Reed Business Information (RBI™)

EDN®

VOICE OF THE ENGINEER

JAN 7

Issue 1/2010

www.edn.com

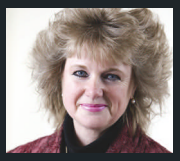

UNDER THE LID

ANALOG TEST IS SUDDENLY THE CRITICAL INGREDIENT

Page 24

INNOVATIVE PACKAGING IMPROVES LEDs' LIGHT OUTPUT, LIFETIME, AND RELIABILITY

Page 21

ANECHOIC CHAMBERS RISE FROM THE PITS

Page 31

CONGESTION MANAGEMENT CLEARS A PATH THROUGH 10 GbE

Page 37

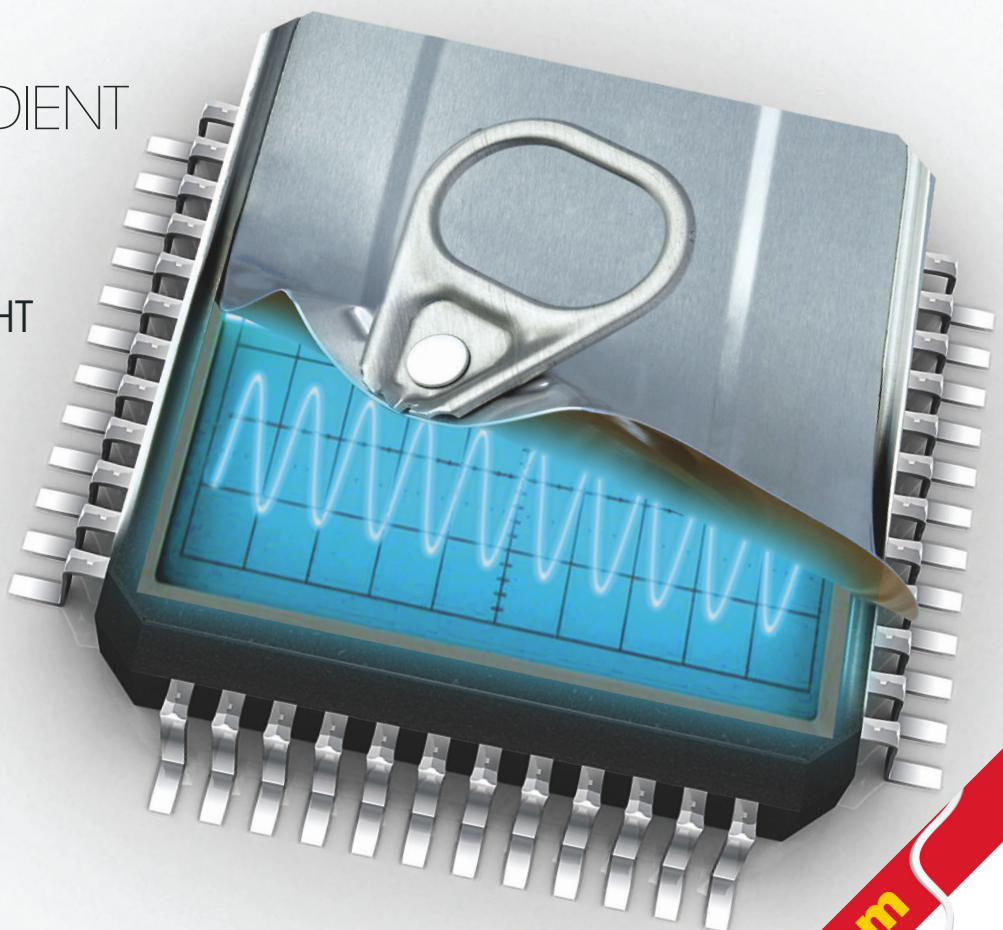

innovation

APEX PRECISION POWER

Product Innovation from Cirrus Logic

DP 12-PIN POWER SIP

(actual footprint 30.99mm X 20.17mm)

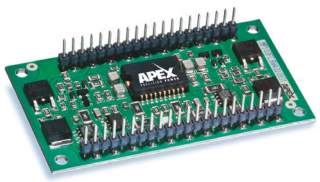

FC MODULAR 42-PIN DIP

Open Frame Product Technology
(actual footprint 65.1mm X 42.5mm)

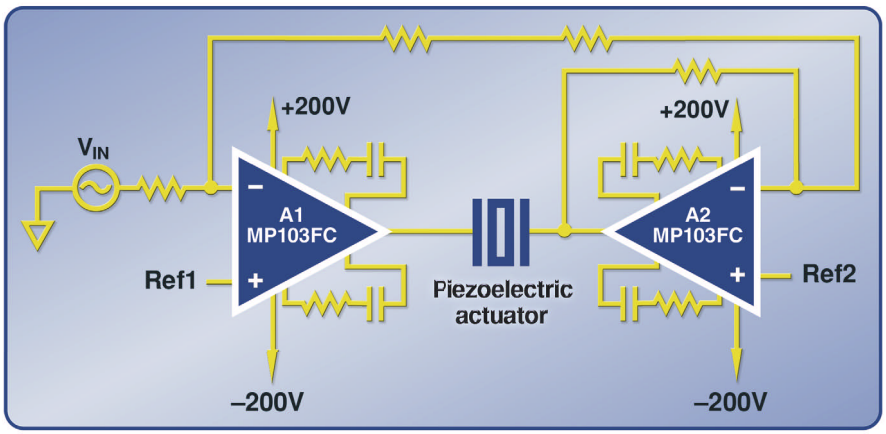

Dual Channel Amplifier is a Space Saver For Piezoelectric Actuator Drives

The MP103FC is a dual channel power amplifier that sets new standards for combining high current, high speed and high voltage for the Apex Precision Power® family of linear amplifiers from Cirrus Logic. The MP103FC has a power bandwidth rated at 230 kHz, and a 180 V/µs slew rate, for driving piezo electric applications that require lower speeds but multiple drivers. The device also features up to 15 A of output current per channel and an operating voltage ranging from 30 V up to 200 V. The MP103FC uses "open frame" modular packaging that is ideal for high speed assembly and provides a low per unit cost in comparison with many in-house discrete designs.

Model	Slew Rate	Output Current	Supply Voltage Operation
MP103FC	180 V/µs	Up To 15 A PEAK	30 V to 200 V Dual Supply
PA107DP	3000 V/µs	1.5 A continuous 5 A Peak	40 V to 200 V Dual Supply

NORTH AMERICA
+1 800-625-4084

ASIA PACIFIC
+852 2376-0801

JAPAN
+81 (3) 5226-7757

EUROPE/UK
+44 (0) 1628-891-300

THE CIRRUS LOGIC V15 APEX PRECISION POWER PRODUCT DATA BOOK IS HERE! REQUEST YOUR FREE COPY, OR DOWNLOAD AT WWW.CIRRUS.COM/EDN100107

For product selection assistance or technical support with Apex Precision Power products call 1-800-546-2739 or apex.support@cirrus.com.

LEARN MORE AT
www.cirrus.com

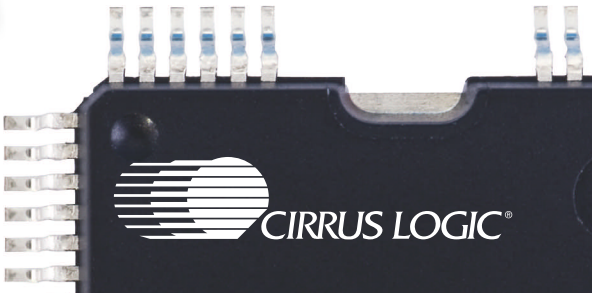

WHERE DESIGN ENGINEERS GO FIRST!

DIGIKEY.COM

Digi-Key® CORPORATION

#1 Breadth of Product Line
#1 Availability of Product
#1 Overall Performance

Source: EETimes Distribution Study, August 2009

The industry's broadest product selection available for immediate delivery

www.digikey.com
1.800.344.4539

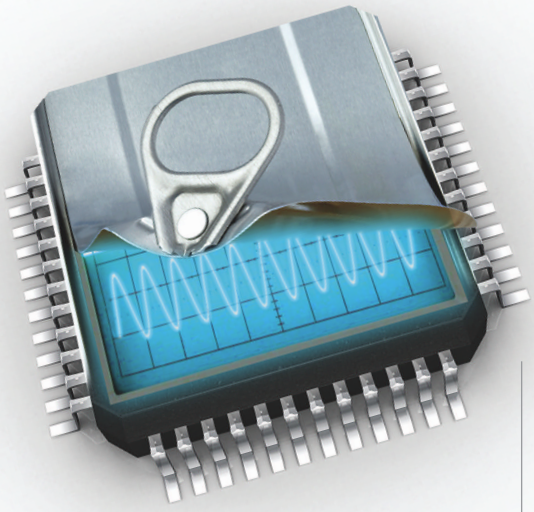

EDN 1.7.10 contents

Under the lid: Analog test is suddenly the critical ingredient

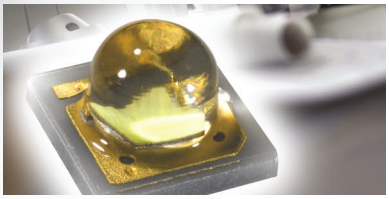

Innovative packaging improves LEDs' light output, lifetime, and reliability

Anechoic chambers rise from the pits

Congestion management clears a path through 10 GbE

pulse

DESIGNIDEAS

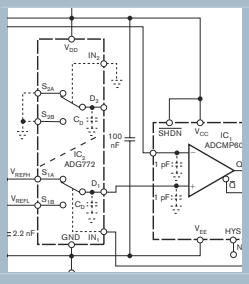

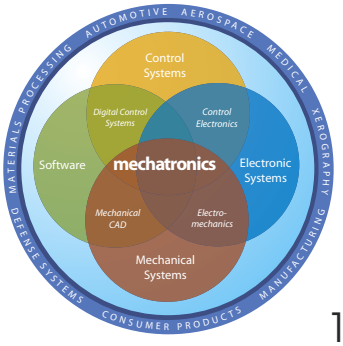

DEPARTMENTS & COLUMNS

EDN online contents www.edn.com

ONLINE ONLY

Check out these Web-exclusive articles:

Clock gating: Smart use ensures smart returns
Clock gating is a must-have technique for managing power, but errors in the design process can cause trouble later on.
➔www.edn.com/article/CA6709780

Mobile video quality: the new media-gateway revolution
Customer-defined applications are now driving telecommunications-product and -service offerings. Video telephony is a case in point, and mobile video is not far behind. Key success factors for video gateways are the quality of experience, ease of use, and affordability.
➔www.edn.com/article/CA6711225

ENGINEERS' TRUE STORIES

In *EDN*'s Tales from the Cube, engineers relate their most vexing design challenges—and how they conquered them. Here are some examples of how your peers solved real-world problems:

Tanks a lot!
➔www.edn.com/article/CA6709576
Keep on truckin'
➔www.edn.com/article/CA6702292
The boss is always right, even when he's wrong
➔www.edn.com/article/CA6662649

Read more in the Tales archive at www.edn.com/tales.

EDN® (ISSN#0012-7515), (GST#123397457) is published biweekly, 24 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. John Poulin, CEO/CFO, RBI-US; Jeff DeBalko, President, Business Media; Jane Volland, Vice President, Finance. Periodicals postage paid 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 2010 by Reed Elsevier Inc. Rates for nonqualified subscriptions, including all issues: US, $179.99 one year; Canada, $229.99 one year (includes 7% GST, GST#123397457); Mexico, $229.99 one year; air expedited, $399.99 one year. Except for special issues where price changes are indicated, single copies are available for $10 US and $15 foreign. Publications Agreement No. 40685520. Return undeliverable Canadian addresses to: RCS International, Box 697 STN A, Windsor Ontario N9A 6N4. 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 55 Number 1 (Printed in USA).

BY RICK NELSON, EDITOR-IN-CHIEF

E-readers and obsolescence

Early adopters often get burned, but usually it's only on price, as I noted in a recent blog (**Reference 1**). As I remarked in the blog post, however, commenting on an article from *The Wall Street Journal*, e-readers could be the eight-track tapes of the 21st century (**Reference 2**). Various e-readers from different generations and different vendors may be incompatible, you can't easily lend an e-book to a friend, and an e-reader is an expensive piece of equipment to inadvertently leave behind somewhere.

The *WSJ* article quotes Bob LiVolsi, the founder of BooksOnBoard, as saying that dedicated e-readers are for people with disposable incomes who love technology—not books. Book lovers might well settle for reading on a laptop or a netbook.

But most of those readers who commented on my blog strongly disagree. One mentions a preference for old-fashioned books but notes that laptops are battery hogs. One touts the adjustable font size as well as a user-friendly page-forward button that suits the device to an elderly parent. Many travelers cite the lightweight devices' ability to carry many books or stacks of data sheets.

One skeptical traveler points out that he would be dragging his laptop along anyway, suggesting that an e-reader would be superfluous. And he's not about to replace his 1000-volume collection of real books with e-versions. But another commenter notes that laptops aren't that easy to use in coach seats. The general consensus is that e-readers won't replace books but are a great addition, especially for travel.

The main criticism seems to be the need to continue buying the same content in different formats. *EDN* Senior Technical Editor Brian Dipert addresses this issue in his "Brian's Brain" blog (**Reference 3**). "Paintings and photographs may fade over time," he writes, "but if you've properly mounted and displayed them, you'll still be able to enjoy them for many decades. … The same goes for film stock and magnetic-based magnetic tape for both audio and video, although you may need to supplement low-humidity and temperature-controlled archiving with an occasional oven bake. And the same goes for books, yellowed and tattered pages aside ... at least in the paper 'analog' generation."

Brian is now contending with being an early adopter of the first-generation Sony PRS-500 eBook reader. He explains, "Sony is transitioning its online eBook store from the proprietary BBeB [broadband-electronic-book] format to industry-standard ePub [electronic publishing]. The firmware upgrade is necessary to add ePub support to the PRS-500. Curiously and for unknown reasons, a user-installable firmware update utility is not available; PRS-500 owners need to send their units to Sony. To the company's credit, it's paying for the upgrade, including free and speedy … round-trip shipping. And folks who've already shipped out their PRS-500s seem generally pleased with the process and outcome."

Brian continues, "But will my upgraded PRS-500 still be able to access the hundred-plus BBeB files stored on my SD [secure-digital] card? Reading between the lines of Sony's instructions, I think so, but I'm not yet positive. In the worst case, I'll need to redownload them from Sony's online store in the new ePub format. More generally, the trail of history is already littered with plenty of examples of obsolete services and associated hardware that left consumers who'd paid plenty of precious money hanging high and dry."EDN

REFERENCES

1 Nelson, Rick, "E-readers obsolete?" *Test & Measurement World*, Dec 2, 2009, www.tmworld.com/blog/Taking_the_Measure/26551-E_readers_obsolete_.php.

2 Fowler, Geoffrey A, "E-Readers: They're Hot Now, But the Story Isn't Over," *The Wall Street Journal*, Dec 7, 2009, www.wsj.com.

3 Dipert, Brian, "Digital-media obsolescence: Sony's eBook acquiescence," *EDN*, Dec 16, 2009, www.edn.com/blog/400000040/post/1530051353.html.

Contact me at rnelson@reedbusiness.com.

EDN

PRESIDENT, BUSINESS MEDIA, REED BUSINESS INFORMATION
Jeff DeBalko, jeff.debalko@reedbusiness.com
1-646-746-6573

PUBLISHER, *EDN* WORLDWIDE
Russell E Pratt, 1-781-734-8417;
rpratt@reedbusiness.com

ASSOCIATE PUBLISHER, *EDN* WORLDWIDE
Judy Hayes, 1-925-736-7617;
judy.hayes@reedbusiness.com

EDITOR-IN-CHIEF, *EDN* WORLDWIDE
Rick Nelson, 1-781-734-8418;
rnelson@reedbusiness.com

EXECUTIVE EDITOR
Ron Wilson, 1-510-744-1263;
ronald.wilson@reedbusiness.com

MANAGING EDITOR
Amy Norcross
1-781-734-8436;
fax: 1-720-356-9161;
amy.norcross@reedbusiness.com
Contact for contributed technical articles

SENIOR ART DIRECTOR
Mike O'Leary
1-781-734-8307;
fax: 1-303-265-3021;
moleary@reedbusiness.com

ANALOG
Paul Rako, Technical Editor
1-408-745-1994;
paul.rako@edn.com

MASS STORAGE, MULTIMEDIA, PCs, AND PERIPHERALS
Brian Dipert, Senior Technical Editor
1-916-760-0159;
fax: 1-303-265-3187;
bdipert@edn.com

MICROPROCESSORS, DSPs, AND TOOLS
Robert Cravotta, Technical Editor
1-661-296-5096;
fax: 1-303-265-3116;
rcravotta@edn.com

NEWS
Suzanne Deffree, Managing Editor
1-631-266-3433;
sdeffree@reedbusiness.com

POWER SOURCES, ONLINE INITIATIVES
Margery Conner, Technical Editor
1-805-461-8242;
fax: 1-805-461-9640;
mconner@reedbusiness.com

DESIGN IDEAS EDITOR
Martin Rowe,
Senior Technical Editor,
Test & Measurement World
edndesignideas@reedbusiness.com

SENIOR ASSOCIATE EDITOR
Frances T Granville
1-781-734-8439; fax: 1-303-265-3131;
f.granville@reedbusiness.com

EDITORIAL/WEB PRODUCTION
Diane Malone,
Manager
1-781-734-8445; fax: 1-303-265-3024
Steve Mahoney,
Production/Editorial Coordinator
1-781-734-8442; fax: 1-303-265-3198
Melissa Annand,
Web Operations Specialist
1-781-734-8443; fax: 1-303-265-3279
Adam Odoardi,
Prepress Manager
1-781-734-8325; fax: 1-303-265-3042

CONSULTING EDITOR
Jim Williams, Staff Scientist,
Linear Technology

CONTRIBUTING TECHNICAL EDITORS
Dan Strassberg,
strassbergedn@att.net
Nicholas Cravotta,
editor@nicholascravotta.com

COLUMNISTS
Howard Johnson, PhD, Signal Consulting
Bonnie Baker, Texas Instruments
Pallab Chatterjee, SiliconMap

PRODUCTION
Dorothy Buchholz,
Group Production Director
1-781-734-8329
Joshua S Levin-Epstein,
Production Manager
1-781-734-8333; fax: 1-781-734-8096

EDN EUROPE
Graham Prophet, Editor, Reed Publishing
+44 118 935 1650;
gprophet@reedbusiness.com

EDN ASIA
Luke Rattigan, Chief Executive Officer
luke.rattigan@rbi-asia.com
Kirtimaya Varma, Editor-in-Chief
kirti.varma@rbi-asia.com

EDN CHINA
William Zhang,
Publisher and Editorial Director
wmzhang@rbichina.com.cn
Jeff Lu, Executive Editor
jefflu@rbichina.com.cn

EDN JAPAN
Katsuya Watanabe, Publisher
k.watanabe@reedbusiness.jp
Ken Amemoto, Editor-in-Chief
amemoto@reedbusiness.jp

EDN. 225 Wyman St, Waltham, MA 02451. **www.edn.com**. Phone 1-781-734-8000. **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

THE EMBEDDED DEVELOPER

Microsoft

Windows Embedded

BUILT WITH WINDOWS® 7 TECHNOLOGIES

VOL. 5

A MALFUNCTION IN THE SYSTEM COULD COST THE PLANT MILLIONS...

THEY'RE COUNTING ON ME TO DELIVER.

THE DEVICE HAS TO WORK PERFECTLY. TO THE MICROSECOND.

AND HAVE THE CONNECTIVITY TO TRACK PERFORMANCE IN REAL TIME.

WINDOWS® EMBEDDED OFFERS A HIGHLY RELIABLE PLATFORM, WITH THE LEVEL OF PERFORMANCE YOU NEED TO HELP DELIVER CONNECTED DEVICES THAT STAND OUT.

WHICH WINDOWS EMBEDDED PLATFORM CAN HELP YOU DELIVER STANDOUT DEVICES?
FIND OUT AT WINDOWSEMBEDDED.COM/DEVICESTORIES

EDITED BY FRAN GRANVILLE

pulse

INNOVATIONS & INNOVATORS

32-bit processor pushes deeper into 8-bit territory

NXP's 50-MHz ARM (www.arm.com) Cortex-M0-based LPC1100 microcontroller family represents the latest 32-bit challenge to 8- and 16-bit processors. The parts are available now with prices starting at 65 to 95 cents (10,000). They have on-chip flash covering 8, 16, 24, and 32 kbytes in a 33-pin package; 48-pin LPQFP and PLCC44 packages are also available. The company claims that the EEMBC (Embedded Microprocessor Benchmark Consortium, www.eembc.org) CoreMark Benchmark measures 40 to 50% better code density for the LPC1100 than that of 8- and 16-bit microcontrollers. Efforts are in progress to further evaluate why the CoreMark code density is better on the 32-bit architecture.

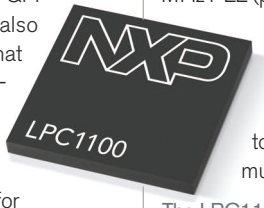

The LPC1100 microcontroller family represents the latest 32-bit challenge to 8- and 16-bit processors.

The LPC1100 supports 32 vectored interrupts with four priority levels, and it includes one UART (universal-asynchronous-receiver/transmitter) interface; as many as two SPIs (serial-peripheral interfaces); FM+ (fast-mode plus) I^2C (inter-integrated-circuit) support; and two 16- and 32-bit timers with PWM (pulse-width modulation), match, and capture. The eight-channel, 10-bit ADC is accurate to ±1 LSB (least-significant bit) DNL (differential nonlinearity). The device supports power-on-reset, multilevel brownout detection, and a 10- to 50-MHz PLL (phase-locked loop). The 12-MHz internal RC oscillator maintains 1% accuracy over the temperature and voltage range. The GPIO (general-purpose-input/output) pins are 5V-tolerant, and some pins support as much as 20-mA drive. Software-development tools are available for the LPC1100 from IAR Systems (www.iar.com), Keil (www.keil.com), Hitex (www.hitex.com), and Code Red Technologies (www.code-red-tech.com). NXP will also offer a development-tool platform for less than $30.**—by Robert Cravotta**

▷**NXP**, www.nxp.com.

FEEDBACK LOOP

"If my microprocessor can evaluate one configuration in one microsecond, it will take 13.3 hours to go through 48 billion if I did the math correctly. Are they using cloud computing?"

—Applications engineer Harry Holt, in *EDN*'s Feedback Loop, at www.edn.com/article/CA6705379. Add your comments.

ICs tackle accuracy in multiphase-metering applications

Addressing the need for accuracy in multiphase metering, Analog Devices recently announced four energy-metering chips that focus on improving the accuracy and performance of commercial, industrial, and residential smart meters. The ADE7878, ADE7868, ADE7858, and ADE7854 chips target use in polyphase configurations, including three- and four-wire wye and delta services. The devices measure both reactive and active energy with 0.1% accuracy. The ADE7878 can measure fundamental-only energies that power-quality measurements use. The ICs feature 0.1% accuracy for both active- and reactive-energy measurements over a dynamic range of 1000 to 1 and 0.2% accuracy over a dynamic range of 3000 to 1. Prices range from $5.34 (1000) for the most basic version to $7.47 for the ADE7878.

The Analog Devices announcement comes as the Smart Grid increases robustness and reliability and integrates renewable energy into the power grid. These goals bring sizable challenges, such as accurately measuring and communicating power-usage data. In the United States, 18 million homes will get smart meters within three years. For more on the challenges and applications of smart meters, go to www.edn.com/100107pa.

Analog Devices' ADE78XX family of smart-meter chips measures both reactive and active energy with 0.1% accuracy.

—by Margery Conner

▷**Analog Devices**, www.analog.com.

Chip provides speedy MPEG transcoding, trans-scaling

Zenverge claims that its new ZN200 chip transcodes MPEG-2 and various MPEG-4 video flavors in high definition at speeds as much as four times greater than real time. The ZN200 also promises to downscale HD (high-definition) video to portable-device resolutions at as much as 40 times real-time rates, according to the company. An accompanying sonic processor handles conversion between AAC (advanced audio coding), AC (audio-codec)-3, MP3, and other audio formats, and the ZN200 also internally converts between various DRM (digital-rights-management) standards. The lower-priced ZN100, whose price Zenverge has not yet announced, has half the video-processing performance of its ZN200 sibling.

According to Tony Masterson, Zenverge's co-founder, chief operating officer, and chief technology officer, the chip enables users to control the format, resolution, size, and speed of transcoding. The chip transcodes US terrestrial-broadcast content in MPEG-2 format into the H.264 format that portable video players and cell phones support. The ZN200 transscales the content from high or main profile to the baseline-profile H.264 that cell phones require and scales the resolution of HD content from either 1920×1080 or 1280×720 pixels into cell-phone resolution of 320×240 pixels. It also converts 30-frame/sec content to the 15-frame/sec content that cell phones can handle, requiring sophisticated conversion between compressed I (intra-coded), P (predicted), and B (bi-directionally predicted) pictures. Although some of these abilities may represent overkill for your application, consider that Zenverge's packet-processing architecture alternatively allows the chips to concurrently handle multiple streams with little to no efficiency loss. Such an approach might, for example, enable a set-top box to record one or a few channels while playing back another in a time-shifted manner.

Recent demos of a ZN200 on a one-lane PCIe (Peripheral Component Interconnect Express) add-in card encompassing video encoded at various bit rates, resolutions, and formats came impressively close to the company's claims, and company officials were quick to point out that the chip was running code that was not performance-optimized and was full of speed-strapping debugging routines. The chip sells for $50 (sample quantities).

—by Brian Dipert

▷**Zenverge**, www.zenverge.com.

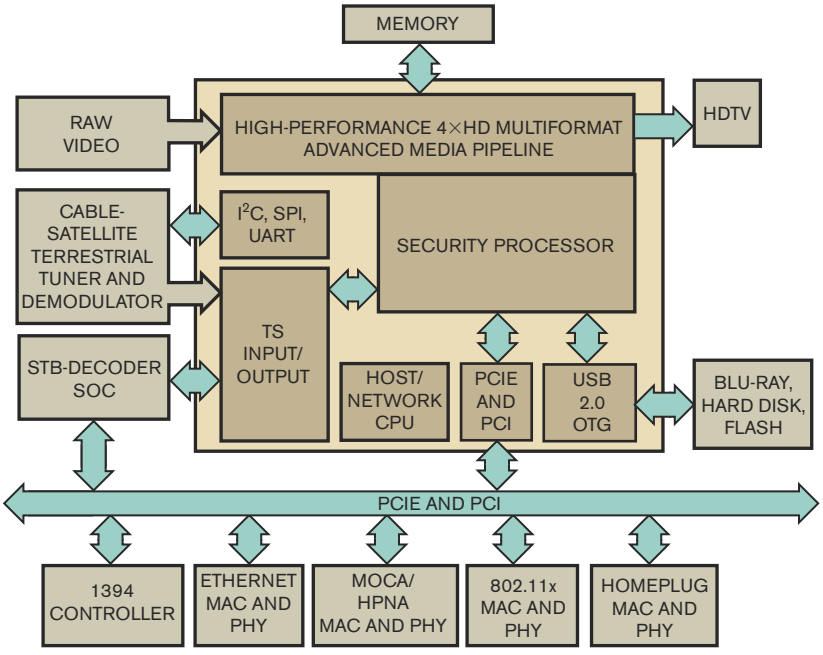

HDTV=HIGH-DEFINITION TELEVISION
HPNA=HOME PHONELINE NETWORKING ALLIANCE
I²C=INTER-INTEGRATED CIRCUIT
MAC=MEDIA-ACCESS CONTROL
MOCA=MULTIMEDIA OVER COAXIAL
OTG=ON-THE-GO
PCIE=PERIPHERAL COMPONENT INTERCONNECT EXPRESS
PHY=PHYSICAL LAYER
SOC=SYSTEM ON CHIP
SPI=SERIAL-PERIPHERAL INTERFACE
STB=SET-TOP BOX
TS=TRANSPORT STREAM
UART=UNIVERSAL ASYNCHRONOUS RECEIVER/TRANSMITTER
USB=UNIVERSAL SERIAL BUS

The ZN200 chip transcodes MPEG-2 and various MPEG-4 video flavors in high definition at speeds as much as four times greater than real time.

DILBERT By Scott Adams

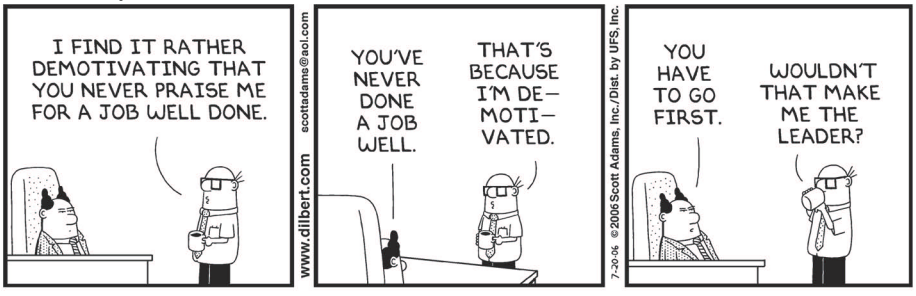

16-bit, dual-channel ADCs span 20M to 80M samples/sec

Analog Devices' new AD9269 family of 20M-, 40M-, 65M-, and 80M-sample/sec pipeline ADCs takes 16-bit samples. The units consume 42 to 93 mW of power, depending on speed, and they integrate both a reference and a sample-and-hold circuit. The differential-input stage features a 700-MHz bandwidth, and output data is on separate parallel interfaces for each channel.

