SPECIAL ANNIVERSARY ISSUE

Reed Electronics Group

EDN®

VOICE OF THE ENGINEER

SEPT 28
Issue 20/2006
www.edn.com

Low-Power DACs in Miniature Packages

8-/10-/12-Bit D/A Converters Provide Seamless Upgradeability

Features

- Pin- and function-compatible across resolutions
- 2- and 4-channel family with smallest package outline in-class (3 mm x 3 mm LLP)
- Rail-to-rail output swing
- Power consumption at 3.6V
  1 channel, 226 µA (max)
  2 channel, 270 µA (max)
  4 channel, 485 µA (max)
- External reference (2- and 4-channel)
- Accepts input clock rates up to 30 MHz over 2.7V to 5.5V
- Operates over -40°C to 105°C

Ideal for use in portable, battery-powered applications in industrial, medical, and consumer designs

D/A Converter Family

Product	Resolution	Channels	Settling Time (typ)	Package
DAC081S101	8-bit	1	3 µsec	SOT-6, MSOP-8
DAC101S101	10-bit	1	5 µsec	SOT-6, MSOP-8
DAC121S101	12-bit	1	8 µsec	SOT-6, MSOP-8
DAC082S085	8-bit	2	3 µsec	MSOP-10, LLP-10
DAC102S085	10-bit	2	4.5 µsec	MSOP-10, LLP-10
DAC122S085	12-bit	2	6 µsec	MSOP-10, LLP-10
DAC084S085	8-bit	4	3 µsec	MSOP-10, LLP-10
DAC104S085	10-bit	4	4.5 µsec	MSOP-10, LLP-10
DAC124S085	12-bit	4	6 µsec	MSOP-10, LLP-10

625,000* PARTS
READY TO ROLL
Rated #1 for Availability of Product!
SAME-DAY SHIPMENT
On Orders Entered by
8:00
PM Central
NEXT-DAY DELIVERY
Digi-Key®
CORPORATION
Quality Components Just a Click Away!
1.800.344.4539 • www.digikey.com
* New Product Added Daily!

I'm Counting On My Prescription To Be Safe.

At Atmel®, we are too. That is why you can count on brand integrity when drug companies use Radio Frequency IDentification (RFID) tags featuring our CryptoRF® integrated circuits to track your medication.

Cryptographic RFID tags provide the pharmaceutical industry with ultra secure technology that ensures product integrity — from the manufacturer to the pharmacy. Safety sealed caps containing RFID tags securely store the history of the product with a digital signature that cannot be duplicated.

When it comes to security, nobody has more industry proven secure solutions than Atmel. With a broad portfolio, we have the expertise and flexibility to provide customer-driven solutions for security in an unsecured world.

When you think security, think Atmel. We're Everywhere You Are.

Check out Atmel's security solutions today at www.atmel.com/ad/pharma

Everywhere You Are®

CRYPTORF®

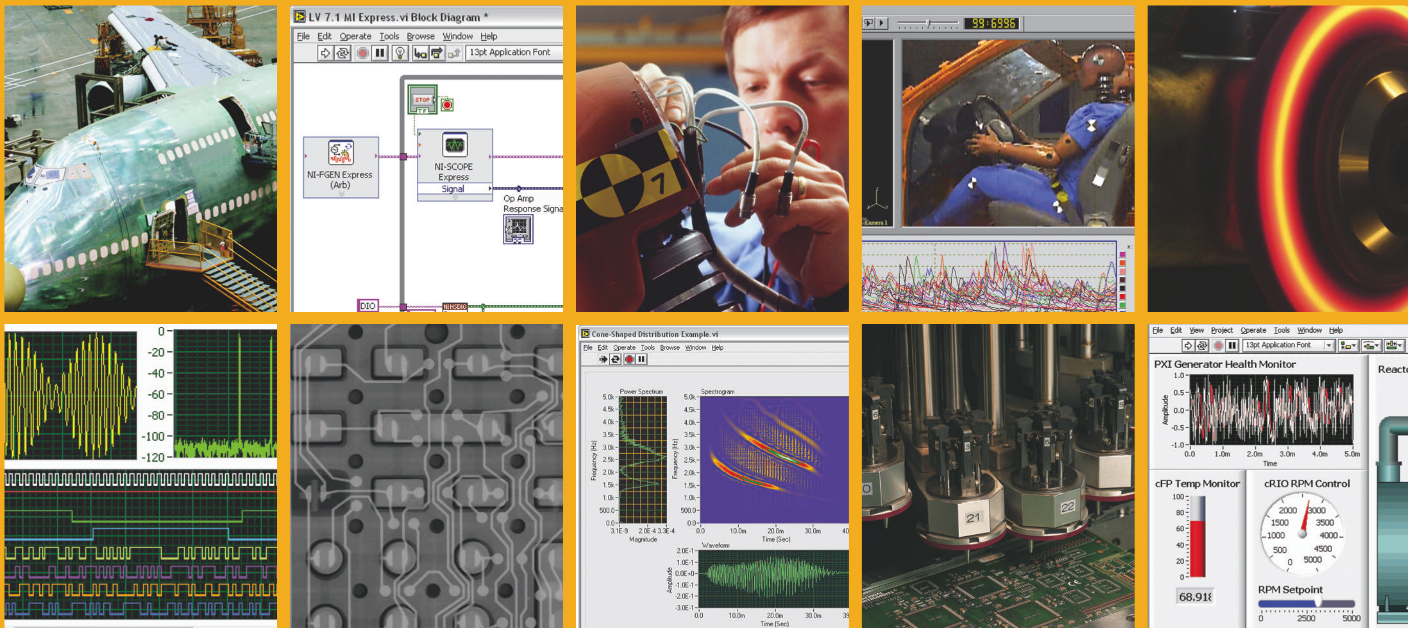

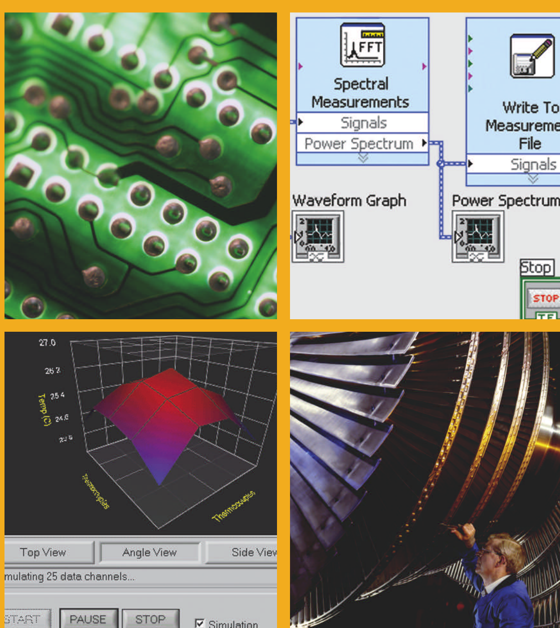

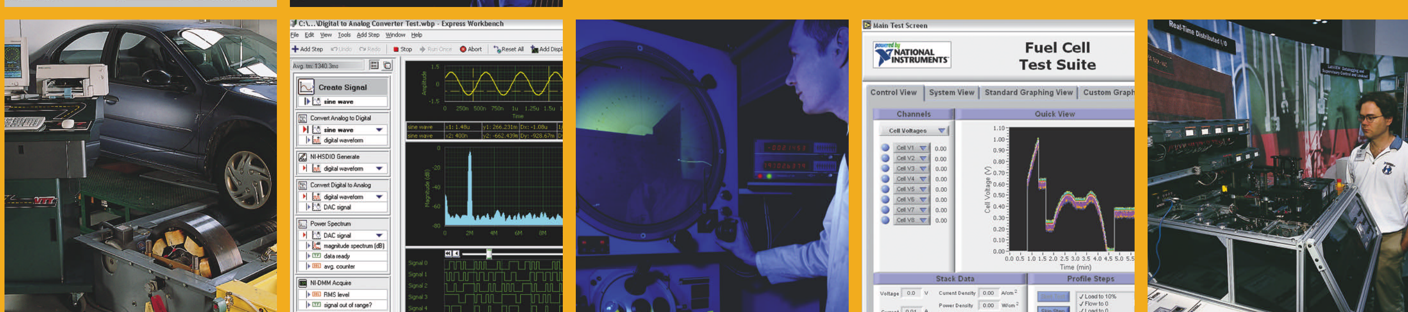

123

EDN

contents

9.28.06

DESIGNIDEAS

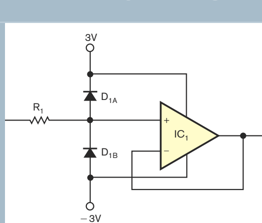

contents 9.28.06

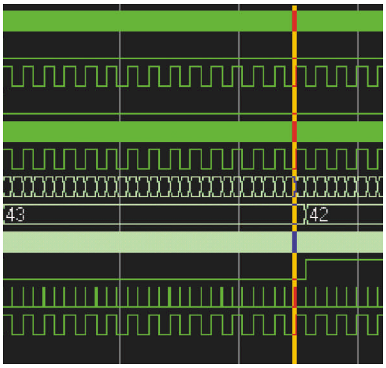

pulse

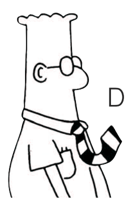

DEPARTMENTS & COLUMNS

PRODUCT ROUNDUP

EDN® (ISSN#0012-7515), (GST#123397457, R.B.I. Intl Pub Mail #0280844) is published biweekly, 26 times per year, by Reed Business Information, 8878 Barrons Blvd, Highlands Ranch, CO 80129-2345. Reed Business Information, a division of Reed Elsevier Inc, is located at 360 Park Avenue South, New York, NY 10010. Tad Smith, Chief Executive Officer; Stephen Moylan, President, Boston Division. Periodicals postage paid, ride along enclosed, at Littleton, CO 80126 and additional mailing offices. Circulation records are maintained at Reed Business Information, 8878 S Barrons Blvd, Highlands Ranch, CO 80129-2345. Telephone (303) 470-4445. POSTMASTER: Send address changes to EDN®, PO Box 7500, Highlands Ranch, CO 80163-7500. EDN® copyright 2006 by Reed Elsevier Inc. Rates for nonqualified subscriptions, including all issues: US, $165 one year; Canada, $226 one year, (includes 7% GST, GST#123397457); Mexico, $215 one year; air expedited, $398 one year. Except for special issues where price changes are indicated, single copies are available for $10 US and $15 foreign. Publications Mail Agreement No. 40685520. Return undeliverable Canadian addresses to: Deutsche Post, 4960-2 Walker Road, Windsor ON N9A 6J3. E-mail: subsmail@reedbusiness.com. Please address all subscription mail to EDN®, 8878 S Barrons Blvd, Highlands Ranch, CO 80129-2345. EDN® is a registered trademark of Reed Elsevier Properties Inc, used under license. A Reed Business Information Publication/Volume 51, Number 20 (Printed in USA).

EDN online contents

www.edn.com

THE CELEBRATION CONTINUES ONLINE

At www.edn.com/50th, you'll find an interactive version of the poster that's bundled with this 50th-anniversary issue. The poster highlights some of the most significant milestones in tech-industry history, and the online edition allows you to dig deeper into these key innovations. Start by listening to an audio introduction from Editor in Chief Maury Wright. Then, click on any of the milestones for an in-depth exploration that includes *EDN*'s original coverage and a present-day analysis.

Also at www.edn.com/50th, you'll find all of our anniversary-themed content (including extras that did not make it into the print edition), an index of all the Milestones That Mattered articles we've been publishing throughout 2006, and the complete contents of *EDN*'s 25th-anniversary issue. In this quarter-century-ago issue, published Oct 14, 1981, *EDN*'s editors made an ambitious attempt to forecast how technology would progress by 2006. Some of their predictions proved wildly incorrect, but others now seem downright spooky. We thought you'd enjoy this unique look at the electronics industry through a 25-year-old lens.

READERS' CHOICE

Recent articles getting high traffic on www.edn.com:

Consumer Court: Portable media players, Wi-Fi Skype phones, more ...
→ www.edn.com/article/CA6366986

Architects adapt to the multicore era
→ www.edn.com/article/CA6366273

Milestones That Mattered: HP-IB revolutionized ATE
→ www.edn.com/article/CA6360312

Ultralow-cost, two-digit counter features few components
→ www.edn.com/article/CA6360316

Howard Johnson: Voltage-regulator model
→ www.edn.com/article/CA6360320

Designing Ethernet into industrial applications
→ www.edn.com/article/CA6360321

Battery-free sensors convert motion into energy
→ www.edn.com/article/CA6360326

Circuit makes simple high-voltage inverter
→ www.edn.com/article/CA419572

Ultrasonic range finder uses few components
→ www.edn.com/article/CA45965

NEW BLOGS

When you hear the word "blog" you probably think of lonely losers posting their ramblings and uninformed opinions. But in *EDN*'s blogs, some of the brightest engineering minds share their highly informed positions and hard-won industry insights. We recently launched these new blogs:

Anablog: Technical Editor Paul Rako delves into analog technology.

Embedded Webblog: Technical Editor Warren Webb covers board-level design.

Practical Chip Design: Executive Editor Ron Wilson explores how IC-design teams really work.

Check out all eight of *EDN*'s topical, technical blogs:
→ www.edn.com/blogs

PRODUCT FEED

Throughout *EDN*'s 50 years, engineers have had a taste for *EDN*'s product briefs: concise, no-nonsense write-ups of new ICs and components. Now, we're proud to introduce *EDN* Product fEEd, a centralized online location for this critical part of our coverage.

We invite you to belly up to fEEd, where you'll find a daily dose of new products, as well as handy category filters that allow you to select the morsels that most interest you, be they analog or digital ICs, power components, sensors, passives, boards and systems, or software (to name just a few).

EDN Product fEEd:
→ www.edn.com/productfeed

RSS feed of the fEEd:
→ www.edn.com/ProductfEEdWebBlog.xml?nid=3066

BY MAURY WRIGHT, EDITOR IN CHIEF

EDN AT 50

Welcome to the 50th-anniversary issue of *EDN*. Although we've been celebrating with our "Milestones That Mattered" columns and other excerpts from the past all year, this issue includes a special section that's dedicated to our anniversary and, more important, to the last five decades of the most exciting industry ever. Of course, you'll find a bit of speculation about the future, as well. And I'd like to share some of our online celebration plans with you, too.

The 50th-anniversay section of this print edition of *EDN* begins after our regular features section, starting on pg 123. As we planned this section, we first discussed a series of articles on technologies that have been driving forces in the industry and that will continue to drive the industry forward. It quickly became apparent, however, that a more meaningful set of articles might focus on higher level trends in the industry that have been with us for more than five decades in one way or another and that will be with us for some time yet.

Picking trends was also tough, but we settled on four: miniaturization, software-ization, virtualization, and thermal efficiency. It's easy to see that smaller is an enduring trend. The second trend focuses on the fact that an ever-increasing number of functions that engineers once implemented in hardware, now migrate to software. Virtualization started at least as early as IBM mainframes and plays key roles today in applications such as compelling video games. Finally, all electronic functions and measures of performance have a thermal cost. I believe you'll find the stories on these trends both compelling and entertaining.

This issue includes a special section that's dedicated to our anniversary and, more important, to the last five decades of the most exciting industry ever.

You'll also find a poster with a time line of "Milestones That Mattered" packaged with this print issued. Some of the milestones have appeared in previous issues, and others will be brand new to you. The poster includes brief descriptions of each milestone. Check out the interactive version of the time line at www.edn.com/timeline to get more detail on each, including stories from our archives going back to 1956.

You may note on the poster that the last of the milestones that we've chosen is from 1996. We chose to invoke Leibson's Law. Steve Leibson is a former *EDN* editor in chief, and he declared that it takes anything 10 years to catch on. In fact, you'll find examples of that law, such as Ethernet, on our time line. Had we created a time line in the early 1980s, a decade after work on Ethernet started, we wouldn't likely have chosen Ethernet as a milestone.

Without question, many innovations have debuted in the past decade. But we found that it was simply too soon to bestow milestone status. The interactive time line, however, will live into the future. We'll add milestones from the last 10 years and perhaps ones we missed from the past, and we will happily take you suggestions.

Other goodies await you in the 50th-anniversary section of our Web site, www.edn.com/50th. You will find the entire contents of our 25th-anniversary issue from Oct 14, 1981. I'd encourage you to spend some time with it. The articles are compelling, and the projections were uncannily accurate.

We'll also offer up some audiocasts from industry luminaries discussing the industry's past, present, and future. The audiocasts will also shed additional light on the four cornerstone trends that we've chosen. You'll even find some reflections from previous *EDN* chief editors.

Our work on the 50th celebration has been a year-long process. And we'll have a few more surprises later in the year. I'd like to thank everyone in the *EDN* family for all of the extra effort. I hope you readers enjoy the results as much as we did in producing the entire 50th project. And I encourage each of you to think about and embrace the fact that there is no more exciting industry to which you could dedicate your life's work. **EDN**

You can reach me at 1-858-748-6785 or mgwright@edn.com.

8MHz
Industry's First 500mA Synchronous Buck Regulator

Inductor Size Gets Leaner By Greater Than 95%

Micrel's MIC2285 is a high efficiency 8MHz pulse width modulated (PWM) synchronous buck (stepdown) regulator. It features a LOWQ™ LDO standby mode that draws only 20µA of quiescent current. The MIC2285 is the industry's breakthrough ultra-low noise, small size, and high efficiency solution for portable power applications.

In PWM mode, the MIC2285 operates with a constant frequency 8MHz PWM control. Under light load conditions, such as those in system sleep or standby modes, the PWM switching operation can be disabled to reduce switching losses. In this light load LOWQ mode, the LDO maintains the output voltage and draws only 20µA of quiescent current. The LDO mode of operation saves battery life and does not introduce spurious noise and high ripple which is common place in pulse skipping or bursting mode regulators. The MIC2285 operates from a 2.7V to 5.5V input voltage and features internal power MOSFETs that can supply up to 500mA output current in PWM mode. It can operate with a maximum duty cycle of 100% for use in low-dropout conditions.

For more information, contact your local Micrel sales representative or visit us at: www.micrel.com/ad/mic2285.
Literature: 1 (408) 435-2452 Information: 1 (408) 944-0800

The Good Stuff:

- 2.7 to 5.5V supply/input voltage
- Light load LOWQ LDO mode
- 20µA quiescent current
- > 90% efficiency
- Low noise, 75µVrms
- 100% maximum duty cycle
- Adjustable output voltage option down to 1V
- Ultra-fast transient response
- Ultra-small 0.47µH inductor
- Fully integrated MOSFET switches
- Micropower shutdown operation
- Thermal shutdown and current limit protection
- Pb-free 10-pin 3mm x 3mm MLF™ package
- –40°C to +125°C junction temperature range

MICREL®
Innovation Through Technology™

www.micrel.com

EDN. 225 Wyman St, Waltham, MA 02451. **www.edn.com**. Phone 1-781-734-8000; fax: 1-781-734-8070. **Address changes or subscription inquiries:** phone 1-800-446-6551; fax: 1-303-470-4280; subsmail@reedbusiness.com. For a free subscription, go to **www.getfreemag.com/edn**. Reed Business Information, 8878 S Barrons Blvd, Highlands Ranch, CO 80129-2345. Include your mailing label.

Reed Electronics Group

Audiophile amplifier offers 0.00003% THD+N

National Semiconductor, long a major player in the audio market, revisits that arena with the high-voltage, high-performance LM4562 dual audio op amp. The device operates from 5 to 34V, and THD+N (total harmonic distortion plus noise) is 0.00003%. It has an equivalent-input-voltage noise of 2.7 nV/$\sqrt{\text{Hz}}$ and a 1/f corner of 60 Hz. The device drives 45 mA and features output-short-circuit and thermal-shutdown protection. The product's bias current is only 10 nA, and bandwidth is 6 MHz. The LM4562 comes in SOIC, DIP, and hermetic-metal-can packages for $2.35, $2.65, and $9.95 (1000), respectively.

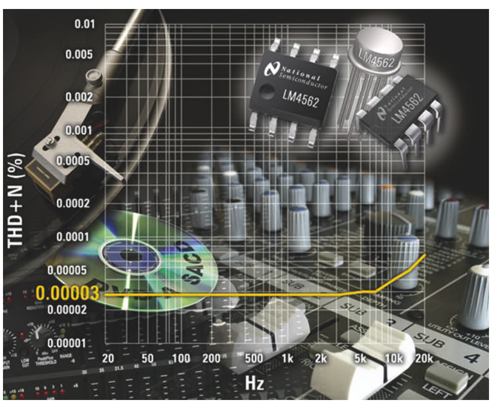

National Semiconductor revisits the audio market with the LM4562 dual audio op amp.

The LM4562 complements the recently released LM4702, which will drive high-voltage transistors to create a 150V audio driver, and the company plans units that will operate from 170V. You can create a 300W/channel audio amplifier by adding a couple of power transistors. The LM4702A sells for $24.95 (100), and the LM4702B sells for $150 (25) and will be available by early 2007.**—by Paul Rako**

▷**National Semiconductor**, www.national.com.

FROM THE VAULT

Because of their enormous appetite for floating-point operations, 3-D-graphics displays were previously limited to high-end workstations. With some powerful new processors, however, you can now bring workstation-quality 3-D graphics to such applications as PC add-in boards and embedded-graphics systems.

—*EDN*, March 30, 1989, pg 97

Mixed-signal module boosts wideband performance

Mercury Computer Systems' new Echotek Series ECV4-RFT (remote-fiber-transceiver) VME module combines high-speed ADC and DAC technology, FPGA silicon, and fiber communications to coordinate data streams to and from distributed sensors. The ECV4-RFT resides near sensors or antennas while communicating with a remote signal-processing subsystem. It accepts two analog inputs and produces two analog outputs. After conversion to a digital signal, the input data stream moves to a Virtex-4 SX55 FPGA, which can function as a digital receiver or preprocessor. The digital IF signal then travels over a fiber-optic connection to the signal-processing subsystem. After analysis, the signal returns over the fiber link to another FPGA and finally through the DACs and into the antenna as an analog signal.

Three onboard Xilinx (www.xilinx.com) FPGAs enable users to deploy custom algorithms, such as frequency conversion, FFTs, or signal filtering, directly on the board. The 6U VME board needs no backplane interface other than input power. The ADC interfaces operate as fast as 1.5G samples/sec at 8-bit resolution, and the DAC interfaces deliver as much as 1.2G samples/sec at 14-bit resolution. The ECV4-RFT is available now, and prices start at $13,300 (OEM quantities).**—by Warren Webb**

▷**Mercury Computer Systems**, www.mc.com.

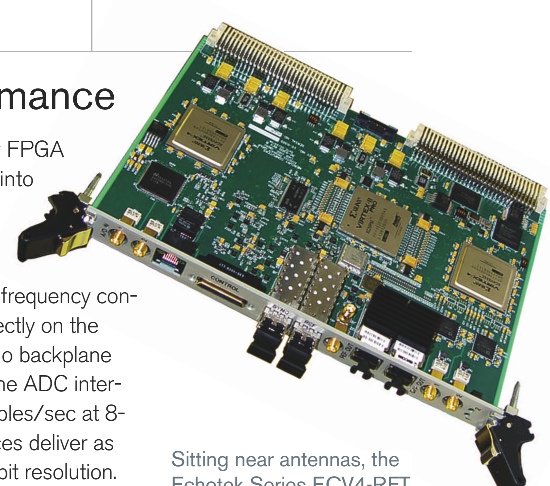

Sitting near antennas, the Echotek Series ECV4-RFT communicates with a signal-processing subsystem as much as 100 ft away.

Software turbocharges IC-logic simulation

With ICs becoming ever more complex and larger in gate counts, an ongoing demand exists for faster and higher capacity verification tools. EDA start-up Liga Systems addresses that need with its new NitroSim hardware-accelerated simulation environment. The company claims that the tool improves simulation performance from 10 to 100 times over single-CPU simulation and can handle designs with as many as 300 million gates.

Most of Liga's creators were previously with Aptix. That company, which Mentor Graphics (www.mentor.com) last year purchased, mainly focused on prototyping boxes. Liga's NitroSim is instead selling "turbocharged software simulation," according to Chief Executive Officer Henry Verheyen.

NitroSim comprises three pieces of software and a custom pc board that users plug into their PCs. The NitroSim PCI, essentially an acceleration card, includes a custom VLIW (very-long-instruction-word) processor and lots of onboard memory with slots for expansion. The software includes the NitroSIM CC RTL (register-transfer-level) and gate-level netlist compiler for running designs on the NitroSIM PCI card; the NitroSIM TB compiler upgrade for running behavioral RTL netlists on NitroSIM PCI; and the NitroSIM RT runtime technology, which allows the NitroSIM to communicate with other simulation and test environments. The RT software also configures the PCI card.

Unlike accelerators, the NitroSim doesn't map the design into an FPGA or a custom processor. Instead, the software compiles the code into memory instructions for execution by the VLIW processor. This approach allows the design to quickly run the design, says Verheyen, and avoids the problem of cache misses that limit the capacity scalability of pure-software simulators. "Software simulators run so slowly mainly because the amount of data the instructions are processing is many times larger than the on-chip cache," says Verheyen. "The simulator is constantly running into cache misses, which means that the processor must stop every time and then go to memory to access the new data and instructions."

A Virtex-4 holds the VLIW-processor grid, which memory surrounds. The system feeds instructions into the grid from all memories in parallel, providing much higher bandwidth than a normal CPU, according to Verheyen. Because the system doesn't rely on a logic-synthesis tool to compile the design under test into an FPGA or a custom processor, users can program behavioral RTL code into a NitroSIM system. NitroSIM supports testbenches, four-state handling, behavioral checks, and assertions. A NitroSIM for a 300 million-gate configuration sells for $50,000.

—by Michael Santarini

▷**Liga Systems**, www.ligasystems.com.

FROM THE VAULT

Logic-synthesis tools are the latest in a long line of CAE (computer-aided-engineering) tools that are destined to shrink the product-development cycle for ICs and ultimately convert gate-level designers to systems engineers.

—*EDN*, March 30, 1989, pg 51

DILBERT *By Scott Adams*

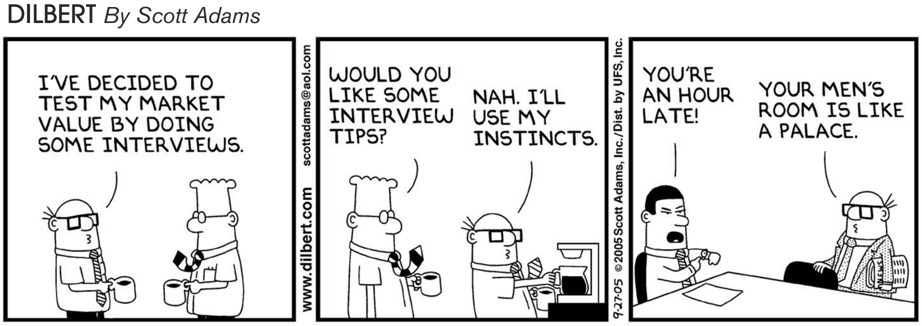

09.28.06

1-MHz gain-selectable amps compensate for bandwidth roll-off

Microchip's new MCP6G01 and MCP6G02 series of amplifiers use one pin to select a gain of one, 10, or 50 at 1 MHz, 350 kHz, and 300 kHz, respectively. The third pin is floating, yielding the third gain setting, and gain selection adjusts compensation.

The MCP6G01 and MCP6G02 amplifiers have gains of one, 10, or 50.

Quiescent current is 100 µA. The internal gain-setting resistors also give a gain accuracy of better than 1%. CMOS-process technology limits the maximum operating voltage to 5.5V, but the part operates on 1.8V for mobile-system and other low-voltage-system applications.

The MCP6G01 and MCP6G02 are available in eight-pin SOIC packages or MSOPs for 34 and 43 cents (10,000), respectively. The quad MCP6G04 comes in 14-pin SOIC packages and TSSOPs for 70 cents.**—by Paul Rako**

▶**Microchip**, www.microchip.com.

RF-signal-analysis instruments: now faster and smarter

The growing complexity of wireless-communication standards is motivating users of RF-signal-analysis instruments to demand units that can make more complex measurements than do existing products. Moreover, the economics of such businesses as mobile phones, whose manufacturers annually produce hundreds of millions of units, necessitate drastic reductions in test times, not only in production, but also in product development. This situation has motivated Agilent and Tektronix to introduce new RF-signal-analysis products. Even if only by accident, the companies invited comparisons by making their announcements on the same day.

At first, though, it might appear that, despite the similarities, comparisons between these products could be unfair because the base prices of Agilent's LXI Class C-compliant N9020A MXA signal analyzers are considerably lower than those of Tektronix's RSA-6100A real-time spectrum analyzers. MXA prices start at $25,900 for a unit that covers 20 Hz to 3.6 GHz. Prices for the widest bandwidth member of the series, which covers 20 Hz to 26.5 GHz, start at $42,900. The company will also continue to offer its higher performance, higher priced PSA spectrum-analyzer line. Meanwhile, RSA6100A prices start at $69,900 for a unit that covers 9 kHz to 6.2 GHz.

Prices enter the same ballpark, however, if you carefully pick units that offer similar frequency ranges and if you equip the selected Agilent product with options that enable it to make many of the measurements that are the Tek products' forte. Also, each product family boasts a dynamic range of approximately 73 dB, and each includes a member whose frequency range extends to approximately 14 GHz. Tek's RSA-6114A covers 9 kHz to 14 GHz. Agilent's N9020A-513 covers 20 Hz to 13.6 GHz. The Tek unit, whose standard measurement bandwidth is 40 MHz (and, optionally, 80 or 110 MHz), carries a base price of $75,900. You can purchase the Agilent unit with an analysis bandwidth of 25 MHz. So equipped and with a hardware connection and signal-analyzer software, the Agilent unit costs $62,420. Additional options can easily increase this price beyond $70,000.

Both analyzer families use heterodyne downconversion to bring the signals of interest to a relatively low frequency and then use high-speed, high-resolution ADCs to convert the signals to the digital domain in which DSP-based techniques enable rapid analysis. Moreover, the instruments can store and recall the digitized data, allowing postprocessing, which can reveal phenomena that you may want to examine further.

The character of the Agilent instruments depends on the analysis software you load into them. They can become vector-signal analyzers, spectrum analyzers, or both. Although Tek does not offer such options, it offers some capabilities that the company calls revolutionary. Potential customers will almost certainly agree.

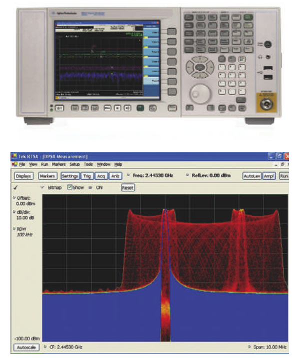

You can endow Agilent's N9020A MXA RF-signal analyzers (top) with the capabilities of a spectrum analyzer, a vector-signal analyzer, or both. Tektronix's DSP-based RSA-6100 real-time spectrum analyzers (bottom) can display how complex spectra change in real time.

The most dramatic new feature is Tek's first-time incorporation into RF-signal-analysis instruments of its proprietary DPX (digital-phosphor) technology, which the company originally developed for use in real-time-sampling digital oscilloscopes. With this technology, you can display rapidly changing broadband spectra as what amount to real-time videos. At the end of each video, a color-keyed image remains on the screen to show the percentage of time the signal occupied each portion of the displayed frequency band. The instruments do not, however, display waterfall diagrams—the traditional presentation of time-varying spectra—although Tek insists that the units display the same information in a more easily understood form. The live spectrum displays are possible not only because of the DPX technology, but also because of the units' DSP-based capability to perform 48,000 spectral analyses/sec. According to Tek, that number is almost 1000 times the analysis speed of the fastest competing instrument.

Another feature of the RSA6100 series is an improved version of FMT (frequency-mask trigger), which Tektronix first offered in earlier real-time spectrum analyzers. FMT lets you trigger measurements based on the occurrence of unique patterns of events in the spectrum. Coupled with the instruments' high dynamic range, FMT allows triggering on weak signals and ignores known strong ones. The company says that legacy swept-frequency spectrum- and vector-signal analyzers lack similar capabilities.

Returning to Agilent, the company also announced a family of RF-signal generators at the same time as it announced the MXA signal analyzers. The MXG family comprises three analog-signal generators—the N5181A MXG series—and two vector-signal generators that make up the N5182A MXG series. Each of the five generators, which are companion products to the MXA family, covers a frequency range that begins at 250 kHz. The analog series includes units whose frequency coverage extends to 1, 3, and 6 GHz at base prices ranging from $6200 to $15,000. Within the vector series are two units whose coverage extends to 3 and 6 GHz and whose base prices are $16,000 and $25,000. Among the family features are simplified self-maintenance, and what Agilent terms the industry's fastest switching speeds and best ACPR (adjacent-channel power ratio).

—by Dan Strassberg

▷**Agilent Technologies**, www.agilent.com.

▷**Tektronix Inc**, www.tektronix.com.

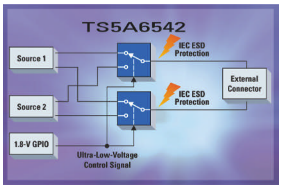

▶ Single SPDT analog switch with input logic translation and IEC ESD (TS5A6542)

– +/- 15kV IEC 61000-4-2 Contact Discharge ESD protection (common pin)

– Low ON-state resistance (0.75Ω max)

– Excellent ON-state resistance matching

– Break-before-make switching

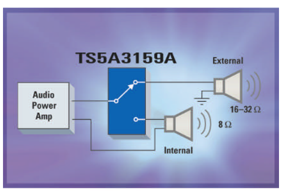

▶ Single 0.9-Ω SPDT analog switch (TS5A3159A)

– 100-MHz bandwidth – highest of any 3159-type in the industry

– Ultra-small wafer chip scale packaging

– Wide low-voltage operating range: 1.65 V to 5.5 V

– Break-before-make switching

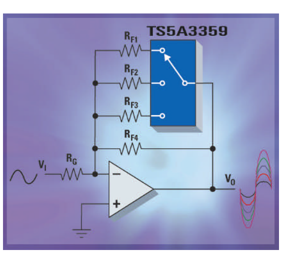

▶ Single, 0.9-Ω SP3T analog switch (TS5A3359)

– Space-saving integration eliminates the need for using serial connected SPDT analog switches

– Isolation in the powered down mode (V+ = 0V)

– 1.65-V to 5.5-V single-supply operation (Control Inputs are 5.5 V tolerant)

– Break-before-make switching

Analog Switches from Texas Instruments

Device	r_{on}* (Ω) (max)	r_{on} Flatness (Ω) (max)	r_{on} Mismatch (Ω) (max)	V_+ (V) (min)	V_+ (V) (max)	ON Time (ns) (max)	OFF Time (ns) (max)	Pins/Packages
SPST								
TS5A3166	0.9	0.15	—	1.65	5.5	7	11.5	5/SC70, SOT-23, WCSP
TS5A3167	0.9	0.15	—	1.65	5.5	7	11.5	5/SC70, SOT-23, WCSP
TS5A4594	8	1.5	—	2.7	5.51	7	14	5/SC70, SOT-23
TS5A4595	8	1.5	—	2.7	5.51	7	14	5/SC70, SOT-23
TS5A4596	8	1.5	—	2.7	5.5	17	14	5/SC70, SOT-23
TS5A4597	8	1.5	—	2.7	5.5	17	14	5/SC70, SOT-23
TS5A1066	10	5	—	1.65	5.5	5.5	4.5	5/SC70, SOT-23, WCSP
SPST x 2								
TS5A23166	0.9	0.25	0.1	1.65	5.5	7.5	11	8/US8, WCSP
TS5A23167	0.9	0.25	0.1	1.65	5.5	7.5	11	8/US8, WCSP
TS3A4741	0.9	0.4	0.05	1.65	3.6	14	9	8/SSOP, MSOP
TS5A2066	10	5	1	1.65	5.5	5.8	3.6	8/SM8, US8, WCSP
SPST x 4								
TS3A4751	0.9	0.4	0.05	1.65	3.6	14	9	14/TSSOP
SPDT								
TS5A6542	0.75	0.25	0.25	2.25	5.5	25	20	8/WCSP
TS5A4624	0.9	0.25	0.1	1.65	5.5	22	8	6/SC70
TS5A3153	0.9	0.15	0.1	1.65	5.5	16	15	8/US8, WCSP
TS5A3154	0.9	0.15	0.1	1.65	5.5	8	12.5	8/US8, WCSP
TS5A3159A	0.9	0.25	0.1	1.65	5.5	30	20	6/SC70, SOT-23, WCSP
TS5A3159	1.1	0.15	0.1	1.65	5.5	35	20	6/SC70, SOT-23
TS5A3160	0.9	0.25	0.1	1.65	5.5	6	13	6/SC70, SOT-23
TS5A3157	10	5	0.2	1.65	5.5	8.5	6.5	6/SC70, SOT-23, WCSP
TS5A63157	10	2	0.14	1.65	5.5	5	3.4	6/SC70, SOT-23
TS5A2053	13.8	4.5	4.5	1.65	5.5	6.8	4.1	8/SM8, US8
SPDT x 2								
TS5A23159	0.9	0.25	0.1	1.65	5.5	13	8	10/MSOP, QFN
TS5A23160	0.9	0.25	0.1	1.65	5.5	5.5	10	10/MSOP
TS5A23157	10	4(typ)	0.15(typ)	1.65	5.5	5.7	3.8	10/MSOP
SPDT x 4								
TS3A5018	10	7	0.8	1.65	3.6	8	6.5	16/SOIC, SSOP (QSOP), TSSOP, TVSOP,QFN
SP3T								
TS5A3359	0.9	0.25	0.1	1.65	5.5	21	10.5	8/US8
TS5A3357	15	6.5(typ)	0.1(typ)	1.65	5.5	6.5	3.7	8/SM8, US8
SP4T x 2								
TS3A5017	12	9	2	2.3	3.6	9.5	3.5	16/SOIC, SSOP (QSOP), TSSOP, TVSOP, QFN

**Data measured under typical conditions with maximum V_+. Data collected as of 7/06*

*New Products are listed in **bold red.***

Technology for Innovators™ TEXAS INSTRUMENTS

5-ball/6-ball WCSP (YZP)
Ball pitch = 0.020 mm (0.50 mm)
Height = 0.020 mm (0.50 mm)
Area = 0.002 mm (1.26 mm)

10-pin QFN (RSE)
Lead pitch = 0.020 mm (0.50 mm)
Height = 0.039 mm (0.60 mm)
Area = 0.005 mm (3.18 mm)

NEW!
Analog Switch Selection Guide
Datasheets and Samples

www.ti.com/switches1

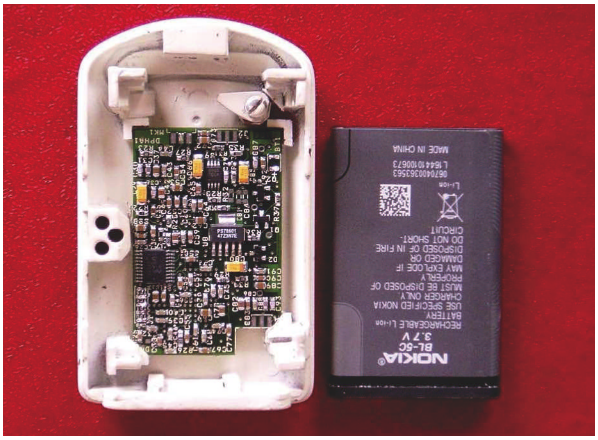

Users wear CDAC's new DPHA on their body, which is more convenient than wearing it behind their ear.

GLOBAL DESIGNER

India's CDAC develops digital hearing aid

The hardware-design group at India's government-funded CDAC (Center for Development of Advanced Computing), in Thiruvananthapuram has developed a low-cost DPHA (digital programmable hearing aid) that a user wears on his body. The DPHA-1 device employs a proprietary ASIC and embedded DSP to deliver stable amplification characteristics over a wide dynamic range.

Unlike conventional analog hearing aids, you can tailor DPHAs to improve clarity of speech, reduce background noise, and help control unwanted loudness. You can also program them to make automatic adjustments in a variety of settings. "Currently, no [other] manufacturer in India offers a DPHA in the body-worn format," says project head R Ravindra Kumar, additional director at CDAC. "They are available only in the more expensive, behind-the-ear format. Our design has all the advantages of DPHAs, such as reprogrammability, stability, and fidelity, but at a lower cost of ownership. This has great benefit for a country like India, which has a huge hearing-deprived populace [about 30 million people] for whom affordability is a significant issue. Also, body-worn DPHAs are more convenient for children."

CDAC's DPHA-1 features a digital volume control to eliminate the crackling-noise characteristic of conventional hearing aids and incorporates frequency-dependent filtering to match its output to the audiogram of the user. You can reprogram the device in the field with PC-based software to tailor the output to match the user's hearing characteristics over a period of time. "All this is made available in a low-power package using CMOS-ASIC implementation," says Kumar. The device works on an easily available, rechargeable lithium-ion battery.

Biju C Oommen, joint director and deputy head at C-DAC, says that the DPHA-1 is a multidisciplinary effort that melds diverse electronic-design technologies, such as DSP design, algorithms and firmware, ASIC- and analog-circuit design, and power management, with the physics of audiology, the anatomy of the human ear, and precision mechanical fabrication.

"We needed to develop the algorithms for the programmable filters and balance the implementation of hardware/software functions in a 32-bit, real-time embedded system using our RISC core," explains lead designer NM Shaji. The pc-board design and engineering were other challenges because the designers had to cram all the components into a 40×30-mm pc board. They also had to carefully design the enclosure to prevent internal mechanical resonance and microphone feedback.

Currently, the DPHA-1 is undergoing formal field trials after CDAC deemed prototype field trials using FPGA implementation successful. CDAC is considering fabricating a 1 million-gate SOC (system-on-chip) ASIC and is scouting for manufacturing partners. The organization estimates that it could eventually offer the DPHA to users for less than \$20. "Our DPHA meets and exceeds the ANSI S3.22 and IEC 60118 standards. We will initiate the certification process after completion of clinical trials," says Kumar.

Winner of *EDN Asia*'s Innovation of the Year award in 1999 for the Oorja ASIC, Kumar and his 30-strong engineering team have a number of successful projects, including energy meters, complex nuclear instruments, compact onboard computers for space vehicles, and smart-card monitoring and control systems, to their credit.

—by Chitra Giridhar, *EDN Asia*

▷**Center for Development of Advanced Computing**, www.cdactvm.in.

Preprocessing switch speeds data in DSP clusters

A new preprocessing switch from IDT incorporates the ability to perform preprocessing in the switch path, leading to an increase of as much as 20% in system efficiency. The switch has 40 bidirectional and configurable SRIO (Serial RapidIO) links, including 10 four-lane-wide ports, as many as 22 one-lane-wide ports, or a combination of four- and one-lane-wide ports. You can independently program each port for 1.25-, 2.5-, or 3.125-Gbps transmission speeds, as well as for short-haul, chip-to-chip-transmission distances or long-haul, backplane-transmission distances.

The chip comes with a simulation tool that allows you to evaluate boards and systems based on RF-card and DSP models, accurately representing channel latency. Within a design, you control the preprocessing switch by selecting and parameterizing a set of its functions at system start-up. Available for sampling now, the switch has a supporting evaluation board, comes in a 676-ball BGA package, and will sell for \$125 (10,000).

—by Graham Prophet, *EDN Europe*

▷**IDT**, www.idt.com.

Solution: Bergquist S-Class Gap Pad®

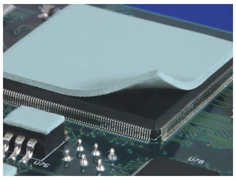

Gap Pad S-Class has natural tack making its application clean and easy to handle.

Gap Pad S-Class is perfectly suited for high performance applications such as VRMs, BGAs and ASICs.

Bergquist's S-Class material is the perfect combination of softness and high thermal conductivity. With a low bulk Shore (00) hardness range from 20 to 30 and thermal conductivity from 2 to 3 W/m-K, these materials conform to demanding contours while maintaining their structural integrity. They are an ideal gap filling solution for applications with fragile components that can be damaged by higher mounting pressures. Gap Pad S-Class is also an excellent solution for CD/DVD drives, memory modules, and PC boards to chassis.

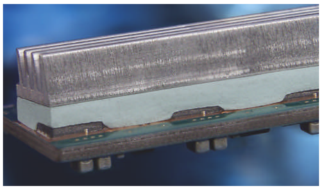

Excellent interfacing and wet-out makes Gap Pad S-Class ideal for fragile components with demanding contours and stack-up tolerances.

Ultra soft S-Class integrity offers easy application.

Gap Pad S-Class has a natural tack that eliminates the need for additional adhesive layers that can inhibit thermal performance. Its super soft, yet elastic nature provides excellent interfacing and wet-out, even to surfaces with high roughness or topography. Gap Pad S-Class features an embedded-fiberglass reinforcement that makes it puncture, shear and tear resistant. No tearing, flaking or crumbling – just clean and easy handling during the assembly process.

FREE S-Class swatch kit.

Visit our web site or call to qualify for your FREE S-Class swatch kit with product details.

FREE Gap Pad S-Class Swatch Kit

Call **1.800.347.4572** or visit **www.bergquistcompany.com/coolrun**

THE BERGQUIST COMPANY

18930 West 78th Street • Chanhassen, MN 55317 • A ISO9001:2000 registered facility
(800) 347-4572 • Phone (952) 835-2322 • Fax (952) 835-0430 • www.bergquistcompany.com

Thermal Products • Membrane Switches • Touch Screens • Electronics • Labels and Graphic Overlays

VOICES

Digi-Key: growing up, privately

Engineer Ronald Stordahl, inventor of the Digi-Key electronic-telegraph key, founded Digi-Key (www.digi-key.com) in 1972. The company went on to become the sixth-largest catalog-electronic-component distributor. *EDN* recently interviewed Mark Larson, president and chief operating officer. The following is an excerpt. Read the entire interview online at www.edn.com/060928p1.

Any thoughts about an IPO (initial public offering)?

A Based on Ron's interest in holding Digi-Key privately and the fact that we have adequate resources to finance our growth initiatives, there is no interest at this time in an IPO or, for that matter, selling to a financial or strategic buyer.

Must you abide by Sarbanes-Oxley accounting regulations? Is Sarbanes-Oxley a good thing or is it a waste of resources?

A Because Sarbanes-Oxley accounting regulations protect the investors in publicly held companies and we are privately held, we are not subject to its cumbersome regulations and high compliance costs.

Protecting investors in public companies is a good thing. Conceptually, that is what Sarbanes-Oxley is supposed to do. Aspects of Sarbanes-Oxley are good, but it is unwieldy and resource-intensive. It remains to be seen if it proves to be an effective tool in protecting investors from dishonest management.

But the existence of Sarbanes-Oxley is not why we are still a privately held company, and I am confident that we wouldn't consider it a show-stopper if at some point in the future we determined that Digi-Key should become a public company.

I heard a rumor that some giant multinational offered to buy Digi-Key and you folks refused the offer since you were happy with the way things were and saw no benefit to a larger company's absorbing you. Is there any truth to that?

A The story you heard is reasonably accurate. We are flattered by the fact that a number of competitors [strategic buyers] have approached Digi-Key over the years. A number of investment companies [financial buyers] have also approached us. But we continue as a privately held company. With the advantage of hindsight, we know we made the right decision. I expect that we will continue to be an acquisition target, given our success in redefining distribution as our strong growth in both customers and sales reflects.

Today, Digi-Key is a household word for engineers, whereas, 20 years ago, they might have never heard of it. What kind of growth has Digi-Key experienced over its lifetime in the tough catalog-electronics-distribution market? What accounts for your unusual success?

A Things have changed rapidly over the last 34 years. When I began managing Digi-Key in 1976, we had sales of about $800,000. This year, our sales should be approximately $830 million. In 1976, we had 14 employees; today, we have more than 1700 employees. In 1976, we had about 1000 square feet of office and warehouse space; today, we have more than 600,000 square feet. In 1976, Digi-Key was probably one of the smallest of more that 600 electronic-component distributors in North America. We have grown to become the sixth-largest electronic-component distributor in North America and the ninth largest in the world.

My stock answer to why we have been so successful is "good management!" But, seriously, I think we have enjoyed "unusual success" because we have been unusually responsive to the needs of the design engineer as well as the purchaser of production quantities.

In the crash of 2000, Digi-Key saw only a 4% reduction in business and did not have to lay off any employees. Has your work force been stable?

A Digi-Key was fortunate in the crash of 2000. When many distributors saw their sales plummet 25 to 40% or even more, our sales dropped only slightly. Our work force has been remarkably stable, with relatively low turnover. This fact has allowed us to invest in their training and strengthen their ability to service our customers. But, just as important as the low turnover is the strong work ethic, which is characteristic of the people in this area (Thief River Falls, MN). They work hard and take pride in their work.

Is your work labor-intensive?

A Although I don't think of Digi-Key, specifically, or electronic distribution, generally, as particularly labor-intensive, having dedicated employees is critical to Digi-Key's success.

***Modern Materials Handling* named you Warehouse of the Month in 1992. Back then, I realized that Digi-Key had some of the most advanced and sophisticated materials-handling and inventory systems on earth. Has Digi-Key taken advantage of high-tech warehousing?**

A Digi-Key's product-distribution center is highly technology-based. Our goal is to ship all orders received before 8 pm Central Standard Time the same day that we receive them. To consistently achieve 99% or better success requires the thoughtful

coupling of sophisticated computer-supervised material-transport systems with a work force that is highly focused on performance. We can process an order with more than 100 line items, from the time it is released until it is on our dock and ready for shipment, in a matter of minutes. This capability is critical in serving a customer base that needs and demands product quickly. Time is money; if we can save our customers time, we are saving them money.

When the Internet boom hit in the late '90s, many people said that paper catalogs were obsolete and everything would take place online. Digi-Key was one of the few companies that used the Web to complement the catalog. Did you foresee how the Web and catalog could complement one another?

A We intuitively believed that the catalog and the Internet would work best in tandem. But, not being particularly trusting of intuition, we tested serving customers with and without catalog support. To our amazement, our intuition was correct. We also test to determine the optimum catalog frequency.

Lately, the search functions in the Digi-Key Web site are so good that I will find a part on the Web, and then see what page in the catalog the part is on to view other parts in the same family. I often prefer leafing through my paper catalog versus the Web site.

A Many of our customers share your preference for the Digi-Key catalog in certain circumstances. I don't foresee the Digi-Key catalog as ever going away. We currently print more than 5 million catalogs a year.

Do manufacturers now realize that having their line in Digi-Key is not a tactical sales decision but a critical strategic-marketing decision?

A Many manufacturers have come to the conclusion that Digi-Key is more that just a sales channel for them. They know that, for many design engineers, we are the critical resource—the "one stop"—that offers the greatest breadth and broadest selection of board-level components for off-the-shelf delivery.

Studies show that most design engineers recognize that dealing with Digi-Key gives them the lowest total cost of acquisition. What do you think?

A Many design engineers find it more cost-effective to purchase from Digi-Key than from suppliers, where they must endure a time-consuming interrogation to qualify for "free" samples. They know that their time is worth more than the samples and that they can easily cost-justify buying from Digi-Key.

Digi-Key handles National Semiconductor's sample program. Would you offer that service to other manufacturers? I also know that National Semi patched Digi-Key into its Ariba ordering system and now uses an internal ordering system that still hooks into Digi-Key. Do you offer this service only to big companies?

A Digi-Key works with many suppliers and many customers on many unique initiatives and in many unique ways and makes every attempt to accommodate their specific needs and wants. Digi-Key typically evaluates the pros and cons of each initiative or suggested operational modification and makes a decision on that basis.

What is the philosophy behind your Web site?

A Digi-Key's Web site is the product of evolution and discipline. We introduced it in 1997, and it is going through constant change. But the underlying philosophy to keep it simple has endured and continues to endure since its inception, and it takes discipline. A number of Digi-Key specialists are now working to create the ultimate searchable database. Customer suggestions have driven the site as it exists today. Some suggestions are in the form of complaints, and some are slightly more constructive.

United Parcel Service and Federal Express run special trucks and planes from Digi-Key to their hubs. Has that Minnesota weather ever caused the company to miss a shipment?

A Surprisingly, weather rarely impacts logistics, even in Minnesota. My guess is that overall service from our product-distribution center in Minnesota, because of the strength of its operations, meets or exceeds that of the best in the industry, regardless of geographic location.

You are about to conquer Europe and Asia in the same way that you did America. Will you outsource your order-taking to India? How do you deal with all the European barriers to entry, such as customs and duty fees?

A Digi-Key's international-sales initiative is growing rapidly. We are now comfortable with our decision to ship all orders from Minnesota but will continue to evaluate this decision on a regular basis, as we do all aspects of our operation. We have no plans to outsource any support to India and have never seriously considered it.

Customs and duties to Europe and Asia are typically low or nonexistent, and we don't see them as barriers to Digi-Key's success in either region. We currently offer our catalogs and Web sites for Europe in English, German, and French, with additional language translations to follow. We also offer live local support in Europe in several languages, with additional languages to follow.

The situation is similar in Asia, where we currently offer our catalog in traditional Chinese, simplified Chinese, Korean, and Japanese and also offer live local support in the key languages of Asia.

Have the ROHS [reduction-of-hazardous-substances] regulations caused you any grief?

A Digi-Key embraced ROHS early and is currently the leader in terms of breadth of ROHS-compliant products available for off-the-shelf delivery. Today, we stock nearly 140,000 unique ROHS components, far more than any of our competitors.

—by Paul Rako

RAQ's

Rarely Asked Questions

Strange but true stories from the call logs of Analog Devices

Resourcefulness, or Why Good Engineers (and Good Cooks) Use the Best Available Components

Q. *Which (amplifier/switch/converter) is best in a (mobile phone/strain gauge/medical ultrasound system)?*

A. Whichever is most efficient and cost-effective, no matter what it was originally designed for. When choosing an IC, ask what it does, and how well it will work, not whether it was designed for the type of system you are building.

While at the butcher's last week to buy ground beef for spaghetti bolognaise, I noticed some cheap stewing venison. So I purchased some, minced it, and made a delicious "spaghetti alla salsa di cervo." Later that week "chile con venado" was equally successful.

As engineers we must use the best available resources to design our systems. The word engineer is derived from the Latin "ingenium," which means ingenious, i.e. resourceful. Bolognaise and chili recipes specify beef, but venison has less fat and improves the flavor.

It is unwise to avoid an IC because its data sheet does not specify a particular application. Recently, we conducted research to find out why some companies that bought a lot of our analog switches for handsets did not buy one particular type which was more ideally suited in both performance and price. The switches that were selling listed "handsets" among possible applications, but the data sheet of the better suited switch did not, so it was not considered. This is not good engineering.

Asking fundamental questions helps us to choose components (ingredients) which, though not intended or specified for our application, are actually well-suited to it.

For example, a useful component that can be easily overlooked is the AD8210. Described as a "current shunt monitor" its data sheet lists its applications as "current sensing" followed by six automotive current sensing tasks. Potential users can be forgiven for assuming that the AD8210 is simply an automotive current sensor. In fact, the AD8210 is an in-amp with a CMR from 0 to +65V when operated with a single +5V supply. Developed for automotive high-side current measurement, it is valuable wherever a small signal with a high positive common-mode voltage and reasonably low source impedance must be measured. It has been used successfully in industrial instrumentation, avionics, battery chargers and innumerable other applications, but only by engineers who see past the description on its data sheet.

To learn more about in-amps and reading data sheets,
Go to: http://rbi.ims.ca/4935-695

Contributing Writer James Bryant has been a European Applications Manager with Analog Devices since 1982. He holds a degree in Physics and Philosophy from the University of Leeds. He is also C.Eng., Eur.Eng., MIEE, and an FBIS. In addition to his passion for engineering, James is a radio ham and holds the call sign G4CLF.

Have a question involving a perplexing or unusual analog problem? Submit your question to:

raq@reedbusiness.com

BY BONNIE BAKER

Changing the PWM signal to dc

You can generate a variable dc reference voltage with a DAC or with a PWM (pulse-width-modulation) signal from your controller (**Figure 1**). The accuracy of the voltage source that the PWM/lowpass-analog-filter combination generates is as accurate as your onboard timer, filter operational amplifier, and power-supply voltage. If your controller has 64 divisions available for the PWM function, this function's accuracy can be as low as 78 mV in a 5V system.

With a microcontroller PWM generator, the clock sets the fundamental frequency. You then adjust the duty cycle by changing the ratio of T_{ON} to T_{OFF}. T_{ON} is the high time of the PWM signal, and T_{OFF} is the low time. $T_{ON}+T_{OFF}=T$, where T is the time of one cycle (or period) from the PWM.

The number of divisions (K) that your clock can produce during the PWM period determines one part of the accuracy and granularity of your PWM reference. The highest granularity that you will get out of your adjustable voltage reference is 1/K of your full-scale range. Based on T, the number of time divisions in the period, the ideal number of bits, or the resolution, of this DAC is:

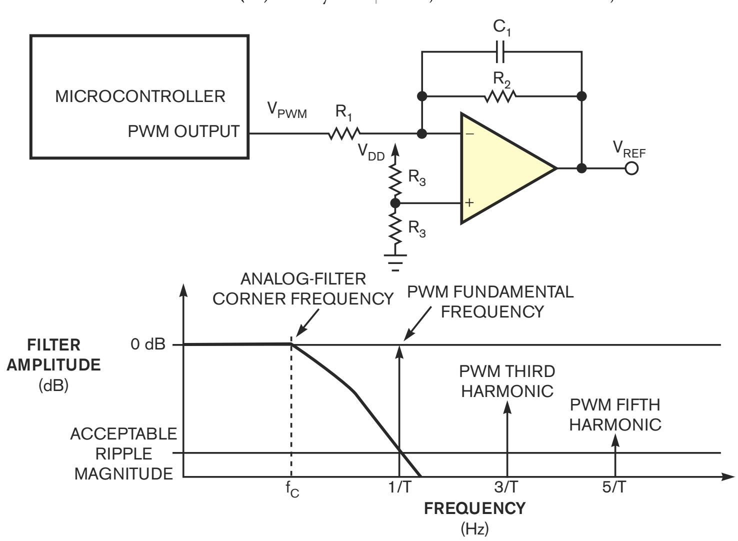

Figure 1 A hardware implementation of a PWM voltage reference uses a controller to generate the PWM signal (a). The analog, first-order, lowpass filter changes the PWM signal to a dc voltage. The primary frequency generated at the output of the PWM generator in the FFT plot is equal to 1/T, where T is the number of time divisions in the period (b). When designing the analog lowpass filter, the fundamental-frequency (f_C) response dominates the calculations and results.

DAC resolution=log (K)/log (2) in bits.

Having an analog filter after the PWM pulse produces a dc voltage, V_{REF}. This voltage's value depends on the ratio of T_{ON} to T_{OFF} and the power-supply voltage. If the PWM signal is on more than it is off, the output voltage after the filter in an inverting configuration will be below midscale, where midscale equals $V_{DD}/2$.

If you properly filter the PWM signal on the controller's output port, the errors in this system are the quantization error from the controller clock, the I/O gate's output-swing range, the lowpass filter's ripple rejection, and any offset errors and output-swing limitations of the lowpass-filter amplifier circuit. In **Figure 1**, the FFT plot separates the PWM's square-wave response into its equivalent frequencies. **Figure 1** also shows a lowpass filter's frequency response.

The calculation of the analog filter's single-pole corner frequency for this circuit is: $f_{C(FIRST\text{-}ORDER\ FILTER)}=f_{PWM}/\sqrt{((10^{-ASTOP/20})^2-1)}$.

If you need your voltage reference to remain stable under transient conditions, you may want to increase the filter corner frequency or filter order. In this case, a higher order filter is a good alternative, because you already have an amplifier in the circuit. Designing active analog filters is easy if you use the free lowpass-filter software from various op-amp manufacturers.

With the design equations in this column, a PWM, and an op amp, you can design a DAC that generates a dc reference voltage. This design's frequency-limiting factors are the clock speed of your controller's fundamental PWM signal and the analog lowpass filter's cutoff frequency. If you want to improve this system's frequency response, you can use a faster clock without compromising the PWM/DAC, or you can use a stand-alone DAC. Using a stand-alone DAC may be attractive if your application requires precision. EDN

Bonnie Baker is a senior applications engineer at Texas Instruments. You can reach her at bonnie@ti.com.

PRY FURTHER AT EDN.COM

Go to www.edn.com/060928pry for expanded coverage of the HD-A1.

PC-centric architecture speeds HD-DVD to market

A 2.5-GHz Pentium 4 and standard Intel (www.intel.com) core logic mean that the HD-A1 could essentially function as a PC. In fact, the system employs the Linux operating system. The standardized design may have reduced time to market, but it also resulted in a sluggish user interface relative to most consumer products.

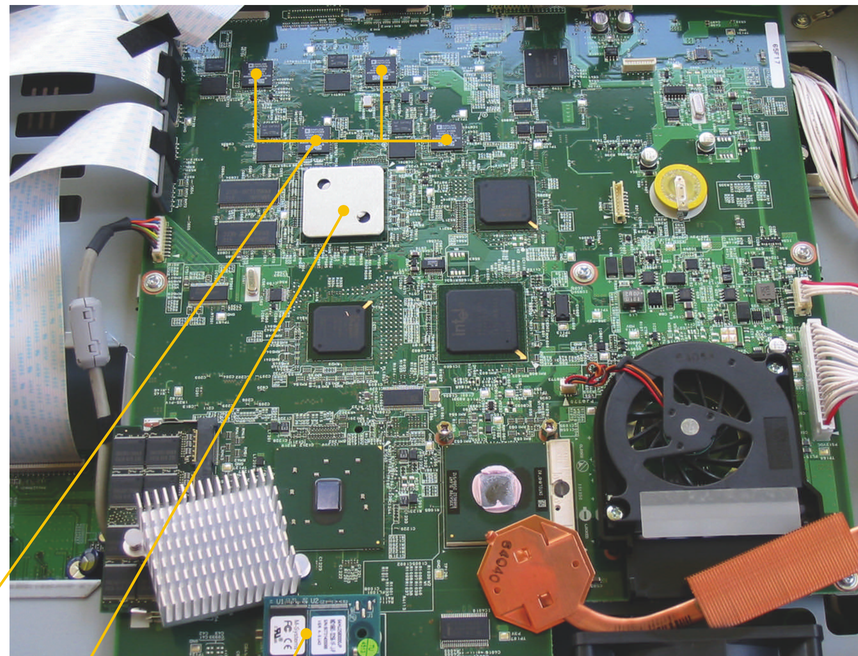

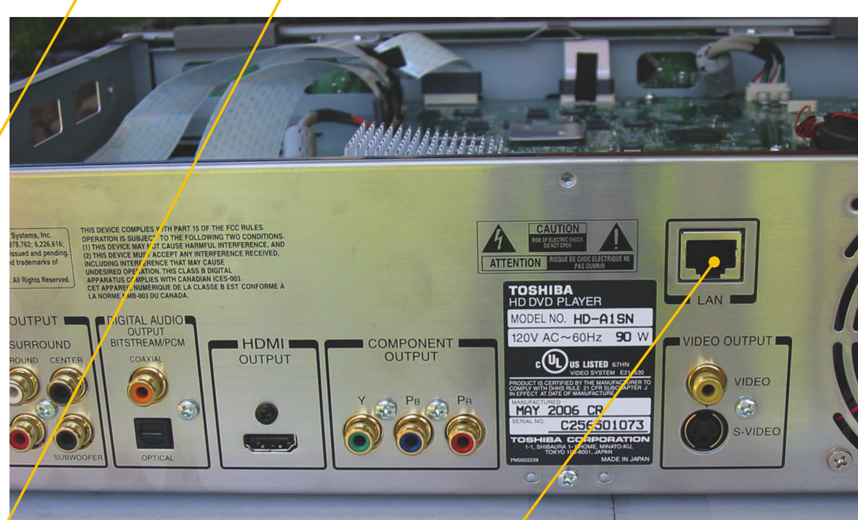

Four Analog Devices (www.analog.com) SHARC DSPs, each with dedicated memory, handle high-end audio processing, including support for Dolby Digital Plus, Dolby TrueHD, and DTS-HD.

A Broadcom (www.broadcom.com) BCM7411 handles video-decoding tasks, supporting the MPEG-2, H.264, and Microsoft (www.microsoft.com) VC-1 video formats. The IC also handles audio decoding, although as mentioned, the HD-A1 depends on DSPs to support high-definition-audio formats.

A module with flash memory and an M-Systems (www.m-systems.com) DiskOnChip IC connects to the system using a USB link, serving in essence as a bootable drive and operating-system store.

An Ethernet connector provides Internet connectivity—not only allowing DVDs to support enhanced downloadable content but also enabling firmware upgrades.

THINK Semiconductors.

THINK Customer Commitment.

THINK Infineon.

YOUR FOCUS is on designing innovative products with a partner you can trust. Our focus is dedicated to providing you smart chip solutions to help you succeed.

TODAY WE DELIVER innovative semiconductors in the areas of energy savings, mobility, networking and security. In energy savings, our power ICs improve efficiency up to 30% in consumer goods and electronic products and we lead in semiconductors for hybrids. Our mobility solutions keep people connected to data and entertainment anywhere, anytime. Our strength in networking ICs translates into innovative designs for the digital home. And in security, we provide trusted chips for computers, credit cards, phone SIM cards and personal identification. This customer-centric way of doing business has made us one of the world's top semiconductor companies with leading positions in automotive electronics and communications ICs.

WHEN YOU think of innovative designs, think about partnering with the company committed to your success. Think Infineon.

www.infineon.com/us/commitment

Infineon

Never stop thinking

Renesas Technology's dedicated register banks minimize interrupt response time to optimize the control performance of high-end, real-time applications.

In high-end industrial applications, real-time motor control by microprocessors is essential, and the need for prompt exception handling is pervasive. Can it really be possible to execute real-time interrupt-driven tasks quickly without compromising overall efficiency?

It can, thanks to a dedicated set of 15 register banks in Renesas' new SH-2A CPU core. Unlike conventional CPUs, which respond to interrupts by saving register content to embedded RAM or other stack memory, the SH-2A takes a much faster approach. Internal banks of special-purpose registers, connected to the CPU by a dedicated 128-bit bus, eliminate the need for processing software to store register content and the time required to execute it. Internal register data is instantly stored in the special interrupt register banks by hardware. As a result, the SH-2A has all but eliminated response time and reduced interrupt response time to one-sixth of what was previously available on a RISC chip.

Moreover, the SH-2A's superscalar architecture executes multiple instructions simultaneously, delivering high performance at relatively low clock speeds and low power consumption. The result is superior real-time control processing without pipeline delays.

By offering solutions for optimum efficiency to industry, Renesas is creating new infrastructure opportunities for the expanding world of ubiquitous networking.

SH-2A Interrupt Latency Improvement

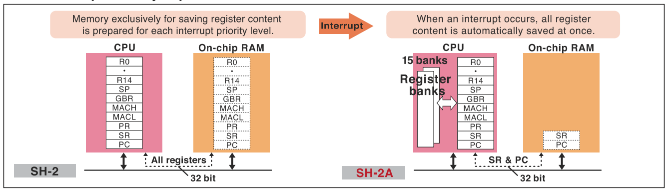

SH-2A Reduction of Interrupt Response Time

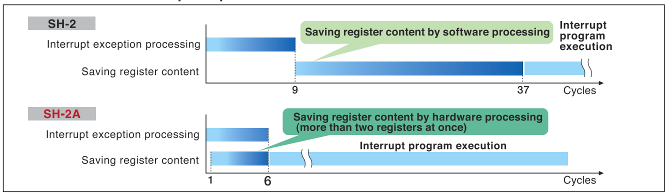

for Industry

Renesas Technology Corp. / www.renesas.com

Everywhere you imagine. RENESAS

POWER designer

Expert tips, tricks, and techniques for powerful designs

No. 114

Best Layout Practices for Switching Power Supplies

— By L. Haachitaba Mweene, Applications Manager

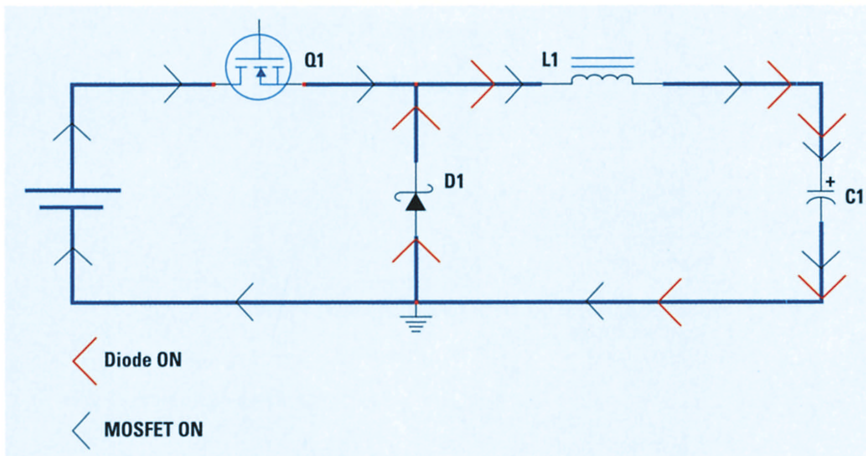

Figure 1. High di/dt Current Loops

Typical power supplies consists of a mixture of power components handling switching voltages and currents of large value and amplitude, and small signal components handling low-level analog signals, all in close proximity. Laying out a power supply board entails positioning and routing the components in such a way that the high-power signals do not corrupt the low-power signals and cause poor performance. A poor layout will lead to the generation of unwanted voltage and current spikes which will cause not only noise to appear on DC voltages in the supply, but also EMI to radiate to adjacent equipment. Thus proper layout techniques are critical to achieving optimal performance of a power supply. This article describes the most important of these techniques.

Placing the Power Components

After importing a power supply schematic into a PCB editing environment, deciding where and how to place and route many discrete components on the board can be confusing.

NEXT ISSUE:
Single-Chip FPGA Power Solutions

National Semiconductor
The Sight & Sound of Information

High Power Density Switching Regulators Deliver Up to 2A Output Current

Tiny LM2830/31/32 Step-Down Regulators Minimize External Components and Shrink Footprint

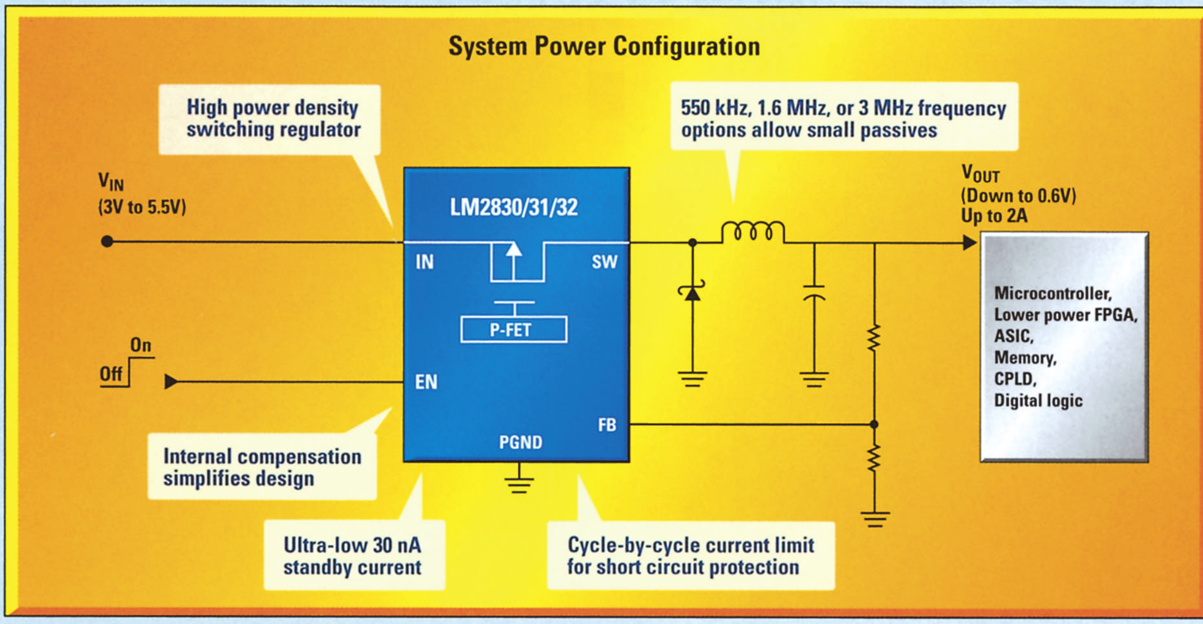

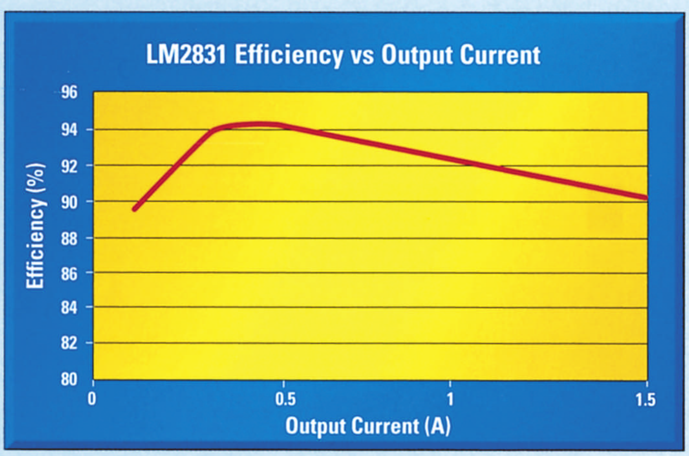

AVAILABLE LEAD-FREE

Product ID	I_{OUT}	Packaging
LM2830	1A	SOT23-5, LLP-6
LM2831	1.5A	SOT23-5, LLP-6
LM2832	2A	eMSOP-8, LLP-6

Ideal for use in multimedia set-top boxes, USB–powered devices, DSL modems, and hard disk drives

For FREE samples, datasheets, and more information, visit

www.national.com/pf/LM/LM2830.html
www.national.com/pf/LM/LM2831.html
www.national.com/pf/LM/LM2832.html

Best Layout Practices for Switching Power Supplies

Most power supplies are laid out on multi-layer boards with four copper layers or more. Most of the board space will be occupied by the power components: input capacitors, MOSFETs, current sense resistors or transformers, rectifiers, inductors, and output capacitors. These components will pass large currents and require thick traces to connect them together. They should be laid out first.

First loops of high di/dt, where large switching currents circulate, should be identified and made as tight and compact as possible to minimize stray inductance that will otherwise lead to the generation of unwelcome voltage spikes. *Figure 1* shows how to identify these loops. In the figure, the small black arrows indicate how the current circulates when the MOSFET is on. The big red arrows indicate the current loop for when the diode is on. All the paths which have either a black or a red arrow (but not both) are the high di/dt paths.

Source currents and their return paths should flow one on top of the other or next to each other to minimize the areas of the loops they form and reduce the generation of magnetic interference. Input power should be taken by the switching circuitry from directly across the input capacitors. Similarly, the load current should be taken from directly across the output capacitor.

Circuit nodes should be sized according to the magnitude and nature of the current that passes through them. High impedance nodes with high di/dt, such as the switch node (the junction in many topologies where the MOSFET, the rectifier, and the inductor meet) should be as small as possible while being adequately large for the current flowing through them. Minimizing the size of such nodes minimizes the EMI generating area. Low impedance quiet nodes, such as ground or the output, should be made as large as possible.

Copper Thickness

The traces and copper pours carrying current from one power component to the next should be made adequately wide.

An approximate formula for the minimum trace width required to carry a given current which is accurate over a current range of 1 to 20A is

$$T = \frac{2}{CuWt}\left(-1.31 + 5.813I + 1.548I^2 - 0.052I^3\right)$$

where T = trace width in mils; I = current in Amperes, and $CuWt$ = copper weight in ounces. The formula assumes that the current causes a temperature rise of 10 degrees Centigrade in the traces.

Using this formula, the minimum trace width for a current of 1A with 1 oz copper is 12 mils; for 5A, ½ oz copper it is 240 mils; and for 20A, ½ oz copper it is 1275 mils. If space allows, and especially where switching currents flow, these widths should be increased. Design goals of 30 mils per amp for 1 oz copper and 60 mils per amp for ½ oz copper should be striven for. Copper pours or floods should be used to connect the high current paths. Pours on multiple layers connected together with vias should be used for currents in excess of 10A.

Placing the Analog Components

Analog control components should be routed last because they take up little space and only need thin traces. One way to organize them is to create component subgroups by function and route the subgroups. For example, all the components that make up the feedback compensation network of the supply can be one subgroup. The bypass capacitors, soft-start capacitor, and frequency-setting resistor of the PWM controller can make up another subgroup. These subgroups typically connect to the PWM controller (or another IC). The subgroups should be placed as close to, and routed as directly as possible to the pin they connect to on the IC.

