-frequency converter (VFC) and a microcomputer is to have the microcomputer time the period of the VFC output frequency. For example, an output frequency of 25kHz has a period of 40μs. If a timer that is incremented once per microsecond is gated to this signal, a total count of 40 will result. An output frequency of 250Hz has a period of 4ms. The same timer gated to this period signal will then produce a total count of 4000.

One of the advantages of period timing over pulse counting is that the count window is dependent upon the output frequency of the VFC; in many cases the count window will be shorter for period timing than for pulse counting. This will be especially important in systems where a number of channels are being converted. Recall that in the Pulse Counter application the count window was 50ms—whether the output frequency was 50kHz or 50Hz. With Period Timing, the count window is the inverse of the output frequency. Thus, a 50kHz signal will have a count window of 20μs, while a 50Hz signal will have a count window of 20ms. In fact, it's not until the output frequency reaches 20Hz that the Period Timing count window is equal to the 50ms Pulse Counter count window.

Figure 3 shows the circuitry necessary to perform an analog-to-digital conversion via period timing using the Analog Devices AD650 voltage-to-frequency converter. The AD650 has a maximum full-scale frequency of 1MHz with a nonlinearity error of 0.1% max at this frequency. The AD650 in Figure 3 is configured such that a 0 to +10V input will result in a 0 to 50kHz output. Because the AD650's output is made up of pulses, the SN7474 D-type flip-flop is used to convert these pulses into a square wave. The SN7474 Pin 3 and Pin 5 waveforms sketched in Figure 3 show that the width of either the high-level or low-level output appearing at Pin 5 is the same as one period of the AD650 output frequency. Note also that when Pin 1 on the SN7474 is held low, Pin 5 is also held low.

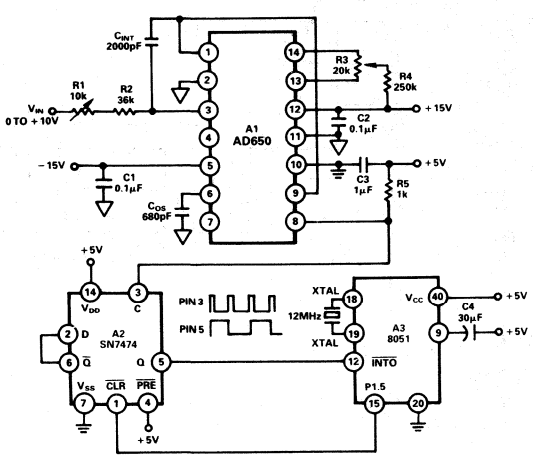

Figure 3. AD650 Period Timing

Recall that the $\overline{INTO}$ pin (pin 12) on the 8051 is the Timer 0 gate pin. (See Pulse Counter discussion). When the "GATE" bit is set in the TMOD register, Timer 0 will run only when $\overline{INTO}$ at Pin 12 is high and the TR0 bit in the TCON register has been set via software. Thus, connecting the "Q" output of the SN7474 to the $\overline{INTO}$ pin on the 8051 will ensure the timer runs for one period of the AD650 frequency.

One problem that could arise is the TR0 bit in software being set in the middle of a high-level edge at the $\overline{INTO}$ pin. In this case, Timer 0 would run for a fraction of a period rather than one full period. This is countered by tying the 8051 Port 1 Bit 5 (P1.5) pin to the SN7474 $\overline{CLR}$ pin. When $\overline{CLR}$ is low and $\overline{PRE}$ is high, Q is low. When $\overline{CLR}$ and $\overline{PRE}$ are both high, Q changes state on every positive edge appearing at the clock (C) pin. Setting P1.5 low, setting TR0 (in software) and then setting P1.5 high will thus ensure that Timer 0 runs for one full period.

Table II shows the software subroutine PCNT, which increments Timer 0 once per microsecond for one AD650 output frequency period. Note that there are two interrupt service programs, one for $\overline{INTO}$ and one for Timer 0. The $\overline{INTO}$ service program is accessed after a negative edge appears at the $\overline{INTO}$ pin (Pin 12), signifying the end of one period. The timer is then stopped and its contents are loaded into RAM. The user may then access the contents at his convenience.

```
        ORG   00H
        AJMP  MAIN
PCNT    ORG   90H          ;PERIOD COUNTER SUBROUTINE
        MOV   TMOD, #05H   ;Put Time 0 in Mode 1, Enable INTO Pin
        CLR   P1.5         ;Initially Set INTO Pin Low
        SETB  IT0          ;Specify Edge Triggered Interrupt
        MOV   TL0, #00H    ;Initialize Timer
        MOV   TH0, #00H
        SETB  EX0          ;Enable INTO Interrupt
        SETB  ET0          ;Enable Timer 0 Interrupt
        SETB  EA           ;Enable All Interrupts
        SETB  TR0          ;Start Timer
        SETB  P1.5         ;Enable Gate INTO Pin
        RET                ;Return to Main Program
        ORG   03H          ;INTO Subroutine Service Program
        CLR   TR0          ;Stop Timer
        CLR   EA           ;Disable Interrupts
        AJMP  COUNT        ;Jump to Count
        ORG   0BH          ;TIMER 0 SUBROUTINE
                            SERVICE PROGRAM
        CLR   TR0          ;Stop Timer
        CLR   EA           ;Disable Interrupts
        AJMP  OFLW         ;Jump to OFLW
        ORG   40H
OFLW    MOV   60H, #FF     ;Load RAM With Overflow
        MOV   61H, #FF     ;Value
        CLR   P1.5         ;Set INTO Pin Low
        RETI               ;Return From Subroutine
        ORG   50H
COUNT   MOV   60H,TH0      ;Load RAM with Counter Contents
        MOV   61H,TL0
        CLR   P1.5         ;Set INTO Pin Low
        RETI               ;Return From Subroutine
        ORG   100H
MAIN    -     -            ;Main Program for Which
                           ;Subroutine Was Written
```

Table II. Period Timing Routine

The Timer 0 service program serves to limit the count window to approximately 65.5ms. It is accessed when the contents of Timer 0 reach 65536. This will occur when the AD650 input voltage is approximately 3.05mV, or the output frequency is approximately 15.26Hz. This service program then loads the overflow value "65535" into RAM. After both interrupt subroutines, and after initialization of the PCNT subroutine, control returns to the main program for which the subroutine was written. Thus the 8051 is not tied up while the period timing is occurring.

One possible source of error in this application is jitter, which is the range of variation in the period of the output frequency. This variation in period would result in a variation in the number of pulses counted from one period to the next. The magnitude of this error can be greatly reduced in software by taking the average of a number of period counts and using this average value for calculations.

COMPLETE 16-BIT RESOLUTION ANALOG-TO-DIGITAL SYSTEM

A complete 16-bit resolution analog-to-digital conversion system is shown in Figure 4. This system uses the Analog Devices AD651 as its VFC. The AD651 is a 2MHz full-scale output synchronous voltage-to-frequency converter, requiring an external clock to generate the full-scale frequency rather than relying on the stability of passive components. This allows the AD651 to achieve linearity and stability far superior to any other monolithic VFC available. Other key elements in this system include an Intersil 7208 single chip counter-decoder-LED driver, a 4MHz TTL oscillator, and two 4020B binary counters.

The AD651 is configured in the 0 to + 10V input, 2MHz output mode. Pull-up resistor, R1, is used to feed the AD651 output frequency directly into the counter input pin of the 7208 counter-decoder-driver. The 4020Bs are 14-stage binary ripple counters which have clock and master reset inputs (Pins 10 and 11), and buffered outputs from the first stage and the last eleven stages. The 4020Bs are used to generate the count window for the 7208. By dividing the 4MHz TTL clock by 2^{18} (two divide by 2^9 stages cascaded) a 15.2588Hz signal results. The 7208 will count the negative edges of the AD651 output pulses when pin 13 is low. With a 15.2588Hz signal, the count window will be 32.768ms.

Besides being tied to the count enable pin, the 15.26Hz signal is also tied to the 74LS221 dual monostable multivibrator.

Figure 5 shows the purpose of the 74LS221 one-shot. After the count window on the 7208 has been closed (that is, the count enable input has gone high), the data must be latched and then decoded in order to drive the LEDs. After this data has been latched, the counter must then be cleared in order to get an accurate count during the next count window. The purpose of the 74LS221 is to provide the required $\overline{\text{STORE}}$ input and $\overline{\text{RESET}}$ input pulses. Although the 7208 only requires pulse widths greater than 50µs, $R_3 \cdot C_2$ and $R_4 \cdot C_3$ have been selected to provide a pulse width of approximately 500µs.

The 7208 automatically handles the decoding and driving of the LEDs. The only other components required are resistors R_5 and R_6 and capacitor C_4. These non-critical passive components control the display multiplex rate of the

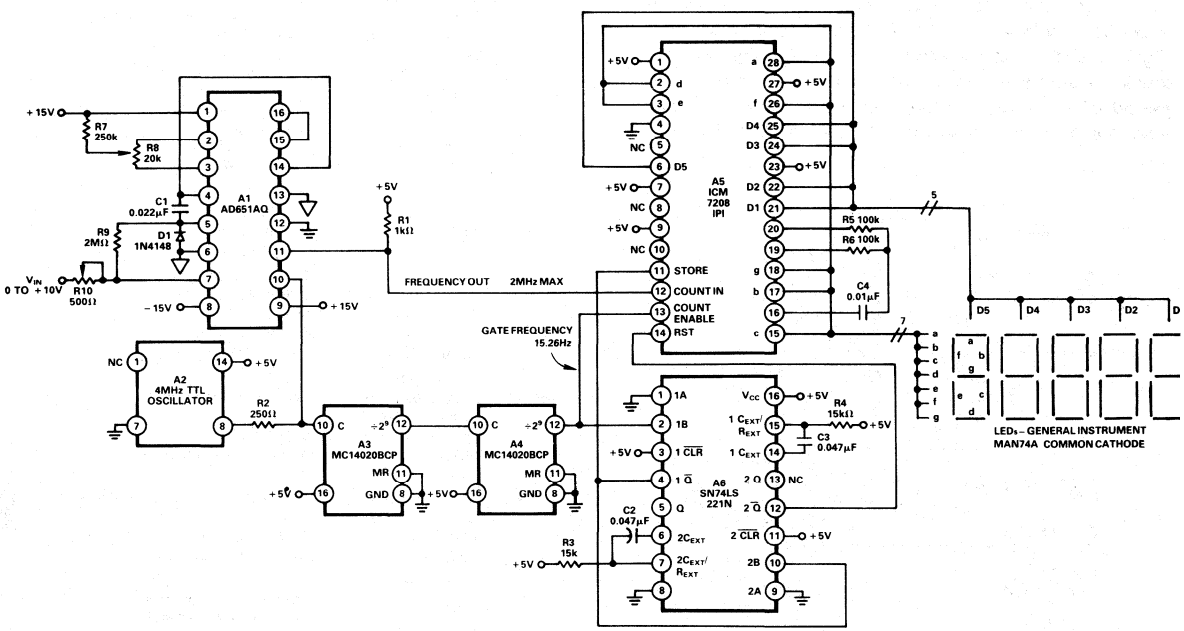

Figure 4. Complete 16-Bit Analog-to-Digital System

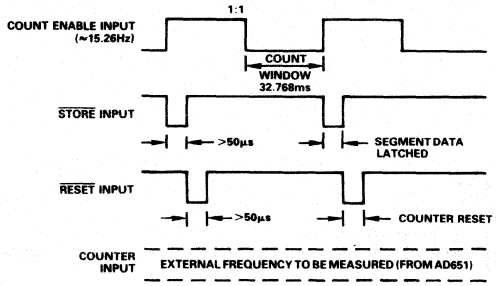

Figure 5. ICM7208 Timing Diagrams

EDs and were selected following the manufacturer's uidelines to provide a multiplex rate between 50 and 00Hz.

o achieve maximum performance with the AD651, the xed gate interval (or the "count window") should be enerated using a multiple of the AD651 SVFC clock input, s was done in this application. Counting in this manner liminates any errors due to the clock (whether it be jitter, rift with time or temperature, etc.) since it is the ratio of he clock and output frequencies that is being measured.

he resolution of the A/D conversion, of course, is deter- nined by the clock frequency and the gate time. If, for nstance, a resolution of 12 bits is desired and the clock requency is 1MHz (resulting in an AD651 FS frequency of 00kHz) the gate time will be:

$$\left(\frac{\text{FS Freq}}{N}\right)^{-1} = \left(\frac{1}{2}\,\frac{\text{Clock Freq}}{N}\right)^{-1} = \left(\frac{1\text{MHz}}{2(4096)}\right)^{-1}$$

$$= \frac{8192}{1\times10^{6}}\ \text{sec} = 8.192\text{ms}:\ \text{where N is the total number of codes for a given resolution}$$

Table III shows the relationship between the AD651 clock frequency and gate time for various degrees of resolution. n this application, 16 bits of resolution were desired using 4MHz clock, making the required gate time 32.77ms.

GATING OFF A KNOWN INTERFERING SIGNAL

One source of error in analog-to-digital systems is an interfering signal that couples into the analog signal to be converted. For example, undesired coupling of power-line energy often appears as a sine wave riding on top of the DC level to be converted, resulting in an erroneous digital output. Since the frequency of this interfering sine wave is well-known (50 or 60Hz), the errors caused by the pickup can be integrated out by using a gating time equal to a multiple of the period of the sine wave. A replica of the interfering signal can be picked off of a nearby transformer and fed into a phase locked loop, a block diagram of which is shown in Figure 6. This loop provides two output signals – a high frequency clock (at a high harmonic of the interference) and a gating clock (at a lower harmonic of the interference).

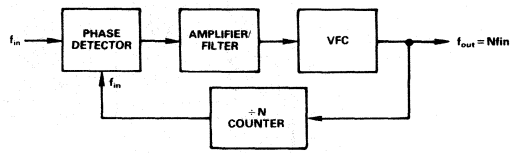

Figure 6. Phase Locked Loop

By using f_{OUT} as the AD651 clock source and by using f_{IN} from the divide by N counter as the count window source, a resolution of one part in 1/2 N is achieved, where N is the "÷N" of the counter. If the count window is level-triggered rather than edge-triggered (as with the 7208), then the resolution will be 1/4 N.

Figure 7 shows the hardware required to implement Figure 6. The MC4044 in Figure 7 contains both the phase detector and the amplifier/filter shown in the block diagram. The interfering signal was 60Hz and was converted into a TTL level signal to feed into the MC4044. Components R_1, R_2 and C_1 were selected to allow for a 50Hz or a 60Hz interfering signal. The error voltage from the MC4044 is then fed into the AD654, which is configured in a 0 to +1V input, 0 to 500kHz output mode. Because the 74LS393

Resolution	N	Clock	Conversion or Gate Time	Typ Lin	Comments
12 Bits	4096	81.92kHz	100ms	0.002%	50,60,400HzNMR
12 Bits	4096	2MHz	4.096ms	0.01%	
12 Bits	4096	4MHz	2.048ms	0.02%	
4 Digits	10000	200kHz	100ms	0.002%	50,60,400HzNMR
14 Bits	16384	327.68kHz	100ms	0.002%	50,60,400HzNMR
14 Bits	16384	1.966MHz	16.66ms	0.01%	60HzNMR
14 Bits	16384	1.638MHz	20ms	0.01%	50HzNMR
4 1/2 Digits	20000	400kHz	100ms	0.002%	50,60,400HzNMR
16 Bits	65536	655.36kHz	200ms	0.002%	50,60,400HzNMR
16 Bits	65536	4MHz	32.77ms	0.02%	

Table III. AD651 Clock Frequency and Gate Time Relationships for Various Resolutions

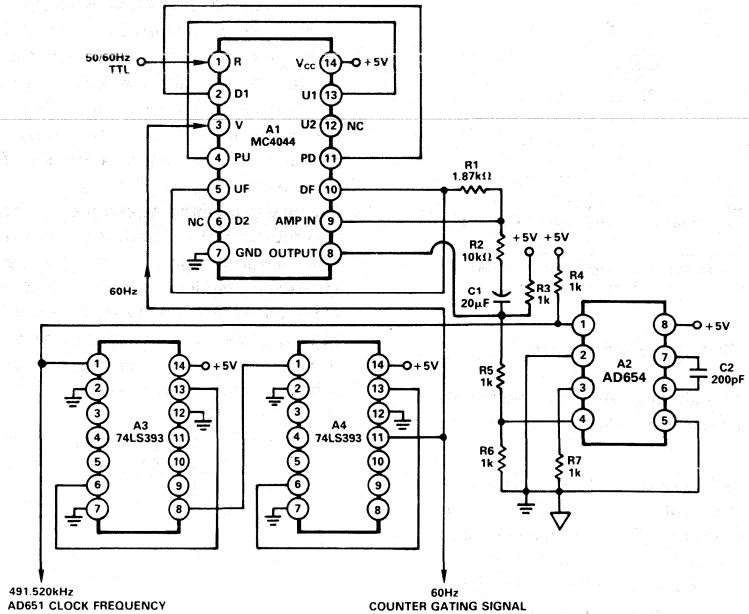

Figure 7. Gating Off a Known Interfering Signal

dual 4-bit binary counters are configured to provide a "÷N" of 8192, the output frequency of the AD654 will be 491520Hz. This signal is used as the clock for the AD651. The output of the second counter is 60Hz and is both fed back into the MC4044 and used as the frequency counter gate signal.

In this system it is also possible to get higher resolution by using Pins 10, 9 or 8 of A4 as the frequency counter gate signal. Pin 11, though, must be fed back into the MC4044. From Pin 1 of A3 to Pin 11 of A4, N = 8192 or 2^{13}. This will supply 12-bit resolution. Using Pins 10, 9 or 8 to provide the gate frequencies will supply resolutions of 13, 14 or 15 bits, respectively.

MC6801-AD650 BASED ANALOG-TO-DIGITAL CONVERSION

In some applications, it may be necessary to perform an analog-to-digital conversion with a microprocessor that either does not have an on-chip timer/counter, or does not have an available on-chip timer/counter. Using some additional hardware, it's still possible to count the output pules of a VFC for a fixed period of time and store this count in the microprocessor's RAM, when the user may access it at his convenience.

Figure 8 shows the circuitry necessary to perform the analog-to-digital conversion using the AD650 VFC and the MC6801 μP, and assumes that the on-chip counter is dedicated to another function. The MC6801 is an 8-bit, single-chip microcomputer with 2048 bytes of ROM, 128 bytes of RAM, a serial communications interface, and a three function programmable timer. The AD650 VFC is configured in the 0 to + 10V input, 0 to 1MHz output mode.

The additional hardware necessary to count the output pulses of the AD650 includes two 74590 8-bit binary counters with output registers, one 4020B 14-stage binary counter, one 7474 dual D-type flip-flop with preset and clear, and one inverter of a 74LS04 hex inverter.

The 4020B and 7474 provide the timing signal which tells the counters when to start and stop counting. The input to the 4020B is tied to the 'E' Pin (Pin 40) of the MC6801. The signal at the 'E' Pin is simply the MC6801 external crystal frequency divided by four, or in this case 1.2288MHz. The 4020B divides the 1.2288MHz signal by 2^{14}, which results in a 75Hz signal being fed into Pin 3 of the 7474. The 7474 further divides this signal by four, to 18.75Hz. This signal will appear at Pin 9 depending upon the signal level at Pins 1 and 10. If Pins 1 and 10 are low, Pin 9 will be held at a TTL high. If Pins 1 and 10 are high, the 18.75Hz square wave will appear at Pin 9. When Pin 9 is high, the 74590 counters are disabled and no counting occurs; if the output of Pin 9 is a square wave, the counters will be enabled during the low level of the period and will thus count for 26.67ms. Note that it is Bit 4 of Port 1 (P14) that controls whether the counters will be enabled or not. If this bit is low, no count will occur. Note also that Pin 9 is connected to the external interrupt request ($\overline{IRQ1}$) pin of the MC6801 through an inverter. This produces two results when Pin 9 changes from low to high. First, the counters are disabled since $\overline{CCKEN}$ of A2 is tied to Pin 9. Secondly, the low-to-high transition is inverted and generates an interrupt request on the $\overline{IRQ1}$ line. The MC6801 then clears P14 to prevent another count from occurring before the 74590 counter values are read.

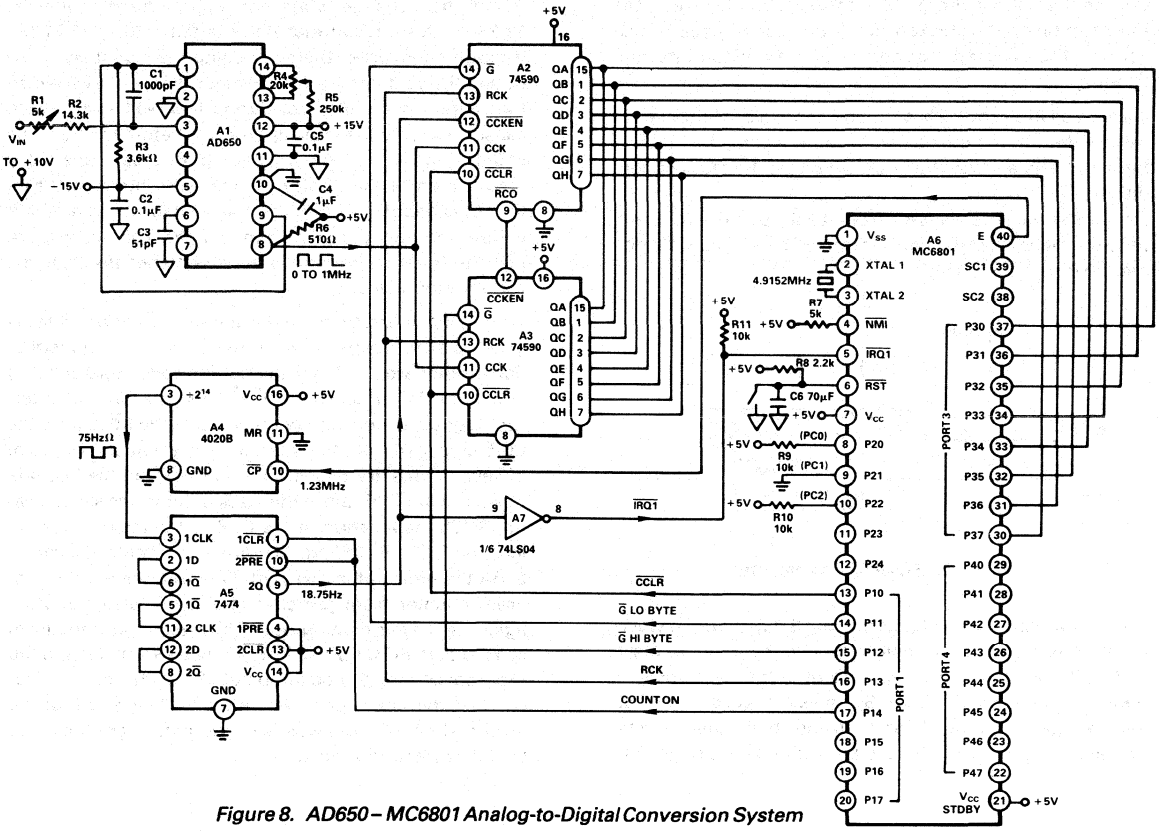

Figure 8. AD650–MC6801 Analog-to-Digital Conversion System

All of the counting events are controlled by the signal levels of the different bits of Port 1. By writing different values to Port 1 the 74590 counters are cleared, enabled, disabled, latched, and read. The values which control these functions are shown in Table IV below.

Port 1 Configuration

Event	P4	P3	P2	P1	P0	Hex
Clear Counters	0	0	1	1	0	06
Enable Counters	1	0	1	1	1	17
Disable Counters	0	0	1	1	1	07
Latch Data	0	1	1	1	1	0F
Output Hi Byte	0	1	1	0	1	0D
Output Lo Byte	0	1	0	1	1	0B

Table IV. Port 1 Event Control Values

Table V shows the software routine that controls the AD650 pulse counter routine. This routine sets the stack pointer, disables the interrupts, clears the 74590 counters and then enables the interrupts after a delay period (this will be discussed later). After the interrupt request is triggered by a low-to-high transition at Pin 9 of the SN7474, the program counter jumps to the interrupt routine. This routine disables any further interrupts, turns off the 74590 counters and reads the lo and high bytes,

```
                        ORG   0100          COUNT ROUTINE
8E 00C0  BEGIN          LDS   #$C0          Set Stack Pointer
0F                      SEI                 Disable Interrupts
86 06                   LDAA  #$06          Clear Counter
97 02                   STAA  $02
86 17                   LDAA  #$17          Turn On Counter
97 02                   STAA  $02
86 2F                   LDAA  #$2F          Insert Delay >13.33ms To
C6 7F           AGN     LDBB  #$7F          Prevent Interrupt Error
5A              CNT     DECB
2E FD                   BGT   CNT
4A                      DECA
2E F8                   BGT   AGN
0E                      CLI                 Allow Interrupt
39                      RTS                 Return From Subroutine
                        ORG   FFF8
0080                    AIRQ1 FDB   LDCNT   Define Interrupt Routine
                                            Start Point
                        ORG   0080          INTERRUPT ROUTINE
0F              LDCNT   SEI                 Disable Interrupt
86 07                   LDAA  #$07          Turn Off Counter
97 02                   STAA  $02
86 0F                   LDAA  #$0F          Latch Data In Counter
97 02                   STAA  $02
86 0D                   LDAA  #$0D          Output Counter Lo Byte
97 02                   STAA  $02
96 06                   LDAA  $06           Read Byte From Port 3
97 A7                   STAA  $A7           Store in Location 00A7
86 0B                   LDAA  #$0B          Output Counter Hi Byte
97 02                   STAA  $02
96 06                   LDAA  $06           Read Byte From Port 3
97 A6                   STAA  $A6           Store In Location 00A6
86 07                   LDAA  #$07          Turn Off Counter
97 02                   STAA  $02
86 10                   LDAA  #$10          Set The Interrupt Bit On The
9A BA                   ORAA  $BA           Stacked Condition Code Register
97 BA                   STAA  $BA
3B                      RTI                 Return From Interrupt
```

Table V. MC6801 Pulse Count Routine

storing them in locations 00A7 and 00A6 respectively. The interrupt bit on the stacked condition code register is then set and the program counter returns from the interrupt routine.

Two areas in the software routine need to be discussed in further detail – the time delay before enabling the MC6801 interrupt and setting the interrupt bit on the stacked condition code register.

Figure 9 shows the waveforms that will appear around the time the 74590 counters are instructed to start counting by setting Pins 1 and 10 high.

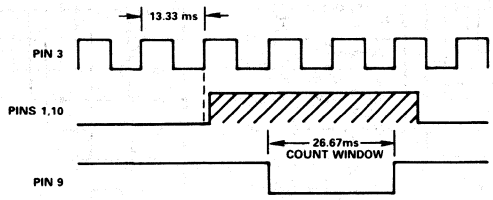

Figure 9. SN7474 Waveforms

Recall that Pin 9 will only output an 18.75Hz square wave when Pins 1 and 10 are held high. Pin 9 will not change state after Pins 1 and 10 go high until the next positive edge occurs at Pin 3. Figure 9 shows the worst case that will happen – Pins 1 and 10 go high just a few nanoseconds after Pin 3 sees a positive going edge. Pin 9 will not change state for approximately another 13.33ms. Also, recall that Pin 9 is fed to the $\overline{IRFQ1}$ pin through an inverter, thereby setting the signal level low during the 13.33ms wait state. Since $\overline{IRFQ1}$ is level sensitive, allowing an interrupt during this wait state would cause a jump to the interrupt routine before the 26.67ms count window has been opened up. By inserting a delay greater than 13.33ms before allowing the interrupt, the interrupt is enabled during the count window (when Pin 9 is low, and thus the $\overline{IRFQ1}$ pin is high). This will ensure that the interrupt routine is accessed after the 26.67ms count window closes.

Setting the interrupt bit on the stacked condition code register is also important. After executing the CLI instruction in software, Bit 4 of the condition code register (CCR) is cleared. This enables the $\overline{IRQ1}$ interrupt. Once an interrupt request is detected (a low level at the $\overline{IRQ1}$ pin), the CCR is pushed onto the stack in its present state. A total of seven bytes are pushed on the stack and the last one pushed is the CCR. Since the stack pointer was set at location 00C0 the location of the CCR is 7 bytes down, or at location 00BA. By OR'ing 10H with the contents of location 00BA the interrupt is disabled. This will prevent the program counter from jumping into the interrupt routine again right after coming out of it, even though the level sensed at the $\overline{IRQ1}$ pin is low. If this bit were not set, the program counter would jump out of the interrupt routine, see that the interrupt enable bit was clear and that the level at the $\overline{IRFQ1}$ pin was low, and jump right back into the interrupt routine again.

ANALOG DEVICES

AN-277
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Applications of the AD537 IC Voltage-to-Frequency Converter

by Barrie Gilbert and Doug Grant

I INTRODUCTION

Product Description

The AD537 is a monolithic voltage-to-frequency converter combining simplicity of use in standard applications with an unusual degree of flexibility and versatility. Designed for either single- or dual-supply operation at low voltages and current (5V and 1mA) the AD537 has the capability for driving high-voltage, high-current loads (36V and 20mA). The chip includes an accurate band-gap reference generator, a low-drift input amplifier capable of operating directly from millivolt signals, a precision current-controlled oscillator and a high-current output stage. It is a complete circuit, using low temperature coefficient silicon-chromium thin-film resistors throughout, except for one external resistor and one capacitor. These provide the user with a means for programming the full-scale input voltage from 100mV to 10V (or greater, depending on the positive supply voltage in some cases), and the full-scale (FS) frequency to any value less than 150kHz. Either positive or negative input voltages can be accepted. The scaling relationship is simply f = V/10RC, simplifying the choice of the external components. Linearity error is as low as ±0.07% for 10kHz FS guaranteed over an 80dB dynamic range. The output frequency is stable with temperature (typically ±50ppm/°C excluding the effects of external components) and supply (typically ±0.01%/V from 5 to 36V).

The specially-designed input amplifier, which actually functions as a voltage to current converter, has a typical drift of only ±1μV/°C when nulled, which permits operation directly from such low-level transducers as strain-gauges, thermocouples, current-shunts, etc., while offering a high (250MΩ) input resistance to positive signals. The output stage, an open-collector NPN circuit, can sink up to 20mA with a saturation voltage less than 0.4V, and withstand a supply of 36V. The Logic Common terminal can be connected to any level between ground (or $-V_S$) and 4V below $+V_S$, permitting easy interfacing with any digital logic family. The high current capability means that LED's, long cables, or up to 12 TTL loads can be driven directly. Unlike most V-F converters, the AD537 is designed to deliver a square-wave output; this has advantages both internally (the power dissipation is essentially independent of frequency, so that self-heating effects do not cause linearity errors) and externally (the average level of the output is constant, useful in ac-coupled links; it also allows operation as a phased-locked-loop in F-V and other applications).

Two auxiliary outputs are provided to enhance the versatility. The first is a fixed voltage of 1.00V, generated by the primary band-gap reference circuit. This is useful in many applications described here; for example, it can be used to power a resistive transducer, or provide the reference voltage for an associated DAC. The second is an output scaled +1.00mV/°K which enables the circuit to be used as a reliable temperature-to-frequency converter; the error due to self-heating can be kept to less than 1°K using a 5V supply and light output loading. Used in conjunction with the fixed reference output, offset scales such as Celsius or Fahrenheit can be accommodated.

Whereas most V-F converters are limited in their range of applications to basic telemetry and slow A-D conversion, the low cost, low power drain and high flexibility of the AD537 open up a wide spectrum of uses. The long-term stability is excellent, and assured by subjecting every device to a stabilization bake and ten temperature cycles of -65°C to +150°C prior to final test. The circuit is available in three performance grades in a 14-pin ceramic package or 10 pin hermetic metal can. For complete specifications see the data sheet.

Theory of Operation

A block diagram of the AD537 is shown in Figure 1. The key to accurate operation is the current-to-frequency converter, which is a very carefully designed multivibrator. The main advantages of the circuit are its simplicity, requiring only a single timing capacitor; the use of push-pull charging, resulting in a large voltage across this capacitor even at low supply voltages, for improved timing accuracy and low cycle-to-cycle jitter; its square-wave output, generally more useful than the narrow pulse generated by charge-dispensing converters; and its good linearity. By using special adaptive biasing techniques operation of this multivibrator is possible over a very large dynamic range, from a maximum control

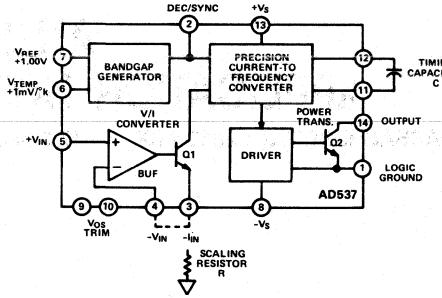

Figure 1. Block Diagram of AD537

current of 2mA to less than 100nA. It can be shown that the basic circuit has a well-defined temperature-coefficient of 300ppm/°K (at all values of control current), and use is made of the tight thermal coupling to the associated band-gap reference generator, which supplies an exactly-proportioned temperature-compensation voltage to the multivibrator. The band-gap cell also provides all the required bias for internal operation of the chip, the fixed reference output and the temperature-proportional output.

The square-wave output of the multivibrator operates the output driver which provides a floating drive current to the large-geometry output transistor, Q2 (Figure 1). Internally, there are actually two transistors (in a thermally symmetric layout for minimal interaction with the remaining circuits) designed for very low saturation voltage and driven to combine low ON voltages with a well-defined current limit.

The SYNC input (pin 2) allows the oscillator to be slaved to a master clock if desired. It also permits control of the state of the output. These uses of the SYNC pin are dealt with in later sections.

A versatile operational amplifier (BUF) serves as the input stage. Its purpose is to convert the applied input voltage to a control current in Q1 (Figure 1); this current is optimally 1mA at the input corresponding to the maximum output frequency. By strapping pins 3 and 4 the input voltage can be impressed across the external scaling resistor, R, which is chosen to provide the needed transconductance for the application. For example, for a FS input voltage of 2.5V the optimum resistor value in most cases will be 2.5kΩ. The $+V_{IN}$ terminal of the op-amp (pin 5) offers a high input resistance (250MΩ) with a bias current of about 100nA. The design of the op-amp ensures that the effect of finite bias current on drift is small, that is, drifts of the order of 1μV/°C can be achieved even with source-resistance/scaling-resistance imbalances of 1kΩ. Consequently, the AD537 can accommodate millivolt signals without the need for a pre-amplifier. The input configuration is also quasi-differential, so that errors due to ground-loops can be avoided by proper choice of signal connections. The common-mode range of the input extends from 4V below $+V_S$ (that is, from +11V for a $+V_S$ of 15V) right down to $-V_S$, so that inputs down to ground potential can be accepted even when operating from a single supply voltage. Negative inputs (voltages or currents) can be accepted, by fixing the voltage on pin 5 (usually to ground potential) and driving pins 3 and 4.

The precision voltage reference is based on Brokaw's cel and incorporates some novel features to ensure a smal spread in output voltage without the need for laser-trim ming. The internal voltage of approximately 1.22V is resis tively-divided to provide the output of 1V at pin 7, whicl therefore has a finite output resistance (380Ω nominal) Loading errors will occur in applications where this voltag is used to drive external components. This is clarified in a later section. Likewise, the temperature-proportional out put at pin 6 has a nominal source resistance of 900Ω, anc loading must be considered. When the AD537 is used in the thermometer mode, the loading of the input op-amp (pin 5) is negligible.

Pins 9 and 10 are provided for trimming the offset voltage of the input op-amp to exactly zero, using an external pot whose slider is connected to $+V_S$. However, in many cases the low initial offset of the AD537 (2mV guaranteed for the AD537K) eliminates the need for trimming (it amounts to a zero-error of only 0.02% for a 10V input) and the device operates correctly without this pot. Note that no damage will result to the circuit if pin 9 or 10 is inadvertently connected to $-V_S$. Offset trimming is not available on the 10-pin metal can version of the AD537 due to pin limitations.

For further details of the internal design of the AD537, refer to B. Gilbert, "A Versatile Monolithic Voltage-to-Frequency Converter", *IEEE J. Solid State Circuits,* Vol SC-11, P852-864, Dec. 1976.

II BASIC OPERATING MODES

Positive Input Voltage

Figure 2 shows the connections for basic operation of the AD537 as a V-F converter with positive voltage inputs. Usually a single supply will be sufficient, with $-V_S$ and LOGIC GND strapped to the supply ground. The lower end of the scaling resistor should be connected to the appropriate signal ground to avoid errors due to the inevitable on-board voltage drops. The nominal value of R is chosen such that the FS input voltage sets up a current of 1mA in it. Thus, for a 0 to +10V input, R would be nominally 10kΩ. It is a simple matter to recalculate R for other FS voltages; note that with $+V_S$ = 5V the maximum input is +1V. However, the specifications allow for the fact that the supply may be 5% low, and in practice reliable operation can be achieved for inputs slightly in excess of +1V. Full-scale currents other than 1mA can be chosen, with some degradation of linearity.

In precision applications sufficient trim range must be provided to accommodate the scaling error of the AD537 (±7% max for the J, ±5% max for the K and S grades) and the tolerance on the timing capacitors, which might be ±3%. In the example just given, the fixed part of R could be 9.09kΩ and the scaling adjustment a 2kΩ pot of good resolution and stability. To accommodate an input of +10V the $+V_S$ supply must be at least +14V ($+V_{IN}$ + 4V); using a standard 15V supply a 10% over-range is possible.

Having chosen the scaling resistor to fit the FS input voltage requirement, the timing capacitor, C, is simply

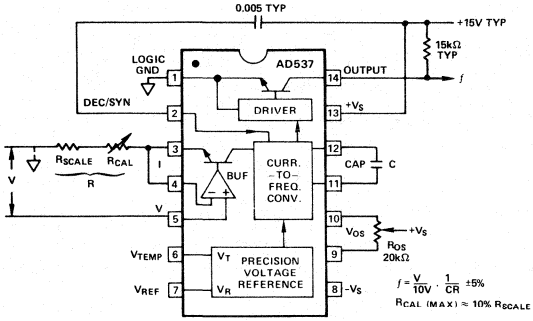

a. 14 Pin Dip

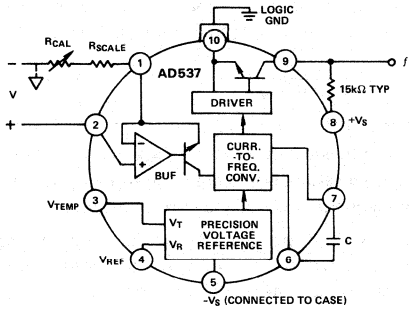

b. 10 Pin Metal Can

Figure 2. Basic (Positive Input) Connections

calculated from the basic relationship

$$f = \frac{V}{10} \cdot \frac{1}{CR}$$

Thus, for a 10V FS input, the FS frequency is just 1/CR.

The timing capacitor will usually be either 0.01μF for 10kHz FS (1Hz/mV) or 1000pF for 100kHz FS (10Hz/mV), assuming R is 10kΩ. However, any convenient value can be used to set up a special scaling relationship. Linearity will be degraded with values below 1000pF, and 100pF is the minimum recommended capacitor size. There is no upper limit, but since the voltage on the capacitor reverses polarity every half-cycle, non-polarized types must be employed.

Good linearity requires the use of a capacitor with low dielectric absorption, and stable operation over temperature calls for a component having a small temperature coefficient. NPO ceramic, Teflon, and polystyrene capacitors best serve these needs. Mica and polycarbonate dielectrics are also acceptable, but linearity and stability will be degraded with most other types.

It is desirable to mount the capacitor as close as possible to pins 11 and 12, to avoid the effects of stray capacitance and pick-up.

Many of the applications found in these notes make use of the high sensitivity and low drift of the input amplifier. Special care should be exercised in returning the SIGNAL LOW input to the appropriate signal ground associated with the source, which can be at any potential from $-V_S$ to 4V below $+V_S$. The input resistance at pin 5 is very high (250MΩ typ.) so errors due to non-zero source resistance will not normally be encountered. The input bias current (100nA typ.) will generate an offset voltage of 100μV/kΩ. Where signals of high source resistance are used this offset may be significant and require the use of the offset trim adjustment, which has a range of approximately ±6mV. The input amplifier is designed so that imbalance between the source resistance and the scaling resistance does not cause appreciable drift after nulling has been effected with this adjustment; however, it is good practice to equalize these resistances whenever possible. The compensating resistor may either be included in series with pin 5 (when $R_S < R$ in which case it should be $R - R_S$) or in series with pin 4 (when $R_S > R$ in which case it should be $R_S - R$).

Calibration

There are two independent adjustments: scale and offset. The first is trimmed by adjustment of the scaling resistor R and the second by the (optional) potentiometer connected to $+V_S$ and pins 9 and 10. Precise calibration requires the use of an accurate voltage standard set to the desired FS value and a frequency meter; a scope is useful for monitoring output waveshape. Verification of linearity requires the availability of a switchable voltage source (or a DAC) having a linearity error below ±0.005%, and the use of long measurement intervals to minimize count uncertainties. Every AD537 is automatically tested for linearity, and it will not usually be necessary to perform this verification, which is both tedious and time-consuming.

Although drifts are small it is good practice to allow the operating environment to attain stable temperature, and to ensure that the supply, source and load conditions are those applicable. Begin by setting the input voltage to zero. Adjust the offset pot until the AD537 just ceases to oscillate; this is most easily seen using a scope connected to the output (pin 14). For a 0.01% error it is permissible to adjust slightly above this point, until the output frequency is 1/10,000 of FS (for example, 1Hz for a FS of 10kHz). Then apply the FS input voltage and adjust R until the desired FS frequency is indicated. In applications where the FS input is small, this adjustment will very slightly affect the offset voltage, due to the small input current at pin 4. A change of 1kΩ in R will affect the input by approximately 100μV, which is as much as 0.1% of a 100mV FS range. Therefore it may be necessary to repeat the offset and scale adjustments for the highest accuracy. The design of the input amplifier is such that the input-voltage drift after offset nulling is typically below 1μV/°C.

Negative Input Voltage

By interchanging the SIGNAL HIGH and SIGNAL LOW inputs the AD537 can operate from negative input voltages, as shown in Figure 3. This connection does not offer a high input resistance, since the control current must now be provided by the source. Consequently the input resistance is simply equal to the scaling resistance, R, which is chosen to set up a nominal 1mA at the FS input voltage. Note that

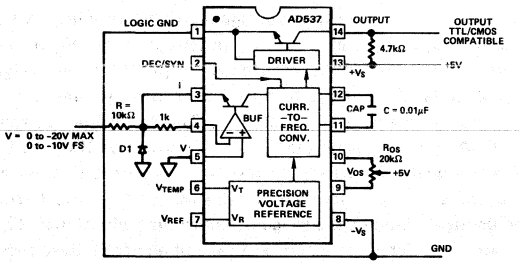

Figure 3. Negative Input Operation

the FS voltage can now be as large as desired, since it is not limited by the input amplifier common-mode range or supply voltage; for example, a -100V FS input would use a scaling resistor of 100kΩ, and the AD537 could operate from a +5V supply. If it is desirable to reduce the loading on the signal a lower FS control current can be used, with some sacrifice in dynamic range. For example, 1MΩ resistor could be used to convert the -100V input to a 100μA control current; the capacitor would then be reduced by a factor of ten to maintain the same FS frequency. This also provides a very large over-range capacity, useful in handling signals which may occasionally have large peak values.

Usually pin 5 will be at ground potential (strictly, at the SIGNAL LOW point), resulting in a current-sum node at pin 3. Thus, any number of signals may be added before conversion, and each source may have any desired scaling by appropriate choice of resistor. It is not necessary for pin 5 to be grounded; offset scales can be generated by setting the voltage on this pin at any value between $-V_S$ and $(+V_S - 4V)$. Diode D1 and the 1kΩ resistor (Figure 3) are added to provide overload and latchup protection.

The linearity in the negative-input mode is better than for positive inputs, since the degrading effect of finite common-mode rejection in the input amplifier is removed. This also allows the use of the minimum supply voltage, resulting in a further improvement in linearity.

The design considerations and trimming procedures are otherwise the same as for the positive-input mode.

Negative Input Current

In some cases the signal may be in the form of a negative current source. This can be handled in a similar way to a negative input voltage. However, the scaling resistor is no longer required, eliminating the capability of trimming full scale in this fashion. Since it will usually be impractical to vary the capacitance, an alternative calibration scheme is needed. This is shown in Figure 4. A resistor-potentiometer connected from the V_R output, pin 7, to $-V_S$ will alter the internal operating conditions in a predictable way, and provide the necessary adjustment range. With the values shown, a range of ±4% is available; a larger range can be attained by reducing R1. This technique does not alter the temperature-coefficient of the converter, and the linearity will be as for negative input voltages. The minimum supply voltage may be used.

Unless it is required to set the input node at exactly ground potential, no offset adjustment is needed. The capacitor C

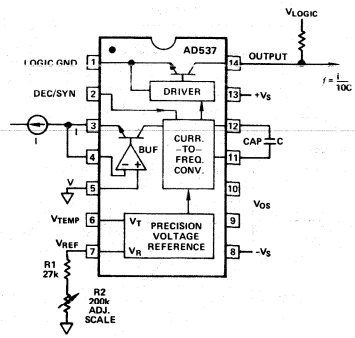

Figure 4. Scale Adjustment for Current Inputs

is selected to be 5% below the nominal value; with R2 in its mid-position the output frequency is given by

$$f = \frac{I}{10.5 \times C}$$

where f is in kHz, I is in mA and C is in μF. For example, for a FS frequency of 10kHz at a FS input of 1mA, C = 9500pF. Calibration is effected by applying the full-scale input and adjusting R2 for the correct reading.

This alternative adjustment scheme may also be used when it is desired to present an exact input resistance in the negative-voltage mode. The scaling relationship is then

$$f = \frac{V}{R_{exact}} \cdot \frac{1}{10.5\ C}$$

The calibration procedure is then similar to that used for positive input voltages, except that the scale adjustment is by means of R2.

Output Interfacing

The AD537 logic output is designed to interface with all digital logic families, LED's (for optical coupling), pulse transformers, and long lines. The open collector (pin 14) and the emitter (pin 1) of the output stage NPN transistor are both available for use; the emitter can be tied to any potential from $(+V_S - 4V)$ down to $-V_S$ and the collector can be pulled up to any voltage up to 36V above the emitter regardless of the $+V_S$ to pin 13.

Figure 5 shows the AD537 connected for a 0 to +10V input with general output interfacing. The required logic common voltage, logic supply voltage, pull-up resistor and $-V_S$ supply are shown in the accompanying table. In the TTL mode, up

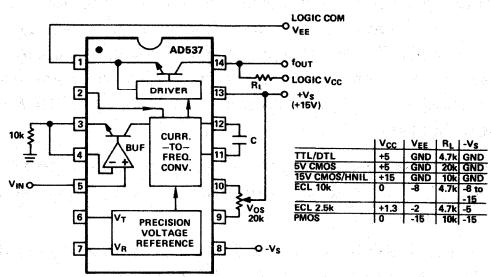

	V_{CC}	V_{EE}	R_L	$-V_S$
TTL/DTL	+5	GND	4.7k	GND
5V CMOS	+5	GND	20k	GND
15V CMOS/HNIL	+15	GND	10k	GND
ECL 10k	0	-8	4.7k	-8 to -15
ECL 2.5k	+1.3	-2	4.7k	-5
PMOS	0	-15	10k	-15

Figure 5. Interfacing Standard Logic Families

to 12 standard gates (20mA) can be driven at a maximum LOW voltage of 0.4V.

The output stage is current-limited and can be shorted indefinitely to pin 1 (normally ground). When shorted to $+V_S$ the current in the ON state is approximately 35mA, and under most conditions this will not result in damage. For example, with pin 1 at ground and pin 14 shorted to +15V, and the oscillator running, the average power in the output stage is 262.5mW (15 x 35/2). However, if high supply voltages are used and the output stage is in the ON state for long periods (low input levels), the peak dissipation of 1260mW (36 x 35) will cause considerable heating, and it is recommended to avoid this situation for prolonged periods.

Signal Isolation

It will often be necessary to couple the frequency output of the AD537 to a system at a different dc level. Where the voltage difference is fixed, or varying relatively slowly, capacitive coupling can be used. Usually, the purpose of the isolation scheme is to reject large amounts of common-mode voltage up to high frequencies, in which case optical or transformer coupling should be used. These also provide a higher degree of isolation safety. Examples are found elsewhere in these notes.

True Two-Wire Data Transmission

Figure 6 shows a useful technique for operating an AD537 at the remote end of a single wire pair, which serves to supply power to the circuit and to transmit the frequency-encoded signal back to a central point. The PNP circuit at the receiving end of the line converts the current modulation imposed by the output of the AD537 back into a voltage signal suitable for driving digital logic; in the example shown the supply current down the line is as follows:

	OUTPUT "OFF"	OUTPUT "ON"
ZERO SIGNAL	1.2mA	5.2mA
FULL-SCALE (1mA)	3.5mA	7.5mA

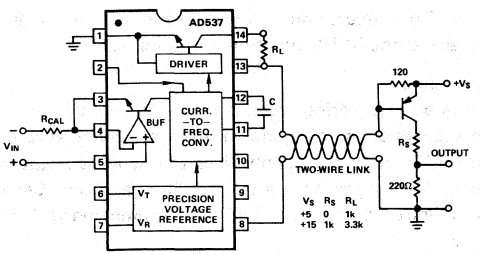

Figure 6. True Two-Wire Operation

The PNP circuit is designed so that the transistor is not appreciably turned on by the supply current to the AD537 in the "OFF" state; a large ratio of currents between FS/OFF and ZERO/ON ensures reliable operation. Where power considerations are of paramount importance a smaller value of FS control current can be employed and a more efficient means of detecting the modulation current would allow the use of a larger R_L. Approximately 500mV of variation appears on the remote end of the supply line; this does not affect operation of the AD537, due to its good supply-rejection ratio.

A useful feature of the AD537 in this application is the provision of a stable 1.00V reference output at the remote location, making it possible to drive various transducers requiring a stable supply. An example might be the conversion of linear position using a resistive-element transducer; other examples will be found later in these notes, including temperature measurement using the output at pin 6.

Signal Multiplexing

Where a variety of transducers, each of different type and sensitivity, and possibly other signal sources, are required to interface with a common data-acquisition system, one solution is the use of separate instrumentation amplifiers for each channel, driving an analog multiplexer which finally drives a V-F converter. The versatility of the AD537 allows it to operate directly from almost any signal source commonly encountered, and a more economic solution in some cases can result by using one V-F per channel and multiplexing their output digitally.

The open-collector feature of the AD537 makes this easy to implement. One method is shown in Figure 7; all V-F circuits are operating continuously, but only the device having its LOGIC COMMON pin grounded, through the open-collector decoder or other digital switching element, transmits its output. An alternative method uses the SYNC input (pin 2). In this method, all V-F's have pin 1 grounded, and all but one have pin 2 shorted to $+V_S$ (using a PNP transistor) which stops the oscillator circuit in these AD537's. This method has the advantage of reducing the risk of crosstalk between converters, but is less convenient to implement.

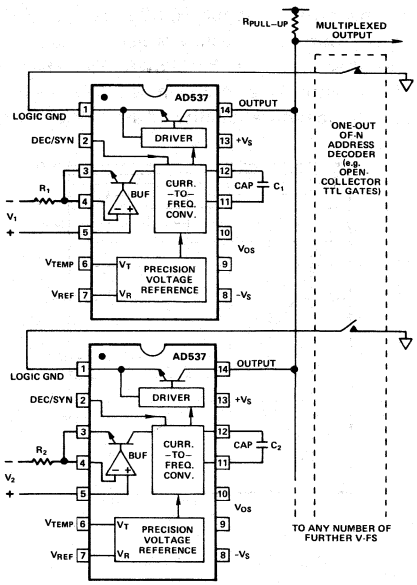

Figure 7. Signal Multiplexing

Two-Phase Output

Using two load resistors, one from the positive logic supply to pin 14 (output) and one from the negative logic supply

to pin 1 (supply ground), a two-phase output can be provided. Assuming equal resistors, and a typical situation where +5V and ground are used as logic supplies, the output at pin 14 is a square-wave from about +2.6V to +5V, and that at pin 1 goes from ground to about +2.4V and has the opposite phase. Unequal resistors can of course be used to alter the amplitude ratio of the two outputs.

Temperature Stability

The stability of the output frequency is determined by several factors. Broadly speaking, they can be grouped into offset drift and scale drift. The design of the AD537 is such that input offset drift is very small—in the 1μV/°C region—and unless the circuit is connected for operation from very low signal levels, this will rarely be a significant source of instability. Scaling drifts arise both in the AD537 itself—less than ±50ppm/°C for the AD537K—and in the external timing resistor and capacitor which must be of adequate quality for the application. In practice, it will be these external components which most seriously degrade stability over temperature; for most applications metal-film resistors are recommended.

Operation with Non-Zero TC

The drift induced by the external components can be corrected by introducing a deliberate drift of the opposite sign into the AD537. For example, if a particular resistor and capacitor cause the AD537 to exhibit an output frequency drift of 150ppm/°C, this drift can be corrected by connecting an 8.3kΩ resistor between the 1mV/°K output (pin 6) and $-V_S$ (pin 8). Other values, giving TC's of +20 to +350ppm/°C, can be found using the graph (Figure 8). Over this range the scaling factor is only slightly affected; the error is about +0.03% per ppm/°C. In the example the output frequency would be nominally 4.5% high, and this error can be allowed for when choosing the scaling resistor. It should be noted that this option also impairs the temperature-stability of the +1V output (pin 7).

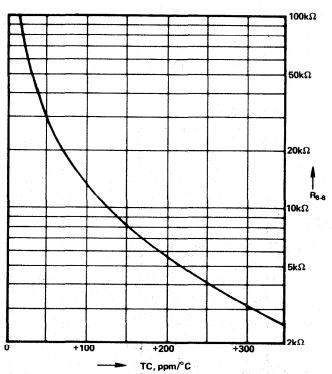

Figure 8. TC Versus R_{6-8}

Some Precautions

The AD537 is intended to be used with a minimum of additional hardware, and the circuits shown in these notes are for the most part complete. However, the successful application of an IC involves a good understanding of possible pitfalls and the use of suitable precautions.

Input Protection

Pins 3, 4 and 5 should not be driven more than 300mV below $-V_S$. This would cause internal junctions to conduct possibly damaging the IC. The AD537 can be protected from "below $-V_S$" inputs by a resistor, R1 and a diode, D1 as shown in Figure 3. It is also desirable not to drive pins 3, 4 and 5 above $+V_S$. In operation, the converter will become very nonlinear for inputs above ($+V_S$ -3.5V). Control currents above 2mA will also cause nonlinearity.

The 80dB dynamic range of the AD537 guarantees operation from a control current of 1mA (nominal FS) down to 100nA (equivalent to 1mV for 10V FS). Below 100nA improper operation of the oscillator may result, causing a false indication of input amplitude. In many cases this might be due to short-lived noise spikes which become added to the input. For example, when scaled to accept a FS input of 1V, the -80dB level is only 100μV, so when the mean input is only 60dB below FS (1mV), noise spikes of 0.9mV are sufficient to cause momentary malfunction.

The problem can be minimized by using a simple low-pass filter ahead of the converter. For a FS of 10kHz a single-pole filter with a time-constant of 100ms will be suitable, but the optimum configuration will depend on the application and type of signal processing. Noise spikes are only likely to be a cause of error when the input current remains near its minimum value for long periods of time; above 100nA (1mV) full integration of additive input noise occurs.

The AD537 is somewhat susceptible to interference from other signals. The most sensitive nodes (besides the inputs) are the capacitor terminals and the SYNC pin. The timing capacitor should be located as close as possible to the AD537 to minimize signal pickup in the leads. In some cases, guard rings or shielding may be required. The SYNC pin should be decoupled through a 0.005μF (or larger) capacitor to pin 13 ($+V_S$). This minimizes the probability that the AD537 will attempt to synchronize to a spurious signal. This precaution is unnecessary on the metal can package, as the SYNC function is not available.

Decoupling

It is good engineering practice to use bypass capacitors on the supply-voltage pins and to insert small-valued resistors (10 to 100Ω) in the supply lines to provide a measure of decoupling between the various circuits in a system. Ceramic capacitors of 0.1μF to 1.0μF are recommended for the AD537.

A decoupling capacitor may also be useful from pin 2 to pin 13 in those applications where very low cycle-to-cycle period variation (jitter) is demanded. The explanation is as follows: the voltage generated by the bandgap circuit is somewhat noisy, and this causes a jitter ratio of about one in 5000 (for example, a period uncertainty of 200ns at 1kHz). By placing a capacitor across pins 2 and 13 this noise is reduced. On the 10kHz FS range, a 6.8μF capacitor reduces the jitter to one in 20,000 adequate for most applications. A tantalum capacitor should be used to avoid errors due to dc leakage.

Analog-to-Digital Conversion

A voltage-to-frequency converter is very useful as part of an analog-to-digital conversion scheme in many applications, for the following reasons. First, unlike converters based on binary-weighted networks, monotonicity is inherent under all supply and temperature conditions. A V-F with a nonlinearity of ±0.01% is thus roughly comparable to a 12-bit A-D with a worst-case differential linearity of ±1/2 Bit. Second, the fact that the signal is converted to a serial bit stream allows the simple analog part (the V-F) to be located close to the signal source, and the digital part (counter, timing gate, display) to be located centrally. As previously mentioned, many AD537's can be multiplexed on to a single transmission path, and share a single digital system. Finally, since the digital number is accumulated over a large number of cycles, integration of unwanted signals is inherent. In fact, the integration period is fixed by the gating interval, so that even for small inputs, when there may be only a few cycles to count, the same amount of signal integration is applied. Such a converter behaves as a filter with attenuation vs. frequency of

$$G(\omega) = \frac{1/\omega T}{\sin(1/\omega T)}$$

where ω is the angular frequency of the interfering signal and T is the digital gating time. Of course, the need to use a long gating interval is a basic limitation of all integrating A-D converters.

Figure 9 shows a typical scheme which delivers a 12-bit buffered output word and depends for its timing on an externally supplied clock. The small amount of logic required performs the reset-counter and load-latches functions automatically. No further control signals are necessary. Obviously, the clock can be generated by another AD537 if ratiometric measurements are to be made.

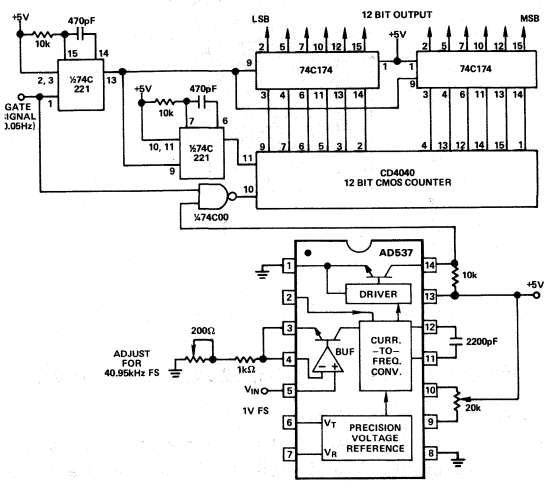

Figure 9. 12-Bit Analog-To-Digital Conversion

Accuracy

The absolute accuracy of an A-D converter using the AD537 will be much poorer than the linearity, since it is not fundamentally free from scaling drift. The scaling is a function of the internal reference voltage, the timing capacitor, the adjusted scaling resistor and the clock frequency, all of which are uncorrelated quantities subject to temperature- and supply-induced errors. An overall temperature stability of about ±100ppm/°C can be achieved in practice, comparable to that found in low-cost panel-meter DVM's.

Resolution

The resolution is determined primarily by the quantization imposed by the counting process, but there will also be a ±1 count uncertainty which can be troublesome for small inputs. This can be alleviated by adding several stages to the counter ahead of the first used stage, so that gating uncertainties are less apparent. Of course, to maintain the same scale, the V-F converter must operate at a higher frequency, and the trade-off in linearity may detract from the improvement in resolution.

Dynamic Range

The AD537 has a guaranteed dynamic range of 10,000:1, which is commensurate with a 13-bit quantization. In the example shown in Figure 9 the full-scale input of +1V generates a binary output of 111111111111 (decimal 4095) and the first bit occurs for an input of 244μV. To fully utilize this dynamic range it is necessary to null the input amplifier as described earlier under Basic Operating Modes.

Signal Isolation

V-F converters are useful for conveying signals between systems at greatly different potentials (for example, from the cathode of an electron microscope operating at several hundred kilovolts below ground) or for providing high-integrity isolation in equipment where ground currents would be hazardous (such as in medical instrumentation). Excellent special-purpose products are available for these requirements, usually in modular form, but in some cases the AD537 may offer an attractive alternative due to its low cost, low power consumption, high sensitivity and small size. Two basic points arise in these applications: first, the choice of the transmission medium; second, the provision of an isolated power supply.

Coupling Devices

The two most popular transmission mediums are magnetic and optical; a third choice, that of acoustical coupling, is applicable in some cases. Capacitive coupling is another possibility, although in many cases the large voltage between sending and receiving points, often accompanied by ripple and noise, excludes this approach. For illustrative purposes, a typical optically-isolated scheme is shown in Figure 10. To conserve power, the LED is driven at the relatively low level of 8mA, and one stage of gain is added on the receiving side. This circuit can operate at any frequency within the AD537 bandwidth and drive at least six standard TTL loads.

Isolated Power Supply

An ordinary 50/60Hz transformer-rectifier circuit can be

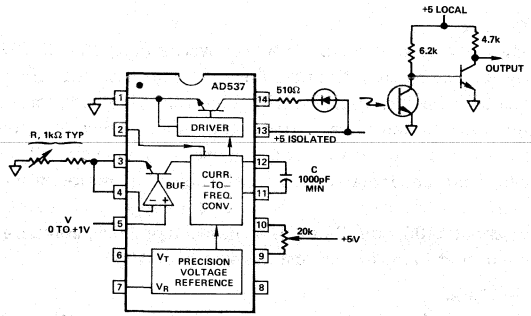

Figure 10. Typical Optical-Coupling Scheme

used to provide the floating voltage for the V-F converter. For high-voltage isolation, a toroidal transformer with an insulating bobbin should be used for the secondary. Alternatively, a dc-dc converter can be employed. However, in this case special care should be taken to ensure that all switching spikes from the dc-dc converter are completely filtered from the V-F circuit to prevent unwanted synchronization. Very little supply regulation is required; note that the current drain is nearly independent of frequency, simplifying the power-supply design.

Fiber-Optic Receiver

The output of a V-F converter is often transmitted over an optical-fiber cable for considerable distances, and detected by a photodiode at the receiving end. The photocurrents are usually small, and a preamplifier is necessary. Figure 11 shows a suitable circuit, having sufficient bandwidth to accept optical inputs up to 20kHz.

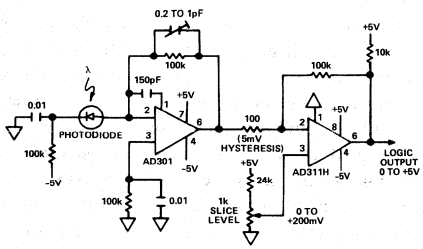

Figure 11. Photodiode Preamplifier

Handling Bipolar Inputs

The AD537 can handle signals of either polarity as has already been shown. Occasionally it is necessary to perform V-F conversion on a signal whose polarity is indeterminate. This can be achieved either by offset operation or by the use of an absolute-value circuit; both techniques will be described. Of course, the single-supply feature cannot be used in these circumstances. However, when the input is an alternating signal and it is only necessary to measure the average value, the AD537 can automatically perform precise half-wave rectification and operate from a single supply.

Offset Operation

By simply returning the timing resistor to a negative voltage (rather than ground), the V-F converter can accept negative and positive input signals, and generate a frequency

$$f = \frac{E_{IN} + E_O}{10V} \cdot \frac{1}{CR}$$

Clearly, the offset voltage E_O must have a magnitude greater than the most negative value of E_{IN}, and usually some overrange is desirable. Thus, a value of -12.5V might be used for E_O when the signal input has a full-scale range of ±10V. Figure 12a shows a simple implementation of this scheme, using the -15V supply as the offset source. To calibrate this circuit, first set the input to zero and adjust R1 for an output frequency of 10kHz. Then apply an input of +10V and adjust R2 for an output of 18kHz. This circuit has the advantage of retaining the high input resistance feature of the AD537. A disadvantage of this scheme is that the center-scale frequency is totally dependent on the -15V supply.

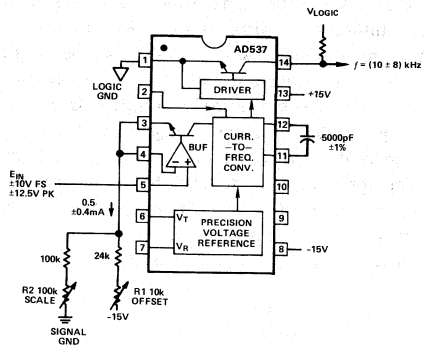

Figure 12a. Offset Operation, ±10V FS

A better arrangement is to make use of a stable reference element such as the AD589, as shown in Figure 12b. For the purpose of illustration a ±0.5V input range is shown. Calibration once again proceeds by first adjusting the zero-input frequency using R1 and then the +FS input using R2. Scaling for other operating conditions is straightforward.

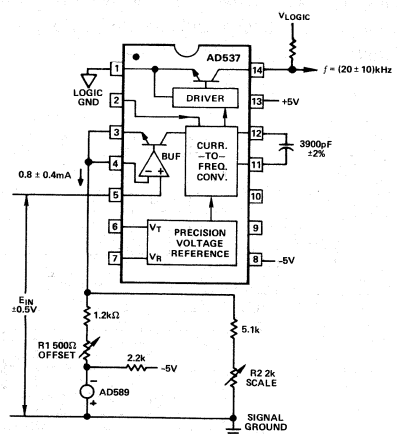

Figure 12b. Offset Operation, ±0.5V FS, Using AD589

Using Absolute-Value Circuit

In some cases the relatively large "zero-error" associated with offset operation may be troublesome. Also, the frequency data is inconveniently formatted if it is to be used for digital display. These problems are overcome by using an

ısolute-value (precision full-wave rectifier) circuit ahead the converter, with the bonus of effectively doubling the nge of linear operation. By utilizing the high input re- ıtance of the AD537 a very simple absolute-value circuit possible (Figure 13). The circuit shown is scaled for an put of ±10V full-scale, but any voltage can be handled ith appropriate change of scaling resistor. Similarly the ıll-scale frequency can be tailored to the application by ıanging the timing capacitor.

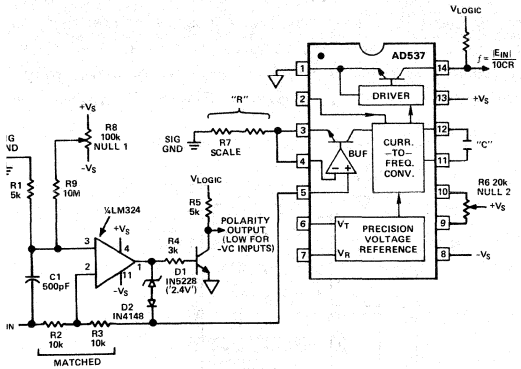

Figure 13. Bipolar Interface Using Absolute-Value Circuit

or positive inputs the op-amp is effectively disconnected ince its output goes low and cuts off D2, and the signal is onnected without attenuation or inversion to the input of he AD537 via 20kΩ. For negative inputs the op-amp ıutput becomes positive and turns on D2 via the low-voltage ener D1, whose purpose is to ensure proper operation of he polarity detector. The input is thus inverted, so once gain the AD537 receives a positive input. The output of 21 is low for negative inputs, and it can be used to indicate ignal polarity. (In a two-wire system such as described in ection II this information can be conveyed as a further nodulation of the supply current, for example.)

For operation from a single supply an op-amp type such ıs the LM324 can be used. A dynamic latching condition an occur with this amplifier if the input switches abruptly. This is prevented by the addition of R1 and C1. These com- ponents can be omitted if the input is slowly-varying or if a standard op-amp type is used. The speed of this absolute value circuit is such that the input to the AD537 is "oscil- oscope clean" for a 20V pk-pk sinewave or triwave input up to at least 2kHz. One caution to be noted is that the in- put resistance is unequal for positive and negative inputs, requiring either that the source resistance be very low (less than 10Ω for a 0.1% reversal error) or constant (in which case the error can be absorbed by subtracting this source resistance from R2. Full calibration requires the elimination of the op-amp offset by adjustment of R8. For most appli- cations this can be achieved by simply grounding the input and observing the polarity output. Better accuracy can be achieved using a DVM to null the input to the AD537. After this step, calibration proceeds as for unipolar operation, except for the need to check the reversal error due to R2/ R3 imbalance.

Measuring AC Signals

The AD537 only responds to inputs which cause pin 5 to be more positive than pin 4. For opposite-polarity inputs, no timing currents are generated. Thus, the device is inherently a half-wave rectifier with a very sharp "knee". This characteristic can be used to perform V-F conversion on an ac input signals without recourse to any further external parts. Filtering is automatically achieved by the integrating action of the timing capacitor for input frequencies above the instantaneous operating frequency. Of course, if the input has a period less than that of the oscillator frequency the output will be frequency-modulated, although in many cases further integration will be provided by the counters which often are used to measure the frequency. Examples of this mode of operation are given in Section V.

Transducer Interface

The AD537 was specifically designed to accept a broad range of input signals, particularly small voltage signals, which may be converted directly unlike many V-F converters which require signal pre-conditioning. The 1.00V stable reference output is also useful in interfacing situations, and the high input resistance allows non-loading interfacing from a source of varying resistance, such as the slider of a potentiometer. Some examples of transducer interfacing are provided by way of suggestion; many others will be obvious.

Thermocouple Input

The output of a Chromel-Constantan (Type E) thermocouple, using a reference junction at 0°C, varies from 0 to 53.14mV over the temperature range 0 to +700°C with a slope of 80.678μV/degree over most of its range and some nonlinearity over the range 0 to +200°C. For this example, we assume that it is desired to indicate temperature in Degrees Celsius using a counter/display with a 100ms gate width. Thus, the V-F converter must deliver an output of 7kHz for an input of 53.14mV. If very precise operation down to 0°C is imperative, some sort of linearizing is necessary (see, for example, Analog Devices, *Nonlinear Circuits Handbook*, pp92-97) but in many cases operation is only needed over part of the range.

The circuit shown in Figure 14 provides good accuracy from +300°C to +700°C. The extrapolation of the temperature-voltage curve back to 0°C shows that an offset of -3.34mV is required to fit the curve most exactly. This small amount of voltage can be introduced without an additional calibration step using the +1.00V output of the AD537. To adjust the scale the thermocouple should be raised to a known reference temperature near 500°C and the frequency adjusted to value using R1. The error should be within ±0.2% over the range 400-700°C.

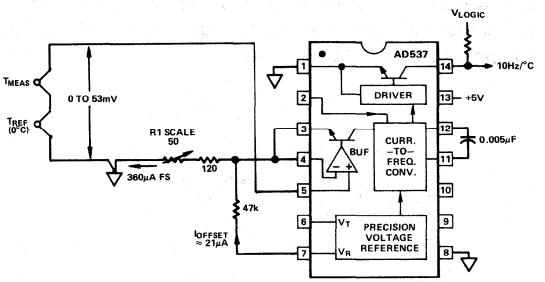

Figure 14. Thermocouple Interface with First-Order Linearization

Strain-Gauge Input

Depending on the application, the output of a strain-gauge consisting of a bridge of four resistors may be either unipolar or bipolar, linear or nonlinear. Thus, only an illustrative example can be given here. In all cases, the bridge requires excitation, and we here assume dc excitation using the 1.00V reference output buffered by an external op-amp (Figure 15). In this way, the calibration is not only independent of the supply voltage, but also very stable over temperature, since the same reference voltage (even though having a finite TC) is also used as the reference for the V-F section of the AD537.

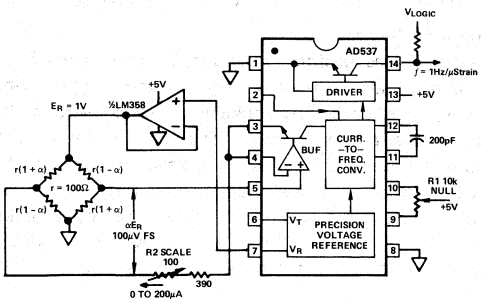

Figure 15. Typical Strain-Gauge Connections

This example assumes that the strain is always of the same polarity (unipolar operation) and that a balanced-force (linear) system is used. The timing resistor is returned to one side of the bridge, both to double the sensitivity and effect offset operation. This connection is capable of introducing a nonlinear error term, and we wish to determine its magnitude. Referring to Figure 16 it is evident that the current in the timing resistor (to which the frequency is proportional) is given by

$$I = \frac{\frac{1+\alpha}{2}E_R - \frac{1+\alpha}{2}E_R}{R + r\,\frac{(1-\alpha^2)}{2}}$$

This can be re-written

$$I = \frac{\alpha E_R}{R\left(1+\frac{r}{2R}\right)}\left\{1+\frac{r\alpha^2}{2R}\right\}$$

for typical small values of α, and this expression clearly reveals that a cubic error term results from this connection. In practice this error is tolerable. For example, for $R \geqslant 5r$ and $\alpha \leqslant 0.1$ the error is less than +0.1% absolute; if the system is calibrated at full-scale strain, this error is reduced to approximately +0.025% peak.

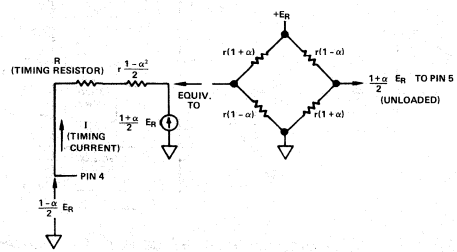

Figure 16. Circuit for Analysis of Nonlinearity

The circuit of Figure 15 is calibrated to generate a scale c 1Hz per microstrain (100kHz at the assumed FS value c $\alpha = 0.1$). Single supply operation allows this circuit to b located at the far end of a two-wire line, as previousl discussed.

Position Transducers

Linear displacement transducers, rotary servopots, man kinds of level transmitters, light-comparators using photc resistors, etc. are readily interfaced with the AD537. I many cases it will be possible to use the reference outpu to power the transducer for greater accuracy and stabilit of scale. For track resistances down to 3kΩ this may b connected directly, as shown in Figure 17; however, th

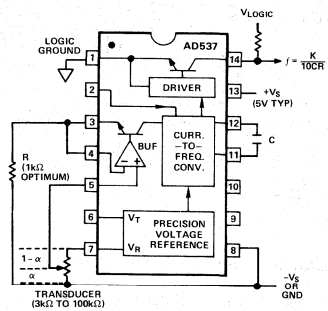

Figure 17. Resistive Transducer Interfacing

loading effects must be considered. First, the finite outpu resistance of the reference source (about 380Ω) causes th voltage across the track to be less than 1.00V, thus reducin the output frequency. Second, the track current flowing i pin 7 causes the voltage across the timing capacitor to b increased. This also decreases the frequency (by as muc as 50% for a resistance of 3kΩ). The effect is quite predic able, and a scaling correction factor, K, can be applied; thi is plotted in Figure 18. Alternately, a buffer amplifier ca

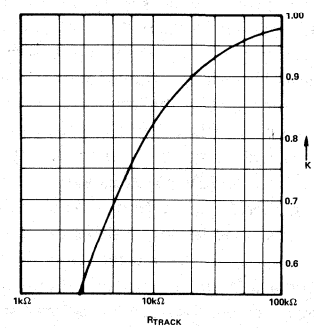

Figure 18. K-Factor Versus Track Resistance

be used, as for the strain-gauge driver shown in Figure 15 in which case no loading correction is necessary and mucl lower track resistances (down to about 100Ω) can be used.

Higher track resistances can also be used, up to at leas 100kΩ. Above this value the parabolic error caused by th input bias current (about 100nA) at pin 5 may becom troublesome. For example, at mid-position of a 100kΩ track, the error voltage is +2.5mV, or +0.25% of the 1V full

cale signal. This problem is mitigated by using a larger excitation voltage or by including a low-bias op-amp buffer. This general configuration may also be used as a manually-adjustable clock generator needing a minimum number of components, having a wide frequency range with very linear control, and capable of operating from a TTL or CMOS supply rail.

nverse-law operation may sometimes be desirable; that is, inear motion of the slider should result in linear control of the period rather than the frequency of the output. Figure 19 shows how this may be accomplished with good accuracy even when there is contact resistance to the slider. With the values shown, a range of 100:1 is provided, but the circuit is easily re-configured for other conditions.

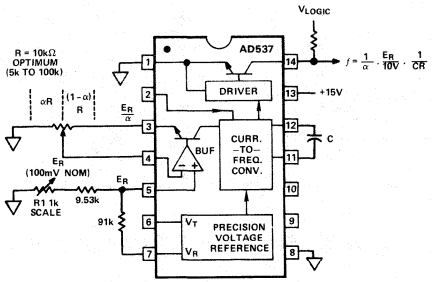

Figure 19. Linear Period Control

Photodiode Input

Some transducers, such as photodiodes and phototransistors, generate a current output and require an operating bias of a few hundred millivolts. Once again, the AD537 can provide the necessary interface conditions, as shown in Figure 20. The scaling accuracy will depend primarily on the stability of the phototransducer. It may be unnecessary to include provisions for scale adjustment; where justified, the scheme previously outlined for current-input operation is used as shown. Using a photodiode, the current will generally be small and direct interfacing may be impractical. A photoresistive (e.g. cadmium sulphide) cell may also be used.

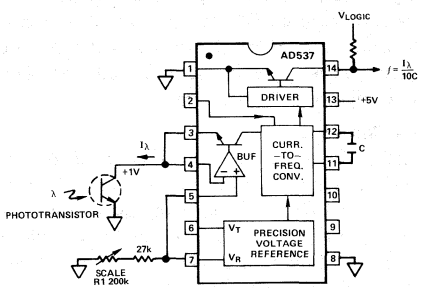

Figure 20. Phototransistor Interfacing

4–20mA Loop Operation

Various ways are available for converting instrumentation signals in the 4–20mA format to a frequency format. The choice will depend on whether the circuit is to be self-powered or powered directly from the current loop, and whether scale offset is to be introduced or not.

Externally-Powered

The method for converting the current-mode signal to a suitable voltage signal is trivial when no offset is needed and an external supply can be used. A somewhat more interesting problem is to provide offset in such a way that the output frequency is zero for a current input of 4mA. Figure 21 shows a suitable scheme, which makes use of the input amplifier sense terminal (pin 4) to introduce the offset.

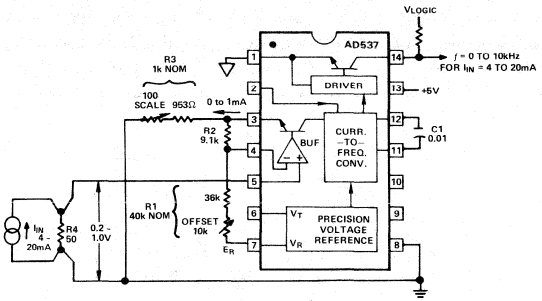

Figure 21. Offset Conversion of 4-20mA Signal

To understand the operation of this scheme, it should be noted that the offset voltage referred to the input is

$$E_O = E_R \frac{R2+R3}{R1+R2+R3}$$

The current in the timing resistor for values of E_{IN} above the offset is given by

$$I = \frac{1}{R3}\left\{E_{IN}\frac{R1+R2+R3}{R1} - E_R\frac{R2+R3}{R1}\right\}$$

Since the input is first converted to a voltage of 0.2 to 1.0V by the 50Ω load, an offset of 0.2V is required. For a nominal value of 1.0V for E_R this determines the ratio (R2 + R3)/(R1 + R2 + R3) = 0.2. The timing current, I, can be written

$$I = \frac{1}{R3}\cdot\frac{R1+R2+R3}{R1}(E_{IN} - E_O)$$

and for $E_{IN} - E_O = 0.8V$ at full-scale, we need

$$1mA = \frac{R1+R2+R3}{R1R3}\cdot 0.8V$$

for a 1mA full-scale current. This is only two equations, but there are three unknowns. A third requirement is that the loading on the reference output be light, 20μA for example. This determines the value of (R1 + R2 + R3) to be 50kΩ.

With these constraints, the resistor values can be calculated as shown. Note, however, that adjustment of offset and scale are unavoidably interactive, and the calibration procedure must be iterative.

Self-Powered Scheme

An interesting applications challenge is to find a way to use the signal current to also power the AD537. This is possible because the minimum value of 4mA is always greater than the quiescent current requirement of the IC. Even greater

utility results if the circuit can be optically coupled out, since the control loop then has complete isolation from the receiving system.

One approach is shown in Figure 22. Note first that there are two return paths for the signal current: through the main load resistor, R3, and through the timing resistor, R1 + R2. Since the current in the latter is exactly proportional to the load current in R3, linearity is preserved; even though the current gets used in a variety of ways by the rest of the circuit. The supply voltage is provided by the 5.6V zener diode, which together with the 1V full-scale voltage across R3 adds up to a worst-case compliance of 7V.

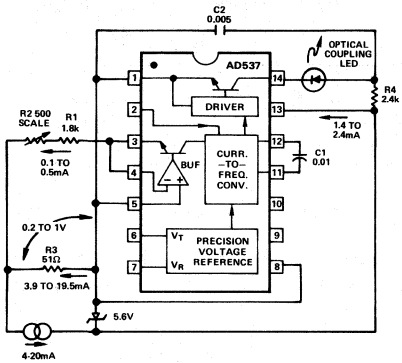

Figure 22. Isolated 4-20mA Converter

Each cycle the capacitor C2 is discharged through the LED at the peak current of approximately 25mA, internally limited in the AD537. The value of C2 is chosen to produce an optical pulse of about 1μs duration. The final diode current is limited to about 2mA by R4, calculated to ensure that even when the signal current is at its minimum value of 4mA there is at least 0.5mA in the zener; thus, the optical output persists at a lower level for a half-cycle. During the next half-cycle C2 recharges via R4, with a time-constant of 12μs, allowing full recharging in the minimum half-cycle time of 100μs.

In many cases the only calibration required will be to set the frequency to 5kHz at the full-scale input of 20mA. However, the V_{OS} of the AD537 may also be nulled for more exacting applications, using the standard technique described in Section I. Other scaling arrangements can readily be devised.

IV TELEMETRY AND COMMUNICATIONS

Frequency-Shift Modulators

A very common requirement is the transmission of binary data encoded as two discrete frequencies, for example, in data modems. The excellent temperature and supply stability of the AD537, its square-wave output and phase-continuous modulation property commend it as an alternative to semi-discrete implementations. Furthermore, its small power requirement allows it to be supplied from available telephone-line current at remote sites. Occasionally more than two levels can be encoded; these notes describe a binary modulator designed to Bell System 202 standards and an eight-level encoder.

Binary Encoder

Figure 23 is a schematic of an FSK encoder with a mark frequency of 1200Hz and a space frequency of 2200Hz. When the data input is low (mark condition) the timing current is supplied by R1. When the data input is high, Q1 saturates and the timing current is supplied by both R1 and R2. The inverted connection of Q1 ensures a very small offset error. Only one adjustment is needed; the voltage at pin 5 sets both mark and space frequencies. Note the very low current consumption of this circuit. The square-wave output must be filtered before transmission over a public telephone line.

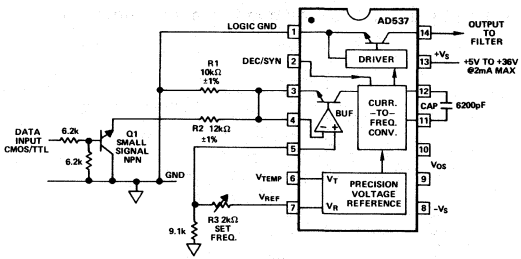

Figure 23. Bell System 202 Data Encoder

Eight-Level Encoder

A simple extension of the above circuit provides for the generation of eight linearly separated frequencies controlled by a 3-bit input word. Three inverted-mode transistors are used with collector loads of 1.8kΩ, 3.6kΩ and 7.2kΩ, which generate currents of 0.5mA, 0.25mA and 0.125mA respectively, when pin 5 is adjusted to 900mV. Using a 0.0125μF timing capacitor, this generates the frequencies zero through 7kHz. By adding a non-switched resistor (as in Figure 23) of 7.2kΩ, the frequencies 1 through 8kHz are generated. A more complex realization of digital frequency control is given under "Clock Generators" in Section V.

Phase-Locked Loops

The AD537 is useful as the VCO in phase-locked loops, due to its high linearity and the square-wave output. One application is as an accurate F-V converter; another is as an FSK demodulator. Both of these employ first-order loops. A linear phase-locked loop is also described, suitable for the recovery of signals buried in noise.

Basic F-V Converter

Figure 24 shows how the AD537 is connected for F-V conversion. The quad-nand open-collector TTL gate serves as a phase comparator, having as inputs the signal and the oscillator output. Its output is a rectangular pulse whose average value is forced to be equal to that voltage which causes the oscillator to run at the same frequency as the input. Since there is no low-pass filter, the circuit can lock on to any frequency from zero to full-scale (10kHz in this example), and this locked condition is acquired within four or five cycles. The dc output is extracted by a low-pass filter outside of the loop. As shown, the output has a FS value of +1V and is unbuffered.

The input duty cycle is not critical, but a square-wave is preferable. Pulses shorter than 50μs will cause errors.

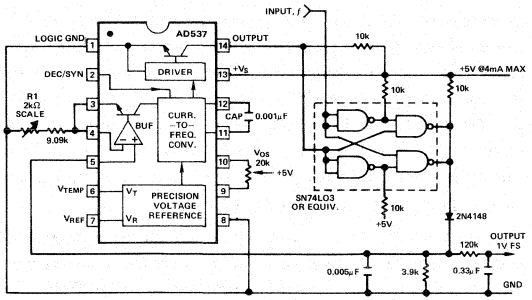

Figure 24. 10kHz F-V Converter

Similar trimming techniques as for V-F operation are employed. With the V_{OS} trimmer at mid-scale, trim the output to 1.00V for a frequency input of 10kHz using R1. Then apply a 10Hz input and trim V_{OS} for a 1mV output. Finally, re-trim R1 at 10kHz.

FSK Demodulator

For demodulation of a frequency-shift signal a narrower operating range is called for. Figure 25 shows a first-order loop similar to that just described, but operating only from 800Hz to 2600Hz. The center of this range is 1700Hz, which is also the center of the 1200/2200 Type 202 Dataset signal. The phase-comparator output is filtered by a simple low-pass circuit, suitable for data rates up to 120 baud. No trimming is required for this applications.

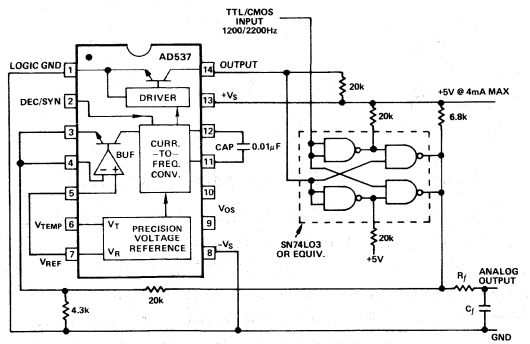

Figure 25. Bell System 202 Decoder

Linear PLL

The phase-locked-loop circuits described above operate from an essentially noise-free binary input. PLL's are also well-known for their value in extracting frequency information from a noisy analog signal. To do this, the digital phase-comparator must be replaced by a linear multiplier. In the implementation shown in Figure 26 the triangular waveform appearing across the timing capacitor is used as one of the multiplier inputs; the signal provides the other input. It is easy to prove that the mean value of the multiplier output is zero when the two signals are in quadrature. In this condition, the ripple in the error signal is also quite small. Thus, the voltage at pin 5 is essentially zero, and the frequency is determined primarily by the current in the timing resistor, controlled either manually or by a control voltage.

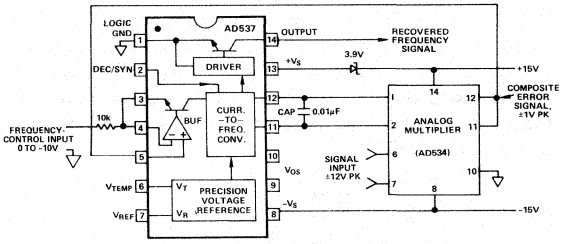

Figure 26. Linear Phase-Locked Loop

Noise on the input signal affects the loop operation only slightly; it appears as noise in the timing current, but this is averaged out by the timing capacitor. On the other hand, if the input frequency changes there is a net error voltage at pin 5 which acts to bring the oscillator back into quadrature. Thus, the output at pin 14 is a noise-free square-wave having exactly the same frequency as the input signal. The effectiveness of this circuit can be judged from Figure 27 which shows the response to an input of 1V rms 1kHz sinusoid plus 1V rms Gaussian noise. The positive supply to the AD537 is reduced by about 4V in order to keep the voltages at pins 11 and 12 within the common-mode range of the AD534.

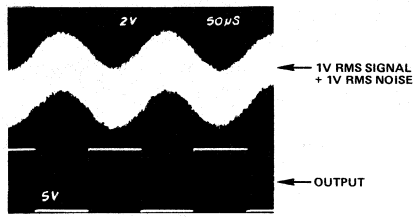

Figure 27. Performance of AD537 Linear Phase-Locked Loop

Since this is also a first-order loop the circuit possesses a very wide capture range. However, even better noise-integrating properties can be achieved by adding a filter between the multiplier output and the VCO input. Details of suitable filter characteristics can be found in the standard texts on the subject.

By connecting the multiplier output to the lower end of the timing resistor and moving the control input to pin 5, a high-resistance frequency-control input is made available. However, due to the reduced supply voltage, this input cannot exceed +6V.

Waveform Generation

The AD537 combines many useful features as a waveform generator, including a very wide (10,000:1) control range with a single capacitor, cycle times up to 100 seconds (C=1μF), external voltage or current control, square-wave or triwave outputs. The capability for asymmetrical waveforms is also available.

Triangle Wave

The circuit shown in Figure 28 generates a loadable triwave output whose mark-space ratio can be set by adjustment of the factor α. For one half-cycle (when pin 14 is open-circuit), the only path for timing current is through R1; for the other (when pin 14 is short-circuit), both R1 and R2 are in

the circuit and the half-cycle is thus shortened. The mark-space ratio is maintained over the frequency-control range. Of course, if a symmetrical waveform is needed, R2 is omitted, releasing pin 14 for use as a square-wave output. The mark-space relationship can be inverted by reversing pins 1 and 3 on the AD521.

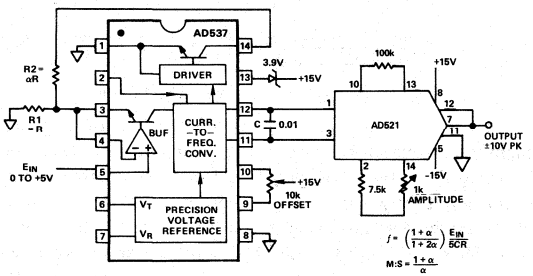

Figure 28. Asymmetric Triwave Generator

Note that the positive supply to the AD537 is reduced by the zener diode in order to remain within the common-mode range of the AD521 inputs. This also reduces the input signal range.