Targeting use in wireless base stations, the ADCs feature a quadrature-decoding error-correction block to improve the performance of an I/Q (in-phase/quadrature) complex-signal receiver system. They operate from a 1.8V analog supply and a 1.8 to 3.3V power supply for the output section. SNR (signal-to-noise ratio) is 77 dBFS (decibels relative to full-scale) at a 10-MHz input and 73 dBFS at a 170-MHz input frequency. The SFDR (spurious-free dynamic range) is 90 dBc (decibels referenced to the carrier) at 10 MHz and 78 dBc at 170-MHz input frequencies. The ENOB (effective number of bits) is 12.4 for inputs of 9.7 to 30.5 MHz, 12.2 for inputs of 70 MHz, and 11.6 for input frequencies of 170 MHz.

The AD9269 is available in a 9×9-mm, 64-pin LF-CSP with suggested retail prices of $49, $59.50, $73.66, and $84.09 (1000) for the 20M-, 40M-, 65M-, and 80M-sample/sec versions, respectively. They operate over a −40 to +85°C temperature range. Production quantities are now available.**—by Paul Rako**

▷**Analog Devices**, www.analog.com.

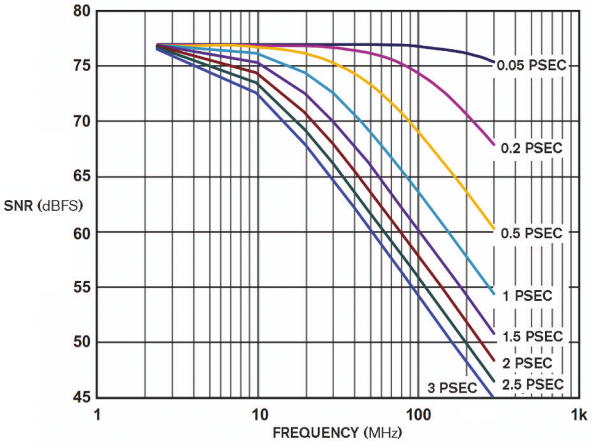

The AD9269's SNR performance depends on the quality of the signal you use to clock it. Low-jitter clock inputs are essential for best performance.

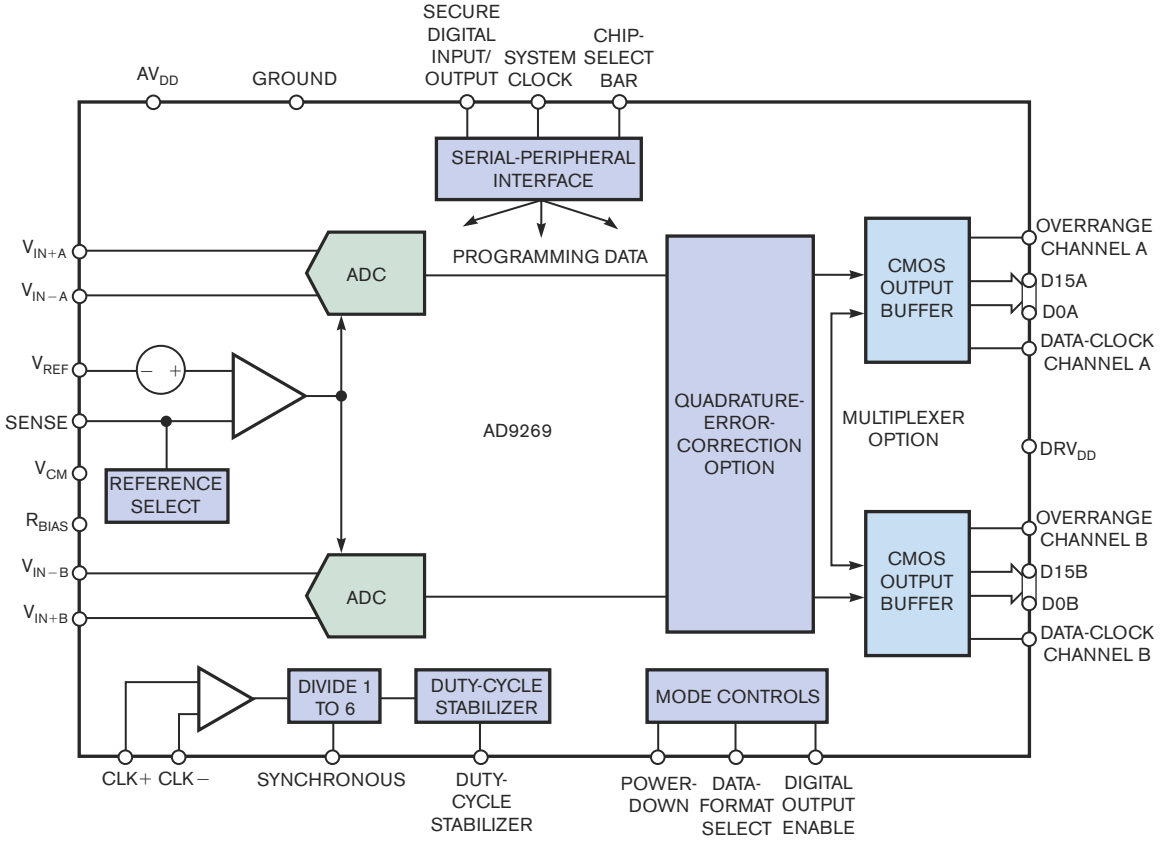

The AD9269 contains a quadrature-decoding error-correction block for wireless-base-station-receiver chains.

Touchscreen controller provides tactile feedback

Maxim Integrated Products recently announced the MAX11810 touchscreen-interface chip. The device generates the output voltage to excite a four-wire resistive touchscreen and includes ADCs to read back a touch position. It also provides haptic feedback by rotating a vibrator motor or pulsing a piezoelectric transducer on the LCD screen. It excites and reads back an IR (infrared) photodiode to sense when the screen is pressed against your face. This feature allows it to ignore the inputs and suppress the haptic outputs. A sister part, the MAX11811, has identical functions but employs a 400-kHz I²C (inter-integrated-circuit) interface instead of a 25-MHz SPI (serial-peripheral interface). Both parts operate from a 1.7 to 3.6V power supply and feature 12-bit ADCs. At a 34.4k-sample/sec rate, the device consumes 246 mW at 1.8V or 698 mW at 3.6V.

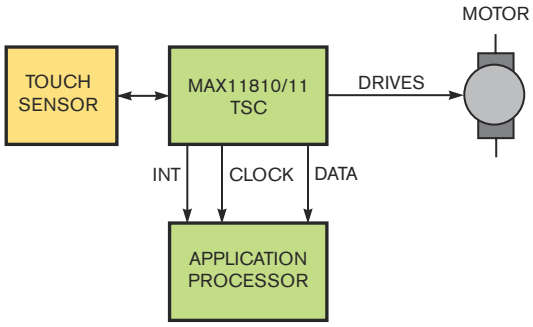

The MAX11810 contains the circuitry to drive a touchscreen, a haptic tactile-feedback device, and a photodiode. The part has an SPI, and the MAX11811 uses an I²C interface.

Applications include cell phones, MP3 players, portable media players, digital photo frames, multifunction printers, point-of-sale terminals, bar-code scanners, card readers, and other industrial equipment. Automotive-qualified-part versions exist for use in car GPSs (global-positioning systems), entertainment head units, and rear-seat entertainment systems.

The product has a state machine for onboard processing of touch events to validate them before sending the data to the system microprocessor. The onboard processing can also supply an improved capture rate of as many as 161 coordinates/sec.

The MAX11810 operates over a −40 to +85°C temperature range and comes in 2.1×2.1-mm, 20-pin TQFN packages and 16-pin WLPs. It sells for $1.81 (1000) and is available now, along with evaluation kits that include a touchscreen and a dc motor.

—by Paul Rako

▷**Maxim Integrated Products**, www.maxim-ic.com.

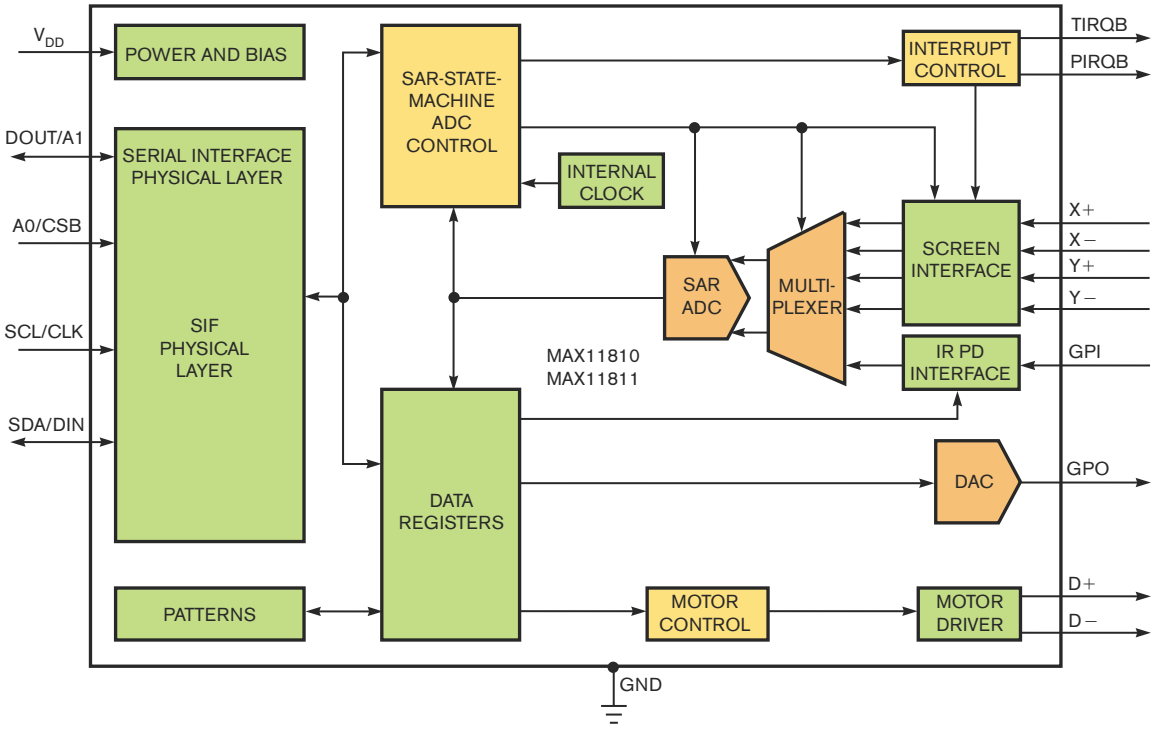

The MAX11810 touchscreen-interface chip for handheld consumer electronics can process touch events and output a tactile feedback without waking up the system application processor.

Analog FastSpice RF delivers noise analysis for RF circuits

Berkeley Design Automation Inc has announced AFS RF (Analog FastSpice radio frequency), which Chief Operating Officer Paul Estrada calls the industry's first true Spice-accurate noise-analysis tool for RF circuits. AFS RF accurately analyzes nanometer-scale device-noise impact for all types of prelayout and postlayout circuits, ensuring early insight into its impact on performance, power, and area.

Before the emergence of AFS RF, designers had to use limited-spectrum RF tools that can only approximate device-noise impact on RF circuits, Estrada explains. Such approximations are increasingly inaccurate with decreasing process geometries, often becoming grossly inaccurate in nanometer-scale circuits. Circuits with sharp transitions, such as switched-capacitor filters, charge pumps, and dividers; high-frequency circuits, such as RF front-end blocks; and oscillators are especially sensitive to these inaccuracies. Without accurate analysis, designers must include expensive design margin or risk missing specifications in silicon.

> "Without accurate analysis, designers must include expensive design margin.

Using the industry's first full-spectrum device-noise-analysis engine, AFS RF provides true Spice accuracy for every run. For complex circuits, it is five to 10 times faster than traditional RF tools that can only approximate device-noise effects. AFS RF features the DNA (device-noise-analysis) Advisor to characterize DNA requirements, high-capacity periodic-steady-state analysis for greater than 100,000-element postlayout circuits, full-spectrum periodic-noise analysis with true Spice accuracy, full-spectrum total oscillator-device-noise analysis with phase and amplitude noise, and harmonic balance for fast single-tone analysis of moderately nonlinear circuits.

"We have been using the AFS platform for the last two years for full-circuit transceiver verification and more recently transient-noise analysis of our analog/RF blocks," says Alan Wong, head of IC design at Toumaz Technology (www.toumaz.com). "AFS RF delivers full-spectrum periodic noise analysis and does not trade off accuracy for performance. AFS RF allows us to analyze our prelayout and postlayout RF blocks, delivering true Spice accuracy and a five- to 10-times speedup over traditional RF-analysis tools."

"Full-spectrum periodic noise analysis is critical for accurate characterization of device noise in nanometer analog/RF designs," says Boris Murmann, PhD, assistant professor in the department of electrical engineering at Stanford University (www.stanford.edu). "Without accurate device-noise analysis, designers need to add significant margin to ensure performance. This [addition] can be very expensive. For example, adding just 0.5 bit of margin—that is, 3-dB SNR (signal-to-noise ratio)—in a noise-limited circuit will double the required power."**—by Rick Nelson**

▷**Berkeley Design Automation Inc**, www.berkeley-da.com.

FUNCTION-LEVEL VISIBILITY IMPROVES APPLICATION OPTIMIZER

As processor architectures enable developers to build ever-more-complex applications, the software-development tools that speed and improve the process increase in value. For example, the Ceva-Toolbox software-development environment now includes an application-optimizer tool chain that provides software optimization for code targeting Ceva's DSP cores. The tool enables developers to generate target code completely in C that performs within 20% of the performance of hand-optimized assembly. The tool-chain optimizers assist development throughout design and building and include a project optimizer, an application profiler, scoring-based compilation, a postlinker optimizer, and automated test generation. The tool chain is now available.

The project-build optimizer enables you to automatically build multiple configurations based on compiler settings at individual function levels. The tool chain supports the automated simulation, profiling, and down-selection of these multiple configuration scenarios to help find the best match of the application requirements and the system resources. The application profiler includes a cycle-accurate C-level application and memory-subsystem profiler that enables the tool chain to detect application- and system-level bottlenecks by taking into account the full memory subsystem. The scoring-based compilation component supports multiple iterative compilation levels to allow trade-offs between cycle count and code size on a C-level, function-by-function basis. The postlinker optimizer works with a global view of the application within the system and resolves resource scenarios that you cannot completely address during compilation and that further enable smaller application code. The automated test generation provides full test suites and supports standard Perl scripts. The tool chain also supports the inclusion and integration of C-callable assembly-optimized functions as a native part of the compiler's standard libraries. The debugger supports connectivity features with external tools to simplify migration of algorithms, such as those from The MathWorks' (www.mathworks.com) Matlab, so that parts of the code can run within Matlab or on the debugger.–by Robert Cravotta

▶**Ceva, www.ceva-dsp.com.**

01.07.10

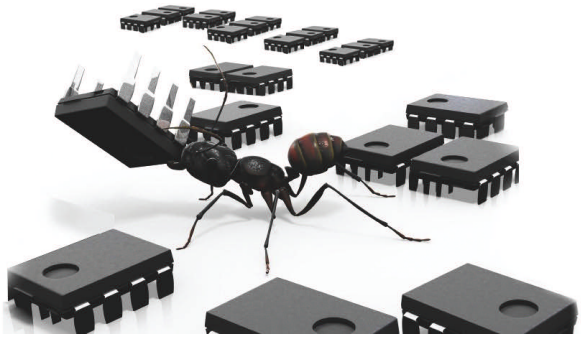

Chains of superparamagnetic colloidal particles rotate to produce flows on length scales much larger than the chain, allowing them to behave like micro-ants that can move large particles (courtesy Charles Sing and Alfredo Alexander-Katz, MIT).

RESEARCH UPDATE

EDITED BY FRAN GRANVILLE

Micro-ants act as tiny conveyor belts inside chips

Researchers at the Massachusetts Institute of Technology's department of materials science, in collaboration with researchers in Germany and from Boston University (www.bu.edu), have devised a microscopic system that could provide a novel method for moving tiny objects inside a microchip and could provide new insights into how cells and other objects move around within the body. Organs such as the trachea and the intestines have cilia, tiny, hairlike filaments that are constantly in motion, beating in unison to create currents that sweep along cells, nutrients, or other tiny particles. The new research uses a self-assembling system to mimic that kind of motion, providing a simple way to move particles around in a precisely controlled way.

Alfredo Alexander-Katz, a professor of materials science and engineering; his doctoral student Charles Sing; and other researchers devised a system that uses tiny so-called superparamagnetic beads made of polymers with specks of magnetic material in them. Due to the heavy magnetic material content, these beads sink to the bottom of the sample. By applying a rotating magnetic field, which caused the beads to spontaneously form short, spinning chains, the researchers created currents that could then carry along surrounding particles as much as 100 times larger than the beads themselves. Alexander-Katz refers to the microscopic assembly of beads—each just a few microns in size—as micro-ants because of their ability to move while "carrying" objects so much larger than themselves (**Reference 1**).

The new method could provide a simpler, less-expensive alternative to currently available microfluidic devices, a field that is still in its early stages of development. Such devices now require precisely made channels, valves, and pumps using microchip-manufacturing methods to control the movement of fluids through them. The new system could offer such precise control over the movement of liquids and the particles suspended in them that it may be possible to dispense with the channels and other plumbing altogether, controlling the movements entirely through variations in the applied magnetic field. In other words, software rather than hardware could control the chip's properties, allowing you to instantly reconfigure it through changes in the controlling software.**—FG**

▷**Massachusetts Institute of Technology**, www.mit.edu.

REFERENCE

1 Sing, Charles, Lothar Schmid, Matthias F Schneider, and Alfredo Alexander-Katz, "Controlled surface-induced flows from the motion of self-assembled colloidal walkers," *Proceedings of the National Academy of Sciences*, Dec 14, 2009, www.pnas.org/content/early/2009/12/17/0906489107.abstract.

Biochip-based test aims at speedy disease diagnosis

IBM scientists have built a biochip-based, one-step point-of-care diagnostic test that requires less sample volume and is significantly faster than currently available techniques. An article in the December 2009 issue of *Lab on a Chip* describes the research. The portable, easy-to-use system can test for many diseases, including cardiovascular disease. IBM Research-Zurich scientists Luc Gervais and Emmanuel Delamarche, in collaboration with the University Hospital of Basel in Switzerland, developed the test, which uses capillary forces to analyze samples of serum, or blood, for the presence of disease markers. Capillary-action force is the tendency of a liquid to rise in narrow tubes or to be drawn into small openings.

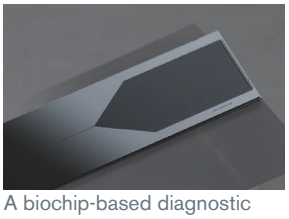

A biochip-based diagnostic test requires less sample volume than other techniques.

IBM encoded the forces of capillary action on a microfluidic chip made of a silicon compound, measuring 135 cm and containing sets of micronwide channels in which the test sample flows through in approximately 15 sec—several times faster than traditional tests. Further, you can adjust the filling speed to several minutes when the chip requires additional time to read a more complex disease marker.

The microfluidic chip works with a 1-microliter sample—50 times smaller than a teardrop—that the researchers move onto the chip using a pipette. The capillary forces then begin to push the sample through an intricate series of mesh structures. The sample passes into a region into which researchers have deposited microscopic amounts of the detection antibody. The test uses only 70 picoliters—1 million times smaller than a teardrop—of these antibodies, making for their fast and efficient dissolution in the passing sample. For more, go to www.edn.com/article/CA6708422.

—by Suzanne Deffree

▷**IBM Research-Zurich**, www.zurich.ibm.com.

BY HOWARD JOHNSON, PhD

It's a Gaussian world

A previous article suggests that most digital output waveforms follow a nearly Gaussian profile (**Reference 1**). Let's test that theory.

Figure 1 depicts the rising edge from a Texas Instruments (www.ti.com) DL100 44T LVDS (low-voltage-differential-signaling) driver. This driver exists on a custom test board with SMA output connectors. A coaxial cable feeds the test outputs directly to the 50Ω terminated inputs of a LeCroy (www.lecroy.com) SDA6000 oscilloscope. A small bit of residual ringing from the previous bit appears in the left side of the **figure**. Then, just before the signal begins its major ascent, you may see a tiny precursor, probably the result of crosstalk from the test equipment and predrivers inside the DL100 44T. The amplitude of these features amounts to only about 1% of the main signal's step size.

Directly after the main edge, the signal overshoots a tiny amount and then rings back approximately 2%. This artifact is more likely due to the layout of the test board than anything related to the DL100 44T. With artifact amplitudes this low, spectral analysis of the waveform should clearly reveal whether the true spectral content of a DL100 44T driver follows the Gaussian theory.

Figure 1 plots the step response of the driver. Spectral calculations require the impulse response. The frequency response associated with that impulse response gives you a filter you can then apply to any square-edged digital signal, making the output look just like the step shapes in **Figure 1**.

To make the necessary spectral calculations, first convert the step waveform to an impulse response by computing, at each point, the slope of the signal:

$$x'_n = (x_{n+1} - x_n)/T,$$

where T equals the sampling interval. Next, truncate the signal to a finite-length region. Make the region wide enough to capture the main features of the signal in question but narrow enough to eliminate unrelated events before and after the main edge. Try using the whole region shown in **Figure 1**, from 0 to 5 nsec. Within this region, 2000 points are sampled at 2.5 psec each.

Simple truncation of a sampled-data sequence can induce Gibbs phenomena. To mitigate these unusual effects, multiply the truncated sequence by a Hamming window of length N=2000:

$$h_n = 0.54 - 0.46\cos(2\pi n/N).$$

A Hamming window smoothly feathers the sequence to zero at both ends, reducing the impact of discontinuities at the endpoints. It accomplishes this task at the expense of a modest loss in frequency-domain resolution.

Finally, use an FFT to compute the frequency response of the windowed impulse response and normalize the result so that it starts at 0 dB. **Figure 2** compares the result against a Gaussian template. It's a match. EDN

REFERENCE

1 Johnson, Howard, "Real signals," *EDN*, Oct 8, 2009, pg 13, www.edn.com/article/CA6699737.

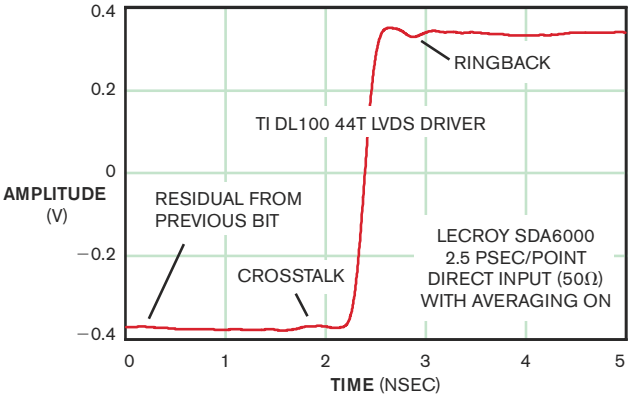

Figure 1 Small artifacts precede and follow each edge.

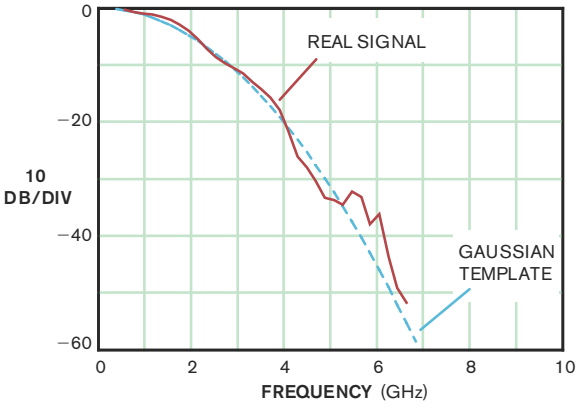

Figure 2 The real signal follows a Gaussian template.

MEET THE GUY that ELIMINATED HIS TEAM'S MANUFACTURING VARIABILITY ISSUES.

FIX YOUR MANUFACTURING VARIABILITY PROBLEMS AND YOUR REPUTATION WILL PRECEDE YOU. If you're designing chips for high functionality, high speed and lower power consumption at the most advanced process nodes, you've got variability issues. We have the solution that will increase your yield, performance and prestige by a wide margin. | Get more information at mentor.com/solutions/manufacturing-variability.

Mentor Graphics®

THE EDA TECHNOLOGY LEADER

BY PALLAB CHATTERJEE, CONTRIBUTING TECHNICAL EDITOR

MEMS, sensors, and nanotechnology

The relentless pursuit of more transistors in a smaller area has driven the semiconductor industry since its start. This drive has now produced sophisticated SOCs (systems on chips) that include billions of transistors per design for memories and microprocessors. The costs of these large designs on small geometry processes make for a $50 million to multibillion-dollar project. Because of these costs, the advanced design work is being performed by significantly fewer development teams and companies and is outside the scope of most start-ups.

In contrast, MEMS (microelectromechanical systems), solid-state sensors, and new applications of nanotechnology are rapidly gaining in importance and diversity. These areas of creative design activity require a stronger interaction between design and processing technology. The result of this interaction is the modification and advancement of the process technology as well as the creation of new designs and design methods.

Recent advances in the data-processing capability in electronics products have given rise to new and smaller MEMS products and to new functions. Some of these new functions and applications that people encounter regularly are air-bag accelerometers, tire-pressure monitors, 2- and 3-D gyroscopes for GPS (global-positioning-system) devices, gyroscopes for video-game motion controllers, and rotation controllers for cell-phone displays. You can see the complexity of some of these new designs from a typical MEMS gyroscope (**Figure 1**).

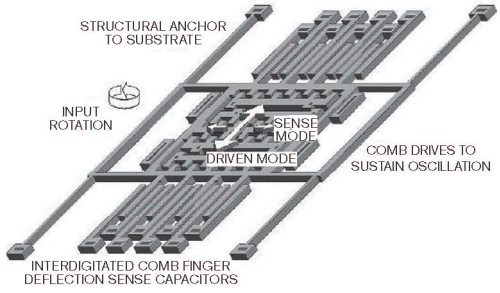

Figure 1 A typical MEMS gyroscope is complex (courtesy Steve Beeby, MEMS Mechanical Sensors).

Sensor technology has long focused on image sensing and pressure sensing as baseline markets. New technologies have additionally brought about huge advances in the design of microfluidic and nanofluidic sensors, chemical and gas sensors, vibration and energy-harvesting sensors, and temperature sensors. Recent advances, such as the November 2009 announcement of technology licensing for audio sensors for the design and manufacture of MEMS microphones between ST-Microelectronics (www.st.com) and Omron (www.omron.com), are bringing about new cell-phone, laptop, mobile-Internet, and health-care-product designs. These new solid-state devices depend on the low cost of semiconductor manufacturing to meet the volume and price points necessary for mass commercialization and the reduction in power and increase in effective signal integrity to be compatible with digital signal processing with either wired or wireless connectivity.

Nanotechnology is a growing area of multiple disciplines and definitions. The general consensus is that the word "nanotechnology" refers to devices and materials that are complete objects smaller than 100 nm. This definition includes semiconductor processes and lithography at the 65- to 22-nm nodes, graphine- and carbon-nanotube structures, and most structures and materials that designers create using catalyzed self-assembly. Nanotechnology design includes active devices, passive devices, and mechanical structures. These design objects find application in solar-energy collection, nanomedicine, photonics, and electronics. Some of the more exotic electronics devices possible are the recently discovered memristor, finfets, and other 3-D transistors.

The new technologies address societal and industrial issues, such as monitoring of vibration and wear on the structural aspects of bridges and buildings, creation of portable and intelligent-clinic-use medical-diagnostics equipment, automated access control and asset management, a more human interface to computing platforms, and increased operation safety in transportation systems. The reduced size, power requirements, and cost of manufacturing along with available distributed-computing resources will make sensors and sensor signal processing ubiquitous in most products in the future. The creation of circuits and systems with these design elements is an iterative task with both design exploration and technology development taking turns as the driver in the loop. EDN

Contact me at pallabc@siliconmap.net.

Renewable Energy Applications: Reliability and Safety with Fiber Optic and Optocoupler Isolation Technology

Introduction

Climate change, rising carbon dioxide levels, environmental pollution and volatile fossil fuel prices have driven electric power generation producers, governments, and individual consumers to adopt alternative energy generation methods such as solar and wind. Governmental renewable energy support has been incorporate into strategic energy plans worldwide. During 2009 at least 73 countries had energy policy targets and more than 64 countries had policies to promote renewable power generation. Renewable energy is no longer an area dominated by research, but rather a large commercial market with high adoption and installation growth rates.

Renewable energy growth in 2008 was substantial and 2009 activity points toward continued progress. The United States ended 2008 with 25 GW of wind power, eclipsing Germany's 24 GW capacity. China's total wind power doubled for the fifth year in a row, ending the year above 12 GW, exceeding China's 2010 development target of 10 GW two years ahead of plan. More than 80 countries around the world now have commercial wind power installations.

