Canon Communications LLC

EDN®

VOICE OF THE ENGINEER

APRIL 8

Issue 7/2010

www.edn.com

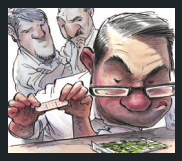

COMING SOON: 3-D TV

Page 24

LTE IN FOCUS

Page 17

COMPLEX MODULATION COMES TO OPTICAL FIBER

Page 32

MEASURE PROPAGATION DELAYS USING TIME-DOMAIN REFLECTOMETRY

Page 36

CONTINUOUS-TIME EQUALIZERS IMPROVE HIGH-SPEED SERIAL LINKS

Page 40

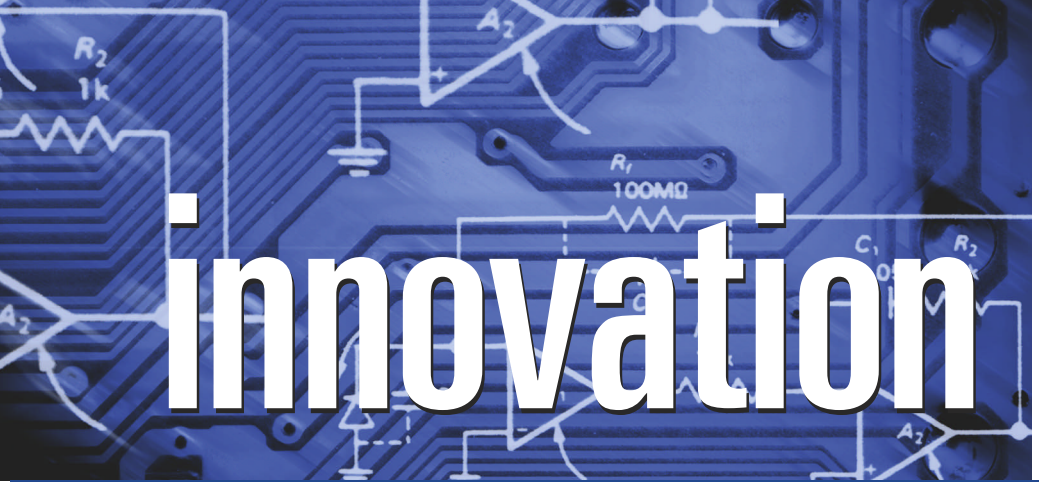

APEX PRECISION POWER

Product Innovation from Cirrus Logic

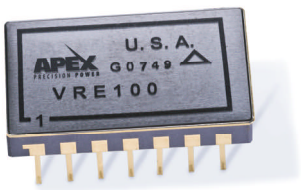

VRE100 SERIES
DIP14 PACKAGE

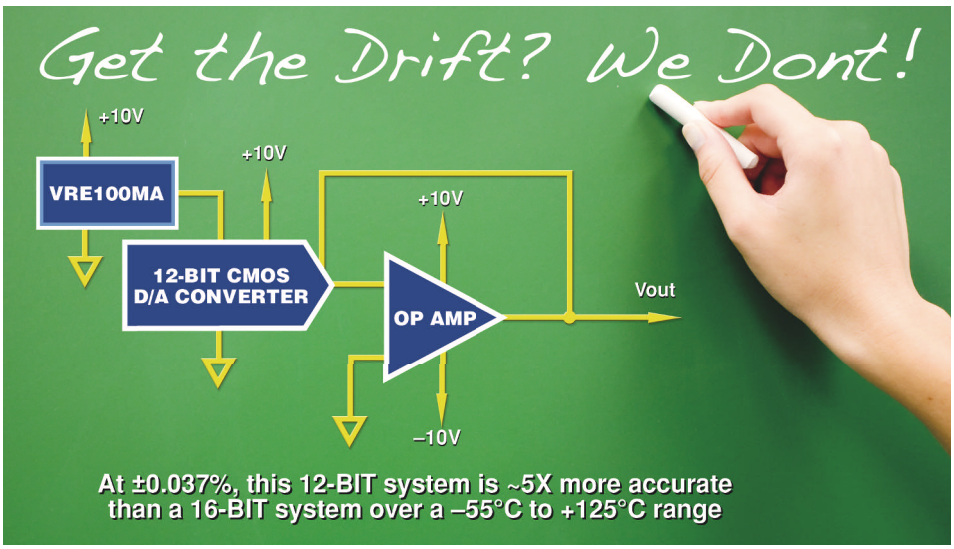

Cirrus Logic Voltage References Deliver Low Temperature Drift for Better Accuracy

LOW INITIAL OFFSET PEFORMANCE ALSO MINIMIZES IMPACT OF OPERATING CONDITIONS

When your system demands voltage supply accuracy, rely on Cirrus Logic to deliver. The VRE series of precision voltage references offer best in class performance for low temperature drift by minimizing the impact of operating conditions on voltage delivery over time. Performance "at temperature" is enhanced further with low initial offsets, improved temperature coefficients and reduced thermal hysteresis. With the quality and reliability of more than 25 models of off-the-shelf solutions, Cirrus Logic can eliminate design cycle headaches and speed time to market.

Model Specifications

Series	Output (V)	TempCo (ppm/°C)	Feature
VRE302	2.5	0.6, 1.0, 2.0	Low Cost
VRE3025	2.5	0.6, 1.0, 2.0	+10V Supply
VRE4125	2.5	1.0, 3.0	Low Cost, 5V Supply
VRE117	+3	0.3, 0.6	Low Output V
VRE205	5	0.4, 0.8	Small Pkg
VRE305	5	0.6, 1.0, 2.0	Low Cost
VRE306	6	0.6, 1.0, 2.0	Low Cost
VRE100	+10, –10, ±10	0.3, 0.5	High Rel
VRE210	10	0.3, 0.5	Small Pkg
VRE310	10	0.6, 1.0, 2.0	Low Cost

TO LEARN MORE ABOUT CIRRUS LOGIC VOLTAGE REFERENCES,OR TO RECEIVE A FREE COPY OF THE LATEST VRE PRODUCT GUIDE, VISIT www.cirrus.com/ednvre

NORTH AMERICA
+1 800-625-4084

ASIA PACIFIC
+852 2376-0801

JAPAN
+81 (3) 5226-7757

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

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

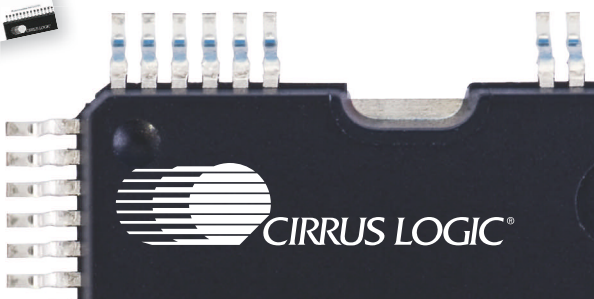

WORLDWIDE DISTRIBUTION... MADE EASY

SAME-DAY SHIPPING on orders entered by 8:00 pm CENTRAL
NEXT-DAY DELIVERY

Digi-Key® CORPORATION

450,000+ PRODUCTS IN STOCK.
440+ SUPPLIER PARTNERS.
45,000+ NEW PRODUCTS ADDED IN THE LAST 90 DAYS.

The industry's broadest product selection available for immediate delivery

www.digikey.com
1.800.344.4539

NECA Supplier Authorized Distributors

EDN 4.8.10 contents

Coming soon: 3-D TV

LTE in focus

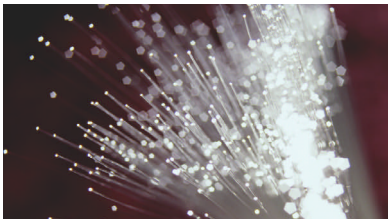

Complex modulation comes to optical fiber

Measure propagation delays using time-domain reflectometry

Continuous-time equalizers improve high-speed serial links

pulse

DESIGNIDEAS

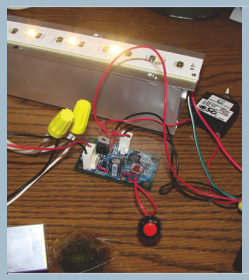

▶ Send your Design Ideas to edndesignideas@reedbusiness.com.

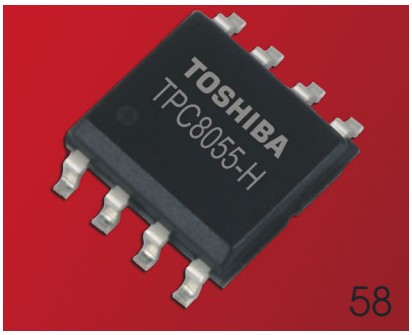

58

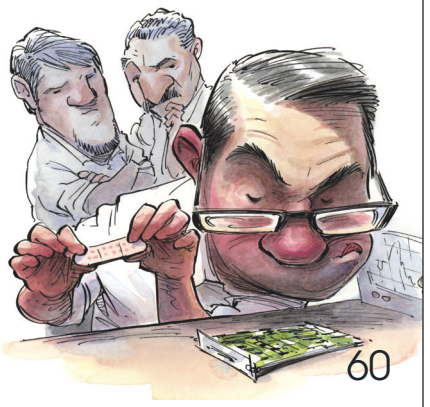

60

DEPARTMENTS & COLUMNS

EDN online contents www.edn.com

ONLINE ONLY

Check out these Web-exclusive articles:

IC reverse engineering—a design-team perspective

The IC Insider explains the why and how of circuit extraction of semiconductor chips.

→www.edn.com/article/CA6722668

Circuit protection for outdoor LED lighting

Outdoor LED lighting can be an efficient, "greener" source of light in addition to promising a long, maintenance-free life. However, outdoor LED-lighting installations can be less reliable than conventional lighting unless designers add the proper circuit protection to guard against the most severe overvoltage conditions.

→www.edn.com/article/CA6721136

EDN® (ISSN#0012-7515) is published semimonthly, 24 times per year, by Canon Communications LLC, 11444 W. Olympic Blvd., Los Angeles, CA 90064-1549; 310/445-4200; FAX 310/445-4299. Periodicals postage paid at Los Angeles, California, and at additional mailing offices. SUBSCRIPTIONS—Free to qualified subscribers as defined on the subscription card. 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. For telephone inquiries regarding subscriptions, call 763/746-2792. Email: EDN@kmpsgroup.com. CHANGE OF ADDRESS—Notices should be sent promptly to PO Box 47461, Plymouth, MN 55447. Please provide old mailing labels as well as new address. Allow two months for change. NOTICE—Every precaution is taken to ensure accuracy of content; however, the publishers cannot accept responsibility for the correctness of the information supplied or advertised or for any opinion expressed herein. POSTMASTER—Send address changes to EDN, PO Box 47461, Plymouth, MN 55447. Canada Post: Publications Mail Agreement 40685520. Return undeliverable Canadian addresses to RCS International, Box 697 STN A, Windsor, Ontario N9A 6N4. Printed in USA.

BY RON WILSON, EXECUTIVE EDITOR

Lifetime and reliability in LEDs and why SOC designers should care

One of the wonderful things about this job is the chance to sit in on presentations unrelated to my beat. The subject of this commentary, a paper at *EDN*'s "Designing with LEDs" seminar, is a case in point. In 45 minutes, I heard some hard-earned lessons from a different world but with implications for SOCs (systems on chips). The presenter, Geof Potter, is a power technologist at Texas Instruments. After a quick review of what goes into an LED-lighting assembly—basically, a power supply and a bunch of LEDs—Potter decomposed the luminaire assembly into its components and discussed the impact of each on lifetime and reliability.

The two terms are different in Potter's world. "Lifetime" refers to the length of service before the luminaire falls below some determined fraction of rated light output. "Reliability" refers to the probability that the luminaire will require repair or replacement during its rated life. It's possible, for example, for a luminaire to have a lifetime of 10 years but still require repairs every three weeks.

It seems obvious that the shortest-lived, least-reliable component in the luminaire would be the LED. It turns out, though, that if you run the little guys conservatively and cool them aggressively, they will be fine. The next most obvious culprit is the fluid-filled aluminum electrolytic capacitors in the power supply. Everybody has had bad experiences with e-caps: aging, electrical degradation, and a distinctively bad aroma as they near end of life. According to Potter, however, when you use them conservatively, extended-life e-caps have a working life that can exceed that of the LED assemblies.

Similarly, optocouplers have been problems in the past. But high-reliability couplers, properly used, will have more than adequate reliability. Cutting costs on these components or allowing an assembly house to make unauthorized substitutions will lead to trouble, however.

Once you've eliminated LEDs and e-caps as the worst problems, only solder joints remain. According to Potter, the 50-year-old problem has reappeared in LED lighting, and the reasons are unfortunately relevant to many SOC designs. The key issue is quality of soldering, Potter said. All solder joints age, and poor mechanical connection, contamination, or insufficient heating greatly accelerates this aging. The lighting market tends to be highly cost-sensitive, and it heavily outsources the supply chain to low-labor-cost areas and those more familiar with mechanical assembly than electronics fabrication.