First 7V to 75V Input, 2.5A Buck Regulator

LM5005 Architecture Simplifies Design and Pushes the Performance Barrier

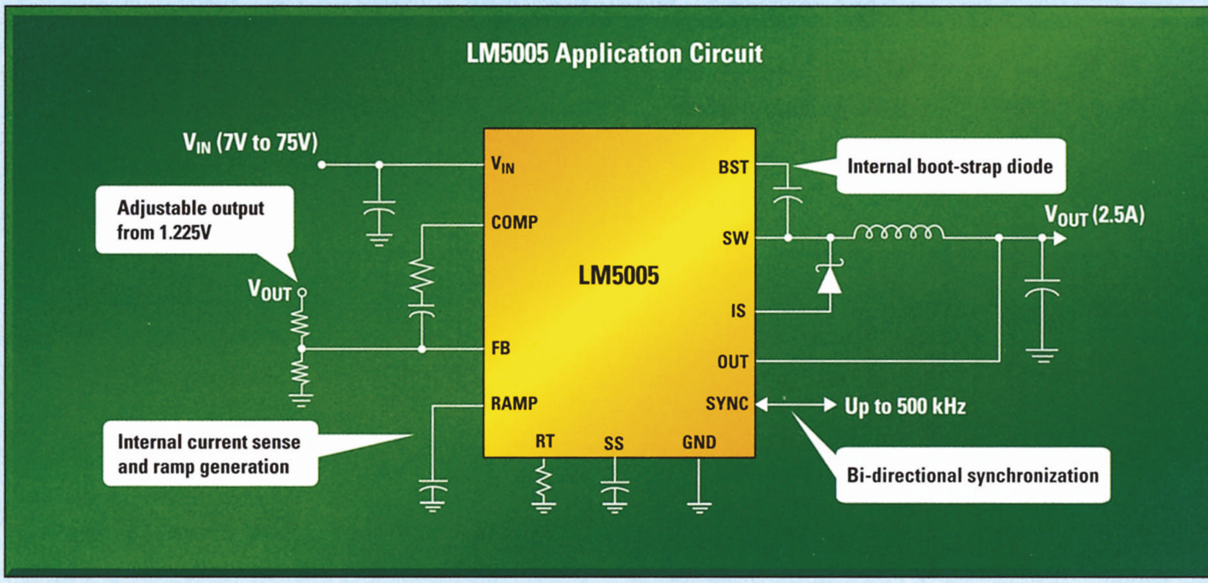

LM5005 Features

- Integrated 75V power MOSFET supports load currents up to 2.5A
- Unique, easy-to-use emulated peak current mode control topology enables high frequency operation at 75V
- Programmable switching frequency with bi-directional synchronization capability simplifies system design
- Highly integrated high-speed full feature PWM regulator reduces overall solution size

Ideal for use in automotive power systems, telecommunications, industrial systems, distributed power applications, data communications systems, and consumer electronics

WEBENCH® Online Tools

Design, build, and test power circuits in this FREE online design and prototyping environment.
webench.national.com

For FREE samples, datasheets, and more information, visit
www.national.com/pf/LM/LM5005.html

Best Layout Practices for Switching Power Supplies

This is especially true of decoupling capacitors which must be right next to the pin that they decouple. The capacitors must connect directly to the pins, and not to any ground or power planes that are electrically part of the pins.

All the big components in the circuit, such as MOSFETs, rectifiers, electrolytic capacitors, inductors, and connectors should be put on the top side of the board so they do not fall off during reflow soldering. The bottom side of the board should contain only small components which can stick to the solder flux on the board by surface tension before they are soldered.

Grounding

When routing the circuitry around the controller IC, the analog small signal ground and the power ground for switching currents must be kept separate. It is suggested to isolate the control circuitry on a local ground island, which can then be connected to the rest of the system at only one point, preferably at the input capacitor. This stratagem helps to keep the analog ground quiet. If the creation of a ground island is not possible for all or some of these components, the ground pins of the components can be connected together as a daisy-chain, but they must still be connected to the main ground at one point.

Components which straddle high impedance and low impedance nodes must be placed close to the high impedance nodes. For example, resistors setting the output voltage will see a low impedance at the output and ground connections, and a high impedance where they connect to the input of the error amplifier. The resistors must be placed as near as possible to the error amplifier. To achieve the best possible load regulation, a separate trace that carries no load current must connect one resistor directly to the load terminal of the supply, and the bottom side of the other resistor must hook directly to the chip analog ground.

Segregating Analog and Switching Signals

Power inductors/transformers, MOSFETs, and rectifiers must be placed away from the traces and circuitry with low level analog signals to minimize the amount of noise from them that the analog circuitry picks up. If power switching and analog components cannot be segregated due to space constraints, they should be placed on opposite sides of a multi-layer board and an inner copper ground plane should be used to shield the two sets of components from each other. The ground plane must be connected to the rest of the circuit in such a manner that little or no current flows in it, so that it is electrically quiet. Only then can it be considered to be a low noise reference node. All high switching currents should be arranged to flow on wide copper pours on the top layer.

For a four layer board the layer stack-up should be as follows: all the power parts should be on the top layer, as well as the copper shapes carrying the large switching currents. This layer can also have small signal components. The second layer should be a quiet ground plane with no large currents flowing through it. Layer three and the bottom layer can have a mixture of power and signal traces, with only small components populating the bottom layer. As much of the board areas possible on all layers should be flooded with copper, to improve the thermal performance.

Vias

Though it is desirable to have all the high current paths on the top layer, this is not always possible because of board size, routing, and component placement constraints. Vias must then be used to make connections between layers and to parallel the layers to allow more current to be carried between components on the board.

Multiple vias should be used to connect high current paths on different layers. Microvias should be designed to pass a current of 1A each; 14 mil

1A, 3 MHz Buck Regulator Provides Industry's Highest Power Density

LM2734—Highly Integrated, Wide Input Range, SOT-23 Regulator Reduces Total Solution Size

Features

- 3.0V to 20V input voltage range
- 0.8V to 18V output voltage range
- 1A output current
- 550 kHz (LM2734Y), 1.6 MHz (LM2734X), and 3 MHz (LM2734Z) switching frequencies
- 300 mΩ NMOS switch
- 30 nA shutdown current
- 0.8V, 2% internal voltage reference
- Internal soft-start
- Current-mode, PWM operation
- Available in Thin SOT23-6 and LLP-6 packaging AVAILABLE LEAD-FREE

Ideal for use in local point of load regulation, printers, scanners, multimedia set-top boxes, point-of-sale devices, battery-powered devices, gaming, USB-powered devices, DSL modems, RFID tags, and notebook computers

750 mA Version Also Available! See LM2736

LM2734 Typical Application

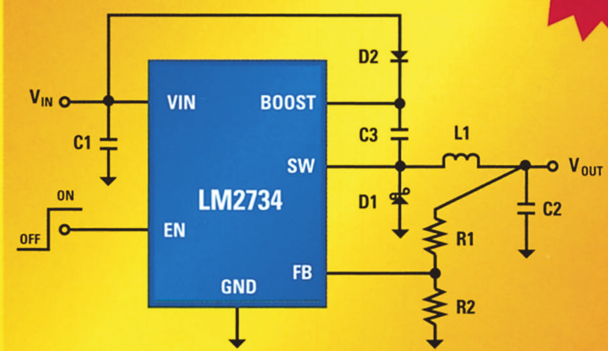

Efficiency vs Load Current
V_{IN} = 5V, V_{OUT} = 3.3V

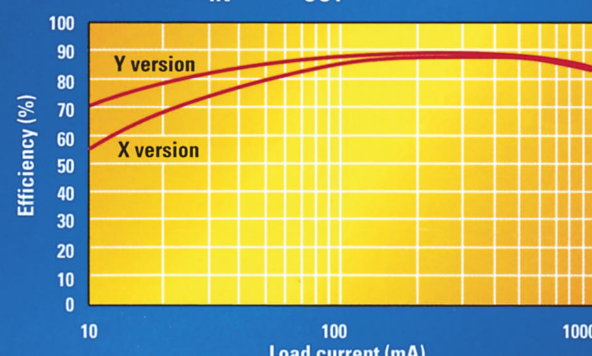

For FREE samples, datasheets, and more information, visit **www.national.com/pf/LM/LM2734.html**

Best Layout Practices for Switching Power Supplies

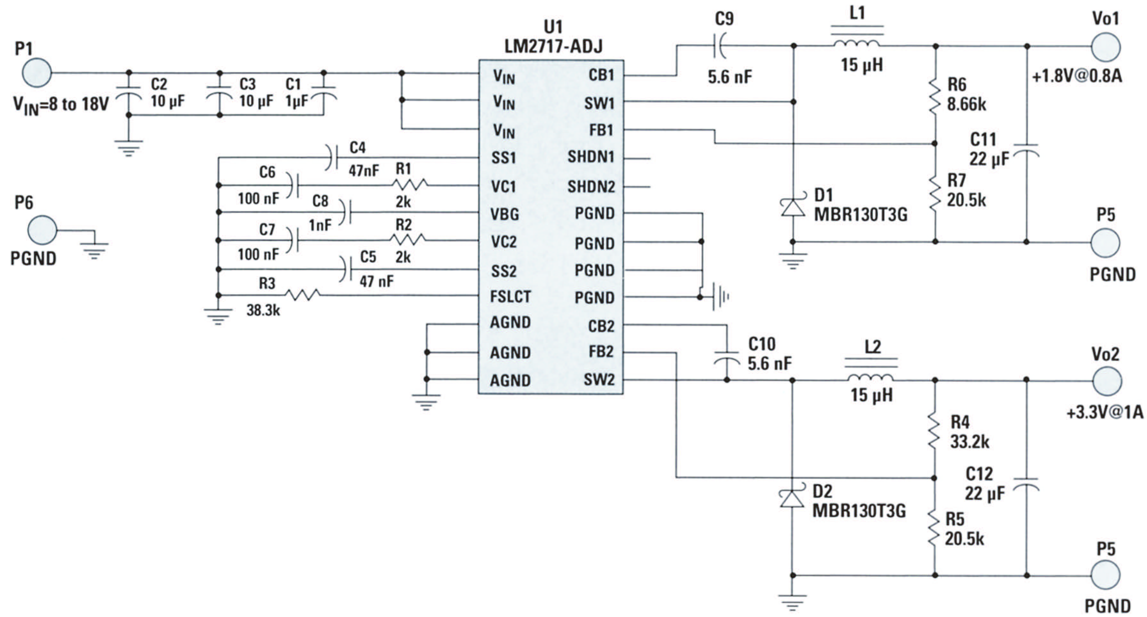

Figure 2. Circuit Schematic of a Dual Buck Converter Using the LM2717

diameter or larger vias should pass up to 2A; and 40 mil or larger vias should see no more than 5A each. Vias should be allowed to fill with solder to spread heat better, and copper alleyways in the direction of current flow should be left between them.

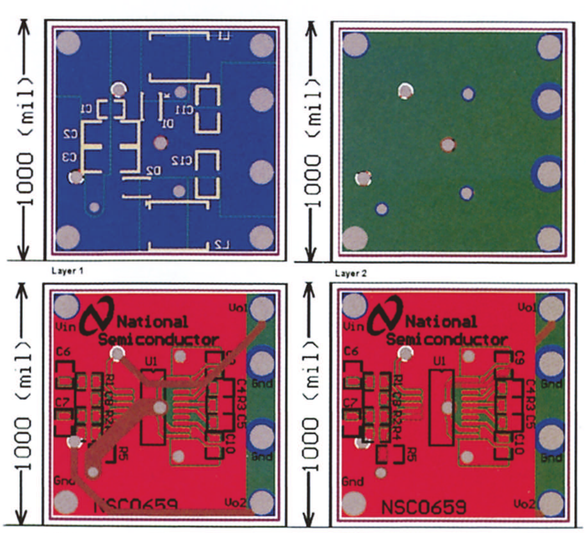

Figure 3. : A Well-Designed Four Layer Board for the Buck Schematic in *Figure 2*

Example Layout

The schematic in *Figure 2* is a dual buck converter based on the LM2717. A printed circuit board for this schematic is shown in *Figure 3* and incorporates the layout practices recommended in this article. Layer 1 contains all the power parts and thick copper pours to pass large currents. Layer 2 is a ground plane which is connected to the rest of the circuit at only one point near the input so it passes no current. Layer 3 and the bottom contain signal and power traces. All the components on the bottom layer are small. All the unused board area is flooded with copper.

More layout recommendations can be found in the references listed below, available on National's website. ■

Acknowledgement

The author wishes to thank Craig Varga for reviewing this article and providing critical background material.

References

[1] "SIMPLE SWITCHER® PCB Layout Guidelines," National Semiconductor Application Note AN1229.

[2] "Layout Guidelines for Switching Power Supplies," National Semiconductor Application Note AN1149.

Power Design Tools

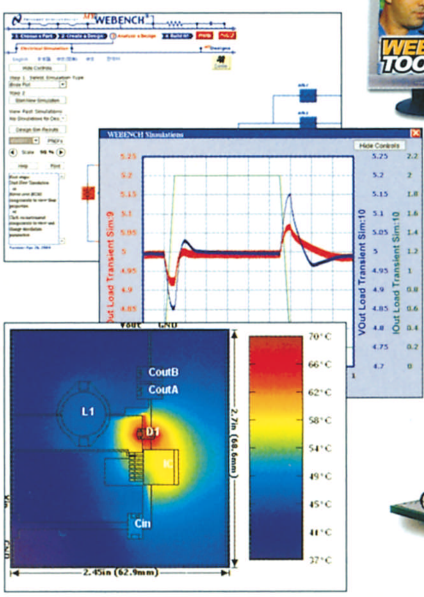

WEBENCH® Online Design Environment

Our design and prototyping environment simplifies and expedites the entire design process.

1. Choose a part
2. Create a design
3. Analyze a power supply design
   – Perform electrical simulation
   – Simulate thermal behavior
4. Build it
   – Receive your custom prototype kit 24 hours later

webench.national.com

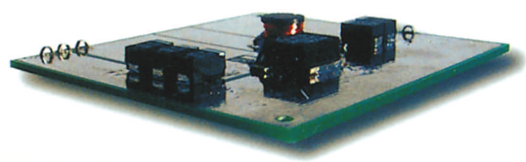

Reference Designs

National's power reference design library provides a comprehensive library of practical reference designs to speed system design and time-to-market.

www.national.com/refdesigns

Don't miss a single issue!

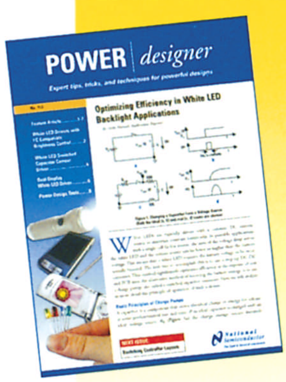

Subscribe now to receive email alerts when new issues of Power Designer are available:

power.national.com/designer

Read our Signal Path Designer® online today at:

signalpath.national.com/designer

National Semiconductor
2900 Semiconductor Drive
PO Box 58090
Santa Clara, CA 95052
1 800 272 9959

Visit our website at:
power.national.com

For more information, send email to:
new.feedback@nsc.com

National Semiconductor
The Sight & Sound of Information

550263-014

SAFETY REGULATIONS FOR MEDICAL DEVICES MANDATE ISOLATED POWER SUPPLIES AND I/O LINES. THE GOAL OF THESE REGULATIONS IS SIMPLE: DON'T ELECTROCUTE THE PATIENT. YOU CAN USE A VARIETY OF TECHNIQUES, INCLUDING TRANSFORMERS, OPTOISOLATORS, AND ADVANCED ISOLATION ICs.

Medical devices demand stringent isolation techniques

BY CHARLES H SMALL • CONTRIBUTING TECHNICAL EDITOR

According to DataBeans (www.databeans.net), medical electronics constitute a small but profitable part of the industrial-semiconductor market. In 2005, medical electronics represented 8% of the $25.9 billion industrial-semiconductor market, or about $2.1 billion. The medical-electronics industry is growing at a yearly rate of 11%, making it the fastest growing segment in the industrial-semiconductor market. The market segment is expected to exceed $4 billion in revenue by 2011.

DataBeans divides medical electronics into three major categories: diagnostics and therapy, home, and imaging. The diagnostics-and-therapy segment contributes roughly 48% of revenue, with the home segment contributing 37% and imaging contributing 15%. As a result of medical care and treatments increasingly moving away from clinical settings and into homes, the home market is growing the fastest at 12% average annual growth, followed by imaging at 11%, and diagnostics and therapy combined at 10%.

The insurance industry is instigating major changes in how patients receive medical care. To reduce costs, hospitals and clinics treat patients for a much shorter time than in the past. Insurance companies can save billions by moving treatment and monitoring to the home, and the increasing use of medical electronics in the home has created multiple opportunities for semiconductor suppliers.

COMMERCIAL SUPPLIES

One way to get an isolated power supply is to simply specify a commercially available medical supply. For example, GlobTek's 350W, 4×8-in. power supplies suit medical applications. The GT(M) 200P350 series delivers as much as 350W of continuous output power, and devices

AT A GLANCE

- Safety regulations mandate isolated power supplies and I/O lines. The goal is simple: Don't electrocute the patient.
- You can use transformers, optoisolators, and advanced isolation ICs to achieve your goals.
- Commercial isolated power supplies for medical applications are widely available.
- You can design your own power supply and submit it for certification yourself.

have regulated output voltages of 3.3 to 48V in 0.1V increments with remote sense available. Efficiency specifications are typically 85%. The supplies feature an aluminum U-channel chassis with optional top cover and are rated for 200W convection cooled or as much as 350W with forced air. The supplies have a dc power-good signal and current-share capability. A feature that is important for European markets is active PFC (power-factor correction) to IEC (International Electrotechnical Commission) 3000-3 2 Class A. The supplies accept 90 to 264V ac and have a Class B EMI filter as well as built-in protection for overcurrent, short circuit, overvoltage, and overtemperature. Footprint sizes start at 196×107×46 mm, and devices meet UL (Underwriters Laboratories and TUV (Technischer Uberwachungs-Verein, or Technical Surveillance Association) 60950/60601.1 standards. They comply with EMC (electromagnetic-compatibility) directives for consumer electronics and FCC (Federal Communications Commission) Class B applications. All models carry UL, DEMKO (Danmarks Elektriske Materielkontrol), PSE (Product Safety Engineering), and CE (Consumer Electronics) logos with reports that an independent certified laboratory generates.

Another supplier of medical-grade power supplies is Condor, which offers the GSM25 series of 25W medical switchers. The company claims that the devices are the industry's smallest 25W switchers, measuring 2.5×4×0.86 in. Conducted EMI exceeds FCC Class B and CISPR (The IEC's International Special Committee on Radio Interference) 11 Class B. Overvoltage protection is standard, and the supplies meet medical standards UL2601-1, IEC60601-1, and CSA (Canadian Standards Association)-C22.2 No. 601-1. Output voltages are 5V main with ±12, ±15, or 12V, and −24V secondary voltages. Input voltage is 90 to 264V ac and 47 to 63 Hz single phase.

Leakage current in the ground-wire connection is 50 μA (nominal)/78 μA (fault condition) measured per UL2601-1 at 132V ac/60 Hz and 94 μA (nominal)/156 μA (fault condition) measured per IEC-60601-1 at 264V ac/50 Hz. All models include built-in EMI filtering to meet the following emissions requirements: conducted emissions EN55011 Class B, FCC Class B; static discharge EN61000-4-2, 6-kV contact, 8-kV air; RF field susceptibility EN61000-4-3, 3V/m; and fast transients/bursts EN61000-4-4, 2 kV, 5 kHz.

Power-One offers the ESM ac/dc modular power-supply series. The devices' input-voltage range is 88 to 264V ac at 47 to 63 Hz. Earth leakage current per EN60601-1 at 250V ac, 60 Hz specifies 300 μA for the ESM4B, ESM4C, ESM-6C, and ESM6D supplies. The ESM4 and ESM6 series of modular ac/dc power supplies' leakage currents conform to IEC601-1. The ESM4 series is available in 400 and 600W configurations, both providing as many as eight outputs from a 2.56×5×10.63-in. chassis. The ESM6 series is available in 600 and 1000W versions, both providing as many as 12 out-

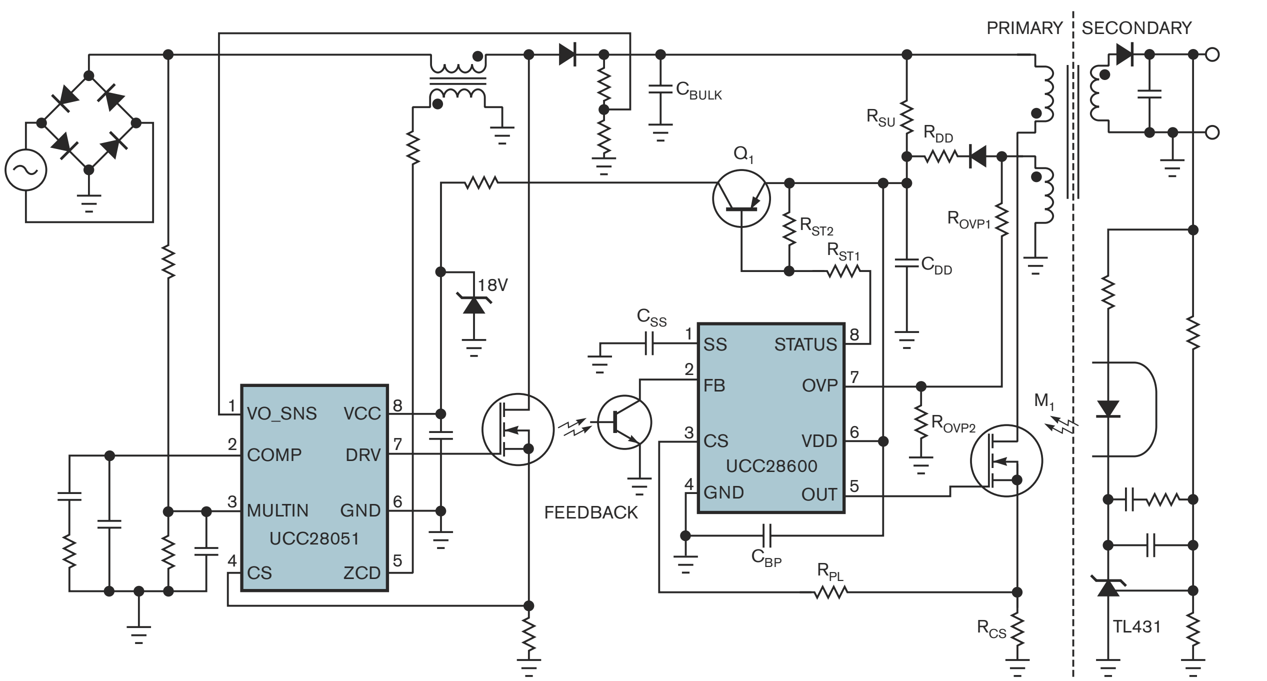

Figure 1 An optoisolator provides voltage feedback for this power-supply circuit. Note the input power-factor-correction circuit that Europe mandates for products that draw more than 60W.

Intersil Battery Charger ICs

High Performance Analog

Unshackle Your Handheld Device

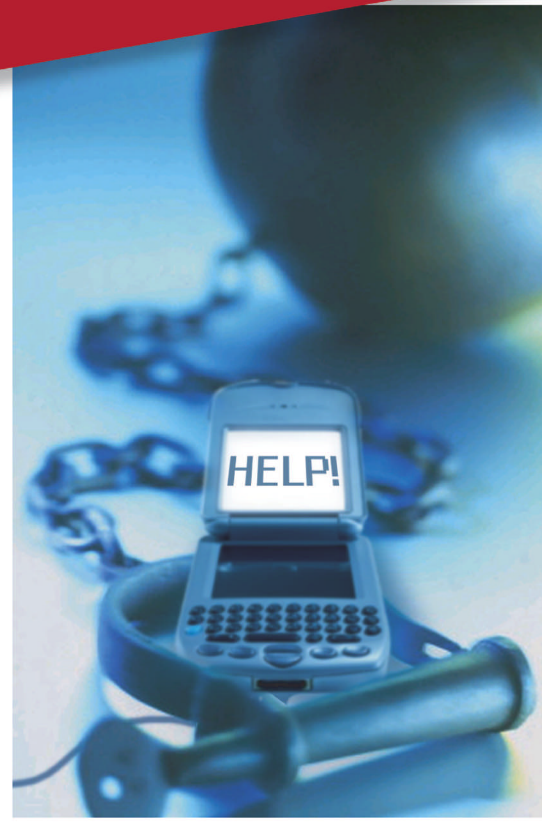

Intersil's ISL6299A is a fully integrated low-cost Li-ion or Li-polymer battery charger that accepts both USB port and desktop cradle chargers.

The ISL6299A is a low component count solution that features programmable cradle charge current, charge indication, adapter present indication, and programmable end-of-charge (EOC) current with latch. All these advanced features, along with Intersil's Thermaguard™ technology for an added measure of thermal protection, are delivered in a single 3x3 mm DFN package.

ISL6299A System

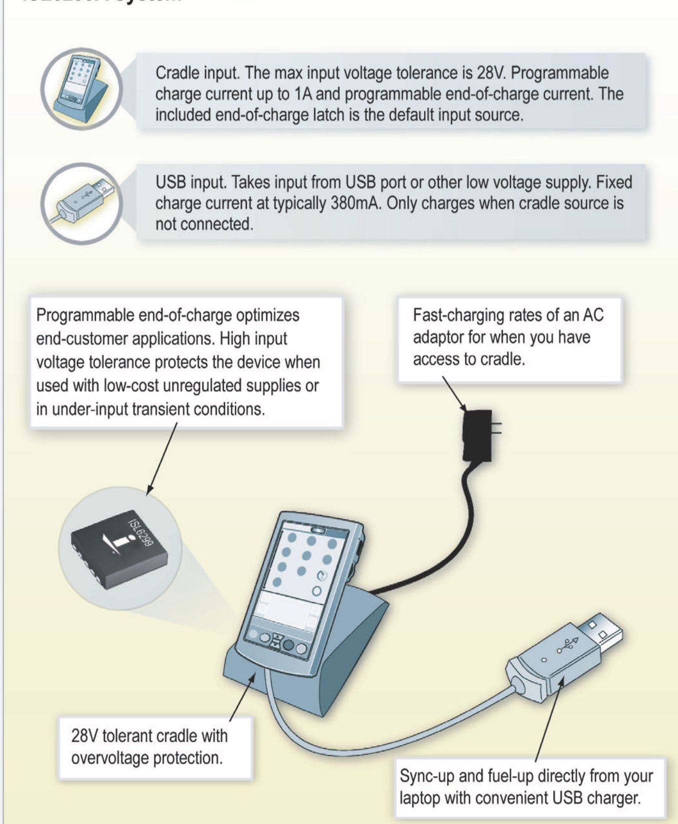

ISL6299A Key Features:

- Dual-input charger for single-cell Li-ion/polymer batteries for cradle and USB
- Low component count
- Integrated pass element
- Fixed 380mA USB charge current
- Programmable cradle charge current
- Charge current Thermaguard™ for thermal protection
- 28V maximum voltage for the cradle input
- Charge and adapter presence indicators
- Less than 0.5µA leakage current off the battery when no input power attached
- Programmable end-of-charge current with latch for cradle input
- No external blocking diode required
- RoHS compliant

Datasheet, eval kit with USB interface, free samples, and more information available at **www.intersil.com**

Intersil – Switching Regulators for precise power delivery.

intersil®

HIGH PERFORMANCE ANALOG

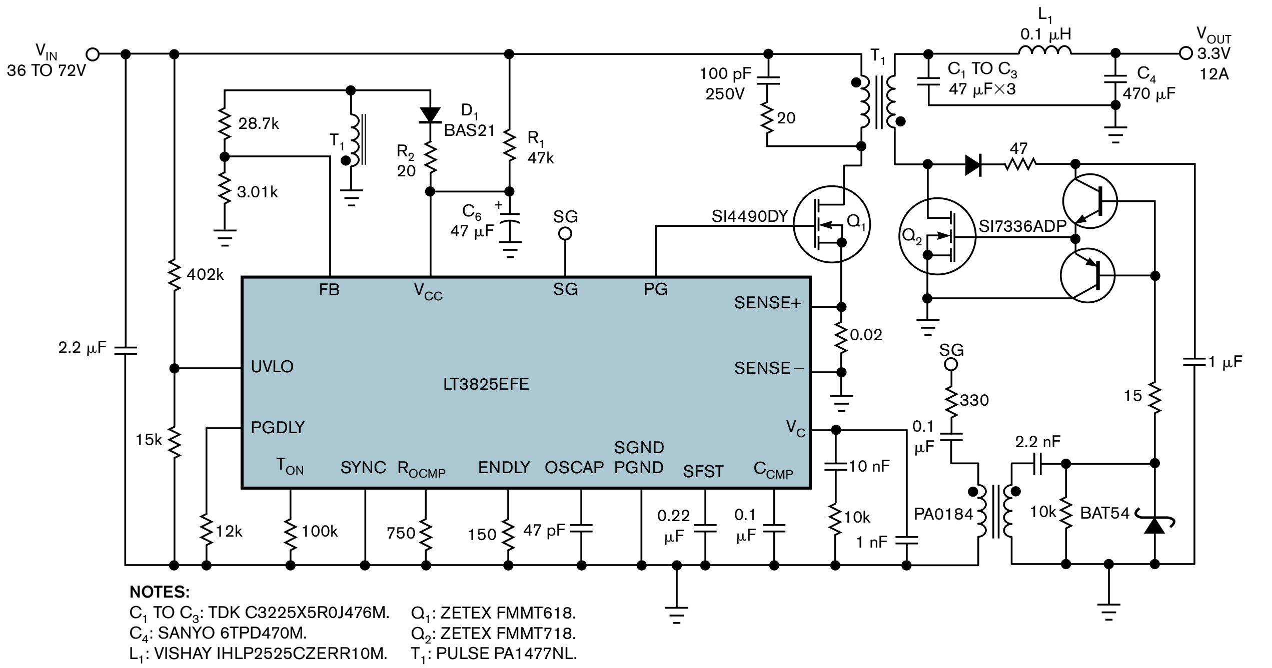

Figure 2 This transformer-isolated power-supply circuit employs transformer feedback for device powering and feedback voltage.

puts. You can configure five single-output and two dual-output modules in series or parallel to provide outputs of 1.45 to 58V dc.

ROLL YOUR OWN SUPPLY

Of course, you can design your own power supply and submit it for certification yourself. The Texas Instruments UCC28600 PWM controller has "green"-energy features to meet stringent worldwide energy-efficiency requirements. The UCC28600 integrates energy-saving features with high-level protection features. The UCC28600 incorporates frequency foldback and burst-mode operation to reduce the operation frequency at light load and no load (**Figure 1**). The UCC28600 is available in eight-pin SOIC and PDIP packages that operate from −40 to +105°C.

Texas Instruments also offers 1W, unregulated dc/dc converters in miniature packages that feature 3-kV isolation. The DCH01 family of plug-in power modules operates from a 5V input and provides single voltages of 5, 12, and 15V and complementary dual-output voltages of ±5, ±12, and ±15V. The modules, which suit medical instrumentation, provide as much as 85% efficiency and meet UL60950 specifications with 3-kV rms isolation and EN55022 Class B EMC performance. Supporting an operating-temperature range of −40 to +85°C, the modules can power TI's ISO721 and ISO721M high-speed digital isolators. In addition, the DCH01 provides highly effective point-of-load power conversion and ground-loop elimination for noise-sensitive applications.

The DCH01 uses reliable pc-board construction and comes in an industry-standard 20×8×10-mm SIP-7 package that is lead-free and ROHS (reduction-of-hazardous-substances)-compliant. The DCH01's open-frame design enhances reliability by eliminating the potting you commonly find in competitive products. The DCH01 series costs $4.25 (1000).

Not all isolated-power-supply ICs require an optoisolator. The Linear Technology LT1725 monolithic switching-regulator controller targets use in the isolated-flyback topology. It drives the gate of an external MOSFET and is powered from a third transformer winding. These features allow for an application input voltage limited only by external power-path components. The third transformer winding also provides output-voltage-feedback information, such that an optoisolator is unnecessary. Its gate-drive capability couples with a suitable external MOSFET to deliver load power that reaches into the 10s of watts.

The LT1725 has a number of features absent from most other switching-regulator ICs. By using current-mode switching techniques, it provides excellent ac- and dc-line regulation. Its unique control circuitry can maintain regulation well into discontinuous mode in most applications. Optional load-compensation circuitry allows for improved load regulation. An optional undervoltage-lockout pin halts operation when the input voltage is too low. An optional external capacitor implements a soft-start function. A 3V output is available at as much as several milliamps for powering primary-side application circuitry.

Targeting applications requiring more output current, the Linear Technology LT3825 isolated switching-regulator controller finds use in medium-power flyback topologies (**Figure 2**). A typical application is 10 to 60W with input voltage limited only by external power-path components. A third transformer winding provides output-voltage feedback. The LT3825 current-mode controller regulates output voltage based on sensing secondary voltage through a transformer winding during flyback. This setup allows for tight output regulation without the use of an optoisolator, improving dynamic response and reliability. Synchronous rectification increases converter efficiency and improves output cross-regulation in multiple output converters. The LT3825 operates in forced continuous-conduction

Intersil Switching Regulators

High Performance Analog

Need a Multiple Output PWM That Can Tackle a Wide Range of Voltages?

Now you can get true 180º out-of-phase PWM performance along with your choice of two or three regulated outputs.

Intersil's new line of wide V_{IN} PWM Controllers offers industry-leading performance and protection, along with unmatched design flexibility. So, no matter what your input voltage, switching frequency, or number of system supply voltage requirements are, we've got the right PWM Controller IC for your design.

Triple Output PWM Controller
4.5V to 5.5V or 5.6V to 24V Input Voltage →

V_{OUT1}: **Adjustable, 0.8V to V_{IN}**
V_{OUT2}: **Adjustable, 0.8V to V_{IN}**
V_{OUT3}: **Adjustable, 0.8V to V_{IN}**

Synchronized 180º out-of-phase, reducing the RMS input current and ripple voltage.

Triple Output PWM Controller
4.5V to 5.5V or 5.6V to 24V Input Voltage →

V_{OUT1}: **Adjustable, 0.8V to V_{IN}**
V_{OUT2}: **Adjustable, 0.8V to V_{IN}**
V_{OUT3}: **Adjustable, 0.8V to V_{IN}**

An adjustable overcurrent protection circuit monitors the output current by sensing the voltage drop across the lower MOSFET.

Dual Output PWM Controller
4.5V to 5.5V or 5.6V to 24V Input Voltage →

V_{OUT2}: **Adjustable, 0.8V to V_{IN}**
V_{OUT3}: **Adjustable, 0.8V to V_{IN}**

Key Features:

- Operates from wide range of input supplies (4.5V to 24V)
- 1.4MHz switching frequency (ISL6441, ISL6445) for smaller passive components or 300kHz switching frequency (ISL6440, ISL6443) for highest efficiency. ISL6442 switching frequency is adjustable from 300kHz to 2.5MHz.
- Dual (ISL6440, ISL6445) and Triple (ISL6441, ISL6442, ISL6443) regulated outputs
- Internal compensation replaces external components, freeing-up valuable board space
- Overcurrent, overvoltage, P_{GOOD} and thermal shutdown
- Out-of-phase operation to reduce input filter requirements and EMI
- Small footprint and excellent thermal resistance in 5mm x 5mm QFN package (ISL6441, ISL6443) and 20-ld QSOP (ISL6440, ISL6442, ISL6445)

Datasheet, free samples, and more information available at **www.intersil.com**

Intersil – Switching Regulators for precise power delivery.

intersil®
HIGH PERFORMANCE ANALOG

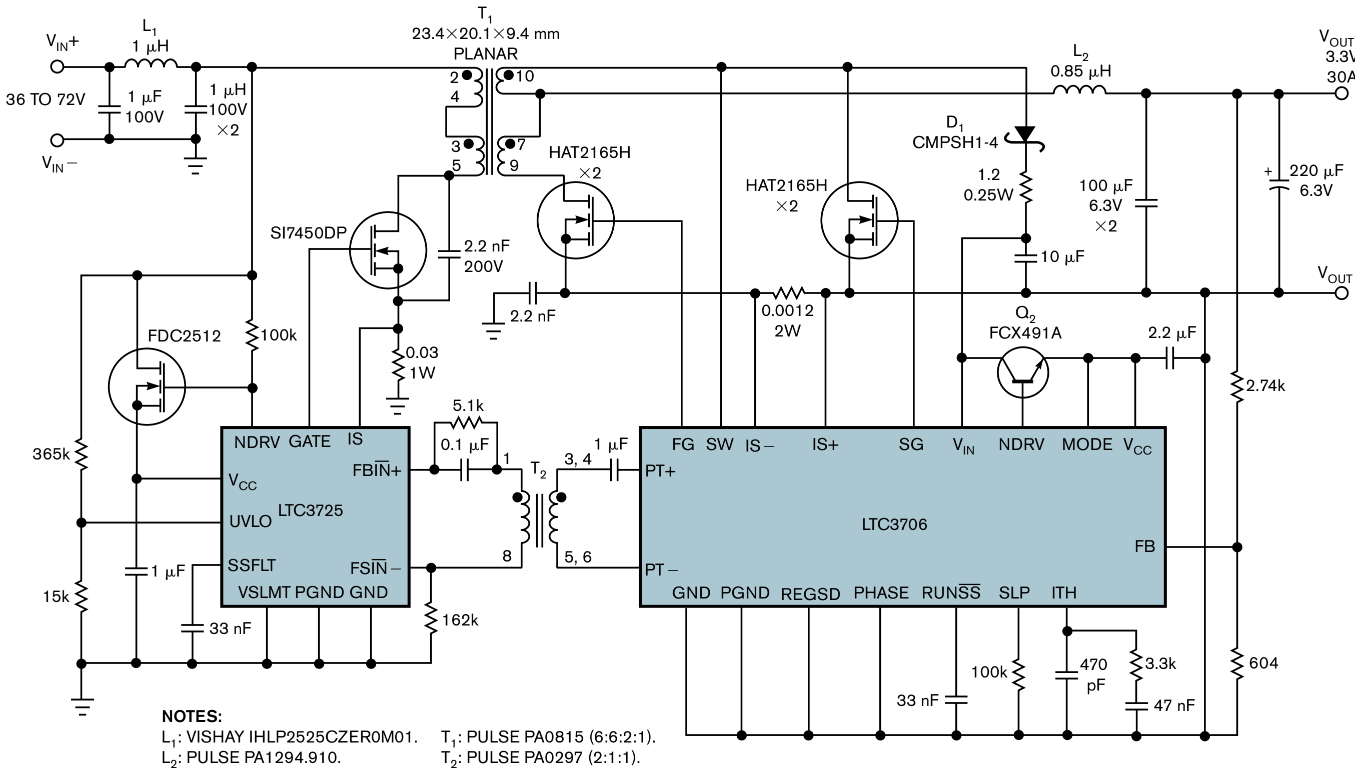

Figure 3 The LTC3706 secondary-side synchronous controller and its companion smart gate driver, the LTC3725, make it possible to create an isolated forward converter with the simplicity and performance of a buck regulator.

mode, which improves cross-regulation in multiple winding applications. Users can program and externally synchronize the switching frequency. The part also has load compensation, undervoltage lockout, and soft-start circuitry.

The circuit in **Figure 2** shows an isolated, synchronous, 48 to 3.3V flyback waveform at a 12A supply. Instead of using a parts-intensive secondary-side voltage reference and error amplifier to drive an optocoupler, it uses the primary bias winding on the flyback transformer, T_1. Proprietary feedback circuitry inside the LT3825 reads the reflected output-voltage information on this winding during the flyback pulse. The circuit then compares this voltage with a precision internal reference and obtains an error signal. You use the error signal to modulate the on-time of Q_1 in such a way as to regulate the output voltage. An important benefit of this technique is that output-voltage information instantly arrives at the controller after the switching cycle ends. In a conventional optocoupler-based design, delays of 10s to 100s of microseconds occur in the optocoupler alone, severely limiting the converter's transient response.

According to Linear Technology's Kurk Mathews, buck-converter designers have long benefited from the simplicity, high efficiency, and fast transient response that the latest buck-controller ICs make possible. These ICs feature synchronous rectification and polyphase interleaved power stages. Unfortunately, these same features have been difficult or impossible to implement in the buck converter's close relative, the forward converter, often used in isolated medical applications. The LTC3706 secondary-side synchronous controller and its companion smart gate driver, the LTC3725, make it possible to create an isolated forward converter with the simplicity and performance of a buck regulator. Many isolated supplies place the controller IC on the input (primary) side and rely on indirect synchronous-rectifier timing and optoisolator feedback to control the output (secondary). The circuit in **Figure 3** offers a more direct approach using fewer components. The secondary side uses the LTC3706 controller, and the primary side uses the LTC-3725 driver with self-starting capability. When you apply an input voltage, the LTC3725 begins a controlled soft-start of the output voltage. As the output voltage begins to rise, the LTC3706 secondary

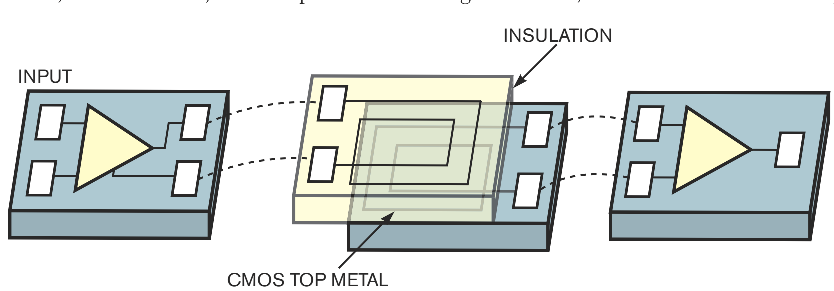

Figure 4 Analog Devices' iCoupler transformers are planar structures that use the CMOS metal layers as well as a gold layer fabricated on top of the wafer passivation.

controller quickly powers up via T_1, D_1, and Q_2. The LTC3706 then assumes control of the output voltage by sending encoded PWM gate pulses to the LTC-3725 primary driver through signal transformer T_2. The LTC3725 then operates as a simple driver, receiving both input signals and bias power through T_2.

ISOLATED I/O

In a variety of medical systems, designers face the challenge of signaling data between two points and preventing the flow of electrical current. The solution to this problem is to employ a galvanic-isolation device, which allows signals to travel between the two points but prevents the flow of electrical current.

In addition to providing isolation in the feedback loop of a power supply, optoisolators can also isolate I/O lines. Avago Technologies' fastest optoisolator is the HPCL-7723, which runs at 50 Mbps and exhibits maximum pulse-width distortion of 2 nsec and a propagation delay of 20 nsec. Avago offers single-, two-, three-, and four-channel devices. The HCN-Wxxx series comes in a wide-body package that withstands a minimum isolation voltage of 5 kV rms for 1 sec per UL 1577.

A GALVANIC-ISOLATION DEVICE ALLOWS SIGNALS TO TRAVEL BETWEEN THE TWO POINTS BUT PREVENTS THE FLOW OF ELECTRICAL CURRENT.

Analog Devices bases its iCoupler isolation technology on chip-scale transformers rather than the LEDs and photodiodes that optocouplers use. By fabricating the transformers directly on-chip using wafer-level processing, you can integrate iCoupler channels with each other and other semiconductor functions at low cost. The iCoupler transformers are planar structures that use the CMOS metal layers as well as a gold layer fabricated on top of the wafer passivation (**Figure 4**). A high-breakdown polymide layer underneath the gold layer insulates the top transformer coil from the bottom. CMOS circuits connected to the top and bottom coils provide the interface between each transformer and its external signals.

The circuitry encodes input-logic transitions using 1-nsec pulses routed to the primary side of a given transformer. These pulses couple from one transformer coil to another, and the circuitry on the secondary side of the transformer detects them. This circuitry then re-creates the input digital signal at the output. In addition, a refresh circuit at the input side ensures that the output state matches the input state even if no input transitions are present.

The ADuM240x product family is Analog Devices' first with an isolation rating greater than 2.5 kV rms. These quad-channel isolators are pin- and specification-compliant with the ADuM140x family but provide double the isolation rating at 5 kV rms. Targeting medical and other safety-critical applications, the ADuM240x isolators are certified to a working voltage of 250V rms (reinforced insulation), according to medical-equipment standard IEC 60601-1. They are also approved for working voltages reaching 500V rms, per the more general standard EN60747-5-2. These quad-channel isolators are available in three channel configurations, each with three performance grades. The ARW grade supports data rates to 1 Mbps, the BRW grade supports data rates to 10 Mbps, and the CRW grade supports data rates to 90 Mbps. All models operate from 2.7 to 5.5V over −40 to +105°C.

Texas Instruments offers a two-member family of high-speed digital isolators, featuring on-chip capacitors to enable faster data transmission with higher signal integrity. These capacitive isolators, which combine the fastest data rates with high reliability, provide six-orders-of-magnitude higher magnetic immunity than inductive devices, and they use 60% less power than high-performance optocouplers.

The ISO721 and ISO721M provide data transmission and circuit protection with isolation of as much as 560V of operating voltage or 4-kV peak overvoltage transient. The ISO721M suits applications that require fast digital-data transmission with low system noise. The ISO721 is more flexible and robust for transmitting data in noisy environments. The devices meet the standards for isolators set by UL 1577, IEC 60747-5-2, and CSA Component Acceptance Notice 5A.

By using on-chip, high-voltage capacitors, the TI isolators transmit data as much as three times faster and use less power than commonly used high-performance optocouplers. In addition, TI's capacitive technology offers immunity from external magnetic fields that frequently occur in the industrial environment and can distort signal integrity. These devices also offer high immunity against data corruption due to fast voltage transients, providing a minimum protection level of 25 kV/μsec.

TI's isolators use a semiconductor-grade silicon-oxide dielectric. This stable insulator provides proven reliability and long operational life. These requirements are critical in industrial applications, in which voltage surges can otherwise degrade device lifetime. At typical operating voltage, each device's life expectancy exceeds 25 years.

The ISO721 features TTL inputs with a noise filter, a 100-Mbps transmission speed, a typical propagation delay of 17 nsec, 2 nsec of jitter (typical), and support for 3 or 5V signals. The ISO721M, on the other hand, has CMOS inputs without a noise filter, a 150-Mbps transmission speed, a typical propagation delay of 10 nsec, and jitter of 1 nsec (typical). The device supports 3 or 5V signals. The ISO721 and ISO721M devices come in an eight-pin SOIC package and cost $1.65 (1000). An evaluation module is available.EDN

FOR MORE INFORMATION

Analog Devices
www.analog.com

Avago Technologies
www.avagotech.com

Condor DC Power Supplies
www.condorpower.com

GlobTek
www.globtek.com

Linear Technology
www.linear.com

Power-One
www.power-one.com

Texas Instruments
www.ti.com

AUTHOR'S BIOGRAPHY

Charles H Small is a contributing technical editor for EDN. *You can reach him at charles.small@att.worldnet.net.*

CIRCULATING CURRENTS: THE WARNINGS ARE OUT

UNDERSTANDING HOW TO AVOID OR MINIMIZE THE EFFECTS OF CIRCULATING CURRENTS CAN MAKE YOUR DESIGNS MORE ROBUST. EVERY ENGINEER SHOULD KNOW THE TECHNIQUES FOR NEUTRALIZING THIS INSIDIOUS PHENOMENON.

BY PAUL RAKO • TECHNICAL EDITOR

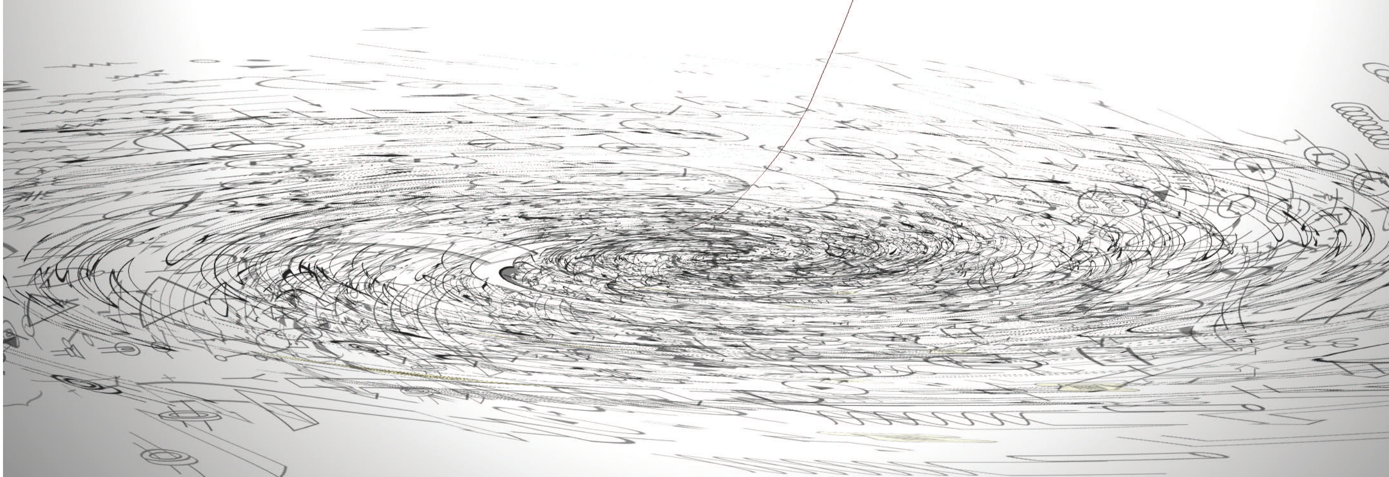

Circulating currents can wreak havoc for a design engineer, no matter whether your application is computers or communications. Some engineers lack an appreciation for circulating currents because of the schematic convention of using a ground or common symbol to show the return path for all the circuitry. Novice engineers often misinterpret this symbol as representing an ocean of zero impedance. Nothing could be further from the truth. That ground symbol represents just another wire in your schematic. If the current in the ground connection is large enough or if it changes fast enough, it generates a significant amount of voltage. That voltage might interfere with the accuracy of a power supply. That voltage can also cause measurement errors in an instrumentation application. Digital-system engineers must grapple with ground bounce. Audio buffs see the effects of circulating currents in the dreaded ground loop that causes buzz and hum. RF engineers always struggle with controlling the flow of ground currents in high-frequency-system applications. Read on to find out the cause of circulating currents, get some real-world examples, and then learn solid design principles to keep these circulating currents from ruining your design.

Even experienced engineers must remember that the ground or common symbols on schematics are just notational conveniences. A ground symbol represents just another wire, albeit a wire that has many connections. Even when the ground symbol represents a ground plane, a finite impedance will still exist, and it may interfere with the proper operation of your circuit. The key word is "impedance," not simply resistance. Resistance of the ground circuit can cause problems in situations having sensitive nodes that microvolt changes affect. A more common problem is the impedance of the circuit, or the resistance it shows over frequency. This problem should be intuitive to even a novice engineer (see **sidebar** "Impedance 101: those old, familiar impedance equations"). A 50Ω coaxial cable shows milliohms of resistance when you measure it with a DVM (digital voltmeter). But at high frequencies, the cable has the advertised impedance of 50Ω.

It might now be helpful to plug and crank a few actual numbers through the old, familiar impedance formulas to demonstrate why ground connections depend so highly on impedance—not just resistance. A capacitance of 25 pF does not sound like a lot, but, at 100 MHz, the impedance formula gives a value of 64Ω. Knowing that video-signal impedances are often 300, 75, or 50Ω should give you pause when you consider that only 25 pF of stray capacitance provides an impedance of only 64Ω. In the realm of circulating current, the inductance is often the cause of the problems. An inductance of 15 nH is a small value. An inch of wire in free space has about 15-nH inductance. Yet, at 100 MHz, that inductance has an equivalent resistance of 9.5Ω. Again, you can see that what appears to be an irrelevant stray has become a significant amount of impedance.

At first blush, these facts wouldn't trouble many engineers. They would think that, because their switching power supply has a 200-kHz clock, they need not worry about impedance. However, they are misinterpreting the fundamental frequency of operation with the highest frequencies of interest in the circuit. Fourier analysis shows that a 200-kHz square wave can have frequency components in the hundreds of megahertz. To better understand these issues, consider the relationship of capacitance and inductance to voltage and current.

Once again, plug a few real numbers into the same familiar impedance equations. Looking at the stray capacitance between a circuit node and the substrate, a value of 2 pF is not uncommon in a semiconductor. ICs often have rise times of 1 nsec. If the part operates on 5V, the rate of change is 5V/nsec, or 5 GV/sec. Multiplying this figure by the admittedly tiny capacitance still yields a current of 10 mA. Every node that slews at this rate will dump 10 mA of supply current into the ground.

This aspect of circulating currents is just one of many. For example, examine the effect of a small inductance on that same IC. If all those stray ground currents add up to 100 mA and that current appears over the same 1 nsec, then the rate of current change is 100 MA/sec. A bond wire in an IC can easily have 2 nH of inductance. That current change across the bond wire creates 0.2V. This amount can affect the logic level or transition time. Don't forget: The same transition occurs on the power rail and may cause its own problems (**Reference 1**).

With circulating currents, even small stray inductances and capacitances can

AT A GLANCE

- Circulating currents can plague power-supply circuits, audio circuits, and RF. Even IC designers must struggle with the effects of circulating currents.
- The ac impedances of your circuits are more important than the resistances when it comes to the trouble that circulating currents cause.
- The common, or ground, symbol on your schematic is just another wire—not an ocean of zero impedance.
- Power supplies have large output currents, as well as internal circulating currents. Keep the reference ground away from these nodes. Connect the supply circuit to your system at one point.
- Cutting up the ground plane usually causes more problems that it solves. However, as with all things analog, there are exceptions to the rule.
- You can avoid ground loops in audio and RF circuits with good design practices and differential-signal chains.

create large currents and voltages if the signals are moving fast enough. Because board-level Spice and other simulations often do not model these strays, the circuit performs flawlessly on the computer. On the breadboard or in production, however, the effects of circulating current can ruin your design.

Consider some real-world examples in which circulating currents can wreak havoc on your designs. Understanding the bad effects of fast-moving voltages and currents on IC designers, consider the grief that mixed-signal-system designers must endure. Digital circuits on silicon die inject amperes of current into the substrate during short transients. The same die may include delicate analog circuits, which have no noise margin. Worse yet, if the simulation and verification tools available do not model the strays to ground or the intercircuit strays, the first silicon will not work properly. It will then take many months if not a year to figure out the problems and redo the circuit. "Constraint rules that allow the communication of critical design intent from upstream circuit designers to downstream layout engineers are critical," says Anthony Gadient, marketing-group director for Virtuoso at Cadence Design Systems. "These constraints will allow designers to unambiguously communicate critical information, such as which devices need to match or which parts of a circuit need special isolation from noise injected into the substrate."

Figure 1 shows the internal circuits in an IC and the inductance of the bond wire, as L_3 represents. The connections for the power and ground also have their respective inductances. **Figure 2** shows the effect of ground bounce on digital-logic signals. Ground currents react to those stray inductances, causing overshoots, undershoots, and ringing. When fast-moving currents react to the ground pin's bond wire, voltage appears on a ground pin. This problem prompted some manufacturers to revise their packaging to minimize the stray inductance in the ground circuit.

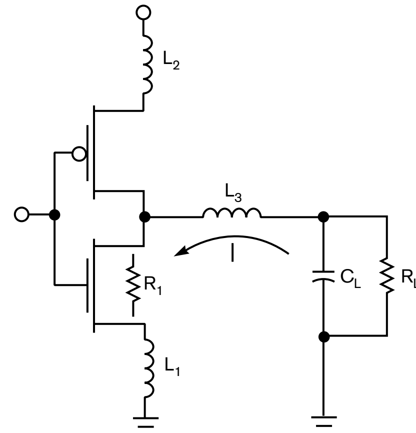

Figure 1 L_1, L_2, and L_3 represent the bond wires inside an IC. Fast-changing currents in these wires create appreciable voltages (courtesy Fairchild Semiconductor).

POWER SUPPLIES

Circulating currents cause problems in power supplies when the large output currents cause a difference between the actual output voltage and the feedback voltage as well as when the large output-ground currents move the power-supply chip's analog ground off its true ground. They also cause problems when the large output currents in the supply's power and ground rails interfere with the delivery of an accurate voltage to the systems and when the circulating currents are in loops that radiate excessive EMI into the system or cause the system to fail FCC (Federal Communications Commission)-compliance tests (**Figure 3**).

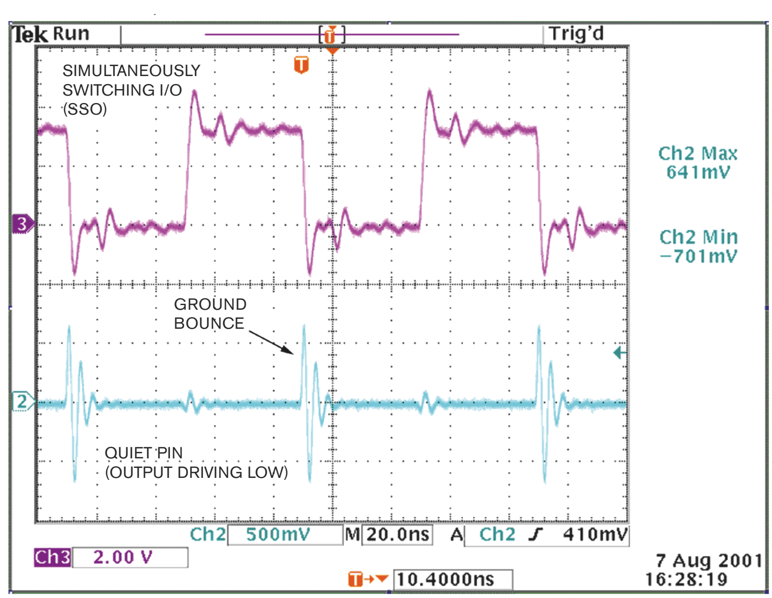

Figure 2 The ground bounce in the lower trace is large enough to cause false triggering of logic circuits (courtesy Altera).

Stackable. Scalable. Flexible.

DC/DC Controller Boosts Efficiency

The TPS40140 turns power supplies in data center and telecommunication equipment into fully scalable, stackable power systems with greater load-handling capability and maximum efficiency. This unique PWM buck controller offers the simplicity of a stand-alone dual or two-phase controller with the ability to "stack" multiple devices together, creating a high-density power supply. Generating from 10 A to 320 A of output current, true interleaved operation enables maximum efficiency up to 16 phases.

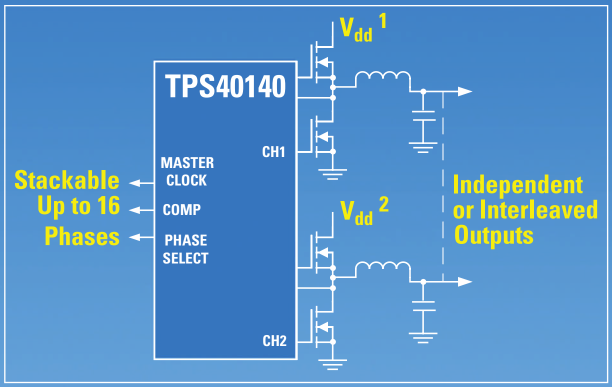

High Performance. Analog. Texas Instruments.

For datasheet, evaluation module and samples visit:
www.ti.com/tps40140

Technology for Innovators™ | Texas Instruments

In the **figure**, the feedback resistors are not to the left of the output capacitor because this placement would put the feedback sampling point along a trace with circulating ac currents. Instead, the feedback resistors sit to the right; alternatively, you could use a four-wire Kelvin connection of the feedback resistors to the output capacitor's pads (**Figure 4** and **references 2** and **3**). You must also be sure to reference the internal voltage reference in the IC to the proper ground. If the reference ties to ground near the input capacitor or the synchronous switch—or diode in an asynchronous design—large,

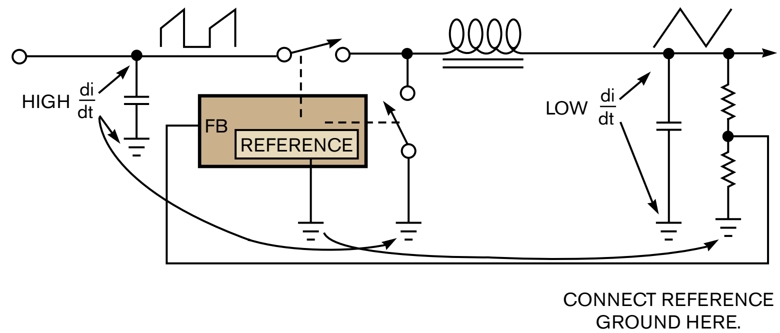

Figure 3 Buck regulators have large circulating currents in the input capacitor and the synchronous switch ground. Keep the reference ground away from those nodes.

IMPEDANCE 101: THOSE OLD, FAMILIAR IMPEDANCE EQUATIONS

Resistance is the well-known property of a resistor, but impedance is a frequency-dependent property that involves the capacitance and inductance of connections, as well. Every engineer should have committed to memory the equations that relate current and voltage in reactive components. Because voltage and current in inductors and capacitors are time-variant phenomena, you must invoke calculus to represent the behavior of these types of components. In calculus, the voltage of a node depends not on the dc current but on the time-variant change in the current, which you express using the differential di/dt, the change in current over time.

The basic nature of capacitors and inductors should allow a forgetful engineer to derive the relationship by inspection. The larger the capacitance and the faster the signal changes, the smaller the voltage across a capacitor gets. This situation demands that you put both the capacitive and the frequency terms in the denominator. Hence, for a capacitor, $Z=1/wC$, and radian frequency, w, is equivalent to $2\pi f$, where f is the frequency in hertz, or cycles per second. You therefore express the relationship between voltage and current for a capacitor as $Z=1/2\pi fC$. Conversely, the intuitive observation that an inductor becomes an open circuit at high frequencies provides the hint that the larger the inductance or the higher the frequency, the greater the effective resistance of an inductive circuit. Hence, for an inductor, the frequency and value appear in the numerator. The impedance of an inductor is therefore $Z=2\pi fL$.

At dc, or zero frequency, your intuitive feelings toward the inductor and capacitor become justified. The zero in the denominator of the capacitive equation means that the impedance of a capacitor at dc is infinite. Similarly, the zero in the numerator of the inductive circuit means the impedance of an inductor at dc is zero. Real-world capacitors have a leakage current that provides a noninfinite impedance, and real-world coils always have a metallic resistance that means that the impedance is not zero. Circulating currents are not dc phenomena, but currents can cause problems at dc. More problems occur at high frequencies.

Engineers analyze and conceive of circuits using the voltage, rather than the current, at the various nodes. They think this way because oscilloscopes and voltmeters are easier to use to probe nodes than they are to determine the current in the branches of a node. Because voltage is so important, you want to solve for voltage in your reactive circuits. Your intuition tells you that the larger a capacitor is, the smaller the voltage change due to a current that you inject into the capacitor. Because a big capacitor implies a small voltage, the capacitance must appear in the denominator. Remember: You are dealing with time-variant components, so the current term is not a dc term, but you express the rate of change of the current as di/dt, the rate of current change over time. Therefore, solving for the voltage across a capacitor yields $V=(1/C)(di/dt)$. Similarly, a larger inductor would have more voltage across it in response to a changing current, so that equation becomes $V=L(di/dt)$.

You can apply your intuition to solve for the current through a capacitor or an inductor. The forgetful need realize only that the current through a capacitor becomes larger with the value of the capacitor or the voltage change across it. This fact yields $I=C(dv/dt)$. Similarly, the larger the value of the inductor is, the smaller the current that forces itself through it will be. The inductance must appear in the denominator. That equation yields $I=(1/L)(dv/dt)$.

Support **Across The Board.**™

Bringing Products to Life with Mixed-Signal FPGAs.

At Avnet Memec, support across the board is much more than a tagline for us. From initial design through end of life — we are deeply committed to driving maximum efficiency throughout the product lifecycle.

The Challenge

How do you launch the first **Mixed-Signal FPGA** to the engineering masses? Actel's Fusion Programmable System Chip integrates configurable analog, large Flash memory blocks, comprehensive clock circuitry, and high performance programmable logic in a single chip. With so much to offer, the company wanted to quickly reach the engineering market. Enter Avnet Memec.

The Solution

Actel's leading distribution partner Avnet Memec underwent rigorous product training on the applications and benefits of Fusion. In the spring of 2006, Avnet Memec rolled out its SpeedWay™ full-day hands-on workshops and labs, highlighting the versatility of Fusion for OEM engineers. All this adds up to quite an education.

As OEMs look to implement Fusion into their designs, Avnet Memec and Actel will be there all the way — providing technical support in every phase of design development. It's true support across the board.

To learn more, request a free Fusion informational CD and enter to win one of **10 free Fusion Starter Kits**, go to: **www.em.avnet.com/actel/satb**

Actel® FUS1ON ONECHIP

AVNET® Memec

Enabling success from the center of technology™

1 800 408 8353
www.em.avnet.com

Avnet Green Initiative

NEDA Supplier Authorized Distributor

fast-moving, circulating currents will cause the ground reference to hop. John Dutra, member of the technical staff for field applications for National Semiconductor, points out that the fastest changing waveform that can exit from the side of the inductor that has no switch or diode is a triangle wave. That triangle wave is less troublesome than the high di/dt currents that circulate between the input capacitor and the synchronous switch. If the IC has a power ground, it can dump to this node, but the analog or reference ground should tie to the output capacitor right at its ground, preferably sharing the Kelvin connection with the feedback resistor.

Once you place the feedback network to the right of the output capacitor and you tie the IC reference to the output-capacitor ground, you must present the power-supply output to the system. To control the currents in the ground plane, some engineers advocate cutting slots into it. However, many engineers find that approach to be a bad idea. If there are currents that flow around the cut, it will create loop areas that radiate and cause other problems. It also limits the use of the plane as an RF shield.

"Controlling noise on a mixed-signal pc board can be a difficult problem," says Henry Ott, president of Henry Ott Consultants. "This situation is especially true on boards with multiple ADCs. Some designers suggest splitting the ground plane to isolate the digital ground from the analog ground. Although the split-plane approach can work, it has many

CIRCULATING CURRENTS IN AUDIO

Engineers often refer to circulating currents in audio as ground loops. Often, large currents in power stages interfere with the audio signal, much as the power current in power supplies interferes with the reference voltage. For example, a power amplifier boosts the output of a CD player, so that it will be strong enough to drive the speakers (Figure A). The application has a voltage regulator to step down the 12V on the board to 5V that the CD player needs, but the figure omits this detail. Many CD players work well with this design, but others exhibit unacceptable noise in the audio. These players had poor internal layout and small capacitors in the power supply.

Figure B rearranges the system to show the flow of large currents through the CD motor. Resistor symbols replace the wires, but ac impedance, rather than simple resistance, causes the signals that interfere with the audio.

Power enters the power-amplifier pc board from the top. The figure also omits the regulator on the pc board. The resistor represents the impedance of the board traces. The next resistor represents the impedance of the cable that brings power to the CD player. Yet another resistor represents the impedance of the traces on the CD player that go to the positive terminal of the motor. There are corresponding impedances below the motor. The figure also shows wires for the ground connections to the amplifiers in the CD and the power-amp board. These wires also have impedances, and you should also consider those impedances, especially the ground from the board amplifier because it carries higher currents to power a speaker. For simplicity's sake, only the motor's power wires show equivalent impedance. Those impedances are all that are necessary to explain the audio noise.

Only certain players had audio-noise problems because they have poor internal grounding and small power-supply capacitors. The figure shows the amplifier ground coming off the bottom of the CD player's internal impedance. But, if the player's designer had laid out the board so that the amplifier ground was at the negative terminal of the motor, things would be even worse. That approach would result in three equivalent impedances reacting against the motor current to add noise to the amplifier's ground reference. Just as deleterious, the use of small capacitors in the CD power supply ensures that ac currents end up in the audio chain.

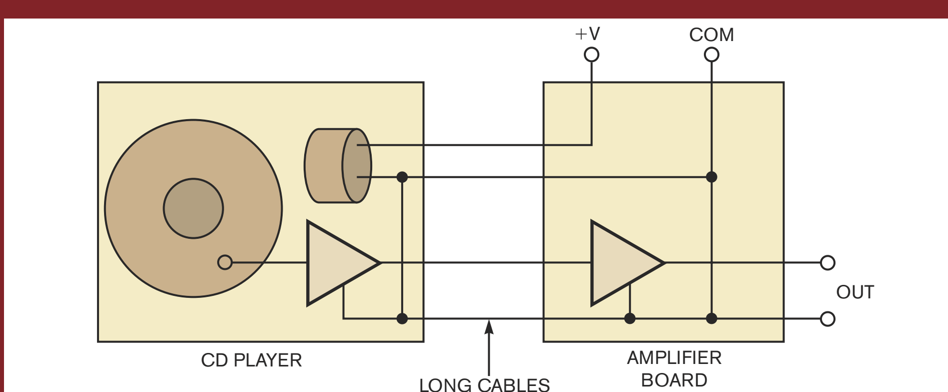

Figure A An amplifier board, which supplies power to the CD player, also creates an audio-ground loop.

potential problems. By understanding how and where high-frequency ground currents flow, it is usually possible to develop an approach to controlling noise and, in most cases, still maintain a single contiguous ground plane. Component placement and partitioning, combined with routing discipline—not splitting the ground plane—are the keys to success in laying out a mixed-signal pc board."

In the case of power supplies, it may be a bad idea to just nail every ground node to the plane even if you have placed the components to reduce the effects of circulating currents. Four-layer boards are now ubiquitous, and many engineers are now working on six- or eight-layer boards. You can take advantage of these layers to keep a ground plane that has minimal circulating currents. A simple technique exists even for two-layer boards: Tie together the ground nodes with ac-circulating currents on the top layer of the board. Alan Martin, principal field-applications engineer at National Semiconductor, reports that the best approach is to then tie that ground to the ground plane with only one via and to place that via at the ground pad of the output capacitor. This approach eliminates all ac-circulating currents from the ground plane and presents a regulated voltage at the output capacitor because the feedback network and the IC reference also connect to this pad.

Figures 5 and **6** show an example of this philosophy in practice. This buck regulator uses a Linear Technology 1.25-MHz LT1767 step-down regulator on a

If you imagine farads of capacitance in the player, then you can see that only dc would be in the cable and pc-board impedances. The CD-player motor could go on and off, the audio could get loud and soft, and the capacitance would smooth those ac pulses out so much that the response would be less than 20 Hz and undetectable to the human ear as well as below the cutoff frequency of the signal chain. A dc error would exist, and current would still flow in the ground loop, but it would be dc current and the signal would lack ac power currents. The decoupling of the amplifier in the CD player doesn't reduce this noise. Instead, the entire CD player must have massive capacitance to smooth out the current pulses in the motor circuit. Remember that CD players often buffer the data, and the motor can turn on and off periodically. This case caused the noise that appeared in the audio chain. In this case, star grounds or 5-oz copper planes will not help, and the impedance of the cable is certain to cause problems. Because a designer cannot dictate which brands of CD players connect to the system, he should make the audio input a differential connection by using an input transformer, a differential input amplifier, or an integrated chip, such as the Rohm BA3121 ground-loop-eliminator IC (Figure C).

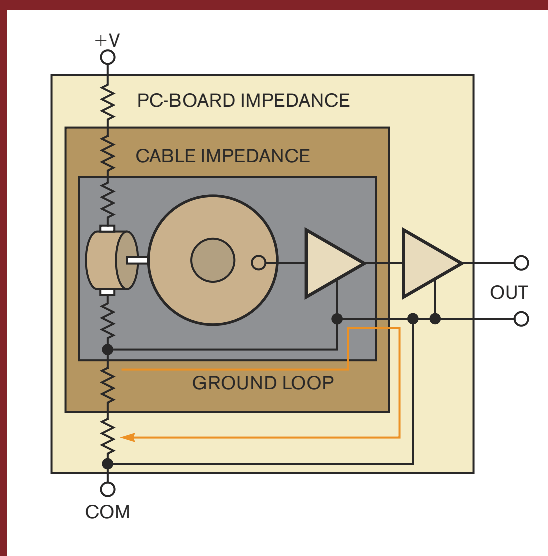

Figure B The pc board encloses the CD player in a ring.

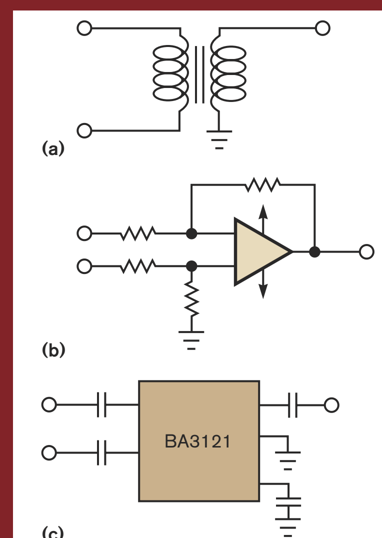

Figure C You can stop ground loops with the transformer (a), the differential amplifier (b), and the ground-loop eliminator (c).

Audio fanatics have long insisted that star grounds are the only acceptable grounding system. The advent of complex audio chains, including DSPs and Class D amplifiers, has taken the luster off the star-grounding system. Although star grounds can eliminate ground loops on a board, they also increase the impedance of all the grounds on that board. There are many star-grounded audio amplifiers that break up into oscillation at the least provocation. Skip Taylor, chief technology officer of D2Audio, recommends that the company's customers "maintain continuous ground planes and use placement, appropriately split power planes, and topside pours to manage the currents' effect on the signal."

four-layer board. In this layout, the bottom of the board is grounded but serves as a shield to other sensitive circuitry that is right below the board. The power supply connects to this circuitry at the ground pad of the output capacitor. The next layer up has just two connections: the output from the ferrite bead to the system and the connection between diode D_1 and capacitor C_5. The next layer up also ties to the ground at the output capacitor. In addition, this plane provides the connection for ground between the IC and the circulating node of C_2, D_2, and C_5. The circulating currents are in close proximity, and the currents enclose a minimal area. The switch in IC_1 is as close as possible to C_2 and D_1, where the worst circulating currents are. The connection between C_2, D_2, and C_5 uses a copper pour to remove heat from diode D_2, which is the hottest component in the system. Similarly, the ground for IC_1 is as large as possible. The via that ties it to the gray plane allows that plane to remove heat, as well. Note that the via for the IC does not connect to the plane. That plane has no circulating currents in it and serves as a shield. This design is currently in production and displays some of the cleanest, sharpest waveforms of any high-speed, high-performance buck regulator.

You can apply the same principles to a boost or SEPIC (single-ended-primary-inductance) regulator. **Figure 7** shows that topology and the commensurate principles for a good design. “You will never regret designing isolation into your power supply,” says Martin. Isolation can help localize circulating currents on the primary and the secondary sides, respectively. **Figure 8** shows that the isolated converter has the same requirements for the feedback network. The feedback-resistor ladder should not connect to the trace between the rectifier and the output capacitor, and a Kelvin connection would be useful. Also note that the reference in the IC and the isolated primary-side feedback should be close to each other and out of the path of the circulating currents in the switch and the input capacitor. The **figure** shows stray capacitors, which cause the additional issue with the isolated supply. The package tab of the switch transistor is the collector or drain, and that node is flying back with the primary winding when the switch opens. This fast dv/dt can inject current into the case of the power supply—that is, earth ground. A stray interwinding capacitance in the transformer allows the injection of current into the secondary, and it can charge the secondary to dangerous voltages if the secondary does not tie to earth or some other reference, at least with a high-value resistor. The **figure** shows the secondary referenced to earth ground, which is often the case in system design. Thus, two sources of current are flowing in the product’s chassis, which can wreak havoc with ground-fault interrupters or cause errors in measurement equipment.

“You have to keep three things in mind,” says Paul Greenland, vice president of marketing at Enpirion. “A wire in space has an inductance of 15 nH/in. Therefore, keep traces short. That inductance will increase if the wire forms a

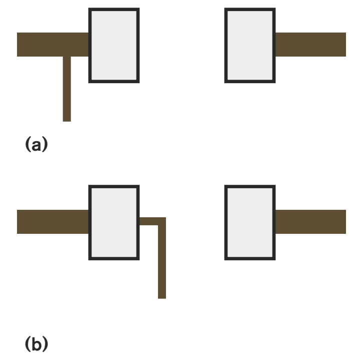

Figure 4 Instead of a conventional connection (a), use a Kelvin connection (b) to ensure that no circulating or dc currents flow in the sensing trace running from the pads of a component.

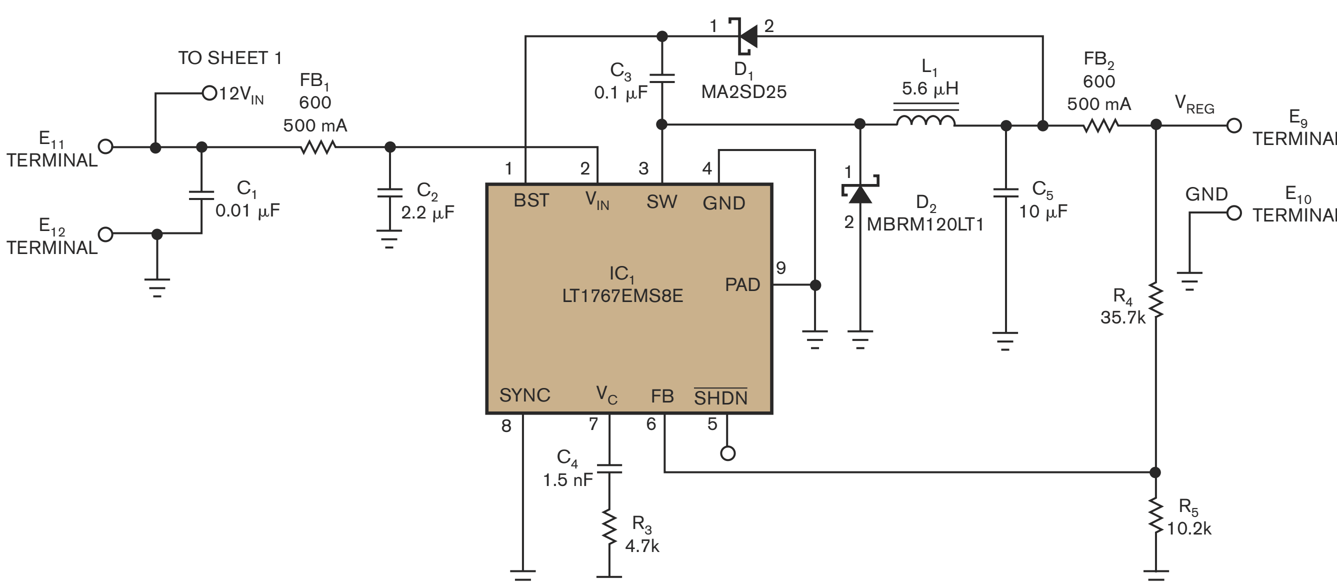

Figure 5 This 1.25-MHz buck-regulator circuit finds use in a high-volume consumer product.

Intersil Real-Time Clocks

High Performance Analog

And the Winner is...

Intersil's low power I²C Real-Time Clocks save costs and board space by integrating 4k Bytes of E²PROM memory **AND** CPU Supervisory Functions.