Voltage Programmable Pulse Generator

The circuit shown in Figure 29 produces a pulse train output with continuously variable output frequency, duty cycle, and output levels. The output waveform is shown below:

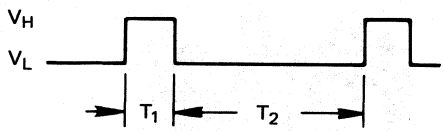

$$V_o = \begin{cases} V_H, \text{ for } T_1 = \dfrac{5RC}{V_1} \\ V_L, \text{ for } T_2 = \dfrac{5RC}{V_2} \end{cases} \qquad V_1, V_2 > 0$$

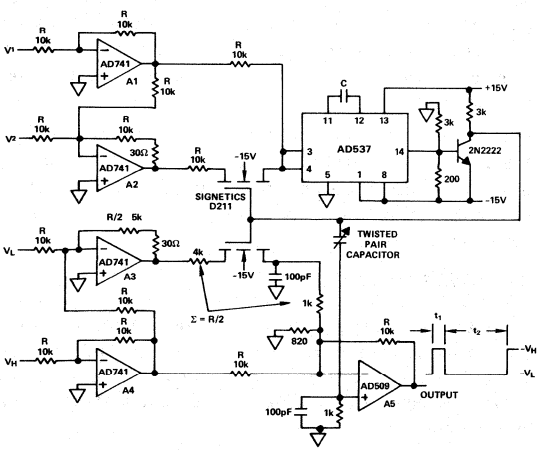

Figure 29. Voltage Programmable Pulse Generator

The circuit uses a 10kΩ timing resistor (R in the above formula), and the C value is chosen for the maximum square-wave running frequency. The minimum square-wave frequency is determined by circuit offsets. Frequency sweeps in excess of 1000 to 1 (corresponding to a 10mV to 10V range of V1 and V2) can easily be achieved. Additional range is possible by trimming offsets. The matching of the resistors around A1 and A2 will influence the accuracy of the duty cycle. If 0.01% resistors are used, a duty cycle sweep of 1000 to 1 with less than 10% error can be obtained. The 30Ω resistors shown in the feedback loops of A2 and A3 compensate for the R_{ON} of the FETs.

The output of the AD537 drives an external transistor which speeds up the transition and minimizes timing errors due to switching. The output amplifier, an AD509, is compensated slightly with the 100pF and 1k at the inputs. This results in a rise time of approximately 300ns with minimal overshoot and noise on the transitions.

Digitally Programmable Clock Generator

By combining a DAC with the AD537, digital control of frequency is achieved, using either an external binary word or manually-set thumbswitches. Figure 30 shows one realization using a 12-bit BCD-coded AD562 to provide 3-digit control of frequency from 100Hz to 99.9kHz.

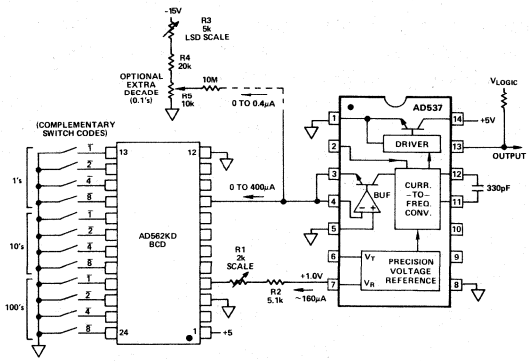

Figure 30. Digitally-Programmable Clock Source

The output current of the AD562 is used directly by the AD537, thus eliminating the need for a V_{OS} trim. Also, the reference output of the AD537 provides the AD562 reference input. This ratiometric operation contributes to overall stability. The DAC exhibits a current-gain of 2.5 at full-scale, so only 160μA of current is needed in the scale-setting resistors R1 and R2. If a straight-binary AD562 is substituted the FS output current is raised to 640μA, and the capacitor should be changed to 510pF for a nominal 100kHz FS.

Calibration is simple. With 999 on the thumbswitches adjust R1 for a frequency of 99.9kHz. The linearity of the AD537 ensures that all other frequencies are correct within the three-digit window. A fourth decade can be added in the LSD position as shown by the dotted lines.

Another digitally programmed clock generator is shown in Section VI of these notes.

V. SPECIAL APPLICATIONS

Long-Term Analog Integration

The integral of a unipolar input voltage is obtained simply by accumulating the output of an AD537 in a suitable counter:

$$N = \frac{1}{T}\int_{o}^{t} V(t)dt \quad , \quad T = 10CR$$

An analog output can be recovered using a DAC. Figure 31 shows a scheme using a 12-bit CMOS counter and a 10-bit CMOS DAC; the first two bits of the counter provide a prescaler, allowing the AD537 to run four times faster.

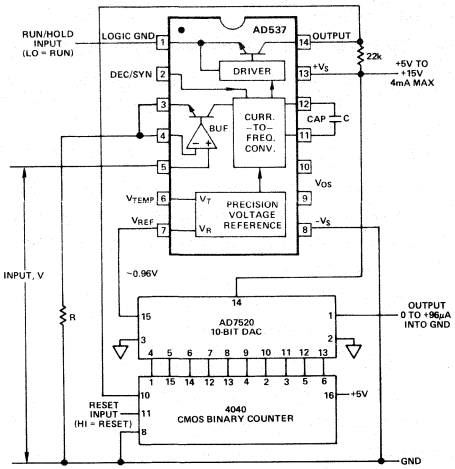

Figure 31. Long-Term Analog Integrator

Note that the 1.0V output of the AD537 is used as reference voltage, but the 10kΩ input resistance of the DAC loads this point, requiring a small correction in the output scaling, now nominally 96μA FS, and to the frequency adjustment which determines T. For a +1V analog input these values are recommended:

T	C	R	f(FS)
1 second	0.01μF	1.8K + 500Ω adj	4096Hz
1 minute	0.5μF	2.2K + 500Ω adj	68.27Hz
1 hour	10.0μF	6.8k + 1kΩ adj	1.138Hz

Longer integration constants should be obtained by interposing further prescaling counter stages ahead of the accumulator. Note that even when the AD537 is operating at low frequencies, wideband voltage fluctuations on the input are correctly integrated by the action of the timing capacitor.

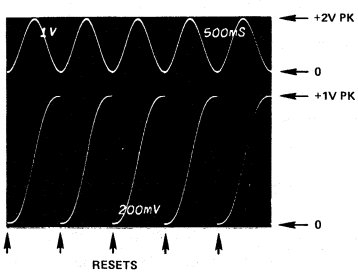

Figure 32. Output of Analog Integrator

Analog Division

In some cases it will be desirable to generate a frequency proportional to the ratio of two voltages. Of course, an analog divider may be used, but where the dynamic range of the denominator is fairly small (about 10:1) very good accuracy may be achieved using a few additional parts with the AD537. Figure 33 shows the scheme. First note that the numerator input voltage is applied to the usual input and scaling of full-scale input voltage and full-scale frequency is accomplished as in the standard modes. Positive inputs of up to (V_S - 4V) can still be accepted as well as negative voltages or currents. The division mode places no additional constraints on the numerator input.

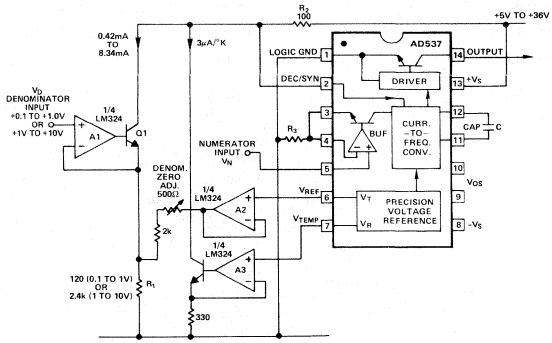

Figure 33. Analog Divider (Frequency Output)

Key to operation in this mode is the fact that the frequency is inversely proportional to the voltage between pin 2 and $+V_S$. Internally, this node has a resistance of 1.8kΩ to $+V_S$ and is supplied with a current of 460μA plus a further 50μA which is proportional to absolute temperature and is required to achieve temperature compensation of the basic current-to-frequency converter. By connecting an external 100Ω resistor, the control voltage at this node is reduced. This causes the nominal frequency to be raised by a factor of 19 times in the absence of any external current in this resistor. To restore temperature-independent operation, it is first necessary to supply 3μA/°K to this node. This is achieved using the thermometer output, the op-amp A3 (one-quarter of an LM324, chosen for its ability to operate from a single supply) and the transistor Q2 (any high-gain NPN). To provide conversion of the denominator signal to a voltage at pin 2 requires two steps: scaling and offset. Scaling for either +0.1V to +1.0V or +1.0V to +10V is demonstrated in the figure; the actual range extends considerably beyond the limits in both directions. Offset is necessary because there is a finite voltage at pin 2 even for zero denominator input, and this is provided by the stable +1.0V reference output. A2 buffers this voltage and A1 provides a high-impedance input point for V_D, the denominator signal. For R1 = 2.4K the output frequency is simply $V_N/V_D CR_3$. To adjust the denominator offset, connect the V_N and V_D inputs together (if they have same FS scaling voltage) and trim to maintain frequency independent of input voltage. Linearity of division is typically ±0.1%.

Temperature-to-Frequency Conversion

The 1mV/°K output of the AD537 is very exactly propor-

tional to absolute temperature, and the specified scaling error of ±0.1% can be absorbed in the overall scale adjustment. By operating at $+V_S$ = 5V and using small timing currents, the chip dissipation can be held to 6.5mW. This results in a temperature offset in free air of about +1°K, roughly equal to the typical input amplifier offset (since 1mV = 1°K). Accuracy in the absolute scale is therefore very high. Figure 34 shows the AD537S scaled for 10Hz/°K, corresponding to an output frequency range of 2180 to 4980Hz over the specified temperature range of -55°C to +125°C. For applications where ambient temperature is to be monitored, calibration to useful accuracy can be achieved by simply measuring the temperature at the package and adjusting the output accordingly. For more precise applications calibration at least two points (e.g., 0 and 100°C) will be necessary to eliminate offset effects.

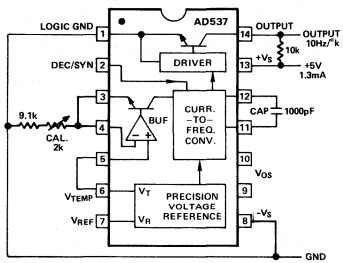

Figure 34. Absolute Temperature Transducer

The 1.00V output of the AD537 may be combined with the 1mV/°K output to realize other temperature scales. For the Celsius scale the lower end of the timing resistor must be offset by +273.15mV. Figure 35a shows how this can be achieved. The component values include corrections for the loading of the 1.00V output. The frequency range is reduced to 0 to 1250Hz, corresponding to 0 to +125°C.

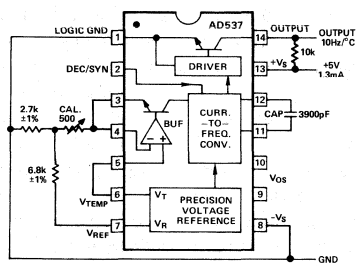

a. Celsius

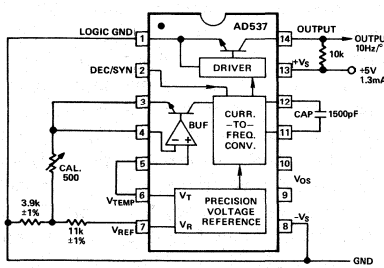

b. Fahrenheit

Figure 35. Connections For Offset Temperature Scales

The Fahrenheit scale requires an offset of +255.37mV, and suitable component values are shown in Figure 35b. The output frequency range is now 0 to 2570Hz for a temperature range of 0 to 257°F (-17.78°C to +125°C). As for the absolute scales, satisfactory calibration is possible by adjusting the frequency to correspond to one locally-measured test temperature.

It should be noted that only the AD537S will perform to +125°C. For applications requiring a maximum temperature of +70°C (+158°F), the J or K grade can be used. The circuit values remain unchanged.

Sound Velocity Monitor

The thermometer and reference outputs of the AD537 can simplify the computation of many physical parameters. As an example, the velocity of sound in air can be computed from the formula:

$$V_S = (331.5 + 0.6T_C)$$

$$= (167.6 + 0.6T_K)$$

where V_S is velocity in m/sec
T_C is Celsius Temperature
T_K is Kelvin Temperature

The circuit of Figure 36 uses two resistors to provide the appropriate weighing of the thermometer and reference outputs as given by:

$$\left(\frac{R_2}{R_1 + R_2}\right) 1V + \left(\frac{R_1}{R_1 + R_2}\right) \quad 0.001V/°K = 167.61 + 0.6T_K$$

Thus, $\frac{R_2}{R_1} = \frac{167.61}{600}$. If we choose a value of 20k for R_2, then R_1 = 71.5k.

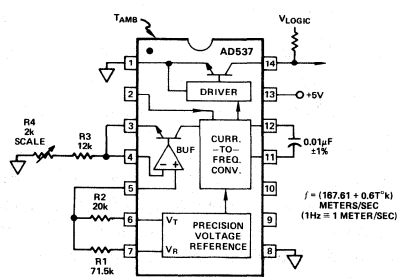

Figure 36. Sound Velocity Monitor

With these values, the voltage on pin 5 will be 452.8mV at 300°K. This must be scaled to an output frequency of 347.6Hz, corresponding to the velocity of sound at 300°K. If a 0.01µF timing capacitor is chosen, the value of R3 is found to be 13kΩ.

VI. MICROPROCESSOR INTERFACING

The AD537 is well suited to microprocessor interface applications due to its low cost, small physical size, and low power requirements. The device can be used as an output port, allowing software control of oscillator frequency, or it

can be used as the basis for a microprocessor-based data acquisition system.

Programmable Frequency Source

Figure 37 shows a circuit using an AD7522 10-bit CMOS digital-to-analog converter and an AD537 in a programmable oscillator circuit. The AD7522 is directly microprocessor-compatible and can be treated as either two output ports or two memory locations. Two locations are required to load 10 bits from 8-bit bus structures. The low order 8 bits are loaded into the DAC holding register on the rising edge of LBS, and the two most significant bits are loaded on the rising edge of HBS. The HBS signal is also connected to LDAC, which strobes the complete 10 bit word onto the actual DAC control lines coincident with loading of the two most significant bits.

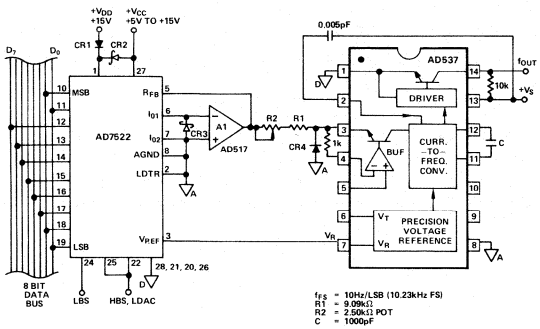

Figure 37. Digital Frequency Control from Microprocessor Bus

The reference for the AD7522 is provided by the 1.00V output of the AD537. The 10kΩ nominal impedance at the DAC reference input will lower the full-scale frequency of the AD537 due to loading on the 1.00 volt output (see Figure 18). Resistor R2 provides sufficient trim range to compensate for this effect. The DAC current output is converted to a voltage by the AD517, chosen for its low offset voltage and low drift. This voltage has a range of 0 to -1 volt F.S. and is applied to the AD537 configured for negative input voltage. As shown, the output frequency is scaled for 10Hz/LSB.

Data Acquisition

In microprocessor-based systems, a V-F converter is often used as a data-conversion element. The AD537 in particular can be used in systems requiring remote sensors reporting to a central processor. The serial outputs of many devices can be easily multiplexed, providing multiple-channel data acquisition at low cost.

The classic A-to-D conversion with a V-F converter involves using a series of counters to count pulses from the converter in a particular time interval. This gate signal must be, of necessity, quite stable, since any drift will cause a gain error in the conversion. Many microcomputers utilize a crystal oscillator for the system clock, and this can be used as the basis for a stable gate signal. Since most processor clocks run at a few megahertz, a divider stage will be necessary to provide a sufficiently long gate interval (usually a few hundred milliseconds). The gate frequency can, of course, also be derived from another AD537, in which case it can be adjusted to provide proper gate signals for several other AD537's operating at different full-scale frequencies.

If the parameter being measured is actually the denominator of the desired result, the role of the AD537 and the processor clock can be reversed. The AD537 output is used as the gate, and pulses of the system clock are counted. In this scheme, the total number of pulses counted bears an inverse relationship to the voltage at the AD537 input, relieving the microprocessor of the task of division. This type of conversion takes less time for a given level of resolution, since the processor clock runs much faster than the AD537 maximum frequency.

Figure 38 shows a generalized microprocessor data acquisition system using the AD537 as the basic analog-to-digital conversion element. The eight AD537's can be located some distance from the processor and linked by two-wire lines.

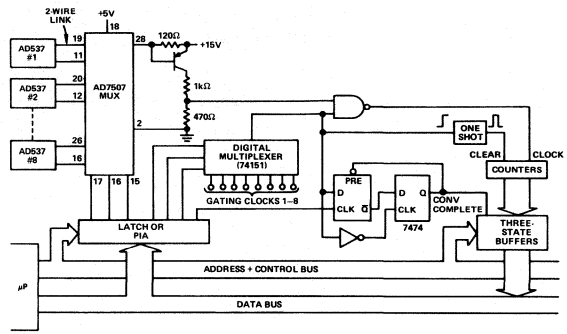

Figure 38. 8-Channel Microprocessor Data Acquisition System Using AD537

A single "write" instruction selects the AD537 to be monitored, the gating clock to be used, and issues the command to initiate a conversion. This byte might, for instance, use the following format:

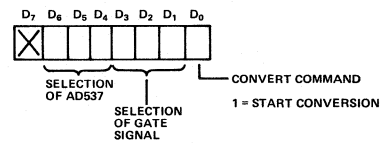

The AD537 to be monitored is selected by the AD7507 analog multiplexer. The level shifter stage converts the current modulation to TTL-compatible pulses as described in an earlier section. A +15 volt supply is suggested, to compensate for the voltage drop due to the ON resistance of the multiplexer. The gating clock is selected by the 74151 digital multiplexer. As shown, any of the gate signals can be used on any of the input AD537's. This permits selectable gain ranging or shortened conversion times for measurements not requiring full resolution. If only one clock is to be used for each V-F channel, then both multiplexers can share the same selection bits.

The rising edge of the gate signal triggers a one-shot of minimal pulse width to clear the counters at the beginning of the conversion. The two flip-flops provide a definite "conversion complete" signal, based on the state of the gate signal at the beginning of the conversion. The gate signal itself is a poor indicator of status, since a conversion may be initiated while the counters are running. The first conversion would then be inaccurate. However, the logic shown provides a clear indication of status. This flag can be read as a memory location, or it can signal an interrupt to the processor.

The latches, buffers, and bus structures have deliberately not been specified. The choice of these components is dependent upon the particular microprocessor used, system configuration, and so forth. The actual implementation is left to the designer, and it is hoped that these notes have provided some insight into the practical considerations of data acquisition with the AD537.

ANALOG DEVICES

AN-278
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Operation and Applications of the AD654 IC V-to-F Converter*

by Walt Jung

INTRODUCTION

DEVICE DESCRIPTION

The AD654 is a monolithic voltage-to-frequency (V/F) converter combining simplicity of use in standard applications with a high degree of flexibility and versatility. It is very much like its predecessor, the popular AD537.[1,2] The AD654 exchanges the all-around flexibility and versatility of the AD537 for maximum efficiency and economy of purpose, within a single 8-pin package. As a significant bonus, it adds an extended frequency range limit, contrasted to the original AD537.

Designed for either single or dual supply operation with low supply voltages and low current drain (5V and 1.5mA), the AD654 nevertheless has the capability of driving high voltage, high current loads (36V and 20mA). The chip includes a low drift input amplifier capable of operating directly from millivolt signals, a precision current controlled oscillator, and a high current output stage. It is a complete circuit, using low temperature coefficient silicon chromium thin-film resistors throughout. Only two application dependent external scaling/timing components are required; one resistor and one capacitor (Figure 1).

These two components provide the user with means for customizing the device for two application-unique parameters; scale of the input voltage and the output frequency. Programming of the full scale (FS) input voltage is possible from 100mV to 10V (or more, in some cases). The FS frequency can be programmed to any value less than 500kHz.

Either positive or negative input voltages can be accepted by the AD654 for a conversion. The basic AD654 V/F scaling relationship is:

$$f = \frac{V_{IN}}{10V \times (R_t C_t)} \qquad (1)$$

This simplifies the choice of the external components.

V/F linearity error is typically ±0.03% for 250kHz FS,

*Portions of this note are revised and expanded from the original "Applications of the AD537," by Barried Gilbert and Doug Grant.

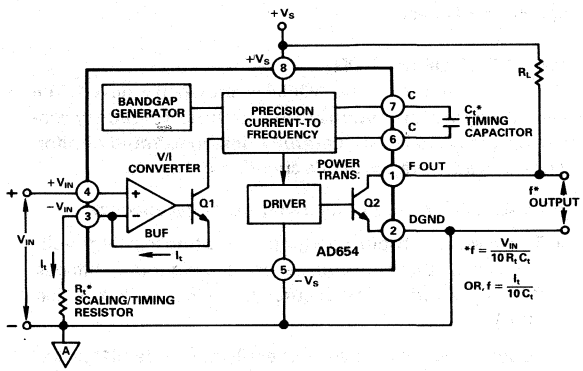

a. Block Diagram

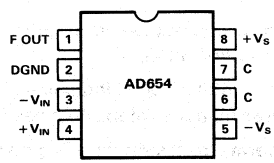

b. Schematic Symbol

Figure 1. AD654 Basics

guaranteed over an 80dB dynamic range. The output frequency is stable with temperature (typically ±50ppm/°C, excluding the effects of external components) and supply (typically ±0.05%/V from 5 to 36V). Complete specifications are included in the supplementary pages at the end of this note.

The input amplifier functions as a voltage-to-current converter and has a typical offset voltage drift of only ±4μV/°C. This low drift permits operation directly from such low level transducers as strain gauges, thermocouples, current shunts, etc, while offering a high (250MΩ) input resistance to positive voltage signals.

The output stage, an open collector NPN circuit, can sink up to 20mA with a saturation voltage less than 0.4V and can withstand a voltage of 36V. The collector and emitter

of this stage can be connected to any level between the limits of ground (or $-Vs$), and 4V below $+Vs$, permitting easy interfacing with any digital logic family. The high current capability means that LED's, long cables, or up to 12 TTL loads can be driven directly.

Unlike most V/F converters, the AD654 is designed to deliver a *square wave* output. This factor has advantages both internally and externally. Internally, it means that the power dissipation is essentially independent of frequency so that self-heating effects do not cause linearity errors. Externally, the average dc level of the output is constant, useful in ac coupled data links. It also allows direct operation as a phased-locked-loop within F-V circuits as well as other applications.

Whereas many V/F converters are limited in their range of applications to basic telemetry and slow A/D conversion, the low cost, low power drain, single-supply operation, and high operating speed of the AD654 open up a wide spectrum of uses.

THEORY OF OPERATION

The block diagram of the AD654 is shown in Figure 1. The key to accurate operation is the current-to-frequency converter, which is a very carefully designed multivibrator. The main advantages of this circuit are:

- Its simplicity—requiring a single timing capacitor,
- The use of push-pull charging—resulting in a large voltage across this capacitor, even at low supply voltages (for improved timing accuracy and low cycle-to-cycle jitter),
- Its square wave output—generally more useful than the narrow pulse generated by charge dispensing converters,
- Its high speed of operation—up to 500kHz, or higher with special operating modes,
- Its application flexibility, allowing cost-effective applications as well as high precision when needed,
- Its ability to operate from a single $+5V$ power supply,
- And, most importantly, its good linearity.

By using special adaptive biasing techniques, operation of this multivibrator is possible over a very large dynamic range, from a maximum control current of 2mA to less than 100nA (a frequency range of 20,000/1 or more).

It can be shown that the basic circuit has a well defined temperature coefficient of 300ppm/°K at all values of control current. Use is made of the tight thermal coupling to the associated band-gap reference generator, which supplies an exactly proportioned temperature compensation voltage to the multivibrator. The band-gap cell also provides the required bias for internal operation of the chip.

The square wave output from the multivibrator operates the output drive circuit, which provides a floating drive current to the large geometry output transistor Q2 (Figure 1). Internally, there are actually two transistors in a thermally symmetric layout for minimal interaction with the remaining circuits. This transistor is designed for very low saturation voltage and is driven to combine low ON voltages with a well defined current limit.

A versatile operational amplifier serves as the input stage. Its purpose is to convert the applied input voltage to a proportional control current in Q1. This current is optimally 1mA at the input signal level corresponding to the maximum output frequency. With the internal strapping of Q1's emitter and the op amp's (−) input, the input voltage is effectively impressed across the user defined external scaling resistor (R_t) by the theory of zero input differential.

This resistor, R_t, is chosen to provide the needed transconductance for the application, consistent with the applied FS input voltage. For example, for a FS input voltage of 1.0V, the optimum resistor value is 1.0kΩ. The $+V_{IN}$ terminal of the op amp (pin 4) offers a high input resistance to the source, with a bias current of about 30nA. The design of this op amp ensures that the effect of finite bias current on drift is small. Drifts of the order of 4μV/°C can be achieved, even with source resistance/scaling resistance imbalances of 1k. Consequently, the AD654 can accommodate millivolt signals without the need for a preamplifier.

The input configuration is quasi-differential, so that errors due to ground loops can be avoided with the proper choice of signal connections. The common-mode range of the AD654 extends from 4V below $+V_S$ (that is, from $+11V$ for a $+V_S$ of 15V), down to $-V_S$, so that inputs down to ground potential can be accepted even when operating from a single supply voltage. This feature maximizes the utility of the chip when operating from a battery or other single supply power source. Negative inputs (voltages or currents) can also be accepted by fixing the voltage on pin 4 (usually to ground potential) and driving a current into pin 3 (either directly from a current source, or via a scaling resistor).

Unlike the 14-pin AD537, the 8-pin packaged AD654 has no provision for trimming the offset voltage of the input op amp. However, since this error voltage is extremely low (1mV max), it is not likely to be a very significant factor in many applications. For example, even for an input FS voltage as low as 1V, this offset amounts to a zero error of only 0.1% of FS. For demanding applications, this residual offset can be trimmed externally at the option of the user.

BASIC OPERATIONAL MODES

Positive Input Voltages

Figure 2 shows the connections for basic operation of the AD654 as a V/F converter with positive voltage inputs. For such operation, a single supply voltage will usually be sufficient, with $-V_S$ (pin 5) and the output emitter (pin 2) strapped to their respective grounds. The lower end of the scaling resistor R_t should be connected to the appropriate signal ground, shown here as *analog ground*. Note that this point is a noise-free, high quality ground, denoted by the "A" in the triangle. Note also that this point is common to the (−) side of the input signal, V_{IN}. This is to minimize conversion errors due to the inevitable voltage drops on PC board traces. The digital current path returns are denoted by the digital ground symbol ("D" in the triangle) and should be maintained separate from analog ground, except at one point. This convention is generally used throughout this note.

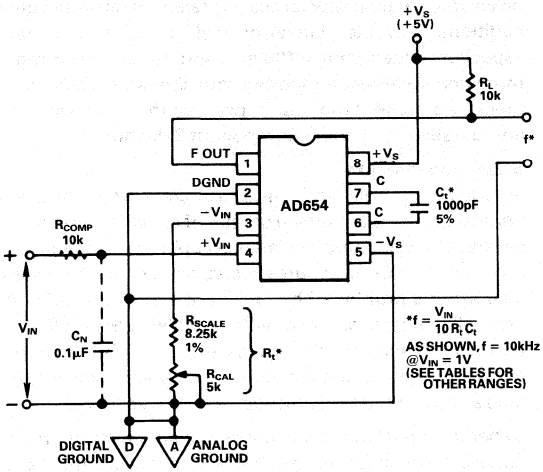

a. Low Power, 1V/10kHz

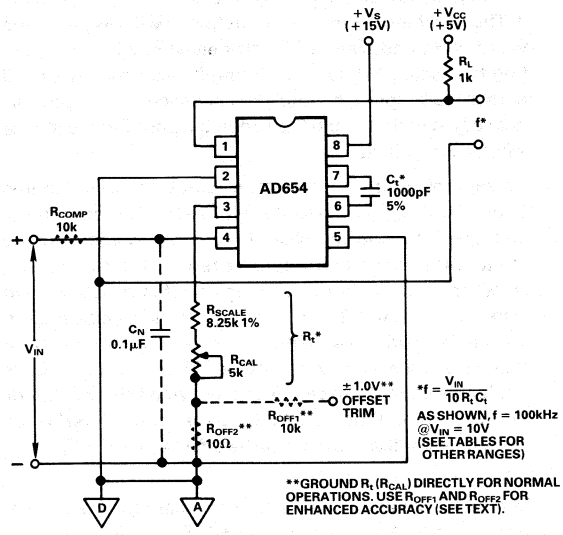

b. Wide Dynamic Range, 10V/100kHz

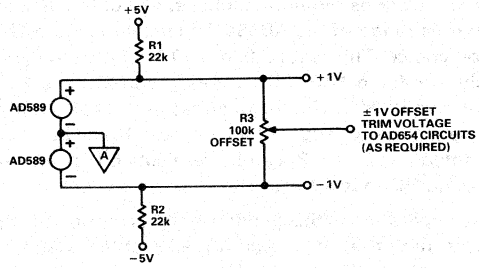

c. Offset Trim Bias Network

Figure 2. Basic Positive V_{IN} Connections

Choosing R_t and I_t

The nominal value of R_t is chosen for the application such that the FS input voltage sets up the desired FS current. There is a tremendous range of possible values here, due to the high dynamic range of the device. The designer can opt for a relatively low value of FS current I_t, and in so doing minimize the overall power consumption of the circuit. An example would be the use of a FS I_t of 100µA. Or operation can be towards a relatively high FS I_t of 1mA, whereby the maximum dynamic range and best linearity will be realized. It is the option of the designer to exploit this flexibility to the best advantage.

Since the input voltage V_{IN} appears across R_t, the relationship between R_t and I_t is simply:

$$R_t = \frac{V_{IN}}{I_t} \qquad (2)$$

Examples of standard operation are contained in Table I, but are discussed here to bring out the "why".

For a 0 to +1V input, where power is to be conserved, the circuit of Figure 2a is appropriate. This is a 10kHz FS circuit operating from a 5V supply. To minimize current drain, I_t is set at 100µA FS by the choice of R_t. R_t is nominally 10k for a 100µA FS current with a 1V V_{IN}. This is noted in the left column. It is a simple matter to recalculate R_t for other FS input voltages and/or currents not contained in the table. For example, for a max V_{IN} of 0.5V, R_t would be 5k for the same I_t. In the interest of simplicity, however, Table I summarizes the most useful operating ranges of the AD654 in terms of FS voltage, FS current, and the appropriate R_t value.

FS frequency is programmed via the choice of C_t, which is discussed latter. The circuit of Figure 2a is useful when overall simplicity is important and a three decade or lower dynamic range is adequate. The next example is slightly more complex, but realizes more dynamic range and higher speed.

FS V_{IN}	R_t FS I_t = 100µA	R_t FS I_t = 1mA
100V*	1 meg	100k
10V	100k	10k
1V	10k	1k
100mV	1k	100Ω

NOTE
*Applies *only* to Figure 3.

Table I. R_t Selection for V_{IN}; I_t

For applications requiring higher speed or the maximum dynamic range of the AD654 the circuit in Figure 2b is a good choice. This is for a 0 to +10V input, with R_t nominally 10k for a 1mA FS operating current; the FS frequency is 100kHz. This V/F operates from a 15V supply because of the 10V range. Other values of R_t can be selected as noted in Table I. R_t can be recalculated for any intermediate FS input voltage.

When operating in this general mode of input, it is important to note that the maximum allowable voltage input range of the AD654 is limited by the power supply voltage. For example, with $+V_S$ at 5V (as in Figure 2a) the maximum signal input is only +1V ($+V_S - 4V$). However, the device specifications do allow for the fact that the supply may be 10% low. In practice, reliable operation can actually be achieved for actual inputs slightly in excess of +1V with a 5V supply. By the same token, to accommodate an input voltage range of +10V, the $+V_S$ supply must be at least +14V ($+V_{IN} + 4V$), as in Figure 2b. Therefore, using a 15V supply a 10% input over range is possible.

Choosing C_t

Having chosen the input scaling resistor and I_t to fit the general requirements, the timing capacitor C_t is then selected. Note that C_t is chosen to accommodate the desired FS frequency, *after* R_t and I_t have been established. The basic relationship of (1) can be re-written in terms of C_t, as:

$$C_t = \frac{V_{IN}}{10R_t f} \quad (3)$$

Or, since V_{IN}/R_t is equal to I_t, this expression can be also be written in an alternate form, as:

$$C_t = \frac{I_t}{10f} \quad (4)$$

Equation 3 can be used to verify the values of the example(s). In the example of Figure 2a, with V_{IN} FS = 1V and R_t = 10k, the value of C_t will be 1000pF for a FS frequency of 10kHz. This can also be noted from the information of Table II (left column, bottom).

The scaling of a V/F is often specified in terms of conversion *sensitivity*, i.e., in Hz/V. In this case the scaling is 10Hz/mV. Table II simplifies the selection of C_t, allowing easy choice as is appropriate to FS frequencies of up to 500kHz. Higher frequencies are also possible, as noted. However, some special constraints are applicable, and are covered with the applications section.

For the 10V FS input voltage example of Figure 2b (again with R_t = 10k nominal, but I_t = 1mA), the value of C_t will be 1000pF for a FS frequency of 100kHz. This can also be noted from the information of Table II (right column). The scaling in this case is again 10Hz/mV.

Note that generally one must select the 1mA current ranges of I_t for the highest frequency ranges. This is to maintain optimum operating speed within the device. Additionally, when operating above 10kHz the value of the load resistor, R_L, should be lowered to minimize delay due to the associated time constant. A value of 1–2 kilohms is appropriate for the highest ranges, asshown in Figure 2b. At FS frequencies of 10kHz or below these measures are not neccessary and power can be conserved by operating at an I_t of 100μA, as in Figure 2a.

FS f	C_t FS I_t = 100μA	C_t FS I_t = 1mA
≥1MHz**	*	≤100pF**
500kHz	*	200pF
250kHz	*	390pF
100kHz	*	1000pF
10kHz	1000pF	10000pF

Notes
*Not recommended, see text.
**"Exalted" operation, see text.

Table II. C_t Selection for FS f; I_t

In general, the recommended C_t values are the most readily available ones. Within limits, any convenient value can be used to set up a special scaling relationship with some additional caveats. Linearity will be degraded with capacitor values below 1000pF, and 100pF is the minimum generally recommended capacitor size. There is no theoretical upper limit, but in practice there is a real one (see the section "Timing Component Selection").

Bias Compensation

In terms of optimizing the interface presented to the signal source, there are only a couple of considerations necessary. Since the input resistance at pin 4 is very high, errors due to non-zero source resistance will not normally be a serious problem. The low input bias current (30nA, typ) will generate an offset voltage of 30μV/k of *differential* source resistance. From this viewpoint, it is then desirable to keep the dc source resistances seen by the two inputs as near equal as is practical to minimize this error.

Generally speaking, bias current compensation resistance can be added to the inputs to minimize the resistance difference. However, with the AD654 it can only be added to the op amp's (+) input; that is, in series with pin 4. The resistance at the (−) input (pin 3) will simply be the value of R_t, and cannot be compensated without upsetting the scaling. Thus, minimizing bias current error will amount to keeping the total dc resistance seen at pin 4 as near equal to that at pin 3 by adding a padding resistance when appropriate.

An example of such a resistance can be seen in Figures 2a and 2b as R_{COMP}, the same value as R_t. The use of R_{COMP} is suggested for any applications where low end dynamic range is important (that is below a few mV), and in particular when R_t is 10k or more. Note that this resistance can also serve as the "R" section of an RC low-pass filter in conjunction with C_N. This filter is recommended to minimize noise disturbance of the V/F at low levels. C_N is not highly critical and can be chosen with this resistor for up to a 100ms time constant for maximum filtering. In general, a small value will also be effective (~0.1μF).

Calibration

In theory, two adjustments are necessary for the calibration of a V/F; FS (full scale) and ZERO scale (offset). In practice, the offset voltage of the AD654 is sufficiently low

ıat in many cases its adjustment is simply not necessary. the very best low end accuracy is required, the device an be offset trimmed as follows. Note that the dynamic ange of the lowest ranges (100-200mV) will be essentially etermined by offset voltage limitations, not current re- ated ones.

When offset voltage is to be trimmed, the trim will be lone *external* to the device using an optional connection, uch as that shown in Figure 2b. Resistors R_{OFF1} and R_{OFF2} ıre additional standard resistors placed to add a variable ıffset in series with R_t. A variable source of ±1.0V is ap- lied to R_{OFF1} to adjust the offset as much as ±1mV. A imple bipolar source of ±1.0V could be two AD589's onnected as shown in Figure 2c or a pair of LED's where he high stability of the AD589's is not necessary.

For most applications of the AD654 the major calibration rim will be for FS, which considerably simplifies its setup. n general, FS trim is the calibration of the circuit so that t produces the desired output frequency with a FS signal nput. Specifically, this is accomplished by adjustment of he scaling resistor R_t, in most cases. The precise calibra- ion of an AD654 requires the use of an accurate voltage standard set to the desired FS value and an accurate fre- quency meter, while a scope is useful for monitoring the output waveshape. Verification of converter linearity re- quires the use of a switchable voltage source or DAC hav- ing a linearity error below ±0.005%, and the use of long measurement intervals to minimize count uncertainties. Since each AD654 is factory tested for linearity, it will not usually be necessary for the end user to perform this tedi- ous and time consuming test on a routine basis.

In application circuits, sufficient FS calibration trim range must be provided to accommodate the worst case sum of all major scaling errors. This includes the AD654 full scale error (±10%), the tolerance of the fixed timing resistor (±1% assumed), and the tolerance on the timing capacitor (±5% is assumed). Therefore, in these discus- sions a trim range of ±16% will be used, since this will ac- commodate all of the above. In the circuit examples, the fixed part of R_t will be 82% of the nominal, and the variable portion chosen to allow 116% of nominal.

For instance, this would be an 8.25k ±1%, ≤100ppm/°C metal film resistor for R_{SCALE}, plus a 5k trimmer for R_{CAL}, for a net R_t of 10k. For the adjustable portion, this should be an infinite resolution film trimmer, with high mechani- cal stability and low TC (≤100ppm/°C), sealed against the environment. For other values of R_t, use the same general proportions.

Although device warmup drifts are small, it is always good practice to allow the circuit's operating environment to stabilize before trim, and to provide the pertinent sup- ply, source and load conditions. The calibration is begun with the offset trim (if applicable). If no provision is made for offset trim (or it is not neccessary to the application), proceed to FS calibration.

ZERO (Offset) Calibration

For ZERO scale cal, accomplished by offset trimming, begin by setting the input signal voltage to zero by short- ing the inputs. Adjust the offset pot until the AD654 just ceases to oscillate. This may be most easily seen using a scope at the output; the adjustment required should be slight. At this point, the device should already be close to calibration at low scale. Apply a known source of 1mV to the V/F input and adjust the offset trimmer further for the correct output according to Table III. For example, if the FS V_{IN} is 1V and the FS frequency is 10kHz then the offset trim frequency will be 10Hz (Table III, left column). Note also that this adjustment will be quite tedious unless the counter is used in the PERIOD mode, adjusting for 0.1S.

With offset properly adjusted, proceed with FS calibra- tion, as follows:

FS Calibration

To trim an AD654 for a desired FS frequency, apply a known FS input voltage at the appropriate level and adjust the R_{CAL} portion of R_t until the desired FS output fre- quency is indicated on the counter (Table III). In some ap- plications where the FS input voltage is small (≃100mV) this adjustment may slightly affect the device offset volt- age due to the presence of bias current at pin 3.

In a case where the offset is being trimmed and where the highest accuracy is desired, it may be appropriate to re- peat the ZERO and FS adjustments until no further im- provement is possible. If offset is not being trimmed, the FS adjustment completes the converter's calibration.

	FS Frequency		
FS V_{IN}	10kHz	100kHz	500kHz
10V	10V → 10kHz 1mV → 1Hz	10V → 100kHz 1mV → 10Hz	10V → 500kHz** 1mV → 50Hz*
1V	1V → 10kHz 1mV → 10Hz	1V → 100kHz 1mV → 100Hz	1V → 500kHz** 1mV → 500Hz*
100mV	100mV → 10kHz 1mV → 100Hz	100mV → 100kHz 1mV → 1kHz	100mV → 500kHz** 1mV → 5kHz*

Notes
*Adjust OFFSET (if used) as noted.
**Adjust FS cal as noted.

Table III. Calibration Tables for Various Ranges

Negative Input Voltage

By interchanging the two input connections the AD654 can operate from negative input voltages, as shown in Figure 3. In contrast to the connections of Figure 2, this cir- cuit does *not* offer a high input resistance since the entire control current must now be provided by the source. Con- sequently the input resistance is simply equal to the scal- ing resistance R_t which is again chosen to set up a desired I_t at the FS input voltage in use. In this example, a −10V V_{IN} is used with a 5V supply and the FS frequency is 10kHz. Note that other values of R_t can be selected from Table I.

In this case, the FS voltage can now be as large as desired since it is limited neither by the input amplifier common mode range, nor by the supply voltage. For example, a −100V FS input would use a scaling resistor of 100k, and the AD654 could still operate from a +5V supply in the cir- cuit as shown, with only this single change.

If it is desirable to reduce both power consumption as well as the loading on the signal, a lower FS I_t can be used. For example, a 1MΩ resistor could be used for R_t to convert

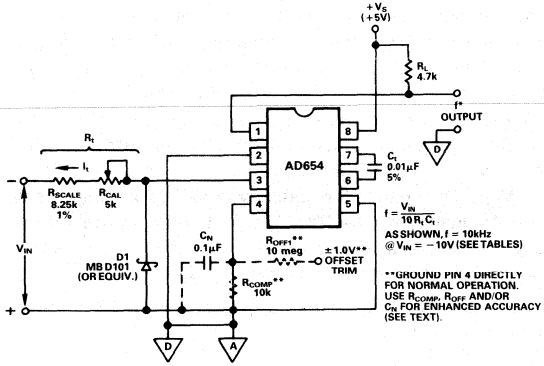

Figure 3. Negative V_{IN} Connection, −10V/10kHz

the −100V input into a 100μA FS I_t as noted in Table I. For the same FS frequency (such as 10kHz in this example) the timing capacitor would then be reduced by a factor of ten, and would be 1000pF (Table II). The same general considerations apply to this circuit with reference to speed as were applicable to the circuits of Figure 2.

A point worth noting is that when lower values of I_t are used, not only is power conserved, but a very large input over-range capacity is also provided (20/1 at I_t = 100μA). This can be useful in handling signals which may occasionally have large peak values.

When operating in this general mode, pin 4 will usually be at ground potential or the signal return level. This results in a current summing node at pin 3. Thus, any number of signals may be algebraically added before conversion and each source may have any desired scaling by appropriate choice of the associated resistor. Note however that it is not fundamentally necessary for pin 4 to be grounded; in fact offset scales can be generated by setting the voltage on this pin to some (stable) level.

For lowest offset voltage error, the optional compensation resistor (R_{COMP}) can be added at pin 4 in lieu of grounding the pin directly. This is not likely to be necessary except on the very lowest voltage scales or with high values of R_t (above 10kΩ). If offset is to be trimmed, one more resistor is added, R_{OFF}. This is selected to provide a ±1mV range at pin 4, with ±1.0V input. If a resistor is used for R_{COMP}, it may lead to noise coupling at pin 4. It should be bypassed for lowest noise, particularly in the higher values, with a capacitor such as C_N.

At any time when the inputs of the AD654 are operated from a negative source, it is imperative that the designer control the potentials at pins 3 and 4 with respect to −Vs (pin 5). *Neither input should be allowed to see voltages of more than −300mV with respect to pin 5.* If such a condition should occur, unpredictable behaviour of the V/F may result, and latchup may occur. With this in mind, the source to be used should be tested thoroughly for any possible over-range transients which could give rise to such input conditions. In the event that transient behaviour cannot be totally characterized as safe, a protective clamp diode can be used. This clamp, D1 in Figure 3, is a low leakage type with a low forward threshold and will provide adequate overload and latchup protection when located *after* R_t, as shown.

For applications where the best possible linearity performance is important, it is worth noting that the linearity in the negative input mode as shown here is inherently better than for the positive voltage mode circuits of Figure 2. This comes about since the degrading effect of finite common-mode rejection inherent to that connection is eliminated. It also allows the use of the minimum supply voltage.

The design considerations and trimming procedures for this circuit are otherwise the same as for the positive input voltage mode. Similar allowances for component tolerances should be provided and stable components used when requirements so dictate.

Negative Input Current

In some cases the input signal may be in the form of a negative *current source*. True current sources can be handled in a way somewhat similar to negative input voltages. However, the scaling resistor is no longer required, eliminating the capability of trimming the FS frequency in this fashion. In fact, trimming presents some special problems when the signal input is a true current source.

Operating the AD654 in a current input mode results in the slightly modified transfer expression:

$$f = \frac{I_t}{10C_t} \qquad (4)$$

As equation (4) states, the output frequency is only dependent on two variables, I_t (the input current to pin 3) and C_t. Since it will not always be practical to smoothly vary the capacitance for calibration purposes, an alternative scheme is needed.

A basic V/F circuit suitable for operation from a current source with a FS level of 1mA is shown in Figure 4. In operation, this circuit divides the input current into two main paths. One path is through the 100Ω resistor R1, and flowing into pin 3; it constitutes the signal current I_t to be converted. The second path, through another 100Ω resistor R2, carries the same nominal current. Two equal valued resistors offer the best overall stability, and should be either 1% discrete film units, or a pair from a common array.

As a result of the above, the 1mA input current is divided into two 500μA legs, one to ground, and one into pin 3 of the AD654. Since the total input signal current (I_s) is divided by a factor of 2 in this network, C_t must be reduced by a factor of 2. This results in a new transfer expression unique to this hookup:

$$f = \frac{I_s}{20C_t} \qquad (5)$$

As can be noted from the circuit values, a 5000pF capacitor will provide a FS frequency of 10kHz with a signal input (I_s) of 1mA.

For calibration purposes, resistors R3 and R4 are added to the network, allowing a ±16% trim of scale factor with the values shown. This will allow resistor-only trim for

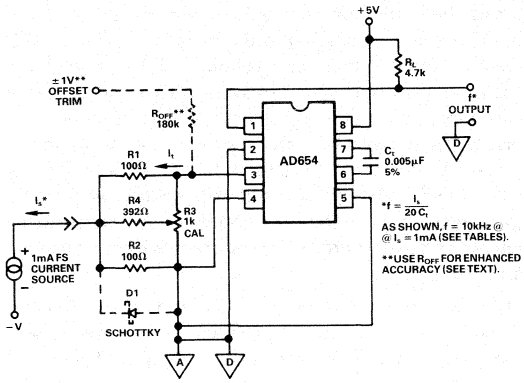

Figure 4. Current Source Input Connection

calibration of all the tolerances previously discussed. If the trimming range needs modification, R4's value can be changed. For example, a current output device such as a temperature transducer will typically have a specified calibration tolerance, and R4 can be lowered to accommodate this range.

For the greatest V/F dynamic range, some consideration to offset voltage should be given when using this circuit. Because of the current fork input network, a finite voltage offset in the AD654 can actually result in low frequency oscillation *without* a signal present! This will be understood when the paths for offset voltage induced parasitic currents are considered. Practically, the solution to minimizing this problem is to use an offset null externally. The optional input shown with the use of R_{OFF} will accommodate ±1mV of device offset.

Care must be taken when dealing with negative current inputs to this circuit for the same general reasons as outlined above under negative voltage inputs. Again, pins 3 and 4 must *not* be allowed to see more than −300mV with respect to pin 5. In this circuit, the low values of R1 and R2 prevent this for normal operation (plus some overrange). However, if transient currents can produce dangerous voltages, the use of a Schottky clamp diode is recommended.

The values of R1-R4 shown are suitable *only* for 1mA sources, but they can readily be scaled upward (proportionally) for lower FS current levels. For example, if one had a 100μA source signal these resistors would be increased by a factor of ten. At a given FS input current the R1-R2 junction should operate at about −50mV. The values of these resistors should never be set to allow more than 100mV drop under maximum input conditions.

Some General Precautions and Housekeeping

The AD654 is intended to be used with a minimum of additional hardware and the circuits shown in these notes are for the most part complete. However, the successful application of any IC involves a good understanding of all possible pitfalls, and the use of suitable precautions and preventions.

Input Protection

In general, pins 3 and 4 of the AD654 should never be driven more than 300mV below $-V_S$ (pin 5). If done, this could cause internal junctions to conduct, possibly damaging the IC or causing latchup. The AD654 can be protected from "below $-V_S$" inputs by suitable choice of input resistance and a clamping diode, as discussed in the basic applications section. The most important requirement of this clamp is that it must have a low forward threshold, ie., 300mV. A Shottky type is generally recommended in the figures. In some instances, where the higher leakage may be not be a factor, germanium types may also be useful, such as 1N67, 1N270, etc.

It is also desirable not to drive pins 3 or 4 *above* $+V_S$. In operation, the converter will become very nonlinear for voltage inputs above ($+V_S$ −4.0V). Input control currents above 2mA will also increase nonlinearity. By carefully designing for operation within the device's recommended limits, the user will avoid these "murky" areas and assure the possible highest performance.

Noise Filtering

The full rated 80dB dynamic range of the AD654 is for operation from a control current of 1mA (nominal FS) down to 100nA (this is equivalent to a 1mV input for 10V FS). For I_t below 100nA improper operation of the oscillator may result, causing a false indication of input amplitude. In many cases, this phenomenon might be due to short duration noise spikes which become added to the input. For example, when scaled to accept a FS input of 1V, −80dB is a voltage of only 100μV. For the same scale, a mean input of −60dB is 1mV and noise spikes of only 0.9mV are sufficient to cause momentary malfunction.

This particular problem can be minimized by using a simple low pass filter ahead of the converter. In fact, in *any* case where low scale operation is expected, filtering/bypassing around the inputs is recommended. For a FS of 10kHz, a single pole filter with a time constant of 100ms will be suitable. As noted in the examples, this resistor can also serve as a bias compensation resistance, enhancing accuracy there as well. Note that filtering within the signal path is also applicable to the negative input voltage mode, simply by splitting R_t into two parts with the midpoint bypassed by a capacitor to analog ground.

Of course, the optimum filtering configuration will depend on the application and the type of signal processing. Noise spikes are only likely to be a cause of error when the input current remains near its minimum value for long periods of time. Above 100nA (1mV) inputs, full integration of additive input noise occurs.

The AD654 is somewhat susceptible to interference from other external signals. The most sensitive nodes (besides the inputs) are the capacitor terminals. The timing capacitor should be located as close as possible to the AD654 package to minimize signal pickup in the leads. In some cases, guard rings or shielding may be required. Of these two precautions, guarding is the easier to implement as the package layout normally will guard pins 6-7 with pin 5 grounded and pin 8 to $+V_S$ (ac ground).

The latter technique (shielding) is not likely to be necessary except with the use of higher value film capacitors which can be quite large physically. It is usually desirable to minimize capacitor size, not only from a noise susceptibility point of view, but also from a cost standpoint.

Decoupling

It is generally good engineering practice to use bypass capacitors on the device supply voltage pins and to insert small valued resistors (10 to 100Ω) in the supply lines to provide a measure of decoupling between the various circuits in a system. For local bypasses at the IC terminals, ceramic, stacked film, or other comparably low-inductance capacitors in the 0.1μF to 1.0μF range are recommended in parallel with a low ESR electrolytic capacitor (0.1Ω or less at 100kHz). Note also that three terminal regulators are excellent spot decouplers and available in high performance-for-cost TO-92 packages.

Finally, it should be understood that proper attention to grounding at the analog and digital common points is a large part of successful operation of AD654 circuits. The digital and analog grounds should be connected at a single point to minimize the cross-coupling of currents and spurious error voltages. The floating output structure of the device goes a long way towards this but careful wiring should always be done, particularly where the accuracy required is high.

TIMING COMPONENT CONSIDERATIONS

Temperature Stability

The stability of the AD654's output frequency is determined by several factors. Broadly speaking, they can be grouped into offset drift and scale drift, or the TC of these parameters. Fortunately, the design of the AD654 is such that the amplifier's input offset drift is very small (in the 4μV/°C region). Unless the circuit is to be connected for operation from very low signal levels, this will rarely be a significant source of instability.

Temperature induced changes in the conversion scale factor include sources of more serious error. Scaling drifts arise both in the AD654 itself and in the external timing resistor and capacitor. Inasmuch as the AD654 commercial temperature range scale drift is typically ±50ppm/°C, the external timing components can in many cases contribute larger TC's. Both R_t and C_t should be selected for a quality consistent to the particular application. In practice, it is more likely to be the external R_t and C_t components which most seriously degrade stability over temperature, particularly the capacitors.

The best choice for a timing capacitor is dependent on a number of parallel performance considerations and of course, cost. As always, careful tradeoffs should be made to match the best performance parameters to the particular application. This is especially true for timing components used with the AD654.

Good V/F linearity definitely requires the use of a capacitor with low dielectric absorption (DA), while the most stable operation over temperature calls for a component having a small temperature coefficient. NPO ceramic, polypropylene, Teflon, and polystyrene capacitors as a rule best serve these two needs. These types should be the components of first choice. Mica and polycarbonate dielectrics are less desirable due to linearity degradation. Linearity and stability are generally degraded with most other types. Capacitor *tolerance* is not a major factor since it should be trimmed out as part of the calibration. 5% units are assumed for these applications.

Polystyrene types have very low DA and reasonably low TC (typically ≃ −120ppm/°C, this varies somewhat with processing), low cost, and reasonable size. The TC of polystyrene is also very *linear*, allowing compensation with a positive TC resistor. Polypropylene is similar, but with a wider range of values. Film types such as polypropylene are available in values up to 10μF or more. Teflon is superior in almost all regards (except cost), and will be unmatched for operation over 100°C. "NPO" or "COG" characteristic ceramic capacitors are the most desirable from the mutual standpoints of size, low TC (nominally zero), and a reasonably low DA. While the DA of this type is not as low as the better films (and it is not manufacturer controlled . . . a potential caveat), is still low enough to be very effective, particularly where the lowest TC is of paramount importance. Any other ceramic type should be avoided for timing uses.

For the scaling/timing resistor, the selection problem is much less severe. For most applications, metal film resistors are recommended, as they are available off the shelf in a wide range of values, and with TC's down to 50ppm/°C (or less).

In any event, when the lowest temperature errors are required of an AD654 circuit, care should be taken to minimize unnecessary heat sources, and keep all temperature sensitive components in close proximity to one another. For the most critical applications, RC components with nominally zero or the lowest available TC's can be used. This would suggest 25 or 50ppm/°C film resistances, and COG ceramic capacitors, which will minimize the timing component error. The ultimate in TC performance can be achieved only through rigorous attention to all temperature related factors. Application examples will illustrate this.

Other Capacitor Considerations

Since the voltage on the AD654's capacitor reverses polarity every half cycle, nonpolarized types must be employed. For use at very low FS frequencies, this becomes a problem in terms of both size and cost. This will generally rule out electrolytic types for precision applications.

As a practical note, if the maximum V/F dynamic range is not absolutely necessary to a given application, the capacitor size can be minimized by operating with a reduced FS current (rather than 1mA). For example, a 10/1 increase in R_t will reduce this current and allow a 10/1 reduction in C_t for the same FS frequency. Tables I and II will be helpful in making this type of tradeoff.

Output Interfacing

The AD654's logic output stage, available at pins 1-2, is designed to interface easily with all digital logic families. This includes LED's (for optical coupling), pulse transformers, and long lines. Due to the internal floating drive scheme, both the collector (1) and the emitter (2) of the stage's NPN transistor are available for external discretionary uses. The emitter can be tied to any potential from ($+V_S$–4V) down to $-V_S$, and the collector can be pulled up to 36V above the emitter (regardless of the $+V_S$ level applied to pin 8).

Logic Driving

Figure 5a shows the AD654 connected for a 0 to +10V voltage input with general output interfacing. The required logic common voltage (V_{EE}), the logic supply voltage (V_{CC}), the pull up resistor (R_L), and the $-V_S$ supply for the AD654 are shown in the accompanying table. The user can design an appropriate interface by selecting the proper set of conditions from this table. In the TTL mode, up to 12 standard gates (20mA) can be driven at a maximum LOW state voltage of 0.4V. This output flexibility is available with all the other possible input connections not explicitly shown.

The AD654 output stage is current limited, and the collector node (pin 1) can be shorted indefinitely to pin 2 (normally logic ground). When this point is shorted to $+V_S$, the current in the ON state (if not limited by any external means) will be approximately 35mA. Under most conditions this will not result in damage.

As an example, with pin 2 grounded, pin 1 shorted to +15V, and the oscillator running, the average power in the output stage is 262.5mW (15V × 35/2 mA; at a 5V supply it will be 1/3 this power). If however, a high supply voltage is used on the logic stage and the output is in the ON state for long periods (i.e., at a low frequency/long period), the peak dissipation of 1260mW (36 × 35) will cause considerable heating. It is recommended that this situation be avoided for prolonged periods of time, particularly with the plastic package device.

While the AD654 has ample drive for lower frequency TTL loads, at frequencies above a few hundred kHz it can begin to lose square wave symmetry and logic saturation. For operation above 200kHz into TTL or other logic type loads a buffer is recommended, as shown in Figure 5b.

Phantom Power Schemes

It is very often desirable to power a remotely located V/F circuit via a simple cable hookup while taking the output signal from the same cable. Whenever possible, it is also useful to use the most simple cabling, such as a plain #22 twisted pair. Concurrent power/signal transmission on the same cable implies suitable separation circuitry so that no loss of performance results. Two circuits suitable for phantom power of AD654 V/F converters are shown in Figure 6.

Figure 6a is a simple but effective hookup useful for interfacing with a variety of logic types at the output. The positive supply line is fed to the remote V/F cable through R1.

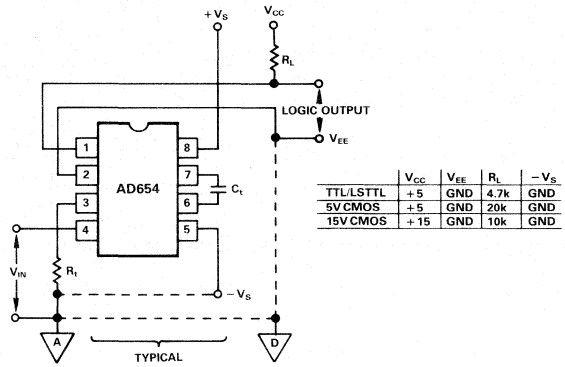

	V_{CC}	V_{EE}	R_L	$-V_S$
TTL/LSTTL	+5	GND	4.7k	GND
5V CMOS	+5	GND	20k	GND
15V CMOS	+15	GND	10k	GND

a. Standard Interfacing

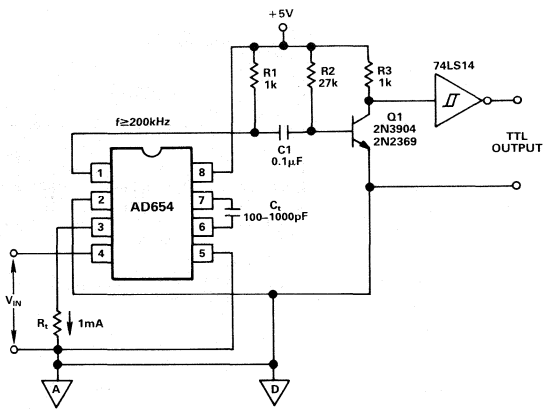

b. High-Speed TTL Output Buffer

Figure 5. Interfacing with Standard Logic

This resistor is selected such that the quiescent current of the AD654 will cause less than one V_{BE} to be dropped. As the V/F oscillates, additional switched current is drawn through R_L at pin 1. The peak level of this additional current causes Q1 to saturate, and thus regenerates the AD654's output square wave at the collector. In the process of supplying power to the AD654, the supply voltage is reduced one V_{BE} plus the line drop.

To set up the receiver circuit for a given voltage the R_S and R_L resistances are selected as appropriate from the table. CMOS logic stages can be driven directly from the collector of Q1 for any supply level. A single TTL load can be driven from the junction of R_S and R2.

At the V/F end it should be noted that the reduction of the supply voltage to the V/F also reduces the input range, by about 0.7V. This will reduce the input range to about 0.3V (max) for a 5V supply. The quiescent power of the V/F should be minimized by operating it with an I_t of 100μA and careful controlling any additional power loading. Because of these power supply restrictions, this type of circuit is best suited for those types of sources or transducers where little or no additional power support is needed.

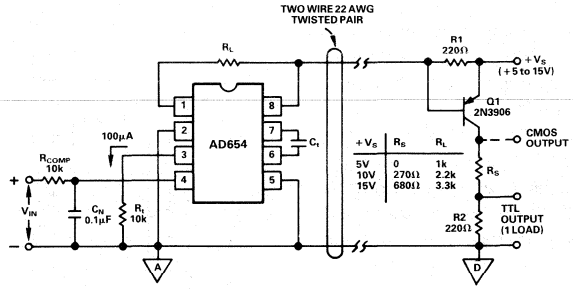

a. Simple Phantom Power Driver

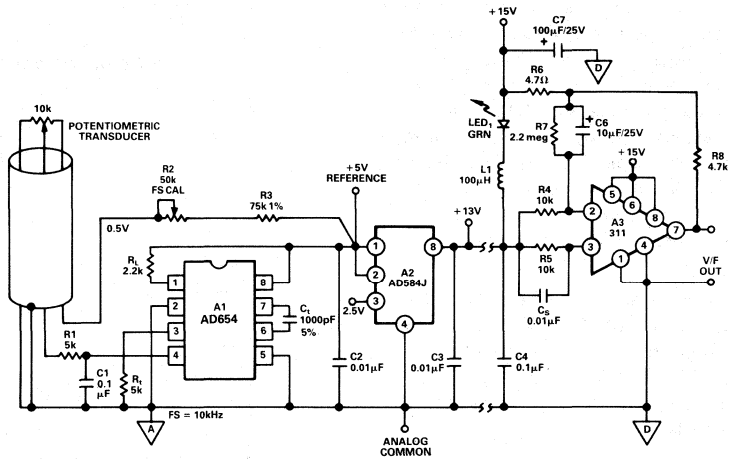

b. Phantom Power Driver/Regulator

Figure 6. Phantom Power Schemes

Figure 6b solves the problem of additional power, and adds considerable flexiblity to the application at both the transmitter and receiver locations.

By inserting a regulator IC (A2) between the V/F output and the receiver input, two things are achieved. First, the supply voltage to the V/F and the surrounding circuitry is buffered and regulated to a reference quality level. This allows additional precision circuitry support to be added with relative freedom. Second, the V/F square wave output delivered to the receiver (in current form) is also regulated providing independence from power supply changes and noise pickup in the cable.

In the receiver circuit several functions are performed. DC power is coupled to the transmitter via LED_1 (green), and low resistance choke L_1. The choke allows slowly varying or static DC curents to pass without developing any sizeable terminal voltage.

On the other hand, the faster edges corresponding to the V/F square wave develop a pulsed waveform across L_1, which is passed to the comparator A3. This stage re-constitutes the original square wave at the output with 15V levels.

The circuit as shown is suitable for operation up to 10kHz, and will operate at least three decades lower. In this example, the 5V reference output is divided down to 0.5V and applied to a 10k linear potentiometric transducer. The variable output of this pot is the V/F input, and FS calibration is done by trimming the terminal voltage via R2.

Load resistor R_L sets the output current level of the V/F to a nominal 2.5mA. If the supply voltage of the V/F should be changed, the resistor should be adjusted to maintain this same current. This point suggests the versatility of the circuit, as the AD584 reference IC can be programmed for either 5V or 10V. It could also be set for 2.5V, but this will not be useful to the AD654 (nevertheless, this tap is still available at pin 3, and can be used for bias or multiple reference uses). For any given output voltage programmed into A2, about 3V of headroom should be maintained for proper regulation.

A very wide range of other circuits can be driven from A2's output, particularly if a buffer output option is used. The output will be 5V ±15mV, at a TC of 30ppm/°C or better using the AD584J. Of course the same general principle of three terminal regulation can be employed with one of the more utilitarian regulator IC's if a precise output is not required. No additional load can be allowed to pulse load A2. If this were to be permitted, the step change in output current would confuse the receiver and generate erroneous output data.

References

There are often instances where a stable reference voltage is necessary for bias, offset, or the establishment of scaling in a V/F application. Figure 7 illustrates several techniques which are useful.

Figure 7a is a simple 1.235V reference using the AD589, a bandgap-based, low power reference diode. The low terminal voltage allows use from supplies down to 5V, therefore this type of circuit is applicable to virtually any AD654 use.

Figure 7b shows how appreciably higher source currents can be provided from a given reference source. It is applicable to either a basic 1.2V source (7a), or an optional 5V source (such as 6b). In either case, the output is 1V at a low impedance level, so several mA can be sourced for driving bridge transducers, servo pots, etc.

Negative bias voltage from a positive supply is a problem that can be handled with a simple dc/dc converter when the overall load power is low. An example is shown in Figure 7c.

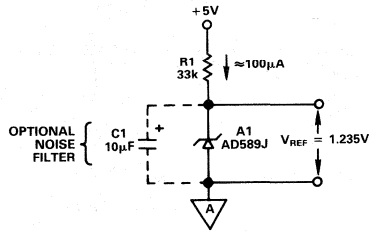

a. Basic Positive Bias Reference Source

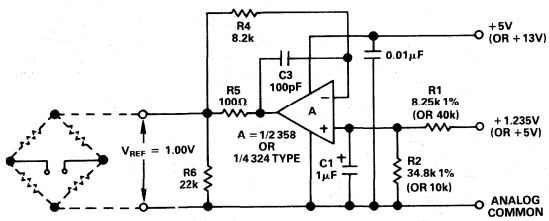

b. Buffered Low Power Reference

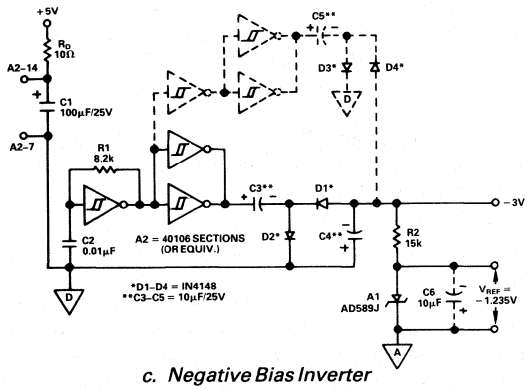

c. Negative Bias Inverter

Figure 7. Reference Circuitry

This circuit is an 8kHz inverter using CMOS inverters and an output p-p detector. With a 5V supply, it will develop a −3V output capable of 1-2mA drain. Regulation as shown is relatively poor, but can be optionally improved by using spare inverter sections to provide full wave rectification (shown dotted). If a negative reference potential is required a 1.2V reference diode can be used, as shown.

An input decoupling filter for the inverter is shown (R_D-C1), which will minimize the potential for noise to be injected back into the +5V supply.

Signal Multiplexing

When a variety of transducers are required to interface with a common data acquisition system, one solution is to use separate instrumentation amplifiers for each channel, followed by an analog multiplexer which drives a (single) V/F converter. However, in many cases a more economic solution may be realized by using a dedicated AD654 per channel and multiplexing these outputs digitally. Such a circuit is shown in Figure 8.

The open collector feature of the AD654 makes this easy to implement. In this application, all the V/F circuits operate continuously and can have a variety of input scales. For a measurement, only the V/F device having its output emitter pin grounded (2) through the "n" decoder (or other digital switching element) actually transmits an output. Note that the common V/F supply voltage must be high enough to accommodate the highest individual positive input voltage. V_{CC} can be any appropriate logic level, and can be different from $+V_S$, if necessary.

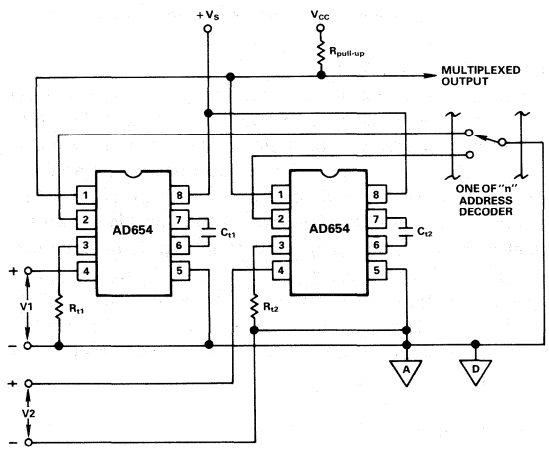

Figure 8. Signal Multiplexing

Transformer Coupling

Coupling a V/F output signal between two isolated logic grounds via a pulse transformer can be both economical as well as useful for frequencies up to about 10kHz. An example of a 1V/10kHz transformer coupled V/F is shown in Figure 9.

In this circuit the AD654 is operated at an I_t of 100μA in order to minimize power consumption, a factor important for many isolated power hookups. T1 is a miniature,

printed circuit mount telecommunications type 1:1 isolation transformer operated in a pulsed mode.

The V/F produces an output square wave, which is shaped into negative going pulses by the low power CMOS logic for each transition. The pulses are easily transmitted across the transformer's isolation barrier, even at the extremely low frequencies which are well outside the transformer's normal bandwidth.

At the secondary side, bias is applied to one end of the transformer and coupled to a CMOS Schmitt type inverter. The hysteresis of this stage provides noise immunity as well as the necessary latching function to transform the alternate (+) and (−) pulses back into a square wave.

The supply voltage shown for the AD654 is +9V and is an example of a low power circuit compatible with a simple isolated dc/dc converter or a low cost battery. Other supply voltages, consistent with the CMOS logic, can be used if the drive to T1 is maintained at a high level.

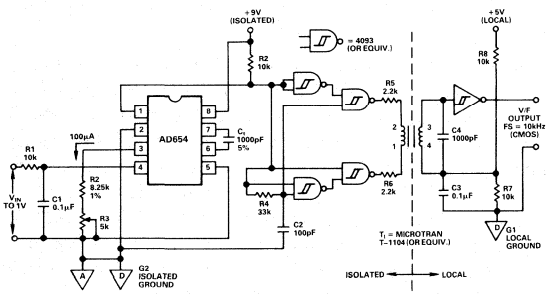

Figure 9. Transformer Signal Isolation

Optoisolator Coupling

A very popular method of isolated signal coupling is via optoelectronic isolators, commonly called optoisolators or optocouplers. In this type of device, the signal is coupled from an input LED to an output photo-transistor or photo-diode, with light as the connecting medium. This technique allows DC to be transmitted, and is extremely useful in overcoming ground loop problems between equipment. It is also applicable over a wide range of speeds and power. In fact, a very large percentage of V/F applications can potentially be useful with opto or other types of isolated coupling, so a number of useful circuit examples are given.

General Purpose, Medium Speed

General purpose versions of optocouplers abound, using infrared LED's for input and photo-transistors for output. This type is very cost-effective for general purpose work and can be useful up to about 100-200kHz with proper interfaces. Higher speed devices are also available for applications up to as high as 1MHz or more. The tradeoff between the two classes is generally cost, which is relatively high for the higher speed units.

Figure 10a shows a general purpose isolated V/F circuit using a 4N37 opto-isolator. A +5V power supply is assumed for both the isolated (+5V ISOLATED) and local

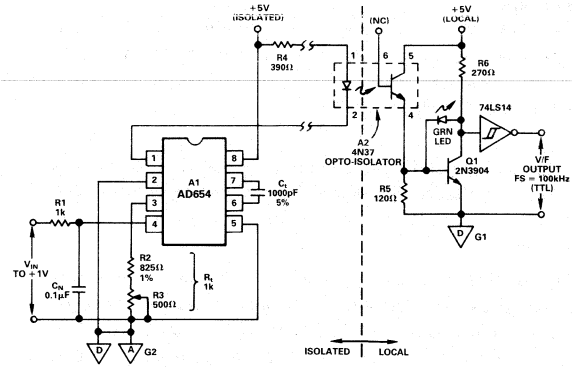

a. Medium Speed, General Purpose (100kHz)

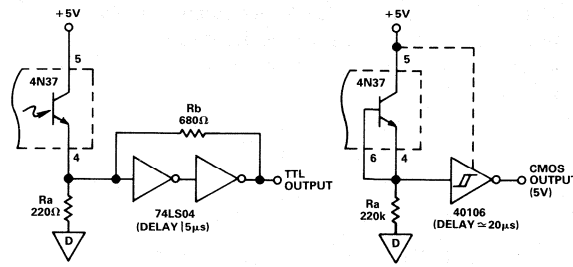

a1. TTL Output Buffer

a2. CMOS Output Buffer (Diode Mode)

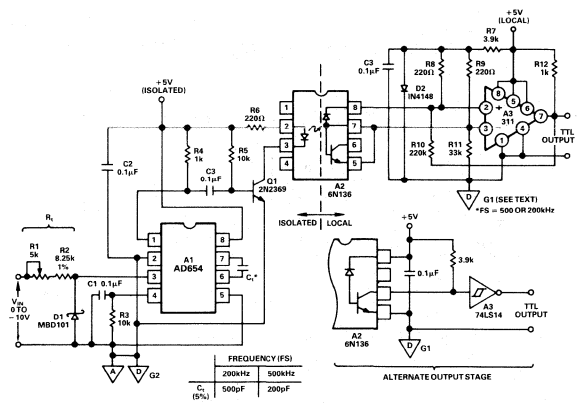

b. High-Speed Isolated V/F

Figure 10. Opto-Coupler Isolation

(+5V LOCAL) supplies. The input LED of the isolator is driven from the collector output of the AD654, with a 9mA current level established by R4 for high speed as well as for a 100% current transfer ratio (CTR).

At the receiver side, the output transistor is operated in the photo-transistor mode; that is, with the base lead (pin 6) open. This allows the highest possible output current. However, for reasonable speed in this mode it is imperative that the load impedance be as low as possible. A single transistor stage current-to-voltage converter does this, providing a dynamic load impedance of less than

10Ω and interfaces with TTL at the output. Q1 is this stage and achieves 2μs (typical) delay times for both the ON and OFF transitions. This allows the stage to be useful to as high as 100kHz. The circuit values around the V/F are for 100kHz, but other input configurations can be used.

In the circuit of Figure 10a, the values around the 4N37 are optimized for speed at the expense of power. If simplicity or lower power operation is desired, alternate output side buffer-interface connections can be used. For a simple (albeit slower) TTL buffer the hookup of Figure 10a1 can be used. If low power is key, Figure 10a2 is suitable. Both stages provide TTL or CMOS 5V output levels, and should use Schmitt type amplifiers for the fastest output transitions and the nonambiguous switching of downstream logic stages.

High Speed

For isolated V/F uses appreciably above 100kHz, an optocoupler specified for high speed is the best choice. This can be achieved several ways, including user-designed high-speed photo-diode type isolators with high-speed output buffers, or complete integral-buffer high-speed TTL output optoisolators such as the 6N137 types. Of these two, the former is illustrated below.

Figure 10b shows how an isolated photo-diode V/F output coupler suitable for a FS frequency of up to 1MHz can be implemented. In this circuit, a high-speed, separate photo-diode/transistor optoisolator is used (the 6N136). For this particular use, the 6N136's internal transistor buffer is not used, but the high-speed diode between pins 7-8 is applied to take advantage of the low inherent delays.

The V/F portion to the left is a −10V, 200 or 500kHz FS circuit using the AD654 from an isolated +5V supply with a 16mA output drive. Other input arrangements can also be used, however the input mode used here allows the maximum V/F linearity to be achieved. The AD654's I_t should be set at no lower than 1mA for best speed and the Q1 output driver stage for the LED is recommended. The V/F can be easily set up for either 200 or 500kHz FS operation as shown in the table. The best overall linearity will be attained at a 200kHz FS and it can be lower than 0.1%.

At the receiver, resistors R8-R9 act as a differential, low impedance load resistance for the photo-diode which is operated in a photo-voltaic mode. A3, a standard 311 type comparator, regenerates the logic state of the diode, producing 5V TTL levels at its output. The stage has ON and OFF delays of about 200ns (or less). Thus, it is useful at speeds up to 1MHz or higher, depending upon the comparator used. The 311 is the upper speed limit to this coupler. However, even the general purpose 311 is capable of tracking the AD654 speed at 500kHz and its single 5V supply operation is attractive for reasons of simplicity.

Note that the 6N136 coupler also allows the more conventional use by applying the photo-diode output current to the built-in NPN transistor as a saturated logic stage. This mode, which is an alternate option for this circuit, is shown in detail in the inset. It is attractive for intermediate speeds, delays on the order of 1μs or less, since it does not require the comparator stage used here. Note that a Schmitt type TTL buffer is still recommended.

Single-Supply Bipolar V/F

In many applications using the AD654, a +5V power supply may be the only power supply available. This makes it a challenge to process bipolar signal inputs without generating additional supply voltage(s). Figure 11 illustrates a novel circuit which can in fact processes inputs of ±10V, with a +5V supply on the AD654.

The bipolar input capability is achieved by A2-a, a precision rectifier stage which scales a 10V input to +1V across R3. A2-b detects the SIGN of the input and is HIGH for positive inputs. The 1V scaled output from A2a is always positive and is fed to the AD654 (+) input, with optional noise filtering by C3.

The V/F is scaled for a 10kHz FS frequency with a 10V input. If desired, compensation for rollover error (gain difference between equal magnitude positive and negative inputs) can be added by trimming R3 for a positive gain equal to the negative gain. R12 sets overall scale factor (adjusted for 10kHz with −10V input).

The SIGN output is handled in a novel manner, allowing both frequency and SIGN to be transmitted on the same serial cable. The output from the AD654 is steered by gates A3-b and A3-c, which are controlled by the sign bit. For a given sign, one of the one shots will be triggered and combined by gate A3-d. The resulting output pulse train follows the input frequency, but with a pulse *width* encoded for the sign. A width of 38μs corresponds to a negative input sign, while 65μs indicates a positive input.

At the decoder, the logic simply clocks a D flop 50μs after the leading edge, decoding a one for a positive sign, and lighting the indicator.

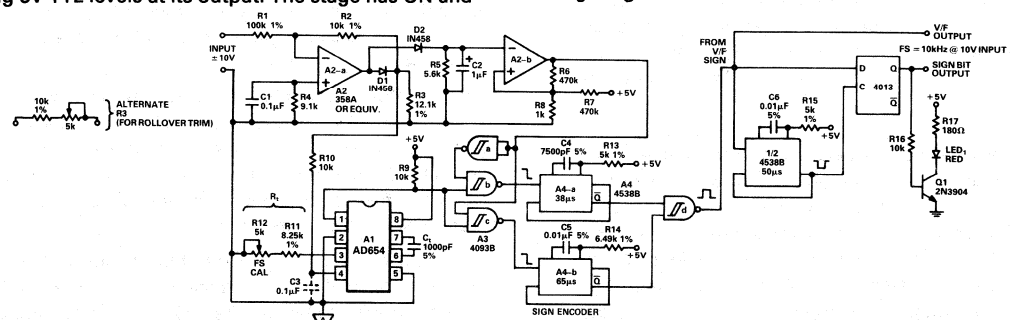

Figure 11. Single-Supply Bipolar V/F with Encoded Sign

Offset Scaled V/F

Figure 12 illustrates a method of offset scale operation with a bipolar input signal. By the very nature of this application bipolar supplies are assumed; in this case they are ±5V.

This circuit is designed for a basic range of ±0.5V, as seen at the AD654's (+) input. This is raised to ±1V by the input divider, allowing trim of span or alternate input scaling (shown here as a 10V option).

The basic frequency for zero input is defined by the quiescent offset current, which is derived from a negative reference of 1.235V, A2. Since the timing resistance will drop 1.235V for a zero voltage input, 1.235V becomes the effective V_{IN} for the calculation of center frequency. C_t is chosen in a conventional manner. Note that for best stability an NPO type should be used. While the linearity will not be as low as that of polystyrene, little difference will be seen in this application, since the oscillator always sees a relatively narrow frequency range.

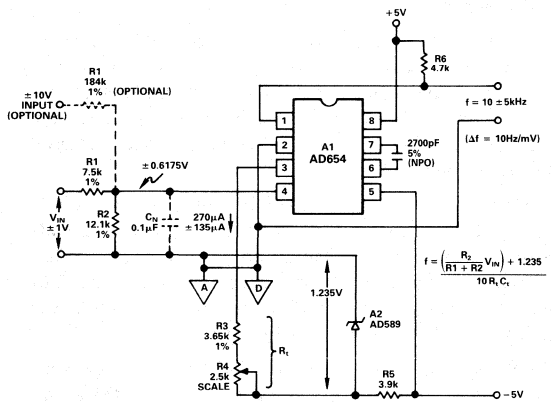

Figure 12. Offset Operation, ± 1V (or ± 10V) FS

In this example, 10kHz is the center frequency and a ±1V input swings the frequency ±5kHz (15kHz to 5kHz). The frequency swing sensitivity is set by the R1/R2 ratio which scales the AD654 input to ±0.6175V. This in turn changes the current in R_t ±135μA for ±5kHz.

An offset scaled V/F such as this is a convenient method of transmitting ac or bipolar data without the complication of sign bits and zero crossing detectors. The data can be recovered at the receiver site via an F/V converter and ac coupling (see Figure 20), or via an FM discriminator system.

Thermocouple Source Temperature-To-Frequency

Figure 13 illustrates another method of offset scale operation, one very much related to the more general scheme of Figure 12. In this case, the signal source to the V/F is the AD594 (AD595) thermocouple signal conditioner, allowing measurement of temperatures with either a type J or type K thermocouple. In this application bipolar supplies are again assumed to be ±5V.

This circuit is designed for a basic range of ±1.0V, as seen at the AD654's (+) input. This signal is produced by A2, which is either an AD594 (type J) or an AD595 (type K), and corresponds to a signal output of ±1V for a basic ±100°C span. The AD594(5) part of the circuit consists only of A2 and the thermocouple ground return resistor, R3. All the amplification and offset necessary is within A2, which operates from ±5V supplies. The A2 output signal is fed directly to the AD654 through a low pass filter, R1-C1.

The basic frequency for a zero °C input is defined by the AD654 quiescent timing offset current, which is derived from a negative reference of 1.235V, A3. The timing resistance will drop 1.235V for a zero voltage V/F input and this part of the circuit is otherwise similar to Figure 12. With the values shown for R2 and C_t, the output frequency becomes 4574Hz, at 0°C.

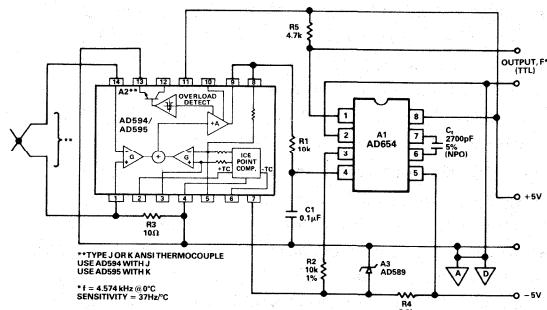

Figure 13. Thermocouple Source Offset V/F

As can be noted, this circuit has no calibration provision. It is intended to be used with an F/V converter in order to trim the output voltages for zero and full scale temperatures. Used with the AD650 or a similar F/V and followed by a differential amp to shift zero scale to zero volts dc, trim can be easily accomplished. The F/V's full scale calibration is used for the high-temperature calibration point and the differential amplifier offset for zero (at 0°C, or freezing).

Resistive Transducer Interfaces

With the wide dynamic range of control inherent to the AD654, it functions very well as an A/D conversion device to read out linear resistive transducers, such as servopots, etc. A basic scheme to accomplish this is shown in Figure 14.

This circuit simply applies a stable dc reference voltage on the order of +1V across the pot. With a stable excitation on the pot, the output voltage will represent the product of the fractional rotation α and V_{REF}. Since V_{REF} is unity, the output frequency is, therefore, directly proportional to α, as shown by the expression.

The major V/F conversion error of this circuit will likely be that of the AD654's basic nonlinearity, which can be 0.1% (or better). With a 50k pot, the worst case bias current induced error will be seen at mid scale, but this still will be

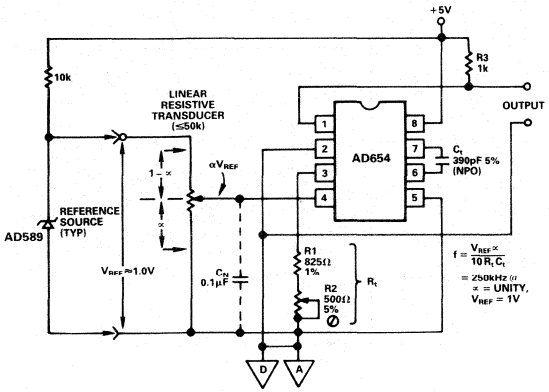

Figure 14. Resistive Transducer Interfacing

typically under 1mV. With the values as shown, the full scale frequency is 250kHz when α is unity.

To implement this type of circuit, virtually any stable 1V reference source can be used but the output impedance should be low. This is important, to minimize the possible variations of the reference source with variable pot loading. A buffered reference (such as Figure 7b) is therefore suggested. Reference voltages *higher* than 1V can be used if the AD654 is powered from a higher supply voltage. For frequencies above 200kHz, output buffering may be used (Figure 5b).

Photosensitive Transducer V/F's

Photosensitive sources such as phototransistors and photodiodes are easily interfaced to the AD654. An example of a phototransistor source is shown in Figure 15. This figure illustrates how a phototransistor can be biased at a defined collector-emitter voltage by the AD654, and the resulting current output converted to a linearly proportional frequency. For a sensor photocurrent of Iλ, this current becomes equivalent to I_t in the frequency expression, as shown.

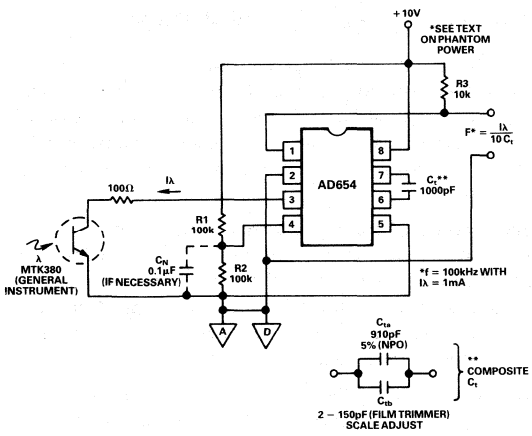

Figure 15. Phototransistor Photosensitive Sources

Since the AD654's internal op amp will bias pin 3 to the potential applied to pin 4, the voltage divider R1-R2 and the supply voltage will determine the effective C-E potential. Here this voltage is 5V, but it could be any voltage within the common-mode limits of the AD654. With a general purpose sensor such as the MTK380, operation is specified at a 1mA current with a 5V bias. An important point for this circuit is that it can easily accommodate any sensor. The divider can be bypassed for noise if necessary and, if maximum bias voltage stability is required, R2 can be made a zener or stable reference diode.

For an output photocurrent of 1mA, this circuit as shown will produce a frequency of 100kHz with a C_t of 1000pF. Calibration can be a problem since not all sensors and/or sources will be controlled as to scaling. The variable C_t shown in the inset will allow for about a 10% scale variance but even this may not be enough. If the sensor and V/F are remote, hardware calibration may not be feasible at all. For either case, scale calibration for a specific light input can be accomplished via software. Note that the circuit can be phantom powered easily, since the transducer is biased completely by the V/F.

Photodiode Sources

Interfacing a photodiode to the AD654 V/F is simple and can be accomplished two ways. The first is the photodiode can be biased to a fixed terminal voltage simply by substituting it for the phototransistor shown in Figure 15. Similar comments apply as to calibration, terminal voltages, etc.

The second is the photodiode can be operated in its photo-voltaic mode, with a basic terminal voltage of zero and a minimum load impedance. A circuit which accomplishes this is shown in Figure 16.

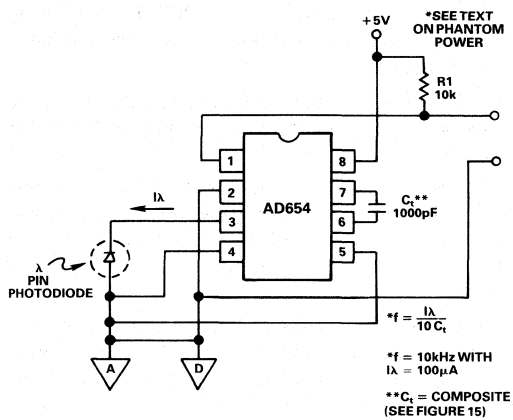

Figure 16. Photodiode Photosensitive Sources

The simplicity of this circuit is appealing, since the diode is merely tied between the two op amp inputs with the (+) input grounded. C_t is selected for the desired frequency and is the only circuit part critical to operation. Calibration in general has the same set of considerations as does the phototransistor, with similar solutions. It can be done by

making a small part of C_t a trimmer (as shown in Figure 15), or via software. Like Figure 15, this circuit is also easily used with phantom powering.

A legion of photo devices are available which are potentially suitable for use in either of these two circuits, operated over the wide range of currents/frequencies accommodated by the AD654.

Temperature-To-Frequency Converters

Low cost IC temperature transducers are both cost effective as well as easy to apply to remote measurements. Two cases are shown in Figure 17.

In Figure 17a, an AD592 IC temperature transducer is interfaced to the AD654 in a manner similar to that of the phototransistor (Figure 15). In the case of the AD592 the minimum bias voltage is 5V, so the circuit as shown will work down to 10V supplies (it will also work with higher supplies).

The AD592 produces a 1μA/K current output which drives pin 3 of the AD654. With the timing capacitor fixed at a value of 0.01μF this produces a frequency which is scaled as 10Hz/K; that is, 298K = 2980Hz. Since temperature is the parameter of interest, a 0.01μF NPO ceramic capacitor is used for lowest V/F TC. Calibration of this circuit can be accomplished either by using an input current divider or simply by subjecting the transducer to an ice reference and measuring the output frequency. If it differs from 2730Hz, software scale corrections can be applied to the measurement. A hardware calibration can be achieved by scaling the output a decade higher using a variable C_t of 1000pF.

Like the photosensitive circuits, this circuit is a good candidate for phantom powering. Optional components include protection resistance R_{PROT} and D1 (which are only necessary if it is possible for A2 to contact dangerous external voltages).

Figure 17b shows how the same basic transducer can also be used with a current offset to produce a Centigrade scaled V/F.

In this circuit, the AD592 drives the AD654 in the current input mode. The output of the AD592 is I_T and is scaled at 1μA/K. However, in this case a fixed offset current I_{OFF} is subtracted from I_T. Thus, what the V/F actually receives is the input $I_T - I_{OFF}$.

The V/F current input is offset by 273μA as established by R3-R4 and corresponds to 273K, or 0°C. Thus, this point is the zero frequency scale of the V/F. Above 0°C (273K) the V/F operates with a sensitivity of 100Hz/°C. An interesting feature of this particular circuit is that the offset current is regulated by the "floating" reference diode A3, even though the common-mode input to the AD654 can vary with supply voltage.

Trimming the circuit for zero scale calibration can be accomplished by placing sensor A2 in an ice bath reference and adjusting R3 (ZERO) until the output frequency just stops. To complete the calibration, the circuit is adjusted for 10,000Hz with A2 at 100°C (boiling water). This

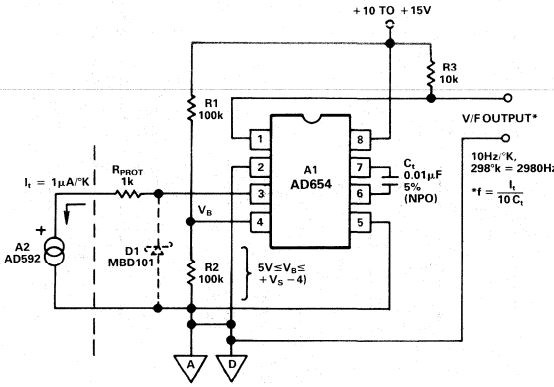

a. Kelvin-Scaled Temperature – to – F

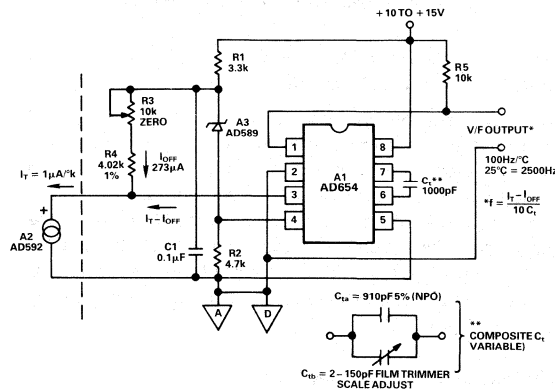

b. Centigrade-Scaled Temperature – to – F

Figure 17. Temperature-Frequency Circuits

is done simply with a hardware trim via C_t with the film trimmer. Here, the trimmer is C_{tb} (SCALE) while the bulk of the timing capacitance (C_{ta}) is an NPO type. Note that the trim is best done with a plastic tuning wand since C_t is floating. As Teflon and polypropylene film trimmers are available up to about 200pF, this type of tuning can be effective for circuits using a C_t of 1000-2000pF. The scale trim should be repeated once for best accuracy.