Furthermore, luminaires can live in high-stress environments. Think of an automobile headlamp's "eyebrow," that cute little string of white LEDs that makes this year's luxury sedans distinctive in the dark. The front of the car may be moving at 100 km/hour into a −10°C evening when the lamps come on, quickly self-heating to 70°C, all while enduring road shocks and engine vibrations.

Maybe Bosch is comfortable assembling a fixture for this environment, but is your low bidder in Southeast Asia? Potter's suggestions are to ensure manufacturing quality and minimize thermal and mechanical stress. These approaches may be largely beyond the control of the design team, however, so he stressed another point: Use the fewest solder joints that you can.

The SOC world also faces growing cost pressures, uncertainties about its customers' supply chains, and the need for designs to operate in hostile environments. Potter's data suggest that, by doing whatever is possible in the architecture and implementation of the system SOC to minimize the number of solder joints in the end system, you can have a dramatic impact on system reliability and on total cost of ownership. Ideas such as using the highest feasible level of functional integration, absorbing or eliminating the need for external passive components, and minimizing both the number of pins on the SOC and the number of lines that must pass through connectors are not just good practice. They are becoming differential advantages that can sell—or doom—an entire SOC family. EDN

Contact me at ronald.wilson@reedbusiness.com.

EDN

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

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
Steve Mahoney,
Production/Editorial Coordinator
Melissa Annand,
Web Operations Specialist
Adam Odoardi, Prepress Manager
1-781-734-8325; fax: 1-303-265-3042

CONSULTING EDITOR
Jim Williams, Staff Scientist,
Linear Technology

CONTRIBUTING TECHNICAL EDITOR
Dan Strassberg,
strassbergedn@att.net

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
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

EXECUTIVE OFFICERS
Charles G. McCurdy,
Chief Executive Officer
Daniel Koskovich,
Chief Financial Officer
Ron Wall,
Senior Vice President, Publications
Kevin O'Keefe,
Senior Vice President, Events Division
Roger Burg,
Vice President, Operations
Jason Brown,
Vice President, E-Media

EDN. 225 Wyman St, Waltham, MA 02451. **www.edn.com**. Phone 1-781-734-8000. **Subscription inquires:** 1-763-746-2792; EDN@kmpsgroup.com. **Address changes:** Send notice promptly to PO Box 47461, Plymouth, MN 55447. Please provide an old mailing label as well as your new address. Allow two months for the change. Canon Communications LLC, 11444 W. Olympic Blvd., Los Angeles, CA 90064-1549; 1-310-445-4200; fax: 1-310-445-4299.

CANON COMMUNICATIONS LLC

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

Expandable 3.6-, 6-, and 13.5-GHz signal analyzers feature optional built-in generator

Anritsu has introduced the MS2830A signal analyzers, which deliver best-in-class speed and accuracy, according to the manufacturer, over a frequency range of 9 kHz to 3.6, 6, or 13.5 GHz in a cost-effective package. Focusing on applications in system integration and production, the new units enable manufacturers of third- and fourth-generation wireless devices and systems to reduce manufacturing costs and increase both product yields and confidence in product performance.

The base model conducts frequency switching and sweeping and transfers batched measurement results in approximately 12 msec. If you install the VSA (vector-signal-analysis) option, the instruments support noise averaging in approximately 0.1 second, several times as fast as do conventional swept-frequency spectrum analyzers. In a 31.25-MHz analysis bandwidth, the VSA option's FFT-based batch-capture capability greatly increases the speed of inband measurements, such as channel power, adjacent-channel power, and one-box tester.

The high speed does not compromise accuracy: The average noise level without preamplification is −153 dBm at 1 GHz; tertiary phase-intermodulation distortion (third-order intercept) is 15 dBm, and typical total level accuracy is ±0.3 dB. The instruments also perform precise distortion and spurious-frequency measurements. The MS2830A has measurement capabilities that previously required both a spectrum analyzer and a signal generator. An optional built-in signal generator creates a one-box tester that performs transmitter and receiver tests. With everything in one chassis, the resulting configuration costs as much as 30% less than a discrete signal analyzer and signal generator. In addition, the MS2830A consumes no more than 110W, which is 45% less than the power consumption of comparable signal analyzers.

You can customize the units to meet your current test requirements and perform simple and cost-effective upgrades when those needs change. The manufacturer offers such options as a VSA and a vector-signal generator. The new units are compatible with the manufacturer's application software that allows you to configure an analyzer to conduct quality analysis on LTE (long-term-evolution), W-CDMA (wireless code-division/multiple-access)/HSDPA (high-speed-downlink packet access), GSM (global-system-for-mobile-communication), and other wireless signals. The base MS2830A sells for $14,200.**–by Dan Strassberg**

▷**Anritsu Co**, www.anritsu.com.

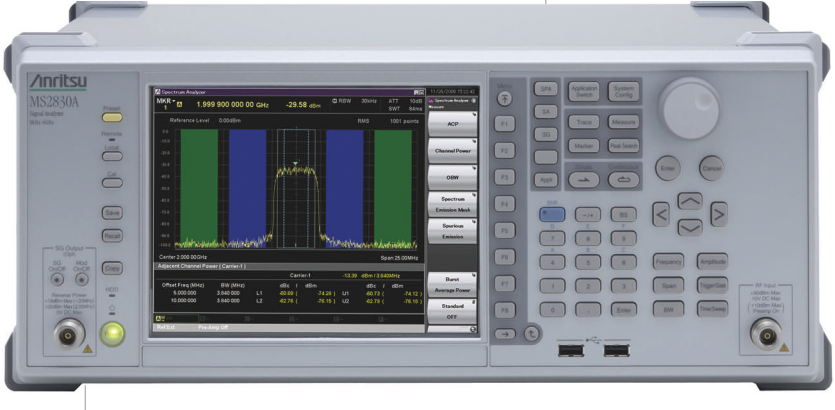

Cost-effective MS2830A signal analyzers, which feature bandwidths from 9 kHz to 3.6, 6, and 13.5 GHz, offer an optional built-in signal generator and a vector-signal-analysis option that batches measurement results, greatly speeding throughput.

FEEDBACK LOOP

"The big, red E-Stop button may not bring about a smooth and dignified stop, but it will cause a stop in spite of any software failure or failures. ... Toyota could have done as well as we do, if it had just been willing to spend a small amount more."

—Engineer William Ketel, in *EDN*'s Feedback Loop, at www.edn.com/article/CA6720350. Add your comments.

Vitesse zeros in on carrier-Ethernet switch needs

The explosion of interest in CE (carrier Ethernet) seems to be at the front of everyone's mind these days. In case you haven't been following the networking market lately, the basic story is this: Just as wireless service providers and conventional telephone-service providers decided to evolve from their legacy switched networks and synchronous rings into Internet-compatible packet networks, they faced a coincidence of three massive trends: the explosion of data traffic in cellular networks, the allure of delivering high-definition television through IP (Internet Protocol) packets to homes, and the rumor of a huge shift toward cloud computing.

The result for both wireless backhaul networks and the wired infrastructure behind all those DSL (digital-subscriber-line) and cable connections was the same. The providers want a packet-based network with enormous bandwidth—such as enterprise Ethernet—but with all the features these providers had from their legacy networks. These services include awareness of the service needs of each flow through the switch, multicast capability, carrier-class reliability and management functions, and support for precise timing. Carriers also want to provide guaranteed QOS (quality of service) for media types such as voice and high-definition video. The answer to all these desires, the industry claims, is CE.

The next question is how to implement CE in a way that can be both fast and cheap. Service providers are blowing right past 40-Gbit switches and asking for 100-Gbit capability, but they are severely financially constrained. The obvious solution is to start with a fast enterprise switch and enhance it to provide the additional services. This approach runs into problems, however.

If your enterprise switch relies on NPUs (network-processing units), you can simply add to the NPU software, but you will almost certainly run out of processing power long before you get all the new features in, even at low wire speeds. If your switch uses an ASSP (application-specific standard product) or an ASIC, it won't be flexible. Either way, you will have to add another ASIC or, more likely, an FPGA or two to the design, running up the BOM (bill-of-materials) cost, power, and design time.

> Service providers are blowing past 40-Gbit switches and asking for 100-Gbit capability, but they are financially constrained.

And, as Morteza Ghodrat, director of CE technology at Vitesse Semiconductor, is quick to point out, adding more packet-processing sites to the design means adding more DRAMs and CAMs (content-addressable memories). Either you put duplicate memory chips around each chip, or you attempt some sort of shared-memory pool with the obvious complications.

To address these problems, Vitesse recently announced three MAC (media-access-controller) and switch chips. Ghodrat argues that all CE services interrelate both architecturally and in efficiency; thus, a single architecture rather than multiple chips should handle them. The new ASSPs, the VSC7460 Jaguar CE switch, the VSC7462 LynX CE switch, and the VSC7364 CE-MaX-24 MAC/switch, bring the full range of CE functions to high-speed Ethernet switches.

Accordingly, each chip has a service-aware classifier that can manage as many as 4000 services, each with its own QOS treatment, DE (discard-eligible) marking, color, policing, OAM (operations/administration/management), performance monitoring, and timing support through IEEE 1588 Version 2. The chips provide advanced QOS and MEF (Metro Ethernet Forum) policing based on a shared 32-Mbit buffer. They also have statistics counters; Ethernet OAM for all 4000 services; and support for 802.1ag, 802.3, Y.1731, and MEF-16 performance assurances.

The chips differ in CPU and I/O complements. The Jaguar has two 10-Gbps XAUI (10-Gbps attachment-unit-interface) and two 10-Gbps VAUI (5-Gbps-attachment-unit-interface) ports on one side and 24 multifunction SGMII (serial-gigabit-media-independent-interface)/SERDES (serializer/deserializer)/100BaseFX ports on the other. The LynX has just half as many of each. Both devices include a 400-MHz MIPS24KEc processor core.

The CE-MaX chip, which operates with an ASIC or an FPGA, uses two XAUI ports and a host interface to attach to the other chip and provides the full 24 SGMIIs plus two more XAUI ports downstream. The CE-MaX has no on-chip CPU. All three chips will be available in 27×27-mm HSBGA packages, and all are scheduled to become available for sampling in the second quarter.

–by Ron Wilson

▷**Vitesse**, www.vitesse.com.

DILBERT By Scott Adams

Single-key capacitive touch controllers target use in portable products

Atmel has incorporated the technology of its QTouch charge-transfer capacitive touch sensors into a series of single-key touch controllers for the portable-device market. Applications extend beyond touch keys for single-button mechanical-switch replacements to a range of body-proximity-sensing scenarios. The chips have a power consumption of less than 17 µA at 1.8V in low-power mode and a fast wake-up time. From the low-power state, when these devices detect a key touch, they temporarily switch to a 12.6-msec fast-response mode, allowing rapid detection of additional touches.

As with other related parts, the AT42QT101X family uses spread-spectrum modulation to ensure good EMC (electromagnetic compatibility). The ICs automatically calibrate when you power them up and remain calibrated, even with moisture buildup or other contaminants on the touch surface.

The AT42QT101X family of capacitive touch sensors target use in body-proximity-sensing scenarios.

> There are few constraints on key shape or size.

The family comprises the AT42QT1010, AT42QT1011, and AT42QT1012. The 1010 part includes a timer to reset a "stuck-key" condition: After 60 seconds, it powers down, resets, and self-calibrates to an assumed untouched state. The 1011 remains in the "touched" condition indefinitely when it senses a contact; this feature allows uses such as detecting when an earpiece of a headset is in the user's ear canal or detection of face proximity in a smartphone. The 1012 includes a touch-on/touch-off, or toggle, key and a configurable power-down timer to shut down devices that users inadvertently leave on.

Users can configure the key sensitivity on all devices for different panel thickness and materials. Electrodes can consist of any conductive material, including transparent indium-tin oxide. There are few constraints on key shape or size, allowing the user a great level of flexibility in the industrial design. The user can also set up the IC for use as a proximity sensor, allowing detection of a nearby hand or object. Toys typically have this function to illuminate "hidden-until-lit" keys when a hand or a finger approaches a consumer appliance. The devices come in a 2.9×1.6-mm SOT-23 package and need only two external components. Prices start at 20 cents (volume quantities).

–by Graham Prophet

▷**Atmel**, www.atmel.com/products/touch.

04.08.10

3-D HOLOGRAPHIC IMAGES DELIVER IMPRESSIVE REALISM, NEED NO GLASSES

Zebra Imaging has developed 3-D holographic images that can replace expensive, bulky, and time-consuming models at construction sites. You can put these displays into a briefcase, hop onto a plane, and show the on-site construction foreman just how a project should look.

Zebra Imaging's 3-D holographic images can fit into a briefcase and replace bulky models at construction sites.

The company burns each image into photopolymer film with intersecting laser beams. Each holographic image starts with a 3-D digital model. If no such model exists, the company can create it by scanning geospatial terrain, digitizing an object, or creating a digital model from scratch using computer graphics or engineering software.

For geospatial applications, terrain data may exist for a region of interest. For architecture and other applications, Zebra accepts many 3-D digital file formats, including those exported by Catia, Maya, Alias, 3D Studio Max, and AutoCAD software. After Zebra receives a 3-D digital model, the customer specifies the size, color, flexibility or rigidity, and opaque- or transparent-mounting options for the hologram.