Switching to Intersil's ISL12027, ISL12028 and ISL12029 can save you money and board space two ways. First, we've integrated 4k of E²PROM memory, Power On Reset and a Watchdog Timer eliminating two external devices. Secondly, we've added crystal frequency trimming capability to deliver high accuracy timekeeping with a low-cost 32.768kHz crystal. The end result is a highly efficient real-time clock you can rely on for >2,000,000 Write Cycles.

Crystal frequency compensation provides initial crystal trimming and subsequent timing correction due to temperature variation, saving you money by delivering accurate timekeeping with less expensive crystal.

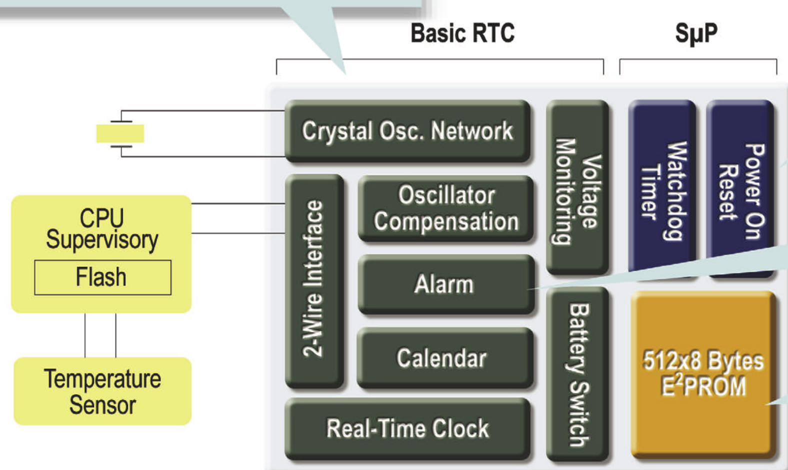

Integrated Power On Reset and Watchdog Timer eliminates external devices.

Two non-volatile alarms can be set to the second, minute, hour, day of the week, day or month.

Integrated 4k Bytes of E²PROM memory reliable for >2,000,000 write cycles eliminates external device.

800nA General Purpose Real-Time Clock Selector Table

	Int. E²PROM (Bytes)	Alarm	CPU Sup.Fx's POR	Wdg Timer	IRQ	FOUT	VTRIP for Rest/Bat Switch	Package
ISL12026	512 X 8	2	N	N		IRQ/FOUT	5 Sel. (2.63V to 4.64V)	8-Ld SO/TSSOP
ISL12027	512 X 8	2	Y	Y		RESET	5 Sel. (2.63V to 4.64V)	8-Ld SO/TSSOP
ISL12028	512 X 8	2	Y	Y		IRQ/FOUT	5 Sel. (2.63V to 4.64V)	14-Ld SO/TSSOP
ISL12029	512 X 8	2	Y	Y		IRQ/FOUT	5 Sel. (2.63V to 4.64V)	14-Ld SO/TSSOP

For datasheet, free samples, and complete line of general purpose Real-Time Clocks go to **www.intersil.com**

Intersil – Amplify your performance with advanced signal processing.

intersil®
HIGH PERFORMANCE ANALOG

loop. Therefore, keep loop areas small. Last, reduce field emission and susceptibility by routing traces so currents cancel." (See **Figure 9**.)

CUTTING THE GROUND PLANE

It is usually a mistake to cut up ground planes (see **sidebar** "Circulating currents in audio"). There are ways to use placement and topside pours to keep nasty circulating currents from the ground plane. Meanwhile, the uniform planes provide the lowest impedance and can provide a valuable shield from RFI. In some cases, cutting the plane can yield good results, but you must be careful to analyze what you have gained and whether your approach is pertinent in a real-world system.

Paul Grohe, applications engineer at National Semiconductor, has developed a pc board that connects directly to the front of an HP3577A network analyzer (**Figure 10**). He uses this board to evaluate the gain and phase properties of the amplifiers that his group designs. This board exhibited good results, exceeding the accuracy, repeatability, and low noise of previous efforts, which used cut-up, copper-clad, handmade boards. Grohe achieved a noise floor of −110 dB using this system. This floor was important at the lower frequencies, when the network analyzer measured an error signal of 1 μV or less. Still not satisfied, Grohe took a knife to the ground plane on the board. By cutting the plane, he improved the noise floor to −130 dB. Anyone familiar with amplifier characterization and performance knows what a remarkable achievement this is. **Figure 11** shows the system.

The HP3577A's chassis common is the reference for the 50Ω source, as well as the input amplifiers. Grohe realized that the circulating currents use the connector ground on the reference input to return to the instrument. In a flash of inspiration, Grohe discovered that these currents could interfere with the measurement of the submicrovolt signals at that input. Because the single-ended HP3577A has a 200-MHz bandwidth, it was important to terminate the source with a 50Ω resistor. Grohe knew that the relatively large currents from this resistor

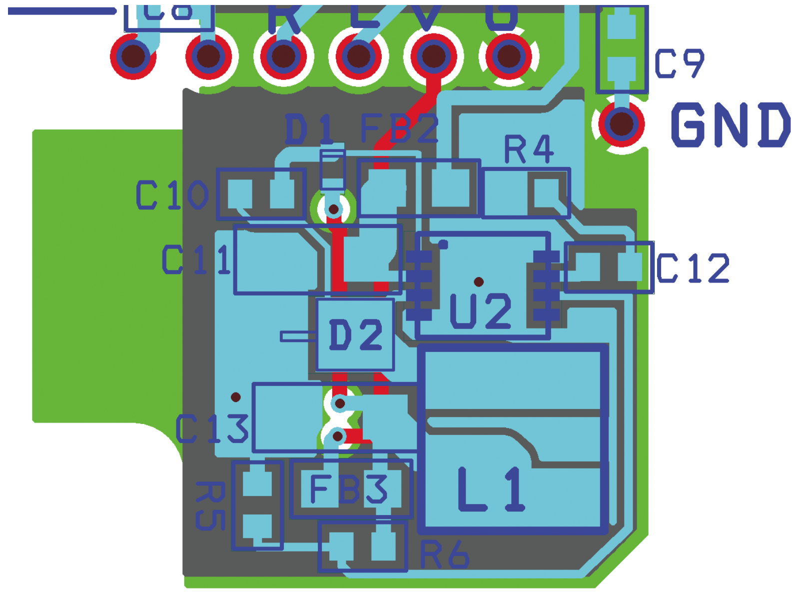

Figure 6 The pc-board layout of the circuit in Figure 5 demonstrates good design practice that minimizes the effects of circulating currents.

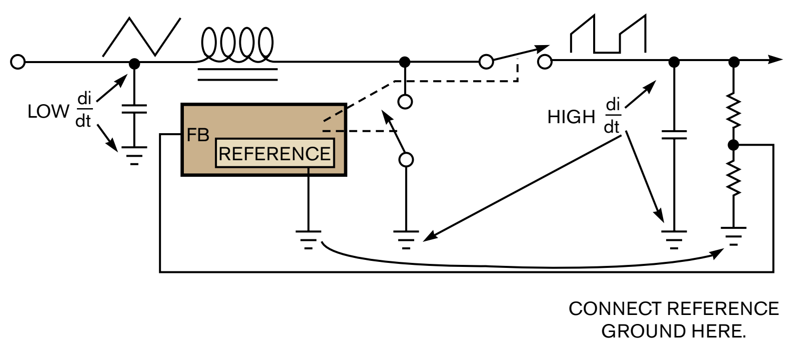

Figure 7 Boost regulators have large circulating currents in the output capacitor and the switch ground. Do not put the reference ground between those nodes.

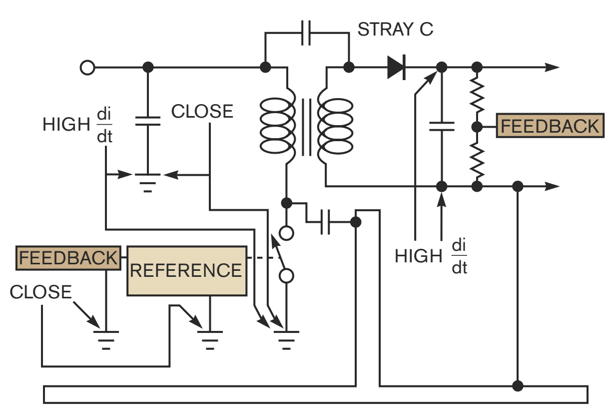

Figure 8 Isolated converters have large circulating currents. Stray capacitance between the transformer windings and between the switch case and the heat sink also cause circulating currents.

Intersil Digital Light Sensors

High Performance Analog

The Best Digital Light Sensor and Its Closest Competitor

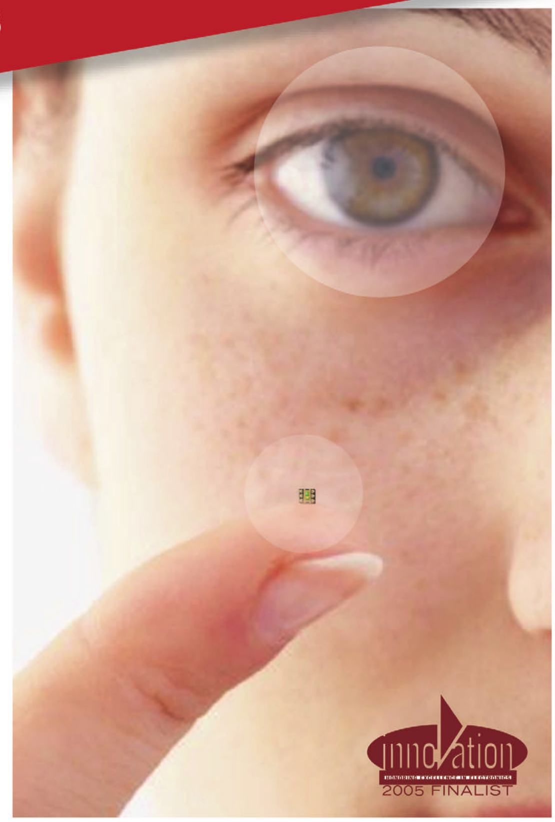

With light sensitivity only matched by the human eye, Intersil's ISL29001 Digital Light Sensor provides simple, pure 15-bit I²C digital data.

Drawing less than 300µA of supply current, the ISL29001 provides 15-bit effective resolution. This state-of-the art device integrates two photodiodes and an ADC into a super small 2.1mm x 2.0mm ODFN package. The digital data in standard I²C format couldn't be simpler to use. It's no wonder **EDN Magazine** has selected one of this family's light sensors as a finalist for this year's **Innovation of the Year Award**.

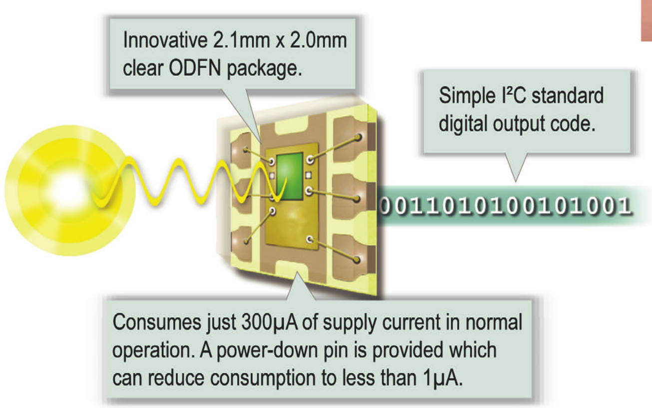

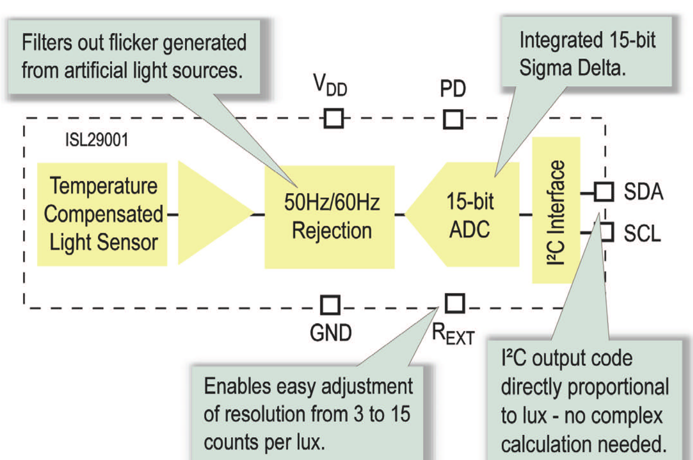

ISL29001 Key Features:

- I²C interface produces simple I²C output code, directly proportional to lux
- 0.3 lux to 10,000 lux range
- 50Hz/60Hz rejection to eliminate artificial light flicker
- Human eye response
- 15-bit effective resolution
- Adjustable resolution: 3 to 15 counts per lux
- 2.5V to 3.3V supply
- Temperature compensated
- 6-pin ODFN (2.1mm x 2mm)
- Pb-Free plus anneal available (RoHS compliant)

Datasheet, free samples, and more information available at **www.intersil.com**

Intersil – Amplify your performance with advanced signal processing.

intersil®
HIGH PERFORMANCE ANALOG

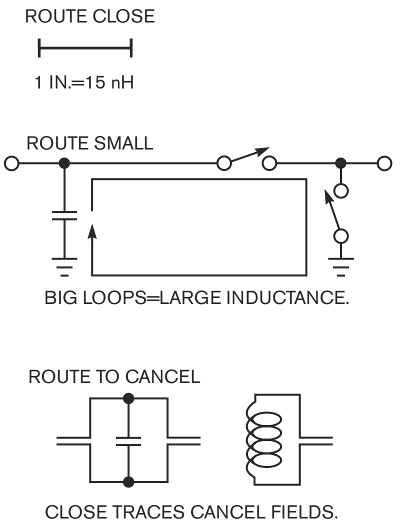

Figure 9 Small, tight layouts with canceling currents exhibit the best performance and provide low EMI/RFI emissions.

would seek "home" back to the analyzer by all possible connections. Because all three connections shared the same chassis ground, the current would return in the delicate R reference connector, as well as the source and A inputs. The interference with the A input was less serious because this input measures the output of the amplifier under test. This signal is large enough to resist interference. By cutting the ground plane around the amplifier under test, Grohe ensured that none of the larger currents from the 50Ω source-terminating resistor would flow through the reference connector. This 20-dB improvement is testament to the principle that an engineer must think of the ground not as an ocean of low impedance, but as simply another connection to keep in mind for maximum performance.

In defense of the rule against cutting ground planes, you should note that cutting improved not the performance of the amplifier, but the measurement of a microvolt signal that is internal to the amplifier's loop. Using a differential-input analyzer, such as a Ridley or a Venable, would also alleviate this issue of circulating currents. Unfortunately, those instruments target use in power-supply analysis and have bandwidths one-tenth that of the HP3577A. Before advocating cutting the ground plane in a production board, make sure that the improvement is not just an improvement in measuring and interfacing with test equipment and rather a genuine improvement in signal-chain performance.

CIRCULATING CURRENTS IN RF

RF engineers are well-aware of the headaches that circulating currents cause. Radio systems must often reference earth ground, and the components are single-ended. **Figure 12** shows a coaxial cable connecting two RF subsystems. Because the cable must serve as an RF shield, it must connect to the cases of the subsystems. If a large power usage in one subsystem injects current into earth ground, that current will travel along the coaxial shield and interfere with the signal on that coaxial cable. Diligent power-supply design can help in this case. Minimize current injection to the frame of the supply. Use electrostatic shields in the switching power transformer and between the FET switch and the supply case. Vicor takes this approach for most of its modules, and its products typically do not inject measurable common-mode currents.

Remember that power-supply return, chassis common, and shielding ground have different requirements. You should avoid using the system's frame to return current, as you would in the design of an automobile. Every power supply should have its own return lines for the delivered current. In this way, the power leads form a loop with a smaller area and reduce EMI radiation. Ensure that this return does not connect to the frame of the system at various places. If that approach is unavoidable, those connections should be through common-mode chokes or, at least, ferrite beads. This approach reduces the higher frequency ac currents that are more likely to interfere with RF signals. Keeping the shield a shield instead of a power-supply return also ensures that the

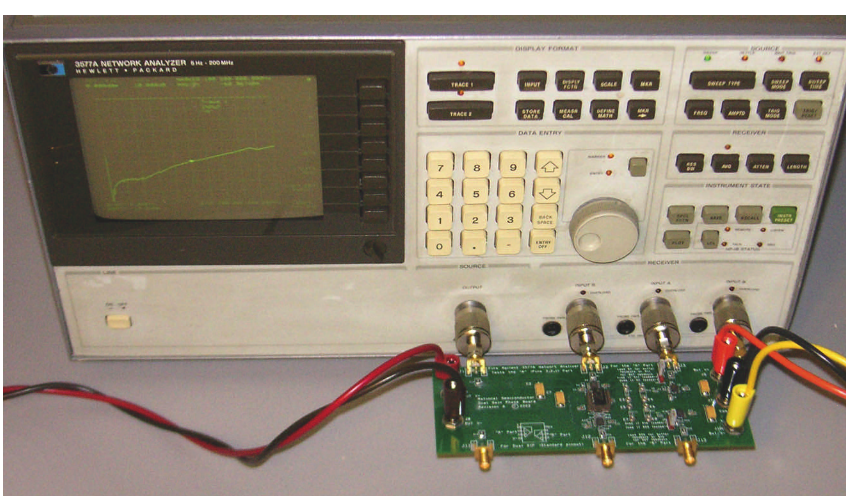

Figure 10 The gain- and phase-characterization pc board of the HP3577A network analyzer mounts directly on the front-panel connectors.

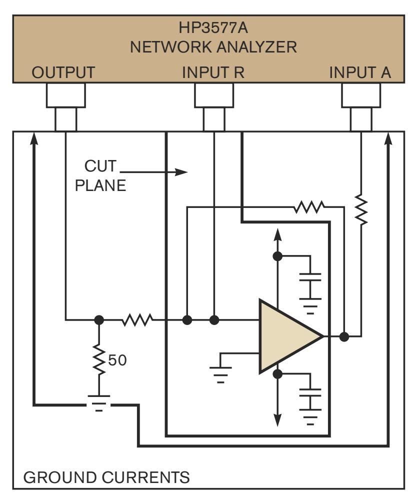

Figure 11 Cutting the ground plane on the gain/phase board routes circulating currents away from the R input, which measures microvolt signals.

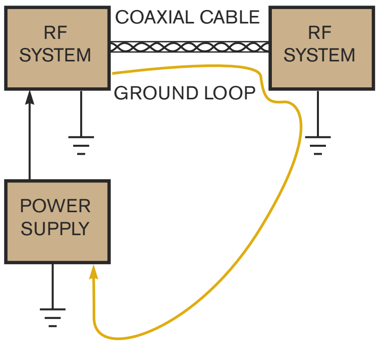

Figure 12 The power-supply current in RF systems often returns over the coaxial-cable shield, causing circulating currents.

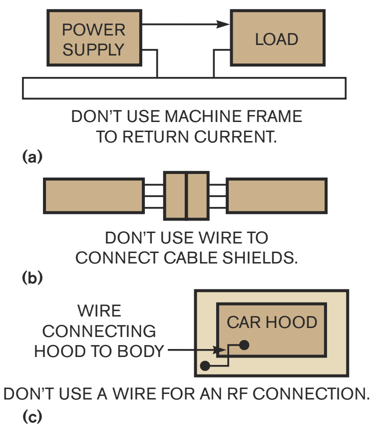

Figure 13 In a classic ground loop, the shield in the coaxial cable carries the power-supply current and interferes with the signal by inducing a series voltage drop in the outer conductor (a). You can also gather up the shields in a cable and solder a wire to them to pass the shield through a conventional amp-pin connector (b). You can use a wire to connect the hood of a car to the body, but RF signals easily pass through the gap between the hood and the fenders (c).

product has minimal RF emission, and, just as important, the machine will be immune to RF interference. Using the machine frame as a power-supply return is one reason that maintenance people cannot use FM radios in a semiconductor fab. Doing so may cause the machinery to reboot or act unpredictably.

More discussion of RF radiation and immunity is beyond the scope of this article. However, if you do not make the errors that **Figure 13** shows, your system will be more robust. The problem is one of a classic ground loop, in which the shield in the coaxial cable carries the power-supply current and interferes with the signal by inducing a series voltage drop in the outer conductor (**Figure 13a**). This approach can cause million-dollar semiconductor machines to crash when you stand next to them and key in a radio. You can also gather up the shields in a cable and solder a wire to them to pass the shield through a conventional amp-pin connector (**Figure 13b**). However, at high frequencies, the inductance of that wire is a high reactive impedance. Also, the loops that form will act like small antennas to radiate and receive EMI. This approach can cause a semiconductor to fail CE (Conformité Européenne)-immunity certification because the RF currents circulate in the sensor side of the cable and give false signals to the microprocessor, causing the wafer-handling system, in turn, to shatter 25 wafers at once.

The situation in **Figure 13c** comes from the auto industry. It demonstrates the difference between a galvanic connection and an RF connection. You can use a wire to connect the hood of a car to the body, but RF signals easily pass through the gap between the hood and the fenders. The hole in a sheet of metal degrades shielding effectiveness and relates to the largest linear dimension of the hole compared with a wavelength. Placing several contacts around the periphery of the hood reduces the length of each hole. Again, at RF frequencies, the wire is a high reactive impedance. Mechanical engineers don't believe this fact because the ohmmeter displays 0Ω between the car body and the hood. This problem combines with the adoption of plastic inner fenders, causing no end of grief to auto engineers trying to prevent ignition noise from interfering with radio reception.

James Long, an analog and RF consultant, advises clients: "Remember that current flows in closed loops, as Kirchoff's current law dictates. Also, current takes the path of least impedance. At RF, this is inductive reactance. It chooses the path in which the inductance of the loop is smallest, which means that it encloses the smallest area. Using these rules, you can visualize where the current will flow and what secondary effects it will have." Long also advises that you can encourage the current to flow where you want by placing the going and coming conductors close together and away from areas in which the return current would cause harm. EDN

MORE AT EDN.COM

+ For a recent *EDN* article that stresses the importance of minimizing loop area, go to www.edn.com/article/CA6347258.

+ For more on *EDN*'s analog editor, Paul Rako, visit www.edn.com/info/1340007045.html.

FOR MORE INFORMATION

Altera Corp
www.altera.com

Analog Devices
www.analog.com

Cadence Design Systems
www.cadence.com

D2Audio
www.d2audio.com

Enpirion
www.enpirion.com

Fairchild Semiconductor
www.fairchildsemi.com

Henry Ott Consultants
www.hottconsultants.com

James Long Consulting
www.analog-rf.com

Linear Technology
www.linear.com

National Semiconductor
www.national.com

Rohm
www.rohm.com

Vicor
www.vicr.com

REFERENCES

1 "Understanding and Minimizing Ground Bounce," Fairchild Application Note AN-640, February 2003, www.fairchild-direct.com/an/AN/AN-640.pdf.
2 Kelvin (4-wire) resistance measurement, All About Circuits, www.allaboutcircuits.com/vol_1/chpt_8/9.html.
3 Varga, Craig, "Improved Kelvin contacts boost current-sensing accuracy by an order of magnitude," *EDN*, Feb 17, 2005, pg 80, www.edn.com/article/CA502424.

You can reach
Technical Editor
Paul Rako
at 1-408-745-1994
and paul.rako
@reedbusiness.com.

BY BRIAN BOORMAN • HARRIS CORP

Designing dual-modulus dividers in an FPGA

A TECHNIQUE FROM DEEP IN THE DIGITAL DESIGNER'S DUSTY BAG OF TRICKS MAY BE JUST THE TICKET FOR GENERATING MODERATE-SPEED CLOCKS IN PROGRAMMABLE LOGIC.

It is often necessary for designers to implement a digital clock divider where the output frequency is not an integer factor of the reference clock. Today's newer FPGA technologies usually contain digital or analog PLLs for frequency synthesis. But because of their specifications, these components are usually applicable only for high clock rates. There is still a need for relatively low-speed clock generation—to connect to peripheral devices, for example.

These peripherals usually connect using either synchronous or asynchronous serial interfaces. An example is the standard asynchronous serial port on virtually all PCs and many embedded products. If your design needs to interface to devices with asynchronous interfaces, synchronous interfaces, or both, then you may need several baud rates available. Often, the synchronous and asynchronous rates are different. Choosing an input clock frequency that can satisfy all of the required output frequencies is frequently impossible. In some cases, you may be able to get close to the desired rate for most frequencies, but the error percentage is too large for other frequencies. Such a scenario is an example of where a dual-modulus frequency-divider circuit may be helpful.

This article presents a method for designing a dual-modulus divider. It includes the math for determining the circuit characteristics and offers an example implementation. Although the use of dual-modulus dividers has been around for a long time, joining the work force every year are new engineers that may be unaware of this technique.

A dual-modulus frequency divider is a counter in which the preloaded value can take on two values depending on the current state of a separate sequence counter. Sometimes the counter counts down from (or up to) one number; other times it counts down from (or up to) a second number. By using the sequence counter to control the ratio of usage of one preload value to another preload value, it performs a fractional division on an average, though not instantaneous, basis.

FIRST, THE MATH

Start by developing the math that will allow you to design dual-modulus frequency dividers. Given the following identifiers: M=real divisor, equivalent to $F_{REF}/F_{DESIRED}$; N=integer divisor, which you find by rounding M up to the next integer; and A=integer, A+1 is the total number of division cycles. Then,

$$(A\times N)+(N-1)=(A+1). \qquad (1)$$

In this form, **Equation 1** is not very useful, because the term you are trying to find, (A), is on both sides of it. Rearranging this **equation** in terms of A yields:

$$A=[(M+1)-N]/N-M. \qquad (2)$$

However, **Equation 2** does not yield an answer for all possible input values. You need a more generalized form of the formula that allows the usage ratio of the two moduli to be other than A-to-1. Thus,

$$(A\times N)+(B\times(N-1))=(A+B)\times M. \qquad (3)$$

If you say that

$$A+B=C, \qquad (4)$$

then

$$(A\times N)+(B\times(N-1))=C\times M. \qquad (5)$$

Because A, B, C, and N must all be positive-integer numbers,

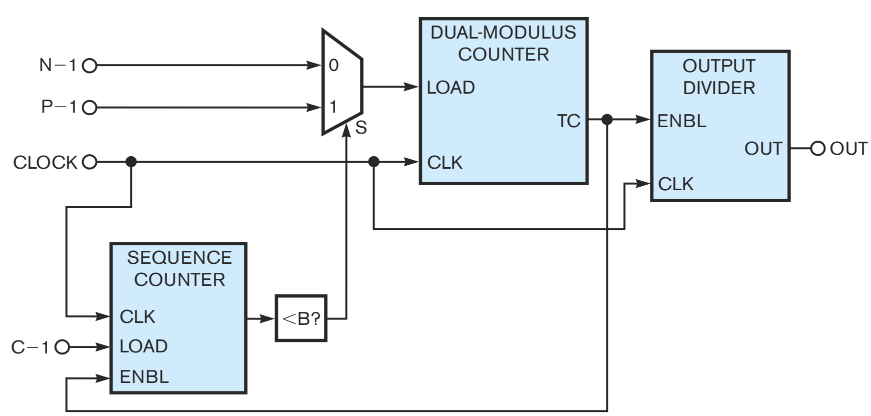

Figure 1 The dual-modulus counter comprises two loadable counters: one to keep track of which modulus to use, and one to do the actual counting.

the value of C must be such that the product C×M is also an integer number. The quickest method for determining the value of C is to examine the fractional portion of the real divisor value M; let's investigate that step further.

If you define P as the greatest integer that is less than the real divisor M (in other words, P=floor(M) in C/C++ vernacular), then you can say that

$$D=(1/M-P), \quad (6)$$

where D is the inverse of the fractional portion, and then

$$C=D\times E, \quad (7)$$

where E is the smallest positive integer that makes C an integer also. You find the correct value for E either by inspection (recognizing common fractional values such as 0.25 and 0.333) or by trial and error. A spreadsheet program is useful for this task.

Once you solve the equations for C, you know the sequence length of the dual-modulus counter. The sequence length represents the total number of times that the two moduli will be used before the cycle repeats. The next step is to determine the specific values for A and B. You can accomplish this task by solving:

$$(A/C)=M-P, \quad (8)$$

and, solving for A,

$$A=(M-P)\times C. \quad (9)$$

Combining **equations 6** and **7** yields

$$C/E=1/(M-P), \quad (10)$$

and, solving for E,

$$E=(M-P)\times C. \quad (11)$$

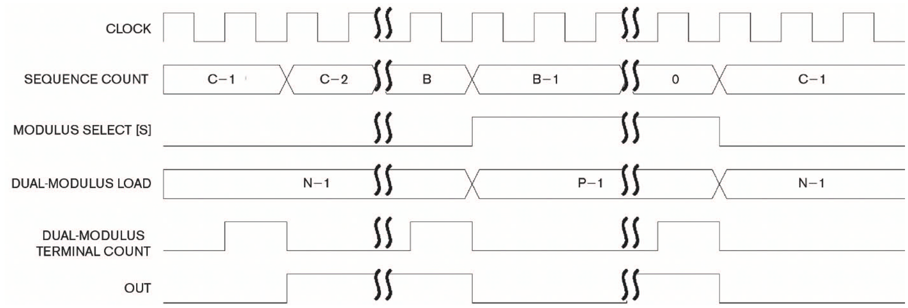

Figure 2 The two counters work together to generate an output at the desired frequency.

As you can see from **equations 9** and **11**, A=E. From **Equation 4**, you know that B=C−A. You now have the necessary information to put the circuit together, at least from a mathematical standpoint. Next, work through a practical example to put this methodology to use.

Figure 1 shows the example circuit for implemention. The entire module is synchronous to the clock input, CLOCK. The other main parameters to the design are the moduli N and P, the sequence length C, and the N-versus-P selection parameter B. The output derives from a divide-by-two D flip-flop circuit.

The sequence counter tracks the total number of divisions that the dual-modulus counter has completed. Because the dual-modulus counter counts from N−1 down to zero A number of times and from P−1 down to zero B number of times, the sequence counter counts the total value C (**Equation 4**). You implement this counter with a downcounter architecture that counts from C−1 to zero. The sequence counter decrements only when the dual-modulus counter has reached its terminal-count value, indi-

	A	B	C	D	E	F	G	H	I	J	K
1	**Dual Modulus Counter Calculation Tool**										
2											
3	**Fref:**	10,000,000	Hertz								
4	**Fdesired:**	115,200	Hertz								
5	**Initial "E" Value:**	21									
6											
7	**Note: Adjust 'E' to give integer result for 'C'.**										
8											
9	**M**	**N**	**P**	**D**	**E**	**C**	**A**	**B**	**Check1**	**Check2**	**C Is INT?**
10	43.40277778	44	43	2.482758621	21	52.137931	21	31.1379	2262.93	2262.93	FALSE
11					22	54.6206897	22	32.6207	2370.69	2370.69	FALSE
12					23	57.1034483	23	34.1034	2478.45	2478.45	FALSE
13					24	59.5862069	24	35.5862	2586.21	2586.21	FALSE
14					25	62.0689655	25	37.069	2693.97	2693.97	FALSE
15					26	64.5517241	26	38.5517	2801.72	2801.72	FALSE
16					27	67.0344828	27	40.0345	2909.48	2909.48	FALSE
17					28	69.5172414	28	41.5172	3017.24	3017.24	FALSE
18					29	72	29	43	3125	3125	TRUE
19					30	74.4827586	30	44.4828	3232.76	3232.76	FALSE
20											

Figure 3 A simple spreadsheet can find a value of E that makes C an integer.

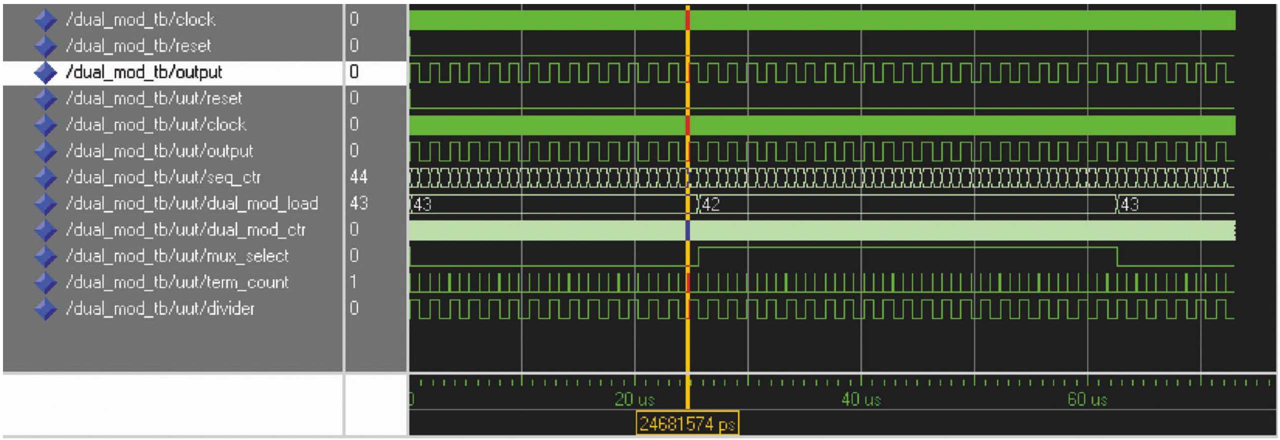

Figure 4 In simulation, the dual-modulus counter behaves as the mathematics behind it predict.

cating that a division cycle has completed. When the sequence counter reaches zero and the next terminal-count pulse is received, the sequence counter reloads to the value of C−1.

The count output of the sequence counter routes to a comparator block that determines whether the sequence-count value is less than parameter B. The output of this block is a 1 when the sequence count is less than B; otherwise, the output is a 0. You use this output to select which modulus the dual-modulus counter will use during the next division cycle using the 2-to-1 bus multiplexer that **Figure 1** shows.

The dual-modulus counter is simply a free-running down-counter. When the counter reaches zero, it reloads itself with whatever value is present on the LOAD input. The output of this counter is a simple 1 pulse whenever the counter is at zero.

You generate the output of the circuit with a simple output divider. This divider consists of a D flip-flop that inverts its out-

put state whenever the ENBL input is a 1 (also referred to as a T, or toggle, flip-flop). The purpose of this step is to generate a square-wave output clock signal, OUT, from the module. The impact of it, however, is that the output rate from the dual-modulus counter must be twice the desired final output rate. If your design does not require a square-wave output, then you could replace the T flip-flop with a single D flip-flop output register.

Figure 2 shows the relative timing relationships between the various signals in the example design.

Once the design architecture is in hand, you can compute the necessary values for generating a 115.2-kHz output from a 10-MHz input suitable for an asynchronous transmitter. Because the example uses a square-wave output, the output of the dual-modulus counter needs to have a rate of 230.4 kHz.

Start by calculating M, N, P, and D:

$$\begin{aligned} M &= \frac{F_{REF}}{F_{DESIRED}} = \frac{10\text{ MHz}}{230.4\text{ kHz}} = 43.402\overline{7}. \\ N &= 44. \\ P &= N-1 = 44-1 = 43. \\ D &= \frac{1}{M-P} = \frac{1}{43.402\overline{7}-43} = 2.482758621. \end{aligned} \qquad (12)$$

You now need to find a value for E that satisfies **Equation 7**. Inspecting the value you calculated for D yields no useful clues as to the value of E, so the method involves trial and error. Using a spreadsheet, you can simply plug in values for E until C is an integer. **Figure 3** shows the output of the spreadsheet calculation. Through this method, you find that E=29 satisfies the criteria:

$$C = D \times E = 2.482758621 \times 29 = 72. \qquad (13)$$

Finally, the last parameter you need is B:

$$B = C - A = C - E = 72 - 29 = 43. \qquad (14)$$

Figure 4 shows a screenshot of the spreadsheet that aids in performing the necessary calculations. The spreadsheet is in Microsoft Excel format and uses the conditional-formatting feature to help visually identify acceptable values.

Now that the detailed design is complete, take a look at an implementation suitable for FPGA development. This implementation uses VHDL (1993-syntax) coding for hardware description.

USING A SPREADSHEET, YOU CAN SIMPLY PLUG IN VALUES FOR E UNTIL C IS AN INTEGER.

VHDL CODE EXAMPLE 1

Listing 1, which is available at the Web version of this article at www.edn.com/ms4194, is the VHDL implementation of the example design. The design has been implemented in a manner that enables reuse, which includes defining constants for


MILL-MAX®
MILL-MAX MFG. CORP. • OYSTER BAY, NEW YORK
Visit our Design Guide Online for
MAXIMUM
Convenience
Maximum Ideas. Minimum Effort.
The Mill-Max MAXIMUM Design Guide is a must for mechanical printed circuit board designers looking to find sockets, connectors and discrete components to fulfill the most demanding applications.
SEARCH
• Download individual product drawings and specifications.
• View 3D images.
• Order free samples.
BROWSE
• 237 pages available for immediate download or by mail.
• The latest product offerings from Mill-Max including new additions to our popular spring-loaded product.
• The standard Mill-Max offerings available with traditional and RoHS compliant platings.
CREATE
• Design in Mill-Max products from choices of standard and application specific parts.
For our new Design Guide and to see more exciting products, go to www.mill-max.com/design
Response Code: EDN563 Phone: 516-922-6000

all the key design parameters. You could further extend this implementation by passing these values into the module using a VHDL generic interface in the entity declaration. The interface is simple for this example, comprising just a synchronous reset input, a clock input, and the output signal.

You then enter the values for N, P, B, and C into the code in the constant declarations in the architecture body. The code uses these constants to constrain the range of the counters that are also declared, which allows the synthesis tool to infer the correct width of the counter in hardware and verifies during simulation that the circuit does not exceed the range of the counters.

The VHDL code is straightforward, contains comments, and is consistent with the design diagram in **Figure 1**.

VHDL CODE EXAMPLE 2

Listing 2, also available at www.edn.com/ms4194, is the VHDL implementation of the testbench for the example design. The testbench is also a simple design, which comprises an instantiation of the example design, labeled UUT (unit under test); a process to generate the input clock; and a process to generate an initial reset to the design. Further verification is by inspection of the simulation-waveform results.

ONCE YOU UNDERSTAND THE CONCEPT AND MATHEMATICS, IT BECOMES EASY TO USE THIS TYPE OF FUNCTION TO SOLVE NONINTEGER-DIVISION PROBLEMS.

SIMULATION RESULTS

Figure 4 shows a screen capture of the simulation-wave window. This window shows both the testbench-level signals and the example-design internal signals over a simulation duration of 75 msec. The output waveform is at a 50% duty cycle, and the logic correctly switches between the two modulus-counter values, 42 and 43. By zooming in on the actual simulation, you can verify correct operation with respect to the sequence and modulus counters.

This article explores the basic math and concepts behind the use of dual-modulus frequency-division techniques. Once you understand the concept and mathematics, it becomes easy to use this type of function to solve noninteger-division problems. This article touches on only the basic implementation, but as you may envision, you could use the technique in more complex ways, such as by multiplexing the parameters to create a more versatile baud-rate generation, for example. EDN

AUTHOR'S BIOGRAPHY

Brian Boorman is a principal engineer with Harris Corp's RF Communications Division (Rochester, NY). He holds a bachelor's degree from the State University of New York at Alfred and a master's degree in electrical engineering from Alfred University (Alfred, NY).

Power Supply Monitoring

High Performance Analog Solutions from Linear Technology

Many modern electronic systems require the presence of proper power levels and stability before start up, ongoing supply monitoring while up and running and warnings for an imminent power failure. Power supply supervisors fill those needs by ensuring all supplies are within valid levels and indicating if and when a power supply has dropped below a minimum threshold in order to prevent unreliable operation and possibly initialize housekeeping operations.

The design of any given system presents its own unique demands and challenges. To ease the systems' design, Linear Technology offers the most complete portfolio of multi-supply supervisors that combine great precision, low power, versatile functions and different threshold levels. Linear Technology's portfolio consists of supervisors that offer fixed thresholds to effectively eliminate the need for external components; supervisors that feature only adjustable inputs in order to accommodate even the most custom voltage levels, and some that combine both. Wrapping up the rich collection are supervisors for under- and overvoltage monitoring, which can also be configured for negative supplies. Regardless of the functions and blend of features, all of Linear Technology's supervisors have a common trait- engineering elegance, while integrating innovative design and excellent performance.

Hex and Octal Supply Supervisors

The new generation of memories, PLDs, ASICs and microprocessors often require their own, unique power supply, increasing the number of voltages to be monitored in a system. Realizing the growing trend in the number of supplies with non-standard levels, Linear Technology offers supervisors for six and eight independent supplies.

Housed in small SOT-23 and DFN packages, the hex supply supervisor, LTC2908, accommodates the latest generation of systems with multiple

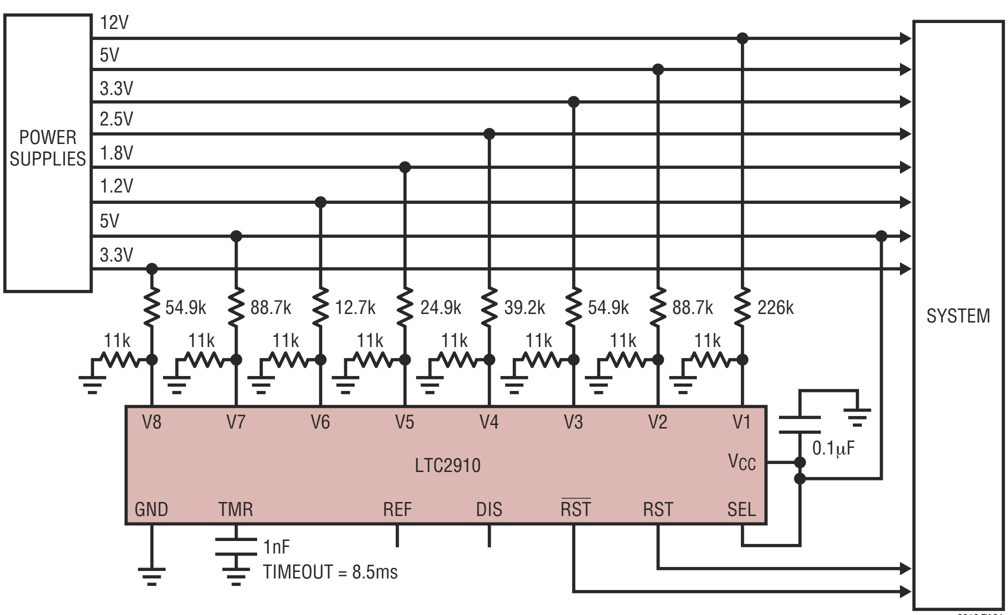

Figure 1. LTC2910 Typical Application Diagram

Power Supply Monitoring

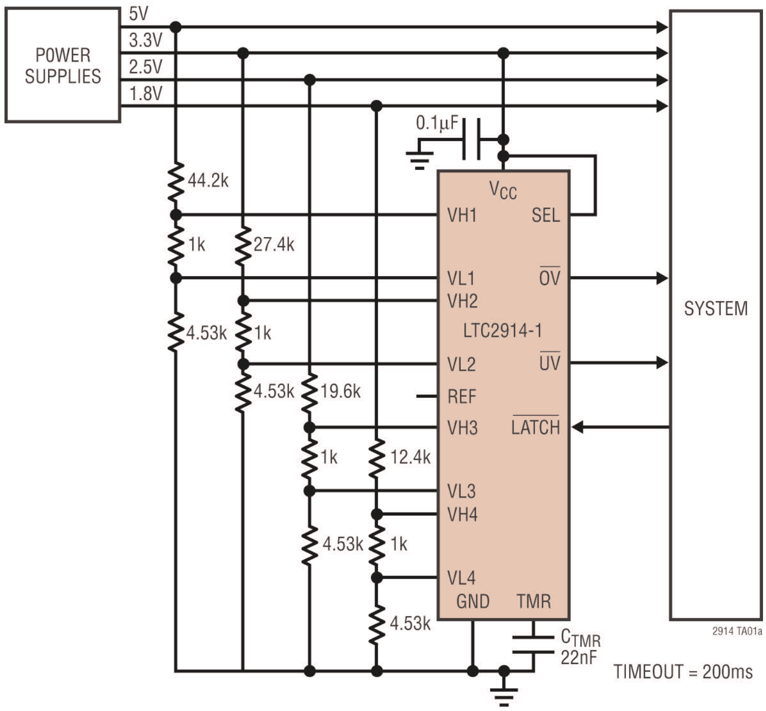

Figure 2. Quad UV/OV Supply Monitor, 10% Tolerance, 5V, 3.3V, 2.5V, 1.8V

supply rails by combining traditional fixed thresholds with low-voltage adjustable inputs (table 1). Allowing for compact and precise monitoring, the IC guarantees that the Reset will be held low with as little as 500mV of the input supply voltage. The tight ±1.5% threshold accuracy over the entire operating temperature range of –40°C to + 85°C gives the designer an extra degree of precision, since it ensures reliable Reset operation without false triggering.

The LTC2910 is a low power voltage monitoring circuit with eight individual adjustable inputs (0.5V) that reduces the board space required for voltage monitoring of densely packed systems with a large number of supplies. Figure 1 shows a typical application diagram for the LTC2910. Two external resistors per input determine the desired trip point for the undervoltage monitor, while a three-state select pin programs the polarity combinations for the input thresholds. Negative supplies can be monitored using the 1.0V reference output. Two complimentary Reset outputs will signal when any of the supplies is below its operating range, while the input glitch filter ensures that there will not be false or noisy triggering. A disable input masks the reset output and is useful during supply margin testing.

The LTC2910 employs an internal shunt regulator for high-voltage operation, guaranteeing that the IC will work at any voltage with just the addition of a single resistor.

Table 1. LTC2908 Voltage Inputs, V1-V6

Part	V1	V2	V3	V4	V5	V6
LTC2908-A1	5 V	3.3 V	2.5 V	1.8 V	Adj	Adj
LTC2908-B1	3.3 V	2.5 V	1.8 V	1.5 V	Adj	Adj
LTC2908-C1	2.5 V	Adj	Adj	Adj	Adj	Adj

Undervoltage, Overvoltage and Negative Voltage Monitoring

Some high-availability systems require that their power supplies be monitored for undervoltage and overvoltage conditions. Overvoltage monitoring has become necessary to

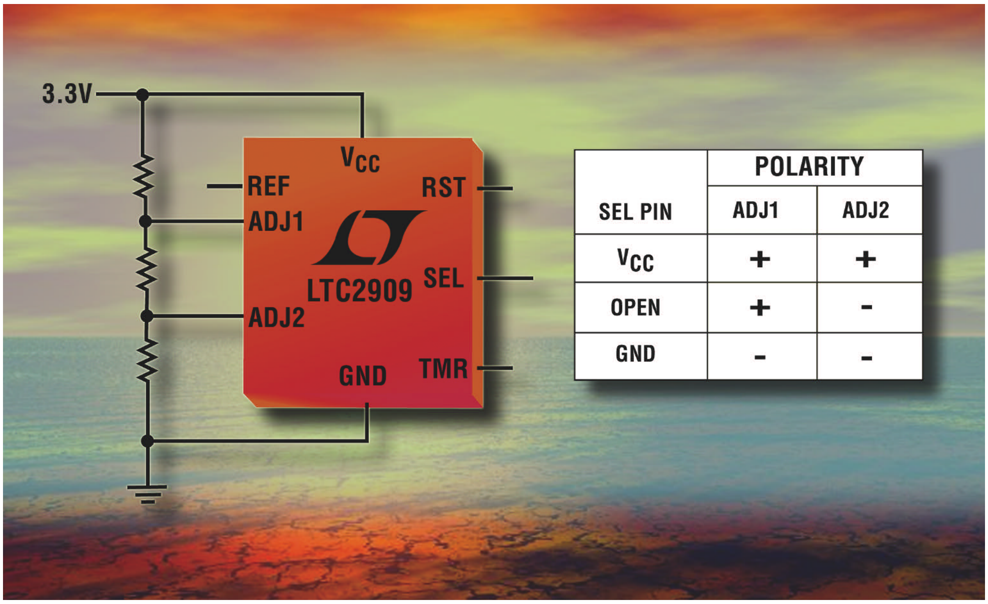

Figure 3. Dual Input UV, OV and Negative Voltage Monitor

prevent memory loss, data errors, equipment turn off or damage to expensive system components.

The LTC2912, LTC2913, and LTC2914 form a family of single, dual and quad voltage monitors with separate undervoltage (UV) and overvoltage (OV) inputs. The ICs are designed to provide the user with the optimum flexibility to set custom UV and OV trip thresholds, while sharing common UV and OV open-drain outputs that can be wired-OR'd together to indicate a single fault. Each of the ICs comes in two versions: the LTC291x-1 has latching capabilities for the overvoltage output, and the LTC291x -2 has an output disable feature to facilitate margining.

The quad LTC2914, shown in figure 2, has an extended feature set. A buffered reference output and a three-state input polarity-select pin allow for monitoring of up to two separate negative voltages without the need of external components, greatly simplifying both design and layout.

With an unlimited choice of threshold selection, simple configuration and low quiescent currents of 40µA, 60µA and 70µA respectively, the LTC2912, LTC2913 and LTC2914 are perfect for systems that require reliable and accurate voltage monitoring. The single and the dual are ideal for portable devices and applications, while the quad can be easily tailored for the larger telecom/network equipment and storage servers.

The versatile LTC2909, a selectable polarity dual supervisor, provides a different approach to under- or over-voltage monitoring for positive or negative supplies. The LTC2909 may monitor two supplies-positive, negative, or both- for UV, OV or one of each. It may also monitor a single supply (positive or negative) for UV and OV simultaneously. The selection is easily configured by a single select pin that chooses one of three possible polarity combinations for the adjustable inputs (figure 3).

The user adjustable inputs of the LTC2909 have a low voltage threshold of 0.5V and an undervoltage lockout allows V_{CC} to be used as an accurate third fixed 10% UV supply monitor. The common reset output delay can be configured to use a preset 200ms timeout, can be programmed by an external capacitor or disabled.

In addition to providing a clever, multipurpose solution to supply monitoring, the low quiescent current of 50µA and the tiny DFN package make the LTC2909 an ideal choice in low-voltage, space-limited applications. The onboard 6.5V shunt regulator,

Power Supply Monitoring

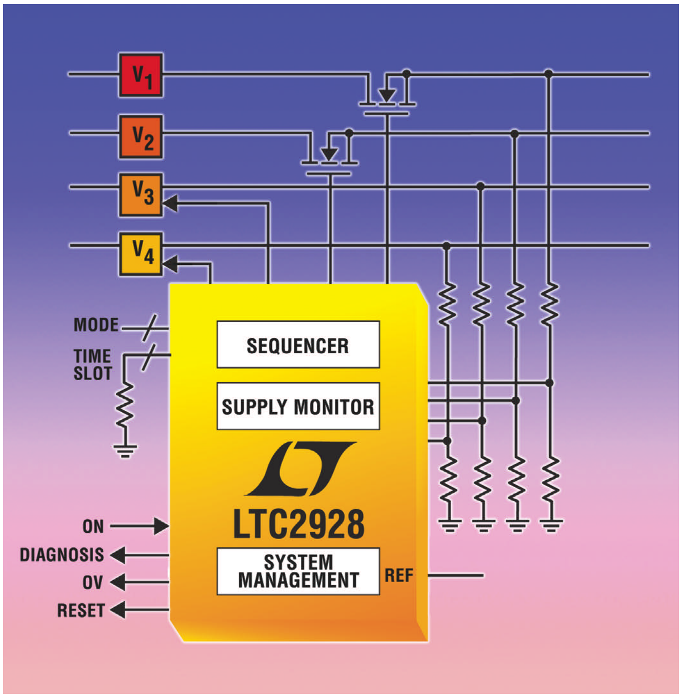

Figure 4. Simplified LTC2928 Functional Diagram

however, permits operation from a high voltage supply. Thus, the LTC2909 fits perfectly in small and portable appliances as well as in network servers and automotive applications.

Voltage Monitoring Combined with Supply Sequencing

Besides ongoing supply monitoring, for many high-end applications it is also critical to ensure proper sequencing of the supplies within the system. The right order of power up will prevent latch-up conditions, which can create system problems or damage important and expensive components. Realizing the importance of the two, Linear Technology introduced the LTC2928, a high accuracy quad supervisor and cascadable power supply sequencer for use with external N-channel FETs or power supplies with shutdown pins.

Precision input comparators with individual outputs monitor power supply voltages to ±1.5% accuracy. Supervisory functions include UV/OV monitoring and reporting, as well as µP reset generation. A reset will be issued should a supply voltage fall below its monitored value. A simplified functional diagram is shown in figure 4. The LTC2928 also features a buffered reference output that permits negative power supply sequencing and monitoring. Application faults will shut down all system supplies, and the type and source of faults are reported, enabling system diagnosis without software.

During the sequence up or down events, a power-good timer acts as a watchdog for stalled supplies. The sequencing order and timing are configured with just a few external components, allowing for effortless design changes during system development. In addition, the LTC2928 is easily cascadable, making sequencing and monitoring of unlimited number of supplies smooth and simple.

The LTC2928 operates from 3.3V to 16.5V, making it suitable in a wide range of applications requiring sequencing and monitoring of multiple I/O and core voltages.

Specified over the commercial and industrial temperature ranges, the LTC2928 is offered in 5mm x 7mm 38-lead QFN package and 36-lead SSOP packages.

BY COLIN MACDONALD AND ANIS JARRAR • FREESCALE SEMICONDUCTOR

Dynamic frequency scaling optimizes SOC performance

A CLOCK-EXTENSION SCHEME ALLOWS A DESIGN TO RUN AT ITS TARGET OPERATING FREQUENCY WHEN THE SYSTEM IS NOT ACCESSING THE KEY SLOW PATH AND TO SLOW DOWN WHEN IT IS.

Most IC designs involve engineering trade-offs. For example, when determining the operating frequency of an SOC (system on chip), the system architect has to consider how the frequency will impact design attributes, such as power consumption and die size. Another common consideration is the timing of a "key" path that significantly affects application or benchmark performance. Sometimes, despite all efforts, the worst-case delay through this path does not fit neatly into an integer number of clock periods. In this case, the architect needs to either accept the slack and any corresponding extra wait states in the path or, if the performance of this key path is of overriding importance, reduce the operating frequency of the SOC so that it requires fewer wait states and optimizes the path timing.

The first scheme has the advantage of the SOC's running at the original target frequency, but, as a result, it could compromise the performance of the KSP (key slow path). For example, if the KSP requires 7 nsec and the corresponding clock period is 5 nsec, the KSP will absorb two clock periods, or 10 nsec, which is 3 nsec greater than necessary. This 3 nsec of slack essentially goes to waste, and the KSP is 42% slower than it needs to be. The advantage of the second scheme is that it improves KSP performance. However, reducing the operating frequency of the SOC reduces some broad measures of performance, such as operating frequency. In addition, the KSP may not be 100% in use, needlessly reducing the operating frequency of the SOC.

The following scheme provides a better option for designs with a KSP that you can decode and that the SOC is not constantly accessing. It allows the design to run at its target operating frequency when the system is not accessing the KSP and to slow down when it is. Therefore, the scheme's biggest benefit occurs either when a system uses the KSP to a significant but not dominating extent or when it uses the KSP in certain applications. Finally, the scheme supports multiple KSPs and does not require the distribution that higher frequency clocks require.

The key to this scheme lies in the clock-generation circuitry. The high-speed clocks in many of today's SOCs originate in an on-chip PLL (phase-locked loop). Typically, multiple clocks, which have integer-multiple relationships, are distributed in the SOC's sea of gates. For example, an SOC may have a 400-MHz CPU clock, a 200-MHz bus clock, and a 100-MHz peripheral clock. Because many paths run between these clock domains,

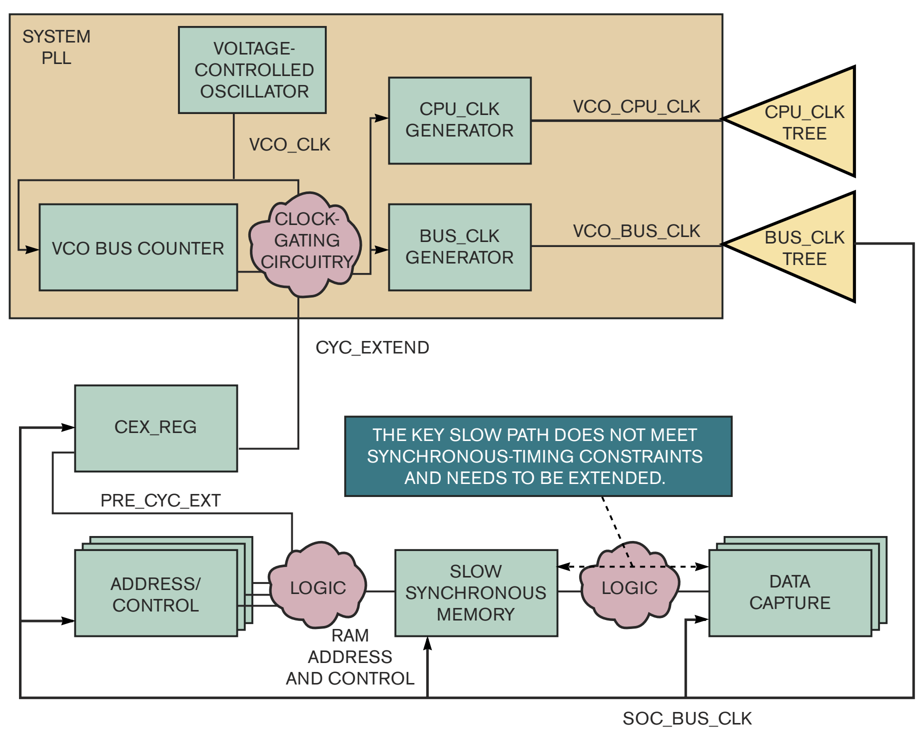

Figure 1 You can implement the clock-extension scheme in an SOC with one or many KSPs.

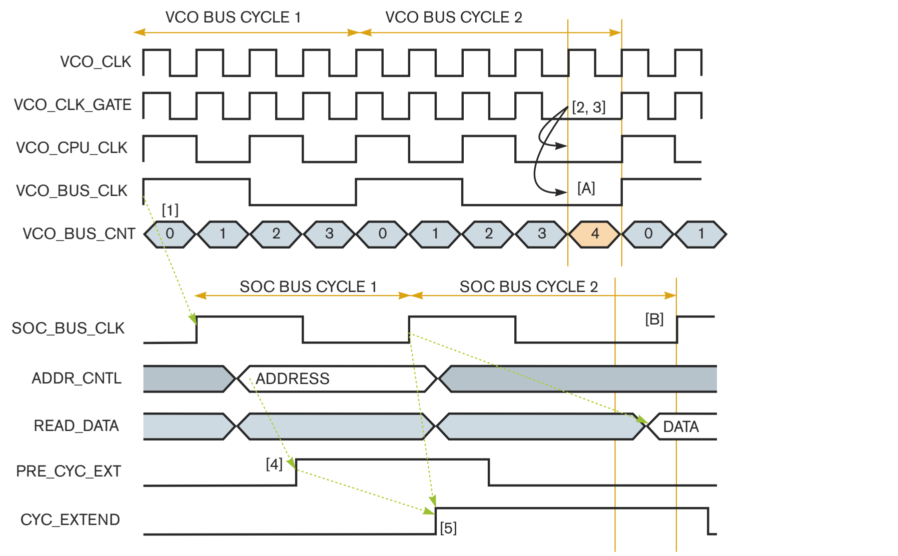

NOTES:
1. CLOCK TREE CAUSES DELAY.
2. VCO_CLK IS GATED, AND CYC_EXTEND IS ASSERTED ON COUNT 3.
3. VCO GATING DELAYS NEXT RISING-EDGE CPU AND BUS CLOCKS. VCO_BUS_CNT INCREMENTS TO 4.
4. PRE_CYC_EXT IS DECODED FROM SLOW RAM ADDRESS AND CONTROL.
5. CYC_EXTEND, THE REGISTERED VERSION OF PRE_CYC_EXT, ASSERTS VCO BUS CYCLE 2.
[A] EXTRA VCO CYCLE IN VCO BUS CYCLE 2 TRANSLATES INTO [B], AN EXTRA VCO CYCLE IN SOC BUS CYCLE 2, GIVING ENOUGH TIME FOR READ DATA FROM THE SLOW MEMORY TO BE LATCHED AT THE RISING EDGE OF BUS_CLK.

Figure 2 Waveforms of a VCO and an SOC signal illustrate how a VCO cycle extends clocks.

these clocks typically match according to insertion delay and align according to skew. Invariably, these clocks derive from a faster clock that the VCO (voltage-controlled oscillator) generates inside the PLL. The VCO clock runs at least two—and typically more than four—times faster than the fastest distributed clock. Although this VCO clock provides higher granularity for timing circuits, its duty cycle is usually uneven, and widely distributing it would significantly increase the power consumption of the SOC.

Frequently, the function of a module can determine some of its clock-domain frequencies. This situation is especially true with industry-standard communication blocks, such as USB (Universal Serial Bus) and PCI (Peripheral Components Interconnect). In addition to these "fixed"-frequency clocks, the SOC also has the CPU, bus, and peripheral clocks and their inverted partner clocks, all of which the system PLL generally drives. IC synthesis optimizes paths that cross these domains and times them to ensure that the system meets all synchronous-timing constraints. It is therefore crucially important, when modifying clock signals to improve performance, that you do not worsen the timing relationships between the synchronized domains; otherwise, the scheme cannot guarantee correct operation of the SOC.

Further, because the CPU and related clocks are divided versions of the VCO clock, a gated cycle of the VCO clock extends all derived clocks by the same amount. This situation does not worsen the timing relationships between these derived clocks, and the period of each distributed clock stretches by a VCO period. So, for example, if there is a 10-nsec bus clock, a 2.5-nsec VCO clock, and a KSP that takes 14.6 nsec, the bus clock (and CPU and peripheral clocks) can extend by two VCO clocks during KSP accesses to allow a 15-nsec cycle access.

In **Figure 1**, the KSP is the datapath from a slow synchronous memory, which bus_clk clocks, to a capture register, which bus_clk also clocks. The memory's chip-enable and read-control line in the previous cycle identify the assertion of the KSP cycle. This cycle identifier needs to reach the VCO-gating logic during the formation of the bus_clk cycle that, when it propagates through the clock tree, latches the read data from the memory. To ensure that the cycle-extend signal reaches the VCO-gating logic during the KSP cycle, the pre_cyc_ext signal uses the soc_bus_clk to create a cyc_extend input to the PLL. A state machine, which vco_clk clocks, qualifies the cyc_extend signal, gating a single VCO-clock pulse.

The cyc_extend signal needs to reach the VCO-clock-gating logic in time to gate the last VCO clock in the targeted bus_clk. This limitation of the scheme is important—especially when, unlike in this example, the system cannot identify the KSP in the previous cycle. Implementations are easier when the clock tree is short, the cycle time is long, and the VCO runs at a high multiple of the KSP clock. However, if the clock-tree

FEASIBILITY AND USEFULNESS

Use the feasibility and usefulness equations to determine whether you can and should implement the dynamic frequency scaling.

Feasibility equation:

clock_period − vco_period > clock_tree_latency + extend_signal_delay

Usefulness equation:

path_slack > vco_period

Definition of terms:

- **clock_period: period of KSP_clock.**
- **KSP_clock: clock that captures data on key slow path.**
- **clock_tree_latency: clock-tree insertion latency of KSP_clock, as measured from the VCO (voltage-controlled-oscillator) divider worst-case PVT (pressure, voltage, and temperature).**
- **extend_signal_delay: cycle-extend signal delay, which is measured from the rising edge of the KSP (key-slow-path)_clock in the extended cycle to the clock-gating circuitry inside the PLL (phase-locked loop).**
- **path_slack: positive_slack on the KSP, as measured to the first KSP_clock that gives a positive_slack.**
- **vco_period: period of VCO that generates KSP_clock.**

insertion delay exceeds the clock period, implementations can be difficult, if not impossible. Also, if the slow path effectively requires almost an entire additional clock, adding wait states is a better approach (see **sidebar** "Feasibility and usefulness").

Figure 2 shows the waveforms associated with the example. The top five waveforms show signals inside the PLL. The gated version of the raw VCO clock, vco_clk_gate, drives the circuitry that generates the CPU clock, vco_cpu_clk, running at half the VCO-clock frequency and the bus clock, vco_bus_clk, running at a quarter of the VCO frequency. A simple state machine produces a vco_bus_cnt output that qualifies the cyc_extend signal; vco_bus_cnt must equal three to gate the clock. The lower five waveforms represent key signals in the SOC. Representing the bus clock at the clock-tree endpoints is soc_bus_clk—which is essentially vco_bus_clk delayed by some PLL circuitry and the bus-clock tree. The soc_bus_clk clocks the address and control signals, addr_cntl, that interface to the slow memory. The read-data bus that forms the KSP is read_data. The KSP-cycle identifier that occurs in the previous bus-clock cycle (SOC Bus Cycle 1) is pre_cyc_ext. The registered version of pre_cyc_ext that feeds into the enhanced PLL is cyc_extend.

Notice the complicating effect of the clock-tree delay, which can vary dramatically across processes, temperatures, and voltages. To ensure the cyc_extend signal does not reach the PLL in time to gate VCO Bus Cycle 1, you must register the pre_cyc_ext signal using soc_bus_clk. Because the bus-clock-tree delay effectively advances vco_bus_clk with respect to soc_bus_clk, there is less time for the registered cyc_extend signal to reach the clock-gating circuitry. This step must occur before rising vco_clk, when vco_bus_cnt equals three. Notice also that extending the vco_bus_clk and vco_cpu_clks does not worsen the timing between these domains. The cycle lengthening only makes the synchronous constraints easier to meet and does not push the capturing clock edge into the next cycle, making its timing more difficult.

SYNOPSYS' PRIMETIME OR CADENCE'S ENCOUNTER MEASURE ALL THE DESIGNATED PATHS WITHIN A DOMAIN TO DETERMINE WHETHER THEY MEET A SYNCHRONOUS-TIMING CONSTRAINT.

Figure 1 illustrates a simple implementation, with only one KSP. However, the scheme can easily accommodate systems with multiple KSPs requiring different amounts of clock extending. For example, in an SOC with VCO (in the system PLL) running at an 8× bus clock and a 4× CPU clock, the PLL could potentially have as many as seven cycle_extend inputs, ranging from one VCO delay to seven VCO delays. KSPs in the CPU-clock domain could use inputs corresponding to delays one to three, and KSPs in the bus-clock domain can use all seven inputs. The assertion of the different cycle_extend signals can overlap; the larger delay dominates.

It's important to expand this scheme beyond certain types of slow paths, such as memory datapaths. The scheme can target any slow path that fits the timing and performance criteria. For example, an SOC may include a multiplier in an execution unit—not necessarily part of the CPU—that the CPU clocks. Now, suppose that this multiplier meets synchronous-timing constraints at the target operating frequency for byte and word operands—for example, 133 MHz—but not for long word operands, such as 117 MHz. Unlike the slow-memory example, in which you can add wait states to meet the target-frequency goal, the only option open to designers that avoids rebuilding the multiplier might be to reduce the SOC's operating frequency. However, doing so would sacrifice the overall device performance for what might be a relatively rare occurrence—long word multiplies in a specific execution unit. Alternatively, if the system can identify long-word-multiply cycles, it can extend the CPU-clock periods only on occasions that merit it.

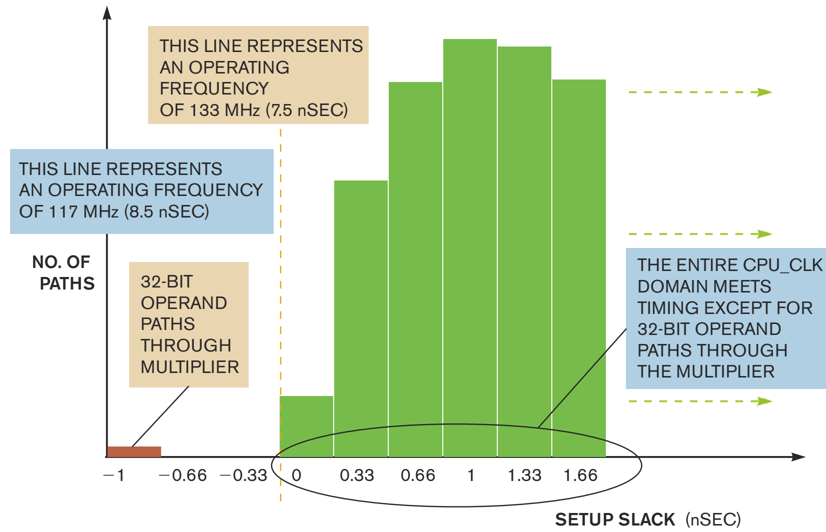

Figure 3 A static-timing report for an SOC with a "slow" multiplier generates a histogram of maximum delay-path slack for the CPU_CLK domain.

Figure 3 shows a representative graph of path slack for the CPU-clock domain. Running an STA (static-timing-analysis) tool, such as Synopsys' (www.synopsys.com) PrimeTime or Cadence's (www.cadence.com) Encounter, produces this type of graph. The tools measure all the designated paths within a clock domain to determine whether they meet a synchronous-timing constraint—input setup for maximum-delay paths or output hold for minimum-delay paths. This graph of setup slack illustrates how such problem KSPs can stand out during STA. The 1 nsec of negative slack on the multiplier paths effectively reduces the operating frequency of the CPU clock from 133 to 117 MHz.

Other considerations are the repercussions of varying the operating frequency. The SOC may contain modules that normally synchronize with the KSP clock. If these modules cannot tolerate a varying-input clock, the system must supply a stable clock, which cycle-extend signals do not affect. As this clock now moves in and out of phase with main SOC clocks, the interface to the module requires synchronizing circuitry.

The decision to implement this clock-extending scheme depends mainly on five factors. First, the paths in question must be sufficiently important that, for the cycles in which they are active, the performance improvements associated with the path more than offset the penalty of running the SOC at a lower operating frequency. Second, it must be possible to determine when the path is being used in sufficient time to extend the clock. Being able to identify the path in the previous cycle helps the system meet setup-timing constraints that can be tight because of clock trees' insertion delays. Having the system VCO run at a higher multiple of the extended clock helps, because there is more time to reach and block the final VCO clock. Having a shorter clock-tree-insertion delay also helps; too long a delay can prohibit use of this technique.

MORE AT EDN.COM

+ Go to www.edn.com/ms4198 and click on Feedback Loop to post a comment on this article.

Third, introducing wait states to achieve the benefit of slowing the clock is generally a preferable option, because wait states do not impact performance of the rest of the SOC. Fourth, you need to determine whether the SOC can tolerate a varying system-clock frequency, and, if a module requires synchronizing circuitry, you also need to determine whether the engineering effort is worth the performance benefits. Finally, path usage needs to be significantly less than 100%, because full usage translates into the SOC's always running at the lower frequency. However, low usage means that the performance benefits are likely to be small.

Although the scheme may not suit every failing timing path, it can yield significant benefits in certain instances. In today's highly competitive IC marketplace, designers need to explore all feasible options when trying to maximize device performance. EDN

AUTHORS' BIOGRAPHIES

Colin MacDonald is a principal design engineer at Freescale Semiconductor (Austin, TX). His main areas of focus are static-timing analysis, synthesis, and image-processing architectures. MacDonald holds a bachelor's degree with honors in natural philosophy from the University of Glasgow, Scotland.

Anis Jarrar is a principal design engineer at Freescale Semiconductor (Austin, TX), where he has worked since 1993. He is currently developing PowerPC- and ColdFire-based chips, focusing on synthesis, timing, and optimization techniques for deep-submicron designs. He has bachelor's and master's degrees in electrical engineering from the Georgia Institute of Technology (Atlanta).

Power and Event Sequencing with PSoC Express™

By Oliver Bailey
President
Timelines Industries, Inc.

The complexity of orderly power-up and shutdown has increased by an order of magnitude. Solving this problem with an off-the-shelf IC can be time-consuming and costly. One alternative – detailed in the article below – is a custom solution facilitated by Cypress's PSoC Express visual embedded design tool, the code-free development environment for the flexible Programmable System-on-Chip (PSoC) mixed-signal array.

PSoC Express Design Process

With PSoC Express a design engineer can define the power sequence visually without the need to write any microcontroller code. To implement a power sequencing component, the following steps are needed for each power sequence event:

1. A power input that fires the startup sequence
2. An output to control the next sequence event
3. A function (called "Valuator" in PSoC Express) that provides the "glue" logic

For the sake of clarity, a single input will be used to monitor and invoke the shutdown sequence. The final project is shown in Figure 1.

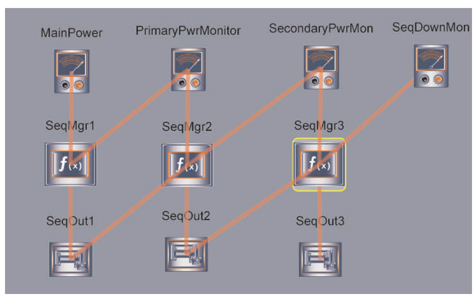

Figure 1

Each column represents one sequenced power rail. At the top is the power rail input. At the bottom is the output that enables (or disables) the next power rail. The object between the input and output objects is the sequence manager, a valuator where all of the logic is performed. The valuators contain a state machine, which sets the state based on changes in the inputs. Figure 2 shows the state machine for SeqMgr.

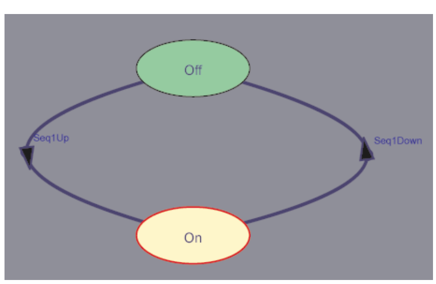

Figure 2

Continued on page 2

From page 1

The transition from Off to On is determined by the input value of MainPower and is a single line expression as shown in Figure 3.

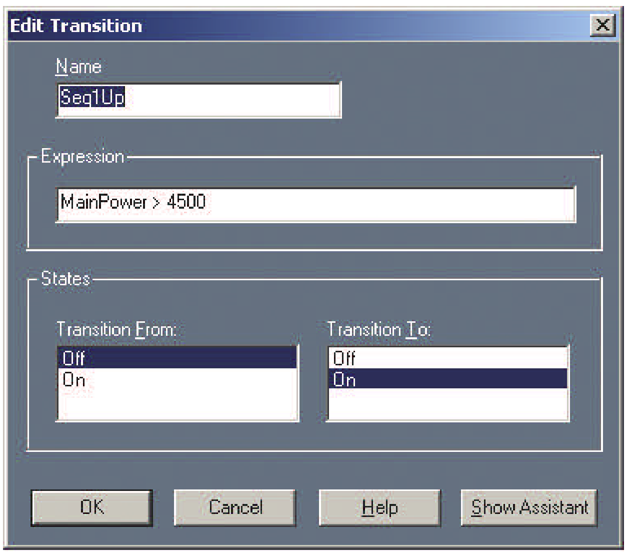

Figure 3

When the MainPower input goes above 4500 or 4.5 Volts DC, a state transition from Off to On will occur. To shut the circuit down in reverse order, the On to Off transition takes into account the power level of PrimaryPower as shown in Figure 4.

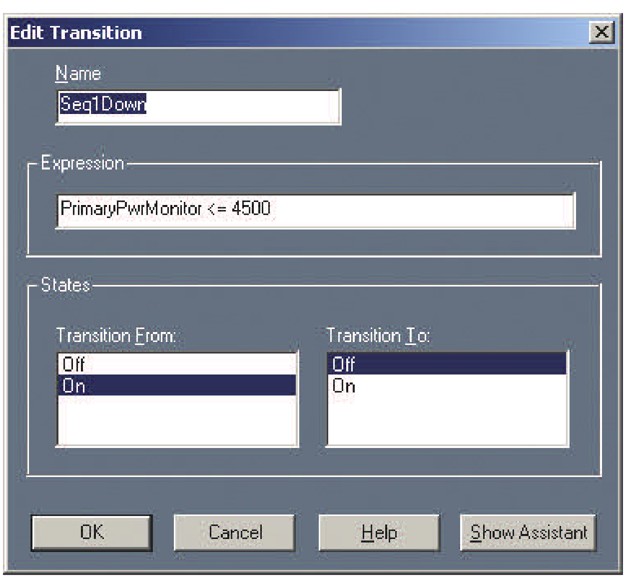

Figure 4

The On to Off transition occurs when PrimaryPowerMonitor falls below a value of 4500 or 4.5 VDC. To control the SeqOut1 output (the power supply enable) a transfer function within the output looks at the state of the State Machine (SeqMgr1) and sets its output to match, as shown in Figure 5.

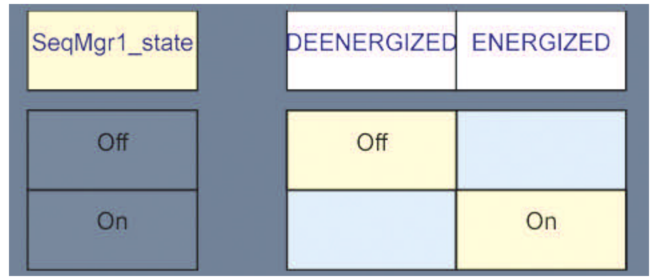

Figure 5

Input State Assignment

Using PSoC Express, we make the assignments of input states to output states by dragging the tiles from the left to the desired output state. During the startup sequence the PrimaryPower and SecondaryPower are dependent upon the prior column having its output turned on (which enables that particular voltage supply). To ensure this is taken into account in each sequence manager, the prior column's output is checked as shown in Figure 6, which shows the expression for the Off to On transition.