4-20mA Current Loop Conversion Circuits

The 4-20mA industrial current loop is a commonly used interface. Converting such input signals into a proportional frequency is relatively easy and allows ac transmission techniques to be used. Two examples are shown in Figure 18.

The most straightforward way to convert a 4-20mA signal is to convert the current to a voltage with a precision load resistor, then convert the voltage to a frequency with a V/F, and then transmit the signal in frequency form to a microprocessor for removal of the offset (i.e., the frequency corresponding to 4mA).

The offset in *current* form can also be removed at the V/F stage as shown in Figure 18a. Here, the load R_L converts the 4-20mA signal to a 40-to-200mV voltage which drives the V/F.

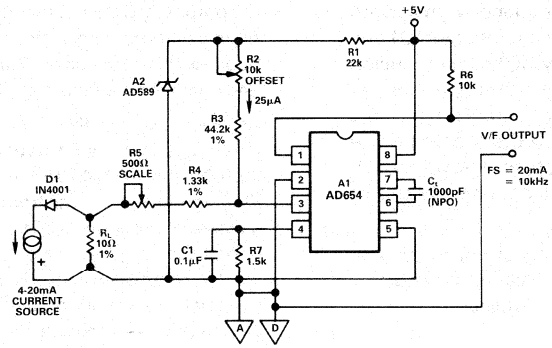

a. Offset Conversion of 4 – 20mA Signal

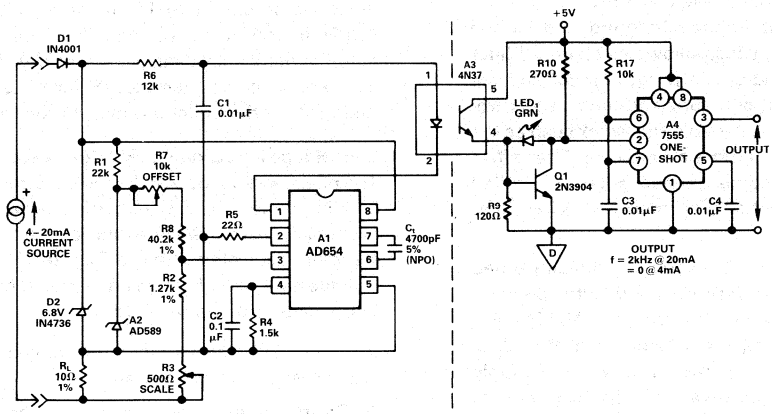

b. 4 – 20mA Loop Powered, Isolated to F

Figure 18. 4 – 20mA Conversion Circuits

The bias network R2 through R5 provides an input voltage offset of 40mV so that the 4mA zero scale will produce a zero output frequency. The basic V/F scale factor is set via R4 + R5. From 40mV to 200mV across R_L the AD654 sees a 160mV change, thus R_t needs to be about 1600Ω for an I_t of 100μA. With a FS I_t of 100μA, C_t is 1000pF for a 10kHz FS frequency.

Calibration is achieved by applying a 20mA input and trimming R5 for 10kHz out. Then, with 4mA applied the output can be checked for zero frequency, and trimmed (if necessary) with R2. Note that the two adjustments (R5 and R2) are interactive. Because of this factor, the calibration process should be repeated at least once. The circuit as shown assumes a floating source, so D1 might be necessary for reverse protection. It can be eliminated for a known input polarity but even if included it will not harm accuracy. Since the R_L terminal sense voltage is very low with this circuit (200mV) line voltage drops are not likely to be a problem.

The circuit of Figure 18b is an interesting variation of the basic one in Figure 18a as it is powered by the *loop current*. This is possible since the lowest loop current of 4mA is higher than that required to bias the AD654 plus the support circuitry. The advantage of this technique lies not just in the fact that the power is "free", but that it allows a completely floating I/F converter to be implemented using isolation techniques.

In this circuit the offset and scaling of the AD654 is very close to that used in Figure 18a. D2 clamps the supply voltage of the V/F to 6.8V and the offset reference diode A2 is biased from this voltage by R1. The AD654 V/F sees virtually a constant supply voltage for loop currents above about 2mA.

Since the desired LED drive current of ≃10mA is greater than the low scale input current an alternate technique must be used; in the case here this is a shift of duty factor for the LED drive. Thus, getting the signal out of the V/F and into logic stages requires a pulse forming network R6-C1. This allows the charge on C1 to be dumped in the opto isolator's LED once per cycle of oscillation, which in turn is coupled thru to Q1 and regenerated by the one shot A4. The technique works over the entire 4-20mA range, with the peak discharge current safely limited by the AD654's output transistors.

This circuit is scaled for a 2kHz FS frequency with a 20mA FS input current. The recharging considerations of C1 prevent it from working at appreciably higher frequencies. While this low frequency would appear to suggest the use of a low-speed optoisolator and buffer, such is not the case. The discharge of C1 is fast; on the order of a few microseconds. This requires the relatively fast buffer, Q1. The use of the following one shot is optional; as shown it produces positive 110μs pulses.

This circuit is quite valuable wherever current source signals of 4mA and up are used and isolation is required, whether or not the offset bias function is used.

Digitally-Tuned Switched Capacitor Filter

Switched capacitor (switchcap) filters using IC packaged state variable building blocks have become quite popular. Two reasons for this fact are the all-around utility of the state variable filter itself, which allows all standard filter modes to be resistor and/or pin programmed, and the ability to easily and precisely tune the filter center frequency via a clock input. A variable rate clock such as the output of a V/F can readily drive a switchcap filter effectively, avoiding cumbersome ganged pots and limited range tuning. This is particularly important when filter sections are cascaded.

Figure 19 is an example of a 12-bit binary, digitally tuned switchcap filter using the MF10C CMOS filter building block driven by an AD654. In this circuit, the AD654 V/F is operated as a single supply digitally programmed clock source. With the V/F device operated from a +10V supply the output will be a 10V square wave (f_{CLK}), which satisfies the clocking requirements of the MF10C filter.

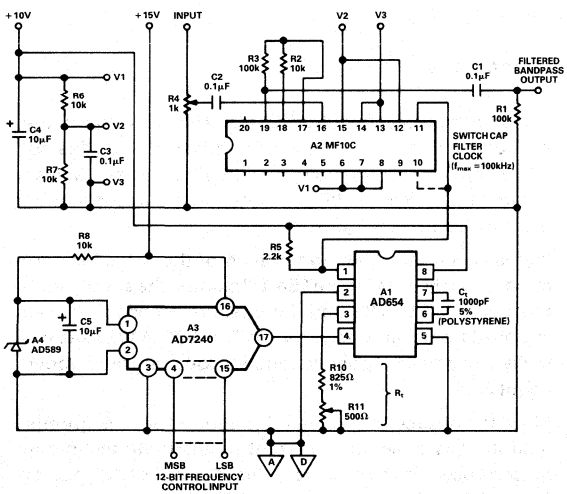

Figure 19. Digitally-Tuned Switched Cap Filter

A2 and the associated components comprise the V/F tuned switchcap filter. As connected here, the MF10C is operated in its "1a" bandpass mode with a gain and Q as determined by R3/R2. The signal to be filtered is applied across input level control R4 and fed via C2 to the filter. The filtered output appears at pin 19 and is ac coupled to the output by C1. Note that the MF10C is comprised of *two* filter sections; only one of which is operated in this hookup. It is also capable of operating in the HP, LP, AP, and notch modes (see manufacturer's literature).

As operated here the filter center frequency is $f_{CLK}/100$, where f_{CLK} is the V/F frequency. This in turn is programmed by the AD7240 D/A with 12-bit control resolution. Thus the filter's center frequency is programmable over a dynamic range of 2^{12}. The precision of tuning will be as good as the f_{CLK}/f_O spec of the MF10C, which is ±1.5% or better. This assumes calibration for an actual filter frequency of 1kHz with an all 1's D/A input. Further discussion of this D/A tuning technique of the AD654 is covered in a later application (Figure 23).

This versatile circuit cannot only achieve all the filtering functions noted above, but it can also be used with the AD654 in any of the standard V/F input control modes, tuned either electronically or manually. A simple, manually variable clock generator (Figure 22 series) will allow tuning of filtering functions over a wide range.

Wide Range Precision PLL F/V Converter

The phase locked loop (PLL) has become a familiar circuit building block, one routinely used for a variety of signal processing tasks. Since a number of readily available, standard, dedicated chips perform the PLL function, a logical question would be why one would need to use a V/F. The answer to this is two-fold:

1) To achieve a greater dynamic range of frequency lock, i.e., more than a single decade.

2) To implement precision F/V circuits. These performance attributes are not available from the dedicated PLL chips, but can be provided by V/F circuits.

The V/F based PLL circuit of Figure 20 uses a wide dynamic range digital frequency/phase detector as the PLL comparator. With this type of circuit the input signal can range over a very wide span and the locking characteristic is not harmonically sensitive. It accepts a TTL compatible square or pulse input, f1. As shown here, f1 can be as high as 100kHz with overranging to 150kHz. The PLL produces a square wave TTL output, f2, at the same nominal frequency as the input. This output is locked to the input in terms of both frequency and phase and is referenced to the positive going edge of f1. In addition, the circuit also produces a low-pass filtered F/V output from stage A4. This signal, a smoothed version of the V/F's control voltage, is precisely proportional to the input voltage. It can be used as a precision F/V output voltage signal or as part of a larger overall loop . . . for example to characterize specific V/F devices at test. The entire PLL circuit operates from a single +5V supply.

The digital frequency/phase detector is made up of the 7474 dual TTL D flop (A2) and the associated gate sections (A3). When the loop is in lock, a positive leading edge at f1 and at f2 clocks to a "1" at the A2-a and A2-b Q outputs, respectively. Gate A3-a resets both flip-flop sections after about 50ns. The $\overline{Q}$ and Q outputs of A2-a and A2-b are, therefore, narrow pulses coincident in time, but opposite

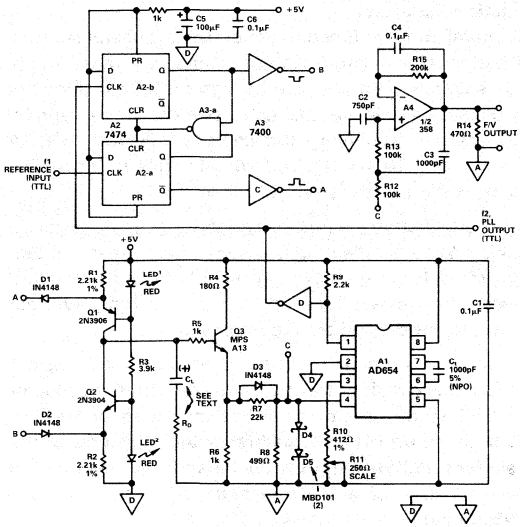

a. Circuit Diagram

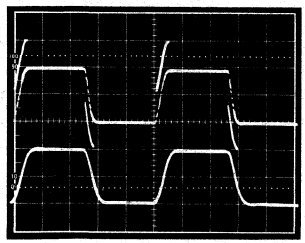

b. AD654 PLL Performance (Fast Response Mode)
Upper Trace: Phase Detector Output (Point "C")
Lower Trace: F/V Filtered Output

Scales: 0.2V/div, 500μs/div
Source: 50/100kHz FSK @ 400Hz Rate
Condition: C_L = 0.1μF, R_d = 500Ω

Figure 20. PLL

in phase. These outputs will only be 50ns in width for the locked condition. They are buffered and used to control the sampling gate at points A and B.

The sampling gate used consists of the current source transistors Q1 and Q2, which are in turn controlled by diodes D1 and D2. For the locked condition the coincident, equal width pulses will turn on the current sources for equal lengths of time, thus no net current will flow into the loop storage capacitor, C_L. However, if Q1 is on longer than Q2, C_L is charged towards +5V; conversely if Q2 is on longer, C_L is discharged towards ground. The voltage which is stored on C_L represents the raw dc error voltage of the loop and ultimately controls the V/F. This voltage is buffered by Q3 and is fed to the V/F through the network R7-R8-D3. The Schottky diodes D4-D5 simply provide an overvoltage input clamp for the V/F to prevent a potential latchup caused by excessive dc input.

The AD654 is set up here as a 100kHz, 0.5V FS single supply V/F with a TTL output buffer A3-d. Calibration for exactly 0.5V at the V/F input at FS is achieved by R11, the SCALE control. This is appropriate for the circuit's use as an F/V, with the A4 two pole low-pass filter stage included. If the F/V output is not required A4 and the associated components are deleted, and R10-R11 can become a single 499Ω, 1% film unit.

The actual loop time constants of any PLL should be optimized towards the final intended use. In this case, the loop can be adjusted for either quick response (with modest dynamic range), or for a wider dynamic range (with slower response). For fast response, the C_L-R_D series combination is 0.1μF and 510Ω. For a wider dynamic range, they become 1μF and 160Ω. In general, faster response will be desirable for such uses as FSK demodulation, etc. On the other hand, a precision PLL F/V would likely use the greatest dynamic range option. It is, of course, possible to switch the filter time constants to allow fast lockup for fast changes, and wide dynamic range once the input frequency is stabilized.

The general method of selecting the loop components is to select C_L consistent with the desired dynamic range and/or speed, and then select R_D for the required damping. This can be done either from the equations below, or by scaling from the values noted above.

The conversion gain, K_d, is set by R1-R2 and is:

$$K_d = \frac{500\mu A}{2\pi} \qquad (6)$$
$$= 8 \times 10^{-5} \text{ amperes/radian}$$

The V/F conversion sensitivity, K_o, is:

$$K_o = \frac{2\pi f_{FS}}{V_{FS}} \qquad (7)$$
$$= 1.26 \times 10^6 \text{ radians/volt-sec}$$

where f_{FS} is the FS frequency and V_{FS} is the FS V/F input voltage. For the conditions shown, these are 100kHz and 0.5V, respectively.

The natural frequency of this second order loop will be ω_n, which is:

$$\omega_n = \left[\frac{(K_o)(K_d)}{C_L}\right]^{1/2} \qquad (8)$$

The damping factor ζ is:

$$\zeta = \frac{R_D[C_L K_o K_d]^{1/2}}{2} \qquad (9)$$

It will usually be desirable to make ζ about 0.8 for mimimum overshoot with a step input; the values listed assume this. Practical ranges of C_L are from just under 0.1μF up to 1μF or more; the smaller values allowing faster settling, the larger greater dynamic range. With C_L selected R_D can be calculated to provide the desired damping, as:

$$R_D = \frac{2\zeta}{[(C_L)(K_o)(K_d)]^{1/2}} \qquad (10)$$

Since for this particular circuit the product of K_o and K_d is simply 100, this can be simplified to:

$$R_D = \frac{2\zeta}{[100\,C_L]^{1/2}} \qquad (11)$$

Figure 20b is a photo illustrating the performance of the PLL used in the fast response mode (see conditions). For this test the input is a 50/100kHz FSK signal, which represents a 1/2 FS change. The two traces represent the V/F control input (upper) and the F/V output (lower). The dc baselines are offset for waveform comparison purposes. In the upper trace the loop is seen to produce the large dynamic errors characteristic of a slewing interval, but is damped and settles quickly to the final dc value in just over 200μs. The lower trace shows a similar waveform, with the fast ac components reduced by the filter. Reliable data recovery to logic levels is possible simply by voltage comparison of this waveform to a fixed dc reference which corresponds to 1/2 the p-p amplitude.

With the loop adjusted for wide dynamic range (C_L = 1μF, R_D = 160Ω), it will maintain a frequency and phase lock from 100kHz down to about 500Hz. Since it can also overrange to 150kHz, the total dynamic range is more than 200/1. With C_L = 1μF, settling is proportionally longer, about 2ms for the same 1/2 scale frequency step (50/100kHz).

An interesting feature of the circuit is that it requires very few precision components. C_t should be an NPO type for best stability and linearity, but R_t (R10-R11) is the only other precision scaling component. R1 and R2 should be stable film units as noted (since they set K_d), but this is not absolutely critical. C_L is best a polypropylene film for fast response applications, where the values are small (≃ 0.1μF). Where space is at a premium and fast settling is not critical, a high-voltage electrolytic can be used (such as 1μF/50V). R_D can be a fixed value type, and if desired it can be easily optimized for damping with the use of a dynamic display such as Figure 20b (top). Power supplies for this circuit should be well bypassed with low impedance at high frequency capacitors near the logic stages. The analog grounds (including R_D) should be returned to the V/F's common point as shown.

There are a number of possible uses for a PLL such as this. Note that f1 can be a divided crystal reference for synthesis uses. Similarly, the f2 output can be divided before the comparision to scale frequencies upward. This PLL circuit capitalizes on the virtues that a wide range V/F device brings to PLL applications circuits.

Function Generator

A general purpose function generator is a useful lab tool, as it is capable of producing the standard waveforms of sine, square, and triangle. Although the AD654 does not produce a triangle wave directly, one can be realized by differentially amplifying the capacitor waveform(s). A sine wave can be synthesized from the triangle wave. The square wave is already available from pin 1.

Figure 21 illustrates a function generator circuit based on an AD654 with a wide range, exponentially generated timing current for frequency control. This allows a standard linear pot to be used for a 4 decade control range without crowding the low end of the control span.

The control current generator is comprised of A2, A3, and Q1-Q2. This circuit produces a current from the collector of Q2 which ranges downward as controlled by the FREQUENCY control R9. R2 calibrates the high end of the scale for a 100kHz output while R6 sets the lower end to 10Hz. Q1 and Q2 are a matched pair, with high gain (2N5089, etc.).

Dual op amp A4 and single device A5 make up a low bias current, high-speed buffer for the timing capacitor. For the input stage, nothing other than a high slew rate FET input device should be used, to preserve the V/F dynamic range. The output of the instrumentation amp is from A5 and is a ground referenced, 3.6V p-p triangle wave. If the resistor ratios are matched as shown, this waveform will have a negligible dc component. R14 serves generally as an amplitude control, but is best adjusted for minimum sine wave THD (below).

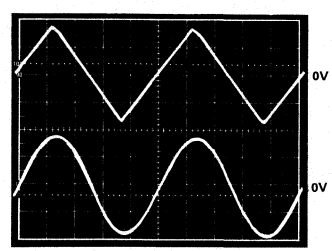

Output Waveforms

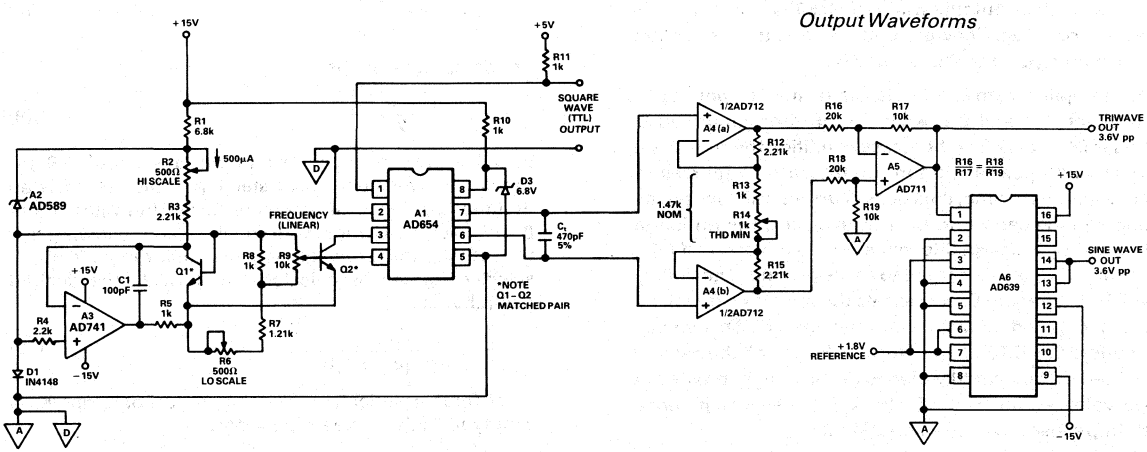

Figure 21. Function Generator

A6 is an AD639, a precision trigonometric function generator capable of (synthesized) sine wave THD as low as 0.02%, given an input triangle of sufficient linearity. With a triangle wave input amplitude of twice the device's 1.8V reference amplitude at pin 6 the chip produces a low distortion sine wave of a 3.6V p-p amplitude. Since the AD654 triangle taken differentially is 1.8V p-p, the IA net gain requirement is 2 for an optimum triangle wave drive to the AD639. The IA gain is best *trimmed* to the nominal gain of 2 via R14. This is done using the lowest observed output sine wave THD as an adjustment criterion. This allows the various circuit tolerances to be adjusted out and guarantees lowest THD. The AD654 triangle wave as it is used here is not actually good enough to take complete advantage of the AD639's low distortion capability. Nevertheless, the measured THD in this circuit is 0.5% over the full 100kHz range (after trim). Some improvement in THD performance can be expected if the AD654 is operated towards a higher I_t, such as 1-2mA. The photo indicates the quality of waveforms attainable with conditions as shown.

This circuit is useful as it is shown for general purpose work, with buffered triangle and sine wave outputs and a TTL square wave output from A1. It can easily be adapted to other means of frequency control such as the many other voltage input options throughout these notes. Also, the log control input can be enhanced further as shown below (Figure 26).

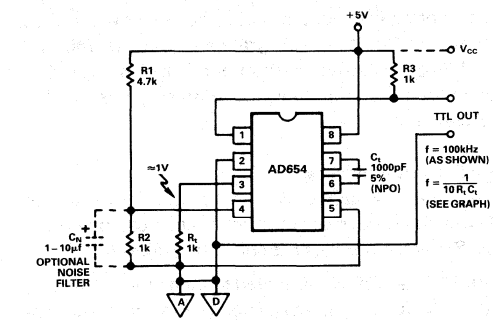

a. Simple TTL Clock Source (Fixed Frequency)

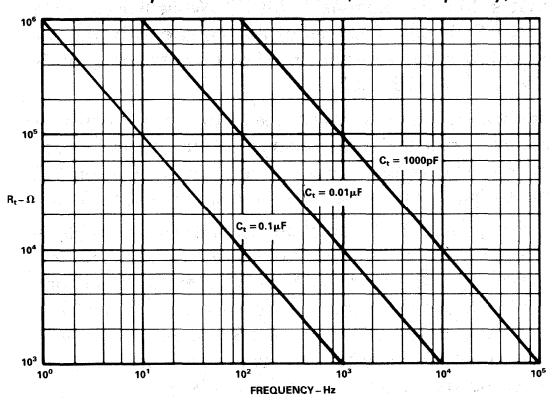

a1. Spot Frequency Clock Source Nomograph

Clock Sources

Not only is the AD654 useful as a wide range controlled frequency source, it is also useful as a stand-alone clock source. Here, the general distinction is that clock sources are not highly dynamic in terms of frequency control. Attributes of the AD654 as a fixed frequency clock source include the ease of logic output adaptation, low power, and the ease of tuning when required. Figure 22 illustrates a number of useful clock circuits.

One of the simplest clock sources possible with an AD654 is shown in Figure 22a and is suitable for TTL or CMOS uses. Note that even if the circuit is powered from +5V the output can be pulled up to *higher* voltages, such as 10 or 15V CMOS. For a control voltage, the 5V bus is simply divided down to 1V, and used as an input (with optional filtering). The output frequency is the expression shown, since V_{IN} is 1V.

This circuit is most useful where a single spot frequency is desired and a fixed R-C combination can be selected. For example, with C_t of 1000pF, an R_t of 1kΩ will yield 100kHz as shown. For lower frequencies R_t can easily be adjusted upward; for example when it is 1MΩ, 100Hz will be produced. The nomograph of 22a1 is useful in selecting R_t values for various frequencies; not only for a C_t value of 1000pF, but for other values as well. Another example is 0.01μF and 100k, which yields 100Hz. It is

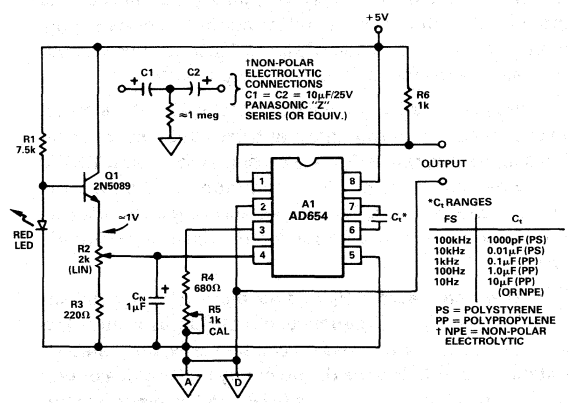

b. 10/1 Manually-Tuned Clock Generator

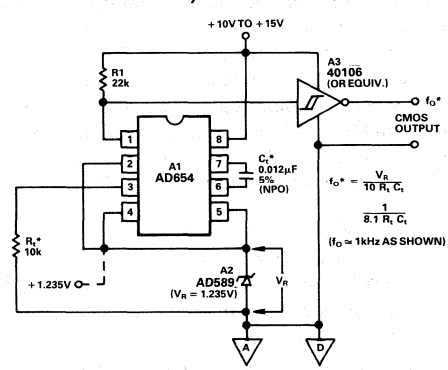

c. Self-Biasing Precision Clock

Figure 22. Clock Circuits

possible to use C_t values below 1000pF, or R_t values above 1MΩ. The tradeoff is simply that some predictability will be sacrificed in doing so. This nomograph is based on a fixed V_{IN} of 1V and is generally applicable to other circuits with such operating conditions. Examples are the two circuits which immediately follow.

With a little more effort, the same circuit can be adapted for manual tuning and switched capacitors. The circuit of Figure 22b provides this, with smooth tuning over a 10/1 range.

Here the simple 1V supply divider of Figure 22a is replaced by an LED reference circuit which improves supply stability and adds coarse temperature compensation for the diode drift. One volt is developed at the emitter of Q1 and drives the linear frequency control R2. Note that this low-cost circuit features utility with reasonable performance, but for the best TC an AD589 1.2V reference will be useful driving R2.

Resistors R4 and R5 here function as R_t, and the nominal value can be seen on Figure 22a1 with the various C_t values (see table). R5 serves to calibrate the circuit for full scale. For this circuit, a table of recommended C_t capacitor types are shown, including a preferred non-polar electrolytic hookup for low frequency (albeit non-precision) scale.

The previous two circuits are very strong on utility but as shown lack high precision and immunity to power supply changes. The clock circuit of Figure 22c improves on both of these weaknesses and adds an output buffer.

This circuit capitalizes on the fact that the AD654's internal quiescent current is regulated, and uses this current at pin 5 to excite a reference diode A2. This establishes a 1.235V reference, V_R, from pin 5 to common. This voltage is highly immune to supply changes, allowing the diode current to be stabilized without the expense of additional parts.

V_R also appears across R_t, and develops a timing current of V_R/R_t making the frequency expression as noted. The general frequency expression can be easily simplified so that the V_R and C_t factors divide out by making C_t a power of 1.2. This allows easy values for R_t, and also permits the use of the frequency nomograph 23a1 for R_t-C_t selection. To do so, simply multiply the capacitor curves as shown by 1.2 (1200pF, 12000pF, etc.) and select R_t. Of course the circuit will still need to be calibrated if exact frequency operation is needed. The output signal from pin 1 is not TTL compatible, but it is CMOS compatible with the use of a CMOS Schmitt buffer as shown. This restores a rail-rail 5V logic swing, which can drive either 5V CMOS, or TTL.

12-Bit D/F Converter

Closely allied to the clock sources, as well as standard V/F circuits is the D/F converter, a digitally programmed frequency source. A useful example of such a circuit is the 12-bit D/F shown in Figure 23.

This circuit uses the AD7240, a 12-bit resolution CMOS D/A available with differential nonlinearity as low as 1/2LSB. It is used here with the AD589 1.235V reference in

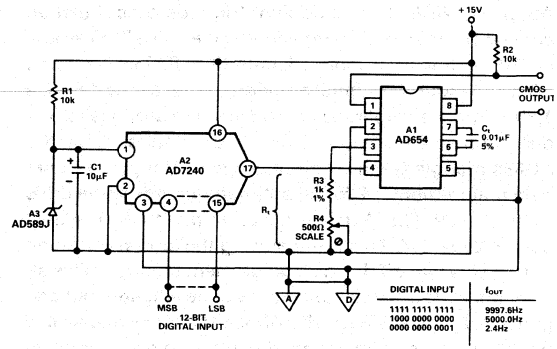

Figure 23. 12-Bit D/F Converter

a voltage output mode, driving the (+) input of the AD654 directly. Since the AD7240 is a multiplying D/A, it multiples the 1.2V reference by the fractional binary input, producing a variable voltage output similar to a pot. The AD654 is scaled here for a 1mA FS current; a C_t of 1000pF yields a FS output (for an "all ones" input) of 99976.Hz. The timing capacitor can be changed for other full scale frequencies. R_t is trimmed via R4 for FS calibration.

Nonlinearity error of this circuit will be predominantly due to the AD654 if the AD7240 is operated from a 15V supply as shown. Inasmuch as this circuit's linearity will be on the order of 0.1%, the best linearity grade for the AD7240 may not be cost effective. Nevertheless, the circuit will still be monotonic for 12-bits for all grades of the AD7240 and have 12-bits of control resolution. Similarly operated D/A devices of lower resolution can also be employed in this general circuit.

Voltage/Current Monitors

Using isolated V/F conversion techniques, a variety of useful voltage and current monitor circuits can be implemented. Several examples are shown in Figure 24.

In Figure 24a is a AD654 V/F scaled to measure a −5V bus, with a 1.0Hz/mV sensitivity. The output from the optoisolator is fed to one of the output buffers of Figure 10 (that is, 10a1 for example) to 5V TTL. While this example is for a 5V bus, a simple adjustment of the R1/R2 divider will allow other voltages. R5 should be selected for 9-10mA of drive in the LED.

Figure 24b illustrates a technique for monitoring negative supply current, in this case on a −20V bus. In this hookup, the negative voltage input mode is used with the V/F measuring the drop across R_S, a series sampling resistor.

In this circuit R1-R3 constitute R_t, and a value of R_S is selected from the table for the desired full scale current. In each case, the sense voltage is 100mV for minimum loss in the supply voltage. The capacitance C_N is a large value, selected for more than 25dB of noise attenuation. This circuit is scaled for the sensitivities shown in the table, and the nominal FS frequency is 1kHz.

A very interesting form of current monitor is shown in Figure 24c. This circuit is self powered and measures the

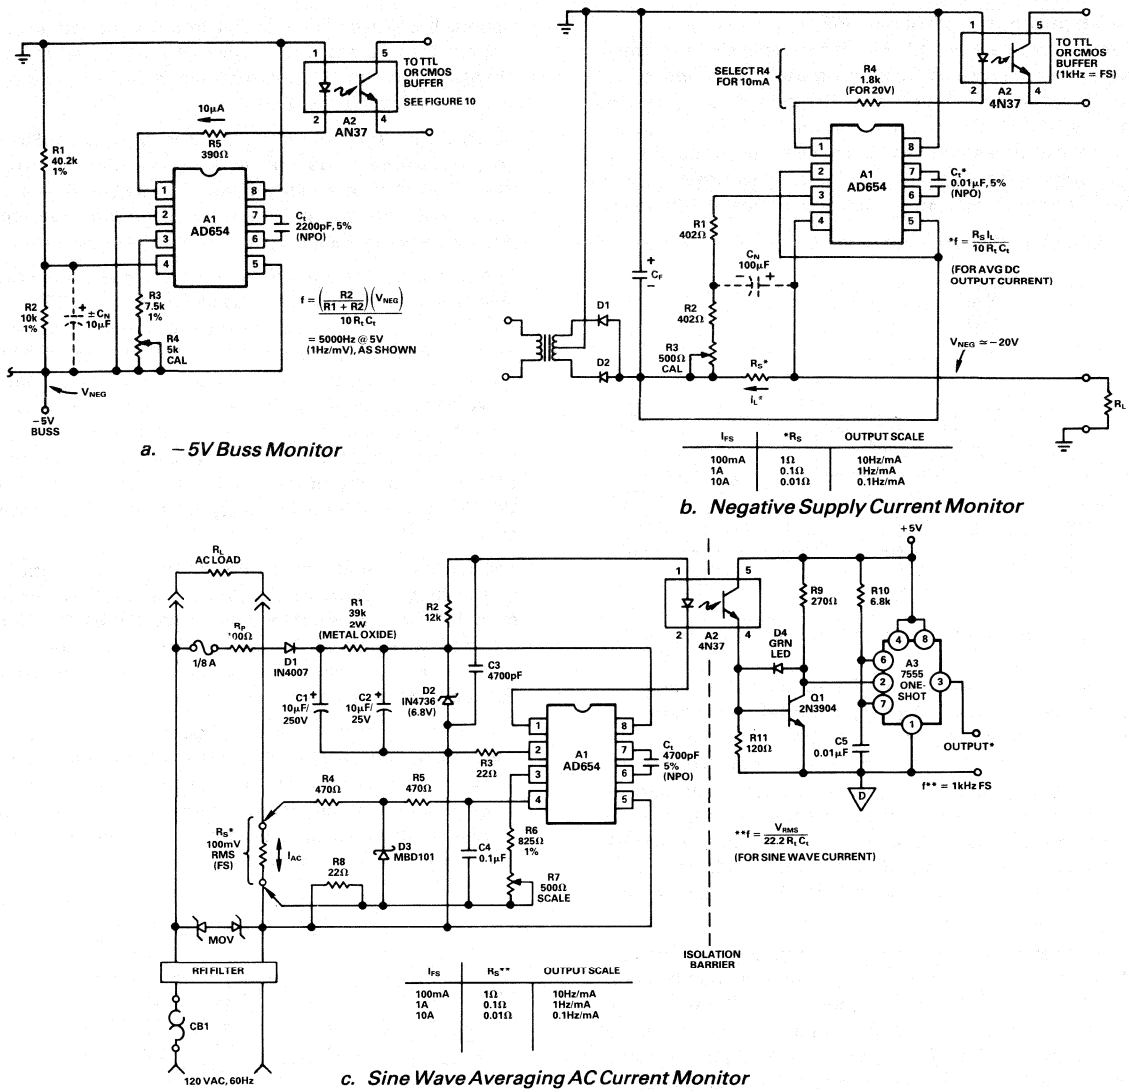

Figure 24. Voltage Current Monitors

average ac line current in the load R_L delivering an isolated output frequency. Several aspects of the circuit are similar to previous applications, such as the pulsed LED drive and the one shot regeneration (i.e., A2 and A3). They will not be discussed extensively here.

The circuit samples the current in R_S by measuring the ac terminal voltage. This is possible with an AD654 if the p-p ac input is only a few hundred mV. The AD654 simply half wave rectifies the ac and then measures the dc average over ½ cycle. Of course the input must be protected against overvoltage, thus the use of D3. C4 is simply a hash filter.

With a 100mV rms FS drop across R_S, the circuit is scaled to produce the frequency noted, 1kHz FS. R_S is selected for various FS currents, so as to produce 100mV rms. Note that since this is an average responding ac measurement, nonsinusoidal currents will not read correctly. For such applications, an rms-dc converter can be used, such as the AD636.

The circuit can be either line powered by a peak detector circuit and series resistor (D1, C1, R1), or can be fed from a small rectifier transformer and simple half-wave supply. The supply as shown functions adequately, but it is not efficient from a power consumption standpoint. When working with circuits operated directly from the line, it is of paramount importance that adequate safety precautions be taken. This would include safety and reliability considerations such as overvoltage clamps, breakers, sufficient derating, etc.

Although the circuit as it is shown measures current, it can also be rewired to measure line voltage by adding a divider so as to produce 100mV rms ac at R4.

High Input Impedance Bipolar V/F

A more precise bipolar V/F with a parallel sign bit output is shown in Figure 25. This circuit uses a low offset voltage, low bias current precision FET input op amp preceding the AD654. This allows operation from very high impedance sources and can achieve very accurate conversion down to millivolt levels with good symmetry about zero. The circuit operates from conventional ±15V power supplies.

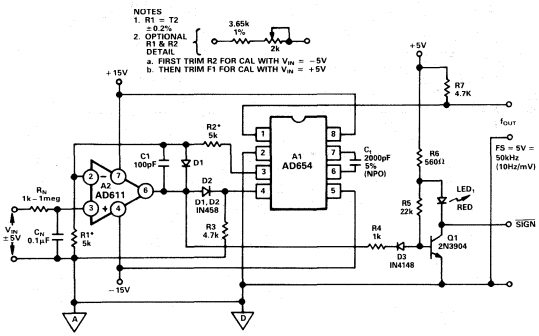

Figure 25. Ultra High Input Impedance Bipolar V/F

In this rather unusual circuit the AD654 is alternately driven in one of its two voltage input modes for the two input polarities. It acts as a positive input voltage V/F (for "positive" signal inputs), and as a negative input voltage V/F (for "negative" signal inputs). A2 is a low input current precision op amp which loads the source with only 10pA of bias current (typical). Low leakage diodes D1 and D2 provide switching at the AD654's inputs, while the precision resistors R1 and R2 set the V/F scale in conjunction with timing capacitor C_t.

Since R2 sets the scale factor for "negative" inputs, and R1 for "positive" inputs, gain calibration can be a bit complex. The most expeditious means of FS calibration is via software, with a known reference input, a fixed C_t, and a closely matched R1-R2. Alternately, C_t can be made 2200pF and R1-R2 trimmed, as noted. The scaling with the values shown is 50kHz for a FS V_{IN} of 5V (or 10Hz/mV), and excellent linearity is achieved for inputs down to 1mV or less. For best accuracy around zero a low offset device should be used for A2, such as the AD611K (500μV max), or A2 should be trimmed.

The circuitry around Q1 is a sign amplifier and produces a TTL compatible SIGN output as well as LED visual indication. Aside from the high precision of this circuit as a bipolar input V/F, the input buffer technique used is also useful for many single polarity sources which might require minimal loading, or the addition of long-time constant input filters.

Log Controlled Oscillator

For manually tuned oscillator applications, the widest range of control comes about with a logarithmic control characteristic. However, a log taper pot with a repeatable 4 decade characteristic is not a real-world item. A better approach is to use a linear pot fed into an exponential control circuit. This has the practical effect of a log control characteristic with a standard pot taper, but one not sensitive to the control's absolute value.

Figure 26 is a log controlled oscillator using the AD654 in conjunction with a dual op amp and a monolithic IC array. The array uses the same general principle discussed above with the Figure 21 function generator, but with some enhancements.

The lower part of the circuit is simply a hot substrate stabilizer for the plastic packaged CA3046 array. Q5 of the array is a diode connected sensor and Q3-Q4 the heaters. The V_{BE} of Q5 is compared against 650mV derived from the +15V supply; the loop forces the chip temperature upwards to about 40°C. This is only slightly in excess of room temperature, but more importantly it will be *constant*. The constant temperature removes the strong temperature sensitivity of the basic exponential generator, Q1-Q2, and it avoids the necessity of a special TC compensation resistor.

The current generator is composed of Q1 and Q2 with Q1 establishing a reference V_{BE} proportional to the current established by the reference diode A2 and R2-R3. R2 is trimmed to set this current to a nominal 500μA, but also to trim out circuit tolerances. With a C_t of 470pF this produces a 100kHz frequency.

The lower end of the operating range is established by the divider across Q1, R6-R8. R8 drops a voltage which is adjusted to be equal to 4 (KT/q ln 10), or about 240mV, as set by R6 for low scale calilbration. With R9 in shunt with R8, a sweep of this control spans 4 decades of current, controlling the oscillator over 4 decades.

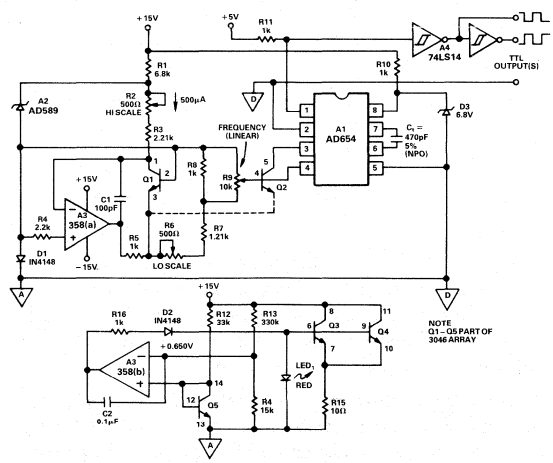

Figure 26. Log Controlled Oscillator 10Hz – 100kHz

Frequency Doubling

Since the AD654's output is a square-wave rather than a pulse train, information about the input signal is carried on both halves of the output waveform. The circuit in Figure 27 converts the output into a pulse train, effectively doubling the output frequency, while preserving the

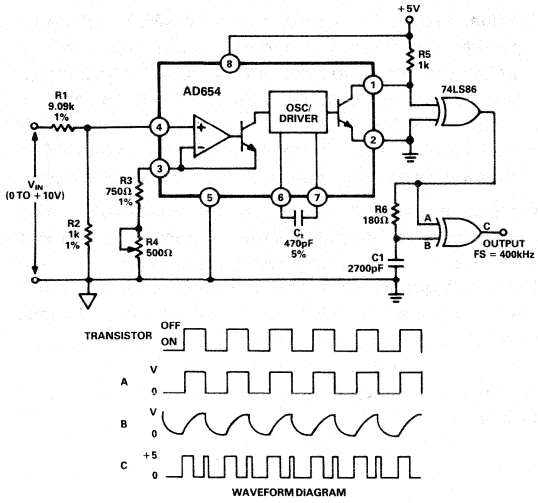

Figure 27. Frequency Doubling

better low frequency linearity of the AD654. This circuit also accommodates an input voltage that is greater than the AD654 supply voltage.

Resistors R1–R2 are used to scale the 0 to +10V input voltage down to 0 to +1V as seen at pin 4 of the AD654. Recall that V_{IN} must be less than V_{SUPPLY} −4V, or in this case less than 1V. The timing resistor and capacitor are selected such that this 0 to +1V signal seen at pin 4 results in a 0 to 200kHz output frequency.

The use of R6, C1 and one XOR gate doubles this 200kHz output frequency to 400kHz. The voltages seen at the input of the second section of the 74LS86 are shown in the waveform diagram. Due to the difference in the charge and discharge time constants, the output pulse widths of the 74LS86 are not equal. The output pulse is wider when the capacitor is charging due to its longer rise time than fall time. The pulses should therefore be counted on their rising, rather than falling, edges.

2MHz, Frequency Doubling V/F

Operation of the AD654 via the conventional output (pins 1 & 2) is speed limited to approximately 500kHz for reasons of TTL logic compatibility. However, even though the output stage of the device may become speed limited at this output the multivibrator core itself will still continue to oscillate to 1MHz or more. Advantage can be taken of this feature to allow even higher frequencies of operation then might be expected, frequencies well in excess of 500kHz. To implement such an "exalted" mode, a basically different signal extraction method is used. Rather than using the internal output stage the timing capacitor signal is buffered and zero crossing level detected to produce a high speed, TTL square wave output.

Figure 28 is a circuit example of this type of operation illustrating a 2MHz full scale V/F. In this circuit, the AD654 is operated at a FS of 1mA with a C_t of 100pF. This achieves a basic device FS frequency of 1MHz across C_t. The P channel JFETs Q1 and Q2 buffer the differential timing capacitor waveforms to a low impedance level where the push-pull signal is then ac coupled to the high-speed comparator A2, an LM360 (or μA760). This device has a low delay of 12ns (typical), and push-pull TTL outputs. Hysteresis is used, via R7, for nonambiguous switching and to eliminate the oscillations which would otherwise occur at low frequencies. While the value of this resistor is not critical, it must *not* be omitted. Without the hysteresis the circuit will generate spurious output frequencies and simply be uncontrollable.

The net result of the above is a very high-speed circuit which does not compromise the AD654 low scale dynamic range as the FET buffers will typically have only a few pA of bias current. At the other extreme, the high end

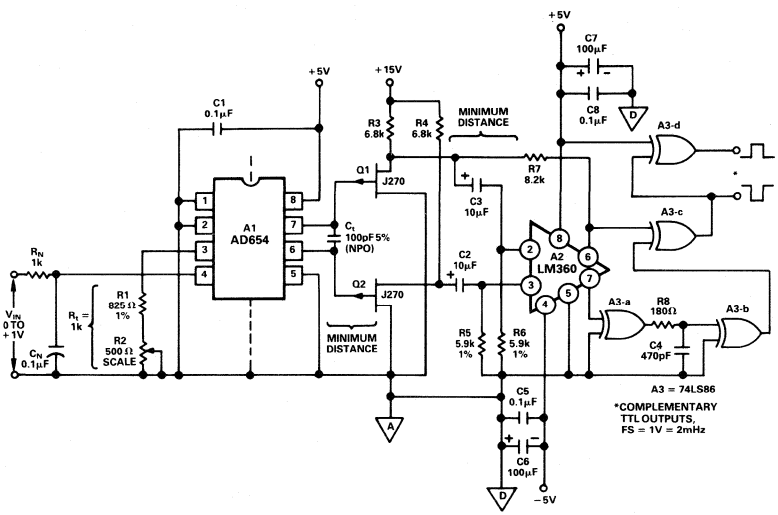

Figure 28. 2MHz, Frequency Doubling V/F (Exalted Mode)

dynamic range is limited essentially by the AD654 parasitic package and layout capacitances in shunt with C_t, and also those from each node to ac ground. A tight, minimum lead length PC layout will help between A1-6/A1-7 and Q1/Q2. When constructed properly the circuit is capable of stable operation to frequencies higher than 1MHz. A ground plane will also help stability, connnected to analog ground at one point. Low impedance bypasses on all supplies are also needed, as noted.

The output of the comparator is a complementary square wave at 1MHz FS. The XOR gate following A2 acts as an edge detector, producing a short (≈50ns) pulse for each input state transition. This effectively doubles the V/F FS frequency to 2MHz with complementary logic outputs at gates A3-c and A3-d. Note that this frequency doubler technique is not at all unique to this particular V/F circuit; it can be applied to a number of other circuits in this note as well.

The final result is a 1V FS input V/F with a 2MHz FS output capability. The increased dynamic range and speed of this circuit can open up many additional categories of applications. For example, wider data bandwidths can now be handled by the V/F since the operating frequency is about an order of magnitude higher than is standard. High speed data communications and video are two examples.

Interestingly, while the linearity related error of this circuit is not as good as it might be if the V/F were to be operated at 100kHz, the relative degradation is actually much less than the frequency scaling would suggest. For example, while a 100kHz scaled V/F might have a nonlinearity on the order of 0.1%, this circuit has a typical nonlinearity of about of 0.5% (or less). This level of performance is useful for many types of applications, particularly so in view of the low cost.

REFERENCES:

1) Gilbert, B.A. "A Versatile Monolithic Voltage-to-Frequency Converter," *IEEE Journal of Solid-State Circuits*, vol SC-11, pp. 852-864, Dec. 1976.

2) Kress, D.; Gilbert, B.A. "Versatile Monolithic V/f or I/f Converter," *Analog Dialogue*, vol 10-2, pp. 6-8, 1976.

3) Jung, W.G. *IC Op Amp Cookbook, 2nd Edition*, Howard W. W. Sams and Co. Inc, Indianapolis, 1980.

4) Jung, W.G. *IC Timer Cookbook, 2nd Edition*, Howard W. Sams and Co. Inc, Indianapolis, 1983.

5) Jung, W.G. "V/F Converter Doubles as Clock and Input of Stable Sine-wave Source," *Electronic Design*, pp. 328, November 15, 1984.

6) Gilbert, B.A. "Monolithic Analog Trigonometric Function Generator," *Analog Dialogue*, vol 18-3, pp. 12-14, 1984.

7) DeVito, L.M. "AD650 Data Sheet," Analog Devices

8) Gardner, F.M. *Phase Lock Techniques, 2nd Edition*, Wiley, 1979.

SPECIAL SOURCES

Listed below are some manufacturers of specialty electronic components useful within V/F circuits.

Manufacturer:	Component(s):
Corning Glass Works 3900 Electronics Drive Raleigh, NC, 27604 (919)-876-1100	NPO (COG) Ceramic Caps Types CAC02/03/04/05
F-Dyne Electronics 449 Howard Ave. Bridgeport, CT, 06605 (203)-367-6431	Polystyrene Caps Type PJ (See Also PSA, PST Types)
Optoelectronics Division General Instrument Corp. 3400 Hillview Ave. Palo Alto, CA, 94304 (415)-493-0400	4N137, 6N135/6N136 (See Also Phototransistors, Photodiodes, and LEDs)
Tel-Labs, Inc. 154 Harvey Road P.O. Box 375 Londonderry, NH, 03053 (603)-625-8994	WW and Metal Film Resistors (See Also Special Linear, Pos TC Resistors; +140ppm/°C)

ANALOG DEVICES

AN-279
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Using the AD650 Voltage-to-Frequency Converter As a Frequency-to-Voltage Converter

By Steve Martin

The AD650 is a versatile monolithic voltage-to-frequency converter (VFC) that utilizes a charge balanced architecture to obtain high performance in many applications. Like other charge balanced VFCs the AD650 can be used in a reverse mode as a frequency-to-voltage (F/V) converter. This application note discusses the F/V architecture and operation, component selection, a design example, and the fundamental trade-off between output ripple and circuit response time.

F/V CIRCUIT ARCHITECTURE

Figure 1 shows the major components of the frequency-to-voltage (F/V) converter. It includes a comparator, a one-shot with a switch, a constant current source, and a lossy integrator. When the input signal crosses the threshold at the comparator input, the comparator triggers the one-shot. The one-shot controls a single pole double throw switch which directs the current source to either the summing junction, or the output, of the lossy integrator. When the one-shot is in its "on" state, there is current injected into the input of the integrator and its output rises. When the one-shot period has passed, the current is steered to the output of the integrator. Since the output is a low impedance node, the current has no effect on the circuit and is effectively turned off. During this time the output falls due to the discharge of C_{INT} through R_{INT}. When there is constant triggering applied to the comparator, the integration capacitor will charge to a relatively steady value and be maintained by constant charging and discharging. The charge stored on C_{INT} is unaffected by loading because of the low output impedance of the op amp.

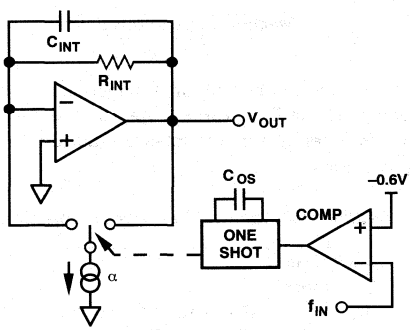

Figure 1. Circuit Architecture

THEORY OF OPERATION

Figure 2 shows a simplified representation of the AD650 in the F/V mode. Figure 3 represents the current i(t) delivered to the lossy integrator. The current can be thought of as a series of charge packets delivered at frequency $f_{IN} = \frac{1}{T}$ with constant amplitude (α) and duration (t_{OS}).

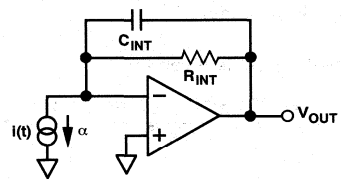

Figure 2. Simplified Schematic

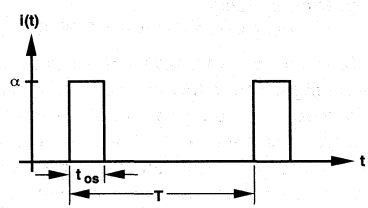

Figure 3. Current i(t) into Integrator

From inspection of Figure 3 the average value of input current is found by dividing the area of the current i(t) by the period T. The dc component of the output voltage is found by scaling the average input current by the feedback resistor R_{INT}.

$$V_{OUTAVG} = \frac{\alpha\, t_{OS}}{T} \cdot R_{INT} \qquad (1a)$$

Equation 1a becomes a linear function of frequency when f_{IN} is substituted for $\frac{1}{T}$.

$$V_{OUTAVG} = t_{OS} \cdot R_{INT} \cdot \alpha \cdot f_{IN} \qquad (1b)$$

Notice that the relationship between the average output voltage and input frequency is a function of the one-shot time constant and the feedback resistor but not of the integration capacitor. This is because the integration capacitor is an open circuit to dc. From Equation 1b it's clear that the most practical way to trim the full-scale voltage is to include a trim potentiometer in series with R_{INT}. Typically a 30% trim range will be required to absorb errors associated with t_{OS} and α.

It is also important to characterize the transient response of the integrator in order to determine settling time of the F/V to a step change of input frequency. The transfer function of the lossy integrator is given in the frequency domain by:

$$\frac{V_{OUT}(S)}{I_{IN}(S)} = \frac{\frac{1}{C_{INT}}}{S + \frac{1}{R_{INT} \cdot C_{INT}}} \qquad (2)$$

which indicates that the natural or step response to a change of input frequency is governed by an exponential function with time constant $\tau = R_{INT} \cdot C_{INT}$.

With the average output voltage and transient response known, the peak-peak output ripple can be determined using Equation 3. Once this is determined, a design algorithm can be developed. The derivation of Equation 3 is provided in Appendix A.[1]

The peak-peak ripple is given by:

$$V_{PP} = \frac{e^{t_{OS}/RC} - e^{T/RC} + e^{(T-t_{OS})/RC} - 1}{1 - e^{T/RC}} \cdot \alpha \cdot R \qquad (3)$$

where:

t_{OS}	= one-shot time constant	[seconds]
T	= period of input frequency ($1/f_{IN}$)	[hertz]
R	= integration resistor	[ohms]
C	= integration capacitor	[farads]
α	= current source value (1mA for AD650)	[amps]

Equation 3 accurately represents the ripple amplitude for a given design. The following section shows how this equation is used as an iterative part of the total solution. Equation 3 can also be used to illustrate how the ripple amplitude changes as a function of input frequency. It is interesting to note that the ripple amplitude changes only moderately with input frequency and has its largest magnitude at the minimum frequency.

DESIGN PROCEDURE

Recall from looking at Figure 3, the one-shot "on" time will be some fraction of the total input period. This is the time that the circuit integrates the current signal α. The output ripple can be minimized by allowing the current source to be on during the majority of this period. This is achieved by choosing the one-shot time constant so that it occupies almost the full period of the input signal when this period is at its minimum (or the input frequency at its maximum). To design safely and allow for component tolerance at f_{max}, make t_{OS} approximately equal to 90% of the minimum period. Given t_{OS}, the value of the one-shot timing capacitor, C_{OS}, is determined from Equation 1 in the AD650 data sheet. This equation has been rearranged and appears here as:

$$C_{OS} = \frac{t_{OS} - 3 \cdot 10^{-7}\ sec}{6.8 \cdot 10^{3}\ sec/F} \qquad (4)$$

[NOTE: For maximum linearity performance use a low dielectric absorption capacitor for C_{OS}.]

where t_{OS} is in seconds and C_{OS} is in farads.

Once C_{OS} is known, the integration resistor is uniquely determined from the full-scale equation (Equation 1b), since t_{OS}, α, f_{IN}, and V_{OUT} are known. This leaves only the integration capacitor as the final unknown.

C_{INT} is chosen by first determining the response time of the device being measured. If, for example the frequency signal to be measured is derived from a mechanical device such as an aircraft turbine shaft, the momentum of the shaft and the blades should be used to determine the response time. The time constant of the F/V is then set to match the time constant of the mechanical system. It may be set somewhat lower depending on the desired total response time of the mechanical and electrical system. Remember to allow several time constants (N) for the F/V to approach its final value. For the first iteration of C_{INT} use the following expression:

$$C_{INT} = \frac{Mechanical\ Response\ Time}{N \cdot R_{INT}} \qquad (5)$$

where N is the number of time constants chosen to allow adequate settling. Table I may be used to determine the number of time constants required for given settling accuracy.

Table I. Settling Accuracy vs. Number of Time Constants

# of Time Constants (N)	# of Bits	% Accuracy
4.16	6	1.6
4.85	7	0.8
5.55	8	0.4
6.23	9	0.2
6.93	10	0.1
7.62	11	0.05
8.30	12	0.024
9.00	13	0.012
9.70	14	0.006
10.4	15	0.003
11.0	16	0.0015

[1]The essence of the solution was captured by McGillem and Cooper in [1].

A larger number of time constants will give a more responsive circuit but will also increase the ripple at the F/V output. A practical approach is to start with 8 bit settling accuracy using N = 6 time constants and increase or decrease N depending on ripple content.

The ripple content is calculated using Equation 3. Remember that the ripple amplitude will change with frequency and will be largest at the lowest frequency. It is also important to note that while in some cases the ripple amplitude may be large, the *average* value of the output voltage will always represent the input frequency (unless the ripple gets too close and "clips" at the positive supply rail). Figure 4 shows an example of how output ripple amplitude will change with input frequency for a typical application. This graph was obtained by plotting Equation 3 over the full range of input

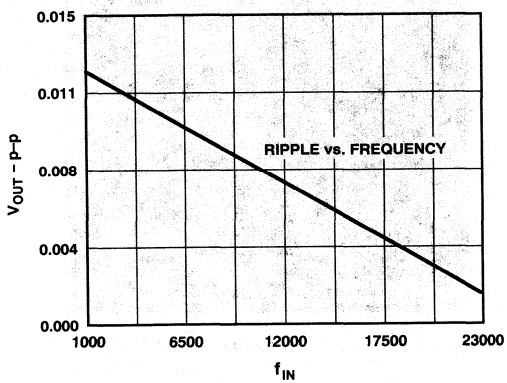

Figure 4. P-P Ripple vs. Frequency (See Design Example)

frequencies. For design purposes it is only necessary to calculate ripple at the worst case frequency (f_{min}). Figure 5 summarizes the design procedure.

DESIGN EXAMPLE

The rpm of an automobiles engine is to be monitored for use by an on-board computer. The rpm signal which will be generated from an F/V converter is to be digitized with an 8-bit A/D converter. The rpm range of the engine extends from 300 rpm to 7000 rpm. A 200 tooth flywheel at these rotational speeds will generate pulses from 1 kHz up to 23 kHz. The response time to a step change in throttle position of the engine has been measured, in neutral, to be 400 ms. The goal is to design an F/V converter that will respond at approximately the same rate as the engine or faster and will have ripple that is undetectable by the A/D converter. The A/D converter has a ten volt full scale.

1. Let t_{OS} be $0.9 \cdot \frac{1}{f_{max}} = 0.9 \times 43.5\ \mu s = \mathbf{39\ \mu s}$.
2. Find $C_{OS} = \mathbf{0.0057}\ \mu F$ (from Equation 4)
   (an impractical value for polystyrene, but tantalum may be used with reduced linearity).
3. $R_{INT} = \frac{10\ V}{1\ mA \cdot 39\ \mu s \cdot 23\ kHz} = \mathbf{11.14\ k\Omega}$ (from Equation 1b).

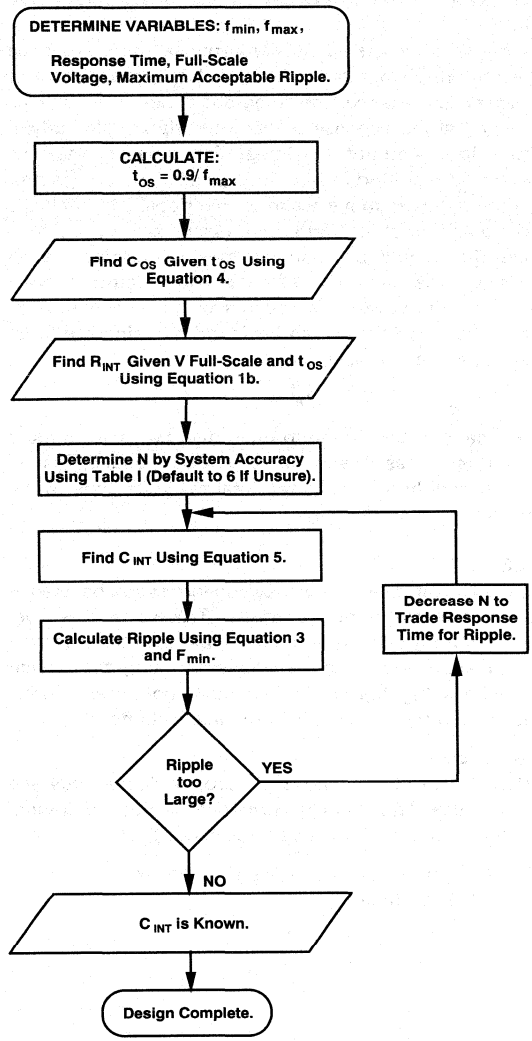

Figure 5. Design Flowchart

(If R_{INT}, the load seen by the amplifier, is less than 1 kΩ, then t_{OS} must be reevaluated.)

4. From Table I an RC network can settle to 8 bits in 6 time constants so: $C_{INT} = \frac{400\ ms}{(6)\ 11.14\ k\Omega} = \mathbf{6\ \mu F}$.
5. Ripple = 6.25 mV @ 300 rpm and 0.67 mV @ 7000 rpm (from Equation 3).
6. 1/2 LSB size for an 8-bit converter with 10 V full scale is 19.5 mV. Fortunately the ripple is below the quantization level on the first iteration. If desired, the integration capacitor may be lowered to reduce response time of the F/V converter.
7. Guessing $C_{INT} = 3.0\ \mu F$ or using an iterative computer program gives a maximum ripple content of 12.5 mV and a response time of 200 ms.

THE TRADE-OFF BETWEEN RIPPLE AND RESPONSE TIME

In many instances some compromise must be made between ripple and response time. If response time is of primary importance, the integration capacitor may be lowered at the expense of increased ripple. Similarly, if ripple is paramount, the integration capacitor must be increased resulting in slower response. The design procedure outlined above assumes that ripple content is the less desirable effect. Rather than increasing C_{INT}, a low-pass filter could be used, but this also slows the response time. An approximation to determine total response time of two cascaded systems, each with separate response times, can be found by using the "root sum of squares" technique.

$$T_{TOTAL} = \sqrt{T_A^2 + T_B^2}$$

This leads to the "three to one" rule, i.e., if T_A is more than three times T_B and their squares are added, T_B may be ignored, hence, T_A is the total response time of the system.

SUMMARY

Low cost voltage-to-frequency converters can be used in the frequency-to-voltage mode. Trade-offs exist between output settling time and ripple with the selection being application specific. However using the design guidelines highlighted in this note, optimized performance can be achieved in many applications.

References

[1] McGillem, C. D., and G. R. Cooper, *Continuous and Discrete Signal and System Analysis,* Second Edition. New York: CBS College Publishing, 1984.

[2] Analog Devices Inc. *AD650 Data Sheet* C795b-10-10/86

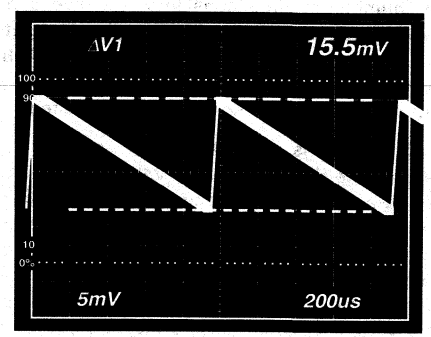

Figure 6. Typical Ripple Output

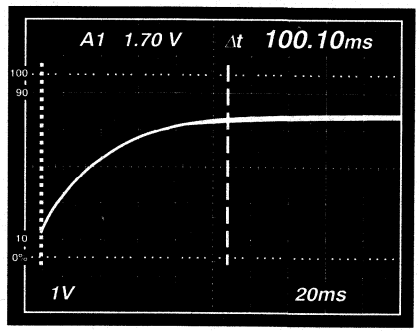

Figure 7. Response to Step Change in Frequency

NOTE

This circuit of Figure 8 should be used in place of the one shown in Figure 6 of the AD650 data sheet if the input signal has slower edges than the one-shot time delay. In such cases the comparator has a tendency to "double trip" possibly giving inaccurate results. This new circuit uses the open collector output to pull down on the input signal until the one-shot period occurs. After this time it gets pulled up by the +5 V supply. This circuit may not be needed for all applications since fast edges such as a TTL signal will give accurate results with the previous configuration. This circuit has the disadvantage of requiring a +5 V supply and increased power consumption.

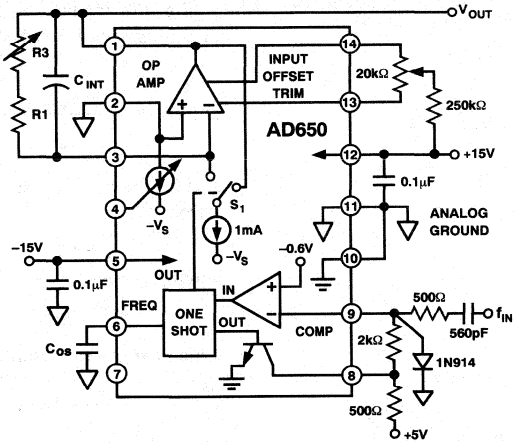

Figure 8. Circuit Showing Modified Input Structure

APPENDIX A: DERIVATION OF RIPPLE EQUATION

From Figure 2, one period I(t) is given by:

$$I_T(t) = \alpha[u(t) - u(t-t_{OS})] \text{ where } \alpha = 1\ mA$$

$$L\{I_T(t)\} = \alpha\left[\frac{1}{S} - \frac{1}{S}e^{-St_{OS}}\right] = \frac{\alpha}{S}(1 - e^{-St_{OS}})$$

Using the rule for repeated cycles gives:

$$I_P(S) = I_T(S)\frac{1}{1 - e^{-TS}} \text{ so } I_P(S) = I(S) = \frac{\alpha}{S}\frac{[1 - e^{-St_{OS}}]}{[1 - e^{-TS}]}$$

From Figure 2 we have:

$$i(t) = \frac{V_{OUT}(t)}{R} + C\frac{dV_{OUT}(t)}{dt} \text{ which gives:}$$

$$I(s) = V_{OUT}(S)\left[\frac{1}{R} + SC\right]$$

$$\text{so } \frac{V_{OUT}(S)}{I(S)} = \frac{\frac{1}{C}}{\left[S + \frac{1}{RC}\right]}$$

but using I(S) from above gives:

$$V_{OUT}(S) = V_{TOTAL}(S) = \frac{1}{S}\frac{\frac{\alpha}{C}}{\left(S + \frac{1}{RC}\right)}\frac{[1 - e^{-St_{OS}}]}{[1 - e^{-TS}]}$$

The transient portion of the solution results from evaluating this expression at the frequency $S = -\frac{1}{RC}$ which gives the residue:

$$V_{TRANS}(S) = \frac{1}{\left(-\frac{1}{RC}\right)}\frac{[1 - e^{t_{OS}/RC}]}{[1 - e^{T/RC}]}\frac{\alpha}{C}\frac{1}{\left(S + \frac{1}{RC}\right)}$$

which can be called $\frac{\beta}{S + \frac{1}{RC}}$ for simplicity

where $\beta = \frac{\alpha}{\left(-\frac{1}{R}\right)}\frac{[1 - e^{t_{OS}/RC}]}{[1 - e^{T/RC}]}$

and so $V_{TRAN}(t) = \beta e^{-t/RC}\,u(t)$

The total response of the system is the sum of the transient response and the steady state response:

or $V_{TOTAL}(S) = V_{TRAN}(S) + V_{SS}(S)$ and so:

$V_{SS}(S) = V_{TOTAL}(S) - V_{TRAN}(S)$

using V_{TOTAL} and V_{TRAN} from above gives:

$$V_{SS}(S) = \frac{\alpha}{CS}\frac{[1 - e^{-St_{OS}}]}{\left(S + \frac{1}{RC}\right)[1 - e^{-TS}]} - \frac{\beta}{S + \frac{1}{RC}}$$

which can be massaged into $\frac{\frac{\alpha}{C} - \frac{\alpha}{C}e^{-St_{OS}} - \beta S + \beta s e^{TS}}{S\left(S + \frac{1}{RC}\right)(1 - e^{TS})}$

If we are interested in just one period of the steady state solution we can drop all periods but the first by simply removing all terms containing e^{TS} (again this goes back to the rule for laplace transform of a periodic function given just one period),

$$\text{so: } V_{SS}(t) = \frac{\frac{\alpha}{C} - \frac{\alpha}{C}e^{-St_{OS}} - \beta S}{s\left(S + \frac{1}{RC}\right)}$$

and for the solution as it exists before t_{OS} drop all $e^{-St_{OS}}$ terms

$$\text{leaving: } \frac{\frac{\alpha}{C} - \beta S}{S\left(S + \frac{1}{RC}\right)}$$

which can be represented as:

$$\frac{\alpha R}{S} - \frac{\alpha R + \beta}{\left(S + \frac{1}{RC}\right)} \text{ by using partial fraction expansion.}$$

So now $V_{SS}(t) = [\alpha R - (\alpha R + \beta)e^{-t/RC}]\,u(t)$ for $0 < t < t_{OS}$. $V_{SS}(t)$ for the second half of the period ($t_{OS} < t < T$) can be found by subtracting all terms containing $e^{-St_{OS}}$ from the solution before t_{OS}.

$$V_{SS}(t) = [\alpha R - (\alpha R + \beta)\,e^{-t/RC}]\,u(t) - L^{-1}\left\{\frac{\frac{\alpha}{C}e^{-St_{OS}}}{S\left(S + \frac{1}{RC}\right)}\right\}$$

Again using partial fraction expansion and manipulating gives:

$$V_{SS}(t) = \left[1 - \frac{\alpha R + \beta}{e^{t_{OS/RC}}}\right]e^{-(t - t_{OS})/RC}\ t_{OS} < t < T.$$

The lowest point on the steady state wave can be found by evaluating:

$V_{SS}(t) \mid_{0<t<t_{OS}}$ @ t = 0 which gives $V_{LOW} = -\beta$.

Likewise the highest value of the wave can be found at t_{OS} which gives:

$V_{HIGH} = \alpha R - (\alpha R + \beta)e^{-t_{OS}/RC}$ and substituting β gives:

$$V_{HIGH} = \alpha R - \alpha R\left[1 - \frac{[1 - e^{t_{OS}/RC}]}{[1 - e^{T/RC}]}\right]e^{-t_{OS}/RC}.$$

Finally the p–p ripple is found by subtracting V_{LOW} from V_{HIGH}. After more manipulation Equation 3 results.

Miscellaneous
Contents

		Page
AN-345	Grounding for Low and High Frequency Circuits	24–3
AN-214	Ground Rules for High Speed Circuits	24–7
AN-346	Understanding Interference-Type Noise	24–11
AN-347	Shielding and Guarding	24–15
AN-348	Avoiding Passive-Component Pitfalls	24–21
AN-349	Keys to Longer Life for CMOS	24–27
AN-311	How to Reliably Protect CMOS Circuits Against Power Supply Overvoltaging	24–29
AN-350	Take the Guesswork Out of Settling-Time Measurements	24–33
AN-351	Ask the Applications Engineer—2: Offset and Gain Trimming	24–43
AN-352	Ask the Applications Engineer—5: High Speed Comparators	24–45
AN-362	Ask the Applications Engineer—9: Decoupling; Component Selection	24–47
AN-353	Ask the Applications Engineer—10: PC Board Layout	24–49
AN-354	Ask the Applications Engineer—1: ADC Problems; Log Compensation Resistors	24–51
AN-280	Mixed Signal Circuit Techniques	24–53
AN-282	Fundamentals of Sampled Data Systems	24–85

ANALOG DEVICES

AN-345
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Grounding for Low- and High-Frequency Circuits

Know Your Ground and Signal Paths for Effective Designs

Current Flow Seeks Path of Least Impedance — Not Just Resistance

by Paul Brokaw and Jeff Barrow

Noise reduction is a significant design issue in most electronic systems. Along with dissipation constraints, ambient temperature changes, size limitations, and speed & accuracy requirements, noise is an omnipresent factor that must be dealt with for a successful final design. We are not concerned here with techniques for reducing *external noise* (which arrives with the signal); since its presence is generally beyond the direct control of the design engineer; it must be dealt with in the operational design of the system by means such as filtering, analog signal processing, and digital algorithms.

In contrast, preventing *internal noise* (noise generated or coupled within the circuit or system) from corrupting the signal is a direct responsibility of the design engineer. Noise sources, if not fully considered *early in the design cycle*, can adversely affect final performance and prevent the high-resolution potential of a system from being realized; at the very least, costly redesign and rework may be required. Some of the design factors that relate noise to system behavior have been treated in earlier articles in these pages[1,2,3,4,5]. Here, we consider the major role that schematic, topology, and final layout of the system "ground" play in minimizing the coupling of internally generated noise.

To deal adequately with noise, we need several perspectives: the actual internal pin connections of a component, versus the conceptual ones; the proposed schematic for ground-referenced signals; and the effects of layout on noise generation and pickup. These subjects divide into two overlapping domains, depending on bandwidth of the noise phenomena; ground-noise sources, problems, and solutions differ at low and high frequencies. Fortunately, good grounding practices in one band are generally compatible with those in the other.

[1]"Noise and Operational Amplifier Circuits," by D. H. Sheingold and L. R. Smith, *Analog Dialogue* 3-1 (1969).
[2]"Understanding Interference-Type Noise," by Alan Rich, *Analog Dialogue* 16-3 (1982).
[3]"Shielding and Guarding," by Alan Rich, *Analog Dialogue* 17-1 (1983).
[4]"Ground Rules for High-Speed Circuits," by Don Brockman and Arnold Williams, *Analog Dialogue* 17-3 (1983).
[5]"Amplifier Noise Basics Revisited," by Al Ryan and Tim Scranton, *Analog Dialogue* 18-1 (1984).

Reprinted from Analog Dialogue 23–3 1989

BASIC OP-AMP INTERCONNECTIONS

Many discussions of op amps present the ideal op amp as a three-terminal device with a pair of differential inputs and a single output (Figure 1). But the output voltage has to be measured with respect to some reference point, and output current from the amplifier must find a closed circuit back to the amplifier. The

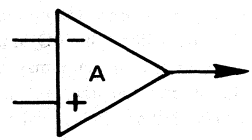

Figure 1. Conventional "three terminal" op amp.

infinite common-mode rejection of the ideal differential op amp disengages the input and output reference potentials, and the high input impedance eliminates the possibility of using an input terminal as an output current return; so there must be a fourth terminal, which some call "ground."

Of course, most IC op amps don't have a "ground" connection; the fourth terminal is generally considered to be the common connection of a dual power supply (which may also be serving other amplifiers and system elements). While it indeed serves this function at low frequencies, it will continue to do so only as long as the supply connections actually present the amplifier with a low (ideally zero) impedance at all frequencies within the amplifier bandwidth. When this requirement is not met, the impedance at the supply terminals affects the signal path and a wide variety of problems will arise, including noise, poor transient response, and oscillation.

An op amp must accept a fully differential signal and convert this to a single-ended output, with respect to the fourth terminal. Figure 2 shows the actual signal flow used by several basic and popular op amp families. Most of the voltage difference between the amplifier output and the negative rail appears across the compensating capacitor of the integrator (which controls the open-loop frequency response); if the negative supply voltage changes abruptly, the output of the integrator amplifier will immediately follow its "+" input. With the op amp in a typical closed-loop configuration, the input error signal tends to restore the output, with recovery limited by the integrator bandwidth.

This type of amplifier may have excellent low-frequency power-supply rejection, but negative supply rejection is limited at higher frequencies. Since the amplifier's gain is what causes the output to be restored, the negative supply rejection approaches zero for signals above the closed-loop bandwidth. The result: high-speed, high-level circuits can interact with the low-level circuits through the common impedance of the negative supply line.

Decoupling often is the recommended solution, but there are many wrong and some better ways. A decoupling capacitor *near the*

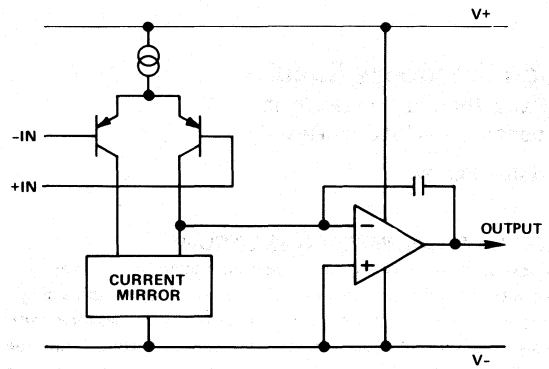

Figure 2. Simplified "real" op amp.

power supply may be separated from the op amp by many centimeters of wire, which looks like a high-Q inductor. Placing the decoupling capacitor near the op amp may still not solve the problem since, for decoupling, the other end of the capacitor must be connected to that mystical somewhere called "ground."

Figure 3* shows how a decoupling capacitor is connected to to minimize disturbances between the negative rail and ground buses. The load current's high-frequency component is confined to a path that doesn't include any part of the ground path. As an example of a more complex case, in Figure 4, the op amp is driving a load that goes to virtual ground (input of the second amplifier) and actual load current does not return to ground. Instead, it must be supplied by the second amplifier via its positive supply. Decoupling the negative supply of the first amplifier to the positive supply of the second one closes the high-speed signal current loop without affecting ground or signal paths.

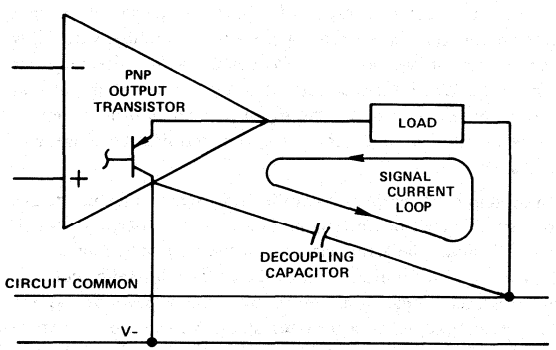

Figure 3. Decoupling of negative supply for a grounded load.

*Many of these illustrations can be found in the free Application Note, "An I. C. Amplifier User's Guide to Decoupling, Grounding, and Making Things Go Right for a Change," by Paul Brokaw.

Allowing ground currents to share a path with a low-level signal can cause problems. Figure 5 shows how careless grounding can degrade the performance of an amplifier driving a load resistor. The load current is supplied by the power supply and controlled by the amplifier. If points A and B are power supply "ground" connections, connecting the supply at A causes the load current to share a segment of wire with the input signal connections.

For example, fifteen centimeters of number 22 wire present about 8 mΩ of resistance to the load current. For a 2-kΩ load, a 10-volt

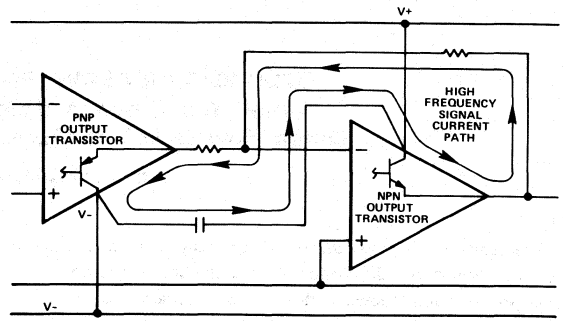

Figure 4. Decoupling of negative supply for "virtual ground" load.

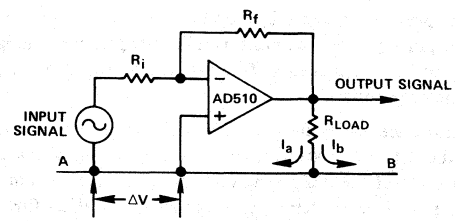

Figure 5. Proper choice of power connections will minimize this problem.

output swing results in about 40 μV between the points marked ΔV. This signal is in series with the non-inverting input and can result in significant errors: for an amplifier with a gain of 8 million, this positive feedback of 1/250,000 introduces a gain error factor 32× worse than that associated with the amplifier's open-loop gain alone. In addition, the positive feedback can cause circuit latchup or oscillation for large closed-loop gains (typically >250 V/mV). But the common feedback impedance can be eliminated by connecting the power supply to point B.

In a real system, the situation is more complicated. The input signal source, shown as floating in Figure 5, may also produce a current which must be returned to the power supply. With the supply's return at B, any current flowing in additional loads other than R_i may interfere with this amplifier's operation. Where amplifiers are cascaded, the scheme in Figure 6 shows how they can still drive auxiliary loads without common-impedance feedback coupling. Output currents flow through auxiliary loads and back to the power supply through the power common. Bypasses are connected as shown in Figure 4 so that currents in the input and feedback resistors are supplied from the power supply via the amplifiers. Only amplifier input current flows in signal common; its effect is usually small enough to ignore.

Understanding where the actual load and signal currents go is essential. The key to optimizing the circuit is to bypass these

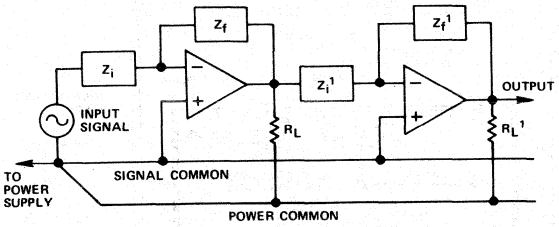

Figure 6. Minimizing common impedance coupling.

currents around ground and other signal paths. The voltage —more accurately called the potential difference—between two points defines such a current flow.

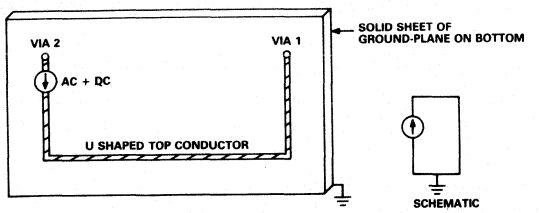

Figure 7. Schematic and layout of current source with U-shaped trace on pc board and return through ground plane.

GROUNDING FOR HIGH-FREQUENCY OPERATION

The "ground-plane" layer is often advocated as the best return for power and signal current, while providing a reference node for converters, references, and other subcircuits. However, even extensive use of a ground plane does not ensure a high quality ground reference for ac circuits.

The simple circuit shown in Figure 7, built on a two-layer printed circuit board, has an ac and dc current source on the top layer connected to via 1 at one end and to a single U-shaped copper trace connected to via 2. Both vias go through the circuit board and connect to the ground plane. Ideally, impedance is zero and the voltage appearing across the current source should be zero.

The simple schematic hardly begins to show the actual subtleties. But an understanding of how the current flows in the ground plane from via 1 to via 2 makes the realities apparent and shows how ground noise in high-frequency layouts can be avoided.

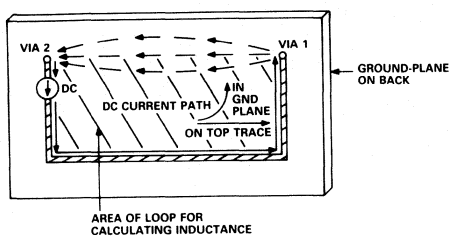

Figure 8. DC current path for Figure 7.

The dc current flows in the manner of Figure 8, as one might surmise, taking the path of least resistance from via 1 to via 2. Some current spreading occurs, but little current flows at substantial distance from this path. In contrast, the ac current does not take the path of least resistance; it takes the path of least *impedance*, which in turn depends on *inductance*.

Inductance is proportional to the area of the loop made by the current flow; the relationship can be illustrated by the right-hand rule and magnetic fields shown in Figure 9. Inside the loop, current along all parts of the loop produces magnetic field lines that add constructively. Away from the loop, however, field lines from different parts add destructively; thus the field is confined principally within the loop. A larger loop has greater inductance; this means that, for a given current level, it has more stored magnetic energy (Li^2), *greater impedance—since* $X_L = jwL$, and hence will develop more voltage at a given frequency.

Which path will the current choose in the ground plane? Naturally the lowest-impedance path. Considering the loop formed by the U-shaped surface lead and the ground plane, and neglecting resistance, high-frequency ac current will follow the path with the least inductance, hence the least area.

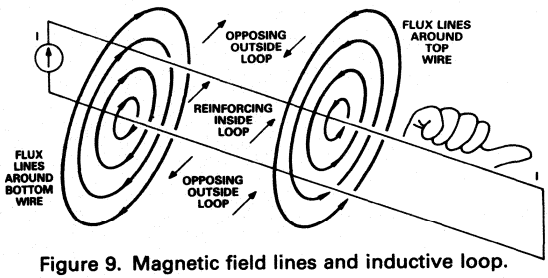

Figure 9. Magnetic field lines and inductive loop.

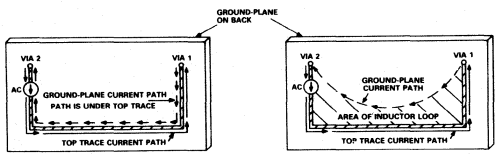

Figure 10. AC current path without (left) and with (right) resistance in ground plane.

In the simple example shown, the loop with the least area is quite evidently formed by the U-shaped top trace and the portion of the ground plane *directly underneath it*. So, while Figure 8 shows the dc current path, Figure 10 (left) shows that the path that most of the ac current takes in the ground plane, where it finds minimum area, directly under the U-shaped top conductor. In practice, the resistance in the ground plane causes the current flow at low- and mid-frequencies to be somewhere between straight back and directly under the top conductor (right). However, the return path is nearly under the top trace even at frequencies as low as 1 or 2 MHz.

Avoiding Layout Problems Once the return current paths in the ground-plane are understood, common layout trouble spots can be identified and corrected. For example, in Figure 11, path A is identified as critical, to be kept short, away from digital lines, and free of vias. Path B is of lesser importance, but needs to cross A. Typically, the ground plane is cut under A, and B is then routed through two vias and under A.

The unfortunate result is that inductance is introduced into the ground returns of both signals, because the interrupted ground plane makes both return loops larger. Since path A conducts a high-frequency signal, an induced voltage drop will appear across the opening of the ground plane. For typical ECL or TTL signals, this drop can be greater than several hundred millivolts, enough

to seriously compromise the performance of a 12-bit, 10-MHz converter or an 8-bit, 20-MHz unit. A simple fix is to install a wire directly across the cut in the ground plane to keep the loop area small.

Power distribution is another area of concern. Power supply lines must be kept at lowest possible characteristic impedance ($\sqrt{L/C}$). To keep this ratio small, inductance is reduced and capacitance increased by maintaining ground planes under the supply lines; capacitance can be further increased by selectively placing bypass capacitors at critical locations, as discussed earlier. If only capacitance is dealt with by, for example, placing 0.1-μF capacitors on supply pins to lower their impedance, a supply line with 30-nH inductance will have damped oscillations at about 3 MHz after every transient.

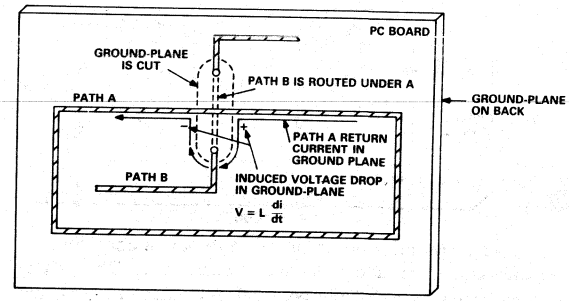

Figure 11. Typical PC layout problem, with paths crossing.

ANALOG DEVICES

AN-214
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Ground Rules for High-Speed Circuits

Layout and Wiring Are Critical in Video-Converter Circuits

How to Keep Interference to a Minimum

by Don Brockman and Arnold Williams

In recent issues, **Analog Dialogue** *has dealt extensively with topics in shielding and grounding,[1, 2] emphasizing the techniques needed to protect the integrity and precision of analog signals in the dc and audio-frequency domain from interfering signals, whether at line frequency or at much higher frequencies. To complement those articles, we suggest here the elements of good practice for high-resolution "video speed" converters, i.e., converters of 10-bit or greater resolution, operating at word rates above 1 MHz.*

Electronics may be frustrating for designers who cross the threshold from low-resolution-low-speed to high-resolution-high-speed designs, or from digital to analog-signal-conditioning circuits. For them, it often seems the "ground rules" have changed.

Experienced designers can readily attest to the difficulty of obtaining consistent grounds. They can relate stories about the ground that wasn't where they thought it was, or the ground that wasn't there at all, despite a conviction that "it has to be." On printed-circuit (p-c) boards, wires and/or runs that seemed to be perfectly good grounds are transformed into inductors or worse in high-speed or high-frequency circuits.

At ADI's Computer Labs Division, where high-speed circuits are its bread and butter, applications engineers have found that grounding is the focus of a large percentage of questions from designers making their initial foray into high-speed circuits. In most cases, the designers encountered difficulties as the result of being unaware of—or ignoring—certain basic ground rules.

BASIC PC-CARD RULES

Knowledgeable high-speed circuit designers have learned that *every* square inch of a printed-circuit board which doesn't contain circuits or conducting runs should be ground plane. Violating that simple rule invites disaster. But sometimes, strict adherence to the rule is still no guarantee of success if circuit density is too high; then one must reduce the density and create more "real estate" for the ground plane.

Our applications engineers strongly recommend that all bread-board designs be done on double-sided copper-clad boards. Although this is not a sure cure for ground problems, it improves the designer's chances.

Another basic rule for working with high-speed and/or high-frequency printed-circuit-board designs is to connect analog ground and digital ground together within the PC board. Connecting the two grounds enhances the performance of the converters when they are operated either by themselves or as tightly knit subsystems. However, it can raise some system-level problems, to be discussed below.

Another rule for printed-circuit-board designs containing analog and digital circuitry is to use every available spare pin for making ground connections, and to use those pins to separate the analog and digital signals entering or leaving the board.

Avoid using purely insulating (e.g., "Vector") bread-boards and small-diameter hookup wire (e.g., #24) for connections, including supply voltages and grounds. The approach will create ground and noise problems if the circuit is intended to operate at 1 MHz or more (it will probably lead to problems at even slower speeds).

To summarize: Use double-sided copper-clad boards with maximum ground area and heavy, well-located power-supply and ground-return leads. Tie rounds together locally.

GENERAL CIRCUIT PRACTICE

Any subsystem or circuit layout operating at high speeds with both analog and digital signals needs to have those signals physically separated as much as possible to prevent possible crosstalk between the two. Digital signals leaving or entering the layout should use runs that have minimum length. The shorter the digital runs, the lower the likelihood of coupling to the analog circuits.

Analog signals should be routed as far from digital signals as size constraints allow; and the two, ideally should never closely parallel one another's paths. If they

[1]Alan Rich, "Understanding Interference-Type Noise," *Analog Dialogue* 16-3, 1982, pages 16-19.
[2]Alan Rich, "Shielding and Guarding," *Analog Dialogue* 17-1, 1983, pages 8-13.

must cross, they should do so at right angles to minimize interference. Coaxial cables may be necessary for analog inputs or outputs—a demanding condition mechanically, but sometimes the only solution electrically.

When combining track-and-hold and a/d-converter hybrids or modules on the same board, keep them as close together as is practical. All grounds need to be connected to the single, low-impedance ground plane: and the connections should be made right at the units themselves (another argument for having large amounts of good, solid ground plane available all over the p-c board).

A suggested practical approach for accomplishing this is illustrated in Figure 1, which shows a flow-chart layout, as the preferred method for combining high-speed analog and digital circuits on a p-c board.

If one assumes a 10-volt input range on the 12-bit a/d converter, the least-significant bit (LSB) of the ADC will have a value of 2.5 mV (10 V/4,096). Assume that a single pin of the p-c connector, which is used for ground, has a resistance of 0.05 ohm—and that the p-c card draws a total of 1.5 amperes.

The voltage drop at the ground pin could be 75 millivolts in these circumstances. If only digital logic were used, this voltage drop would be minuscule, hardly worth considering. However, the hypothetical real-world situation being considered here is a mixture of both analog and digital circuits, and the 75 mV can have a significant impact on the subsystem's performance.