Imaging times for 2×2- and 2×3-foot tiles are approximately two and three hours, respectively. The cost of a hologram depends on color, quantity, and size. Prices start at $499 for a 12×12-in. hologram, $699 for a 12×18-in. version, and as much as $1999 for a 24×36-in. hologram. Zebra can also mass-produce the images.

–by Margery Conner

▶**Zebra Imaging, www.zebraimaging.com.**

GaN power transistors take aim at MOSFETs

IR (International Rectifier), the first to build GaN (gallium-nitride) devices on silicon, has now introduced GaNpowIR, its first product on a GaN platform. GaN-device structures are not new; using a sapphire substrate, they've been around for 20 years in the RF realm. IR has, however, figured out how to make the ultrafast switching of GaN commercially viable for power electronics.

Consider the impact that hexagonal MOSFETs, the first commercially viable MOSFETs, had on the power semiconductor: According to Mike Briere, PhD, the company's technology consultant, the devices made possible the commercial manufacture of high-volume switching power supplies. The figure of merit for the devices—on-resistance times gate charge—has improved approximately two orders of magnitude over the past 30 years. This rate of improvement can't go on forever, Briere says. He envisions a scenario in which GaN-on-silicon technology will take over.

Enter IR's new iP2010 and iP2011 GaN-on-silicon-technology devices, which operate in native depletion mode, although they can operate in enhanced mode and switch at speeds as high as 5 MHz. The company's road map calls for devices to span 20 to 1200V. The devices integrate a power-stage device that includes a PowIRtune driver IC and a multiswitch monolithic GaN-based power device, all in a flip-chip package, that the company claims more than doubles the switching frequency of the devices.

The iP2010 operates at speeds as high as 3 MHz and features an input voltage range of 7 to 13.2V, an output voltage range of 0.6 to 5.5V, and an output current as high as 30A. Operating at speeds as high as 5 MHz, the pin-compatible iP2011 features the same input and output voltage ranges and an output current as high as 20A. By offering multiple current-rating devices in a common footprint, the devices meet customer requirements in current level, performance, and cost. The parts target the server power-supply market, in which some customers are willing to pay a hefty premium for power efficiency. Prices for the iP2010 and iP2011 begin at $9 and $6 (2500), respectively.

Hot on the heels of the IR announcement comes a new family of power transistors from start-up EPC (Efficient Power Conversion). The company based the devices on its proprietary GaN-on-silicon technology. Their drain-to-source voltages range from 40 to 200V, and on-resistances range from 4 to 100 mΩ. The company claims that GaN devices can reduce total server-power requirements, including ac/dc and dc/dc losses and fans, by about 18%.

EPC developed the GaN technology with a normally off enhancement mode explicitly to replace power MOSFETs. According to Alex Lidow, EPC's co-founder and chief executive officer, enhancement mode, rather than depletion mode, is essential for GaN to become a broad-scale silicon-power-MOSFET replacement. "It's meant to imitate the functionality of a power MOSFET on supersteroids," he says. Because EPC can lay down the GaN structures on standard 6-in. wafers, its prices—80 cents and $5 (1000)—are comparable to those of high-end MOSFETS. The parts are available from Digi-Key (www.digikey.com).

The on-resistance for a device area is a key determinant of cost, and a GaN transistor is markedly smaller than an equivalent MOSFET. Unlike MOSFETs, EPC's transistors lack a packaging structure. The hermetically sealed GaN elements lie atop the silicon layer, which acts as an insulator. You can mount the GaN transistor directly on a heat sink with no surrounding packaging and no thermal resistance between the transistor and its package.

–by Margery Conner

▷**International Rectifier**, www.irf.com.

▷**Efficient Power Conversion Corp**, www.epc-co.com.

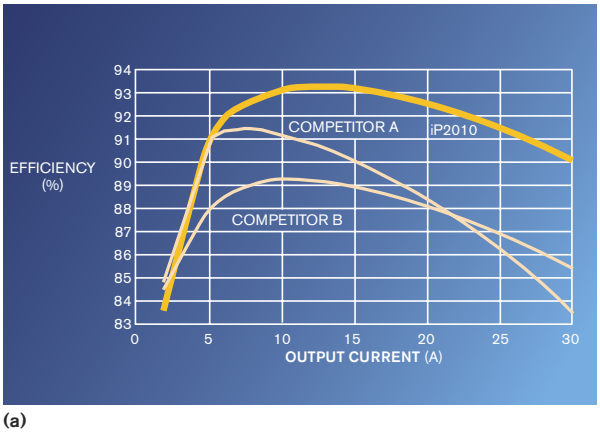

(a)

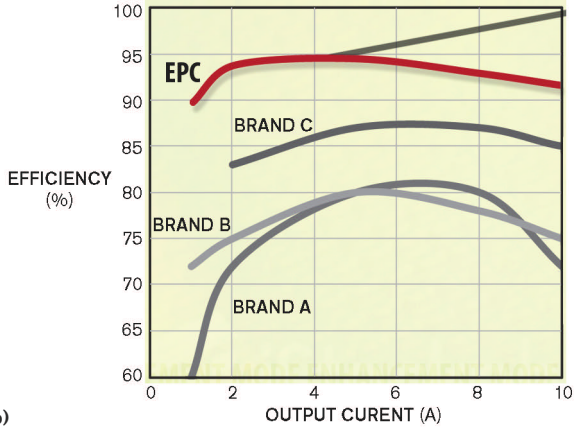

(b)

(c)

Employing International Rectifier's GaNpowIR GaN-on-silicon platform, the iP2010 and iP2011 achieve switching speeds as great as 5 MHz (a). EPC's GaN-on-silicon power transistor (b) is much smaller than a comparable power MOSFET because of its smaller die and lack of a packaging structure (c).

Smart FET driver emulates rectifier in secondary-side applications

International Rectifier's new IR11672 synchronous-rectification IC drives a MOSFET in the secondary side of an ac/dc flyback and resonant half-bridge switched-mode power supply, providing the function of a rectification diode. The FET's forward-voltage drop is smaller than that of a Schottky diode. You connect an input comparator across the power MOSFET, enabling it to sense the direction of the rectified current. The IC operates from an 11.3 to 20V power supply, and its drain-sense pin can withstand 200V. Maximum switching frequency is 500 kHz, turn-off propagation delay is 50 nsec, and turn-on delay is 60 nsec. The device provides peak FET-gate turn-off drive current of 7A, and the unit has both an enable pin and a gate-output pin with a 10.7V clamp circuit.

In addition to conventional gate-drive circuitry, the IR11672 contains minimum-on-time circuitry that blanks the comparator's input to prevent spurious ringing and oscillation from turning the part off. This feature guarantees proper operation in continuous-, discontinuous-, and critical-conduction modes. The IR11672 provides secondary-side synchronous rectification in switch-mode power supplies. It comes in an eight-pin SOIC package, operates in the −40 to +150°C temperature range, and sells for 92 cents (10,000).**–by Paul Rako**

▷**International Rectifier**, www.irf.com.

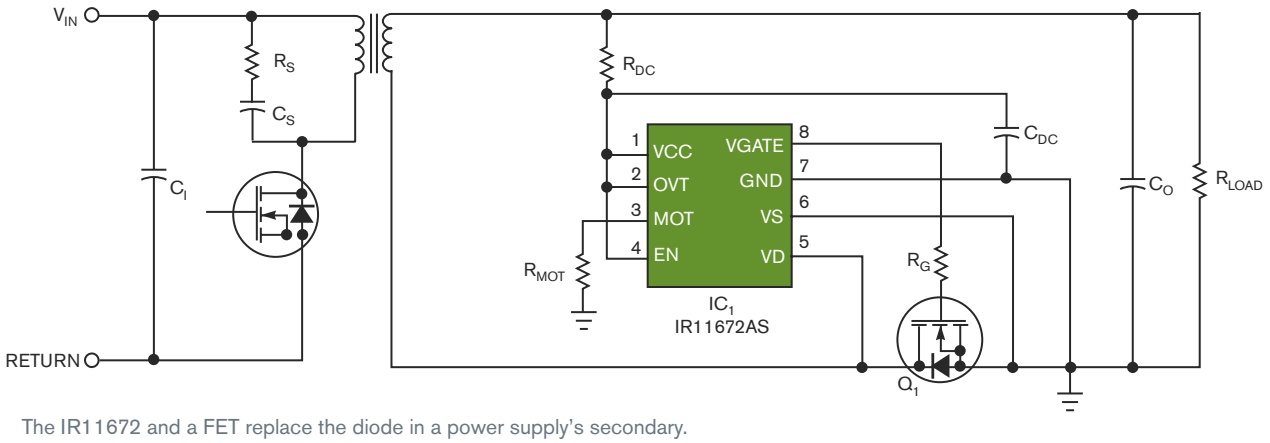

The IR11672 and a FET replace the diode in a power supply's secondary.

SHARP EXPANDS ITS REPERTOIRE OF LCDs FOR IN-CAR USE

The display-research arm of Sharp Electronics has revealed some of its advanced work exploring options in LCD technology for automotive use, including an enhanced version of its dual-view LCD panel and a dual-depth panel that shows information on two visual planes. Sharp supplies its dual-view panels to two car manufacturers, which install them in the center console of the dashboard. The dual-view panel shows two separate images when you view them off-center from either side of the display. In the automotive application, the driver can see only relevant information, such as navigation or car-system menus, while the front-seat passenger views entertainment content. In the first generation of panels, now in production, Sharp achieves the image separation with a parallax barrier: The two images appear on interleaved vertical pixel strips, and a vertical linear grating in front of the screen obscures the set of pixel strips corresponding to the "hidden" image on the respective side.

Now, Sharp has developed an enhanced version of the screen, which may use microlens arrays rather than an aperture grating. This approach allows the company to tune the performance to different use cases; you can have twice the brightness of the previous version at the cost of a somewhat wider region of "mixing"–that is, the viewer sees both images superimposed. Alternatively, you can achieve almost no region of mixed images but with similar brightness to that of the first-generation panels.

The second innovation is a dual-depth panel, which has two distinct visual-image planes that appear at different distances. The optics involve a standard panel with a layered arrangement of optical planes, including a partially reflecting mirror and polarizers, in front of it. The light path through the extra optics is electrically switchable. In one state, light passes straight through, and the viewer sees the LCD screen directly. In the alternative state, light from the panel internally reflects within the arrangement of polarizers before exiting in the direction of the viewer. The optical-path length is increased by twice the thickness of the polarizer/mirror-assembly layer, and the image therefore appears to be displaced backward by that amount.–by Graham Prophet

▶**Sharp, www.sharpsme.com.**

04.08.10

Rethinking static-timing analysis

STA (static-timing analysis) was nearly an instant success at timing closure 15 years ago, and nothing much has changed since except for creating partitioning/scheduling algorithms to parallelize the algorithms for multicore CPUs. This stasis has allowed an increase in the number of instances in a design, the number of modes in which you must analyze a design, and the number of process corners. Consequently, runtimes for full designs across modes and corners have become enormous—days, in some cases.

That situation has turned STA from an elegant, fast tool into a powerful and trusted, but ponderous, necessity, consuming licenses and days of precious schedule with abandon. If there were a way to dramatically speed up STA, users would need fewer licenses, could save the real-estate and power costs for huge server collections, and could employ the tool in situations in which it has become impractical today, such as checking timing constraints or evaluating ECOs (engineering change orders).

Several alternatives exist for accelerating STA. To begin with, the task is parallelizable. Most nets are independent with regard to delay, so you can organize the nets into independent sets and dispatch them to independent threads. That fact makes STA inherently friendly to multicore computing and, more practically, to execution on graphics processors.

According to Dan Blong, technical-marketing manager at Magma Design Automation, 15 years have elapsed since anyone looked at the underlying algorithms and the code to see whether there were other ways to accelerate the analysis. A team headed by Pathmill pioneer Jacob Avidan set out two years ago to accomplish

> An incremental mode lets you work directly on particular blocks of IP.

that goal. The result is Tekton, a new STA/extraction/Spice environment. Magma claims that Tekton runs significantly faster as a timing analyzer on a single CPU and dramatically faster in a multimode/multicorner analysis on a multicore machine. The company claims that it can perform these tasks on any design, using one machine, and in less than an hour, further claiming a speed of 1 million nets per minute and near-linear scaling for as many as 24 CPUs.

The improved performance comes primarily from careful organization of the work to avoid the need for repeating calculations—that is, recognizing tasks that are unnecessary and saving intermediate results for use in future calculations. It does not mean compromising timing accuracy. Blong says that Tekton correlates well with PrimeTime delay results and approximates PrimeTime crosstalk results.

Tekton drops into existing flows, accepting PrimeTime tcl and Perl scripts. The complete Tekton environment includes the Tekton QCP extraction tool, which correlates to QuickCap, and crosstalk analysis, on-chip variation analysis, and an integrated Spice engine. In addition to timing an entire design in less than an hour, the package includes an incremental mode that lets you work directly on particular blocks of IP (intellectual property) or on ECOs.

–by Ron Wilson

▷**Magma Design Automation**, www.magma-da.com.