With most 2008 offshore wind generation capacity in Europe, capacity reached nearly 1.5 GW. The United Kingdom became the offshore wind power leader in 2008 but other countries, such as the United States, have plans for new offshore facilities.

Grid-connected solar photovoltaic plants are the fastest growing power generation technology, with a 70-percent increase in existing capacity to 13 GW in 2008. This was a six fold increase in global capacity since 2004. Including off-grid applications, total 2008 photovoltaic generation increased to over 16 GW. In 2008, utility-scale solar PV power plants—plants over 200 kW—were estimated to have grown by 80%, to 1,800, since 2007. The added plants totaled over 3 GW, a tripling of existing 2007 capacity.

Utility photovoltaic plant additions were a major 2008 focus in Spain but facilities also went into the Czech Republic, France, Germany, Italy, Korea, and Portugal. New solar photovoltaic plants are planned and under development throughout Europe as well as China, India, Japan, and the United States.

Given the high growth rates and importance of renewable energy, Avago focused its fiber optic and isolation product portfolios and product development efforts toward renewable energy applications. Digital optocouplers, gate drivers, isolation amplifiers and solid state relays provide safety isolation and insulation along with their basic function. Fiber optic components, besides providing isolation, provide EMI resistant, long distance communication channels in wind turbine farms and panel-to-panel communications in solar photovoltaic systems. These control and communication channels are critical for safety and optimal power efficiency in distributed solar systems and especially for offshore wind farm installations.

Avago Renewable Energy Solutions

Avago Technologies offers fiber optic transmitters, receivers and transceivers, as well as IGBT and Power MOSFET gate drivers and optocoupler isolation products for wind turbine, wind farm and solar electric power generation.

Avago fiber optic components, isolated gate drivers, isolation amplifiers and digital optocouplers are deployed in wind turbines, wind farms and solar photovoltaic energy farms today.

Applications include:

- Power rectifiers and inverter power assemblies
- Control and communication boards
- Turbine control unit condition monitoring systems
- Wind farm networks
- Sun tracking control
- Communications boards
- Solar farm substation automation and protection relays
- Single- and three-phase inverters
- Battery chargers
- Motor control
- DC-to-DC converters,
- System control and monitoring

Solar Energy: Fiber Optic Solutions

Avago fiber optic transceivers provide control and communication links within a solar energy system and within the solar farm. Resistant to EMI and able to withstand rugged environments, Avago transceivers support POF, HCS and Multimode fiber. Plastic cable supports transmission up to 50 m. For longer distances 200 µm HCS and 62.5/125 µm fiber will support up to 2,700 m links. Maximum transmission distances, depending on data rates, are up to 80 km.

Data rates from 1/5/10/20/160 MBd are available with many different connector styles: LC, SC, SFP and MT-RJ.

Figure 1. Fiber data links in solar energy farm

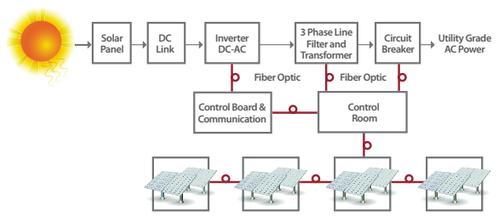

Wind Turbine Networking, Control and Communications: Fiber Optic Solutions

In wind turbine applications access for maintenance and repair is more limited than in land-based solar farms. This is especially true for off shore wind generation installations. In addition, large mechanical loads and changing local weather conditions must be monitored in near real time for optimal power generation, safety and control. Long term reliability and quality are very important component selection factors.

Inside the turbine nacelle, short range fiber optic connections connect power generation subsystems and the turbine pitch controller to the main system controller. Avago offers a broad portfolio of fiber-optics transmitter and receiver like the HFBR-0500Z family which are used in Wind-Turbines since many years. For Wind-Turbines >2 MW the main choice inside the turbine is fiber optics because of EMI, lightning and galvanic isolation. Longer range fiber links connect to tower base and the wind farm network.

In turbine monitoring and solar systems, Fast Ethernet transceivers with DMI and SC-RJ connectors, such as the Avago AFBR-5978Z, are popular as they operate at 650 µm over POF and HCS cable and feature a –25°C to 85°C temperature range. A demonstration kit for the AFBR-5978Z (AFBR-0978Z) is available from any Avago sales offices or representative.

Our new multimode, small form factor (SFP) HFBR-57E5APZ Fast Ethernet transceiver with DMI requires only 611 mW and operates from 3.3 V.

Figure 2. Fiber optic communication links in wind turbines

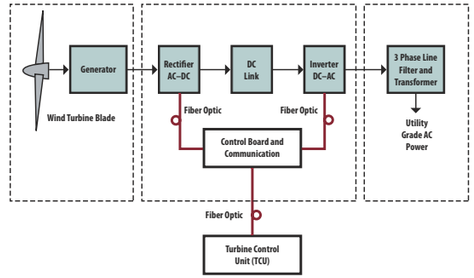

Solar and Wind Power: Application Specific Optocouplers, Gate Drivers and Isolation Amplifiers

As the leading optocoupler supplier, Avago develops application specific isolated power MOSFET gate and IGBT drivers for power applications. Featuring high peak output current, high CMR, wide temperature range, low supply current, surface mount packages and integrated protection circuits they are ideal for motor control, inverter and battery charging circuits.

The ACNW3190 isolated gate driver has a 5 A peak drive current and under voltage lockout (UVLO) protection. Other devices have 0.4 A to 2.5 A peak drive current. For maximum reliability and design flexibility, Avago devices have protection features such as integrated VCE detection, UVLO, "soft" IGBT turn-off, isolated open-collector fault feedback and active Miller clamping for maximum design flexibility and circuit protection. A Miller clamp controls the Miller capacitance current during high dV/dt transitions and can eliminate the need for a negative supply voltage.

Figure 3. Isolation and insulation solutions in wind-solar power generation systems include inverters, charge controllers and communications.

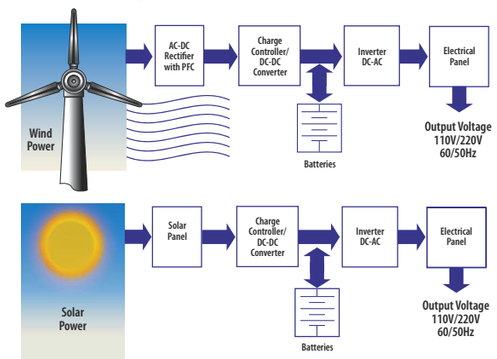

With the Avago analog isolation amplifiers, current and voltage monitoring in motor control, inverter and battery applications is simple, safe and accurate. Phase currents and DC link voltages are easily measured while maintaining isolation. The amplifiers have working voltages of up to 1140 V for safety. Gain accuracy ranges from ±1% to ±5%.

The ACPL-796J 1-bit, second-order sigma-delta modulator converts an analog input signal into a high-speed data stream with galvanic isolation based on Avago's optical coupling technology. The ACPL-796J operates from a 5 V power supply with dynamic range of 80 dB with an appropriate digital filter. The differential inputs of ±200 mV, full scale ±320 mV, are ideal for direct connection to shunt resistors or other low-level signal sources in motor phase current measurement applications.

Solar and Wind Power: Digital Optocouplers

Wind control systems must accommodate varying wind conditions and avoid run-to-failure by monitoring subsystems accurately. High voltage isolation, current isolation and immunity to EMI are critical for reliable and robust control systems. Solar and wind farms face harsh environmental conditions and operate over large physical areas susceptible to lightning and extreme temperatures.Avago digital optocouplers have the insulation and isolation specifications to make command and control systems safe and reliable. All our digital optocouplers are qualified to all international safety organization insulation and isolation specifications. Many devices have working voltages that exceed 1768 Vrms and isolation voltage up to 7500 Vrms are offered.

Operating at 10 MBd, multi-channel and bi-directional optocouplers (ACSL-6210/6310/64x0 in Dual/Triple/Quad channel) maximize design options and conserve PCB area. For extreme environments, the ACPL-M61U features –40°C to 125°C operation. The ACPL-M61L/061L/064L/W61L/K64L are ultra low power optocouplers suitable for energy efficient systems. Higher speed devices such as the ACPL-772L/072L operate at 25 MBd and from –40°C to 105°C.

Our photo MOSFETS, or solid state relays (SSRs), operate with voltages from 60 V to 600 V and with currents up to 2 A for energy storage management systems.

Figure 4. Isolated gate drive, digital optocoupler, level shifter and isolation amplifier applications.

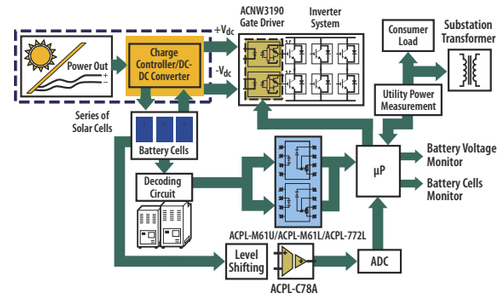

Summary

Innovative fiber optic and optical isolation solutions have been available from Avago for over three decades. As a major contributor to the work of international standards organizations, device reliability, quality and safety are always part of any Avago design. We offer the industry's best isolation technology, specifically designed and manufactured to meet the stringent requirements of applications in power generation systems.

Around the world our customers value what we do. We provide the components and technology to make their visions for the renewable energy market expand and prosper.

CONTROL ENGINEERING | design news | EDN | Test & Measurement World

MECHATRONICS IN DESIGN

FRESH IDEAS ON INTEGRATING MECHANICAL SYSTEMS, ELECTRONICS, CONTROL SYSTEMS, AND SOFTWARE IN DESIGN

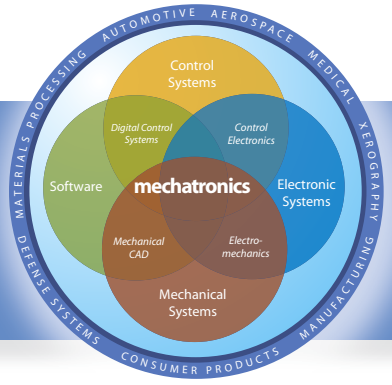

What's your point of view?

Time, frequency, and modal domains offer complementary views, insight.

Each of us is a critical-thinking problem solver. We have to be, as society's problems are mounting, getting harder, broader, deeper, and multidisciplinary. As basic engineering skills—analysis, hardware, and software—have become commodities worldwide, America's competitive advantage comes from being immediate, innovative, integrative, and conceptual. Our innovation must be local; you can't import it, you create it! It is a way of thinking, communicating, and doing. It differentiates us from other engineers around the world.

As multidisciplinary teams are formed to solve these problems, usually with a core group comprising mechanical, electronic, computer, and controls engineers together with problem-specific experts in, for example, combustion, chemistry, structures, materials, anatomy, and physiology, insight and communication are of utmost importance. We have all witnessed how engineers from different backgrounds describe the same concepts using different language and different points of view, which often can lead to confusion and, ultimately, design errors. Being able to describe concepts with clarity and insight in a variety of ways is essential for the mechatronics engineer as the multidisciplinary team leader.

The two domains, time and frequency, represent different perspectives. They are interchangeable, complementary points of view—that is, no information is lost in changing from one domain to another—and together lead to better understanding and insight.

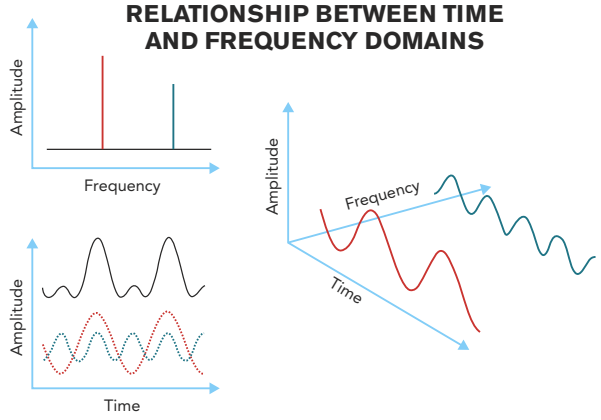

Time and frequency domains together give insight and enhance communication.

Most signals and processes involve both fast and slow components happening at the same time. In the time, or temporal, domain, we measure how long something takes, whereas in the frequency, or spectral, domain, we measure how fast or slow it is. No one domain is always the best answer, so the ability to easily change domains is quite valuable and aids in communicating with other team members.

A third domain, the modal domain, is particularly valuable in analyzing the behavior of mechanical structures. It breaks down complicated structural vibration problems into simple vibration modes. Unique insight into the use of the modal domain in mechatronic system design has been provided in the work of Dr Adrian Rankers, manager of mechatronics technologies at Philips Applied Technologies.

The time domain is a record of the response of a dynamic system as indicated by some measured parameter, as a function of time. More than 100 years ago, Jean Baptiste Fourier showed that any real-world signal can be broken down into a sum of sine waves, and this combination of sine waves is unique. By picking the amplitudes, frequencies, and phases of these sine waves, one can generate a waveform identical to the desired signal. To show how the time and frequency domains are the same, the figure shows three axes: time, amplitude, and frequency. The time and amplitude axes are familiar from the time domain. The third axis, frequency, allows us to visually separate the sine waves that add to give us the complex waveform. Note that phase information is not represented here.

If we can predict the response of a system to a sine wave input—that is, the frequency response—then we can predict the response of the system to any real-world signal once we know the frequency spectrum of that signal. The system's frequency-response curves are really a complete description of the system's dynamic behavior.

Engineers who can bridge gaps among disciplines and articulate complementary points of view clearly and insightfully will certainly have a competitive advantage. **EDN**

Kevin C Craig, PhD, is the Robert C Greenheck chair in engineering design and a professor of mechanical engineering, College of Engineering, Marquette University. For more mechatronic news, visit mechatronicszone.com.

More measurements. More data. Less cost.

With up to 560 channels of data logging at a maximum scan rate of 1000 channels/sec, an internal 6 ½ digit DMM with 22-bit resolution, and 0.03% accuracy, it's no mystery why Agilent's data acquisition systems are among the top selling in the world. The 34970A and 34980A are both at home with either data acquisition or functional test, with a broad selection of plug-in modules. Now with Agilent Data Logger Pro software, you'll be able to collect data, check limits, perform extensive branching, and view and store the data without authoring the programming yourself. It's an ideal expansion at a significant savings.

Agilent 34970A switch/measure unit

Agilent 34980A switch/measure unit

www.agilent.com/find/daqswitch

*Pricing in USD and subject to change

u.s. 1-800-829-4444 canada: 1-877-894-4414

Agilent Technologies

Figure 1 Manufacturers typically use SMT packages for 1 to 3W and higher-power HB LEDs because the close contact with the underlying PCB dissipates heat faster (left). Manufacturers typically use through-hole packages for 0.5W LEDs (above). An SMT HB LED includes the LED chip, the substrate, an encapsulant, and a primary lens.

BY MARGERY CONNER • TECHNICAL EDITOR

INNOVATIVE PACKAGING

improves LEDs' light output, lifetime, and reliability

HB-LED PACKAGING DESIGN CONTINUES TO ADVANCE, CONTRIBUTING ALMOST AS MUCH TO PERFORMANCE AS THE LED CHIP ITSELF.

HB-LED (high-brightness-light-emitting-diode) efficacy may grab the headlines, but packaging is where the action is in this market. Getting the light out of the die, through the encapsulant and lens, and onto the application's surface at a consistent, reliable color for a long lifetime is a challenge, and LED manufacturers accomplish this task through the use of innovative packaging for the devices. HB-LED users may be unaware that Big 4 manufacturers Cree, Nichia, Osram, and Philips Lumileds control almost all HB-LED chips. Although China is home to hundreds of LED-chip manufacturers, those vendors produce devices with only 10 to 70 mA of current, which eventually find use as indicators rather than illumination devices. The remaining HB-LED vendors are for the most part packagers of the tiny HB-LED chips that they purchase from the Big 4 manufacturers.

R&D efforts to coax more lumens from an LED die are enabling the use of LEDs in new markets and applications. Just as important, though, are manufacturers' packaging methods. These methods enhance efficiency to meet the needs of unique applications, such as solid-state-lighting, automotive, signage, and medical applications. To select the best HB LED for your application, you

need to understand its packaging as well as some of the methods its manufacturer used to squeeze out and precisely place as many photons as possible. Packaging affects light extraction, heat extraction, and lumen maintenance—important characteristics of LEDs. Light extraction determines how much emitted light will reach the intended application; thermal extraction affects how much heat the LED chip will experience and, hence, its lifetime and performance; and lumen maintenance indicates the lifetime of the LED. The HB-LED package determines—or at least strongly influences—all three of these characteristics, which in turn affect the LED's performance, lifetime, and reliability.

HB LEDs currently have no standardized package footprint. Unlike for other electronic components, such as ICs, MOSFETs, and most passive components, every HB-LED manufacturer uses a unique package. This practice can be annoying to LED designers who want to line up second sources for their products, but LED packages comprise intricate combinations of the LED chip, encapsulant, primary optics, and substrate. The package for a 20- to 70-mA indicator-type LED dissipates heat through its leads. The epoxy encapsulant serves as a lens and provides a rigid protective structure for the die and leads. Compared with low-power LEDs, HB LEDs must rely on greater heat-dissipation capability for their packaging, and they typically come in SMT (surface-mount technology) or through-hole packages (**Figure 1**). SMT tends to find use in HB LEDs of 1W or higher power because it mounts directly onto a heat sink, whereas through-hole packages generally suit midrange HB LEDs of 0.5 to 1W.

An LED chip generates light in an epitaxial structure. Ideally, all of the hole/electron combinations would result in photons. Material characteristics and imperfections, however, cause some of the combinations to produce heat rather than light. The next challenge is getting the photons out of the die because of the difference in index of refraction between the layers of the chip, the encapsulant, and the lens. Air has an index of refraction of approximately one, whereas some die materials have an index of approximately 1.4. A difference in indexes for two materials causes light to reflect back at the boundary when the light exceeds a certain angle. Total internal reflection is the most significant barrier in extracting the light from an LED chip.

AT A GLANCE

- Relatively few manufacturers provide most HB-LED (high-brightness-light-emitting-diode) chips to the industry.
- Although Big 4 vendors manufacture and package their own LED chips, they also sell to other LED vendors that specialize in packaging HB LEDs for specialized applications.
- Light extraction, thermal extraction, and lumen maintenance affect packaging goals for HB LEDs.

To manufacture white HB LEDs, manufacturers cover blue LEDs with a phosphor that emits light in the white region when it irradiates with blue light. Developing the phosphor that most efficiently converts blue light to white is part of the IP (intellectual property) of some LED vendors that purchase LED chips from other vendors and then package them with their own phosphor mixture in a package. Because of the LED chips' high index of refraction, the package needs a transition layer between the air and the chip/phosphor combination. This layer also serves as a protectant and a lens. Low-power indicator LEDs use an epoxy encapsulant, but epoxy tends to yellow when you expose it to high heat and UV (ultraviolet) radiation, according to Kee Yean Ng, product-marketing manager for solid-state lighting and displays in Avago's optoelectronics-products division. "Silicone is now almost universal as both the lens and the encapsulant for HB-LED packages," he says. The silicone helps both to extract the light and to protect the chip from air and moisture.

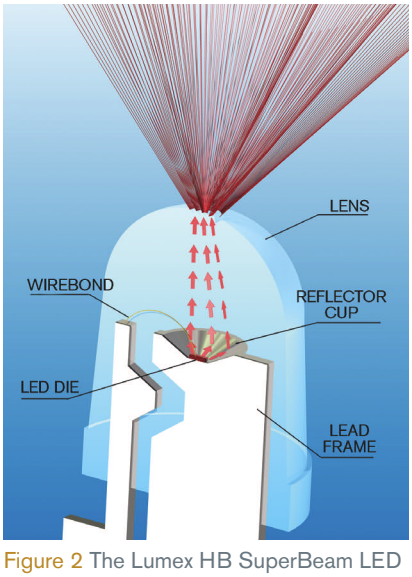

Figure 2 The Lumex HB SuperBeam LED uses an internal primary optical reflector to direct edge-emitted light from the die out the top of the package.

Some HB LEDs are surface-emitting, meaning that most of the light comes from the top surface of the die. However, many HB LEDs, especially those that HB-LED packagers purchase, are edge-emitting, and the package must redirect the light so that it exits the top of the package. Several proprietary methods exist for redirecting edge-emitted light. One method is surface roughening, which is a secondary manufacturing process that roughens the surface to make it less reflective to the internal light. Lumex uses primary optics—those that are part of the HB-LED package—that redirect the light from the top in a pinpoint pattern (**Figure 2**).

Most of the electron-hole combinations that don't result in photons create heat—the biggest drawback to LED life. LEDs generate less light as they get hotter. In addition, all manufacturers specify and test their LEDs at room temperature with a test pulse rather than a continuous current, so the LED never heats up. In practice, however, an LED's junction temperature is never at room temperature when the LED is on. Thermal performance also affects LED lifetime: The hotter the junction temperature, the shorter the LED lifetime will be. Cooler LEDs lose less light due to thermal inefficiency, have a more efficient light output, and last longer.

You can compare the thermal performance of various LEDs by looking at the thermal-resistance rating. The lower it is, the less temperature difference there will be between the junction and the solder point. The thermal path for any LED system starts at the solder point and includes the case temperature of the LED out to the ambient temperature, the PCB (printed-circuit board) that it's mounted on, and its thermal-interface material. These thermal-design steps are the same ones engineers use to determine the operating points for power MOSFETs. Even though you can't measure the junction temperature of the LED chip, you can measure the solder-point temperature to

determine the power going through the LED and use it to calculate the junction temperature. The junction temperature is useful for understanding lumen maintenance, or how long you can expect the LED to remain operational at its operating temperature.

Lumen maintenance is the percentage of initial light output an LED can maintain over its operational life. The predominant failure mode of an LED is to get dimmer over time rather than to burn out as an incandescent bulb does. Lumen maintenance is an agreed-upon number that defines end of life for an LED and generally is L70, which means that the LED is emitting 70% of the light it did at its maximum output when it was new. Energy Star requirements for solid-state lighting set the lifetime at L70, as well. "We've done a lot of higher-temperature testing and found out that the primary mode of degradation for HB LEDs is the package itself," says Paul Scheidt, product-marketing manager at Cree. "If you keep the package cool, the chip doesn't degrade all that much over time—maybe 2 to 3%—not that much of a contributor compared with a total 30% degradation over years." The predominant failure mode is due to degradation of the materials, including silicone and plastic, in the package. Both deteriorate over time in different amounts at different temperatures and light intensities.

The question of aging, reliability, and lifetime for LEDs is not a simple one: High-power LEDs have not been in existence long enough to have life-testing numbers. A common number for LED life is 50,000 hours, or almost six years, which is longer than the most recently introduced high-power LEDs and their packages have existed. But how useful is it to know that an LED has a 50,000-hour lifetime if a manufacturer measured that lifetime in an environment that differs from the one in which your application must operate? For example, a flashlight may require a lifetime of only 1000 to 2000 hours. Users who don't need the full 50,000 hours want to know how to calculate lifetime numbers under different operating conditions. "We have what we call the four critical parameters: junction temperature, drive current, solder-point temperature, and ambient temperature," says Scheidt. "Knowing these parameters, users can get a good estimate of how long the lifetime is going to be, not just of whether they're going to achieve 50,000 hours." EDN


You can reach Technical Editor Margery Conner at 1-805-461-8242 and mconner@connerbase.com.

For links to the Web sites of companies this article mentions, go to www.edn.com/100107df.

For more technical articles, go to www.edn.com/features.

BY RON WILSON • EXECUTIVE EDITOR

UNDER THE LID

ANALOG TEST IS SUDDENLY THE CRITICAL INGREDIENT

AFTER YEARS OF ATTENTION TO DIGITAL-TEST TECHNOLOGY, ANALOG TEST IS EMERGING AS A ROADBLOCK TO COST REDUCTION.

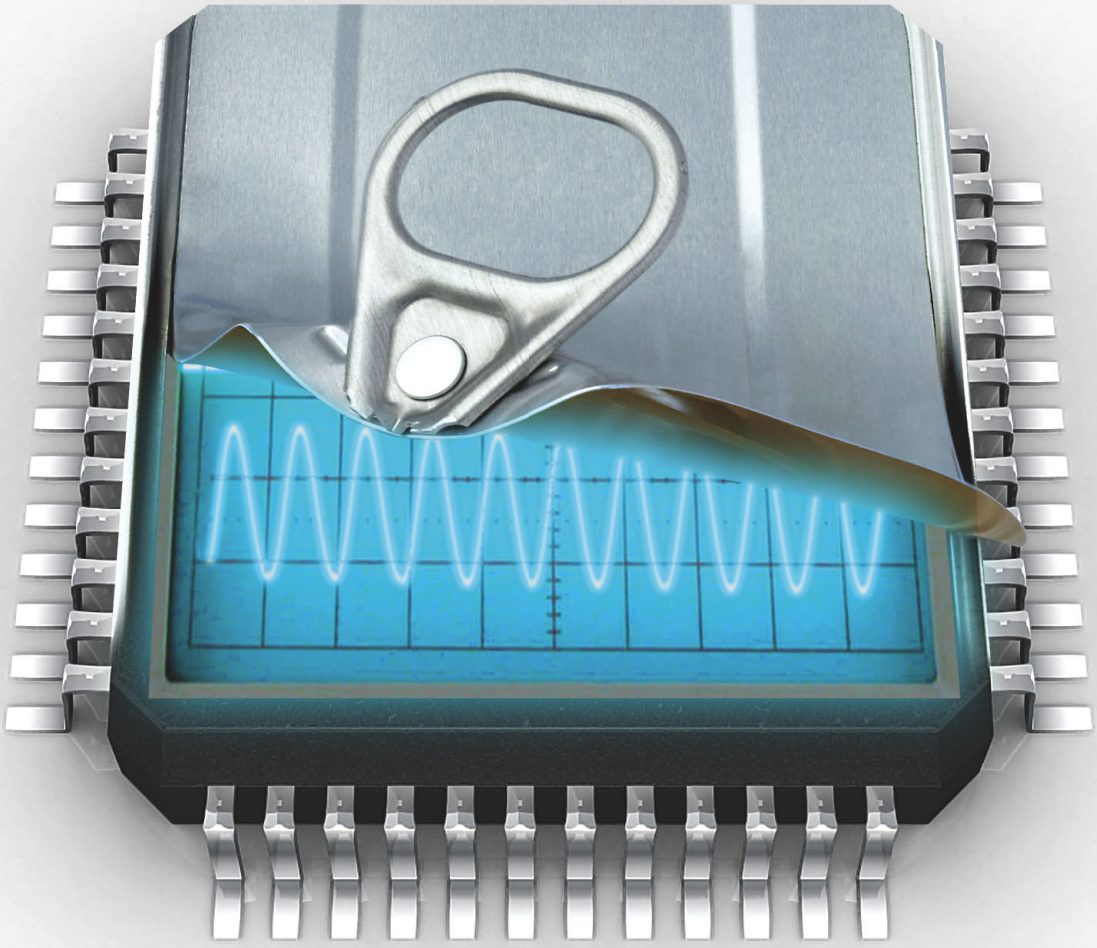

ATPG (automatic-test-pattern generation), BIST (built-in self-test), and structural-test techniques have kept digital-test costs nearly constant during the explosion in digital complexity. Without these tools, however, as analog complexity starts to grow rapidly, analog-test cost is growing, too. "AMS [analog and mixed-signal] circuits account for 70% of SOC [system-on-chip]-test cost and 45% of test-development time, even though they make up a small fraction of the chip complexity," said Karim Arabi, senior director of engineering at Qualcomm, speaking at a panel on analog-IC test at the ITC (International Test Conference) in Austin, TX, last November. "There is no ATPG for AMS circuits. There is no practical fault model. And what DFT [design-for-test] and BIST efforts we use are purely custom." Arabi's complaint neatly summarizes the situation.

Three broad strategies for analog test now exist. The first and most traditional approach is characterization—sweeping the inputs through their full allowable range and measuring all the pins to verify each parameter on the data sheet. This approach still dominates—for good reason—in the world of discrete analog components. The second approach, increasingly common in the SOC world, is to use hardware in the chip itself to test the AMS circuitry. The ideal of this school is to achieve something like BIST for AMS: functional blocks that test themselves on command. The third approach is the Holy Grail of AMS test: structural test.

"We need to redefine the problem," said Craig Force, a test-engineering-platform manager at Texas Instruments, speaking to his fellow ITC panelists. "We don't want to do parametric test of everything on the data sheet. We want to do the minimum number of tests to prove that the circuit is built right." Despite years of effort, the goal remains elusive, but hope seems to be rekindling among some experts.