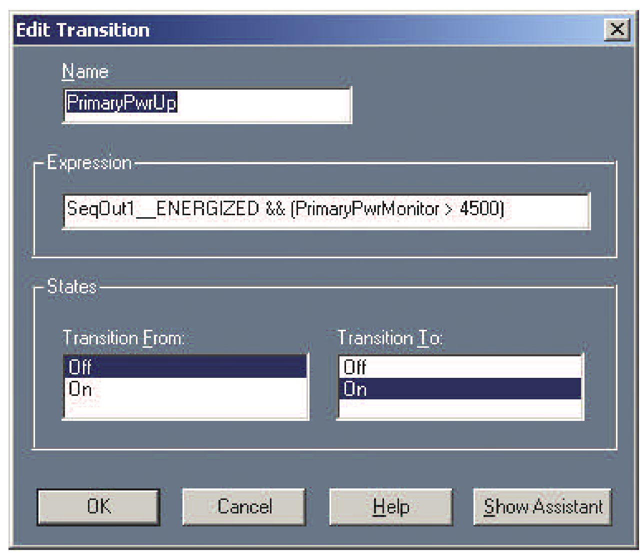

Figure 6

The PrimaryPower input and the state of SeqOut1 output are both required to change state. For columns two and three, the prior columns output must be powered for the Off to On transition. For the proper shutdown sequence to occur, the next column's Input value must drop below the threshold (here 2450 or 2.45VDC). Figure 7 shows the On to Off transition to trigger the shutdown of SeqOut3 (SecondaryPower).

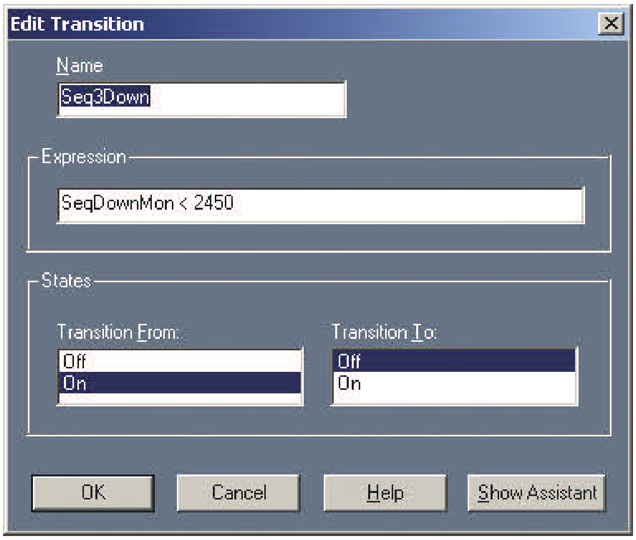

Figure 7

The Input SeqDownMon could be a drop in voltage on MainPower or a different power source or a condition that would trigger the cascading shutdown. In this demo the outputs were all defined as 0-5V relays, which provides a lot of flexibility. PSoC Express allows a wide variety of input and output definitions. LEDs, motors, FET's, relays, logic level, and many other types of inputs and outputs exist. In the end, a programmable device requires program code, and PSoC Express generates the code for you, using C and assembly language as appropriate. Listing 1 shows C code generated by PSoC Express that executes during the SeqMgr1 Off state.

```
case ID_SeqMgr1_state_Off:
{
  if (SystemVars.ReadOnlyVars.MainPower > 4500)
  {
     SystemVars.ReadOnlyVars.SeqMgr1_transition
     = SeqMgr1_transition__Seq1Up;
     SystemVars.ReadOnlyVars.SeqMgr1_state
     = SeqMgr1_state__On;
  }
  else if (1)
  {
     SystemVars.ReadOnlyVars.SeqMgr1_transition
     = SeqMgr1_transition__NoTransition;
  }
  break;
}
       Listing - 1
```

PSoC Express generates the code, then compiles and links this code to a hex file used to program the target processor. As a part of this process, the developer chooses which PSoC device (from a list of all PSoC devices that support the design) and can assign functions to specific pins or allow PSoC Express to assign them automatically. Figure 8 shows the screen displayed after building the project, which also provides access to all schematics and a BOM as well.

PSoC Express generates a Bill-Of-Materials (BOM) based on the inputs and outputs chosen. The BOM for our project is listed in Figure 9.

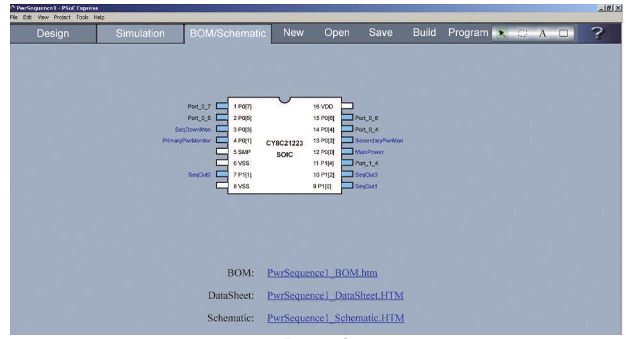

Figure 8

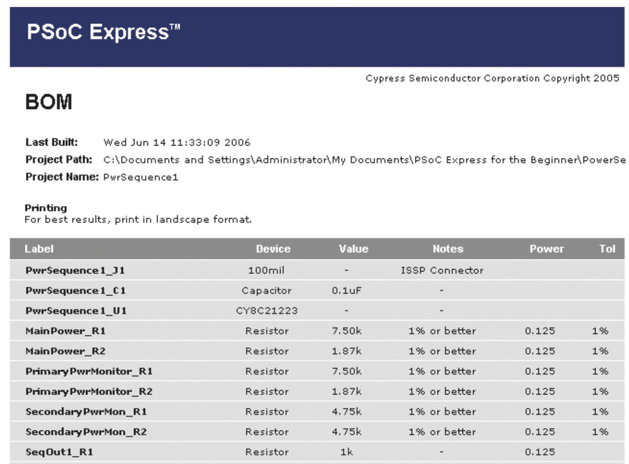

PSoC Express™

Cypress Semiconductor Corporation Copyright 2005

BOM

Last Built: Wed Jun 14 11:33:09 2006
Project Path: C:\Documents and Settings\Administrator\My Documents\PSoC Express for the Beginner\PowerSe
Project Name: PwrSequence1

Printing
For best results, print in landscape format.

Label	Device	Value	Notes	Power	Tol
PwrSequence1_J1	100mil	-	ISSP Connector		
PwrSequence1_C1	Capacitor	0.1uF	-		
PwrSequence1_U1	CY8C21223	-	-		
MainPower_R1	Resistor	7.50k	1% or better	0.125	1%
MainPower_R2	Resistor	1.87k	1% or better	0.125	1%
PrimaryPwrMonitor_R1	Resistor	7.50k	1% or better	0.125	1%
PrimaryPwrMonitor_R2	Resistor	1.87k	1% or better	0.125	1%
SecondaryPwrMon_R1	Resistor	4.75k	1% or better	0.125	1%
SecondaryPwrMon_R2	Resistor	4.75k	1% or better	0.125	1%
SeqOut1_R1	Resistor	1k	-	0.125	
SeqOut1_Q1	Transistor	2N2222	NPN		

Figure 9

PSoC Express generates a detailed schematic of interface circuitry for the selected drivers. Figure 10 shows the schematic for our project.

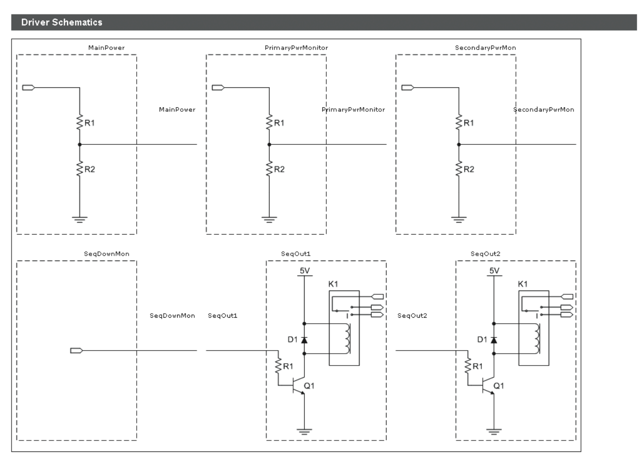

Figure 10

PSoC Express Benefits

This power-sequencing project took less than 30 minutes to complete. Building a power sequencer or some other specialized ladder logic component using PSoC Express cuts the design time considerably, PSoC Express makes any project based on state machines and logic tables quick and easy. A Designer need only specify design requirements visually and PSoC Express does the rest, generating the processor-specific C and assembly code, and a project specific schematic with a bill-of-materials.

Powerful Embedded Design Tools That Speed Time-to-Market

Software Tools

PSoC Designer and PSoC Express provide a choice in Windows®-based integrated development environments.

PSoC Designer™

A complete development environment, including a C Compiler and Assembler.

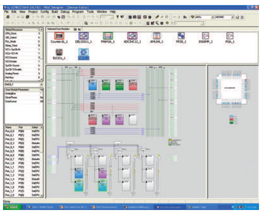

Device Editor: Select, Place, and Configure User Modules, and Resources, and Pins

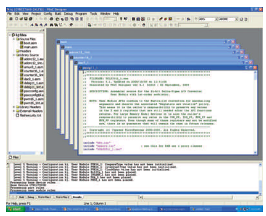

Application Editor: C Compiler, Assembler, Librarian

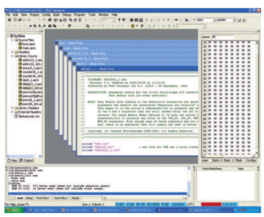

Debugger: In-Circuit Emulation, Break/Event Points, Trace

PSoC Express™

The Industry's first virtual embedded design tool, allowing designs to be completed without writing a single line of code.

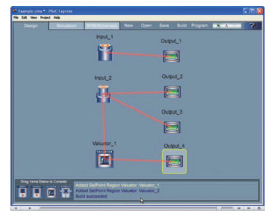

1. Select Inputs and Outputs

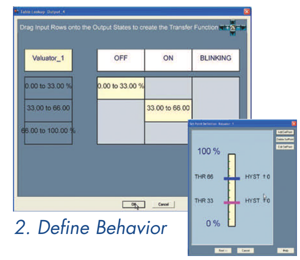

2. Define Behavior

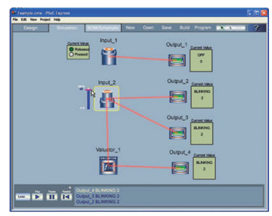

3. Simulate and Verify

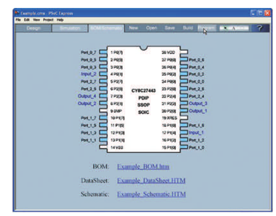

4. Build: Generate Hex Files, Produce Project Reports, BOMs, and Schematics

Hardware Tools

For evaluation to prototyping and emulation, Cypress offers development kits to suit your design requirements.

Evaluation Kits

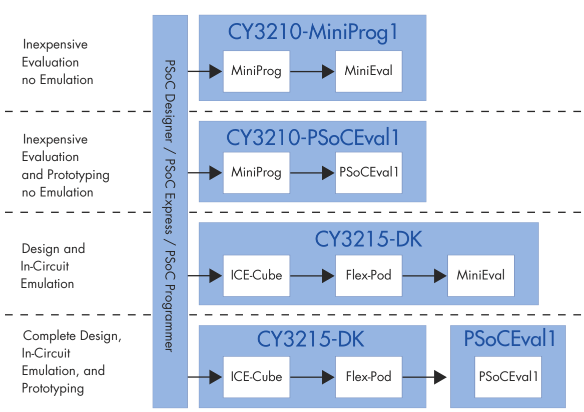

CY3210-MiniProg1

CY3210-PSoCEval1

CY3215-DK

Development Kits

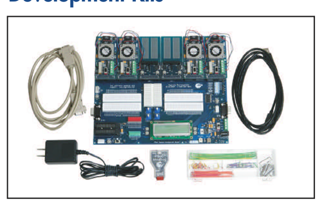

PSoC Express – CY3210-ExpressDK

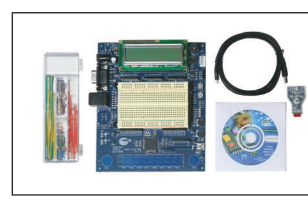

PSoC w/USB – CY3214-PSoCEvalUSB

PSoC CapSense – CY3212-CapSense

PRoC – CY3653

In-System Programming Tools

- In-System or In-Socket
- Prototype or Production
- Protection Circuitry
- Industrial Grade

CYPRESS
PERFORM

TI MS430 Series Microcontrollers

Ultra-low power microcontrollers with a variety of peripherals. Includes16-bit RISC CPU, 16-bit registers, constant generators for maximum code efficiency. Digitally controlled oscillator allows wake-up from low-power to active mode in less than 6µS.

Texas Instruments mouser.com/ti/a

Tyco/AMP PCI Express

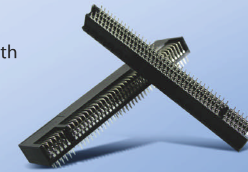

New serial I/O technology compatible with the current PCI software environment. Provides low-cost, scalable performance for the next generation of computing and communications platforms, while also providing a long-term path to improved I/O slots.

tyco Electronics Authorized Distributor mouser.com/tyco/a

The Mouser Advantage: Faster Time to Market for YOUR New Designs!

- Most Rapid New Product Introduction in the Industry!
- The ONLY New Catalog Every 90 Days!
- Easy to Browse-and-Buy over 700,000 Products Online!
- New Products Added Daily

Mouser's Time-to-Market Advantage!

Mouser's rapid introduction of new products and new technologies is faster than any other distributor in the industry. We are the only distributor to publish a 1,800+ page print catalog every 90 days! And with daily new product updates to over 700,000 products online, you can depend on Mouser to save you critical time to market for your new design projects.

Mouser Delivers!

The NEWEST products, NEWEST technologies, best-in-class technical customer service with NO minimums, and streamlined warehouse systems that deliver same-day shipping on most orders. We make it easy to do business with Mouser!

mouser.com **(800) 346-6873**

MOUSER ELECTRONICS
a tti company

Newest Products for Your New Designs

Omron Non-Amplified Photomicrosensors

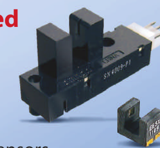

Several transmissive photomicrosensors from Omron. These photomicrosensors eliminate most alignment and aiming problems. The emitter and receiver are in a fixed "U-shape" housing with a wide range of mounting styles and slot widths. Transistor outputs include photo-IC, photo-transistor, and photo-Darlington.

OMRON ELECTRONIC COMPONENTS mouser.com/omron/a

Vishay/Sprague 94SVP Vertical Surface Mount Capacitors

Features superior heat-proof characteristics. Effective when used as surface mount devices for switching power supplies. Rated ripple current values guaranteed at 105°C.

VISHAY mouser.com/sprague/a

ON Semiconductor MC33 Series AC-DC Switching Off-line Controllers

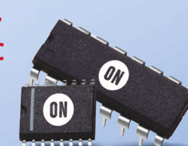

Specifically designed for AC-DC switching off-line power management and DC-DC converter applications. High-speed, fixed frequency, double-ended pulse width modulator controllers for high-frequency operation and easy configuration.

ON Semiconductor mouser.com/onsemiconductor/a

The NEWEST Semiconductors | Passives | Interconnects | Power | Electromechanical | Test, Tools & Supplies from Mouser Electronics

BY JIM WILLIAMS • LINEAR TECHNOLOGY CORP

Load-transient-response testing for voltage regulators

VARIATIONS OCCUR IN VOLTAGE REGULATORS' TRANSIENT LOADS; THUS, THE DEVICES REQUIRE CAREFUL EVALUATION AND TESTING.

Semiconductor memory, card readers, microprocessors, disk drives, piezoelectric devices, and digital systems create transient loads that voltage regulators must service. Ideally, regulator output would be invariant during a load transient. In practice, however, some variation occurs, and this variation becomes problematic if a system exceeds its allowable operating-voltage tolerances. This problem mandates testing the regulator and its associated support components to verify desired performance under transient-loading conditions. You can use various methods to generate transient loads and allow observation of regulator response.

Figure 1 shows a conceptual load-transient generator. The regulator under test drives dc and switched resistive loads, which may be manually variable. The device monitors its switched current and output voltage, permitting comparison of the output voltage and the load current under static and dynamic conditions. The switched current is either on or off; there is no electronically controllable linear region.

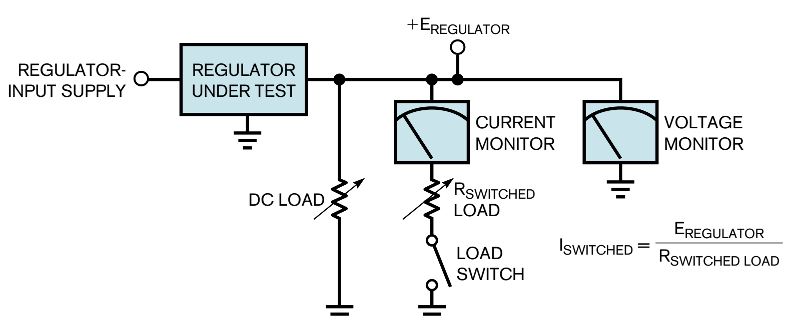

Figure 1 This conceptual regulator-load tester includes switched and dc loads and voltage and current monitors. The resistor values set dc and switched-load currents. The switched current is either on or off; there is no controllable linear region.

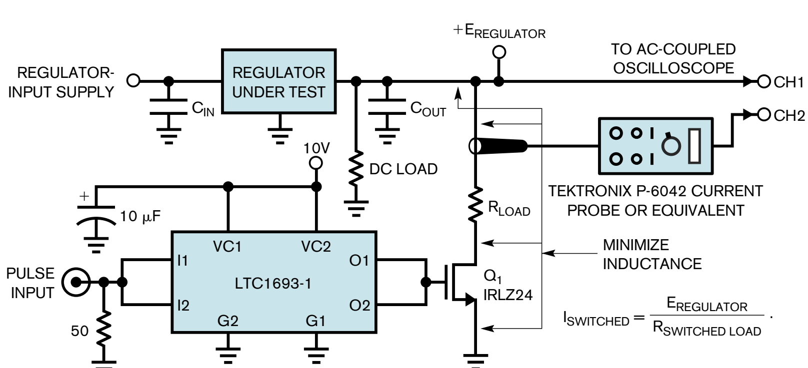

Figure 2 A practical regulator-load tester includes a FET driver and Q_1 switch. The oscilloscope monitors the current-probe output and regulator response.

Figure 2 shows a practical implementation of the load-transient generator. Capacitors augment the voltage regulator under test; these capacitors provide an energy reservoir, similar to a mechanical flywheel, to aid transient response. The size, dielectric, and location of these capacitors, particularly C_{OUT}, have a pronounced effect on transient response and overall regulator stability (**references** 1 and 2). The input pulse triggers the LTC1693 FET driver to switch Q_1, generating a transient-load current from the regulator. An oscilloscope monitors the instantaneous load voltage and, through a "clip-on" wideband-probe, current (see **sidebar** "Probing considerations for load-transient-response measurements"). **Figure 3** provides an evaluation of the circuit's load-transient-generating capabilities by substituting a low-impedance power source for the regulator. The combination of a high-capacity power supply, low-impedance connections, and generous bypassing maintains low impedance across frequency. **Figure 4** shows the circuit in **Figure 3**'s response to the LTC1693-1 FET driver (Trace A) by cleanly switching 1A in 15 nsec (Trace B). Such speed is useful for simulating many loads but has restricted versatility. Although fast, the circuit cannot emulate loads between the minimum and the maximum currents.

CLOSED-LOOP TESTERS

Figure 5's conceptual closed-loop load-transient generator linearly controls Q_1's gate voltage to set instantaneous transient current at any desired point, allowing simulation of nearly any load profile. Feedback from Q_1's source to the A_1 control amplifier closes a loop around Q_1, stabilizing its operating point. Q_1's current assumes a value that depends on the control-input voltage and the current-sense resistor over a wide bandwidth. Once A_1 biases to Q_1's

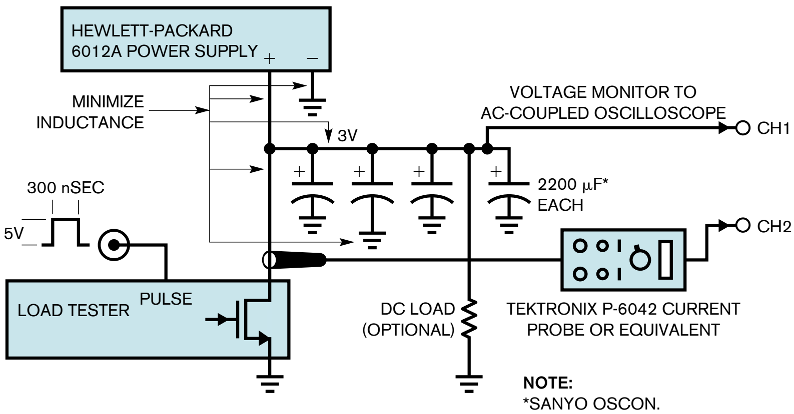

Figure 3 Substituting a well-bypassed, low-impedance power supply for the regulator lets you determine the load tester's response time.

conductance threshold, small variations in A_1's output result in large current changes in Q_1's channel. As such, A_1 need not output large excursions; its small signal bandwidth, rather than its slew rate, is the fundamental speed limitation. Within this restriction, Q_1's current waveform is the same shape as A_1's control-input voltage, allowing linear control of load current. This versatile capability permits a variety of simulated loads.

FET-BASED CIRCUIT

Figure 6 shows a practical incarnation of a FET-based closed-loop load-transient generator, including dc-bias and waveform inputs. A_1 must drive Q_1's high-capacitance gate at high frequency, necessitating high peak A_1 output currents and attention to feedback-loop compensation. A_1, a 60-MHz current-feedback amplifier, has an output-current capacity exceeding 1A. Maintaining stability and waveform fidelity at high frequency while driving Q_1's gate capacitance necessitates settable gate-drive-peaking components, a damper network, feedback trimming, and loop-peaking adjustments. You make the required dc trim first. Without applying an input, trim the 1-mV adjust for 1 mV dc at Q_1's source. You make the ac trims using **Figure 7**'s arrangement. Similar to the circuit in **Figure 3**, this "brick-wall"-regulated source provides minimal ripple and sag when the load-transient generator step-loads it. Apply the inputs as the **figure** shows and trim the gate drive, feedback, and loop-peaking adjustments for the cleanest square-cornered response on the oscilloscope's current-probe-equipped channel.

BIPOLAR TRANSISTORS

The circuit in **Figure 8** considerably simplifies the previous circuit's loop dynamics and eliminates all ac trims. The major trade-off is a halving of speed. The circuit is similar to the one in **Figure 6**, except that Q_1 is a bipolar transistor. The bipolar's greatly reduced input capacitance allows A_1 to drive a more benign load. This approach permits you to use an amplifier with lower output current and eliminates the dynamic trims necessary to accommodate **Figure 6**'s FET-gate capacitance. The sole trim is the 1-mV adjustment, which you accomplish as described. You can eliminate this trim at the cost of circuit complexity (see **sidebar** "A trimless, closed-loop-transient-load tester"). Aside from the twofold speed decrease, the bipolar transistor also introduces a 1% output-current error due to its base current. You add Q_2 to prevent excessive Q_1 base current when the regulator supply is absent. The diode prevents reverse-base bias under any circumstances.

CLOSED-LOOP-CIRCUIT PERFORMANCE

Figures 9 and **10** show the two wideband circuits' operation. The FET-based circuit (**Figure 9**) requires only a 50-mV A_1 swing (Trace A) to enforce Trace B's flat-topped current pulse with 50-nsec edges through Q_1. **Figure 10** details the bipolar-transistor-based circuit's performance. Trace A, taken at Q_1's

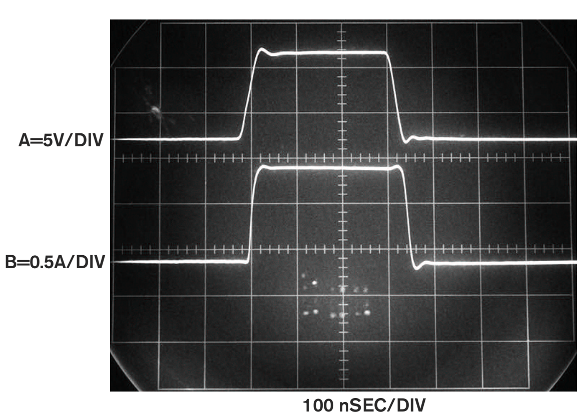

Figure 4 Figure 2's circuit responds to the FET driver's output (Trace A), switching a 1A load (Trace B) in 15 nsec.

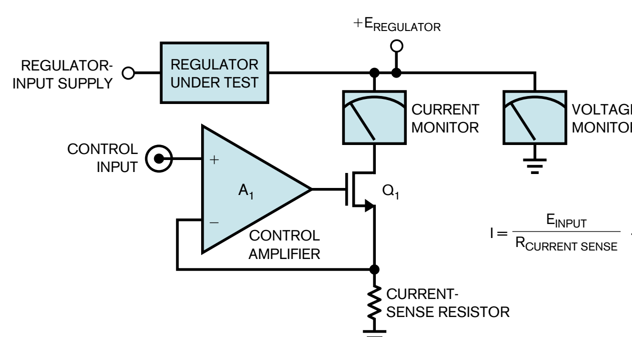

Figure 5 In this conceptual closed-loop-load tester, A_1 controls Q_1's source voltage, setting the regulator's output current. Q_1's drain-current waveshape is identical to A_1's input, allowing linear control of the load current. The voltage and current monitors match those in Figure 1.

DAVINCI™

TEXAS INSTRUMENTS

PROCESSORS | SOFTWARE | TOOLS | SUPPORT

Now that DaVinci products are here, your digital video innovations are everywhere.

That's the DaVinci Effect.

DaVinci™ Technology makes astounding creativity possible in digital video devices for the hand, home and car. The DaVinci platform includes digital signal processor (DSP) based SoCs, multimedia codecs, application programming interfaces, application frameworks and development tools, all of which are optimized to enable innovation for digital video systems. DaVinci products will save OEMs months of development time and will lower overall system costs to inspire digital video innovation. So what are you waiting for? You bring the possibilities. DaVinci will help make them real.

What is DaVinci?

Processors: Digital Video SoCs:

- TMS320DM6446 – Video encode/decode
- TMS320DM6443 – Video decode

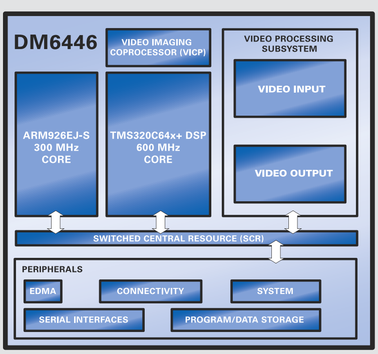

Performance Benchmarks:

STANDALONE CODECS	DM6446	DM6443
MPEG-2 MP ML Decode	1080i+ (60 fields /30 frames)	720p+
MPEG-2 MP ML Encode	D1+	n/a
MPEG-4 SP Decode	720p+	720p+
MPEG-4 SP Encode	720p+	n/a
VC1/WMV 9 Decode	720p+	720p+
VC1/WMV 9 Encode	D1+	n/a
H.264 (Baseline) Decode	D1+	D1+
H.264 (Baseline) Encode	D1+	n/a
H.264 (Main Profile) Decode	D1+	D1+

+ denotes available processor headroom for analytics and/or other features

Tools: Validated Software and Hardware Development

- DVEVM (Digital Video Evaluation Module)
- MontaVista Development Tools
- Code Composer Studio IDE

Software: Open, Optimized and Production Tested

- Platform Support Package
- MontaVista Linux Support Package
- Industry-recognized APIs
- Multimedia frameworks
- Platform-optimized, multimedia codecs:

- H.264	- AAC	- G.729ab
- MPEG4	- WMA9	- WMV9/ VC1
- H.263	- MP3	
- MPEG2	- G.711	
- JPEG	- G.728	
- AAC+	- G.723.1	

>>> For complete technical documentation or to get started with our Digital Video Evaluation Module, please visit **www.thedavincieffect.com**

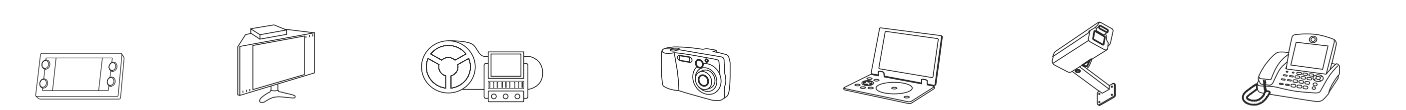

Portable Media Player | IP Set-Top Box | Automotive Infotainment | Digital Still Camera | Digital Video Innovations | Video Surveillance | Video Phone & Conferencing

Technology for Innovators™ TEXAS INSTRUMENTS

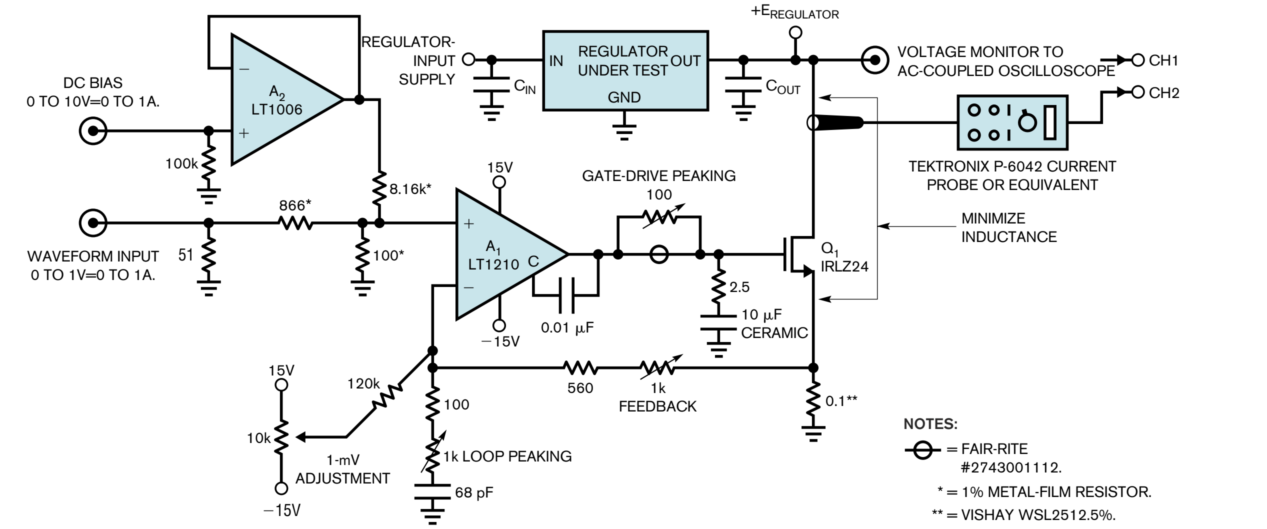

Figure 6 In a detailed closed-loop-load tester, the dc-level and pulse inputs feed A_1 to the Q_1 current-sinking-regulator load. Q_1's gain allows a small A_1 output swing, permitting wide bandwidth. The damper network, feedback, and peaking trims optimize edge response.

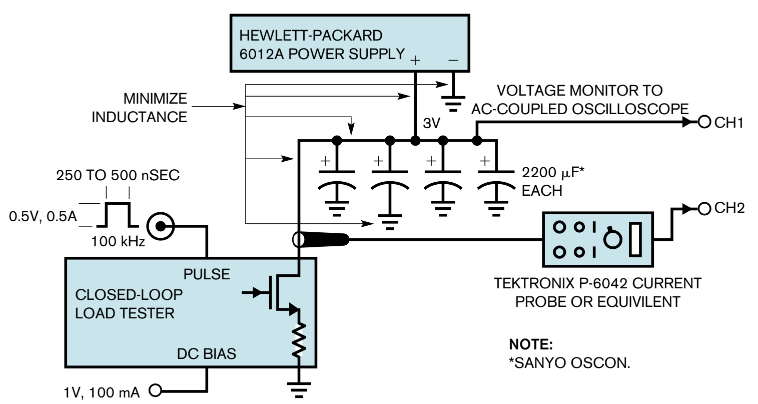

Figure 7 Determining the closed-loop-load-tester response time occurs as in Figure 3. A "brick-wall" input provides a low-impedance source.

base, rises less than 100 mV, causing Trace B's clean, 1A current conduction through Q_1. This circuit's 100-nsec edges, about two times slower than the more complex FET-based version, are still fast enough for most practical transient-load testing.

LOAD-TRANSIENT TESTING

These circuits permit rapid and thorough voltage-regulator load-transient testing. **Figure 11** uses **Figure 6**'s circuit to evaluate an LT1963A linear regulator. **Figure 12** shows regulator response (Trace B) to Trace A's asymmetrically edged input pulse. The ramped leading edge, within the LT1963A's bandwidth, results in Trace B's smooth 10-mV p-p excursion. The fast trailing edge, well outside the LT1963A's passband,

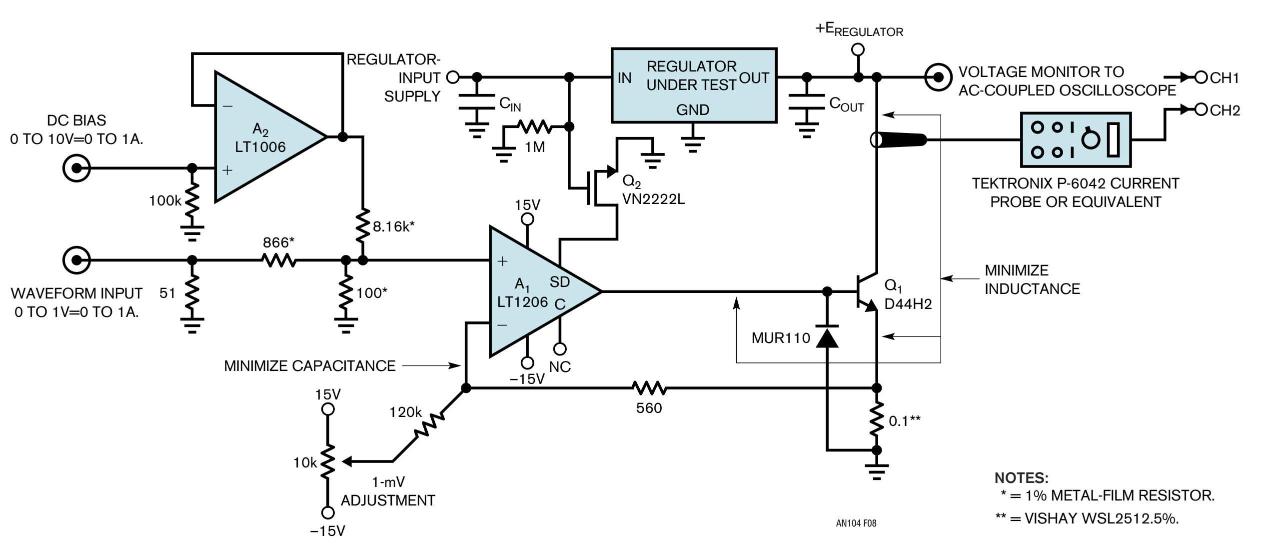

Figure 8 This circuit matches that of Figure 6 but with a bipolar transistor. Q_1's reduced input capacitance simplifies loop dynamics, eliminating compensation components and trims. The trade-off is a halving of speed and a base-current-induced 1% error.

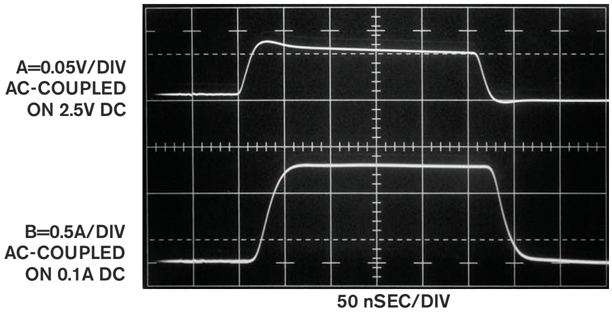

Figure 9 Figure 6's closed-loop-load-tester step response is quick and clean, showing 50-nsec edges and a flat top. (Q_1's current is Trace B.) A_1's output (Trace A) swings only 50 mV, allowing wideband operation. Trace B's presentation is slightly delayed due to voltage and current-probe time skew.

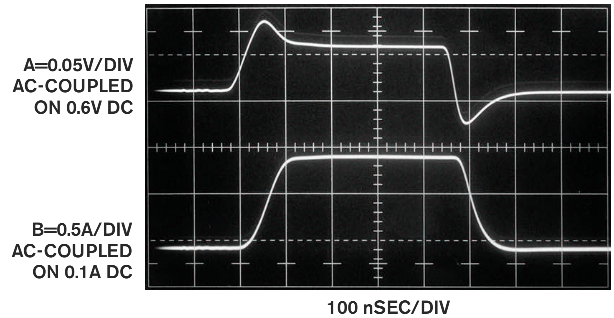

Figure 10 Figure 8's bipolar-output-load-tester response is two times slower than the FET version, but the circuit is simpler and eliminates compensation trims. Trace A is A_1's output, and Trace B is Q_1's collector current.

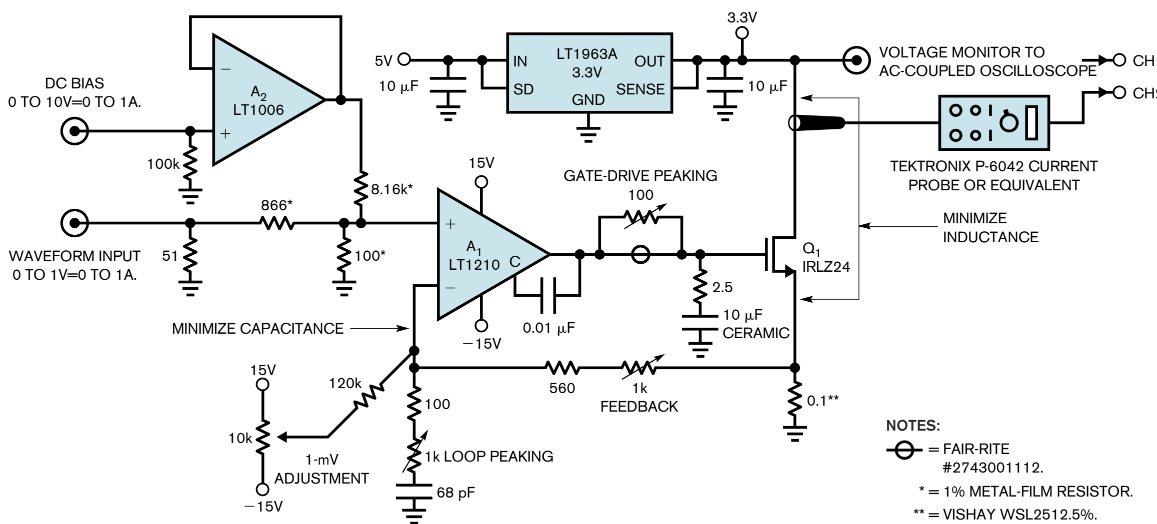

Figure 11 This closed-loop-load tester with an LT1963A regulator provides load testing for a variety of current and load waveshapes.

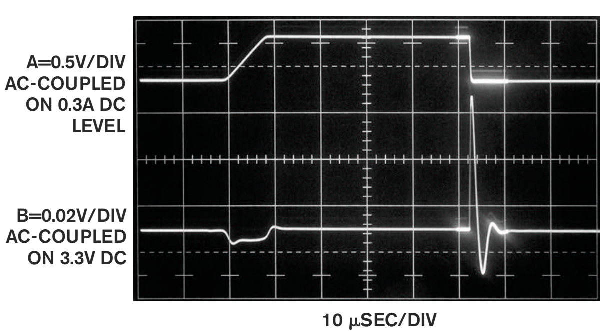

Figure 12 The circuit in Figure 11 responds (Trace B) to an asymmetrically edged pulse input (Trace A). A ramped leading edge within the LT1963A's bandwidth results in Trace B's smooth, 10-mV-p-p excursion. A fast trailing edge outside the LT1963A's bandwidth causes Trace B's abrupt 75-mV-p-p disruption. The photo intensifies the trace's latter portion for clarity.

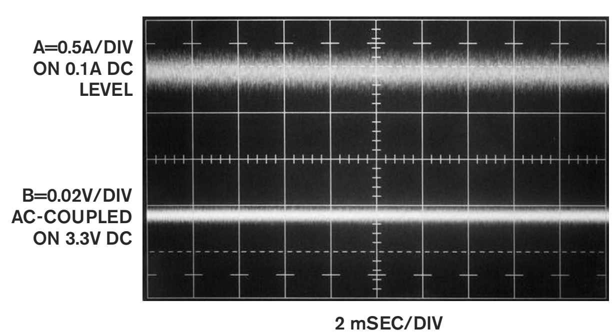

Figure 13 A 500-mA-p-p, 500-kHz noise load (Trace A) within the regulator's bandpass produces only 6-mV artifacts at Trace B's regulator output.

NXP– the new company born out of Philips Semiconductors
With all the energy of a new beginning, fuelled by 50 years of wisdom, NXP Semiconductors is ready to answer all your questions – especially the impossible ones. The ones that are still dreams but soon will be transformed into amazing sensory experiences.
Driven by the open mind and the curiosity of a newborn, NXP is already leading the world of vibrant media technologies. This leadership position is reflected in billions of dollars invested in research projects, the registration of more than 25000 patents and the numerous state-of-the-art solutions in the ID, Automotive, Mobile and Home sectors. Discover how your innovations can be seen, heard and felt in a completely new way at www.what-if-you-could.com, because the question from now on is: what would you like to question?
What if you could
NXP
NXP
founded by Philips
What if you could be born twice?
founded by
PHILIPS

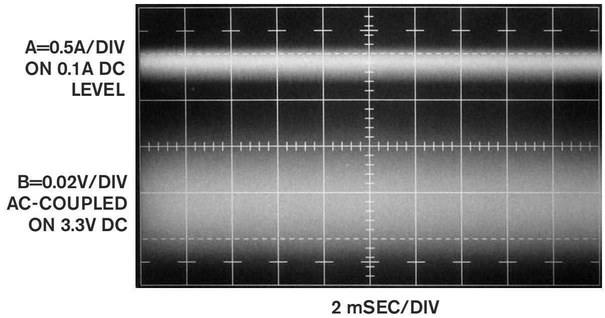

Figure 14 This waveform has the same conditions as Figure 13, except with increased noise bandwidth of 5 MHz, exceeding the regulator's bandwidth and resulting in 50-mV-p-p output error.

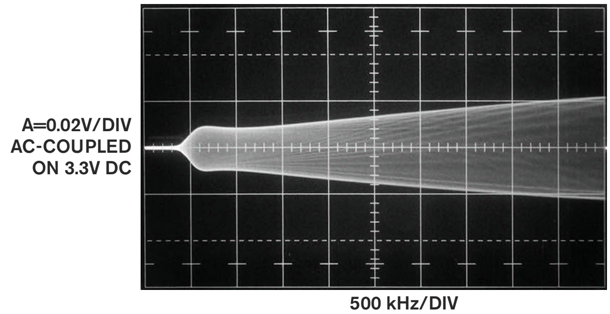

Figure 15 A swept, dc to 5-MHz, 0.35A load on 0.2A dc causes the regulator's output impedance to rise with frequency and correspondingly increases output error.

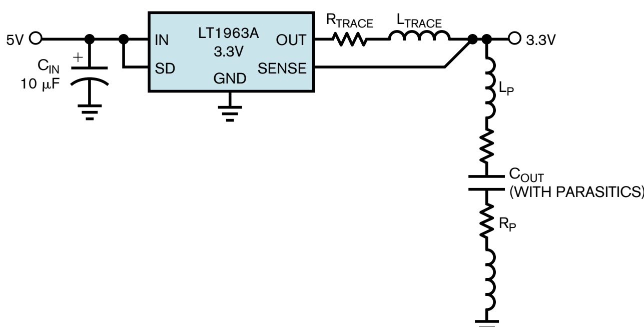

Figure 16 C_{OUT} dominates the regulator's dynamic response; C_{IN} is much less critical. Parasitic inductance and resistance limit the capacitor's effectiveness at frequency. The capacitor's value and dielectric significantly influence the load-step response. Excessive trace impedance is also a factor.

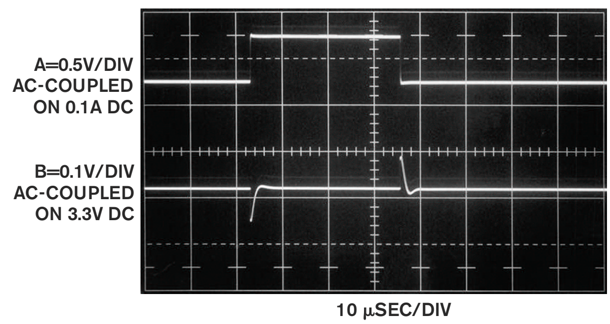

Figure 17 A stepped 0.5A load to Figure 16's circuit (Trace A) with C_{IN}=C_{OUT}=10 µF results in Trace B's regulator output. The use of low-loss capacitors promotes controlled output excursions.

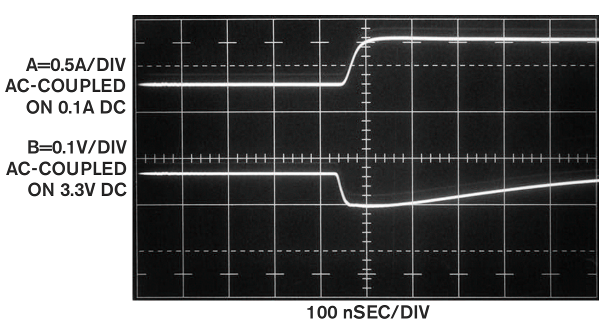

Figure 18 The expanding horizontal scale shows Trace B's smooth regulator-output response. Mismatched current- and voltage-probe delays account for slight time skewing.

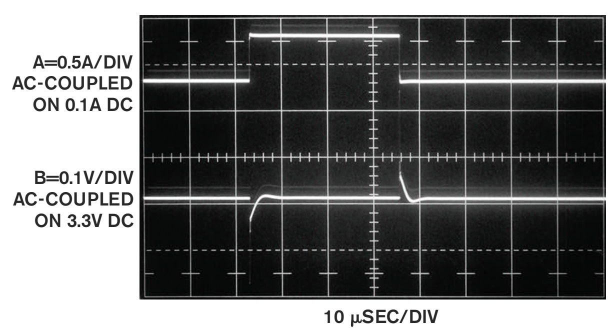

Figure 19 An "equivalent" 10-µF C_{OUT} capacitor to the one in Figure 17 shows performance that appears similar at 10 µsec/division.

causes Trace B's abrupt disruption. C_{OUT} supplies too little current to maintain output level, and a 75-mV-p-p spike results before the regulator resumes control. In **Figure 13**, a 500-mA p-p, 500-kHz noise load, emulating a multitude of incoherent loads, feeds the regulator in Trace A. This frequency is within the regulator's bandwidth, and only 6 mV p-p of disturbance appears in Trace B, the regulator output. **Figure 14** maintains the same conditions, except that noise bandwidth increases to 5 MHz. This increase exceeds regulation bandwidth, resulting in more than 50-mV p-p error, an eightfold increase.

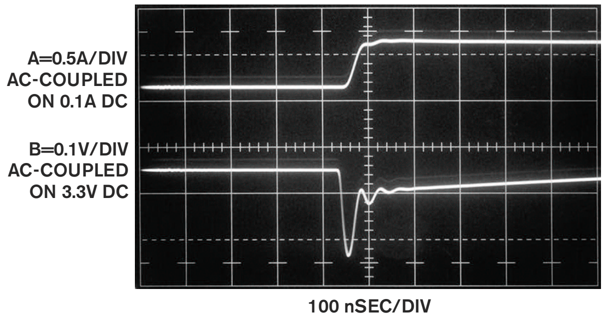

Figure 20 The horizontal-scale expansion reveals that the "equivalent" capacitor produces two times more amplitude error than the one in Figure 18. Mismatched probe delays cause time skewing between traces.

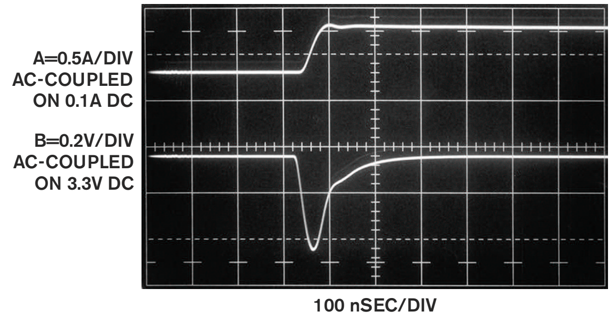

Figure 21 An excessively lossy 10-μF C_{OUT} allows a 400-mV excursion—four times Figure 18's amount. The time skewing between the traces derives from probe mismatch.

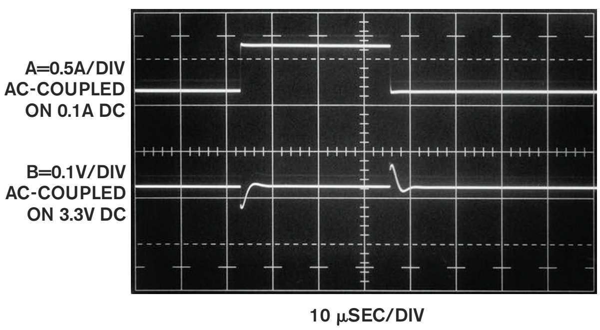

Figure 22 Replacing C_{OUT} with a low-loss, 33-μF unit yields a 40% smaller output-response transient than that of Figure 17.

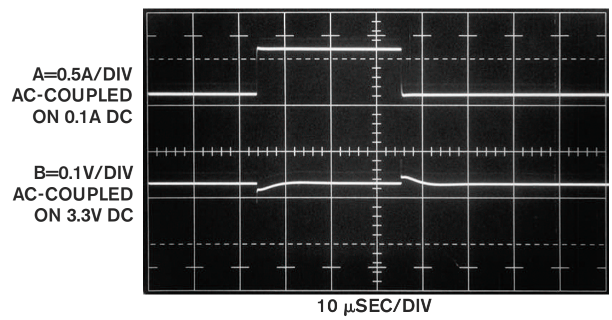

Figure 23 A low-loss, 330-μF capacitor keeps output-response transients to less than 20 mV—four times lower than Figure 17's 10-μF capacitor.

Figure 15 shows what happens when you present a 0.2A, dc-biased, swept, dc to 5-MHz, 0.35A load to the regulator. The regulator's rising output impedance versus frequency results in ascending error as frequency scales. This information allows determination of regulator output impedance versus frequency.

CAPACITOR'S ROLE IN REGULATOR RESPONSE

The regulator employs capacitors at its input (C_{IN}) and output (C_{OUT}) to augment its high-frequency response. You should carefully consider the capacitor's dielectric, value, and location because they greatly influence regulator characteristics (**references 1**, **2**, and **3**). C_{OUT} dominates the regulator's dynamic response; C_{IN} is much less critical, as long as it does not discharge below the regulator's dropout point. **Figure 16** shows a typical regulator circuit and emphasizes C_{OUT} and its parasitics. Parasitic inductance and resistance limit capacitor effectiveness at frequency. The capacitor's dielectric and value significantly influence load-step response. A "hidden" parasitic, impedance buildup in regulator-output-trace runs, also influences regulation characteristics, although you can minimize the parasitic's effects by remote sensing and distributed capacitive bypassing.

Figure 17 shows **Figure 16**'s circuit responding (Trace B) to a 0.5A load step biased on 0.1A dc (Trace A) with $C_{IN}=C_{OUT}=10$ μF. The circuit employs low-loss capacitors, resulting in Trace B's well-controlled output. **Figure 18** greatly expands the horizontal time scale to investigate high-frequency behavior. Regulator-output deviation (Trace B) is smooth with no abrupt discontinuities. **Figure 19** runs the same test as **Figure 17** using an output capacitor claimed as "equivalent" to the one that **Figure 17** employs. At 10 μsec/division, the scope photos seem similar, but **Figure 20** indicates problems. This photo, taken at the same higher sweep speed as the one in **Figure 18**, reveals the "equivalent" capacitor to have twice as much amplitude error, higher frequency content, and higher resonances than the one in **Figure 18**. (Always specify components according to observed performance, rather than salesmen's claims.) **Figure 21** substitutes a lossy 10-μF unit for C_{OUT}. This capacitor allows a 400-mV excursion (note Trace B's vertical-scale change), greater than four times **Figure 18**'s amount. Conversely, **Figure 22** increases C_{OUT} to a low-loss, 33-μF type, decreasing Trace B's output-response transient by 40% versus **Figure 18**. **Figure 23**'s further increase, to a low-loss, 330-μF capacitor, keeps transients inside 20 mV: four times lower than **Figure 18**'s 10-μF value.

The lesson is clear: Capacitor value and dielectric quality have a pronounced effect on transient-load response. Try before specifying!

RISE TIME VERSUS REGULATOR RESPONSE

The closed-loop-load-transient generator also allows investigating load-transient rise time on regulation at high speed. **Figure 24** shows **Figure 16**'s circuit ($C_{IN}=C_{OUT}=10$ μF), respond-

PROBING CONSIDERATIONS FOR LOAD-TRANSIENT-RESPONSE MEASUREMENTS

Signals of interest in load-transient-response studies occur within a bandwidth of approximately 25 MHz and a rise time of 14 nsec. This modest speed range eases probing techniques, but high-fidelity measurement requires some care. You measure load current with a dc-stabilized, Hall-effect, clip-on current probe such as the Tektronix (www.tektronix.com) P-6042 or A6302/AM503. The conductor loop in the probe jaws should encompass the smallest possible area to minimize introduced parasitic inductance, which can degrade measurement. At higher speeds, grounding the probe case may slightly decrease measurement aberrations, but this effect is usually small.

You perform voltage measurement, typically ac-coupled and ranging from 10 to 250 mV, using the arrangement in Figure A.

This arrangement feeds the measured voltage to a BNC 50Ω, back-terminated cable, which drives the oscilloscope through a dc-blocking capacitor and a 50Ω termination. The back termination is strict practice, enforcing a true 50Ω signal path. You can eliminate the unit's 6-dB attenuation if it presents problems with only minor signal degradation in the 25-MHz measurement passband. The termination at the oscilloscope end is not negotiable. Figure B shows a typical observed load transient with no back termination but 50Ω at the oscilloscope. The presentation is clean and well-defined. Figure C removes the cable's 50Ω termination, causing a distorted leading edge, ill-defined peaking, and pronounced postevent ringing. Even at relatively modest frequencies, the cable displays unterminated-transmission-line characteristics, resulting in signal distortion.

In theory, a 1× scope probe using a probe-tip coaxial connection could replace the described circuit, but such probes usually have bandwidth limitations of 10 to 20 MHz. Conversely, a 10× probe is wideband, but the oscilloscope's vertical sensitivity must accommodate the introduced attenuation.

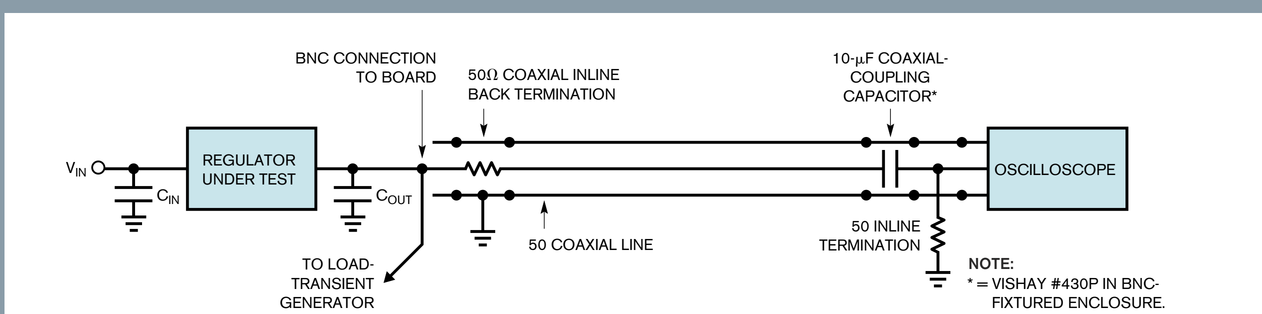

Figure A A coaxial-load-transient voltage-measurement path promotes observed signal fidelity. You can remove the 50Ω back termination with minimal impact on the 25-MHz signal path's integrity.

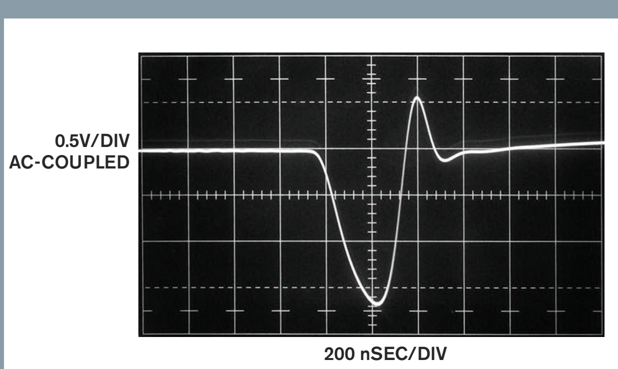

Figure B Observing a typical high-speed transient through Figure A's measurement path presents a clean and well-defined signal.

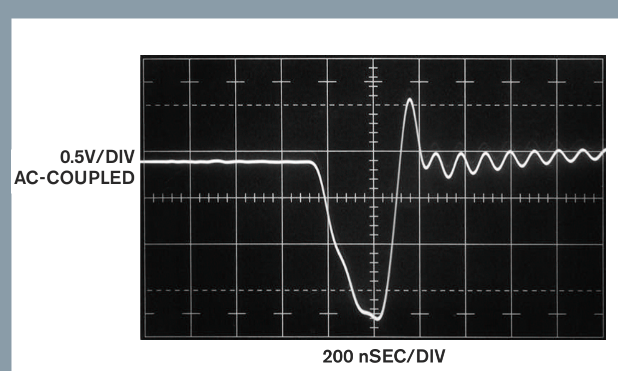

Figure C Measuring Figure B's transient without the 50Ω oscilloscope's termination shows results in waveform distortion and postevent ringing.

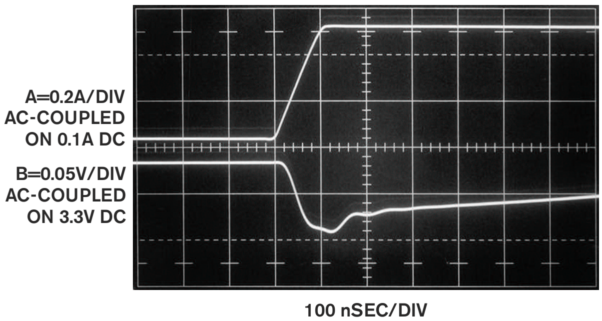

Figure 24 The regulator's output response (Trace B) to a 100-nsec rise-time current step (Trace A) for C_{OUT} is 10 μF. The response decay peaks at 75 mV.

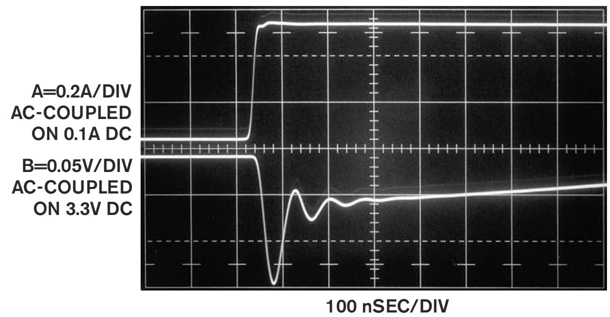

Figure 25 Faster rise-time current step (Trace A) increases response-decay peak (Trace B) to 140 mV, indicating increased regulation loss versus frequency.

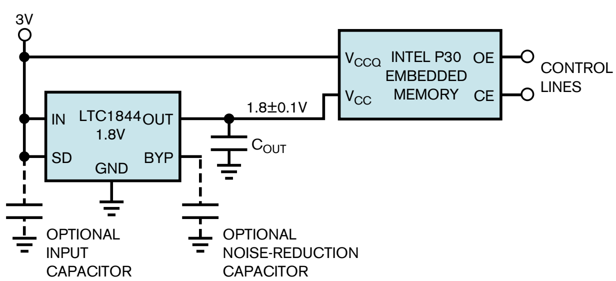

Figure 26 The P30 embedded-memory voltage regulator must maintain a ±0.1V error band. Control-line movement causes 50-mA load steps, necessitating attention to C_{OUT} selection.

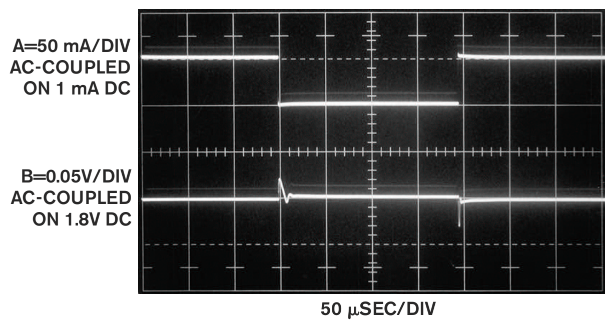

Figure 27 A 50-mA load step (Trace A) results in 30-mV regulator-response peaks, two times better than error-budget requirements. C_{OUT} is a low-loss, 1-μF capacitor.

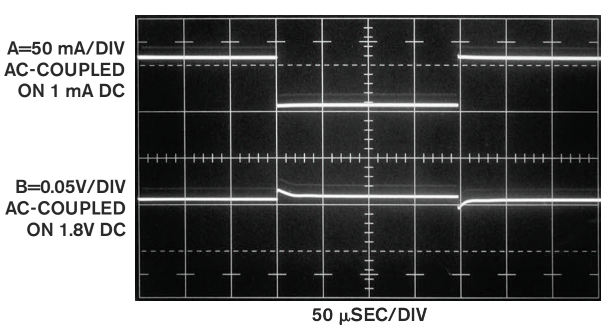

Figure 28 Increasing the value of C_{OUT} to 10 μF decreases regulator-output peaks to 12 mV, almost six times better than required.

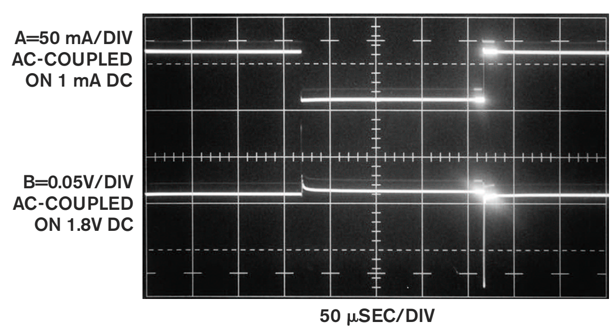

Figure 29 A low-grade, 10-μF C_{OUT} causes 100-mV regulator-output peaks (Trace B), violating the P30 regulator's memory limits. The scope photo intensifies the trace's latter portion for clarity.

ing to a 0.5A, 100-nsec rise-time step on a 0.1A dc load (Trace A). Response decay (Trace B) peaks at 75 mV with some following aberrations. Decreasing Trace A's load-step rise time (**Figure 25**) almost doubles Trace B's response error, with attendant enlarged following aberrations. This scenario indicates increased regulator error at higher frequency.

All regulators present increasing error with frequency—some more than others. A slow load transient can unfairly make a poor

TABLE 1 INTEL P30 EMBEDDED-MEMORY VOLTAGE-REGULATOR ERROR BUDGET

Parameter	Limits
Intel-specified supply limits	1.8V±0.1V
LTC1844-regulator initial accuracy	±1.75% (±31.5 mV)
Dynamic-error allowance	±68.5 mV

A TRIMLESS, CLOSED-LOOP-TRANSIENT-LOAD TESTER

Eliminating a FET-based design's ac trims is an attractive option; however, eliminating the dc trim is also important. The circuit in Figure A trades circuit complexity to achieve this goal. The circuit includes two amplifiers, A_1 and A_2.

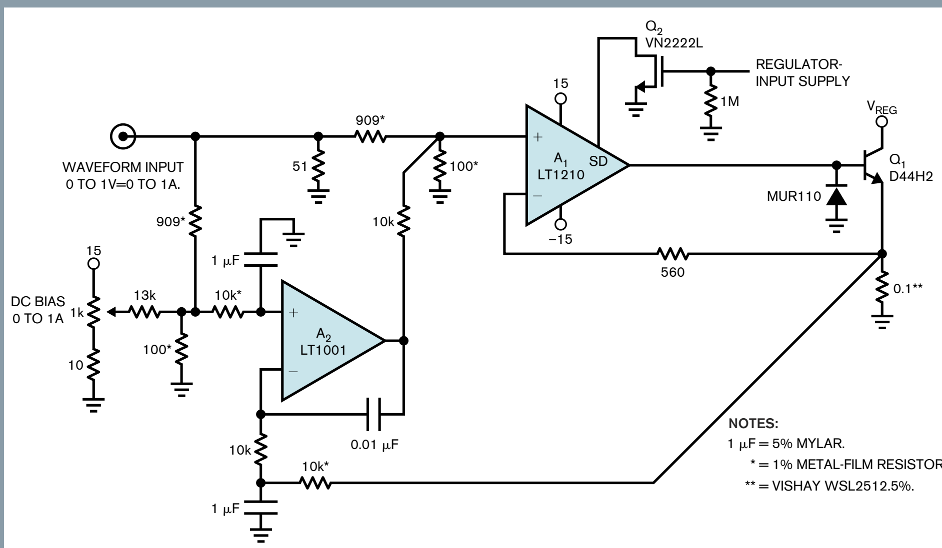

Figure A A_2's feedback controls A_1's dc errors, eliminating trim. Filtering restricts A_2's response to dc and low frequency.

A_2 replaces the dc trim by measuring the circuit's dc input and comparing it with Q_1's emitter dc level and controlling A_1's positive input to stabilize the circuit. The system filters high-frequency signals at A_2's inputs, and these signals do not corrupt A_2's stabilizing action. A useful way to consider circuit operation is that A_2 balances its inputs and, hence, the circuit's input and output, regardless of A_1's dc-input errors. You can set the dc-current bias to any desired point by directing a variable reference source to A_2's positive input. The arrangement of the network's resistors yields a minimum load current of 10 mA, avoiding loop disruption for currents near zero.

regulator look good. Transient-load testing that does not indicate some response outside regulator bandwidth is suspect.

The Intel (www.intel.com) embedded-memory voltage regulator furnishes a good, practical example of the importance of voltage-regulator-load-step performance. The memory requires a 1.8V supply, typically regulated down from 3V. Although current requirements are relatively modest, supply tolerances are tight. **Table 1** shows only 0.1V allowable excursion from 1.8V, including all dc and dynamic errors. The LTC1844-1.8 regulator has a 1.75% initial tolerance at 31.5 mV, leaving only a 68.5-mV dynamic-error allowance. **Figure 26** shows the test circuit. Memory-control-line movement causes 50-mA load transients, necessitating attention to capacitor selection. (The LTC1844-1.8's noise-bypass pin works with an optional external capacitor to achieve low output noise. This application, however, does not require it, and remains unconnected.) If the regulator is close to the power source, C_{IN} is optional. If not, use a high-grade, 1-µF capacitor for C_{IN}. C_{OUT} is a low-loss, 1-µF type. In all other respects, the circuit appears deceptively routine. A load-transient generator provides **Figure 27**'s output-load test step (Trace A). This test uses **Figure 8**'s circuit and changes Q_1's emitter-current shunt to 1Ω. Trace B's regulator response shows just 30-mV peaks, more than two times better than necessary. Increasing C_{OUT} to 10 µF (**Figure 28**) reduces peak output error to 12 mV, almost six times better than specification. However, a low-grade 10-µF—or 1-µF, for that matter—capacitor produces **Figure 29**'s unwelcome surprise. Severe peaking error on both edges occurs with 100 mV observable on the negative-going edge. (The photograph shows an intensified version of Trace B's latter portion to aid clarity.) This figure is well outside the error budget and would cause unreliable memory operation (**references 4**, **5**, and **6**). EDN

MORE AT EDN.COM

Go to www.edn.com/ms4200 and click on Feedback Loop to post a comment on this article.

REFERENCES

1. Bonte, Tony, "Capacitor Parasitic Effects on Load Transient Response," Linear Technology Corp, Application Note 104, Appendix A.
2. O'Neill, Dennis, "Output Capacitors and Loop Stability," Linear Technology Corp, Application Note 104, Appendix B.
3. LT1584/LT1585/LT1587 Fast Response Regulators Data Sheet, Linear Technology Corp.
4. LT1963A Regulator Data Sheet, Linear Technology Corp.
5. Williams, Jim, "Minimizing Switching Residue in Linear Regulator Outputs," Linear Technology Corp, Application Note 101, July 2005.
6. Migration Guide for Intel Strata Embedded Memory (J3) to Intel StrataFlash Embedded Memory (P30), Application Note 812, Intel Corp, April 2006, http://download.intel.com/design/flcomp/applnots/30666703.pdf.

AUTHOR BIOGRAPHY

Long-time EDN contributor Jim Williams, staff scientist at Linear Technology Corp (Milpitas, CA), has more than 20 years' experience in analog-circuit and instrumentation design.

September 2006

Agilent
Basic Instruments Catalog

Troubleshoot the toughest problems faster

U1251A and U1252A handheld DMMs

- More tools in a single package for solving more problems quickly
- Built-in data logger to catch intermittent failures
- The performance you need to investigate marginal failures and verify close tolerances

	U1251A	U1252A
Resolution	4½ digits	4½ digits
Basic DCV accuracy	0.03%	0.025%
Displays (50,000 count)	Dual	Dual
True RMS	AC	AC & DC
Math functions	6	6
Data logger	100 points (manual) 200 points (interval) Unlimited (PC)	100 points (manual) 200 points (interval) Unlimited (PC)
Frequency counter	N/A	0.5 Hz to 20 MHz
Programmable square wave generator	N/A	0.5 Hz to 4,800 Hz
Operating temperature	-20 °C to +55 °C	-20 °C to +55 °C
Safety compliance	Cat III 1000 V	Cat III 1000 V
Battery	9 V Alkaline	9 V NiMH rechargeable, with built-in charger

With more measurement functions and higher performance than comparable handheld DMMs, you'll solve more problems in less time.

For more information, visit
www.agilent.com/find/quick

Agilent Technologies

www.agilent.com/find/quick

Get a Quick Quote in 2 minutes!

34405A 5½ digit bench DMM

- 16 measurement functions, including temperature and capacitance
- 70 readings/sec @ 4½ digits
- USB 2.0 standard
- Dual display
- Comes with IntuiLink datalogging software

Price: From $745

34410A enhanced-performance DMM, 6½ digit

- 10,000 readings/sec @ 5½ digits continuous
- LAN, USB, and GPIB standard
- Data logger
- 50,000 reading NV memory
- Capacitance and temperature measurements

Price: From $1,595

34411A high-performance DMM, 6½ digit

All the features of the 34410A plus:

- 50,000 readings/sec @ 4½ digits continuous to the PC
- 1 M reading memory
- Pre/post and level triggering

Price: From $1,995

34401A DMM, 6½ digit

- 12 measurement functions, plus limit testing and min/max/average
- 0.0035% 1 year dcV accuracy
- 0.06% 1 year acV accuracy
- 1,000 readings/sec in ASCII format across the GPIB bus @ 4½ digits continuous
- True RMS AC volts and current

Price: From $1,172

Free

Reduce system cost and complexity with instruments that do more for less.

New 2006 System and Bench Instruments Catalog

Profiles more than 180 instruments, from power products to data acquisition systems

To get your copy, visit www.agilent.com/find/quick

U1251A handheld DMM 4½ digit

New

- 4½ digit resolution with 50,000-count dual display
- Basic DC accuracy - 0.03%
- True RMS - AC
- Capacitance range 10 nF to 100 mF, resistance 50 MΩ, current 10 A
- Six math functions plus temperature
- Built-in data logger to catch intermittent failures

Price: $399

U1252A handheld DMM 4½ digit

New

The U1252A expands on the U1251A with key performance enhancements and capabilities:

- Basic DC accuracy - 0.025%
- True RMS - AC and DC
- Built-in frequency counter and programmable square wave generator
- Rechargeable battery with built-in charger

Price: $439

3000 Series oscilloscopes

- 60 MHz to 200 MHz bandwidth
- Color displays, standard
- 2 channels
- Up to 1 GSa/s sampling rate
- 4 kpts per channel memory
- GPIB, RS-232 and USB connectivity option (USB connectivity available with optional Scope Connect software)

Price: From $995

Meet today's challenges and prepare for tomorrow's with more capability and more performance.

Agilent Direct

www.agilent.com/find/quick

Get a Quick Quote in 2 minutes!

6000 Series oscilloscopes

- Integrated insights for analog, digital, and serial data
- Fully synchronized measurements of the analog, digital and serial bus data found in most of today's embedded systems
- Unique mixed-signal measurement capability combined with responsive MegaZoom III deep memory
- Models with 2 or 4 scope channels or the unique combination of 2+16 or 4+16 scope and logic channels, from 100 MHz to 1 GHz

Price: From $4,595

Features

Bandwidth	100 MHz, 300 MHz, 500 MHz, and 1 GHz
Channels	2 or 4 scope channels on DSOs 2 or 4 scope + 16 time-correlated logic channels on MSOs
Sample rate	2 GSa/s to 4 GSa/s
Memory	1 Mpts standard, up to 8 Mpts optional
Display resolution	XGA (1024x768) color, 256 levels of intensity
Waveform update rate	100,000 waveforms/sec
Connectivity	Standard: LAN, USB, GPIB, and XGA video
Serial bus trigger	Standard: I²C, SPI, USB, CAN, and LIN trigger

New 6000 Series enhancements

Extend the power and versatility of your Agilent 6000 Series scope with these enhancements:

I²C and SPI serial decode option	N5423A
CAN/LIN automotive triggering and decode option	N5424A
89601A vector signal analysis software	89601A
PGA dynamic probe for mixed signal scopes	N5406A

8 MB Memory Free

Through February 15, 2007

With the purchase of an Agilent 6000 Series 4-channel / 4+16-channel oscilloscope and an analysis package, Agilent will provide the maximum memory – 8 Mpts – ***free of charge.***

Qualifying application solutions:

- N5423A (Opt LSS) I²C/SPI triggering and decode package
- N5424A (Opt AMS) CAN/LIN triggering and decode package
- N5406A (Opt 001) FPGA dynamic probe for 6000 Series oscilloscopes

For more information visit
www.agilent.com/find/quick

33220A function/arbitrary waveform generator, 20 MHz

- Sine, square, ramp, triangle, pulse, noise and DC waveforms
- 14-bit, 50 MSa/s, 64 K-point arbitrary waveforms
- AM, FM, PM, FSK, and PWM modulation types
- Linear and logarithmic sweeps and burst operation modes
- USB, GPIB, and LAN interfaces included
- LXI class-C compliant

Price: From $1,853

33250A function/arbitrary waveform generator, 80 MHz

- Sine, square, ramp, triangle, pulse, noise and DC waveforms
- 12-bit, 200 MSa/s, 64 K-point arbitrary waveforms
- AM, FM, and FSK modulation types
- Linear and logarithmic sweeps and burst operation modes
- GPIB and RS-232 interfaces included

Price: From $4,553

82357A USB/GPIB interface

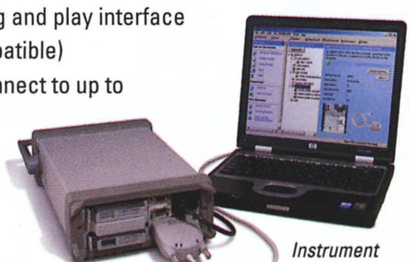

Instrument and PC not included

- Easy connection – plug and play interface
- USB 1.1 (USB 2.0 compatible)
- IEEE 488 interface (connect to up to 14 instruments)
- Includes industry standard libraries – now compatible with NI-488.2
- Convenient built-in 2.5 m USB cable allows direct connect from PC to instrument

Price: From $499

82350B PCI high-performance GPIB interface card

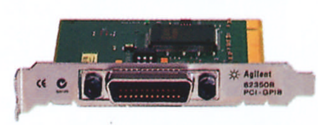

- Built-in buffering for speeds up to 900 KB/s
- Dual processor support on Windows® 2000/XP
- IEEE 488 interface (connect to up to 14 GPIB instruments)
- Includes industry standard libraries – now compatible with NI-488.2
- Easy connection – plug and play interface

Price: From $499

N1996A Agilent CSA spectrum analyzers

Breaking the low-cost barrier with high-quality spectrum analysis

- Today's unbeatable combination of portability, low price, and substantial performance
- Two models: 100 kHz to 3 or 6 GHz
- Best-in-class usable dynamic range

Features

Frequency range	100 kHz to 3 or 6 GHz
Resolution bandwidth	10 Hz to 5 MHz
Third order intercept	+18 dBm
Displayed average noise level valid at 1 GHz normalized	-138 dBm/Hz -156 dBm/Hz with preamp
Overall amplitude accuracy	±0.5 dB
Display	21 cm XGA color
Connectivity	USB and LAN standard
Key options	Built-in bridge and tracking generator, cable fault / return loss measurements, preamp, spectrogram

Price: From $9,800

E4418B/19B EPM Series power meters

- Average power measurements from 9 kHz to 110 GHz, -70 to +44 dBm (sensor dependent)
- Fast measurement speed (up to 200 readings per second, with E-Series sensors)

Cables not included

- GPIB and RS-232 connectors for remote control of all functions
- Internal rechargeable battery option
- Code compatibility with 438A, 437B, and 436A power meters

Price: From $3,464

E4416A/17A EPM-P Series power meters

- Peak, average, and peak-to-average ratio power measurements
- 5 MHz video bandwidth (maximum) with 20 Msample-per-second continuous sampling

Cables not included

- Free 'analyzer software' for complete pulse and statistical analysis
- Eight pre-defined wireless configurations (GSM, EDGE, NADC, iDEN, Bluetooth,™ IS-95 CDMA, W-CDMA, and cdma2000
- Time-gated and free-run measurement modes

Price: From $4,497

Multiport microwave switches – 19 models available

- Magnetic latching
- Repeatability for > 5 million cycles
- Typical isolation > 90 dB to 26.5 GHz
- Position indicator, internal control logic

Price: From $1,023

Model	*Frequency range*	*Number of ports*	*Type*	*Connector*
87104A	DC - 4 GHz	SP4T	Terminated	SMA (f)
87106A	DC - 4 GHz	SP6T	Terminated	SMA (f)
87104B	DC - 20 GHz	SP4T	Terminated	SMA (f)
87106B	DC - 20 GHz	SP6T	Terminated	SMA (f)
8766K	DC - 26.5 GHz	SP3T	Unterminated	3.5 mm (f)

SPDT microwave switches – 15 models available

New 85331B / 85332B solid state switches

- Maximize your operating frequency range from 45 MHz to 50 GHz
- Port-to-port isolation > 90 dB
- SPDT and SP4T models available

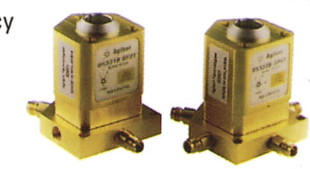

Electromagnetic switches

- Magnetic latching
- Repeatability for > 5 million cycles
- Break-before-make connections
- Isolation > 100 dB to 4 GHz

Price: From $190

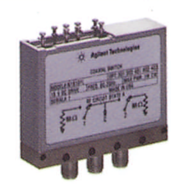

Model	*Frequency range*	*Type*	*Connector*
N9397A	300 kHz to 8 GHz	Absorptive	SMA (f)
N9397C	300 kHz to 18 GHz	Absorptive	SMA (f)
N1810TL	DC to 26.5 GHz	Terminated	SMA (f)
N1810UL	DC to 26.5 GHz	Unterminated	SMA (f)
8765A	DC to 4 GHz	Unterminated	SMA (f)
8765B	DC to 20 GHz	Unterminated	SMA (f)

Free advice for selecting the right solid state switch

- See how solid state switches compare to electromechanical alternatives
- Explore the relative advantages of FET, PIN, and hybrid switches

www.agilent.com/find/quick

Agilent Authorized Distributors

The confidence of Agilent quality and support with the convenience of local, personalized service

- **Work with** your favorite distributor to make all of your everyday equipment purchases
- **Enjoy** the benefits of in-stock delivery, flexible financing, and personalized local attention
- **Count on** in-depth expertise and the information you need to make smart decisions
- **Rest easy,** knowing you're always covered by Agilent warranty, maintenance, and calibration services

Agilent Authorized Distributors meet rigorous qualifications for business ethics, financial stability, customer-driven culture, and technical expertise. Many offer focused expertise in specific industries and technical sectors, and all are backed with training and support from Agilent factory experts.