INTEL UNVEILS 32-nm DATA-CENTER MICROPROCESSORS

Aiming to provide data centers with a stronger foundation for cloud security, Intel recently announced its Xeon 5600 series, which combines security, performance, and energy-efficient features, according to the company. The processors contain Intel's AES-NI (Advanced Encryption Standard-new instructions), which Intel introduced with the Core processor family. These instructions accelerate AES performance to allow faster data encryption and decryption for applications such as database encryption, full disk encryption, and secure Internet transactions. The processors also feature TXT (trusted execution technology), which allows faster encryption and decryption performance for more secure transactions and virtualized environments. These workstation and server chips use the company's 32-nm logic technology, which employs Intel's second-generation, high-k-metal-gate transistors to increase speed and decrease energy consumption.

The devices allow as many as six cores to reside on each processor and can deliver as much as 60% greater performance than the 45-nm 5500 series. Intel claims that data centers using the new devices can replace 15 single-core servers with one. A two-socket server using the new series' low-voltage L5640 can deliver the same performance as but consume as much as 30% less power than a server using the previous generation's X5570 series. The processor also integrates TXT-hardware features that prohibit intrusion from malicious software, allowing applications and data to run more securely in a virtualized environment. Together, these features ensure that virtualized environments experience better performance and are more secure when users deploy or migrate them.

The frequency-optimized quad-core version of the 5600 series offers a maximum speed of 3.46 GHz with a total power dissipation of 130W, and the six-core version reaches 3.33 GHz with a power dissipation of 130W. Advanced six-core versions will top out at 2.93 GHz and 95W, and the standard quad-core processor will reach 2.66 GHz and 80W. Low-voltage versions of the chip will have power dissipations as low as 60 and 40W and feature six and four cores, respectively.

–by Ann Steffora Mutschler

▶**Intel Corp, www.intel.com.**

04.08.10

Avago Fiber Optics: Breaking Bandwidth and Performance Barriers in Supercomputing

Introduction

Modern datacenters, high-performance computing (HPC) centers, and supercomputers demand dense high-bandwidth cost-effective interconnect solutions for the fabric between various computing resources and network elements. The bisection bandwidth of this fabric is a key factor in determining the overall performance of the system. As applications (internet search, cloud applications, complex computations, virtualization) and hardware (faster multicore processors, more I/O per server) put more demands on resources, the fabric bandwidth must scale as well.

Parallel interfaces, which can be either electrical or optical, provide a convenient means to achieve these interconnects. In contrast to single-channel connectors, parallel interfaces contain multiple lanes in a single electrical or optical connector. Logically, the interface can be a collection of coupled lanes that provide an aggregate channel bandwidth that is not achievable with serial technology (4xQDR Infiniband provides 40Gbps of aggregate channel bandwidth using four 10Gbps lanes) or a collection of individual single-channel interfaces (such as the fabric of a multistage switch/router). The interface can also be a combination of the two (three 4xQDR channels on one 12-lane interface).

As lane speeds move to 10 Gbps and beyond, fiber-optic technologies are displacing copper-based solutions that have long dominated these parallel interfaces. Fiber optics offer distinct advantages in density, power, weight, and link length. However, a thorough understanding of the benefits of fiber optics over copper goes beyond these simple metrics. Architects can choose where to deploy dense parallel optics solutions to achieve the optimal I/O bandwidth density: rack-to-rack, within a rack, and within board. In this technical note, we review various parallel optical technologies and then focus on embedded solutions developed by Avago that uniquely enable our customers to bring differentiated systems to market.

Pluggable parallel optics

Today, parallel-optical modules are widely available in two formats: pluggable and embedded. Pluggable modules are applicable when the aggregate I/O requirements of a system can be accommodated by the limited front-panel area. Common packages include QSFP (4-lane interface) and CXP (12-lane interface) with line rates running up to 12.5 Gbps. Examples of these form factors are shown in Figure 1. The advantages of the pluggable solutions include pay-as-you-grow deployment, choice of copper and optical options, and ease of assembly. However, density, thermal, EMI, and signal integrity challenges sometimes lead to tradeoffs in system performance in order to accommodate the pluggable format.

Figure 1: Avago AFBR-79Q4Z QSFP and AFBR-83BDZ CXP pluggable parallel transceiver modules.

Embedded parallel optics

Avago Technologies is the market leader in delivering embedded parallel optics for applications such as the fabric between chassis in multi-chassis routers/switches and high-performance computing clusters. Historically, the dominant form factors have been SNAP12 and POP4, with speeds up to 3.125Gbps per lane. Avago has extended the bandwidth capability to as high as 10 Gbps per lane. These modules can be mounted at the card edge as well as mid-board to optimize their placement for signal integrity and air-flow. In the latter case, fiber jumper cables are used to provide the optical interface to the panel. The use of high-density optical connectors, such as 24-, 48-, and 72-ferrule MTP™, enables single ports with extreme amounts of bandwidth. For example, a 72-ferrule MTP connector with each lane running at 10 Gbps can provide a 360 Gbps bidirectional interface in less than 1 square inch of panel area, equivalent to the connectivity provided by three CXP modules.

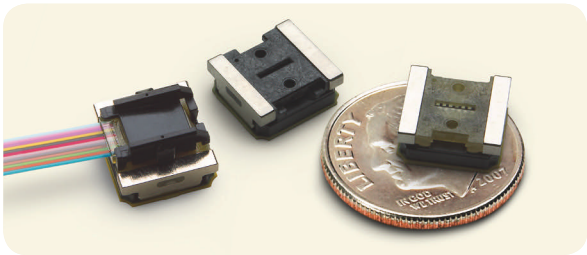

Figure 2. The MicroPOD, Avago's next-generation embedded parallel-optics module and optical connector shown with a dime for scale.

Next-generation computing, switching, and routing systems will require Terabytes of interconnect between nodes. In order to accommodate pluggable solutions, all of the I/O must be routed to the front panel. Even if edge-mounted optics could provide sufficient density, significant signal integrity and thermal challenges would need to be overcome and the number of individual connections would be cumbersome. Embedded parallel optics can provide a solution, but legacy form factors (POP4 and SNAP12) are large and require keepout areas for the MTP connectors. To address higher density applications, Avago has introduced a novel parallel optics module, called MicroPODs (Figure 2), with unprecedented density (7.8 mm x 8.2 mm for 12 channels). The 12-channel transmitter and 12-channel receiver modules nominally operate at 10 Gbps per lane (120 Gbps aggregate bandwidth). A top-attached optical connector with a convenient slot for fiber routing permits dense tiling of the MicroPODs (Figure 3). At the card edge, high-density optical connectors can be used for consolidated I/O. For high-speed signal integrity, a µLGA electrical connector is used to mate the module to the host. The module includes such features as programmable per-lane equalization and de-emphasis as well as a full set of diagnostic monitoring capabilities.

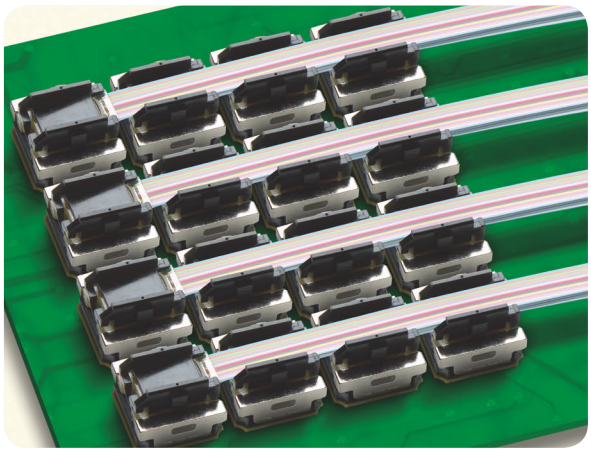

Figure 3. A densely tiled group of 8 MicroPOD Tx/Rx pairs provides a total bidirectional bandwidth of (960 + 960) Gbps in less than 3 square inches of board area.

Application example: High-performance computing

High performance computing has allowed advances in a wide variety of areas, including climate modeling, drug design, car crash studies, oil reserve discovery and computation-based business strategies. Businesses, education and research institutions, and governments increasingly rely upon HPC to provide timely answers to computationally intensive analysis and design problems. HPC and supercomputer manufacturers are striving to improve the performance of these machines to address these applications.

HPC and supercomputer systems are clustered networks of computing resources. The connectivity between computing nodes, or interconnect bandwidth, is an important factor in determining the useful computational power of the system. As individual processors increase in computation capability, the interconnect bandwidth between nodes of the cluster must scale to achieve the desired system performance. Where individual processor cores today may have 1 - 10 Gpbs of connectivity bandwidth to the fabric, next generation systems will require 10s to 100s of Gbps per core. Furthermore, many processor cores can reside in one node, requiring 10's of Tbps of connectivity from each node. Pluggable solutions cannot achieve this scale of bandwidth, considering the difficulties of routing 1000's of high speed lanes to the card edge as well as real-estate limitations.

Avago's MicroPOD technology alleviates the I/O bandwidth escape limitations imposed by front-panel mounted parallel optics. Supercomputer architects have freedom to scale the bandwidth of the inter-processor communications to optimize system performance with limited restrictions on the I/O bandwidth. For example, a 1 Teraflop node with ~2 Tbps/sec of connectivity to the network would require 16 Tx/Rx MicroPOD pairs. These modules would be positioned near the host IC for optimal signal integrity while also arranged in a tiling pattern for density. Using high-density MTP connectors at the front-panel, the I/O could be consolidated into ports that have the ideal bandwidth for node-to-node communications. The area required for these ports would be significantly less than that required for a similar CXP-based interface.

Additionally, whereas the CXP must be mounted in the plane of the host PCB, the passive MTP connectors could be arbitrarily positioned for optimal density and fiber management. These benefits of using MicroPOD can enable more processing power to be packed onto one physical compute node, thereby conserving valuable space and power in the data center.

Summary

Increasing demands on data centers and computational facilities are driving requirements for connectivity between computing resources. Pluggable solutions such as QSFP and CXP can address some of these applications, but very dense I/O requirements are best met with embedded optical solutions. Avago has addressed this market need with a novel parallel-optics module, the MicroPOD. The MicroPOD enables true system differentiation by relieving performance-restricting bandwidth bottlenecks.

BY HOWARD JOHNSON, PhD

Manager's guide to digital design

This article is for all those hardworking engineering managers who just want a simple, one-page summary of everything they need to know about digital design.

GROUND: mythical electrical object that absorbs unlimited quantities of electrical current. Ground exists in Spice but nowhere else. Radar engineers in the 1930s discredited the concept of ground as anything more than a good place to grow carrots and potatoes.

PARASITIC EFFECT: a mysterious circuit effect that persists despite your best attempts to either eradicate it or blame it on a previous designer.

RISE TIME: rate of ascent through management ranks. Rise time may be sluggish or, in the case of an overambitious candidate, subject to wild overshoot and gyration.

EQUALIZER: the Chuck Norris of serial-transmission circuits. An equalizer improves the odds of success for all good bits by knocking out the bad artifacts. Just saying you have an equalizer makes investors swoon.

ADAPTIVE EQUALIZER: Chuck Norris with brains.

DECISION-FEEDBACK LOOP: a critical management decision made, then retracted, and then made again in a repeating loop. This loop knocks productivity to zero.

HEAT SINK: a small metallic device attached to your CPU that, like the cooling tower at a nuclear-power plant, is the only device standing between safe, reliable system operation and total core meltdown.

Figure 1 The author visualizes his next great product, or is he just fooling?

NETWORK ANALYZER: an expensive piece of gear that always reports bad news. When interviewing new hardware-engineering candidates, always say that you have one. Later, if that employee misbehaves, threaten to make him use it.

INDUCTOR: a two-terminal component, which, like a financial derivative, causes huge spikes followed by systemic crashes. Avoid speaking with engineers for at least two days if they mention the words "inductor" and "Spice" in the same sentence.

DE-EMBEDDING: the insomnious effect of a difficult measurement problem that consumes copious amounts of overtime, often late at night, preventing normal sleep.

POWER-SUPPLY DROOP: a diminution in the output of a healthy power supply when engaged in vigorous activity. An insufficiently turgid power supply droops to the point of ineffectiveness. No pill cures that condition.

BOOT TIME: the interval of time between your decision, based on how slowly your product runs under real-life conditions, to fire your chief software engineer and the moment his feet hit the pavement outside your building.

SSO (simultaneous-switching-output) noise: a feature. IC manufacturers believe that only by skimping on the number of power and ground pins can they offer high-speed IC products in inexpensive packages, thereby transforming SSO noise from a problem into a feature.

ROHS (restriction-of-hazardous-substances) Lead-Free Solder Initiative: Evil plot by Luddites to rid the world of computers by first rendering all electronic products flaky and unreliable. The initiative may precipitate the collapse of Western civilization. Until then, just smile and go along with the scheme like everyone else.

VISUALIZATION: a mental process. More than any other group, hardware engineers must visualize a solution to every problem, which they do with their eyes closed. You may hear noises that sound like ordinary snoring but actually indicate a deep state of complete concentration. Never interrupt an engineer engaged in visualization (**Figure 1**).EDN

Howard Johnson, PhD, of Signal Consulting, frequently conducts technical workshops for digital engineers.

"How can I tell if a power supply is reliable?"