PROVING THE DATA SHEET

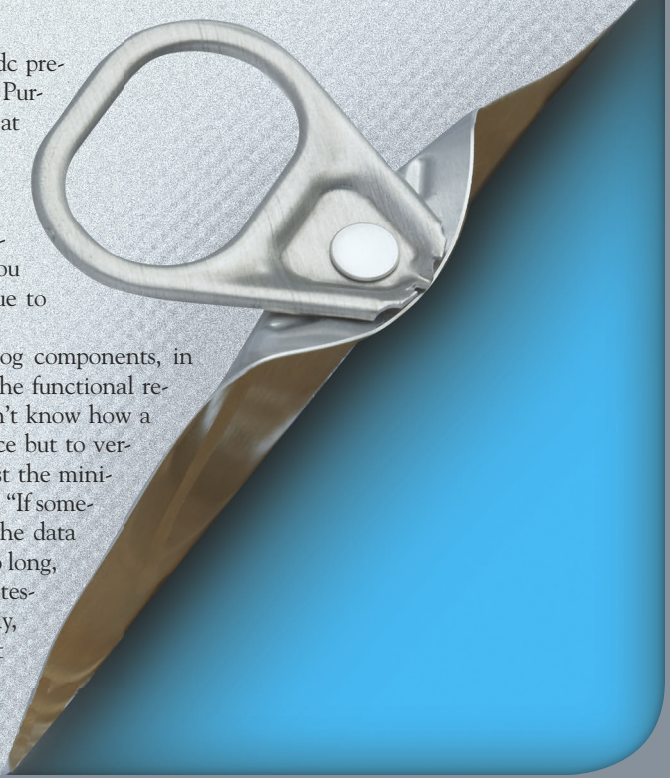

"In our world, we are always fighting for dc precision and for signal fidelity," says Michael Purtell, product- and test-engineering manager at analog-components vendor Intersil. "You are always asking yourself if you are measuring the part or the equipment. For instance, you look at a spectrum going through an analog-to-digital conversion. What part of what you see is really the device, and what part is due to your measuring technique?"

Welcome to the world of precision analog components, in which the claims on the data sheet define the functional requirements. Because the manufacturer doesn't know how a customer will use the chip, there is no choice but to verify every number on the data sheet. "We test the minimum/maximum numbers 100%," Purtell says. "If something is untestable, we will write it up on the data sheet as 'typical.'" Such thoroughness leads to long, complex test routines on expensive analog testers with elaborate load boards. Fortunately, however, there is no internal state—at least in purely linear circuits—and it is often suffi-

cient to examine just the pins of a device. "Usually, once you set up the configuration registers for a test mode, you can figure out what you need to know from just the inputs and outputs," Purtell says.

Such exhaustive testing is not always necessary. "Different markets have different test needs," says Mark Hemming, director of product development at audio-chip vendor Nuvoton. "An automotive customer can be demanding and insist that your test procedure cover everything on the data sheet. For the same part, a commercial customer might just want you to check that the chip works."

THE SOC SIMPLIFICATION

Although discrete analog components often require full testing of data-sheet parameters, AMS IP (intellectual-property) blocks in SOCs benefit from a defined environment and usually a fixed set of functions. These constraints reduce the test challenge from full characterization to simply verifying some functions. The almost-ideal case is an IP block that implements a standard interface for which the standard defines the test requirements.

Another considerable simplification for SOC designers is the availability of lots of digital logic. Virtually free gates make it easy to build mode-control registers and to digitally configure the AMS block for test modes: Shorting inputs and enabling loop-back paths are both examples of test modes. The digital wealth also makes possible the latest trend in SOC test: on-chip instrumentation. "We are deep in the throes of on-chip instrumentation," Force says. "Everyone puts a widget in their SOC."

The trend has gathered the most momentum in digital blocks in which, Force says, off-the-shelf IP often implements test functions. In more purely linear blocks, test instruments tend to be ad hoc designs. Instruments for each design incur additional design time and die area. These costs don't mean you can't build a signal generator, a voltmeter, and an oscilloscope into your SOC. They do mean that you must do as much of the work as possible in the digital domain and reuse test circuitry for other purposes.

"In an SOC there's a high comfort level with handling digital signals," Force says. "Digital is a great medium for passing around data, so anything we can digitize is good. But you can't always just hang a 1-bit ADC on a sensitive signal path. Sometimes you are still going to have to touch down and test."

AT A GLANCE

- Test cost for analog circuits is now significant.
- Traditional techniques may be inappropriate for analog IP (intellectual property) in SOCs (systems on chips).
- On-chip instruments and test buses are now in wide use.
- The goal is to find the analog equivalent of ATPG (automatic-test-pattern generation) and scan testing.

One approach to on-chip instrumentation then is to pick up analog signals at key nodes—inputs, outputs, and internal nodes that might not be readily inferable from the outputs—and route them to an ADC. This converter could be a centralized resource if the signals are robust enough to route across the block, or it could be a simple sigma-delta converter, comprising little more than a comparator and a control register, right on the node. Once the converter digitizes the signals, the SOC usually has enough processing power to turn the raw samples into sophisticated measurements.

As for reuse, one example is an SOC's built-in circuitry that allows customers to perform board-level tests. For instance, if a high-speed-serial-interface specification requires an automatic-calibration mode or a training mode, the hardware you need to implement that function may also be useful in test mode. In this respect, the evolution of signal-path design, as analog designers face the lower voltages and greater process variations at advanced process nodes, is providing resources for test engineers.

"There are two trends here that are really helpful," says Sanjiv Taneja, vice president for Encounter Test at Cadence. "Increasingly, the analog functions that you have to test at speed are both self-calibrating and self-adapting. Often, you can reuse those capabilities to implement on-chip test functions."

AN INSTRUMENTED SERDES

To see this principle in action, Navraj Nandra, director of product marketing for AMS IP at Synopsys, provides a guided tour through a high-speed SERDES (serializer/deserializer) IP block. The block, an AMS IP function available from Synopsys and other IP vendors, serves as a bidirectional interface between a parallel datapath in an SOC and a multigigahertz serial bus that goes to the outside world. A CDR (clock- and data-recovery) circuit lies at the heart of the SERDES. The CDR circuit reconstructs a clock from the incoming signal and then uses that clock to extract raw data bits from the signal. To perform these functions, circuits in the CDR unit must find the phase of the clock that created the data stream, essentially by matching it up with a locally generated clock. In the Synopsys design, that task employs a digitally controlled phase rotator. The phase rotator also performs a second job: By sweeping the phase in test mode and by capturing the output of the CDR circuit with a comparator, designers can build a table of phase versus voltage for an input signal. Sending that data from the chip to a workstation for graphing produces an eye diagram. This debug approach collects the data inside the CDR circuit so that it represents the signal the SERDES must use to extract the data.

Because most serial-interface standards define minimum signal quality in terms of a polygon that fits into the opening in the eye diagram, the internal tool collects the data you need to generate a compliance eye mask. "This technique forms the basis for the production test routines, as well," Nandra says. If you equip an ADC on the output of the CDR circuit with limit registers and if you provide a flexible sample clock to the ADC, you can collect data and send it to external display routines to form, in effect, a sampling oscilloscope (**Figure 1**). With the scope, you can see not only eye diagrams but also any repetitive waveform you choose to direct into the ADC's inputs. Nandra points out that your customers can use these capabilities to measure their designs' link performance in-system. By simply adding one or two pieces of hardware, the Synopsys designers are getting multiple capabilities for several audiences.

TOWARD ANALOG BIST

The SERDES design employs implementation foresight to permit creative

reuse of functional components in test mode, but its techniques are implementation-specific. Another SERDES designer might take a different approach. Designers of a sigma-delta converter, for example, can't directly apply any of the circuits the SERDES uses. Although on-chip widgets have rapidly and pervasively proliferated, this growth has not resulted in a library of reusable instruments or in general principles that you can use to automate the design of self-test into analog circuits.

Engineers at the former LogicVision, now a part of Mentor Graphics, have put a lot of thought into this problem. Some of the thinking Mentor received with the acquisition showed up in the ITC panel discussion. "We are looking for the analog equivalent of scan," said Stephen Sunter, an engineer with Mentor. "But first we need an accepted fault model. After years of trying, people worry that accuracy is impossible. So, fine. Let's get a good-enough model. We need to separate the issues of highly accurate characterization from our problem, which is simply to find random defects in our circuits. In the past, we have come up against complex waveforms with huge dynamic ranges, and we have faced complex analyses. But I think we can move forward if we agree on a few techniques. First, use loop-back modes.

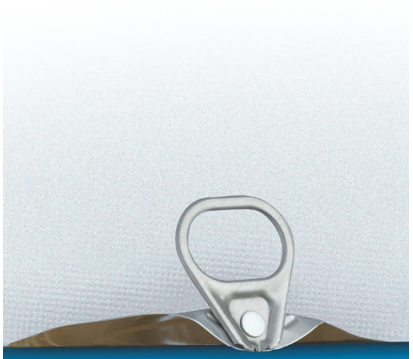

WE ARE STILL A LONG WAY FROM THE EMERGENCE OF A TOOL THAT CAN AUTOMATICALLY GENERATE INSTRUMENTS AND TESTS IN A DESIGN.

Second, employ a low-frequency analog bus, controlled by 1149.1, to move analog signals from their sources to an ADC. Third, use a high-frequency serial bus to extract data from a sigma-delta converter. The principle here is to convert voltage into time to manage the huge analog dynamic range and then to do 90% of the work in digital."

Sunter's thinking matches what is happening in the SOC area, describes the approach Synopsys designers used on their SERDES, and describes Nuvoton's strategy. "To speed parametric testing, you can use an internal analog bus to interrogate analog nodes with high accuracy, but you always wonder if you are accurate enough," says Nuvoton's Hemming. "In our market, you worry about cost. So either you try to find circuits that are useful in regular operation, or you design them so that your customers can use them too." Hemming illustrates this concept with the rather considerable telephone-line measurement capability in the company's Pro-X (programmable extended) codecs (**Figure 2**). It appears that the use of on-chip instrumentation is beginning to coalesce into the sort of general principles that Sunter advocates. We are still a long way, however, from the emergence of a tool that can automatically generate instruments and tests in a design.

ANALOG STRUCTURAL TEST

With all the progress in on-chip instrumentation, a fundamental difference still exists between analog widgets on chips and digital scan or BIST. Engineers manually insert analog instruments into a block as an ad hoc means of measuring its behavior. In contrast, automated tools create synthesized digital scan chains or BIST hardware after analyzing the structure of the logic.

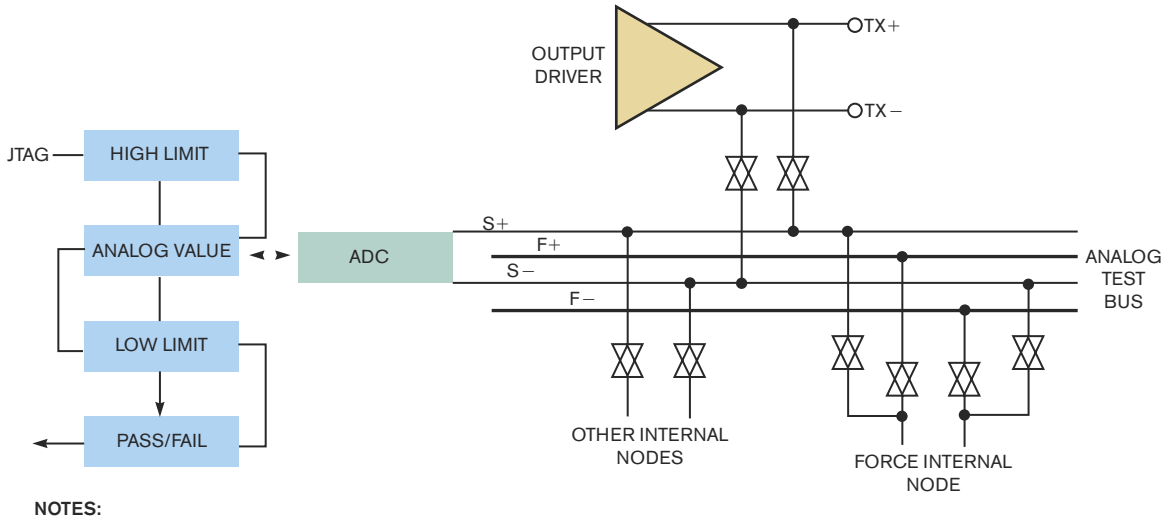

NOTES:
YOU ACCESS ALL REGISTERS THROUGH THE JTAG INTERFACE.
THE ANALOG TEST BUS IS UNDER REGISTER CONTROL.
THE JTAG INPUT VECTOR CONTAINS HIGH AND LOW LIMITS.
THE JTAG OUTPUT VECTOR CONTAINS A SIMPLE PASS/FAIL BIT FOR COMPARISON TO THE "EXPECTED VECTOR."

Figure 1 Limit registers and an analog bus in an ADC make a flexible on-chip test instrument.

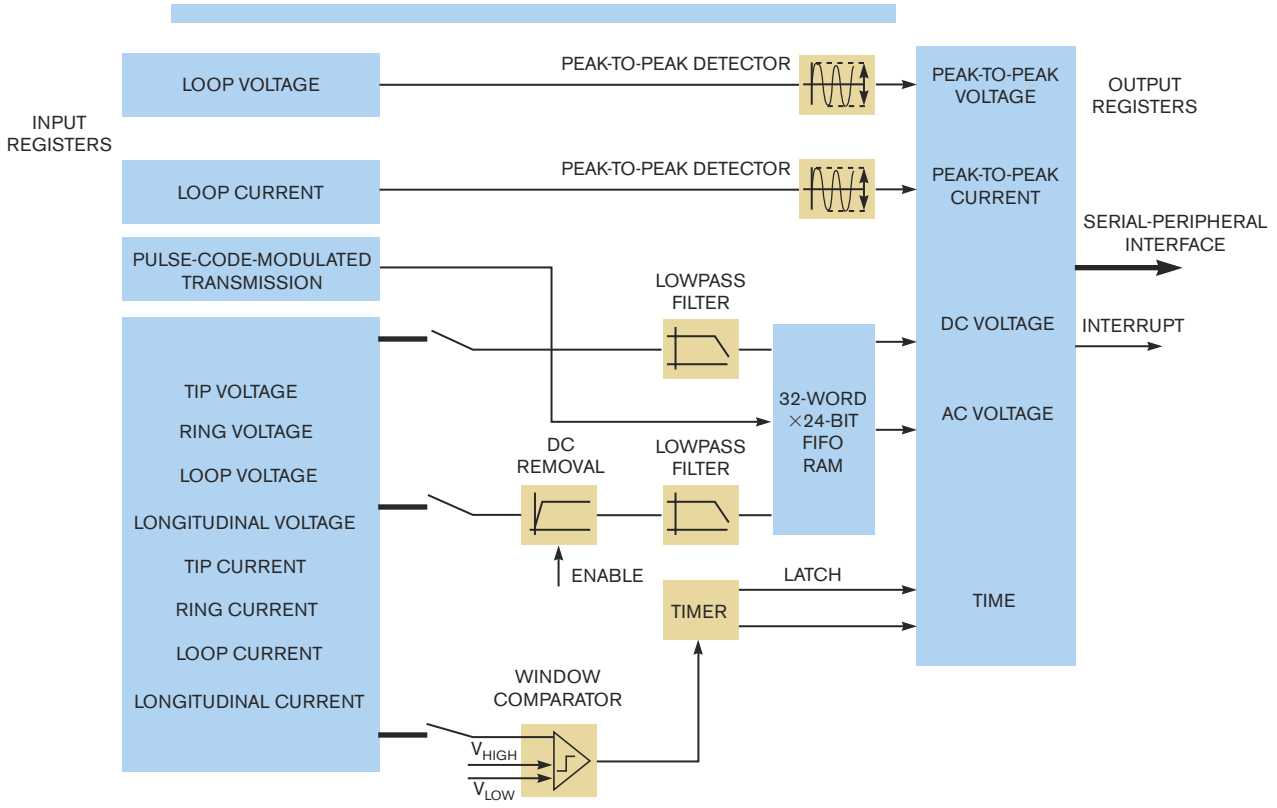

Figure 2 On-chip diagnostics both assist in testing a Nuvoton codec and allow the customer to examine the line connection.

These tools verify the integrity of the structure; they don't measure the behavior of the logic.

The path from behavioral to structural test runs directly through fault models. If you know what defects can occur on a wafer and you can model the effect of each defect on the circuit, you can deduce how to look in one spot to find a large number of possible problems. In the digital world, ATPG software, scan insertion, and DFT hardware work together to do the minimum number of tests necessary to see all the possible faults. No accepted fault models exist in the analog world, however, so there is no automated way of collapsing an analog network into a minimum set of paths and tests that would prove the structure of the entire circuit.

There is a glimmer of hope—adaptive test—but it comes from statistical, not structural, analysis. Adaptive test collects all the test data as a design runs through production test, applies statistical analysis to all of the data, and identifies tests that detect no unique faults. For instance, a test may find only those faults that another test also finds. If that other test finds additional faults, as well, adaptive-test programs would remove the first test and keep the second. So without fault models or even a knowledge of the circuit topology, adaptive testing can in principle narrow down the test suite until only the minimum necessary number of tests remain.

This scenario is complex, however. It requires a huge amount of data to make confident decisions. You don't want to eliminate a test only to discover that it was the only way to spot a rare failure. "Most analog parts never reach the kind of volume you'd like for those kinds of decisions," says Intersil's Purtell.

Adaptive test, like analog structural test, is still in its childhood. It's probably not yet possible to estimate the potential for either technique. Someone could propose a powerful set of analog fault models tomorrow, for example, and suddenly make structural approaches much more attractive. But such discoveries are unpredictable. So the industry now faces a problem. Some test techniques are becoming too expensive to use at advanced process nodes. On-chip instruments are becoming more popular but are expensive to design and require verification for accuracy. Both structural and statistical analytical techniques for varying reasons don't yet help much with analog. AMS test is an important problem with no easy solution on the horizon. EDN

FOR MORE INFORMATION

Cadence Design Systems
www.cadence.com

International Test Conference
www.itctestweek.org

Intersil
www.intersil.com

Mentor Graphics
www.mentor.com

Nuvoton
www.nuvoton.com

Qualcomm
www.qualcomm.com

Synopsys
www.synopsys.com

Texas Instruments
www.ti.com

You can reach Executive Editor Ron Wilson at 1-510-744-1263 and ronald.wilson@reedbusiness.com.

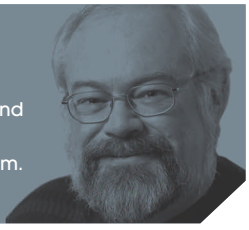

Powerful, Versatile and Affordable...

100 kHz Dynamic Signal Analyzer

- ***DC to 100 kHz frequency range***
- ***100 kHz real-time bandwidth***
- ***Dynamic range***
  - ***90 dB (FFT)***
  - ***145 dB (swept-sine)***
- ***Low-distortion (–80 dBc) source***
- ***Up to 32 Mbyte memory***
- ***GPIB and RS-232 interfaces***

FFT analyzers starting at $4950

The SR785 Dynamic Signal Analyzer offers state-of-art performance at a fraction of the cost of competitive analyzers. Standard features include swept-sine mode, order tracking, octave analysis, curve fit/synthesis, and arbitrary waveform source.

When compared to the Agilent 35670A, the SR785 comes out on top. It offers twice the frequency range (2 ch.) and 10 times the real-time bandwidth at less than half the price of the 35670A.

The SR785 is ideal for filter design, control systems analysis, noise measurement, audio or acoustical test, and mechanical systems analysis.

Take a close look at the SR785.

Stanford Research Systems
1290-D Reamwood Ave., Sunnyvale, CA 94089 • e-mail: info@thinkSRS.com
Phone (408) 744-9040 • Fax (408) 744-9049 • **www.thinkSRS.com**

More power. Less rack space.

How about up to 5200W in 2U?

Limited rack space doesn't have to limit power. Introducing the new Agilent N8700 Series DC power supplies, offering the highest power density available. Imagine up to 5200W in a mere 2U of rack space. Or consider the compact and reliable Agilent N5700 Series (1U). Both include standard GPIB, LAN (LXI C) and USB interfaces, not to mention built-in voltage and current measurements for simplified test system set up. So you can accomplish more, in less space.

DC Output Ratings	N8700 Series	N5700 Series
	(21 models)	(24 models)
Max. Voltage	600V	600V
Max. Current	400A	180A
Max. Power	5200W	1560W
Rack Height	2U	1U

Increase your power—download the latest application notes today!
www.agilent.com/find/morepower

u.s. 1-800-829-4444 canada: 1-877-894-4414

Agilent Technologies

BY MARTIN ROWE • SENIOR TECHNICAL EDITOR, *TEST & MEASUREMENT WORLD*

ANECHOIC CHAMBERS RISE FROM THE PITS

THE DESIGN, CONSTRUCTION, AND QUALIFICATION OF A CHAMBER DEPEND ON THE PRODUCTS IT WILL TEST.

Tests for electromagnetic immunity and emissions work best when located far from intentional transmitters such as broadcast stations and cell towers. Testing products for emissions or immunity in populated areas requires an anechoic or semianechoic chamber to keep ambient signals from interfering with measurements. Every chamber installation is unique and is influenced by many factors, some of which can change over time.

Major factors affecting chamber design include the size of the EUT (equipment under test), the relevant standards for EMC (electromagnetic compatibility), and the building facilities. A company that manufactures handheld devices won't need a chamber as large as one that makes rack-mounted equipment or vehicles. Chambers also differ depending on whether they must comply with commercial or with military standards. Chambers exclusively for emissions tests differ from those that also host immunity tests. Chambers for precompliance tests and troubleshooting may also differ in size and con-

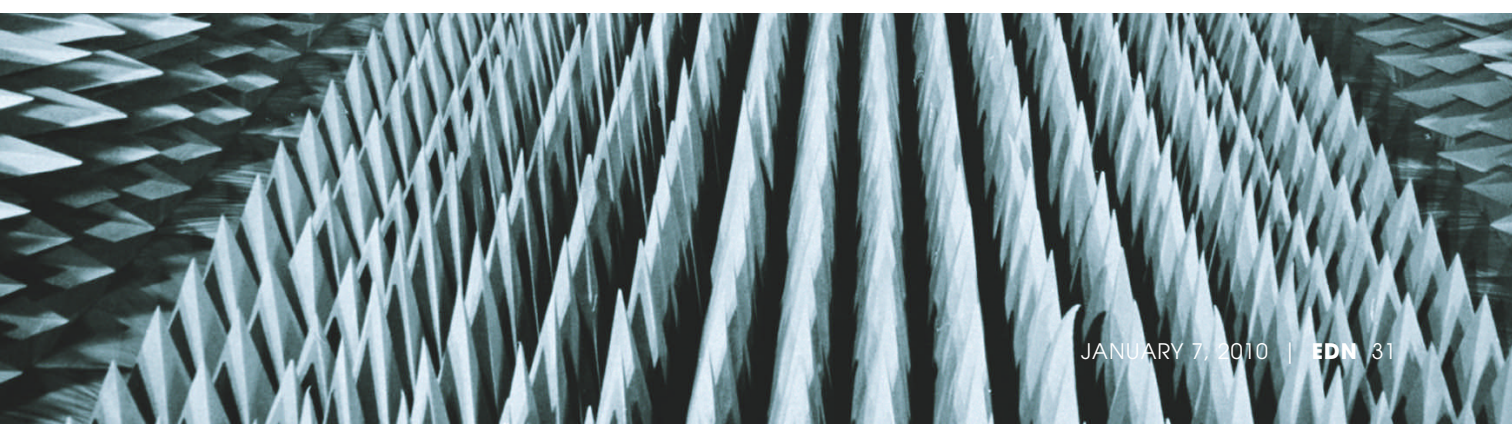

struction from those for full-compliance tests. The **sidebar** "What is an anechoic chamber?" explains how chamber components minimize signals inside the room.

KNOW WHAT GOES IN

Proper chamber design starts with a knowledge of the products that you need to test. "Have a good idea of your product road map over the next several years because test requirements tend to grow over time," advises Bryan Sayler, senior vice president and general manager of ETS-Lindgren. Planning for future products eases decisions about a chamber's size, turntable size, cable routing, power provisioning, and instrumentation.

National and international EMC standards significantly influence chamber size because they often specify the distance from the antenna to the EUT: 1, 3, or 10m. CISPR 22, for example, requires a 10m distance for EUTs larger than $2m^3$ (**Reference 1**). Product size affects chamber size because a chamber must be large enough to accommodate the distance from the EUT to the antenna, plus the size of the EUT, plus the size of any internal absorber materials. Absorbers such as cones that line a chamber's walls significantly reduce the usable area.

Chambers may also accommodate other distances, such as 5m. "Some test labs will use a 5m distance for precompliance tests because it better correlates to how an EUT will perform at 10m than a 3m distance," notes Peggy Girard, president of Panashield. Some chambers accommodate a 5m path, but the FCC (Federal Communications Commission) currently accepts only the chamber based on the 3m test data because the agency doesn't formally recognize 5m data in its comparison with an OATS (open-area test site).

The construction of the host building also affects chamber size, particularly height. When engineers at Northwest EMC opened a facility in June 2009 in Brooklyn Park, MN, they wanted a location close to potential customers, which ruled out an OATS. Tim O'Shea, operations manager at the company, explains that ceiling height was a critical factor because many cities have building-height restrictions, and he needed a building with ceilings high enough to house a 10m chamber, which can typically be 30 to 32 feet tall.

AT A GLANCE

- Anechoic chambers shield a test setup from ambient signals and absorb reflected signals generated inside.
- Factors that affect chamber design include the size of the EUT (equipment under test), relevant standards for EMC (electromagnetic compatibility), and building facilities.
- A typical chamber has a raised working floor above a pit, which provides space for EUT-support equipment and cable routing from the control room to the center of the floor's turntable.
- Because anechoic chambers are shielded rooms, they must prove that they sufficiently attenuate outside signals before a manufacturer installs internal absorbers.

RISING FROM THE PIT

Height isn't the only consideration for a building to house an anechoic chamber. Depth counts, too. **Figure 1** shows a typical chamber with a raised working floor. The chamber floor is supported on legs because there's a pit under the floor. The pit, typically 18 to 24 in. deep, provides space for EUT-support equipment and cable routing from the control room to the center of the floor's turntable. The raised floor is flush with the building floor outside the chamber, which lets technicians roll equipment into the chamber.

Girard describes an unusual chamber that, because of low ceilings in the building that housed it, required a "pit within a pit." **Figure 2** shows that the chamber's raised floor was actually 6 feet below the building's floor. The 6-foot pit was larger than the chamber, providing working space outside the chamber at the chamber floor's level. A second pit extended 18 in. below the chamber's floor to provide space for cables and turntable motors.

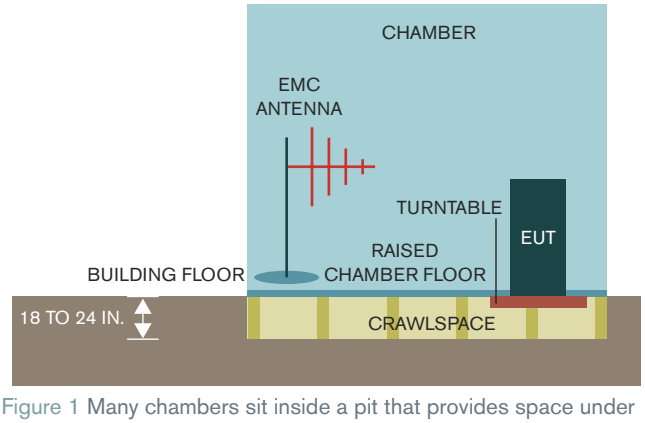

Figure 1 Many chambers sit inside a pit that provides space under the working floor.

Pits under chambers need not be as large as the entire chamber. To save on construction costs, some companies and test labs choose to leave much of the area under the chamber floor solid and provide just conduits for cables to the EUT, turntable, and antenna. This design costs less to dig; however, it minimizes flexibility. Chambers with conduits also need metal pipes around the cables for electromagnetic shielding.

Instead of having just the 18- to 24-in. space under their raised floors, some chambers have deeper pits under the turntable or antenna. Ghery Pettit, EMC regulatory compliance manager at Intel, chose to construct a chamber with a 6-foot pit under the antenna. That depth provides space for RF amplifiers used in immunity tests. Keeping amplifiers close to an antenna results in shorter cables that carry RF signals. "We really wanted a 7-foot pit because 6 feet isn't deep enough for some people to stand," says Pettit. "When we built a chamber in Dupont, WA, we dug a 7-foot pit so that anyone could stand upright in it."

Other chambers have no pits under them at all. Instead, the base of the chamber rests on the building floor, and the chamber has a raised floor under which the cables run. In that case, an equipment ramp lets technicians roll equipment into the chamber (**Figure 3**).

Cables under a chamber floor provide power, control, and I/O signals to an EUT. RF cables connect antennas in the chamber to RF amplifiers for immunity tests or to EMI (electromagnetic-interference) receivers and spectrum analyzers for emissions tests. Other cables provide power and control signals for turntables and antenna masts. Test equipment usu-

ally resides in a shielded control room close to the chamber.

CONNECTING THE CABLES

Because test equipment resides in a separate control room, all chambers need penetration panels that hold cable connectors. Cables on either side of the panels connect equipment in the control room to the EUT, turntable, and antenna inside the chamber. The control room should be close to the portion of the chamber that holds the panels to minimize cable length.