TESTEQUITY
TESSCO TECHNOLOGIES
STANLEY Supply & Services
TRANSCAT CALIBRATION SERVICES • TEST & MEASUREMENT INSTRUMENTS
NEWARK inone
FUTURE ELECTRONICS
Davis inotek INSTRUMENTS & CALIBRATION
ALLIED ELECTRONICS
CONTINENTAL RESOURCES, INC.
TUCKER ELECTRONICS
TECHNI-TOOL Tomorrow's Products Shipped Today
METRICTEST
testmart
HENSLEY TECHNOLOGIES
Instrument Engineers
electro sonic

www.agilent.com/find/distributors

- To find your local distributor of Agilent products
- And get a quick quote from your distributor

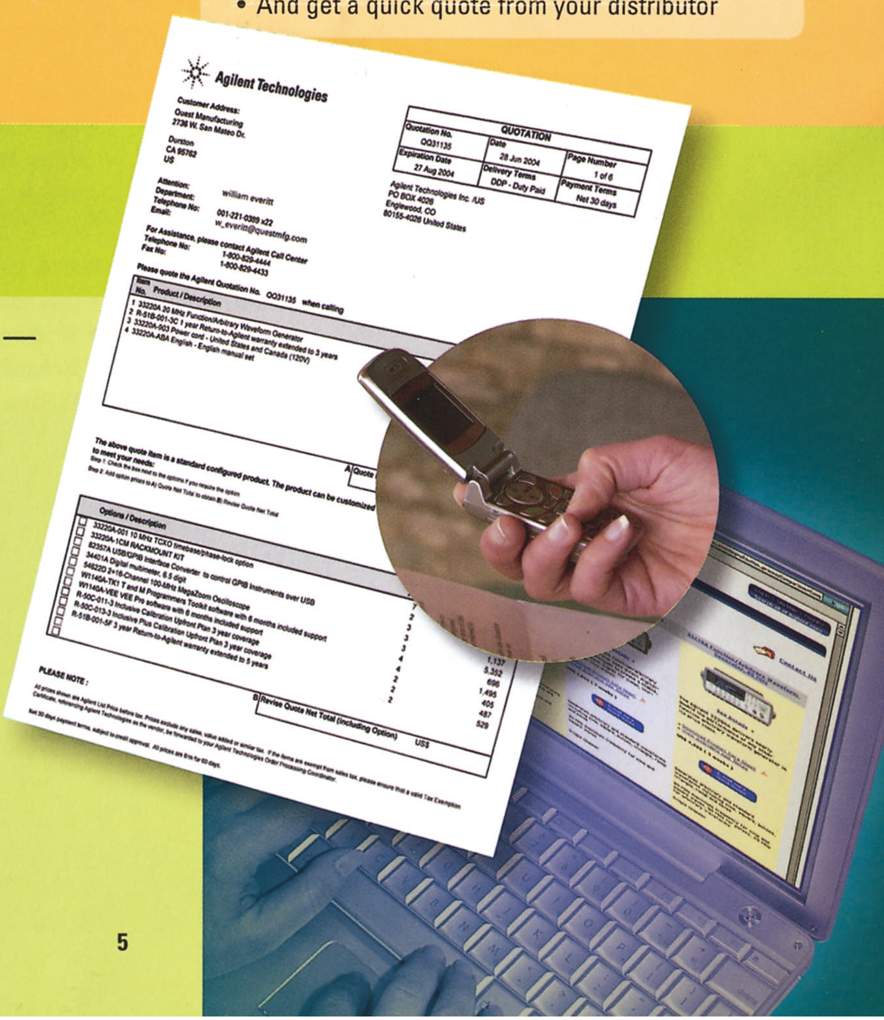

Agilent Direct

Find the right solution for your application — in minutes, not days

- **Compare** up to four products side-by-side.
- **Get** immediate pricing information with Quick Quote.
- **Call** Agilent Direct to discuss your application needs via the Web, e-mail, or phone.

get quick quote

Get a quick quote in 2 minutes

Contact us at **1-800-829-4444**
or **www.agilent.com/find/quick**

www.agilent.com/find/quick

Get a Quick Quote in 2 minutes!

81101A / 04A pulse generator 50 / 80 MHz

- Single-/dual-channel pulse pattern generator
- 50/80 MHz, 10 V from 50 Ω into 50 Ω
- Triggerable / PLL operation selectable
- Pulse / burst / pattern modes selectable
- Variable transition times
- 81104A: requires 1 or 2, 81105A output modules

Price: From $6,436

81110A* pulse pattern generator, 165 / 330 MHz

- Single-/dual-channel pulse pattern generator
- ***Either,*** 165 MHz, 10 V, 50 Ω into 50 Ω (81111A) ***or,***
- 330 MHz, 3.8 V, 50 Ω into 50 Ω (81112A)
- Triggerable / PLL operation selectable
- Pulse / burst / pattern modes selectable
- Self-cal for improved timing accuracy

Price: From $13,511

** Each 81110A 165 MHz single / dual-channel pulse pattern generator requires one/two of the 81111A output channels (165 MHz, 10 V, 50 Ω into 50 Ω).*

Each 81110A 330 MHz single / dual-channel pulse pattern generator requires one/two of the 81112A output channels (330 MHz, 3.8 V, 50 Ω into 50 Ω).

Trade Up

Pulse pattern generators:
Trade up to higher performance

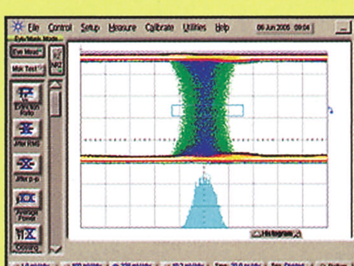

- Get the newest technology for characterizing today's designs
- Recover up to 25% of your original purchase price (varies by model)
- From low-cost, sub-100 MHz to 13 GHz high-speed solutions – all with high-quality waveforms and flexible data stream settings
- Offer valid through December 31, 2006

Learn more at www.agilent.com/find/quick

N5700 Series DC power supplies, 750 W and 1500 W

- Up to 600 V and up to 180 A
- Small high density 1U package
- Full protection from over-voltage and over-current
- 85-265 V_{ac} universal AC input
- LAN, USB, and GPIB interfaces standard
- LXI class-C compliant

Price: From $2,386

E3634A 200 W power supply, 25 V, 7 A or 50 V, 4 A

- 175 W/200 W single output dual range
- Low noise and excellent regulation
- Front and rear output terminals
- GPIB and RS-232 standard
- Remote sense, OVP and OCP
- SCPI compatible

Price: From $1,304

E3640A 30 W power supply, 8 V, 3 A or 20 V, 1.5 A

- 30 W single output, dual range
- Low noise and excellent regulation
- Front and rear output terminals
- GPIB and RS-232 standard
- Over-voltage protection and remote sensing
- SCPI compatible

Price: From $705

E3631A 80 W triple output power supply, 6 V, 5 A & ±25 V, 1A

- 80 W triple output
- Low noise and excellent regulation
- GPIB and RS-232 standard
- Isolated 6 V output
- SCPI compatible

Price: From $1,234

E3632A 120 W power supply, 15 V, 7 A or 30 V, 4 A

- 105 W/120 W single output dual range
- Low noise and excellent regulation
- GPIB and RS-232 standard
- Remote sense, OVP and OCP
- SCPI compatible

Price: From $1,079

Do you have other power supply needs?

Agilent has over 200 other power supplies to choose from.

www.agilent.com/find/quick

www.agilent.com/find/quick

Get a Quick Quote in 2 minutes!

34970A data acquisition switch unit, 6½ digit

- 3-slot mainframe with built-in GPIB and RS-232 interfaces
- Built-in signal conditioning measures temperature, AC/DC volts, current; resistance; frequency and period
- 50 k readings of non-volatile memory
- HI/LO alarm limits on each channel, plus 4 TTL alarm outputs
- Agilent BenchLink Data Logger software included

Price: From $1,521

34970A plug-in modules

Model	Description	Speed (ch/sec)	Max voltage	Starting at
34901A	20 ch mux (2/4-wire)	60	300 V	$469
34902A	16 ch mux (2/4-wire)	250	300 V	$553
34903A	20 ch actuator/GP switch	120	300 V	$404
34904A	4x8 matrix	120	300 V	$469
34905A	2 GHz dual 1:4 RF mux, 50 Ω	60	42 V	$647
34906A	2 GHz dual 1:4 RF mux, 75 Ω	60	42 V	$647
34907A	Two 8-bit digital I/O ports 26-bit, 100 kHz event counter Two 16-bit analog outputs	N/A N/A N/A	42 V 42 V ±12 V	$351
34908A	40 ch single-ended mux	60	300 V	$469

34980A multifunction switch/measure unit

- 8-slot mainframe with built-in LAN, USB 2.0 and GPIB interfaces
- Extended functionality offers temperature, AC/DC volts, current, resistance, digital I/O, D/A converters, frequency and period
- LXI class-C compliant, high-performance but low-cost alternative to PXI and VXI

Priced from $2,315

34980A plug-in modules – 19 modules available

	Key specifications	Starting at
34921A-25A multiplexers	Up to 300 V/1 A	$790
34931A-33A matrix switches	Up to 128 crspts	$889
34937A/38A GP switches	1 A and 5 A	$889
34941A/42A RF switches	50 or 75 Ω	$1,571
34945A µW switch/atten driver	Drive 64 coils	$691
34946A/47A µW switches	SPDT sw to 20 GHz	$691
34950A-59A system control	D/A, DIO, counter	$1,177

For a complete list go to: **www.agilent.com/find/quick**

53131A / 32A universal frequency counters, 10 / 12 digit/sec

- Measurements include frequency, frequency ratio, time interval, rise/fall time, phase, and much more
- Automated limit tests and one button measurement setups for fast, easy operation
- Two 225 MHz input channels, plus optional third channel from 3, 5 or 12.4 GHz
- Built-in statistics feature lets you simultaneously measure average, min/max, and standard deviation

Price: From $2,113

53181A RF frequency counter, 10 digit/sec

- Measurements include frequency, frequency ratio (with optional channel 2), period and peak voltage
- Automated limit tests and one-button measurement setups for fast, easy operation
- Built-in statistics feature lets you simultaneously measure average, min/max, and standard deviation
- One 225 MHz input channel, plus optional 2nd channel from 3, 5 or 12.4 GHz

Price: From $1,841

Free

Agilent Basic Instruments CD (4th Edition)

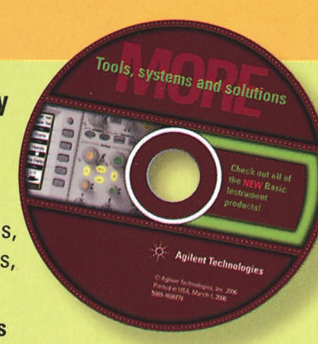

In-depth information ready to print or view:

- **Technical data sheets, video demonstrations, and application notes** on switches, power supplies, oscilloscopes, and more
- **Test system developer guides** to help you integrate systems using open connectivity standards
- **New catalogs:** for RF & Microwave Test Accessories and System and Bench Instruments
- Information on many **recently introduced products,** including the 34410A/11A digital multimeters, 3000 and 6000 Series oscilloscopes, and more

Order your FREE COPY today at
www.agilent.com/find/quick

Agilent Authorized Distributors

Agilent Authorized Distributors meet the highest standards for integrity, expertise, and customer service. Turn to page 5 for more information on the advantages of working with a distribution partner in your local area.

TestEquity
Techni-Tool — Tomorrow's Products Shipped Today
Future Electronics
Transcat — Calibration Services • Test & Measurement Instruments
Tucker Electronics
Tessco Technologies
Stanley Supply & Services
MetricTest
Davis Inotek Instruments & Calibration
Newark inone
Allied Electronics
Instrument Engineers
testmart
electrosonic
Hensley Technologies
Continental Resources, Inc.

Agilent Authorized Distributors deliver responsive, local support backed by Agilent's full technical resources.

Agilent Technologies

Authorized Distributor

New handheld DMMs with more ways to solve more problems.

See page 2.

Agilent Direct

Learn more at:

www.agilent.com/find/quick
1-800-829-4444

Agilent Technologies

Department DMEP3464606
P.O. Box 3828
Englewood, CO 80155-3828

Presorted
Standard
U.S. Postage
PAID
Minneapolis, MN
Permit 7499

CHANGE SERVICE REQUESTED

Technical data and pricing subject to change.

Printed in U.S.A. September 1, 2006

5989-5239ENUS

NUS

Low Power, Rail-to-Rail Input/Output, Single Supply Op Amps

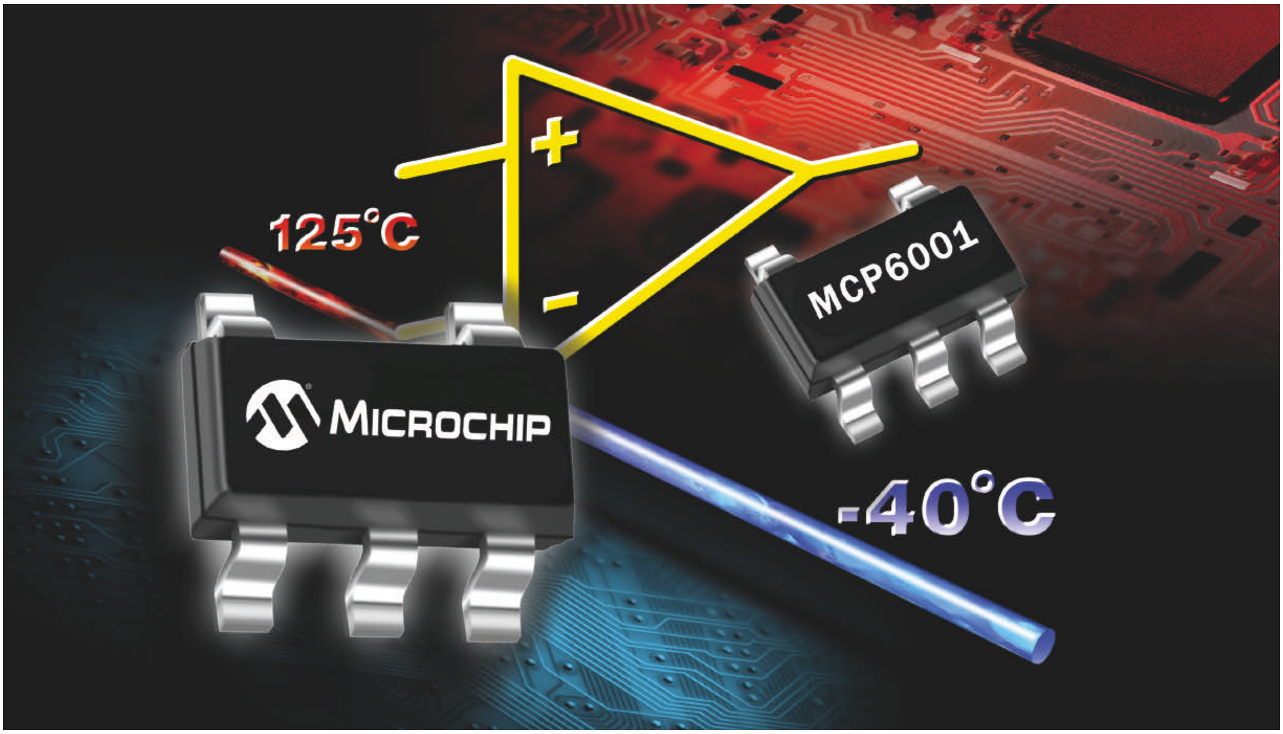

Select Standard Op Amps

Part #	GBWP	Iq Typical (µA)	Vos Max (mV)	Input Voltage Noise Density @ 1 kHz (nV/√Hz)	Operating Voltage (V)
MCP6041/2/3/4	14 kHz	0.6	3.0	170	1.4 – 5.5
MCP6141/2/3/4	100 kHz	0.6	3.0	170	1.4 – 5.5
MCP6231/2/4	300 kHz	20	5.0	52	1.8 – 5.5
MCP6241/2/4	550 kHz	50	5.0	45	1.8 – 5.5
MCP6001/2/4	1 MHz	140	4.5	28	1.8 – 5.5
MCP6271/2/3/4/5	2 MHz	170	3.0	20	2.0 – 5.5
MCP6281/2/3/4/5	5 MHz	445	3.0	16	2.2 – 5.5
MCP6291/2/3/4/5	10 MHz	1100	3.0	8.7*	2.4 – 5.5
MCP6021/2/3/4	10 MHz	1000	0.5	8.7*	2.5 – 5.5

* Value is typical at 10 kHz

- Select devices available in PDIP, SOIC, MSOP, TSSOP, SOT-23, and SC-70
- Select devices offer a Chip Select pin for additional power savings
- The MCP62X5 offers dual amplifiers with a Chip Select pin in an 8-pin package
- **FREE!** Download the FilterLab® Active Filter Design Tool at www.microchip.com
- **FREE!** Order device samples at www.microchip.com

Purchase your op amps at...

microchip DIRECT

www.microchipdirect.com

Now Pb-free!

RoHS Compliant

MICROCHIP

www.microchip.com/OpAmps

Microcontrollers • Digital Signal Controllers • Analog • Serial EEPROMs

BY CECIL DEISCH • TELLABS OPERATIONS INC

Modeling skin effect in Spice

SKIN EFFECT CAUSES INCREASED LOSSES AS FREQUENCY INCREASES. IT ALSO CAUSES CHANGES IN SIGNAL VELOCITY, DEGRADING SIGNAL FIDELITY, ESPECIALLY THE EYE OPENING OF HIGH-SPEED DATA SIGNALS ON LONG SIGNAL PATHS ON PC BOARDS AND BACKPLANES.

Skin effect involves the interplay between flux linkages and currents in and on conductors. At dc and low frequencies, current density is essentially uniform throughout the cross section of a conductor. The fact that current flows on the interior of the conductor implies that a magnetic flux also exists within the conductor. Thus, at low frequencies, the inductance of a conductor is higher because magnetic flux is both inside and outside the conductor.

However, as frequency increases, current density within the conductor varies in such a way that it tends to exclude magnetic flux inside the conductor. This situation results in an apparent increase in resistance of the conductor because more of the current is concentrated near the surface and edges of the conductor, and it also causes the effective inductance of the conductor to decrease as frequency increases. These two effects become especially important when modeling the performance of high-speed data signals. A Spice model can easily accommodate the frequency variation of the skin resistance, which usually increases as the square root of frequency, such that simulation losses accurately agree with measured losses. However, it's also important to model the change in conductor inductance with frequency because this change affects the signal velocity and, hence, time of flight. This effect is especially important when modeling gigahertz data signals as they travel across the backplane.

SIGNAL DEGRADATION

Digital data signals of several gigahertz greatly degrade after they have traveled, say, 30 in. from one end of a backplane to the other end. The travel causes the "eye" opening of the signal to close severely and become corrupt, making decoding error-prone. The fact that the higher frequency signal components have greater attenuation than the low-frequency components causes much of this eye closure. Another factor is the time of arrival of the various frequency components. Because the conductor's inductance is slightly less at higher frequencies, the velocity of propagation is slightly greater, which implies that the higher frequency signal components arrive at the far end slightly before the bulk of the lower frequency components. This varying time-of-arrival, or dispersion, factor affects the zero-crossing times and further degrades the eye opening. Both the skin effect's attenuation and the time of arrival of the high-frequency components can dramatically corrupt the signal and lead to a nearly closed eye pattern.

Because skin effect and the change in inductance are inextricably linked, a Spice model also links these two effects. The simplest model for an incremental length of transmission line is the basic series L_1 and a lossless shunt, C, with some resistor R_1 in series with the inductance (**Figure 1**). R_1 can be either fixed or frequency-dependent to account for skin-effect losses. However, this simple model has no provision to change signal velocity. In another model, part of the series inductance has a shunt resistance across it (**Figure 2**). It's easy to see that, at low frequencies, the total inductance is simply the sum of L_1 and L_2, and the loss due to shunt resistor R_2 is negligible because its impedance is so much higher than that of L_2. As frequency increases, R_2 comes more into play and causes more losses as the impedance of L_2 increases. Thus, L_2 and R_2 perform the same function that occurs in a conductor—that is, as currents inside the conductor exclude some of the internal magnetic-flux lines, the apparent inductance decreases, and the losses increase. By juggling the values of L_1, L_2, R_1, and R_2, this simple circuit mod-

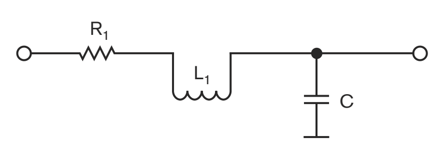

Figure 1 This basic model has no provision to change signal velocity with frequency.

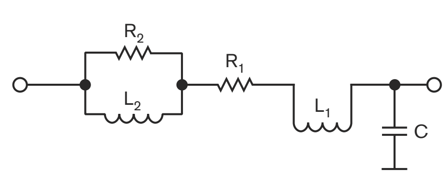

Figure 2 L_2 and R_2 increase losses and velocity as frequency increases.

els the resistance and inductance of an actual conductor over a frequency range greater than 25-to-1.

If accuracy over a wide frequency range is insufficient, you can add L_3 and R_3 (**Figure 3**). The sum of L_1, L_2, and L_3 represents the low-frequency inductance of the conductor, and you can adjust the values of all six of the series components to give an even better model. For most applications, one or two extra series sections give accurate models, but you can add series sections to achieve the required accuracy. Modeling a long signal path that carries high-frequency signals requires many of these incremental sections. A good rule is to account for at least the third harmonic and sometimes as high as the seventh harmonic of the highest data rate. This rule implies that the delay through each incremental section should be no greater than one-fifth the period of the highest harmonic. It's also a good idea not to make all incremental sections of equal length because this step can lead to certain high-frequency anomalies in the simulation waveforms. Vary the length of each incremental section, but make none greater than one-fifth the period.

TABLE 1 APPROXIMATE RESISTANCE OF THE CENTER CONDUCTOR

Frequency (MHz)	Handbook resistance values (ohm/ft)	Model resistance values (ohm/ft)	Model inductance values (nH)
1	0.13	0.15	83+7.7
5	0.24	0.22	83+6.8
10	0.32	0.34	83+5.1
50	0.7	0.67	83+1.7
100	0.96	0.89	83+1.3
500	2.2	2.4	83+0.31
1000	3.1	2.8	83+0.09

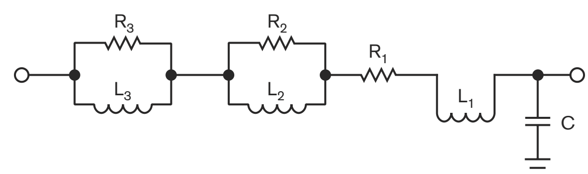

Figure 3 If accuracy over a wide frequency range is insufficient, you can add L_3 and R_3.

A SIMPLE EXAMPLE

Although designers have much interest in how data signals degrade as they travel along a backplane, the mathematics are intractable, so another model uses a simpler coaxial cable (**Figure 4**). The center conductor is #30 AWG, the relative dielectric constant of the interior insulation is 2.46, and the inside diameter of the outer sheath is 37 mils (0.037 in.). This arrangement gives a characteristic impedance of 50Ω. At dc and low frequencies, #30 AWG wire has a resistance of approximately 0.1Ω/ft. **Table 1** shows the approximate resistance of the center conductor (**Reference 1**). A quick calculation shows that the resistance increase is approximately the square root of frequency above 10 MHz.

According to standard tables, the center conductor of this

SIGNAL
INTEGRITY
FROM
THE
OUTSIDE
IN
PANEL-TO-BOARD
PANEL-TO-PANEL
CABLE-TO-BOARD
BOARD-TO-BOARD
samtec
www.samtec.com/si

coaxial cable has an inductance of approximately 83 nH/ft. But this inductance occurs at high frequencies at which virtually all current flows on the surface of the inner conductor and all magnetic flux lines are outside the conductor. However, at low frequencies, the inductance of the center conductor is significantly higher.

For a given total conductor current, flux lines outside the conductor are constant, independent of frequency or current distribution within the conductor. However, at low frequencies, current flows throughout the conductor, and a substantial number of flux lines are inside the conductor. For this model, more than 25% of the flux lines are inside the conductor. Because total inductance relates to the number of flux lines encircling the current, inductance is higher at low frequencies. However, the low-frequency inductance is not 25% higher because these internal flux lines don't enclose the entire conductor current. Low-frequency inductance is typically 10 to 15% higher. These circumstances imply that the inductance of the center conductor cannot be greater than its dc value of about 92 nH and cannot be less than 83 nH/ft, which you can model as 83 nH in series with approximately 9 nH. You can subdivide this 9 nH into several sections with appropriate shunt resistors to give skin-effect resistance that varies close to the square root of frequency.

THE SPICE MODEL

The model in **Figure 5** applies to the earlier-described coaxial cable and provides excellent agreement with the calculated skin resistance over the 1-MHz to 1-GHz frequency range. This model requires just two sections of resistance-shunted inductance and has an "error" of only ±12% from **Reference 1**'s values over this range (**Figure 6**). Although the skin-effect resistance of actual conductors, including backplane conductors, closely approximates the square root of frequency, the relationship isn't exact. This model uses fixed-value resistors and has no need for frequency-dependent resistance. Although this model gives per-foot values, the incremental sections for high-speed signals need to be 0.3 in. or less, which implies that you must divide the values by 40. A 30-in. length would require 100 of these incremental sections.

This small change in effective inductance as frequency increases is not insignificant. Consider a data rate of 2 Gbps over a 30-in. length. Using inductance numbers from **Table 1** indicates that the 1-GHz-frequency components arrive at the far end approximately 125 psec sooner than the 10-MHz components. Because the bit interval is 500 psec, this difference in arrival time amounts to an approximately 45° phase

MORE AT EDN.COM

Go to www.edn.com/ms4204 and click on Feedback Loop to post a comment on this article.

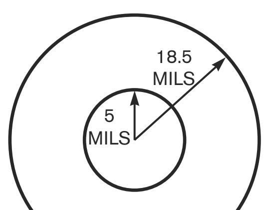

Figure 4 This model uses a simple coaxial cable.

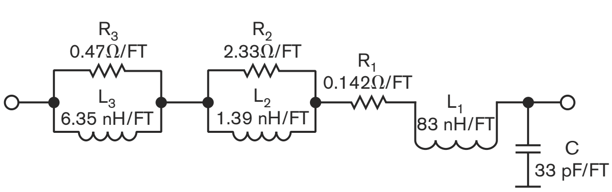

Figure 5 This model applies to the earlier-described coaxial cable and provides excellent agreement with the calculated skin resistance over the 1-MHz to 1-GHz frequency range.

shift, which greatly degrades the eye. **Figure 7** shows an alternative model with a slightly different topology and the same accuracy and error-versus-frequency curve as the previous model.

This article covers only the change in inductance of a conductor that causes high-frequency components of the data signal to arrive sooner than the low-frequency components. This model assumes a constant shunt capacitance, but some published data indicates that the relative dielectric constant of backplane material decreases slightly as frequency increases. This effect would also cause high-frequency components to travel faster, further degrading the eye opening. EDN

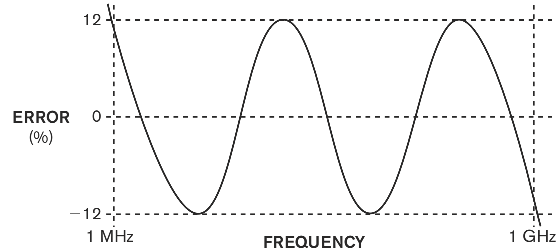

Figure 6 The Spice model in Figure 5 agrees with Table 1 within 12% over a 1000-to-1 frequency range.

REFERENCE

1. *Reference Data for Radio Engineers, Fifth Edition*, Howard W Sams and Co, 1968.

AUTHOR'S BIOGRAPHY

Cecil Deisch is a staff engineer with Tellabs; he was previously at Bell Labs. He holds a bachelor's degree in electrical engineering from the South Dakota School of Mines and Technology (Rapid City) and a master's degree in electrical engineering from New York University (New York). His interests include thermodynamics, economics, and mathematics.

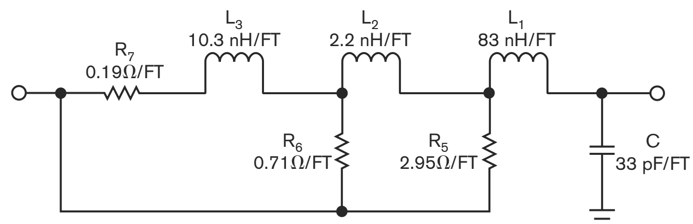

Figure 7 This model has a slightly different topology but the same error-versus-frequency (Figure 6) curve as the model in Figure 5.

Whatever direction you choose, Gore gets you optimal performance without speed limits.

We've got all the preferred high-speed assemblies ready-to-roll. So, fasten your seat belt!

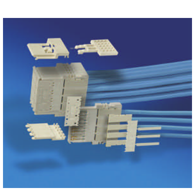

GORE™ EYE-OPENER® ZD Cable Assemblies

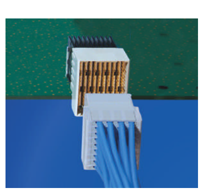

GORE™ EYE-OPENER® AirMax VS® Cable Assemblies

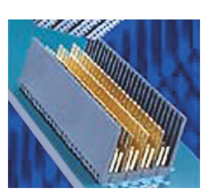

GORE™ EYE-OPENER® GBX Cable Assemblies

GORE™ INFINIBAND™/CX4 Cable Assemblies

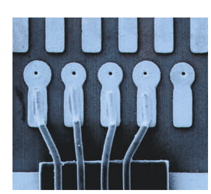

GORE™ Micro-miniature Flat Cable Assemblies

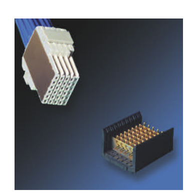

GORE™ EYE-OPENER® VHDM™ Cable Assemblies

W. L. Gore & Associates

North America:
1 (800) 445-4673
International:
1 (302) 292-5100

Europe:
+49/91 44/6010
+44/1382-561511
gore.com

GORE
Creative Technologies Worldwide

ExpressFabric

Ultimate Performance…

VIRTEX™ 5

Achieve highest system speed and better design margin with the world's first 65nm FPGAs.

Virtex™-5 FPGAs feature ExpressFabric™ technology on 65nm triple-oxide process. This new fabric offers the industry's first LUT with six independent inputs for fewer logic levels, *and* advanced diagonal interconnect to enable the shortest, fastest routing. Now you can achieve 30% higher performance, while reducing dynamic power by 35% and area by 45% compared to previous generations.

Design systems faster than ever before

Shipping now, Virtex-5 LX is the first of four platforms optimized for logic, DSP, processing, and serial. The LX platform offers 330,000 logic cells and 1,200 user I/Os, plus hardened 550 MHz IP blocks. Build deeper FIFOs with 36 Kbit block RAMs. Achieve 1.25 Gbps on all I/Os without restrictions, and make reliable memory interfacing easier with enhanced ChipSync™ technology. Solve SI challenges and simplify PCB layout with our sparse chevron packaging. And enable greater DSP precision and dynamic range with 550 MHz, 25x18 MACs.

Visit ***www.xilinx.com/virtex5***, view the TechOnline webcast, and give your next design the ultimate in performance.

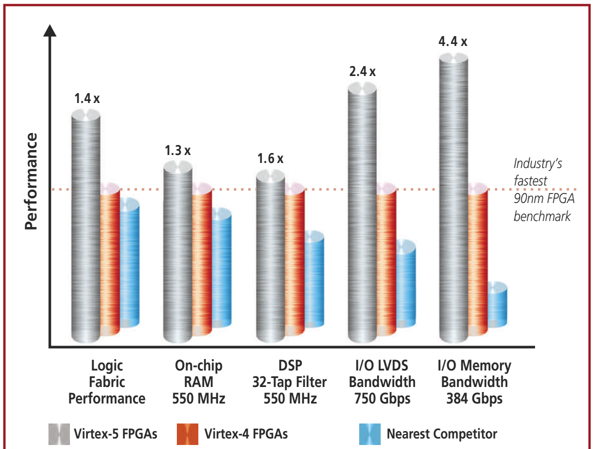

Numbers show comparision with nearest competitor

Based on competitor's published datasheet numbers

XILINX®

The Programmable Logic Company℠

www.xilinx.com/virtex5

The Ultimate System Integration Platform

BY GENE HOWELL • TT ELECTRONICS, IRC BOONE DIVISION

Five questions about resistors

UNDERSTANDING THIS UBIQUITOUS PART CAN HELP YOU AVOID COMMON CIRCUIT PROBLEMS.

The resistor is one of the simpler electronic components. Engineers seldom examine resistor characteristics until there's a resistor-related problem with a circuit design. You can probably resolve 90% of these problems after answering these five key questions:

- How much voltage can I put on the resistor?
- What will be the temperature of the resistor in my circuit?
- How much surge will the resistor withstand?
- What makes a resistor fail?
- How much change in resistance can I expect?

HOW MUCH VOLTAGE?

On the surface, this question appears to be the simplest. If, for example, you have a 0.25W resistor, you look on the appropriate product sheet for the part and find that it lists 250V as the maximum voltage. However, your engineering instincts tell you that, because the part has 10Ω, you better not put 250V onto a 10Ω part that measures 0.250 in. long by 0.090 in. in diameter unless you first get behind a concrete wall.

You can find the correct value of the maximum continuous voltage to put onto the 10Ω part by remembering the following: For capacitors, voltage rules; for inductors, current rules; and, for resistors, power rules. The primary parameter that you do not want to exceed for a resistor is the continuous-power rating, P_{RATED}. Because the rated power is the voltage squared divided by the resistance, the maximum continuous-voltage rating, V_{RATED}, is the square root of the rated power times the resistance: $V_{RATED}=\sqrt{P \times R}$. For the 10Ω, 0.25W example, the rated voltage is 1.58V, far less than the 250V that the spec sheet lists as the maximum voltage. You must make this calculation for each resistor value you use.

You determine when to use the specified maximum voltage or the calculated rated voltage using the critical resistance, $R_{CRITICAL}$. Resistor manufacturers define the critical resistance as the resistance at which the part dissipates rated power at the maximum continuous-voltage rating: $R_{CRITICAL}=(V_{RATED})^2/P_{RATED}$. For the 0.25W part with a maximum voltage of 250V, the critical resistance is 250 kΩ. At resistances higher than the critical resistance, the maximum continuous rated voltage is 250V. At resistances lower than the critical resistance, you must calculate the square root of the power times the resistance to determine the maximum continuous-voltage rating.

If you do not apply the voltage continuously, you can put a higher voltage on a resistor for a short duration. Resistor manufacturers define an STOL (short-time-overload) condition, the degree of which varies depending on the type of resistor. For power wirewound resistors, STOL can be two to 10 times the rated power for 5 or 10 sec. For most film resistors, STOL is 2.5 times the rated voltage, or 6.25 times the rated power, for 5 sec. For high-voltage resistors, STOL is commonly 1.5 times the rated voltage, or 2.25 times the rated power, for 10 sec. Refer to your resistor's product sheet to see how the manufacturer defines STOL and to find the maximum percentage of change in resistance that can occur when you apply the STOL voltage.

You should heed three notes of caution. First, STOL is a nonrepetitive surge or overload condition. Second, putting two to 10 times the rated power on a resistor for more than 5 or 10 sec can cause permanent damage and can melt the solder joints that hold the part in place. Third, there is a maximum allowable STOL voltage: typically, two times the maximum continuous rated voltage. You should refer to the data sheet or call an application engineer if the data sheet lacks this information. In your calculation of the STOL voltage, do not let it exceed the maximum STOL voltage that the data sheet specifies. For example, for a 0.25W, 10-MΩ film resistor, the STOL-voltage calculation is 2.5 times the square root of the power times the resistance, or 1581V. However, the maximum allowable STOL voltage is two times the maximum continuous rated voltage of 250 or 500V.

Before you can establish the maximum voltage that you can apply to a resistor by considering the maximum continuous- and maximum STOL-voltage ratings, you must consider another set of conditions. When a resistor "sees" a STOL condition, the amount of power you apply may be as much as 10 times rated power for 5 sec. If you apply power for only 1 msec or 1 µsec, you should be able to apply even more than 10 times the rated power. Under a surge condition for a given type of resistor, you may be able to apply 100 or even 1000 or more times the rated power. A manufacturer may allow the voltage you apply to exceed the maximum STOL voltage of resistors that can handle high power surges. However, do not make this assumption. See the product data sheet or speak with the application engineer for information on the maximum power and voltage that you can apply for a given resistance value and pulse duration.

TEMPERATURE OF THE RESISTOR?

In the past, most resistors had solder-coated copper leads, so determining the "hot-spot" temperature was relatively straightforward. A standard test procedure allowed a specified lead length between the resistor and the pc board, and a fine-wire thermocouple determined the hot-spot temperature at various power loadings of the resistor. The design permitted no air movement across the part. From this data, you could make a temperature-versus-power-applied plot. For many resistors of 1W or less, the plot was a straight line from ambient temperature with no power to the hot-spot temperature at rated power. You expressed the slope of the line as the temperature rise in degrees Celsius per watt. For power resistors with a hot spot that exceeded 175°C, radiation caused a heat transfer, making the plot deviate somewhat from a straight line at temperatures higher than 175°C. Even so, engineers normally used a typical temperature-rise figure. For example, if a 2W resistor had a temperature rise of 80°C/W, then you could calculate the hot spot at 1.5W at an ambient temperature of 50°C as 1.5W(80°C/W)+50°C, or 170°C.

Surface-mount components have complicated the temperature question. Without leads, the surface-mount resistor transfers more of its heat directly onto the pc board. Because the parts are smaller, the density of heat-producing parts is higher. The type of pc board, the number of layers, and the weight of the copper plate for the traces all become important factors. A thermal-imaging camera rather than the fine-wire thermocouple often makes temperature measurements at all points on the board. The temperature of the solder joints may become more important than the hot-spot temperature.

The best information that the resistor manufacturer can provide is from powering a single component onto a given size and type of pc board. No universal test standards exist. They may provide this data as temperature rise or thermal resistance and express both in degrees Celsius per watt. Manufacturers sometimes provide two values—one for determining the hot-spot temperature and the other for determining the solder or terminal temperature. In recent years, heat-transfer software has become available to assist in predicting temperature of components at various points on a pc board. To use these programs, you must know the size of the component, the thermal resistance, and perhaps the weight of the component in grams. Just remember that thermal resistance is the largest variable due to the way manufacturers of resistor components make their measurements.

THE TYPE OF PC BOARD, THE NUMBER OF LAYERS, AND THE WEIGHT OF THE COPPER PLATE FOR THE TRACES ALL BECOME IMPORTANT FACTORS.

You will not get an answer about resistor temperature from the resistor manufacturer. The application engineer can provide data that can assist in the thermal analysis of your circuit. You can use this data to determine the size or family of products from a manufacturer to limit the maximum temperature to a desired value. If you compare similar components from different manufacturers, be aware that the conditions they use to obtain the thermal data may differ. When you prototype the pc board, the application engineer can assist in getting samples of one or more resistor products to evaluate. The question of operating temperature has always been difficult to answer. Unfortunately, the answer has become more important as the size and weight of the finished product has decreased.

HOW MUCH SURGE?

A surge condition for a resistor is the application of a power level that exceeds the continuous-power rating of the part for a defined length of time or pulse width. The pulse width is normally 25% or less of the thermal time constant of the resistor. For example, if a resistor has a thermal time constant of 20 sec—time to reach 63% of the final temperature—then the application of a power pulse, exceeding the continuous-power rating, of 5 sec or less would meet your definition. If you apply the same pulse for 60 sec, the resulting temperature would exceed the continuous-rated-power temperature. Applications that have surges that last for seconds are rare. Most surges last for milliseconds or microseconds.

You need to consider both repetitive and nonrepetitive surge conditions. A repetitive surge applies power for a given pulse width and then repeats at a regular interval or time period. You can usually easily measure the period from the beginning of the power pulse to the beginning of the next power pulse because the leading edge of the pulse is often the most defined. For repetitive surges, the average power dissipation over the period of the pulses must not exceed the continuous-power rating of the resistor. To determine average power, first determine the rms power within each power pulse. For a rectangular pulse, this figure is the voltage squared divided by the resistance. For a half-sine-wave pulse, the power is $0.707V^2/R$. For common exponential-capacitor-discharge pulses, a conservative estimate of the rms power over one time constant of the pulse is $0.5V^2/R$. For other pulse shapes, power is normally 0.5 to 1 times the voltage squared divided by the resistance. To choose the proper resistor, be conservative; if anything, overestimate the power. The second step is to find the average power over the period. This figure is the pulse's power, P_{PULSE}, times the ratio of the pulse width to the period: $P_{AVG}=P_{PULSE}(PW/T)$, where P_{AVG} is the average power, PW is the pulse width, and T is the time.

For example, a rectangular pulse applies 100V to a 50Ω resistor for 5 msec. The pulse repeats every 0.75 sec. The average power is $V^2/R(PW/T)$ or $(100)^2/50\ (0.005/0.75)=1.33$W. A 2W resistor is probably the smallest resistor that meets this surge condition. However, you need to quickly check two things. First, if the resistor will be operating in an unusually hot environment, then make sure that the average power is less than the power rating of the resistor after derating it to the temperature of your environment. For example, if the resistor has a power rating of 2W at 70°C and ambient temperature is 100°C, then go to the derating curve and make sure that its rating is still higher than 1.33W. Second, refer to the data sheet for information on repetitive pulses, or contact the application engineer to make sure the resistor can withstand the high power applied for a given pulse width. In this example, the resistor must handle $(100)^2/50$, or 200W for 5 msec.

A nonrepetitive surge is the application of a single high-power pulse to a resistor. The resistor then has sufficient time to cool to the ambient or initial temperature that preceded the pulse. The statement "power rules for resistors" becomes a little shaky under these conditions, when the energy rather than the power you deliver to the resistor during the surge becomes paramount. With a short pulse, the temperature of the resistor material may reach hundreds of degrees Celsius by the end of the pulse. The substrate or case of the resistor remains cool because insufficient time remains for the resistor material to transfer heat during the short pulse. Hence, the resistor material cools down within a few seconds or less as heat flows to the cooler and larger parts of the resistor and, ultimately, to the pc board and the air.

If you apply too much energy, the resulting high temperature destroys the resistor material. Whether the resistor is metal-film, wire, glass or glass-ceramic, its material melts. Resistor manufacturers must accumulate much data to determine the amount of energy that you can apply during a nonrepetitive surge for a given pulse width. They may tell you how many millijoules or joules of energy that you can apply to a resistor. In this case, convert your pulse to energy by multiplying the rms power of the pulse time by the pulse width in seconds and look for a part that has an energy rating that exceeds your calculated value. You should call an application engineer if the duration of your pulse does not match the time or range of times that correspond to the published energy rating.

DETERMINE THE VOLTAGE AND CURRENT FOR THE RESISTOR YOU EMPLOY AND THEN MAKE SURE THAT BOTH FALL WITHIN THE SPECIFIED MAXIMUMS.

Another way of presenting data for a nonrepetitive surge is a plot of maximum pulse power versus the pulse width. A plot of maximum power versus pulse shows that the power level is higher for a nonrepetitive surge than for a repetitive. However, make sure that you read all the fine print! For example, expect also to see a maximum permissible voltage and perhaps even a maximum permissible current. So, in addition to calculating the power you apply to the resistor, you must determine the voltage and current for the resistor you employ and then make sure that both fall within the specified maximums.

If a surge condition falls between repetitive and nonrepetitive, you may have to ask an application engineer about how the manufacturer defines nonrepetitive surges for your selected resistor. For example, three equally spaced pulses that occur over a three-minute period upon initialization of a circuit may not recur until you turn the circuit off and back on again. You may have to check with an application engineer to see how the manufacturer defines a nonrepetitive surge for the resistor product you are considering. As a rule of thumb, treat any surge condition that is between repetitive and nonrepetitive as a repetitive surge.

WHAT CAUSES FAILURE?

Assume that a resistor has no defects, perfectly terminates, and attaches to the pc board with an ideal solder joint. You can bet that such a resistor can fail, and many conditions can cause it to do so. "Power rules" for resistors because the power you apply relates to operating temperature, which relates to oxidation. Resistor materials show little oxidation effects at temperatures lower than a threshold temperature. At temperatures higher than this temperature, oxidation effects typically translate into a positive change in resistance over time. The resistor manufacturer tests its resistors at a number of temperatures with power applications that vary from full load to no load. These test results translate into a maximum percentage of change in resistance when you operate the part at rated power at a given temperature. For example, a resistor may have a rating of 1W at 70°C.

The data sheet also provides a maximum-storage-temperature figure with no applied power. For a 1W part, 150°C is a common temperature where the power applied must be derated to zero. The degradation of the encapsulation material or the solder temperature of the terminal or leads may also influence the zero-power storage temperature of 150°C. Film and wirewound resistors at 0.5W and operating at 70°C may have internal resistance temperatures of 150 to 200 and 200 to 300°C, respectively.

Oxidation leading to a change in resistance can result when you exceed these temperatures by applying more than 1W.

Because oxidation is a temperature-time phenomenon, putting 1.5W on the 1W resistor may result in little change in 24 hours, but the part will then fail when time at the elevated temperature reaches hundreds of hours. A failure occurs when a part exceeds the maximum specified percentage of change in resistance. For example, 1W might be a failure when the percentage of change in resistance exceeds 0.5% when you operate the resistor at 70°C for 2000 hours. On the product sheet, look for the maximum load-life change.

A circuit fault may sometimes cause a resistor to run for an extended time at greater than rated wattage. Underestimating the ambient temperature is a more common problem than overestimating it. A 1W resistor may operate at a maximum of 0.9W, but, if it sees a 110°C ambient temperature, a failure may occur. You may also have to derate a resistor that sits next to another power resistor or a power transistor to prevent a failure.

Excess energy can result in a resistor failure. When this energy level produces a high enough temperature to destroy the resistor material, a catastrophic change in resistance may occur. Instead of a resistance that increases slowly with time for a moderate power overload, the resistance may increase many times or go to an open- or high-resistance state. When evaluating a resistor in a surge application, contact the application engineer and find out what percentage of change signifies a danger signal. Most resistors do not significantly change until the temperature spike from the surge approaches a problem area. You cannot measure that temperature for a 1-µsec surge, but you can measure the change in resistance before and after the surge. For some resistors, even a few tenths of a percentage point of change may be a tip-off that the surge may be excessive.

Sometimes, you base your resistor selection on the surge that a computer-assisted analysis predicts. You prototype the circuit in the lab, and all is well. However, testing the initial production circuit shows failures of the resistor. Often, you can trace the cause to an unexpected surge condition, which may occur when you turn on a circuit or when someone rapidly turns it on and off. A resistor in a motor-control circuit may have a high turn-on surge due to a rare mechanical load condition. You may need a storage oscilloscope to document the maximum possible surge voltage that appears across the resistor.

Voltage stress can also make a resistor fail. Normally, this stress comes into play only on resistors with resistance of more than 100 kΩ and voltage of more than 500V. Lower resistance values can also undergo high-voltage stress conditions during surge conditions. Manufacturers use lasers to trim film resistors to remove a 0.001- to 0.005-in. path. A voltage difference arises across this narrow path. Even if the difference is only 50V, the corresponding voltage stress is 50V divided by 0.005 in., or 0.001 in., which translates to 10,000 to 50,000V/in. stress. These levels would be a problem in the air. However, if the encapsulating material is a good dielectric, the part will not fail.

For a cylindrical-film resistor, the trim cut is a helix cut over perhaps 75% of the resistor's length. Two additional voltage stresses are present. First is the turn-to-turn

UNDERESTIMATING THE AMBIENT TEMPERATURE IS A MORE COMMON PROBLEM THAN OVERESTIMATING IT.

stress, or the voltage you apply divided by the number of resistor turns, divided by the distance between turns. Second is the overall stress. Here we must find the length of the resistor path; "uncoiling" the helix, this figure translates to resistor turns times diameter times π. The voltage stress is the applied voltage divided by the length of the resistor path. For a wirewound resistor, the volts-per-inch stress between wire turns can be excessive. Both film and wirewound resistors targeting high-voltage applications have enough turns to keep the voltage stress at a safe level.

You will realize that voltage stress has become a problem if a modest voltage-stress overload results in a negative change in resistance that exceeds the value for the maximum-STOL-percent change. Also, for a thick-film or composite resistor consisting of metal particles in a glass or ceramic matrix, particles sep-

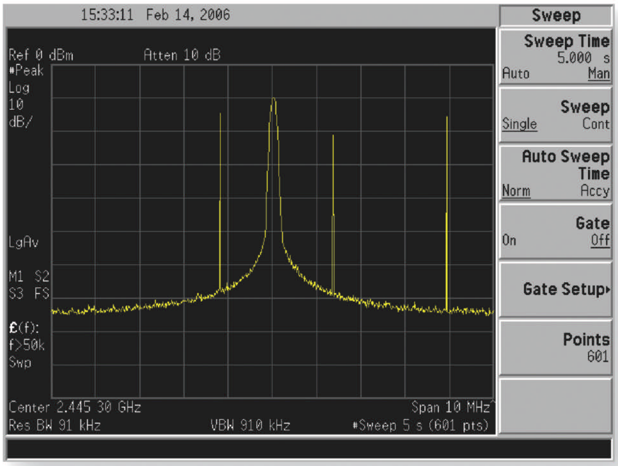

Agilent 5 seconds

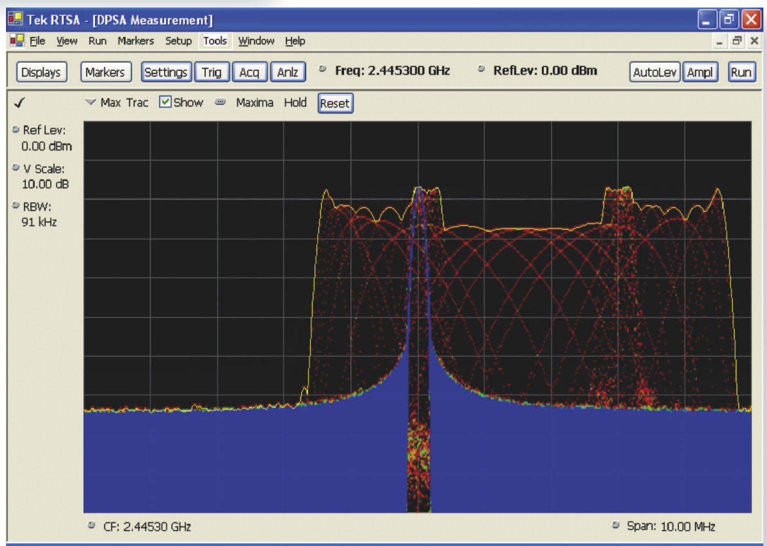

Tektronix 5 seconds

The two screens represent an Agilent E4440A PSA and a Tektronix RSA6114A running DPX™, each detecting an identical signal that changes frequency every 1.28 seconds.

It's time to get real.

Tektronix RSA6100A Series Real-Time Spectrum Analyzers. Nothing else comes close. Troubleshooting transient glitches? Use the best tool. One with greater capture bandwidth (110 MHz) than Agilent, uniquely combined with 73 dB spurious free dynamic range (SFDR) for complete analysis. And nearly 1000 times more spectrum measurements per second than existing products (48,000 spectrum measurements/sec vs. Agilent's 50). And an intuitive color display that shows a live view of your signal changing over time. You'll see RF signal characteristics that are practically invisible with any other instrument.

See for yourself.
www.tektronix.com/getreal

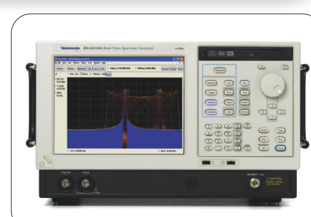

Tektronix®
Enabling Innovation

arated by a thin layer of angstroms-thick dielectric may become microwelded together due to a voltage stress that exceeds the breakdown voltage of the dielectric. For a metal-film or wirewound resistor, a small arc may transport metal across a laser cut or between wires, again lowering the resistance. A negative change in resistance can also occur on any resistor type that has an organic encapsulation as the heat an arc generates can carbonize the encapsulation, providing a shunt path for current flow. High-voltage stresses that produce a high-energy arc can vaporize enough material to cause the resistive value to become significantly positive. A visual examination of the affected resistor often shows evidence of an arc. If you feel that your design is failing due to voltage stress, call the application engineer. Modifying the product or using a product designed for high-voltage stress should solve the problem.

MORE AT EDN.COM

+ Go to www.edn.com/ms4203 and click on Feedback Loop to post a comment on this article.

Although rare, current density, the current flowing through a resistor divided by the cross-sectional area of the resistance material, can cause a resistor to fail. There is a limit to how much current can flow through a given area without causing damage. To visualize this scenario, imagine substituting a 22-gauge wire for a 12-gauge wire in a circuit carrying 20A. Like the finer gauge wire, a resistor fails due to a positive change in resistance if the current density is too high. Generally, the failure mechanism occurs only in resistors of 1Ω or less. Again, if you expect current density on a failure mechanism, you can change to a resistor with greater cross-sectional area or use a material with a higher current-density rating.

CHANGE IN RESISTANCE?

Assume that a critical part of your circuit requires a 1-kΩ resistor. Computer-aided analysis indicates that, if the resistor remains within 4%—960 to 1040Ω—over the life of the product, then you will meet the equipment specifications. You choose a resistor with a tolerance of 1% and a TCR (temperature coefficient of resistance) of 100 ppm/°C; that is, a 0.01%/°C change in resistance for each degree the temperature deviates from 25°C, or room temperature. This figure means that the resistors will all have 990 to 1010Ω of resistance at room temperature. Next, you solder the resistor onto the pc board, so you must allow for the additional resistance of the solder connection. The solder product sheet might specify a maximum of 0.1%. Now, you have used 1.1% of the 4%.

Now, check the circuit at the lower and upper limits of temperature operating range for the product: −50 to +75°C, for example. At this point, the TCR comes into play. A TCR of 100 ppm/°C is equivalent to 0.01%/°C, so a 50°C change in temperature means that the part can change in resistance by 0.5%. You have now used 1.6% of the 4%. Next, you begin long-term testing or field trials of the equipment.

Over an extended period of time, what other changes in resistance come into play? The major sources for change are load life, moisture, high-temperature storage, STOL, thermal shock or cycling, and mechanical shock and vibration. From knowledge of how and where your circuit will find use, you can select the sources of change that best fit your situation. For example, assume that the resistor will operate at approximately 80% of rated wattage and that some of the equipment will work on oil rigs in the Gulf of Mexico. For this application, you might select STOL because power surges are common in almost all applications. You might also select load life and moisture as two criteria. If the application were automotive, you might consider thermal shock and vibration instead. If the maximum change for STOL, load life, and moisture were 0.2, 0.5, and 0.5%, respectively, then you add 1.2% to our previous total of 1.6% to obtain a

THE MAJOR SOURCES FOR CHANGE ARE LOAD LIFE, MOISTURE, HIGH-TEMPERATURE STORAGE, STOL, THERMAL SHOCK OR CYCLING, AND MECHANICAL SHOCK AND VIBRATION.

possible change in resistance of 2.8%.

Engineers often overlook one additional change: resistive change due to self-heating. This change usually comes into play when a resistor dissipates more than 50% of the rated wattage. For a resistor operating at 80% of the rated wattage, look at the data sheet. You'll find that the part has a temperature rise of 100°C at 80% of its rated wattage. The resistive material with a maximum TCR of 100 ppm/°C or 0.01%/°C will go from a 25°C room temperature at no load to 100°C+75°C, or 175°C, when you operate it at the maximum ambient temperature of 75°C at 80% of rated wattage. This operation could produce a maximum change of (175°C−25°C)(0.01%/°C), or 1.5%. You have allowed for the 25 to 75°C ambient change, causing a 0.5% shift, so you must include an additional 1.0% due to internal resistor-temperature change.

This calculation brings your total possible change to 3.8%, or a bit less than the maximum permissible change of 4%. If you had started with a 2%-tolerance resistor, which might have seemed an illogical choice at the beginning of the design, you would have faced some trouble. When change due to TCR is a large portion of overall change, then you might want to consider a resistor with a lower TCR. For example, in the above example, you could have used an initial tolerance of 2% if the TCR had been 50 ppm/°C, or 0.005%/°C.

An expert in statistics would have found a flaw in your arriving at a total of 3.8% change. If you have a number of events that can result in both plus and minus change, and there is an equal opportunity for each event to affect the outcome, and the standard deviations associated with each event are equal, you should calculate the total change as: $\sqrt{a^2+b^2+c^2+...+n^2}$, where a, b, c, and n are the individual percentage-point changes from tolerance, solder effects, TCR, STOL, load life, moisture, and self-heating. Applying this formula to your example yields a total percentage change of 1.67%. A problem can arise from these assumptions, however.

For example, for a given resistor type, our major causes of change may typically cause all positive changes. Also, depending on the environment, one or more of the changes may dominate and be more likely to affect the outcome. The standard deviations are not equal, and the data for standard deviation is sometimes unavailable. In actual practice, the maximum or worst-case total change would probably fall at 1.67 to 3.8%. Another source of error may result from using the data sheet to define maximum changes. For mature resistor-product lines, the application engineer may be able to provide more realistic maximum-change data. For example, whereas the product's data sheet may give 0.5% for load life, an accumulation of load-life data may yield an average change of 0.18% with a standard deviation of 0.05%. Using, for example, ±4 standard deviation would give 99.99% certainty that the maximum change would be 0.18%+4(0.05%), or 0.38%. Most resistor applications do not merit this much attention. The major point is that you must account for more than just initial tolerance and TCR in allowing for resistance change in a critical circuit application. EDN

YOU MUST ACCOUNT FOR MORE THAN JUST INITIAL TOLERANCE AND TCR IN ALLOWING FOR RESISTANCE CHANGE IN A CRITICAL CIRCUIT APPLICATION.

AUTHOR'S BIOGRAPHY

Gene Howell is an electrical engineer at TT Electronics, IRC Boone Division (Boone, NC), where he has worked in the resistor-component business since 1965. He has also taught at Appalachian State University (Boone, NC). He is currently responsible for thick-film-development projects at IRC. Howell holds a bachelor's degree in electrical engineering from North Carolina State University—Raleigh and a master's in education from Appalachian State University. He is a registered professional engineer in North Carolina and a member of the IEEE.

50 YEARS OF SEMICONDUCTOR INNOVATION

50 YEARS IN THE MOST EXCITING INDUSTRY

50 YEARS OF DESIGN INGENUITY

50 YEARS OF INCREASINGLY COMPELLING PRODUCTS

50 YEARS OF *EDN*

VIRTUALIZATION

SOFTWARE-IZATION

THERMAL EFFICIENCY

MINIATURIZATION

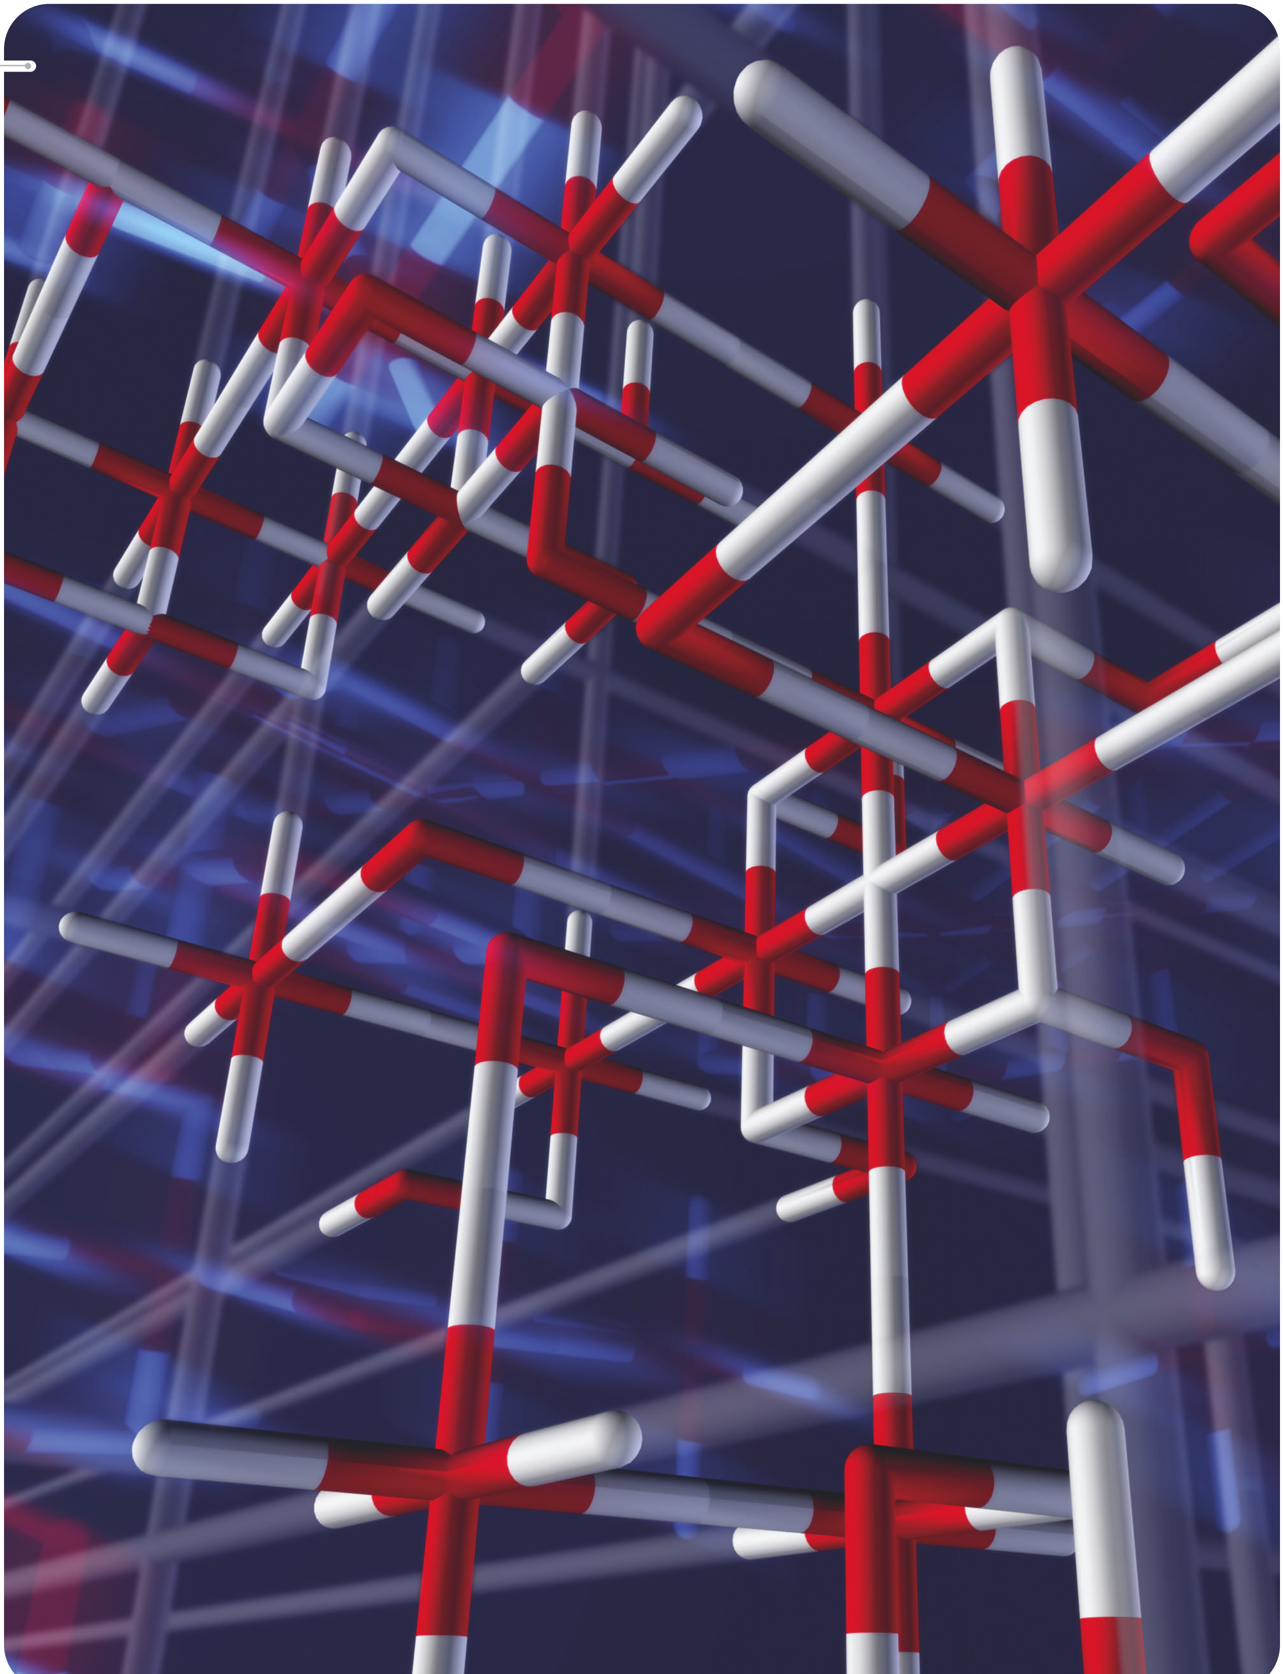

ROB MAGIERA / NOUMENA DIGITAL

BY RON WILSON, EXECUTIVE EDITOR

DOES VIRTUALIZATION DRIVE THE FUTURE?

THE ABILITY OF ELECTRONIC SYSTEMS TO SIMULATE REALITY HAS MADE THEM MORE INTELLIGENT. COULD IT MAKE THEM SELF-CREATING?

Over the 50 years since the first tabloid issue of *Electrical Design News*, underlying forces have driven the evolution of the electronics industry. Since manufacturers first fabricated transistors using photolithography, the inexorable shrinking that process enabled has changed the world. Since it first became viable to put a stored-program computer on a few chips, the shift of functions from hardware to software has changed the world, as well.

These forces have been obvious—with obvious results. But consider a more abstract and less obvious driving force that has, arguably, also been important, and in the future may unleash a revolution as great as those of the IC and microprocessor.

You could call that underlying trend "virtualization." The word is not susceptible to a one-line definition, so let's digress for a moment. When you use electrical quantities to perform physical work or release light, you say the system is electrical. When you use the same quantities—charge, current, or voltage—to convey information rather than merely to do work, you say that the system is electronic. Virtualization is, in this sense, a step beyond electronics. A system—be it a physical process, an object in the real world, or an imaginary person—is virtualized when it has undergone three key steps. First, a boundary must isolate the system from its environment. Second, designers identify the inputs and outputs that cross the boundary, along with the transforms that produce the outputs, thus modeling the system. Third, designers produce a functionally equivalent block—one that accepts the same inputs and produces the same outputs under the same circumstances—with an electronic system.

From at least the mid-1960s, engineers have used electronic systems to virtualize physical things—either components of the electronic system itself or objects in the outside world—and incorporate those models in place of the real objects. This virtualization has made it possible for electronic systems to behave as if they had hardware that they did not have. It has also allowed systems to behave as if they were interacting with a world from which they were isolated—either by distance or by the fact that that world didn't exist. These capabilities have accelerated the growth of electronics and in the future are likely to lend electronic systems capabilities that in the past were the province of humans alone.

IN THE BEGINNING

One of the earliest exercises in virtualization was—like many a breakthrough that later became standard practice in computer architecture—an advance in the IBM (www.ibm.com) 360 mainframe family. Before that time, machine-code instructions in computers referenced memory through physical addresses—numbers representing physical locations on the surface of a drum or, after the development of magnetic cores, physical locations within the array of small ferrite cores.

The development of virtual memory rested on the idea that the addresses that machine instructions created needn't be the last word. If arithmetic hardware were fast enough, it could translate addresses from the computer program on the fly into physical addresses using some prearranged mapping. This feature allowed a program that was assembled to run at one location in physical memory to execute in another location—even if the programmer had not made all the memory

SOFTWARE I/O, VIRTUAL I/O, OR SOFTWARE-ASSISTED I/O?

By David Fotland, Ubicom Inc

Typical microprocessor families have dedicated hardware for each I/O function. This difference leads to families of chips with the same CPU but different I/O mixes. Cost is higher because the semiconductor company must make many chip versions, and mask costs are high in state-of-the-art process technology. The alternative is to create an SOC (system on chip) with all of the I/O hardware on the same chip. This approach also leads to higher cost, because the customer is paying for silicon to implement I/O that he won't use in his application.

The solution to this problem is software I/O. Some 8-bit microcontrollers use this technique, called "bit-banged" I/O. If the microcontroller has on-chip memory and deterministic execution, the software can directly control I/O pins to implement the I/O protocol. A simple example is a UART. The start bit causes an interrupt, and software reads the input pin to receive the data. While the data is arriving, the CPU cannot do anything else, so this technique is useful only for I/O that is infrequent or intermittent. The interrupt-response time limits the use of this technique to low-speed I/O.

Some 32-bit processors, such as those from ARM (www.arm.com) or MIPS (www.mips.com), can't use software I/O because code execution is far from deterministic. Pipeline hazards and cache misses make it impossible to use instructions for accurately timed external events. Operating systems such as Linux turn off interrupts for milliseconds at a time, making real-time I/O response impossible.

Ubicom (www.ubicom.com) has the only 32-bit CPU that uses software I/O. The multithreaded CPU has a hardware scheduler that can select a thread for execution during every clock. Real-time threads have a fixed schedule and deterministic execution, even if other threads have mispredicted branches or cache misses. The unit has 10 threads, so it can allocate one real-time thread to each I/O port to manage that port.

The instruction set supports software I/O and packet processing. An instruction can move data between memory and I/O. MIPS and ARM CPUs, in comparison, need two instructions: a load and a store. Single instructions can set, clear, or test any I/O bit. Interrupt-response time from an I/O event to scheduling instructions in the managing thread takes only a few CPU clocks. When an I/O port is idle, it suspends its managing thread, using no CPU resources.

The high-performance, 32-bit CPU can use software-I/O for functions more complex than a UART. Ubicom has implemented a full PCI bus at 27 MHz, MPEG Transport Stream, IDE, and Utopia in software. It has also implemented MII (media-independent interface) for 10/100-Mbps Ethernet, USB, SPI, GPSI (graphics-processor software interface), and other serial interfaces with a combination of hardware and software. By spending 10 to 20% of the CPU throughput on software I/O, the company dramatically reduced the die area necessary for I/O, resulting in a flexible single chip to cover a wide range of applications.

AUTHOR'S BIOGRAPHY

David Fotland is chief technology officer of Ubicom, where he led the architecture of the Ubicom multithreaded processor for packet processing and software I/O. He is responsible for all aspects of architecture and definition of processors and solutions. Previously, Fotland spent 21 years at Hewlett-Packard Co (www.hp.com) as lead engineer, project manager, and system architect. He was a key participant in the PA-RISC instruction-set definition and was lead designer of the first PA-RISC processor and system. He was a leader in the development of the HP-Intel (www.intel.com) Itanium instruction set.


CRYSTALS & OSCILLATORS
SEMICONDUCTORS & OPTOELECTRONICS
COILS & INDUCTORS
CAPACITORS & RESISTORS
ONE
STOP
SHOP
SWITCHES & RELAYS
INTERCONNECT
POWER SUPPLIES & TRANSFORMERS
RF SEMICONDUCTORS
Digi-Key®
CORPORATION

references relative to the contents of a base register. More important, it also meant that programs could run on a machine whose physical memory was many times smaller than the virtual-address range the program used. A portion of the operating system—a well-developed concept by this time—could allocate portions of the physical memory as necessary for the immediate needs of the program, swapping blocks of information onto and off of disk drives as necessary. By extension, this ability meant that a modest-sized machine could concurrently run a number of large programs, convincing each of them that it had access to all the physical memory it wanted when in fact it was borrowing small blocks of memory on an as-needed basis. Severing the link between the program's address space—which now became virtual—and the physical-address space was a huge and vital step in the creation of modern data processing. But it was far from the last one.

At first glance, it might appear that this scenario has little to do with the definition of "virtualization." However, the process executes all of the steps in virtualization. IBM's designers isolated physical memory from the rest of the mainframe system. They identified the inputs—addresses and write data—and the outputs—timing signals and read data—that characterized the system. And they constructed a combination of hardware, which would become a memory-management unit, and software, which would become the virtual-memory manager that created a virtual main memory.

VIRTUALIZING THE IT WORLD

Rich Lechner, vice president of IBM Virtualization (www-03.ibm.com/systems/virtualization), defines the term as "the logical representation of resources not constrained by physical devices." He points out that, when you use virtualiza-

IMMERSED IN ENGINEERING, ADVANCED 3-D VISUALIZATION PROMOTES INSIGHT

By Jeff Brum, Fakespace Systems

As electronics speed into an era in which manufacturers fabricate not just circuits, but also physical structures themselves in nanoscale geometries, the role of computer-based simulation as a design tool is increasingly important. Correspondingly, the benefits of visualization in the review and analysis of simulations play a growing role. Looking to the future, immersive stereoscopic display tools will amplify the power of visualization.

The adoption of immersive visualization in electronics design revolves around several factors. Atomic-scale phenomena, which are major concerns as the semiconductor road map extends beyond the 65-nm-process node, have been major players in advanced visualization techniques. Scientists at NIST (National Institute for Standards and Technology, www.nist.gov), for example, use a stereoscopic display–two walls and a floor–to gain insight into the molecular bonding of "smart-gel" polymers (Figure A and Reference A).

Similarly, researchers at LANL (Los Alamos National Laboratory, www.lanl.gov) use a range of immersive environments–from wall-sized to a 43-million-pixel, five-walled projected room–to view terascale data sets (Figure B). Bob Green, visualization specialist at LANL, notes that the researchers "are viewing simulations based on computations that generate more data than is contained in the entire print collection of the Library of Congress in one calculation."

In engineering, visualization has had its largest impact to date in macroscale CAD programs, such as automotive and aerospace design, and the evaluation of complex structures for interferences that are not readily apparent in simpler graphical representations.

Figure A NIST research chemist Carlos Gonzalez uses a 3-D immersive environment to study shake gels. Comprising two walls and a floor, the system displays stereoscopic images with more than 3.1 million pixels of resolution (courtesy Robert Rathe).

As simulation and visualization data accumulate in MEMS (microelectromechanical-system) design, this type of modeling and simulation will grow in importance. The ability to visualize and "fly through" transistor-scale structures and even large segments of a complex microcircuit design will also benefit from advanced 3-D visualization that blends multiple streams of data,

continued on pg 130

THINK Wireline and Wireless ICs.

THINK Seamless Communications.

THINK Infineon.

OUR MISSION is to help your customers achieve seamless communications. Count on Infineon to deliver the architectural building blocks.

AT INFINEON, we never stop thinking about end-to-end semiconductor solutions. Our broad portfolio in wireline and wireless technology enables us to deliver innovations that help you build tomorrow's converged networks. In wireline, we lead the way in broadband access, CPE, core access and VoIP system-on-chip products. In wireless, we're number one in RF technologies with the solutions you need from 2G, 3G, WLAN to an entire line of flexible mobile phone platforms including entry phones, feature phones and mobile software solutions. We also provide complete solutions, including application software and ready-to-use platforms, to help you get to market faster.

WHEN YOU think about seamless communications, turn to the company that delivers tomorrow's wireline and wireless innovations. Think Infineon.

www.infineon.com/us/comm

Infineon

Never stop thinking

tion in this way, it can "either treat one physical resource as if it were many or treat many, possibly dissimilar, resources as if they were one." Lechner traces the beginnings of virtualization to the 360 family's virtual-memory-system debut 40 years ago. But he says the practice has become far broader today than simply virtualization of memory, which is now a common practice even in microprocessors. "At one level, we can gather all of the storage assets available to a network into a single pool of storage," Lechner says. In this way, any program executing in the environment has access to all of the storage assets of the network through a single interface.