There's an indicator on the front.

It says "Agilent." With a typical MTBF of 40,000 hours, over half-a-century of experience, and with more than 250 models to choose from, Agilent's power supplies are the ones you can count on. In fact the array of our power supplies is so extensive, it wouldn't fit on this page. For clean, low-noise, programmable power to countless DUTs, there's an Agilent power supply with your name on it. Actually, it's our name on it, but you know what we mean.

For free measurement tips and the Agilent Power Products brochure go to www.agilent.com/find/powertips

Agilent and our Distributor Network
Right Instrument. Right Expertise. Delivered Right Now.

Buy from an Authorized Distributor
www.agilent.com/find/distributors

Agilent Technologies

AS LTE NETWORKS ROLL OUT THIS YEAR, ENGINEERS WILL BE APPLYING VARIOUS TEST SCENARIOS IN LAB AND SMALL-SCALE FIELD TESTS TO THE LTE DEVICES AND EQUIPMENT THEY ARE DEVELOPING.

LTE in FOCUS

BY REINER GOETZ, ANNE STEPHAN, AND MEIK KOTTKAMP, ROHDE & SCHWARZ

As part of the worldwide 3GPP Release 8 standard, LTE has been fully defined since March 2009. This new technology is an essential enhancement of classic mobile-communications technologies, such as GSM/EDGE, WCDMA/HSPA(+), and CDMA2000/Ev-Do Rev A. It promises end users significant data-rate, network-capacity, and latency improvements. LTE will also enable network operators to more effectively and inexpensively deliver services, such as Web browsing, gaming, and video streaming, and it will open the door for new mobile services.

Consequently, no fewer than 59 operators in 28 countries have announced plans to deploy LTE on their networks. A sizable number of commercial networks will roll out in 2010. Therefore, vendors of LTE products and infrastructure are performing extensive testing ranging from lab tests on individual devices to small-scale field tests with just a few base stations and user devices from a variety of vendors. Major test networks, too, with numerous base stations and a significant number of initial LTE user devices, are undergoing user trials.

What requirements does this new technology have to meet, and how can you assess and effectively verify LTE performance in a laboratory environment by means of suitable tests? Consider the differences between LTE and other technologies before proceeding to examine the various phases in the development of LTE-enabled infrastructure and end-user equipment.

SIMILAR BUT DIFFERENT

The 3GPP participants developed LTE on the basis of available technologies. Therefore, it is hardly surprising that, despite its numerous differences, LTE has a lot in common with them. Consider HSPA(+), an established technology in packet-oriented services. In LTE, just as in HSPA(+), the allocation of resources for the transmission of data from the base station to user equipment relies on a rapid feedback mode in the user device. The device ascertains the quality of the transmission channel and informs the base station what maximum resource size to allocate. The only difference is that LTE offers a much faster feedback mode than does HSPA(+): LTE uses a 1-msec time-transmission interval, whereas HSPA(+) uses a 2-msec TTI and WCDMA uses a 10-msec TTI, allowing the data rate to adapt to current transmission conditions practically every millisecond.

The biggest differences between LTE and current 3GPP standards lie in the technologies manufacturers use to implement the air interface. LTE employs OFDMA and MIMO. In addition, LTE works with a flat IP-based network architecture. OFDMA enables granular resource allocation because LTE uses a large number of narrowband subcarriers with a bandwidth of 15 kHz, compared with 200 kHz in GSM and 5 MHz in WCDMA (**Figure 1**). This enhancement combines with a maximum chan-

nel width of 20 MHz and the ability to work with as many as four transmitting and receiving antennas to create the basic conditions necessary for meeting high-data-rate and high-capacity requirements.

LTE also greatly simplifies the mapping of logical channels onto physical channels. Shared channels have replaced dedicated channels, and LTE has fewer MAC-layer entities and RRC states for greater simplicity. In contrast, the number of parallel processes in the protocol has increased, and you can use MIMO to combine multiple data streams. The encryption function has also changed: In LTE, the eNodeB and MME use different keys. The data in the PDCP layer and the NAS layer are encrypted differently; in WCDMA, the NAS layer is not encrypted at all.

HIGH SPEED, COMPLEXITY

The way in which LTE integrates into networks also plays a crucial role. From the outset, developers of LTE user equipment have been under enormous time pressure. Even before LTE's core specifications were complete, equipment makers augmented the then-unfinished specifications based on their own assumptions and implemented their own protocol "dialects." Their aim was to be in a position to demonstrate LTE's ability to function and its benefits as soon as possible. Now that the developers have finalized the specifications to a sufficient degree, the focus is on reducing development time to be able to deliver LTE products to the market as swiftly as possible.

LTE's higher complexity poses significant challenges because, for the most part, LTE resides in user devices alongside technologies such as WCDMA, CDMA2000, and GSM. The presence of multiple technologies entails a variety of hand-over scenarios, all of which designers must test. In addition, LTE user devices must support other noncellular standards—Wi-Fi, GPS, and Bluetooth, for instance—that complement the wide operational coverage that cellular technologies afford.

Ideally, to enable steps in the development process for a mobile phone to take place concurrently, developers pursue an approach employing reusable modules. However, the developers must test the various components as early in the process as possible to minimize the number of potential errors during integration and to avoid creating problems in subsequent field tests. The test-and-measurement equipment they use must therefore be able to separately drive or bypass layers or functional modules. The modular approach, for example, enables developers to test the functions they loaded in software onto a baseband chip in a virtual environment without relying on the availability of the hardware.

Other LTE features at the center of development efforts are the data rates on uplinks and downlinks. Just as in

AT A GLANCE

- LTE promises end users significant data-rate, network-capacity, and latency improvements.
- The biggest differences between LTE and 3GPP standards lie in the technologies to implement the air interface.
- LTE will find use in user devices alongside technologies such as WCDMA, CDMA2000, and GSM.
- In RF-signaling tests, testers examine the user equipment's transmitter and receiver in combination with all of the signaling layers.
- The use of OFDM, which enables TTI-based allocation of resource blocks, leads to significant changes in testing requirements.
- MIMO requires extended antenna systems at the base station that you must verify with the aid of signal analyzers.

LTE GLOSSARY

ACLR: adjacent-channel-leakage ratio
AWGN: additive white gaussian noise
BLER: block-error rate
CDMA: code division/multiple access
CPRI: common public-radio interface
CQI: channel-quality indicator
CS: circuit-switched
EDGE: enhanced data rates for GSM evolution
eNodeB: base station
ETSI: European Telecommunications Standards Institute
Ev-Do: evolution-data optimized/evolution-data only
EVM: error-vector magnitude
GCF: Global Certification Forum
GPS: global positioning system
GSM: global system for mobile communications
HSPA: high-speed packet access
I/Q: in-phase/quadrature
IP: Internet Protocol
LTE: long-term evolution
MAC: medium-access control
MIMO: multiple input/multiple output
MME: mobility-management entity
NAS: nonaccess stratum
OBSAI: Open Base Station Architecture Initiative
OFDMA: orthogonal frequency-division multiple access
PCS: personal communications service
PDCP: Packet Data Convergence Protocol
PMI: precoding matrix indicator
PRBS: pseudorandom bit sequence
PS: packet-switched
PTCRB: PCS Type Certification Review Board
QAM: quadrature-amplitude modulation
QPSK: quadrature phase-shift keying
RI: rank indicator
RLC: radio-link control
RRC: radio-resource control
RRH: remote radio heads
16QAM: 16-phase QAM
64QAM: 64-phase QAM
SMS: short-message service
3GPP: Third Generation Partnership Project
TTCN2: Testing and Test Control Notation Version 2
TTCN3: Testing and Test Control Notation Version 3
TTI: time-transmission interval
UL CQI: uplink CQI
U plane: user plane
WCDMA: wideband CDMA

HSPA(+), they place considerable demands on the user equipment. To assess performance, manufacturers must run tests that evaluate the acknowledgments and negative acknowledgments on the RLC and MAC layers. Designers can achieve the specified data rates using the wide variety of MIMO modes. However, the test-and-measurement equipment must support these modes. Besides the ability to connect multiple antennas, the test equipment must also be able to simulate fading channels. Only then is it possible to test the functioning of receivers under realistic conditions.

Scalable bandwidths and the 17 regional frequency bands further increase the breadth of required testing. And developers must also address limitations, such as different possible power levels in tests at high bandwidths on frequency bands with low separation between receiving and transmitting frequencies.

This high complexity calls for frequent regression tests, such as testing daily software builds or the performance of endurance-testing scenarios in a realistic signal environment. Test equipment with advanced automation and remote capabilities makes it easier to efficiently conduct these tests. The range of measurements is not confined to verifying that developers have correctly implemented the specifications; it also extends to modules' stability and robustness when facing varying interpretations of specifications (**Figure 2**).

At the end of the development process comes conformity testing on certified test systems, during which developers run a selected number of tests from the 3GPP conformance-test specification on the finished user device. Unlike functional tests, this type of testing formally verifies the protocol layers and the requisite RF performance.

MOBILE USER DEVICES

Testers primarily conduct pure RF-performance testing using signal generators and analyzers; combined RF-signaling tests are also important. These tests examine the user equipment's transmitter-and-receiver combination with all of the signaling layers. The tests are simulations that closely approximate signaling procedures and scenarios in live use under realistic conditions, with possible interfering signals, and during continuous operation of the device. The prima-

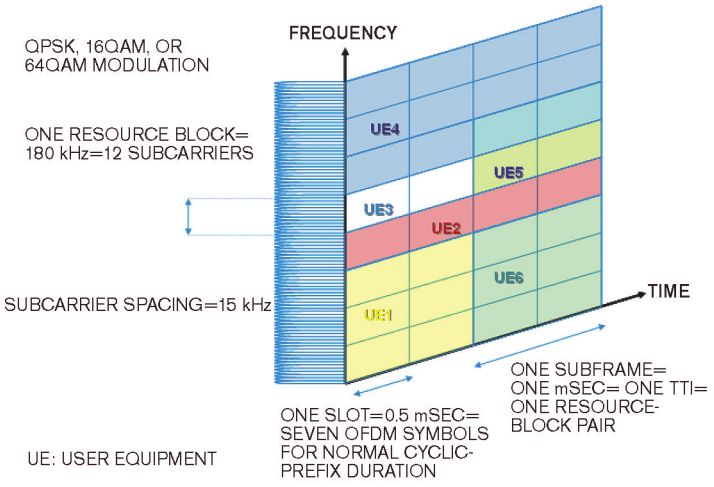

Figure 1 LTE employs an OFDMA-modulation scheme that employs a large number of narrowband, 15-kHz subcarriers to enable granular resource allocation.

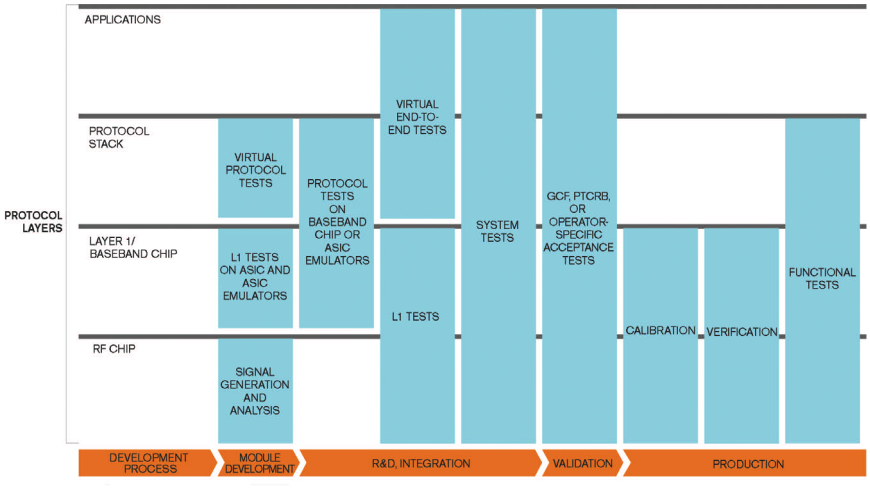

Figure 2 LTE development involves various module and integration testing procedures as the product moves toward production.

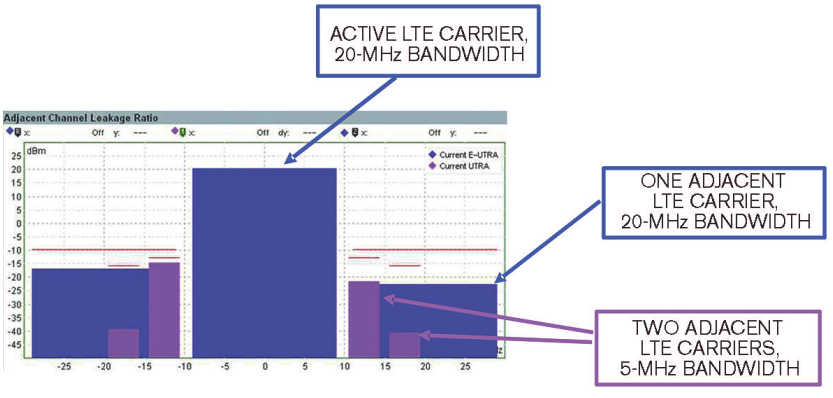

Figure 3 ACLR measurements on an LTE signal help to prevent interference between neighboring WCDMA and LTE systems.

ry emphasis is not on testing the actual signaling procedures; rather, the signaling serves as a means to an end for performing realistic tests on the entire device. Testers perform separate tests for the transmitter and the receiver.