For commercial EMC tests, cables that come from outside the chamber generally run under the floor. The penetration panels may be part of a chamber's walls, or they can be on the floor next to the chamber, provided that there's space under the panel for cable runs. At Northwest EMC's Brooklyn Park facility, penetration panels are in the chamber's wall next to the control room but below the raised floor's surface (**Figure 4**). At Intertek's facility in Boxborough, MA, connector panels are on the floor of the control room (**Figure 5**). Cables then run through the pit under the floor to the center of the turntable. The raised chamber floor has panels through which

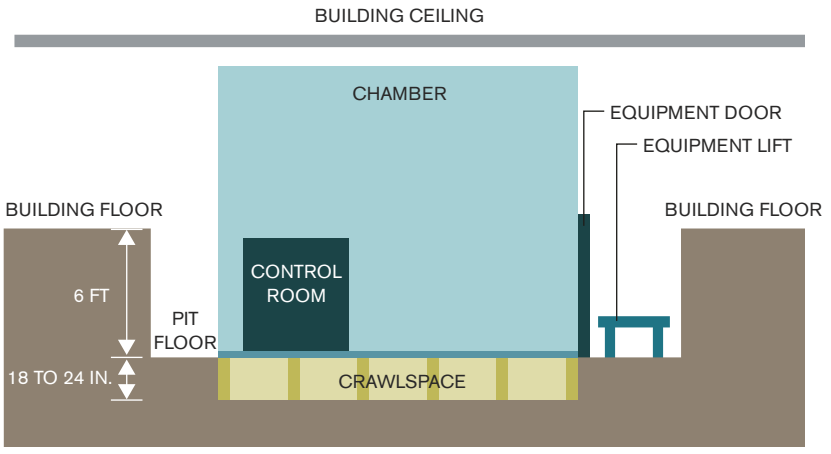

Figure 2 One chamber design needed a 6-foot pit to clear a building's ceiling.

WHAT IS AN ANECHOIC CHAMBER?

Anechoic chambers perform two basic functions as part of an overall EMC (electromagnetic-compatibility) measurement system. They shield a test setup from ambient signals, and they absorb reflected signals generated inside.

The walls and ceiling of a chamber are lined with metal that connects to a grounded metal floor, which serves as the ground plane that many EMC standards require. The metal lining, both inside and outside the walls, attenuates signals by providing a low-impedance path to the earth ground. A typical chamber can attenuate signals by at least 100 dB.

Unintentional emissions, radiated either by the EUT (equipment under test) or by the signals intentionally transmitted from an antenna, bounce off the shielded walls, creating undesirable fields inside the chamber. Thus, anechoic chambers need absorbing materials to minimize reflections.

A typical chamber wall consists of a wood panel sandwiched between two metal layers (Figure A). The inside metal layer is lined with ferrite tiles and absorber cones, typically referred to as hybrid absorbers, which are impedance-matched to the ferrite to allow for testing across the full range from 30 MHz to greater than 40 GHz.

The ferrite tiles absorb signals up to about 1 GHz. If a chamber won't see higher frequencies, then it won't need the hybrid cones on top of the tiles. The cones, up to about 4 feet long, significantly reduce chamber size. They're typically made of carbon-loaded polyurethane foam, polystyrene foam, or fibrous material.

A chamber's raised floor is an integral part of the chamber's shielding and grounding. For radiated emissions tests, the floor is grounded and is attached to the chamber walls. Radiated emissions standards such as ANSI C63.4 and CISPR 22 require a grounded metal floor that simulates signals bouncing off the ground of an outdoor facility. Because the floor isn't covered with absorbing materials, the chamber is called semianechoic rather than fully anechoic.

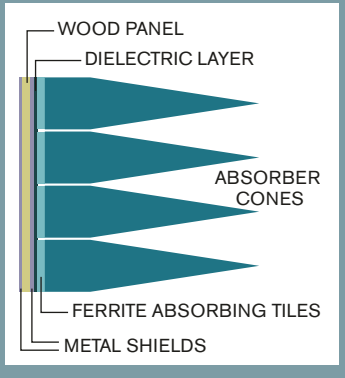

Figure A Chamber walls consist of shielding layers and absorbing materials.

Radiated immunity standards, such as EN 61000-4-3, require a partial floor coverage of absorbers between the front of the uniform field and the antenna for a 3m path length, typically 10×11 feet. Thus, absorbers must be on all six surfaces of the chamber. If a chamber is for both commercial emissions and immunity tests, it will need removable absorbers for the floor.

Anechoic chambers used for commercial EMC emissions and immunity tests need a turntable that rotates the EUT, exposing all sides to an EMC antenna. EUT size and weight dictate the size of the turntable. Ghery Pettit, EMC regulatory compliance manager at Intel, has built three anechoic chambers in his career. The first, built in 1989 and still in use, has an 18-foot-diameter turntable that can support 20,000 lbs.

Standards also require tests over different frequency ranges. ANSI C63.4 and CISPR 22 originally called for frequencies from 30 MHz to 1 GHz. CISPR 16-1-4 covers frequencies from 1 GHz to 6 GHz and up to 18 GHz. Military standards, such as MIL-STD (military standard) 461 call for conducted tests from 30 Hz to 80 MHz and for radiated tests from above 80 MHz to 40 GHz.

Figure 3 A chamber without a pit requires an equipment ramp. A swinging door lined only with ferrite tiles clears the door opening (courtesy Panashield).

Figure 4 A penetration panel sits below the building floor, providing access to cables that run under the chamber's raised floor. A pipe opening alongside the panel provides access for additional cables and hoses (courtesy Northwest EMC).

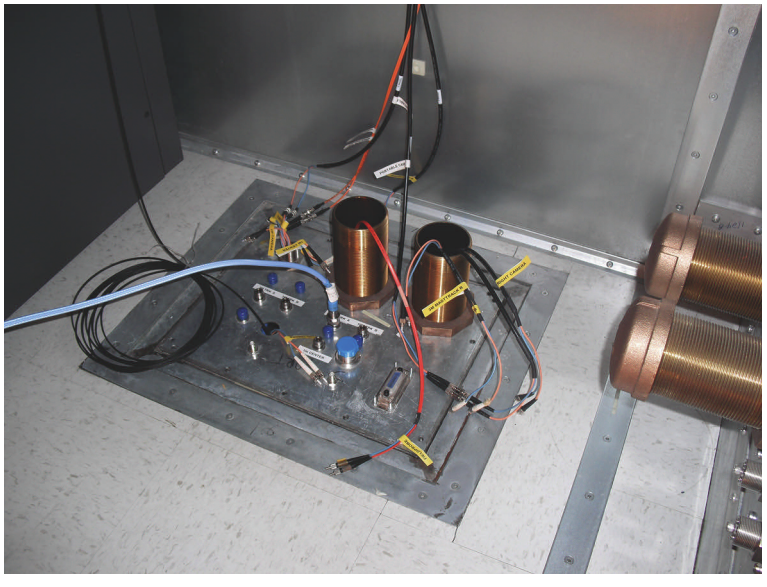

Figure 5 A floor-mounted panel provides power and signal cables with access to the inside of a chamber (courtesy Intertek).

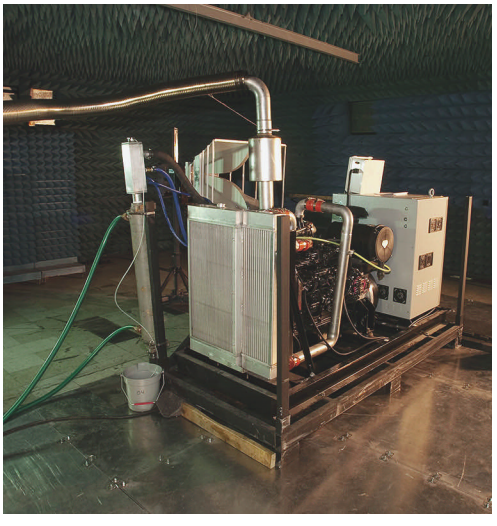

Figure 6 Chambers for MIL-STD 461-compliance tests don't need turntables; cables for these tests can run on the chamber floor (courtesy Dayton T Brown).

the antenna cables can emerge to reach the antenna and its mast.

For military tests, cables must run on the chamber floor because that arrangement simulates actual use. In addition, EUTs need not rotate. Tom Arcati, an engineering specialist at Dayton T Brown, thus uses chambers without pits or turntables (**Figure 6**). That setup simplifies the chamber's design and makes it less expensive than chambers for testing to commercial standards. The penetration panel is on the chamber wall, slightly above the floor. "MIL-STD [Military Standard] 461 requires at least 10m of cable between the antenna and EUT," says Arcati. "Power cables from the LISN [line-impedance-stabilization network] and the equipment must be at least 2m long."

EUTs often need external ac or dc power, so a chamber must be wired to make power available. The ac or dc power that a chamber must support varies with how the chamber is used. Independent EMC test labs test a wide range of powered products and require

more forms of power than do chambers dedicated to one company. The type of power also depends on national compliance standards. For example, a chamber might need 100V for Japan, 110V for Taiwan, 120V for North America, and 220V for Europe. It may need to provide single- or three-phase power frequencies of 50, 60, or 440 Hz.

The ac mains lines that enter a chamber require filtering to minimize noise. Some chambers filter their ac voltages using a single filter for each voltage and then distribute the clean power to the equipment inside the chamber. Some EMC engineers, however, prefer to distribute the power to its load and filter it there. The choice depends on cost versus flexibility.

Because power and signal cables must connect to an EUT on a turntable, the cables must have enough slack to accommodate turntable rotations. "Most turntables have 380° of rotation," says ETS-Lindgren's Sayler. "The turntable typically rotates ±180° during a test." A more expensive approach and one that he rarely sees uses slip rings that conduct power and signals while permitting continuous rotation.

DESIGNING THE DOORS

Regardless of whether a chamber is for commercial or military tests, it needs a door large enough to get an EUT in and out. The size and design of a door also depend on the products tested in the chamber. Size and location are important. "The door locations should minimize the distance to the turntable if the chamber will test large, heavy EUTs," notes Intel's Pettit.

Door location matters for more than just convenience. "There are electrically good and bad places for chamber doors," Sayler warns. "If the door lacks cones, it should be as far away from the turntable as possible because it will affect the quiet zone around the turntable."

Engineers at Northwest EMC use their chamber for radiated emissions only, and the EUTs are typically small enough to fit onto a table. Thus, the chamber has two 4×8-foot swinging doors for equipment entry and exit. The insides of the doors don't have absorbing cones because they would interfere with the door opening in the chamber wall. The door has ferrite tiles lining its inside surface. **Figure 3** shows a single door lined with ferrite tiles,

PLAN AHEAD. LOOK AT THE WHOLE FACILITY AND HOW THE CHAMBERS WILL FIT INTO THE BUILDING.

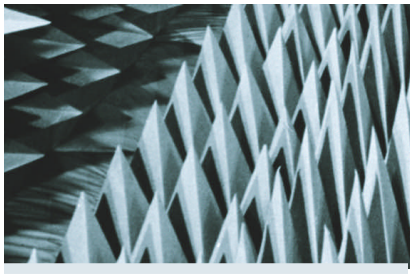

but no cones, so the door can clear the door opening. In contrast, the chamber at Intertek's Boxborough facility has a retractable door and thus has cones but costs considerably more than a chamber with a swinging door. Swinging door leaves can have cones over the ferrites provided that the cones can clear the door opening (**Reference 2**).

QUALIFYING A CHAMBER

Because anechoic chambers are shielded rooms, they must prove that they sufficiently attenuate outside signals before a manufacturer installs internal absorbers. Most EMC engineers specify that chambers attenuate outside signals by at least 100 dB. "We can build chambers that attenuate up to 120 dB for areas high in ambient signals," says Sayler. "For most applications, 80 dB of attenuation is enough." The ambient noise levels inside the chamber are those that count. If a chamber is in a region of relatively low ambient signals, then 80 dB of attenuation may be sufficient.

A technician tests the chamber's shielding effectiveness by generating a signal at a known power and frequency and measuring signal strength on the other side of the wall. O'Shea explains the test procedure at Northwest EMC. "We had the transmit antenna at 26 locations outside the room—including the top—and moved the receive antenna along all inside seams closest to the transmit position," he says. "We checked every seam in the room."

Technicians may also perform shielding-effectiveness tests with the transmitting antenna outside the chamber and the receiving antenna inside. Placing the receiving antenna in the chamber reduces ambient signals at the receiver, which makes the transmitted signals easier to see on a spectrum analyzer.

Following a shielding-effectiveness test, a chamber is ready for the absorbing materials. After installation, the chamber needs additional measurements around the turntable to verify its quiet zone for radiated emissions tests per ANSI C63.4 from 30 MHz to 1 GHz and per CISPR 16-1-4 (**references 3** and **4**). For radiated immunity tests, a chamber must undergo a field-uniformity test to comply with EN 61000-4-3 (**Reference 5**).

DO THIS, NOT THAT

When building an anechoic chamber, Panashield's Girard urges an understanding of your needs. "Start by knowing what your present and future testing needs will be for your product, per international standards," she says. "Will you require radiated emissions, immunity, or both?"

"Plan ahead. Look at the whole facility and how the chambers will fit into the building," says O'Shea of Northwest EMC. "Evaluate customer needs to customize room size, turntable size, door size, and power requirements. Don't take short cuts in the building process or during shielding-effectiveness testing because they could cause problems that are much more difficult to fix once the room is fully constructed."

"Get on the good side of your facilities people," says Sayler of ETS-Lindgren. "You'll need them."

And Intel's Pettit expresses concern for ac mains power. "Have plenty of power, at least 100A service for each voltage," he says. "Separate power feeds from the turntable to the EUT so you won't get interference." EDN

For a list of the references cited in this article, as well as more information on the companies it mentions, go to www.edn.com/article/100107df1.

For more technical articles, go to www.edn.com/features.

A version of this article appeared in the December 2009/January 2010 issue of *EDN*'s sister publication *Test & Measurement World*.

Are you up for a challenge?

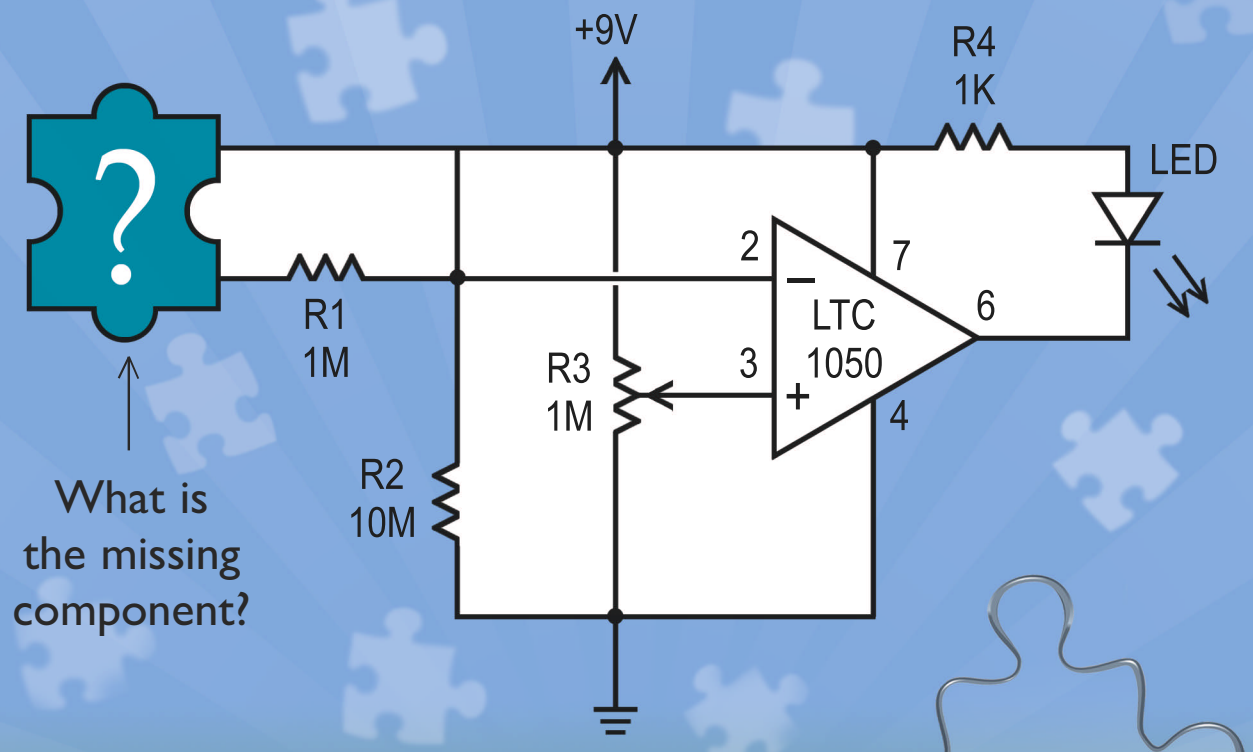

What is the missing component?

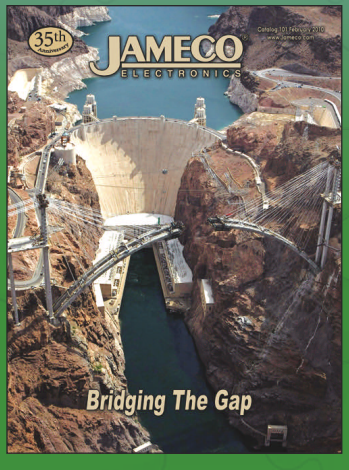

Industry guru Forrest M. Mims III has created yet another stumper. The Ultra Simple Sensors Company assigned its engineering staff to design a circuit that would trigger an LED when a few millimeters of water is present in a basement or boat. What is the water sensor behind the puzzle piece? Go to www.Jameco.com/brain5 to see if you are correct and while you are there, sign-up for our free full color catalog.

JAMECO® ELECTRONICS

1-800-831-4242 | www.Jameco.com

BY ZHI-HERN LOH • FULCRUM MICROSYSTEMS

Congestion management clears a path through 10 GbE

WITH THE NEW IEEE STANDARDS FOR CONSOLIDATION AND THE ABILITY TO BUILD HIGHLY SCALABLE HIGH-SPEED ETHERNET NETWORKS, ETHERNET WILL BECOME THE ONLY FABRIC YOU NEED FOR CONVERGED DATA CENTERS, AND IT WILL LOWER THE COST OF SUPERCOMPUTER NETWORKS.

Ethernet first emerged in the 1970s, and the IEEE standardized it in 1985. Over the years, the technology has undergone numerous updates. In today's data centers, the newest trend is network convergence using Ethernet as a consolidated fabric replacing traditionally separate special-purpose fabrics, such as FC (Fibre Channel) for storage, InfiniBand or Myrinet for computer-cluster IPC (interprocess communication), and Ethernet for LAN (local-area-network) traffic. With the introduction of 10 GbE (10-Gbps Ethernet), link speed has become competitive with these other technologies, but speed alone is not sufficient for using Ethernet as a converged fabric because Ethernet does not usually supply other necessary data-center features. For example, FC provides a lossless fabric for the SAN (storage-area network), and InfiniBand provides advanced congestion management, neither of which Ethernet supports.

The IEEE DCB (Data Center Bridging) Task Group is developing three new standards—802.1Qaz ETS (enhanced transmission selection), 802.1Qbb PFC (priority flow control), and 802.1Qau QCN (quantized congestion notification)—to address network-convergence issues (**Reference 1**). Although Ethernet switches have historically supported 802.1Q traffic classes, the IEEE has not standardized the administration of the traffic classes, so vendors have implemented proprietary methods of defining scheduling policies for each traffic class.

ETS strives to provide one policy that the administrative domain will consistently apply. ETS enables an administrator to define per-traffic-class bandwidth allocation and strict prioritization. An expected usage of ETS is for consolidating IPC, SAN, and LAN traffic classes. The IPC traffic is latency-sensitive, and the IEEE would configure it as strict high priority. You can limit the bandwidth of bursty SAN traffic to prevent starvation of other traffic classes. LAN traffic, which is typically best-effort traffic, would use strict low priority to minimize the latency of other types of traffic.

Ethernet has for years supported lossless operation; however, the Ethernet Pause standard lacks the per-priority flow-control ability that today's data centers require. Whereas all protocols running over Ethernet must handle drops, some protocols have high drop-recovery penalties and need lossless operation for performance. For example, the designers of the FCOE (Fibre Channel Over Ethernet) Protocol intended it for a lossless FC fabric, and it thus lacked the ability to retransmit single frames. Instead, on drops, FCOE retries an entire transaction, which could be megabits long. In a data center, cluster computing also suffers from high performance penalties because of dropped frames. For example, if the system drops interprocess-synchronization frames, the whole cluster stalls while the parallel processes wait for a time-out and retransmission of the dropped message. Moreover, traditional computer networks, such as InfiniBand and Myrinet, ensure reliable delivery, so most computing applications do not handle dropped frames. With the new PFC proposal, Ethernet can selectively provide flow control to the traffic classes serving these applications.

Issues arise with flow-control operation when you enable it by interchangeably using Pause and PFC. Pause prevents buffer overflows by stopping the flow of traffic on a port's incoming link, whether on a switch or a NIC (network interface card). On a switch, you typically implement Pause with thresholds on queue occupancy per ingress port. If an ingress port uses too much memory, a Pause-on frame will wait in a queue for transmission to the link partner after the current frame on the wire. After the short delay, the link partner will receive the Pause-on command and stop sending new frames.

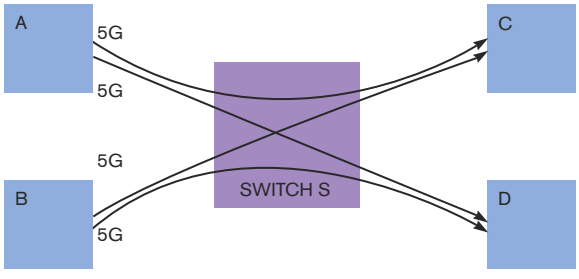

Figure 1 If sources A and B send data at 10 Gbps, or 5 Gbps per flow, and the sinks drain data at 10 Gbps, flow control is unnecessary. If, however, D can receive at only 1 Gbps, then the A-to-D and B-to-D flows must each slow down to 0.5 Gbps.

When the link partner stops, the frames in the switch memory eventually drain, and the queue occupancy for the paused ingress port drops below the Pause-off threshold, causing the transmission of a Pause-off frame to the link partner to unpause. On average, the Pause-on/Pause-off cycles result in the matching of the ingress rate and the maximum egress rate and thus make the switch lossless.

When you enable Ethernet flow control, you slow the link speed of the entire network to the speed of the slowest link. A common misconception is that the degradation is due to the fact that Pause is not working; however, it is actually due to congestion spreading and the penalization of innocent flows.

CONGESTION SPREADING

Congestion spreading can cause reduced performance in data-center networks (**Figure 1**). Consider a case in which two flows are on each source: A to C, A to D and B to C, B to D. Initially, each source sends data at 10 Gbps, or 5 Gbps per flow, and the sinks drain data at 10 Gbps. Clearly, there is no need for flow control because C and D can both fully drain the arriving traffic. Now, replace D with a 1-Gbps NIC. The A-to-D and B-to-D flows then must slow down to 0.5 Gbps each, a reduction of one-twentieth link capacity. Because Pause operates on the entire ingress port, however, the innocent A-to-C and B-to-C flows also slow down to 0.5 Gbps, even though they do not have to.

If the network includes multiple switch hops instead of one hop, such as replacing A and B in **Figure 1** with more switches, a congestion tree could form (**Figure 2**). Each of the upstream switches will see congestion on the downstream switches and propagate backward flow control on ingress links that talk to congested egress links. The flow control thus causes a congestion tree to appear, with Port D of S being the root of the tree. Eventually, all the flows passing through the congested links (the solid arrows in the **figure**) would slow down, even if some of them do not go through the root.

Usually, flow control handles short bursts of traffic with transient congestion patterns. When the bursts are short, a congestion tree does not have time to propagate. In the previous example, the upstream link's capacity decreases due to the long-lived congestion that occurs when you plug in the low-speed NIC in a misuse of Pause. However, legitimate cases exist in which multiple sources send to one sink for a long time. If this scenario were to happen in a 1 million-node data center, a congestion collapse could occur when 10-Gbps links decrease their speed to kilobits per second due to the possibility of high fan-ins to the congested port (**Reference 2**). Besides using PFC to limit the congestion spreading to select traffic classes, a CN (congestion-notification) mechanism could further limit the spread of congestion.

The DCB Task Group has introduced the IEEE 802.1Qau QCN project as a more precise flow-control mechanism. Cisco first presented QCN as BCN (backward congestion notification). In 2007, a number of proposals emerged, which the InfiniBand congestion-control architecture influenced (**Reference 3**). The latest version, QCN, is a hybrid of the previous ideas (**Reference 4**).

CN BACKGROUND

CN is a control mechanism like the TCP (Transmission Control Protocol). Because TCP and CN are similar in principle, you may wonder why it is not sufficient to run TCP. One reason is that TCP without a hardware-offload engine requires high CPU usage to maintain maximum throughput and minimize latency (**Reference 5**). TCP's developers intended it for congestion control in the Internet, in which the network's diameter is orders of magnitude larger than that of a data center. With the longer round-trip time and lower bandwidth of the Internet, keeping the link running at capacity requires less CPU usage. Another difference between TCP and CN is that TCP guarantees reliable and in-order delivery of frames. In the presence of congestion, this requirement causes large latency jitter due to retransmissions and time-outs. Reliable and in-order delivery is not always necessary; some data-center applications, such as streaming audio, video, and multiplayer games, favor timely delivery of data grams instead. A competing proposal from the IETF (Internet Engineering Task Force) that satisfies this purpose is the DCCP (Data-gram Congestion Control Protocol).

With respect to congestion signaling, CN has advantages over TCP and DCCP. In TCP, the senders back off by sensing drops due to overflows. However, backoffs are not immediate due to the time lag for either the receiver to notify the sender of missed sequence numbers or the sender to detect a drop due to a retransmission time-out. CN also has congestion predictors like TCP's RED (random early detection). CN relies on the congestion point's queue length and rate of change as predictors of future congestion.

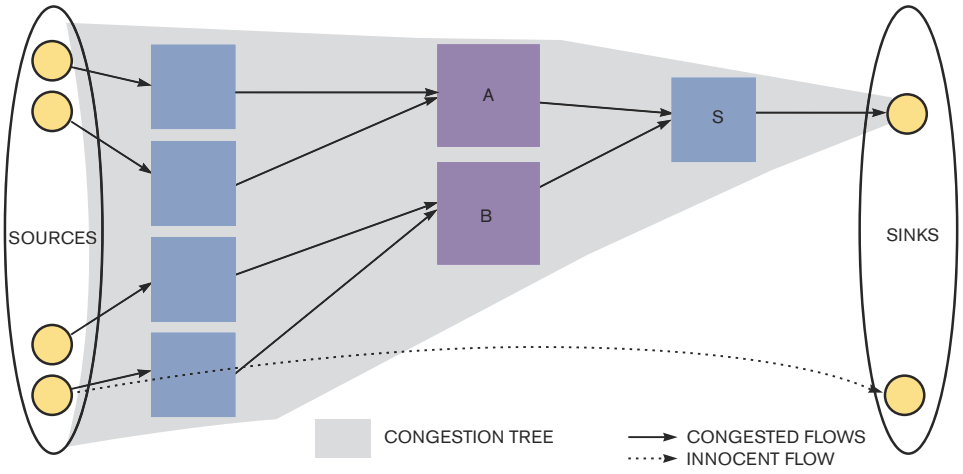

Figure 2 If the network of Figure 1 included multiple switch hops, a congestion tree could form.

Also, CN's multilevel feedback from congestion points allows more precise congestion response than does the 1-bit ECN (explicit-congestion-notification) marking that TCP and DCCP use. In addition, some important data-center protocols, such as FCOE, can benefit from CN because they usually run without a congestion-control algorithm, causing poor network performance due to congestion spreading.

BASIC OPERATION

For CN, each source has a rate limiter, which is usually a token bucket; sensing the congestion in the network controls the rates of these sources. Congestion points detect congestion by monitoring the queue length and report congestion by sending CN messages back to the sources that have the rate-limiting reaction points. Sampling the frames going through the congestion point generates the CN messages. By using the source address in the sample, you can send a CN message back to a reaction point that is contributing to congestion. The sampling interval at the congestion point is also a function of the congestion. For example, in times of congestion, the sampling occurs more frequently to increase the back-off signaling.

Ideally, you implement reaction points in NICs, with each flow having its own reaction point with a dedicated congestion-control state. For rate decreases, you compute the decrease amount from the feedback in the CN message that the congestion point returns. QCN uses self-increases to autonomously compute rate increases.

CN is essentially a control loop that maintains maximum throughput at the congestion point. As in control theory, the CN control loop uses a proportional-integral controller in which the control feedback is proportional to the difference in the arrival and departure rates of the congestion point's queue and the integral of the rate difference.

In control theory, a control loop aims to keep the process variable—the traffic rate at the congestion point, in this case—at the setpoint. In normal operation, the congestion point's queue length never actually stays on the setpoint but oscillates about it. In QCN, these oscillations arise because of a cycle in which the reaction points autonomously perform self-increases, increasing their transmission rates until the congestion point detects the onset of congestion and signals the reaction points to slow down.

Due to the CN message packet propagation delay, a lag occurs before the congestion point observes a new rate from feedback events. Increasing delay increases the magnitude of the oscillations. To compensate, you can tune the gain parameters of the CN's control loop. For example, many gain parameters in QCN control the amplitude of the reaction to congestion. The feedback calculation at the congestion point uses some of these parameters, and the rate-change calculations at the reaction point use others. By tuning the gain, you can optimize the control loop's response.