In this example, the digital circuits are TTL. Since TTL is a saturated logic, ground currents vary widely, and varying current flowing through the ground often produces noise signals which modulate the ground plane. This noise, created by digital switching, can couple into the analog portion of the circuit and have an important effect on performance, even at low digital levels. For example, if only 10% of the 75-millivolt I-R drop cited here couples into the analog signal, that would represent 3 LSBs.

The result? The circuit intended for operation as a 12-bit system is now reduced to a system of 10 to 11 bits, because of noise masking the 2.5-millivolt level of the desired 12-bit LSB. The recommended solution? Assign multiple pins for ground connections, to reduce the total contact resistance. As Figure 1 shows, those pins are also used to separate the analog and digital signals.

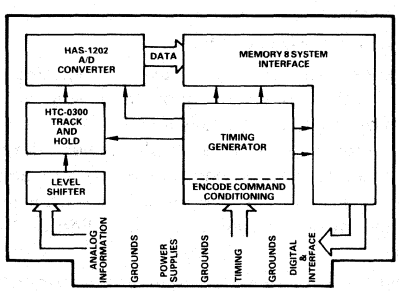

Figure 1. "Flow Chart" Layout for Logical Separation of Functions.

This design approach may seem unnecessarily rigorous and time-consuming but can prove rewarding when the p-c board is installed in its final system location.

Locate the timing circuits near the center of the board (Figure 1) because the timing is at the heart of the circuit, being connected to all of the major circuit components of the board. A central location helps assure minimum paths for the digital signals.

Variations of this theme may not use the exact same components or functions, but the same basic techniques should be applied in any design containing analog and digital circuits. For cards with all connections at one end, avoid configurations which have analog circuits near the p-c connector, and digital signals at the opposite end of the card—or vice versa; either situation will cause analog and digital paths to pass in close proximity to one another.

SYSTEM GROUNDING

Although local ties for analog and digital grounds help the performance of a card, they can cause problems for the system designer using ADCs and DACs. In systems, data converters should be considered as *analog* (not digital) components; the system design must be assigned to experienced and capable analog engineers, who are used to defending millivolt signals against interference.

Place ADCs and DACs (like other analog devices) near other parts of the analog section, because: (1) reflections make it hard to transmit analog signals more than a short distance without loss of bandwidth and amplitude; and (2) noise generated by the digital section can couple into the analog through the ground plane or power supplies, or radiate to nearby analog components.

Each card in the system should be returned directly to the power supply common, using heavy wire. Where it is mandatory that a card's analog and digital grounds be separated, each should be separately returned to the power supply; don't connect the two grounds and return a single ground line to the power supply.

POWER SUPPLIES

Besides ground rules, designers of high-speed circuits must also consider the rules about power supplies to obtain best results.

Every power-supply line leading into a high-speed p-c card or data-acquisition circuit must be carefully bypassed to its ground return to prevent noise from entering the card. Ceramic capacitors, ranging in value from 0.01 to 0.1 μF, should be used generously in the layout, mounted as closely as possible to the device or circuit being bypassed; and at least one good-quality tantalum capacitor of 3 to 20 μF should be assigned to each power-supply voltage, mounted as near as possible to the incoming power pins to keep potentially high levels of low-frequency ripple off the card.

To some extent, the p-c's power-supply connector pins can introduce noise problems. If their contact resistance is sufficiently high, and a varying current is flowing, the varying IR drop which results is noise and can be cou-

led into parts of the card. This caution applies especially to +5-volt supplies used to power TTL systems, but the problem can be alleviated with a variation of the rule about multiple pins for making ground connections. Parallel the I-R drops by also using multiple pins for power connections.

Low-noise, low-ripple temperature-stable linear power supplies are the preferred choices for high-speed circuits. Switching power supplies often seem to meet those criteria, including ripple specifications. *But ripple specs are generally expressed in terms of rms*—and the spikes generated in switchers may often produce hard-to-filter, uncontrollable noise peaks with amplitudes of several hundred millivolts. Their high-frequency components may be extremely difficult to keep out of the ground system.

If switchers cannot be avoided for high-speed designs, they should be carefully shielded and located as far away from the "action" as possible, and their outputs should be filtered heavily.

ABOUT IC DESIGNS

There is often a difference in implementing designs using high-precision IC circuits vis-à-vis p-c card designs using modules or hybrids. Some ICs are specifically designed to keep analog and digital grounds separated within the device, because they would be unable to perform their functions properly without the separation.

Recognizing this, IC manufacturers are generally very careful in detailing how to obtain optimum performance from their devices. Those details of the application notes frequently instruct the user to connect analog and digital grounds for the device together externally; when they do, the connection needs to be made as closely as possible to the device. In other, much rarer, instances, the characteristics of an individual device—or system—may require some remote connection of the grounds.

The best approach for getting optimum performance from any device is to follow diligently the recommendations of the manufacturer. If the recommendations are missing or vague, ask for them.

Logical signal flow generates logical treatment of ground paths and ground connections— a logical way to prevent potential problems.

ANALOG DEVICES

AN-346
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Understanding Interference-Type Noise

How to Deal with Noise without Black Magic

There Are Rational Explanations for — and Solutions to — Noise Problems

by Alan Rich

If the circuit doesn't work, add a decoupling capacitor anywhere—a 0.01μF ceramic disc, of course; they'll fix anything! Or when your circuit is broadcasting its noise, a shield will cure it; just wrap a piece of metal around the circuit, connect that shield to "ground," and watch the noise disappear!

Unfortunately, Nature is not that kind to us in real life. That 0.01μF disc you added only increased the noise; and the shield you added was totally ineffective—or, worse yet, the noise reappeared in a remote part of the circuit.

This article is the first of a two-part series to help you understand and deal effectively with interference noise in electronic systems. We will consider here the mechanism that causes noise to be picked up, since the first step in solving any noise problem is to identify the source of the noise and the coupling mechanism; only then can an effective solution be implemented.

The second article will suggest specific techniques and guidelines for effective shielding against electrostatic and magnetically coupled noise.*

WHAT KIND OF NOISE ARE WE TALKING ABOUT?

Any electronic system contains many sources of noise. Three basic forms in which it appears are: *transmitted noise*, received with the original signal and indistinguishable from it, *intrinsic noise*, (such as thermally generated Johnson noise, shot noise, and popcorn noise) originating within the devices that constitute a circuit, and *interference noise*, picked up from outside the circuit. This last may either be due to natural disturbances (e.g., lightning) or be coupled in from other electrical apparatus in the system or its vicinity, for example computers, switching power supplies, SCR controlled heaters, radio transmitters, switch contacts, etc.

This article will consider only the last category, man-made noise, the most pervasive form of system noise in data-acquisition or test systems. Although it is most annoying in low-level circuits, no part of the system is immune to it. But it is the only form of noise that can be influenced by choices of wiring and shielding.

ASSUMPTIONS AND ANALYTICAL TOOLS

Although Maxwell's equations—with all the mathematical agony that they imply—are necessary for a complete and accurate description of how electrical systems behave, conventional circuit analysis is a useful tool in most cases. The assumptions that permit circuit analysis to be valid in solving these problems are:

1. All electric fields are confined to the interior of capacitors.
2. All magnetic fields are confined to the immediate vicinity of inductors.
3. Dimensions of the circuits are small compared to the wavelengths under consideration.

Using these assumptions, we can model noise-coupling channels as lumped circuit elements. A magnetic field coupling two conductors is modeled as a mutual inductance. Stray capacitance can be modeled as two conductors with an electric field between them. Figure 1 shows an equivalent circuit of a situation where two short wires are adjacent to one another over a system ground.

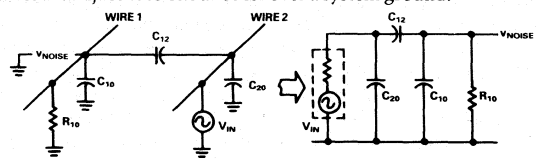

Figure 1. Noise-equivalent circuit of two adjacent wires and a ground plane.

Once the complete noise equivalent-circuit is obtained for a system, the problem becomes one of solving network equations for a desired parameter. All standard linear circuit analysis techniques can be applied, including node equations, loop equations, matrix algebra, state variables, superposition, Laplace transforms, etc. When circuits exceed 5 or 6 nodes, manual calculation becomes difficult; at this point, computer-aided programs, such as SPICE, and other CAD techniques become necessary. Experienced designers can make appropriate simplifying assumptions; but their validity should always remain in question until proven.

The lumped-element approach will not always give an accurate numerical answer, but it will show clearly how noise depends on system parameters. Just the act of drawing a reasonably faithful equivalent circuit may offer clues to methods to reduce noise levels. Once network equations or CAD programs are written, the quantitative effects of noise-suppression techniques can be studied.

In spite of all the modern technical advances, such as microprocessors and switching power supplies, wires still have resistance and inductance, capacitance still exists in the real world, and such phenomena must be reckoned with.

THE BASIC PRINCIPLE

There are always three elements involved in a noise problem: a *noise source* (line transients, relays, magnetic fields, etc.), a

*Another helpful and relevant article that appeared in these pages was "Analog Signal Handling for High Speed and Accuracy," by A. Paul Brokaw, *Analog Dialogue* 11-2, 1977, pp. 10–16.

Reprinted from Analog Dialogue 16–3 1982

coupling medium (capacitance, mutual inductance, wire), and a *receiver*, a circuit that is susceptible to the noise (Figure 2).

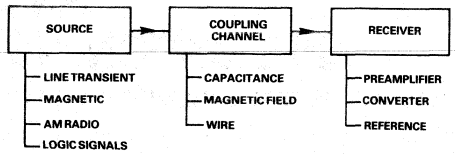

Figure 2. Noise pickup always involves a source, a coupling medium, and a receiver.

To solve the problem, one or more of these three elements must be removed, reduced, or diverted. Their role in the problem must be thoroughly understood before the problem can be solved. If the solution is inappropriate, it may only make the noise problem worse! Different noise problems require different solutions; adding a capacitor or a shield will not solve every such problem.

TYPES OF SYSTEM NOISE

Noise in any electronic system can originate at a large number of sources, including computers, fans, power supplies, adjacent equipment, test devices; noise sources can even include improperly connected shields and ground wires that were intended to combat noise. Our discussion of noise sources and coupling mechanisms will include the following topics:

Common-impedance noise
Capacitively coupled noise
Magnetically coupled noise
Power-line transients
Miscellaneous noise sources

Common-Impedance Noise. As the name implies, common-impedance noise is developed by an impedance that is common to several circuits. Figure 3 shows the basic configuration, which might occur when a pulse output source and an op amp's reference terminal are both connected to a "ground" point having tangible impedance to the power-supply return terminal. The noise current (the noisy return current of Circuit 1) will develop across impedance, Z, a voltage, V_{noise}, which will appear as a noise signal to Circuit 2.

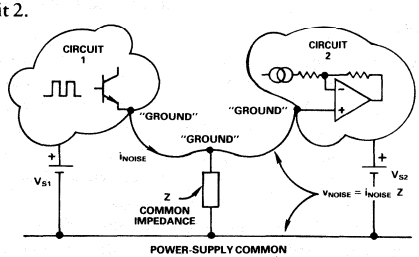

Figure 3. How noise is developed by a common circuit impedance.

Typically, this type of noise has a repetition rate that is set by the rate of the noise source. The actual waveshape is determined by the characteristics of the impedance, Z. For example, if Z is purely resistive, the noise voltage will be proportional to the noise current and of similar shape (Figure 4a). If Z is an R-L-C, the noise voltage will ring at a frequency, $1/(2\pi\sqrt{LC})$ and decay exponentially at a rate set by L/R (b).

If noise of this kind is found in a circuit, its origin may be readily deduced from the repetition rate and waveshape. The *repetition rate* will point to the source of noise, since the noise and its source are synchronized. For example, a noise waveform like that shown

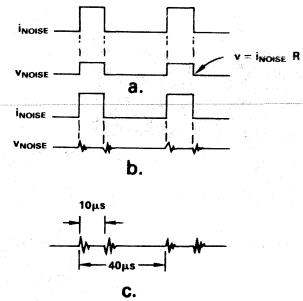

Figure 4. Noise effects in a common impedance. (a) Resistance. (b) An R-L-C circuit. (c) Switching-supply noise response.

in (c), at a 25kHz repetition rate and a 25% duty cycle, might be typical of a switching power supply containing a regulating loop using pulse-width modulation.

The *waveform* will help identify the impedance that is actually generating the undesired noise. If, for example, the waveform of the noise is the simple damped sinusoid shown in Figure 5, the following features allow us to deduce the nature of Z:

•A constant resistance, R, is in series with the line. The voltage change, V_1, is the product of R and a current step, I_1.

•The natural frequency of the oscillation, f_1, is determined by the series L and shunt C, $f = 1/(2\pi\sqrt{LC})$.

•The damping time constant, τ, is determined by L/R.

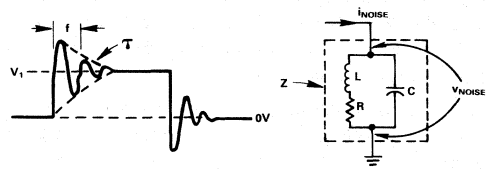

Figure 5 Waveshape for an underdamped R-L-C circuit.

Capacitively Coupled Noise. Noise is also produced by capacitive coupling from a noise source to another circuit. This type of noise is often seen when signals with fast rise-and-fall times or high frequency content are in close proximity to high-impedance circuits. Stray capacitance couples the fast edges of the signal into adjacent circuits, as the circuit model of Figure 6 shows. The nature of the impedance, Z, determines the shape of the response. Typical capacitances are listed in Table 1.

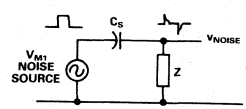

Figure 6. Stray capacitance couples noise into high-impedance circuits.

Capacitive pickup can occur in many ways, shapes, and sizes. Here are a few examples:

•A TTL digital signal produces fast edges, with a typical rise time of 10 nanoseconds and voltage swings of 5 volts. If Z is a 1-megohm resistor, even 0.1pF will produce 5-volt spikes with decay time constants of 100 nanoseconds.

•Crosstalk may result between two adjacent wires. For example, if two wires in a 10-foot (3-meter) length of cable have a capacitance of 40 pF/ft, the total capacitance is 400 pF. If a test voltage

f 10 V at 1 kHz is on one conductor, 250 mV at 1 kHz will be oupled into the adjacent wire if Z is a 10 k resistance.

•Noise on the ac power line, developed through common im-›edances, will couple into other circuits. A common case is when ransients couple through the interwinding capacitance of power-upply transformers.

Table 1. Typical capacitances.[1]

Condition	Capacitance
Human standing on an insulator to earth	700 pF
Power input (ac) to output (dc) of ± 15-V dc supply	100 pF
Two-conductor shielded cable:	
Conductor to conductor	40 pF/ft
Conductor to shield	65 pF/ft
RG58 coaxial cable, center conductor to shield	33 pF/ft
Connector, pin to pin	2 pF
Optical isolator, LED to photodetector	2 pF
½-watt resistor (end to end)	1.5 pF

It is amazing how little capacitance can cause serious problems. For example, consider the situation where high noise-immunity CMOS logic is used in an industrial circuit where 2500-volt, 1.5 MHz noise transients (IEEE Standard 472-1974) are present. Suppose that stray capacitance of only 0.1 pF exists between a CMOS input and the noise source, as shown in Figure 7. The calculated noise voltage, V_c, will be 2.4 volts, steady state, with an initial 50-V transient, which will cause improper logic operation or worse!

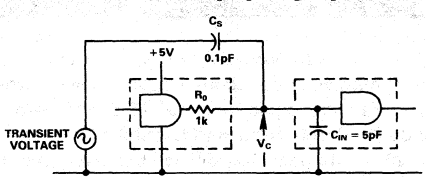

Figure 7. Coupling of high-voltage transients from test generator to logic.

Magnetically Coupled Noise. Strong magnetic fields are found where cables carry current, where ac power is distributed, and near machinery, power transformers, fans, etc. There is an analogous relationship between circuits coupled magnetically and those coupled capacitively, as shown in Figure 8 and Table 2.

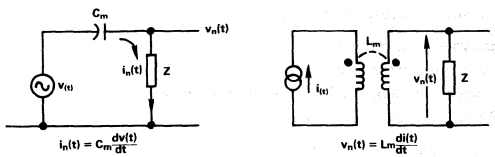

Figure 8. Comparison of magnetic and capacitive noise coupling.

Table 2. Characteristics of capacitive and magnetic coupling.

	Capacitive Coupling	Magnetic Coupling
Noise Source	Voltage change (dV/dt)	Current change (dI/dt)
Coupling Medium	Mutual capacitance	Mutual inductance
Coupled Noise	Current (frequently converted to voltage by Z)	Voltage

This analogy helps us consider some differences between capacitively and magnetically coupled noise:

•When the noise is magnetically coupled, voltage noise (V_n) appears in series with the receiver circuit; in the capacitive situation, the voltage noise produced between the receiver and ground is the voltage in Z caused by the noise current, i_n.

•Reducing the receiver impedance, Z, will reduce capacitively coupled noise. This is not the case in magnetically coupled circuits; lowering Z will not dramatically reduce voltage noise.

The voltage, V_n, induced in a closed loop (single turn) by a magnetic field is given by

$$V_n = 2\pi f B A \cos\theta \times 10^{-8} \quad (1)$$

volts, where f is the frequency of the sinusoidally varying flux density, B is the rms value of the flux density (gauss), A is the area of the closed loop (cm^2), and θ is the angle of B to area A.

For example, consider the circuit of Figure 9. It shows the calculation for two one-foot conductors, separated by 1 inch, in a 10-gauss 60-Hz magnetic field (typical of fans, power wiring, transformers). The maximum voltage induced in the wires is 3 mV.

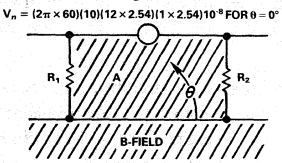

Figure 9. Example demonstrating magnitude of magnetic pickup.

The equation tells us that the noise voltage can be reduced by reducing B, A, or cosθ. The *B term* can be reduced by increasing the distance from the source of the field or—if the field is caused by currents flowing through nearby pairs of wires—twisting those wires to reduce the net field to zero by alternating its direction.

The loop area, A, can be reduced by placing the conductors closer together. For example, if the conductors in the example were placed 0.1" apart (separated only by insulation), the noise voltage would be reduced to 0.3mV. If they can be twisted together, the area is, in effect, reduced to small positive and negative increments that cancel, practically nullifying the magnetic pickup.

The cos θ term can be reduced by proper orientation of the receiving wires to the field. For example, if the conductors were perpendicular to the field, the pickup would be minimized, while if they were run together in the same cable (θ = 0), pickup would be maximized.

The rms induced voltage, V_n, in a conductor in parallel with a second conductor, carrying a current I_2 at an angular frequency $\omega = 2\pi f$, with a given mutual inductance, M, is

$$V_n = \omega M I_2 \quad (2)$$

The application of this relationship shown in Figure 10 illustrates why only one end of a shield should be grounded. A 100-ft length of shielded cable is used to carry a high-level low-impedance signal

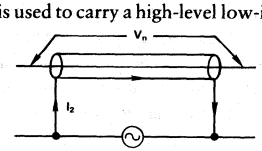

Figure 10. Magnetic pickup from current flowing through a cable shield.

(10V) to a 12-bit data-acquisition system (1 LSB = 2.4 mV). The shield, which has series resistance of 0.01 ohms per foot and mutual inductance to the conductor of 0.6μH/ft, has been grounded at both the source and the destination. A potential of 1 volt at

[1]Sources: Excerpts from Ralph Morrison, *Grounding and Shielding Techniques in Instrumentation*, Second Edition (New York: John Wiley & Sons, 1977), p.30, and actual measurements.

60 Hz exists between the two ground points, causing a current of 1 ampere to flow in the 1-ohm total resistance of the shield. By (2), the noise voltage induced in the conductor is

$$V_n = (2\pi\, 60\,\text{Hz})(100 \times 0.6 \times 10^{-6}\,\text{H})(1\,\text{A}) = 23\,\text{mV},$$

or 10 LSBs, thereby reducing the effective resolution of the system to less than 9 bits. This noise voltage is a direct consequence of the large current flowing in the shield because it is grounded at both ends. And the 1-volt potential assumed between the grounds was conservative! In heavy-industry environments, 10 to 50 volts between earth grounds is not uncommon.

Power-Line Transients. Another type of system noise is that generated by high-voltage transients in inductive circuits, such as relays, solenoids, and motors, when they are turned on and off. When devices having high self-inductance are turned off, the collapsing fields can generate transients of the order of kilovolts, with frequencies from 0.1 to 3 megahertz, that appear on the power line.

Besides creating noise in sensitive circuitry, via capacitive and conductive coupling and radiated energy, these transients are hazardous to equipment and people. Standards exist to characterize certain transient waveforms for the purpose of protection; however, besides being designed to withstand them, systems should also be designed to deal with their potential interference with signals. Figure 11 shows 4 typical waveforms existing in industry standards.

Miscellaneous Noise Sources Finally, there is a group of noise sources that can be considered as miscellaneous—or just "flakey."

For low-level signals at high impedance, the cable itself can become a noise source. A charge can be produced on the dielectric material within the cable; if the dielectric does not maintain contact with the conductors, this charge will act as a noise source within the cable, unless the cable can be kept rigid. This noise is highly dependent on any motion of the cable; noise levels of 5 to 100 mV were reported by Belden Corporation. Noise of similar character (5 to 25 mV) was observed in the laboratory for RG188 coaxial cable, as it was moved and flexed.

Another type of motion-related noise occurs when a cable is moved through a magnetic field. Voltage will be induced in the cable as the cable cuts fixed flux lines or the flux density, B, changes. This kind of noise is troublesome in a high-vibration environment, where the cables can be in rapid motion. If the cable can be kept from vibrating relative to the field, this noise will not occur.

Finally, if instrumentation is operating in close proximity to a radio or television station, signals may be picked up from the transmissions. In addition to AM, FM, and television transmitters, the RFI may come from CB radios, amateur radios, walkie-talkies, paging systems, etc. High-frequency noise should be considered as a possible source of mysterious drifts in dc circuitry, due to rectification of picked-up rf; investigations of drift should always be conducted with a wideband oscilloscope.

SUMMARY

We have described here the different types of interference noise that will exist in any electronic system. Table 3 lists the noise sources discussed above and some effective approaches to solving the pickup problem. It is important to understand the complete noise system (source, coupling medium, receiver, and relationships) before noise-reduction techniques are employed.

Noise reduction is not a mystical job for wizards; it is a practical and analytical job for engineers. Needless to say, the most effective approach is *prevention*—applying noise-reduction analysis and minimization techniques *before the system is built.*

In part 2 of this article, we will describe the proper application of shielding and guarding techniques for noise reduction.

Further Reading:

Ralph Morrison, *op. cit.*

Henry W. Ott, *Noise Reduction Techniques in Electronic Systems* (New York: John Wiley & Sons, 1976).

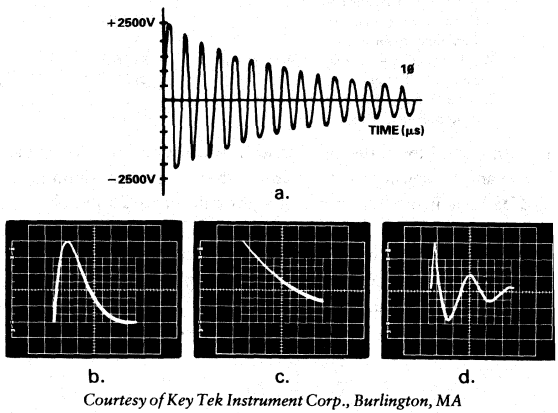

Courtesy of Key Tek Instrument Corp., Burlington, MA

Figure 11. Examples of transients existing in standards for industrial power-line equipment. (a) IEEE Standard 472-1974 "Guide for Surge Withstand Capability." (b) Impulse wave, 8 × 20, 1000V peak, 5μs/div. (c) Impulse wave, 10 × 1000, 1500V peak, 0.2ms/div. (d) 100kHz ac surge, 6kV peak (500kHz leading edge); successive peaks down by 40% (1kV/div, 2μs/div).

Table 3. Noise sources and possible solutions.

Common-Impedance Noise
- Proper circuits for distributing power
- Isolation transformers, optical isolators, analog isolators
- Shielding of sensitive circuits

Capacitively Coupled Noise
- Reducing noise sources
- Properly implemented shields (very effective)
- Reducing stray capacitance

Magnetically Coupled Noise
- Careful routing of wiring
- High-permeability (mumetal) shields (the most effective)
- Reducing area of receiver circuit (twisted pairs, physical wire placement)
- Reducing the noise source (twisted pairs, driven shields to cancel field)

Power-Line Transients
- Coil suppression on relays, solenoids, etc.
- Zero-crossing turnoff for relays, solenoids, etc.
- Shielding
- Reducing stray capacitance

Miscellaneous
- Rigid wiring
- Low-noise cable
- Shielding from RFI source

ANALOG DEVICES

AN-347
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Shielding and Guarding

How to Exclude Interference-Type Noise

What to Do and Why to Do It — A Rational Approach

by Alan Rich

This is the second of two articles dealing with interference noise. In the last issue of *Analog Dialogue* (Vol. 16, No. 3, pp. 16-19), we discussed the nature of interference, described the relationship between sources, coupling channels, and receivers, and considered means of combatting interference in systems by reducing or eliminating one of those three elements.

One of the means of reducing noise coupling is *shielding*. Our purpose in this article is to describe the correct uses of shielding to reduce noise. The major topics we will discuss include noise due to capacitive coupling, noise due to magnetic coupling, and driven shields and guards. A set of guidelines will be included, with do's and don'ts.

From the outset, it should be noted that shielding problems are always rational and do not involve the occult; but they are not always straightforward. Each problem must be analyzed carefully. It is important first to identify the noise source, the receiver, and the coupling medium. Improper shielding and grounding, based on faulty identification of any of these elements, may only make matters worse or create a new problem.

You can think of shielding as serving two purposes. First, shielding can be used to confine noise to a small region; this will prevent noise from extending its reach and getting into a nearby critical circuit. However, the problem with such shields is that noise captured by the shield can still cause problems if the return path the noise takes is not carefully planned and implemented by understanding of the ground system and making the connections correctly.

Second, if noise is present in a system, shields can be placed around critical circuits to prevent the noise from getting into sensitive portions of the circuits. These shields can consist of metal boxes around circuit regions or cables with shields around the center conductors. Again, where and how the shields are connected is important.

CAPACITIVELY COUPLED NOISE

If the noise results from an electric field, a shield works because a charge, Q_2, resulting from an external potential, V_1, cannot exist on the interior of a closed conducting surface (Figure 1).

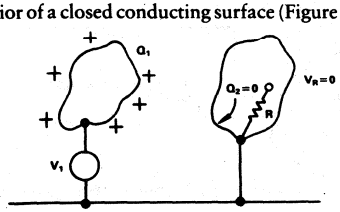

Figure 1. Charge Q_1 cannot create charge inside a closed metal shell.

Reprinted from Analog Dialogue 17-1 1983

Coupling by mutual, or stray, capacitance can be modeled by the circuit of Figure 2. Here, V_n is a noise source (switching transistor, TTL gate, etc.), C_s is the stray capacitance, Z is the impedance of a receiver (for example, a bypass resistor connected between the input of a high-gain amplifier and ground), and V_{no} is the output noise developed across Z.

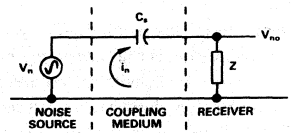

Figure 2. Equivalent circuit of capacitive coupling between a source and a nearby impedance.

A noise current, $i_n = V_n/(Z + Z_{Cs})$, will result, producing a noise voltage, $V_{no} = V_n/(1 + Z_{Cs}/Z)$. For example, if C_s = 2.5 pF, Z = 10kΩ (resistive), and V_n = 100mV at 1.3 MHz, the output noise will be 20 mV (0.2% of 10V, i.e., 8 LSBs of 12 bits).

It is important to recognize the effect that very small amounts of stray capacitance will have on sensitive circuits. This becomes increasingly critical as systems are being designed to combine circuits operating at lower power (implying higher impedance levels), higher speed (implying lower nodal stray capacitance, faster edges, and higher frequencies), and higher resolution (much less output noise permitted).

When a shield is added, the change to the situation of Figure 2 is exemplified by the circuit model of Figure 3. With the assumption that the shield has zero impedance, the noise current in loop A-B-D-A will be V_n/Z_{Cs1}, but the noise current in loop D-B-C-D will be zero, since there is no driving source in that loop. And, since no current flows, there will be no voltage developed across Z. The sensitive circuit has thus been shielded from the noise source, V_n.

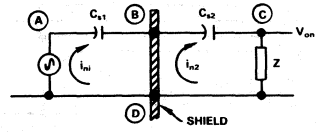

Figure 3. Equivalent circuit of the situation of Figure 2, with a shield interposed between the source and the impedance.

Guidelines for Applying Electrostatic Shields

•An electrostatic shield, to be effective, should be connected to the reference potential of any circuitry contained within the shield. If the signal is earthed or grounded (i.e., connected to a metal chassis or frame, and/or to earth), the shield must be earthed or grounded. But grounding the shield is useless if the signal is not grounded.

•The shield conductor of a shielded cable should be connected to the reference potential at the signal-reference node (Figure 4).

•If the shield is split into sections, as might occur if connectors are used, the shield for each segment must be tied to those for the

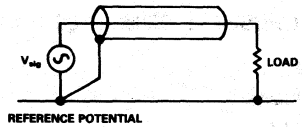

Figure 4. Grounding a cable shield.

adjoining segments, and ultimately connected (only) to the signal-reference node (Figure 5).

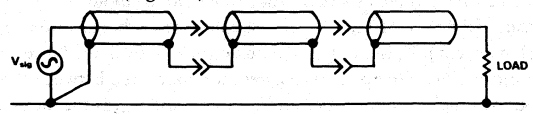

Figure 5. Shields must be interconnected if interrupted.

•The number of separate shields required in a system is equal to the number of independent signals that are being measured. Each signal should have its own shield, with no connections to other shields in the system, unless they share a common reference potential (signal "ground"). If there is more than one signal ground (Figure 6), each shield should be connected to its own reference potential.

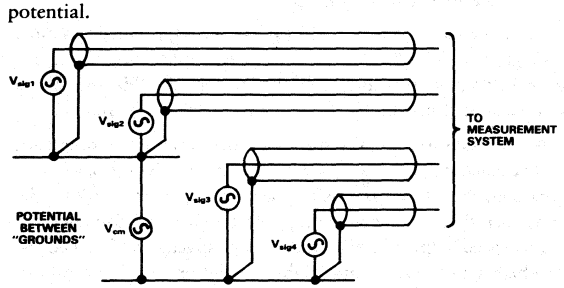

Figure 6. Each signal should have its own shield connected to its own reference potential.

•*Don't connect both ends of the shield to "ground".* The potential difference between the two "grounds" will cause a shield current to flow (Figure 7). The shield current will induce a noise voltage into the center conductor via magnetic coupling. An example of this can be found in Part 1 of this series, *Analog Dialogue* 16-3, page 18, Figure 10.

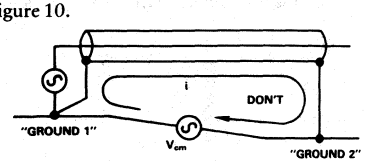

Figure 7. Don't connect the shield to ground at more than one point.

•*Don't allow shield current to exist* (except as noted later in this article). The shield current will induce a voltage in the center conductor.

•*Don't allow the shield to be at a voltage with respect to the reference potential* (except in the case of a guard shield, to be described). The shield voltage will couple capacitively to the center conductor (or conductors in a multiple-conductor shield). With a noise voltage, V_s, on the shield, the situation is as shown i Figure 8.

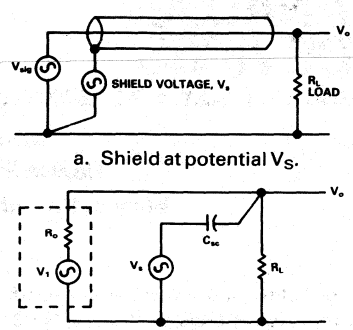

Figure 8. Don't permit the shield to be at a potential with respect to the signal.

The fraction of V_s appearing at the output will be

$$V_o = \frac{V_s}{\sqrt{1 + \frac{1}{(2\pi f R_{eq} C_{sc})^2}}} \quad (1)$$

where V_1 is the open-circuit signal voltage, R_o is the signal's source impedance, C_{sc} is the cable's shield-to-conductor capacitance, and R_{eq} is the equivalent parallel resistance of R_o and R_L. For example, if V_s = 1V at 1.5MHz, C_{sc} = 200pF (10 feet of cable), R_o = 1000 ohms, and R_L = 10kΩ, the output noise voltage will be 0.86 volts. This is an often-ignored guideline; serious noise problems can be *created* by inadvertently applying undesired potentials to the shield.

•*Know by careful study how the noise current that has been captured by the shield returns to "ground."* An improperly returned shield can cause shield voltages, can couple into other circuits, or couple into other shields. The shield return should be as short as possible to minimize inductance.

Here is an example that illustrates the problems that can arise in relation to these last two guidelines: Consider the improperly configured shield system shown in Figure 9, in which a precision voltage source, V1, and a digital logic gate share a common shield connection. This situation can occur in a large system where analog and digital signals are cabled together.

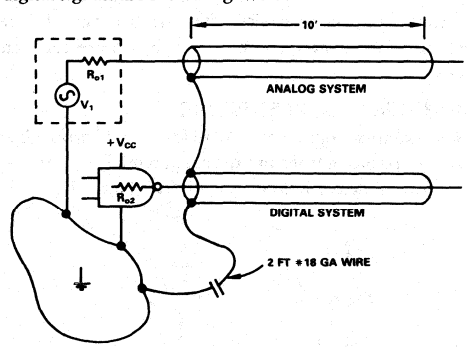

Figure 9. A situation that generates transient shield voltages.

A step voltage change in the output of the logic circuit couples capacitively to its shield, creating a current in the common 2-foot

hield return. This, in turn, develops a shield voltage common to oth the analog and digital shields. An equivalent circuit is shown n Figure 10, in which V(t) is a 5-volt step from a TTL logic gate, R_{o2} is the 13-ohm output impedance of the logic gate, C_{ws} is the 470-pF capacitance from the shield to the center conductor of the hielded cable, and R_s and L_s are the 0.1-ohm resistance and 1-nicrohenry inductance of the 2-foot wire connecting the shield to he system ground.

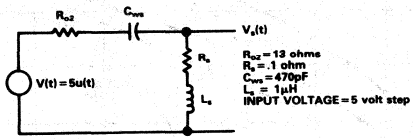

Figure 10. Equivalent circuit for generating shield voltage.

The shield voltage, $V_s(t)$, can be solved for by conventional circuit-analysis techniques, or simulated by actually building and carefully making measurements on a circuit with the given parameters. For the purpose of demonstration, the calculated response waveform, illustrated in Figure 11, with a 5-volt initial spike, resonant frequency of 7.3 MHz, and damping time constant of 0.15μs, is sufficient to illustrate the nature of the voltage that appears on the shield and is capacitively coupled to the analog input. If the voltage is looked at with a wideband oscilloscope, it will look like a noise "spike." We can see that this transient will couple a fast damped waveform of significant peak amplitude to the analog system input.

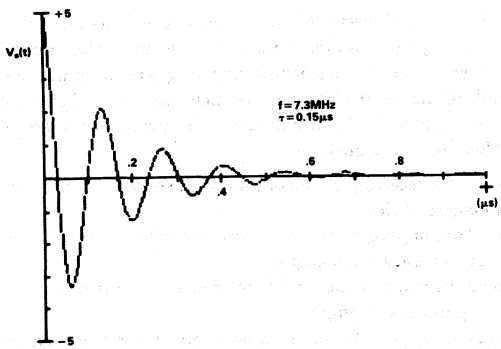

Figure 11. Computed response of circuit of Figure 10.

Even in a purely digital system, noise glitches can be caused to appear in apparently remote portions of a system having the kind of situation shown. This can often explain some otherwise inexplicable system bugs.

In quite a few cases, the proper choice of shield connection among the many possibilities may not be immediately obvious, and the guidelines may not provide us with a clear choice. There is no alternative but to analyze the various possibilities and choose the approach for which the lowest noise may be calculated.

For example, consider the case illustrated in Figure 12, in which the measurement system and the source have differing ground potentials. Should we connect the shield to A: the low side at the measurement-system input, B: ground at the system input, C: ground at the signal source, or D: the low side at the source?

A is a poor choice, since noise current is allowed to flow in a signal

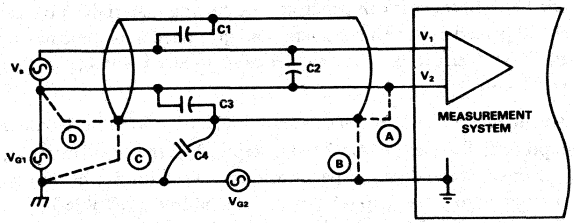

Figure 12. Possible grounds where system and source have differing ground potentials.

conductor. The path of the noise current due to V_{G1}, as it returns through C4, is shown in Figure 13a.

B is also a poor choice, since the two noise sources in series, V_{G1} and V_{G2}, produce a component across the two signal wires, developed by the source impedance in parallel with C_2, in series with C_1, as shown in Figure 13b.

C is poor, too, since V_{G1} produces a voltage across the two signal wires, by the same mechanism as (B), as Figure 13c shows.

D is the best choice, under the given assumptions, as can be seen in Figure 13d. It also tends to confirm the grounding guideline to connect the shield at the signal's reference potential.

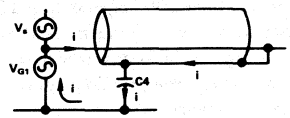

a. Return path A.

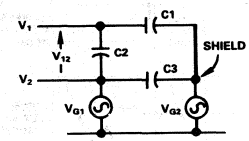

b. Return path B.

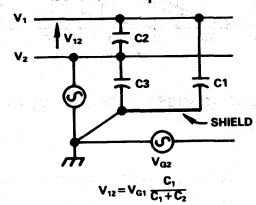

c. Return path C.

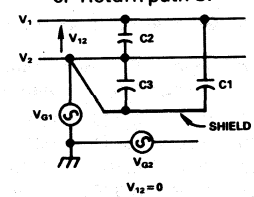

d. Return path D.

Figure 13. Equivalent circuits.

NOISE RESULTING FROM A MAGNETIC FIELD

Noise in the form of a magnetic field induces voltage in a conductor or circuit; it is much more difficult to shield against than elec-

tric fields because it can penetrate conducting materials. A typical shield placed around a conductor and grounded at one end has little if any effect on the magnetically induced voltage in that conductor.

As a magnetic field, B, penetrates a shield, its amplitude decreases exponentially (Figure 14). The skin depth, δ, of the shield material, is defined as the depth of penetration required for the field to be attenuated to 37% (exp (-1)) of its value in free air. Table 1[1] lists typical values of δ for several materials at various frequencies. You can see that any of the materials will be more effective as a shield at high frequency, because δ decreases with frequency, and that steel provides at least an order of magnitude more effective shielding at any frequency than copper or aluminum.

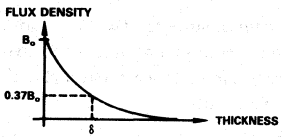

Figure 14. Magnetic field in a shield as a function of penetration depth.

Figure 15 compares absorption loss as a function of frequency for two thicknesses of copper and steel. ⅛-inch steel becomes quite effective for frequencies above 200 Hz, and even a 20-mil (0.5 mm) thickness of copper is effective at frequencies above 1 MHz. However, all show a glaring weakness at lower frequencies, including 50-60-Hz line frequencies—the principal source of magnetically coupled noise at low frequency.

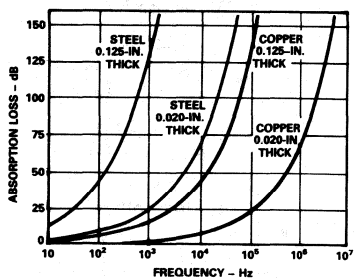

Figure 15. Absorption loss vs. frequency for two thicknesses of copper and steel.

For improved low-frequency magnetic shielding, a shield consisting of a high-permeability magnetic material (e.g., Mumetal)

Table 1. Skin depth, δ, vs. frequency

Frequency	δ for Copper (in.)	δ for Copper (mm)	δ for Aluminum (in.)	δ for Aluminum (mm)	δ for Steel (in.)	δ for Steel (mm)
60Hz	0.335	8.5	0.429	10.9	0.034	0.86
100Hz	0.260	6.6	0.333	8.5	0.026	0.66
1kHz	0.082	2.1	0.105	2.7	0.008	0.2
10kHz	0.026	0.66	0.033	0.84	0.003	0.08
100kHz	0.008	0.2	0.011	0.3	0.0008	0.02
1MHz	0.003	0.08	0.003	0.08	0.0003	0.008

[1]Table 1 and Figures 15 and 16 are from Ott, H.W., *Noise Reduction Techniques in Electronic Systems* (New York: John Wiley & Sons, © 1976).

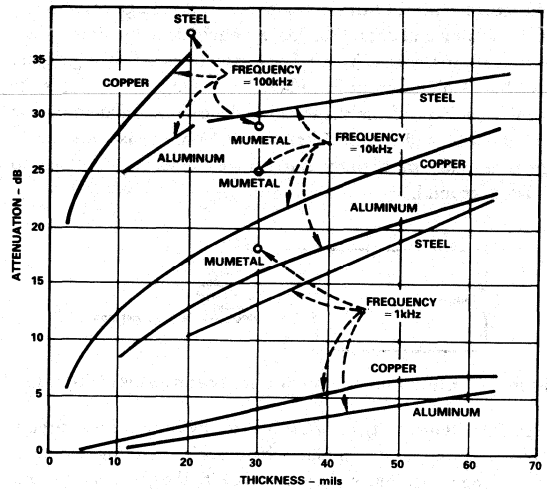

Figure 16. Shielding attenuation of Mumetal and other materials at several frequencies.

should be considered. Figure 16 compares a 30-mil thickness of Mumetal with various materials at several frequencies. It shows that, below 1 kHz, Mumetal is more effective than any of the other materials, while at 100kHz it is the least effective. However, Mumetal is not especially easy to apply, and if it is saturated by an excessively strong field, it will no longer provide an advantage.

As you can see, it is very difficult to shield against magnetic fields, i.e., to modify the coupling medium by shielding. Therefore, the most effective approaches at low frequency are to minimize the strength of the interfering magnetic field, minimize the receiver loop area, and minimize coupling by optimizing wiring geometries. Here are some guidelines:

- Locate the receiving circuits as far as possible from the source of the magnetic field.
- Avoid running wires parallel to the magnetic field; instead, cross the magnetic field at right angles.
- Shield the magnetic field with an appropriate material for the frequency and field strength.
- Use a twisted pair of wires for conductors carrying the high-level current that is the source of the magnetic field. If the currents in the two wires are equal and opposite, the net field in any direction

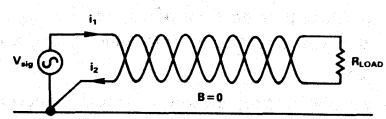

a. Correct connection with balanced currents.

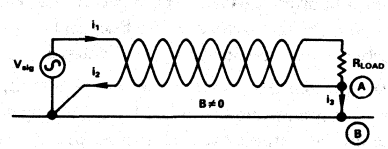

b. Incorrect connection forming ground loop.

Figure 17. Connections to a twisted pair.

ıver each cycle of twist will be zero (Figure 17a). For this arrangeıent to work, none of the current can be shared with another conluctor, for example, a ground plane. Figure 17b shows what can ıappen if a ground loop is formed; if part of the current flows hrough the ground plane (depending on the ratio of conductor reistance to ground resistance), it will form a loop with the twisted ›air, generating a field determined by i_3 ($= i_1 - i_2$).

The ground connection between A and B need not be as simple as a short circuit to cause trouble. Any stray unbalanced capacitance or resistance from R_{load} circuits to the ground plane will also unbalance the currents and produce a net current through the wires and the ground plane, producing a ground loop and a related magnetic field. For this reason, it is also good practice to run the twisted pair close to the ground plane to tend to balance the capacitances from each side to ground, as well as to minimize loop area.

●Use a shielded cable with the high-level source circuit's return current carried in the shield (Figure 18). If the shield current, I_2 is equal and opposite to that in the center conductor, the center-conductor field and the shield field will cancel, producing a zero net field. In this case, which seems to violate the "no shield current" rule for receiver circuits, the concentric cable is not used to shield the center lead; instead, the geometry produces cancellation.

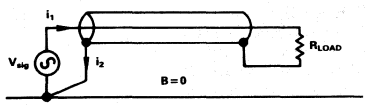

Figure 18. Use of shield for return current to noisy source.

This scheme can be usefully employed in an ATE system where accurate measurements must be performed on devices with high power-supply currents that may be noisy. For example, Figure 19 shows the application of this technique to the connections for the high-current logic supply for an a/d converter under test—at the end of a test cable.

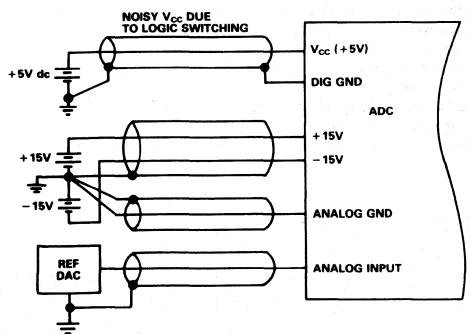

Figure 19. Application of circuit of Figure 18 in a test system.

●Since magnetically induced noise depends on the area of the receiver loop, the induced voltage due to magnetic coupling can be reduced by reducing the loop's area. What is the receiver loop? In the example shown in Figure 20, the signal source and its load are connected by a pair of conductors of length L and separation D. The circuit (assuming it has a rectangular configuration) forms a loop with area D · L.

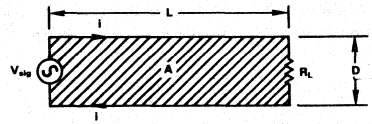

Figure 20. Area of a loop that receives magnetically coupled noise.

The voltage induced in series with the loop is proportional to the area and the cosine of its angle to the field. Thus, to minimize noise, the loop should be oriented at right angles to the field, and its area should be minimized.

The area can be reduced by decreasing the length of and/or decreasing the distance between the conductors. This is easily accomplished with a twisted pair, or at least a tightly cabled pair, of conductors. It is good practice to pair conductors so that the circuit wire and its return path will always be together. To do this, the designer must be certain of the actual path that the return current takes in getting back to the signal source. Quite often, the current returns by a path not intended in the original design layout.

If wires are moved (for example, by a technician troubleshooting some other problem), the loop area and orientation to the field may change, so that yesterday's acceptable noise level may be transformed to tomorrow's disastrous noise level. Which may lead to a service call . . . and another repetition of the cycle. The bottom line: Know the loop area and orientation, do what must be done to minimize noise—*and permanently secure the wiring*!

DRIVEN SHIELDS AND GUARDING

We have discussed the role of a current-driven shield carrying an equal and opposite current to reduce generated noise by reducing the magnetic field around a conductor.

Guarding is similar, in that it involves driving a shield, at low impedance, with a potential essentially equal to the common-mode voltage on the signal wire contained within the shield. Guarding has many useful purposes: It reduces common-mode capacitance, improves common-mode rejection, and eliminates leakage currents in high-impedance measurement circuits.

Figure 21 shows an example of an op amp with negligible bias current connected as a high-impedance non-inverting amplifier with gain. The purpose of the cable is to shield the high input-impedance signal conductor from capacitively coupled noise and to minimize leakage currents. The signal comes from a 10-megohm source, and the cable is assumed to have 1000 megohms of leakage resistance (which may change as a function of temperature, humidity, etc.) from conductor to shield. If connected as shown, the equivalent input circuit is an attenuator which loses 1% of the

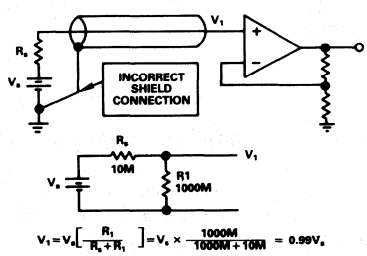

Figure 21. Op amp connected as high-impedance non-inverting amplifier with gain, with shielded input lead.

signal at the time it is measured, and an unknown fraction at other times. Also, the cable capacitance produces a substantial lag time constant, R_sC_c.

Figure 22 has the same players, but the shield is connected to the tap of the gain divider (usually at low impedance). Being connected to the inverting input of the op amp, it should be at the same potential as the amplifier's non-inverting input. Since there is no voltage across the cable's leakage resistance, there is no current through it and its resistance value doesn't matter; $\dot{V}_1$ must therefore be equal to V_s, since bias current was assumed negligible.

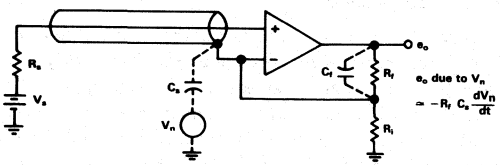

Figure 22. Same as Figure 21, but cable shield connected as a guard.

Also, there is no voltage across the cable capacitance, hence no charging or discharging of the cable; thus the lag time constant depends mainly on circuit strays and the amplifier's input capacitance. For stability, capacitance should be connected between the output and the negative input, such that $C_fR_F = C_sR_i$, where C_s is sum of the stray capacitance between shield and ground and the input capacitance.

There must be no noise voltage applied to the guard. In noisy systems, as Figure 22 shows, capacitively coupled noise will be differentiated, emphasizing the higher-frequency components. This can be avoided (Figure 23) by either using a buffer follower with fast response and low output impedance to drive the guard (a) or a second shield, around the guard, grounded to the signal common (b).

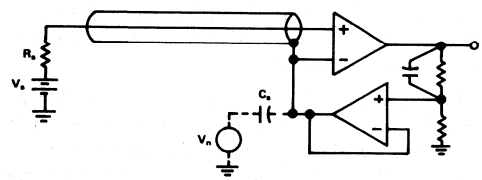

(a) Driven guard.

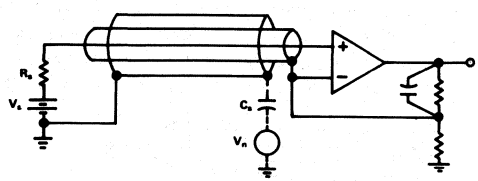

(b) Shielded guard.

Figure 23. Avoiding noise pickup on the guard.

In high-impedance current-input inverting configurations, where a length of shielded wire is used to guard the lead from the current source to the amplifier's inverting input, the guard should either be driven by a buffer at the same potential as the non-inverting input (and connected nowhere else), or be tied directly to the non-inverting input, with a second outer shield connected to the signal's reference point.

SUMMARY

Table 2 summarizes the important points made in this article. All are important to maintaining a high-integrity shield system. However, we cannot emphasize too strongly the two subjects that are most-often ignored: appearance of noise voltage on signal shields and proper disposition of shield noise currents. *Noise voltage must not exist on the shield*; shield-to-conductor capacitance will couple the noise directly to the center conductor. *If shield currents are not returned properly, they can show up in a remote part of the system* and perhaps cause trouble in a location totally unrelated to the shielding problem that was "solved."

Table 2. Applicability of shielding considerations

Consideration	Universal	Electric	Magnetic
Know the noise source, coupling medium, and receiver.	X	X	X
Different shielding techniques are required for different noise sources, coupling channels, and receivers.	X	X	X
In most situations, conventional circuit analysis using lumped elements can be used.	X	X	X
Connect the shield at the signal-source end only.		X	
Carry shields through connectors.		X	
Individual shields should not be tied together.		X	
Do not ground both ends of a shield.		X	
Do not allow shield current to flow, except for driven shields - to cancel magnetic fields		X	X
Do not allow voltage on a shield, except for guarding.		X	
Know *exactly* where noise current from the shield will flow.		X	
Use short connections to return noise current from the shield.		X	
Electrostatic shields have little effect in reducing noise resulting from magnetic fields.			X
Reduce magnetic fields by physical separation proper orientation, twisted pairs, and/or driven shields.			X
Know the receiver loop area and orientation to the field. Keep field at right angles and reduce the loop area by using paired conductors, preferably twisted pairs, and minimize wire lengths.			X
Use guarding in high-impedance circuits	X	X	
In high-impedance circuits, be extremely careful of shield noise	X	X	

BRIEF BIBLIOGRAPHY

For further reading, see:

Brokaw, A. Paul. "Analog Signal Handling for High Speed and Accuracy," *Analog Dialogue* 11-2, 1977, pp. 10-16.

Brokaw, A. Paul. "An I.C. Amplifier Users' Guide to Decoupling, Grounding, and Making Things Go Right for a Change," *Analog Devices Data-Acquisition Databook 1982*, Volume I, Pages 21-13 to 21-20.

Morrison, Ralph. *Grounding and Shielding Techniques in Instrumentation* Second Edition. (New York: John Wiley & Sons, 1977).

Ott, Henry W. *Noise Reduction Technique in Electronic Systems.* (New York: John Wiley & Sons, 1976).

ANALOG DEVICES

AN-348
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Avoiding Passive-Component Pitfalls

The Wrong Passive Component Can Derail Even the Best Op Amp or Data Converter Here Are Some Basic Traps to Watch for

by Doug Grant and Scott Wurcer

You've just spent $25 or more for a precision op amp or data converter, only to find that, when plugged into your board, the device doesn't meet spec. Perhaps the circuit suffers from drift, poor frequency response, oscillations—or simply doesn't achieve the accuracy you expect. Well, before you blame the device itself, you should examine your passive components—including capacitors, resistors, potentiometers, and yes, even the printed circuit boards themselves. Subtle effects of tolerance, temperature, parasitics, aging, and user assembly procedures can unwittingly sink your circuit. And these effects all too often go unspecified or underspecified by manufacturers.

In general, if you use data converters having 12 bits or more of resolution, or op amps that cost more than $5, you should pay particularly close attention to passive-component selection. To put the problem in perspective, consider the case of a 12-bit digital-to-analog converter (DAC). One half LSB (least-significant bit) corresponds to 0.012% of full scale, or only 122 parts per million (ppm)! The host of passive-component phenomena can quickly accumulate errors far exceeding this level.

Buying the most-expensive passive components won't necessarily solve your problems either. Often, the correct 25-cent capacitor will yield a better-performing, more cost-effective design than the premium-grade $8 part. Although not necessarily easy, understanding and analyzing passive-component effects may prove quite rewarding, once you understand a few basics.

CAPACITORS

Most designers are generally familiar with the range of capacitors available. But the mechanisms by which both static and dynamic errors can occur in precision circuit designs are easy to forget because of the tremendous variety of capacitor types, e.g.: glass, aluminum foil, solid tantalum and tantalum foil, silver mica, ceramic, Teflon, and the film capacitors, including polyester, polycarbonate, polystyrene, and polypropylene types.

Figure 1 is a workable model of a non-ideal capacitor. The nominal capacitance, C, is shunted by a resistance R_p, representing insulation resistance or leakage. A second resistance, R_s—equivalent series resistance, or ESR—appears in series with the capacitor and represents the resistance of the leads and capacitor plates.* Inductance, L—the equivalent series inductance, or ESL, models the inductance of the leads and plates. Finally, resistance R_{da} and capacitance C_{da} together form a simplified model of a phenomenon non known as dielectric absorption. Dielectric absorption can ruin the dynamic performance of both fast and slow circuits.

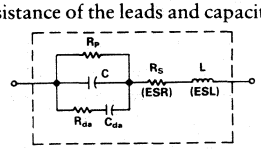

Figure 1. Capacitor equivalent circuit.

*Capacitor phenomena aren't that easy to separate out. This matching of phenomena and models is for convenience in explanation.

Reprinted from Analog Dialogue 17–2 1983

Dielectric Absorption

We begin with dielectric absorption, also known as "soakage" and sometimes as "dielectric hysteresis"—perhaps the least understood and potentially most damaging capacitive effect. Upon discharge, most capacitors are reluctant to give up all of their former charge. Figure 2 illustrates the effect. After being charged to V volts at time t_0, the capacitor is shorted by the switch at time t_1. At time t_2, the capacitor is open-circuited; a residual voltage slowly builds up across its terminals and reaches a nearly constant value. This voltage is due to "dielectric absorption."

Figure 2. Residual voltage characterizes capacitor dielectric absorption.

Standards techniques for specifying or measuring dielectric absorption are few and far between. Measured results are usually expressed as the percentage of the original charging voltage that reappears across the capacitor. Typically, the capacitor is charged for more than 1 minute, then shorted for an established time between 1 and 10 seconds. The capacitor is then allowed to recover for approximately 1 minute, and the residual voltage is measured (see reference 10).

In practice, dielectric absorption makes itself known in a variety of ways. Perhaps an integrator refuses to reset to zero, a voltage-to-frequency converter exhibits unexpected nonlinearity, or a sample-and-hold exhibits varying errors. This last manifestation can be particularly damaging in a data-acquisition system, where adjacent channels may be at voltages which differ by nearly full scale. Figure 3 illustrates the case in a simple sample-hold.

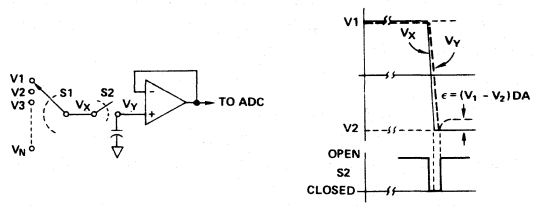

Figure 3. Dielectric absorption induces errors in sample-and-hold application.

The dielectric absorption is a characteristic of the dielectric material itself, although it can be affected by inferior manufacturing processes or electrode materials. As a percentage of the charging voltage, dielectric absorption specifications range from a low of 0.02% for Teflon, polystyrene, and polypropylene capacitors to a high of 10% or more for some aluminum electrolytics. For some time-frames, the D.A. of polystyrene can be as low as 0.002%.

Common ceramic and polycarbonate types display typical dielectric absorptions of 0.2%; this corresponds to one-half of an LSB at only 8 bits! Silver mica, glass, and tantalum capacitors typically exhibit even larger dielectric absorptions, ranging from 1.0% to 5.0%, with those of polyester devices falling in the vicinity of 0.5%. As a rule, if your capacitor's spec sheet does not discuss dielectric absorption *in your time frame and voltage range*, exercise caution.

Dielectric absorption can produce long tails in the transient response of fast-settling circuits, such as those found in high-pass active filters or ac amplifiers. In some devices used for such applications, Figure 1's R_{da}-C_{da} model of dielectric absorption can have a time constant of milliseconds.* In fast-charge, fast-discharge applications, the dielectric absorption resembles "analog memory"; the capacitor tries to remember its previous voltage.

In some designs, you can compensate for the effects of dielectric absorption if it is simple and easily characterized, and you are willing to do custom-tweaking. In an integrator, for instance, the output signal can be fed back through a suitable compensation network, tailored to cancel the circuit equivalent of the dielectric absorption by placing a negative impedance effectively in parallel. Such compensation has been shown to improve sample- and-hold circuit performance by factors of 10 or more (Reference 7).

Parasitics and Dissipation Factor

In Figure 1, a capacitor's leakage resistance, R_p, the effective series resistance, R_s, and effective series inductance, L, act as parasitic elements which can degrade an external circuit's performance. The effects of these elements are often lumped together and defined as a dissipation factor, or DF.

A capacitor's leakage is the small current that flows through the dielectric when a voltage is applied. Although modeled as a simple insulation resistance (R_p) in parallel with the capacitor, the leakage actually is nonlinear with voltage. Manufacturers often specify leakage as a megohm-microfarad product, which describes the dielectric's self-discharge time constant, in seconds. It ranges from a low of 1s or less for high-leakage capacitors, such as aluminum and tantalum devices, to the 100's of seconds for ceramic capacitors. Glass devices exhibit self-discharge time-constants of 1,000 or more; but the best leakage performance is shown by Teflon and the film devices (polystyrene, polypropylene), with time constants exceeding 1,000,000 megohm-microfarads. For such a device, leakage paths—created by surface contamination of the device's case or in the associated wiring or physical assembly—can overshadow the dielectric's leakage.

Effective series inductance, ESL (Figure 1) arises from the inductance of the capacitor leads and plates, which, particularly at the higher frequencies, can turn a capacitor's normally capacitive reactance into an inductive reactance. Its magnitude depends on construction details within the capacitor. Tubular wrapped-foi devices display significantly more lead inductance than molded ra dial-lead configurations. Multilayer ceramic and film-type device typically exhibit the lowest series impedances, while tantalum an aluminum electrolytics typically exhibit the highest. Con sequently, electrolytic types usually prove insufficient for high speed local bypassing applications.

Manufacturers of capacitors often specify effective series induc tance by means of impedance-versus-frequency plots. These show graphically, and not surprisingly, that the devices display predo minantly capacitive reactance at low frequencies, with rising im pedance at higher frequencies because of their series inductance.

Effective series resistance, ESR (resistor R_s of Figure 1), is made up of the resistance of the leads and plates. As noted, many manufacturers lump the effects of ESR, ESL, and leakage into a single parameter called *dissipation factor*, or DF. Dissipation factor measures the basic inefficiency of the capacitor. Manufacturers define it as the ratio of the energy lost to energy stored per cycle by the capacitor. The ratio of equivalent series resistance to total capacitive reactance—at a specified frequency—approximates the dissipation factor, which turns out to be equivalent to the reciprocal of the figure of merit, Q.

Dissipation factor often varies as a function of both temperature and frequency. Capacitors with mica and glass dielectrics generally have DF values from 0.03% to 1.0%. For ceramic devices, DF ranges from a low of 0.1% to as high as 2.5% at room temperature. And electrolytics usually exceed even this level. The film capacitors usually are the best, with DF's of less than 0.1%.

Tolerance, Temperature, and Other Effects

In general, precision capacitors are expensive and—even then—not necessarily easy to buy. In fact, choice of capacitance is limited by the the range of available values and tolerances. Tolerances of ± 1% for some ceramics and most film-type devices are common, but with possibly unacceptable delivery times. Most film capacitors can be made available with tolerances of less than ± 1%, but on special order only.

Most capacitors are sensitive to temperature variations. Dissipation factor, dielectric absorption, and capacitance itself are all functions of temperature. For some capacitors, these parameters vary approximately linearly with temperature; in others they vary quite nonlinearly. Although not usually important for sample-and-hold applications, an excessively large temperature coefficient (ppm/°C) can prove harmful to the performance of precision integrators, voltage-to-frequency converters, and oscillators. NPO ceramic capacitors, with temperature-drift as low as 30 ppm/°C, usually do the best. On the other hand, aluminum electrolytics' temperature coefficients can exceed 10,000 ppm/°C.

A capacitor's maximum working temperature should also be considered. Polystyrene capacitors, for instance, melt near 85°C, compared to Teflon's ability to survive temperatures up to 200°C.

Sensitivity of capacitance and dielectric absorption to applied voltage can also hurt capacitor performance in a circuit application. Although capacitor manufacturers do not always clearly specify voltage coefficients, the user should always consider the possible effects of such factors. For instance, when maximum voltages are applied, some high-density ceramic devices can experience a decrease in capacitance of 50% or more!

*Much longer time constants are also quite usual. In fact, some devices can be modeled by several paralleled R_{da}–C_{da} circuits, with a wide range of time constants.

Similarly, the capacitance and dissipation factor of many types vary significantly with frequency, mainly as a result of a variation in dielectric constant. In this regard, the better dielectrics are polystyrene, polypropylene, and Teflon.

Assemble Critical Components Last

The designer's worries don't end with the design process. Commonly used printed-circuit-board assembly techniques can prove ruinous to even the best of designs. For instance, some commonly used p-c board cleaning solvents can infiltrate certain electrolytic capacitors—those with rubber end caps are particularly susceptible. Even worse, some of the film capacitors, polystyrene in particular, actually melt when contacted by some solvents. Rough handling of the leads can damage still other capacitors, creating random or even intermittent circuit problems. Etched-foil types are particularly delicate in this regard. To avoid these difficulties, it may be advisable to mount especially critical components as the last step in the board assembly process—if possible.

Designers should also consider the natural failure mechanisms of capacitors. Metallized film devices, for instance, often self-heal. They initially fail due to conductive bridges that develop through small perforations in the dielectric films. But the resulting fault currents can generate sufficient heat to destroy the bridge, thus returning the capacitor to normal operation (at slightly lower capacitance). Of course, applications in high-impedance circuits may not develop sufficient current to clear the bridge.

Tantalum capacitors also exhibit a degree of self-healing, but—unlike film capacitors—the phenomenon depends on the temperature at the fault location rising slowly. Therefore, tantalum capacitors self-heal best in high impedance circuits which limit the surge in current through the capacitor's defect. Use caution, therefore, when specifying tantalums for high-current applications.

Electrolytic capacitor life often depends on the rate at which capacitor fluids seep through end caps. Epoxy end seals perform better than rubber seals, but an epoxy sealed capacitor can explode under severe reverse-voltage or overvoltage conditions.

RESISTORS AND POTS

Designers have a broad range of resistor technologies to choose from, including carbon composition, carbon film, bulk metal, metal film, and both inductive and non-inductive wire-wound types. As perhaps the most basic—and presumably most trouble-free—of components, the resistor is often overlooked as a potential source of errors in high-performance circuits. Yet, an improperly selected resistor can subvert the accuracy of a 12-bit design by developing errors well in excess of 122 ppm, (½ LSB). When did you last take the time to actually read a resistor data sheet? You'd be surprised at what can be learned from an informed review of the data.

Consider the circuit of Figure 4, which amplifies a 0-to-100-mV input signal 100 times for conversion by a 12-bit ADC with a 0-to-10-volt input range. The gain-setting resistors can be bought in initial tolerances of as low as ±0.001% (10 ppm) in the form of precision bulk metal-film devices. Alternatively, the initial tolerance

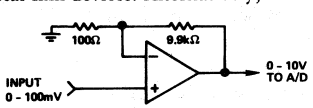

Figure 4. Temperature changes can reduce amplifier accuracy.

of the resistors may be corrected through calibration or selection. Consequently, the initial gain accuracy of the circuit can be set to whatever tolerance is required, limited perhaps by the accuracy of calibration instrumentation.

Temperature changes, however, can limit the accuracy of the amplifier of Figure 4 in several ways. The absolute temperature coefficients of the resistors are unimportant, as long as they track. Even so, carbon composition resistors, with temperature coefficients of approximately 1,500 ppm/°C, would not suit the application. Even if the tempcos could be matched to an unlikely 1%, the resulting 15 ppm/°C differential would prove inadequate—a shift of as little as 8°C would create a ½-LSB error of 120 ppm.

Manufacturers do offer metal film and bulk metal resistors with absolute temperature coefficients ranging between ±1 and ±100 ppm/°C. Beware, though; temperature coefficients can vary a great deal, particularly among resistors from different batches. To avoid this problem, expensive matched resistor pairs are offered by a few manufacturers, with temperature coefficients that track one another to within 2 to 10 ppm/°C. Low-priced thin-film networks are good and are widely used.

Unfortunately, even matched resistor pairs cannot fully solve the problem of temperature-induced resistor errors. Figure 5a illustrates error-inducing through self-heating. The resistors have identical temperature coefficients but dissipate considerably different amounts of power in this circuit. With an assumed thermal resistance (data sheet) of 125°C/W for ¼-watt resistors, resistor R1's temperature rises by 0.0125°C, while resistor R2's temperature rises by 1.24°C. With a temperature coefficient of 50 ppm/°C, the result is an error of 62 ppm (0.006%).

Even worse, the effects of self-heating create nonlinear errors. In the example of Figure 5a, with half the voltage input, the resulting error is only 15 ppm. Figure 5b graphs the resulting nonlinear transfer function for the circuit of Figure 5a. This is by no means a worst-case example; smaller resistors would give even worse results due to their higher thermal resistance.

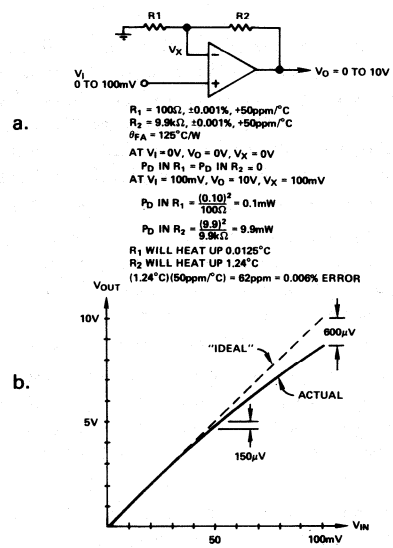

Figure 5. Resistor self-heating leads to nonlinear amplifier response. (a) Anatomy of temperature-induced nonlinearity. (b) Nonlinear transfer function (scale exaggerated).

The use of higher-wattage resistors for those devices that dissipate the greatest power can minimize the effects of resistor self-heating. Alternatively, thin- or thick-film resistor networks minimize the effects of self heating by spreading the heat more evenly over all the resistors in a given package.

Often overlooked as a source of error, the temperature coefficient of resistance of typical wire or pc-board interconnects can add to a circuit's errors. Metals used in p-c boards and for interconnecting wires (e.g., copper) have a temperature coefficient as high as 3,900 ppm/°C. A precision 10-ohm, 10 ppm/°C wirewound resistor, with 0.1-ohms of interconnect resistance, for instance, effectively turns into a 45 ppm/°C resistor. The temperature coefficients of interconnects play a particularly significant role in precision hybrids, where thin-film interconnects have non-negligible resistance.

One final consideration applies mainly to designs that see widely varying ambient temperatures: a phenomenon known as temperature retrace describes the change in resistance which occurs after a specified number of cycles of exposure to low and high ambients with constant internal dissipation. Temperature retrace can exceed 10 ppm, even for some of the better metal-film components.

In summary, to design resistance circuits for minimum temperature-related errors, consider the following (along with their cost):

- Closely match resistance-temperature coefficients.
- Use resistors with low absolute temperature coefficients.
- Use resistors with low thermal resistance (higher power ratings, larger cases).
- Tightly couple matched resistors thermally; (use standard resistance networks or multiple resistors in a single package).
- For large ratios, consider using stepped attenuators.

Resistor Parasitics

Resistors can exhibit significant levels of parasitic inductance or capacitance, especially at high frequencies. Manufacturers often specify these parasitic effects as a reactance error, in % or ppm, based on the ratio of the difference between the impedance magnitude and the dc resistance, to the resistance, at one or more frequencies.

Wirewound resistors are especially susceptible to difficulties. Although resistor manufacturers offer wirewound components in either normal or noninductively wound form, even noninductively wound resistors create headaches for designers. These resistors still appear slightly inductive (of the order of 20 μH) for R values below 10,000 ohms. Noninductively wound resistors that exceed 10,000 ohms actually exhibit about 5 pF of shunt capacitance.

These parasitic effects can raise havoc in dynamic circuit applications. Of particular concern are applications using wirewound resistors with values both greater and less than 10,000 ohms. Here it is not uncommon to see peaking, or even oscillation. These effects become evident at frequencies in the low-kHz range.

Even in low-frequency circuit applications, parasitic effects in wire wound resistors can create difficulties. Exponential settling to 1 ppm takes 20 time constants or more. Parasitics associated with wire wound resistors can increase settling time beyond the length of those time constants significantly.

Unacceptable amounts of parasitic reactance are often found even in resistors that aren't wirewound. For instance, some metal-film types have significant interlead capacitance, which shows up at high frequencies. Carbon resistors do the best at high frequencies.

Thermoelectric Effects

The junction between any two dissimilar metals creates a therma EMF. In many cases, it can easily produce the dominant error in a precision circuit design. In wire wound resistors, for instance, the resistance wire generates a thermal EMF of 42 microvolts/°C when joined to the leads (A typical lead material is Alloy 180, consisting of 77% copper and 23% nickel). If the resistor's two terminations see the same temperature, the EMFs cancel and no net error results However, if the resistor is mounted vertically a temperature gradient may exist between the bottom and top of the resistor because of air flow past the long lead and its lower heat capacity.

For a temperature difference of as little as 1°C, an error voltage of 42 microvolts results, a level which easily overwhelms the 25 microvolt offsets of typical precision op amps! A horizontally mounted resistor (Figure 6) can resolve the difficulty. Alternatively, some resistor manufacturers offer, on special order, tinned copper leads, which reduce the thermal EMF to 2.5 microvolts/°C.

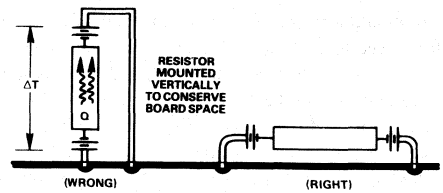

Figure 6. Thermal gradients create significant thermoelectric circuit errors.

In general, designers should strive to avoid thermal gradients on or near critical circuit boards. Often this means thermally isolating components that dissipate significant amounts of power. Thermal turbulence created by large temperature gradients can also result in dynamic noise-like low-frequency errors.

Voltage, Failure, and Aging

Resistors are also plagued by changes as a function of applied voltage. The deposited-oxide high-megohm type components are especially sensitive, with voltage coefficients ranging from 1 ppm/volt to more than 200 ppm/volt. This is another reason to exercise caution in precision applications, such as high-voltage dividers.

Resistors' failure mechanisms can also create circuit difficulties if not carefully considered. Carbon-composition resistors fail safely by turning into open circuits. Consequently, in some applications, these components play a useful secondary role as a fuse. Replacing such a resistor with a carbon-film type can lead to trouble, since carbon-film devices can fail as short circuits. (Metal-film components usually fail as open circuits.)

All resistors tend to change slightly in value with age. Manufacturers specify long-term stability in terms of change—ppm/year. Values of 50 or 75 ppm/year are not uncommon among metal film resistors. For critical applications, metal-film devices should be burned-in for at least one week at rated power. During burn-in, R values can shift by up to 100 or 200 ppm. Metal film resistors may need 4,000 or 5,000 operational hours for full stabilization, especially if they are deprived of a burn-in period.

Resistor Excess Noise

Most designers have some familiarity with thermal, or Johnson, noise in resistors. But a less widely recognized second noise phenomenon, called excess noise, can prove particularly trouble-

some in precision op amp and converter circuits. Excess noise becomes evident only when current passes through a resistor.

To review briefly, thermal noise results from the thermally induced random vibration of charge carriers in a resistor. Although the average current from these vibrations remains zero, instantaneous charge motions result in an instantaneous voltage across the resistor's terminals.

Excess noise, on the other hand, occurs primarily when dc flows in a discontinuous medium, such as a carbon composition resistor. The current flows unevenly through the compressed carbon granules, creating microscopic particle-to-particle "arcing". The phenomenon gives rise to a 1/f noise-power spectrum, in addition to the thermal noise spectrum. In other words, the excess spot noise voltage increases as the inverse square-root of frequency.

Excess noise often surprises the unwary designer. Resistor thermal noise and op amp noise set the noise floor in typical op-amp circuits. Only when voltages appear across input resistors and cause current to flow does the excess noise become a significant—and often dominant—factor. In general, carbon composition resistors generate the most excess noise. As the conductive medium becomes more uniform, excess noise becomes less significant. Carbon film resistors do better, and metal film resistors do better yet.

Manufacturers specify excess noise in terms of a noise index—the number of microvolts of rms noise in the resistor in each decade of frequency per volt of dc drop across the resistor. The index can rise to 10dB (3 microvolts per dc volt per decade of bandwidth) or more. Excess noise is most significant at low frequencies. Above 100 kHz, thermal noise predominates.

Potentiometers

Trim potentiometers can suffer from most of the phenomena that plague fixed resistors. Users must also remain vigilant against additional hazards unique to these components.

For instance, many trim potentiometers are not sealed and can be severely damaged by board-washing solvents, and even by excessive humidity. Vibration—or simply extensive use—can damage resistive-element and wiper terminations. Contact noise, tempcos, parasitic effects, and limitations on adjustable range can all hamper circuit operation. Furthermore, the limited resolution of wirewound types and the hidden limits to resolution in cermet and plastic types (hysteresis, incompatible material tempcos, slack) make the obtaining and maintaining of precise settings anything but an "infinite resolution" process. Rule: Use infinite care and infinitesimal adjustment range to avoid infinite frustration.

PRINTED-CIRCUIT BOARDS

Printed-circuit boards act as "unseen components" in all precision circuit designs. Since designers rarely consider the electrical characteristics of PC boards as additional circuit components, the circuit's performance usually ends up worse than predicted.

Printed-circuit-board effects that are harmful to precision circuit performance include leakage resistances, voltage drops in ground foils, stray capacitances, dielectric absorption and related "hook" (a salient feature of the circuit's step-response waveform). In addition, the tendency of p-c boards to absorb atmospheric moisture ("hygroscopicity") means that changes in humidity often cause the contributions of some parasitic effects to vary from day to day.

In general, printed-circuit-board effects can be divided into two categories—those that most noticeably affect the static or dc operation of the circuit, and those that most noticeably affect dynamic or a-c circuit operation.

Static PC-Board Effects

Leakage resistance is the dominant static circuit board effect. Contamination of the board's surface, in the form of flux residues, deposited salts, and other debris can create leakage paths between circuit nodes. Even on well cleaned boards, it is not unusual to find 10 nA or more of leakage to nearby nodes from 15-volt supply rails.* Nanoamperes of leakage current into the wrong nodes often cause volts of error at a circuit's output; for example, 10 nA into a 10-megohm resistance causes 0.1 V of error.

To identify nodes sensitive to the effects of leakage currents, ask the simple question: If a spurious current of a few nanoamperes or more were injected into this node, would it matter?

If the circuit is already built, you can localize moisture sensitivity to a suspected node with a classic test. While observing the circuit's operation, blow on potential trouble spots through a simple soda straw. The soda straw focuses the breath's moisture, which, with the board's salt content in susceptible portions of the design, disrupts circuit operation upon contact.

There are several means of eliminating simple surface leakage problems. Thorough washing of circuit boards to remove residues helps considerably. A simple procedure includes vigorously brushing the boards with isopropyl alcohol, followed by a thorough washing with deionized water and an 85°C bakeout for a few hours. Be careful when selecting board-washing solvents, though. If cleaned with Freon-based solvents, some water-soluble fluxes create salt deposits, exacerbating the leakage problem.

Unfortunately, if a circuit displays a sensitivity to leakage, even the most rigorous cleaning can offer only a temporary solution. Problems soon return upon exposure to handling, foul atmospheres, and high humidity. *Guarding*, on the other hand, offers a fairly reliable and permanent solution to the problem of surface leakage. Well-placed guards can eliminate leakage problems, even for circuits exposed to harsh industrial environments.

Guarding principles are simple: Surround sensitive nodes with conductors that can readily sink stray currents, and maintain those conductors at the exact potential of the sensitive node. The guard potential must be maintained close to the potential of the sensitive node, otherwise the guard will serve as a source rather than a sink. For example, to keep the leakage current into a node below a picoampere, assuming 1000-megohm leakage resistance, the guard and the node must be within 1.0 millivolts of one another.

Figures 7a and 7b illustrate the guarding principle as applied to typical inverting and non-inverting op-amp applications. Figure 7c illustrates an actual circuit-board layout for a guard. Note that, to be most effective, the guard pattern should appear on both sides of the circuit board. Try to include the guards when first laying out a new board pattern, from the beginning of the layout process. At later stages, there is usually insufficient space left to locate them optimally—if at all.

Dynamic PC-Board Effects

Although static pc board effects can come and go with changes in humidity or board contamination, problems that most noticeably

*Unfortunately, the standard op-amp pinout places the −15V supply pin right next to the + input, which is often hoped to be at high impedance.

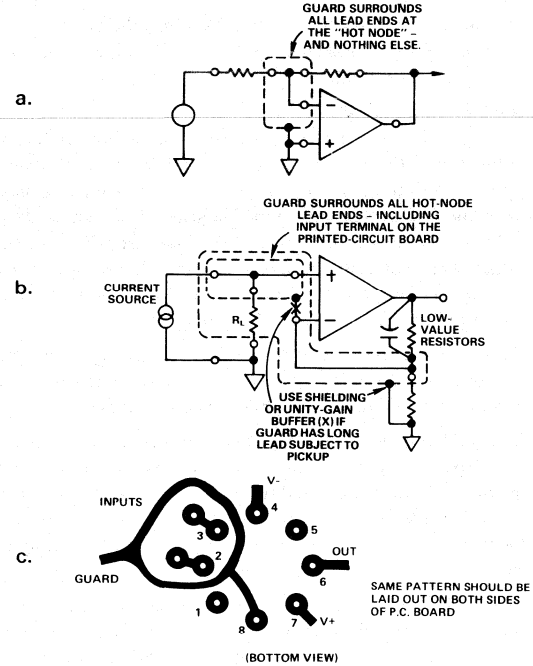

Figure 7. Proper circuit guarding resolves both static and dynamic pc-board induced errors. (a) Use of guard in inverting application. (b) Use of local guard in non-inverting application. Voltage buffer would help in guarding cable. (c) Printed-circuit board guard pattern for op amp.

affect the dynamic performance of a circuit usually remain relatively constant. Short of a new design, they can't be fixed by washing or any other simple fixes. As such, they can permanently and adversely affect a design's specifications and performance.

The problems of stray capacitance, linked to lead and component placement, are reasonably well known to most circuit designers. Since lead placement can be permanently dealt with by correct layout, any remaining difficulty is solved by training assembly personnel to orient components or bend leads in an optimal way.

Dielectric absorption, on the other hand, represents a more troublesome and still poorly understood circuit-board phenomenon. Like dielectric absorption in capacitors, dielectric absorption in a printed-circuit board can be modeled by a series resistor and capacitor connecting two closely spaced nodes (Figure 8). Its effect is inverse with spacing and linear with length. The model's effective capacitance ranges from 0.1 to 2.0 pF, with the resistance ranging from 50 to 500 MΩ. Values of 0.5 pF and 100 MΩ are most common. Consequently, circuit-board dielectric absorption

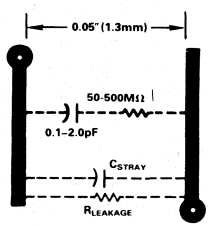

Figure 8. Dielectric absorption plagues dynamic response of pc-based circuits.

interacts the most with high-impedance circuits.

Such dielectric absorption most noticeably influences dynamic circuit response, for example, settling time. Unlike circuit leakage, the effects are not usually linked to humidity or other environmental conditions, but rather, are a function of the board's dielectric properties. The chemistry involved in producing plated-through holes seem to exacerbate the problem. If your circuits do not meet expected transient response specs, you should consider circuit-board dielectric absorption as a possible cause.

Fortunately, there are solutions. As in the case of capacitor dielectric absorption, external components can be used to compensate for the effect. More importantly, surface guards that totally isolate sensitive nodes often completely eliminate the problem (The guards must be duplicated on both sides of the board).

Circuit-board "hook", similar if not identical to dielectric absorption, is characterized by a variation in effective circuit-board capacitance with frequency. In general, it affects the transient response of high-impedance circuits where the board's capacitance is an appreciable portion of the total circuit capacitance. Circuits operating at frequencies below 10 kHz are the most susceptible. As in circuit-board dielectric absorption,the board's chemical makeup very much influences its effects.

DON'T OVERLOOK ANYTHING

Remember, if your precision op-amp or data-converter-based design does not meet specs, try not to overlook anything in your efforts to find the error sources. Analyze both active and passive components, and try to identify and challenge any assumptions or preconceived notions that may blind you to the facts. Take nothing for granted.

For example, when not tied down to prevent motion, cable conductors, moving within their surrounding dielectrics, can create significant static charge buildups that cause errors, especially when connectd to high-impedance circuits. Rigid cables, or even costly low-noise Teflon-insulated cables, are an expensive alternative.

As more high-precision op amps and higher resolution data converters become available, and system designs call for higher speed and accuracy, a thorough understanding of the error sources described in this article becomes more important.

REFERENCES

The following references may help in locating additional information. Available from Analog Devices *only if starred (*)*.

1. Buchanan, James E., "Dielectric Absorption—It Can Be a Real Problem In Timing Circuits." *EDN*, January 20, 1977, page 83.
2. Counts, Lew, and Wurcer, Scott, "Instrumentation Amplifier Nears Input Noise Floor." *Electronic Design*, June 10, 1982.
3. Doeling, W., Mark, W., Tadewald, T., and Reichenbacher, P., "Getting Rid of Hook: The Hidden PC-Board Capacitance." *Electronics*, October 12, 1978, page 111.
4. Fleming, Tarlton, "Data-Acquisition System (DAS) Design Considerations." *WESCON '81 Professional Program Session Record No. 23.*
5. Jung, Walter C., and Marsh, Richard, "Picking Capacitors, Part I." *Audio*, February, 1980.
6. Jung, Walter C., and Marsh, Richard, "Picking Capacitors, Part II." *Audio*, March, 1980.
7. Pease, Robert A., "Understand Capacitor Soakage to Optimize Analog Systems." *EDN*, October 13, 1982, page 125.
8. Rappaport, Andy, "Capacitors." *EDN*, October, 13, 1982, page 105.
9. *Rich, Alan, "Shielding and Guarding." *Analog Dialogue* 17-1, 1983, page 8.
10. Specification MIL C-19978D: *Capacitor, fixed, plastic (or paper and plastic) dielectric, general specification for.*

ANALOG DEVICES

AN-349
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Keys to Longer Life for CMOS

Here's How CMOS Can Be Protected Against Abuses

by Jerry Whitmore

The two principal dangers to analog CMOS (Complementary-symmetry Metal-Oxide Semiconductors) are static electricity and overvoltage (signal voltages exceeding the supply). Both can be effectively dealt with by the aware user.

STATIC ELECTRICITY

The danger from electrostatic voltage buildup is that of "punch-through" of the thin oxide (or nitride) layer that insulates the gate from the substrate, due to accumulation of static charge (V = q/C = 1kV/nC/pF). This danger is minimal in working circuits, because the gate is protected by zener diodes on the chip, which permit depletion of the charge at safe levels. However, during socket insertion, it is possible for a large static charge to exist between the CMOS device and the socket. If the first pin plugged into the socket happens not to be common to the zener-diode protection circuit, the charge on the gate will discharge through the oxide layer, destroying the device. These four steps will help the device survive the system assembly stage.

1. Keep unused CMOS devices in the black conductive foam in which they were shipped to prevent charge buildup between pins.

2. Ground the operator, who is inserting the devices, to the system power ground with a plastic ground strap.

3. Before pulling the CMOS device from the protective foam, ground the foam to power common to deplete accumulated charge.

4. After the circuit has been inserted into a circuit board, keep the board grounded or shielded when carrying it around.

THE SCR "LATCH"

When working with analog CMOS circuits, the safest practice is to make sure that no analog or digital voltage applied to the device exceeds the supply voltage, and that the supply voltage itself is within ratings. Nevertheless, occasions do arise where it is necessary to tolerate overvoltage. Protection is possible in most cases, if the mechanism of failure is understood.

Reprinted from Analog Dialogue 8–2 1974

Figure 1 shows the circuit and cross-section of a typical CMOS output switching element. From the connections and associations of the various elements and regions, we can draw an equivalent diode circuit (Figure 2). If the analog input voltage at either the S or the D terminals exceeds the power-supply

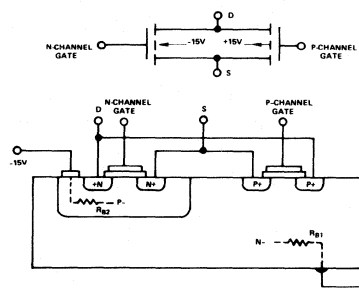

Figure 1. Cross-section of CMOS switching element

voltages, the parasitic transistors formed by the various diode junctions are placed in the forward bias mode. These parasitic NPN and PNP transistors appear to form the SCR ("silicon controlled-rectifier") circuit shown in Figure 3. Overvoltage

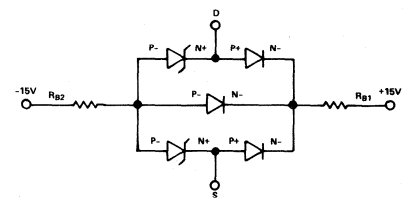

Figure 2. Diode equivalent circuit of CMOS switching element

can cause excessive current and metallization failure. Normally, the outputs of op amps are used as the voltage sources feeding the S or D terminals, so the currents cannot exceed the op amps' dc output current limitations. Nevertheless, it is still possible for transient induced currents to destroy the CMOS device; protection is therefore desirable.

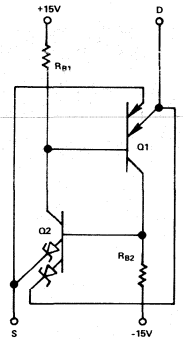

IF, FOR EXAMPLE, A POSITIVE OVER-VOLTAGE IS APPLIED TO THE DRAIN TERMINAL, THE BASE OF Q1 WILL CONDUCT WHEN ITS EMITTER VOLTAGE EXCEEDS V_{DD} BY ONE BASE-EMITTER DROP. THE COLLECTOR CURRENT WILL INCREASE TO A LEVEL LIMITED ONLY BY THE V_{SS} AND V_D CURRENT LIMITATIONS. SINCE THE METAL INTERCONNECT TO THE SUPPLY TERMINAL IS NORMALLY DESIGNED TO HANDLE SMALL CURRENTS, THE CURRENT DENSITIES INVOLVED CAN CAUSE DEVICE FAILURE.

Figure 3. Parasitic transistor action in CMOS switch

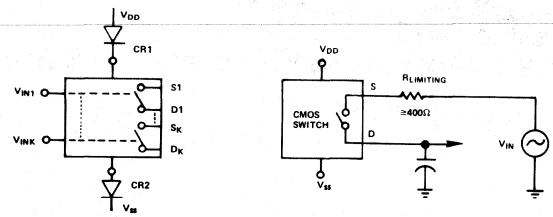

a. Diode protection

b. Current-limiting protection

Figure 4. Circuit protection schemes

Figure 4 illustrates a means of preventing turn-on of the parasitic transistors by means of diodes (say 1N459's) in series with the supply leads. If the S or D terminal is at a higher-than-supply voltage, CR1 and/or CR2 are reverse-biased and base drive is unavailable to turn the transistors on. A separate pair of diodes should be used for each CMOS device to be protected. Though powerful, the method is not infallible. If one terminal of the switch is tied to a negative potential (e.g., a charged capacitor), and the other terminal is raised above V_{DD}, the avalanche diode at one emitter of Q2 is sufficient to supply enough base drive to turn Q2 on, despite the protective diodes. For such a situation, a current-limited supply, or resistance in series with the capacitor is necessary.

If transient overvoltages are expected at the S or D terminals, 300–400Ω in series with the terminal to be fed by the voltage source is suggested (Figure 4b).

ANALOG DEVICES

AN-311
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

How to Reliably Protect CMOS Circuits Against Power Supply Overvoltaging

by Mike Byrne

INTRODUCTION

All IC processes have an intrinsic breakdown voltage associated with them, and this results in a maximum voltage stress which can be applied to any device fabricated on that process. As a result, all IC manufacturers give an Absolute Maximum Ratings specification for their devices. This generally gives the maximum voltage which can be applied to any of the pins on the device. Overvoltaging a device means that stresses or voltages in excess of the absolute maximum ratings are applied to the device. This application note deals in particular with overvoltage of the power supply inputs to CMOS and linear-compatible CMOS devices.

The intrinsic breakdown voltage associated with an IC process means that a transistor, buried Zener or other such element on the process will have a defined breakdown voltage. Obviously, if only one such element appeared between the positive supply input (V_{DD}) and negative power supply input (V_{SS}) of a device, the absolute maximum rating voltage, for $V_{DD} - V_{SS}$, would be the breakdown voltage of the element. It is not always possible to ensure that $V_{DD} - V_{SS}$ does not appear across any single element within the IC as it is often compromised by the required IC functionality, die size constraints and other factors. This means that the device manufacturer will be left with a finite voltage which can be applied across the supplies before destruction of the device results. Thus the device manufacturer will determine the limit and specify an absolute maximum rating on the data sheet for the part which is safely inside the breakdown voltage. Users of the part must ensure that the operating voltages applied to the device are within the absolute maximum ratings. So where's the problem?