Testers apply a variety of measuring methods during the transmitter tests. First, they test LTE signals using proven methods—power and EVM measurements, for example—they adopted from other mobile-communications technologies. Second, they verify extensive procedures, such as power control based on profiles in LTE and WCDMA. Many of the measurements may resemble well-known procedures. With LTE, they are more complex, however. A spectrum measurement is a case in point: The fact that LTE and WCDMA frequency bands may be adjacent to each other places exceptional demands on the user equipment. To help prevent interference between neighboring WCDMA and LTE systems, the transmitting power in adjacent bands must not exceed either LTE- or WCDMA-specific limits (**Figure 3**). An extended ACLR test can check that it doesn't.

The use of OFDM, which enables TTI-based allocation of resource blocks, has led to significant changes in testing requirements. The measuring equipment must flexibly configure the requisite assignment tables and scheduling parameters for the uplink and downlink and send these tables and parameters to the user equipment. Meanwhile, the testers must check the correct allocation of resource blocks and the transmitting characteristics of the user equipment on the uplink (**Figure 4**).

Given that multiple user devices can concurrently use the available bandwidth, testers must measure inband emissions to determine whether the user device complies with allocation and transmitting power requirements on the uplink. This approach ensures that the device does not interfere with other uplink signals outside its allocated resource blocks. Measurement equipment that can flexibly set limits and independently check limits greatly simplifies testing (**Figure 5**).

Because of the breadth of allocation options available, testing generates a large number of results. These results depend extensively on the location and size of the allocated resource blocks within the time and frequency domains, and developers must therefore interpret them in context. In addition, some RF impairments have an effect only on certain allocations.

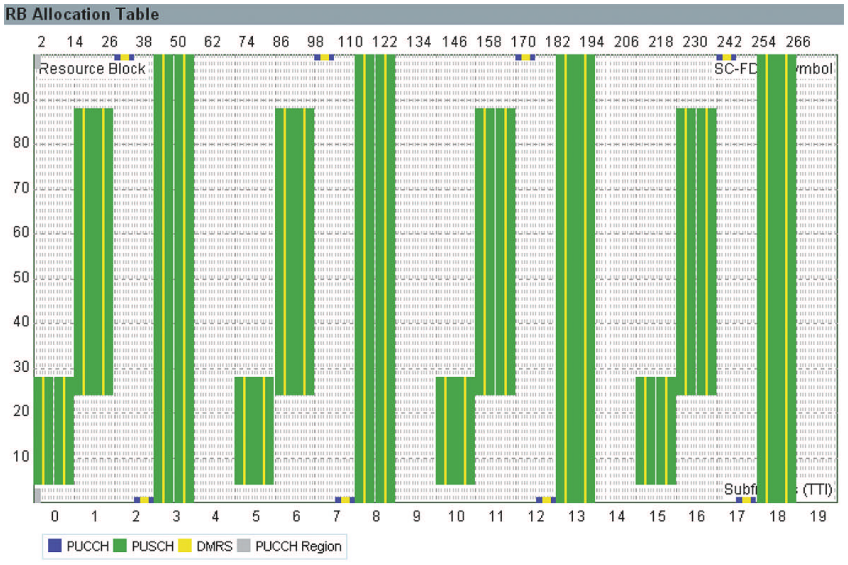

Figure 4 The use of OFDM enables TTI-based allocation of resource blocks; this measurement shows partial allocation of resource blocks.

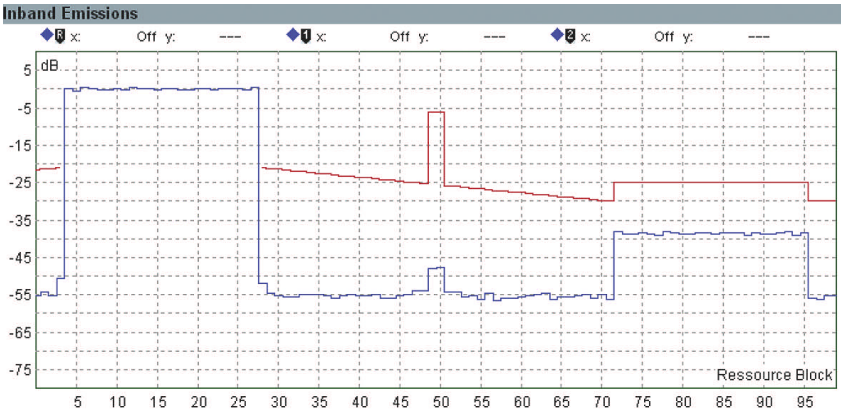

Figure 5 You must measure inband emissions to determine whether the user device complies with allocation and transmitting power requirements on the uplink.

The distribution of transmitting power across multiple subcarriers can lead to power differences between subcarriers. You can examine transmitting power at the subcarrier level by testing spectrum flatness, thereby enabling users to identify potential fluctuations with exceptional precision.

In receiver tests, the MAC layer uses localized acknowledgment/negative-acknowledgment-based BLER methods. These methods for analyzing uplink signals are familiar from HSPA. With LTE MIMO, the focus is on a scenario in which you apply various fading profiles to the downlink signal. To reduce development time and costs, you can use static channel models that simulate a static fading profile instead of dynamic fading profiles. The static models enable you to analyze the effects on receiver behavior using the BLER methods. In HSPA, too, you measure the downlink signal with fading and AWGN. In LTE, other technologies inside and outside the LTE band cause additional interference signals, calling for wider blocking tests and adjacent-channel tests.

The follow-UL-CQI test, again familiar from HSPA, is an important means of adjusting the signaling parameters and thus optimizing the receiving-signal quality that a user device reports through the CQI. Several values affect the quality in LTE, including Rank 1 or Rank 2 CQI, PMI, and RI values. Dynamically changing parameters on active connections in the measuring equipment can help to save time during these tests.

Testing user devices calls for a range of measurement methods suitable for checking the transmitter's RF in combination with the allocated resource blocks on the uplink. Ideally, data for computing transmitter measurements should originate from a test sample and appear simultaneously in a clearly structured form (**Figure 6**).

Figure 6 A test instrument's multievaluation mode presents all RF-transmitter measurements at a glance.

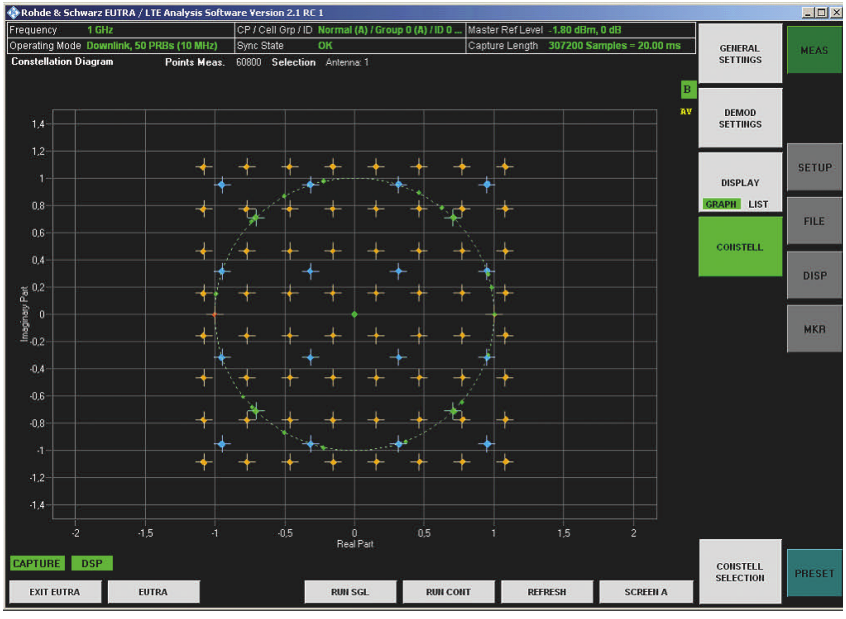

Figure 7 This constellation diagram for an eNodeB signal shows how multiple elements representing, for example, different data rates or modulation schemes combine into a composite signal.

PERFORMANCE TESTING

All LTE user devices to date have been data devices—in other words, USB sticks and PC cards. The data-services sector has mainly driven the motivation for introducing LTE. Nevertheless, the continuing debate over technical alternatives for the voice service and the fact that a number of special-interest groups focusing on it have formed within the industry substantiate the emphasis on equally efficient voice support. It remains to be seen which of the alternatives will gain primacy. However, data services pose greater requirements than the voice service as far as protocol test-and-measurement systems are concerned.

In LTE, only the PS domain exists, not the CS domain. In general, multiple services with different bearers operate in parallel in a manner comparable with WCDMA multicall services. Furthermore, once users power up and register their devices, the devices immediately have always-on status and can almost instantly issue requests to transmit data on the uplink or the downlink. Consequently, functional tests spanning all layers always require the measuring equipment to provide a service that delivers data through the U plane.

A mobile device's performance is of direct relevance for end users. How fast is the data rate, and what is the latency when starting services? To what extent does performance degrade when reception is poor? Does the manufacturer guarantee that the device will interoperate with different base stations? When finding answers to these questions and optimizing a data device, it is not enough simply to check signaling procedures or individual values you measure at the IP level. Rather, it is important to analyze bottlenecks in the protocol layers: Which level is causing unnecessary retransmissions? Why does the BLER increase under certain conditions? Protocol test-and-measurement equipment must answer these questions and verify signaling procedures. Thus, the lines between classic application tests and protocol tests are becoming increasingly blurred.

If you test modules during develop-

ment, the test equipment must provide the necessary interfaces. In the past, it may generally have been sufficient to connect to the user equipment through RF, but it is now essential to provide interfaces on the I/Q baseband because the protocol software runs on the baseband chip or on a chip emulation. It is even possible to completely test a protocol stack without hardware if you replace the physical layer with emulation software. Access to details of lower-level protocol-layer configurations is essential for meeting all these requirements.

Manufacturers will only gradually roll out LTE networks. The rollout will require thorough testing of the handover signaling, as it is important to ensure that user equipment can transition smoothly between technologies. Thus, it is essential for test-and-measurement equipment to provide a basic implementation of all technologies and support for synchronization with LTE. Given that MIMO plays a central role in enabling higher data rates, you must test complex signaling procedures and user-device feedback.

CONFORMANCE TESTING

GCF certification of LTE user equipment should begin in late 2010. Validation of the first test cases employing 3GPP specifications 36.521 and 36.523 has occurred. ETSI has chosen TTCN3 as the language for describing the tests. TTCN3 is an enhancement of TTCN2, the language used with WCDMA, and now has more in common with a traditional programming language such as C++. TTCN3 is therefore easier to learn, and adjacent areas of development, in addition to certification tests, will likely adopt it.

In addition to the essential user-equipment certification that GCF and PTCRB perform, network operators must also perform certification to their own high standards. These tests place greater emphasis on the characteristics of the network infrastructure and on optimization of the available mobile-communications services.

A CHANGE OF PERSPECTIVE

The swift and efficient development of LTE base stations represents a core challenge for infrastructure vendors. As a rule, they deploy test systems well before the commercial rollout of networks. If possible, these systems should run on the network infrastructure alongside commercial platforms during the test phase. This is why vendors conducted numerous LTE field trials in 2009. To determine which tests to perform during the base-station-development cycle, vendors drew on experience they gathered over many years of successfully operating mobile-communications systems, such as GSM and WCDMA. **Table 1** lists the measurements they performed on transmitters and receivers.

TABLE 1 MEASUREMENTS ON TRANSMITTERS AND RECEIVERS

Transmitter	Receiver
Output power and dynamic range of output power	Sensitivity and receiver dynamic range
Signal quality, such as frequency errors and EVM	Channel selectivity
Unwanted emissions	Immunity to interference with adjacent interfering signals on frequency band (blocking)
Transmitter intermodulation	Interfering emissions in the receiver
	Receiver intermodulation
	Performance tests under various channel conditions

THE SWIFT AND EFFICIENT DEVELOPMENT OF LTE BASE STATIONS REPRESENTS A CORE CHALLENGE FOR INFRASTRUCTURE VENDORS.

Certain technical aspects of LTE are of special importance in the development of infrastructure products. MIMO requires extended antenna systems at the base station that you must verify with the aid of signal analyzers. To measure a precoded MIMO signal from two transmitting antennas, you must record both data streams at once. LTE technology uses complex precoding matrices on the transmitting branch. Measuring module parameters, such as EVM, for example, generally requires information from both transmitting signals. This requirement calls for a modular-test-equipment setup with two connected analyzers. The measured values from the first signal analyzer transfer to the second on a master/slave basis.

In addition, LTE uses shared frequency channels that multiple user devices use jointly. These devices may use different data rates and, therefore, different modulation types, including QPSK, 16QAM, and 64QAM. In addition, a base station's transmitting signal comprises user data, reference- and channel-estimation information, and signaling data, and these elements combine into a composite signal (**Figure 7**).