As in any control loop, instability can occur. The amplitudes of the oscillations grow instead of attenuating, eventually resulting in loss of throughput due to underflows and overflows. Decreasing the gain may improve stability but at the cost of reduced responsiveness. The IEEE's CN simulations show that instability can occur (**Reference 6**). However, they also show that, within their data-center parameters, you can use one set of static-gain parameters to stabilize the CN loop (**references 7** and **8**).

Other possible causes of instability may be mixed-speed networks. Optimal parameters for 1 and 10 Gbps differ, so you must make trade-offs. With the upcoming 40- and 100-Gbps networks, interoperating with 10-Gbps CN devices represents a big unknown. Simulations have yet to show the effect of these higher-speed links on CN performance. CN also has yet to show that it will not negatively affect TCP flows. TCP friendliness is a good metric for measuring new CN protocols. DCCP explicitly states it as a goal. CN has also made this feature a goal, but it is yet to be proved in real-world applications. The issue is that TCP requires drops for sensing congestion and applying the back-off mechanism. Some TCP flavors, such as New Reno, can use other metrics, such as round-trip time, for sensing congestion and thus interact better with CN. You can also use ECN marking to notify TCP of congestion. Some IEEE simulations have illustrated these scenarios.

DEALING WITH SPECIFICATION INSTABILITIES REQUIRES A HYBRID APPROACH OF MIXED FIRMWARE AND HARDWARE SUPPORT.

IMPLICATIONS FOR SWITCHES

Traditionally, switches have offloaded control-plane-processing operations to a multipurpose CPU. Unfortunately for CN, increasing stability margins require CN messages to provide quick feedback; thus, you must accelerate frame generation as much as possible with hardware. Switches implement the line-rate data-forwarding logic in silicon, whereas software implements the lower event rate but more functionally complex routing-table maintenance functions. The flexibility of software lowers the cost of maintenance and problem corrections. One issue with implementing CN with fast fixed logic is that the IEEE specification is new and likely to change. Dealing with these specification instabilities requires a hybrid approach of mixed firmware and hardware support to allow flexibility for handling specification changes and the ability to operate necessary features at line rate. For example, in most IEEE specifications, the frame formats are usually the last to stabilize. By offloading some of the congestion point's functions to a programmable device, you can maximize compliance with the standard. Moreover, the increased flexibility allows experimentation with the algorithm for improving performance.

Although the 802.1Qau QCN standard is new, it has been in progress for several years, leading to many prestandard implementations. Interoperability between vendors' implementations will be a big hurdle in the near future. Before the IEEE ratifies the standard, using programmable devices will improve interoperability. An example was a recent IEEE QCN demonstration of an NEC NIC with a prestandard Fulcrum Microsystems switch. Although the switch implemented much of the congestion point's logic in silicon, the FPGA-based NIC pro-

vided enough performance and flexibility to implement QCN.

Another step toward deployment involves a switch-only QCN network. With more head-of-line blocking than a NIC-based reaction point, switches could implement the reaction point on-chip and convert the flow control into 802.1Qbb PFC signaling. A small data center could be a good place to evaluate this function because it would be easy to limit the switches to a single vendor.

The DCBX (Data Center Bridge Exchange) Protocol is yet another feature that will aid deployment. With DCBX, a network can automatically negotiate advanced capabilities, such as Qau for CN, Qbb for Pause, and Qaz for ETS. Each link would use the DCBX Protocol to discover its neighbor's capabilities and automatically turn on the common features. When you replace switches and NICs in the cloud with Qau-capable devices, the CN domain would automatically expand.

DEPLOYMENT IN HPC NETWORKS

Another type of likely deployment outside the scope of the standard is for the switches to enable congestion points but not to enable the reaction points in the NICs. HPC (high-performance-computing) users have indicated that they would welcome the congestion point's feedback for use in their applications. One use of the feedback would be for adaptive routing, enabling load balancing over multiple paths to avoid dynamic hot spots and for application tuning to avoid static hot spots. Adaptive routing can improve performance, but Ethernet does not natively support it. By default, Ethernet runs the spanning-tree protocol to cut loops in the connectivity graph. Loops can cause frame duplication when flooding occurs. While cutting loops, however, spanning trees also cut off alternative paths between sources and sinks, reducing the usable network bandwidth. Adaptive routing can recover these additional paths.

In a random network, loop prevention can be difficult because loops can form over multiple hops. However, with the fat-tree topology, or Clos network, that you commonly find in large data centers and supercomputers, loop prevention is simpler because of the fat tree's regularity. The IBM MareNostrum uses a multistage Clos built on special-purpose Myrinet switches. With the recent enhancements of Ethernet and the availability of low-latency Ethernet switches, however, new supercomputers can deploy Ethernet instead.

An important property of the fat tree is the constant bisectional bandwidth over each stage. Ideal fat trees are nonblocking, meaning that you can connect any free input port to any free output port without affecting connections. However, due to the packet nature of Ethernet and flow-ordering constraints, an Ethernet fat tree is not exactly nonblocking. **Figure 3** illustrates a three-stage fat tree using switches of degree N. A stage is one switch hop. The fat tree comprises line and spine switches. The line switches are externally facing; the spines handle internal connections. When traffic ingresses onto a line switch, it transmits onto one of the N/2 uplinks to a spine switch. Although all of the spine switches connect to the egress line switch, picking one to send the traffic is not trivial. Poor spine-selection algorithms can lead to load imbalance and create congestion on the uplinks in the line switches.

CRC HASHING DISTRIBUTES A RANDOM IDENTIFICATION WELL, BUT PEARSON'S HASHING BETTER SUITS CONSECUTIVE IDENTIFICATIONS.

You can easily scale fat trees by recursively using smaller fat trees as the spine switches and adding line switches on the bottom externally facing layer. The number of ports grows exponentially with the number of stages. For example, using a 24-port switch, you can build a 288-port, three-stage; 3456-port, five-stage; or 41,472-port, seven-stage fat tree.

Unlike with fat trees, you cannot easily scale spanning trees because doing so turns off many uplinks, motivating a need for a different uplink-routing algorithm. The routing protocol must preserve frame ordering, though, because reordering can cause accuracy problems in applications that assume ordered frames. Even in applications that do not assume in-order delivery, such as those that run over the IP (Internet Protocol) and TCP, reordering can confuse the congestion-control algorithm, resulting in loss of performance (**Reference 9**).

The simplest form of such an algorithm that also performs load balancing is a hash function. A flow identifier hashes to an uplink port, thus ensuring that all frames of a flow follow the same path. Multiple flows typically use all uplinks, improving load balancing. Common flow identifiers include the five-tuple address, source IP address and port, destination IP address and port, and protocol number.

ADAPTIVE MULTIPATH NETWORKS

Although hashing may seem like a good approach to building fat trees, the technique involves some subtleties. For example, the various ingress stages to the spine in a multistage fat tree must use independent hash functions to avoid hash polarization (**Reference 10**). Hash polarization occurs when the hash function produces the same results on multiple stages, such as when Stage S1 sends to its uplink P1 to a switch on Stage S2, which in turn must use only one uplink, P1, because the hash function is the same. As a result, the second stage uses only one of the N/2 possible uplinks. For an M×2+1-stage network, this approach reduces the effective bandwidth by a factor of (N/2)M. Even if subsequent stages use different hash functions, correlations would also cause some polarization.

To provide the hash-function performance necessary for load balancing within the fat tree, the latest Fulcrum Microsystems switch uses a combination of CRC (cyclic-redundancy-check) functions and Pearson's hash to yield nearly perfect load balancing for large numbers of both randomly and consecutively numbered flow identifications. CRC hashing distributes a random identification well, but Pearson's hashing better suits consecutive identifications. Although most data-center applications have large numbers of flows and thus hash evenly, some applications, such as HPC, have few high-bandwidth flows. With fewer flows to hash, the hash function is

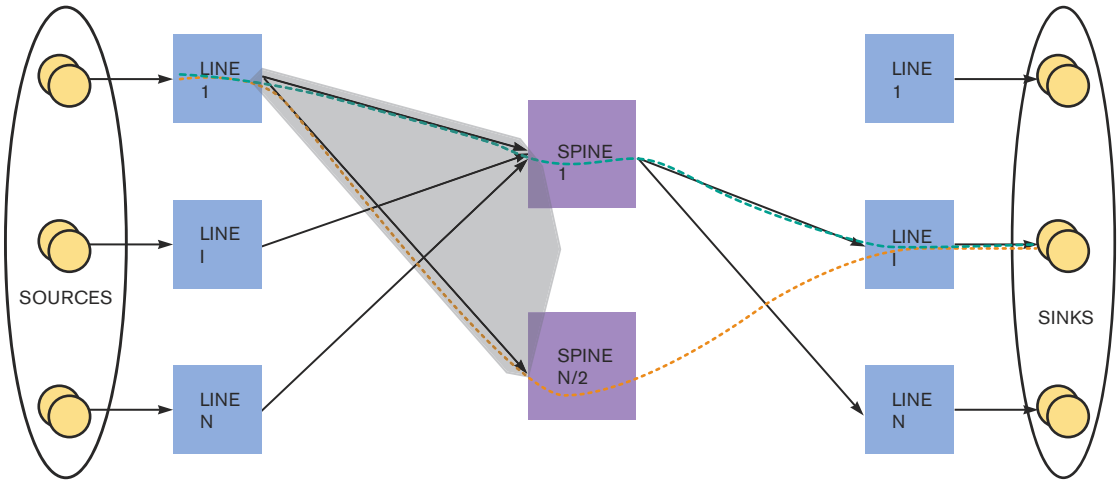

Figure 3 In this three-stage fat tree, line switches pick from N/2 uplinks for each flow. Spine switches have only one destination line.

unlikely to perfectly distribute the flows. If the arrival rate is high, persistent congestion will occur. This issue has led some vendors to add adaptive-routing capability, such as the Fortinet VScale technology, which monitors fabric latency and assigns flows to paths in a manner that minimizes the switching latency through the fat tree.

Adaptive routing can also act as a control loop, meaning that it is possible to observe oscillations in the system. When flows move from congested uplinks to uncongested ones, the load-balancing algorithm can overcompensate and overload the previously uncongested uplink, leading to oscillations in the queue length and frame latency, which increase jitter. Another difficulty with adaptive routing is that routing cycles may form when you use it in conjunction with link-level flow control, causing deadlock situations in the network.

Large fat trees have posed a challenge to CN, as some IBM-Zurich lab simulations have shown (**Reference 8**). The congestion-point-to-reaction-point distance varies from one to five hops, resulting in a range of delays and making it hard to choose a set of gain parameters to stabilize the system. Additionally, a fat tree's fan-in to the hot spot grows exponentially, further exacerbating the instability.

Going to 40- and 100-Gbps ports for uplinks could improve the efficiency and alleviate hashing issues. For example, instead of hashing to one of four 10-Gbps ports, you could replace the four uplinks with one 40-Gbps port. The flows using that port would enjoy better bandwidth sharing than they would if you statically allocated them to four discrete 10-Gbps ports. EDN

REFERENCES

1 "Data Center Bridging Task Group," IEEE 802. 2009. www.ieee802.org/1/pages/dcbridges.html.

2 Katz, Randy H, "Tech Titans Building Boom," *IEEE Spectrum*, February 2009, www.spectrum.ieee.org/feb09/7327.

3 Pfister, G, M Gusat, W Denzel, D Craddock, N Ni, W Rooney, T Engbersen, R Luijten, R Krishnamurthy, and J Duato, "Solving Hot Spot Contention Using InfiniBand," Center for Experimental Research in Computer Systems, Georgia Institute of Technology, 2005, www.cercs.gatech.edu/hpidc2005/presentations/GregPfister.pdf.

4 Prabhakar, Balaji, "An overview and a proposal," Jan 24, 2007, www.ieee802.org/1/files/public/docs2007/au-prabhakar-monterey-proposal-070124.ppt.

5 "TCP/IP Offload Engine," 10 Gigabit Ethernet Alliance, 2002, www.10gea.org/tcp-ip-offload-engine-toe.htm.

6 Wadekar, Manoj, "CN-SIM: Topologies and Workloads," 2007, ieee802.org/1/files/public/docs2007/au-sim-wadekar-reqd-extended-sim-list020807.pdf.

7 Matthews, Brad, Bruce Kwan, and Ashvin Lakshmikantha, "Initial QCN Serial HAI Results: Effects of PAUSE and RTT Delay," January 2008, www.ieee802.org/1/files/public/docs2008/au-sim-kwan-qcn-pause-rttdelay-0108.pdf.

8 Gusat, M, C Minkenberg, R Birke, and R Luijten, "QCN Stability Study," 2007, www.ieee802.org/1/files/public/docs2007/au-ZRL-triple-case-for-RTT-probing-r1.01.pdf.

9 Blanton, Ethan, and Mark Allman, "On Making TCP More Robust to Packet Reordering," *ACM Computer Communication Review*, January 2002, www.icir.org/mallman/papers/tcp-reorder-ccr.ps.

10 "100Gigabit and Beyond: Increasing Capacity in IP/MPLS Networks Today," Foundry Networks, www.ripe.net, May 2008, www.ripe.net/ripe/meetings/ripe-56/presentations/Vir-100Gigabit_and_Beyond.pdf.

AUTHOR'S BIOGRAPHY

Zhi-Hern Loh is a senior IC-design engineer at Fulcrum Microsystems, where he has worked for five years. He architects and designs Ethernet system-on-chip ASICs and develops the next generation of tightly integrated high-performance routers, specializing in implementing the latest data-center protocols and standards. Loh has bachelor's degrees in both electrical and computer engineering and computer science and a master's degree in engineering from Cornell University (Ithaca, NY). In his spare time, he enjoys cycling with his wife in the East Coast Park of sunny Singapore.

More Bang for Your Boost

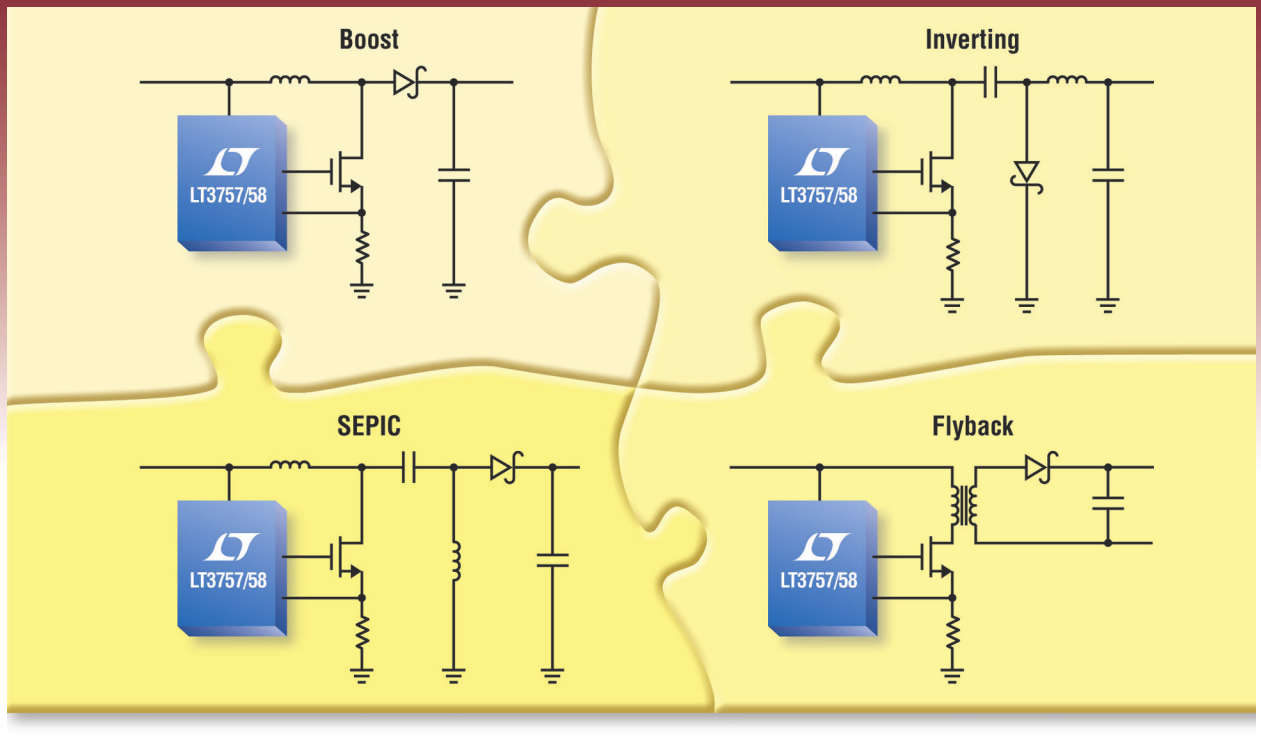

Multitopology, High Power, V_{IN} to 100V, Accurate I_{SENSE}

Linear's next generation boost controllers and regulators are so versatile, they can easily be configured for boost, flyback, SEPIC, LED driving and inverting power supply topologies. Features include input voltages up to 100V and very accurate current sense thresholds. These controllers can regulate a positive or negative output voltage via a single feedback pin and have overload protection.

Boost Controllers & Regulators

Part Number	V_{IN} Range	Topology	Switch Current Rating	Package
LT®3580	2.5V to 32V	Boost, Flyback, SEPIC, Inverting	Monolithic, 2A	3mm x 3mm DFN-8, MSOP-8E
LT3581	2.5V to 22V		Monolithic, 3.3A	3mm x 4mm DFN-14, MSOP-16E
LT3579	2.5V to 16V		Monolithic, 6A	4mm x 5mm QFN-20, TSSOP-20E
LT3755	4.5V to 40V	Boost, Flyback, SEPIC, Inverting, LED Driver	Controller*	3mm x 3mm QFN-16, MSOP-16E
LT3756	6.0V to 100V		Controller*	3mm x 3mm QFN-16, MSOP-16E
LT3757	2.9V to 40V	Boost, Flyback, SEPIC, Inverting	Controller*	3mm x 3mm DFN-10, MSOP-10E
LT3758	5.5V to 100V		Controller*	3mm x 3mm DFN-10, MSOP-10E
LT3956	6.0V to 80V	Boost, Flyback, SEPIC, Inverting, LED Driver	Monolithic, 3A	5mm x 6mm QFN
LT3957	3.0V to 40V	Boost, Flyback, SEPIC, Inverting	Monolithic, 4.5A	5mm x 6mm QFN
LT3958	5.5V to 80V		Monolithic, 3A	5mm x 6mm QFN

*Depends on MOSFET selection.

Info & Free Samples

www.linear.com/3757

1-800-4-LINEAR

Free DC/DC Controllers Brochure

www.linear.com/dcdcsolutions

LINEAR TECHNOLOGY

EDITED BY MARTIN ROWE
AND FRAN GRANVILLE

designideas

READERS SOLVE DESIGN PROBLEMS

Circuit uses two reference voltages to improve hysteresis accuracy

Marián Štofka, Slovak University of Technology, Bratislava, Slovakia

In advanced IC comparators, programmable hysteresis eliminates 0V-centered differential-input-voltage noise (**Reference 1**) and improves the comparator's response if its differential-input voltage is low or varies slowly over time. For example, the ADCMP609 comparator from Analog Devices (www.analog.com) lets you program its hysteresis from 0 to 160 mV with a single resistor that connects between the HYS (hysteresis) pin and ground. That voltage range may be too narrow for some applications, however. The circuit in **Figure 1** lets you widen the hysteresis by using two reference voltages applied to the noninverting input. The circuit uses IC_3, an Analog Devices ADR390B, to generate a high reference voltage, V_{REFH}, of 2.048V. Resistor divider R_1/R_2 produces a low reference voltage, V_{REFL}, of 0.2048V, or a difference of 1.8432V. Thus, the hysteresis equals the high reference voltage minus the low reference voltage. IC_2, an Analog Devices ADG772 dual-SPDT (single-pole/double-throw) switch, routes the voltages to the comparator's noninverting input (**Reference 2**).

DIs Inside

▶To see all of *EDN*'s Design Ideas, visit www.edn.com/designideas.

Whenever the input voltage at the comparator's inverting input exceeds the high reference voltage, the output

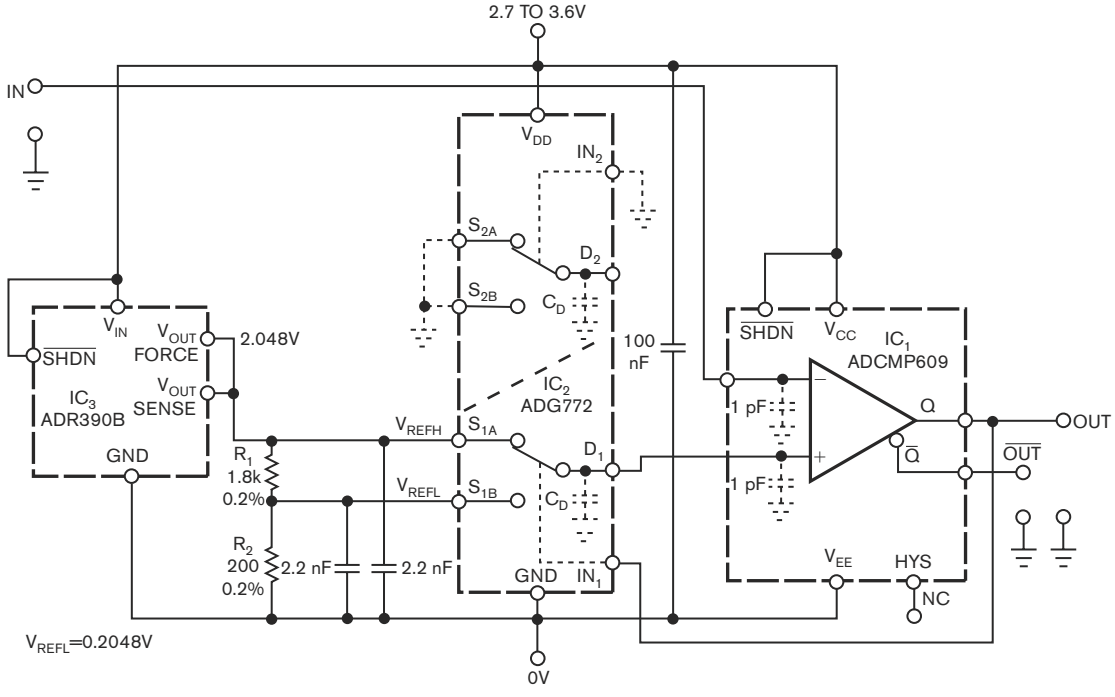

Figure 1 An analog switch changes the comparator's reference voltage, thus increasing hysteresis.

goes to a logic low. That action forces S_{1B} to connect the low reference voltage to the comparator's noninverting input. The circuit remains in that state until the input voltage drops below the low reference voltage. At that point, the switch connects the high reference voltage back to the comparator. For fast-ramping waveforms at the input, the hysteresis increases because of signal-propagation delays in both IC_1 and IC_2. The 35-nsec propagation delay in the ADCMP609 occurs at approximately 10 mV of input overdrive, and this overdrive roughly doubles in effect as an addition to the hysteresis voltage, increasing it by approximately 1%.

Because ADCMP60x comparators are rail-to-rail I/O devices, the low reference voltage could be 0V. In this case, however, the value of 0.2048V lets the comparator cooperate with other rail-to-rail I/O ICs from the same supply voltage. The outputs of these ICs can swing between 0V and the power-supply voltage with a margin of millivolts to tens of millivolts, depending on the load. An ADR390B provides the reference voltages (**Reference 3**). The 2.2-nF decoupling capacitors suppress variations of these voltages during voltage transitions at the Q output. These values are sufficient because the parasitic-charge transfer from the switch's control input, IN_1, to the respective channel's source electrode is typically 0.5 pC. The short-term variation of the decoupled reference voltages is less than 250 µV. EDN

REFERENCES

1. "ADCMP609 Rail-to-Rail, Fast, Low-Power 2.5 V to 5.5 V, Single-Supply TTL/CMOS Comparator," Analog Devices, 2007 to 2008, www.analog.com/static/imported-files/data_sheets/ADCMP609.pdf.
2. "ADG772 CMOS Low Power Dual 2:1 Mux/Demux USB 2.0 (480 Mbps)/USB 1.1 (12 Mbps)," Analog Devices, 2007 to 2008, www.analog.com/en/switchesmultiplexers/analog-switches/ADG772/products/product.html.
3. "ADR390/ADR301/ADR392/ADR395 Micropower, Low Noise Precision Voltage References with Shutdown," www.analog.com/static/imported-files/data_sheets/ADR390_391_392_395.pdf.

Resistor network sets gain for fixed-gain differential amplifier

Miles Thompson, Maxim Integrated Products, Sunnyvale, CA

When a fixed-gain amplifier doesn't offer the optimum gain for your application, you can adjust the gain to a lower value by adding an external resistor network. This attenuation circuit works like a voltage divider but with a key difference: Resistors inside the fixed-gain amplifier load down the external network (**Figure 1**). For the differential-input configuration, you can reduce this system to an equivalent half-circuit for analysis (**Figure 2**).

You determine the voltage gain for the differential configuration as follows:

$$\text{GAIN (dB)} = 20 \times \log \left(\frac{1}{1 + R_1/R_1 + 2R_1/R_2} \right) + \text{GAIN},$$

where gain is the amplifier's fixed-voltage gain in decibels. Maxim's (www.maxim-ic.com) MAX9705 differential-audio amplifier, for example, is available in fixed-gain versions of 6, 12, 15.5, and 20 dB, and its input resistance is typically 20 kΩ. For single-ended configurations, the gain is:

$$\text{GAIN (dB)} = 20 \times \log \left(R_2 \times \frac{1 - \frac{R_{1A}}{R_{1A} + (R_2 + R_{1B} \| R_{I2}) \| R_{I1}}}{R_2 + R_{1B} \| R_{I2}} \right) + \text{GAIN}.$$

These gain **equations** assume that

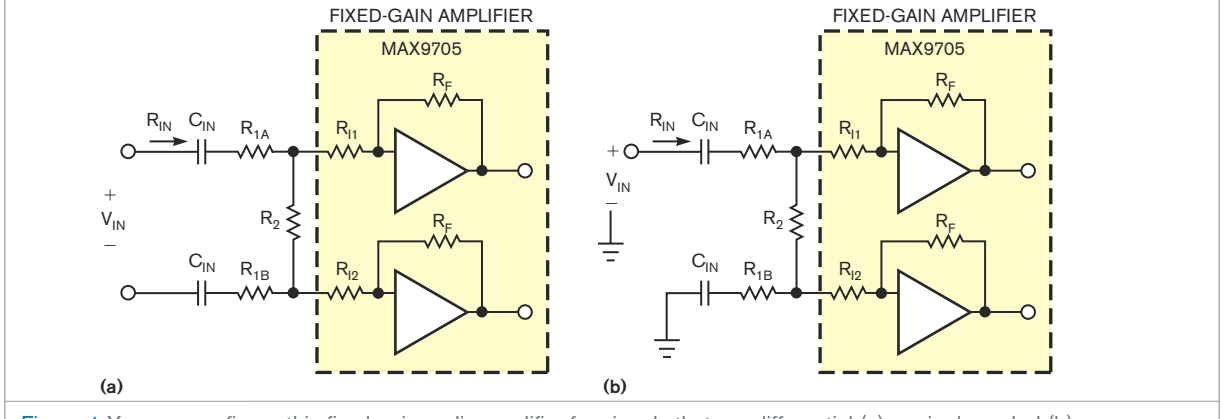

Figure 1 You can configure this fixed-gain audio amplifier for signals that are differential (a) or single-ended (b).

the frequency is much higher than the cutoff frequency of the highpass filter comprising C_{IN} and the equivalent input resistance of the circuit.

The fixed gain of the MAX9705 amplifier is typically within 5%, but the internal input resistors have an absolute tolerance of ±40%, which you must consider when calculating the resulting system gain. You must also account for the tolerance of the external resistors, using a worst-case approach for calculating the gain tolerance. Replacing the resistor values with their maximum deviations from normal—that is, the extreme values at the positive and negative tolerance limits—yields the worst-case scenario (**Table 1**).

For the tightest tolerance, choose a smaller value for R_1 than for R_I and remember to account for the source characteristics of the input voltage, which must be able to drive the equivalent load of the network. When determining the final gain for the system, remember that the output impedance of the input voltage forms a voltage divider with the input impedance of the attenuation circuit. You can determine this load, the input resistance, as follows: For the differential configuration, $R_{IN}=R_1+R_I\|(R_2/2)$. R_{IN} and C_{IN} form a highpass filter whose cutoff frequency, $F_{-3\,dB}$, is $1/(2\times\pi\times R_{IN}\times C_{IN})$. For the single-ended configuration, $R_{IN}=R_1+R_I\|(R_2+R_1\|R_I)$. Input impedance for the attenuation network is:

$$Z_{IN}=\frac{1}{2\pi\times F\times C_{IN}}+R_1+R_I\left\|\left(R_2+\left(R_I\left\|\left(R_1+\frac{1}{2\pi\times F\times C_{IN}}\right)\right)\right)\right.,$$

where Z_{IN} is the input impedance and F is the input-signal frequency.