POWER SUPPLY SPIKES

The problem arises, not with the steady-state value of the power supplies which is easily controlled, but with voltage spikes on the power supply lines. The most likely place for these voltage spikes to occur in most systems is during turn-on and turn-off of the power supplies. Other potential sources of voltage spikes are switching mode power supplies or when operating the devices in noisy environments, such as in the presence of large motors. During these times, depending on the output impedance of the power supplies, the load presented to the power supplies and the overall design of the power supply, the power supply voltage may significantly overshoot its nominal value and in doing so exceed the absolute maximum ratings of the device (see Figure 1).

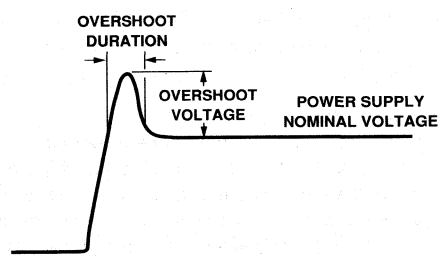

Figure 1. Power Supply Turn-On Spikes

Historically, external clamp elements such as Zener and Schottky diodes have been used to limit voltage spikes to a duration short enough to prevent any damage to the device. However, as geometries of CMOS and linear-compatible CMOS processes shrink, the devices fabricated on them are getting faster. This means that in normal operation there is a significant benefit to the user in terms of speed, bandwidth, etc. However, it also means that elements fabricated on the faster process will respond to much shorter power supply transients. It is not uncommon for devices to respond to power supply transients which are of the order of 1μs duration or less. This means that traditional methods of protecting devices using Zener and Schottky diodes no longer reliably protect the device. This is because their response time to a voltage spike or transient is now, in many cases, slower than what the device's response time is, and therefore they do not provide any protection.

TRANSIENT VOLTAGE SUPPRESSORS

In order to provide adequate protection, the protecting element needs to have a very fast response time to voltage spikes or transients. Devices which are ideal for this function are transient voltage suppressors or TransZorbs.* TransZorb suppressors are PN Silicon transient voltage suppressors that are characterized by their surge handling capability, their extremely fast response time and low series resistance. Their response time can be as low as 1ns while their clamping ratio (ratio of clamp voltage to nominal voltage) is low. The manner in which the TransZorb is used to protect the device depends on the actual application, and a number of different situations are discussed in the following sections. In all cases, the TransZorb should be placed as close as possible to the device it is protecting to reduce the resistance between the TransZorb and the device.

Single Supply Systems

The first case is where the CMOS device to be protected is operated from a single power supply rail. Obviously, in this case, the breakdown path is between this single power supply rail and the ground rail of the device. In such cases, a single transient suppressor, connected between the rail and ground as per Figure 2, is sufficient to reliably protect the device from turn-on and turn-off transients. This single TransZorb will protect the device from transients on either the power supply rail or on the ground rail by clamping the voltage differential between them. The TransZorbs come in a variety of different voltage ratings, and Table I gives the recommended TransZorb part number for some commonly used power supply voltages.

Dual Supply Systems

The second case is where the device is powered from two supply rails. In this case, a number of potential breakdown paths exist. The first is between the positive and negative supply rails, the second between the positive supply rail and ground and the third between the negative supply rail and ground. For dual supply systems there are two possible protection schemes.

For some cases, a single TransZorb, connected between the two supply rails as in Figure 3, is sufficient to protect the device. In this case, the TransZorb voltage rating is equal to the sum of the two power supply voltages. This single TransZorb arrangement assumes that when a

*TransZorb is a registered trademark of General Semiconductor Industries, Inc.

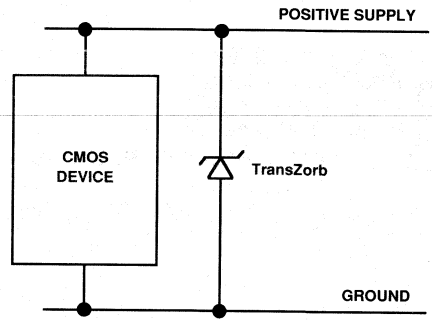

Figure 2. Using TransZorbs in Single Supply Systems

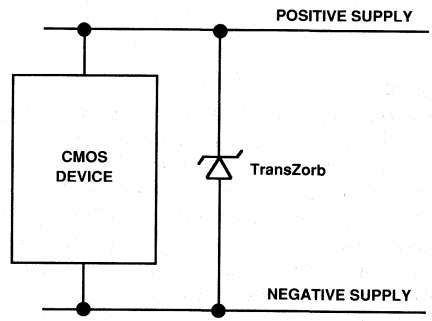

Figure 3. Dual Supply, Single TransZorb Arrangement

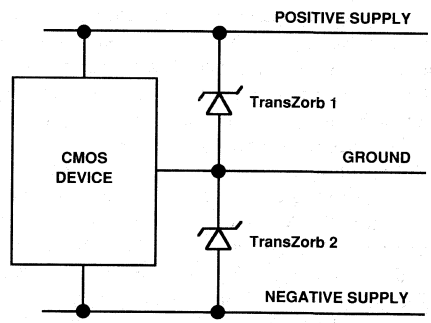

Figure 4. Dual Supply, Two TransZorb Arrangement

Power Supply Voltage	TransZorb Part Number (JEDEC Type Number)	Reverse Stand-Off Voltage	Maximum Clamping Voltage @ 1A†	Maximum Clamping Voltage @ 10A†
+5 V Supply	1N6373	+5 V	+7.1 V	+7.5 V
−5 V Supply	1N6373	−5 V	−7.1 V	−7.5 V
+12 V Supply	1N6376	+12 V	+16.1 V	+16.5 V
−12 V Supply	1N6376	−12 V	−16.1 V	−16.5 V
+15 V Supply	1N6377	+15 V	+20.1 V	+20.6 V
−15 V Supply	1N6377	−15 V	−20.1 V	−20.6 V

†This is the peak pulse current.

Table I. Recommended TransZorbs for Figure 2 and Figure 4

Power Supply Voltage	TransZorb Part Number (JEDEC Type Number)	Reverse Stand-Off Voltage	Maximum Clamping Voltage @ 1A†	Maximum Clamping Voltage @ 10A†
±5 V	1N6375	10 V	13.7 V	14.1 V
+5 V, −12 V	1N6378	18 V	24.2 V	25.2 V
+5 V, −15 V	1N6379	22 V	29.8 V	32.0 V
+12 V, −5 V	1N6378	18 V	24.2 V	25.2 V
±12 V	1N6282A	25.6 V	32.0 V‡	33.0 V‡
+15 V, −5 V	1N6379	22 V	29.8 V	32.0 V
±15 V	1N6484A	30.8 V	38.5 V‡	39.5 V‡

†This is the peak pulse current.
‡These are typical numbers. TransZorbs are available with lower clamping voltages than the JEDEC part numbers given in the table above.

Table II. Recommended TransZorbs for Figure 3

spike occurs on one of the supplies that the other supply presents a low output impedance. It also assumes that the supplies are capable of sinking or sourcing the additional current which flows when a spike occurs. If these assumptions hold, then all three breakdown paths are protected when the spikes occur on either of the power supplies. For example, if the positive supply is +5 V and the negative supply is −15 V, the TransZorb rating will be 20 V. Assuming there is a +15 V spike on the +5 V supply, the TransZorb will react to this by absorbing the spike and sinking the current down to the −15 V supply. This means that the −15 V supply must be able to sink the additional current. If it can, the −15 V supply will not move from its −15 V level during the spike, and the TransZorb will ensure that the positive supply will not be more than 20 V above this. This means that the positive supply does not move above +5 V, the negative supply does not move from −15 V, and all three breakdown paths are protected. Table II gives some examples of recommended TransZorbs for a number of common dual power supply voltages.

There are two situations where the single TransZorb scheme, across the positive and negative supplies, will not always protect the device. The first of these, already mentioned above, is where one of the supplies cannot sink or source the additional current which flows during a spike or when either of the supplies does not present a low output impedance during the spike. It should be noted here that while the power supply may present a low output impedance under steady-state conditions, its impedance during turn-on or turn-off may be different. Using the same example as before, if the −15 V supply cannot sink the required spike current, it means that the negative supply is no longer at −15 V. The TransZorb will still clamp the two supplies together so the breakdown path between the two supplies is protected. However, the absolute value of either supply is no longer clearly defined. If we assume that the −15 V goes to, say, −5 V because it cannot sink the current, it means that the positive supply goes to +15 V. This will possibly exceed the breakdown path between the positive supply and ground and lead to damage to the device.

The second situation is where the spikes occur on the ground line and not on either of the power supply rails. A positive spike of +10 V on the ground line, for example, would not cause the TransZorb to turn on, yet the voltage difference between ground and the negative supply is now 25 V, and this could exceed the breakdown voltage between the negative rail and ground.

In these situations, the two TransZorb scheme, shown in Figure 4, is recommended. This two TransZorb scheme protects against voltage spikes on the ground and also protects the device in cases where either of the supplies is not capable of sinking or sourcing the additional current which flows during the spike. This scheme ensures that each supply is independently protected and the spike current flows to ground. This scheme protects the breakdown path from each supply in the same manner as outlined for single supply systems. By protecting each supply, the breakdown between the two supplies is also protected, and so all three possible breakdown paths are guarded against. Suitable TransZorb values for common power supply voltages are as per Table I.

SWITCHING-MODE POWER SUPPLIES/NOISY ENVIRONMENTS

Another situation, other than power on, where spiking of the supplies can occur is in applications using switching-mode power supplies. Although the switching mode power supply may be regulated, the regulation may not be sufficient to remove spikes as low as 1 μs duration. If these spikes have sufficient magnitude and energy, they also can cause damage to the device. Other potential problem areas are applications in environments which are inherently noisy and produce spikes on power supply and ground rails. Examples of this type of application are operating devices in the presence of large motors or operating the devices in industrial environments. The schemes recommended in Figures 2, 3 and 4 for turn-on/turn-off spikes are equally applicable to protecting against switching-mode power supply spikes or power supply spikes generated in noisy environments.

The TransZorb schemes shown in this application note protect the device against overvoltaging of the power supplies. It will not prevent damaged being caused to the device when recommended power supply sequencing is not obeyed. Consult the absolute maximum ratings section in the manufacturer's data sheet to see if there is a particular power supply sequence for a device or if the digital inputs cannot be powered before the supplies.

ANALOG DEVICES

AN-350
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Take the Guesswork Out of Settling-Time Measurements

by Barry Harvey

By building an oscilloscope input amplifier and a step generator, you can use off-the-shelf laboratory oscilloscopes to make precise measurements of the settling times of fast linear circuits.

Because linear circuits are getting faster and faster, it's becoming increasingly difficult to accurately measure such circuits' dynamic performance. Available test equipment is inadequate to measure the dynamic performance of these circuits, which have settling times as fast as 100 nsec. Such equipment either lacks the requisite sensitivity or has poor overload characteristics, which mask the effects you're trying to observe (see **box,** "Sources of error in fast-linear-IC measurement").

To test a device accurately, you must use test equipment that has much better resolution than the device under test will ever require. To test fast linear circuits, you have to observe small effects superimposed on a very large signal change. But if you increase oscilloscope sensitivity enough so you can see these small effects, you will almost certainly produce overload conditions in the oscilloscope. Off-the-shelf oscilloscopes do not provide the resolution you need at the extreme ends of a large step signal, and they won't be able to handle the overload conditions.

To make an accurate determination of settling time—especially for op amps and fast D/A converters—you need to build two items of special test equipment: a step-function generator (**Fig 1**) and an oscilloscope input amplifier (**Fig 2**).

Fig 1's step-function generator lets you evaluate the recovery characteristics of the available oscilloscopes. Further, by providing a reference signal that has precisely known characteristics, the generator will eliminate guesswork in settling-time measurement.

The oscilloscope input amplifier in **Fig 2** provides a resolution to 200 μV/div on a voltage-input range of ±10V, or 200 nA/div on a current-input range of ±10 mA. This input amplifier will allow you to make 0.01% settling determinations without having to estimate the magnitude of small and questionable oscilloscope images. This test system provides settling-time measurements to 50 nsec for 8-bit measurements and will allow you to perform dynamic tests on 12-bit devices at full operating speeds.

Test fixture removes sources of error

Fig 2 shows the complete test fixture. This test fixture, which will eliminate most sources of error, consists for the most part of discrete components and CA3127 transistor arrays. The main amplifier chain consists of cascaded differential amplifiers IC_1 through IC_4, with emitter followers between each stage to

Reprinted from EDN – September 19, 1985

Off-the-shelf scopes don't provide the resolution you need at the ends of a large step signal.

perform buffering and level shifting. Although these amplifiers overload in the presence of large inputs, they do not saturate, and they recover from overload within approximately 15 nsec.

It's essential that you place each amplifier section in a separate, well-shielded compartment. If the shielding is not adequate, the entire amplifier chain may oscillate at several hundred megahertz.

The V_{IN} path has a gain of 0.5 between the input terminal and the summing junction. IC_1 and IC_2 provide most of the system gain; you set the product of their gains to a value between 4 and 100 by varying the bias currents fed into the differential amplifiers. IC_3 and IC_4 together form a linear output driver with a gain of 2 and a differential output swing of no more than ±160 mV; thus, an oscilloscope set for an A+B display of vertical inputs with B inverted and a sensitivity of 20 mV/div needs to handle only two screens' worth of signal; this small input range does not overload the oscilloscope input circuitry. All signals are measured relative to ground, not to some bias point.

The offset voltage is provided by transistors Q_1 through Q_7 and their related components. Transistor Q_2, the offset voltage source, operates at a bias of 25 mA to reduce its output impedance, and it has a feedback loop to reduce the output impedance still further at high frequencies. Transistor Q_8 acts as a buffer between Q_2 and amplifier IC_9, which, with three analog switches of IC_7, constitutes an S/H amplifier that stores the offset level.

The test fixture has an autozero feature, which offsets the final output value of the device under test so that it falls on the oscilloscope baseline. The test fixture cycles continuously through the four time states shown

Continued on pg 182

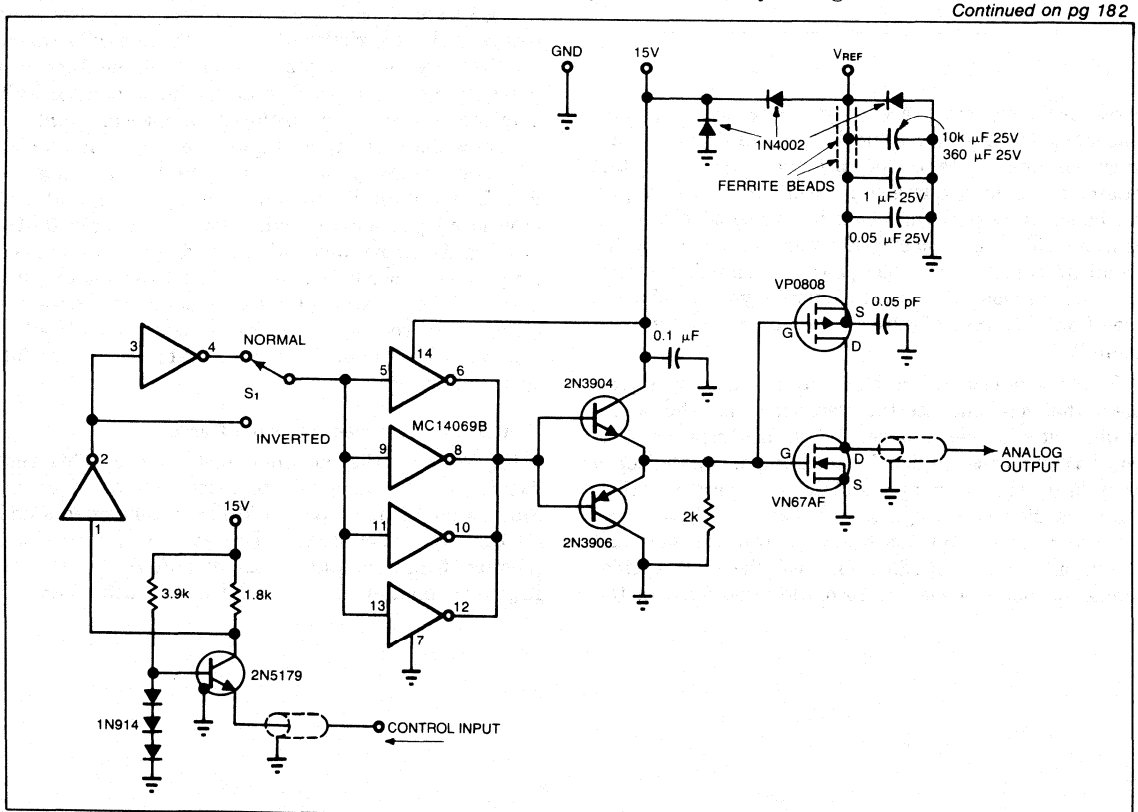

Fig 1—Use this perfect-step generator *for checking your oscilloscope. Output power VMOS transistors Q_4 and Q_5 exhibit no saturation effects and switch the analog output cleanly to ground or to V_{REF} (nominally 10V). The step rate is greater than 1500 V/µsec.*

Sources of error in fast-linear-IC measurement

It's often difficult to measure the performance of fast linear ICs, mainly because off-the-shelf laboratory oscilloscopes either can't handle the sensitivity you require or mask the effects you're trying to observe.

When you use off-the-shelf equipment to measure the performance of a fast linear circuit, your measurements may be inaccurate because such equipment doesn't allow you to see the actual shape of a waveform. And you can't accurately estimate the actual shape of a waveform by assuming that it follows an exponential curve; if you do so, your estimation may be wildly inaccurate.

Although many components of analog waveforms are exponential—overshoot, ringing, and the response produced by feed-forward compensation fall into this category—waveshapes are often composed of components that are not exponential, although they may seem to be. Overload recovery, for example, is almost never exponential, and thermal tails can have almost any shape.

Fig A shows the difference between an actual waveform and an expected exponential curve. The lower diagram in **Fig A** shows an analog signal moving from some initial value to a higher level and then settling toward a final value. The area within the dashed lines is expanded by a factor of 2000 in the upper diagram. You'll see that overshoot and recovery components cause the waveform to have a tail that crosses the lines representing the ±1-mV acceptable settling window more than once.

The time that elapses between the start of the rising edge and the last entry into the settling window is called T_{SETTLE}. The dashed curve in **Fig A** approximates the sum of all exponential components of the waveform; if the real waveform followed this exponential curve, you could easily determine settling time.

However, most electronic circuits exhibit, after slewing, a slow-settling component that's not exponential. One of the most troublesome components of this kind is the tail generated by a device input stage that's driven into overload by a step function. It's hard to determine the difference the tail makes, because a wide range of overloads will produce the same recovery characteristics—for example, a 10V step and a 20V step might both produce an identical 1-mV settling component, which would appear as a 0.01% tail on the 10V step, but a 0.005% tail on the 20V step.

Although a 0.005% tail on the 20V step appears to indicate better system performance, the performance is actually unchanged. Other causes of tails are thermal changes, the destabilization of a power supply or of bias, and the characteristics of a device that's recovering from saturation; tails due to these causes are not symmetrical for input steps of opposite slope.

The pole/zero pairs in op amps that use feed-forward compensation are particularly troublesome. Many such amplifiers have settling times of 1 μsec or less for settling accuracies of 1 or 0.1%. However, when these amplifiers must settle to an accuracy of 0.01%, the settling time increases to 3 μsec or more.

This phenomenon is significant in oscilloscope probes that attenuate, because such probes often use pole/zero pairs to smooth out response aberrations over a wide frequency range, but in the process they produce linear tails of 1% or more.

Noise, too, can be a source of confusion. Noise peaks that constantly cross the settling window may make you think the device under test hasn't settled, even when the device's contribution to the waveform is constant. Actually, 16-bit D/A converters never really do settle; to estimate their performance, you just have to smooth the output wave and ignore the noise.

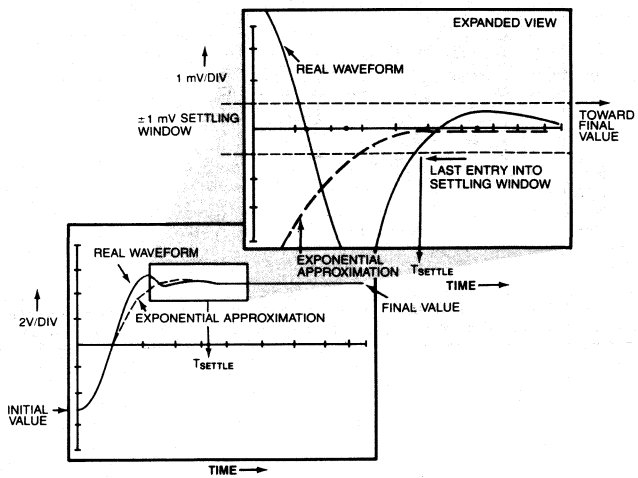

***Fig A—This typical settling waveform** shows what you can expect to see on a low-sensitivity oscilloscope (bottom) and an expanded view (top) of the details that are difficult to measure.*

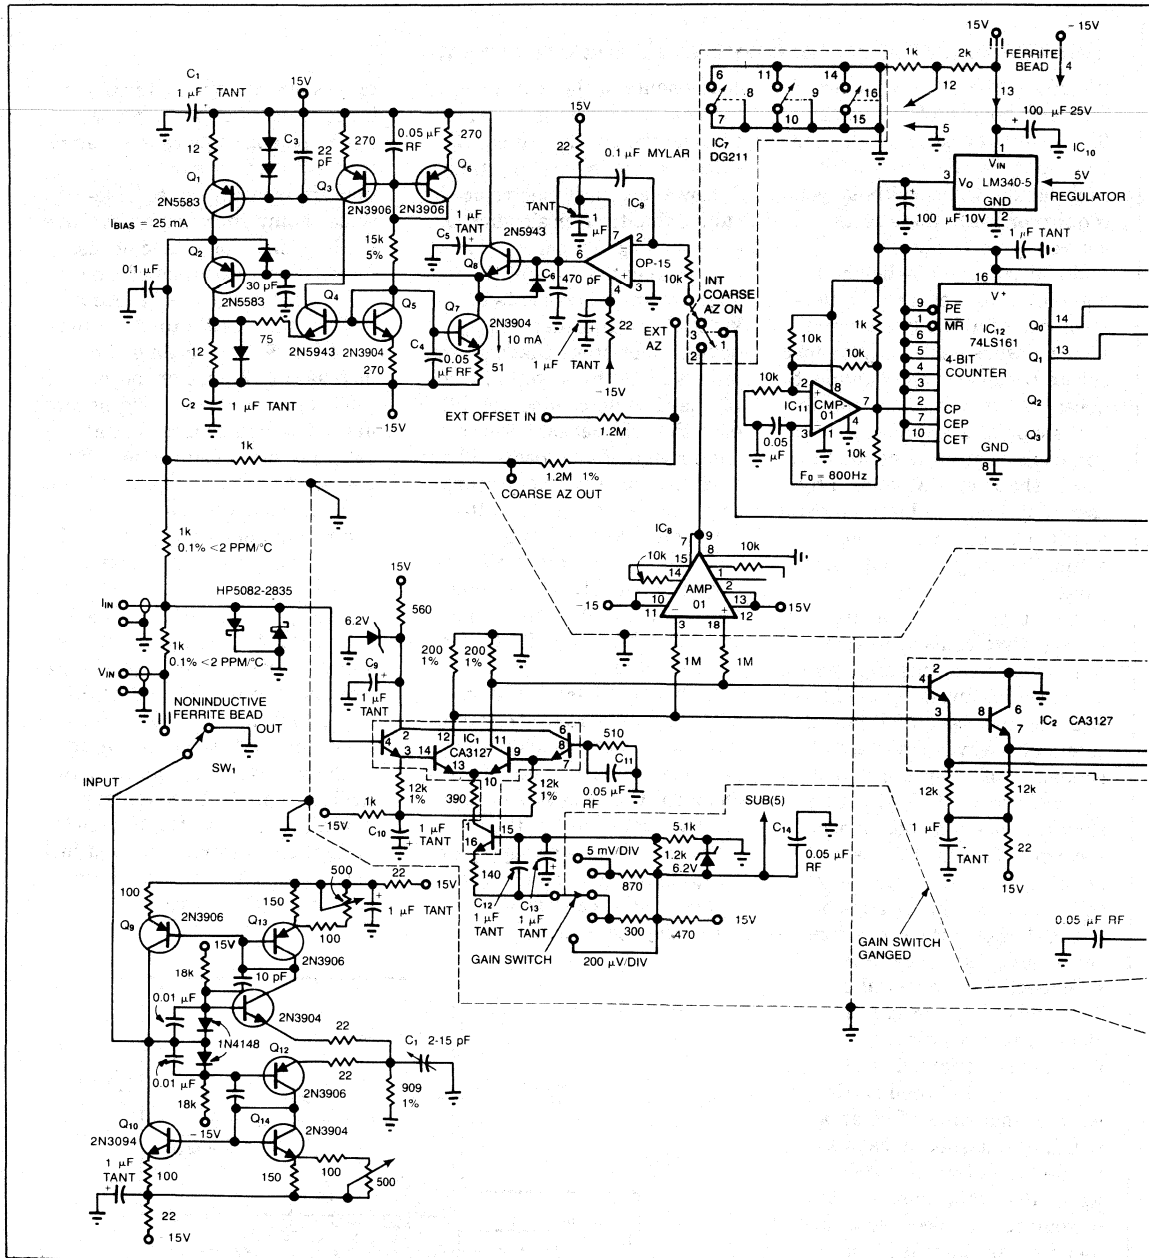

Fig 2—This settling-time test fixture *has an autozero feature, which ensures that the final output of the device under test will fall on the oscilloscope baseline. The fixture has special protection against overloading of the oscilloscope input circuits.*

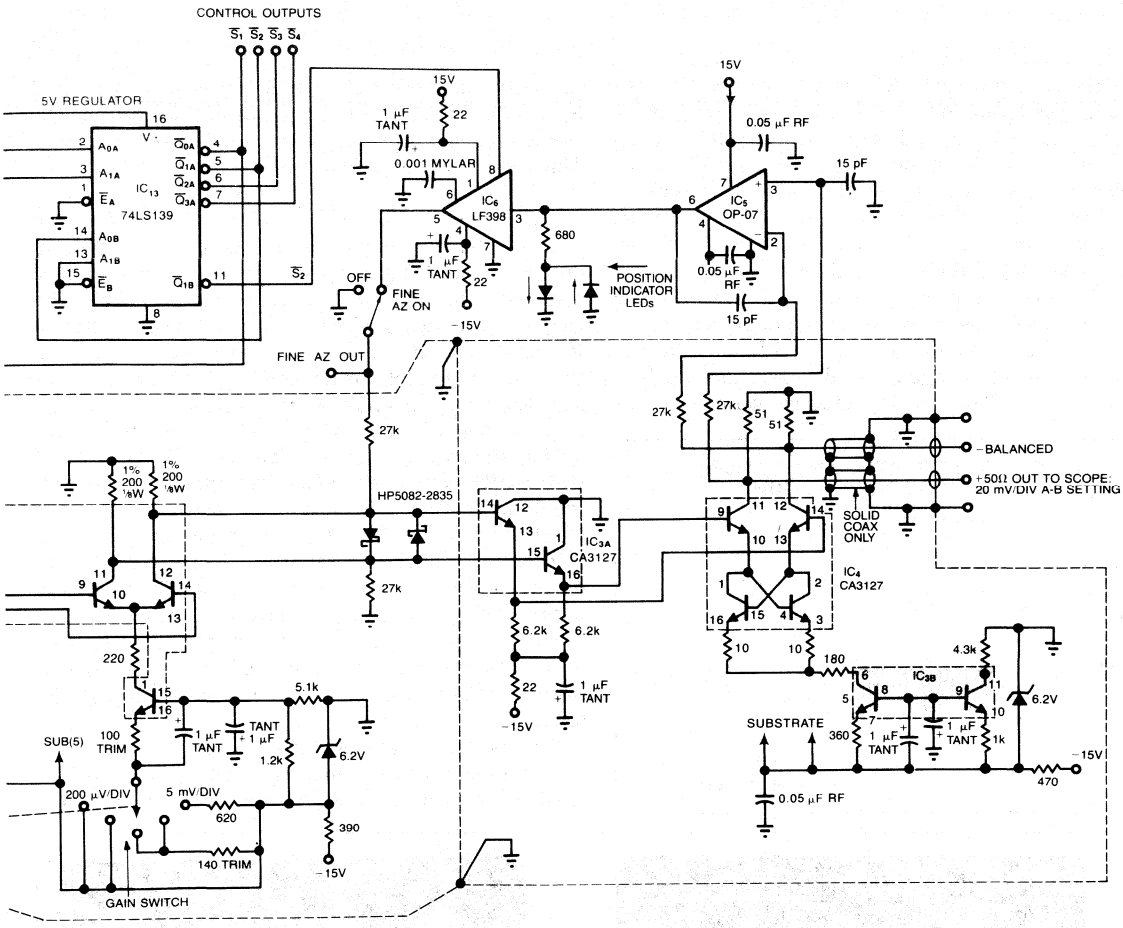

NOTES:
1. AMPLIFIER SECTIONS MUST BE ENCLOSED IN SEPARATE, WELL-SHIELDED COMPARTMENTS.
2. ALL DIODES ARE SIGNAL DIODES EXCEPT WHERE MARKED.
3. RF—ANY GOOD HIGH-FREQUENCY CAPACITOR.

If you assume that settling tails too small to see clearly are exponential, your estimate may be wildly inaccurate.

in **Table 1.** Comparator IC_{11} acts as an 800-Hz clock oscillator that generates the time states with the aid of IC_{12}, a 4-bit counter, and IC_{13}, a decoder. IC_{12} divides and squares the clock signal; IC_{13} provides four nonoverlapping but closely synchronized state signals.

The repetition rate of the complete cycle (S_1 through S_4) is set at approximately 200 Hz to allow the oscilloscope to develop sufficient brightness. During the first state (S_1), when the device under test is at its final value, IC_8 measures the unbalance of amplifier IC_1.

The output of IC_8 drives IC_9 (an S/H amplifier) to change the offset voltage in the direction that tends to rebalance IC_1. This coarse autozero voltage settles toward the end of time state S_1. Exactly at the end of S_1, IC_9 (in conjunction with IC_7) samples and holds the voltage on capacitor C_7 (C_7 is in IC_9's feedback loop). This voltage acts as a reference voltage for the succeeding time states.

Although it's stable to several tens of microvolts, the coarse autozero voltage is not accurate enough to ensure that the final output value of the device under test will be offset onto the oscilloscope baseline. The largest error occurs when the fourth analog switch of IC_7 disconnects IC_8 from IC_9, delivering a charge from the switch to the integrator input and causing a step of approximately 1 mV. To correct for the effects of this step, IC_5 and IC_6 force IC_3 and IC_4 to come to balance during state S_2 by injecting feedback. At the end of state S_2, IC_6 samples and holds the voltage that brought the fixture output to balance; input baseline variations of ±10V cause less than 20 μV of output baseline shift.

In state S_3, the fixture sends the device under test to its initial value; at this time, the amplifiers of the fixture will probably be overloaded, but will recover rapidly. In state S_4, the device under test settles back toward its final value and the oscilloscope is triggered so that you can observe the settling event through the fixture.

The test fixture accepts either a current input (connected directly to the summing junction) or a voltage input (connected via a 1000Ω resistor). The effective

TABLE 1—TIME STATES AND ACTIONS FOR AUTOZERO FUNCTION

TIME →

STATE	S_1	S_2	S_3	S_4
VALUE FOR DEVICE UNDER TEST	FINAL	FINAL	INITIAL	FINAL
SYSTEM FUNCTION	COARSE AUTOZERO	FINE AUTOZERO	STARTING MARK (TEST FIXTURE IS OVERLOADED)	DISPLAY SETTLING CHARACTERISTIC

↑ SCOPE TRIGGER

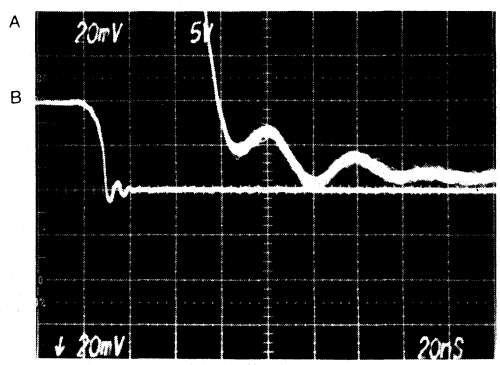

(a)

TRACE	VERTICAL	HORIZONTAL
A	1 mV DIV	20 nSEC DIV
B	5V DIV	20 nSEC DIV

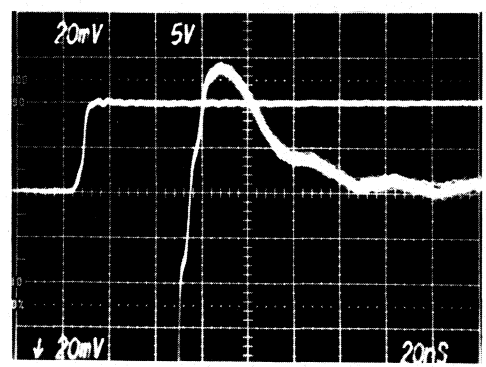

(b)

TRACE	VERTICAL	HORIZONTAL
A	5V DIV	20 nSEC DIV
B	1 mV DIV	20 nSEC DIV

***Fig 3—The step-function generator settles to 0.01%** within 70 nsec toward ground **(a)** and toward 10 V **(b)**. These scope photos show settling times of the step generator and the test fixture alone.*

impedance of R_{IN} and R_{OFF} in parallel is 500Ω, and the time constant is approximately 7 nsec. (A lower impedance would improve settling speed, but some sources, which have difficulty in driving even the 1000Ω input impedance at the voltage input, would not tolerate any decrease.)

Transistors Q_9 through Q_{14} alleviate this difficulty by simulating a negative resistance of −1000Ω, which can be switched in parallel with the 1000Ω voltage-input resistor. When actuated, the negative-resistance simulator raises the impedance at the voltage input to approximately 30,000Ω and provides whatever current is needed to assist the device under test in placing its normal output current through the input resistor.

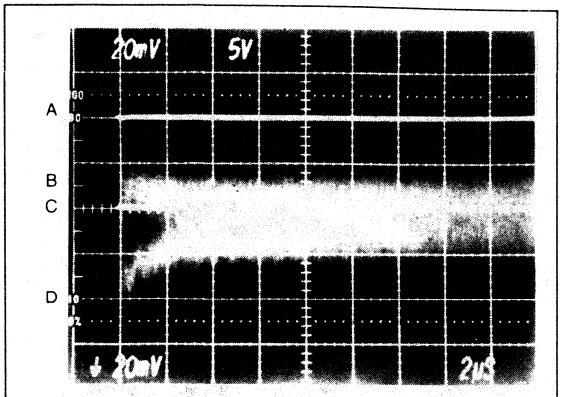

TRACE	VERTICAL	HORIZONTAL
A	5V/DIV	2 μSEC/DIV
B	200 μV DIV	2 μSEC/DIV
C	5V DIV	2 μSEC/DIV
D	200 μV DIV	2 μSEC/DIV

***Fig 4—The long-term step-generator tail is approximately 80 μV** when settling to ground and approximately 300 μV when settling to 10V.*

Determine settling times accurately

Fig 3 shows the settling times—both toward 10V and toward ground—that you obtain when the edge generator is driving the test fixture. The combined settling time of both units to within 0.01% of 10V is 70 nsec or less, measured from the point where ringing is no longer visible after slewing (at the second major division on curve A) to the last crossing of the 1-mV division (curve B). It's difficult to know what each unit actually contributes to the combined settling time; however, it seems probable that the fixture settles within 50 nsec, and that, beyond the 40-nsec mark, the display shows the real settling tail of the edge generator plus a small, slow, thermal response from the fixture.

The size of the tail, as you can see from **Fig 4,** is approximately 80 μV when the signal is settling to

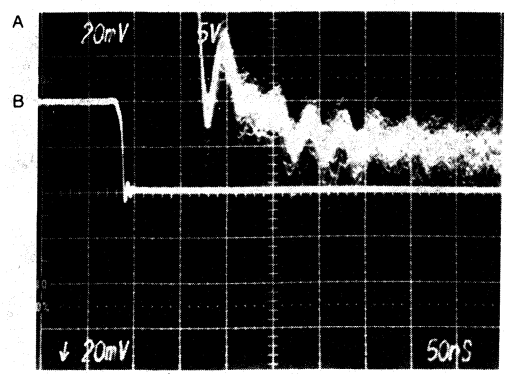

(a)

TRACE	VERTICAL	HORIZONTAL
A	200 μV/DIV	50 nSEC/DIV
B	5V/DIV	50 nSEC/DIV

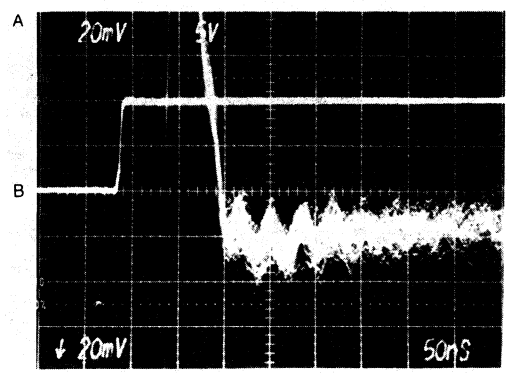

(b)

TRACE	VERTICAL	HORIZONTAL
A	200 μV/DIV	50 nSEC/DIV
B	5V/DIV	50 nSEC/DIV

***Fig 5—Even at 0.002% (16-bit) resolution,** settling time for the step-function generator is only 150 nsec toward ground **(a)** or toward 10V **(b)**.*

Overload recovery is almost never exponential.

ground and 300 μV when it's settling to 10V. A supply-voltage glitch occurs on each transition of the output, so the edge generator probably creates some of the tail.

When you're applying a large signal to the input of the test fixture, you may find that the signal source develops a long settling tail in driving the 1-kΩ resistors. In that case, switch in the negative-resistance simulator.

When it's driven by a fast signal source, the negative-resistance-simulator circuit can cause short-term ringing. Thus, when you're testing very fast devices,

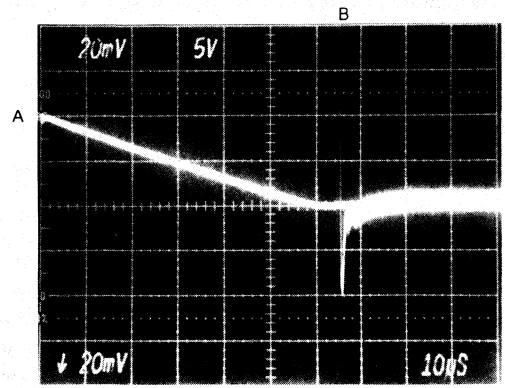

(a)

TRACE	VERTICAL	HORIZONTAL
A	5V/DIV	10 μSEC/DIV
B	1 mV/DIV	10 μSEC/DIV

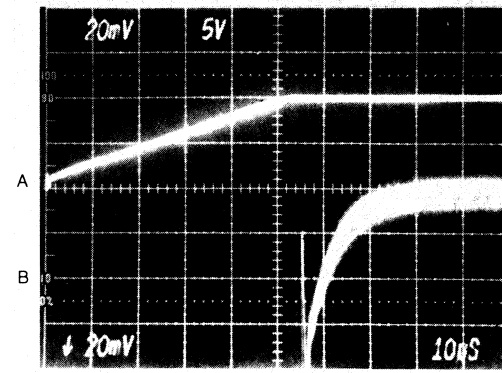

(b)

TRACE	VERTICAL	HORIZONTAL
A	5V DIV	10 nSEC DIV
B	1 mV DIV	10 nSEC DIV

Fig 6—A standard high-accuracy op amp settles *in 65 μsec when operating as a unity-gain follower. Settling is complete 8 to 10 μsec after slewing is finished.*

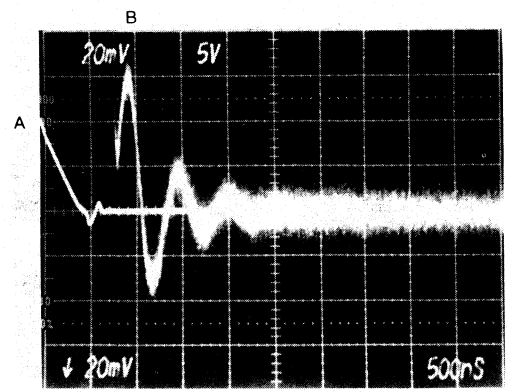

(a)

TRACE	VERTICAL	HORIZONTAL
A	5V/DIV	500 nSEC/DIV
B	1 mV/DIV	500 nSEC/DIV

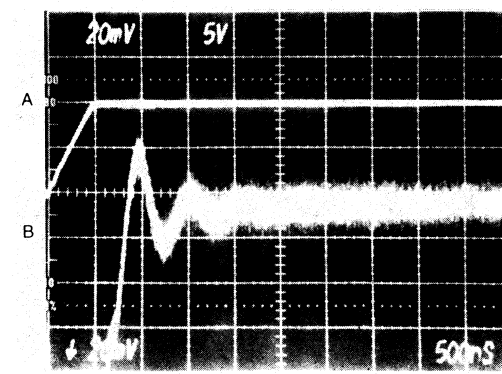

(b)

TRACE	VERTICAL	HORIZONTAL
A	5V DIV	500 nSEC DIV
B	20 mV DIV	500 nSEC DIV

Fig 7—A fast FET op amp settles to 0.01% *within 1.5 μsec toward ground* ***(a)*** *and toward 10V* ***(b)****. The last ringing peak that grazes the ±1 major vertical division is taken as the 0.01% settling point.*

It's essential that you place each amplifier section in a separate, well-shielded compartment.

you may have to use the normal 1000Ω input for the short-term measurements and then switch in the impedance buffer for an accurate display of the long-term effects. **Fig 5**'s photos, which were taken with the negative-resistance simulator inactive, show that even with 16-bit resolution, the test system's settling time is less than 150 nsec.

You can use **Fig 2**'s test fixture to test op amps as well as D/A converters. The scope photos in **Figs 6** and **7** show the settling times of a standard, high-accuracy op amp (OP-07) and a fast FET op amp (OP-16). **Fig 8** shows the settling times of two 565A 12-bit D/A con-

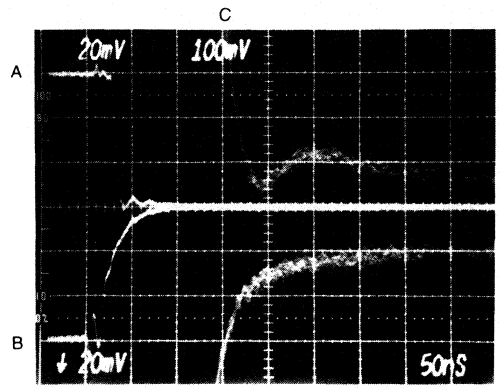

TRACE	VERTICAL	HORIZONTAL
A	0.5 mV/DIV	50 nSEC/DIV
B	0.5 mV/DIV	50 nSEC/DIV
C	0.5 mV/DIV	50 nSEC/DIV
D	0.5 mV/DIV	50 nSEC/DIV

(a)

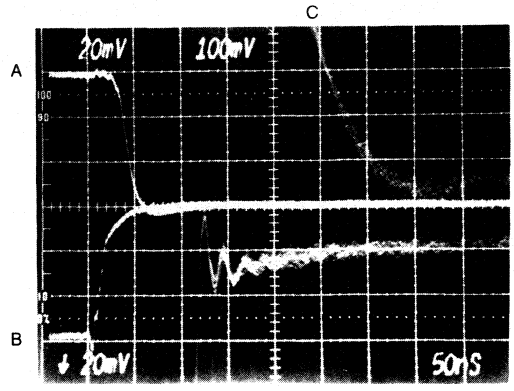

TRACE	VERTICAL	HORIZONTAL
A	0.5 mV/DIV	50 nSEC/DIV
B	0.5 mV/DIV	50 nSEC/DIV
C	0.5 mV/DIV	50 nSEC/DIV
D	0.5 mV/DIV	50 nSEC/DIV

(b)

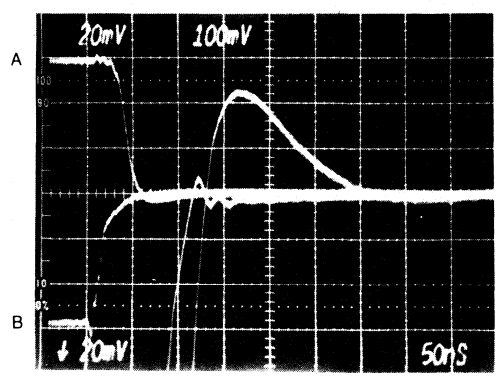

TRACE	VERTICAL	HORIZONTAL
A	2 mV/DIV	50 nSEC/DIV
B	2 mV/DIV	50 nSEC/DIV
C	2 mV/DIV	50 nSEC/DIV
D	2 mV/DIV	50 nSEC/DIV

(c)

Fig 8—Two 565A D/A converters (a and b) *exhibit similar settling times but one* ***(b)*** *generates a large glitch at the off-to-on transition. A compression of the glitch appears in* ***c.***

verter devices from different manufacturers.

The 565A devices require a special test setup, because their full-scale output is only −2 mA, in contrast with the fixture's ±10-mA input range. This sensitivity penalty degrades the signal-to-noise ratio. You should ground the feedback and offset resistor of the 565A devices to prevent their capacitance from slowing the settling. You should also bypass the reference output with a 1-μF tantalum capacitor to prevent the reference from destabilizing.

To measure current, ground the voltage input terminal of the fixture, or terminate it with a 50Ω resistor, and connect the converter output to the current-input terminal of the fixture. The input impedance is now 500Ω, but by reckoning it as 1000Ω, you can maintain the calibration of the gain switch. Thus, 1 mV/div on the gain switch equals 1 μA/div of input, or approximately 2 LSB of the converter output. The first device seems to settle within its 250-nsec specification, even though it has a slow tail (approximately ±1 LSB) (**Fig 8a**). This tail is probably real, but you should, of course, always be suspicious of any measurements at these speeds and sensitivities. The second device exhibits a large glitch on the off-to-on transition (**Fig 8b**); this glitch is probably caused by destabilization of the device's reference input amplifier. **Fig 8c** shows a compression of **Fig 8b**'s glitch so that the whole glitch is visible.

ANALOG DEVICES

AN-351
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Ask the Applications Engineer—2

by James Bryant

WHEN IT COMES TO TRIMMING . . .

Q: *I need some advice about trimming offsets and gains.*

A. Don't!—unless you must. Good alternatives include (a) using headache-free devices, components, and circuits that meet the specs without trimming; (b) taking advantage of digital technology in system applications to make trim corrections in software. Savings provided on occasion by trim potentiometers, in conjunction with loosely spec'd devices, can turn out to be illusory when you consider the effects of circuit design, temperature, vibration, and life on performance and stability—as well as additional paperwork and complexity trimming entails.

Q: *Nevertheless, how do I trim the offset and gain errors in analog circuitry?*

A: In the correct order and with the correct inputs. If you consider the transfer characteristic of the circuit being trimmed the method to use is generally straightforward.

The simplified ideal transfer characteristic of a linear analog circuit (such as an amplifier, ADC or DAC) is given by the equation:

$$OP = K \times IP \quad (1)$$

where OP is output, IP is input, and K is a scale factor (Note that this simplification hides an enormous number of issues: quantization error in an ADC; dimensionality of K if the input and output are in different forms [e.g. voltage in / current out]; intentional offsets; and many others.)

In a real (non-ideal) circuit, offset and gain errors, OS (referred to the input) and ΔK, respectively, also appear in the equation, which becomes:

$$OP = (K + \Delta K) \times (IP + OS) \quad (2)$$

$$OP = (K \times IP) + [(K \times OS) + (\Delta K \times IP) + (\Delta K \times OS)] \quad (3)$$

Equations (2) and (3) are incomplete in that they assume only one offset—at the input—but this is the most-common case. Systems with separate input and output offsets will be considered later.

From (3) we see that it not possible to trim gain directly when an unknown offset is present. Offset must be trimmed first. With IP set at 0, the offset trim is adjusted until OP is also 0. Gain may then be trimmed: with an input near to full scale (FS), the gain trim is adjusted to make the output obey equation (1).

Q: *But what about bipolar ADCs and DACs?*

A: Many ADCs and DACs may be switched between unipolar and bipolar operation; such devices, wherever possible, should have their offset and gain trimmed in the unipolar mode. Where it is not possible, or where the converter is to operate only in the bipolar mode, other considerations apply.

A bipolar converter may be considered as a unipolar converter with a large offset (to be precise, an offset of 1 MSB—one-half of full-scale range). Depending on the architecture used, this bipolar offset (BOS) may or may not be affected by the gain trim. If it is so affected, equation (1) becomes:

$$OP = K \times (IP - BOS) \quad (4)$$

In this case offset is trimmed at analog zero, after which gain is trimmed near FS—positive or negative, but usually positive. This is normally the method used for DACs where the bipolar offset is within the DAC.

If the bipolar offset is not affected by the gain trim:

$$OP = K \times IP - BOS \quad (5)$$

Here offset is trimmed at FS negative and gain is trimmed at (or very near to—see below) FS positive. This method is used for most ADCs and for DACs where bipolar offset is obtained by the use of op amps and resistors external to the DAC.

Naturally, the method suggested on the data sheet should always be followed, but where a data sheet is unobtainable, in general, offset should be trimmed at analog zero for DACs and FS negative for ADCs—and near FS positive for both.

Q: *Why do you keep saying "near" to full scale?*

A: Amplifiers and DACs may be trimmed at zero and full scale. In the case of a DAC, all-1's—the largest digital input possible—should produce an output 1 LSB below "full scale," where "full scale" is considered as some constant times the reference; this follows since the output of a DAC is the normalized product of the reference and the digital input.

ADCs are not trimmed at zero and FS. The output of an ideal ADC is quantized, and the first output transition (from 00 . . . 00 to 00 . . . 01) takes place 1/2 LSB above the nominal value of all 0's. Thereafter transitions take place every 1-LSB increase in analog input until the final transition takes place 1 1/2 LSB below FS. A non-ideal ADC is trimmed by setting its input to the nominal value of a desired transition and then adjusting until the ADC output flickers between the two values equally.

The offset of an ADC is therefore trimmed with an input corresponding to the first transition (i.e., 1/2 LSB above zero or above FS negative—which is "near" zero or "near" FS negative); and the gain is then trimmed at the last transition

Reprinted from Analog Dialogue 23–1 1989

(i.e. 1 1/2 LSB below FS positive—which is "near" FS positive). This procedure results in an interaction between the gain and offset errors during offset trim but it should be too slight to be significant.

Q: *Are there any other anomalies resulting in a need to trim "near", rather than at full scale?*

A: Synchronous voltage-to-frequency converters (SVFCs) are liable to injection locking phenomena when their output frequency is harmonically related to their clock frequency, i.e., when their output is very close to 1/2, 1/3 or 1/4 of clock frequency. FS for an SVFC is 1/2 clock frequency. The presence of a trim tool can exacerbate the problem. It is therefore advisable to trim the gain of an SVFC at around 95% of FS.

Q: *What about circuits requiring both "input" and "output" offset trim?*

A: Circuits such as instrumentation and isolation amplifiers often have two stages of dc gain, and the gain of the input stages can be variable. Thus a two stage amplifier, with an input offset, *IOS*, an output offset, *OOS*, a first stage gain of *K*, and a unity-gain output stage, has (for zero input) an output, *OP*, of:

$$OP = OOS + K \times IOS \qquad (6)$$

From (6) it is evident that if the gain is constant we need only adjust either *IOS* or *OOS* to null the total offset (although if the input uses a long-tailed pair of bipolar transistors we will get a better offset temperature coefficient if we trim both—for FET long-tailed pairs this is not necessarily the case). If the first stage gain is to be varied, both offsets must be trimmed.

This is done by an iterative procedure. With zero input, and gain set to maximum, the input offset is adjusted until the output is also zero. The gain is then reduced to its minimum value and the output offset adjusted until the output is zero again. The two steps are repeated until no further adjustment is necessary. Gain trimming should not be done until both *IOS* and *OOS* are nulled; the actual values of the high and low gains used in offset trim are unimportant.

Q: *What circuitry should I use for gain and offset trims?*

A: Many amplifiers (and a few converters) have terminals for trimming gain and offset. Many more do not.

Offset trim is normally performed with a potentiometer connected between two assigned terminals, and its wiper is connected (sometimes via a resistor) to one of the supplies. The correct connections and component values will be given on the device data sheet. One of the commonest differences between op-amps is the value of offset correction potentiometer and which supply it should be connected to.

Where separate terminals are not provided for offset trim, an offset-adjusting constant can usually be added to the input signal. Two basic possibilities are shown in Figures 1a and 1b. Where the correction is being made to a system where a differential input op amp is used as an inverter (the commonest case) the method of 1a is best to correct for device offsets—but not system offsets. In the single-ended connection, method 1b will work for system offsets but should be avoided where possible for small device offsets, because it often requires a very large value of summing resistance, compared to the signal-input resistances, in order to (i) avoid loading the summing point excessively, (ii) scale the correction voltage properly and produce enough attenuation to minimize the effects of differential supply-voltage drifts. It is often helpful to use resistances between the supplies and the potentiometer to increase trim resolution and reduce dissipation.

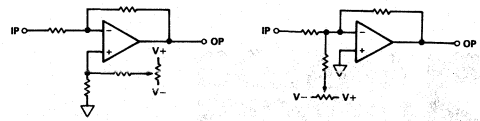

a. Voltage applied to non-inverting input.

b. Current summed at inverting input.

Figure 1. Two connections for offset adjustment.

Where gain trim is provided for in a circuit, it will generally consist of a variable resistor. Details of its value and connection will appear on the data sheet of the device. Where gain trim is not required, this resistor may be replaced by a fixed resistor having half the resistance of the maximum value of the recommended trim potentiometer.

Where gain trim is not provided it is not always achievable externally without an additional variable-gain stage. For example, consider a DAC using a ladder network. If the ladder network is used in the current mode (Figure 2a), the input impedance at the reference terminal does not vary with digital code, and the gain of the DAC may be trimmed with a small variable resistor in series with either the reference input or the feedback resistor. However, if the DAC is used in the voltage mode (Fig 2b), then the reference input impedance is code dependent, and gain may only be trimmed by varying the reference voltage—which is not always possible—or the gain of the buffer amplifier.

The possibility of trimming gain in circuits not furnished with gain-trim circuitry, therefore, will depend on individual cases; each must be assessed on its own merits.

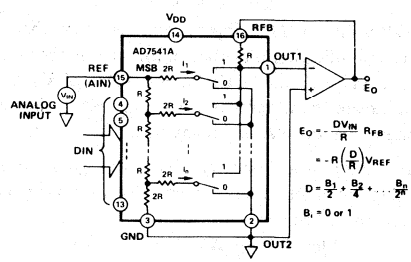

a. CMOS DAC connected for current steering. Input impedance is constant.

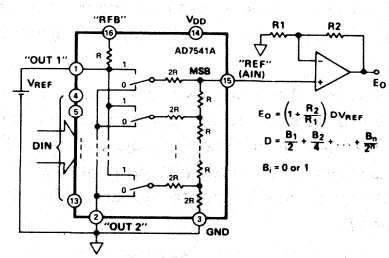

b. The same DAC connected for voltage output.

Figure 2. Comparing basic DAC circuits.

)NE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Ask the Applications Engineer—5

High-Speed Comparators Provide Many Useful Circuit Functions When Used Correctly

by John Sylvan

uestion: Why can't I just use a standard op amp in a high-gain or en-loop configuration as a voltage comparator?

ɔu can—if you are willing to accept response times in the tens of icroseconds. Indeed, if in addition you require low bias-rrents, high-precision and low offset voltages, then an op amp ight be a better choice than most standard voltage comparators. ıt since most op amps have internal phase/frequency compensa-ɔn for stability with feedback, it's difficult to get them to spond in nanoseconds. On the other hand, a low-cost popular ɔmparator, the LM311, has a response time of 200 ns.

lso the output of an operational amplifier is not readily matched standard logic levels. Without external clamping or level-ıifting, an op amp operating as a comparator will swing to within few volts of the positive and negative supplies, which is incom-ıtible with standard TTL or CMOS logic levels.

y comparator oscillates uncontrollably. Why does this happen?

ĸamine the power-supply bypassing. Even a few inches of PC ace on the supply lines can add unacceptable dc resistance and ductance. As a result, transient currents while the output is vitching may cause supply-voltage fluctuations, which are fed ıck to the input through the ground and supply lines. Install w-loss capacitors (0.1μF ceramic capacitors) as close as possible the supply pins of the comparator to serve as a low-impedance servoir of energy during high-speed switching.[1]

ve installed bypass capacitors, but I still can't keep my high-speed mparator from oscillating. Now what's the problem?

could be the comparator's ground connection. Make sure that ıe ground lead is as short as possible and connected to a low-ɲpedance ground point to minimize coupling through lead ıductance. Use a ground plane if possible and avoid sockets.

nother cause of the oscillation may be a high source impedance ıd stray capacitance to the input. Even a few thousand ohms of ɔurce impedance and picofarads of stray capacitance can cause nruly oscillations. Keep leads short, *including the ground clip of ıur scope probe*. For best measurement results use the shortest ossible ground lead to minimize its inductance (<1").

Vith a slowly moving input signal, my comparator seems to "chatter" s it passes through the transition voltage. Why can't I obtain a single 'ean transition from the device?

. comparator's high gain and wide bandwidth are typically the ɔurce of this problem. Any noise is amplified, and as the signal asses through the transition region, the noise can cause a fast-ɛsponding amplifier's output to bounce back and forth. Also, ince the device's sensitivity (i.e., *gain*) is higher during a transi-on, the tendency to oscillate due to feedback increases. If ossible, filter the signal to minimize noise accompanying it.

A useful discussion of comparator foibles can be found in *Troubleshooting Techniques Quash Spurious Oscillations*, by Bob Pease, *EDN*, September 14, 1989, pp. 152-6

Reprinted from Analog Dialogue 23–4 1989

Then try using hysteresis which, like backlash in gear trains, requires the input to change by a certain amount before a reversal occurs. For example, after a high-to-low transition on the AD790, its built-in hysteresis requires the input voltage (positive input) to increase by 500 μV to produce a low-to-high transition.

If my comparator does not have internal hysteresis, can I add it externally?

Yes, with external positive feedback. This is done by feeding a small fraction of the output of the comparator back to the positive input. This simple technique is shown in Figure 1. The hysteresis voltage from the lower transition point to the upper transition point will depend on the value of the feedback resistor, R_F, the source resistance, R_S, low output level, V_{low}, and high output level, V_{high}. The low and high transition points are set by:

$$V_{low} \times \frac{R_S}{R_S + R_F} \text{ and } V_{high} \times \frac{R_S}{R_S + R_F}$$

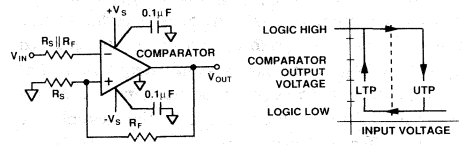

Figure 1. Applying external hysteresis to a comparator.

Figure 2 shows how adding external hysteresis can "clean up" a comparator's response. Figure 2a shows the response of a comparator with bipolar output swing without hysteresis. As the triangular-wave input (trace A) passes through the transition point (ground), the device oscillates vigorously (and couples a portion of the oscillation into ground and the signal-source). Figure 2b depicts the response of the same comparator with 5 mV of external hysteresis applied; it shows a much cleaner transition.

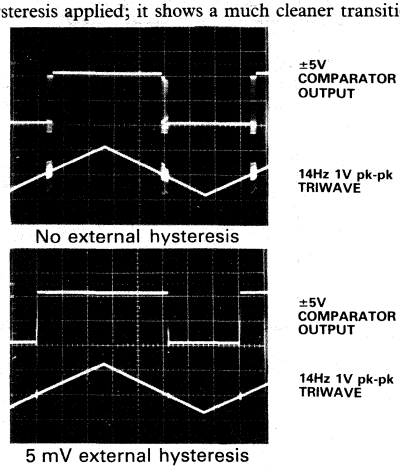

Figure 2. Hysteresis helps clean up comparator response.

A problem encountered with external hysteresis is that output voltage depends on supply voltage and loading. This means the hysteresis voltage can vary from application to application; though this affects resolution, it need not be a serious problem, since the hysteresis is usually a very small fraction of the range and can tolerate a safety margin of two or three (or more) times what one might calculate. Swapping in a few comparators can help confidence in the safety margin. Don't use wirewound resistors for feedback; their inductance can make matters worse.

What's the difference between propagation delay and prop-delay dispersion? Which of the two specifications is of most concern?

Propagation delay is the time from when the input signal crosses the transition point to when the output of the comparator actually switches. Propagation-delay dispersion is the variation in prop delay as a function of overdrive level. If you're using a comparator in pin-drive electronics in an automatic test system, then prop-delay dispersion will determine the maximum edge resolution. In contrast, propagation delay can be considered as a fixed time offset and therefore compensated for by other techniques.

I have a +5-volt system and don't want to add an additional supply voltage; can I use my comparator with a single supply?

Yes, but to establish the threshold use an adequately bypassed stable reference source well within the common-mode range of the device. The signal level is also to be referenced to this source.

I sometimes see unexpected behavior in my comparator. What could be the cause of this problem?

Examine the common-mode range of the input signal. Unlike operational amplifiers, that usually operate with the input voltages at the same level, comparators typically see a large differential voltage swings at their inputs. If the inputs exceed the device's specified common-mode range (even though within the specified *signal* range), the comparator may respond erroneously. For proper operation, ensure that *both* input signals do not exceed the common-mode range of the specific comparator. For example, the AD790 has a $\pm V_S$ differential input range, but its common-mode range is from $-V_S$ to 2 volts below $+V_S$.

Can you suggest a circuit that performs autozeroing when the comparator is off-line to minimize drift?

Try the circuit shown in Figures 3 and 4. In the Calibrate mode, the input is disconnected and the positive input of the comparator is switched to ground. The comparator is connected in a loop with a pair of low-voltage sources of opposing polarity charging a buffered capacitor in response to the comparator's output state.

If the comparator's minus input terminal is above ground, then the comparator output will be low, the 1-μF capacitor will be connected to the negative voltage (–365 mV) and the voltage from the buffer amplifier will ramp down until it is below the plus input (ground)—plus hysteresis and any offsets—at which point the comparator switches. If it is below ground, the comparator's output will be high, the capacitor will be connected to the positive voltage (+365 mV), the output from the buffer amplifier ramps up. In the final state, each time the comparator switches (when the ramped change exceeds the hysteresis voltage), the polarity of the current is reversed; thus the capacitor voltage averages out the offsets of the buffer and comparator.

At the end of the Calibrate cycle, the JFET switch is opened, with the capacitor charged to a voltage equal to the offsets of the comparator and buffer ± the hysteresis voltage. At the same time, the Calibrate signal goes low, disabling the feedback to the polarity switch and connecting the input signal to the comparator.

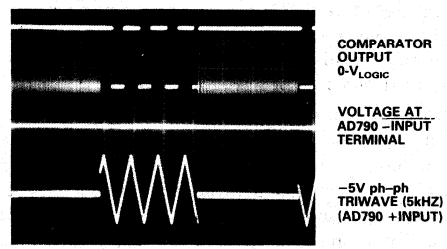

Figure 4. Comparator output, buffer output, and comparator input.

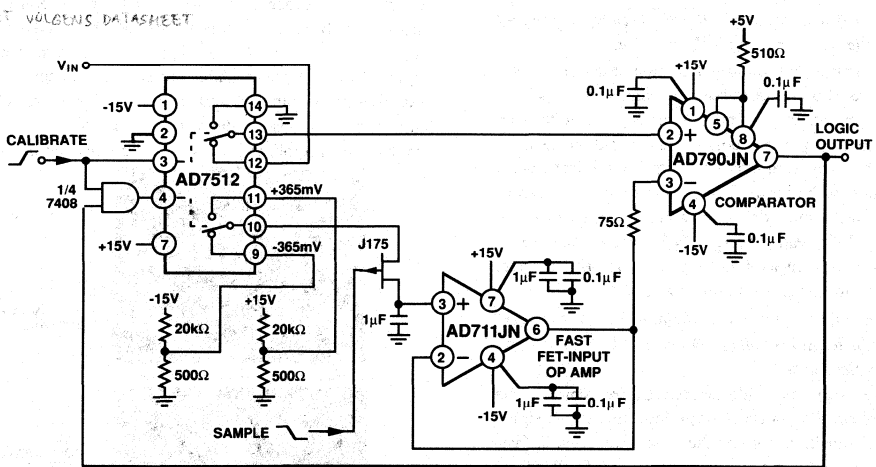

Figure 3. Autozeroed comparator integrates out offsets during calibration cycle.

ANALOG DEVICES

AN-362
APPLICATION NOTE

NE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Ask the Applications Engineer—9

by James Bryant

\RIOUS TOPICS

/ James Bryant

. *Tell me something about supply decoupling.*

. All precision analog integrated circuits, even low-frequency ones, contain transistors having cutoff frequencies of hundreds of MHz; their supplies must therefore be decoupled to the ground return at high frequency—as close to the IC as feasible to prevent possible instability at very high frequencies. The capacitors used for such decoupling must have low self-inductance, and their leads should be as short as possible (surface-mounted chip ceramic capacitors of 10- to 100 nF are ideal, but leaded chip ceramics are generally quite effective if the lead length is kept to less than 2 mm.

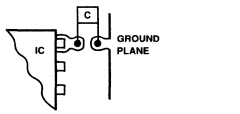

IDEAL HF DECOUPLING HAS
1. LOW INDUCTANCE CAPACITOR (MONOLITHIC CERAMIC)
2. MOUNTED VERY CLOSE TO THE IC
3. WITH SHORT LEADS
4. AND SHORT, WIDE PC TRACKS

IT MAY BE SHUNTED WITH A TANTALUM BEAD ELECTROLYTIC TO PROVIDE GOOD LF DECOUPLING AS WELL.

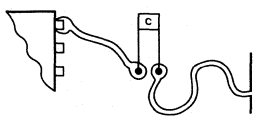

THIS SORT OF THING IS USELESS!

Low-frequency decoupling is also important, since the PSR (power-supply rejection) is normally specified at dc and will deteriorate appreciably with increasing power-supply ripple frequencies. In some high-gain applications, feedback through the common power-supply impedance can lead to low-frequency instability ("motorboating"). However, low-frequency decoupling at each IC is not often necessary.

Supply decoupling does more than prevent instability. An op-amp is a *four-terminal* device (at least), since there must be a return path for both input signals and the output circuit. It is customary to consider the common terminal of both op-amp supplies (for op-amps using ± supplies) as the output signal-return path, but in fact, one of the supplies will be the de facto return path at higher frequencies, and the decoupling of the amplifier's supply terminal for this supply must take into consideration both the necessity of normal high-frequency decoupling and the routing of the output ground.†

Q. In "Ask the Application Engineer," you're always describing non-ideal behavior of integrated circuits. It must be a relief to use a simple component like a resistor and know that you have a near-ideal component.

A. I only wish that a resistor was an ideal component, and that that little cylinder with wire ends behaved just like a pure resistance. Real resistors also contain imaginary resistance components—in other words they're reactive. Most resistors have a small capacitance, typically 1-3 pF, in parallel with their resistance, although some types of film resistors, which have a spiral groove cut in their resistive film, may be inductive, with inductances of a few tens or hundreds of nH.

RESISTORS ARE REACTIVE:

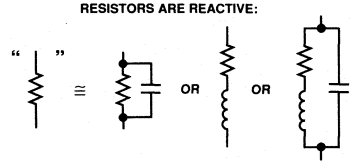

Of course, wirewound resistors are generally inductive rather than capacitive (at least, at the lower frequencies). After all, they consist of a coil of wire. It is commonplace for wirewound resistors to have inductances of several microhenrys or tens of microhenrys, and even so-called "non-inductive" wirewound resistors, which consist of N/2 turns wound clockwise and N/2 turns wound anticlockwise, so that the inductances of the two half windings cancel out, have a residual inductance of a microhenry or even more. (For higher-resistance-value types, above 10 kΩ or so, the residual reactance may be capacitive rather than inductive, and the capacitance will be higher—by up to 10 pF—than a standard film or composition resistor.)

These reactances must be considered carefully when designing high frequency circuits which contain resistors.

†This issue is developed in detail in the free application note, "An IC amplifier user's guide to decoupling, grounding, and making things go right for a change," by Paul Brokaw.

Reprinted from Analog Dialogue 25-1 1991

Q. *But many of the circuits you describe are for making precision measurements at DC or very low frequencies. Stray inductance and capacitance don't matter in such applications, do they?*

A. They actually do. Since transistors (either discrete or within ICs) have very wide bandwidths, if such circuits are terminated with reactive loads, they may sometimes oscillate at frequencies of hundreds or thousands of MHz; bias shifts and rectification associated with the oscillations can have devastating effects on low-frequency precision and stability.

Even worse, this oscillation may not appear on an oscilloscope, either because the oscilloscope bandwidth is too low for such a high frequency to be displayed, or because the scope probe's capacitance is sufficient to stop the oscillation. It is always wise to use a wideband (LF to 1.5 GHz or more) spectrum analyzer to verify the absence of parasitic oscillations in a system. Such checks should be made while the input is varied throughout its whole dynamic range, since parasitic oscillations may sometimes occur over a narrow range of inputs.

Q. *Are there any problems with the resistance of resistors?*

A. The resistance of a resistor is not fixed but varies with temperature. The temperature coefficient (TC) varies from a few parts per million per degree Celsius (ppm/°C) to thousands of ppm/°C. The resistors with the best stability are wirewound or metal film types, and the worst are carbon composition.

Large temperature coefficients are sometimes useful (an earlier "Ask the Applications Engineer"* mentioned how a +3,500-ppm/°C resistor can be used to compensate for the kT/q term in the equation for the behavior of a junction diode). But in general, the variation of resistance with temperature is likely to be a source of error in precision circuits.

If the accuracy of a circuit depends on the matching of two resistors having different TCs, then, no matter how well-matched at one temperature, they will not match at another; and even if the TCs of two resistors match, there is no guarantee that they will remain at the same temperature. Self-heating by internal dissipation, or external heating from a warm part of the system, will result in a mismatch of temperature, hence resistance. Even with high quality wirewound or metal-film resistors these effects can result in matching errors of several hundred (or even thousand) ppm. The obvious solution is to use resistors which are fabricated in close proximity on the same substrate whenever good matching is necessary for system accuracy. The substrate may be the silicon of a precision analog IC or a glass or metal thin-film substrate. In either case, the resistors will be well-matched during manufacture, will have well-matched TCs, and will be at nearly the same temperature because of their proximity.

ANALOG DEVICES

AN-353
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Ask the Applications Engineer—10

by James Bryant

In the last issue of Analog Dialogue you told us about some of the problems of a simple resistor. [More will appear in a future issue.] *Surely there must be some component that behaves exactly as I expected it to. How about a piece of wire?*

Not even that. You presumably expect your piece of wire or length of PC track to act as a conductor. But room-temperature superconductors have not yet been invented, so any piece of metal will act as a low-valued resistor (with capacitance and inductance, too) and its effect on your circuit must be considered.

Surely the resistance of a short length of copper in small-signal circuits is unimportant?

Consider a 16-bit a/d converter with 5-kΩ input impedance. Suppose that the signal conductor to its input consists of 10 cm of typical PC track—0.25 mm (0.010″) wide and 0.038 mm (0.0015″) thick. This will have a resistance of approximately 0.18 Ω at room temperature, which is slightly $<2 \times 2^{-16}$ of 5 kΩ; this introduces a gain error of 2 LSB at full scale.

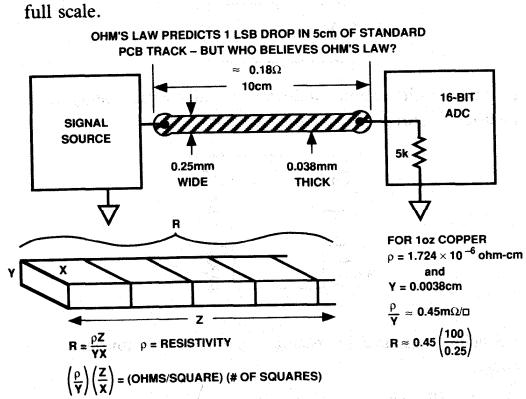

One might argue that the problem would be reduced if PC tracks were made wider—and indeed, in analog circuitry it's almost always better to use wide tracks; but many layout drafters (and PC Design programs) prefer minimum-width tracks for signal conductor. In any case it's especially important to *calculate* the track resistance and its effect in every location where it might cause a problem.

Q. Doesn't the capacitance of the extra width of track to metal on the board's underside cause a problem?

A. Rarely. Although the capacitance of PC tracks is important (even in circuits designed for low frequencies, since LF circuits can oscillate parasitically at HF) and should always be evaluated, the extra capacitance of a wider track is unlikely to cause a problem if none existed previously. If it is a problem, small areas of ground plane can be removed to reduce ground capacitance.

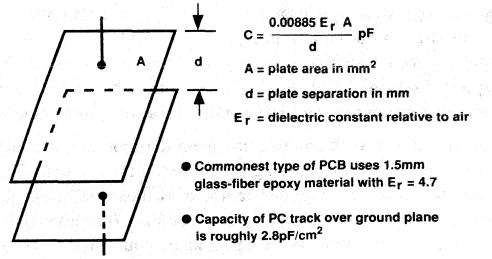

Q. Hold it! What's a ground plane?

A. If one entire side of a PCB (or one entire layer, in the case of a multi-layer PCB) consists of continuous copper which is used as ground this is known as a "ground plane." It will have the least possible resistance and inductance of any ground configuration. If a system uses a ground plane, it is less likely to suffer ground noise problems.

Q. I have heard that ground planes are hard to manufacture.

A. Twenty years ago there was some truth in this. Today improvements in PC adhesives, solder resists and wave-soldering techniques make the manufacture of ground-plane PCB's a routine operation.

Q. You say that a system using a ground plane is "less likely" to suffer ground noise problems. What remaining ground noise problems does it not cure?

A. The basic circuit of a system having ground noise is shown in the diagram. Even with a ground plane the resistance and inductance will not be zero—and if the external current source is strong enough it will corrupt the precision signal.

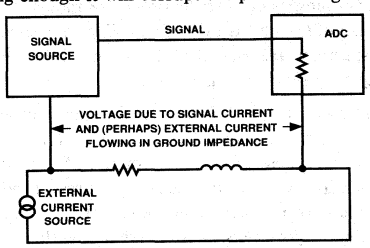

The problem is minimized by arranging the PCB so that high currents do not flow in regions where ground voltages can corrupt precision signals. Sometimes a break or slot in a ground plane can divert a large ground current from a sensitive area—but breaks in a ground plane can also reroute signals into sensitive areas, so the technique must be used with care.

Reprinted from Analog Dialogue 25–2

Q. How do I know what voltage drops are present in a ground plane?

A. They should generally be measured; however, it is sometimes possible to calculate them from the resistance of the ground plane material (standard 1 oz copper has resistance of 0.45 mΩ/square) and the length through which currents flow, but the calculation can be complicated. At DC and low frequencies (dc-50 kHz), voltage drops can be measured with an instrumentation amplifier such as the AMP-02 or the AD620.*

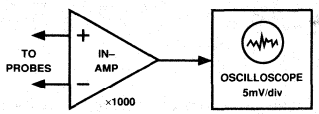

The amplifier is set to a gain of 1,000 and connected to an oscilloscope with a gain of 5 mV/div. The amplifier may be powered from the same supply as the circuit being tested—or from its own supply—but the grounds of the amplifier, its supply if separate, and the oscilloscope must be connected to the power ground of the circuit under test at the power supply.

The voltage between any two points on the ground plane may then be measured by applying the probes to those points. The combination of the amplifier gain and oscilloscope sensitivity give a measurement sensitivity of 5 μV/div. Amplifier noise will swell the oscilloscope trace to a band about 3 μV wide but it is still possible to make measurements with about 1-μV resolution—sufficient to identify most low-frequency ground noise problems; and identification is 80% of a cure.

Q. Are there any cautions about performing this test?

A. Any alternating magnetic fields which thread the probe leads will induce voltages in them. This can be tested by short-circuiting the probes together (and resistively to ground to provide a bias current path) and observing the oscilloscope trace; ac waveforms observed that result from inductive pick-up may be minimized by repositioning the leads or taking steps to eliminate the magnetic field. It is also essential to ensure that the ground of the amplifier is connected to the system ground; without this connection the amplifier, with no return path for bias current cannot work; grounding also ensures that this connection does not disturb the current distribution that is being measured.

Q. What about measuring HF ground noise?

A. It is hard to make a suitable instrumentation amplifier with wide bandwidth, so at HF and VHF a passive probe is more suitable. This consists of a ferrite toroid (6-8 mm OD) wound with two coils of 6-10 turns each. One coil is connected to the input of a spectrum analyzer, the other to the probes, to make a high-frequency isolating transformer.

The test is similar to the LF one but the spectrum analyzer displays noise as an amplitude-frequency plot. While this differs from time-domain information, sources of noise may be easier to identify by their frequency signatures; in addition, the use of a spectrum analyzer provides at least 60 dB more sensitivity than is possible with a broadband oscilloscope.

Q. What about the inductance of wires?

A. The inductance of wire- and PC-track leads should not be overlooked at higher frequencies. Here are some approximations for calculating the inductance of straight wires and runs. For example, 1 cm of 0.25-mm track has an inductance of 10 nH.

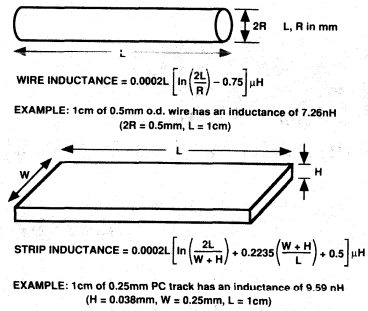

But inductive reactance is generally much less of a problem than stray flux cutting inductive loops and inducing voltages; loop area must be minimized, since voltage is proportional to it. In wired circuits this is easily done using twisted pairs.

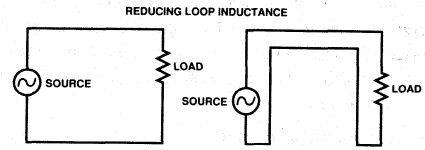

In boards, leads and return paths should be close together; quite small changes in layout will often minimize the effect.

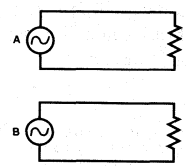

In this circuit, mutual inductance will couple energy from high-level source A into low-level circuit B.

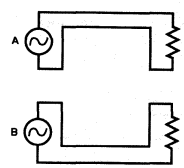

Reducing area and increasing separation will minimize the effect.

Usually, all that is necessary is to minimize loop area and maximize the distance between potentially interfering loops. Occasionally magnetic shielding is required, but it is expensive and liable to mechanical damage; avoid it whenever possible.

REFERENCES

The Best of Analog Dialogue 1967-1991. Norwood MA: Analog Devices (1991), pp. 120-129, 193-195. Contains many additional references.

Mixed-Signal Design Seminor Notes. Norwood MA: Analog Devices (1991). Contains additional References.

ANALOG DEVICES

AN-354
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Ask the Applications Engineer—1

Multi Troubles

by James Bryant

MULTI TROUBLES

Q: *My multiplexed ADC system is misbehaving . . .*

A: Before you go any further, have you grounded all unused multiplexer channels?

Q: *No. But how did you know?*

A: Because the floating terminal is one of the commonest causes of problems in systems containing CMOS multiplexers. Unused MUX inputs and outputs (whether integrated into a multiplexed ADC or part of a self-contained MUX chip) can pick up signals from stray fields and inject them into the device's substrate, turning on spurious substrate devices. Then, even when the unused channel is turned off, the performance of the *on*-channel may be badly degraded (at the unlikely extreme, the injection may turn on a spurious four-layer device and destroy some chips).

Whenever a MUX is used, all its inputs and outputs must be connected to a potential between its supply rails. The best way to deal with unused channels is to ground them, but they may be connected to a more-convenient potential within the rails.

TROUBLE FROM THE START

Q: *To save power, my ADC is powered up only to make a measurement. The system is very accurate in continuous operation, but unpredictable when power is strobed. Why?*

A: When an ADC's power is switched on only to perform a conversion, it may misbehave for three reasons: slow reference turn-on, random initial logic states, and system latch-up.

For various reasons—thermal stabilization, capacitance charging, slow starting of regenerative current mirrors using PNP transistors in band-gap references— it is not uncommon for some voltage references to have relatively large errors for many milliseconds after power-up. Such errors in an ADC's external or internal reference during conversion lead to inaccurate results.

At turn-on, a typical ADC's logic will be in a random state; for a conversion triggered at that time, the ADC may not be able to perform correctly. With one conversion triggered, the logic should return to its correct pre-conversion state—but cases exist where two conversion cycles are necessary before the ADC is certain to perform a valid conversion. Hence, a good general rule is to perform two "dummy" conversions after powerup before relying on the results. (It is also well to recall that some ADCs react badly to having a conversion triggered before the previous conversion is complete; when this happens, one or two "dummy" conversions may be needed to return the logic to a known state.)

If an ADC's external logic is arranged so that the end of the ADC "Busy" signal starts a delay which ends with the start of the next conversion, it is important to realize that if the converter powers up in the Busy state, the Busy signal may remain latched up until a conversion Start pulse has been received. In this case, such a system cannot self-start. If the Busy signal is always present on power-up the problem is almost certain to be recognized—and addressed—during the design of the system; but if the Busy signal is only occasionally present on power-up the system may latch unpredictably. As a rule, control signals to an ADC during start-up should not depend on the logical state of Busy.

ABOUT LOG COMPENSATION RESISTORS

Q: *Designs of logarithmic circuits, including those using the AD538 Y[Z/X]^m unit: (For example, Figure 6 from the AD538 Multifunction Unit data sheet†) call for "kT/q compensation resistors." What are they and where do I get them?*

A: The V_{BE} difference across two opposed silicon junctions, one carrying a current, I, and the other a current, I_{REF}, *is* (kT/q) *ln* (I/I_{REF}). *Here,* k/q is the ratio of Boltzmann's constant to the charge on an electron (about 1/11,605 K/V), and T is the absolute temperature in kelvins.

*Much useful information about logarithmic and other analog function circuits can be found in the *Nonlinear Circuits Handbook*, published by Analog Devices ($5.95), P.O.Box 9902, Norwood MA 02062.

†See also *Analog Dialogue* 19-1 (1985), pp. 3–6.

Reprinted from Analog Dialogue 22–2 1988

Although employing similar junctions in isothermal pairs eliminates the effects of temperature-sensitivity of reverse saturation current, the kT/q term is still temperature-dependent. To eliminate this dependency in the application, the logarithmic voltage must be used in a circuit whose gain is inversely proportional to the absolute temperature of the junctions. Over a reasonable range of temperatures near 20°C, this may be arranged by the use of a gain-setting 1-kΩ resistor having a positive temperature coefficient of approximately 3,400ppm/°C—and keeping it at the same temperature as the junctions.

A 3,500 ppm/°C resistor is available from Tel Laboratories, 154G Harvey Road, Londonderry, New Hampshire 03053 (603)–625-8994, Telex: (710)–220-1844, designated Q-81, and from the Precision Resistor Co. Inc., 10601, 75th. St., Largo, Florida 33543 [(813)–541-5771 Telex: 821788], as the PT146. Analog Devices offices in most European countries are aware of local suppliers of these resistors.

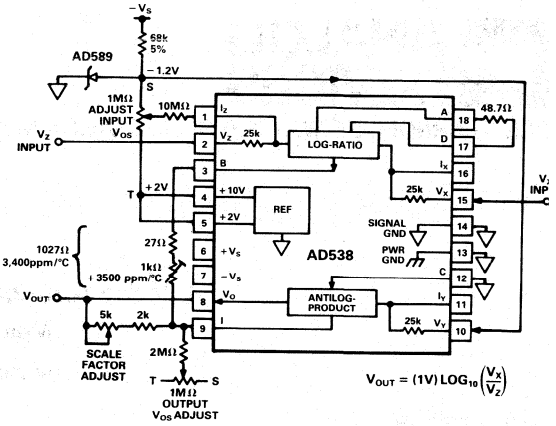

ANALOG DEVICES

AN-280
APPLICATION NOTE

)NE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Mixed Signal Circuit Techniques

[NTRODUCTION

There are considerably more problems nvolved in the successful design of Mixed Signal circuitry than mere circuit design. If ve design an electronic circuit as a diagram, vhether we use an old-fashioned pencil and)aper or, as is the modern fashion, a com-)uter and SPICE or some similar software, ve are overlooking one of the most important 'actors in the design of successful hardware, ıamely that what we are designing is HARD-vare, and until it has been shown to work ;uccessfully in fact, rather than in simula-ion, our design is not complete.

This section of our seminar considers the)roblems which arise when reality reacts on ı design which theory and modelling have ;hown to be satisfactory. What, in fact, has ıappened is that our model probably does ıot consider the effects of non-ideal compo-ıents and of spurious or parasitic compo-ıents resulting from the circuit layout which ıas been used. It is not, perhaps, too fanciful ;o describe this as the section of the Seminar lealing with Murphy's Law.

MURPHY'S LAW

IN ANY SET OF CIRCUMSTANCES THE WORST THING THAT CAN HAPPEN - WILL

- **Any effect which you think can be disregarded, can't.**
- **Nature always sides with the hidden flaw.**

Figure 11.1

Murphy's Law, though frequently expressed humorously, is not entirely a joke. It is a recognition of the complexity of physical systems and a warning against over-simplification and is comparable with Einstein's warning that "Everything should be made as simple as possible - but no simpler".[1]

IMPORTANT COROLLARIES TO MURPHY'S LAW

- **After it has worked successfully for two weeks it will fail during the first public demonstration.**
- **Equipment blows to protect fuses.**
- **Interchangeable parts aren't.**
- **Fail-safes don't.**

Figure 11.2

This section of the seminar discusses the various physical effects which must be considered in the design of the hardware of mixed signal systems. Often such consideration will amount to a quick calculation to demonstrate that further consideration is not necessary, but sometimes extensive analysis, or even actual experiments, will be necessary. However, the quick calculation must not be omitted, since problems are rarely obvious and often unexpected. The effects which must be considered will include many basic laws of physics.

BASIC LAWS INVOLVED IN THE DESIGN OF MIXED SIGNAL CIRCUITRY

- **Ohm's Law**
- **Kirchoff's Law**
- **Faraday's Laws**
- **Lenz's Law**

Figure 11.3

BIT SIZES FOR 10 V FULLSCALE CONVERTERS

RESOLUTION	1 LSB	0.5 LSB	% FS	ppm FS	dB FS
4-bit	625mV	313mV	6.25	62500	-24
6-bit	156mV	78mV	1.56	15625	-36
8-bit	39mV	19.5mV	0.39	3906	-48
10-bit	9.76mV	4.88mV	0.098	977	-60
12-bit	2.44mV	1.22mV	0.024	244	-72
14-bit	610μV	305μV	0.0061	61	-84
16-bit	153μV	76μV	0.0015	15	-96
18-bit	38μV	19μV	0.0004	4	-108
20-bit	9.5μV	4.8μV	0.0001	1	-120
22-bit	2.4μV	1.2μV	0.000024	0.24	-132
24-bit	0.6μV	0.3μV	0.000006	0.06	-144

Figure 11.4.

We therefore shall use as a section heading the major phenomenon considered in the section, but in the most general sense (for example, under "Resistance" we shall consider the non-ideal behavior of resistors, including noise, thermo- electric and inductive effects, which are not strictly issues of Ohm's Law).

When considering the effects of circuit conditions we are, of course, interested in their effects on the performance of the system as a whole. Failure to allow for this is at the root of many of the problems which this section considers. For example, a 16-bit system divides its full-scale (FS) range into 2^{16} or 65536, which means that 1 LSB in a 10 V FS system is only 153 μV. If we assume that we can tolerate errors of no more than 0.5 LSB, this calculation tells us that in a 16-bit system with 10 V FS we must keep the total error to less than 76 μV, which is approximately equal to the thermoelectric voltage in a nichrome wirewound resistor with copper/nickel leads having about 2°C temperature difference between its ends.

Binary logic circuitry, on the other hand, has only two states, logic 0 and logic 1, and noise immunity of hundreds or thousands of millivolts. This is why circuit designers who have only worked with digital circuitry tend to overlook the sources of error which we are considering in this section of the seminar.

Figure 11.4 lists the sizes of 0.5 LSB at various resolutions (the values are given for 10 V fullscale since this is a classical converter range and where LSBs are given a mV value in this section of the seminar a 10 V FS is assumed unless explicitly stated otherwise - scaling to other values of FS is a trivial operation). Every analog designer should be familiar with this table, since not only does it allow the comparison of converters which are specified in different ways but it also indicates whether a design is reasonable or not - if noise or system errors amount to 1 mV there is little point in designing a system with more than 12-bits resolution.

Resistance

Resistance of Conductors

Every engineer is familiar with resistors - little cylinders with wire ends - although perhaps fewer are aware of all their idiosyncrasies. Far too few engineers consider that all the wires and PC tracks with which their systems and circuits are assembled are also resistors.

At 25°C the resistivity of pure copper is 1.724E-6 ohm cm. The thickness of standard (1 ounce) PCB foil is 0.038 mm (0.0015"). The resistance of standard PCB copper is therefore 0.45 milliohms/square, which implies a resistance for the 0.25 mm track frequently used in computer designed digital circuitry of 18 milliohms/cm, which is quite large. Moreover the temperature coefficient of resistance for copper is about 0.4% /°C around room temperature, which can be a further inconvenience.

As an illustration of the effect of PCB track resistance consider a 16-bit ADC with a 5k ohm input resistance which has 5 cm of 0.25 mm PCB track between it and its signal source. This track has a resistance of approximately 0.09 ohms and introduces a gain error of 0.09 ohms / 5000 ohms (0.0018%) which is well over 1 LSB (0.0015% for 16 bits).

PRINTED CIRCUIT BOARD TRACK RESISTANCE

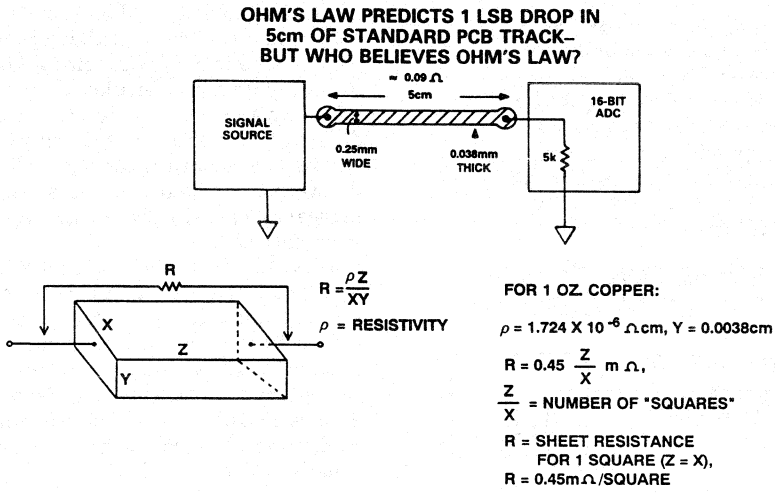

Figure 11.5

Skin Effect

This, of course, is a DC effect. At high frequencies we must also consider the "skin effect" where inductive effects cause currents to flow only in the surface of conductors. This has the effect of increasing the resistance of a conductor at high frequencies (note that this effect is separate from the increase in impedance due to the effects of the self-inductance of conductors as frequency is increased - that will be dealt with later). Skin effect is quite a complex phenomenon and detailed calculations are beyond the scope of this seminar. However a good approximation for copper is that the skin depth in centimeters is

$$\frac{6.6}{\sqrt{f}}, \text{ (f in Hz).}$$

Assuming that skin effects become important when the skin depth is less than 50% of the thickness of the PC foil this tells us that for normal 0.038 mm PC foil we must be concerned about skin effects at frequencies above approximately 12 MHz.