Furthermore, MIMO requires precise time alignment for transmitting signals. The 3GPP test specifications therefore now include a test to ensure that the signals of two or more antennas are time-synchronized with an accuracy of at least 90 nsec. Once you verify this requirement, you combine the MIMO signal from each base-station antenna on the RF and apply it to a signal analyzer's input.

In receiver testing, users need to apply standards-compliant LTE signals, along with various propagation models, to the base-station receiver path. As a rule, this approach employs PRBSs, and the base station can reconstruct these PRBSs if they are of a known length. By means of a simple comparison, you can obtain the error rates that enable the base-station receiver's performance for verification under a variety of simulated propagation conditions.

Because the receiver tests cover a diversity of interference scenarios and propagation conditions, signal generators must be able to generate reference signals. It can be advantageous if the equipment can flexibly combine signals and if you implement reference channels and specified propagation-channel models, thereby enabling users to quick-

ly and flexibly configure scenarios and greatly simplifying error detection.

In recent years, RRHs have become common features in the design of base stations. In designs of this kind, the RF components and the base-station amplifier reside in a remote front end directly on the antenna. The advantage of this approach is that it avoids line loss on the RF cabling connecting the antenna, thus increasing the base station's available output power. The baseband signals proceed to the RRH over an optical connection. Two digital-interface standards, CPRI and OBSAI, exist for WCDMA base stations. Besides increasing modularity in designs, the standardization of this interface eases error identification during the development cycle. It also allows you to combine modules from different vendors. Furthermore, this setup allows you to connect multiple RRHs to a base station's baseband, which can help to achieve optimum coverage in buildings, for example.

Preparations are currently under way to standardize an optimized interface format for LTE. The digital interface between baseband and RF requires test-and-measurement equipment—in particular, signal generators and signal analyzers—to support this format. This requirement means that you can individually verify baseband or RF modules. As a rule, a converter module translates the measuring equipment's digital baseband language into the standardized format.

To support rapid and flexible resource allocation, the base station assigns the user device a certain channel capacity, including both bandwidth and modulation type, based on various parameters, such as the network cell's available capacity. The base station also acknowledges each packet that it receives correctly; in other words, it tells the user device whether it must retransmit a data packet or can proceed to transmit a new packet. To verify correct operation of the control mechanism in the base station's receiving branch during receiver testing, signal generators must emulate the user device's transmitting signal and provide a means of interpreting the base station's feedback. The feedback proceeds to a separate input on the signal generator; the signal generator then decides in real time whether to request retransmission of the same packet or request transmission of a new packet.

Although LTE is simpler in some ways than its predecessor technology, WCDMA, LTE user equipment is more complex overall because it incorporates additional procedures, such as MIMO, and this complexity places tougher demands on test-and-measurement systems. LTE can adopt a number of types of measurements from the world of WCDMA, but LTE involves new measurements and a wide range of parameterization, on both the RF side and the protocol side. LTE base-station testing uses expanded connection options, for example, to verify fast feedback procedures. These procedures are just as necessary as MIMO signal generation and signal analysis, which vendors implement using multipath approaches. When testing hand-over to earlier technologies, multitechnology platforms offer huge benefits; with a view to the future, these platforms are sound investments.**EDN**

AUTHORS' BIOGRAPHIES

Reiner Goetz is a product manager at Rohde & Schwarz. His responsibilities include marketing and sales of mobile-radio-communication testers. Goetz holds a degree in communications engineering from the University of Applied Sciences (Esslingen, Germany) and a master's degree in business administration in international marketing from the University of Reutlingen (Germany).

Anne Stephan is a project manager in the test-and-measurement division of Rohde & Schwarz, where she is responsible for the development of WCDMA and LTE-signaling software for mobile-radio-communication testers. She holds a degree in computer science from the University of Saarbruecken (Germany).

Meik Kottkamp is technology manager in the test-and-measurement division of Rohde & Schwarz, where he is responsible for strategic marketing and product-portfolio development covering the enhancements in 3GPP technologies. He holds a degree in high-frequency electrical engineering from the University of Hannover (Germany).

BY BRIAN DIPERT • SENIOR TECHNICAL EDITOR

COMING SOON: 3-D TV

THE FACT THAT 3-D IS A HOT TICKET IN TECH MAY NOT SURPRISE YOU. THE FACT THAT IT'S HEADED SOON TO CONSUMERS' HOMES, HOWEVER, MAY SURPRISE—AND DISMAY—CINEMA OWNERS.

AS THOSE of you who've closely followed my online editorial coverage of recent years know, the booming recent interest in 3-D video content is no surprise. It didn't take the impressive success of *Avatar* and other 3-D movies to capture my attention, and I was following the embryonic 3-D industry long before being exposed to the diversity of hardware, software, services, and content at January's CES (Consumer Electronics Show) in Las Vegas. More simply, all it took for me to realize a few years ago that 3-D would be the next big thing was a history lesson.

Black-and-white movies began their transition in the late 1920s from silent films to "talkies." Viewer demand, along with movie studios' desires to enhance content appeal and expand the market, drove this migration. By the end of World War II, however, black-and-white televisions were becoming commonplace in homes, and studio and theater owners alike consequently saw TV as a potential distraction to potential movie viewers' eyeballs and wallets. As such, they competitively accelerated what had previously been a somewhat-leisurely transition from black-and-white to color cinema.

The first NTSC (National Television System Committee) broadcast occurred in late 1953, with standardization at the end of that year. Seeing the writing on the wall, movie studios and directors responded by further upping the ante versus TV. Their competitive response was twofold: Wide-screen films almost immediately gained traction and quickly became commonplace once the industry resolved the disparities between competitive wide-screen technologies. Widespread use of surround sound in cinema is a more recent phenomenon, although it dates from 1940's *Fantasia* (**Reference 1**).

Flash-forward to the present, and high-definition wide-screen displays and high-fidelity surround-sound-audio systems are now pervasive in homes (**Reference 2**). Feeding these technologies are high-quality sound and video content, both residing locally on optical discs and transported to the living room through various wired and wireless broadcast channels (**refer-**