Deriving the highpass cutoff frequency for the single-ended case is less straightforward. To find the cutoff frequency, you must know the resistor values. Once you know them, you can solve the following equation for the cutoff frequency: $Z_{IN}(f)=\sqrt{2}\times Z_{IN}(f=5000)$.

A spreadsheet is useful for determining the correct resistor values and system tolerances. Remember that resistors come in discrete values of resistance, and you must match rather than simply calculate them. A spreadsheet for determining custom gains and gain tolerances for single-ended and differential configurations of the MAX9705 is available at www.maxim-ic.com/an-4559-supplement. EDN

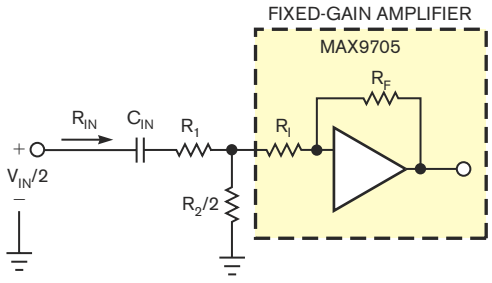

Figure 2 For the differential configuration of Figure 1a, an equivalent half-circuit simplifies analysis.

TABLE 1 WORST-CASE-TOLERANCE CALCULATION

Lowest gain	Single-ended	Differential	Highest gain	Single-ended	Differential
R_{1A}	Positive tolerance	Positive tolerance	R_{1A}	Negative tolerance	Negative tolerance
R_{1B}	Positive tolerance	Positive tolerance	R_{1B}	Negative tolerance	Negative tolerance
R_2	Negative tolerance	Negative tolerance	R_2	Positive tolerance	Positive tolerance
R_{I1}	Negative tolerance	Negative tolerance	R_{I1}	Positive tolerance	Positive tolerance
R_{I2}	Positive tolerance	Negative tolerance	R_{I2}	Negative tolerance	Positive tolerance

LED flasher checks fiber-optic strands

Edwin A Mack, Port Republic, NJ

The circuit in **Figure 1** allows you to verify fiber-optic strands, especially in cases in which a link goes through several patch panels. It uses two high-intensity LEDs that you can see at the far end. The flasher shuts itself off after about 50 minutes and has minuscule power drain when off. The circuit works on multimode fiber at distances greater than 1 km. It also works with single-mode fiber, but is more difficult to see on the receiving end. With short fiber, it is best to look at the far end at a slight angle due to the LED brightness.

IC_{1B}, a Schmitt-trigger oscillator running at approximately 5 Hz, drives IC_2, a 4020 binary divider. IC_{3A} is the control flip-flop. Pressing pushbutton switch S_1 sets the flip-flop, which starts the oscillator and enables the 4020 to start counting from its all-zero state. It also enables gates IC_{1A} and IC_{1D}, which control the PNP LED-driver transistors. Pressing pushbutton switch S_2 resets the control flip-flop; alternatively, the 4020 reaches the end of its count sequence, resetting the flip-flop. IC_{3B} divides the oscillator frequency by two to provide a 50% duty cycle to the LED drivers. The T1-¾ LEDs fit nicely into ST barrel connectors with some glue

to hold them in place. You can use a patch cable to match the fiber connectors in your network. Using a yellow LED works, considering the fiber attenuation. If you choose to use two red LEDs, however, then drive one of the LEDs at the oscillator frequency and drive the other at half the oscillator frequency. The component values and placement are not critical; the circuit in **Figure 1** resides in an old, surface-mount fiber-outlet box. EDN

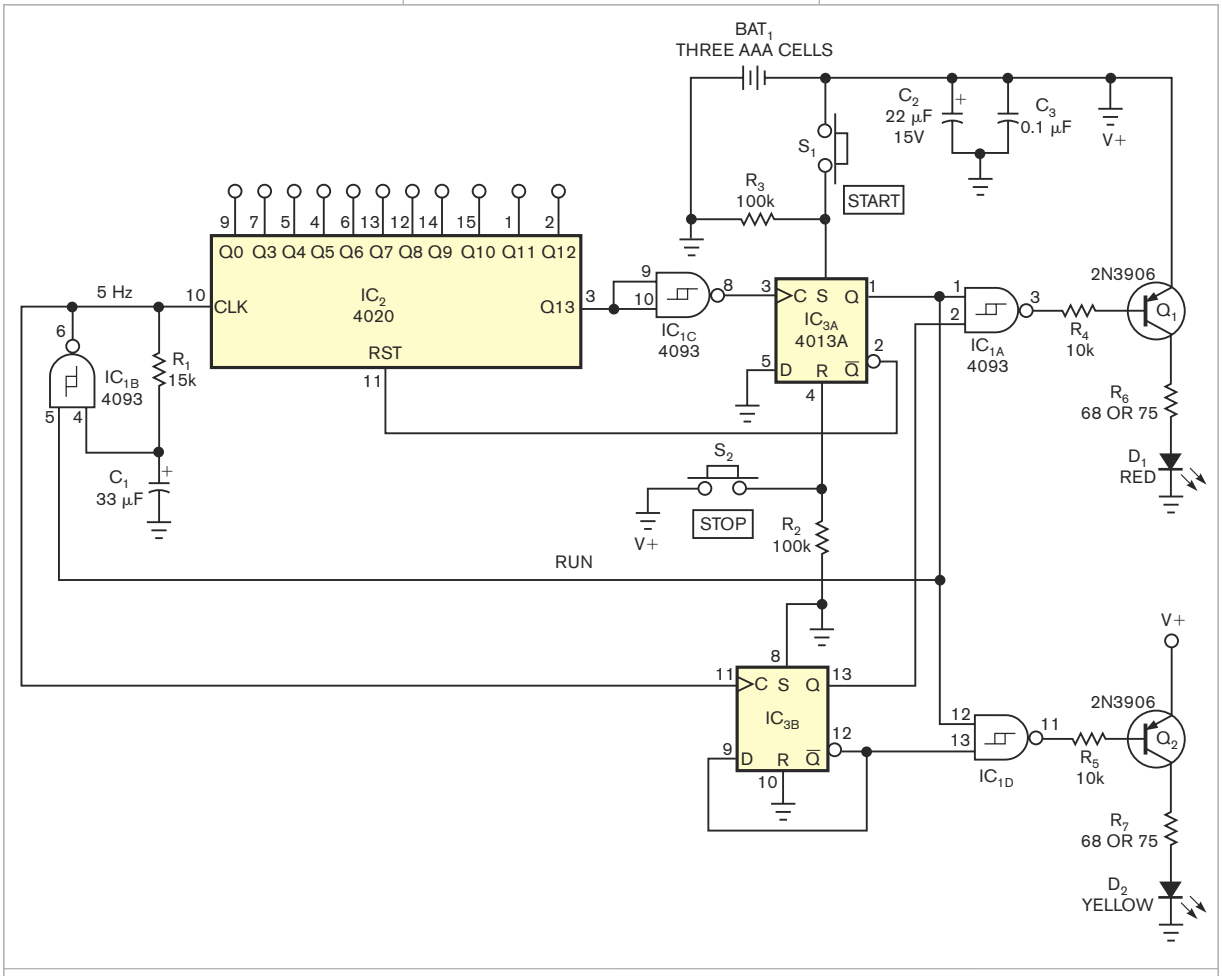

Figure 1 This circuit allows you to verify fiber-optic strands when a link goes through several patch panels.

Circuit routes audio signals between equipment

Roger Khourey, Los Angeles, CA

The audio-signal router in this Design Idea can automatically redirect cable connections between the audio equipment and an amplifier/receiver for dubbing or playback purposes without causing distortion or noise to the signal (**Figure 1**). It can be useful for older-model amplifier/receivers and any type of tape-deck-playback or -recorder equipment. The circuit uses four Maxim (www.maxim-ic.com) 4606 quad-SPST (single-pole/single-throw) CMOS analog switches to accommodate as many as eight playback and record signals for two pieces of audio equipment. Each MAX4606 contains two normally open and two normally closed logic switches. These switches have a maximum on-resistance of 5Ω and an on-resistance match of 0.5Ω between channels to minimize distortion. They can also handle rail-to-rail analog signals.

The circuit consumes little current and therefore requires no heat sinks on the power supplies. Because this circuit targets use with audio signals, it needs positive and negative supplies to pre-

(continued on pg 49)

LINEAR TECHNOLOGY

DESIGN NOTES

Dual µModule DC/DC Regulator Produces High Efficiency 4A Outputs from a 4.5V to 26.5V Input

Design Note 474

Alan Chern

Dual System-in-a-Package Regulator

Systems and PC boards that use FPGAs and ASICs are often very densely populated with components and ICs. This dense real estate (especially the supporting circuitry for FPGAs, such as DC/DC regulators) puts a burden on system designers who aim to simplify layout, improve performance and reduce component count. A new family of DC/DC µModule® regulator systems with multiple outputs is designed to dramatically reduce the number of components and their associated costs. These regulators are designed to eliminate layout errors and to offer a ready-made complete solution. Only a few external components are needed since the switching controllers, power MOSFETs, inductors, compensation and other support components are all integrated within the compact surface mount 15mm × 15mm × 2.82mm LGA package. Such easy layout saves board space and design time by implementing high density point-of-load regulators.

The LTM®4619 switching DC/DC µModule converter regulates two 4A outputs from a single wide 4.5V to 26.5V input voltage range. Each output can be set between 0.8V and 5V with a single resistor. In fact, only a few components are needed to build a complete circuit (see Figure 2).

Figure 2 shows the LTM4619 µModule regulator in an application with 3.3V and 1.2V outputs. The output voltages can be adjusted with a value change in R_{SET1} and R_{SET2}. Thus, the final design requires nothing more than a few resistors and capacitors. Flexibility is achieved by pairing outputs, allowing the regulator to form different combinations such as single input/dual independent outputs or single input/parallel single output for higher maximum current output.

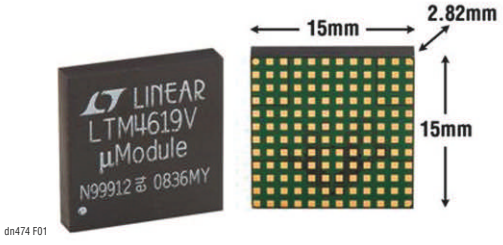

Figure 1. The LTM4619 LGA Package Is Only 15mm × 15mm × 2.82mm and Houses Dual DC/DC Switching Circuitry, Inductors, MOSFETs and Support Components

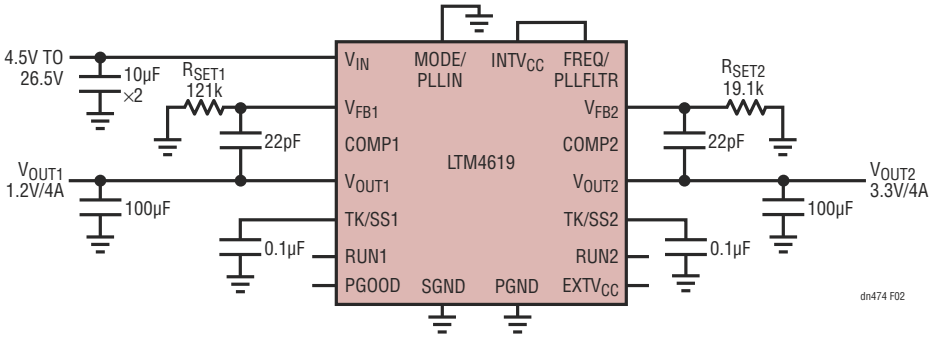

Figure 2. 4.5V to 26.5V Input to Dual 3.3V and 1.2V Outputs with 4A Maximum Output Current Each

The efficiency of the system design for Figure 2 is shown in Figure 3 and power loss is shown in Figure 4, both at various input voltages. Efficiency at light load operation can be improved with selective pulse-skipping mode or Burst Mode® operation by tying the mode pin high or leaving it floating.

Multiphase Operation for Four or More Outputs

For a 4-phase, 4-rail output voltage system, use two LTM4619s and drive their MODE_PLLIN pins with a LTC®6908-2 oscillator, such that the two µModule devices are synchronized 90° out of phase. Reference Figure 21 in the LTM4619 data sheet. Synchronization also lowers voltage ripple, reducing the need for high voltage capacitors whose bulk size consumes board space. The design delivers four different output voltage rails (5V, 3.3V, 2.5V and 1.8V) all with 4A maximum load.

Thermal Performance

Exceptional thermal performance is shown in Figure 5 where the unit is operating in parallel output mode; single $12V_{IN}$ to a single $1.5V_{OUT}$ at 8A. Both outputs tied together create a combined output current of 8A with both channels running at full load (4A each). Heat dissipation is even and minimal, yielding good thermal results. If additional cooling is needed, add a heat sink on top of the part or use a metal chassis to draw heat away.

Conclusion

The LTM4619 dual output µModule regulator makes it easy to convert a wide input voltage range (4.5V to 26.5V) to two or more 4A output voltage rails (0.8V to 5V) with high efficiency and good thermal dissipation. Simplicity and performance are achieved through dual output voltage regulation from a single package, making the LTM4619 an easy choice for system designs needing multiple voltage rails.

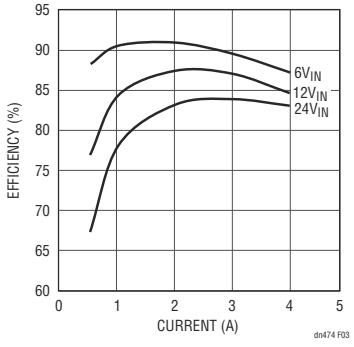

Figure 3. Efficiency of the Circuit in Figure 2 at Different Input Voltage Ranges for 3.3V and 1.2V Outputs

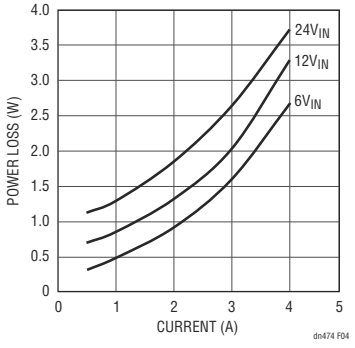

Figure 4. Power Loss of the Circuit in Figure 2 at Different Input Voltages for 3.3V and 1.2V Outputs

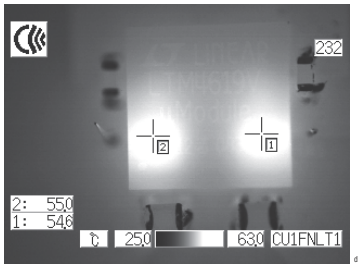

Figure 5. LTM4619: Exceptional Thermal Performance of a Paralleled Output µModule Regulator ($12V_{IN}$ to Paralleled $1.5V_{OUT}$ at 8A Load)

Data Sheet Download

www.linear.com

For applications help,
call (408) 432-1900, Ext. 2593

Linear Technology Corporation
1630 McCarthy Blvd., Milpitas, CA 95035-7417
(408) 432-1900 • FAX: (408) 434-0507 • www.linear.com

dn474 LT/TP 1110 116K • PRINTED IN THE USA

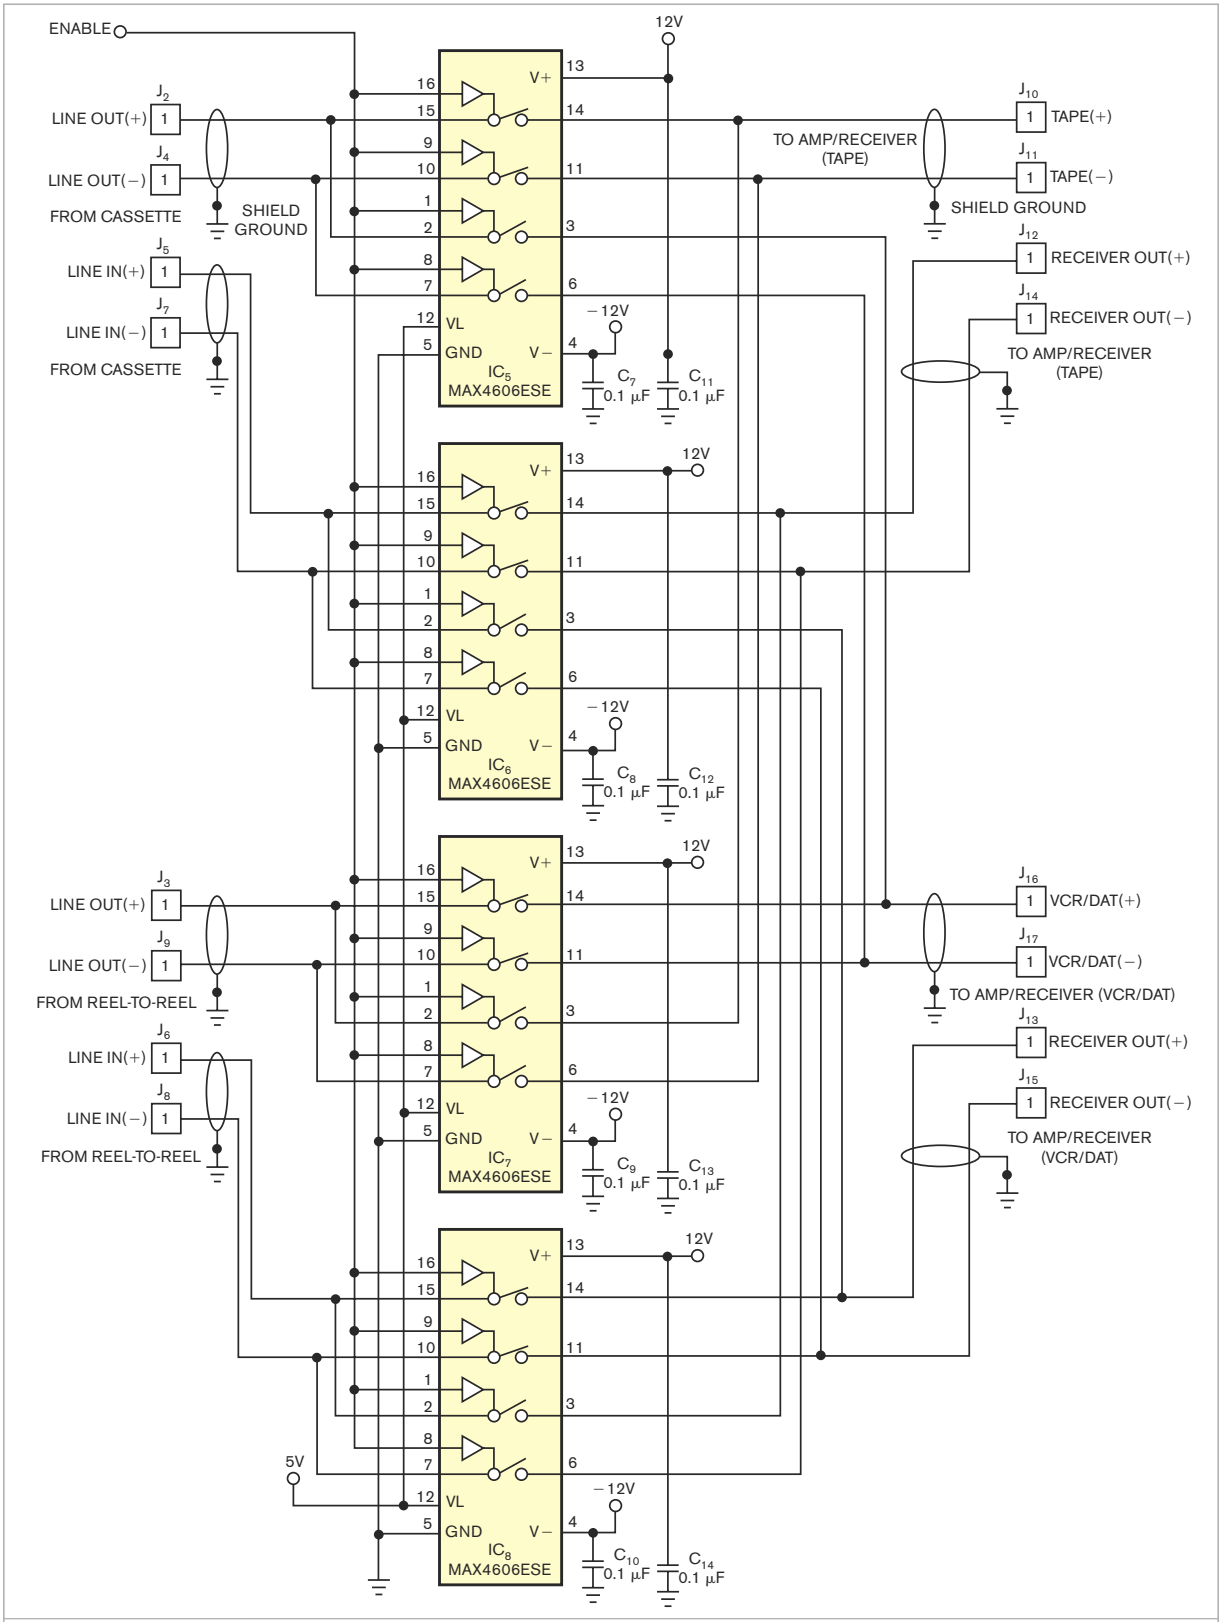

Figure 1 This audio-signal router automatically redirects cable connections between the audio equipment and an amplifier/receiver for dubbing or playback purposes without causing distortion or noise to the signal.

vent imbalance of the audio signal. It also needs a 5V supply for the logic input (VL), which lets you control the switching of the logic gates. **Figure 2** depicts the power-supply circuit, which requires a 10V-ac, 300-mA input. You can connect the enable input to a set/reset circuit or a flip-flop to toggle the internal switches from normally open to normally closed or from normally closed to normally open.

Assume that you have a cassette-tape deck and a reel-to-reel tape deck, in which the cassette deck connects to Tape 1 and the reel-to-reel deck connects to Tape 2 or another input in your amplifier/receiver. You want to record from the reel-to-reel to the cassette tape, but you can instead only record from the cassette deck to the reel-to-reel tape. When you apply a logic level, you will be able to switch your recording and playback direction to the target equipment with a single push of a button.

When wiring the circuit, you should shield and ground all wires from the equipment to the input switches and out to the equipment with reference to the audio cables' return path. This approach minimizes noisy signals due to a lack of proper grounding from one channel to another. Most of the circuit uses SMD-technology devices and can fit into a 2.5×3.5-in. enclosure and minimal installation space. EDN

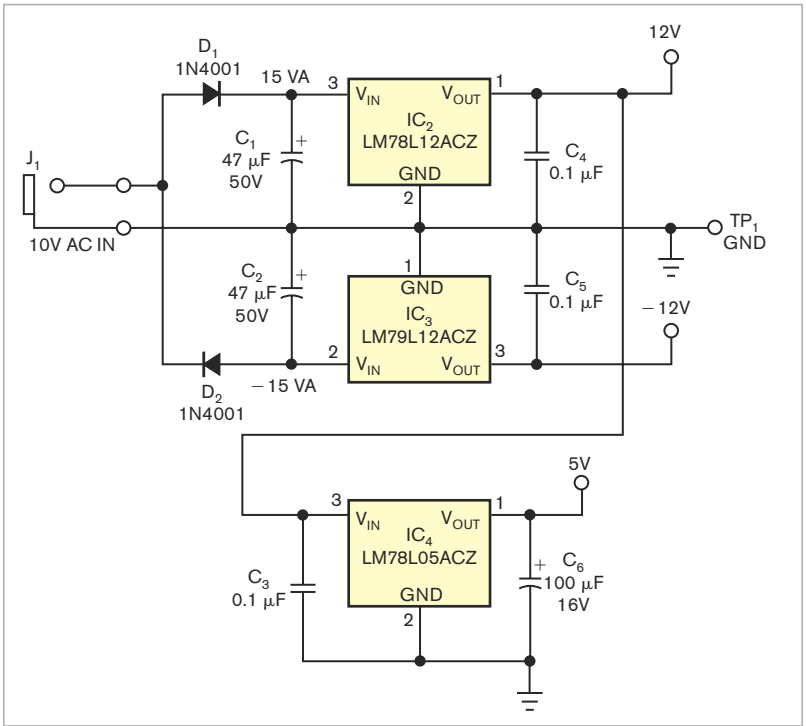

Figure 2 This ac/dc power-supply circuit powers the audio-signal router.

Microcontroller provides an alternative to DDS

Daniele Danieli, Eurocom-Pro, Venice, Italy

Audio and low-frequency circuit systems often require a signal source with a pure spectrum. DDS (direct-digital-synthesis) devices often perform the signal generation by using these specialized integrated circuits. A DDS device uses a DAC but often with no more than 16-bit resolution, limiting the SNR (signal-to-noise ratio). You can perform the same task with a microcontroller programmed as a DDS and use an external high-resolution DAC. To achieve 18 to 24 bits of resolution requires a large memory table containing the cosine function for any values of phase progression.

An alternative approach lets you use a standard microcontroller with a small memory and still implement an effective synthesizer. You can design a circuit to produce a sine wave using a scalable digital oscillator built with adder and multiplier block functions in a simple structure.

Figure 1 shows a microcontroller driving an audio DAC. To develop

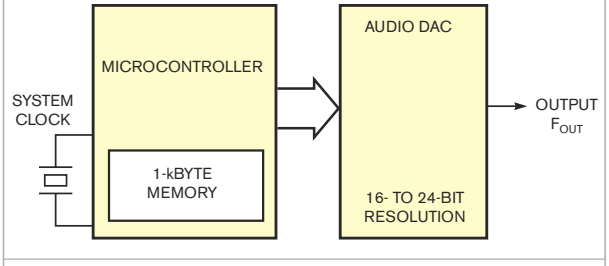

Figure 1 A microcontroller with a small memory can drive a DAC to generate audio-frequency signals.

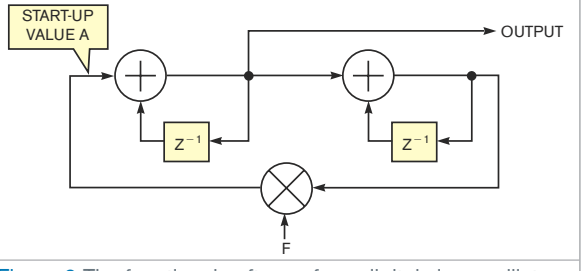

Figure 2 The functional software for a digital sine oscillator uses only two adders, one multiplier, and storage registers.

2A, DC-DC, Li+ battery charger reduces charge time and heat by 50%

Instant on: operates with no battery or dead battery

The energy-efficient MAX8903A Li+ battery charger utilizes its Smart Power Selector™ circuitry to operate with no battery or a deeply discharged battery when connected to external power. It optimizes limited USB or adapter power to utilize input power not used by the system to charge the battery. This charger thus eliminates the need for external FETs, diodes, and sense resistors to simplify design and reduce board space.

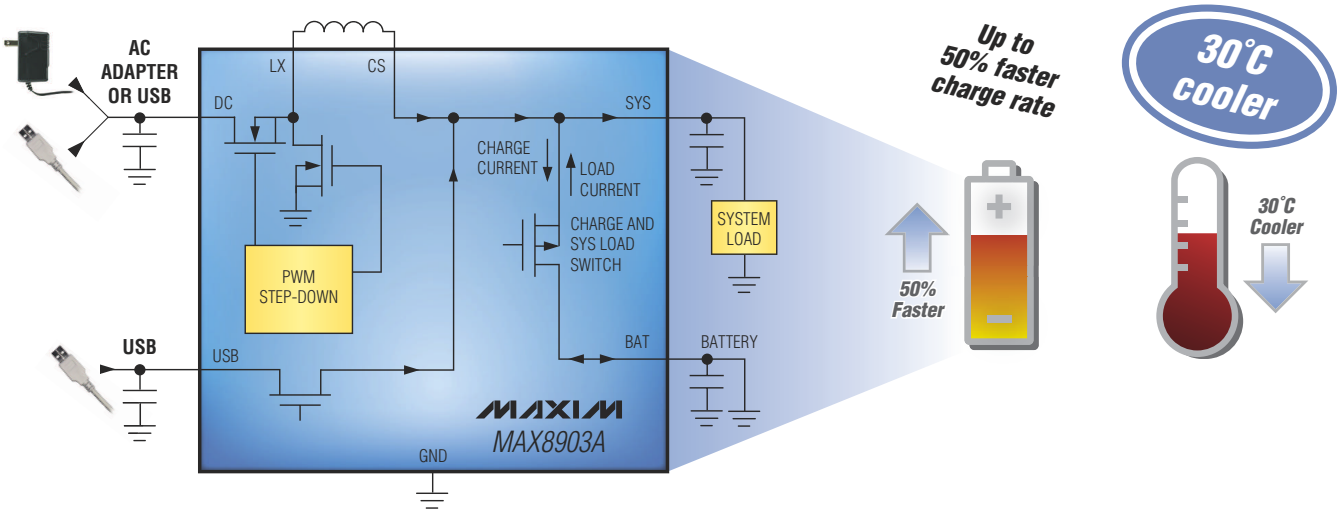

Feature	Competition	MAX8903A	MAX8903A Advantages
Input type (USB/DC)	Single	**Dual/single**	**Greater design flexibility**
Input range, absolute max	4.35V to 5.5V, 7V	**4.15V to 16V, 18V**	**No external overvoltage protection IC needed**
UVLO/OVLO	Yes/no	**Yes/yes**	**Better design robustness**
Switching frequency	2.25MHz	**4MHz**	**Smaller external components**
BAT to SYS (R_{DSON})	180mΩ	**50mΩ**	**Longer battery life and less heat**
Output current (max)	1.2A	**2A**	**Faster charge time**

Smart Power Selector is a trademark of Maxim Integrated Products, Inc.

www.maxim-ic.com/MAX8903A-info

MAXIM DIRECT
www.maxim-ic.com/shop

AVNET electronics marketing
www.em.avnet.com/maxim

MAXIM INNOVATION DELIVERED™

For free samples or technical support, visit our website.

your code to generate a sine wave, the circuit in **Figure 2** comprises two integrators with an analog feedback loop equivalent to that of an ideal resonator. Parameter F defines the frequency and ranges from 0 to −0.2, and Parameter A sets the amplitude of the output signal with a single initial pulse at start-up. The following **equation** derives the frequency of generated signals: $F_{OUT}=(\sqrt{|F|})/(2\times\pi\times T)$, where T denotes the time for computing an entire sequence to obtain output data. The firmware for implementing this system is relatively straightforward. It requires just a few additions and one multiplication. Thus, you can use a slow microcontroller. Remember, though, that the precision of every operation must be adequate to warrant a complete signal reconstruction. Processing data with 8 or 16 bits isn't sufficient. You must write your firmware to emulate a greater number of bits, which requires accurate code implementation.