But if you are not careful, you will face chaos. The location of stored data does make a difference—in access time, cost, persistence, and coherence. So, for storage, virtualization has come to mean more than just providing a mapping from a single mass-storage API (application-programming interface) to a diverse set of storage devices. "The next level involves the routine cleansing and deduplicating of the storage system," Lechner explains. The virtualization system must make sure to remove stale copies of data, propagate updates, eliminate duplicate data sets, and keep data in a place that is most convenient to its clients. "This all by itself is a significant benefit," Lechner says. "Our field studies indicated that, before virtualization, the average midsized enterprise stores the same data in at least 20 places around the network."

The process does not stop with storage systems. In just the same way, system programmers can identify the computing resources in a network, give them wrappers that present a common API, and hence virtualize them. In this way, the computing power available to an appli-

50TH ANNIVERSARY ONLINE
THE CELEBRATION CONTINUES
For more on virtualization, read "Raising the bar with hardware-accelerated physics" by Ageia Technologies' chief architect Sanjay J Patel at www.edn.com/50th.

continued from pg 128

including device characteristics, interconnects, and material behavior.

Wall- or even room-sized immersive displays support a more collaborative, more intuitive, style of work than do graphics processors. A major benefit in the review of complex structures is the ability to "move" freely along all three axes of an image and still maintain the visual acuity of high-resolution desktop processors. (Today's projectors support resolution of 1400×1050 pixels or higher.) Viewing nanoscale design can feel like taking a helicopter tour over a dense city center with the ability to identify and zoom into landmarks. Improved insight and collaboration lead to faster and better decisions, speeding time to market and reducing development costs.

Although researchers are just beginning to measure the benefits of immersive visualization (Reference B), they generally agree that they bring new levels of insight to engineers and scientists. In immersive visualizations, groups of users can view the inner workings of devices and gain deeper understanding of electromagnetic effects and the relationships between elements of a design. With the availability of dual PC clusters and advanced graphics cards, these types of virtual environments no longer require specialized graphics supercomputers. The increased accessibility of immersive visualization makes its addition to the tool kit of electronic engineers practical and inevitable.

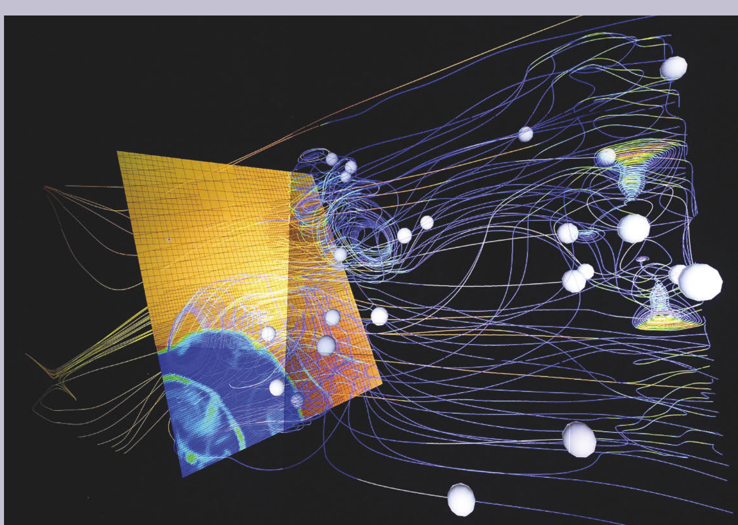

Figure B When researchers view phenomena in "La Cueva Grande" at LANL, they can better understand, for example, this 3-D, hydrodynamic simulation, showing the interface between two gases of different densities. The 33-projector, five-sided immersive environment produces 43 million-pixel imagery for ultraprecise detail (courtesy Presley Salaz, IM-9, Los Alamos National Laboratory).

REFERENCES

A "Recipe for a Shake Gel," National Institute for Standards and Technology, Aug 27, 2003, www.nist.gov/public_affairs/newsfromnist_smartgels.htm.

B Kraak, Menno Jan, "Beyond Geovisualization," *IEEE Computer Graphics and Applications*, May/June 2006, Volume 26, No 3, pg 6.

AUTHOR'S BIOGRAPHY

Jeff Brum is a vice president at Fakespace Systems (www.fakespace.com), a division of Mechdyne Corp with a 15-year history of developing and supporting advanced systems for immersive visualization in science, engineering, and public-exhibition applications.

THINK Automotive Electronics.

THINK Zero Defects.

THINK Infineon.

YOU MAKE no compromises when it comes to quality. That's why we created Automotive Excellence, the industry's most rigorous zero-defect program.

AUTOMOTIVE EXCELLENCE redefines reliability and quality for automotive semiconductors. We don't settle for merely detecting defects. We are in the process of developing a holistic system to prevent defects. Implementing the industry's most rigorous zero-defect initiative requires new ways of thinking and a major commitment of time and capital. But the benefits in reduced cost and increased customer satisfaction are worth the investment.

INFINEON HAS been exceeding customer expectations in automotive electronics for 35 years. Today, our broad technology portfolio includes IC solutions for powertrain, safety management, and body and convenience. Our focus on your success and the Automotive Excellence program have earned us the #2 position in worldwide automotive semiconductor market share.

WHEN YOU think of zero defects, think of the company dedicated to pursuing perfect quality with no compromises. Think Infineon.

www.infineon.com/us/auto

Infineon

Never stop thinking

cation becomes a slice of the entire computing community on the network, not simply the power of the machine on which the application happens to be running.

The next step in this process is even a bit more abstract. You can virtualize storage devices and servers, but what about applications? In exactly the same way, system programmers can draw a wrapper around each application, provide it with a set of APIs, and make it available to the network as a virtual application—say, a virtual database. Thus, when an application program executes in the network, it may be interacting with virtual-storage devices on a number of storage networks, running on a virtual processor when parts of the code are executing on a half-dozen servers around the network, and calling virtual applications that may be databases from different vendors with different organizations.

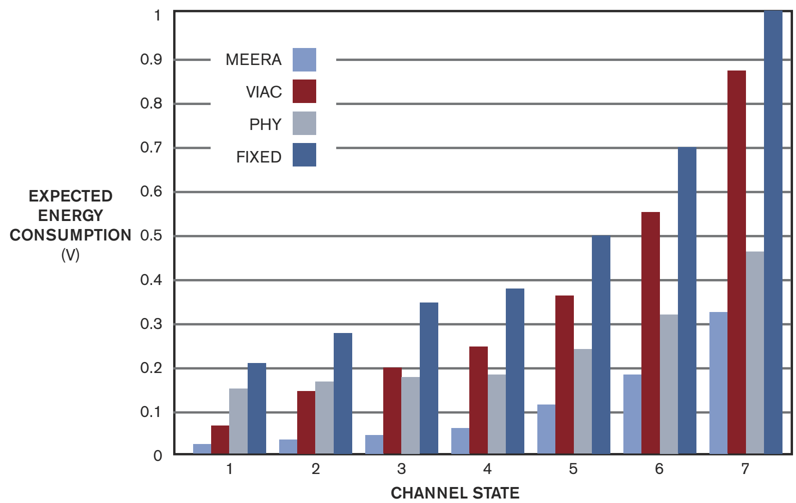

NOTE: MEERA=METHOD FOR ENERGY-EFFICIENT RESOURCE ALLOCATION.

Figure 1 Software-defined radios yield energy-management savings. The channel-state axis indicates increasingly poor radio reception, with State 1 being the clearest channel. The colored bars represent energy-management schemes. MEERA is the scheme using the virtual model of the radio.

IT TO EMBEDDED COMPUTING

All this virtualization might seem to be irrelevant to anything outside the IT (information-technology) world. But if you think of a multicore embedded system—based on the IBM Cell processor, for example—as a heterogeneous collection of computing resources, storage assets, and interconnects, perhaps the relevance becomes more clear. As has so often been the case, the IT solution of today is the SOC (system-on-chip) solution of tomorrow. Virtualization may be the concept that makes SOCs with diverse computing sites usable in real applications.

In fact, we are already seeing indications that this is the case. At IMEC (Interuniversity Microelectronics Center, www.imec.be) in Leuven, Belgium, researchers have created a virtual model of a runtime-configurable, multicore-computing system for software-defined radio. Antoine Dejonghe, a principal scientist at IMEC, describes the situation: "IMEC's M4 project is creating a radio system that can be agile in real time across wireless protocols and media types," he says. "We believe that technology scaling by itself won't be enough to bridge the energy gap between what our configurable architecture requires and what batteries will be able to provide. So, the viability of the system depends on being able to model the entire system, from analog front end through the media-access-controller software, in terms of its energy consumption. We will use this model to establish the optimum trade-off between energy and quality of user experience at runtime—perhaps as often as every 10 msec."

Today, the models are under construction—using the same sequence of steps—to create a virtual software-defined radio. Designers have theoretically modeled the analog front end, for instance, with approximately 11 control bits as inputs, along with theoretically calculating output power, distortion, and energy consumption. The designers are now calibrating these models against actual silicon measurements. They will construct similar models to predict the performance and energy consumption of the configurations of IMEC's ADRES (architecture-for-dynamically reconfigurable-embedded-system)-computing cores in the radio project.

VIRTUALIZATION MAY BE THE CONCEPT THAT MAKES SOCs WITH DIVERSE COMPUTING SITES USABLE IN REAL APPLICATIONS.

The result will be a virtualization of the radio that should execute in software, consuming less than 5% of an ARM9 and continually optimizing the radio for current traffic, link quality, and application variables (**Figure 1**). This process, Dejonghe believes, can bridge the gap between what batteries can deliver and what battery life users will demand.

CREATING A VIRTUAL WORLD

IMEC's work creates a virtualization of a physical system. However, designers are virtualizing larger and more complex systems. Another easy place to find excitement in virtualization is in the world of electronic games. Traditionally, that virtualization did not exist. Most video games today have a lot more in common structurally with a Walt Disney cartoon than with a simulation of the physical world. Games, like cartoons, have story lines. The choices of the player do not directly interact with the world of the game; players merely choose story lines, and the game proceeds down one of the paths the scriptwriters designed for it.

This scenario is true even at the macro level. When you slip around the corner and level the four-eyed alien from the

planet Grr with your super halitosis blaster, you unleash a sequence of animation frames, in which the Grrien satisfyingly rips open, tumbles through the air, and lands as a puddle of unfamiliar proteins. This sequence is almost the reverse of motion estimation in video-compression algorithms: It begins with a set of key frames and animates incremental motions of objects against a relatively fixed back ground. To the player, the result is much the same no matter which way he triggers the sequence, unless he misses.

As games become richer and players become more sophisticated, this situation causes a problem, according to Manju Hegde, chief executive officer and chairman of gaming "physics-engine" developer Ageia. He has a vested interest in this problem because Ageia supplies software and hardware accelerators for performing the dynamic computations necessary for eliminating the animation sequences and computing the trajectories of the spinning Grrien bits.

No one—least of all, the animators—would dispute that it is better to do a dynamics-based simulation of the objects in the game world than to rely on animation sequences. However, the greatest benefit would be to the game architects. Because of the labor involved in producing animation sequences from key frames, a game can have only a limited number of animation sequences. So, architects must confine the action of the game so that only a limited number of outcomes are possible at any time. The art of game design today is to make the game feel rich and unscripted without causing an exponential explosion in the number of key frames that developers must prepare.

Hegde illustrates with a rather less violent example: a basketball game. "If you want a realistic slam-dunk sequence in your game, you start out by rigging lights all over your star player and then you record him doing some dramatic dunks," Hegde explains. "Then, back in the lab, you extract from the recording key frames and interpolate the movement of the image edges to produce animation. This animation can look natural on the screen, but, any time you push the slam-dunk button in the right context, you are going to see the same sequence."

The alternative approach is to build a physics-based model of our hero—that is, to virtualize him. "We create a physics-based model of the body," Hegde says. "Today, it can be as detailed as having approximately 200 bones connected by six kinds of joints. We then model each of these bones and joints according to the laws of physics. You apply forces to them, and they respond. Now, the dunk becomes the dynamics of the individual bones and joints in the model person. If you view a game from a distance during fast action, a game might use just 20 bones to reduce the calculations, but it looks 'right.'" Hegde thus describes a scenario that fits this article's definition of virtualization—in this case, of a basketball forward.

This virtualization has so many advantages over conventional animation that manufacturers would—except for a couple of issues—produce all games in this way. One of these issues goes back to scripting. If the outcome of the player's input depends on both physics and the game script, the sequence of play can quickly become unmanageable. What if physics dictates that the player breaks his wrist on the rim of the basket and leaves the court writhing in pain? Game architects who employ physics-based models often intervene in the simulation to direct it to allowable outcomes, so that the game stays within its script. This delicate business blends physical simulation and animation.

A more brutal problem is the amount of computing requirements. "Today, an AMD FX-62 dual-core CPU running at 2.8 GHz may be able to handle a couple of characters with extensive bone models," Hegde says. "But you couldn't do physics-based simulation with a large number of characters on the screen at once. The bursts of computation necessary to support a number of characters all running, for instance, would overwhelm the processors."

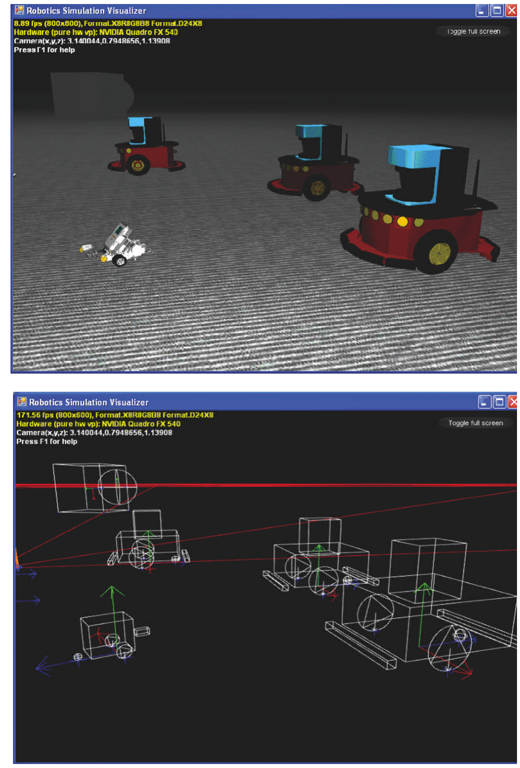

Figure 2 The Microsoft robotics-development kit renders three MobileRobots Pioneer3DX robots with lasers plus the new Lego Mindstorms NXT. Meshed and wire-frame views illustrate the simplicity with which the physical model represents entities.

Microsoft's Robotics Studio (www.microsoft.com/robotics) is developing a similar application, also using Ageia's technology. Rather than playing games, the Microsoft group provides a virtual-development environment for the programming—and, eventually, the design—of robots. Tandy Trower, general manager of the group, says that the need for such an environment is obvious across the spectrum—from industry to education. On one end, with a KUKA (Keller und Knappich Augsburg) Robotics (www.kuka.com) robotic arm selling for more than $100,000, industrial developers need a low-cost environment in which to develop and test programs. At the other extreme, robotics has proved to be

Accelerating the **Speed of Design.**™

Avnet, your one solution for design creation.

Avnet Memec focuses on a select number of leading suppliers to provide original equipment manufacturers (OEMs) with the technical expertise you need in designing cutting edge products.

With factory trained and certified technical support, Avnet Memec specialists offer a level of technical depth that is unmatched in the industry. That translates into faster time to market for you, our valued customer. And it keeps getting better. **Avnet is ranked #1 for having knowledgeable FAEs and well-informed sales reps.* We deliver:**

> Proactive Engineering Support

> Factory Certified FAEs

Accelerate Time to Market.™
Visit the Avnet Design Resource Center™ at:
www.em.avnet.com/drc

AVNET® Memec

1 800 408 8353
www.em.avnet.com

Avnet Green Initiative

Supplier Authorized Distributor

*As rated by Electronics Supply & Manufacturing, 2006 Distributor Customer Evaluation Study.

one of the few endeavors that can attract US students to engineering and mathematics. Even simple robots, however, are far beyond the reach of most secondary schools, even though students have proved capable of programming—and even designing—them. So an affordable virtualization of a robot and its surroundings would be a big win, and Microsoft is attempting to achieve that goal.

Building a virtual world behind the Direct-X graphics environment, the company has provided libraries of models for popular robots; primitives for constructing physical objects, machines, and other robots; and a physics-driven simulation engine to animate them. "There are two ways of creating entities," Trower explains. "Developers can write them directly as code—in C, Visual Basic, Python, and the like—through a managed code interface or assemble them from basic geometric shapes and assign them physical characteristics, such as mass, hardness, and so on" (**Figure 2**). Once you set them in motion, the robots, which are themselves entities, can interact with other entities, including other robots, in a world that physical laws govern.

Trower points out that, although Microsoft's work in this area is on one level similar to the physics-based modeling that is starting to appear in games, it differs dramatically at another level. "Games take place in a well-defined environment," he says. "Robot simulation does not. You have to work out everything that happens based on physics, because there is no script." Trower says that the virtualization technology could move from the development tool into the robots themselves. For example, you can blend the simulation with real-world sensor data. A robot that can include an optional—and expensive—laser range finder, for example, might instead have a virtualized range finder; fusing other sensor inputs to generate the range data. In the future, you may see the next step: robots virtualizing the world around them, a scenario that Brooke Williams, DSP-automotive-vision-marketing manager at Texas Instruments (www.ti.com), sees before his own eyes.

"We are starting to see automobile-safety systems fusing sensor data to create a virtual model of the car's surroundings," Williams says. "This model can either be presented to the driver as warning information, or it can be used to take control of the vehicle." For example, TI is combining radar, a good tool for detection and ranging but poor for forming images, with machine-vision systems, which are great at finding edges and bearings but poor at ranging or detection. This combination of tools will allow the creation of virtual models of the objects surrounding a car. "The object is to predict a crash; prepare the vehicle by tightening seat belts, arming air bags, and closing the windows, for instance; and attempting to take evasive action," Williams says.

But to achieve this goal, simple proximity warning is insufficient. The tools must identify objects from their surroundings and track and categorize them; the insurance industry must know that the system can distinguish between pedestrians and shrubbery, for instance. You want to neither turn away in panic from a car that can't physically reach your trajectory nor mow down a pedestrian to avoid destroying a hedge. "Manufacturers are even talking about external air bags that could deploy to protect pedestrians in collisions," Williams says. Such decisions require not just a measurement, but also an understanding of the car's surroundings.

Therein lies a possible endpoint for the trend of virtualization: electronic systems that can not only sense, but also model and predict their environment. Such systems exhibit artificial intelligence and also express virtualization. These systems, protecting drivers from their own folly, exploring the otherwise-inaccessible reaches of the physical world, and using their virtual models to reason about their surroundings, represent another major expansion in the role of electronics. EDN

You can reach Executive Editor Ron Wilson at 1-408-345-4427 and ronald.wilson@reedbusiness.com.

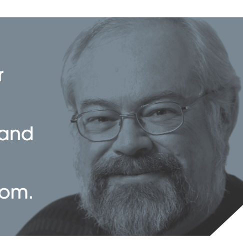

A new standard is open

See it fully realized through Agilent Open

The light's gone on for even easier system integration—and Agilent has flipped the switch. Agilent Open now includes devices based on LAN eXtensions for Instrumentation (LXI), and you can use this new standard to accelerate your system development and deployment.

The versatility of LXI utilizes Ethernet to enable local and remote measurements, monitoring, and control—and eliminates costly cardcages and proprietary interfaces. With more than 40 member companies supporting the LXI Consortium, the growing number of LXI devices will have you well-equipped to take the next step beyond GPIB.

LXI is one more way Agilent Open continues to simplify system creation. Our products are system ready, can be integrated into your test-software environment, and provide USB, GPIB, and LAN I/O options. Enlighten your whole testing process—get details about LXI and Agilent Open at www.agilent.com/find/openlxi.

Agilent Open

Open new possibilities with our growing line of LXI products:

- *34980A Multifunction Switch/Measure Unit*
- *34410A & 34411A Digital Multimeters*
- *N6700 Low-profile Modular Power System*
- *N5700 Series System DC Power Supplies*
- *N8201A 26.5 GHz Performance Downconverter*
- *N8211A 20/40 GHz Performance Analog Upconverter*
- *N8212A 20 GHz Performance Vector Upconverter*
- *N8221A 30 MS/s IF Digitizer*
- *N8241A 15-bit Arbitrary Waveform Generator*
- *N8242A 10-bit Arbitrary Waveform Generator*

u.s. 1-800-829-4444
canada 1-877-894-4414

www.agilent.com/find/openlxi

Agilent Technologies

PolarPro™ Low Power FPGAs–Specifically Architected for Portable Electronics

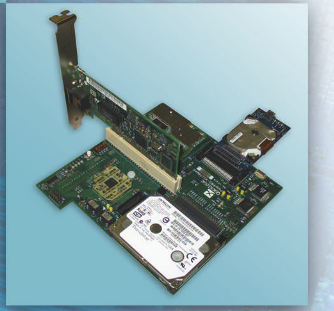

The PolarPro FPGA family is the lowest power programmable logic available in the industry today. Drawing less than 10 uA standby current, it is an ideal solution for power sensitive applications in the portable electronics market.

PolarPro can exhibit ASIC-like power consumption in static, idle and dynamic modes significantly extending battery life.

QuickLogic's unique Very Low Power mode (VLP) allows I/O states, RAM and register values to be retained, and isolates the logic array and I/O in order to further reduce power. In addition, the transition between VLP mode and normal operation takes only microseconds. Altogether, these features result in an energy efficient, cost effective and flexible solution, all offered on a single chip.

Emerging wireless, mass storage technologies such as Wi-Fi, miniature Hard Disk Drives (HDD), and mobile TV, are converging in the portable consumer electronics market. However, adopting these new features can have a significant impact on battery life.

PolarPro was specifically architected to both enable the adoption of these technologies and minimize energy consumption over competing technologies. Using PolarPro FPGAs, designers can extend battery life by as much as 4x over competing technologies.

FEATURES

- **Very Low Power Mode (VLP)**
- **Small Form Factor Packaging**
- **Embedded SRAM/FIFO Blocks**
- **Configuration Clock Manager (CCM)**
- **Native Support for DDR SDRAM**
- **Enhanced Logic Cell Architecture**

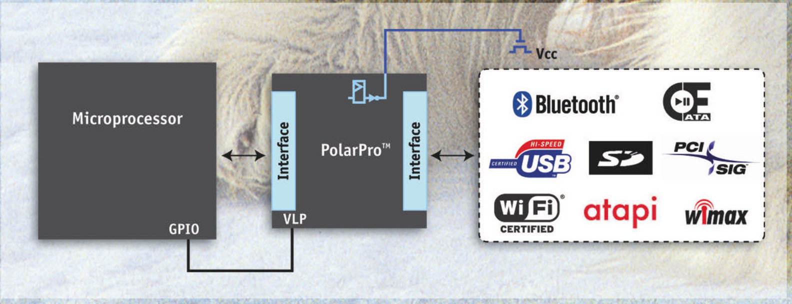

QUICKLOGIC®

POLARPRO™

www.quicklogic.com
info@quicklogic.com
408.990.4000

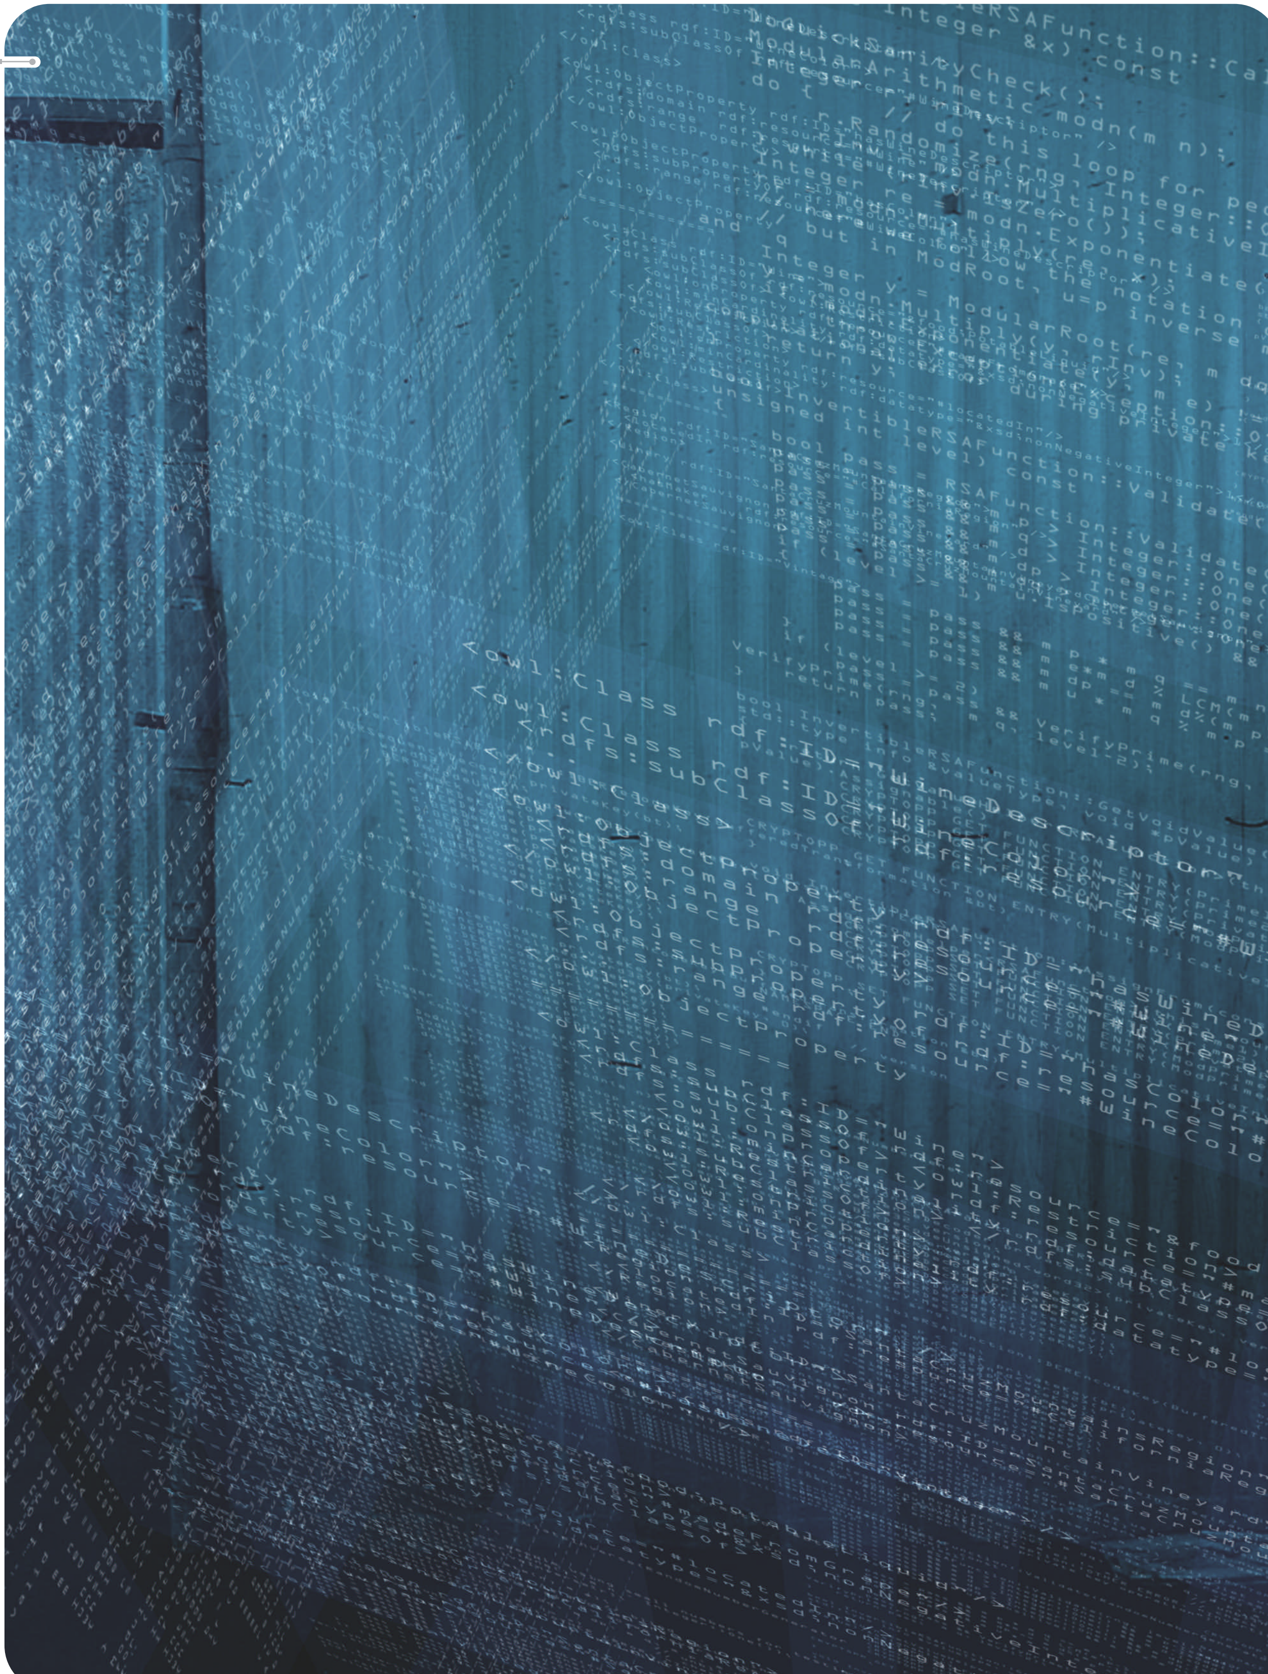

ROB MAGIERA / NOUMENA DIGITAL

BY STEVE LEIBSON, TENSILICA INC

SOFTWARE AND PROCESSORS: HERE, THERE, AND EVERYWHERE

SOFTWARE WAS ALMOST NONEXISTENT AT *EDN*'S DEBUT 50 YEARS AGO. NOW, SOFTWARE AND PROCESSORS ARE UBIQUITOUS.

During *EDN*'s inaugural publication year in 1956, hardly any software existed. Although the University of Cambridge's EDSAC—the first stored-program computer—became operational in 1949, computers and programmers remained rare until 1954, when IBM (www.ibm.com) introduced the model 704. In 1956, the first PC ("personal" computer) and the first significant HLL (high-level-language) compiler appeared, helping to spark the explosive growth of computers and software. Systems based on the DNA-like intertwining of software and stored-program processors have now covered the planet and ventured into interstellar space (see **sidebars** "Software-ization" on pg 144 and "Hardware morphs into software" at www.edn.com/50th).

THE FIRST COMPUTER PROGRAM

Manhattan Project physicists Nick Metropolis and Stanley Frankel created and ran the world's first computer program on ENIAC (electronic numerical integrator and computer)—the first programmable electronic computer—two months after the end of World War II (**Reference 1**). ENIAC lacked stored-program abilities; plug wires, patch-panel configurations, and switch settings comprised its programs. During most of the war, Frankel ran a computing-service bureau at Los Alamos, NM (birthplace of the atomic bomb), staffed with calculator operators (called "hand computers") and punch-card tabulating-machine operators. Frankel and Metropolis developed complicated numerical-analysis algorithms for tough nuclear-physics problems and "ran" these algorithms on the tabulators and hand computers. Frankel became so obsessed with the automated tabulating machines' abilities that he ignored his managerial duties. In January 1945, management transferred Frankel to Enrico Fermi's F Division to work with Edward Teller on the thermonuclear "super" bomb. That move took Frankel to ENIAC.

Metropolis and Frankel arrived at ENIAC's home in the Moore School at the University of Pennsylvania with 1 million punch cards containing their initial program data. The "Los Alamos Problem" that ran on ENIAC in late 1945 and early 1946 was a 1-D nuclear-fusion simulation. Details remain classified, but this calculation was the most complex that anyone in the history of science ever attempted when it ran as ENIAC's "shakedown" program, flushing out many hardware bugs. ENIAC's limited numeric storage—20 10-digit accumulators—forced the physicists to oversimplify their equations, so the numeric result itself wasn't especially meaningful. The real goal of the exercise was to see whether ENIAC or a much bigger version of ENIAC could help develop Teller's super bomb. The conclusion? Yes.

After the war, Frankel consulted for defense contractors but lost his security clearance in 1949. California Institute of Technology (www.caltech.edu) hired Frankel to run a computing-service bureau, where he spent most of his time studying the nascent fields of digital logic and computer design. He dreamt of small, affordable computers that schools could purchase. Frankel designed and breadboarded a computer he called MINAC (minimal automatic computer). It used only 15 flip-flops and stored its register values and memory contents on a rotating-drum memory that Caltech physics major James Cass built by hand.

THE FIRST PC

MINAC was slow, cheap, reliable, and marketable. Librascope—a nearby military contractor wanting entry into the commercial computer market—licensed MINAC from Caltech, hired Cass, and engaged Frankel as a consultant. Cass' engineering group "productized" MINAC. The resulting LGP-30 computer—a compact, desk-sized machine that employed a Flexowriter (an electromechanical typewriter with paper-tape reader and punch) for a programmer's console—required only 115 vacuum tubes and 1450 germanium diodes. It fit into a small room and needed no special climate control. A programmer could enter the room, close the door, and be alone with the machine.

Librascope announced the LGP-30 in mid-1956. It was the first PC. The company built more than 500 LGP-30s—a phenomenal number in an era in which most computers were one-of-a-kind machines. The pool of LGP-30 programmers grew large enough to form the Pool user group, and one-third of the LGP-30s went into business applications. Librascope even transformed three LGP-30s into early embedded systems for industrial-process control.

50TH ANNIVERSARY ONLINE
THE CELEBRATION CONTINUES
For more on software-ization, read Jack Ganssle's article, "Hardware morphs into software," at www.edn.com/50th.

Librascope shipped the first LGP-30 in September 1956. One month later, IBM published the first Fortran programmer's reference manual for the model 704, IBM's first mass-produced computer with magnetic-core memory. IBM programmer John Backus had assembled an HLL-development team shortly before the machine's introduction in May 1954. He sought a more abstract, algebraic language that would make programming faster, cheaper, and more reliable to stem rising assembly-language-programming costs—which were already approaching hardware costs. Backus unveiled Fortran in a paper he presented at the Western Joint Computer Conference in February 1957.

As Frankel had foreseen, many colleges and universities purchased LGP-30s. Dartmouth College (www.dartmouth.edu) acquired an LGP-30 in 1959. John Kemeny, a hand computer at Los Alamos, NM, during the war; Thomas Kurtz; and their students developed many simplified HLLs on the LGP-30, including Algol 30 and Scalp (self-contained Algol processor), which led to the development of Basic on a time-shared GE-225 computer in 1964.

Computer and software development exploded throughout the 1950s and 1960s. Large-scale software development became practical as HLLs, including Fortran and Basic, spread. Minicomputers, starting with the PDP-1, which DEC (Digital Equipment Corp) introduced in 1959 (www.computerhistory.org/pdp-1), lowered the cost of computing. Fortran II, III, IV, and more powerful programming languages, such as Algol 60 and Pascal, fueled a growing army of programmers. However, computers and software still failed to proliferate widely. Hardware costs remained too high, but the situation was about to change.

THE KITCHEN COMPUTER

A glimpse of the near future appeared in an unlikely place. The 1969 Neiman Marcus catalog featured the $10,600 "kitchen computer"—a Honeywell H316 minicomputer souped up in a futuristic, bright-red fiberglass cabinet, supposedly to increase its appeal to women. Honeywell built only a prototype. Neiman Marcus sold none of the units, but falling hardware costs would soon allow computers to penetrate low-cost markets, including the home. Honeywell advertised its H316 minus the glitzy cabinet as the first mini-

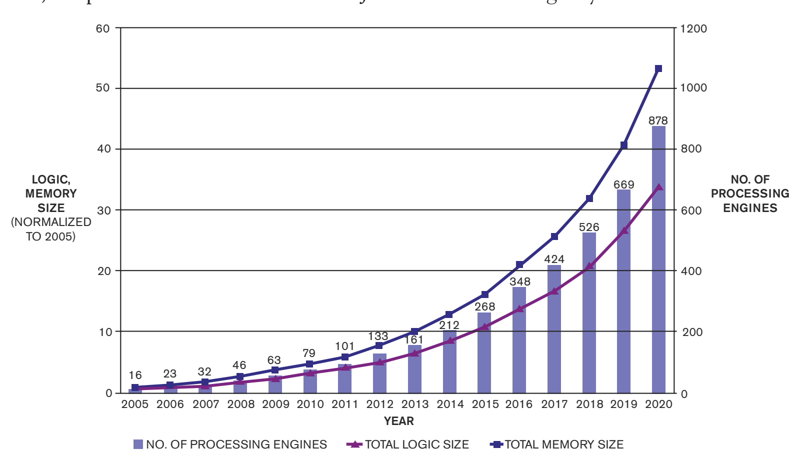

Figure 1 The ITRS (International Technology Roadmap for Semiconductors) predicts that future SOCs will incorporate dozens and then hundreds of processors.

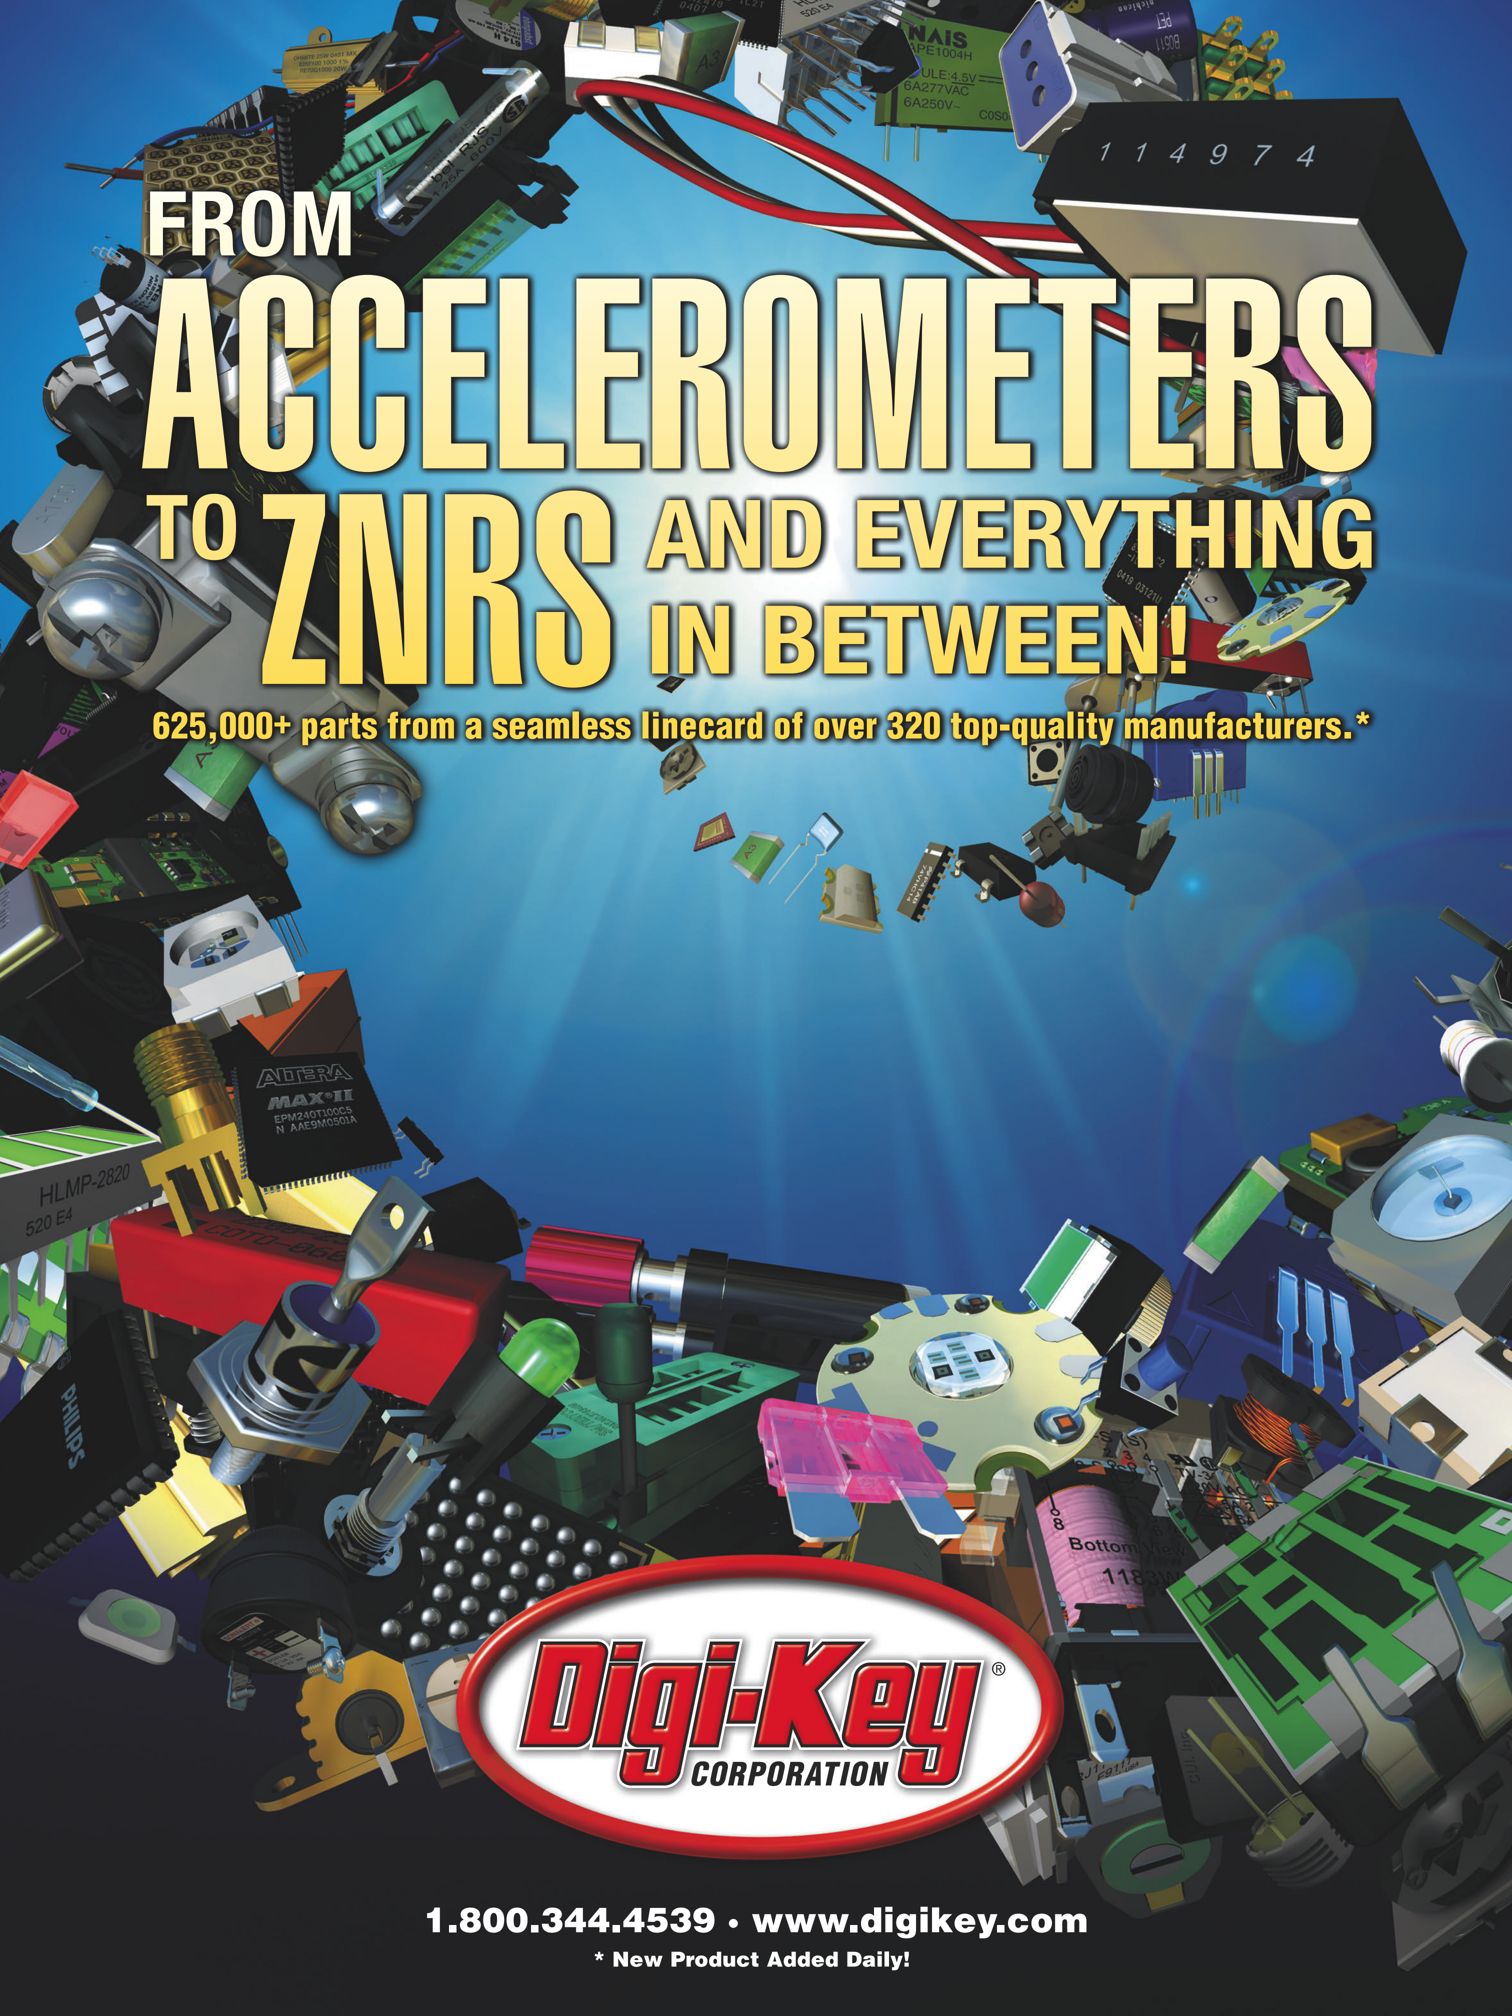
FROM
ACCELEROMETERS
TO ZNRS AND EVERYTHING IN BETWEEN!
625,000+ parts from a seamless linecard of over 320 top-quality manufacturers.*
Digi-Key®
CORPORATION
1.800.344.4539 • www.digikey.com
* New Product Added Daily!

computer to cost less than $10,000. For computers to truly proliferate, industry needed another breakthrough.

In 1969, the same year that the kitchen computer appeared, Japanese calculator manufacturer Busicom sought a semiconductor company to fabricate chips for a family of calculator-based products. Busicom's Osaka and Tokyo factories separately approached Mostek and fledgling Intel Corp (www.intel.com) because only these vendors had high-density silicon-gate-MOS processes. The Tokyo factory struck a deal with Intel for the more complex calculator-chip set. The chip set, which Busicom's Masatoshi Shima designed, was based on ROM-driven decimal state machines. Intel President Bob Noyce assigned the job of Busicom liaison to Ted Hoff, the company's application-research manager (**references** 2 and 3).

Busicom's engineers planned several chips, each to come in expensive (for the time) 40-pin DIPs. Once he fully understood Busicom's plan, Hoff knew Intel couldn't deliver the chips at the agreed price. The logic required large chips, and the 40-pin packages were expensive. Hoff alerted Noyce and suggested an alternative: a CPU-like logic chip with ROM-based software. Noyce, a mathematician and device physicist, didn't understand how a computer on a chip running software could replace logic, but he got the drift and encouraged Hoff to pursue the concept. Busi-

"SOFTWARE-IZATION"

By James Truchard, PhD, National Instruments

"Software-ization" may be a made-up word, but it accurately describes the trend in design. Today, hardware is what is left over when you finish with the software. With the growing popularity of reprogrammable hardware, such as FPGAs, and software techniques, such as digital-filter design and modulation, the software logic is becoming the most important aspect of a design, making the world increasingly "software-ized."

The increased focus on software design is evident in the design teams of today. According to Venture Development Corp (www.vdc-corp.com), the average software-design team now numbers 3.9 members versus 2.3 for hardware-design teams (Reference A). The trend toward software-ization does not end there: The need for programmable hardware is evident in the increasing number of programmable devices, such as FPGAs (Reference B). The need is clear: Designers want to reprogram their applications, which makes economic and efficiency sense. Managers prefer this approach because it means that you get more out of the hardware you purchase because you can reprogram it as specifications change, efficiently develop prototypes, and add features to deployed systems.

Another important angle emphasizing the need for software-ization is related to the "long-tail" business model (Figure A). The long tail describes how future business focus should be on "odd jobs" or millions of niche markets versus a few high-volume vertical markets. Examining vertical markets, such as cellular-telephone or plasma-television designs, you find plenty of reference designs specific to those markets. Those working in the long tail on applications from laser wrinkle removers to underwater surface crawlers have no access to reference designs; nevertheless, the expectations of products delivered in the long tail do not diminish. You still have the same pressure to deliver high quality and innovation but without the tools or reference designs available to help. At this point, software-ization couples with an integrated platform to help design engineers. Software-ization dynamically adjusts to the changing world of odd jobs. You need an integrated, off-the-shelf platform, so that when you implement the design in hardware, you need not hire any more hardware engineers (Figure B).

Another benefit of using software-ization with off-the-shelf hardware platforms becomes evident when examining test challenges with products in lower volumes. Tests on high-volume cell phones are thorough and costly, but high sales volume offsets this time and expense. However, with lower volumes of, say, 10 to 100 annually, it no longer makes economic sense to perform corner-case

continued on pg 146

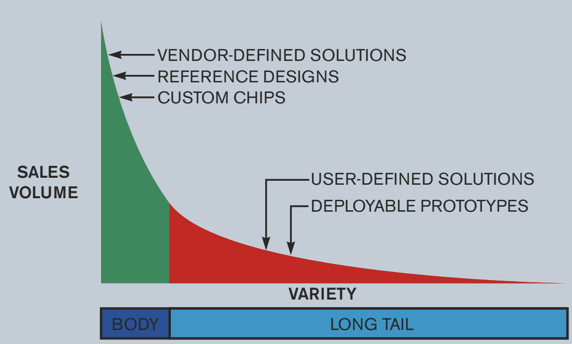

Figure A The long tail describes how future business focus should be on "odd jobs" or millions of niche markets versus a few high-volume vertical markets.

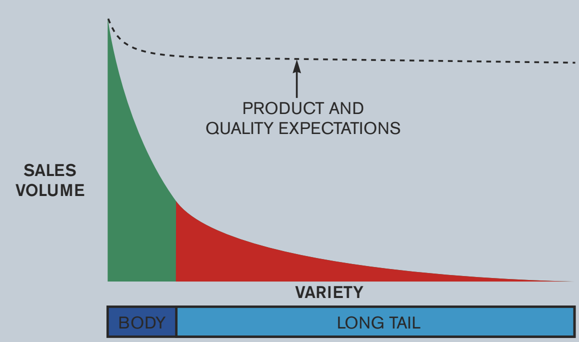

Figure B You need an integrated, off-the-shelf platform, so that when you implement the design in hardware, you need not hire any more hardware engineers.


AUTOMATICALLY-GENERATED CODE FLYING AT MACH 9.8.
THAT'S MODEL-BASED DESIGN.
When NASA made history by launching the X-43A, automatically-generated flight code was at the controls for the vehicle's propulsion and stability systems. Engineers developed the autopilot within a radically reduced timeframe using Model-Based Design and Simulink. To learn more, go to mathworks.com/mbd
MATLAB® & SIMULINK®
The MathWorks
Accelerating the pace of engineering and science
©2006 THE MATHWORKS, INC.

com eventually agreed to this approach.

Hoff, Shima, and Stan Mazor, newly arrived from Fairchild, developed a four-chip set—the CPU, a 320-bit RAM, a 256-byte masked ROM, and a 10-bit shift register for I/O—linked by a multiplexed 4-bit bus. (The RAM and ROM chips included 4-bit I/O ports.) This approach reduced IC-die cost and package size, sharply cutting the cost of the chip set. Intel packaged the RAM and ROM in 16-pin plastic DIPs and the CPU in a 16-pin ceramic DIP. These plastic and ceramic packages cost 18 cents and less than $1, respectively. Intel snagged Federico Faggin from Fairchild in March 1970 to translate the design into silicon. With Shima's help, he'd developed working silicon for the four chips by January 1971—just nine months later.

Busicom, hurt by falling calculator prices, renegotiated chip prices during 1971. Noyce exchanged a price cut for the right to sell the chips into noncompeting applications, and Busicom's CPU debuted as Intel's 4004 microprocessor at the Fall Joint Computer Conference in November 1971. Since then, system designers have incorporated billions of microprocessors, microcontrollers, and DSPs from dozens of vendors into many applications.

RISE OF THE COMPILERS

The microprocessor's debut allowed engineers to exploit software's problem-solving abilities in a vast array of electronic products. At first, machine code and assembly language were the only microprocessor-programming alterna-

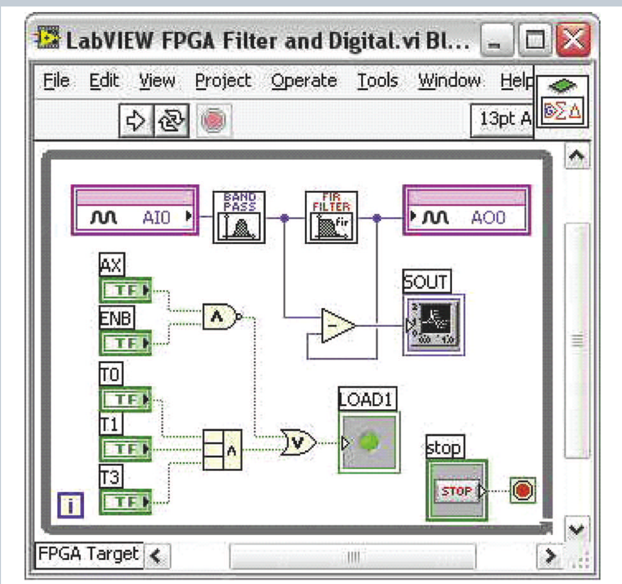

Figure C You implement a software-defined filter, such as National Instruments' LabView, in hardware.

continued from pg 144

testing, but you still need a high-quality product. The more efficient approach is reuse. Tested, off-the-shelf hardware platforms often come with industrial certifications, including international safety, EMC (electromagnetic-compatibility), shock, vibration, and environmental ratings. So, you can reuse this tested equipment and incorporate certified products into your design process for a much lower cost.

What does software-ization look like, and what does it mean to have software-program hardware? Software to program hardware must include syntax and semantics with explicit notations for expressing time and concurrency–primary attributes of hardware. Traditional, sequential programming languages do not exhibit these behaviors, making these applications difficult to develop.

Using the right software tool to program hardware is happening in many applications, including custom-test digital protocols with FPGAs; communications logic running in a DSP; and remote updating of deployed, embedded systems on a real-time microprocessor.

One of the most common design tasks–digital-filter modeling and design–takes advantage of software-ization. The use of digital filters eliminates a number of problems that their analog counterparts face. In software, you can attenuate unwanted signal elements such as noise caused by electrical components and environmental effects, apply antialiasing algorithms to test data, and reduce sample sets with decimation. Digital filters find use in a variety of applications ranging from machine-condition monitoring and animal-vocalization detection to seismic signal decomposition and audio special effects (Figure C). Once you design your filter with an integrated hardware platform, you can implement the filter without knowing the hardware details–another example of the beauty of software-ization.

Software-ization is a good thing for everyone. It means more manageable projects, lower hardware costs due to more reuse, and an easier path to more innovation due to the inherent ability to more quickly experiment, iterate, and implement.

REFERENCES

A Lanfear, Chris, "Embedded software strategic market intelligence program, 2005, Volume 8: Embedded Systems Market Statistics," Venture Development Corp, www.vdc-corp.com/Purchase.asp?viewtype=detail&id=1823.

B Bartos, Frank, "ASICs Versus FPGAs," *Control Engineering*, June 2005, www.manufacturing.net/ctl/article/CA607224.

C Anderson, Chris, "The Long Tail," *Wired* Magazine, October 2004, www.wired.com/wired/archive/12.10/tail.html.

AUTHOR'S BIOGRAPHY

James Truchard, PhD, is co-founder, president, and chief executive officer of National Instruments.


90% CUT IN EMISSIONS
AND 50% CUT IN DEVELOPMENT TIME.
THAT'S MODEL-BASED DESIGN.
To meet a tough performance target, the engineering team at Nissan used dynamic system models instead of paper specifications. The result: 50% time savings, the first car certified to meet California's Partial Zero Emissions Vehicle standard, and a U.S. EPA award. To learn more, visit mathworks.com/mbd
MATLAB® &SIMULINK®
The MathWorks
Accelerating the pace of engineering and science
©2005 The MathWorks, Inc.

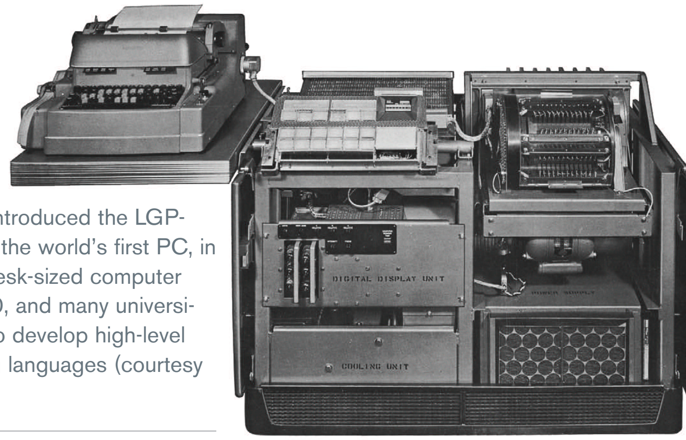

Librascope introduced the LGP-30, arguably the world's first PC, in 1956. The desk-sized computer cost $27,000, and many universities used it to develop high-level programming languages (courtesy Bob Lilley).

tives. Early HLL compilers for microprocessors appeared, but they generated slow, memory-hungry object code rather than well-written, handcrafted assembly code. Early microprocessor-based systems were already slow and memory-starved, so early compilers were nearly useless, but they evolved.

The same year that Intel introduced the 4004 microprocessor, Dennis Ritchie started extending Bell Labs' B programming language for the PDP-7 minicomputer by adding a character type (**Reference 4**). Ritchie called this new language NB (new B). NB was an embryonic version of C, and the language had become recognizable as C by 1973. Ritchie continued to work on C and published *The C Programming Language* with Brian Kernhigan in 1978 (**Reference 5**). This book served as a de facto C standard throughout the 1980s. The embedded-development world dabbled with many HLLs, including C, Basic, Forth, Pascal, and a microprocessor version of PL/I called PL/M, but superior performance and memory efficiency kept assembly-language coding in the lead.

Throughout the 1980s, embedded systems became increasingly complex, compilers improved, and the cost of 32-bit microprocessors decreased. Rising software costs and complexity eventually curtailed assembly-language programming on microprocessors, mirroring events that spurred Backus to develop Fortran for mainframes. C's popularity as an embedded programming language climbed. ANSI established the X3J11 committee to produce a C standard in the summer of 1983. When C became both an ANSI and an ISO standard in 1990, C became king.

By 1990, microprocessor-centric design was the first choice of board-level-system designers. Logic design became a last resort, which designers used only when microprocessors were too slow. Around 1995, microprocessors became cores—just part of an IC—transforming ASICs into SOCs (systems on chips). Early SOCs incorporated only one microprocessor; then, two; then, many. The ITRS (International Technology Roadmap for Semiconductors) predicts that this trend will continue (**Figure 1**). Like board-level-system design before it, SOC design became processor-centric (**Reference 6**). Many SOCs today incorporate dozens or hundreds of interconnected processors. Microprocessors and software are the very fabric of contemporary electronic design. That fabric now covers the planet, as many embedded-microprocessor applications demonstrate (**Table 1**). In fact, processors and software now reach beyond the earth.

The first microprocessor in space, an

TABLE 1 EVERYDAY USES FOR EMBEDDED MICROPROCESSORS AND SOFTWARE

Office and retail	**Home**
Telephones and PBXs	Conventional and microwave ovens
Printers	Food processors, mixers, and blenders
Copiers and faxes	Refrigerators and dishwashers
Postal scales and shipping management	Climate and lighting control
Fire and intrusion alarms	TVs, cable and satellite boxes, VCRs, and DVRs
Lighting and HVAC controls	Home-entertainment systems
Elevator and automatic-door controls	DVD, CD, and MP3 player/recorders
Bar-code and RFID readers	Clothes washers, dryers, and irons
Video security and monitoring	Corded and mobile telephones
Energy and utilities monitoring and billing	Toys and games
Record keeping and management	Digital cameras and camcorders
Inventory management	Barbecue grills
Civil ground transportation	**Space, aviation, Naval, and military**
Drive-train control	Engine control and management
Passenger-cabin climate control	Guidance and attitude control
Entertainment systems	Onboard systems monitoring
GPS and compass navigation	Radar and collision avoidance
Collision-avoidance systems	Global and celestial navigation
Mobile-phone and satellite communications	Radio and satellite communications
Fare collection (public transport)	Passenger-cabin climate control
Traction control and automatic braking	Fuel management
Active shock absorbers	Weapons management and control
Manufacturing	**Medical**
Process control	Diagnostic, imaging, and treatment systems
Robotic assembly and transport	Patient record keeping
Inventory management and tracking	Robotic surgery
Energy and load management	Pharmaceutical dispensing and inventory control
Security, safety, and access management	Therapeutic and rehabilitation systems
The Internet	
Routers and switches	
Network interfaces and bridges	
Cable/DSL modems and gateways	
Web, mail, search, and storage servers	
Firewalls	
Wireless access points and repeaters	

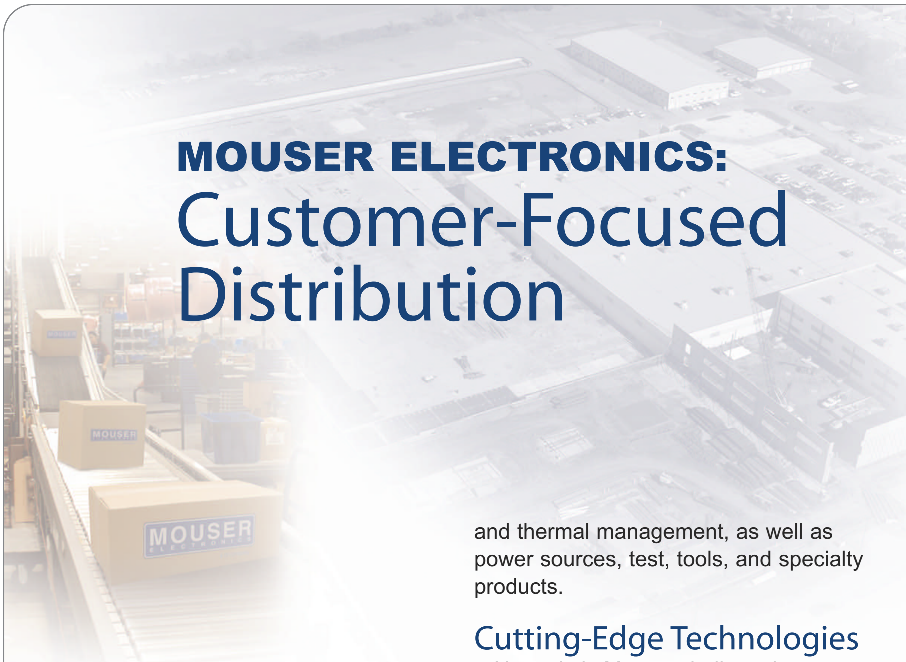

MOUSER ELECTRONICS: Customer-Focused Distribution

Mouser Electronics is the fastest growing catalog distributor in the electronics industry and is dedicated to supplying customers with the NEWEST products and the LATEST technologies. Known for its rapid new product introduction, Mouser delivers a ***time-to-market advantage*** with daily updates to its online catalog and the only NEW 1,800+ page print catalog every 90 days!

Corporate Mission

Mouser's mission is to be ***the preferred source*** for electronic components. Diligent pursuit of customer-focused distribution by delivering customers a ***time-to-market advantage*** is a common thread that runs through every facet of company operations.

The number one priority is making it easy to do business with Mouser! The company's business model of continuously offering customers the most innovative products, the latest technologies, and the best customer service available in the industry, has proven successful.

Broad-Based Product Line

Mouser's explosive growth is due in part to its broad-based line card of electronic components from over 325 premiere industry suppliers of semiconductors, optoelectronics, embedded modules, passives, circuit protection, interconnects, wire and cable, electromechanical, sensors, enclosures and thermal management, as well as power sources, test, tools, and specialty products.

Cutting-Edge Technologies

Not only is Mouser dedicated to providing its customers a vast array of new products, but special focus is placed on suppliers whose products are based on the latest cutting-edge technologies.

Customers choose Mouser because the company provides the newest products and the latest technologies from the leading manufacturers in the industry. In the last year alone, hundreds of leading-edge technology products have been added to the linecard, including ZigBee™, Ethernet, Bluetooth, RF, Embedded, Wireless, and any other new technology on the market.

On-line Browse & Buy

Continuous improvements and new products added to Mouser's website make it easier for customers to browse and buy online. Daily updates allow customers to access over 700,000 products, view more than 490,000 downloadable data sheets, 330 supplier catalogs, and over 1.5 million cross-referenced parts. The constant refinement to the site includes the addition of numerous user-friendly tools, such as automatic order confirmations, Project Manager with auto reorder, and Bill of Materials import tools.

Unique Customer Base

Mouser is considered the small order logistics expert in the distribution industry. Focusing primarily on the design-in and prototyping stages of the design cycle, customers represent a wide range of small, medium, and large companies, as well as purchasing agents and individual consultants who recommend, specify, and purchase board level components for manufacturers' product designs. Customers who require immediate availability of low-volume, off-the-shelf components benefit from Mouser's broad-based product lines. Suppliers seek out Mouser's unique customer base because other distributors do not have the same customer penetration.

Unsurpassed Customer Service

The company believes in providing its customers the very best possible customer service — regardless of the size of the customer or the size of the order. In fact, breaking packs and no minimum order size is especially attractive to engineers working in the early design and prototyping stages of the design cycle. Mouser's excellent technical personnel and well-seasoned sales staff are proficient in assisting customers solve their design challenges.

Expanded Warehouse Operations

Due to explosive growth, Mouser has nearly completed the expansion of its corporate headquarters from 173,000 sq.ft. to 432,000 sq. ft. The company is expanding both office and warehouse space, as well as upgrading numerous systems. Mouser's integrated state-of-the-art wireless and automated warehouse management system has streamlined its pick-and-ship operations to near perfection. These high-tech improvements and the efficient order fulfillment process allow same-day shipping on most orders received by 8:00 p.m. Central Standard Time. Customer-focused distribution makes it easy to do business with Mouser.

MOUSER ELECTRONICS®

a tti company

Newest Products for Your New Designs

mouser.com

(800) 346-6873

THE CASE FOR APPLICATION-SPECIFIC PROGRAMMABLE LOGIC

By Alex Lidow, International Rectifier

The "software-ization" of an enormous range of applications has brought more features and lower product costs to end users and shorter development times and lower product-lifetime-management costs to manufacturers. By and large, these applications exploit the processing power of microprocessors, microcontrollers, and DSPs. They depend on the availability of off-the-shelf programmable logic that economically provides greater processing power than the application's peak processing requirements demand.

Real-time applications as disparate as motion control, communications, and image processing test the limits of software implementations running on general-purpose programmable hardware. In these applications, complex processing requirements combine with high data rates or signal bandwidths to drive hardware usage and code complexity to uneconomic extremes.

In motion control, for example, high-speed control loops depend on real-time vector-transform calculations to determine power-switching times within a rotating frame of reference. A second, inverse, set of transforms operates on the motor's feedback signal. The complexity of these calculations exceeds the processing bandwidth of low-cost processors based on traditional architectures. The "bigger hammer" approach–using more powerful general-purpose processors fabricated on faster semiconductor processes–doesn't sufficiently improve computational efficiency and can drive both silicon and software-development costs to noncommercial levels.

Still, the allure of programmable loops and the accompanying benefits to OEMs of reduced development cost and time savings merits the rethinking necessary to achieve high bandwidth and low cost. Though traditional microprocessor, microcontroller, and DSP architectures may not be best for these applications, custom-hardware implementations often poorly support rapid product-development cycles or easy design maintenance.

As high-speed communications and image processing have done before them, motion-control applications are taking advantage of computational engines–application-specific logic–that serve as a programmable interface between an application-management layer and a real-time process. Computational engines offer a middle ground between software-ization and fixed-topology hardware that benefits OEM designers by reducing both development time and project risk in these computationally intensive real-time applications.

At the real-time-process interface, computational engines efficiently support high-bandwidth control loops through application-specific hardware peripherals, low-level firmware, and vendor-provided algorithms. OEM-configurable computational resources allow product developers to customize the engine's resources in ways that parallel the high-level coding of general-purpose programmable logic. In the case of the engine-based implementation, however, the development environment can support programming in application-relevant terms.

Visual-programming environments can go a step further, allowing product developers to specify signal flows between computational blocks without coding per se. Such high-level constructs do not preclude OEMs from customizing the lower level operating algorithms or developing new algorithms from scratch, but they do provide for significantly shorter development and substantially improved code reuse. This balance between hardware and software resources promotes rapid prototyping, quick development of derivative products, and low-cost design maintenance.

On the application-control end of the design, the engine can present an intelligent interface to either a co-integrated or co-embedded microcontroller. The small microcontroller can efficiently manage the user interface–panel switches and displays–and direct the outermost application-control loop. The advantage of this functional segmentation between the microcontroller and the real-time engine is that the interface between them can use physical quantities instead of coded coefficients. Again using the motion-control example, this interface can express shaft speeds in rpm or torque limits in newton-meters. This layer of abstraction separates the application-control development from all of the details of managing the motor in real time.

Integrated design platforms, comprising the real-time engine, vendor-provided routines, hardware peripherals, and development environment, capture highly specialized hardware-engineering expertise. They also allow the OEM to build upon the resource by customizing or developing in-house functional and algorithm libraries. Most important, this approach reduces the OEM's need to reinvent the technological wheel and instead promote a focus on those design areas that best differentiate the product.

AUTHOR'S BIOGRAPHY

Alex Lidow is chief executive officer of International Rectifier.

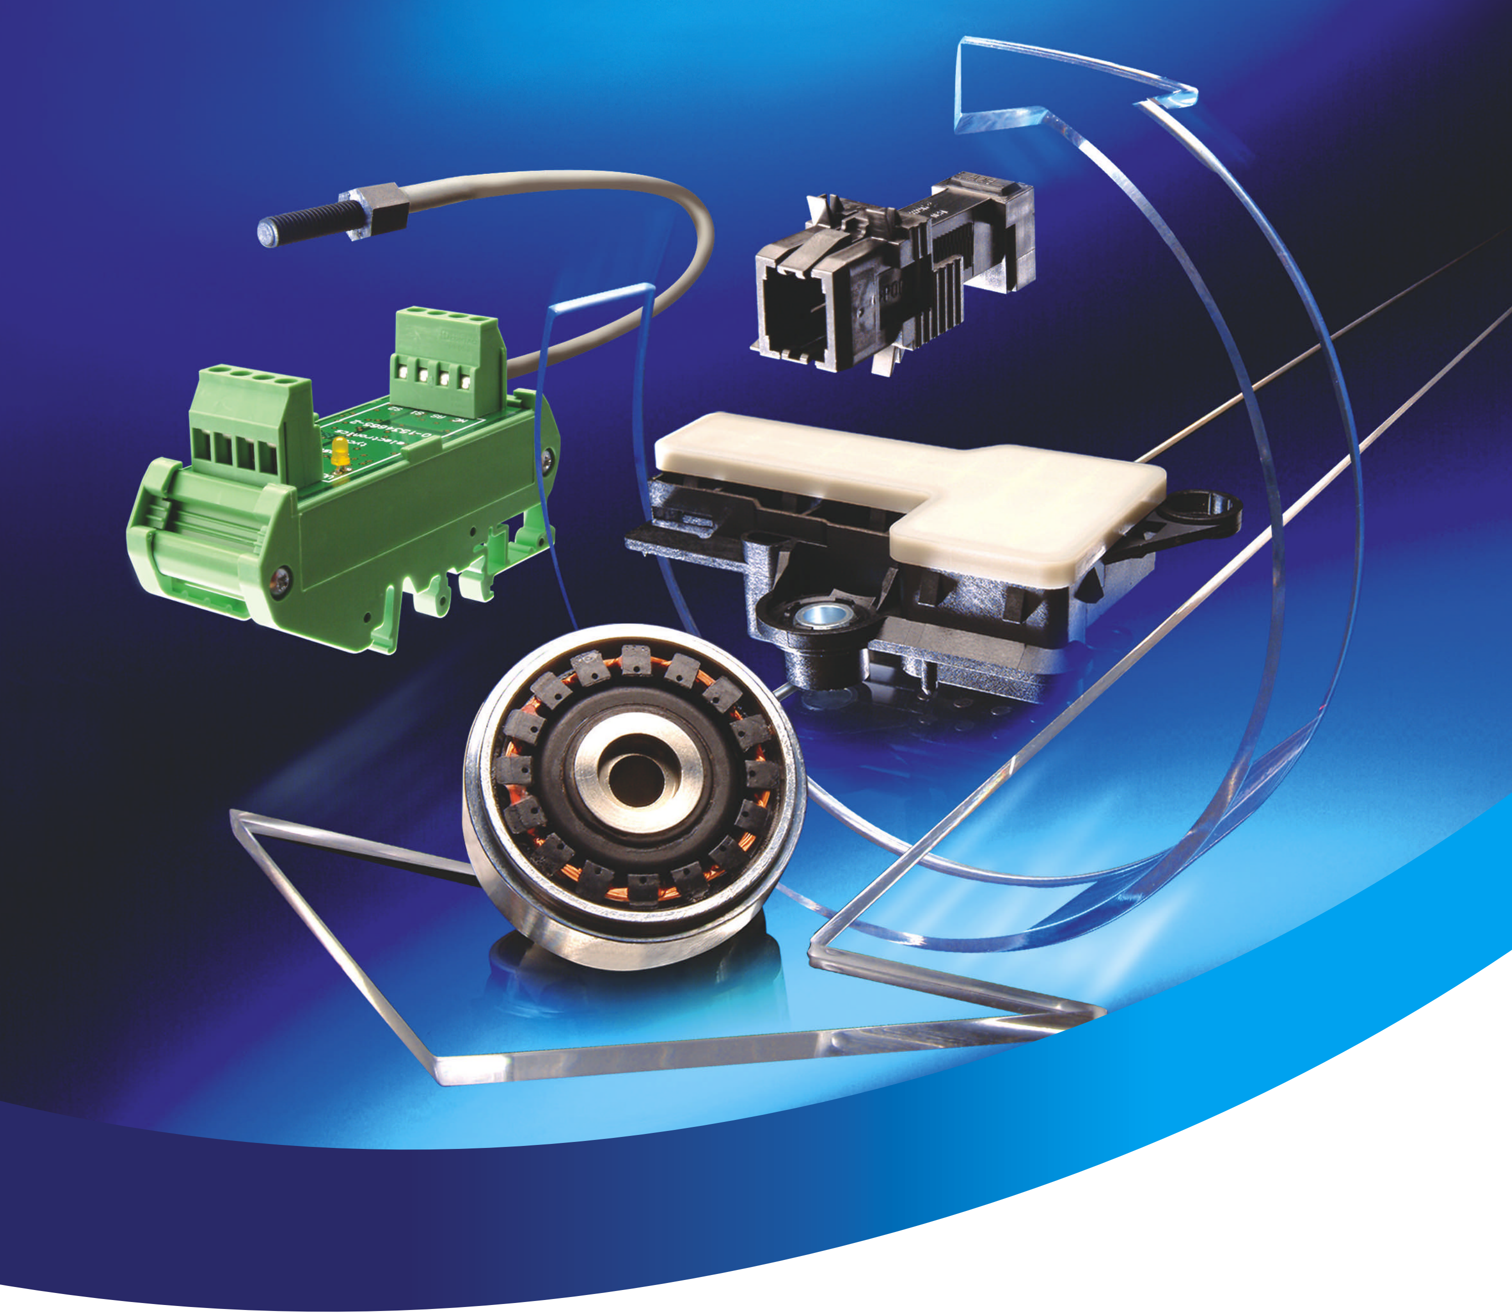

Neiman Marcus unveiled the Honeywell "kitchen computer," which it based on a stock H316 minicomputer, in Neiman's 1969 Christmas catalog. Linux kernel developer Val Henson re-created the original model's catalog pose in this photo taken at the Computer History Museum in Mountain View, CA. Although the $10,000 computer failed as a product because it was ill-conceived and far too expensive, it predicted the successful and pervasive penetration of processors and software into many home products (courtesy Val Henson, more at http://infohost.nmt.edu/~val/kitchen.html).

Intel introduced its 4004 microprocessor—a calculator chip originally designed for Busicom—with this ad in the November 15, 1971, issue of *Electronic News*. The lower right corner includes an invitation to readers to see the new CPU on a chip in the company's hospitality suite at the Fall Joint Computer Conference in Las Vegas, held later that month. This ad, minus the invitation, ran in *EDN* the following month.

RCA Cosmac (CDP1802), launched into polar orbit 30 years ago on Sept 11, 1976, aboard the DMSP's (Defense Meteorological Satellite Program) 5D-1 F1 satellite (**Reference 7**). A high-pressure nitrogen-supply-line leak caused the spacecraft to tumble. The leak depleted the satellite's thrusters and saturated its reaction-wheel attitude-control system. The satellite lost all attitude control. Uploading new on-orbit software allowed direct ground control of the satellite's onboard magnetic torquing coils, which eventually stabilized the satellite and saved the mission. A similar upload fixed another attitude problem on the next satellite in the series, the DMSP 5D-1 F2.