Where skin effect is important the resistance per square for copper is

$$2.6 \times 10^{-7} \sqrt{f} \text{ Ohms per square, (f in Hz)}$$

When calculating skin effects in PCBs it is important to remember that current flows in both sides of the PC foil (this is not necessarily the case in microstrip lines) so the resistance per square of PC foil is half the above value.

SKIN EFFECT

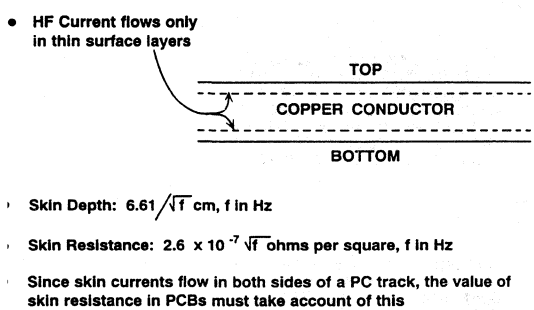

Figure 11.6

SKIN EFFECT

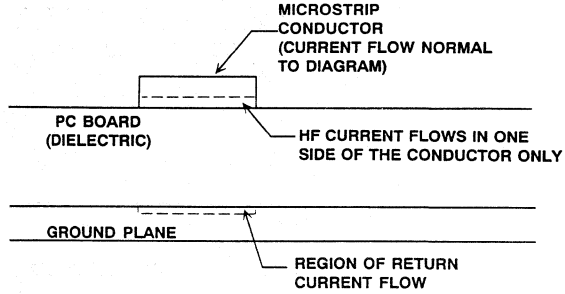

Figure 11.7

VOLTAGE DROP IN SIGNAL LEADS - "KELVIN" FEEDBACK

The gain error resulting from resistive voltage drop in signal leads is important only at high resolutions (as in the example) or where large signal currents flow. Where the load impedance is constant and resistive it can be compensated by adjusting the overall system gain. In other circumstances it may often be removed by the use of "Kelvin" or "voltage sensing" feedback.

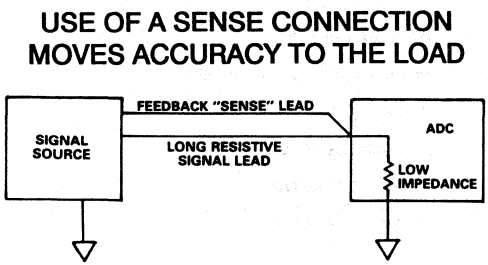

Figure 11.8

Separate force and sense connections at a load remove any errors resulting from voltage drops in the force lead, but, of course, may only be used in systems where there is negative feedback. It is also impossible to use such an arrangement to drive two or more loads with equal accuracy since feedback may only be taken from one point.

LEAKAGE IN INSULATORS

Just as conductors are improperly viewed as superconductors, so are insulators often mistakenly treated as perfect insulators, rather than very high resistances, which is the more accurate model.

Most printed circuit board materials are very good insulators, but they are not perfect, and inadequately cleaned PCB material may be quite a poor insulator. Furthermore, PCBs are anisotropic - even on a clean PCB different parts of the surface may have different resistivities, and the bulk resistance (between two plated through holes, for instance) is generally lower than the surface resistance between two tracks.

Since the insulation resistance is so variable (and it will vary further with temperature and humidity) it is hard to predict in any particular circumstances but it is safe to assume that it is unlikely that the resistance between two conductors on a clean PCB will drop below 10^{10} - 10^{11} ohms, and with teflon PCB material (which is very expensive) will usually be over 10^{12} ohms.

GUARD RINGS

In applications where high impedances and very low currents are involved a guard ring may be used to minimize the effects of low insulation resistance. If critical high impedance nodes are surrounded by a ring of conductor which is at (or very close to) the potential of the node itself then the leakage current at the node will be minimized. If the node is at, or near to, ground then a grounded guard ring will be appropriate, if it is at some other potential it may be necessary to use a high input impedance buffer amplifier, with its input connected to the node, to force the guard ring to the node

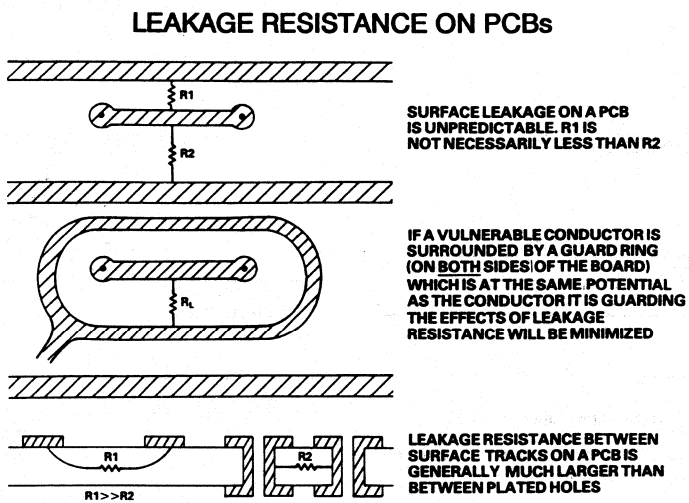

Figure 11.9

otential. It is obvious that, in general, uard rings should be on both sides of the 'CB with plated-through holes.

Nodes which are sufficiently sensitive to equire guard rings should not contain plated nrough holes (unless the PCB is made of eflon) because, as mentioned above, the bulk esistivity of PCB material is less than the urface resistivity.

An alternative to the use of a guard ring is o use teflon stand-off insulator(s) to support he high impedance point(s). If virgin teflon s used insulation resistance of around 10^{15} hms is possible ("Virgin teflon" is a solid iece of new teflon material which has been nachined to shape and has not been welded ogether from powder or grains). The mate-ial of the rest of the circuit board need not .ave particularly high insulation resistance.

A VIRGIN TEFLON STANDOFF INSULATOR HAS MUCH LOWER LEAKAGE THAN A PCB TRACK

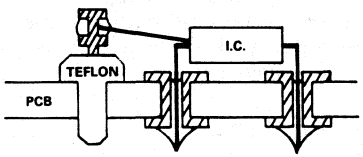

Figure 11.10

Electrostatic damage (ESD)

Where resistances are very high, especially in conditions of low humidity, there is always the possibility of electrostatic charge and electrostatic damage. A full discussion of electrostatic damage (ESD) and its prevention will be found in Analog Devices' Application Note on the subject, which is available free of charge from Analog Devices.[2]

This application note describes procedures to minimize the risk of electrostatic damage to sensitive devices. The basic principle of all ESD protection is to prevent a vulnerable item from being in the path of a discharge. Many of the precautions used in factories are designed to minimize the possibility of any damaging discharge, even in the event of carelessness. When experienced engineers handle ICs they may dispense with most of the ESD protection apparatus and merely ensure that the IC is never in any potential discharge path: when taking a circuit from conductive foam, touch the foam to equalize charge before touching the circuit, similarly touch the foam with the hand before inserting the circuit in it, and hold your colleague's hand BEFORE passing the IC.

ELECTROSTATIC DISCHARGE (ESD)

ESD PREVENTION MANUAL

ANALOG DEVICES

Figure 11.11

All integrated circuit structures are vulnerable to damage from the high voltages and high peak currents involved in even small electrostatic discharges but precision analog circuits suffer from a special disadvantage - the circuitry used to protect integrated circuit structures from ESD can often degrade the analog accuracy of the circuit where it is employed. Thus we have the choice between high performance and a high degree of protection. Which we choose will depend upon individual circumstances but it is essential to realize that the choice must be made - and if it is made in favour of accuracy then the circuit involved must not be exposed to electrostatic discharge.

A precision analog circuit exposed to ESD may not fail totally, but merely suffer degradation of its analog performance, and possible reduction of life expectancy. When an IC is returned to Analog Devices for failure analysis of inadequate performance the first check that is made when the package is opened is a visual inspection for evidence of electrostatic damage - and this is found in a large percentage of cases.

An interesting example of an unobvious effect of ESD occurred in Finland, where very cold winters produce very low humidity and particularly severe electrostatic problems. A customer complained that the AD549 low bias current BIFET op-amp had poor long-term reliability and that its noise performance deteriorated over a few years of use.

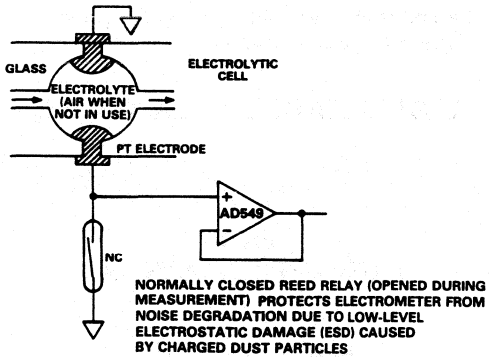

Figure 11.12

The amplifier was being used as a unity gain buffer with an electrochemical cell and the non-inverting input was connected to a platinum electrode and to nothing else. In use this electrode was immersed in electrolyte but after use it was washed (automatically) in deionized water and air dried. It was then left unconnected until the machine was next used.

Although there was no possibility of the electrode being touched at this time (it was in the very center of the machine) it could encounter random particles of electrostatically charged dust - and the pulse currents as these dust particles discharged were sufficient to cause gradual deterioration of the noise figure. As soon as arrangements were made to ground the electrode when it was not in use (with an NC reed relay for minimum leakage) the problem disappeared.

PARASITIC EFFECTS IN RESISTORS

When we model a circuit, either informally or with a program such as SPICE, we generally assume that a resistor is a simple resistance. In fact any resistor is a much more complex device containing, at the very least, an inductance, a noise source, a capacitor and two thermocouples.

THE EQUIVALENT CIRCUIT OF A RESISTOR IS NOT

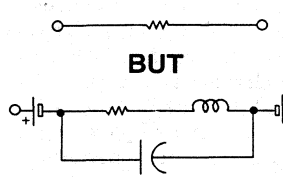

BUT

Figure 11.13

Inductive resistors

All resistors have some inductance (as we shall see, a straight piece of wire has some inductance) but wirewound resistors actually consist of a coil of wire, which must inevitably be inductive. Even if the coil is "non-inductive" and consists of N clockwise turns and N anticlockwise turns there will still be some mismatch and residual inductance. Residual inductance values of up to 20 μH can be expected in "non-inductive" wirewound resistors with values below 10 k ohms, although above 10 k ohms the reactance of such a resistor is more likely to be a capacitance of around 5 pF.

Some film resistors are also inductive, consisting of a spiral of resistive material on a cylindrical ceramic body. Again values of a few μH are typical. High frequency circuits must not use inductive resistors since their impedance is not equal to their resistance and, indeed, varies with frequency. Even low frequency circuitry, where the inductance of the resistors would not seem to be a problem, may suffer from instability arising from unforeseen HF effects of resistor inductance (the transistors used in low frequency op-amps frequently have Ft of up to 1 GHz).

THERMO-ELECTRIC EFFECTS

Wirewound resistors have another problem. The junction of the resistance wire and the lead forms a thermocouple which has a thermoelectric EMF of 42 μV/°C for the standard "Alloy 180"/Nichrome junction of an ordinary wirewound resistor. If a resistor is chosen with the [more expensive] copper/nichrome junction the value is 2.5 μV/°C. ("Alloy 180" is the standard component lead alloy of 77% copper and 23% nickel.)

Such thermocouple effects are unimportant at AC or where a resistor is at a uniform temperature but if the dissipation in a resistor, or its location with respect to heat sources, can cause one of its ends to be warmer than the other then there will be a net thermoelectric EMF which will introduce a dc error into the circuit. With a normal wirewound resistor a temperature differential of only 4°C will introduce a dc error of 168 μV - which is greater than 1 LSB in a 10 V/16-bit system.

The problem may be minimized by mounting wirewound resistors to ensure that temperature differentials are minimized. This may be done by ensuring that both leads are of equal length to equalize thermal conduction through them, by making any airflow (whether forced or natural convection) nor-

MINIMIZING THERMOCOUPLE EFFECTS IN WIREWOUND RESISTORS

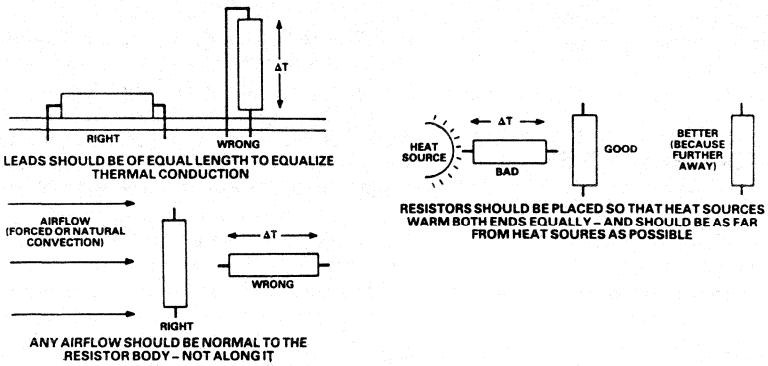

Figure 11.14

nal to the resistor body, and by taking care hat both ends of the resistor are the same listance from any heat source on the PCB. Notwithstanding these precautions it is wiser o use resistors with copper, rather than 'Alloy 180" leads, and to site them as far as possible from any heat source.

Stability & matching

Thermal effects other than thermocouple effects will also affect the accuracy of circuits using resistors. Resistors are never completely stable with temperature and if either the temperature coefficients, or the actual temperatures of two resistors in a precision circuit are mismatched then the performance of the circuit will suffer. Temperature mismatch of two identical resistors in similar environments may arise from differences in self-heating or other causes.[3]

GAIN OF 100 STAGE

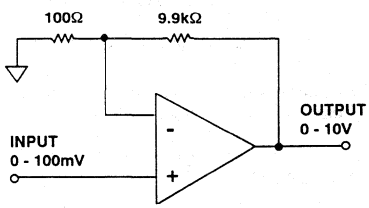

- Resistor mismatch due to mismatch of temperature coefficients, mismatch of temperature (possibly due to self-heating), or both, can cause gain errors.
- Ideally, all resistors whose matching can affect accuracy should be fabricated on a single substrate.

Figure 11.15

Typical temperature coefficients of discrete resistors are apt to be around 100 ppm/°C or more. The best way to minimize the effects of resistor temperature coefficients and to eliminate the effects of different resistor temperatures is to ensure that all resistors whose resistor matching affects the accuracy of a system are built on a single substrate. This substrate may be the glass or ceramic substrate of a thin film resistor network.

A better alternative, when possible, is to use an integrated circuit having laser trimmed thin film resistors on the silicon substrate of the IC. The temperature coefficient of such resistors can be well below 20 ppm/°C, and the differential temperature coefficient between two resistors on the same substrate is of the order of 0.5 ppm/°C or less.

Voltage variation of resistance

It is not possible to fabricate very high resistances on thin film or IC substrates, and high value discrete resistors are considerably less stable than lower value ones. It is inadvisable, therefore, to rely on the stability of high value resistors for the performance of a system. Some types of high value resistor have another imperfection: they have a slightly non-linear voltage/current curve and do not obey Ohm's Law accurately.

HIGH VALUE RESISTORS

- **Likely to be Less Stable**

and

- **Non-Linear With Voltage**

Figure 11.16

A final "imperfection" of resistors is an inconvenience but cannot properly be regarded as an imperfection as it is a fundamental property of all resistors: thermal or Johnson noise.

RESISTOR JOHNSON NOISE

- **All Resistors Have Noise: $V_N = \sqrt{4kTBR}$**
  - **T is Absolute Temperature**
  - **B is Bandwidth in Hertz**
  - **R is the Resistance in Ohms**
  - **k is Boltzmann's Constant (1.38E-23 J/K)**
- **It is possible to reduce the noise of a resistor by reducing T, B, or R but it is *NOT* possible to reduce k because Boltzmann is dead.**

Figure 11.17

At any temperature above absolute zero all resistors have noise due to thermal motion of their structure. This noise, which is described by

$$V_n = \sqrt{4kTBR}$$

(Where k is Boltzmann's constant: 1.38E-23 J/°K)

Johnson noise is present in ALL resistors and can only be reduced by reducing R, the resistance itself, B, the bandwidth of interest, or T, the temperature. Since the function involves a square root, the noise improvement for a drop in temperature from room temperature (298 °K) to liquid nitrogen (77 °K) is only of the order of 50%, so cooling a resistor, unless liquid helium is involved, is unlikely to be very profitable.

Johnson noise is purely an effect of resistance. The Johnson noise of complex impedances consists only of the Johnson noise of the resistive part of the impedance, so pure capacitance or inductance does not have Johnson noise, even though it has an impedance.

CAPACITANCE

STRAY CAPACITANCE

Where two conductors are not short-circuited together, or totally screened from each other by a conducting (Faraday) screen, there is a capacitance between them. There will therefore be a large number of capacitors associated with any circuit, which may or may not be considered in models of the circuit. Where high frequency performance matters (and even DC and VLF circuits may use devices with high Ft and therefore be vulnerable to HF instability) it is very important to consider the effects of stray capacity.

Any basic textbook will provide formulas for the capacitance of parallel wires, concentric spheres and cylinders, and many other configurations.[4] The only example we need consider in this seminar is the parallel plate capacitor, which is formed by conductors on opposite sides of a PCB.

Neglecting edge effects, the capacitance of two parallel plates of area A mm² and separation d mm in a medium of dielectric constant E_r relative to air is

$$0.00885\ E_r\ A/d\ \text{pF}.$$

CAPACITANCE

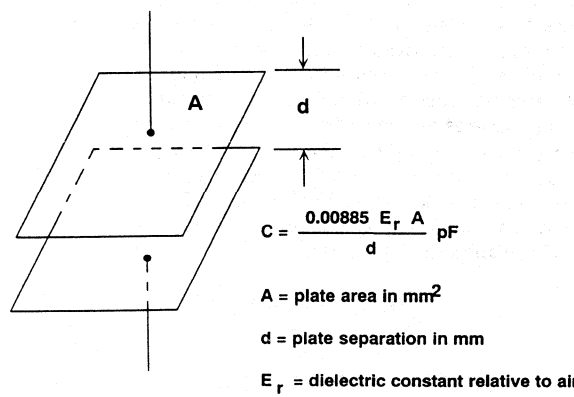

- **Commonest type of PCB uses 1.5mm glass-fiber epoxy material with $E_r = 4.7$**
- **Capacity of PC track over ground plane is roughly 2.8pF/cm²**

Figure 11.18

From this formula we can calculate that for general purpose PCB material ($E_r = 4.7$, d = 1.5 mm) the capacitance between conductors on opposite sides of the board is just under 3pF/cm². In general such capacitance will be parasitic, and circuits must be de-

signed so that it does not affect their performance, but it is possible to use PCB capacitance in place of small discrete capacitors. However the dielectric properties of common PCB materials (teflon is an expensive exception) cause such capacitors to have a rather high temperature coefficient and to have poor Q at high frequencies, which makes them unsuitable for many applications.

Capacitive noise & Faraday shields

There is a capacitance between any two conductors separated by a dielectric (air or vacuum is a dielectric). If there is a change of voltage on one there will be a movement of charge on the other. The basic model is shown in Figure 11.19.

CAPACITIVE COUPLING EQUIVALENT CIRCUIT

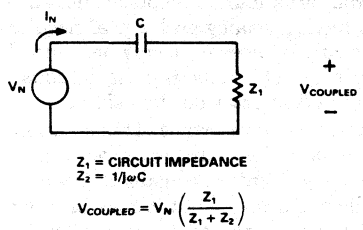

Figure 11.19

It is evident that the voltage coupled into Z1 may be reduced by reducing the signal voltage, Vn, the frequency involved, the capacitance, or Z1, but frequently none of these can be changed. The best solution is to insert a grounded conductor (known as a Faraday shield) between the noise source and the circuit which it affects.

The Faraday shield is easily implemented and almost invariably successful. For this reason capacitively coupled noise is rarely an intractable problem. However, to be effective the shield must completely block the electric field between the noise source and the shielded circuit and must be connected so that the noise current returns to its source without flowing in any part of the circuit where it might introduce conducted noise. A conductor intended as a Faraday shield must never be left unconnected as this almost always increases capacity and exacerbates the problem.

An example of this problem is seen in sidebrazed ceramic IC packages. These DIP packages have a small square conducting

CAPACITIVE SHIELDING

CAPACITIVE SHIELD INTERRUPTS THE COUPLING ELECTRIC FIELD

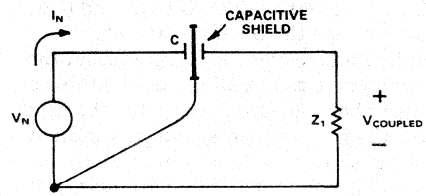

EQUIVALENT CIRCUIT ILLUSTRATES HOW A CAPACITIVE SHIELD CAUSES THE NOISE CURRENTS TO RETURN TO THEIR SOURCE WITHOUT FLOWING THROUGH Z_1

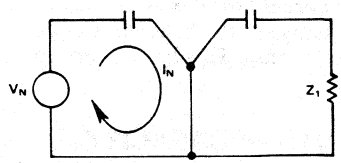

Figure 11.20

kovar lid soldered onto a metallized rim on the ceramic package top. Package manufacturers offer only two options: the metallized rim may be connected to one of the corner pins of the package or it may be left unconnected. Most logic circuits have a ground pin at one of the package corners and therefore the lid is grounded. Many analog circuits do not have a ground pin at a package corner and the lid is left floating - such circuits turn out to be far more vulnerable to electric field noise than the same chip in a plastic DIP package where the chip is completely unshielded.

CAPACITIVE EFFECTS DUE TO METAL LIDS

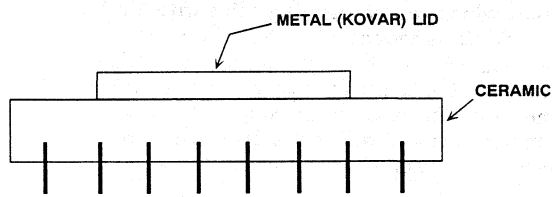

- SIDEBRAZE CERAMIC D.I.L. PACKAGES SOMETIMES HAVE ISOLATED METAL LIDS
- THESE ARE VULNERABLE TO CAPACITIVE INTERFERENCE AND SHOULD BE GROUNDED (IF POSSIBLE)

Figure 11.21

Whatever the environmental noise level, it is good practice for the user to ground the lid of any sidebrazed ceramic IC where the lid is not grounded by the manufacturer - this can be done with a wire soldered to the lid (this will not damage the device as the chip is thermally and electrically isolated from the lid). If soldering to the lid is unacceptable a grounded phosphor-bronze clip may be used to make the ground connection, or conductive paint from the lid to the ground pin. Never attempt to ground such a lid without verifying that it is, in fact, unconnected, as occasionally device types will be found with the lid connected to a power supply rather than to ground!

One case where a Faraday shield is impracticable is between the bondwires of an integrated circuit chip. This has important consequences.

STRAY CAPACITY BETWEEN CHIP BONDWIRES

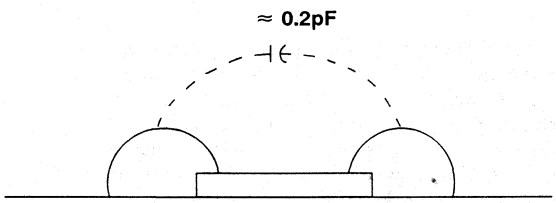

Figure 11.22

The stray capacitance between two chip bondwires and their associated leadframes is of the order of 0.2 pF. (Note this is "of the order" NOT "of the close order" - observed values generally lie between 0.05 and 0.6 pF.) If we have a high resolution converter (ADC or DAC) which is connected to a high speed data bus then each line of the data bus, which will be carrying noise with 2-5 V/ns dV/dT, is connected to the converter analog port via this stray capacitance. Whenever the bus is active intolerable amounts of noise will be capacitively coupled to the analog port and will seriously degrade the performance of which the converter is capable.

WITH A HIGH PERFORMANCE CONVERTER ON A HIGH SPEED DATA BUS, IT IS NOT POSSIBLE TO SHIELD THE ANALOG PORT FROM THE DIGITAL NOISE

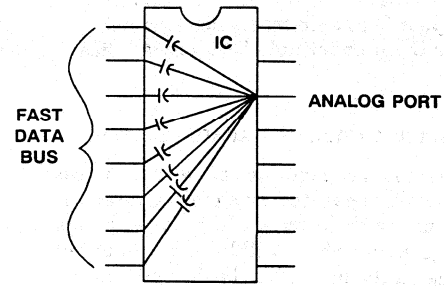

Figure 11.23

Present technology offers no cure for this problem, which also limits the performance possible from broadband monolithic mixed signal ICs having analog and digital circuitry on a single chip. However, it may be avoided quite simply by not connecting the databus directly to the converter but by using a latched buffer as an interface. This solution costs money, occupies board area, reduces reliability (very slightly), consumes power and complicates design - but it does improve the signal-to-noise ratio of the converter. The designer must decide whether it is worthwhile in individual cases.

BUFFER LATCH USED AS FARADAY SHIELD

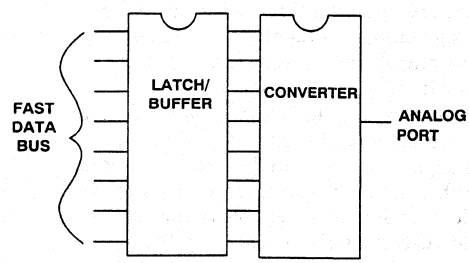

- A BUFFER/LATCH CAN ACT AS A FARADAY SHIELD BETWEEN A FAST DATA BUS AND A HIGH PERFORMANCE CONVERTER
- IT ADDS COST, BOARD AREA, POWER CONSUMPTION, RELIABILITY REDUCTION, DESIGN COMPLEXITY AND IMPROVED PERFORMANCE

Figure 11.24

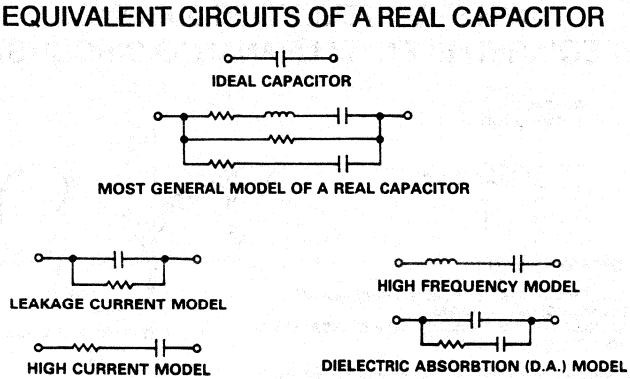

Figure 11.25

Parasitic effects in capacitors

Just as we are too willing to assume that a resistor is a perfect resistor, so do we underestimate the parasitic components associated with a capacitor. Fig 25 shows the ideal, the general model of a real capacitor, and the simplified models which are adequate for the analysis of non-ideal behavior in most applications.

Capacitors are used for coupling (passing AC signals while blocking DC), for decoupling (removing AC superimposed on DC in both power and signal circuitry), for building filters or frequency-selective networks, and for storing charge in "sample and hold" circuits (also known as "track and hold" circuits or SHAs, SAHs or THAs).

Capacitor Leakage

In coupling and SHA applications the leakage of the capacitor can be important. Electrolytic capacitors, where the dielectric is formed by an electrochemical reaction, have relatively high leakage currents of microamperes or even more and so are not used in applications where leakage matters. The leakage of electrolytic capacitors is greater during the first few minutes of operation after a period of storage (the leakage current while the capacitor is in use keeps the dielectric in good condition and it may deteriorate slightly in storage) - this feature can be important in equipment which must perform correctly after a long quiescent period.

The leakage of tantalum electrolytic capacitors is lower than that of aluminium ones and so in applications where capacitances of tens of microfarads or more (which can be easily achieved only with electrolytic capacitors) are required tantalum ones are used, despite their extra cost, if particularly low values of leakage current are necessary. At room temperature the leakage of aluminium electrolytic capacitors is of the order of 20 nA/μF and that of tantalum ones is 5 nA/μF.

Another feature of electrolytic capacitors, both aluminium and tantalum, is that most of them are polarized and require a DC bias for correct operation - a reverse bias may do damage and will certainly increase leakage (unpolarized electrolytic capacitors, which may be biased in either direction, do exist but they are uncommon, and considerably larger than the polarized variety).

Most other types of capacitor have leakage resistances in excess of hundreds of gigohms so that for most applications their leakage currents can be disregarded.

Series/loss resistance

The series resistance of capacitors causes them to dissipate power when high AC currents are flowing in them. This can have serious consequences at RF and in supply decoupling capacitors carrying high ripple currents but is unlikely to have much effect in precision analog circuitry. The series inductance, however, can have very inconvenient consequences.

Inductance of Capacitors

The transistors used in precision analog circuits have transition frequencies (Ft) of hundreds of MHz or even several GHz, even though the precision circuitry itself may be operating at DC or low frequencies. This makes it essential that the power supply terminals of such circuits should be decoupled properly at high frequency.

A common structure for capacitors is two sheets of metal foil separated by sheets of plastic or paper dielectric and formed into a

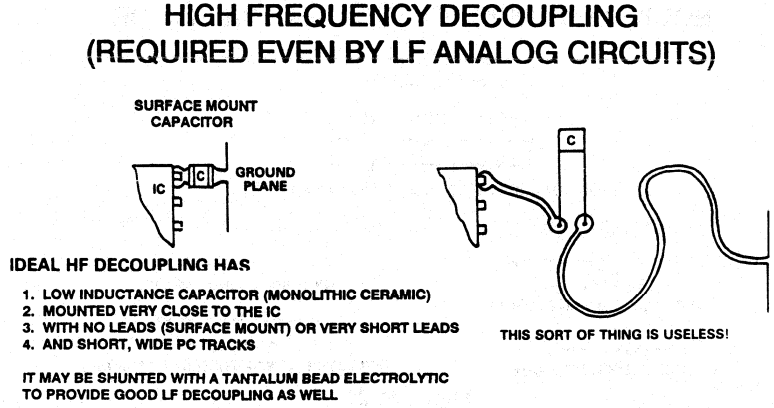

Figure 11.26

roll. Such a structure has considerable self inductance and behaves as an inductance rather than a capacitor at frequencies of more than a few MHz. It is therefore inadvisable to use electrolytic , paper or plastic film capacitors for decoupling at high frequencies.

Monolithic ceramic capacitors have very low series inductance (they are formed of a multilayer sandwich of metal films and ceramic dielectric and all the films are joined to a bus-bar rather than being connected in series). They are therefore ideal for high frequency decoupling. However, monolithic ceramic capacitors can be microphonic, and some types may be self-resonant with comparatively high Q. Disc ceramic capacitors, on the other hand, are sometime quite inductive, although less expensive.

The best way of ensuring that an analog circuit is adequately decoupled at both high and low frequencies is to use a tantalum bead capacitor in parallel with a monolithic ceramic one. The combination will have high capacitance but will remain capacitive at VHF frequencies. It is generally unnecessary to have a tantalum capacitor on each individual IC, if there is less than 10 cm of reasonably wide PC track between each IC and the tantalum capacitor it is possible to share one tantalum capacitor among several ICs.

There is little point in taking great care in the choice of a non-inductive capacitor if it is then unsuitably mounted. Short lengths of wire have appreciable inductance so HF decoupling capacitors must be mounted as close as possible to the points that they are decoupling with short, wide PC tracks. Ideally HF decoupling capacitors should be surface-mount parts to eliminate lead inductance, but wire-ended capacitors are permissible provided the device leads are no longer than 1.5 mm. It is also important to understand where HF decoupling currents should flow and why HF decoupling is more important at some points than at others - the subject is covered at some length in an Analog Devices Application Note.[5]

HF instability in analog circuits is more common than is realized. Oscillation at hundreds of MHz will cause serious malfunction of precision circuitry but may not affect an oscilloscope (indeed the presence of an oscilloscope probe may damp the oscillation, so that the circuit works only when an oscilloscope is attached to it - this is an important diagnostic clue). It is quite good practice to use a broadband spectrum analyzer (say 1-1500 MHz) and a low capacity FET probe to check for parasitic oscillation any analog circuit which is malfunctioning for no obvious reason. This test will also show if the malfunction is due to the presence of a strong RF field from an external source.

Dielectric absorption

Monolithic ceramic capacitors are excellent for HF decoupling but they have considerable dielectric absorption, which makes them unsuitable for use as the hold capacitor of an SHA. Dielectric absorption causes a capacitor which is quickly discharged and then open-circuited to recover some of its charge. Since the amount of charge recovered is a function of its previous charge this is, in effect, a charge memory and will cause errors in any SHA where dielectric absorption is present in the hold capacitor.

CAPACITORS HAVING SIGNIFICANT D.A. ARE USELESS FOR SAMPLE AND HOLD APPLICATIONS

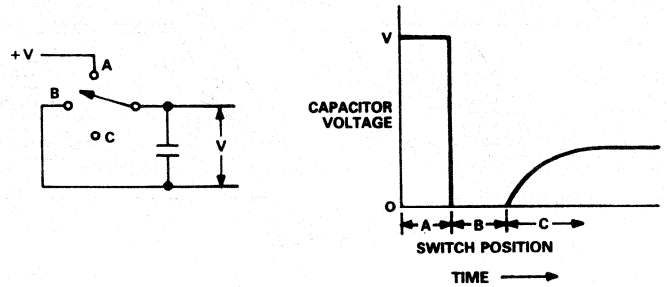

DIELECTRIC ABSORBTION CAUSES A BRIEFLY DISCHARGED CAPACITOR TO RECOVER A PERCENTAGE OF ITS PREVIOUS CHARGE ON BEING OPEN CIRCUITED.

Figure 11.27

Capacitors for this application should therefore be selected to have minimal dielectric absorption. The best strategy is to use a SHA which is supplied with an internal capacitor or where the SHA manufacturer supplies the capacitor with the SHA. If this is not possible (sometimes one may require a longer hold time - and hence extra capacity) a capacitor should be chosen which has its low dielectric absorption (DA) specified on its data sheet.

Such capacitors are normally plastic dielectric types (polystyrene, polypropylene or teflon) but it is not safe to use just any plastic dielectric capacitor with a SHA as special processing and testing is necessary to ensure that it has low DA. For use with a SHA a capacitor should be chosen which is specified for low DA applications.

INDUCTANCE

STRAY INDUCTANCE

All conductors are inductive and at high frequencies the inductance of even quite short pieces of wire may be important. The inductance of a straight wire of length L mm and circular cross-section with radius R mm in free space is

$$0.0002L\left[\ln\left(\frac{2L}{R}\right) - 0.75\right]\mu H.$$

The inductance of a strip conductor (an approximation to a PC track) of width W mm and thickness H mm in free space is

$$0.0002L\left[\ln\left(\frac{2 - L}{W + H}\right) + 0.2235\left(\frac{W + H}{L}\right) + 0.5\right]\mu H.$$

In real systems these formulas both turn out to be approximate but they do give some idea of the order of magnitude of inductance involved. They tell us that 1 cm of 0.5 mm o.d. wire has an inductance of 7.26 nH and 1 cm of 0.25 mm PC track has an inductance of 9.59 nH - these figures are reasonably close to measured results.

At 10 MHz an inductance of 7.26 nH has an impedance of 0.46 ohm and so can give rise to 1% error in a 50 ohm system.

INDUCTANCE

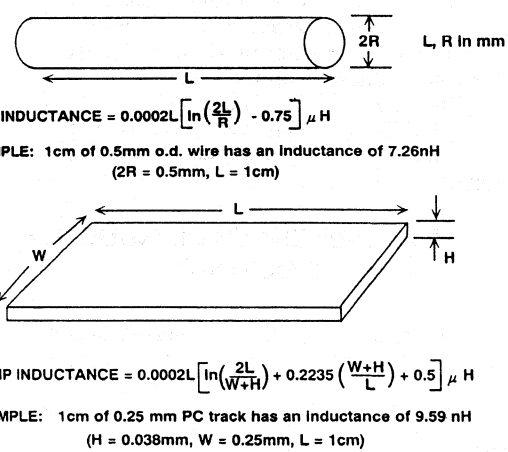

Figure 11.28

NONIDEAL AND IMPROVED SIGNAL ROUTING

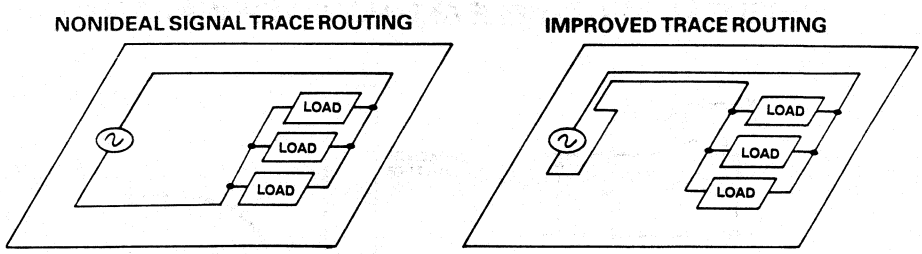

Figure 11.29

MUTUAL INDUCTANCE

Another consideration regarding inductance is the separation of outward and return currents. As we shall discuss in more detail later, Kirchoff's Law tells us that current flows in closed paths - there is always an outward and return path. The whole path forms a single-turn inductor. If the area enclosed by the turn is large the inductance, and hence the AC impedance, will also be large, whereas if the outward and return paths are close together the inductance will be much smaller. The principle is illustrated in Fig 11.29.

The nonideal routing in Figure 11.29 has another drawback - the large area enclosed by the conductor produces extensive external magnetic fields, which may interact with other circuits and cause unwanted coupling. Similarly the large area is more vulnerable to interaction with external magnetic fields, which can induce unwanted signals in the loop. The basic principle is illustrated in Figure 11.30 and is a common mechanism for the transfer of unwanted signals (noise) between circuits.

BASIC PRINCIPLES OF INDUCTIVE COUPLING

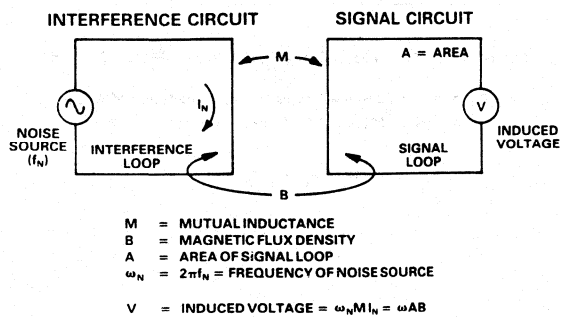

Figure 11.30

As with most other sources of noise, as soon as we define the principle at work we can see ways of reducing the effect. In this case reducing any or all of the terms in the equations in Figure 11.30 will reduce the coupling. Reducing the frequency or amplitude of the current causing the interference may be impracticable but it is frequently possible to reduce the mutual inductance between the interfering and interfered with circuits by reducing loop areas on both sides and, possibly, increasing the distance between them.

PROPER SIGNAL ROUTING REDUCES MUTUAL INDUCTANCE

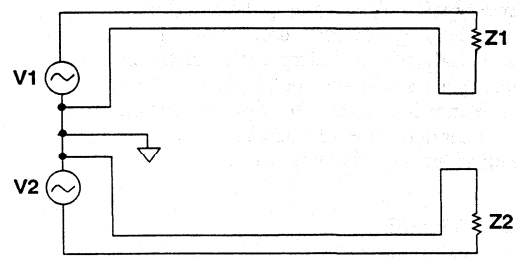

Figure 11.31

Mutual inductance is a common problem in ribbon cables, especially when a single return is common to several signal circuits. Separate signal and return lines for each signal circuit reduces the problem, and using a cable with twisted pairs for each signal circuit is even better (but more expensive and often unnecessary).

MUTUAL INDUCTANCE AND SIGNAL COUPLING IN RIBBON CABLE

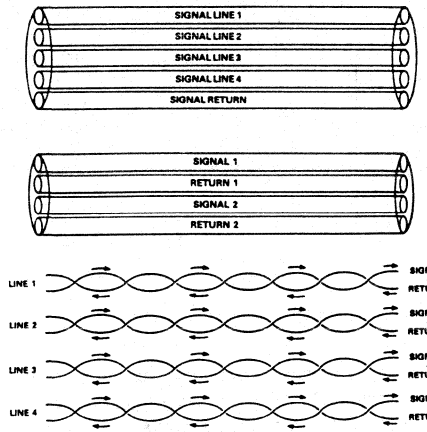

FLAT RIBBON CABLE WITH SINGLE RETURN HAS LARGE MUTUAL INDUCTANCE BETWEEN CIRCUITS

SEPARATE AND ALTERNATE SIGNAL AND RETURN LINES FOR EACH CIRCUIT REDUCE MUTUAL INDUCTANCE

TWISTED PAIRS REDUCE MUTUAL INDUCTANCE STILL FURTHER

Figure 11.32

Shielding magnetic fields to reduce mutual inductance is sometimes possible but is by no means as easy as shielding electric fields with a Faraday shield. HF magnetic fields are blocked by conductive material, while LF and DC fields may be screened by a shield made of mu-metal sheet. Mu-metal is an alloy having very high permeability, but it is expensive, its magnetic properties are damaged by mechanical stress, and it will saturate if exposed to too high fields. Its use, therefore, should be avoided where possible.

MAGNETIC SHIELDING

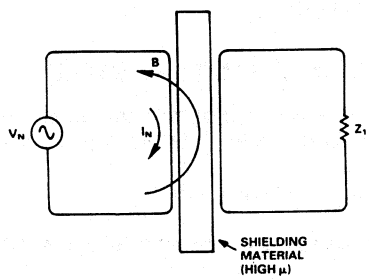

- **Magnetic shielding is not as easily accomplished as electrostatic shielding, but may be done at HF with a simple conducting screen, and at LF and DC with a screen of high permeability material such as Mu-metal.**

Figure 11.33

RINGING

An inductor in series or parallel with a capacitor forms a resonant, or "tuned", circuit, whose key feature is that it shows marked change in impedance over a small range of frequency (how sharp the effect is depends on the Q of the tuned circuit). The effect is widely used to define the frequency response of narrow-band circuitry but can also be a source of problems.

If stray inductance and capacitance (which may or may not be stray) in a circuit should form a tuned circuit then that tuned circuit may be excited by signals in the circuit and ring at its resonant frequency. A common example is shown in Figure 11.34 where the resonant circuit formed by an inductive power line and its decoupling capacitor may be excited by pulse currents drawn by the IC.

RESONANT CIRCUITS FORMED BY DECOUPLED POWER LINES

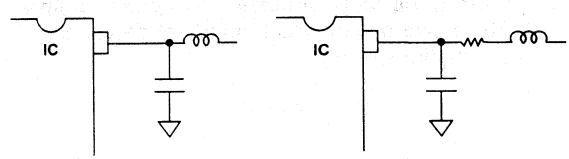

EQUIVALENT CIRCUIT OF DECOUPLED POWER LINE - RESONANT AT

$$f = \frac{1}{2\pi\sqrt{LC}}$$

SMALL SERIES RESISTANCE CLOSE TO THE IC REDUCES THE Q

Figure 11.34

The effect may be minimized by lowering the Q of the inductance, which is most easily done by inserting a small resistance in the power line, close to the IC.

PARASITIC EFFECTS IN INDUCTORS

Although inductance is one of the fundamental properties of an electronic circuit, inductors are less common as precision components than resistors and capacitors. This is because they are harder to manufacture, less stable, and less physically robust than resistors and capacitors. It is relatively easy to manufacture stable precision inductors with inductances from nH to tens or hundreds of μH, but larger valued devices tend to be less stable, and large.

As we might expect in these circumstances, circuits are designed, where possible, to avoid the use of precision inductors. We find that stable precision inductors are relatively rarely used in precision analog circuitry, except in tuned circuits for high frequency narrow band applications.

Of course they are widely used in power filters, switching power supplies and other applications where lack of precision is unimportant. The important features of inductors used in such applications are their current carrying and saturation characteristics, and their Q. If an inductor consists of a coil of wire with an air core its inductance will be essentially unaffected by the current it is carrying, but if it is wound on a core of a magnetic material (magnetic alloy or ferrite) its inductance will be non-linear since at high currents the core will start to saturate.

SATURATION

- **Inductors with solid cores (magnetic alloy or ferrite)**
  will behave non-linearly
  if required to carry too much current.
- **This is unlikely to be a direct problem in precision circuitry but may affect power supply noise performance and thus affect precision circuitry indirectly.**

Figure 11.35

Such saturation will reduce the efficiency of the circuitry employing the inductor and is liable to increase noise and harmonic generation.

As mentioned above, inductors and capacitors together form tuned circuits. Since all inductors will have some stray capacity, all inductors will have a resonant frequency (which will normally be published on their data sheet) and should only be used as precision inductors at frequencies well below this.

STRAY CAPACITANCE MAKES ALL INDUCTORS INTO TUNED CIRCUITS

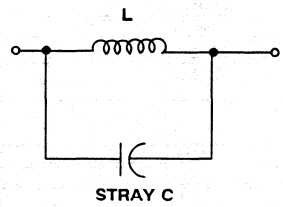

RESONANT FREQUENCY =

$$f = \frac{1}{2\pi\sqrt{LC}}$$

Figure 11.36

Q OR "QUALITY FACTOR"

The other parasitic characteristic of inductors is their Q (or "Quality Factor"), which is the ratio of their reactive impedance to their resistance.

$$Q = 2\pi f L / R$$

It is rarely possible to calculate the Q of an inductor from its DC resistance since skin effect (and core losses if the inductor has a magnetic core) ensure that the Q of an inductor at high frequencies is always lower than that predicted from DC values.

Q is also a characteristic of tuned circuits (and of capacitors - but capacitors generally have sufficiently high values of Q that it may be disregarded for most practical purposes). The Q of a tuned circuit, which is generally very similar to the Q of its inductor (unless it is deliberately lowered by the use of an additional resistor), is a measure of its bandwidth around resonance.

Q or "QUALITY FACTOR"

- **The Q of an inductor or resonant circuit is a measure of the ratio of its reactance to its resistance.**
  $Q = 2\pi fL/R$
- **The resistance is the HF and *NOT* the DC Value.**
- **The 3dB bandwidth of a single tuned circuit is Fc/Q where Fc is the center frequency.**

Figure 11.37

LC tuned circuits rarely have Q of much more than 100 (3 dB bandwidth of 1%) but ceramic resonators may have Q of thousands and quartz crystals have Q of tens of thousands.

IGNAL RETURN CURRENTS

Kirchoff's Law tells us that at any point in circuit the algebraic sum of the currents is ero. This tells us that all currents flow in rcles and, particularly, that the return rrent must always be considered when nalyzing a circuit.[6]

KIRCHOFF'S LAW

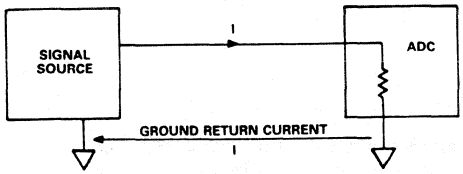

AT ANY POINT IN A CIRCUIT THE ALGEBRAIC SUM OF THE CURRENTS IS ZERO OR WHAT GOES OUT MUST COME BACK
WHICH LEADS TO THE CONCLUSION THAT ALL VOLTAGES ARE DIFFERENTIAL (EVEN IF THEY'RE GROUNDED)

Figure 11.38

Most people consider the return current vhen considering a fully differential circuit, ıut when considering the more usual circuit vhere a signal is referred to "ground" it is :ommon to assume that all the points on the :ircuit diagram where the ground symbol is :o be found are at the same potential. This is ınwise.

THE IDEAL GROUND

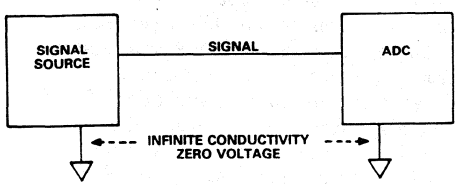

Figure 11.39

GROUND NOISE & GROUND LOOPS

A more realistic model of ground is shown in Figure 11.40. Not only does the return current flow in the complex impedance which exists between the two "ground" points shown in Figure 11.39, giving rise to a voltage drop in the total signal path, but external currents may also flow in the same path, generating uncorrelated noise voltages which are seen by the ADC.

A MORE REALISTIC GROUND

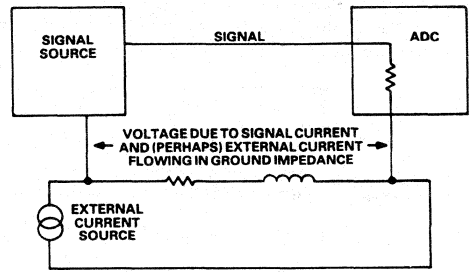

Figure 11.40

It is evident, of course, that other currents can only flow in the ground impedance if there is a current path for them. Figure 11.40 shows such a path at "ground" potential, which is the notorious "Ground Loop", but equally severe problems could be caused by a circuit sharing an unlooped ground return with the signal source but drawing a large and varying current from its supply and ground return.

ANY CURRENT FLOWING IN A COMMON GROUND MAKES NOISE; A GROUND LOOP IS NOT NECESSARY

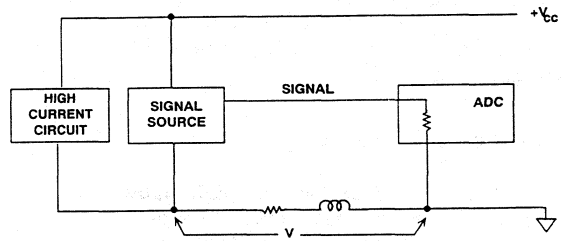

V = VOLTAGE DUE TO SIGNAL CURRENT AND CURRENT FROM HIGH CURRENT CIRCUIT FLOWING IN GROUND IMPEDANCE

Figure 11.41

It is evident from Figure 42 that if a ground network contains loops there is a greater danger of it being vulnerable to EMFs induced by external magnetic fields, and of ground current "escaping" from high current areas to cause noise in sensitive regions. For these reasons ground loops are best avoided.

However, there are situations where looped grounds are unlikely to cause unacceptable noise and the configuration may actually offer benefits in the form of safety or reduced impedance. In such circumstances the optimum ground arrangement may

GROUND LOOP

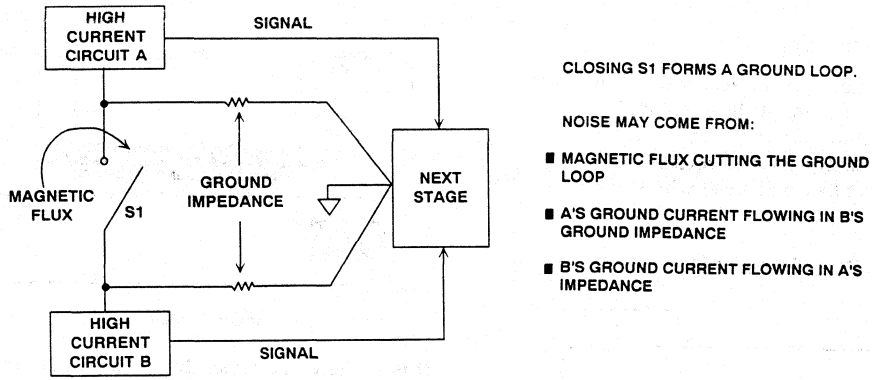

Figure 11.42

contain loops. Sensible engineers should not allow the almost superstitious dread inspired by the term "ground loop" to prevent the adoption of such designs, if careful analysis and experiment has shown that they actually are optimum.

There are a number of possible ways of attacking the problem of ground noise, apart from the (presently) impracticable one of using superconducting grounds. It is rare for a single method to be used to the exclusion of all others, and systems generally contain a mixture of approaches. For the purposes of description, however, it is better to describe each approach separately.

Star (Mecca) Grounds

The "star" or "Mecca" ground philosophy builds on the theory that there is a single point in a circuit to which all voltages are referred. This is known as the "star" or "Mecca" point.

STAR (MECCA) GROUNDS

If all signal voltages in a system are measured with respect to a single point, that point is said to be the *star* ground of the system.

Figure 11.43

This philosophy is reasonable but frequently encounters practical difficulties. For example if we design a system with a star ground, drawing all the signal paths to minimize signal interaction and the effects of high impedance signal or ground paths, we frequently find, when the power supplies are added to the circuit diagram, that the power supplies either add unwanted ground paths or that supply currents, flowing in existing ground paths, are sufficiently large, or noisy, or both, as to corrupt the signal transmission. This problem may often be avoided by having separate power supplies for different parts of the circuit - separate analog and digital supplies, and separate analog and digital grounds joined at the star point, are common in mixed signal applications.

Separate analog and digital grounds

Digital circuitry is noisy. Saturating logic draws large fast current spikes from its supply during switching and, having noise immunity of hundreds of millivolts or more, has little need of high levels of supply decoupling.

Analog circuitry, on the other hand, is very vulnerable to noise in supplies or grounds. It is therefore sensible to separate analog and digital circuitry to prevent digital noise from corrupting analog performance. Such separation will involve separation of both grounds and power supplies, which may be inconvenient in a mixed signal system. Nevertheless, if a system is to give the full performance of which it is capable it is often essential to have separate analog and digital grounds and power supplies. The fact that some analog circuitry will operate from a single +5 V supply does NOT mean that it may safely be operated from the same noisy +5 V supply as the microprocessor and dynamic RAM, the electric fan, and the solenoid jackhammer!

SUPPLY & GROUND NOISE

- **Digital circuitry is noisy**
- **Analog circuitry is quiet**
- **Circuit noise from digital circuitry carried by power and ground leads can corrupt precision analog circuitry**
- **It is advisable to separate the power and ground of the digital and analog parts of a system**
- **Analog and digital grounds must be joined at ONE point**

Figure 11.44

However, analog and digital ground in a system must be joined at some point to allow signals to be referred to a common potential. This star point, or analog/digital common point, is chosen so that it does not introduce digital currents into the ground of the analog part of the system - it is often convenient to make the connection at the power supplies.

Many ADCs and DACs have separate "analog ground" and "digital ground" pins, and users are advised, on the data sheets, to connect these pins together at the device package. This seems to conflict with the advice to connect analog and digital ground at the power supplies, and, in systems with more than one converter, with the advice to join the analog and digital ground at a single point.

ANALOG GROUND & DIGITAL GROUND

- **Monolithic & hybrid ADCs frequently have separate AGnd & DGnd pins which must be joined together at the device.**
- **This is not done from a desire to be difficult, but because the voltage drop in the bondwires is too large to allow the connection to be made internally.**
- **The best solution to the grounding problem arising from this requirement is to connect both pins to system "analog ground".**
- **It is likely that neither the digital noise so introduced in the system AGnd, nor the loss of digital noise immunity, will seriously affect the system performance.**

Figure 11.45

There is, in fact, no conflict. The labels "analog ground" and "digital ground" on these pins refer to the parts of the converter to which the pins are connected, and not to the system grounds to which they must go.

In general these two pins should be joined together and to the analog ground of the system. It is not possible to join the two pins within the IC package because the analog part of the converter cannot tolerate the voltage resulting from the digital current flowing in the bond wire to the chip.

If these pins are connected in this way the digital noise immunity of the converter is diminished by the amount of common- mode noise between the digital and analog system grounds. Since digital noise immunity is of the order of hundreds or thousands of millivolts this is unlikely to be important.

ANALOG GROUND (AGND) AND DIGITAL GROUND (DGND) OF ADCs/DACs SHOULD BE RETURNED TO SYSTEM ANALOG GROUND

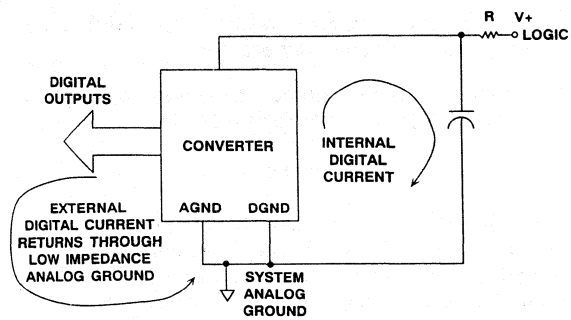

Figure 11.46

The analog noise immunity is diminished only by the external digital currents of the converter itself flowing in the analog ground. These currents should be quite small, and can be minimized by ensuring that the converter outputs do not drive large fanouts. If the logic supply to the converter is isolated with a small resistance and decoupled to analog ground with a 0.1µF capacitor sited as close to the converter as possible all the internal digital currents of the converter will return to ground through the capacitor and will not appear in the external ground circuit. If the analog ground impedance is as low as it should be for adequate analog performance the additional noise due to the external digital ground current should rarely present a problem.

GROUND PLANES

Related to the star ground system is the use of a ground plane. One side of a double-sided PCB, or one layer of a multi-layer one, is made of continuous metal, which is used as ground. The theory behind this is that the large amount of metal will have low resistance and as low inductance as is possible.

GROUND PLANES

- **One entire side or layer of a PCB is continuous grounded conductor.**
- **This gives minimum ground resistance and inductance but is not always sufficient to solve all grounding problems.**
- **Breaks in ground planes can improve or degrade circuit performance - there is no general rule.**
- **Twenty years ago ground planes were difficult to fabricate. Today they are not.**
- **If your PCB facility objects to fabricating ground planes - GET A NEW PCB FACILITY!**

Figure 11.47

It is sometimes argued that ground planes should not be used because they are liable to introduce problems in manufacture and assembly. Such an argument may have had a limited validity twenty years ago when PCB adhesives were less well developed, wave-soldering less reliable, and solder resist techniques less well understood, but today it should not be tolerated.

Ground planes solve many ground impedance problems, but not all. Even a continuous sheet of copper foil has residual resistance and inductance and in some circumstances they can be enough to prevent proper circuit function. Figure 11.48 shows such a problem - and a possible solution.

Consider a ground-plane PCB 100 mm wide with a ground connection at one end and a power amplifier at the other drawing 15A. If the ground plane is 0.038 mm thick and 15 A flows in it there will be a voltage drop of 68 μV/mm. This voltage drop would cause quite serious problems to any ground-referenced precision circuitry sharing the PCB. However, if we slit the ground plane so that high current does not flow in the region of the precision circuitry we can possibly solve the problem - even though the voltage gradient will increase in those parts of the ground plane where the current does flow.

A SLIT IN A GROUND PLANE CAN RECONFIGURE CURRENT FLOW FOR BETTER ACCURACY

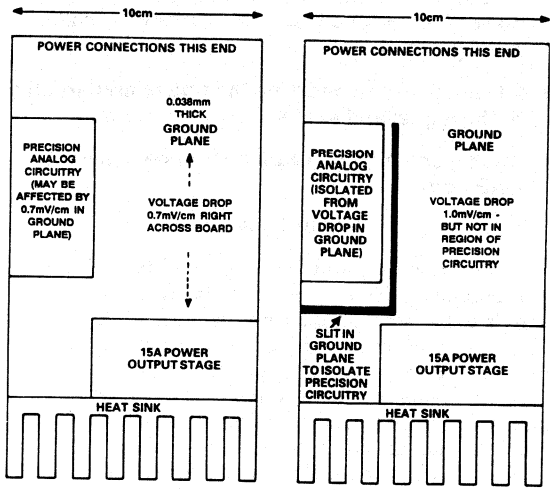

Figure 11.48

TRANSMISSION LINES

A break in a ground plane is not always a good thing. We earlier considered the benefits of outward and return signal paths being close together so that inductance is minimized. As we saw in Figure 11.7, when an HF signal flows in a PC track running over a ground plane the arrangement functions as a microstrip transmission line and the majority of the return current flows in the ground plane underneath the line.

The characteristic impedance of the line will depend upon the width of the track and the thickness and dielectric constant of the PCB material. For most lower frequency applications the characteristic impedance will be unimportant, as the line will not be correctly terminated, but at UHF and higher

MICROSTRIP TRANSMISSION LINE

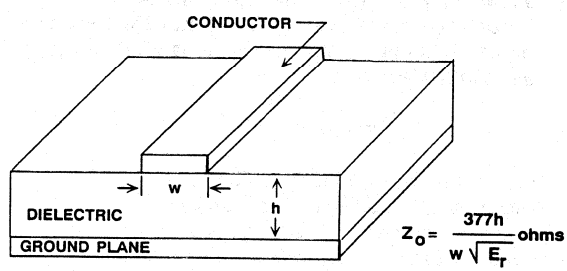

Figure 11.49

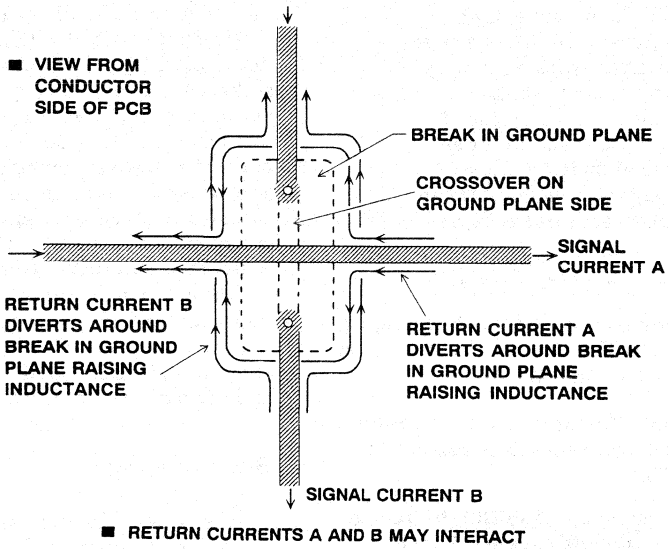

Figure 11.50

it is possible to use PCB tracks as microstrip transmission lines in properly terminated systems. If losses in such systems are to be minimized the PCB material must be chosen for low high frequency loss. This usually means the use of expensive teflon PCB material.

Where there is a break in the ground plane under a conductor the return current must flow around the break and both the inductance and the vulnerability of the circuit to external fields are increased.

Where such a break is made to allow a crossover of two perpendicular conductors it would be far better if the second signal were carried across both the first and the ground plane by means of a piece of wire. The ground plane then acts as a shield between the two signal conductors, and the two ground return currents, flowing in opposite sides of the ground plane as a result of skin effects, do not interact.

With a multi-layer board both the crossover and the continuous ground plane can be accommodated without the need for a wire link. Multi-layer PCBs are expensive and harder to trouble-shoot than simple double-sided boards but do offer even better shielding and signal routing. The principles involved remain unchanged but the range of layout options is increased.

Use of double-sided or multi-layer board with at least one continuous ground plane is undoubtedly one of the most successful approaches to the design of high performance mixed signal circuitry. Often the impedance of the ground plane is sufficiently low to permit the use of a single ground plane for both analog and digital parts of the system, but this does depend upon the resolution and bandwidth required and the amount of digital noise in the system.

System grounds

In systems where there are several PCBs grounding may be more of a problem. At first sight it would appear that the problem is similar to that of a single PCB where particular subsystems must be positioned so that large ground currents do not flow where ground noise must be minimized - in a multi-card system the grounds of individual PCBs must be interconnected so that such harmful interactions are minimized.

There are three problems with this. First of all there is far less opportunity for rearranging the physical layout of a system consisting of a few cards connected to a common backplane. Secondly many multi-card systems are designed to be reconfigured in a "mix 'n' match" arrangement to allow large numbers of system options - it can be impossible to predict what systems are going to be required and to ensure that all of them are noise free. Finally, multicard systems are likely to have higher ground currents than occur on single, relatively simple, PCBs - but these currents must flow in the higher impedances which are associated with the intercard connectors even when multiple ground pins are used.

The basic principles still apply: ground impedance must be as low as possible, high level and low level signals must be separated so that they do not interfere with each other, and capacitance and mutual inductance coupling must be avoided. Nevertheless, it must be accepted that situations can arise where it is not possible to transfer a high speed, high accuracy signal from one PCB to another without unacceptable signal degradation.

MULTIPLE CARD SYSTEMS

- **Multiple card systems are likely to have higher ground currents and higher ground impedances than are found on a single PCB.**
- **It is therefore more difficult to transfer ground-referenced signals accurately between cards than across a PCB.**
- **In some cases it will be IMPOSSIBLE to transfer ground-referenced signals between PCBs without unacceptable loss of quality.**

Figure 11.51

The best way of minimizing ground impedance in a multicard system is to use another PCB as a backplane and have a ground plane (or even two - one analog, one digital) on that mother card. If the earlier advice about multiple ground pins has been observed this arrangement is capable of excellent performance. Where there are several card cages (racks for PCBs) the ground planes of the several mother boards must be tied together and, probably, to the metal chassis holding the card cages - the exact layout of the interconnections will depend on the overall system architecture.

If a mother board with a ground plane is not possible then the ground pins of the PCB sockets must be wired together, with due attention to probable current flows and common ground impedances, with heavy, multi-strand wire, having as low resistance as possible. In many cases the resulting ground screen will be tied to chassis ground at a number of points but it will sometimes be better to join them at a single star point.

It is not just the ground layout that is important in high performance mixed signal systems, the siting of different subsystems and the routing of signals is most important in determining overall system performance.

SIGNAL ROUTING

It is evident that we can minimize noise by paying attention to the system layout and preventing different signals from interfering with each other. High level analog signals should be separated from low level analog signals, and both should be kept away from digital signals. We have seen elsewhere that in waveform sampling and reconstruction systems the sampling clock (which is a digital signal) is as vulnerable to noise as any analog signal, but is as liable to cause noise as any digital signal, and so must be kept isolated from both analog and digital systems.

STAR ANALOG GROUND IN A MULTICARD SYSTEM

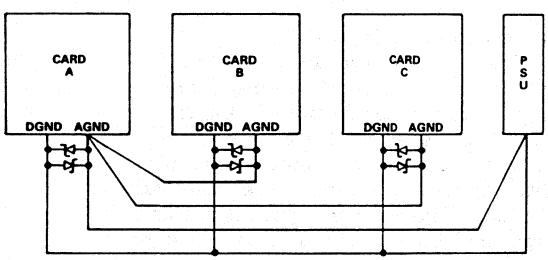

- **Schottky diodes protect cards in the event of loss of analog ground**
- **This grounding system may be inadequate at high resolution or where large ground currents flow**
- **This MAY permit accurate intercard transmission of ground referenced signals**

Figure 11.52

SIGNAL ROUTING IN MIXED SIGNAL SYSTEMS

- **Physically Separate Analog and Digital Signals**
- **Avoid Crossovers Between Analog and Digital Signals**
- **Be Careful with Sampling Clock and A/D Converter Analog Input Runs**
- **Be Careful with High Impedance Points**
- **Use Lots of Ground Plane**
- **Use Microstrip Techniques for Controlled Impedances**

Figure 11.53

If a ground plane is used, as it should in be most cases, it can act as a shield where sensitive signals cross. Figure 11.54 shows a good layout for a data acquisition system

here all sensitive areas are isolated from ıch other and signal paths are kept as short ; possible. While real life is rarely as tidy as ıis the principle remains a valid one.

PCB FLOWCHART

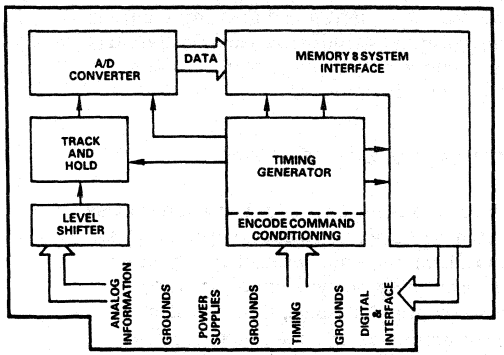

Figure 11.54

There are a number of important points to ɔe considered when making signal and power ɔonnections. First of all a connector is one of the few places in the system where all signal conductors must run parallel - it is therefore a good idea to separate them with ground pins to reduce coupling between them.

EDGE CONNECTIONS

- **Separate sensitive signals by ground pins**
- **Keep down ground impedance with multiple (20-30% of total) ground pins**
- **Have several pins for each power line**
- **Critical signals may require a separate connector (possibly co-ax)**

Figure 11.55

Multiple ground pins are important for another reason: they keep down the ground impedance at the junction between the board and the backplane. The contact resistance of a single pin of a PCB connector is quite low (of the order of 10 mohms) when the board is new - as the board gets older the contact resistance is likely to rise, and the board's performance may be compromised. It is therefore well worthwhile to afford extra PCB connector pins so that there are many ground connections (perhaps 20-30% of all the pins on the PCB connector should be ground pins). For similar reasons there should be several pins for each power connection, although there is no need to have as many as there are ground pins.

Modern high performance mixed signal systems handle signals with resolutions of 8 bits at sampling rates of over 500 MHz and resolutions of 14 bits sampled at more than 10 MHz. Preserving signal integrity between cards in a multi-card system is extremely difficult at such performance levels and may be impossible.

DIFFERENTIAL TRANSMISSION MINIMIZES GROUND ERRORS

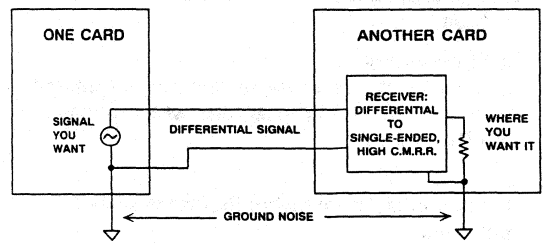

- **At DC and LF the receiver will be an instrumentation amplifier**
- **At HF the receiver will be a transformer**
- **There is no ideal receiver for video signals which have components from DC to HF**

Figure 11.56

The use of balanced transmission lines can help but if the signal bandwidth extends to DC there will be a need for a very high performance instrumentation amplifier at the receiving end to restore a ground referenced signal.

VIDEO SIGNAL TRANSMISSION

- **It is often impossible to transmit very broadband high accuracy signals between the PCBs of a multicard system without unacceptable loss of quality.**
- **In such cases the system must be reconfigured to allow all the analog processing to take place on a single PCB.**
- **It may be inconvenient, but it's the only way you'll get it to work!**

Figure 11.57

The best, and in many cases the only, solution to problems of this sort, is to partition the system so that the highest quality signals are not transferred between boards.

Power Supplies

When we design an electronic circuit we generally assume that the power supplies provide noise-free power, at exactly the nominal voltage, with zero source impedance at all frequencies. This is rarely the case.

We also assume that the published power supply rejection figures (PSRR) for the devices which we use are valid at all frequencies from DC to light. This is rarely the case either.

POWER SUPPLY NOISE

- **Long-term voltage variation (Long-term variations in voltage or AC line voltage)**
- **AC Line noise (Both 100/120 Hz ripple on rectifier output and transient noise on the AC line which passes to the DC output)**
- **Switching noise (Digital noise from switching-mode power supplies)**
- **Power line noise transfer (Unwanted signals which pass from one part of a circuit to another via the common power supply)**

Figure 11.58

Power supply noise

Every power supply is noisy. This noise may contain long-term voltage drift, line ripple at 100 or 120 Hz, high frequency spikes from switching regulators, or all of these at once. Power supplies also have finite output impedance, so that if a circuit draws a varying current the supply voltage will vary with the current - if two circuits are supplied from a common supply this provides a mechanism whereby one circuit may affect the other. Once we appreciate all these effects we can attempt to quantify them, and take steps to minimize their adverse effects on our systems.

Long-term supply voltage changes, whether due to battery voltage drop during life or line voltage variations, are rarely a problem since where such variations might cause difficulties the system will incorporate a supply voltage regulator to keep variations within acceptable limits. Similarly ripple at twice the AC line frequency, and any spikes or HF noise which may enter the system via the AC supply, should not cause degradation of performance in a well-designed system: if the decoupling capacitors in the rectifier circuitry do not adequately minimize the effect, the series regulator almost certainly will. It is, however, always worthwhile to have a surge eliminator on the AC line input to any system - while such a circuit is unlikely to be needed in preventing normal line noise from currupting system performance, it is essential to prevent occasional large surges (from lightning or similar causes) from doing actual damage to the power supply or the system that it is powering.

Switching-mode power supplies

The commonest type of power supply noise is switching noise. Switching power supplies are small, cheap, efficient and, in too many cases, extremely noisy! Not only do they generate conducted noise, they are also efficient producers of capacitively coupled noise, magnetically coupled noise, and electromagnetically coupled noise. The best possible advice is not to use them.

SWITCHING-MODE POWER SUPPLIES

- **Generate every imaginable type of noise and some inconceivable ones as well!**
- **DO NOT USE THEM WHERE NOISE IS IMPORTANT.**
- **If their use is unavoidable do not relax and enjoy it, *but* take extreme precautions against all forms of noise.**
- **Remember that a manufacturer's design change in a bought-in switching-mode power supply may alter its effects on your system noise without altering its published specification.**

Figure 11.59

It is, unfortunately, not always possible to avoid the use of switching power supplies. Where they must be used they must be treated with the gravest suspicion and all possible precautions should be taken to prevent their noise from corrupting the analog circuits that they power. Their input and output lines should be decoupled at all frequencies, they should be shielded to prevent external electric and magnetic fields from causing interference, and they should be sited as far as possible from sensitive circuits so that residual electric and magnetic fields are prevented by distance from doing serious damage.

Where switching supplies are used it is lways worthwhile to remove them tempoarily and supply the system with batteries r a low noise bench supply in order to deterine if the system performance is being ompromised by the switching supply. It ften is.

The noise transients on the output lines of witching supplies consist of voltage spikes of ery short duration. As we have pointed out bove, large capacitors, such as electrolytic r plastic film types, have quite considerable nductance and too high an impedance at HF o decouple such spikes satisfactorily. The est output filter for a switching supply will ave high value capacitors to remove the low requency noise which will also be present, nd a pi filter using ceramic capacitors, with hort leads having low impedance at HF, plus series inductor (which may be a ferrite ead on the output line) to provide inductive locking of the spikes. It is possible to buy uch a pi filter as a single bulkhead mounted eedthrough component.

Electromagnetic Interference

Radio frequency interference

Noise can enter a circuit as electromagnetic radiation. Circuits can also generate electromagnetic radiation which can interfere with electronic devices at quite considerable distances away. Recent legislation in the United States, the European Community and many other countries sets limits on the amount of interference generated and the vulnerability of circuits to such interference.[7]

This legislation, and the techniques needed to comply with it, are the subjects of many seminars and training courses, and an Analog Devices Application Note.[8] It is not proposed to cover the topics in detail in this seminar.

ELECTROMAGNETIC NOISE GENERATION

- **Circuits must be designed so that external E/M fields are minimized.**
- **This is done by shielding, decoupling, minimizing the area of HF current loops and designing circuits which generate as little EMI as possible.**
- **IT'S NOT JUST A GOOD IDEA**
- **IT'S THE LAW!**

Figure 11.60

However, the principles of minimizing external radiation are closely related to the principles of low noise design which we have already discussed: high frequency and high dV/dT signals should be screened with Faraday shields, the area of current loops should be minimized, conductors should be decoupled at HF wherever unnecessary HF signals might otherwise occur, and external wires should be isolated with inductors or ferrite beads.

It is still too common at seminars like this to encounter skepticism about the need to protect circuitry from external electromagnetic fields. Even twenty years ago such skepticism was unjustified but today, when transmitters are ubiquitous, it is folly. Besides the more obvious broadcast, emergency and mobile radio services there are cellular and cordless telephones, radar, garage door openers and other remote controls, telemetry, and amateur and CB radio. For any designer to imagine that his circuit will never encounter a radio transmitter during its lifetime is folly on a grand scale.

ELECTROMAGNETIC NOISE INTERFERENCE

- **The World is full of radio transmitters.**
- **Police, taxis, broadcast, amateur, CB, cellular and cordless telephones, telemetry and garage door openers.**
- **Do not imagine that your circuit will never encounter one.**

Figure 11.61

This is particularly so because the design of circuits which are immune to electromagnetic radiation of reasonable levels is not particularly difficult. If every conductor which leaves a PCB can be decoupled with a ceramic capacitor and a ferrite bead, it is probable that no further precaution is necessary.

EMI PREVENTION

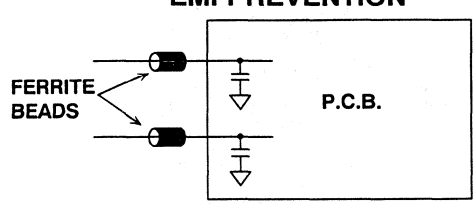

IN MANY CASES, ALL THAT IS REQUIRED IS AN L FILTER, CONSISTING OF A FERRITE BEAD AND A CAPACITOR, ON EACH EXTERNAL CONNECTION TO THE BOARD

Figure 11.62

A few ports may be more vulnerable and require a pi filter rather than an L filter, and, of course, ports where an HF signal must actually enter or leave the board must be filtered to suppress other EMI but allow the signal to pass unaffected.

Boards which may be required to work in areas of high RF field should be screened with a conducting Faraday shield.

PHOTOELECTRIC EFFECTS

Light is also a form of electromagnetic radiation and can effect semiconductor devices. Every silicon P-N junction is a photodiode, although their efficiencies vary widely. Wherever devices are not screened from ambient light, photoelectric effects may be observed.

Nearly all integrated circuits are encapsulated in light-tight packages (EPROMs are an exception, and it is possible to measure threshold changes in EPROMs as light intensity is varied, but since they are digital devices, and remain in specification despite light level changes, the effect is unimportant).

Diodes, on the other hand, are frequently encapsulated in translucent glass packages. When illuminated by light from fluorescent lamps, modulated at 120 or 100 Hz, they can act as a source of hum.

When the signal source of an op-amp contains an energizing voltage which is much higher than the op-amp supply it is common to use a diode and a current limiting resistor to protect the op-amp in the event of a sensor short-circuit. In normal operation the diode is reverse biassed and contributes only its (low) leakage current to the circuit but should the sensor be short-circuited the resulting current will flow through the diode to the op-amp supply rather than destroy the op-amp. It is, of course, important to choose the resistor so that it neither degrades the noise performance of the system nor allows too much current to pass under fault conditions.

The European Applications Department of Analog Devices encountered such a system where about 10% of all the amplifiers built suffered from severe hum at twice the power line frequency. The customer, of course, blamed the op-amp for poor supply rejection but analysis showed that even when the circuit was powered from batteries the problem persisted. The cause eventually turned out to be fluorescent lighting affecting the protective diode - a 1N914 in a glass case.

About 10% of diodes from the particular manufacturer were quite active as photodiodes and when illuminated by fluorescent lights their leakage current was modulated at 100 Hz (this was a European problem) - and the 100 Hz was, of course, amplified with the sensor signal. Use of a black epoxy packaged diode provided a complete cure.

UNEXPECTED PHOTOELECTRIC EFFECTS IN SILICON JUNCTIONS CAN DEGRADE CIRCUIT PERFORMANCE

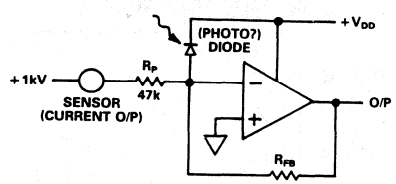

THE DIODE PROTECTS THE OP AMP UNDER FAULT CONDITIONS BY DIVERTING FAULT CURRENT (LIMITED BY R_P) TO THE SUPPLY RAIL. THE DIODE SHOULD NOT BE PHOTO-RESPONSIVE, OTHERWISE FLUORESCENT LIGHTING MAY MODULATE ITS LEAKAGE CURRENT AT 100/120Hz AND CAUSE HUM. USE A PLASTIC DIODE – NOT A GLASS ONE.

Figure 11.63

LOGIC

The majority of this section of our seminar has considered problems within the analog parts of mixed signal systems. Despite their much greater noise immunity, the digital parts of these systems can also suffer from designers' lack of consideration of basic laws. Common problem areas include bus interface issues, including fan-out and timing, for both converters and DSP processors, the care and feeding of sampling clocks, and the design of systems which generate minimum noise (we have already discussed how to keep logic noise from affecting the analog parts of a system - this task becomes easier if the logic noise is minimized in the first place).

DIGITAL PROBLEM AREAS IN MIXED SIGNAL SYSTEMS

- Bus interface - fan-out
- Timing variations
- Sampling clock jitter
- Logic noise

Figure 11.64

FAN-OUT

All Analog Devices' DSP processors, and most of their DSP ADCs, have TTL-compatible CMOS logic ports. The inputs have $V_{il(max)}$ of 0.8 V and $V_{ih(min)}$ of 2.0 V, while the outputs have $V_{ol(max)}$ of 0.4 V and $V_{oh(min)}$ of 2.4 V at particular currents. The DSP processors are also rated for the capacitive load that they will drive without degradation of their timing.

In order to determine the fan-out of such devices it is necessary to consider the curren that they are called upon to source and sink and the capacitance that they will see. This is done from the data sheets of the devices that they will be called upon to drive.

FACTORS LIMITING LOGIC FAN-OUT

- **Maximum available source current (logic high): Dominant factor for *resistive* loads**
- **Maximum available sink current (logic low): Dominant factor for *TTL* loads**
- **Maximum permitted node capacitance: Dominant factor for *CMOS* loads**
- **Node capacitance has contributions both from the input capacitances of gates on the node and from wiring and PC tracks associated with the node.**

Figure 11.65

Consider a typical fast TTL gate, such as the 74F32 OR-gate. Its maximum input high current ($I_{ih(max)}$) is 20 µA, its maximum input low current ($I_{il(max)}$) is 0.6 mA and its maximum input capacitance is 5 pF. An ADSP-2100 will source 1 mA when its output is high, it will sink 4 mA when its output is low, and it will drive capacitance of up to 100 pF.

The ADSP-2100 will therefore drive the capacitance of 20 74F32 gates, it will drive the input current of fifty such gates in the logic 1 (high) state, but it will sink the input current of only 6.7 (in practical terms, 7) such gates. The lowest of these is evidently the fan-out which it will drive.

In typical systems it is likely that a device will be called upon to drive a mixture of devices, so the calculations will be more complex - but the basic principle will be the same. In most systems involving TTL the fan-out will be limited by the sink current, but in CMOS systems the node capacitance is likely to be the limiting factor. The above calculations do not consider the capacitance of the PC tracks and any cables which the

ADSP-2100 DRIVE CAPABILITY

- **Will drive 100 pF**
- **Will drive 1 mA at logic 1 (≥2.4 V)**
- **Will sink 4 mA at logic 0 (≤ 0.4 V)**

Therefore it will drive:

- **22 74ACT CMOS Gates (= 99 pF, ±22 µA)**
- **10 74LS Schottky TTL Gates (= –4 mA)**
- **7 74F Schottky TTL Gates (= –4.2 mA)**
- **1 Grounded 2.4 k resistor (= 1 mA)**
- **Or any combination of loads which does not exceed a total capacitance of 100 pF, a total drain of 1 mA at logic 1, and a total source of 4 mA at logic 0. (Remember to allow for the capacitance of PCB tracks and wiring.)**

Figure 11.66

device may be called upon to drive but such capacitance can sometimes be a limiting factor, and should always be considered, if only to be eliminated.

Most data converters have less powerful output stages than processors and their fan-out is lower. Additionally the return current of the output drive from a converter will flow in the system analog ground (for reasons discussed earlier in this section) and should therefore be kept as low as possible in order to minimize digital noise in the analog part of the system.

This is best achieved by using CMOS, rather than TTL logic. The DC input currents of CMOS are orders of magnitude lower that those of TTL. However, the input capacitances are comparable so the switching transients are not much lower. It is therefore advisable to buffer ADC outputs with an external buffer to minimize digital output currents from the ADC. Such a buffer will also help to isolate the ADC from digital noise in the rest of the system.

TIMING VARIATIONS

A common cause of malfunctions in digital systems, and particularly in the digital parts of mixed signal systems, is timing error, which often arises from failure to consider the effects of temperature variations on the system.

The specifications of converters, memories and processors all contain such parameters as "set-up" and "hold" times. These are the times, respectively, that data must be present before a clock edge may occur, or that it must remain valid after the edge. At room temperatures many digital circuits are quite

tolerant of operation with set-up and hold times which are shorter that the specified minimum - but at extremes of temperature they may be more demanding.

LOGIC TIMING VARIES WITH TEMPERATURE

- **Specifications such as "set-up" & "hold" (the time a signal must be present before a strobe and the time that it must remain after one, respectively) can vary widely with temperature.**
- **A system designed with room temperature "typical" values may only perform properly at room temperature, if then.**
- **Designers MUST use min/max specifications at temperature extremes to ensure correct operation at all times.**

Figure 11.67

Where a system consists only of digital circuitry it is likely (but not certain) that changes in input and output timing will behave similarly so that systems continue to function over temperature. Where ADCs or DACs are interfacing with digital systems the very different processes used for the converters may result in timing changes not tracking and performance, or even functionality, suffering.

Engineers designing mixed signal systems should always ascertain that the maximum and minimum timing specifications of all the circuits in their systems are compatible over the full temperature range of intended operation. Where there is any doubt buffers or monostables should be used as pulse extenders to ensure that all set-up and hold specifications are complied with.

Sampling clock noise

As has been mentioned elsewhere in this seminar, phase noise on the clock of a sampled data system is indistinguishable from phase noise on the signal itself and it is therefore of critical importance to ensure that the sampling clock has sufficient spectral purity that its phase noise is less than the smallest component to be detected in the signal under analysis.

To achieve this the sampling clock should be isolated as much as possible from the noise present in the digital parts of the system. In particular, buffers used for the sampling clock should, ideally, be on separate chips, with separately decoupled supplies, from the remainder of the digital system, and the sampling clock signal lines should not be sited where they can pick up digital noise from the rest of the system.

SAMPLING CLOCK NOISE

- **Phase noise of the clock must be less than the minimum signal to be detected in the system.**
- **Therefore the sample clock signal must be protected from digital noise.**

BUT

- **Clocks are digital and can corrupt the analog part of the system.**
- **Therefore sampling clock lines must be kept separate from both the analog and the digital parts of the system.**
- **The sampling clock must use an oscillator with low phase noise.**

Figure 11.68

Of course the sampling clock is itself a digital signal. It has as much potential for causing noise in the analog part of the system as any other digital signal. In fact, due to its presence in the converter and SHA sections of a system, it is generally the leading suspect for noise. We therefore see that a sampling clock is very inconvenient as it must be isolated from both the analog and digital parts of the system.

The sampling clock generator must also have adequate spectral purity. RC and other relaxation oscillators just will not do since amplitude noise in whatever circuit functions as a comparator will appear as phase noise on the output signal. LC oscillators have better phase noise, but the lowest noise is obtained with the use of a quartz crystal oscillator. For very high speed clocks a SAW (surface acoustic wave) oscillator is preferable.

A popular design of quartz crystal oscillator uses a resistor, one or more logic gates, a quartz crystal and a couple of capacitors. The design is not popular with engineers who understand quartz crystals or oscillators - such designs have bad phase noise and are liable to overdrive the quartz crystal (not enough to shatter it, as sometimes happened with self-excited crystal-controlled transmitters using vacuum tubes, but enough to affect its long-term stability). The only justi-

ication for the use of such oscillators is in watch and clock circuits where the low voltages involved minimize the overdrive and the phase noise is integrated over long periods and so is unimportant.

OSCILLATORS

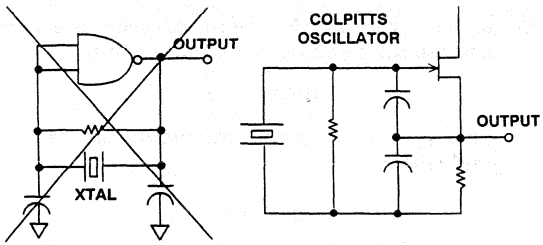

Crystal oscillators built with logic gates have large phase noise.

Low noise crystal oscillators are easily designed with discrete components.

Figure 11.69

Ideally quartz oscillators should use discrete bipolar and FET devices in the circuits recommended by the crystal manufacturers. These circuits are optimized for both crystal drive level and phase noise. The output signal can then be amplified (possibly with a logic gate at this point) to drive the converters.

LOGIC NOISE

One of the most common causes of loss of performance in mixed signal systems is degradation of analog performance by noise from the digital parts of the system. We have already discussed at some length how this digital noise may be isolated from the sensitive analog parts of the system, but it is also worthwhile considering how this noise may be diminished at its source.

It is well-known that TTL is noisy. This is partly because the "totem pole" output stage structure acts as a short-circuit on the supply for a nanosecond or so during switching - giving rise to a large current spike, partly because the current flowing in the input changes, and changes quickly, between logic 0 and logic 1, and partly because the output swing, which takes place in a few nanoseconds, is several volts.

High speed CMOS does not have the change in input current (although there is a capacitance charging current pulse during switching, this is smaller) but may draw a supply current pulse during switching and certainly has a large output swing with a large dV/dT.

4000-Series CMOS is almost 20 years old and slow. It is also widely available, cheap, resistant to RFI, and quite remarkably noise free, since it has low output dV/dT and does not generate a supply current pulse.

LOGIC NOISE

- **TTL has large voltage swings; large, fast I/P & O/P current pulses and asymmetrical circuitry.**
- **HCMOS has large voltage swings; large, fast O/P current pulses and symmetrical circuitry.**
- **4000-Series CMOS is old, slow, cheap and very quiet.**
- **ECL has smaller voltage swings and smaller current surges than TTL & HCMOS even though it is faster.**
- **There is no single ideal logic family.**

Figure 11.70

ECL also draws almost constant current during switching (unless it is driving asymmetrical loads) and has much smaller output voltage swings than TTL or CMOS. Thus, although ECL is faster than TTL and CMOS, it tends to generate less noise.[9]

No single logic family is ideal for all applications (otherwise there would only be one logic family) but it is safe to conclude that TTL should not be used where its noise can corrupt precision analog circuitry but should be replaced by CMOS.

Where only low speeds are necessary 4000 CMOS has overwhelming noise advantages but may not be available in all necessary configurations, and does not interface well with TTL (although it will interface with high speed CMOS families).

In high speed systems where noise is important ECL may offer noise advantages at the interface between the analog and digital parts of the system, even though high speed CMOS is capable of the speeds being used. It is not necessary to use ECL throughout the system - just where its lower noise is advantageous.

Problem Areas

Limitations of SPICE modelling

As we have seen, real electronic circuits contain many "components" which were not present in the circuit diagram but which are there because of the physical properties of conductors, circuit boards, IC packages, etc. These components are difficult, if not impossible, to incorporate into computer modelling software and yet they have substantial effects on circuit performance at high resolutions, or high frequencies, or both.

It is therefore inadvisable to use SPICE modelling or similar software to predict the ultimate performance of such high performance analog circuits. After modelling is complete the performance must be verified by experiment.

This is not to say that SPICE modelling is valueless - far from it. Most modern high performance analog circuits could never have been developed without the aid of SPICE and similar programs, but it must be remembered that such simulations are only as good as the models used and these models are not perfect. We have seen the effects of parasitic components arising from the conductors, insulators and components on the PCB, but it is also necessary to appreciate that the models used within SPICE simulations are not perfect models.

Consider an operational amplifier. It contains some 20-40 transistors, almost as many resistors, and a few capacitors. A complete SPICE model will contain all these components and probably a few of the more important parasitic capacitances and spurious diodes formed by the diffusions in the op-amp chip. This is the model that the designer will have used to evaluate the device during his design. In simulations such a model will behave very like the actual op-amp, but not exactly.

However, this model is not published, as it contains too much information which would be of use to other semiconductor companies who might wish to copy or improve on the design. It would also take far too long for a simulation of a system containing such models of a number of op-amps to reach a useful result. For these, and other, reasons

SPICE MODELLING

- **SPICE modelling is a powerful tool for predicting the performance of analog circuits.**

HOWEVER

- **Models omit real-life effects.**
- **No model can simulate all the parasitic effects of discrete components and a PCB layout.**

THEREFORE

- **Prototypes must be built and proven before production.**

Figure 11.71

the SPICE models of analog circuits published by manufacturers or software companies are "macro" models, which simulate the major features of the component but lack some of the fine detail. Consequently SPICE modelling does not always reproduce the exact performance of a circuit and should always be verified experimentally.

Sockets

It is tempting to mount expensive ICs in sockets rather than soldering them in circuit - especially during circuit development. Engineers would do well not to succumb to this temptation.