If you properly develop your code, then you should generate the DAC output codes that produce a sine wave

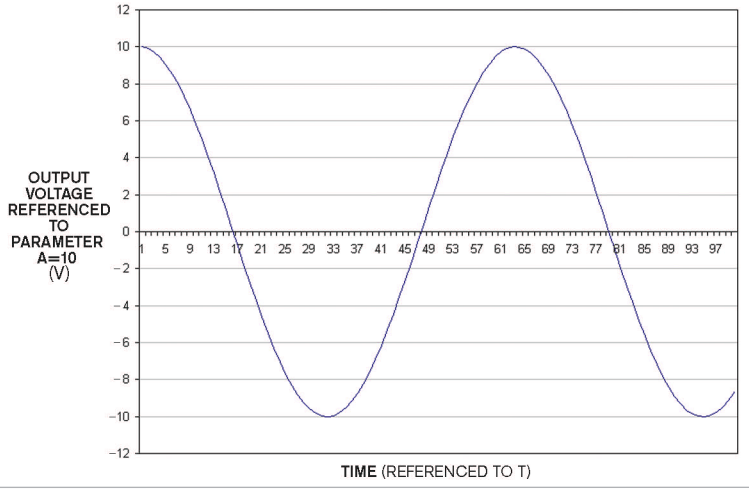

Figure 3 Use DDS to generate a sine wave.

(**Figure 3**). Remember that Parameter F is nonlinear with respect to the output frequency. If you need a directly proportional rate, you can square the value of F before applying it to the input. You'll find it useful when you need to make an easy frequency setting.

You can use just about any microcontroller to implement the oscillator, together with a high-performance DAC. You can achieve an output SNR greater than 110 dB. Many audio DACs operating in monophonic mode have 20- to 24-bit resolution at a 192-kHz sampling rate. They also offer a dynamic range of 120 dB or more. EDN

Precision clock solutions speed time to market

Low-jitter, high-frequency clock generators provide frequency flexibility

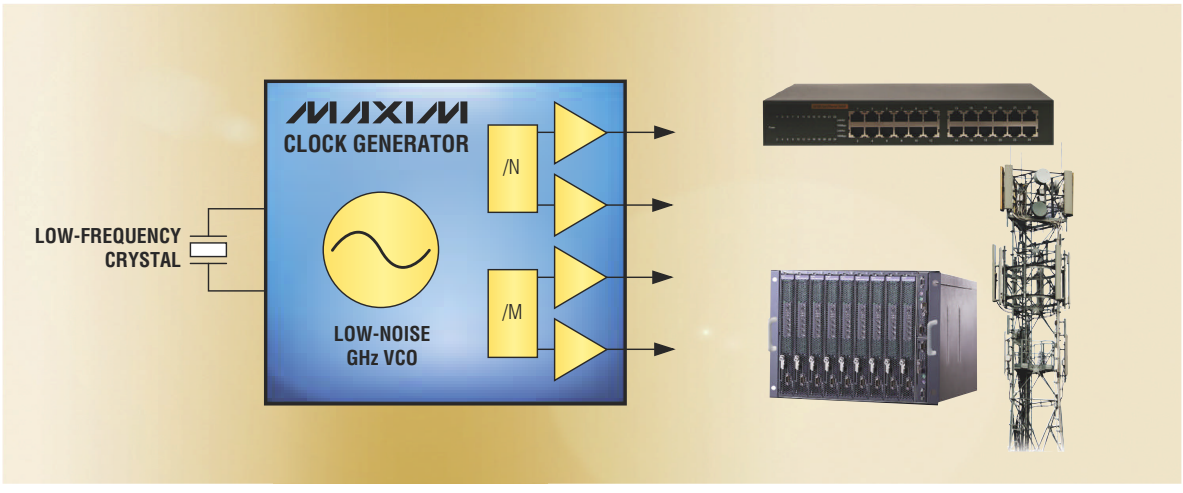

- Low-noise VCO achieves 0.36ps$_{RMS}$ (12kHz to 20MHz) jitter at high frequencies from inexpensive crystal resonator
- -57dBc PSNR minimizes jitter degradation, easing system timing design for noisy environments
- GHz VCO meets multiple frequency requirements

Part	SW/Router	Wireless BTS	SONET/SDH	Storage
MAX3624A/MAX3637/ MAX3638/MAX3639	✓	✓	✓	✓
MAX3625A/MAX3627/ MAX3629/MAX3679A	✓			
MAX3671/MAX3673/ MAX3678	✓	✓		✓
MAX3674/MAX9471/ MAX9472/MAX9491	✓	✓	✓	✓
MAX9450/MAX9451/ MAX9452/MAX3670	✓	✓	✓	✓

Find your low-jitter clock generator at: www.maxim-ic.com/clock-info

MAXIM DIRECT.
www.maxim-ic.com/shop

AVNET electronics marketing
www.em.avnet.com/maxim

MAXIM
INNOVATION DELIVERED™

For free samples or technical support, visit our website.

BY SUZANNE DEFFREE

supplychain

LINKING DESIGN AND RESOURCES

Distributor taps CAD to meet designer needs

Savvy distributors are always on the lookout for new ways to supply design engineers with needed technology, information, tools, and software to help complete the design chain. Toward that end, Premier Farnell plc (www.premierfarnell.com) recently announced its acquisition of CadSoft Computer GmbH, an established player in the CAD-software market, with its Eagle brand of electronic-design-engineering CAD software. "We had been looking at the whole portfolio of products to design engineers, certainly starting with products," says Marianne Culver (**photo**), a group director at Premier Farnell, who is serving as general manager for the CadSoft business. "While we have done a lot of work with bundles of suppliers, we thought more creatively beyond the components and the products."

The distributor had been looking at a global tools strategy. "All engineers are using a CAD tool of some kind to move their design into the real world," she says. "It seemed to us that we should be partnering with some CAD providers, initially on an authorized franchised route. … In the course of conversations, we talked to CadSoft. We were looking for those types of pieces of software that engineers have told us … are highly prized by engineers."

The company expects a positive outcome from the expansion. "The people who are buying these CAD tools are exactly the people we target: design engineers," says Randall Restle, Premier Farnell's director of global technical marketing. "A part by itself is not a solution, nor is a block diagram. A solution is an interconnection of parts. Now, with element14 [the company's online design community] and the CAD tool, we hope to have customers collaborate and suppliers collaborate to have more solutions available for design engineers. Having a common CAD tool facilitates that [goal]."

The product provides a necessary PCB (printed-circuit-board) approach for engineers and is in use with more than 40,000 designers.

GREEN UPDATE

MORE POSSIBLE CHANGES TO ROHS

The EU (European Union) has adopted three new regulations aiming to revise the NRF (New Regulatory Framework) for directives within its member states. The regulations, which the EU adopted in 2008, were all in force as of Jan 1, 2010, and require a common framework for the marketing of products. The new regulations specify the format for any new directives the EU adopted, which require CE marking (**image**). CE marking is a mandatory conformity mark on many products placed on market in the EEA (European Economic Area), certifying that the product has met EU consumer-safety, health, or environmental requirements.

As of December 2009, the ROHS (restriction-of-hazardous-substances) directive was not a CE-mark directive. However, the European Commission's proposals to amend ROHS include changes that could make ROHS a "new-approach" NRF directive requiring CE marking of products.

These regulations aim to provide more effective market surveillance, make the CE mark a community trademark, and introduce common definitions and procedures.

OUTLOOK

MULTICORE MOBILE PROCESSOR SHIPMENTS ON RISE

The year ahead will be key for the adoption of multicore processors in netbooks and smartphones, according to In-Stat (www.in-stat.com). The market-research company further projects that MIDs (mobile Internet devices) and UMPCs (ultramobile personal computers) will see an uptick in 2011 and that consumer-electronics devices will follow that trend in 2012 and 2013.

According to In-Stat, ARM (www.arm.com) architecture will see a higher growth rate as the architecture moves to multicore configurations. By 2013, almost 88% of the processors sold into the mobile market segment will be multicore processors, the company reports. "Integration of graphics/multimedia acceleration is a key trend, as well," says Jim McGregor, an In-Stat analyst. "While multicore dominates in the high-performance mobile-computing segments, integration of graphics/multimedia acceleration favors the smartphone, MID/UMPC, and mobile-entertainment-device segment," he adds. In-Stat expects the market for mobile merchant processing systems to grow at a 22.3% CAGR (compound annual growth rate) through 2013. At that time, the market should grow to about 775 million units.

Test & MEASUREMENT WORLD

TMWorld.com

- Industry News
- Blogs
- Contests
- Video
- White Papers & Design Guides
- Webcasts
- e-Newsletters

Compelling, Current, Original Content.

Looking to buy?
With BuyerZone, every big deal is no big deal at all.

BuyerZone is the trusted, unbiased site that brings buyers and sellers of business purchases together to make the purchasing process easier while saving both parties time and money. Credit Card Processing, Building Services, Home Security, Trade Show Displays, Forklifts and Construction Equipment, Business Purchases and more... Let us do the work so you don't have to.

Join the millions who've already saved time and money on many of their purchases by taking advantage of BuyerZone's FREE, no obligation services like:

- Quotes from multiple suppliers
- Buyer's guides
- Supplier comparisons
- Pricing articles
- Supplier ratings

REQUEST FREE QUOTES NOW!
Call (866) 623-5525 or visit
BuyerZoneQuotes.com

BuyerZone
Where Smart Businesses Buy and Sell

A division of
Reed Business Information.

productroundup

DISCRETE SEMICONDUCTORS

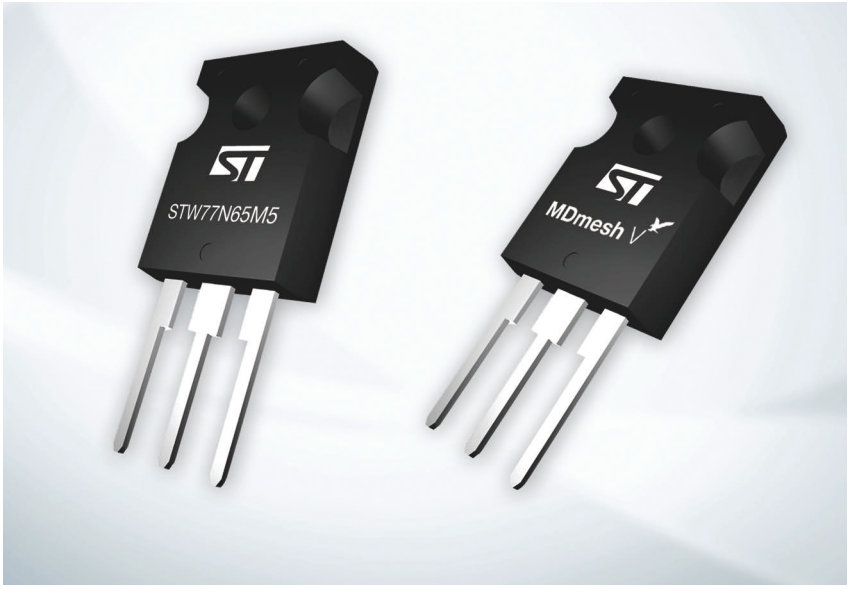

MDmesh MOSFET comes in a TO-247 housing

Part of the vendor's line of MDmesh V power MOSFETs, the 650V STW77N65M5 provides an on-resistance per die area of 38 mΩ. The vendor also plans to introduce a version with 22-mΩ on-resistance in the STY112N65M5 MAX247 package. The MDmesh V power-MOSFET packaging comes in TO-220, TO-220FP, I2Pak, DPak, and IPak options. The STW77N65M5 comes in a TO-247 package and costs $10 (1000).

STMicroelectronics, www.st.com

20V P-channel MOSFET has low on-resistance

Part of the vendor's PowerPak SC-70 product line, the SiA433EDJ MOSFET uses TrenchFET Generation II P-channel technology. The 20V MOSFET suits use in load, battery, and sharing switches in handheld devices, including cell phones, smartphones, PDAs, and MP3 players. Features include 18-mΩ on-resistance at 4.5V, 26-mΩ on-resistance at 2.5V, and 65-mΩ on-resistance at 1.8V. In an attempt to reduce field failures due to ESD, the device provides a built-in zener diode for ESD protection as high as 1800V. The 20V devices have a 12V gate-to-source voltage and a 1.8V on-resistance rating. Fitting into a 2×2-mm footprint, the SiA433EDJ PowerPak SC-70 costs 16 cents (100,000).

Vishay Intertechnology, www.vishay.com

N-channel MOSFETs suit VOIP applications

Suiting VOIP-communications equipment, including voice-over-broadband systems, PBX systems, and cable and DSL gateways, the ZXMN15A27k and the ZXMN20B28K meet requirements as primary-switch positions in transformer-based subscriber-line-interface-circuit dc/dc converters. The ZXMN15A27K and the ZXMN20B28K have 150 and 200V breakdown voltages, respectively, and withstand high-pulse avalanche energy and commutation modes in interface-circuit environments with no additional protection circuitry. Combining the circuits with specific transformers enables them to drive line voltages larger than 150V and supplies multiple subscriber lines at loop lengths exceeding 6 km. Available in a TO252-3L package, the ZXMN15A27k and the ZXMN20B28K cost 28 cents (10,000).

Diodes Inc, www.diodes.com

High-density MOSFET series comes in a new package size

The NTD49xx, NTMFS49x, NTTFS49x, and NTMS49x N-channel synchronous-buck MOSFET series target use in computing and game-console applications. Aiming at servers, VRM, and POL applications requiring synchronous dc/dc converters, the MOSFETs are drop-in replacements for the vendor's previous generations. The 34 to 79A, 30V NTTFS49x devices come in 3.3×3.3-mm M8FL packages, and prices range from 30 to 60 cents (10,000). The 32 to 79A NTD49x series comes in 6.7×10.4-mm DPak packages, and prices range from 28 to 50 cents. The 35 to 166A NTMF49x series comes in 5×6-mm SO8FL packages and costs 26 to 50 cents. The 12.5 to 17A NTMS-49x series comes in 5×6-mm SOIC-8 packages and costs 30 to 80 cents.

On Semiconductor, www.onsemi.com

productroundup

COMPUTERS AND PERIPHERALS

LCD consumes 42% less energy than previous generations

Containing 50% less mercury than comparable displays, the eco-friendly, 17-in. LCD175M LCD, part of the MultiSync 5 series, consumes 42% less energy than its predecessors. Features include a 1000-to-1 contrast ratio, 5-msec response time, 250-cd/m^2 brightness, and 1280×1024-pixel resolution. The LCD175M LCD costs $179 and includes a three-year parts and labor warranty.

NEC Display Solutions of America, www.necdisplay.com

Isolation-memory buffer improves server and workstation outputs

The iMB01-GS02 isolation-memory buffer has four times greater memory capacity and two times greater bandwidth than typical registered DIMMs. The buffer drives the creation of load-reduced DIMMs, enabling workstations and servers to increase the output of multicore processors. Available in ROHS-compliant, flip-chip BGA packages, the iMB01-GS02 isolation-memory buffer cost $26 for initial engineering samples.

Inphi Corp, www.inphi.com

40-Gbyte boot drive boosts desktop systems

Using the industry-standard PCMark Vantage Advanced HDD Suite, the SSDNow V series 40-Gbyte boot drive received a score of 13,883 compared with a 7200-rpm hard-disk drive, which attained a score of 3708. An optional bundle includes cloning software, 2.5- to 3.5-in. brackets, and SATA data and power-cable extenders. The boot drive has 2.17g vibration operation from 7 to 800 Hz and 20g nonoperating vibration from 20 to 2000 Hz. Typical power specifications include 0.15W active-mode and 0.06W sleep-mode power consumption. Additional features include −40 to +85°C storage temperature, 0 to 70°C operating temperature, and a 1 million-hour MTBF. The drive costs $115, including a three-year warranty.

Kingston Digital LLC, www.kingston.com

Micro SATA storage devices work with iPods and subnotebooks

The Pocket Micro SATA UDD and the Internal Micro SATA UDD storage devices enable the use of any 1.8-in. Micro SATA hard-disk or solid-state drive as removable-drive cartridges. The Micro SATA CF adapter allows the use of Compact Flash media as a replacement for 1.8-in. Micro SATA hard drives in iPods or subnotebooks. The Pocket Micro SATA UDD and the Internal Micro SATA UDD cost $45 and $39.95, respectively. The Micro SATA CF adapter costs $23.99.

Addonics Technologies, www.addonics.com

Portable network-storage enclosure supports multiple protocols

The Mini NAS portable storage enclosure supports SMB and open-source Samba network protocols for as many as 64 clients. The device features a 10/100-Mbps Ethernet connection and allows the installation of any 2.5-in. SATA hard drive or solid-state drive, enabling instant sharing over the LAN. The enclosure provides FTP access for as many as eight simultaneous users with an Internet connection. The device also allows configuration as a bit-torrent-downloading appliance or an iTunes media server and suits use as a UPnP AV server. The Mini NAS enclosure costs $69.

Addonics Technologies, www.addonics.com

External hard drive provides an eSATA interface

The 3-Gbps Story Station Plus external hard drive features an eSATA interface, allowing fast data transfers. The 3.5-in. drives feature 1- and 2-Tbyte densities and 0.1W standby-power consumption. Available with a three-year limited warranty, the 2-Tbyte Story Station Plus external hard drives cost $299.

Samsung Electronics Co, www.samsung.com

Graphics cards support multiple displays at 2560×1600-pixel resolution

Supporting as many as eight DisplayPort or DVI single-link outputs, the M9188 PCIe ×16 octal-graphics card has 2 Gbytes of memory and 2560×1600-pixel resolution per output. The 2560×1600-pixel M91287 LP PCIe ×16 DualHead DisplayPort dual-monitor add-in graphics card suits business, industrial, and government applications. The M9128 and M9188 feature 1 and 2 Gbytes of memory, respectively. The single-slot graphics cards enable combination with other M-series products and provide a unified driver package. The devices allow stretched or independent desktop modes and support 32- and 64-bit versions of Windows 7, Vista, XP, and Server 2003/2008, as well as Linux. The M9128 LP PCIe ×16 and the M9188 PCIe ×16 graphics cards cost $259 and $1995, respectively.

Matrox Graphics, www.matrox.com

INTEGRATED CIRCUITS

12-bit magnetic rotary encoder aims at automotive applications

Based on the vendor's stacked-die technology, the 12-bit AS5245 fail-safe magnetic-rotary-encoder IC targets use in pedals, transmission, steering wheels, and other automotive applications. Providing improved phase-matching capabilities, the device requires no additional temperature compensation, calibration, or shielding against external magnetic stray fields. The digital-angle data of the encoder comes in the form of a serial interface or as a PWM signal. Designers can enable an incremental quadrature-output mode with an improved interpolation algorithm and permanently program it with 10- or 12-bit resolution. Standard features include user-specified zero programming, Chip-ID, diagnostic functions for the correct positioning of the magnet, operation over a 3.3 or 5V supply, and an ambient-temperature range as high as 150°C. Available in a QFN-32 package, the AS5245 costs $12.22 (1000).

austriamicrosystems, www.austriamicrosystems.com

Differential audio preamplifier combines with digital preamplifier controller

Combining a low-noise differential audio preamplifier and a high-performance digital preamplifier controller, the That1570/5171 paired chip set provides as much as 130-dB dynamic range when operating from ±17V supplies. The digitally programmable microphone preamplifier enables a programmable gain over 56 1-dB steps from 13.6 to 68.6 dB. An additional 5.6-dB-gain setting allows acceptance of line inputs as high as 22 dBu. Additional specifications include 1.5-nV/$\sqrt{\text{Hz}}$ noise at 68.6-dB gain and 20-nV/$\sqrt{\text{Hz}}$ at 5.6-dB gain, 0.0003% THD+N at 21.6-dB gain, and 1.8-MHz bandwidth at 40.6-dB gain. The device integrates a servo for minimizing dc offset, a zero-crossing detector, four general-purpose digital outputs, and techniques for minimizing gain-switching noise. The 1570 comes in a 4×4-mm QFN-16 package, and the 5171 comes in a 7×7-mm QFN-32 housing. The chip set costs $8.91 (1000). A demo board costs $250.

That Corp, www.thatcorp.com

EDN ADVERTISER INDEX

EDN provides this index as an additional service. The publisher assumes no liability for errors or omissions.

The nose knows

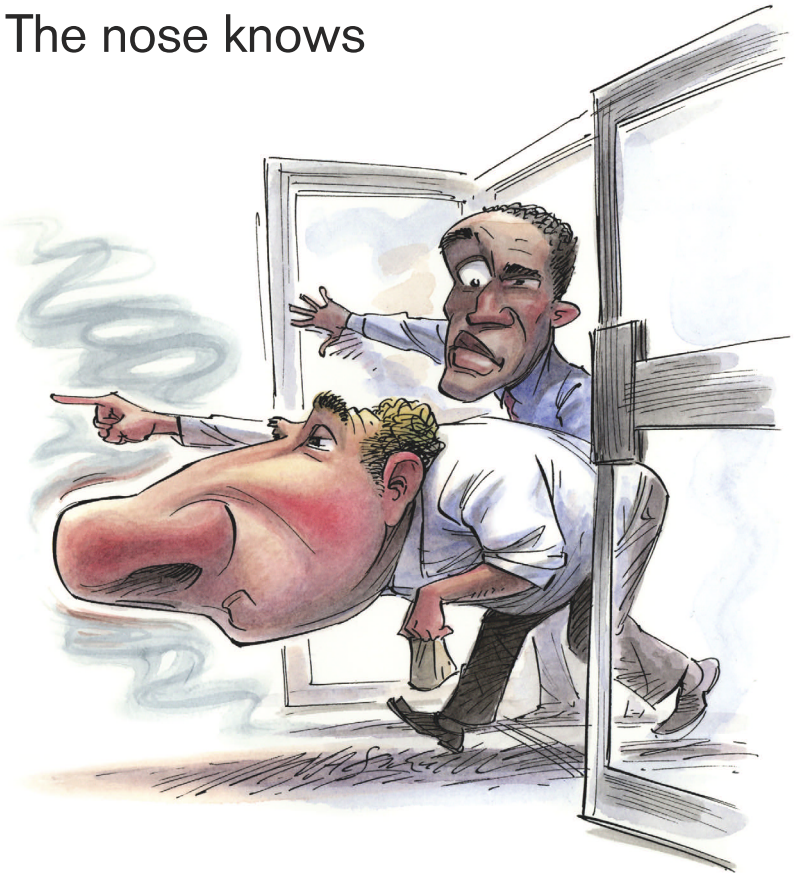

DANIEL VASCONCELLOS

Engineers and technicians are often quick to obtain and use the best tools available. The right tool helps you to efficiently design or repair your circuit. Occasionally, a makeshift tool can also do surprisingly well. All designers may have one or two of these favorite tools in their little bag of tricks. Yet a useful tool that often goes unnoticed is as plain as—and, in fact, *is*—the nose on your face.

We all use it every now and then. How many times have we sniffed around a PCB (printed-circuit board), detecting the smell of heated or overheated components? I remember when I made this not-so-amazing discovery for myself. I was in Santa Barbara, CA, working for Sloan Technology, a manufacturer of vacuum-deposition systems for optics and chip manufacturing.

A team doing research into high-efficiency solar cells at the University of Colorado—Boulder called me out for an urgent repair. The team was using a Sloan vacuum system with full monitoring-and-control equipment; an electron-beam evaporation hearth; and an old, 10-kV Sloan Omni electron-beam power supply. The call came from a rather anxious researcher, stating that the high-voltage power supply had failed and needed repairing. I threw a few key power transistors, components, and tools into my little bag of tricks—in this case, a small, paper lunch bag—and set off for Boulder.

Upon my arrival, a researcher led me to the entrance of a laboratory with the proportions of a basketball stadium. Through the closed glass doors, I could see the Sloan installation off in the distance, some 40 or 50m away. Noticing my brown bag, the researcher asked, "Are those the only tools you have?"

"Yes," I replied. "I have everything I need in here."

As soon as the door opened, I turned and said to the researcher, "I have just now found the problem; there is nothing wrong with your power supply."

Eyeing me suspiciously, he asked, "How can you possibly reckon that, since the equipment is way off in the distance and you are standing next to me?"

"The power supply is working perfectly fine," I stated. "Take a whiff of the air. What do you smell?"

"Ozone," he replied.

I reminded him that ozone forms as a result of intense electrical arcing. In this case, a short circuit of some kind must be inside the vacuum chamber, in which the high-voltage cable connects through the high-voltage feedthrough. Electron-beam power supplies can have built-in output protection, periodically folding back the high voltage for protection. If there were high voltage, then the supply should be working OK. Any subsequent safety arcing would generate ozone. The laboratory just reeked of it. The electron-beam power supply must have been cycling on and off for some time. Although more conventional methods would have located the problem, it's unlikely that they could have pinpointed it as quickly.

We powered down the equipment and found the problem to be carbon tracking in the high-voltage electrical feedthrough that went into the vacuum chamber. We subsequently made the necessary replacements, and everyone went away happy. I then realized that I really did arrive with everything I needed. In addition to my components and tools, I had one more thing in my little bag of tricks: my nose! And I didn't even have to pack it. EDN

Jim Zannis is an analog-design engineer in Baulne-en-Brie, France.

Introducing EDN's New Microsite on LEDs

Your complete source for up-to-the-minute information on LEDs and their applications.

- Latest News from the Editors at EDN, Design News and other Engineering Magazines
- Blogs, including Margery Conner's Popular "Power-Source" Blog
- Resource Center, Including White Papers, eBooks, Tutorials, Webcasts and More
- Videos, Including Design Tips, Application Demos, and More
- Updates on Future In-Person Workshops and Seminars, like this one!
- Fun Challenges to Test Your Knowledge of LEDs

Check It Out Today at www.edn.com/leds

Cool it.

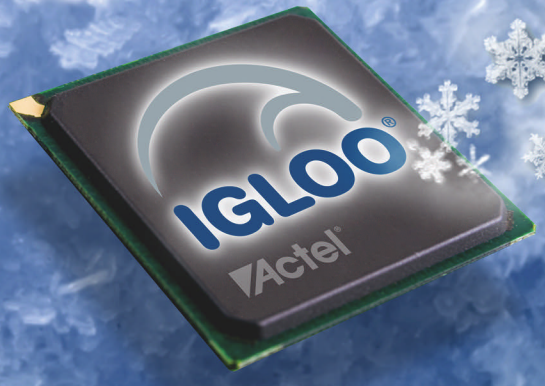

At 2 µW, the Industry's Lowest Power FPGAs.

You can't get any cooler than Actel's IGLOO® FPGAs. And with the industry's widest range of packages—as small as 3x3 mm—you also can't get any smaller. IGLOO FPGAs are perfect for all of your portable, battery-powered applications, and anywhere you need to stretch your power budget to the maximum.

Designing products for the consumer, industrial, medical, automotive or communications markets? We've got you covered with a full range of solutions. As long as power matters, Actel FPGAs should be at the top of your list. Find out more at actel.com.

Lowest Power				Low Power		
IGLOO				ProASIC3		
IGLOO/e	IGLOO nano	IGLOO PLUS		ProASIC3/E	ProASIC3 nano	ProASIC3L
Power	Power and Size	Power and I/Os		Performance	Cost	Power and Performance
●	●	●	Low Power	○	○	●
●	●	●	Small Footprint	○	●	○
●	○	○	Logic Density	●	○	●
○	○	●	I/O to Logic Ratio	○	○	○
○	○	○	Speed	●	●	●
○	○	○	Low Cost	●	●	●

● Best ○ Good

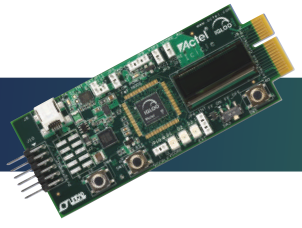

Get your battery-powered IGLOO Icicle Evaluation Kit and start designing today: actel.com/IcicleKit

Actel®
POWER MATTERS