Spacecraft designers favored RCA's Cosmac because it was available in radiation-hardened CMOS on SOS (silicon on sapphire). Many other spacecraft, including Viking, Galileo, the US space shuttle, and the two Voyager interplanetary probes, employed Cosmacs. When NASA launched them in 1977, the Voyagers each contained three Cosmacs, which have far outlived RCA Semiconductor. The two Voyagers have taken their microprocessors and software several billion miles from earth, beyond the solar system, and into interstellar space where they continue to gather and transmit data three decades after launch. EDN

REFERENCES

1 Leibson, Steve, *Stanley P. Frankel: Unrecognized Genius*, www.hp9825.com/html/stan_frankel.html.

2 Aspray, William, "Masatoshi Shima, An Interview," IEEE History Center, May 17, 1994.

3 Berlin, Leslie, *The Man Behind the Microchip: Robert Noyce and the Invention of Silicon Valley*, Oxford University Press, 2005.

4 Ritchie, Dennis, "The Development of the C Language," http://cm.bell-labs.com/cm/cs/who/dmr/chist.html.

5 Kernhigan, Brian W and Dennis M Ritchie, *The C Programming Language*, ISBN 0-13-110370-9 and 0-13-110362-8, Prentice Hall, 1978.

6 Leibson, Steve, *Designing SOCs with Configured Cores: Unleashing the Tensilica Xtensa and Diamond Cores*, Morgan Kaufmann, 2006.

7 Strom, Steven R and George Iwanaga, "Overview and History of the Defense Meteorological Satellite Program," *Crosslink*, The Aerospace Corp, Winter 2005, www.aero.org/publications/crosslink/winter2005/02.html.

AUTHOR'S BIOGRAPHY

Steve Leibson spent the first 10 years of his career as an engineer, then spent 15 years as an award-winning journalist, publishing more than 200 articles and serving as editor in chief of both EDN *(1993 to 1996) and* Microprocessor Report *(2000 to 2001). Now he is the technology evangelist for Tensilica.*

A DMM that works like you do

in the 21st century.

26% off through 30 November 2006.

Agilent 34410A digital multimeter

- *>150x faster reading rate at 6.5 digits**
- *Improved accuracy**
- *Open I/O with LAN (LXI Class C compliant), USB, and GPIB*
- *100x more reading storage**
- *Dual display*

**Compared to Agilent's industry leading 34401A digit multimeter.*

Buy an Agilent 34410A for the price of an Agilent 34401A

The Agilent 34410A digital multimeter is built for the 21st century. Take readings at 10,000 per second. Log data automatically and with more reading storage. Detect peak inputs as short as 20 microseconds. View two measurements at once via the dual display. Take advantage of expanded measurement ranges. Even measure temperature or capacitance.

Based on years of working closely with our customers, this new DMM is pure Agilent; designed to work the way you do. And because it is Agilent, the 34410A is the only DMM with a fully open configuration, so you can easily integrate it into your workflow, regardless of I/O connectivity to your PC.

Take 26% off your purchase of a 34410A
Go to www.agilent.com/find/dmmpromo. Offer ends 30 November 2006.

This offer is only available from Authorized Agilent Distributors.
Go to the website below to locate an approved distributor.
www.agilent.com/find/DMMpromo

Agilent Technologies

Optimizing system power.

Driving energy efficiency. Every moment. Everywhere.

In every corner of the globe, regardless of the day or time, Fairchild's products drive energy efficiency. Fairchild optimizes power in a world of electronic products that help make our lives better, from lightbulbs, plasma screens, washing machines and PDAs to the car you drive. Whether you're in Singapore or San Francisco, Fairchild touches your life at every level, every second of the day, with power analog, power discrete and standard products designed to optimize power in almost every electronic system imaginable. To learn more, visit **www.fairchildsemi.com**

9:30 am central time
Chicago, Illinois

Fairchild parts*inside* washing machine
Diode
Fairchild Power Switch (FPS™)
Insulated Gate Bipolar Transistor (IGBT)
Light Emitting Diode (LED)
Metal Oxide Semiconductor
Field-Effect Transistor (MOSFET)
Motor Driver - SPM™
Optocoupler
Triac

7:30 am pacific time
San Francisco, California

Fairchild parts*inside* LCD television
Audio Amplifier
Bridge Rectifier
Fairchild Power Switch (FPS™)
Inverter Controller
Load Switch
Low Dropout Regulator (LDO)
Power Factor Correction (PFC)
Quantum Field-Effect Transistor (QFET) MOSFET
Stealth Diode
Video Filter

3:30 pm standard time
London, England

Fairchild parts*inside* digital camera
Analog Switch
DC-DC Converter
Insulated Gate Bipolar Transistor (IGBT)
Light Emitting Diode (LED) Driver
Load Switch
- IntelliMAX™
Logic
µSerDes™
Serializer/Deserializer

10:30 am eastern time
New York, New York

Fairchild parts*inside* laptop
DC-DC Controller
Light Emitting Diode (LED) Driver
Linear Regulator
Low Drop Out Regulator (LDO)
Metal Oxide Semiconductor
Field-Effect Transistor (MOSFET)
Pulse Width Modulation (PWM)
Controller
Schottky Diode
Zener Diode

4:30 pm standard time
Paris, France

Fairchild parts*inside* cell phone
Analog Switch
Audio Amplifier
DC-DC Converter
Light Emitting Diode (LED) Driver
Load Switch
- IntelliMAX™
µSerDes™
RF Power Amplifier Module (PAM)
Serializer/Deserializer
Video Filter

11:30 pm standard time
Hong Kong, China

Fairchild parts*inside* plasma TV
Bridge Rectifier
Diode
Fairchild Power Switch (FPS™)
High Voltage IC (HVIC)
Insulated Gate Bipolar Transistor (IGBT)
Optocoupler
Power Factor Correction (PFC)
Pulse Width Modulation (PWM)
Controller
Quantum Field-Effect Transistor (QFET)
Video Filter
Voltage Regulator

11:30 pm standard time
Singapore

Fairchild parts*inside* CFL bulbs
Ballast Control IC
Bipolar Transistor
Bridge Rectifier
Quantum Field-Effect Transistor
(QFET) MOSFET

Jianhong Ju
Technical Marketing Manager

greennote

Power used by a mobile phone/PDA on standby: **1.84 watts**
Power used by a mobile phone/PDA on standby with Fairchild's Green FPS components: **0.29 watts**
Savings with Fairchild components: **1.55 watts**

If every mobile phone sold worldwide in 2005 was on standby 50% of the time and realized this 1.55 watt savings, enough power would be saved on an annual basis to supply power to approximately 500,000 households in the United States.

Power used annually by an average washing machine: **698 kilowatt-hours**
Power used annually by the same washing machine using Fairchild components: **565 kilowatt-hours**
Savings with Fairchild components: **133 kilowatt-hours a year**

If every one of the 70 million washing machines produced globally in 2005 used Fairchild energy efficient products, the savings would be equal to the approximate amount of power generated by the Hoover Dam in over two years. This is no small feat. The Hoover Dam generates more than 4 billion kilowatt-hours a year, enough to serve the electrical needs of 1.3 million people.

Power used by a standard incandescent light bulb: **60 watts**
Power used by a fluorescent (CFL) bulb with Fairchild components: **15 watts**
Savings per bulb with Fairchild components: **45 watts**

If each of the 923,000 households in Singapore used an average of five 60 watt lightbulbs and switched from incandescent bulbs to CFL bulbs that include Fairchild components, the city would save enough power, just from the bulbs, to light more than 69,225 homes!

the power® franchise

www.fairchildsemi.com

FAIRCHILD SEMICONDUCTOR®

ROB MAGIERA / NOUMENA DIGITAL

BY JOSHUA ISRAELSOHN, CONTRIBUTING TECHNICAL EDITOR

THE THERMAL COST OF PERFORMANCE

ELECTRIC-ENERGY EFFICIENCY SERVES AS ONE MEASURE OF HOW FAR THE ELECTRONICS INDUSTRY HAS COME. EXPLORE HOW LIGHTING, MEASUREMENT INSTRUMENTATION, AND AUDIO AMPLIFICATION HIGHLIGHT THE THERMAL CHALLENGES THAT ENGINEERS FACE.

Few aspects of the electronics industry offer milestones that mark the entirety of *EDN*'s half-century of publication. For example, with the exception of only a few limited niches, we no longer use thermionic vacuum tubes (glass FETs, for the modernists) nor are Bakelite boxes now much in vogue. On the other end of the historic interval, though the digital abstraction that dominates our modern design practice was known at the time, its practical application was only barely evident in 1956. Indeed, circuit architectures that are now as common as hands simply could not have been imagined five decades ago. As for the means of their physical implementation, what now conveniently fits into those common hands could not then have been realized on the footprint of a typical house ... if at all.

One of the few themes that does connect the dots from the time of this magazine's comparatively ancient beginnings to the current day is our use of electrical energy. In particular, our electric-energy efficiency—largely a measure of how much of the stuff we must convert into heat in the process of completing a useful task—serves as one measure of how far we've come as an industry.

A proper recounting of our industry's progress in this regard over the half-century span of *EDN*'s existence would result in a book-length work; a summary would occupy the whole of the current issue. Instead, let's take a glimpse into that progress. Analog Devices' Barrie Gilbert offers three applications—lighting, measurement instrumentation, and audio amplification—to demonstrate some of the challenges practitioners of electronics design must face.

As with any complex discipline, electronics designers build on what precedes them. That task, however, is not unidisciplinary. On the contrary, it requires an understanding of materials, processes, device behaviors, topological idioms, and system structures as they pertain to parametric performance.

Paradoxically, perhaps, significant improvement over prior art sometimes requires a dramatic departure from the foundation practice that calls into question our assumptions, habits, and design prejudices. As Gilbert points out, after a century of hot-wires in glass bottles, they are still the dominant interior light source. The contender that holds the greatest promise as a replacement technology looks nothing like the object it will replace, save the presence of a window and a pair of contacts for electrical connection. Just as unlikely and just as true is the fact that the replacement technology did not derive from lighting-device engineering but evolved from a path that begins with signal-processing and ends through a branch of materials science and III-Vs semiconductor processes.

Unfortunately, as a general approach to our business, waiting for a century of progress is not a winning strategy. Accelerating that progress requires a mindfulness of fundamental physics and a willingness to challenge the conventional conclusions built upon those axioms: The occasional heretical thought is good for the designer's soul. In keeping with our theme, Gilbert asks why there is a thermal cost—a use of electrical power—to perform functions at all and, in the asking, suggests that we consider those issues that set the minimum thermal cost of performance.

Lastly, for those of you who do not typically toil at IC design in high-speed processes, Gilbert gives insight into one challenge that SOI (silicon-on-insulator) semiconductor processes pose to the circuit designer: There is not only a thermal cost of performance, but also a performance cost of thermals—in this case, one of the fundamental underpinnings of traditional IC-design practice. EDN

50TH ANNIVERSARY ONLINE
THE CELEBRATION CONTINUES

For more on the thermal cost of performance, read Aengus Murray's article on the evolution in the motion-control sector, which has resulted in significant energy savings. Visit www.edn.com/50th.

AUTHOR'S BIOGRAPHY

Joshua Israelsohn is director, technical information at International Rectifier Corp. Formerly, he worked as EDN*'s analog editor, and he continues to contribute a regular column to* EDN*. You can reach him at jisrael1@irf.com.*

MINIMUM ENERGY AND POWER DEMANDS IN ANALOG ICs AND THE IMPACT OF SELF-HEATING EFFECTS IN TINY TRANSISTORS

By Barrie Gilbert, Analog Devices

You need a certain minimum energy to perform any practical operation. Raising a 12-oz can of beer from belly to lips consumes about 1J. The energy you need to execute a one-time function in electron-based systems (more familiarly stated as the power you need to repeat it—or sustain it continuously) is generally not so clear-cut. It comes down to a matter of where you stand along history's slender arrow, and the state of the art in the many relevant and intertwining technologies of an era.

In Edison's time, the minimum necessary temperature-rise of frail loops of tungsten wire in glass bottles dictated the "house current" power drain for the drawing-room chandelier. Because most of their output power fell into the infrared region, these glowing wires mainly warmed the ladies' wigs and gentlemen's bald pates. Today, we have a long list of devices for lighting homes and workplaces; they are rugged, durable, and efficient. Yet, 100 years later, the use of "hot-wires" still tops the list.

The direct conversion of an electron source to visible light—exploiting new properties of materials and devices—has made great advances in recent years. Ultimately, the electrical power necessary to generate a continuous photon-flux density limits the efficiency of these devices. No technology permits us to realize its theoretical potential; nonetheless, electron-to-photon efficiency in semiconductor emitters, such as recent white-light LEDs, is rapidly climbing toward its asymptote. These solid-state photon sources are poised to eclipse the wire in a bottle as certainly as transistors ousted vacuum tubes: tentatively at first, inevitably in the end.

But even the most ingenious technologies must also be cost-effective. "Sure" says the skeptic, "those new LED lamps are really bright, and they run stone cold! But are they as cheap as a six-pack of 60-watters from Wal-Mart? If they were to slash lifetime to one-tenth, would they cost one-tenth as much? I've heard that CPU manufacturers play that game. They choose either their 'three-year process,' whose narrow interconnects eventually fail as electromigration creates lateral filaments of metal that short to adjacent traces, or their 'seven-year process,' whose wider metal and spacing rules increase the lifetime of the product but at the cost of a general increase in the *inertia* of the interconnects and a larger die."

POWER TRADE-OFFS

Questions about the minimum energy necessary to perform a unit function (such as a single AND decision) or how much continuous power you must supply to a functional block for it to perform a certain repetitive function are among the most intriguing topics in looking toward the future of electronic signal processing. They are readily tractable in the domain of binary signaling. The energy needs in executing logical functions are usually couched as some voltage V_0 on nodal capacitance C, being CV_0^2. A gate output must swing to its only other value, V_1. For a capacitance of 50 fF to swing through 1V, an energy source must provide 50 fJ (femtojoules) at each rising or falling edge. When this operation repeats at a clock rate of 1 GHz (2×10^9 edges per sec), an average current of 50 fJ$\times$(2×10^9)/sec results; that is, each "action node" sips 100 μA of continu-

It's just this simple

or

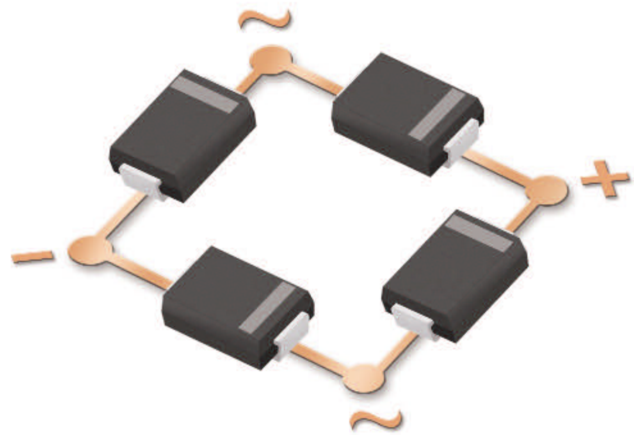

New
Tiny Schottky Bridge Rectifier

4 individual SMA Schottky Rectifiers in a bridge configuration

CBRHDSH1-40L and CBRHDSH2-40

Central's New High Density Schottky Bridge Rectifiers

HD DIP
Actual Size

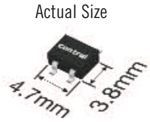

The new Central Semiconductor CBRHDSH1-40L (1A, 40V, Low V_F) and CBRHDSH2-40 (2A, 40V) are full wave glass passivated Schottky bridge rectifiers manufactured in a durable HD DIP surface mount package. Designed for applications requiring a smaller and more energy efficient alternative to a standard bridge rectifier, these new devices are ideal for today's latest electronic product designs. 60V and 100V devices are under development.

Free Samples

Samples available from stock.
To order, visit: www.centralsemi.com/sbr

RoHS Compliant

Pb-Free Plating

EED ENERGY EFFICIENT DEVICE

Typical Applications

- Voice over IP (VoIP)
- Power over Ethernet (PoE)
- Networking equipment
- Modems
- Laptops
- Data line protection
- Any circuit requiring a small energy efficient Schottky bridge rectifier

Features

- Low V_F (0.39V typ. for CBRHDSH1-40L)
- Pb Free and RoHS compliant
- HD DIP utilizes 50% less board space compared with 4 individual SMA devices

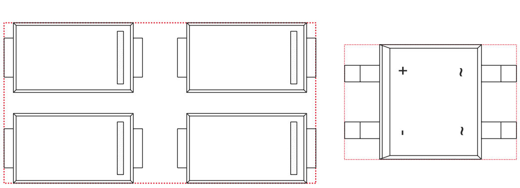

4 SMA packages HD DIP package

Central™ Semiconductor Corp.
www.centralsemi.com

Small Signal Transistors | Bipolar Power Transistors | Energy Efficient Devices | MDMs
Schottky Bridge Rectifiers | Rectifiers | Thyristors | Transient Voltage Suppressors | TLMs

145 Adams Avenue, Hauppauge, NY 11788 USA • Tel: (631) 435-1110 • Fax: (631) 435-1824

ous current. A million elements of this sort will happily drink 100A all day long.

In the analog-IC domain, people rarely ask questions about such issues as the minimum energy for a given function; when they do, the approach is often hopelessly academic, purely theoretical, totally out of touch with the practical world, and thus useless for most mere mortals. From an engineer's perspective, many important questions of this genre remain unanswered. However, it is also apparent that analog signals have vastly more variety and complexity of form. They receive support from deeply recursive meshes of plesiolinear elements, often deliberately using the specific nonlinearities of special elements. Each of these elements has desired or incidental inertia (energy-storage aspects) and boasts an imposingly lengthy list of parameter values. It is hardly surprising that, after an hour or two, minimum-energy considerations end up in one's recycling basket.

Analog-IC designers from 1960 to 1980, who predominantly based their designs on junction-isolated, bipolar-junction-transistor processes, judged and juggled many trade-offs, but, with a few obvious exceptions, power efficiency rarely concerned them. The emphasis was largely on maximizing performance until it was comfortably beyond the competitive limit. Power consumption was whatever you needed to meet those objectives. Few in the industry appreciated the growing importance of ICs that frugally used power. Most regarded *low-power design* as a sideshow, useful for providing thesis projects and interesting enough to justify an occasional specialist session at the ISSCC (International Solid-State Circuits Conference). Today, low-power design is at center stage.

Preconceptions about the power necessary to achieve Function X arise from the norms of present-day designs. For example: What is the minimum power a circuit must dissipate to measure a voltage applied to a probe tip? Reviewing prevalent instrumentation techniques, you might mentally list the power each major section consumes, starting with some sort of input range selector and buffer. Then you'd move on to the ADC—perhaps one of the old dual-slope variety, a charge-dispensing voltage-to-frequency converter, or a modern sigma-delta type—and its indispensable voltage-reference cell. Finally, you'd consider the matter of display elements—Nixie tubes; seven-segment, 30-ft-high Times Square illuminators; rolling metal flaps; or LEDs?

But look again at that question: As a design objective, it is incomplete. Consider what's missing: Is the source a pure-dc voltage, V_X(all t)? Or, are you chasing $V_X(t)$, a complex waveform? If you are, do you wish to determine its upper peak, lower peak, or both? Do you need to know its mean value, its rms value, and its long-term statistics? Will the touch of that probe tip seriously affect this voltage—possibly annihilating it? Should V_X be a few electrons stored on the subfemtofarad capacitance of some fragile nanogizmo? What accuracy do you need? How long can you wait for a result?

AN IC DESIGNER'S RELENTLESS DEMAND FOR CLARITY AND COMPLETENESS IN THE OBJECTIVES IS AN ESSENTIAL PRECURSOR TO THE EVENTUAL SUCCESS OF ANY NEW PRODUCT.

An IC designer's relentless demand for clarity and completeness in the objectives is an essential precursor to the eventual success of any new product. The refining of objectives does not have to come from a formal *proposal*: Experience in your domain is invariably enough to recognize an incomplete specification and fill in the gaps. Here, in tightening the net, you might expand the question as follows: Using the most efficient technologies, what minimum operating power does a handheld DVM need to visibly display the value of fixed dc voltages, from a source of less than 100V, with an error of less than 0.1%, allowing 3 sec of processing time?

With this much information, and using today's low-inertia IC processes and zero-power (although not zero-energy) LCDs, we are now listing microwatts rather than milliwatts of total power—tiny, but not zero. So you ask: Why not? Does the function of converting a voltage to a visible number fundamentally require the expenditure of any power at all? Why? We can accept that that circuit needs a lump of energy whenever you request a reading to change the state of hundreds of elements. System inertia (due to the charge-based nature of the transistors, the capacitance of the display elements, sometimes stray inductances, and other factors) is unavoidable.

But, suppose the requester really meant: "What power do you need to display the value of a fixed dc voltage for just one reading?" The thinking about this teaser is now more closely bounded, and in turn, the options that spring to mind become more specific, keyed to our familiar technologies. We wonder what to use for a voltage reference. It could still be a bandgap cell, operating at an internal bias of only 1 nA. And we can stomach this much wastage, because the READ button is pressed only once, starting the 3-sec measurement. For a given topology, the reference's native noise is non-negotiable, having its roots in transistor shot noise and the resistors' thermal noise. This characteristic does not mean that every bandgap topology will exhibit the same noise at this bias level.

The curious and industrious may wish to attempt the following exercises: First, determine what reference-noise spectral density is commensurate with a reading error of 0.1% that you attain in a 3-sec interval. (Assume no 1/f component.) Second, determine how you can reduce this noise without increasing the bias current. Third, determine how you can add the 3-sec time-out feature with the same stricture. Then, using the information you collect, calculate the one-shot energy usage. Finally, determine how you can trim the voltage to less than 0.1% absolute error, over a "handheld" temperature range—say, −5 to +45°C—and for all process corners.

Of course, a bandgap isn't the only choice. You could, import the pristine electrochemical voltage of an NBS

Digital Hot Swap

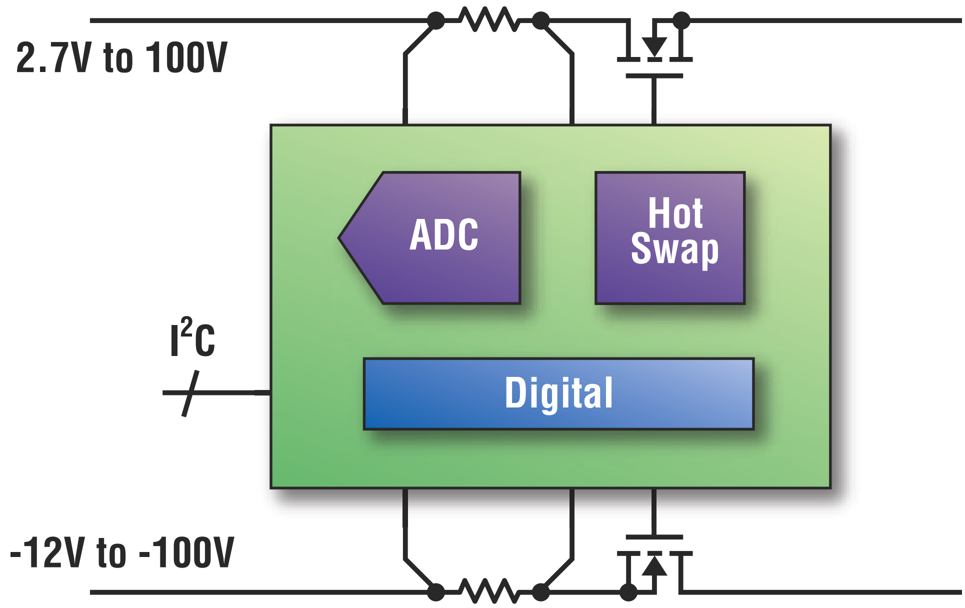

Hot Swap Controller Family Digitally Monitors Current, Voltage and Fault Status

Our Hot Swap™ controller family is optimized for 12V, 24V or 48V applications, including AdvancedTCA® and proprietary, single-supply implementations, and a quad supply controller for CompactPCI® or PCI-Express®. Each includes an integrated ADC and I²C™ interface, allowing digital power monitoring of voltage and current at the point of entry of power to high availability cards. The digital interface also allows fault status monitoring and ON/OFF control without the need to power up any additional components on the card.

Features

- ADC Monitors Current and Voltages
- I²C/SMBus Interface
- LTC®4215 for up to 12V Supply Applications
- LTC4245 Configurable for CompactPCI or PCI-Express
- LTC4260 for 24V or 48V Systems
- LTC4261 for AdvancedTCA, -48V Systems

Digital Hot Swap Controller Family

Part No.	Supply Range	ADC
LTC4215	2.9V to 15V	8-Bit
LTC4245	3.3V, 5V & ±12V	8-Bit
LTC4260	8.5V to 80V	8-Bit
LTC4261	-12V to -100V	10-Bit

Info & Free Samples

www.linear.com/digitalhs
Literature: 1-800-4-LINEAR
Support: 408-432-1900

(National Bureau of Standards) Weston cell into your handheld device and store it in an IC analog memory using well-known floating-gate techniques.

Next, consider the ADC. What power will generating the clock require? Do you even need a clock? The ADC can be asynchronous. It accumulates a debt of energy during its 3 sec of activity. But, if you take only a single reading, the power averages to zero (or to some tiny value if you take a reading once every blue moon). What about that display: How much power does it need, if any? Using LCD light modulators requires another lump of energy to charge its capacitive elements but with an average power approaching zero.

Oh! The objectives conveniently fail to mention that those dc voltages fall in the range of 100V to 1 kV. Does this news affect the design of the low-power dc voltmeter that's beginning to take shape in your head? Can you still make this voltmeter work from a 1.5V supply while measuring 1 kV with no increase in consumption? Can you imagine a way of dispensing with electronics altogether in meeting that objective?

WHEN FACED WITH THE NEED TO PROVIDE HIGH-QUALITY AUDIO FROM CMOS AMPLIFIERS, ENGINEERS DUSTED OFF AND TRIED A VERY OLD IDEA—CLASS D.

CHEATIN' THE POWER DEMON

High current consumption is often truly unavoidable, but at other times, engineers widely accept it as the sad truth until a new paradigm appears. Consider an IC-audio-power amplifier. Work backward from the speaker with a load impedance (casually assume it to be a pure resistance) of R_L and the desired maximum rms power, P_{MAX}. You can now calculate the peak output current $\sqrt{(2P_{MAX}/R_L)}$ and the peak output voltage $\sqrt{(2P_{MAX}R_L)}$. The former dictates the minimum size of the output transistors, with margins for process variations; the latter determines the minimum permissible supply voltage after deciding on the output-stage topology—single-sided or bridged—with adequate allowances for headroom. The required breakdown voltage of the transistors and such other considerations as frequency response and dynamic range narrow down the choice of IC process, then the output-stage bias mode (Class A, AB, and others), and finally, other detailed aspects of this fine architecture.

But design basics and trade-offs change over time. When faced with the need to provide high-quality audio from CMOS amplifiers, engineers dusted off and tried a very old idea—Class D. This approach didn't change the essentials of peak load current and voltage, but it drastically impacted other issues of output-stage design and eventually the entire amplifier. Most obviously, the transistors were now operating in the mode CMOS likes best: on/off switching. One consequence of this major difference is that the overall power efficiency becomes much higher. Just as the hot-wire light bulb turns most of the power it consumes into useless heat, so do classic analog-output stages. (It's what those monster-scale heat sinks are for.) Not surprisingly, the process worked, and, after a bit of learning and refining, it worked rather well. The new *binary amplifier* had emerged from one long-ago-discarded and crumbling cocoon.

First, the amplifier received a facelift by combining its core Class-D nature with other lessons from IC-switching-regulator design and sigma-delta data converters. Pulse-density methods replaced its simplistic duty-cycle modulation; the use of pseudostochastic "carrier" frequencies to broadband the EMI spectrum and other proprietary advances further augmented its sophistication. The "analog" audio amplifier has become a very-large-scale-integration digital engine.

THE RELEVANCE OF INERTIA

The ongoing development of IC-fabrication technologies led first to the significant benefits of well-balanced complementary-bipolar processes using standard junction isolation and, later, to significantly faster SOI complementary-bipolar processes using bonded wafers. In the early days, manufacturers made these SOI wafers by bringing a pair of standard 3-in. wafers into intimate contact, whereupon they would voluntarily "weld"—native oxide to native oxide. A laborious process of grinding and polishing removed all but a few microns of silicon from one of these wafers. This layer became the pure-crystal starting material on which to form transistors, starting with epitaxial deposition, followed by masking, ion implantation and drive-in, and, finally, a solitary metal layer for connections. The other wafer became simply a mechanical handle; the thin oxide layer between them became an important insulator.

Today, engineers widely use SOI processes, so these sandwiches are available commercially. SOI provides several crucial advantages. The transistors are true three-terminal devices: The absence of the usual parasitic transistor that the base, collector, and substrate layers form ensures that layout-level latch-up cannot occur. There is zero leakage current from a collector to its substrate layer. The collector-substrate capacitance, C_{JS}, is much smaller, and it is a pure, voltage-independent capacitance, unlike the varactor C_{JS} of a junction-isolated process.

However, these benefits come with one substantial setback: The thermal resistance of these devices, from the intrinsic transistor to the handle wafer, is very high, due mainly to the low conductivity of silicon dioxide (only 1/100 that of silicon). Thus, self-heating effects are pronounced. The bottom line: It is no longer possible to make the assumption of isothermal operation. This assumption has for decades been critical to monolithic analog design, as has the assumption of reliable matching in the key parameters of a transistor pair (notably, of the V_{BE} via I_S).

AUTHOR'S BIOGRAPHY

Barrie Gilbert is an IEEE fellow and director of the Northwest Labs at Analog Devices. EDN *named him Innovator of the Year for 1999.*

National Congratulates EDN on 50 Years of Industry Success!

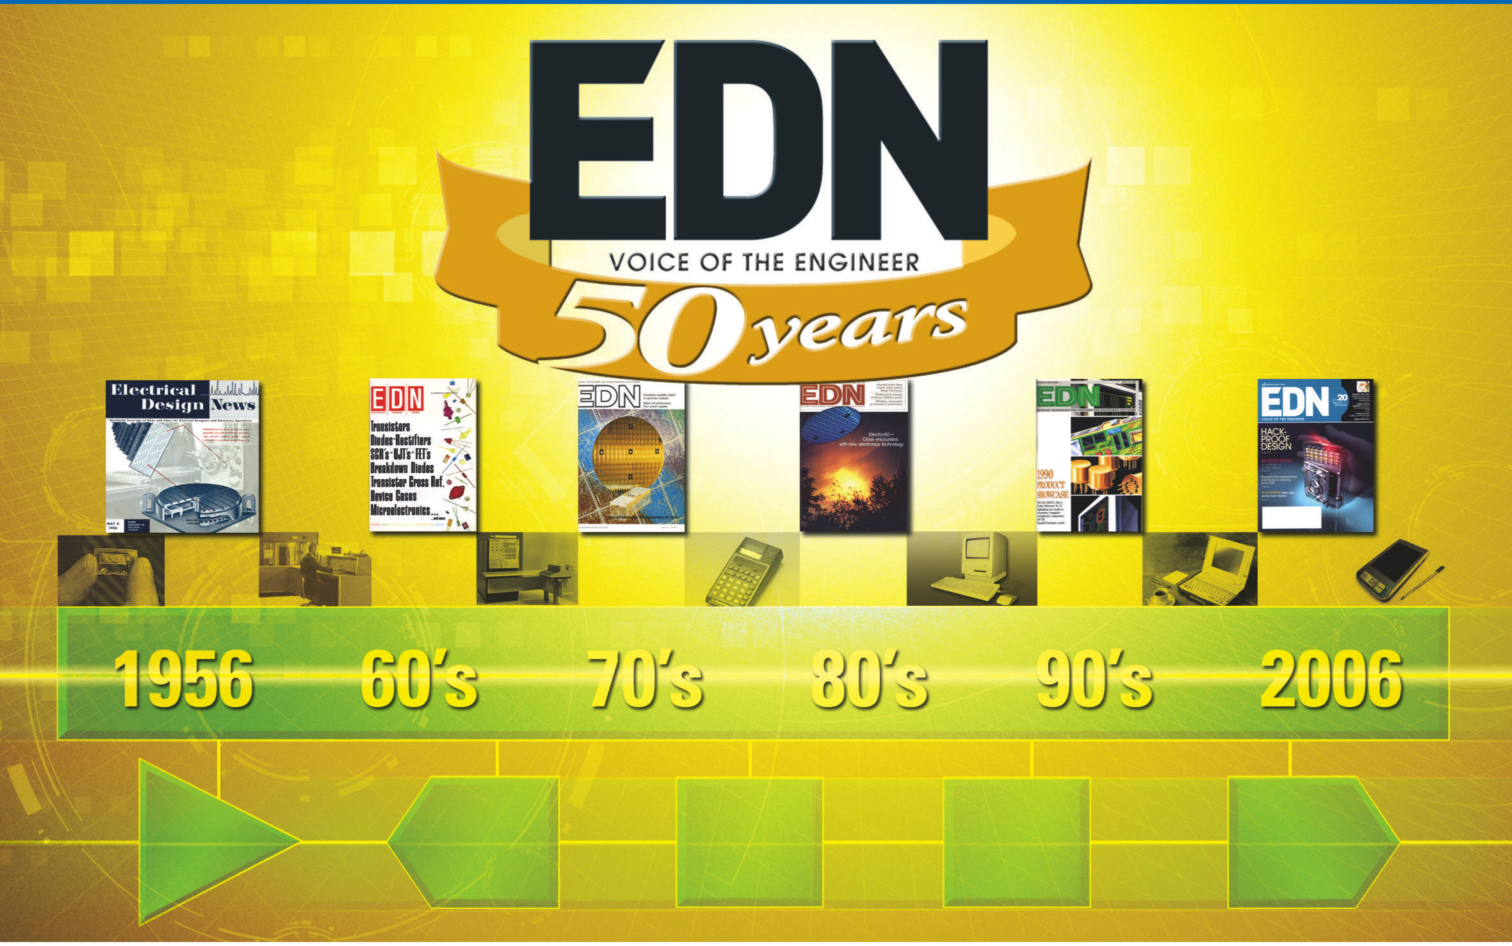

50 Years of Exciting Innovations that Drive the Electronics Industry

For the last half-century, EDN has reported the new technologies that drive the electronic engineering community. And National has helped to power those innovations, from developing the industry's first op amp and SIMPLE SWITCHER® products to the industry's LVDS standard and GHz-speed data converters.

As National also approaches its 50th year, we look back on these analog technologies that are the essential link between the digital world of ones and zeroes and the "real world" of light, sound, temperature, speed, and pressure. And we look forward to the innovations ahead—with EDN by our side reporting the exciting new technology stories that inspire us all.

www.national.com

National Semiconductor
The Sight & Sound of Information

Powering Digital Subsystems… and Simplifying Design

Easy-to-Use Power Management Units for Digital Subsystems

96% Efficient LP3906 Provides Flexibility with Digital Programmability

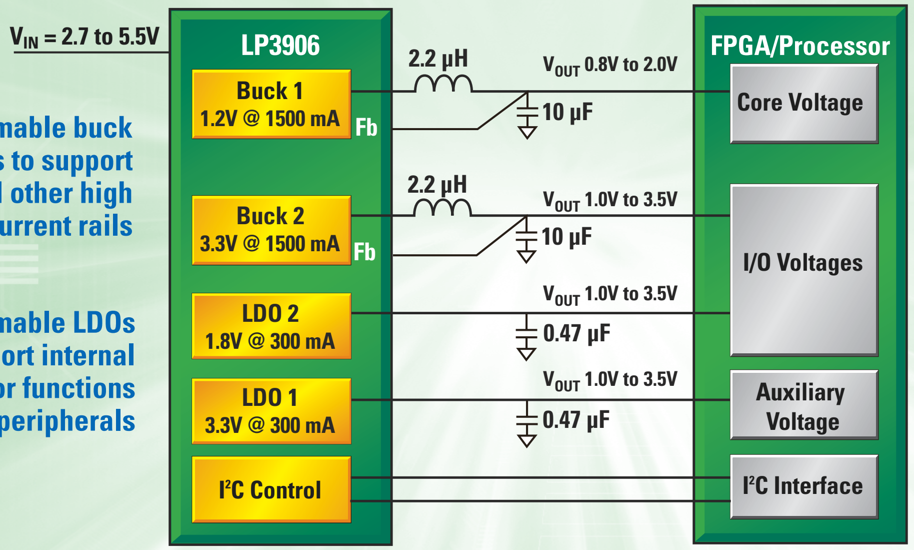

2 Programmable buck regulators to support core and other high current rails

2 Programmable LDOs to support internal processor functions and peripherals

I²C for independent control of LP3906 and peripherals

AVAILABLE LEAD-FREE

Product ID	Digitally Programmable	Efficiency	Regulator Output Current	LDO Output Current	Packaging	Solution Size
LP3906	I²C	Up to 96%	1.5 A	300 mA	LLP-24	20 mm x 20 mm
LP3905	N/A	Up to 90%	600 mA	150 mA (low noise)	LLP-14	15 mm x 10 mm

Ideal for powering application processors, FPGAs, and DSPs where size and efficiency are important

www.national.com

1.5A, 36V SIMPLE SWITCHER® Synchronous Step-Down Regulator

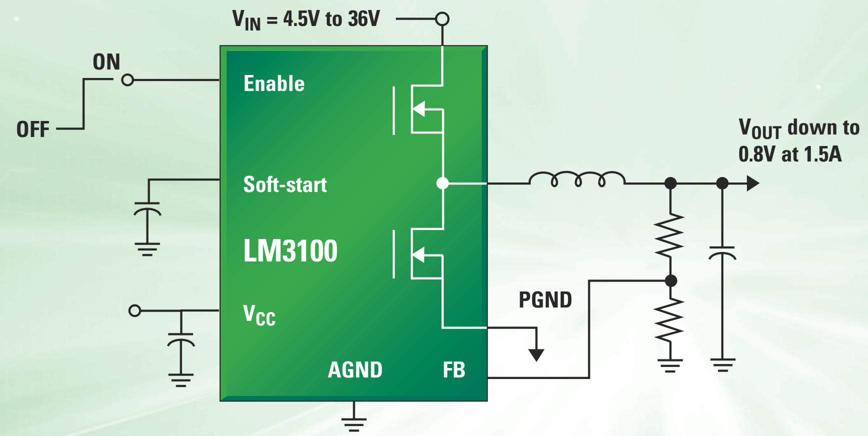

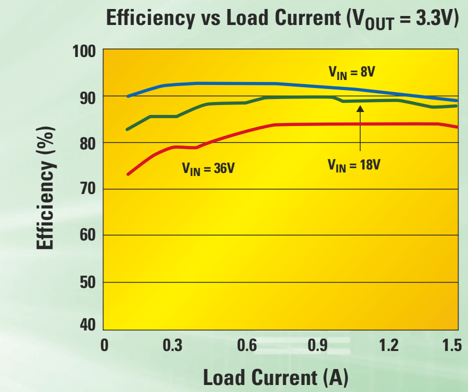

Shown at actual size 4.40 x 6.50 x 0.90 mm [ ] TSSOP-20

AVAILABLE LEAD-FREE

LM3100 Features

- Synchronous conversion for efficient operation below 3.3 V_{OUT}
- COT architecture provides lightning-fast transient response
- Stable with ceramic capacitors
- Near-constant frequency operation from unregulated supplies
- No external compensation reduces external component count
- Frequency adjustable up to 1 MHz

Ideal for use in embedded systems, industrial controls, automotive telematics and body electronics, point-of-load regulators, storage systems, and broadband infrastructure

WEBENCH® Online Development Tools

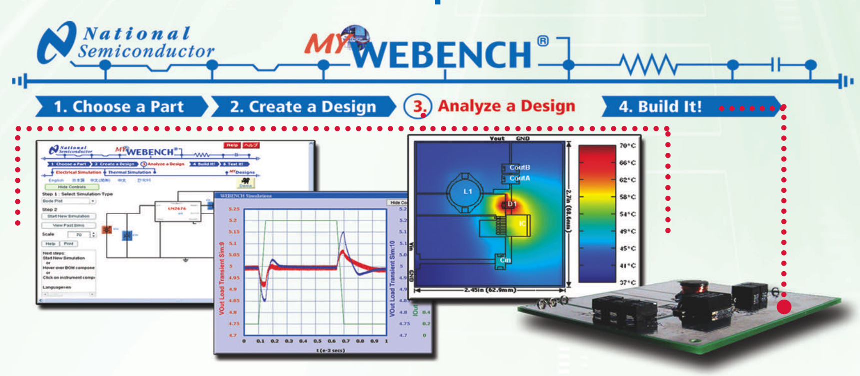

National Semiconductor
The Sight & Sound of Information

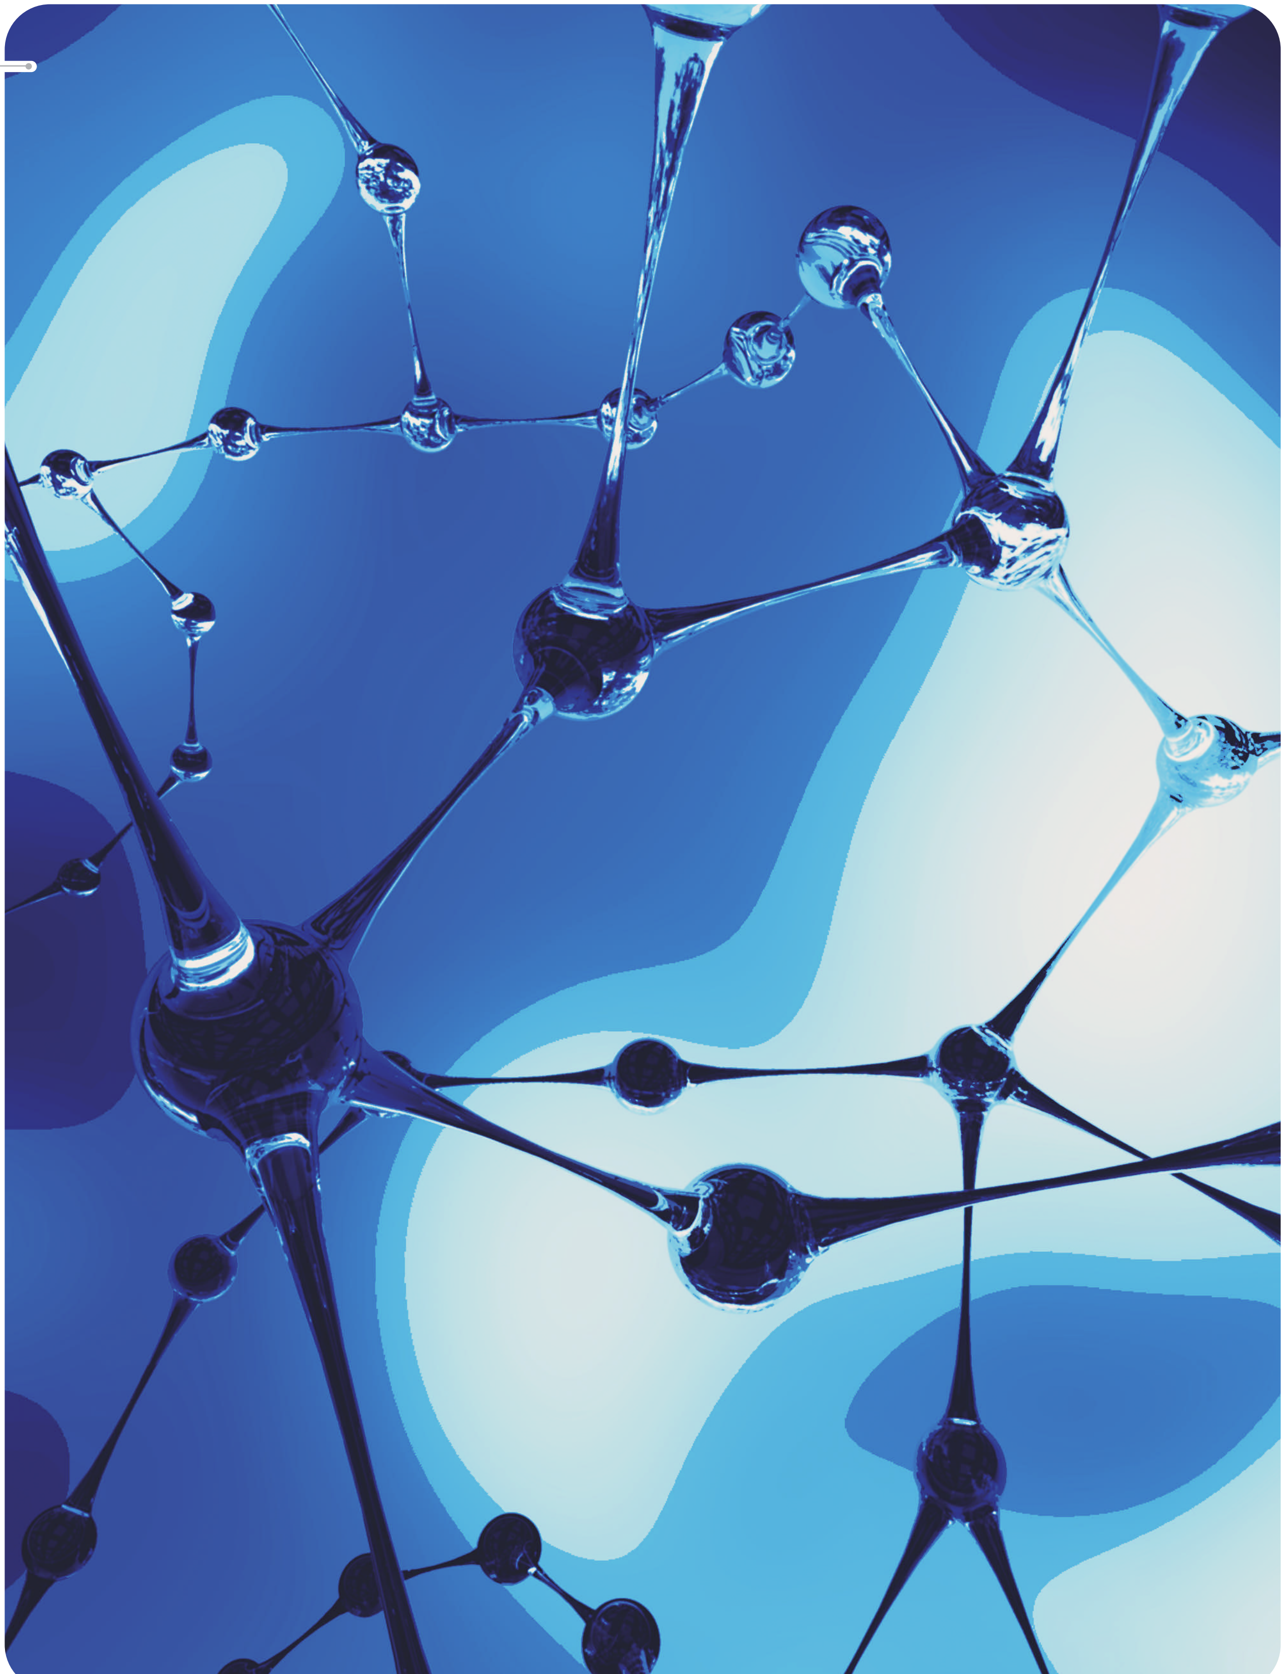

ROB MAGIERA / NOUMENA DIGITAL

BY MAURY WRIGHT, EDITOR IN CHIEF

MINIATURIZATION ENABLES INNOVATION—PAST, PRESENT, AND FUTURE

SHRINKING SEMICONDUCTORS STAND OUT IN THE RACE TO SMALLER PRODUCTS.

Ask someone on the street what miniaturization means to them, and they will most likely indicate a mobile handset or perhaps an MP3 player. Ask an engineer, and you'll probably get a Moore's Law-centric answer. Clearly, those answers are intertwined because one enables the other. And, in fairness, compelling portable consumer devices require smaller everything—speakers, microphones, disk drives, batteries, connectors, and so forth. Complex technologies, such as disk drives, require advancements that rival the innovation in ICs. Still, the big digital IC that lies at the heart of these products has been in the miniaturization spotlight at least since the Intel 4004 debuted, and the SOC (system on chip) promises to continue as the most important enabler of cool things for some time to come.

Over the 50 years of *EDN*'s history, few constants have remained in the tech industry, but there has been a constant march in shrinking end products, driven by the shrinking enabling technologies inside those products. We always hear that everything is smaller, faster, and cheaper, but even that saying is not necessarily true. Early transistors sold for far more than the technologies that they would usurp. Increased integration in ICs always makes for smaller products, albeit not necessarily higher performance ones. Only the smaller angle is almost universally correct.

On miniaturization, Texas Instruments (www.ti.com) Principal Fellow Gene Frantz says, "You can almost say that we are on the path to the vanishing product—where the product will be so small and insignificant in size but so significant in capability that we really don't know where we have it; we just know we have it."

Dean Kamen, founder and head of DEKA (Dean Kamen) Research (www.dekaresearch.com), has leveraged the miniaturization trend in everything from the Segway to computerized medical instruments. "We've come to expect that electronics have gotten so small, computing power has gotten so cheap, memory has gotten so plentiful, and power consumption has gotten so reasonable," says Kamen. "You put all of that together, and, literally, we're now at a point with electronics, which we aren't with any of the other fields of engineering, where it isn't a question of what you can do; it's a question of what should you do because you know it can be done. What was unthinkable 10 years ago is virtually trivial now."

Looking back, as we at *EDN* have done for our Milestones That Mattered series and interactive time line, it's amazing that inventors in our industry often didn't realize the true significance of their work—

HOW MINIATURIZATION BEATS THE HEAT

By Lewis Counts, Analog Devices Inc

Without advances in mixed-signal technology, a cell phone would be too hot to hold; you would need a backpack to carry it. Instead, today's mixed-signal ICs allow cell phones to feature color displays with still cameras and videocameras and even allow for TV broadcasts, especially if you live in Asia. A phone's size might make you forget it is in your pocket while you are on a commuter train, hiking in the mountains, or swooshing down ski slopes. Oh, and by the way, it also lets you make phone calls.

How did technology advancements take us from the op amps of the 1970s that were in dual-in-line packages to the miniature surface-mount ICs in today's cell phones? Moore's Law played a part, yes. But much more was involved than mere reductions in package size. For example, nanometer geometries have resulted in unheard-of digital-performance levels in laptop microprocessors. In addition, process enhancements and clever design also played big roles in improving analog-circuit performance–with a corresponding decrease in power consumption.

GAUGE IT WITH AN OP AMP

Op amps are everywhere, especially in SOCs (systems on chips). Today's op-amp circuit measures less than half the size of a 1970-era bond pad, or between 1/400 and 1/1000 the die area of one of the industry's first op amps, Texas Instruments' μA741.

The size means there is less capacitance to push around, so a microprocessor can perform more digital functions at higher speeds with the same amount of power. The analog world also benefits from less capacitance. At least two ways exist to view this phenomenon.

First, lower on-chip capacitance translates to less current to achieve the same or more bandwidth and slew rate. You can compensate amplifiers to maintain closed-loop stability at higher frequencies because lower capacitance also raises the frequencies of parasitic poles, the dominant cause of instability.

Second, less current means less heat. The μA741, which Analog released in 1968, typically draws 1.7 mA from ±15V supplies. That amount may not seem like much; it is only 51 mW. Today, that same 1.7-mA figure powers two video-speed op amps with more than 400 times the μA741's gain-bandwidth product. And 1.7 mA are more than enough for one precision op amp with 50 times less offset voltage, seven and a half times less drift, and 36 dB more common-mode rejection.

Because of their output stages, stand-alone op amps draw more current than their integrated cousins. System designers need stand-alone circuits for design flexibility in one-of-a-kind applications, and they provide a valuable gauge of bandwidth per milliamp from a standard function.

MORE BANDWIDTH PER MILLIAMP

Some interesting trends have emerged in the gain-bandwidth-per-supply-current metric (Figure A). Analog Devices built the μA741 using Fairchild's planar-bipolar process, then the granddaddy of all IC manufacturing. Its internal compensation set it apart from earlier competition.

BiFET op amps, such as Linear Technology's LF356 and Analog Devices' AD711, began to appear in 1978. These circuits offered a 1.4 to four times bandwidth increase for each milliamp of supply current. Input-bias currents decreased dramatically, thanks to integrated JFETs: The AD711's 15 pA is more than 5000 times lower than the μA741's.

But the situation for wideband-system applications was even better. Adding true complementary PNPs to go with the NPNs, Analog Devices' AD847 of 1988 had more than 10 times the μA741's bandwidth per milliamp.

A later silicon-on-insulator-process technology resulted in the development in 1995 of the Analog Devices AD8011, with more than 400 times the bandwidth performance per milliamp. The demands of high-definition video would have been impossible with the μA741, and the AD847 would have been hard-pressed to meet the performance requirements. But video op amps such as the AD8011 make it easy.

Bandwidth for the supply current has increased rough-

continued on pg 176

Outstanding Performance

Central's New Series: Low Voltage Zener Diodes

Central's New Multi Chip Stabistor Diode

The Central CMXSTB series of multi-chip stabistors are very low voltage zener diodes with an extremely stable forward voltage over a wide current and temperature range. Manufactured in a space saving SOT-26.

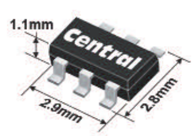

Typical Applications

Low-voltage stabilization:

- Bias stabilizer in class-B output stages
- Clipping
- Clamping
- Meter protection.

Voltage vs. Current

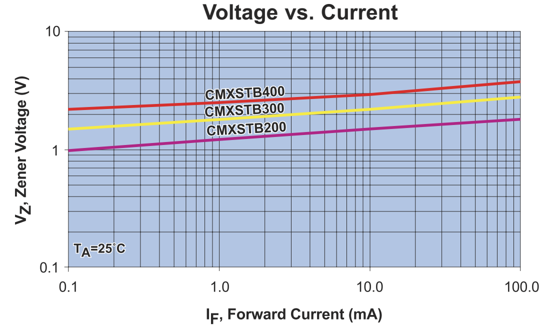

Voltage vs. Temperature

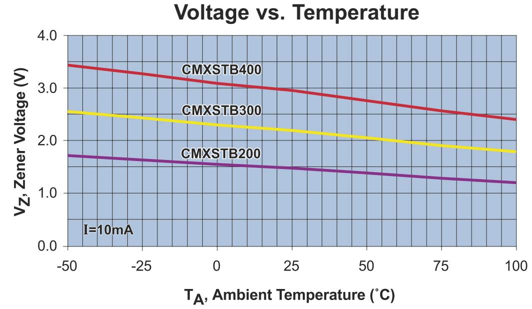

Need Samples?

CMXSTB200
1.2V

CMXSTB300
1.5V

CMXSTB400
2.5V

Visit: www.centralsemi.com/stabistors

Small Signal Transistors | Bipolar Power Transistors | Energy Efficient Devices | MDMs | Stabistors | Rectifiers | Thyristors | Transient Voltage Suppressors | TLMs

Central™ Semiconductor Corp.

www.centralsemi.com

145 Adams Avenue, Hauppauge, NY 11788 USA • Tel: (631) 435-1110 • Fax: (631) 435-1824

at least at the time of their inventions. And third parties, including *EDN* editors, were often no more perceptive. For instance, in our look back at Fairchild's (www.fairchildsemi.com) introduction of Micrologic Elements (see "The planar IC—revolution underestimated" at www.edn.com/article/CA6325586), *EDN* originally discounted the significance of the ICs, reasoning that batteries and other components in aerospace applications would render miniature logic insignificant. Of course, that assumption ignored the multitude of consumer applications that would come to rely on digital logic.

In discussing milestones from the technology industry history, Walden Rhines, chairman and chief executive officer of Mentor (www.mentor.com), states, "In most cases, the actual components were developed initially without knowledge of what the killer application would be." Rhines claims that such lack of foresight is a typical trait of revolutionary developments.

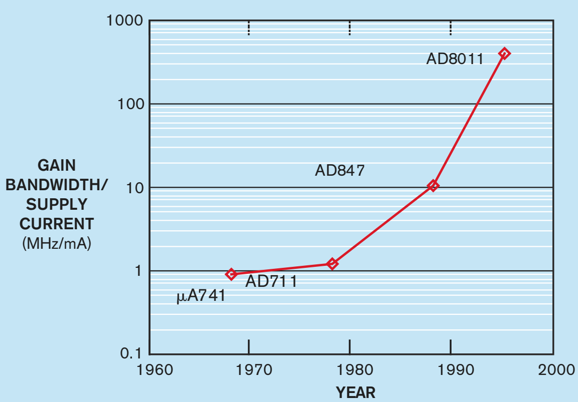

Figure A Some interesting trends have emerged in the gain-bandwidth-per-supply-current metric.

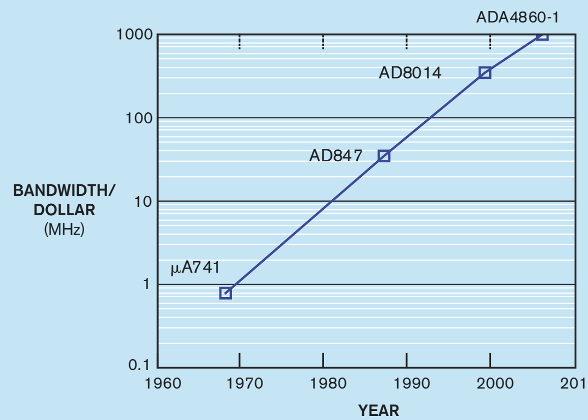

Figure B Bandwidth per dollar has increased about 10 times per decade.

continued from pg 174

ly 10 times every decade since 1980. And the AD8011 costs less than the μA741 did when it debuted.

MORE BANDWIDTH PER DOLLAR

Figure B illustrates that bandwidth per dollar has increased about 10 times per decade. Today, you can purchase 1-GHz devices for $1–and that dollar is inflated compared with a 1968 dollar. The μA741 came in an eight-pin TO-99 or mini DIP. For designers requiring high-speed op amps, the ADA4860-1 is available in a six-pin surface-mount package almost nine times smaller than the μA741 with lower parasitics and no holes to drill into a pc board.

Do you need high precision rather than speed in that tiny space? Op amps such as the AD8698 have just 20-μV offset, 50 times less than a μA741. Analog Devices' AD8698 also has 36 dB more power supply and common-mode rejection, making it useful in noisy environments. It also touts 30,000 times more gain and 1/100 the input-bias current–for lower cost.

The last 35 years have not brought just increasing bandwidth and precision for lower current, cost, and size. Shrinking processes also made possible cool analog- and mixed-signal ICs in small spaces. These ICs in turn made possible cell-phone still cameras and videocameras and wireless headsets. Cell phones are low-power enough to allow several hours of call time on a small battery.

Without advances in signal-processing technology, you would not be able to carry a backpack-sized cell phone far, and it would burn your hand to hold it, even if you could. Instead, you would have to find an ac outlet just to power it. Now, you can transmit video to colleagues, family, or friends as you talk long-distance on your cool-running, pocket-sized gadgets on your commute to work, hike in the mountains, or ski run down the slopes.

REFERENCES

1 **Lee, TH, "IC Op-Amps through the Ages," from course notes, EE214, Stanford University, 2002.**
2 **Aurgarten, Stan, *State Of The Art,* Ticknor & Fields, New York, 1983.**
3 **μA741 data sheet, Texas Instruments Inc, 1970.**
4 **LF356 data sheet, National Semiconductor Corp, 1978.**
5 **AD711 data sheet, Analog Devices Inc, 2002.**
6 **AD847 data sheet, Analog Devices Inc, 1992.**
7 **AD8011 data sheet, Analog Devices Inc, 2003.**
8 **AD8014 data sheet, Analog Devices Inc, 1999.**
9 **AD8698 data sheet, Analog Devices Inc, 2004.**
10 **ADA4860-1 data sheet, Analog Devices Inc, 2006.**

AUTHOR'S BIOGRAPHY

Lewis Counts is vice president of technology and fellow at Analog Devices.

He continues, "If you float a technology out there and make it easy enough to use, someone will discover it and match it up with the end application."

Before working at Mentor, Rhines worked at TI managing the semiconductor group. And both Rhines and Frantz relate similar stories about the birth of the DSP. In the early 1980s when the TMS32010 came to market, TI expected—but didn't find—success in the speech market. The two markets in which they would find early success—PC-graphics acceleration and hard-disk-motor control—weren't on the radar. Later, the DSP would enable the dial-up modem and the digital cellular handset, among other products.

THE ONGOING TECHNICAL REVOLUTION

By Gene Frantz, Texas Instruments

We are in the middle of a technological revolution. Innovation is rampant as a result of advances in IC technology, which has improved end products by bringing us higher performance, lower cost, longer battery life, and decreased size. Each of these aspects has had a part in fueling this technological revolution. Things that you now carry in your pockets used to require large, air-conditioned rooms. This revolution has come in three waves. You have seen one successful wave of innovation in telecommunications. Now, you are in the middle of a second wave, entertainment, and are at the beginning of the third wave.

WAVES OF INNOVATION

Next year will mark the 25th year that Texas Instruments has been in the DSP (digital-signal-processing) business, and, during those years, TI has been a key driver of the first two waves of the revolution. The first wave, telecommunications, digitized the phone system to allow phone lines to transmit data. Modem rates rose from 1.2 to 56 kbps, and DSLs (digital-subscriber lines), cable modems, and Ethernet carried data rates into the megabit-per-second area. Meanwhile, digital cell phones evolved into multimedia appliances that integrate cameras, MP3 players, and more.

The second wave, entertainment, started in 1983 with the introduction of the CD format. Entertainment demands not only performance, but also portability–or, in IC terms, low power dissipation.

Today, products that are not only portable, but also personal define digital entertainment. They record, store, and play back your music play lists, photos, videos, and games.

WHAT'S NEXT?

With the first two waves well under way, what is the next wave? It could be transportation, quality of life, security, or education. In transportation, improving vehicle safety will be a strong technology force over the next 20 years, with imaging, sensing, and in-vehicle safety systems playing important roles. Engineers have worked to make cars safer except for removing the obvious problem: the driver.

Medical applications are among the leaders in improving quality of life. An eye-opening example is an artificial-vision system developed at the University of Southern California (Los Angeles, www.usc.edu). A miniature camera attaches to a wireless receiver behind your ear. The receiver carries signals that travel along a wire under your skin to an electronic electrode array that attaches to the eye's retina. Resolution at this point is minimal–a 4×4-pixel image.

Speech patterns, fingerprints, facial recognition, and retinal scans can all identify people, but security needs 100% accuracy. For this reason, transitions will occur in this area. First, the camera, a source of information for humans, will find use as a second source of opinion, carrying information to the human in charge. In another transition, information from the camera provides the first opinion, and the human acts as the backup. Finally, information from the camera becomes the *only* opinion.

In education, most people accept that technology can improve the learning process. The sticky problems of determining which technologies to deploy and how to measure their effect counter this positive view, however. Much like physicians, educators often believe that their first duty is to do no harm. This conservative approach has allowed students to sprint ahead of their teachers in adopting technology in the classroom. Sometimes, their innovations are unnerving, such as using the SMS (short-messaging service) of mobile phones to cheat on tests. Typically, educators respond to this bright idea by barring cell phones.

Better ways to use this technology should emerge. With all of the computing, communications, and entertainment capabilities available, a revolution in this market is about to happen. Innovation will drive this revolution.

For the next wave, education may well be the sleeper of the four candidates, but transportation, security, and quality of life are not. If you think the last 25 years have been exciting with new innovations, just wait for the next 25 years.

AUTHOR'S BIOGRAPHY

Gene Frantz is technical fellow at Texas Instruments (Houston).

MINIATURIZATION HISTORY IS NOT THE FULL STORY

By Walden C Rhines, Mentor Graphics

At first glance, the progress of electronics seems to have been a steady march toward miniaturization. Certainly, the path along major milestones–from vacuum tubes to transistors to ICs to ASICs to SOCs (systems on chips)–generally tells a tale of more on less. But the cost reduction has been the driving factor toward miniaturization.

Almost all cost reduction for ICs has historically come from shrinking features and increasing wafer diameters. But miniaturization has occurred as a special case of the learning curve, which mandates that the cost per unit of everything produced decreases by a fixed percentage every time total cumulative volume doubles. This law applies to all goods and services produced over the centuries when measured in constant currency. Throughout their history, transistors and ICs have followed a relatively steep learning curve.

Gordon Moore last predicted in 2003 that Moore's Law, which postulates that the number of transistors on the most advanced ICs will double every 18 months, would be good for another decade. But even Moore acknowledges that "no exponential is forever." According to the ITRS (International Technology Roadmap for Semiconductors, www.itrs.net), by 2010 we will have only eight electrons available to switch a transistor. By 2020, only one electron will be available, which, although theoretically possible, stretches the laws of physics. So, it is clear that scaling has physical limits. Shrinking feature sizes and increasing wafer diameters will no longer maintain the learning curve. Miniaturization brings its own set of serious problems. Nevertheless, the learning curve will continue. But the methodologies to stay on the curve are changing. The ongoing drive to achieve reduced cost per function and improved quality will require other innovations besides miniaturization.

These new developments include innovative assembly and packaging of multiple chips. Despite the historical drive toward the single chip as the ultimate approach for any application, it is becoming increasingly apparent that the single chip is not always the lowest cost, most reliable, or highest quality. When vendors integrate semiconductor technologies on a single chip, the final chip favors the lowest cost and takes place in the least expensive technology, rather than a superset of all the technologies. Again, cost reduction is the driving factor, and reduced size matters only in so far as it helps to reduce costs.

As geometries shrink, the fixed costs of design, verification, software, and tooling have increased to the point at which it's common to see companies spend $20 million to design a complex SOC. And the revenue per ASIC is increasing more slowly than the fixed costs of cutting-edge design. The learning curve is seriously out of alignment here, and a change in ASIC methodologies is essential to reducing the fixed costs. Attempts include structured ASICs and richer design libraries, but the most promising approaches include a move to higher levels of abstraction and to new techniques for system verification.

Whenever size decreases, power dissipation and power density increase. Indeed, power issues constitute major challenges for today's designers. To increase gate count without increasing power density, designers need to turn to innovations in advanced process technology, such as high-k dielectric gates; lower channel-leakage structures, such as FinFET and oxide isolation; improvement in circuit-design techniques, such as variable voltage and power-down techniques; and, especially, architectural trade-offs in the system design.

Semiconductor-manufacturing costs have soared with miniaturization. When these two vectors collide, further miniaturization depends on cost containment. For example, resolution-enhancement software is beginning to offset lithography-equipment costs, which were increasing in step with Moore's Law. Techniques such as embedded deterministic test are dramatically reducing test costs, which are sometimes more than half the manufacturing cost.

The number of RET (resolution-enhancement-technology) layers increases every process generation, and physical features dramatically impact yield, which has major cost ramifications. DFM (design-for-manufacturing) and design-for-yield tools are addressing the problem at the design stage, with DFM guidelines to avoid feature-limited yield problems and lithography-friendly tools that allow designers to analyze the effects of manufacturing variability.

Long after Moore's Law becomes invalid, we will continue to reduce the cost of logic gates, just as we did before the development of solid-state electronics. However, we won't achieve this indefinitely by miniaturization, because miniaturization must have a physical limit. We will continue to develop other techniques to reduce costs, as the learning curve ensures that innovation continues.

AUTHOR'S BIOGRAPHY

Walden C Rhines is chairman and chief executive officer of Mentor Graphics Corp.

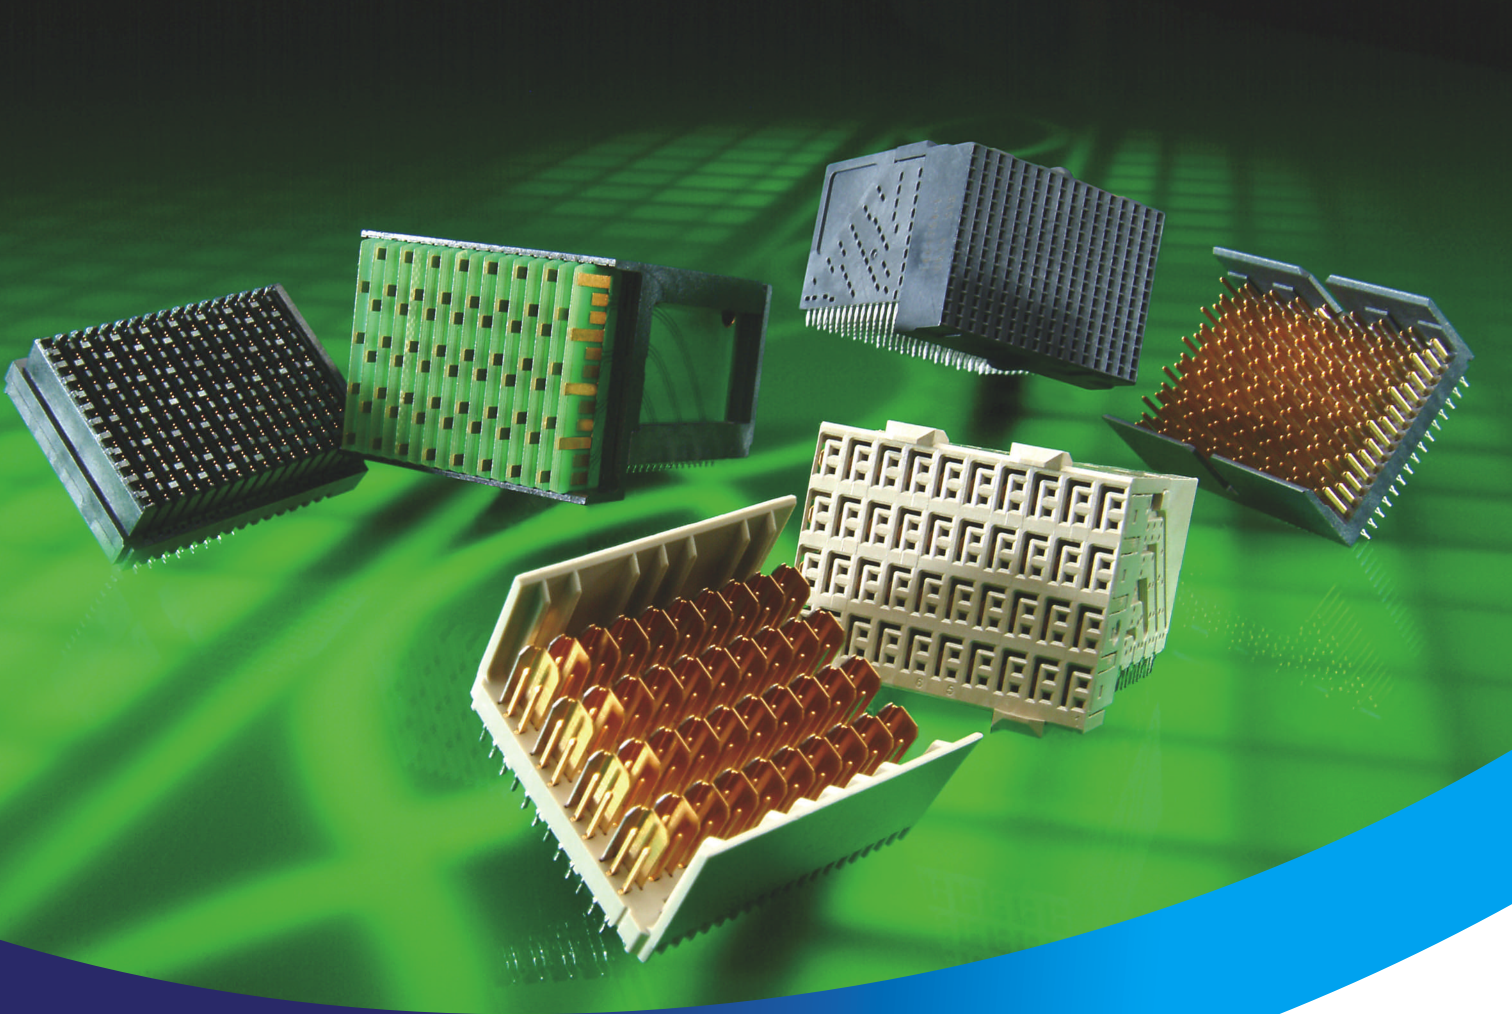

You might wonder about recent examples of technology's success in unforeseen ways, because developers are increasingly designing products for target applications. "As an industry matures, you have less of that," Rhines admits.

Rhines is enthusiastic in looking back at his experience, both at Mentor and at TI, and offering opinions on key milestones in the miniaturization march. He points out the significance of the first germanium and then the first silicon transistors. He claims that the silicon version was especially significant because it was stable over temperature changes and therefore suitable for consumer products such as transistor radios.

Rhines also credits Intel (www.intel.com) for commercializing microprocessors and memory and especially for moving to NMOS technology, thereby enabling denser memories. And he believes the industry's move to CMOS was ultimately even more significant. Rhines relates that TI had to use low-power CMOS for its work with battery-powered devices, such as watches and calculators. But the challenge was immense. "Fundamentally, you were saying, 'I'm going to put two transistors down everywhere there is one today, and I'm somehow going to make the die smaller, faster, and lower power,'" he says. Rhines credits the Japanese conglomerates for wading through the CMOS challenge and delivering on its promise. And CMOS was clearly the best technology choice for digital ICs and even increasingly for analog ICs.

At the same time that microprocessors and memory were matriculating at Intel and elsewhere, Bob Metcalfe was already thinking about the value in connecting discrete computers. The miniaturization trend would make Ethernet possible. Metcalfe developed the LAN while working at Xerox PARC (Palo Alto Research Center) in the early to mid-1970s and went on to found 3Com and make Ethernet a commercial success.

Metcalfe recounts that the first version of Ethernet in Xerox's labs operated at 2.94 Mbps. After five years of work with Ethernet in a laboratory-type environment, his team evaluated the 2.94-Mbps network and found it to be nearly optimally loaded. The group foolishly thought that it had arrived at the exactly correct speed choice. Looking back, Metcalfe admits, "The fact was that any application that couldn't run at 2.94 Mbps didn't catch on. All the applications that could squeeze through 2.94 Mbps did." Soon, the first standardized version of Ethernet would emerge at 10-Mbps rates. "We settled on 10 Mbps for the first standard because that's how fast the Intel chips could run," he says. And Metcalfe garnered the credit for Metcalfe's Law, which suggests that the value of a network is proportional to the square of the number of users on the network.

Although silicon advancements are fairly easy to chart along Moore's Law and to apply in computers, real-world products require motors, valves, and the like. Disk drives are some of the few electromechanical products to truly ride the miniaturization curve. "Those of us in the semiconductor industry tend to try not to pay attention to the fact that the density of bits on a hard disk have increased at a slightly greater or at least at the same rate as the density of bits on an IC," says Mentor's Rhines.

Even in the real world, though, it's tough to move a miniaturization discussion away from semiconductors. Kamen of DEKA has worked on everything from tiny, implantable medical devices to his iBot, which is somewhat akin to a Segway but can also climb stairs and helps the disabled regain mobility. "The big things for us were DSPs, but high-power control stuff was also important," he says, noting that a device such as the iBot needs both powerful motors and power semiconductors and control devices that can drive kilowatt-sized loads in small packages.

The meaning of "miniaturization" also depends on the context. For a disabled person, an iBot is a miniature transportation device. For dialysis patients, the portable Homechoice dialysis machine, which DEKA designed and Baxter International (www.baxter.com) sells, frees patients to travel and receive needed medical care. Kamen also notes that the lowering of the quiescent power in semiconductors is key to small medical devices, such as implantable products, which must run for extended periods.

"ANY APPLICATION THAT COULDN'T RUN AT 2.94 MBPS DIDN'T CATCH ON. ALL THE APPLICATIONS THAT COULD SQUEEZE THROUGH 2.94 MBPS DID."

Regarding electromechanical products, such as motors, Kamen says that performance is improving but that cost has also increased, compared with the costs in the IC market. Kamen makes the same generalization about batteries. He advocates moving mechanical functions into the electronics whenever possible.

Battery and power issues seem to resonate with everyone from users facing frustratingly short handset-battery life to design engineers managing expectations and capability. The road map of process technology for years has guaranteed power savings. These days, however, leakage current means that savings come less freely. Designers have had to get serious about power management.

"We have just figured out that power dissipation is important," says TI's Frantz, who points out that you must attack the problem both by using better batteries and by increasing the power efficiency of the circuits. "If I reduce the power dissipation of my electronics by half, I can either, with the same battery, double the

PIC12HV615/PIC16HV616 Closed Loop Fan Speed Control

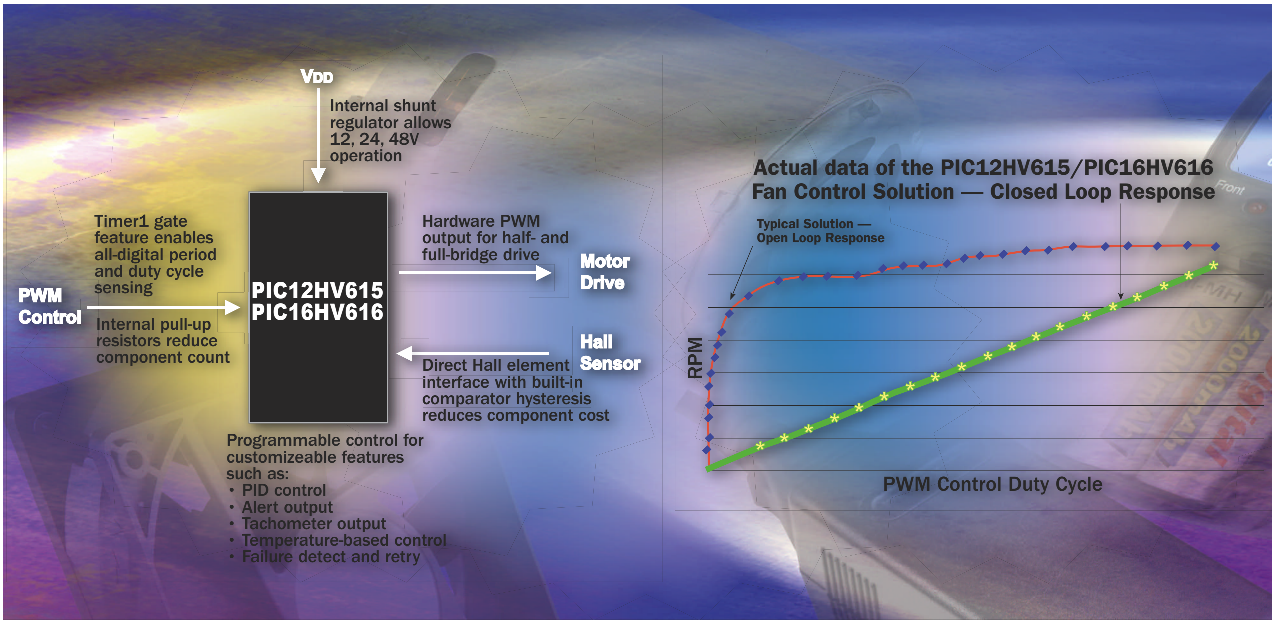

Advantages of the PIC® Microcontroller Fan Control Solution:

- Meets or exceeds the latest industry specifications
- Provides closed loop linear control
- Software-programmable dynamic response
- High level of integration provides simplicity and cost savings
- Offers high voltage support with integrated shunt regulator
- User-enabled comparator for hysteresis direct interface with a Hall Element

Now available for purchase on....
microchip DIRECT
www.microchipdirect.com

Get started with the MPLAB® In-Circuit Debugger (ICD) 2 Module for only $159.99 USD.