USE OF SOCKETS WITH HIGH PERFORMANCE ANALOG CIRCUITS

- **DON'T! (If at all possible)**
- **Use "Pin sockets" of "Cage jacks" such as Amp Part No: 5-330808-3 or 5-330808-6 (Capped & uncapped respectively).**
- **Always test the effect of sockets by comparing system performance with and without the use of sockets.**
- **Do not change the type of socket used without evaluating the effects of the change on performance.**

Figure 11.72

Sockets add resistance, inductance and capacitance to the circuit and may degrade performance to quite unacceptable levels. When this occurs, though, it is always the IC manufacturer who is blamed - not the use of a socket. Even low profile, low insertion force sockets cannot be relied upon not to degrade the performance of high performance (high speed or high precision or, worst of all, both) devices, and as the socket ages and the board suffers vibration the contact resistance of low insertion force sockets is very likely to rise. Where a socket must be used the least loss of performance is achieved by using individual pin sockets (sometimes called "cage jacks") to make up a multi-pin socket in the PCB itself.

It really is best not to use IC sockets with high performance analog and mixed signal circuits. If their use can be avoided it should be. However at medium speeds and medium resolutions the trade-off between performance and convenience may fall on the side of convenience. It is very important, when sockets are used, to evaluate circuit performance with and without the socket chosen to ensure that the type of socket chosen really does have minimal effect on the way that the circuit behaves. The effects of a change of socket on the circuit should be evaluated as carefully as a change of IC would be and the drawings should be prepared so that the change procedures for a socket are as rigorous as for an IC - in order to prevent a purchase clerk who knows nothing of electronics from devastating the system performance in order to save five cents on a socket.

Prototyping high performance analog circuitry

As we have seen, circuit board layout is part of the circuit design of all high performance analog circuits. Prototyping techniques derived from the "node" theory, while ideal for logic breadboarding at low and medium speeds, are quite unsuitable for any analog circuits, or even for very fast digital ones. Vector board and wire wrap prototyping will tell an engineer nothing about the behavior of a properly laid out version of the analog circuit.

The best technique for analog prototyping is to use a prototype of the final PCB - certainly no design is complete until the final PCB layout has been proved to give the required performance. Nevertheless this approach may be a little limiting where a number of different possibilities are to be evaluated, or for a multicard system.

PROTOTYPING MIXED SIGNAL CIRCUITRY

- **NEVER use vector boards or wire-wrap for the analog parts of the system (they can be invaluable for data buses and address lines in the digital part).**
- **Wherever possible avoid the use of sockets for analog ICs.**
- **Use a prototype of your final PCB layout as early as possible.**

Figure 11.73

In this case components should be mounted on a board having a continuous copper ground plane (ideally on both sides of the board, though while convenient this is not essential), with ground connections made to the plane and short point to point wiring made above and below it. The overall component placing and signal routing should be as close as possible to the planned final layout.

As we have already indicated, IC sockets can degrade the performance of analog ICs. While directly soldered components are ideal for prototyping, an IC socket made of pin sockets mounted in the ground plane board may be acceptable (clear the copper, on both sides of the board, for about 0.5 mm around each ungrounded pin socket - solder the grounded ones to ground on both sides of the board).

Allowing wiring to float in the air can be a little tricky. There is a breadboarding system which is conceptually very similar to that described above but which provides adhesive PC pads which stick to the ground plane and allow more rigid component mounting and wiring. This system is manufactured by Wainwright Instruments and is known as "Minimount" in Europe and "Solder Mounts" in the USA. The manufacturer's and distributors' addresses are given in the references at the end of this section.[10]

Manufacturer's evaluation boards are also useful in system prototyping since they have already been optimized for best performance. Analog Devices offers many evaluation boards for a wide array of products. They offer the designer an excellent starting point for the layout.

When the prototype layout is transferred to a CAD system for PCB layout it is important to disable, or at any rate override where necessary, any automatic routing or component placing software. The criteria used by such software are more closely related to "node" theory and aesthetically pleasing rows of components (which, admittedly, are also easier on automatic component placing machinery) than to optimizing stray inductance and capacitance and minimizing common ground impedances.

ADDITIONAL PROTOTYPING HINTS

- **Pay*equal* attention to signal routing, componen placing and supply decoupling in *both* the prototype and the final design.**
- **Verify performance as well as functionality at each stage of the design.**
- **For "freehand" prototyping use a copper-clad board, mount components to it by their ground pins and wire the remaining connections point-to-point (use Wainwright Instruments' Minimount/Solder Mountadhesive PC pads if aerial point-to-point wiring seems too fraught with peril).**

Figure 11.74

REFERENCES

1. Albert Einstein
2. "E.S.D. Prevention Manual" Available free from Analog Devices.
3. For a discussion of these effects see pp I-11 et seq. "Linear Design Seminar" notes. Analog Devices, October 1987.
4. "Electricity & Magnetism" by B.I. & B. Bleaney. OUP 1957, pp 23,24, & 52.
5. "An I.C. Amplifier User's Guide to Decoupling, Grounding and Making Things Go Right for a Change" by Paul Brokaw. Available free of charge from Analog Devices.
6. "Avoiding Ground Problems in High Speed Circuits", Jeff Barrow "R.F. Design" July 1989. AND "Grounding for Low- and High-Frequency Circuits", Paul Brokaw & Jeff Barrow, "Analog Dialogue" 23-3 1989. Free from Analog Devices.
7. International EMI Emission Regulations

Canada	CSA C108.8-M1983	FDR	VDE 0871/VDE 0875
Japan	CISPR (VCCI)/PUB 22	USA	FCC-15 Part J

8. "Design & Layout of a Video Graphics System for Reduced EMI" Bill Slattery & John Wynne (E1309-15-10/89). Free from Analog Devices.
9. "MECL System Design Handbook" 4th Edition, 1983 - Motorola
10. Wainwright Instruments Inc.
P.O.Box 2203, Capistrano Beach, CA 92624. Tel: 714-248-2490

Wainwright Instruments GmbH, Widdersberger Strasse 14,
DW-8138 Andechs-Frieding, Germany. Tel: +49-8152-2245

ANALOG DEVICES

AN-282
APPLICATION NOTE

ONE TECHNOLOGY WAY • P.O. BOX 9106 • NORWOOD, MASSACHUSETTS 02062-9106 • 617/329-4700

Fundamentals of Sampled Data Systems

A Typical DSP Sampled Data System

A block diagram of a typical sampled data DSP system is shown in Figure 3.1. Prior to the actual analog-to-digital conversion, the analog signal usually passes through some sort of signal conditioning circuitry which performs such functions as amplification, attenuation, or filtering. If the analog signal originates as a temperature, pressure, flow-rate, or force, then an appropriate sensor and transducer is required to first convert the physical quantity into an electrical voltage or current.

There are two key concepts involved in the actual analog-to digital conversion process: *discrete time sampling* and *finite amplitude resolution due to quantization.* An understanding of these concepts is vital to DSP applications.

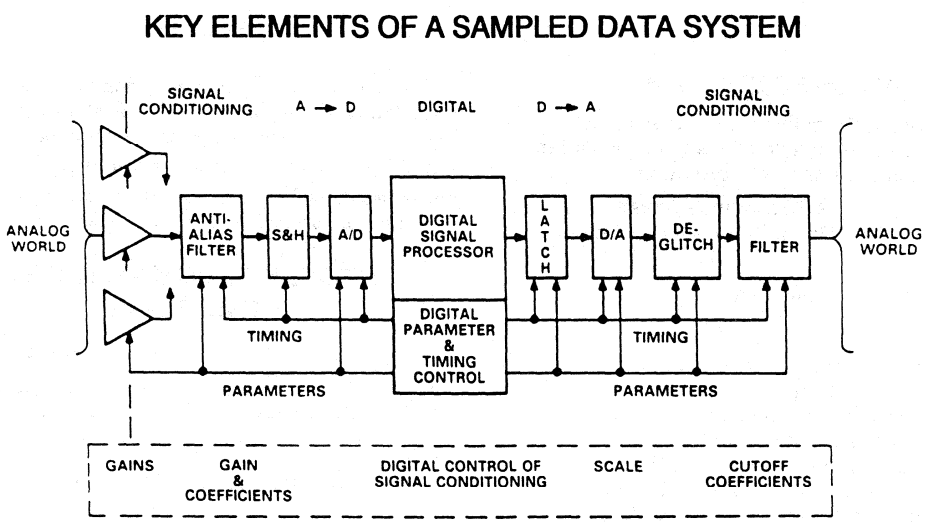

Figure 3.1

Discrete Time Sampling of Analog Signals

The concept of discrete time and amplitude sampling of an analog signal is shown in Figure 3.2. The continuous analog data must be sampled at discrete intervals, t_s, which must be carefully chosen to insure an accurate representation of the original analog signal. It is clear that the more samples taken (faster sampling rates), the more accurate the digital representation, but if fewer samples are taken (lower sampling rates), a point is reached where critical information about the signal is actually lost. This leads us to the statement of Nyquist's criteria given in Figure 3.3.

DISCRETE SAMPLING OF AN ANALOG SIGNAL

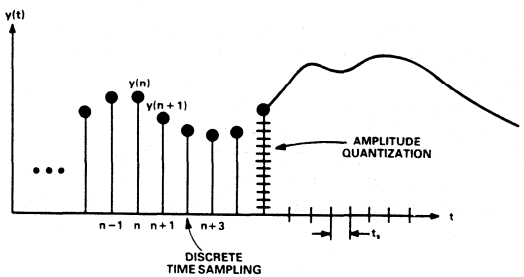

Figure 3.2

NYQUIST'S CRITERIA

- **An Analog Signal with a *Bandwidth* of f_a Must be Sampled at a Rate $f_s > 2f_a$ in Order to Avoid the Loss of Information**
- **If $f_s < 2f_a$, then a Phenomena Called Aliasing Will Occur in the Analog Signal Bandwidth**

Figure 3.3

In order to understand the implications of *aliasing* in both the time and frequency domain, first consider the four cases of a time domain representation of a sampled sinewave signal shown in Figure 3.4. In the Case 1, it is clear that an adequate number of samples have been taken to preserve the information about the sinewave. In Case 2 of the figure, only four samples per cycle are taken; still an adequate number to preserve the information. Case 3 represents the ambiguous limiting condition where $f_s = 2f_a$. If the relationship between the sampling points and the sinewave were such that the sinewave was being sampled at precisely the zero crossings (rather than at the peaks, as shown in the illustration), then all information regarding the sinewave would be lost. Case 4 of Figure 3.4 represents the situation where $f_s < 2f_a$, and the information obtained from the samples indicates a sinewave having a frequency which is lower than $f_s/2$, i.e. the out-of-band signal is *aliased* into the Nyquist bandwidth between dc and $f_s/2$. As the sampling rate is further decreased, and the analog input frequency f_a approaches the sampling frequency f_s, the aliased signal approaches dc in the frequency spectrum.

TIME DOMAIN EFFECTS OF ALIASING

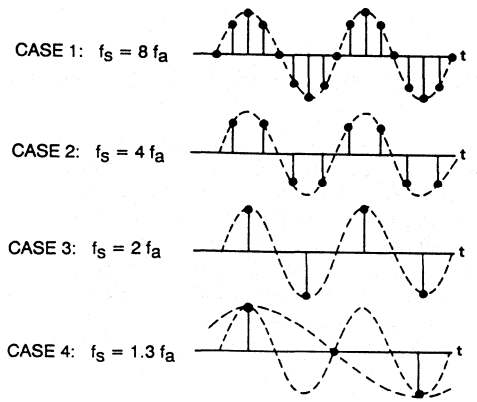

Figure 3.4

The corresponding frequency domain representation of the above scenario is shown in Figure 3.5. Note that sampling the analog signal f_a at a sampling rate f_s actually produces two alias frequency components, one at f_s+f_a, and the other at f_s-f_a. The upper alias,

FREQUENCY DOMAIN EFFECTS OF ALIASING

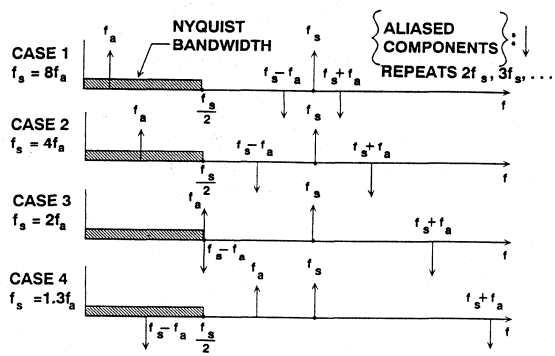

Figure 3.5

f_s+f_a, seldom presents a problem, since it lies outside the Nyquist bandwidth. It is the lower alias component, f_s-f_a, which causes problems when the input signal exceeds the Nyquist bandwidth, $f_s/2$.

It is clear from the above discussion that the ADC must be preceded by an anti-aliasing filter which has sufficient stopband attenuation at $f_s/2$ and above to prevent unwanted in-band aliasing. Aliasing may also occur from harmonics of the fundamental signal which fall outside the Nyquist bandwidth, or from unfiltered broadband noise at the ADC input.

The effects of aliasing on the dynamic range of a sampled data system are shown in Figure 3.6. The top part of the figure illustrates the desired condition at the Nyquist point, where the aliased component intersects the input signal at a point below the desired dynamic range. The lower part of the figure shows the condition where the upper-frequency dynamic range is limited by the aliased components. This condition will result in a reduction in overall signal-to-noise ratio at the higher frequencies, and could result in the distortion due to aliased out-of-band tones or harmonics as shown in Figure 3.7.

FREQUENCY DOMAIN EFFECTS OF ALIASING ON DYNAMIC RANGE

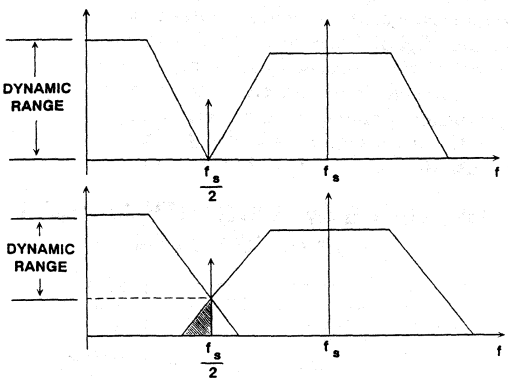

Figure 3.6

UNWANTED TONES DUE TO ALIASING

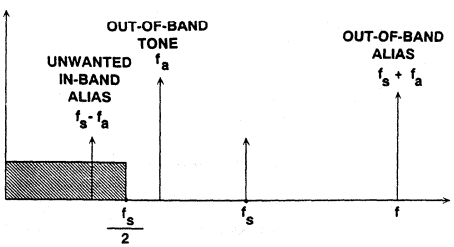

Figure 3.7

SELECTION OF ANTIALIASING FILTERS

It should be clear by now, that for a given analog input bandwidth, f_a, the requirements of the antialiasing filter are related not only to the sampling rate, f_s, but also to the desired system dynamic range. Simply stated, *dynamic range* is the ratio of the largest expected signal to the smallest signal which must be resolved, and is usually expressed in dB. At this point, we are concerned with dynamic range limitations due to aliasing. The limiting effects of ADC quantization noise and other non-linearities will be discussed shortly. The following rules of thumb will result in a filter which is somewhat overspecified, but the concepts are valid and can be refined to fit the actual system requirements.

ANTIALIASING FILTER REQUIREMENTS

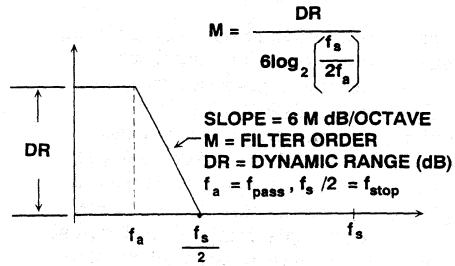

Figure 3.8

First, set the corner frequency of the antialiasing filter equal to the desired analog input bandwidth, f_a. This defines the passband of the filter, $f_{pass}=f_a$. Define the beginning of the filter's stopband, $f_{stop}=f_s/2$. Let the filter stopband attenuation be the desired upper-frequency dynamic range, DR, expressed in dB. These parameters define the transition band characteristics of the filter, i.e., it must achieve a stopband attenuation equal to the dynamic range over $\log_2(f_{stop}/f_{pass})$ octaves. The approximate order of the filter, M, (the number of poles) required to achieve this transition band slope can then be determined, since the filter rolloff is approximately 6M dB per octave. A simple example calculation is shown in Figure 3.9, where the signal bandwidth, f_a, is 3kHz, the sampling rate, f_s, is 12kHz, and a dynamic range of 60dB is required. This implies that a 10 pole filter is needed. Remember that in practice, any analog filter with more than 8 poles becomes a real design challenge, and a filter with more than 12 poles becomes almost an impossibility except

ANTIALIASING FILTER EXAMPLE

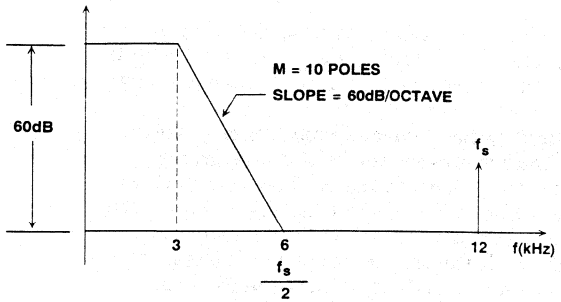

Figure 3.9

for the experienced filter designer. These considerations so far have neglected the filter's phase characteristics, and also the in-band and out-of-band ripple requirements. The addition of these parameters can make antialiasing filter design a truly formidable task.

The above rules-of-thumb for determining the complexity of the antialiasing filter assume that fullscale signals can occur at essentially all input frequencies above Nyquist. In actual practice, this is not usually the case, and there is some natural attenuation of the signal being processed at the higher input frequencies. For instance, in the previous example, if signals at Nyquist and above were already attenuated by 12dB, then a filter stopband attenuation of only 48dB would be required at the Nyquist frequency of 6kHz. This would imply that only an 8 pole filter would be needed. This situation is illustrated in Figure 3.10.

EFFECTS OF OUT-OF-BAND ATTENUATION ON ANTIALIASING FILTER

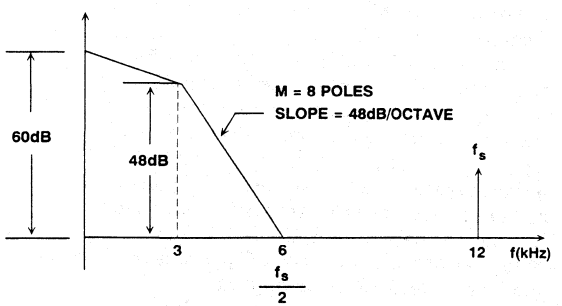

Figure 3.10

From the above discussions, it is clear that the requirements on the antialiasing filter can be relaxed at the expense of higher sampling rates (called *oversampling*). Later in the seminar, we will see that a particular class of ADCs and DACs , called *Sigma-Delta* ($\Sigma\Delta$) are inherently oversampling converters and greatly reduce the complexity of the antialiasing filter.

OVERSAMPLING AND DECIMATION

As previously discussed, the major advantage of oversampling the input signal is the resulting simplification in the antialiasing filter requirements. Of course the downside of oversampling is that it also increases the ADC output data rate, and the DSP must be able to keep up in order to maintain real-time operation. If the data is to be transmitted in serial form, then it consequently will occupy more of the frequency spectrum. An attractive alternative makes use of both analog *and* digital filtering techniques, oversampling, and a process called *decimation*. Figure 3.11 shows the traditional case, where all the antialiasing burden lies with the analog input filter preceding the ADC. In Figure 3.12, however, the oversampling ratio, K (K is an integer), relaxes the rolloff requirement of the input analog filter by increasing the Nyquist frequency to $Kf_s/2$. The digital filter following the ADC (digital filtering will be discussed at length in Section VII) implements the antialiasing function with respect to f_s, and has sufficient stopband attenuation at $f_s/2$ to achieve the desired dynamic range. As we will see later in the seminar, digital filters having sharp cutoff characteristics with good phase response are much more easily implemented than their corresponding analog counterparts (assuming sufficient speed in the DSP). Finite Impulse Response (FIR) filters can be de-

NYQUIST SAMPLING WITH ANALOG LOWPASS FILTER

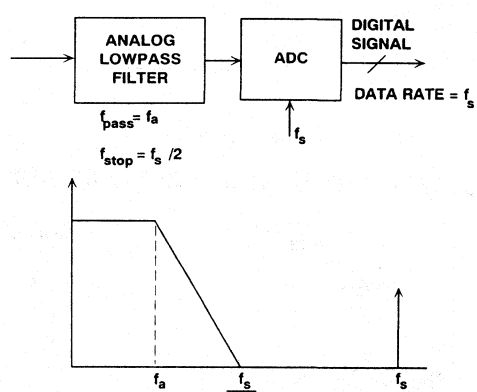

Figure 3.11

OVERSAMPLING WITH ANALOG AND DIGITAL FILTERING

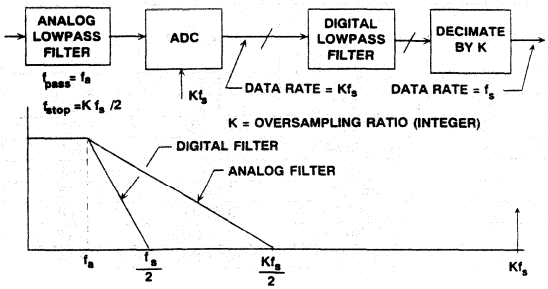

Figure 3.12

signed which have linear phase characteristics. Since the bandwidth has been reduced to $f_s/2$ by the digital antialiasing filter, the data coming out of the digital filter actually contains redundant information, and there is no need to look at every sample. In fact, it is only necessary to look at every Kth sample. This process is called *decimation*, and will be discussed in much more detail in the section of the seminar on Sigma-Delta converters (Section VI). In addition, the actual decimation can be performed by the FIR filter itself by computing a single output sample for every K input samples. This concept of oversampling and decimation is one of the most powerful concepts in real-world DSP.

UNDERSAMPLING AND ITS APPLICATIONS

In this section we will see that there are some applications in DSP where aliasing is perfectly acceptable and can be used advantageously. When the analog signal being digitized by the ADC exceeds $f_s/2$, the condition is often referred to as *super-Nyquist*, or *undersampling*. Nyquist's criteria states that the *bandwidth* (not the actual frequency) of the signal being digitized should not exceed $f_s/2$ for information to be preserved. As an example, consider a telecommunications transmultiplexer application where Frequency Division Multiplexed (FDM) data occupying the bandwidth of 60 to 108kHz is sampled at a frequency of 112kHz. Figure 3.13 shows the spectrum of the signal and the location of the aliased components. At the receiving end of the system, the filter which follows the reconstruction DAC is a bandpass rather than a lowpass and must filter out the aliased components falling between 4kHz and 52kHz as well as the component located at the sampling frequency of 112kHz.

SUPER-NYQUIST SAMPLING OF FDM SIGNAL

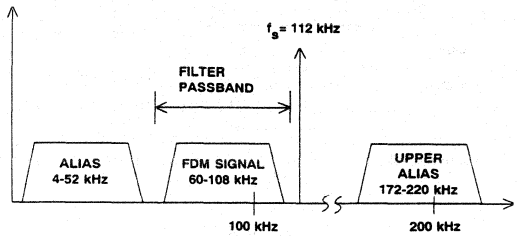

Figure 3.13

Another application for super-Nyquist operation is in the direct conversion of IF signals to baseband. Most traditional communication and radar receivers employing ADCs and DSP utilize a system in which the intermediate frequency (IF) from the front end of the receiver is down-converted or demodulated to a baseband signal by a mixer and a lowpass filter as shown in Figure 3.14. This final IF stage uses a local oscillator which is phase coherent with the signal carrier frequency. The mixer output contains a baseband signal which is proportional to the phase difference between the two inputs. Following the mixer is a lowpass filter, amplifier, and an ADC. Typical mixers have a conversion loss ranging from 4 to 6dB. In cases when the signal-to-noise ratio is limited by the front end, elimination of the mixer will improve the overall noise figure of the receiver.

ANALOG DOWNCONVERSION OR DEMODULATION

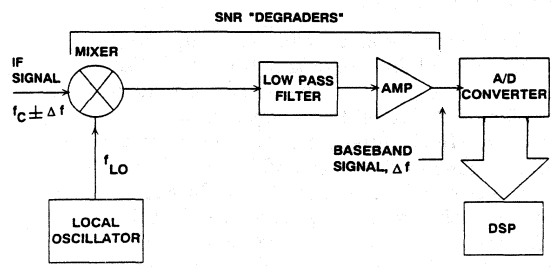

Figure 3.14

This can be accomplished (as shown in Figure 3.15) if the IF frequency is sampled at a rate equal to the local oscillator frequency. The ADC now functions as a demodulator. If the ADC samples an analog signal of the same frequency as the sampling frequency, the digitized output is a dc value. Any deviation in the analog signal from the sampling

DIRECT IF TO DIGITAL DOWN CONVERSION OR DEMODULATION

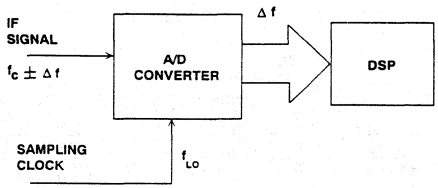

Figure 3.15

frequency looks like a *beat* frequency, Df, and the demodulation process is thereby achieved.

The data from the ADC must be processed by the DSP using an FFT which computes both the real and imaginary components of the digitized signal. This is necessary in order to preserve the phase information contained in the demodulated signal.

Operation of ADCs in a super-Nyquist environment obviously requires that the dynamic performance of the converter be known for input frequencies *above* Nyquist. The signal-to-noise ratio and harmonic distortion performance of an ADC typically degrades at higher input frequencies, so ac performance for the input frequency desired must be adequate to meet system requirements. Super-Nyquist operation typically requires an ADC which is more robust to high frequency input signals than an ADC which is specified for strictly sub-Nyquist applications.

Effects of Finite Amplitude Resolution Due to Quantization

The second major effect to be considered in a sampled data system is that of the finite amplitude resolution caused by the analog-to-digital or digital-to-analog conversion process. In this discussion, we will refer to the number of bits of the ADC (or DAC) as the converter's *resolution,* N. In the case of an ADC, the input range is divided into 2^N discrete levels, each represented by an N-bit binary word. For a DAC, the input consists of an N-bit binary word, and there are 2^N possible discrete output levels. Figure 3.16 shows the number of bits, N, the corresponding number of levels, 2^N, and the weight of the *least significant bit* (LSB) expressed as a percentage and a ratio in dB [$20 \log_{10}(2^N)$], or 6.02N dB. This ratio (whether expressed as a percentage or in dB) represents the *dynamic range* of the converter, i.e., the ratio of the largest resolvable signal to the smallest

RESOLUTION AND DYNAMIC RANGE OF ADCs AND DACs

#BITS, N	#LEVELS, 2^N	%, $100/2^N$	dB, 6N
8	256	0.4	48
10	1024	0.1	60
12	4096	0.025	72
14	16384	0.006	84
16	65536	0.0015	96
18	262144	0.0004	108
20	1048576	0.0001	120
22	4194304	0.000025	132
24	16777216	0.000006	144

Figure 3.16

resolvable signal. At this point, we should point out that the dynamic range values in Figure 3.16 represent *ideal* ADCs and DACs does not consider such ac performance limitations such as harmonic and intermodulation distortion. Neither do these values represent the theoretical signal-to-quantization noise. These topics will be discussed shortly.

Quantization Theory, Signal to Noise Ratio, and Effective Bits

The finite resolution of ADCs and DACs gives rise to a theoretical limitation to the signal-to-noise ratio (SNR) which is a function of the number of bits, N. In order to make a meaningful measurement, the ADC is stimulated with a fullscale sinewave input which is slightly below the clipping range of the converter. This gives rise to a sample-to-sample error which produces *quantization noise.* It can be shown mathematically that the rms noise voltage produced by quantization measured *within the Nyquist bandwidth* is given by the familiar expression $q/\sqrt{12}$, where q is the weight of the least significant bit (LSB) of the converter. The value for the LSB,q, can be calculated by dividing the fullscale range of the ADC or DAC by 2^N. In an ideal converter with no error sources, the theoretical rms quantization noise voltage is also independent of both the input signal amplitude and frequency. The derivation for this simple expression is given in the following reference:

W.R. Bennett, *Spectra of Quantized Signals*, BSTJ 27, pp. 446-472, July 1948

For a fullscale sinewave input, it can further be shown that the theoretical rms signal to quantization noise ratio is given by SNR = 6.02N + 1.76dB.

QUANTIZATION THEORY BASICS

- **RMS Quantization Noise in Nyquist Bandwidth, $f_s/2$:**

  $q/\sqrt{12}$

- **Fullscale Sinewave RMS Signal to RMS noise ratio in Nyquist Bandwidth:**

  $SNR = 6.02N + 1.76dB$

- **Effective Number of Bits (ENOB):**

$$ENOB = \frac{SNR_{ACTUAL} - 1.76dB}{6.02}$$

Figure 3.17

It should be noted that the rms quantization noise generally approximates broadband noise across the Nyquist bandwidth. There are certain conditions, however, where this is not true. If there is correlation between the quantization error signal and the signal being digitized, then the quantization noise may be concentrated at harmonics of the input signal rather than being spread uniformly across the bandwidth. This is most likely to occur if the input signal is a sinewave which is a subharmonic of the sampling frequency.

In testing ADCs, the SNR is usually calculated using DSP techniques while applying a pure sinewave signal to the input of the ADC as shown in Figure 3.18. The Fast Fourier Transform (FFT) processes a finite number of time samples and converts them into the frequency spectrum such as that shown in Figure 3.19 for the AD678 12-bit 200kSPS sampling ADC. The frequency spectrum is then used to calculate the SNR as well as harmonics of the fundamental input signal, very similar to an analog spectrum analyzer. The rms value of the signal is first computed. Then the rms value of all other frequency components over the Nyquist

ADC DYNAMIC TESTING

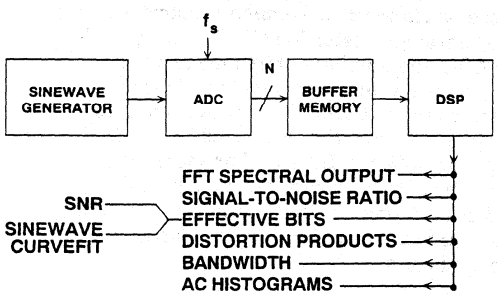

Figure 3.18

2048 POINT FFT OUTPUT FOR AD678 12-BIT, 200 kSPS ADC

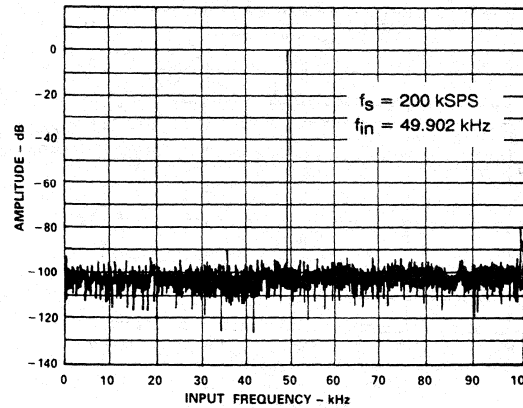

Figure 3.19

bandwidth (this includes not only noise but also distortion products) is computed. The ratio of these two quantities, expressed in dB is the SNR. Various error sources in the ADC cause the measured SNR to be less than the theoretical value, 6.02N + 1.67dB. These errors occur due to integral and differential nonlinearities, missing codes, and internal ADC noise sources. In addition, the errors typically are a function of input slewrate and therefore increase as the input frequency gets higher. In calculating the rms value of the noise, it is customary to include harmonics of the fundamental signal. This is sometimes referred to as the signal-to-noise-plus-distortion, S/(N+D), but is usually called simply SNR. A typical plot of S/(N+D) for the AD678 sampling ADC (12 bit, 200kSPS) is shown in Figure 3.20.

Another way to interpret SNR is in terms of *effective number of bits*, or *ENOBs*. The effective-bit calculation is performed by solving the SNR equation for N, given the measured value of SNR. (See Figure 3.17). For instance, a perfect 12 bit ADC would have a theoretical SNR of 74dB, corresponding to 12 effective bits. A measured SNR of 68dB, however, would correspond to 11 effective bits. This says that the performance of the actual 12 bit ADC is equivalent to that of a perfect 11 bit ADC. Figure 3.20 also shows the ENOB performance of the AD678 on the same graph as the SNR. Note that at low frequencies, the AD678 exceeds 11.4 effective bits.

Effective bits can also be measured using the *sinewave curvefit* method. In this method, a sinewave is applied to the ADC, and a number of samples are collected.

S / (N + D) AND EFFECTIVE BITS FOR AD678 12-BIT, 200 kSPS ADC

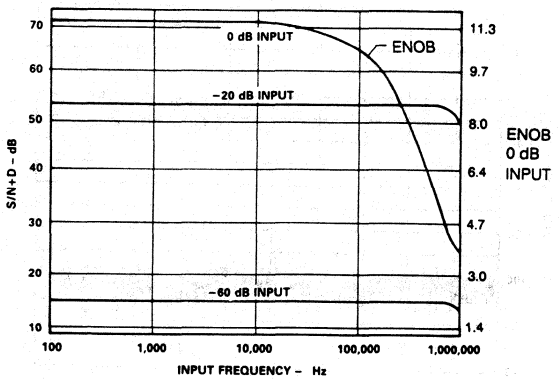

Figure 3.20

Instead of performing an FFT on the time samples, the *best-fit* sinewave to fit the data points is calculated. The sinewave amplitude, offset, frequency, and phase are chosen to minimize the rms error between the actual sinewave data points and the theoretical sinewave. Again, the theoretical rms error for a perfect ADC is $q/\sqrt{12}$. The rms error between the actual sinewave and the theoretical sinewave is computed, and the effective bits are calculated using the formula shown in Figure 3.21. The ENOB measurement using the sinewave curvefit method correlates well with that obtained using the SNR technique. If the SNR calculation is made with a signal which is less than fullscale, then a correction factor must be added as shown in order for the two methods to correlate.

CALCULATING ENOB USING SINEWAVE CURVE FITTING

- Q_A = Actual RMS Error from Fit Sinewave
- Q_T = Theoretical N-Bit RMS Error from Best Fit Sinewave

  $= q/\sqrt{12}$
- $ENOB = N - \log_2\left[\frac{Q_A}{Q_T}\right]$ Correlates to:
- $SNR = \frac{SNR_{ACTUAL} - 1.76dB + \text{Level of Signal Below FS}}{6.02}$

Figure 3.21

SELECTION OF ADC RESOLUTION BASED ON SIGNAL DYNAMIC RANGE

Selection of the proper ADC for a given application involves much more than just determining the number of bits required and the sampling rate. The dc and ac characteristics of the ADC must be examined with respect to the analog signal being processed and a proper match must be found. Inevitably, this process involves certain tradeoffs in performance and cost.

DSP APPLICATIONS AND DYNAMIC RANGE REQUIREMENTS

APPLICATION	SIGNAL BANDWIDTH	DYNAMIC RANGE	ADC # BITS
Seismology	10Hz	146dB	24
Digital Audio	20kHz	100dB	18
Echo Cancelling	4kHz	84dB	14
Speech Processing	4kHz	74dB	12
V.32 Modems	4kHz	74db	14
Ultrasound	15MHz	60dB	10
Radar	5MHz	74dB	12
Broadband Receivers	5MHz	86dB	14

Figure 3.22

Figure 3.22 shows a number of applications which are suitable for DSP processing. The approximate bandwidth and dynamic range of the corresponding signal is given. There are actually two aspects to dynamic range: *dc* and *ac*. The dynamic range corresponds to the values given in Figure 3.22 (neglecting ADC static errors). AC dynamic range, on the other hand, is related to the harmonic distortion performance of the ADC. For instance, in a digital spectral analysis application, the harmonics of a fullscale sinewave input signal limits the system's ability to resolve small signals in the presence of large signals. AC linearity is usually expressed in terms of harmonic distortion, or total harmonic distortion (THD). In a practical ADC, the number of bits may not be a good indicator of the harmonic distortion performance of the converter. AC dynamic range is less than that predicted by the SNR equation, 6.02N + 1.76dB. For these reasons, the data sheet must include both dc and ac performance specifications in order for the user to make an intelligent selection for the application.

ADC Static Transfer Characteristics

The basic specifications which describe the static performance of an ADC are given in Figure 3.23.

ADC STATIC PERFORMANCE SPECIFICATIONS

- Differential Non-Linearity (DNL)
- Integral Non-Linearity (INL)
- Missing Codes
- Gain Error
- Offset Error

Figure 3.23

In the ideal transfer function for a 3 bit ADC (Figure 3.24), the analog input signal is on the horizontal axis and the digital output is on the vertical axis. The digital output of the ADC is valid over a range of input signal. The quantum of input for a given output code is called the *width* of the code. The ideal width is exactly 1 LSB (least significant bit), but, in practice, each code-width is different from its neighbors. The deviation in the code-widths from the ideal 1 LSB value is called differential non-linearity, or DNL. A 3-bit ADC with various errors is shown in Figure 3.25. Note that the code 100 is missing because of the large DNL associated with the adjacent codes. Missing codes can produce oscillation and hunting in a closed-loop system, thus making this an important parameter to consider for ADC selection in this application.

Integral non-linearity, or INL, is usually measured with respect to the code centers. A

TRANSFER FUNCTION FOR IDEAL 3-BIT ADC

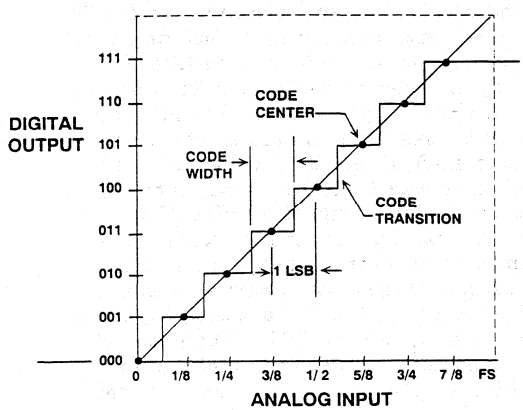

Figure 3.24

TRANSFER FUNCTION FOR NON-IDEAL 3-BIT ADC

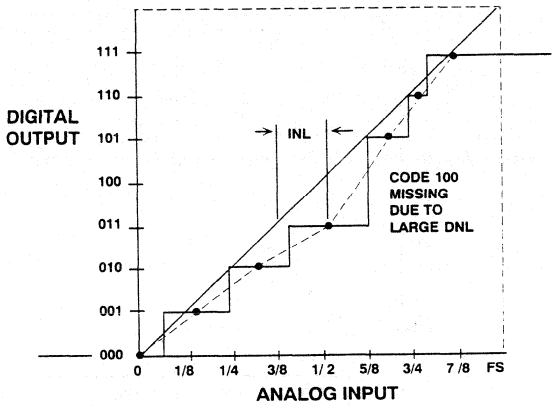

Figure 3.25

straight line is drawn through the end-points, and the worst deviation of any code center from this ideal straight line is the INL as shown in Figure 3.25. In some cases, integral non-linearity is defined with respect to a *best-fit* straight line which is typically calculated using the least-squares method.

Gain and *offset* errors apply to all codes equally and are usually trimmed out in a system using fairly traditional techniques.

DAC Static Transfer Characteristics

The basic specifications which describe the static transfer characteristics of a DAC are given in Figure 3.26.

DAC STATIC PERFORMANCE SPECIFICATIONS

- Differential Non-Linearity (DNL)
- Integral Non-Linearity (INL)
- Non-Monotonicity
- Gain Error
- Offset Error

Figure 3.26

The static transfer function for an ideal 3-bit DAC is shown in Figure 3.27. The digital input values are plotted on the horizontal axis and the corresponding analog output values on the vertical. Unlike an ADC, a DAC cannot have a missing code. There will be a discrete analog output voltage produced for each digital input code. Differential non-linearity is defined as the variation in the spacing between adjacent analog output

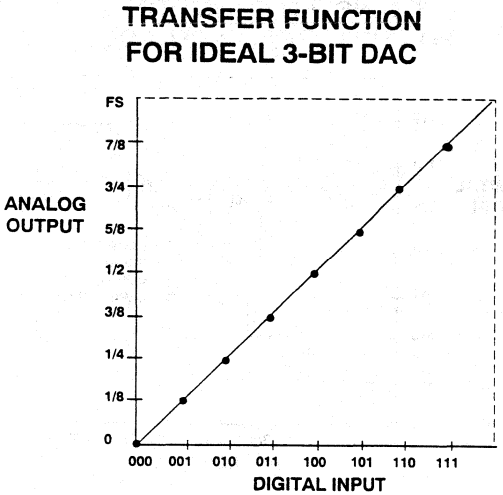

Figure 3.27

values from the ideal 1 LSB value. Excessive DNL errors can result it non-monotonic conditions as shown in Figure 3.28. A DAC is said to be non-monotonic if an increase in the digital code input causes a decrease in the analog output value. Conversely, a DAC is said to be monotonic if the slope of its transfer characteristic has the same sign over its entire range. Non-monotonic conditions can produce oscillations in a closed loop system; therefore, this specification is important in the selection of a DAC for such applications.

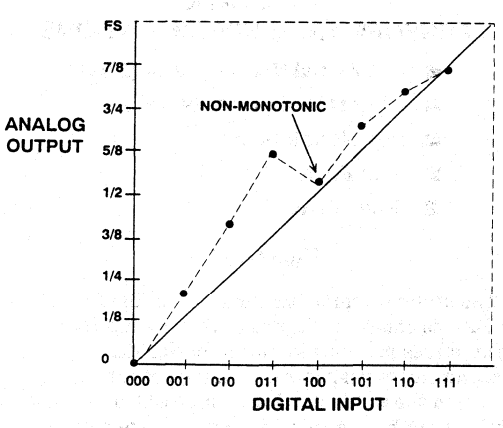

Figure 3.28

Integral non-linearity is defined as the worst case variation in any of the analog output values with respect to an ideal straight line drawn through the end points. As with an ADC, INL may also be defined with respect to a best-fit straight line.

Gain and offset definitions are similar to those for ADCs and affect each analog output value equally.

ADC Dynamic Performance

In order to be useful in most DSP applications, the ADC must have acceptable dc *and* ac performance characteristics. A listing of the most important dynamic ADC characteristics is given in Figure 3.29.

ADC DYNAMIC SPECIFICATIONS

- **Signal-to-Noise Plus Distortion (S/N + D) Ratio and Effective Number of Bits**
- **Peak Spurious, Peak Harmonic Content, and Spurious Free Dynamic Range (SFDR)**
- **Total Harmonic Distortion (THD)**
- **Full-Power Bandwidth (FPBW)**
- **Full-Linear Bandwidth**
- **Intermodulation Distortion (IMD)**
- **Aperture Delay Time and Aperture Jitter**
- **Transient Response**
- **Overvoltage Recovery**

Figure 3.29

As we will see in a later section, there are a number of architectures which are suitable for DSP ADC designs, and most require a sample-and-hold amplifier (SHA) ahead of the actual converter as shown in Figure 3.30. Notable exceptions are flash converters and, in particular, sigma-delta converters. To fully characterize the dynamic performance of a SHA-ADC pair, they must be integrated onto the same chip, or at least offered as a complete functional unit. Otherwise, it is almost impossible to determine the overall dynamic performance of the SHA-ADC combination from the specifications on the individual devices. The requirement for complete dc and ac characterization of ADCs has led to the introduction of *sampling* ADCs which have on-board SHAs. These converters eliminate the problems of interfacing SHAs to ADCs and provide users with complete dc and ac specifications.

ADC WITH TRACK-AND-HOLD

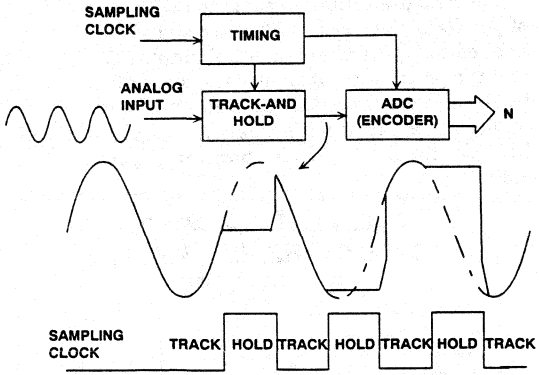

Figure 3.30

Signal-to-Noise Ratio and Effective Bits

As has been previously discussed, the signal-to-noise ratio specification is probably the most all-inclusive ac specification used in the industry today. Since it is common practice to include the effects of harmonic distortion in this measurement, S/N+D is defined as the ratio of the rms value of the measured input signal to the rms sum of all other spectral components below the Nyquist frequency, including harmonics, but excluding dc. A typical plot of S/N+D is shown in Figure 3.31 for three high speed flash ADCs. The harmonic distortion performance of the AD9617 current feedback op amp is shown on the same graph for comparison. The SNR measurement can also be expressed in effective bits, or ENOBs, as is also shown in Figure 3.31.

FLASH ADC AND OP AMP DYNAMIC PERFORMANCE

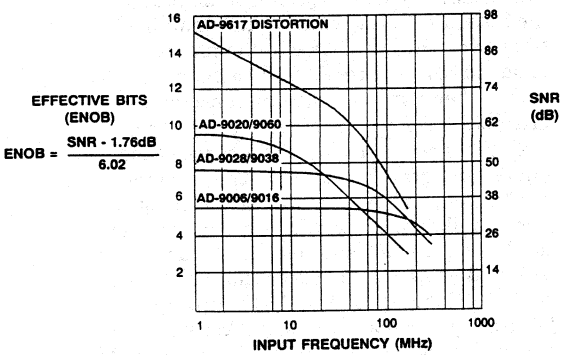

FLASH ADC	RESOLUTION	SAMPLING RATE
AD9020	10 BITS	40 MSPS
AD9060	10 BITS	60 MSPS
AD9028/9038	8 BITS	250 MSPS
AD9006/9016	6 BITS	400 MSPS

Figure 3.31

Peak Spurious, Peak Harmonic Content, and Spurious Free Dynamic Range (SFDR)

The peak spurious or peak harmonic component is the largest spectral component excluding the input signal and dc. This value is expressed in dB relative to the rms value of a fullscale input signal. The peak spurious specification is also occasionally referred to a spurious free dynamic range (SFDR). A typical plot showing the peak spurious performance for the AD678 is shown in Figure 3.32.

PEAK SPURIOUS RESPONSE FOR AD678 AT 200 kSPS, NONAVERAGED 2048 POINT FFT

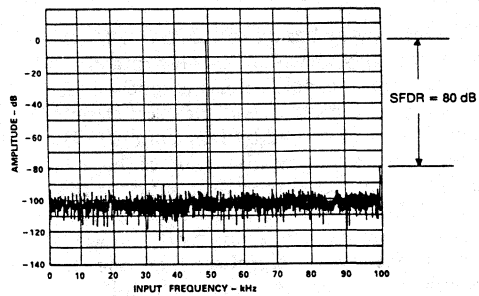

Figure 3.32

Total Harmonic Distortion (THD)

Total harmonic distortion (THD) is the ratio of the rms sum of the first six harmonic components to the rms value of a fullscale input signal and is expressed in a percentage or in dB. For input signals or harmonics that are above the Nyquist frequency, the aliased component is used. Typical THD performance for the AD678 is shown in Figure 3.33.

TOTAL HARMONIC DISTORTION, FULL-POWER BANDWIDTH, AND FULL LINEAR BANDWIDTH FOR AD678

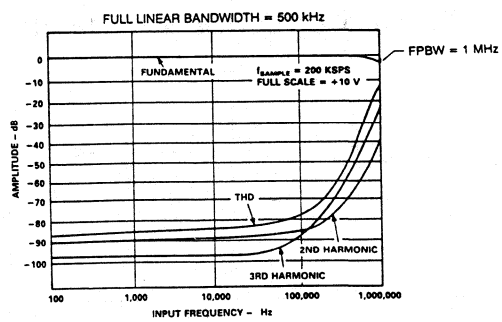

Figure 3.33

Full-Power Bandwidth

The full-power bandwidth (FPBW) of an ADC is that input frequency at which the amplitude of the reconstructed (using FFTs) *fundamental* is reduced by 3dB for a fullscale input. As can be seen from Figure 3.33, the full-power bandwidth of the AD678 is approximately 1MHz. In order to be meaningful, however, FPBW must be examined in conjunction with SNR, ENOB, and harmonic distortion in order to determine the true dynamic performance of the ADC at the FPBW frequency.

Full-Linear Bandwidth

The full-linear bandwidth of an ADC is the input frequency at which the slewrate of the input sample-and-hold (SHA) is reached. At this point, the amplitude of the reconstructed sinewave has degraded by less than -0.1dB. Beyond this frequency, distortion of the sampled input signal increases significantly. The AD678 ADC has been designed to optimize input bandwidth, allowing it to undersample input signals significantly above the converter's Nyquist frequency. The full-linear bandwidth specification is 500kHz for the AD678 and is also shown in Figure 3.33.

Intermodulation Distortion (IMD)

Intermodulation distortion (IMD) occurs when the inputs consist of sinewaves at two frequencies, F1 and F2. Any device with nonlinearities will create distortion products, of the order (m+n), at sum and difference frequencies of mF1± nF2, where m, n = 0,1,2,3 ... Intermodulation terms are those for which m or n is not equal to zero. For example, the second order terms are (F1 + F2) and (F1 - F2), and the third order terms are (2F1 + F2), (2F1 - F2), (F1 + 2F2), and (F1 - 2F2) (see Figure 3.34). The IMD products

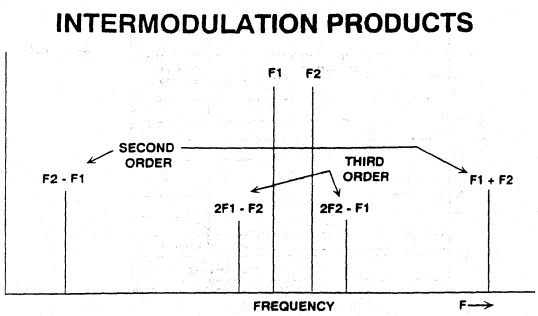

Figure 3.34

are expressed as the dB ratio of the rms sum of the measured input signals to the rms sum of the distortion terms. The two signals applied to the ADC are of equal amplitude and the peak value of their sum is -0.5dB from fullscale. The IMD products are normalized to a 0dB input signal. A typical IMD FFT plot for the AD678 is shown in Figure 3.35.

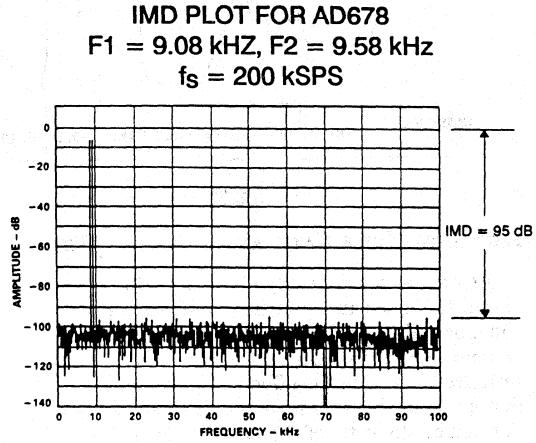

Figure 3.35

AC Linearity Plots Using Histograms

For this measurement, a fullscale sinewave is applied to the ADC, and a large number of samples are taken. The number of occurrences of each code is recorded on a histogram plot as shown in the top left-hand curve in Figure 3.36. In the case of a 12-bit converter, several million samples are required in order to achieve statistically significant results. The histogram should follow the ideal probability density distribution of a sinewave, which is shown in the top right-hand curve in Figure 3.36. The histogram data is then normalized using the sinewave probability density function to obtain the DNL plot shown in the bottom curve of the figure. Integral non-linearity can be determined by compiling a cumulative histogram. The cumulative bin widths are the transition levels. Figure 3.37 shows an ac linearity plot obtained using the histogram method for the AD7870 12-bit 100kSPS ADC digitizing a 25kHz input signal at a 100kSPS rate. The absence of large spikes in the plot shows good differential linearity. More details of the mathematics involved is given in the AD7870 data sheet.

AC LINEARITY USING HISTOGRAMS

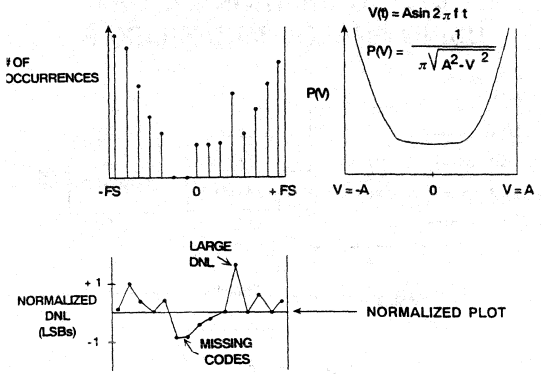

Figure 3.36

AC LINEARITY OF AD7870 12-BIT, 100kSPS ADC WITH 25 kHZ INPUT

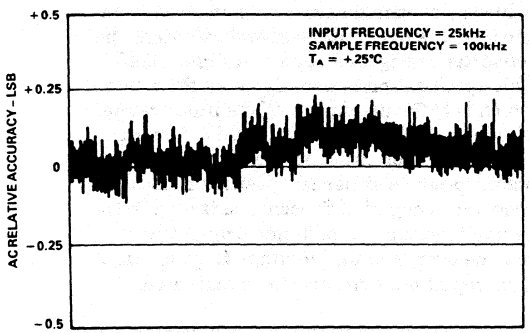

Figure 3.37

Aperture Delay Time (or Effective Aperture Delay Time)

Aperture delay time (sometimes called aperture time) is the amount of time from the leading edge of the sampling clock until the ADC actually takes the sample (see Figure 3.38). This specification is important because it helps the user to know when to apply the sampling clock with respect to the input signal timing. The variation or tolerance placed on this parameter from part to part is important in simultaneous sampling applications or other applications where the ADCs are required to track each other when processing dynamic signals.

MEASUREMENT OF EFFECTIVE APERTURE DELAY TIME

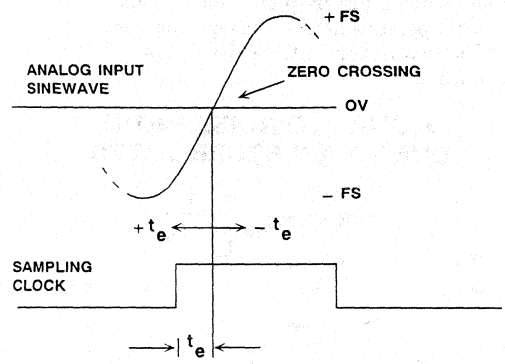

Figure 3.38

Aperture Jitter

Aperture jitter is the sample-to-sample variation in the effective point in time at which the actual sample is taken as shown in Figure 3.39. These errors generally emanate from several sources. In a practical ADC, the sampling clock is often phase-modulated by some unwanted source; the source can be wideband random noise, power line noise, or digital noise due to poor layout, bypassing, and grounding techniques. The resulting error can be expressed in terms of an rms time jitter. The corresponding rms voltage error caused by rms aperture jitter decreases the overall ADC signal-to-noise ratio. Phase jitter on the input sinewave can produce the same effect as jitter on the sampling clock.

EFFECTS OF APERTURE JITTER

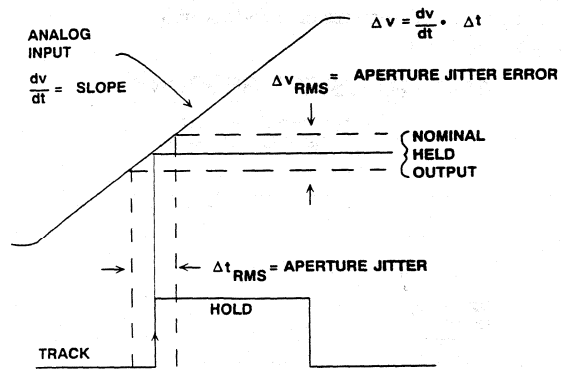

Figure 3.39

The SNR due exclusively to aperture jitter is plotted in Figure 3.40 as a function of fullscale sinewave input frequency for various values of aperture jitter. The equation for SNR due to aperture jitter is derived in Reference 1.

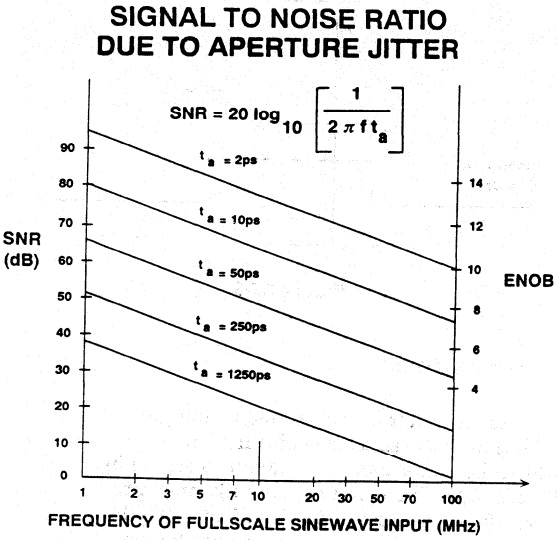

Figure 3.40

Transient Response or Settling Time

The transient response (or settling time) of an ADC is the time required for the ADC to settle to rated accuracy after the application of a fullscale step input (see Figure 3.41). This specification is critical in applications where the ADC is being driven by an analog multiplexer as shown in Figure 3.42. The multiplexer output can deliver a fullscale sample-to-sample change to the ADC input. If both the multiplexer and the ADC have not both sufficiently settled to the required accuracy, dc channel-to-channel crosstalk will result.

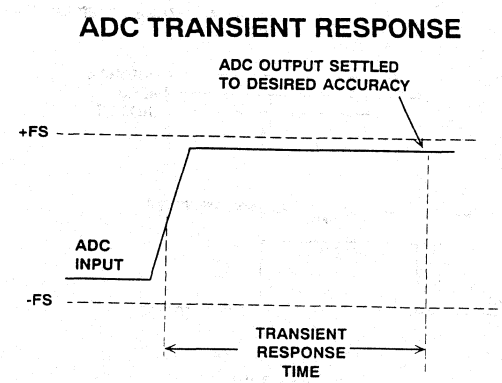

Figure 3.41

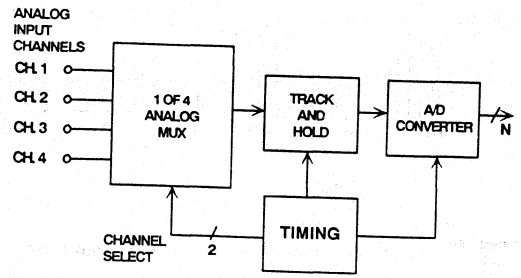

Figure 3.42

Overvoltage Recovery

Overvoltage recovery time is defined as that amount of time required for the ADC to achieve a specified accuracy, measured from the time the overvoltage signal re-enters the converter's range, as shown in Figure 3.43. This specification is usually given for a signal which is 50% outside the ADC's input range. Needless to say, the ADC should act as an ideal limiter for out-of-range signals and should produce either the positive fullscale code or the negative fullscale code during the overvoltage condition. Some converters provide over- and under-range flags to allow gain-adjustment circuits to be activated.

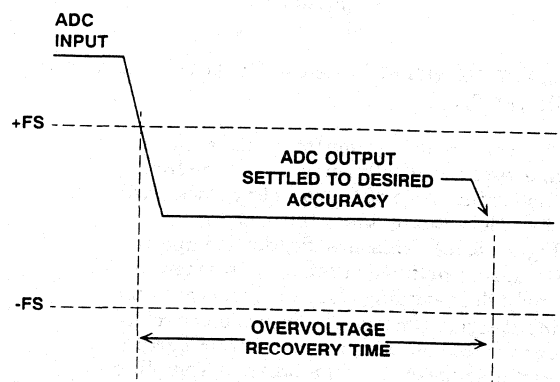

Figure 3.43

DAC Dynamic Performance

Since most DSP applications involve the eventual reconstruction of a dynamic analog signal, ac performance of DACs has become as important as ADC performance. Key DAC ac performance characteristics are given in Figure 3.45.

DAC DYNAMIC SPECIFICATIONS

- Settling Time
- Glitch Impulse Area
- Harmonic Distortion
- Signal-to-Noise Ratio
- Audio-Specific Specifications

Figure 3.44

Settling Time

Settling time of a DAC is traditionally defined as the time from the digital input transition (usually measured from the 50% point) until the DAC output settles to within a certain error band (usually 1/2 LSB) which is centered around the final value. As shown in Figure 3.45, a portion of the settling time may be due to a fixed propagation delay through the switches. If the DAC has a set of input latches or registers, the settling time should be measured from the 50% point of the latch strobe or register clock. Fullscale DAC settling time is measured for a digital input transition from 000...0 to 111...1. Midscale settling time is measured for a digital transition from 011...1 to 100...0 or 100...0 to 011...1.

It is entirely correct to define DAC settling time with respect to the output alone as shown in Figure 3.46. Settling time is measured from the time the output leaves a $\pm$1/2 LSB error band centered around the initial value until the time the output remains within a $\pm$ 1/2 LSB error band centered around the final value. The maximum DAC update rate allowable for $\pm$1/2 LSB fullscale settling time then becomes $f_{max} = 1/t_s$. Faster update rates can be used if sample-to-sample changes in the DAC input are limited to values less than fullscale.

DAC SETTLING TIME WAVEFORM

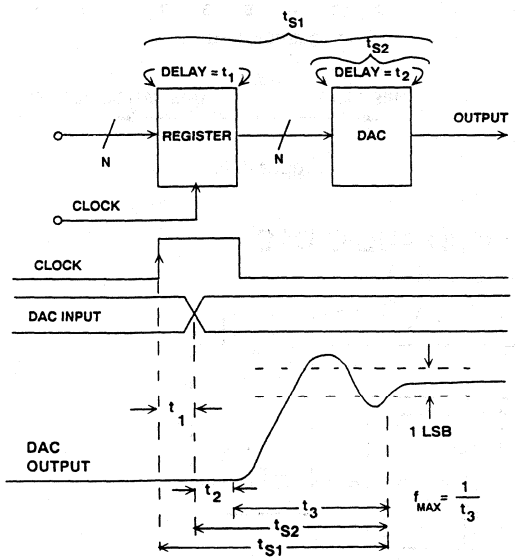

Figure 3.45

SETTLING TIME DEFINED WITH RESPECT TO DAC OUTPUT

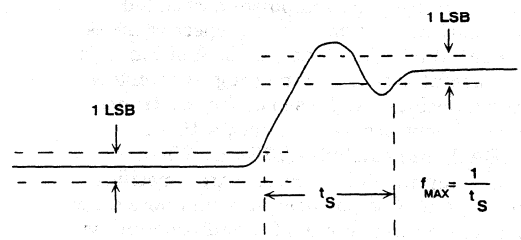

Figure 3.46

Glitch Impulse Area

Glitch impulse area is best understood by examining the waveform shown in Figure 3.47. DAC glitches occur because of digital input logic skew and unequal propagation delays through the DAC switches (a noteworthy exception to this is the sigma-delta DAC architecture to be discussed later in this seminar). The glitches are usually the largest at the midscale transition because all bits in the DAC are changing at this point. The glitch produced by the 011...1 to 100...0 transition is usually different from that produced by the 100...0 to 011...1 transition, so each must be analyzed. Glitch impulse area is simply the area of a particular glitch, and is usually measured in the units of pV-sec, therefore the fullscale output voltage of the DAC must be known in order to make meaningful comparisons between DACs. The term *glitch energy* is incorrect since the unit pV-sec is *not* a measure of energy.

From Figure 3.47 it is clear that there are six possible glitch impulse areas to deal with.

GLITCH IMPULSE WAVEFORMS

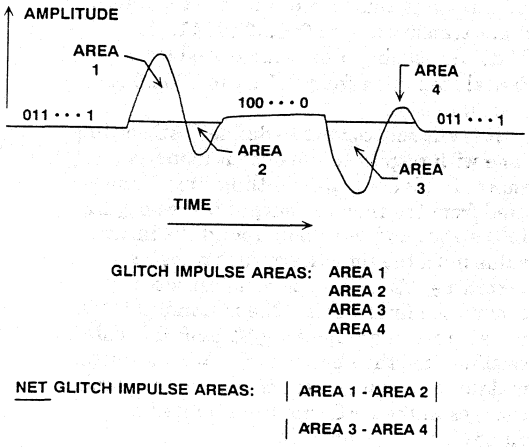

Figure 3.47

There are two glitch impulses associated with each transition. Their respective areas are designated 1,2,3, and 4. In addition, it is also useful to consider the *net glitch impulse* area associated with each of the two transitions. There are, respectively, AREA 1 - AREA 2, and AREA 3 - AREA 4. When examining the glitch impulse area specification on a DAC data sheet, it is therefore clear that there is much room for confusion unless a considerable amount of clarification is provided by the manufacturer.

Glitch impulse area remains constant regardless of filtering. Fast settling time specifications do not always imply low glitch impulse areas. The desirable situation is for the DAC to have a net glitch impulse area of zero for each of the two transitions, i.e., AREA 1 - AREA 2 = AREA 3 - AREA 4 = 0. In the ideal case, of course, each of the four areas would be zero.

HARMONIC DISTORTION

Because the net glitch impulse area is code-dependent, it will produce harmonics when the DAC is reconstructing a sinewave. A net midscale glitch occurs twice during a single cycle of the reconstructed sinewave (at each zero crossing) and, therefore, will produce a second harmonic of the sinewave as shown in Figure 3.48. Note that higher order harmonics of the sinewave which alias back into the Nyquist bandwidth are not filterable. It is difficult to predict the harmonic distortion caused by a specified net glitch impulse area, therefore, both specifications are required to adequately evaluate the dynamic performance of a reconstruction DAC.

Total harmonic distortion (THD) can be measured using DSP techniques as shown in Figure 3.49 for the AD1860 18-bit audio

EFFECTS OF DAC GLITCHES

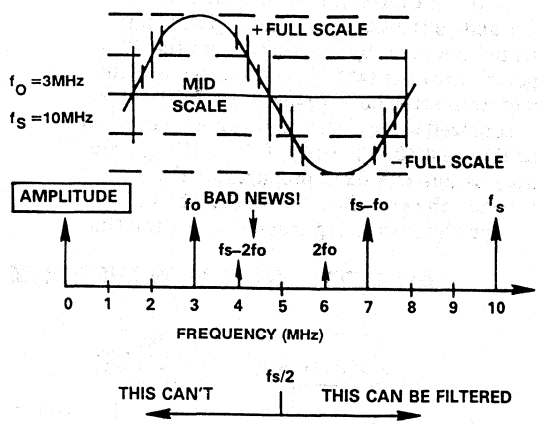

Figure 3.48

FFT TESTING OF AD1860 18-BIT AUDIO DAC

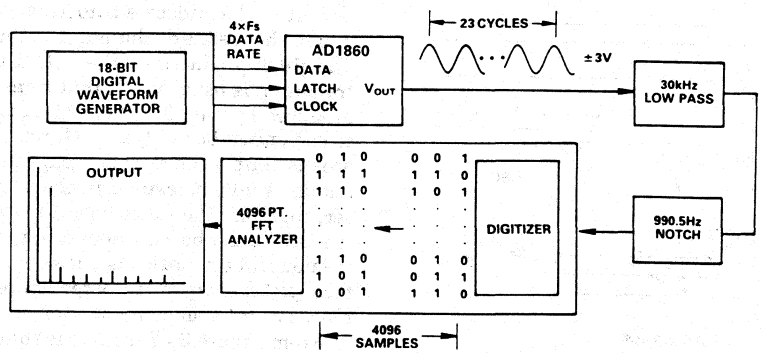

Figure 3.49

DAC. The DAC is driven with an 18-bit digital sinewave having a frequency of 990.5Hz, and the DAC update rate is 176.4kHz. The DSP digitizes 4096 samples of the output test waveform, incorporating 23 complete cycles of the sinewave. A 4096 point FFT is performed on the results of the test. The total harmonic distortion and the SNR is then calculated from the FFT results. The notch filter prevents the large-amplitude fundamental component at 990.5Hz from entering the digitizer, thereby allowing the entire digitizer range to be dedicated to processing the noise and harmonic components. Figure 3.50 shows a typical THD + noise plot for both a fullscale input and a -20dB input. It should be noted that neither a deglitcher nor an MSB trim are used in these measurements.

FFT-MEASURED
TOTAL HARMONIC DISTORTION
FOR AD1860 18-BIT AUDIO DAC

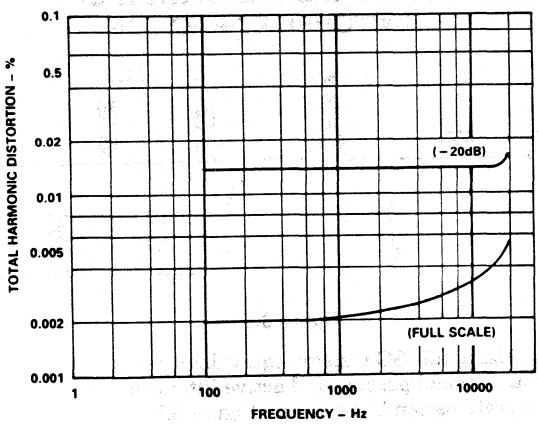

Figure 3.50

Deglitching DACs Using SHAs

SHAs can be used to deglitch DACs as shown in Figure 3.51. Just prior to latching new data into the DAC, the SHA is put into the hold mode so that the DAC switching glitches are isolated from the output. The switching transients produced by the SHA are code-independent and occur at the update frequency, hence, they are easily filterable.

SHA USED AS DEGLITCHER

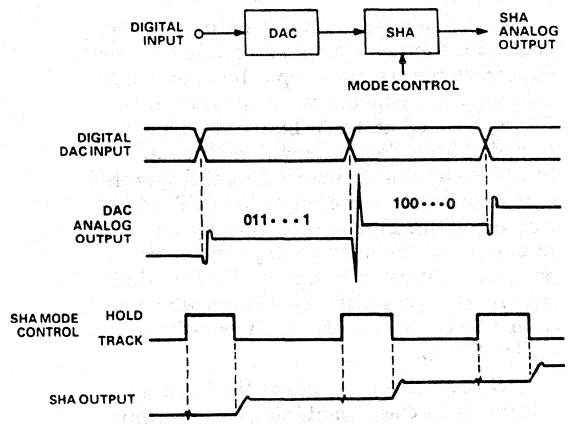

Figure 3.51

Sin(x)/x Frequency Rolloff Effect

The output of a reconstruction DAC can be visualized as a series of rectangular pulses whose width is equal to the reciprocal of the update rate as shown in Figure 3.52. Note that the reconstructed signal is down 3.92dB at the Nyquist limit with respect to the low frequency value. An inverse sin(x)/x filter is sometimes placed after the DAC to correct for this effect.

SIN X/X ROLLOFF

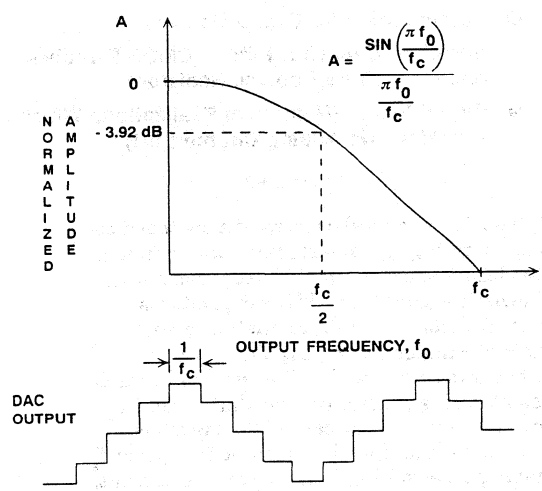

Figure 3.52

Signals were once filtered entirely in the continuous analog domain by configurations of passive components (typically inductors, resistors, and capacitors). Later, active filters, with op amps for buffering and gain, provided filter designers with additional flexibility and performance, but still operated continuously on analog signals. DSP led to stable and flexible discrete-time digital filters, where sampled analog signals are processed entirely by numerical calculations with filtering algorithms—some of which cannot be realized with continuous-time analog filters.

Switched-capacitor filters (SCFs) are an intermediate class, combining both continuous- and discrete-time aspects. They are usually implemented using CMOS switches and capacitors to simulate the behavior of resistors, therefore, many filter architectures can be realized entirely by a monolithic device without the need for external components. SCFs are particularly useful for voice and audio bandwidth signal applications in conjunction with DSP technology. Since SCFs are *sampling* devices, all the concepts of discrete time sampling apply to their use: Nyquist's theorem, aliasing, etc.

FILTERING TECHNIQUES

- **Crystals, SAWs**
- **Passive Components (R, L, and C)**
- **Active Filters (R, C, and Op-Amps)**
- **Switched Capacitor Filters (CMOS Switches and Capacitors Replace Resistors)**
- **Digital Filters (Numerical Realizations Which May Have No Analog Counterpart)**

Figure 3.53

The fundamental concept of a switched capacitor acting as a resistor can best be understood from the concept of charge transfer as shown in Figure 3.54. If the capacitor is switched from V_1 to V_2, an instantaneous charge transfer occurs, $\Delta Q = C(V_1 - V_2)$, either into or out of V_2. This assumes that C has no series resistance, and that V_1 and V_2 are ideal voltage sources. If the switch is thrown back and forth at a clock frequency, f_s (having a period T), then an average current, i, flows between V_1 and V_2 having a value $i = \Delta Q/T = C\Delta V/T$. The equivalent resistance, "R", that would give the same average current is given by :

$$\text{"R"} = \Delta V/i = T/C = 1/(Cf_s).$$

SWITCHED CAPACITOR "RESISTOR"

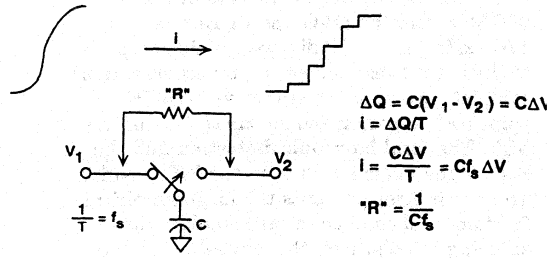

Figure 3.54

In an integrated circuit, the single-pole double-throw switch is implemented using CMOS switches driven by a non-overlapping two-phase clock as shown in Figure 3.55. A requirement for this technique to work is that the switches have very low on resistance and very high off resistance. This is precisely what CMOS technology offers.

CMOS IMPLEMENTATION OF SWITCHED CAPACITOR

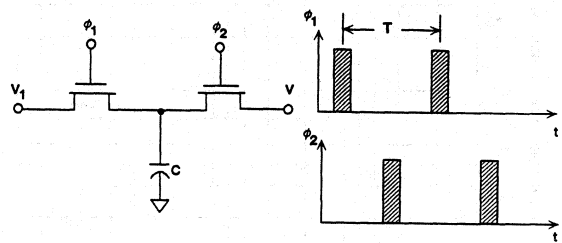

Figure 3.55

Using this SC resistor-equivalent, many conventional passive and active filter configurations can be realized. Figure 3.56 shows a single-pole passive RC filter and its SCF equivalent. The -3dB frequency of the RC filter is $1/(2\pi R_1 C_1)$. For the SCF version,

$$f_{3dB} = f_s C_1/(2\pi C_2).$$

SC EQUIVALENT OF PASSIVE RC NETWORK

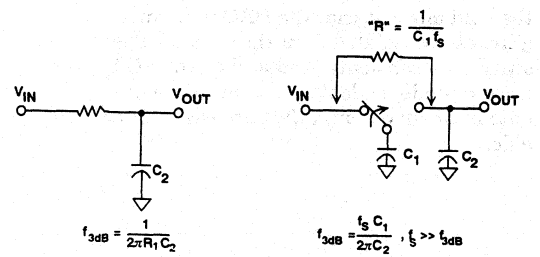

Figure 3.56

Note that for the SCF version, the *bandwidth depends on the sampling rate and the ratio of the capacitor values.* A major assumption which must be made is that $f_s >> f_{3dB}$ (typically 50 to 100) to minimize the effects of time-sampling and charge-sharing. Using the SCF concept, critical frequencies are therefore determined by capacitor ratios and the sampling clock frequency, both of which can be made precise and drift free.

Audio and voiceband filtering with SC filters can greatly reduce passive component physical size. To implement audio filters, a resistance on the order of 10MΩ is required if a monolithic capacitor of reasonable size (~ 10pF) is to be used. This value of resistance is easily achieved by switching a 1pF capacitor at a 100kHz rate, requiring a silicon area of approximately 0.01mm². If the 10MΩ resistor were implemented using polysilicon or diffusion, the area required would be at least 100 times larger.

SWITCHED CAPACITOR FILTER ADVANTAGES

- **Filter Bandwidths Proportional to Capacitance Ratios Not Absolute Values**
- **Filter Bandwidths Variable with Clock Frequency**
- **Defined Like Classic Analog Filters**
- **Low Values of Capacitance Required for Audio Frequencies: 1pF Capacitor Switched at 100kSPS = 10 MΩ "Resistance"**
- **SCFs Ideally Suited to DSP CMOS Processes**

Figure 3.57

By using SC resistors in conjunction with other capacitors and op amps, it is possible to realize many of the circuit configurations used in conventional RC active filters. Unlike digital filters, SC filters may be defined exactly like analog filters. A first-order continuous-time active lowpass RC filter and its SC counterpart are shown in Figure 3.58.

Since they sample analog signals, SC filters must usually be preceded by a continuous-time antialiasing prefilter to eliminate spectral components above the Nyquist frequency. Since the SC filter sampling rate is usually much higher than its passband, a single or double pole RC filter is usually sufficient for this purpose.

FIRST ORDER ACTIVE LOWPASS RC FILTER AND SCF EQUIVALENT

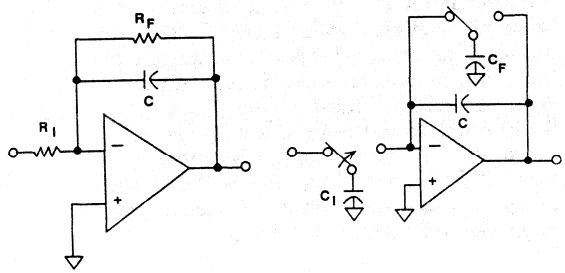

Figure 3.58

Differential amplifiers are often used in analog circuits to achieve good common mode rejection of unwanted signals such as power line noise, etc. The same principles can be used in designing switched capacitor filters. Figure 3.59 shows an active differential integrator and its switched capacitor equivalent. In addition to providing good CMRR to noise, the differential configuration also provides common mode rejection to the transients caused by the operation of the switches. Switched capacitor integrators are often used in the modulator circuits of a Sigma-Delta ADCs as will be discussed later in this seminar.

ACTIVE DIFFERENTIAL INTEGRATOR AND SCF EQUIVALENT

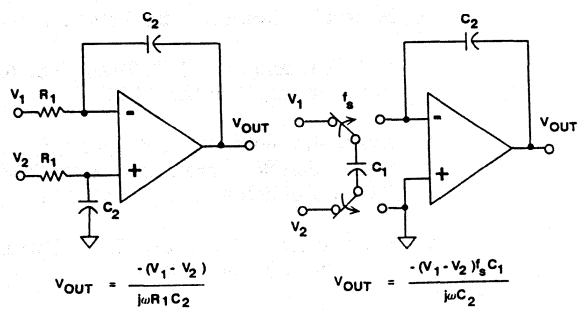

Figure 3.59

Switched capacitor filters are subject to several limitations and error sources. Their usefulness is limited to frequencies in the audio bandwidth, since sampling rates greater than a few hundred kilohertz cannot be readily achieved with current CMOS technology. The switched capacitors and op amps introduce random noise, and leakage currents can produce offset errors. Clock feedthrough from the switches themselves can produce synchronous errors. Finally, since SCFs are sampling devices, large oversampling ratios are usually required in order to prevent errors due to aliasing.

SWITCHED CAPACITOR LIMITATIONS AND ERROR SOURCES

- **Limited to Lower Frequencies**
- **Noise, Offset, and Distortion**
- **Clock Feedthrough from Switches**
- **Must Obey the Laws of Nyquist (Requires Antialiasing Filter)**

Figure 3.60

REFERENCES

1. Frederic J. Harris, "On the Use of Windows for Harmonic Analysis with the Discrete Fourier Transform", **IEEE Proceedings**, Vol. 66, No. 1, Jan. 1978, pp. 51-83.

2. Joey Doernberg, Hae-Seung Lee, David A. Hodges, "Full Speed Testing of A/D Converters", **IEEE Journal of Solid State Circuits**, Vol. SC-19, No. 6, Dec. 1984, pp. 820-827.

3. James R. Andrews, Barry A. Bell, Norris S. Nahman, and Eugene E. Baldwin, "Reference Waveform Flat Pulse Generator", **IEEE Transactions on Instrumentation and Measurement**, Vol. IM-32, No. 1, March 1983, pp. 27-32.

4. Brendan Coleman, Pat Meehan, John Reidy and Pat Weeks, "Coherent Sampling Helps When Specifying DSP A/D Converters", **EDN**, October 15, 1987, pp. 145-152.

5. Howard K. Schoenwetter, "A Programmable Voltage Step Generator for Testing Waveform Recorders", **IEEE Transactions on Instrumentation and Measurement**, Vol. IM-33, No. 3, Sept. 1984, pp. 196-200.

6. Robert W. Ramirez, **The FFT: Fundamental and Concepts**, Prentice-Hall, 1985.

7. R.B. Blackman and J.W. Tukey, **The Measurement of Power Spectra**, Dover Publications, New York, 1958.

8. James J. Colotti, "Digital Dynamic Analysis of A/D Conversion Systems Through Evaluation Software Based on FFT/DFT Analysis", **RF Expo East 1987 Proceedings**, Cardiff Publishing Co., pp. 245-272.

9. Sid Kaufman, "Multistage Error Correcting A/D Converters", **Electronic Products**, April 18, 1983, pp. 103-110.

10. **HP Journal**, Nov. 1982, Vol. 33, No. 11

11. HP Product Note 5180A-2.

12. **HP Journal**, April 1988, Vol. 39, No. 2.

13. **HP Journal**, June 1988, Vol. 39, No. 3.

14. Dan Sheingold, Editor, **Analog-to-Digital Conversion Handbook, Third Edition**, Prentice-Hall, 1986.

15. W.R. Bennett, "Spectra of Quantized Signals", **Bell System Technical Journal**, No. 27, July 1948, pp. 446-472.

16. G.A. Gray and G.W. Zeoli, "Quantization and Saturation Noise Due to Analog-Digital Conversion", **IEEE Transactions on Aerospace and Electronic Systems**, Jan. 1971, pp. 222-223.

17. M.J. Tant, **The White Noise Book**, Marconi Instruments, July 1974.

18. W.A. Kester, "PCM Signal Codecs for Video Applications", **SMPTE Journal**, No. 88, November 1979, pp. 770-778.

19. Lawrence Rabiner and Bernard Gold, **Theory and Application of Digital Signal Processing**, Prentice-Hall, 1975.

20. Matthew Mahoney, **DSP -Based Testing of Analog and Mixed-Signal Circuits,** , IEEE Computer Society Press, Washington, D.C., 1987.

21. IEEE Trial-Use Standard for Digitizing Waveform Recorders, No. 1057-1988.

22. Richard J. Higgins, **Digital Signal Processing in VSLI**, Prentice-Hall, 1990.

23. **High Speed Design Seminar**, Analog Devices, 1990.

24. M. S. Ghausi and K. R. Laker, **Modern Filter Design: Active RC and Switched Capacitors**, Prentice Hall, 1981.

25. Brodersen, Gray, and Hodges, *MOS Switched-Capacitor Filters,* **Proc. IEEE Vol. 67,** January 1979, pp. 61 - 65.

26. J. T. Caves, et. al., *Sampled Analog Filtering Using Switched Capacitors as Resistor Equivalents*, **IEEE J. Solid State Circuits Vol. SC-12**, pp. 592 - 599.

Appendix Contents

Page

Analog-to-Digital and Digital-to-Analog Conversion Factors 25–2
Digital Codes 25–3
Settling Time (Expressed in Time Constants) Required to Achieve Specified Accuracy 25-4
Voltage Gain to dB Conversion 25-5
Worlewide Sales Directory 25-6

25

A/D and D/A Conversion Factors

No. of Bits	2(n)	Resolution %	ppm	Volts	dB
0	1	100.0	1,000,000	10.0000000	0.0
1	2	50.0	500,000	5.0000000	−6.0
2	4	25.0	250,000	2.5000000	−12.0
3	8	12.5	125,000	1.2500000	−18.1
4	16	6.25	62,500	0.6250000	−24.1
5	32	3.125	31,250	0.3125000	−30.1
6	64	1.563	15,625	0.1562500	−36.1
7	128	0.781	7,813	0.0781250	−42.1
8	256	0.391	3,906	0.0390625	−48.2
9	512	0.195	1,953	0.0195313	−54.2
10	1,024	0.0977	977	0.0097656	−60.2
11	2,048	0.0488	488	0.0048828	−66.2
12	4,096	0.0244	244	0.0024414	−72.2
13	8,192	0.0122	122	0.0012207	−78.3
14	16,384	0.0061	61	0.0006104	−84.3
15	32,768	0.00305	31	0.0003052	−90.3
16	65,536	0.00153	15	0.0001526	−96.3
17	131,072	0.00076	8	0.0000763	−102.4
18	262,144	0.00038	4	0.0000381	−108.4
19	524,288	0.00019	2	0.0000191	−114.4
20	1,048,576	0.00010	1	0.0000095	−120.4
21	2,097,152	0.00005	0.5	0.0000048	−126.4
22	4,194,304	0.00002	0.24	0.0000024	−132.5

Decimal	Octal	Hexi-Decimal	Natural Binary	Comp Binary	Offset Binary	Sign + Magnitude	Twos Complement	Ones Complement
15	17	F	1111	0000				
14	16	E	1110	0001				
13	15	D	1101	0010				
12	14	C	1100	0011				
11	13	B	1011	0100				
10	12	A	1010	0101				
9	11	9	1001	0110				
8	10	8	1000	0111				
7	7	7	0111	1000	1111	0111	0111	0111
6	6	6	0110	1001	1110	0110	0110	0110
5	5	5	0101	1010	1101	0101	0101	0101
4	4	4	0100	1011	1100	0100	0100	0100
3	3	3	0011	1100	1011	0011	0011	0011
2	2	2	0010	1101	1010	0010	0010	0010
1	1	1	0001	1110	1001	0001	0001	0001
+0	+0	+0	0000	1111	1000	0000	0000	0000
−0					1000	1000	0000	1111
−1					0111	1001	1111	1110
−2					0110	1010	1110	1101
−3					0101	1011	1101	1100
−4					0100	1100	1100	1011
−5					0011	1101	1011	1010
−6					0010	1110	1010	1001
−7					0001	1111	1001	1000
−8					0000		1000	

Settling Time (Expressed in Time Constants) Required to Achieve Specified Accuracy

Accuracy Bits	Resolution %	Number of Time Constants
0	100.0	0.00
1	50.0	0.69
2	25.0	1.39
3	12.5	2.08
4	6.25	2.77
5	3.125	3.47
6	1.563	4.16
7	0.781	4.85
8	0.391	5.55
9	0.195	6.24
10	0.0977	6.93
11	0.0488	7.62
12	0.0244	8.32
13	0.0122	9.01
14	0.0061	9.70
15	0.00305	10.40
16	0.00153	11.09
17	0.00076	11.78
18	0.00038	12.48
19	0.00019	13.17
20	0.000095	13.86
21	0.000048	14.56
22	0.000024	15.25

Voltage Gain to dB Conversion

dB	Gain	dB	Gain	dB	Gain
0	1.000	10	3.162	20	10
1	1.122	11	3.548	30	32
2	1.259	12	3.981	40	100
3	1.413	13	4.467	50	316
4	1.585	14	5.012	60	1,000
5	1.778	15	5.623	70	3,162
6	1.995	16	6.310	80	10,000
7	2.239	17	7.079	90	31,623
8	2.512	18	7.943	100	100,000
9	2.818	19	8.913	110	316,228
				120	1,000,000

Worldwide Sales Directory

North American Sales Offices and Representatives

Alabama: *404-263-3719, 205-534-0044
Alaska: *206-575-6344, *714-641-9391
Arizona: 602-949-0048
Arkansas: *214-231-5094
California: *714-641-9391, *408-559-2037
Colorado: 303-443-5337
Connecticut: *516-366-0765
Delaware: *215-643-7790
Florida: *407-660-8444
Georgia: *404-263-3719, 404-263-0320
Hawaii: *714-641-9391
Idaho: 303-443-5337, *206-575-6344
Illinois: *708-519-1777
Indiana: *219-489-3083
Iowa: *708-519-1777
Kansas: 913-829-2800
Kentucky: 205-534-0044
Louisiana: *214-231-5094
Maine: *617-461-3000
Maryland: *410-992-1995
Massachusetts: *617-461-3000
Michigan: 313-489-1500, *313-348-5795
Minnesota: 612-835-9234
Mississippi: *404-263-3719, 205-534-0044
Missouri: 314-521-2044, 913-829-2800
Montana: 801-466-9336, *714-641-9391
Nebraska: 913-829-2800
Nevada: *408-559-2037, *714-641-9391
New Hampshire: *617-461-3000
New Jersey: *516-366-0765 (north), *215-643-7790 (south)
New Mexico: 303-443-5337
New York: *315-682-5571, *516-366-0765 (metro), 716-425-9100 (upstate)
North Carolina: *404-263-3719, 919-460-7771
North Dakota: 612-835-9234
Ohio: 216-247-0060, 614-487-1882
Oklahoma: *214-231-5094
Oregon: *206-575-6344
Pennsylvania: *215-643-7790 (east), 412-745-8441 (west)
Rhode Island: *617-461-3000
South Carolina: 919-460-7771
South Dakota: 612-835-9234
Tennessee: *404-263-3719, 404-263-0320 (east), 404-534-0044 (west)
Texas: *214-231-5094
Utah: 801-466-9336
Vermont: *617-461-3000
Virginia: *410-992-1995
Washington: *206-575-6344
West Virginia: 412-745-8441
Wisconsin: *708-519-1777
Wyoming: 801-466-9336
Puerto Rico: *407-660-8444, 809-753-8905
Canada: 416-821-7800, 613-564-0014, 514-697-0801, 604-465-6892
Mexico: *617-461-3000

North American Distributors

Alabama: *Alliance* 505-292-3360; *Allied* 205-721-3500, 205-476-1875; *Bell* 205-837-1074; *Future* 205-830-2322; *Hall-Mark* 205-837-8700; *Pioneer* 205-837-9300
Arizona: *Alliance* 505-292-3360; *Allied* 602-437-9080; *Bell* 602-966-7800; *Future* 602-968-7140; *Hall-Mark* 602-437-1200; *Wyle* 602-437-2088
California: *Alliance* 505-292-3360; *Allied* 415-770-0590, 818-882-6981, 213-373-8828, 714-727-3010, 916-632-3I04, 619-279-2550, 415-952-9599; *Bell* 310-826-6778, 916-652-0414, 408-734-8570, 818-879-9494, 714-895-7801, 619-268-1277; *Future* 818-772-6240, 619-625-2800, 714-250-4141, 408-434-1122; *Hall-Mark* 818-773-4500, 714-727-6000, 619-268-1201, 408-432-4000, 916-624-9781; *Pioneer* 800-535-1430, 213-320-4820, 619-434-4107, 408-954-9100; *Semi-Dice* 310-594-4631, 714-952-2216, 408-980-9777, 800-325-6030 (in CA), 800-626-4411, 800-345-6633 (outside CA, MA); *Wyle* 714-863-9953, 818-880-9000, 800-288-9953, 916-638-5282, 800-627-9953, 408-727-2500, 619-565-9171
Colorado: *Alliance* 505-292-3360; *Allied* 303-790-1664; *Bell* 303-691-9010; *Future* 303-421-0123 *Hall-Mark* 303-790-1662; *Wyle* 303-457-9953
Connecticut: *Alliance* 617-756-1910; *Allied* 203-272-7730; *Bell* 203-269-6801; *Future* 203-743-9594; *Hall-Mark* 203-271-2844; *Pioneer* 203-929-5600
Florida: *Alliance* 505-292-3360; *Allied* 305-978-3008, 904-346-0761, 407-831-3331, 813-541-4660; *Bell* 407-339-0078, 305-421-1997, 813-541-4434; *Chip Supply* 407-298-7100; *Future* 407-767-8414, 813-530-1222; *Hall-Mark* 407-830-5855, 305-971-9280, 813-541-7440; *Pioneer* 407-834-9090, 305-428-8877
Georgia: *Alliance* 505-292-3360; *Allied* 404-497-9544; *Bell* 404-662-0923; *Future* 404-441-7676; *Hall-Mark* 404-623-4400; *Pioneer* 404-623-1003
Illinois: *Alliance* 505-292-3360; *Allied* 708-918-0250, 708-535-0038, 708-860-0007; *Bell* 708-640-1910, 217-328-1077; *Future* 708-882-1255; *Hall-Mark* 708-860-3800; *Newark* 312-784-5100; *Pioneer* 708-495-9680
Indiana: *Alliance* 505-292-3360; *Allied* 317-875-0448; *Bell* 317-875-8200, 219-423-3422; *Hall-Mark* 317-872-8875; *Pioneer* 317-573-0880
Iowa: *Alliance* 505-292-3360; *Allied* 515-223-6488; *Bell* 319-395-0730
Kansas: *Alliance* 505-292-3360; *Allied* 913-599-6227; *Hall-Mark* 913-888-4747
Kentucky: *Alliance* 505-292-3360; *Allied* 502-452-2293; *Bell* 800-582-7492, 800-445-8795
Louisiana: *Alliance* 505-292-3360; *Allied* 504-885-1195
Maryland: *Alliance* 505-292-3360; *Allied* 301-381-1560; *Bell* 410-290-5100; *Future* 410-290-0600; *Hall-Mark* 301-988-9800; *Pioneer* 301-921-0660
Massachusetts: *Alliance* 617-729-5800, 617-756-1910; *Allied* 508-667-7921, 617-255-0361; *Bell* 508-474-8880; *Future* 508-779-3000; *Hall-Mark* 508-667-0902, 508-935-9777; *Pioneer* 617-861-9200; *Semi-Dice* 508-238-8344; *Wyle* 617-272-7300
Michigan: *Alliance* 505-292-3360; *Allied* 313-462-0190, 616-949-8200; *Bell* 800-582-7492, 313-347-6633; *Future* 313-261-5270; *Hall-Mark* 313-416-5800; *Pioneer* 616-698-1800, 313-525-1800
Minnesota: *Alliance* 505-292-3360; *Allied* 612-881-0838; *Bell* 612-941-1493; *Future* 612-944-2200; *Hall-Mark* 612-881-2600; *Pioneer* 612-944-3355
Missouri: *Alliance* 505-292-3360; *Allied* 314-291-7031; *Bell* 314-647-2355; *Future* 314-469-6805; *Hall-Mark* 314-291-5350; *Pioneer* 314-542-3077
Nebraska: *Alliance* 505-292-3360; *Allied* 402-697-0038
New Hampshire: *Alliance* 617-756-1910; *Allied* 603-883-3006
New Jersey: *Alliance* 505-292-3360; *Allied* 201-428-3350, 609-234-7769; *Bell* 201-227-6060; *Future* 201-299-0400; *Hall-Mark* 609-235-1900, 201-515-3000; *Pioneer* 201-575-3510
New Mexico: *Alliance* 505-292-3360; *Bell* 505-292-2700
New York: *Alliance* 617-756-1910; *Allied* 516-248-2360, 908-572-9600, 914-452-1470, 716-425-3850; *Future* 716-272-1120, 315-451-2371, 516-234-4000; *Hall-Mark* 716-425-3300, 516-737-0600; *Pioneer* 716-381-7070, 516-921-8700, 607-722-9300, 212-971-9560; *Semi-Dice* 516-234-7676
North Carolina: *Alliance* 505-292-3360; *Allied* 919-876-5845; *Bell* 919-821-4041; *Future* 704-455-9030, 919-790-7111; *Hall-Mark* 919-872-0712; *Pioneer* 919-544-5400, 704-527-8188, 919-460-1530
Ohio: *Alliance* 505-292-3360; *Allied* 513-771-6990, 216-248-4200, 614-785-1270; *Bell* 800-348-4619, 513-435-8660, 513-434-8231; *Future* 216-449-6996; *Hall-Mark* 216-349-4632, 614-888-3313; *Pioneer* 216-587-3600, 513-236-9900, 614-221-0043
Oklahoma: *Alliance* 505-292-3360; *Allied* 918-664-0844; *Hall-Mark* 918-254-6110; *Pioneer* 918-665-7840
Oregon: *Alliance* 505-292-3360; *Allied* 503-245-4225; *Bell* 503-644-1500; *Future* 503-645-9454; *Wyle* 503-643-7900
Pennsylvania: *Alliance* 505-292-3360; *Allied* 717-975-3570, 412-367-4124; *Bell* 215-953-2800, 800-525-6666; *Future* 609-596-4080; *Hall-Mark* 215-355-7300; *Pioneer* 215-674-4000, 412-782-2300
South Carolina: *Alliance* 505-292-3360; *Allied* 803-288-8835; *Bell* 800-322-7908
Tennessee: *Alliance* 505-292-3360; *Allied* 901-367-2895; *Bell* 615-367-4400
Texas: *Alliance* 505-292-3360; *Allied* 817-595-3500, 512-389-3771, 214-553-4370, 214-570-5196, 915-779-6294, 713-446-8005, 713-455-3993; *Bell* 214-690-0466; *Future* 713-556-8696, 214-437-2437; *Hall-Mark* 214-343-5000, 512-258-8848, 214-553-4300, 713-781-6100; *Pioneer* 512-835-4000, 214-386-7300, 713-495-4700, 512-377-3440; *Wyle* 512-345-8853, 214-235-9953, 713-879-9953, 800-888-9953
Utah: *Alliance* 505-292-3360; *Allied* 801-261-5244; *Bell* 801-255-9611; *Future* 801-972-8489; *Hall-Mark* 801-973-4737; *Wyle* 801-974-9953
Vermont: *Alliance 617-756-1910;* Allied *603-883-3006*
Virginia: *Alliance 505-292-3360;* Allied *804-363-8862, 703-644-9515*
Washington: *Alliance* 505-292-3360; *Allied* 206-547-2827; *Bell* 206-867-5410, 213-826-6778; *Future* 206-489-3400; *Hall-Mark* 206-547-0415; *Wyle* 206-881-1150
Wisconsin: *Alliance* 505-292-3360; *Allied* 414-796-1280; *Bell* 414-547-8879; *Future* 414-786-1884; *Hall-Mark* 414-797-7844; *Pioneer* 414-784-3480
Canada: *Future* 514-694-7710, 613-820-8313, 416-638-4771, 604-294-1166, 403-436-2858, 416-612-9200; *Zentronics* 403-482-3038, 604-273-5575, 403-295-8838, 416-564-8600, 613-226-8840, 418-654-1077, 204-694-1957, 514-737-9700, 416-798-7710, 416-507-2600

*Analog Devices, Inc., Direct Sales Offices

5/93

International Direct Sales Offices

Austria: (222) 88 55 04-0
Belgium: (03) 237 16 72
Denmark: (42) 84 58 00
France: (1) 46 66 25 25, 76 41 91 83
Germany: 089/57 005-0, 030/391 90 35, 04181/80 51, 0221/68 60 06, 0711/88 11 33
India: (80) 26 72 72
Israel: (052) 911 415
Italy: (2) 665 00 120, (11) 24 87 789, (6) 86 200 306
Japan: (3) 3263 6826, (6) 372 1814
Korea: (2) 554-3301
Netherlands: (1620) 815 00
Sweden: (8) 282 740
Switzerland: (1) 820 01 02, (21) 803 25 50
Taiwan: (02) 740 0099
United Kingdom: 0932 25 33 20, 0506 30306, 0932 246 200

International Distributors

Australia: *N.S.D. Australia* (3) 8900970, (2) 646-5255, (7) 845-1911, (8) 211-8499
Austria: *A.D.E.C.* (0222) 88 92 876-0, *RS-Componenets* (02852) 505
Belgium: *ACAL NV/SA (Betea Components Division)* (02) 720 59 83, *TEXIM ELECTRONICS* (02) 460 52 82
Brazil: *Hicad Sistemas Ltda.* (11) 531-9355
Bulgaria: *K2 ELECTRONICS* (92) 890412
Czechoslovakia: *AMTEK SPOL. S.R.O.* (5) 750 988, *TESLA PIESTANY* (838) 522306
Denmark: *MER-EL A/S* (42) 57 10 00
Finland: *Oxxo Oy 90*-3455 377
France: *DIMACEL COMPOSANTS* (1) 40 87 70 00, *AVNET Diffusion* (1) 49 65 26 26, RS *COMPONENTS* 44 84 72 72, *SCAIB* (1) 46 87 23 13
Germany: *SASCO* 089/46 11-0, *Semitron* 07742/80 01-0, *SPOERLE Elektronic* 06103/30 4-0, *JERMYN* 06431/508-0
Greece: *MICRELEC* (1) 539 02 69
Hong Kong: *General Engineers* 503-2212, *Tektron Electronics (HK) Ltd.* 388-0629
Hungary: *SMD Technology KFT* (1) 189 58 15
India: *Analog Sales* (212) 342150, (11) 6862460, (80) 560506
Italy: *DeMico* (2) 95343600, *Hellis S.A.S.* (536) 804104, *LA Tecknica Due* (11) 2425905, *Pantronic S.R.L.* (6) 6276209, *Special-Ind* (2) 66805177, *Tecknica Due* (438) 555447
Korea: *Tong Baek Electronics* (2) 715-6625, *Analog World* (2) 7015993
Malaysia: *Excelpoint Systems (PTE) Ltd. 03-2448929*
Mexico: *Canadien* (83) 652020
Netherlands: *TEXIM ELECTRONICS B.V.* (05427) 333 33, *TME Components B.V.* 073-211010
New Zealand: *N.S.D. Australia* (61 3) 8900970
Norway: *BIT ELEKTRONIKK* 02 98 13 70
People's Republic of China – Beijing: *Excelpoint Company Limited* (1) 500 7788, Ext 1625
Poland: *ALFINE* (61) 301 450
Romania: *TOP 9+* (400) 12787
Portugal: *ATD Electronica* (1) 847 2200, *Comelta* (1) 942 4106
Russia: *ARGUSSOFT Co.* (7095) 2883602, *ARGUS TRADING Ltd.* (7095) 945 2777, *AUTEX Ltd.* (7095) 3347741, *TESLA PIESTANY* (838) 522306
Singapore: *Excelpoint Systems PTE* 741-8966
South Africa: *Analog Data Products CC* (11) 444-8160
Spain: *Comelta* (1) 7543001, *Selco* (1) 326 4213
Sweden: *Integrerad Elektronik* 08-80 46 85
Switzerland: *Elbatex AG* 056/275 111, Fabrimex Distribution AG 01/386 86 86
Taiwan: *Andev Technology Co., Ltd.* (02) 763-0910, *Chieftron Enterprise Co., Ltd.* (02) 722 3570, *Jeritronics Ltd.* (02) 585-1636, *Lite-On Inc.* (02) 918 6699
Turkey: *ELEKTRO* (1) 337 22 45
United Kingdom: *Arrow Electronics (UK) Limited* 0234 270 777, *Avnet-Access* 0462 480888, *Jermyn Distribution* (0732) 743 743, *Phoenix Electronics Limited* 0555 515 66, *Polar Electronics Limited* 0525 377 093
Yugoslavia: *ALMA-ELECTRONIC* (41) 428 678, *ELECTRA* (40) 826545

EUROPEAN HEADQUARTERS
Edelsbergstraße 8-10, W-8000 München 21, Germany; Tel: 089/57 005-0; Fax: 089/57 005-257

JAPANESE HEADQUARTERS
Daini Jibiki Building, 7-8 Kojimachi 4 Chome, Chiyoda-ku, Tokyo, Japan; Tel: (3) 3263 6826; Fax: (3) 3221 5187

WORLDWIDE HEADQUARTERS
One Technology Way, P.O. Box 9106, Norwood, MA 02062-9106, U.S.A.
Tel: (617) 329-4700, (1-800) 262-5643 (U.S.A. only); Fax: (617) 326-8703

5/93

Notes

Notes

Notes

Notes

Notes

Notes

Notes

Index
Contents

Page

Product Index 26–3
Subject Index 26–7
Application Note (AN) Number Index 26–17
Author Index 26–21

Product Index

	Page
AC1226	11–5
AD2S80	16–17, 21
AD2S80A	16–11-16, 21, 53
AD2S81	16–17
AD2S81A	16–11-16, 21
AD2S82	16–17
AD2S82A	16–21
AD2S83	16–21, 30, 53-55
AD2S90	16–45-47, 49-51
AD2S100	16–33, 40-41
AD2S110	16–40-42
AD208	8–71-72; 11–3, 5-6, 8-9
AD230	5–3
AD232	5–6
AD233	5–4
AD235	5–4
AD241	5–3, 7-8
AD246	11–5-6, 8-9
AD301	23–20
AD311	8–118, 166; 23–20, 42, 44
AD346	3–29-32
AD380	8–112; 13–24
AD390	13–24
AD394	13–24
AD395	13–24
AD396	13–24
AD507	6–34-36; 13–20, 24
AD508	13–24
AD509	6–31; 8–20; 13–20-21; 23–26
AD510	13–21, 24; 24–4
AD517	8–72; 13–24; 23–29
AD518	13–24
AD521	8–73; 13–22-23, 23-24; 23–26
AD522	10–8-10
AD524	6–11, 21, 24; 8–7-8; 10–11, 12, 13-18; 13–24; 19–7-10
AD526	13–24
AD532	13–24
AD533	13–24
AD534	6–39; 23–25; 13–24
AD535	13–24
AD536	6–22-23
AD536A	13–24; 17–3-4
AD537	8–71-73; 23–3, 13-32
AD538	6–3-9; 13–24; 24–51-52
AD539	6–29, 31-37, 39; 13–93-94
AD542	13–24
AD544	6–18; 8–5-6, 8-9, 118, 182; 13–24
AD545A	13–24
AD546	13–24
AD547	8–166; 13–24
AD548	13–24
AD549	13–24; 24–57-58
AD557	13–24; 14–4
AD558	6–36; 13–24
AD561	8–71, 73-74; 13–24
AD562	23–26
AD565A	13–24
AD566A	13–24
AD568	13–24
AD572	8–20
AD574A	18–8-9
AD580	3–8; 8–209-210; 13–24; 19–7, 10
AD581	6–27; 8–71; 13–24
AD582	13–24
AD584	6–13, 22-23, 26; 8–112, 117, 119-120, 182, 203; 13–24; 23–42
AD585	3–29; 18–8-10
AD586	3–65; 8–208; 13–24
AD587	13–24
AD588	3–30; 11–8; 13–24
AD589	6–52-53; 23–43, 53, 56; 24–52
AD590	19–3-10
AD594	21–4-8; 23–46
AD595	21–4-8; 23–46
AD611	6–20, 23, 33, 36; 23–56
AD624	10–12; 13–24
AD625	13–24
AD630	6–11, 13, 17-21, 27
AD636	13–24; 17–3-4
AD637	13–24; 17–3-5
AD639	23–52
AD642	13–24
AD644	6–22-24; 8–113, 118, 163; 10–18; 13–24

	Page
AD645	13–24
AD647	8–166; 13–24
AD648	13–24; 16–17
AD650	13–24; 23–3, 11, 61-62, 65
AD651	23–8
AD652	13–24
AD654	23–4, 5-8, 10, 33-60
AD662	13–24
AD664	13–24
AD667	13–24
AD668	13–24
AD670	3–3-9
AD671	3–59-65; 18–13-14
AD678	24–91-92, 95-96
AD679	3–96, 100-101
AD684	3–29
AD688	13–24
AD689	13–24
AD693	22–3-10
AD704	13–24
AD705	11–6; 13–24
AD706	13–24
AD707	6–27; 13–24; 14–4
AD708	6–11, 27; 13–24
AD711	3–76; 5–23-25; 6–41; 8–199, 208-209; 13–24; 23–52
AD712	8–208, 210; 13–24; 16–17; 23–52
AD713	13–24; 16–17
AD734	6–39
AD736	13–24; 17–3-9, 12, 15
AD737	13–24; 17–3-11, 13-14
AD741	8–71-74, 109; 10–10, 14; 13–24; 23–26, 52
AD741L	10–10
AD744	4–74; 17–13-14, 24
AD746	13–24
AD766	9–51; 13–24
AD767	13–24
AD779	3–96, 100-101, 105, 106
AD781	3–30
AD811	4–74; 6–49
AD827	13–93
AD829	8–199
AD834	6–39, 49
AD840	13–24
AD841	18–11, 13-14, 24
AD842	13–24
AD843	3–62, 64; 13–24; 18–14-15
AD844	13–24
AD846	13–24
AD845	3–29, 33; 13–24; 18–11, 13
AD847	13–24
AD848	13–24
AD849	13–24
AD1139	8–6-7, 9
AD1334	3–35
AD1408	8–71-73
AD1851	8–197, 199
AD1856	8–197-199; 13–24
AD1860	8–189-190, 192; 13–24; 24–100-101
AD1862J	8–192
AD1864	13–24
AD1879	20–3, 9-10, 16
AD2700	13–24
AD2701	13–24
AD2702	13–24
AD2710	13–24
AD2712	13–24
AD5245	10–13
AD5539	6–43-45
AD7111	8–80-81
AD7115	8–80-81
AD7118	8–80-81
AD7224	8–103-105; 13–24
AD7225	13–24
AD7226	8–107-112, 149-151; 13–24
AD7228	13–24
AD7237	8–35-37; 13–24
AD7240	23–50, 54
AD7243	8–35-37
AD7245	8–209, 211-212; 13–24
AD7247	13–24
AD7248	8–209, 211-212; 13–24
AD7341	5–19-21, 23-25

Page

AD7371 . . . . . 5–19-21, 24-25
AD7507 23–29
AD7510DI 12–3
AD7512 . . . . . 6–18; 8–118
AD7519 8–118
AD7520 . . . . . 8–71-72, 101; 23–27
AD7522 23–29
AD7528 . . . . . 3–8; 8–113-120, 127, 131, 161-166, 171, 193
AD7534 . . . . . 8–43-47, 127, 132-133, 135, 137, 139, 141, 144
AD7535 8–43-47
AD7537 . . . . . 8–127, 131, 153-158, 179, 193, 208-209, 211
AD7538 . . . . . 8–127, 132, 139, 177, 210-211
AD7541A 24–44
AD7542 . . . . . 8–121-123, 125, 201-204
AD7545 . . . . . 8–127, 130, 131, 133, 134-135, 137, 139, 141, 144, 177
AD7545A 8–35-37
AD7547 . . . 8–36, 127, 131, 140, 143-144, 147, 153-154, 158, 178-179, 193, 208-209, 211
AD7548 8–36
AD7549 . . . 8–36, 127, 131, 153-154, 158, 181-182, 184, 185-187, 193,
AD7560 8–111
AD7568 8–39-42
AD7569 . . . . . 2–5-6, 11-14, 15-20;13–24
AD7569 13–24
AD7572A 3–17-19
AD7574 3–21-25
AD7575 3–15
AD7578 3–11-12
AD7581 7–3-7
AD7582 3–11-12
AD7590DI 12–3
AD7592 . . . . . 3–9; 8–116-117
AD7628 8–171-172
AD7669 . . . . . 2–5-6, 11-14; 8–171, 173; 13–24
AD7672 3–27
AD7701 . . . . . 3–121; 5–7-8; 20–11-12
AD7703 . . . . . 3–121; 20–13
AD7710 20–13
AD7711 20–13
AD7712 20–13
AD7769 . . . . . 2–3-4, 6, 7-10, 11-14; 13–24
AD7770 13–24
AD7772 9–51
AD7773 8–31-32
AD7774 2–3-4
AD7775 8–31-32
AD7820 . . . . . 3–67; 8–171
AD7821 . . . . . 3–37, 67
AD7824 . . . . . 3–67, 69
AD7828 3–67
AD782X 3–67-69
AD7837 13–24
AD7840 13–24
AD7845 13–24
AD7846 13–24
AD7847 13–24
AD7848 13–24
AD7868 9–51
AD7870 24–97
AD7878 . . . . . 3–13-14; 9–51
AD7884 3–109
AD7885 3–115
AD9006 . . . . . 3–74; 24–95
AD9016 24–95
AD9020 24–95
AD9028 . . . . . 3–74; 24–95
AD9038 24–95
AD9060 . . . . . 3–74; 24–95
AD9101 8–292
AD9300 3–62
AD9500 . . . . . 14–3-8, 9-12, 14-15
AD9501 14–3-8
AD9617 . . . . . 3–74-75; 6–46; 13–171-174; 24–95
AD9618 13–171-174
AD9713B 8–292
AD9721 8–292-294
AD9955 8–292
AD22050 17–10-12
ADC908 8–97
ADC-910 . . . . . 8–249, 252
ADC-912A 18–23
AD OP-07 . . . . . 8–130, 134, 144; 13–24

Page

AD OP-27 . . . . . 8–135, 144, 182, 184; 13–24
AD OP-37 13–24
ADG201A 3–69; 8–208; 12–5
ADG202A 12–5
ADG221 . . . . . 12–5, 6
ADG222 . . . . . 12–5, 6
ADG506A 12–11-12, 21, 24-26
ADG507A 12–11, 21, 22, 26
ADG508A 12–11, 21, 22-26
ADG509A 12–11, 21, 22
ADG526A 12–11, 21 22, 24-26
ADG527A 12–11, 21, 26
ADG528A 12–11, 21, 22-26
ADG529A 12–11, 21 22
ADSP-28msp02 20–14-15
ADSP-1010 9–24-25
ADSP-2100 . . . . 3–17; 9–27-28, 32-34, 36, 38-39, 48, 49, 54, 58;24–79
ADSP-2100A 3–14; 9–39, 54, 58
ADSP-2101 . . . . . 3–106, 109; 9–8-10, 12, 14-16, 39, 54, 58, 59, 61-64, 65-76, 77-83, 85-92; 20–15
ADSP-2102 . . . . . 9–39, 50-51
ADSP-2105 . . . . . 9–54, 58, 59; 20–15
ADSP-2111 20–15
ADSP-21010 . . . . . 9–93, 101
ADSP-21020 9–93-101
ADV453 8–269
ADV471 . . . . . 8–254, 257, 259, 265, 269, 274, 276-278, 282
ADV476 . . . . . 8–257, 269, 277
ADV478 . . . . . 8–254, 257, 259, 269, 274, 276-279, 282
BUF-03 13–67-69
CMP-01 . . . . . 5–15; 8–213, 216-217, 221-222, 239, 242, 246; 18–18, 21; 24–36
CMP-404 13–55
COMDAC 86/89 . . . . . 8–213, 216, 221-222, 238, 246
COMDAC 88/89 . . . . . 8–239, 242, 244
DAC-08 . . . . . 8–8, 33, 49, 51-61, 63-70, 75-76, 78; 13–68
DAC-20 . . . . . 8–33, 77-78
DAC-86 . . . . . 8–213, 216-217, 221-223, 225, 238
DAC-88 . . . . . 8–239, 241, 242-243, 247, 248
DAC-89 . . . . . 8–213, 216-217, 221-222, 225, 238, 239, 241, 242, 248
DAC-8221 8–83
DAC-8228 18–20
DAC-8408 13–55
DAC-8426 18–22
DAC-8800 . . . . . 8–85-97, 99, 101-102
DAC-8821 13–91
DAC-8840 8–100-101
DAC-8841 8–100-101
DMX-88 . . . . . 8–239-240, 242, 246
DSP56000 3–14
ICM7208 23–9
IPA1764 16–14
LM-10 8–104-105
MAT-01 . . . . . 13–52, 53, 89; 15–4
MAT-02 13–51
MAT-04 . . . . . 13–61, 111
MUX-08 . . . . . 12–9; 13–68, 107; 18–23
MUX-16 . . . . . 12–9; 18–21
MUX-24 12–9
MUX-28 12–9
MUX-88 . . . . . 8–213, 216-217, 221, 239, 242, 246; 12–9
OP-01 8–251
OP-02 8–49
OP-06 13–71-74
OP-07 . . . . . 8–49, 130, 131, 134; 24–37
OP-10 13–56
OP-10A 13–62
OP-14 8–76
OP-15 . . . . . 12–17; 24–36
OP-16 . . . . . 8–213, 216-217, 221-222, 239, 242, 244, 246, 247; 12–17; 13–52, 69
OP-17G 8–77-78
OP-20 15–4-5
OP-22 . . . . . 13–51, 52, 58, 77, 81, 87, 89
OP-27 . . . . . 4–55; 13–35-37, 57, 64; 8–143, 145
OP-37 . . . . . 13–55, 58
OP-41 . . . . . 4–3; 13–50, 53, 56, 107; 13–53
OP-42 . . . . . 8–95; 13–107, 109-113, 131-135, 162-168; 18–21, 28
OP-43 13–54
OP-44 . . . . . 13–107, 113
OP-50 . . . . . 13–50-51, 57, 63

Page

OP-61 . . . . . 8–83; 13–91, 166-170
OP-64 13–125-129
OP-77 . . . . . 8–5, 49; 13–49-50, 53, 56, 59-65; 13–61
OP-90 . . . . . 13–111; 15–3-4
OP-177 8–97
OP-215 . . . . . 4–13, 36, 39; 13–85
OP-220 15–6
OP-221 . . . . . 8–75-76; 13–83
OP-227 . . . . . 13–64, 79
OP-260 13–143-148
OP-292 8–32
OP-400 13–137-141
OP-420 . . . . . 13–55, 59
OP-421 13–65
OP-470 . . . . . 13–43, 45, 149-154
OSC1758 . . . . . 16–5-6, 14
PM-139 18–23
PM-311 13–52
PM-7545 8–250-251
REF-01 . . . . . 8–75-76, 239, 250-251; 13–62, 85, 111
REF-02 . . . . . 8–213, 216-217, 221, 239, 242, 246; 18–20
REF-02A 13–75, 87
REF-43 15–4-6
RPT-82 . . . . . 5–9, 11-12, 15
RPT-83 . . . . . 5–9, 11-14, 15-16
RPT-86 5–17
RPT-87 5–17
SMP-04 . . . . . 18–19, 22-26, 27-28
SMP-08 . . . . . 18–19-21, 24-26, 27-28
SMP-11FY 18–17-18
SMP-18 . . . . . 18–19, 24, 27-28
SMP-81 . . . . . 8–213, 216-217, 221-222, 239, 242, 246
SSM-2013 4–13
SSM-2014 . . . . . 4–47, 52; 8–94
SSM-2015 . . . . . 4–3, 5, 9, 11, 23, 35-36
SSM-2016 4–9
SSM-2110 . . . . . 4–13, 47-48, 50
SSM-2120 . . . . . 4–29-30, 33, 37, 42, 44
SSM-2122 4–39-40
SSM-2131 . . . . . 4–52, 59
SSM-2134 . . . . . 4–7, 13, 19, 21, 29, 32, 34, 37, 39, 52
SSM-2210 . . . . . 4–50; 13–35-37
SSM-2220 4–55-56
SSM-2402 . . . . . 4–19, 31
SW-01 18–18
TSL1612 16–15
450J 23–32
451J 23–32
452J 23–32
453J 23–32
454J 23–32
456J 23–32
458J 23–32
460J 23–32
5S70 16–11
5S72 16–11

A

AC induction motor 16–34-35
ADC 3–3-128
  AC linearity plots 24–96-97
  aliasing 24–86-87
  aperture delay time 24–97
  aperture jitter 24–97-98
  asynchronous clock interfacing 3–13-14
  data acquisition 3–30-31
  decimation 24–88-89
  digital signal processing 3–13, 17, 111; 9–3-18
  discrete time sampling 24–85-86
  dynamic performance 24–94-98
  effective bits 24–90-92, 95
  evaluation 8–249
    code width 8–250
    cross-plot generator 8–249-250
    dynamic testing 8–250
    missing codes 8–250
    system noise 8–251-252
  fast transient capture 3–37
  fast transient detection 3–42
  finite amplitude resolution 24–85-86
  flash,
    aperture-jitter 3–90-92
    buffer 3–72
    characteristics 3–71-74
    composite video, DAC 3–93
    digital signal processing testing 3–79-84
    errors 3–73-74, 92-93
    FFT 3–79-80
      two-tone intermodulation 3–86
      verification 3–83
    histogram 3–86-87
    input drive amplifier 3–75
    noise-power ratio,
      digital signal processing 3–88-89
      test 3–87-89
    output 3–71-72
    S/N ratio 3–74, 79, 85
    sampling 3–80-81
      Blackman-Harris window 3–82
      Hanning window 3–81
      weighting 3–81
    signal coupling 3–72
    sine-wave curve fitting 3–83-84
    sparkle codes 3–72-74
    testing 3–71-93
    THD 3–85-86
    track-and-hold amplifier 3–76-78
    transient-response testing 3–89-93
  full-linear bandwidth 24–96
  full-power bandwidth 24–96
  high resolution, noise 3–18
  input level shifting circuits 2–8
  input multiplexing 3–24
  intermodulation distortion 24–96
  memory mapped input 3–21
  microprocessor interface 3–6, 21
  monitoring voltages 3–68-69
  multiple synchronous sampling 3–35
  noise 3–27-28
  noise-power ratio testing 3–53-54
  Nyquist's criteria 24–86
  oversampling 3–38; 24–88-89
  overvoltage recovery 24–98
  power strobing 24–51
  quantization 24–90-92
    finite amplitude resolution 24–90
  RAM mode 3–22-23
  range switching 3–23-24
  ROM mode 3–21-22
  S/N plus D 3–109, 115
  sample-and-hold amplifier 18–11
    harmonic amplitudes 3–32
  sampling 3–105
    aliasing 3–104
    disk drive application 3–15
    IMD 3–98
    signal/noise 3–32
    THD 3–98
  settling time 24–98
  sigma-delta 3–121
  signal dynamic range, resolution 24–92-94
  signal/noise 24–90-92, 95
    analysis 3–53
    digital measurement 3–51-52
    measurement 3–50-51
  slow memory mode 3–23
  spurious free dynamic range 24–95
  static transfer 24–93
  subranging 3–59-65
  super-Nyquist 24–89-90
  switching gain-range 3–9
  total harmonic distortion 24–95
  transient playback 3–44-48
  transient recorder 3–38-42, 43-44
  transient response time 24–98
  undersampling 24–89-90
    automatic gain control 3–97
    calculated noise flux 3–101
    digital radio 3–96, 103-107
    FFT 3–100
    full power bandwidth 3–100
    maximum headroom 3–97, 105
    Nyquist frequencies 3–96
    S/N plus D 3–98
    signal/noise 3–105
      aperture jitter 3–98-99
      IMD 3–99
      internal jitter 3–99
      THD 3–99
      theoretical limits 3–99
    SSB filter 3–106
  video,
    aperture jitter 3–57-58
    aperture-time 3–54-56
    bandwidth clipping 3–55-56
    burst mode 3–55
    differential nonlinearity 3–54
  voltage-to-frequency converter 23–5-12
Amplifier:
  audio 4–37
    performance 4–38
  automatic gain control 4–13
    performance 4–14
  balanced,
    audio 4–5-6
    performance 4–6
    true differential input 4–5-6
  commutating 6–27
  current feedback, frequency response 4–70-71
  inverting 2–18
  low noise 4–9, 11
    differential 4–9
    performance 4–10, 12
  power 4–62, 64-67
    characteristics 4–57-58
    circuit design 4–59-61, 61-64
    feedback 4–58-59
    performance 4–67-69
  programmable gain 8–118
  sign-bit 6–18
  single channel, components 4–72
  summing 4–3-4
    audio 4–3-4
    unbalanced virtual ground 4–7
  unbalanced virtual ground,
    audio 4–7
    low noise 4–7
  voltage-controlled 4–29-30
    dual channel 4–39

performance . . . 4–40
voltage-controlled, See also VCA
Analog-to-digital converter. See: ADC
Analog input/output . . . 2-3-20; 8–181-183
performance . . . 8–184
software . . . 8–183-184
Antialiasing filter:
design . . . 3–33
Nyquist limit . . . 3–32
oversampling . . . 3–38
Applesoft Basic program . . . 3–7-8
Audio compressor/limiter:
characteristics . . . 4–47
compressor control . . . 4–50-51
performance . . . 4–53
response . . . 4–47-48
variable time integrator . . . 4–48-50
voltage control amplifier . . . 4–52
Audio control:
audio mixing . . . 4–21
performance . . . 4–22
Audio matrix:
sum and difference . . . 4–35
performance . . . 4–36
Audio switch:
dual . . . 4–15-16, 19
high performance . . . 4–15
mute circuits . . . 4–19
performance . . . 4–17, 20
dual channel . . . 4–32
control interface . . . 4–32
output summing amplifier . . . 4–32
performance . . . 4–34
unbalanced . . . 4–31
Automatic line build-out. See: ALBO

B

Balun-coupled output . . . 6–48
Blackman-Harris window . . . 3–82
Bridge transducer, oscillators . . . 6–23-24
Buffer, simple, noninverting . . . 2–18

C

Capacitance:
noise . . . 24–61-63
PC board . . . 24–49
stray . . . 24–48, 60-62, 68
Faraday shield . . . 24–61-63
Capacitive noise . . . 24–12-13, 15-17, 61-63
Capacitor:
dielectric absorption . . . 24–64-65
errors . . . 24–21-23
inductance . . . 24–63-64
leakage . . . 24–63
parasitic effects . . . 24–63-65
series/loss resistance . . . 24–63
CCD:
linear array,
ADC digitizer . . . 3–59
circuits . . . 3–62
digitized signal balance . . . 3–64
odd/even pixel multiplexing . . . 3–64
output . . . 3–63
output saturation voltage . . . 3–64
prescan, digitizing . . . 3–63
overview . . . 3–59-60
Charge coupled device. See: CCD
Circuit resistance . . . 24–49
Clock-recovery . . . 5–14-16
CMOS analog switch . . . 12–3-4
Codec . . . 8–219
companding DAC . . . 8–227
four-channel . . . 8–221-226
operation and design . . . 8–214-217
shared . . . 8–213, 239
system tests . . . 8–217-218
COMDAC . . . 8–213-219, 221-246
Common-impedance noise . . . 24–12
Companding . . . 8–227-229
Companding DA converter. See: COMDAC
Compandor:
high performance,
limiter . . . 4–42
performance . . . 4–44
receiver expander . . . 4–43
transmitter compressor . . . 4–42
wireless audio . . . 4–41
Comparator, high speed . . . 24–45-46
Complementary-symmetry metal-oxide semiconductor. See: CMOS
Conductor:
Kelvin feedback . . . 24–56
resistance . . . 24–54
skin effect . . . 24–55
voltage drop in signal leads . . . 24–56
Constant linear velocity . . . 8–175-178
Converter:
flash . . . 3–71-93
RMS-to-DC . . . 6–42
Cross-plot generator, ADC testing . . . 8–249-250
Crosstalk . . . 12–13-20

D

DAC . . . 8–189-190
A/X function . . . 8–77-78
AC multiplying mode . . . 8–87
AC time domain specifications . . . 8–291-295
AC-coupled multiplication . . . 8–66-67
ADC,
input impedance . . . 8–20-21
buffer . . . 8–21
reducing effects . . . 3–51
ADC adjustable reference . . . 8–97
ADC log function . . . 8–82
algebraic digital computation . . . 8–69-70
amplifier,
programmable . . . 8–127
programmable gain . . . 8–118
analog input/output . . . 8–181
antilog . . . 8–79, 81
applications . . . 8–51-61, 63-70, 71-74, 85-97, 94
4–20mA current controller . . . 8–73-74
8-bit programmable square-wave oscillator . . . 8–72-73
algebraic manipulations . . . 8–72
digital audio . . . 8–197
microcomputer interfacing . . . 8–185
modems . . . 8–207-212
ohms-to-frequency converter . . . 8–71-72
programmable gain trimming/calibration . . . 8–73
programmable offset . . . 8–73
quotients of digital inputs . . . 8–72
UART . . . 8–201
attenuation control, programmable . . . 8–83
BCD-coded . . . 8–77
bias voltage, programmable . . . 8–41
bipolar references . . . 8–64
bridge transducer null . . . 8–68
bypassing . . . 8–17-18
coarse-fine control . . . 8–96-97
Codec system . . . 8–213, 239
common ground problems . . . 8–19-20
companding . . . 8–227-238, 239
accuracy . . . 8–236, 238
for codec . . . 8–227
gain tracking . . . 8–236-237
outputs . . . 8–232-236
power levels . . . 8–237
signal-to-distortion . . . 8–237-238
cross-plot generator . . . 8–249-250
current-steering . . . 8–39-40, 208, 210

decoupling 8–17-18, 20
deglitcher 8–211-212
  using sample-and-hold amplifier 24–101
differential A/D conversion 8–67-68
differential amplifier 8–19
digital audio circuits 8–189
digitally controlled VCA 8–94
direct digital synthesis 8–285-286
  selection 8–289
  specifications 8–289-291
dual, potentiometers 8–161-162
eight-channel 8–213, 240
failure analysis 8–46
fast pulsed operation 8–64
filter 8–153-154, 190-192
  control 8–153
  finite impulse response 8–191
  oversampling 8–191-192
filtering 24–102-104
four-quadrant multiplication 8–65-66, 70
four-quadrant ratiometric A/D conversion 8–68
full-scale level, programmable 8–41
function fitting sine wave synthesizer 8–115-117
gain control, programmable 8–83
glitch impulse area 24–99-100
grounding 8–18-19
harmonic distortion 24–100-101
high resolution,
  for instrumentation 8–3-4, 6-7
  integral nonlinearity 8–4-5
  measurement programming 8–5-6
  summation error, suppression 8–6
high temperature 8–44-45
highest speed four-quadrant multiplying 8–66
hyperbolic function generation 8–77
instrumentation amplifier 8–20
interface 8–121-126, 197, 198
  Bessel filter 8–198-199
  passband attenuation 8–198-199
  software 8–203-205
life test 8–45-46
linearity 8–37-38, 42
logic input, adjustable input logic threshold 8–54-56
low leakage 8–43
microcomputer interface 8–121-126
microprocessor 8–61
motor speed control 8–175, 177-178
multiplication and function fitting 8–162-163
multiplying 8–63, 79
noise sources 8–270-271
nulling problems 8–16-17
offset adjust, programmable 8–107-108
offset voltage trimming 8–95-96
output 8–202-203
  dual complementary 8–52-53
  high speed 8–54
  high voltage compliance current 8–51-52
output level shifting 2–9
output noise 8–88
panning 8–193-194
performance 24–99-104
PGA,
  dual DAC 8–139
    AC coupling output 8–143
    compensation 8–141-142
    dynamic gain error 8–142-143
    equations 8–139-140
    error 8–140-141
    fixed gain 8–145
    gain error 8–146-148
    noise and distortion 8–143
    small-signal bandwidth 8–142
    stability 8–141-142
    test results 8–143-145
  equations 8–127-128
    digital input code 8–138
    system gain error 8–136-137
  error,
    definition 8–128-130
    gain 8–131
    offset 8–132
    reduction 8–131-132
    sources, DC offset 8–130-131
  noise gain 8–133-134
  single DAC 8–132-133
  test results 8–134-135
power monitor 8–68-69
power splitter 8–194-195
  applications 8–196
power supply, requirements 8–60
problems 8–15-16
process current 8–75-76
  calibration 8–76
  output voltage 8–76
  scale modification 8–76
quantization 8–189
reconstruction filter 8–190
rectangular waveforms 8–163-166
reference amplifier compensation 8–64
reference input, multiplying capability 8–56-60
sampling 8–189
SCR 8–23
  analog ground and output 8–25-26
  digital input/output 8–24-25
  latch-up, causes 8–26-29
  prevention,
    Schottky diode 8–28-29
    silicon diode 8–28-29
self-programmable reference 8–109-110, 151
set-point controller 8–108-109
settling time 3–50; 24–41, 99
signal-generation 8–161
sine wave generation 8–149-151
sine wave oscillator, programmable 8–114-115
single supply 8–33
  operation 8–112, 119-120
sin(x)/x frequency rolloff effect 24–101
stand-alone operation 8–90-91
state-variable filter 8–113
static transfer 24–93-94
switched capacitor filter 24–102-104
test fixture 8–93-94
testing ADC 3–50
threshold testing 8–110-111
triangular waveforms 8–163-166
two-phase permanent magnet stepper motor 8–167
two-quadrant multiplication 8–64-65
two-wire interfaces 8–89-90
and variable resistors 8–101
VCA 8–101-102
video applications 8–253, 269-284
voltage, programmable 8–103-105
voltage doubling 8–36
voltage reference trimming 8–96
voltage-mode 8–172
voltage-output 8–173
voltage-switching 8–35-38, 40-41, 209-210
voltage/current source, programmable 8–117-118
V_{SS} generation 8–111-112
waveform generator 8–149-151
  programmable 8–118-119
wire bond strength 8–47
Data-acquisition system. See: DAS
Detector, mean-squared 6–41-42
Digital-to-analog converters. See: DAC
Digital-to-time converter 14–12
  analog signal settling time 14–15
  timing characteristics 14–13
Digital filter 9–3-12
  circular buffering 9–6-7
  decimation 20–7-8
  finite impulse response 9–4-5, 20-21
    audio filter 9–8-10
    CAD design 9–8

design . . . 9-7-10, 21-26
hardware . . . 9-23-25
linear phase . . . 9-11
Remez Exchange . . . 9-22-23
software . . . 9-25-26
windowing . . . 9-21-22
infinite impulse response . . . 9-11-13, 20-21, 37-39
cascaded biquad . . . 9-38-39
direct form . . . 9-37
multirate . . . 9-27-36
applications . . . 9-27
decimation . . . 9-27-32
interpolation . . . 9-32-36
theory . . . 9-19-20
time and frequency domain . . . 9-6
types . . . 9-20-21
Digital regenerator . . . 5-9-10
bipolar code . . . 5-10
integrated circuit . . . 5-9
Digital signal processing . . . 9-3-18
arithmetic ability . . . 9-65-70, 77-82, 93, 94-96
assembly language . . . 9-100-101
benchmarks . . . 9-101
control system model . . . 9-50-51
controller . . . 9-49, 51-58
monitor . . . 9-61, 61-64
notch filter . . . 9-58-59
Nth order . . . 9-54-58
PID . . . 9-51-54
data addressing . . . 9-82-84, 96-98
device interfacing . . . 9-74-76
digital filtering . . . 9-19-26, 37-39
digital radio receiver . . . 3-107
external memory . . . 9-100
FFT . . . 9-13-17
block floating point . . . 9-17
data scaling . . . 9-17
design . . . 9-15-16
hardware . . . 9-15
spectral leakage . . . 9-16-17
windowing . . . 9-16-17
filtering . . . 9-3-12
image data compression . . . 9-41-43, 45-48
in/out handling . . . 9-89-90
JPEG compression . . . 9-45-48
multirate filters . . . 9-27-36
performance . . . 9-90-92
processor selection . . . 9-65-76
program sequencing . . . 9-70-74, 84-89, 98-100
programmable . . . 9-47-48
sampled data system . . . 24-85
sampling . . . 3-18
system design . . . 9-49-59
Direct digital synthesis . . . 8-285
performance . . . 8-286-289
Disk drive . . . 2-5-6
motor speed control . . . 8-175
offset signal grounds . . . 2-5-6
DSP. See: Digital signal processing
DTC. See: Digital-to-time converter

E

Electromagnetic interference . . . 24-77-78
causes . . . 8-272-278
filtering components . . . 8-280-283
PCB design . . . 8-272
photoelectric . . . 24-78
radio frequency . . . 24-77-78
reduced . . . 8-269-284
regulations . . . 8-270
ENOB. See also Effective bits
Evaluation board . . . 3-109-121

F

Faraday shield . . . 24-61-63
Fast Fourier transform. See: FFT
Fast linear circuit, errors . . . 24-35
Fast memory device . . . 7-7
FCC certification . . . 8-278
FFT . . . 3-31
digital signal processing . . . 9-13-17
Filter:
analog versus digital . . . 9-19
antialiasing . . . 3-32-33; 24-87-88
Nyquist limit . . . 3-32
oversampling . . . 3-38
receive channel . . . 5-19-20
Bessel . . . 8-156, 158
Butterworth . . . 8-154, 158
amplitude response . . . 8-160
digital interface . . . 8-155-156, 157
gain bandwidth . . . 8-155
performance . . . 8-156-158
phase response . . . 8-160
Chebyshev . . . 8-156-157
nonsynchronous sampling . . . 5-23-24
program . . . 8-159
low pass,
cutoff frequency, digital control . . . 8-153-155
DAC . . . 8-153
programmable cutoff frequency . . . 8-153
noise spectrum . . . 8-160
reconstruction, transmit channel . . . 5-19-20
Sallen-Key, in modem . . . 6-22
state-variable . . . 8-113
switched capacitor,
communications applications . . . 5-19
cutoff frequency . . . 5-24
nonsynchronous sampling, design . . . 5-21-23
output spectrum . . . 5-20-21
sampling theory . . . 5-19-20
synchronous sampling . . . 5-21
two-channel,
dynamic . . . 4-29
performance . . . 4-30
universal active . . . 8-153
variable-Q voltage-tunable . . . 6-35-36
FIR. See: Finite impulse response
Flash converter . . . 3-71-93
Frequency-to-voltage. See: Voltage-to-frequency
Frequency dependent negative resistor. See: FDNR
Function generator . . . 2-15-20

G

Generator, voltage-reference . . . 3-76
Gilbert gain-cell . . . 6-40
Graphics image scanner . . . 3-61
output digitizing . . . 3-61-63
Grounding . . . 24-3-6, 7-9, 69-75
analog . . . 24-70-71
digital . . . 24-70-71
ground planes . . . 24-72
high-frequency . . . 24-3-6
high-speed . . . 24-7-9
Mecca ground . . . 24-70
noise . . . 24-69
signal routing . . . 24-74-75
star ground . . . 24-70
system grounds . . . 24-73-74
transmission lines . . . 24-72-73
Ground loops . . . 24-69-70
Ground noise . . . 24-69
Ground plane . . . 24-49-50, 72-73

H

Hanning window 3–81
HDD servo applications, power-down circuits 2–3
HF ground noise 24–50
High-speed comparator 24–45-46

I

IIR. See: Infinite impulse response
Image data compression:
  applications 9–41-43
  flexible bandwidth coding 9–42
  JPEG compression 9–42, 45-48
    applications 9–46-47
    entropy coding 9–46
    lossless coding 9–46
    programmable digital signal processors 9–47-48
    progressive coding 9–45-46
    sequential coding 9–45
    spectral selection 9–42
    successive approximation 9–42
    types 9–45
  picture transmission 9–42-43
    motion compensation 9–43
  progressive coding 9–42
  run-length coding 9–41
  sub-band coding 9–41-42
  techniques 9–41-42
  telephone transmission 9–42
    enabling factors 9–42
  transform coding 9–41
Inductance 24–50
  mutual 24–66-67
  ringing 24–67
  stray 24–48, 65
Inductor, parasitic effects 24–68
Input/output, channels 2–12
Input/output level shifting, signal swing levels 2–7-9
Input signal, demodulating 6–16
Instrumentation amplifier 10–3-10; 13–22-23
  applications 10–9-10
    analog delay 10–14-15
    current summer 10–13-14
    distortion measurement 10–15-16
    microphone preamp 10–11-13
    output buffer 10–17-19
    phono preamp 10–16-19
  components 10–4-6
  description 10–3-4
  as op amp 10–17
  specifications 10–6-9
Insulator:
  guard rings 24–56-57
  leakage 24–56
Interference noise 24–11-14
  driven guard 24–19-20
  guarding 24–19-20
  miscellaneous 24–14
  principles 24–11-12
  shielded guard 24–19-20
  shielding 24–15-19
  types 24–12-14
Intermodular distortion. See: IMD
Isolation amplifiers:
  applications 11–3-9
  current-loop 11–8-9
  transducers 11–3

J

Joint Photographic Experts Group. See: JPEG

L

Latch-up:
  causes 8–26-29
  in CMOS devices 8–23
  and SCR 8–23-24
Law of Intermediate Metals, temperature measurement 21–4
Linearization circuit 2–15-20
  principles 2–19
Linear variable differential transformer. See: LVDT
Linear voltage regulator:
  battery 15–5
  current limiting 15–5-6
  power transistor 15–3-4
  reference voltage 15–4
    op amp 15–4-5
Logarithmic function 6–4-5
LOGDAC 8–79, 81
  resolution 8–79-81
  transfer function 8–80-81
Logic 24–78-81
  fan-out 24–79
  noise 24–81
  sampling clock noise 24–80-81
  timing variations 24–79-80
LVDT 6–20-22

M

Magnetically coupled noise 24–13-14, 17-19
MASH. See: Multi-stage noise shaping
Mecca ground 24–70
Microcomputer:
  analog input/output 8–181
  DAC interfacing 8–185-187
  interface 8–121-126
Microphone:
  amplifier, low noise 4–9
  mixer, audio-activated 4–45
Microstepping:
  circuits 8–171
  control 8–170
  DAC errors 8–169-170
  fixed-frequency PWM drive 8–171-172
  load torque effects 8–170
  phase gain errors 8–168
  rate 8–173
  static position errors 8–168
Mixed signal system 24–53-84
  design 24–53-54
  logic 24–78-81
  prototyping 24–83-84
  socket limits 24–82-83
  SPICE modeling limits 24–82
Modem:
  commutating 6–15-17, 17
    dual-channel gain block 6–25-26
    precision timing 6–26
    synchronous detection 6–13
  DACs 8–207-212
  deglitcher 8–208
  loads 6–9
  performance 8–207-212
  signal processing, oscillator 6–22
Morgan compressor/limiter 4–47-54
Motor speed control:
  constant linear velocity 8–175
  feedback/DAC resistance ratio 8–178-179
Multiplexer:
  adjacent channel crosstalk 12–15-16
    measurement 12–17-19
    prediction 12–19-20
  bandwidth 12–21, 21-22
  circuit protection 12–11-12
  CMOS 12–10
    distortion 12–28
    R_{ON} modulation 12–27-28
  crosstalk 12–25-26
  distortion nomograph 12–28

dynamic crosstalk . . . 12–14-15
measurement . . . 12–16-17
equivalent circuit . . . 12–21
external protection . . . 12–12
grounding . . . 24–51
JFET . . . 12–9
OFF isolation . . . 12–22-25
overvoltage protection . . . 12–11
single-supply operation . . . 12–9-10
static crosstalk . . . 12–13-14, 16
uses . . . 12–21
Multiplier:
AC-output coupling . . . 6–47
analog . . . 6–29
balun-coupled output . . . 6–48
commutating amplifier . . . 6–27
controlled-cascode . . . 6–32
DC-coupled VCA applications . . . 6–43-46
development history . . . 6–30
4-quadrant . . . 6–39
Gilbert gain-cell . . . 6–40
mean-squared detector . . . 6–41-42
RMS-to-DC converter . . . 6–42
transformer-coupled output . . . 6–47
translinear circuit . . . 6–40
2-quadrant . . . 6–29-33
VCA . . . 6–49-52
circuit analysis . . . 6–51
feedback keyer . . . 6–53-55
video keyer . . . 6–52
video switch . . . 6–47
for voltage-tunable filter . . . 6–34-37
Multiplier-divider . . . 6–5-6
Multirate filter:
decimation . . . 9–27-32
structure . . . 9–27-32
interpolation . . . 9–32-36
algorithm . . . 9–33-34
structure . . . 9–32-36

N

Noise, crystal oscillators . . . 8–270-271
Noise-power ratio, testing . . . 3–53-54
Nyquist limit . . . 3–32; 24–86, 89

O

Offset adjust, programmable . . . 8–107-108
Operational integrator:
AC . . . 13–15-16
equivalent circuit . . . 13–11
errors . . . 13–11-14
feedback capacitor . . . 13–14-15
leakage resistance . . . 13–15
Op amp:
36 V low noise . . . 13–55
200 mA servo motor amplifier . . . 13–49
200 V low offset . . . 13–63
absolute value amplifier . . . 13–61
active feedback . . . 13–41-48
additional input . . . 13–23
amplifier with active output clipping . . . 13–65
analog multiplier/divider . . . 13–61
applications . . . 13–49-65, 67-69, 75, 77, 79, 81, 83, 85, 87, 89, 91, 93-94
bandgap reference . . . 13–89
bilateral current source . . . 13–50
Boyle model . . . 13–155-158
buffer amplifier . . . 13–67-69
common-mode rejection . . . 13–9-10
comparator . . . 13–71-74
noise . . . 13–73-74
performance . . . 13–71
response time . . . 13–71-72
compensation . . . 13–17-21, 41-48
current feedback macro-model . . . 13–171-174
current monitor . . . 13–60
current pump . . . 13–50
current regulator . . . 13–87
current sink . . . 13–64
current transmitter . . . 13–75
differential to single-ended conversion . . . 13–18-19
drift vs. time . . . 13–8
dual,
current feedback . . . 13–143-148
Spice macro-model . . . 13–143-148
dual programmable window comparator, 13–55
dual tracking voltage reference . . . 13–56
equivalent circuit . . . 13–3
feed-forward . . . 13–20
filter . . . 13–177-184
active . . . 13–177-178
allpass . . . 13–177
antialiasing . . . 13–180-182
bandpass . . . 13–177
bandstop . . . 13–177
Butterworth . . . 13–178
Chebyshev . . . 13–178
designs . . . 13–178-184
elliptical . . . 13–178
FDNR . . . 13–179-180
high Q notch . . . 13–56
highpass . . . 13–177
lowpass . . . 13–177
multiple feedback . . . 13–179
passive . . . 13–177
programmable state variable . . . 13–182
Sallen-Key . . . 13–179
seven-pole FDNR antialiasing . . . 13–183
specifications . . . 13–178
state variable . . . 13–179
wideband Sallen-Key . . . 13–184
free-running square-wave generator . . . 13–60
frequency stability . . . 13–19
full power response . . . 13–5-6
grounding errors . . . 13–21-22
headphone amplifier . . . 13–57
high speed . . . 13–115-123, 125-129, 131-135
bandwidth . . . 13–121-122
capacitance . . . 13–116-120
fast-settling . . . 13–131-135
noise gain . . . 13–122-123
settling time . . . 13–121-122
Spice macro-model . . . 13–125-129, 131-135
unity-gain stability . . . 13–120-121
high stability voltage reference . . . 13–56
high-sensitivity voltage comparator . . . 13–51
impedance transforming amplifier . . . 13–63
initial bias current . . . 13–7
initial difference current . . . 13–7
initial offset voltage . . . 13–7
input design . . . 13–175-176
input impedance . . . 13–8-9
input noise . . . 13–10
input offset voltage nulling . . . 8–49
instrumentation amplifier . . . 13–22-23
instrumentation amplifier (2-op amp) . . . 13–62
interconnections . . . 24–3-5
isolation amplifier . . . 13–54
JFET . . . 4–73; 13–107-114
applications . . . 13–107
design precautions . . . 13–112-113
errors . . . 13–113-114
logarithmic . . . 13–110-111
varying compensation . . . 13–112-113
latchup . . . 13–176
low noise AGC amplifier . . . 13–64
low power amplifier with squelch . . . 13–65
maximum common mode voltage . . . 13–9
maximum voltage between inputs . . . 13–9
microphone pre-amplifier . . . 13–58
micropower 1–23 V bandgap reference . . . 13–52

micropower 5 V regulator . . . 13–51
micropower instrumentation amplifier . . . 13–59
micropower Wien-bridge oscillator . . . 13–58
NAB tape head pre-amplifier . . . 13–58
negative supply decoupling . . . 13–18-19
noise . . . 13–25-34
  calculation . . . 13–39-40
  characteristics . . . 13–27-28
  density spectrum . . . 13–30-31
  inherent . . . 13–28-29
  interference . . . 13–28-29
  random . . . 13–29-30
  sources . . . 13–33
  voltage/current . . . 13–26
noise figure . . . 13–34
noise gain . . . 13–27, 32-33
nonlinear input capacitance . . . 4–73-74
open loop gain . . . 13–3-4
operational integrator . . . 13–11-16
overload recovery . . . 13–6
phase accuracy, active feedback . . . 13–41-48
phase inversion . . . 13–176
piecewise-linear amplifier . . . 13–59-60
piezoelectric transducer amplifier . . . 13–56
pink noise tangent . . . 13–32
positive supply decoupling . . . 13–19-20
programmable attenuation . . . 13–91
programmable gain . . . 13–91
programmable gain amplifier . . . 13–50
quad,
  low offset/low power . . . 13–137-141
  Spice macro-model . . . 13–137-141, 149-154
  very low noise . . . 13–149-154
ramp generator . . . 13–85
rated output voltage and current . . . 13–4
rectifier . . . 13–77, 79
RIAA phone pre-amplifier . . . 13–57
settling time . . . 13–95-103, 105-106; 24–40
signal gain . . . 13–27
slewing rate . . . 13–6
specifications . . . 13–3-10
Spice-compatible macro-model . . . 13–155-170
  Boyle . . . 13–155-158
  design . . . 13–155-159
standard process junction isolated FET . . . 4–73, 74
  design . . . 4–75
  performance . . . 4–75
  unbuffered bootstrap . . . 4–75-76
subtractor . . . 13–22
supply voltage sensitivity . . . 13–8
temperature drift . . . 13–7
terminals . . . 13–17
thermocouple amplifier with cold-junction
  compensation . . . 13–62
threshold detector/analyzer . . . 13–53
triangular wave generator . . . 13–52
types . . . 13–175
unity gain small signal response . . . 13–5
very low noise . . . 13–35-37
  applications . . . 13–35
  performance . . . 13–37
voltage-controlled . . . 13–93-94
wide range low current ammeter . . . 13–53
wide-dynamic-range light detector . . . 13–53
Wien bridge oscillator . . . 13–81, 83
Oscillator:
in modem . . . 6–22
Wien bridge . . . 6–22-23
Oscilloscope . . . 24–33, 36-37
Oversampling . . . 20–3-4
antialiasing . . . 3–38

P

Parasitic transistor action . . . 24–28
Passive circuit components . . . 24–21-26
PCB:
  design . . . 8–272-278
  electromagnetic interference . . . 8–272
PCM repeater . . . 5–9, 13, 17-18
  ALBO diodes . . . 5–17
Photoelectric interference . . . 24–78
Pixel processing unit:
  components . . . 8–258-260
  video animation . . . 8–261
    C computer program . . . 8–262-264
Pixel read mask register . . . 8–258
Polarity programmable peak detector . . . 18–17-18
Portable instruments, battery . . . 15–5
Potentiometer, errors . . . 24–25
  Power-down circuit . . . 2–3-4
  MOSFET performance . . . 2–4
Power-line transient noise . . . 24–14
Power supply:
  grounding . . . 24–8-9
  noise . . . 24–76
  spikes . . . 24–29-30
  switching-mode . . . 24–76-77
Preamplifier:
  low noise . . . 4–11, 55
    performance . . . 4–55-56
  microphone . . . 4–46
    performance . . . 4–45
  phono . . . 4–23-25
    moving coil . . . 4–23
    moving magnet . . . 4–23
    passive . . . 4–25
    passive multifilter . . . 4–24
      design . . . 4–26-27
      performance . . . 4–27
    performance . . . 4–26
    playback characteristics . . . 4–26
    uncomplicated passive, performance . . . 4–27
Precision ramp, op amp . . . 13–85
Printed-circuit board, errors . . . 24–25-26
Programmable delay generator:
  digital,
    description . . . 14–3-8
    enable . . . 14–6
    full-scale range setting . . . 14–4-5
    laser applications . . . 14–7
    layout . . . 14–9
    multichannel deskewing . . . 14–7
    pulse delay . . . 14–7
    pulse width modulation . . . 14–7-8
    timing . . . 14–5-6
    TTL signals . . . 14–6
  digital-to-time converter . . . 14–9-16
  oscillator . . . 14–11-16
  pulse generator . . . 14–10
  system timing . . . 14–9-16
Programmable gain amplifier. See: PGA
Protection for ICs . . . 24–27-28, 29-31
Prototyping . . . 24–83-84
Pulse attenuation, ALBO . . . 5–12
PWM drive, frequency modulation . . . 8–172-173
PWM motor driver . . . 8–174

R

Radio:
  digital,
    with ADCs . . . 3–95-101
    receiver design . . . 3–96, 103
    SSB filter . . . 3–106
    SSB receiver . . . 3–103
    undersampling . . . 3–98
Radio frequency interference . . . 24–77-78
RAM-DAC . . . 8–278
  load terminations . . . 8–254-256
  PCB layout . . . 8–267-268
  video applications . . . 8–254-256, 257-264, 267-268, 269-284

Reference, self-programmable 8–109-110
Resistance 24–54-60
electrostatic damage 24–57-58
parasitic effects 24–58-60
voltage variation 24–59-60
Resistor:
circuit wiring 24–49
errors 24–23-25
inductive 24–58
Johnson noise 24–60
log compensation 24–51-52
non-ideality 24–47
resistance 24–48
stability and matching 24–59
temperature varying 24–48
thermo-electric effects 24–58-59
Resolver-to-digital converter 16–3-6, 11-16, 45-47, 51
absolute position data,
RS-485 communication 16–46
serial-to-parallel 16–45-46
acceleration constant 16–8
angle measurement 16–3
applications 16–5, 11-16
bandwidth 16–7-8
as control transformer 16–13
dynamic modeling 16–21-32
external counters 16–13-14
with 15-volt supply 16–12-13
4096 line incremental encoder emulation 16–47
with geared system 16–15-16
high voltage interfacing 16–11-12
high voltage synchros 16–19-20
incremental encoder 16–14
with inductosyn 16–14, 16
integration 16–49-51
linear measurement 16–3
passive component selection 16–23-24
reference frequency 16–8-9
resolution 16–9
resolvers and synchros 16–18
Scott-T circuit 16–17-20
settling time 16–9
silicon tachogenerator 16–53-55
software 16–21-23
step response 16–26-29, 31-32
tracking 16–5, 9
transfer function 16–24-26
vector control 16–33-43
velocity scaling 16–9
RMS-to-DC converter:
applications 17–3, 13-14
function 17–3-4
selection 17–5
theory 17–6-13
voltage comparator, hysteresis 17–15-16
RS-232 interface device:
applications 5–3, 6-8
charge pump 5–3-6
latch-up immunity 5–4
noise immunity 5–5
overvoltage protection 5–4
voltage doubler 5–3-6
voltage inverter 5–3-6
RTD. See: Resistance temperature detector

S

Sample-and-hold amplifier:
ADC 3–29-30; 18–11
applications 18–8-10, 17-28
capacitors 18–15
CMOS switches 18–12-14
deglitching DAC 24–101
design 18–8, 11-15
dielectric absorption 18–7
drift and noise 18–7
errors 18–4-7
feedthrough 18–12-13, 14-15
FFT 3–31
operation 18–3-4
peak detector 18–17-18
Sample-and-hold circuit:
infinite 2–11-14
dual channel 2–13-14
single channel 2–12-13
Sampled data system 24–85-105
Schottky diode 3–89
Scott-T circuit 16–17
operation 16–19
SCR:
in CMOS devices 8–23
latch 24–27-28
latch-up, causes 8–26-29
principles 8–23-24
Sensor:
accuracy 19–3-6
application 19–7
noise 19–7-10
radio-frequency 19–9
reduction 19–8-10
shielding 19–8-9
temperature errors 19–3-6
thermometer circuit 19–7
worst-case nonlinearities 19–3-6
Set-point controller 9–108-109
Settling time:
definition 13–95
fast circuits 24–33-42
measurement 13–100-101
op amp 13–95-103, 105-106
pole-zero mismatch 13–102-103
testing 13–105-106
Sigma-delta ADC:
audio,
description 20–9-10
digital filter 20–10
specifications 20–10
higher order modulator loops 20–9
idling pattern 20–8-9
low frequency 20–11-13
digital filter 20–12-13
specifications 20–12
multi-stage noise shaping 20–16
tonal considerations 20–8-9
Sigma-delta codec 20–14-15
Sigma-delta converter:
architecture 20–3
digital filtering and decimation 20–7-8
multi-bit 20–16-17
oversampling 20–3-4
Sigma-delta DAC 20–13-14
Sigma-delta modulator, quantization noise shaping 20–4-7
Signal, input, demodulating 6–16
Signal-noise ratio video A/D 3–51
Signal return current 24–69
Signal routing 24–69-75
Signal sampling, theory 5–19-20
Silicon controlled rectifier. See: SCR
Silicon tachogenerator 16–53-54
performance comparison 16–54-55
Sine-wave generator 8–122-123, 247-248
function fitting 8–115-117
Sine-wave oscillator, programmable 8–114-115
Single-sideband. See: SSB
Spice:
macro-model for op amp 13–125-129, 131-135, 137-141, 143-148, 149-154
modeling, mixed signal systems 24–82
op amp, performance 13–128-129, 131-135, 137-141, 143-148, 149-154
Spice-compatible macro-model:
gain stage 13–159-165
simulation accuracy 13–165-170

Spreadsheet, op amp noise calculation . . . 13–39-40
Square roots . . . 6–8-9
Star ground . . . 24–70
Static electricity . . . 24–27
Step-function generator . . . 24–33-34, 38-39
Stepper motor:
  microstepping . . . 8–168
  permanent magnet . . . 8–167-168
Stray capacitance . . . 24–48
Stray inductance . . . 24–48
Summing amplifier . . . 4–3
  performance . . . 4–4
  unbalanced virtual ground . . . 4–7
Supply decoupling . . . 24–47
Switch . . . 12–3-4
  charge injection . . . 12–7
  CMOS, R_{ON} modulation . . . 12–27-28
  digital input trigger levels . . . 12–6
  latch timing . . . 12–6-7
  leakage currents . . . 12–6
  on-resistance . . . 12–5-6
  power supply current . . . 12–7-8
Switched capacitor, errors . . . 24–103-104
Switching-mode power supply . . . 24–31, 76-77
Synchronous detection, resistance measurement . . . 6–11

T

Temperature:
  integrated circuit temperature transducer . . . 21–3
  Law of Intermediate Metals . . . 21–4
  measurement . . . 21–3, 8
  resistance temperature detector . . . 21–3
  thermistor . . . 21–3
  thermocouple . . . 21–3
Temperature transducer, remote sensing . . . 19–7
Thermocouple:
  characteristics . . . 21–3
  circuit trimming . . . 21–6
  cold junction errors . . . 21–5
  direct mean temperature . . . 21–7
  error calculations . . . 21–5-6
  Fahrenheit output . . . 21–6
  gain offsetting . . . 21–8
  ice point referencing . . . 21–3-4
  loop . . . 21–3
  measurement . . . 21–4
  noise minimizing . . . 21–5
  optimizing . . . 21–5
  Seebeck coefficient . . . 21–4
  temperature multiplexing . . . 21–7
  types . . . 21–4
Threshold testing . . . 8–110-111
Total harmonic distortion. See: THD
Transducer:
  pressure . . . 11–7-8
  resistance temperature detector . . . 11–6-7
  strain gage . . . 11–7-8
  thermocouple . . . 11–3-6
Transformer-coupled output . . . 6–47
Transient recorder, ADC . . . 3–38-42
Transient voltage . . . 24–30-31
Transmission line . . . 24–72-73
Transmitter chip . . . 22–3-10
  RTDs . . . 22–9
  strain gages . . . 22–9-10
TrimDAC:
  multichannel DAC . . . 8–85, 99
    applications . . . 8–101-102
    gain addition . . . 8–100-101
    and variable resistors . . . 8–101
  voltage adjustment . . . 8–99-100
Trimming . . . 24–43-44

U

UART:
  DAC driver . . . 8–201
  serial data format . . . 8–201-202
  serial out generator . . . 8–205-206

V

VCA:
  DC-coupled, in AD834 . . . 6–43
  DC-to-60 MHz, passive level shifting . . . 6–43-45
  DC-to-480 MHz, active level shifting . . . 6–45-46
  design . . . 6–33-37
  dual-channel . . . 6–31
  feedback keyer . . . 6–53-55
  ultralow-distortion . . . 6–33
  video applications . . . 6–49-52
Vector rotation:
  device description . . . 16–41-43
  induction motors . . . 16–33, 35-37
  mathematics . . . 16–38-40
  synchronous motors . . . 16–33
VGA, design . . . 8–265
Video graphics system, RAM-DACs . . . 8–269-284
Video standards . . . 8–253
Video switch . . . 6–47
Voltage, programmable . . . 8–103-105
Voltage-controlled amplifier. See: VCA
Voltage-to-current converter . . . 22–3-10
Voltage-to-frequency converter . . . 23–3-4, 33-60
  16-bit analog-to-digital . . . 23–8-9
  ADC . . . 23–5-12
  applications . . . 23–3-4
  device description . . . 23–33-34
  frequency-to-voltage . . . 23–61, 65
    application . . . 23–63
    design . . . 23–62-63
    operation . . . 23–61-62
    ripple . . . 23–64, 66
  gating . . . 23–9-10
  instrumentation applications . . . 23–19-24
    4–20mA loop operation . . . 23–23-24
    analog-to-digital conversion . . . 23–19
    bipolar inputs . . . 23–20-21
    signal isolation . . . 23–19-20
    transducer interface . . . 23–21-23
  microprocessor interfacing . . . 23–28-30
  microprocessor-based ADC . . . 23–10-12
  operation . . . 23–13-14, 34
    calibration . . . 23–15
    modes . . . 23–14-19, 35-40
    negative input current . . . 23–16
    negative input voltage . . . 23–15-16
    output interfacing . . . 23–16-18
    positive input voltage . . . 23–14-15
    precautions . . . 23–18
    temperature stability . . . 23–18
  period timing . . . 23–7-8
  pulse counting . . . 23–5-6
  special applications . . . 23–26-28
    analog division . . . 23–27
    long-term analog integration . . . 23–26-27
    temperature-to-frequency conversion . . . 23–27-28
  specifications . . . 23–31-32
  telemetry and communications . . . 23–24-26
    frequency-shift modulators . . . 23–24
    phase-locked loops . . . 23–24-25
    waveform generation . . . 23–25-26
  timing component . . . 23–40-59
    capacitor . . . 23–40-41
    clock . . . 23–53-54
    coupling . . . 23–43-45
    4–20 mA current loop . . . 23–48-50
    frequency doubling . . . 23–56-58
    function generator . . . 23–52-53

high-input impedance bipolar 23–56
logic 23–41
oscillator 23–56
output interfacing 23–41
phantom power 23–41-42
phase-locked loop 23–50-52
photodiode 23–47-48
photosensitive transducer 23–47
power supply 23–45-46
references 23–43
resistive transducer interface 23–46-47
signal multiplexing 23–43
switched capacitor filter 23–50
temperature stability 23–40
temperature-to-frequency 23–48
thermocouple 23–46
12-bit converter 23–54
voltage/current monitor 23–54-55
Voltage/current source, programmable 8–117-118
V_{SS} generation 8–111

W

Waveform generator, programmable 8–118-119
Wien bridge oscillator 6–2-23; 13–81, 83

Application Note (AN) Number Index

		Page
AB-3	Digital Nulling of Precision Op Amps	8–49
AB-6	Single Supply Operation of DAC-08 and DAC-20	8–33
AB-107	Two-Wire, 4–20 mA Current Transmitter	13–75
AB-108	A Micropower Single-Supply Precision Rectifier	13–77
AB-109	High Speed Precision Rectifier	13–79
AB-111	Single-Supply Operation Wien Bridge Oscillator	13–81
AB-112	Single Resistor Controls Wien Bridge Oscillator Frequency	13–83
AB-113	Precision Ramp Generator	13–85
AB-114	Precision Current Regulator	13–87
AB-115	Micropower 1.23 V Bandgap Reference	13–89
AN-17	DAC-08 Applications Collection	8–51
AN-19	Differential and Multiplying D/A Converter Applications	8–63
AN-21	4–20 mA Digital to Process Current Transmitter	8–75
AN-23	D/A Converter Generates Hyperbolic Functions	8–77
AN-25	Applying the OP-06 Op Amp As a High Precision Comparator	13–71
AN-27	Polarity Programmable Peak Detector	18–17
AN-32	Single Supply Operation of JFET Multiplexers	12–9
AN-35	Understanding Crosstalk in Analog Multiplexers	12–13
AN-37	Eight-Channel Codec Applications	8–213
AN-38	Four-Channel Shared Codec	8–221
AN-39	Companding D/A Converter	8–227
AN-40	A Buffer Amplifier Applications Collection	13–67
AN-41	Improved Shared-Channel Codec Design with Analog Devices' Companding DACs	8–239
AN-42	A 1 kHz, 0 dBm0 Standard Signal Generator	8–247
AN-48	Designing Digital Repeaters with ICs	5–9
AN-50	A Variable-Frequency, Clock-Recovery Circuit Using the RPT-82 or RPT-83	5–15
AN-101	Cross-Plot Generator Allows Quick A/D Converter Evaluation	8–249
AN-102	Very Low Noise Operational Amplifier	13–35
AN-106	A Collection of Amp Applications	13–49
AN-107	Active Feedback Improves Amplifier Phase Accuracy	13–41
AN-108	JFET-Input Amps Are Unrivaled for Speed and Accuracy	13–107
AN-109	Understanding and Preventing Latch-Up in CMOS DACs	8–23
AN-110	OP-64 Advanced SPICE Macro-Model	13–125
AN-111	A Balanced Summing Amplifier	4–3
AN-112	A Balanced Input High Level Amplifier	4–5
AN-113	An Unbalanced, Virtual Ground Summing Amplifier	4–7
AN-114	A High Performance Transformer-Coupled Microphone Preamplifier	4–9
AN-115	Balanced Low Noise Microphone Preamplifier Design	4–11
AN-116	AGC Amplifier Design with Adjustable Attack and Release Control	4–13
AN-117	OP-42 Advanced SPICE Macro-Model	13–131
AN-119	Improved T148 (CEPT) PCM Repeater for #22AWG (0.7 mm) Unshielded Twisted-Pair Wire	5–17
AN-120	OP-400 SPICE Macro-Model	13–137
AN-121	High Performance Stereo Routing Switcher	4–15
AN-122	A Balanced Mute Circuit for Audio Mixing Consoles	4–19
AN-123	A Constant Power "Pan" Control Circuit for Microphone Audio Mixing	4–21
AN-124	Three High Accuracy RIAA/IEC MC and MM Phono Preamplifiers	4–23
AN-125	A Two-Channel Dynamic Filter Noise Reduction System	4–29
AN-126	OP-260 Advanced SPICE Macro-Model	13–143
AN-127	An Unbalanced Mute Circuit for Audio Mixing Channels	4–31
AN-128	A Two-Channel Noise Gate	4–33
AN-129	A Precision Sum and Difference (Audio Matrix) Circuit	4–35
AN-130	A Two-Band Audio Compressor/Limiter	4–37
AN-131	A Two-Channel VCA Level (Volume) Control Circuit	4–39
AN-132	OP-470 SPICE Macro-Model	13–149
AN-133	A High Performance Compandor for Wireless Audio Systems	4–41
AN-134	An Automatic Microphone Mixer	4–45
AN-135	The Morgan Compressor/Limiter	4–47
AN-136	An Ultralow Noise Preamplifier	4–55
AN-137	A Digitally Programmable Gain and Attenuation Amplifier Design	8–83
AN-138	SPICE-Compatible Op Amp Macro-Models	13–155
AN-142	Voltage Adjustment Applications of the DAC-8800 TrimDAC, an Octal, 8-Bit D/A Converter	8–85
AN-202	An I.C. Amplifier User's Guide to Decoupling, Grounding, and Making Things Go Right for a Change	13–17
AN-204	Applications of the SMP-04 and the SMP-08/SMP-18, Quad and Octal Sample-and-Hold Amplifiers	18–19
AN-205	Video Formats and Required Load Terminations	8–253
AN-206	Analog Panning Circuit Provides Almost Constant Output Power	8–193
AN-207	Interfacing Two 16-Bit AD1856 (AD1851) Audio DACs with the Philips SAA7220 Digital Filter	8–197
AN-208	Understanding LOGDACs	8–79
AN-209	8th Order Programmable Low-Pass Analog Filter Using Dual 12-Bit DACs	8–153
AN-210	Adding Additional Input Channels to the AD7773/ AD7775	8–31
AN-211	The Alexander Current-Feedback Audio Power Amplifier	4–57
AN-212	Using the AD834 in DC to 500 MHz Applications: RMS-to-DC Conversion, Voltage-Controlled Amplifiers, and Video Switches,	6–39
AN-214	Ground Rules for High Speed Circuits	24–7
AN-215A	Designer's Guide to Flash-ADC Testing—Part 1: Flash ADCs Provide the Basis for High Speed Conversion	3–71
AN-215B	Designer's Guide to Flash-ADC Testing—Part 2: DSP Test Techniques Keep Flash-ADCs in Check	3–79
AN-215C	Designer's Guide to Flash-ADC Testing—Part 3: Measure Flash-ADC Performance for Trouble-Free Operation	3–85
AN-216	Video VCAs and Keyers Using AD834 and AD811	6–49
AN-217	Audio Applications of the ADSP Family (IIR Filters)	9–37
AN-218	DSP Multirate Filters	9–27
AN-219	Electronic Adjustment Made Easy with the TrimDAC	8–99
AN-221	Power-Down Circuit Cuts Power to AD7769 and AD7774	2–3
AN-222	AD7569/AD7669 Operation with Offset Signal Grounds for Disk Drive Applications	2–5
AN-223	AD7575 Operation with an Offset Signal Ground for Disk Drive Applications	3–15
AN-224	Input/Output Level Shifting with the AD7769	2–7
AN-225	12-Bit Voltage-Output DACs for Single Supply 5 V and 12 V Systems	8–35
AN-227	Digital Control System Design with the ADSP-2100 Family	9–49
AN-228	Evaluation Board for the AD7884 16-Bit A/D Converter	3–109
AN-229	Evaluation Board for the AD7885 16-Bit A/D Converter	3–115
AN-230	Circuit Applications of the AD2S90 Resolver-to-Digital Converter	16–45
AN-232	Bootstrapped IC Substrate Lowers Distortion in JFET Op Amps	4–73
AN-233	Considerations for Selecting a DSP Processor (ADSP-2101 vs. TMS320C50)	9–77
AN-234	Digital Resolver Integration	16–49
AN-235	Considerations for Selecting a DSP Processor (ADSP-21020/ADSP-21010 vs. TMS320C30)	9–93
AN-236	Speed Measurement and Control with the AD2S83, a "Silicon Tachogenerator"	16–53

		Page
AN-237	Choosing DACs for Direct Digital Synthesis	8–285
AN-238	AD2xx Family for RS-232 Communications	5–3
AN-240	Considerations for Selecting a DSP Processor ADSP-2101 vs. WE DSP16A	9–65
AN-244	A User's Guide to IC Instrumentation Amplifiers	10–3
AN-245	Instrumentation Amplifiers Solve Unusual Design Problems	10–11
AN-246	Isolation Barrier Precisely Passes Transducer Signals	11–3
AN-248	ADG201A/202A and ADG221/222 Performance with Reduced Power Supplies	12–5
AN-249	Overvoltage Protection for the ADG5xxA Multiplexer Series	12–11
AN-250	Bandwidth, OFF Isolation and Crosstalk Performance of the ADG5xxA Multiplexer Series	12–21
AN-251	R_{ON} Modulation in CMOS Switches and Multiplexers; What It Is and How to Predict Its Effect on Signal Distortion	12–27
AN-252	Using the 2S80 Series Resolver-to-Digital Converters with Synchros: Solid-State Scott-T Circuit	16–17
AN-253	Find Op Amp Noise with Spreadsheet	13–39
AN-254	Add Programmable Gain, Attenuation	13–91
AN-255	Voltage-Controlled Amp Covers 55 dB Range	13–93
AN-256	Accurately Testing Op Amp Settling Times	13–105
AN-257	Careful Design Tames High Speed Op Amps	13–115
AN-259	AD9617/AD9618 Current Feedback Amplifier Macro-Models	13–171
AN-260	Using Digitally Programmable Delay Generators	14–3
AN-261	Programmable-Delay ICs Control System Timing	14–9
AN-262	A Low Noise, Low Drop-Out Regulator for Portable Equipment	15–3
AN-263	Resolver-to-Digital Conversion—A Simple and Cost Effective Alternative to Optical Shaft Encoders	16–3
AN-264	Dynamic Characteristics of Tracking Converters	16–7
AN-265	Circuit Applications of the AD2S81A and AD2S80A Resolver-to-Digital Converters	16–11
AN-266	Passive Component Selection and Dynamic Modeling for the 2S80 Series Resolver-to-Digital Converters	16–21
AN-267	Vector Control Using a Single Vector Rotation Semiconductor for Induction and Permanent Magnet Motors	16–33
AN-268	RMS-to-DC Converters Ease Measurement Tasks	17–3
AN-269	Extend RMS-to-DC Converter's Range	17–15
AN-270	Applying IC Sample-Hold Amplifiers	18–3
AN-271	Build Precise S/H Amps for Fast 12-Bit ADCs	18–11
AN-272	Accuracies of the AD590	19–3
AN-273	Use of the AD590 Temperature Transducer in a Remote Sensing Application	19–7
AN-274	Taking the Uncertainty Out of Thermocouple Temperature Measurement (With the AD594/AD595)	21–3
AN-275	Versatile Transmitter Chip Links Strain Gages and RTDs to Current Loop	22–3
AN-276	Analog-to-Digital Conversion Using Voltage-to-Frequency Converters	23–5
AN-277	Applications of the AD537 IC Voltage-to-Frequency Converter	23–13
AN-278	Operation and Applications of the AD654 IC V-to-F Converter	23–31
AN-279	Using the AD650 Voltage-to-Frequency Converter As a Frequency-to-Voltage Converter	23–59
AN-280	Mixed Signal Circuit Techniques	24–53
AN-281	Passive and Active Analog Filtering	13–177
AN-282	Fundamentals of Sampled Data Systems	24–85
AN-283	Sigma-Delta ADCs and DACs	20–3
AN-284	Implement Infinite Sample-and-Hold Circuits Using Analog Input/Output Ports	2–11
AN-285	A Function Generator and Linearization Circuit Using the AD7569	2–15
AN-289	8-Bit A/D Converter Mates Transducers with µPs	3–3
AN-290	AD7578 and AD7582 Performance with Reduced V_{DD} Supply (12 V ± 10%)	3–11
AN-291	Asynchronous Clock Interfacing with the AD7878	3–13
AN-292	Interfacing the AD7572A to High Speed DSP Processors	3–17
AN-293	The AD7574 Analog to Microprocessor Interface	3–21
AN-294	AD7672 Converter Delivers 12-Bit 200 kHz Sampling Systems	3–27
AN-295	Using Multiple AD1334s in Many-Channel Synchronous Sampling Applications	3–35
AN-296	Build a Single-Shot Recorder to Catch Fast Transients	3–37
AN-297	Test Video A/D Converters Under Dynamic Conditions	3–49
AN-298	AD671 12-Bit, 2 MHz ADC Digitizes CCD Outputs for Imaging Applications	3–59
AN-299	Simple Circuit Provides Ratiometric Reference Levels for AD782x Family of Half-Flash ADCs	3–67
AN-301	Undersampling Techniques Simplify Digital Radio	3–95
AN-302	Exploit Digital Advantages in an SSB Receiver	3–103
AN-303	Understanding and Applying the AD7341/AD7371 Switched Capacitor Filters	5–19
AN-304	One-Chip "Slide Rule" Works with Logs, Antilogs for Real-Time Processing	6–3
AN-306	Synchronous System Measures µΩs	6–11
AN-307	Modem-Circuit Techniques Simplify Instrumentation Designs	6–13
AN-308	Commutating Amp Multiplies Precisely	6–27
AN-309	Build Fast VCAs and VCFs with Analog Multipliers	6–29
AN-310	Single Supply 12-Bit DAC Circuits Using the AD7568	8–39
AN-311	How to Reliably Protect CMOS Circuits Against Power Supply Overvoltaging	24–29
AN-312	DAS Packs Memory for Easy µP Interface	7–3
AN-313	Getting the Most from High Resolution D/A Converters	8–3
AN-314	14-Bit DACs Maintain High Performance Over Extended Temperature Range	8–43
AN-316	AD7224 Provides Programmable Voltages Over Varying Ranges,	8–103
AN-317	Circuit Applications of the AD7226 Quad CMOS DAC	8–107
AN-318	AD7528 Dual 8-Bit CMOS DAC Application Note	8–113
AN-319	Simple Interface Between D/A Converter and Microcomputer Leads to Programmable Sine-Wave Oscillator	8–121
AN-320A	CMOS DACs and Op Amps Combine to Build Programmable Gain Amplifiers—Part I	8–127
AN-320B	CMOS DACs and Op Amps Combine to Build Programmable Gain Amplifiers—Part II	8–139
AN-321	3-Phase Sine Wave Generation Using the AD7226 Quad DAC	8–149
AN-322	Improve Function Generators with Matched D/A Converters,	8–161
AN-323	Microstepping Drive Circuits for Single Supply Systems	8–167
AN-324	Simple DAC-Based Circuit Implements Constant Linear Velocity (CLV) Motor Speed Control	8–175
AN-325	12-Bit Analog I/O Port Uses AD7549 and 8051 Microcomputer	8–181
AN-326	Interfacing the AD7549 to the MCS-48 and MCS-51 Microcomputer Families	8–185
AN-327	DAC ICs: How Many Bits Is Enough?	8–189
AN-328	Generate 4 Channels of Analog Output Using AD7542 12-Bit D/A Converters and Control the Lot with Only Two Wires	8–201
AN-329	Dynamic Performance of CMOS DACs in Modem Applications	8–207
AN-330	Animation Using the Pixel Read Mask Register of the ADV47x Series of Video RAM-DACs	8–257
AN-331	Changing Your VGA Design from a 171/176 to an ADV471	8–265

		Page
AN-332	Improved PCB Layouts for Video RAM-DACs Can Use Either PLCC or DIP Package Types	8–267
AN-333	Design and Layout of a Video Graphics System for Reduced EMI	8–269
AN-334	Digital Signal Processing Techniques	9–3
AN-336A	Image Compression: Spelling Out the Options	9–41
AN-336B	JPEG Compression	9–45
AN-338	Loading an ADSP-2101 Program via the Serial Port	9–61
AN-342	Analog Signal-Handling for High Speed and Accuracy	8–15
AN-343	Circuit Ideas for IC Converters	8–71
AN-344	Digital FIR Filters Without Tears	9–19
AN-345	Grounding for Low and High Frequency Circuits	24–3
AN-346	Understanding Interference-Type Noise	24–11
AN-347	Shielding and Guarding	24–15
AN-348	Avoiding Passive-Component Pitfalls	24–21
AN-349	Keys to Longer Life for CMOS	24–27
AN-350	Take the Guesswork Out of Settling-Time Measurements	24–33
AN-351	Ask the Applications Engineer—2: Offset and Gain Trimming	24–43
AN-352	Ask the Applications Engineer—5: High Speed Comparators	24–45
AN-353	Ask the Applications Engineer—10: PC Board Layout	24–49
AN-354	Ask the Applications Engineer—1: ADC Problems; Log Compensation Resistors	24–51
AN-355	Behind the Switch Symbol Use CMOS Analog Switches More Effectively When You Consider Them As Circuits	12–3
AN-356	User's Guide to Applying and Measuring Operational Amplifier Specifications	13–3
AN-357	Operational Integrators	13–11
AN-358	Noise and Operational Amplifier Circuits	13–25
AN-359	Settling Time of Operational Amplifiers	13–95
AN-360	Ask the Applications Engineer—6: Op Amp Issues	13–175
AN-361	Ask the Applications Engineer—3: V/F Converters	23–3
AN-362	Ask the Applications Engineer—9: Decoupling; Component Selection	24–47
AN-368	Evaluation Board for the AD7701/AD7703 Sigma-Delta A/D Converters	3–121

Author Index

Page

Adams, Robert 8–189
Alexander, Mark 4–57; 8–23; 13–155
Ash, Dan 9–93
Barrow, Jeff 14–9; 24–3
Bowers, Derek F. 13–35, 155
Brockman, Don 24–7
Brokaw, A. Paul 8–15; 13–17; 22–3; 24–3
Brunner, Everhard 6–49
Bryant, James M. 2–15; 13–175; 23–3; 24–43, 47, 49, 51
Buchanan, David 8–285
Buxton, Joe 8–85, 99; 13–115, 125, 131, 137, 143, 149
Byrne, Mike 2–11; 3–37; 8–103, 107, 149; 24–29
Cavigioli, Chris 9–41, 45
Ciarlone, Frank J. 6–27
Clarke, Bob 6–49; 13–39, 93; 17–3
Counts, Lew 6–3; 8–71
Curtin, Mike 3–109, 115, 121; 5–19; 8–43, 181, 185
Davies, Elwyn 3–59
Deevy, Ken 3–37
Demrow, Robert I. 13–95
Fazio, Mark 17–3
Fine, Bob 9–77
Flett, F.P. 16–33
Fu, Dennis Shi-hsiung 16–45, 49, 53
Garcia, Adolfo A. 18–19
Gasking, John 16–3
Gerstenhaber, Moshe 6–11, 27
Gilbert, Barrie 6–49; 23–13
Goldstein, Stephan 3–35
Gormley, Eamonn 8–257
Grant, Doug 3–3; 23–13; 24–21
Greenwald, Jeffrey A. 11–3
Groshong, Richard 3–95, 103
Harvey, Barry 24–33
Heinzer, Walter 8–99
Henry, Peter 13–107
Hill, Allen 14–3
Hilton, Craven 14–9
Hunt, Bill 8–113
Jachowski, Mike 4–55; 5–17
Jung, Walter 4–73; 8–71; 10–11; 18–3; 23–33
Karagozyan, Kapriel 9–49
Kester, Walt 3–49, 71, 79, 85
Kitchin, Charles 6–3; 13–105
Klonowski, Paul 19–7; 23–5

Page

LeFort, Bob 21–3
McGuire, Gerald 9–61, 77
Martin, Steve 23–61
Mercer, Doug 3–3
Murphy, Mark 6–11; 17–15
Patel, Hitesh 16–11
Reidy, John 3–13, 17, 27; 5–19; 12–5
Rich, Alan 24–11, 15
Ries, Bob 21–3
Riskin, Jeffrey R. 8–71; 10–3
Ruscak, Stephen 3–95, 103
Schirmer, Mark 16–17, 21
Scott, Dave 17–3
Sheehan, Dan 3–11, 37; 12–11, 21
Sheingold, D.H. 13–25
Sherman, Steve 6–3
Slattery, Bill 8–153, 253, 257, 265, 267, 269
Smith, Jeff 17–15
Smith, Lewis 13–25
Smith, Matt 2–5; 3–15; 5–3; 12–21
Stakelum, Robert 8–43
Stata, Ray 13–3, 11
Sylvan, John 18–11; 24–45
Thomas, Mark 16–7
Toldalagi, Paul 9–19
Tolley, William E. 13–171
Toomey, Paul 3–21; 8–113, 161
Tucholski, Hans 7–3
Wayne, Scott 8–3
Whitmore, Jerry 12–3; 24–27
Williams, Arnold 24–7
Windsor, Bill 9–19
Wolfeld, Bruce 9–65
Wong, James 8–83; 13–41, 49, 75, 77, 79, 81, 83, 85, 87, 89, 91; 15–3
Wurcer, Scott 10–11; 13–105; 24–21
Wynne, John 2–3, 7; 3–67; 7–3; 8–35, 39, 121, 127, 139, 167, 175, 193, 201, 269; 12–27

WORLDWIDE HEADQUARTERS

One Technology Way, P.O. Box 9106, Norwood, MA 02062-9106, U.S.A.
Tel: (617) 329-4700, Fax: (617) 326-8703

Complete Worldwide Sales Office Directory Can Be Found on Pages 25–6 and 25–7.

PRINTED IN U.S.A.
G1799–40–5/93
ISBN:0—916550–12–5