	High Voltage Support	Program Memory	10-bit ADC (channels)	Comparators	Enhanced CCP*	Internal Oscillator
PIC12F615	—	1.75 KB/1 Kw	4	1	Half bridge	8 MHz
PIC12HV615	Yes	1.75 KB/1 Kw	4	1	Half bridge	8 MHz
PIC16F616	—	3.5 KB/2 Kw	8	2	Full bridge	8 MHz
PIC16HV616	Yes	3.5 KB/2 Kw	8	2	Full bridge	8 MHz

* Capture, Compare, PWM

For fan control application notes, reference designs and related code examples, visit www.microchip.com/fancontrol. Microchip also offers a full line of stand-alone fan controllers and fan fault detectors.

Now Pb-free! RoHS Compliant

MICROCHIP
www.microchip.com/fancontrol

Microcontrollers • Digital Signal Controllers • Analog • Serial EEPROMs

battery life or, with a battery half the volume, keep the same battery life," he says. The factor of two that Frantz mentions can be significant, but he also suggests that a reduction in power requirements by an order of magnitude could make plausible a battery that is one-tenth the size of today's batteries.

You can save power by increasing the number of transistors on a chip, he points out. The savings can come not only through better power management, but also through more basic chip operation. A dual-processor chip with each processor running at 0.5 GHz from a 1V supply, he claims, offers significantly lower power dissipation than a single-processor chip running at 1 GHz from a 1.5V supply.

Even Intel has low-power-system fever. When the company launched the Intel Core 2 Duo in August, President and Chief Executive Officer Paul Otellini claimed that the company began to focus more heavily on performance per watt when it launched the Centrino in March 2003. Intel still pushes performance. On the newest, second-generation, 65-nm strained-silicon process, Otellini says, " We've seen stunning improvement in the transistor-level performance from generation to generation." The Yonah core that Intel announced in January of this year reduces leakage current by a factor of 25 and boosts performance by 40%, he says.

The jury is still out on how much progress the industry has made on low power. Luis Pineda, senior vice president of Qualcomm Chip Technology's (www.qualcomm.com) marketing and product-management group, is more bullish than TI's Frantz, and both work at companies with a major handset focus in which battery life matters. "We're very advanced with our power-management techniques," says Pineda. Referring to a CDMA-chip set, he says, "One of the chips is a dedicated power-management IC." The power manager can shut down portions of the chip set when necessary and controls the power amplifier, he says. Moreover, it integrates many discrete functions, such as low-voltage regulators and peripheral drivers, from earlier designs.

Pineda also claims that process technology is still helping to reduce power. He points out that Qualcomm is now shipping 65-nm ICs and working toward 45-nm ICs. "There is a lot more work than we can do," he says. He believes that the demands from the consumer—especially to play video—will escalate. Pineda claims, for instance, that the iPod Video can play at most 90 minutes and that consumers will demand more. Referring to power management, he states, "We're probably three-quarters to where the industry can be."

Turning to today and the future, where do we stand overall with enabling technologies for miniature end products, and what are the future trends? Some of the relatively new enablers are amazing. For instance, Knowles Acoustics (www.knowlesacoustics.com) approximately two years ago announced tiny MEMS (microelectromechanical-system)-based microphones for use in handsets and other small products. The company just followed with the MEMS-based Digital SiSonic microphone, which outputs a pulse-density-modulated bit stream. The company also offers miniature balanced-armature speakers that fit into headsets that rival the best noise-canceling headphones on the market.

THE JURY IS STILL OUT ON HOW MUCH PROGRESS THE INDUSTRY HAS MADE ON LOW POWER.

Even cabling and connectors may surprise you. Molex (www.molex.com), for instance, is shipping coaxial connectors with 0.4-mm pitch. The 40-pin connectors find use in connecting through hinges in clamshell handsets.

Grand surprises are still to come. These days, Ethernet inventor Metcalfe works as an entrepreneur and a venture capitalist at Polaris Venture Partners (www.polarisventures.com). He also serves as chairman of Ember Corp (www.ember.com), which focuses on the ZigBee wireless standard.

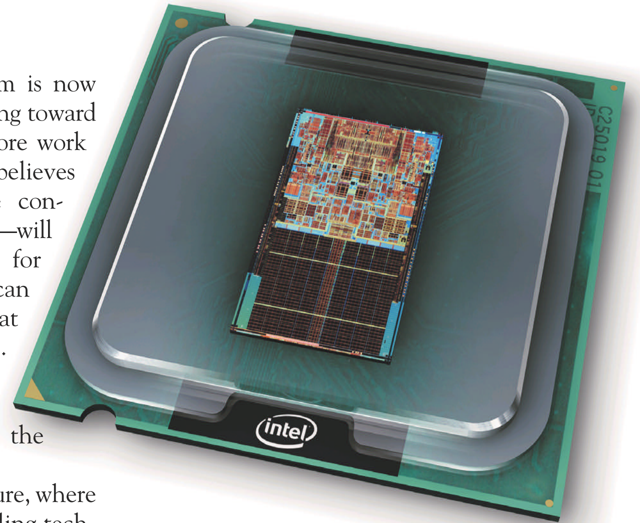

While still tracking Moore's Law with a move to multiple cores, Intel has recently shown an increased emphasis on power efficiency. Both trends are evident in the recent Core 2 Duo launch.

Other investments connect Metcalfe to the miniaturization angle. For example, Metcalfe invests in Unison Products (www.thinfilmaudio.com/), which plans to integrate speaker technology on the surface of a display. The scheme relies on tiny piezoelectric elements along the edges of a display that move a membrane stretched across the screen. Presumably, the scheme will conserve space, and the sound comes from the best possible location facing the user.

Perhaps the biggest advancement might be in energy "harvesting," or gathering energy to power a device from the immediate environment. Researchers are working on converting body heat, vibrations, stray RF energy, and light of all types into energy that could keep a relatively small battery charged. Remember Frantz's comment about a vanishing product? He finished by saying, "As much as possible, it would be good to run a product off body heat or to have no power or battery at all." But is energy harvesting feasible, and, if so, when? "If you think solar power is one of those solutions, it's here," he says, pointing out that TI has long built solar-powered calculators that use a relatively small battery that charges sporadically in lighted environments.

It's probably no surprise that Metcalfe has an energy investment and an interest in harvesting. He states, "Batteries are

50TH ANNIVERSARY ONLINE
THE CELEBRATION CONTINUES

Visit www.edn.com/50th for an interactive time line of the last 50 years, for *EDN*'s Milestones That Mattered series, and for commentary from past *EDN* chief editors.

a big venture-capitalist kind of thing; there are a million battery deals going around." His investment, Infinite Power Solutions (www.infinitepowersolutions), focuses on tiny batteries that generate tiny amounts of power at a low duty cycle. The batteries might find use in applications such as implantable medical devices. A patient might sporadically—say, once a year—visit a doctor to receive an inductive battery charge through the skin, for example. Someday, energy harvesting might continuously charge that battery. Infinite Power is building a semiconductorlike fab to build the flat batteries using vapor deposition on a flexible substrate and is offering a 1-in.^{2} battery for sampling. You can solder the battery onto a pc board or build it into the board and ultimately perhaps into an IC substrate.

The discussion returns to semiconductor processes even when a battery is front and center. After all, process technology has been the constant enabler of technology. "All new fields lead off with innovations in material science," says Mentor's Rhines, pointing out that mastering silicon-purification techniques was the obstacle that TI, Fairchild, Intel, and others cleared in launching the IC industry, and material scientists drove the innovation.

Meanwhile, other engineering disciplines await such a breakthrough. "We're starting to enter an age now in which electronics might have been a couple of decades ago, when material science, surface finishes, nanostructures, building compounds, materials, and surfaces are getting so extraordinary that we are going to see some big changes in everything from electrical properties to thermal properties to strength and structural properties of materials that are going to be pretty exciting," says Kamen.

He still counsels that you should leverage the maturity of electronics in every way possible. "There are opportunities to move some of the issues out of the world of the mechanical into the world of the electronic and take advantage of the high-performance electronics," he says, citing the transition from the phonograph to the CD to the MP3 player as an example. The CD eliminated contact with the media, and a flash-based MP3 player eliminates the mechanics.

Metcalfe is also a staunch believer in the IC as the key enabler. "Moore's Law is the fundamental; all of the rest comes along," he says. His biggest interest these days is in embedded microcontrollers and especially connecting those microcontrollers, again exploiting Metcalfe's Law. He defines a layer of embedded networking below Wi-Fi, Bluetooth, or cellular in which machines communicate with machines.

Ember, with ZigBee, is a player in embedded networking, and Metcalfe believes that wireless is the key. "You can't run cable very well among lots of embedded micros; you had better do it wireless; hence, the development of CMOS radios," he says, claiming that vendors annually ship approximately 10 billion microcontrollers. "That number is going to go up thanks to the arrival of embedded networking," he says.

Not surprisingly, with his semiconductor and EDA background, Rhines believes the future of innovation is in chip design and the number of design engineers that have access to chip-design technology. He points out that, when ASICs debuted, custom-IC folks laughed off the upstart technology because ASICs would offer lower performance and be bigger than custom ICs. During the custom-IC era, however, only a few thousand IC designers existed, he points out. ASIC technology allowed tens of thousands of engineers to design ICs. Rhines claims that FPGAs or another disruptive technology will result in the emergence of 5 million chip designers by 2010 to 2020. "In itself, that should generate more innovation," he says. EDN

Digital Power

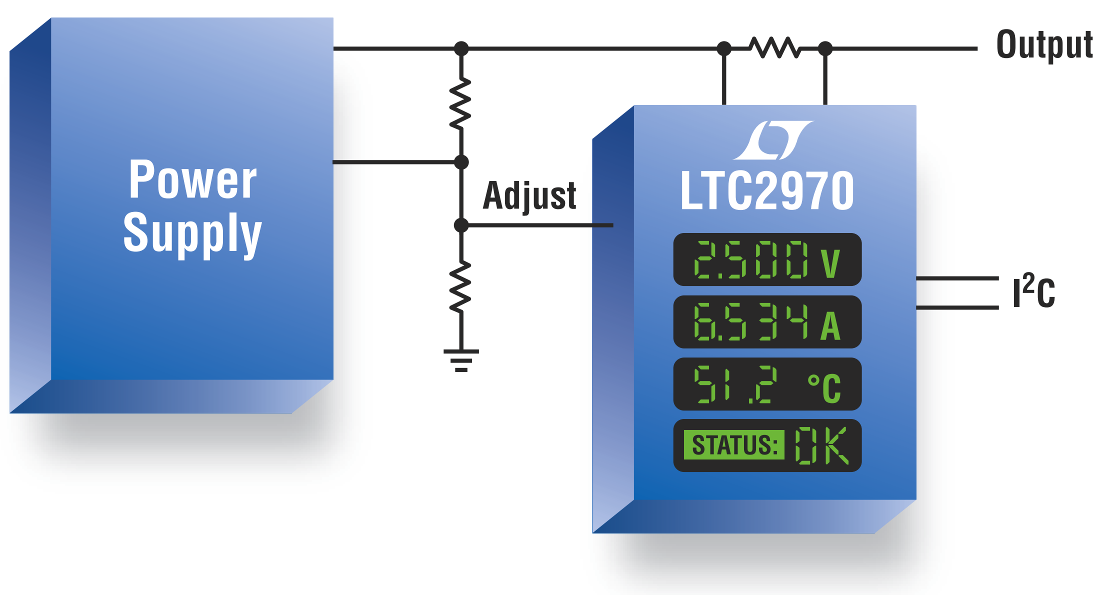

Precisely Monitor and Adjust Your Power Supply with I^2C Interface

The LTC®2970 combines an extremely accurate reference, 14-bit ΔΣ ADC and 8-bit DACs to monitor and control the output voltages of two DC/DC converters. The ADC achieves better than ±0.50% total unadjusted error over temperature. Voltage programming, accurate voltage margining and voltage or current measurement functions are at your command over the I^2C bus. You can even adjust for temperature variations using the internal temperature sensor.

Features

- 14-Bit ΔΣ ADC with On-Chip Reference
- Less Than ±0.5% Total Unadjusted Error Over Temp.
- Dual, 8-Bit DACs
- Automatic Linear Voltage Servo
- I^2C Bus Interface (SMBus Compatible)
- Extensive, User-Configurable Fault Monitoring
- On-Chip Temperature Sensor
- 24-Lead 4mm x 5mm QFN Package
- LTC2970: $3.99 each in 1k Qty.

Functional Block Diagram

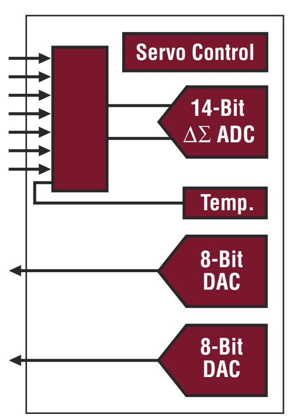

Info & Free Samples

www.linear.com/2970
Literature: 1-800-4-LINEAR
Support: 408-432-1900

LEAD-FREE RoHS COMPLIANT

LINEAR TECHNOLOGY

DESIGN NOTES

Low Noise Amplifiers for Small and Large Area Photodiodes

Design Note 399

Glen Brisebois

Introduction

Photodiodes can be broken into two categories: large area photodiodes with their attendant high capacitance (30pF to 3000pF) and smaller area photodiodes with relatively low capacitance (10pF or less). For optimal signal-to-noise performance, a transimpedance amplifier consisting of an inverting op amp and a feedback resistor is most commonly used to convert the photodiode current into voltage. In low noise amplifier design, large area photodiode amplifiers require more attention to reducing op amp input voltage noise, while small area photodiode amplifiers require more attention to reducing op amp input current noise and parasitic capacitances.

Small Area Photodiode Amplifiers

Small area photodiodes have very low capacitance, typically under 10pF and some even below 1pF. Their low capacitance makes them more approximate current sources to higher frequencies than large area photodiodes. One of the challenges of small area photodiode amplifier design is to maintain low input capacitance so that voltage noise does not become an issue and current noise dominates.

Figure 1 shows a simple small area photodiode amplifier using the LTC6244. The input capacitance of the amplifier consists of C_{DM} (the amplifier's differential mode capacitance) and one C_{CM} (the common mode capacitance at the amplifier's - input only), or about 6pF total. The small photodiode has 1.8pF, so the input capacitance of the amplifier is dominating the capacitance. The small feedback capacitor is an actual component (AVX Accu-F series), but it is also in parallel with the op amp lead, resistor and parasitic capacitances, so the total real feedback capacitance is probably about 0.4pF. This is important because feedback capacitance sets the compensation of the circuit and, with op amp gain bandwidth, the circuit bandwidth. This particular design has a bandwidth of 350kHz, with an output noise of $120\mu V_{RMS}$ measured over that bandwidth.

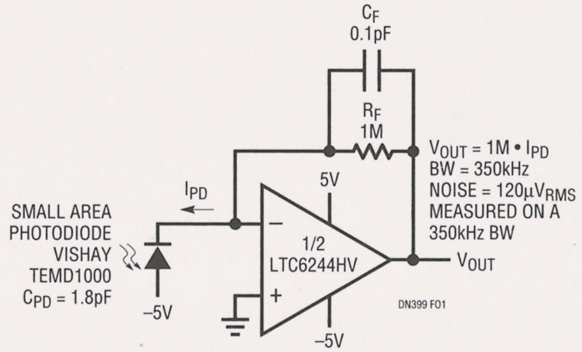

Figure 1. Transimpedance Amplifier for Small Area Photodiode

Large Area Photodiode Amplifiers

Figure 2a shows a simple large area photodiode amplifier. The capacitance of the photodiode is 3650pF (nominally 3000pF), and this has a significant effect on the noise performance of the circuit. For example, the photodiode capacitance at 10kHz equates to an impedance of 4.36kΩ, so the op amp circuit with 1MΩ feedback has a noise gain of NG = 1 + 1M/4.36k = 230 at that frequency. Therefore, the LTC6244 input voltage noise gets to the output as NG • $7.8nV/\sqrt{Hz} = 1800nV/\sqrt{Hz}$, and this can clearly be seen in the circuit's output noise spectrum in Figure 2b. Note that we have not yet accounted for the op amp current noise, or for the $130nV/\sqrt{Hz}$ of the gain resistor, but these are obviously trivial compared to the op amp voltage noise and the noise gain. For reference, the DC output offset of this circuit is about 100µV, bandwidth is 52kHz, and the total noise was measured at $1.7mV_{RMS}$ on a 100kHz measurement bandwidth.

An improvement to this circuit is shown in Figure 3a, where the large diode capacitance is bootstrapped by a $1nV/\sqrt{Hz}$ JFET. This depletion JFET has a VGS of about –0.5V, so that RBIAS forces it to operate at just over 1mA of drain current. Connected as shown, the photodiode has a reverse bias of one VGS, so its capacitance will be slightly lower than in the previous case (measured 2640pF), but the most drastic effects are due to the bootstrapping. Figure 3b shows the output noise of the new circuit. Noise at 10kHz is now 220nV/$\sqrt{Hz}$, and the 130nV/$\sqrt{Hz}$ noise thermal noise floor of the 1M feedback resistor is discernible at low frequencies. What has happened is that the 7.8nV/$\sqrt{Hz}$ of the op amp has been effectively replaced by the 1nV/$\sqrt{Hz}$ of the JFET. This is because the 1M feedback resistor is no longer "looking back" into the large photodiode capacitance. It is instead looking back into a JFET gate capacitance, an op amp input capacitance, and some parasitics, approximately 10pF total. The large photodiode capacitance is across the gate-source voltage of the low noise JFET. Doing a sample calculation at 10kHz as before, the photodiode capacitance looks like 6kΩ, so the 1nV/$\sqrt{Hz}$ of the JFET creates a current noise of 1nV/6k = 167fA/$\sqrt{Hz}$. This current noise necessarily flows through the 1M feedback resistor, and so appears as 167nV/$\sqrt{Hz}$ at the output. Adding the 130nV/$\sqrt{Hz}$ of the resistor (RMS wise) gives a total calculated noise density of 210nV/$\sqrt{Hz}$, agreeing well with the measured noise of Figure 3b. Another drastic improvement is in bandwidth, now over 350kHz, as the bootstrap enabled a reduction of the compensating feedback capacitance. Note that the bootstrap does not affect the DC accuracy of the amplifier, except by adding a few picoamps of gate current.

For more details on photodiode circuits, download the LTC6244 data sheet. To discuss your particular amplifier requirement, contact the author at the number below.

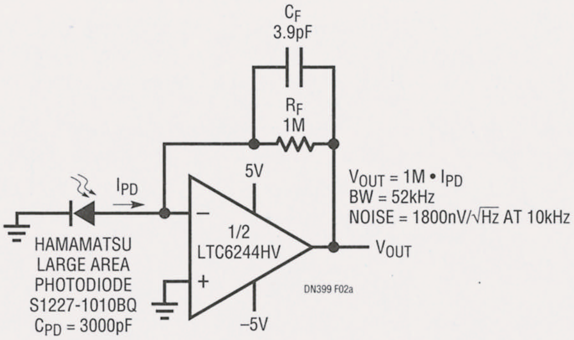

Figure 2a. Large Area Photodiode Transimpedance Amp

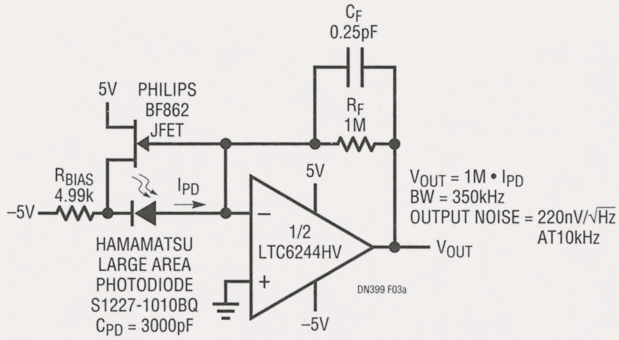

Figure 3a. Large Area Diode Bootstrapping

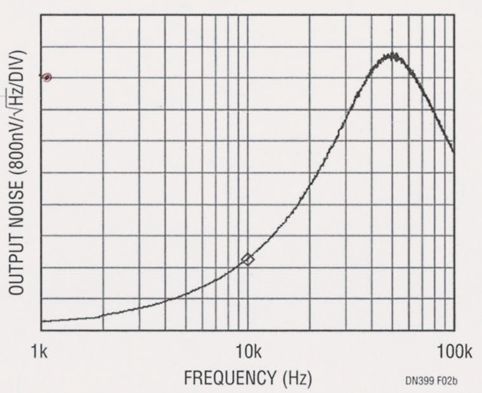

Figure 2b. Output Noise Spectral Density of the Circuit of Figure 2a. At 10kHz, the 1800nV/$\sqrt{Hz}$ Output Noise is Due Almost Entirely to the 7.8nV Voltage Noise of the LTC6244 and the High Noise Gain of the 1M Feedback Resistor Looking into the High Photodiode Capacitance

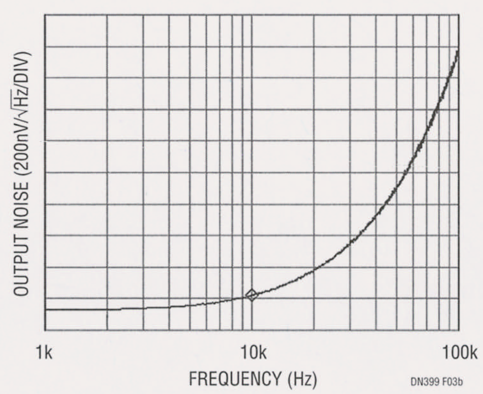

Figure 3b: Output Noise Spectral Density of Figure 3a. The Simple JFET Bootstrap Improves Noise (and Bandwidth) Drastically. Noise Density at 10kHz is Now 220nV/$\sqrt{Hz}$, About a 8.2x Reduction. This is Mostly Due to the Bootstrap Effect of Swapping the 1nV/$\sqrt{Hz}$ of the JFET for the 7.8nV/$\sqrt{Hz}$ of the Op Amp

Data Sheet Download

www.linear.com

For applications help,
call (408) 432-1900, Ext. 3755

Linear Technology Corporation
1630 McCarthy Blvd., Milpitas, CA 95035-7417
(408) 432-1900 • FAX: (408) 434-0507 • www.linear.com

dn399f LT/TP 0906 409K • PRINTED IN THE USA

LINEAR TECHNOLOGY

High-impedance buffer amplifier's input includes ESD protection

Eugene Palatnik, Waukesha, WI

Certain measurement applications, such as for pH (acidity) and bio-potentials, require a high-impedance buffer amplifier. Although several semiconductor manufacturers offer amplifier ICs featuring low bias and offset-input currents, attaching a sensor cable to an amplifier circuit can inflict damage from ESD (electrostatic discharge). **Figure 1** shows one

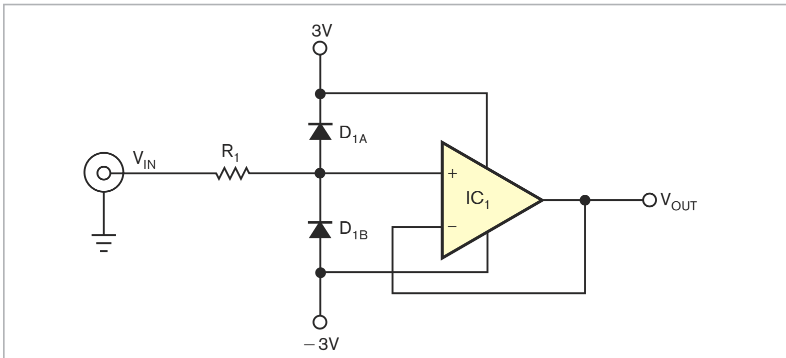

Figure 1 In a conventional ESD-suppression circuit, diodes clamp an amplifier's input voltage to its power-supply rails but introduce unwanted leakage currents.

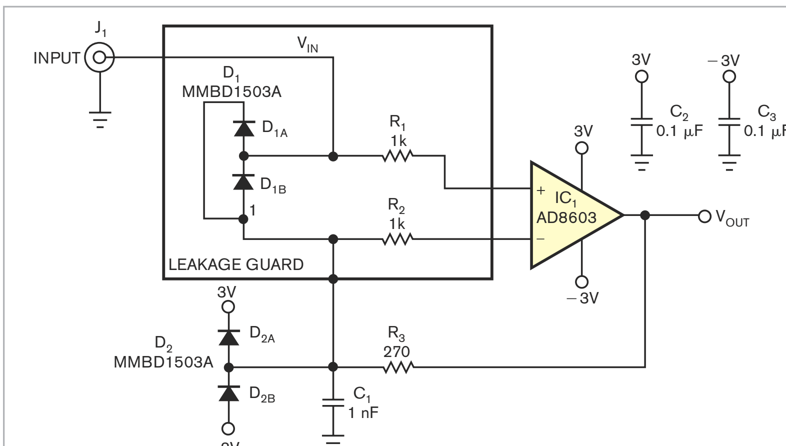

Figure 2 In this alternative design, voltage across both halves of D_1 normally approaches 0V and introduces no leakage currents. During an ESD event, both D_1 and D_2 conduct to protect IC_1's inputs.

DIs Inside

▶What are your design problems and solutions? Publish them here and receive $150! Send your Design Ideas to edndesignideas@reedbusiness.com.

unsatisfactory approach to ESD protection. Resistor R_1 limits an ESD event's discharge current, and diodes D_{1A} and D_{1B} clamp amplifier IC_1's input to its power-supply rails. Unfortunately, when shunting a pH sensor's 400-MΩ input impedance, even low-leakage diodes, such as Fairchild Semiconductor's (www.fairchildsemi.com) MMBD-1503A, introduce significant offset voltages.

The circuit in **Figure 2** offers an alternative approach. An Analog Devices (www.analog.com) low-input-bias, low-offset-current AD8603 amplifier, IC_1, serves as a unity-gain input buffer. For any normal input, the circuit's output voltage, V_{OUT}, equals its input voltage, V_{IN}. Thus, the voltage across ESD-protection diode D_{1A} or D_{1B} approaches 0V, and neither diode's leakage current affects the sensor's output signal. Depending on the polarity of an ESD event you apply to the circuit's input connector, its high-voltage spike discharges through diode D_{1A} or D_{1B} into the positive or the negative power-supply rail. Capacitor C_1 acts as an intermediate "charge reservoir" that slows the ESD spike's rate of rise and protects IC_1's output stage from latching until diode D_{2A} or D_{2B} begins diversion of the ESD transient into the positive or the negative supply rail. In effect, C_1 compensates for D_1's parasitic capacitance. Resistor R_3 allows IC_1 to drive the capacitive load that C_1 presents without going into oscillation.

During an ESD event, both D_1 and D_2 can conduct, but the voltage at V_{IN} exceeds the power-supply-rail voltage by only two forward-biased diode voltage drops. Resistors R_1 and R_2 limit the amplifier input's currents below the manufacturer's recommended 5-mA maximum rating.

When packaging the circuit, pay special attention to the pc board's layout. Imperfections in the board's dielectric properties can provide parasitic-leakage-current paths. Adding copper traces on both sides of the board to form guard rings around the circuit's high-impedance nodes diverts leakage currents (**Figure 3**).EDN

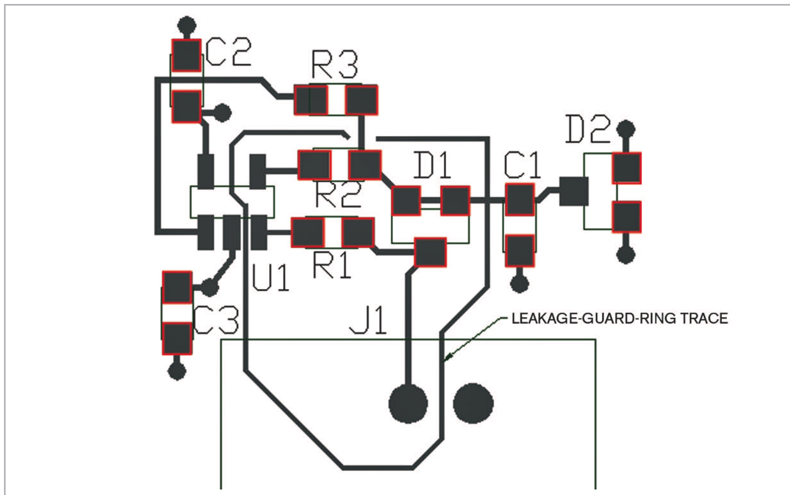

Figure 3 For best performance, place copper traces around the amplifier's high-impedance points to intercept leakage currents.

Composite-VGA encoder/decoder eases display upgrade

Werner Schwiering, Joystick Scoring Ltd, Whitby, ON, Canada

An older computer system fed RGB video and composite-synchronization signals through four 75Ω coaxial cables to an RGB color monitor 150 feet away. To upgrade it, the replacement VGA video cards could directly drive the 75Ω loads that the VGA monitors' internal terminations presented. However, the VGA standard uses separate horizontal and vertical positive-going synchronization signals. Adding an extra coaxial cable to the original cables to carry the separate synchronization signals presented a difficult and expensive proposition. An obvious solution would be to combine the separate synchronization signals into a composite format.

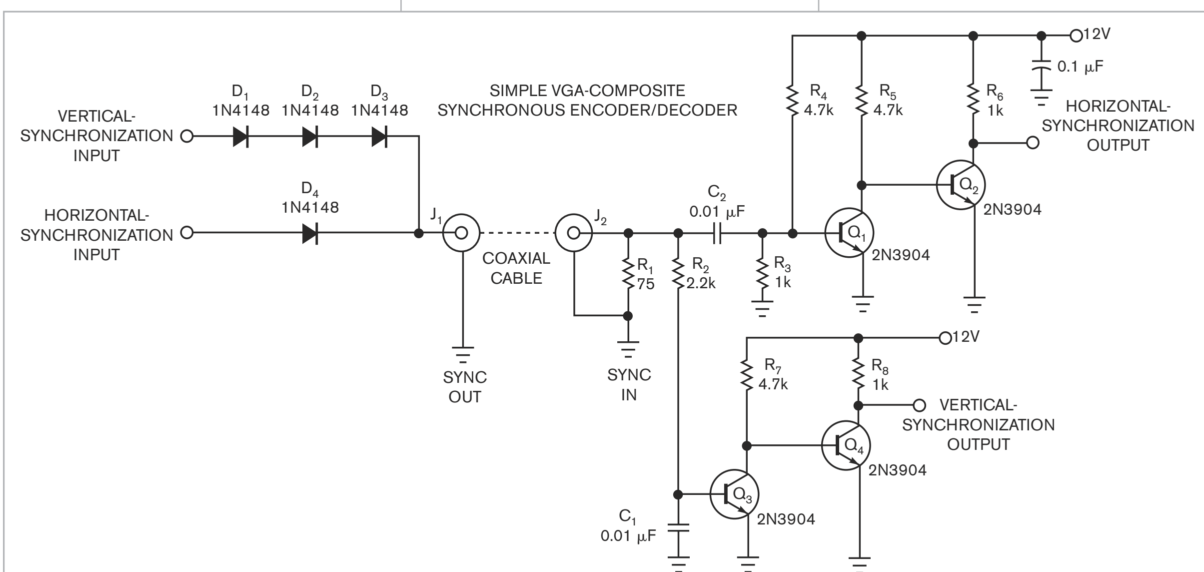

Figure 1 The synchronization-pulse combiner and recovery circuits comprise readily available and inexpensive components.

Cool & Compact 2A Charging

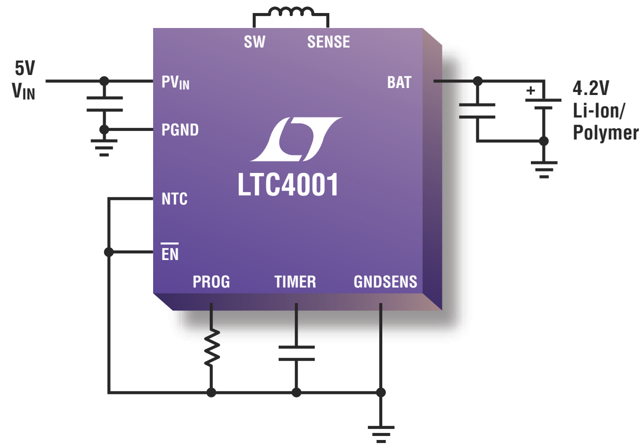

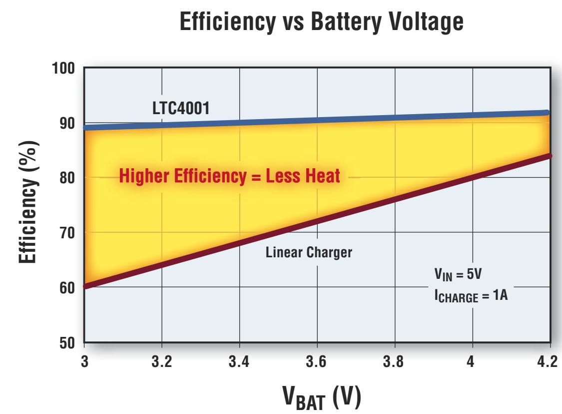

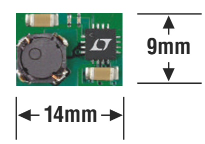

Actual Size
Demo Circuit

Less Heat & Higher Current than Linear Chargers

Is too much heat from your high current linear regulator-based battery charger killing your system? If so, try our new LTC®4001, a monolithic synchronous switching step-down Li-Ion/Polymer charger, capable of delivering 2A of charge current with over 90% efficiency. The highly integrated LTC4001 includes on-chip MOSFETs and current sense resistor. With its 4mm x 4mm QFN package and 1.5MHz switching frequency operation, only five external components are needed for a simple, compact and cost-effective solution.

Features

- Low Power Dissipation
- 2A Maximum Charge Current
- No External MOSFETs, Sense Resistor or Blocking Diode Required
- Synchronous Rectification
- 1.5MHz Switching Frequency
- Input Voltage: 4V to 5.5V
- Only 5 External Components Needed
- Low Profile 16-Lead (4mm x 4mm) QFN Package

Power Solutions Brochures

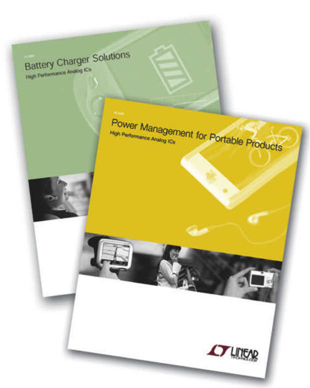

www.linear.com/ad/batsolutions
www.linear.com/ad/portsolutions

Info & Free Samples

www.linear.com/4001
Literature: 1-800-4-LINEAR
Support: 408-432-1900

LEAD-FREE RoHS COMPLIANT

LINEAR TECHNOLOGY

The combiner circuit in **Figure 1** offers simplicity, low cost, and rapid assembly from readily available spare parts.

In operation, two 1N4148 diodes, D_1 and D_2, attenuate the VGA signal's 5V logic-level vertical-synchronization pulses by 1.4V, and diodes D_3 and D_4 form a diode-logical-OR gate to combine the vertical- and horizontal-synchronization pulses. The resultant output signal comprises an approximately 4.3V horizontal-synchronization signal superimposed on a 2.9V vertical-synchronization signal.

At the receiving end, a capacitively coupled highpass filter extracts the horizontal-synchronization signal, and a simple RC (resistor-capacitor) lowpass circuit removes horizontal-synchronization pulses from the directly coupled vertical-synchronization signal. Transistors Q_1 and Q_2 amplify the recovered horizontal-synchronization pulses, and transistors Q_3 and Q_4 amplify the vertical-synchronization pulses. The circuit's resulting outputs consist of clean synchronization pulses that closely approximate those of the original and provide extremely stable synchronization pulses for a VGA monitor operating at 640×480-pixel resolution (**Figure 2**).EDN

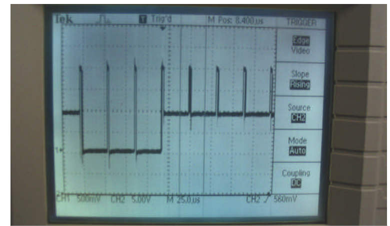

Figure 2 Applying the diode-gated composite-synchronization waveform to a 75Ω load results in clean synchronization pulses.

Solenoid-protection circuit limits duty cycle

Panagiotis Kosioris, Inos Automation Software, Stuttgart, Germany

Several safety-critical solenoids in a laser-measurement system on an automotive-assembly line required protection from internal overheating during normal operation. After a 60-sec activation, the solenoids required 180 sec to cool before their next activation. One apparently straightforward protection circuit would comprise a timer based on a microcontroller, some support components, and a short program written in C++. However, the project would require evaluation and selection of a suitable microcontroller, purchase or rental of a device programmer, and considerable time in programming the microcontroller and evaluating its operational hazards.

As an alternative, I recalled the words of my tutor: "Decrease the number of dangerous components to decrease the risk of danger." A simple analog circuit would be safer, smaller, and easier to maintain. The circuit in **Figure 1** uses a traditional analog method of measuring time: the charge and discharge behavior of a resistance-capacitance circuit.

Figure 2 highlights the circuit's timing components. Capacitor C_2, a tantalum electrolytic with ±10% tolerance, diode D_1, and resistors R_2 and R_5 constitute a double-RC (resistor-

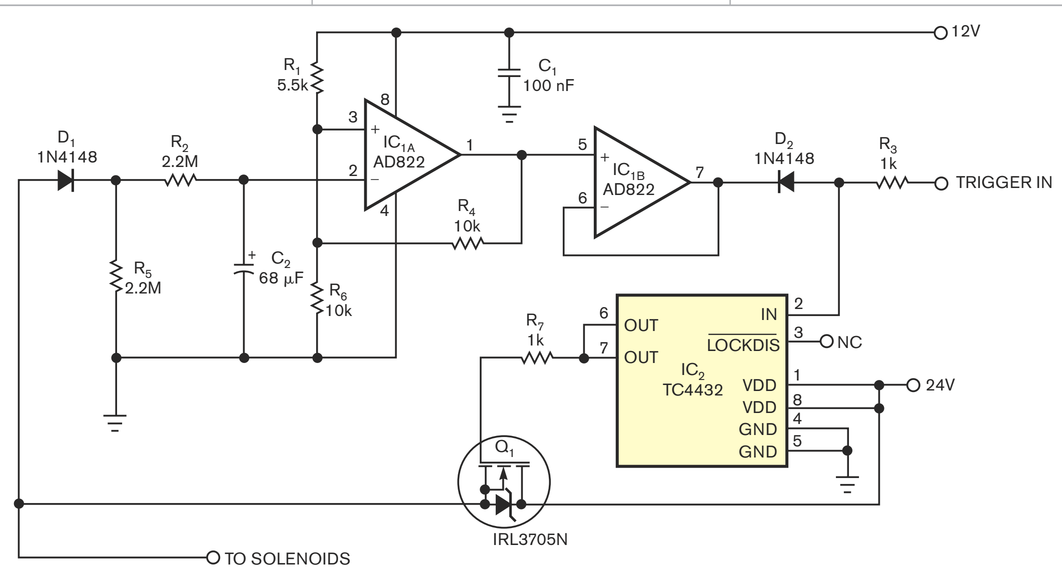

Figure 1 An externally triggered solenoid driver features an analog duty-cycle limiter.

FIRST 3A/6A, STEP-DOWN REGULATORS WITH INTEGRATED SWITCHES AND VID INPUTS

Over 95% Efficiency, Up to 2.4MHz Switching Frequency, Low On-Resistance, and Sync Input—All in a Small 4mm x 4mm TQFN

The MAX8643/MAX8646 step-down regulators provide the highest integration with on-board low on-resistance switches and full overcurrent and thermal protection. Frequencies are programmable and synchronizable from 250kHz to 2.4MHz. A two-pin VID input provides one of nine output preset voltages that can be selected by the user. This eliminates the need for an external resistor-divider, thus providing higher overall accuracy. PWRGD and EN facilitate sequencing, where REFIN can be used for tracking and/or termination, such as in DDR applications. This set of features makes these parts easy to use and suitable for any power-supply application.

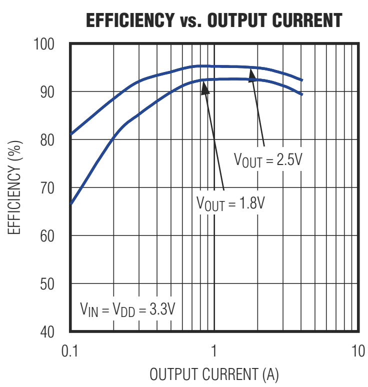

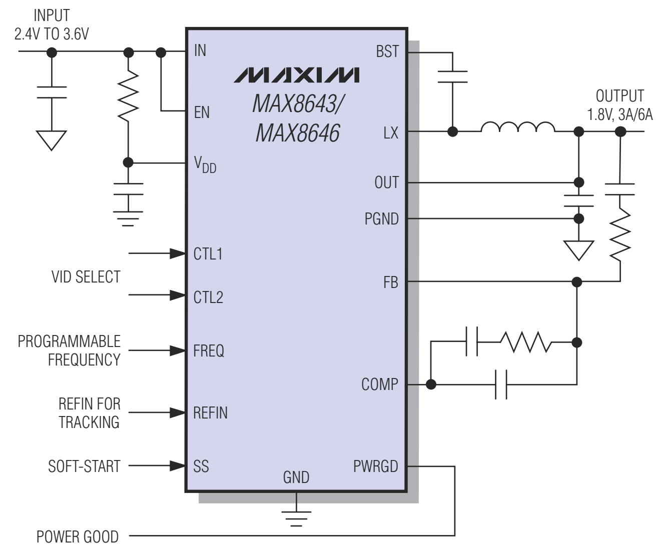

- ♦ VID-Set Output Voltages: 0.6V, 0.7V, 0.8V, 1.0V, 1.2V, 1.5V, 1.8V, 2.0V, and 2.5V
- ♦ Continuous 3A/6A Output (MAX8643/MAX8646)
- ♦ ±1% Output Accuracy Over Temperature
- ♦ Input Supply from 2.3V to 3.6V
- ♦ Soft-Start Reduces Inrush Supply Current
- ♦ 250kHz to 2.4MHz Adjustable Switching or Sync Input
- ♦ All-Ceramic Capacitor Design Reduces Space
- ♦ REFIN for Tracking and DDR Termination
- ♦ PWRGD and EN for Sequencing
- ♦ Available in 24-Pin TQFN Package with Exposed Pad
- ♦ Applications
  - ♦ ASIC/DSP/CPU Core Voltages
  - ♦ DDR Termination
  - ♦ Memory Power Supply
  - ♦ Telecom/Datacom Power Supply
  - ♦ Industrial Power Supply
  - ♦ Point-of-Load Power Supply

www.maxim-ic.com

FREE* Power Supplies Design Guide—*Sent Within 24 Hours!

CALL TOLL FREE 1-800-998-8800 (7:00 a.m.–5:00 p.m. PT)
For a Design Guide or Free Sample

DISTRIBUTION 1-888-MAXIM-IC | ARROW ELECTRONICS, INC. 1-800-777-2776 | electronics marketing 1-800-332-8638

Distributed by Maxim/Dallas Direct!, Arrow, Avnet Electronics Marketing, Digi-Key, and Newark.

capacitor) circuit. During solenoid activation, R_2 provides a charging path for C_2, and diode D_1 prevents C_2 from discharging through the solenoids. When the solenoids are off, the discharge path comprises R_2 plus R_5, which provides a longer time constant. The difference between the two time constants determines the solenoids' activation and recovery periods. A Schmitt trigger designed around one-half of IC_1, an Analog Devices (www.analog.com) AD822 dual operational amplifier, senses the voltage across C_2 and defines the solenoids' cutoff- and turn-on-timing intervals. An intermediate buffer stage, IC_{1B}, drives a Microchip (www.microchip.com) TC4432 MOSFET driver, which in turn controls the gate of Q_1, an N-channel power MOSFET that drives the solenoids from 24V.

When Q_1 switches on, the voltage level across C_2 increases, and, after 60 sec, the output of the Schmitt trigger falls from 12 to 0V. The buffer stage drives the cathode of diode D_2 to 0V. The voltage at D_2's anode reaches 0.7V and is insufficient to trigger MOSFET-driver IC_2. Q_1 now switches off, removing supply voltage from the solenoids and reverse-biasing diode D_1. Capacitor C_2 starts to discharge through R_2 and R_5, and the input voltage you apply to the Schmitt trigger falls at a slower rate than during the charging interval. After 180 sec, the Schmitt trigger's output rises to 12V, and the circuit awaits arrival of another external trigger pulse through resistor R_3. EDN

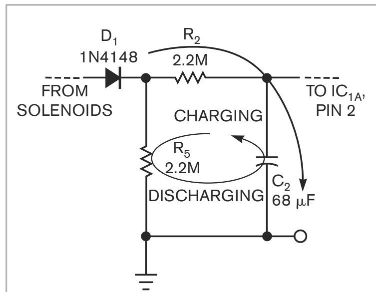

Figure 2 A resistance-capacitance circuit determines on- and off-time intervals.

SPST pushbutton switch combines power-control, user-input functions

Eugene Kaplounovski, Vancouver, BC, Canada

This Design Idea describes an enhancement to a previous one (**Reference 1**). The circuit in **Figure 1** uses a normally open SPST pushbutton switch, S_1, instead of the SPDT switch that the original design required. You can use a membrane switch to significantly simplify the industrial design of the device and enhance its ergonomics. In addition, this circuit slightly reduces the current drain in active mode by eliminating current flow through the unactuated switch.

In standby mode, MOSFET Q_1 remains off and consumes less than 1 µA of leakage current from the battery. Pressing switch S_1 turns on Q_1 by pulling its gate to ground through diode D_1. Voltage regulator IC_1 turns on and supplies power to microcontroller IC_2. The microcontroller boots up and asserts its P1.1 output high, turning on transistor Q_2 and latching on the system's power to allow release of S_1. Meanwhile, resistor R_3 pulls the microcontroller's input, P1.2, to V_{CC}. Pressing the switch a second time pulls the microcontroller's P1.2 input low through diode D_2 and signals the button-pressed event to the firmware. After completing its program, the microcontroller asserts its output P1.1 low to turn off Q_2 and, consequently, Q_1, removing power from the system until the user presses S_1 and restarts the process.

When selecting components, ensure that Q_1's gate-source breakdown voltage exceeds the highest possible input voltage; otherwise, use a zener diode to limit Q_1's applied gate-source voltage. You can omit Q_1 if voltage regulator IC_1 includes an on/off-control pin. To replace Q_1 with a different power-switching device, such as an NPN bipolar transistor or a relay, specify Q_2 to provide the control current that the switching device requires. To further reduce the circuit's component count, replace diodes D_1

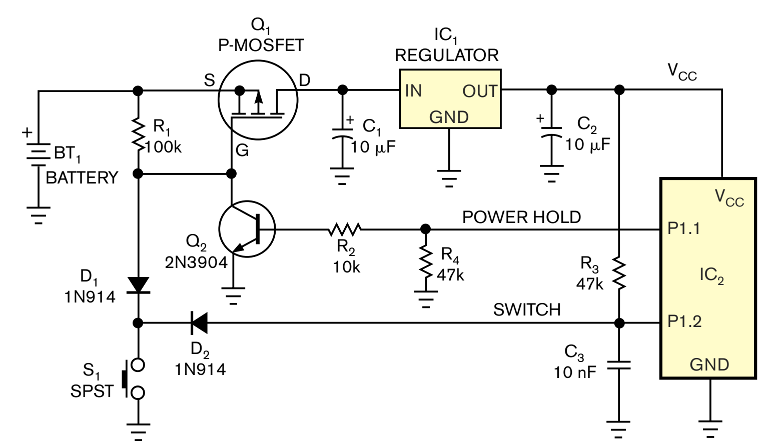

Figure 1 One switch can provide power control and user inputs to a microcontroller-based system.

ONLY STRATUM 3-/3E-COMPLIANT TIMING-CARD IC WITH INTEGRATED DS1/E1/2048kHz TRANSCEIVERS

Ideal for Flexible Synchronization and Timing-Card Applications

The DS3100 integrates all timing-card functions into a single-chip solution. A typical SONET/SDH timing-card solution has multiple discrete analog and digital components, including PLLs, ASICs, FPGAs, microprocessors, and DS1/E1 transceivers. These solutions tend to be very expensive and require extensive DSP and PLL design knowledge. The DS3100 highly integrated, single-chip solution saves design time, board space, and cost. The DS3101 is available without the BITS/SSU transceivers.

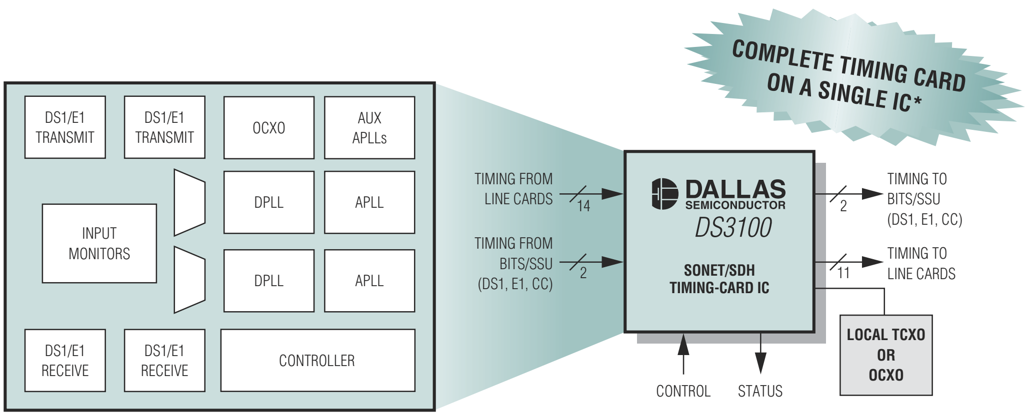

Fourteen Input Clocks

- CMOS/TTL and Differential LVDS/PECL Inputs Support Rates Up to 155.52MHz
- Two 64kHz Composite Clock Receivers
- Supports 2kHz, 4kHz, and Any Multiple of 8kHz Up to 155.52MHz

Eleven Output Clocks

- CMOS/TTL and Differential LVDS Outputs Support Rates Up to 311.04MHz
- Composite Clock and Sync Pulse Outputs
- Supports 2kHz, 8kHz, NxDS1, NxE1, Nx19.44MHz, DS3, E3, and 125MHz

General Features

- Stratum 3E Holdover Accuracy
- Hitless Reference Switching on Loss of Input
- Phase Buildout and Transient Absorption
- 125MHz Support for Timing Over GbE

Two BITS/SSU Transceivers (DS3100)

- DS1, E1, 2048kHz, and 6312kHz Timing Signals
- Insert and Extract SSM Messages (DS1, E1)

Part	Package	Temp Range (°C)	Price† ($)
DS3100GN#	256-CSBGA	-40 to +85	60.94
DS3101GN#	256-CSBGA	-40 to +85	49.40

*Requires a local TCXO or OCXO.
†10k-up recommended resale, FOB USA. Prices provided are for design guidance and are for the lowest grade, commercial temperature parts. International prices will differ due to local duties, taxes, and exchange rates. Prices are subject to change. Not all packages are offered in 1k increments, and some may require minimum order quantities.

DALLAS SEMICONDUCTOR MAXIM

www.maxim-ic.com/DS3100

FREE* Communications Design Guide—*Sent Within 24 Hours!

CALL TOLL FREE 1-800-998-8800 (7:00 a.m.–5:00 p.m. PT)
For a Design Guide or Free Sample

MAXIM/DALLAS DIRECT! DISTRIBUTION 1-888-MAXIM-IC
ARROW ELECTRONICS, INC. 1-800-777-2776
AVNET electronics marketing 1-800-332-8638

Distributed by Maxim/Dallas Direct!, Arrow, Avnet Electronics Marketing, Digi-Key, and Newark.

and D_2 with a suitable common-cathode dual-diode array, such as the BAV70. Omit resistor R_3 if IC_2 includes built-in pullup resistors, as do many modern microcontrollers. EDN

REFERENCE

1 Hageman, Steve, "Single switch serves dual duty in small, microcontroller-based system," *EDN*, March 30, 2006, pg 96, www.edn.com/article/CA6317068.

Electronic circuit replaces mechanical push-push switch

Donald Schelle, Maxim Integrated Products Inc, Sunnyvale, CA

Mechanical push-pushbutton switches (also known as alternate-action or push-on/push-off switches) can be bulky and expensive. As an alternative, an electronic version uses a cheaper, NO (normally open), momentary-on switch (**Figure 1**). A supervisory microprocessor, IC_1, serves as a combination switch debouncer and intelligent controller. Applying power holds IC_1's $\overline{LBO}$ output (Pin 4) low, which in turn resets flip-flop IC_2's output to a logic-low state (off) (**Figure 2**). Pressing the NO momentary-contact switch, S_1, evokes a pulse from the $\overline{RESET}$ output (IC_1, Pin 5), which triggers IC_2's CK input (Pin 1) and toggles IC_2's output to a logic-high state (on). Pressing the switch a second time triggers another $\overline{RESET}$ pulse that toggles flip-flop IC_2's output to a logic-low state (off).

You can add an optional watchdog timer, IC_3, to reset IC_2's output to the logic-low state after a user-selectable interval as long as 60 sec. You can select shorter reset times using IC_3's programming pins: SET0, SET1, and SET2. The entire circuit costs about \$2 (1000) and occupies a pc-board area that's no larger than its mechanical counterpart. EDN

MORE AT EDN.COM

For more Design Ideas, visit www.edn.com/designideas.

For our best entries, go to www.edn.com/bestofdesignideas.

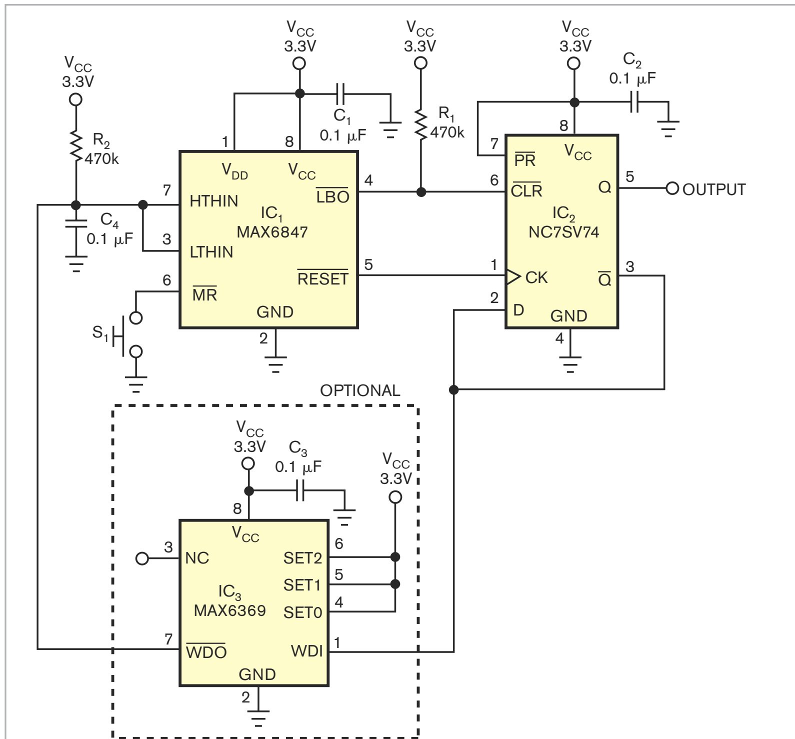

Figure 1 This simple electronic circuit uses a momentary-contact pushbutton switch, S_1, to replace a more expensive mechanical push-on/push-off switch.

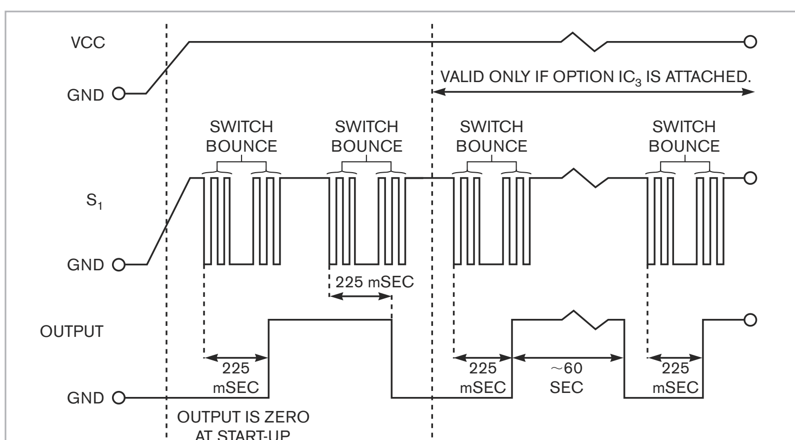

Figure 2 Repeatedly pressing the circuit's momentary-contact switch toggles the circuit's output on and off. After a preselected interval, an optional watchdog timer resets the output to the logic-low state.

PUSHBUTTON AUDIO PROCESSOR WITH SUBWOOFER OUTPUT AND AMBIENCE CONTROL ELIMINATES MICROCONTROLLER

The MAX5406 provides a complete audio-processing solution without the need for a microcontroller (µC). The debounced pushbutton interface controls volume, balance, and tone. The MAX5406 operates with momentary contact switches for docking-station and desktop-speaker applications.

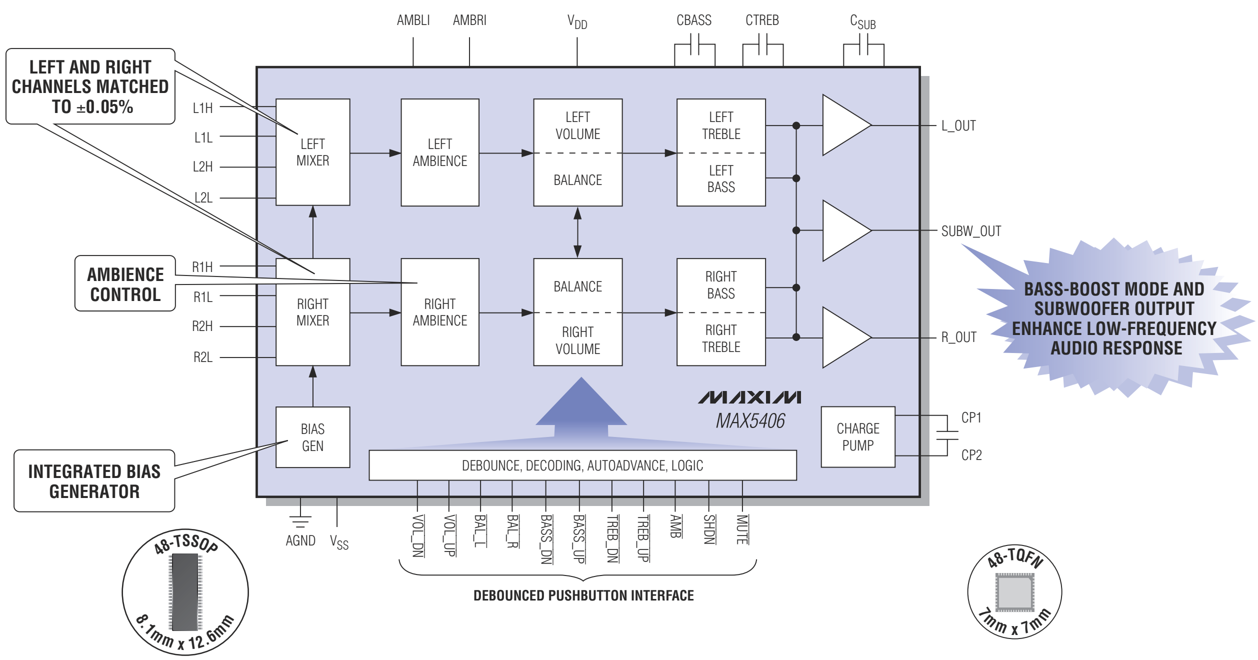

- Ambience Control Enhances Separation of Left and Right Outputs
- Pseudostereo Feature Approximates Stereo Sound for Monophonic Signals
- SmartWiper™ Circuitry Automatically Controls Rate of Wiper Advance
- 32-Tap Volume Control in 2dB Steps
- 0 to +2.7V, 0 to +5.5V, or ±2.7V Operating Voltage
- Clickless Switching and Power-Up
- Mute < -90dB, Channel Isolation > -70dB (typ)
- Differential Inputs Accept Unipolar or Bipolar Signals
- Shutdown Stores Last Wiper Settings, Reduces Supply Current to 0.5µA (typ)
- THD < 0.05% (max), Output Noise < $35\mu V_{RMS}$ (max)
- Passive-Filtered Analog Inputs Reject RF Noise
- µP-Compatible Interface

www.maxim-ic.com

FREE* Digital Potentiometers Design Guide—*Sent Within 24 Hours!

CALL TOLL FREE 1-800-998-8800 (7:00 a.m.–5:00 p.m. PT)
For a Design Guide or Free Sample

1-888-MAXIM-IC

ARROW ELECTRONICS, INC.
1-800-777-2776

electronics marketing
1-800-332-8638

Distributed by Maxim/Dallas Direct!, Arrow, Avnet Electronics Marketing, Digi-Key, and Newark.

productroundup

AMPLIFIERS, OSCILLATORS, AND MIXERS

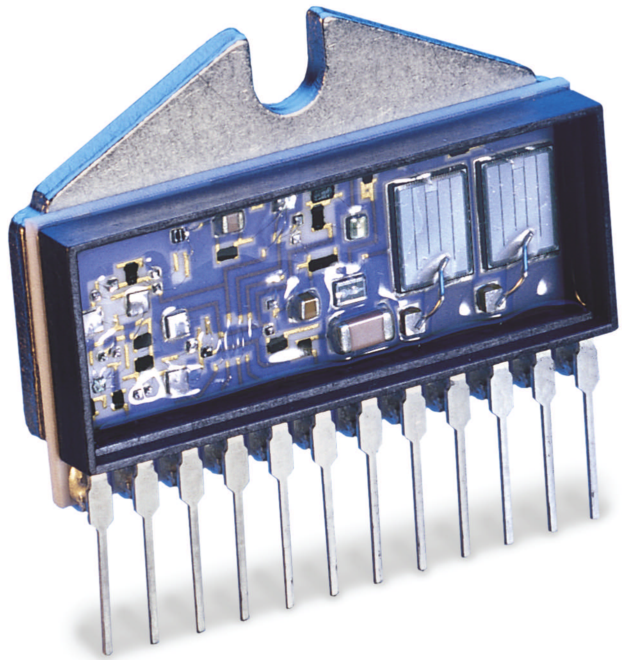

High-voltage amplifier supplies 8A

The hybrid PA93 amplifier/high-voltage output transistor combines a 400V supply voltage with an 8A output current. Features include a 10-mA quiescent current, a 12-MHz bandwidth, and a 4-pF input capacitance. The part's SIP is electrically isolating, meaning you can mount it directly to heat sinks with Grafoil or another low-resistance thermal compound. Thermal resistance is 0.7°C/W ac, junction to case, and the part works to 85°C. Peripherals include a 10-mV maximum offset and a typical 200-pA input-bias current. Aiming at piezo positioning, electrostatic transducers, and programmable power supplies to 390V, the PA93 costs $98.90 (1000).
Apex Microtechnology, www.apexmicrotech.com

Rail-to-rail input amplifier has no crossover distortion

The OPA365 rail-to-rail input amplifier achieves rail-to-rail input swings using a charge pump to create a higher internal voltage for the input stage, in contrast to the more common method of using a dual-input pair. The amplifier has a 50-MHz bandwidth and a 0.0006% THD+N to 100 kHz. At higher than that value, the charge pump contributes a fixed-frequency noise that you can filter out. Voltages include a 2.2 to 5.5V supply voltage and a 200-µV maximum-offset voltage. Settling time is 300 nsec to 0.01%. This part targets multiplexers and ADCs. It settles quickly due to its 50-MHz bandwidth and can input large rail-to-rail signals into the converter with less distortion than a dual-input rail-to-rail amplifier. The OPA365 comes in an SOT-23-5 package, costs 95 cents (1000), and will be available in late October 2006.

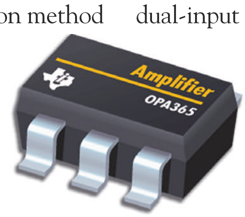

Texas Instruments, www.ti.com

Low-noise CMOS amplifier runs input pairs at higher currents

The low-noise LT6244 amplifier features the low 1-pA input-bias current of a CMOS or JFET amp but has $8\ nV/\sqrt{Hz}$ and $0.5\ fA/\sqrt{Hz}$ noise characteristics. It achieves these noise characteristics by running the input pair at higher currents, offering a 6.25-mA quiescent current. Features include a 1/f corner at 200 Hz, a 40V/µsec slew rate, a 120-dB voltage gain, and a 2.1-pF input capacitance. The LT6244 has an input offset of 100 µV maximum and a 300 µV maximum over temperature. Bandwidth is 50 MHz. The LT6244 uses a 6.35-mA supply current. It targets photodiodes and other transconductance-signal chains, as well as any application with high source impedance. Operating from 2.8 to 12V power supplies, the LT6244 costs $1.65 (1000).
Linear Technology, www.linear.com

800- to 1000-MHz mixer features a 2-dB conversion gain

The MAX9982 mixer has a 26.8-dBm input-third-intercept point. The device features a 70- to 170-MHz intermediate-frequency range, a 725- to 1085-MHz range form, and a 2-dB conversion gain. Using a silicon-germanium process, the mixer works off 5V and sports a 12-dB noise figure. The MAX9982 also features an integrated LO (local-oscillator) switch. The mixer costs $7.90 (1000).
Maxim Integrated Products, www.maxim-ic.com

Chopper amplifier features 150-dB gain

The CS3003 chopper amplifier has 10-µV-maximum offset voltage with 0.05 µA/°C drift. The part operates

productroundup

AMPLIFIERS, OSCILLATORS, AND MIXERS

from 2.7 to 5V and has a 150-dB voltage gain. Features include a 17-nV/$\sqrt{\text{Hz}}$ noise flat beyond 0.08 Hz, a 1.5-pF input capacitance, a 1-MHz bandwidth, and a 0.25V/µsec slew rate. A 150-kHz chopper frequency makes it easy to filter out, even with a single-pole filter. This amplifier features rail-to-rail inputs and outputs. It swings to 10 µV of the rails with a 200-kΩ load. Available in SOIC-8 and leadless QFN packages, the device suits measurements of slow-moving inputs, such as gas detectors, load cells, and thermopile signals. The CS3003 costs 99 cents (1000).

Cirrus Logic, www.cirrus.com

JFET amplifier drives 4000 pF

Aiming at IC, memory, and logic test equipment, the uPC835 JFET amplifier features a 2.8-MHz bandwidth and works on ±16V rails. The device can drive a 4000-pF capacitive load and has a 9V/µsec slew rate. The amp directly targets Texas Instruments' updated TLE-2062, which can drive 900 pF of capacitive load. Available in a TSSOP-8 package, the uPC835 costs $1.20.

NEC Electronics America, www.am.necel.com

Ultralow-noise amplifier is unity-gain-stable

The ADA4899-1 ultralow-noise amplifier is unity-gain-stable and suits use as a buffer. The device features a 250-MHz bandwidth; a 310V/µsec slew rate; and a 600-MHz, −3-dB bandwidth. The vendor's proprietary pinout provides output on both sides of the part to minimize distortion. This eight-pin part does not interchange with conventional pinout parts. Features include a −117-dB distortion at 1 MHz; a low-noise, typical 14.7-mA supply current; and an active-low disable pin. The device has a 117-dBc spurious-free dynamic range at 1 MHz. In addition, the device offers a 0.23-mV offset and a 1-µA input-bias current. The ADA-4800-1 suits use as an analog-to-digital driver, an IF or baseband amplifier, or a DAC buffer. It also targets instrumentation and optical electronics. The amplifier costs $1.89 (1000).

Analog Devices, www.analog.com

Programmable crystal oscillators come in full and half DIPs

Adding to the vendor's programmable-clock-oscillator line, the ECS-UPO-5X7 series comes in a 5×7-mm SMD package, as well as a smaller 3×5-mm package. Full- and half-DIP package series are also available. The programmable crystal oscillators come in 3.3 and 5V versions. The 3.3V parts have a programming range of 1 to 125 MHz, and you can program the 5V parts from 1 to 150 MHz. Stability is 50 ppm over a 0 to 70°C temperature range, with a 100-ppm part series also available. The ECS-UPO-5X7 costs $1.99 (1000). Digi-Key sells the programmer for $1521, or the vendor or distributor can preprogram the units.

ECS International, www.ecsxtal.com

INTEGRATED CIRCUITS

DAC provides high SFDR using minimal power

Providing a 1200-MHz output frequency, the 2.3G-sample/sec MAX19692 DAC features a 68-dBc SFDR (spurious-free dynamic range), claiming a

14-dB improvement over competing devices at the high-output frequency. Operating on 1.8 and 3.3V supplies, this current-steering DAC includes a 4-to-1 multiplexed LVDS (low-voltage-differential-signaling) input and a 12-bit converter core. A 1.5G-sample/sec update rate allows the machine to consume 950 mW of power. Additional features include an internal 50Ω differential-output termination as well as synthesizing, high-frequency and wideband signals in multiple Nyquist zones. The device suits industrial use with a −40 to +85°C temperature range and comes in an 11×11-mm CSBGA-169 package. The MAX19692 is in full production, with evaluation kits available.

Maxim Integrated Products, www.maxim-ic.com

Ceramic amplifier uses bond-wire compensation

Operating at 2.3 to 6V, the MCP-1727 1.5A low-dropout amplifier is available in an adjustable-output version and a fixed-output version with a sense pin compensating for losses in the bond wire and traces. This ceramic output-capacitor-stable amplifier includes a 300-mV dropout voltage, a 140-μA quiescent current, a power-good output with user-programmable delay, and a shutdown feature. Available in an SOIC-8 package or a 3×3-mm DFN-8 package, the MCP1727 costs $1.26 and $1.34 (10,000), respectively.

Microchip Technology, www.microchip.com

Amplifier provides a low A-weighted noise floor

Targeting portable applications requiring high performance in limited board space, the TPA203xD1 Class D audio-power amplifier provides 88% efficiency and drives 2.75W from a 5V power supply into 4Ω with a 27-μV-rms A-weighted noise floor. Featuring a −69 dB CMRR (common-mode-rejection ratio) and a PSRR (power-supply-rejection ratio) at 217 Hz of −73 dB, the device is pin-compatible with the TPA2010D1. Available in a 1.5×1.5-mm WCSP (wafer-chip-scale package), the device encompasses a 2.5×1.7-mm pc-board area. The TPA203xD1 costs 60 cents (1000), with an evaluation module available.

Texas Instruments, www.ti.com

RF IC comes in SOIC and die-form packages

Operating from 50-Hz to 7-MHz received frequencies, the MSR-FIF RF-interface front-end IC runs at 5-kbps-upconverting and 120-kbps-downconverting data rates. The device targets remote medical-sensing devices,

INTEGRATED CIRCUITS

theft-detection products, security systems, inventory-control devices, access cards, automotive-key identifications, and other RFID systems. Features include the ASK (amplitude-shift-keying) and PSK (phase-shift-keying) demodulation and ASK modulation circuitry; a charge pump; a duplexer; and a voltage output to power a microcontroller, a preamplifier, and other system functions operating at less than 50 μA. The RF ICs also include an RSSI (received-signal-strength-indicator) function as well as a power-down function. Available in a plastic SOIC-8 or in die-form packages, the MSR-FIF costs 95 and 40 cents, respectively.

Mixed Signal Integration, www.mix-sig.com

Logic-level translators have low quiescent current

Requiring a direction pin to determine the path of logic translation, the 16-channel MAX13101E, MAX13102E, MAX13103E, and MAX13108E bidirectional logic-level translators operate from a 1.65 to 5.5V power supply and a 1.2V to V_{CC} logic supply. Features include a 20-Mbps transfer rate, a 0.3-μA quiescent current, ±15-kV ESD protection on the I/O V_{CC} lines, and a shut-down-state option. Two output configurations are available during shutdown. Available in 5×5-mm TQFN-40 and 3×3-mm UCSP-36-bump packaging, the devices have a −40 to +85°C temperature range. Prices for the MAX1310x devices start at \$2.50 (1000).

Maxim Integrated Products, www.maxim-ic.com

EDN ADVERTISER INDEX

This index is provided as an additional service. The publisher does not assume any liability for errors or omissions. For immediate information on products and services, go to Reader Service under Tools & Services at www.edn.com.

EDN product mart

This advertising is for new and current products.

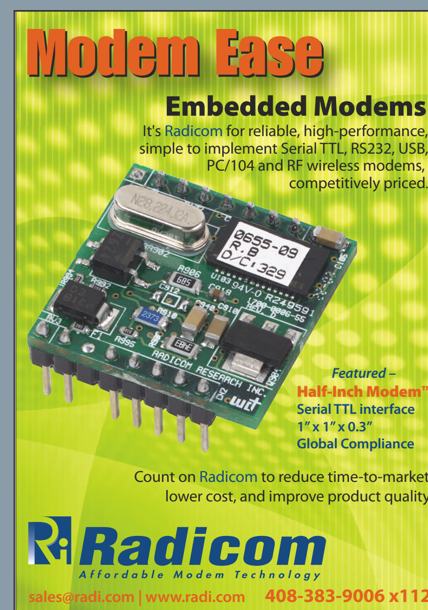

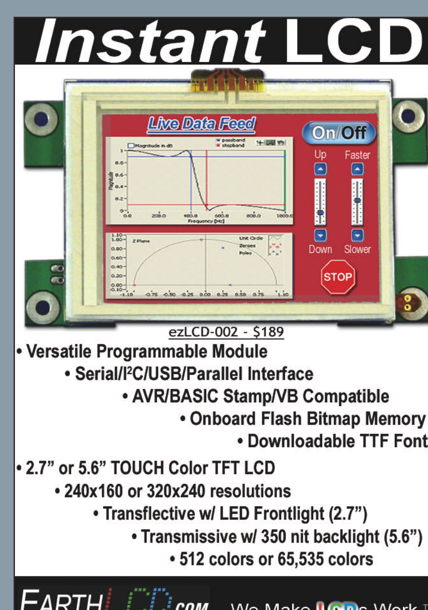

To advertise in Product Mart, call call Judy Keseberg at 800-417-5370

EDITED BY RON WILSON

scope

CHART YOUR COURSE

LOOKING AHEAD

It's going to be all about energy at the Fall Microprocessor Forum (www.in-stat.com/FallMPF/06/index.htm) in San Jose, CA, Oct 9 through 11. Keynote speakers during the three days of presentations will include Intel (www.intel.com) Architect at Large Dileep Bhandarkar on core-architecture power management, Cadence (www.cadence.com) Chief Technology Officer Ted Vucurevich on the intrusion of power into the chip-optimization process, and ARM (www.arm.com) Vice President John Cornish and Microsoft (www.microsoft.com) Development Leader Kurt Kennett, each speaking on energy optimization in hardware/software systems. We've long recognized that serious energy efficiency requires the cooperation of architects, chip designers, and software designers. Now, we are beginning to see all of these camps speaking in the same venue on the subject.

LOOKING AROUND

AND SEEING WIMAX EVERYWHERE

Sprint Nextel (www.sprint.com) announced last month that it would spend as much as $3 billion over the next two years deploying WiMax in the United States. This move could change the whole picture for wireless and mobile communications in this country. The company's enormous commitment—apparently motivated in part by a Federal Communications Commission requirement that it deploy broadband by 2007 or give up its 2.5-GHz spectrum—is likely a big win for Intel, which has been waiting a long time for good news outside the PC business. It also means good tidings for Motorola and Samsung, which Sprint has tapped to provide equipment to the new network, and to broadband users, who have been waiting a long time for this feature. It is probably bad news for Qualcomm (www.qualcomm.com), whose only answer is the EV-DO (evolutionary data-only) service superimposed on CDMA (code-division multiple access).

LOOKING BACK

50 YEARS AGO IN *EDN*

A MILESTONE ON THE ROAD TO THE HANDSET

Transistorized Two-Way Radio Fits Inside Army Helmet

FT MONMOUTH, NJ—An Army experimental helmet contains the smallest known two-way military radio. Weighing only 1 lb, this helmet-radio utilizes transistors for minimum size and weight. Developed by the US Army Signal Corps Engineering Laboratories, this new radio is battle-rugged and allows two-way conversations, bringing walkie-talkie communications to the individual riflemen. At full capacity, the helmet set can reach radios up to a mile away, and can hear powerful stations at even greater distances. The new combat FM set operates continuously for half a day on a single set of dry batteries.

—*EDN*, September 1956