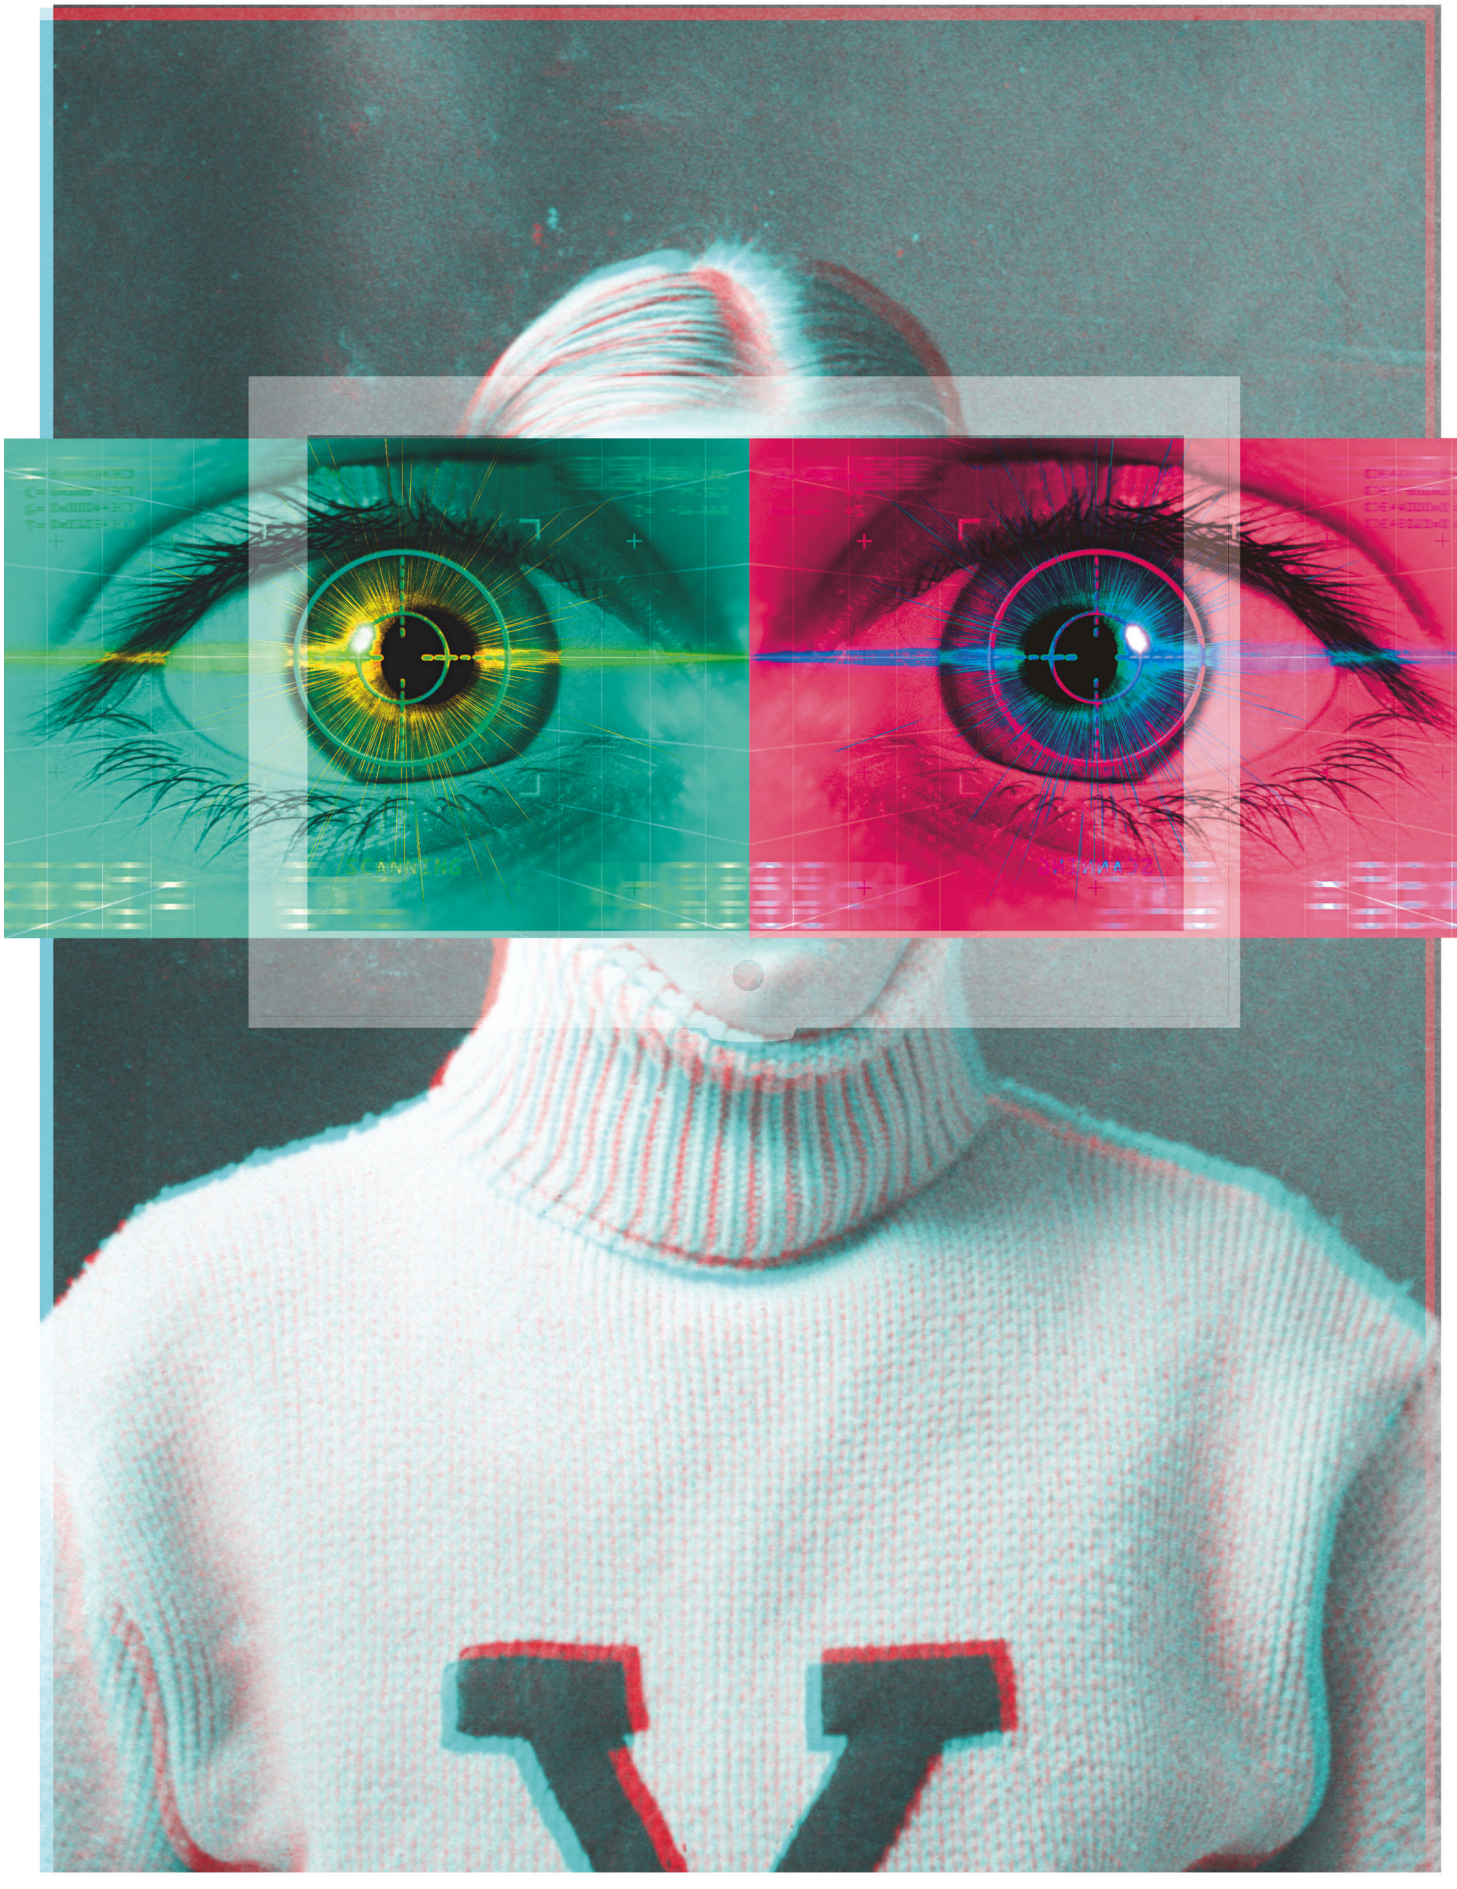

ences 3 and 4). As such, the movie-theater industry has dusted off the other competitive technology it first tried out back in the 1950s: 3-D. Digital cinema is by itself insufficient to ensure continued moviegoer loyalty, in part because the benefits versus the silver-halide predecessor are mostly relevant to the theaters and studios: ease of distribution and accounting, along with no media degradation through repeated showings. But digital technology *does* enable more robust 3-D projection and viewing implementations than the anaglyph—that is, bicolor-lens glasses—approach allowed for. The industry introduced that approach more than a half-century ago and largely discarded it soon afterward (**Reference 5**).

Unfortunately for theater owners, 3-D is coming to living rooms faster than the cinema industry probably had hoped (see **sidebar** "Theater transformations"). Although last year's NTSC-to-ATSC (Advanced Television Systems Committee) conversion in the United States encouraged widespread consumer transitions from standard-definition CRTs (cathode-ray tubes) to high-definition LCDs (liquid-crystal displays) and plasma displays, the consumer-electronics industry's attempts to encourage an evolution in both hardware and content libraries from DVD (digital versatile disc) to Blu-ray disc were less successful (**Reference 6**).

More generally, the last several years' worth of economic downturn has encouraged potential purchasers to keep their wallets in their pockets, to the widespread detriment of the consumer-electronics industry, which now views 3-D as the spark that might reignite consumers' interest and acquisition habits. Ultimately, studios seem loyal only to their investors; a recent dispute over the timing of Walt Disney Co's plans to bring Tim Burton's *Alice in Wonderland* to DVD suggests that Disney and its peers are fundamentally motivated by profit targets, not continued partnership with any link in the historical content-distribution chain (**Reference 7**).

AT A GLANCE

- The latest historical step in the evolution of cinema, 3-D, is one that may quickly sidestep cinemas.
- Various 3-D technologies deliver varying realism results and trade-offs, but all share the same fundamental shortcoming: They require glasses.
- Glasses-free 3-D sounds great in theory, but in reality it is both visually and fiscally unappealing.
- Squeezing two eyes' worth of images through a "straw" designed for only one is a challenge that the industry cannot simply or ideally solve.
- The incremental data payload for 3-D also affects the required storage capacity.

Figure 1 The cover art from a nine-plus-year-old *EDN* article showcases the inexpensive but nonoptimal anaglyph approach to creating 3-D images.

THEORY, IMPLEMENTATION

Although various companies, research laboratories, and academic institutions are investigating 3-D holographic setups like the one that the original *Star Wars* memorably showcased, 2-D displays will constitute the dominant means of viewing 3-D content for years to come. As such, how do you trick a viewer's eyes and ears into extracting a 3-D presumption from a flat-screen presentation? Various 3-D approaches all start from the same premise: Present perspective-corrected views of each frame of a scene to the viewer's left and right eyes, either simultaneously or sequentially and at a sufficiently high rate that the cadence is imperceptible, and then rely on the brain to stitch them together as in real life.

The devil is in the details, however. The anaglyph approach that dates from the 1950s typically relied on red and blue filters, although more modern variants use different patterns (see **sidebar** "Glasses alternatives"). These filters allow the brain to differentiate between right- and left-eye variants of an image within the same frame. Nearly a decade ago, *EDN* published an example of anaglyph 3-D and bundled paper glasses with the issue (**Figure 1** and **Reference 8**).

Anaglyph 3-D is relatively inexpensive to implement, but it suffers from several notable shortcomings that resulted in significant consumer backlash when the industry introduced it. For one thing, the color filters substantially attenuate the amount of light reaching viewers' eyes and degrade the color gamut of the image each eye receives. Second, the technology suffers from image "bleed-through"—that is, the partial presentation of one eye's intended image to the other, and vice versa. This bleed-through distorts the 3-D presentation. Third, the glasses'

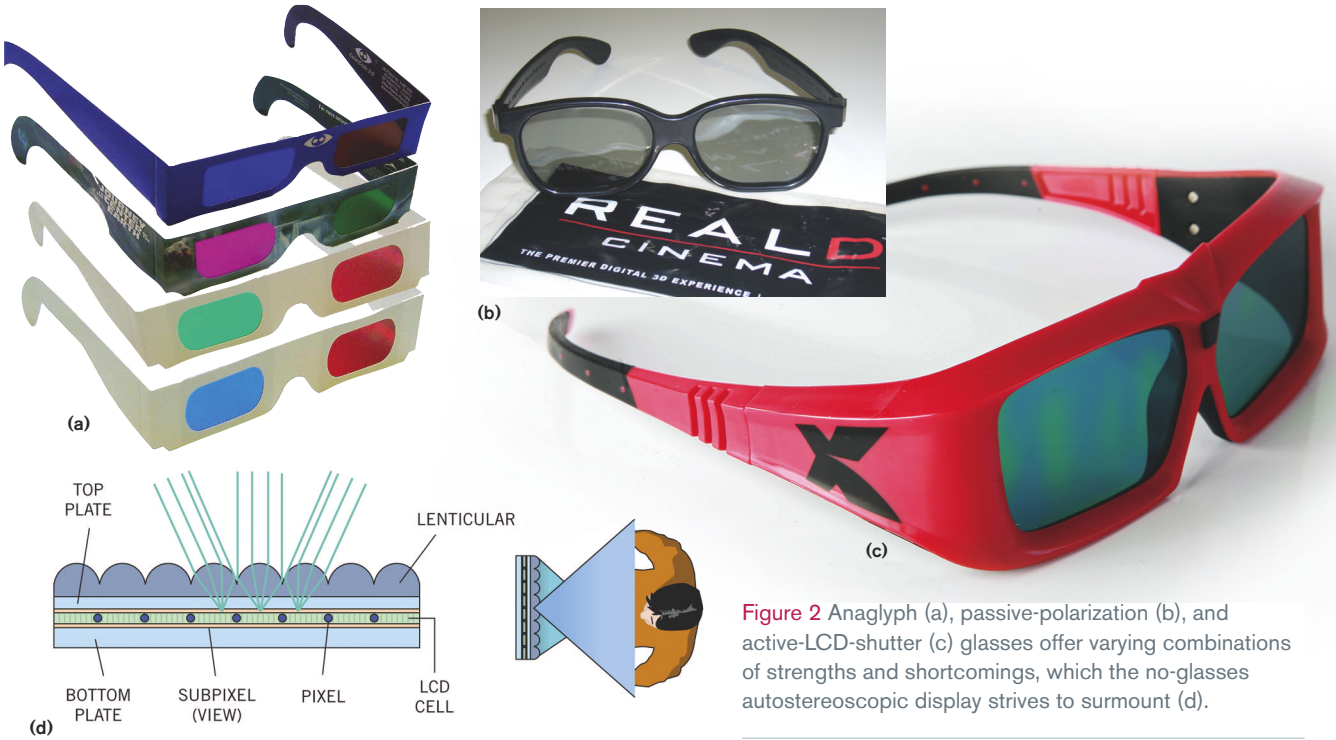

Figure 2 Anaglyph (a), passive-polarization (b), and active-LCD-shutter (c) glasses offer varying combinations of strengths and shortcomings, which the no-glasses autostereoscopic display strives to surmount (d).

dimensions are often incompatible with viewers' head sizes, eye-to-eye spacing, and distance from screen and viewing angles to the screen. These disparities can result in headaches, nausea, dizziness, and other issues.

The industry has since developed several other glasses formats, along with the no-glasses autostereoscopic display (**Figure 2**). One glasses format employs polarization, a variant of the anaglyph approach, with the same luminance-attenuation issue but without anaglyph's chroma-shift problems. One eye's perspective-corrected image leverages light that polarizes differently from that of the image the other eye receives. Matching polarization in the glasses' lenses passively routes the correct image to the correct eye. The perceived success of traditional polarization, as measured by the absence of image "leakage," for example, depended highly on how straight and still viewers held their heads through the presentation. The more modern circular-polarization technique alleviates most of the orientation and immobility requirements. The other now-common glasses approach leverages an LCD shutter in each lens. This technology sequentially projects left- and right-perspective images timed to match the cadence of sequential left and right active passage and blockage of light transmission to each eye.

Passive-polarization systems in theaters can take the form of either a single projector with a precisely paced spinning polarizer disc in front of the projection lens or a sequentially timed dual-projector arrangement. Both cases require the installation of a special "silvered" screen to preserve the projected light's polarization characteristics. Conversely, with active-LCD-shutter glasses, theaters can employ a conventional screen and single-projector setup. This approach commonly uses an infrared beam that comes from the projector, bounces off the screen, and floods the audience to control the switching rate of viewers'

TABLE 1 FULL-FRAME BANDWIDTH REQUIREMENTS OF 3-D VIDEO FORMATS (GBPS)

	24-bit (8 bits/pixel)	**30-bit** (10 bits/pixel)	**36-bit** (12 bits/pixel)
720p24	1.1	1.3	1.6
720p30	1.3	1.7	2
720p60	2.7	3.3	0.3
1080i60	3	3.7	4.5
1080p60	6	7.5	9

THEATER TRANSFORMATIONS

If 3-D technology does rapidly invade the home, it may quickly obviate any meaningful differentiation between movie theaters and living-room theaters. This step does not necessarily herald the death of cinemas, however. Look, for example, at the largely successful live-broadcast trials of concerts, sports venues, and other events to theaters in several cities, states, and countries. Most of these presentations have been in 2-D, but there's no reason that they couldn't quickly migrate to 3-D given the infrastructure in many venues. The 2012 Summer Olympics, the next Super Bowl, or even a gig by a band that otherwise wouldn't come to my home town presented live in large-screen 3-D and surround sound: I'd buy that.

glasses. However, active LCD glasses are substantially more expensive and bulky than their passive-polarizer counterparts, and theaters must regularly recharge their embedded batteries. All these factors necessitate their collection and cleaning before distributing them for reuse, and some consumers voice concerns about sanitation.

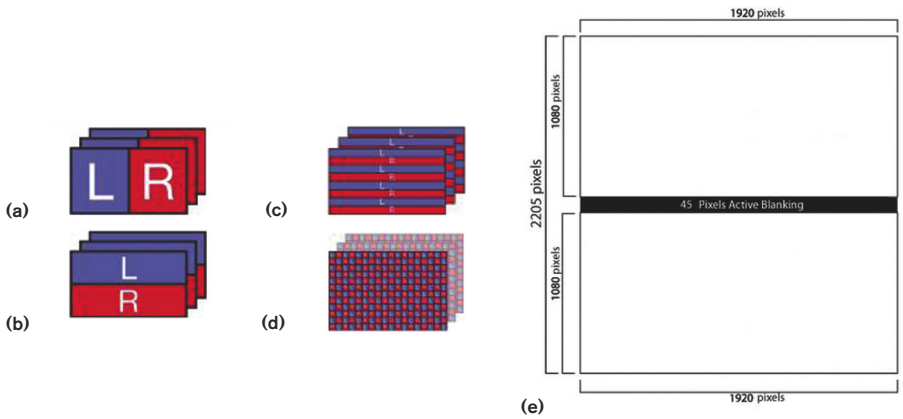

Figure 3 One approach to squeezing two eyes' worth of information into the bandwidth previously tailored for one frame of data is to diminish the resolution of each eye's image using a variety of pixel-combining techniques: side by side (a), over and under (b), line by line (c), and checkerboard (d). Even if there's sufficient payload available for two full frames' worth of pixels per 3-D image, there's no guarantee that hardware down the chain will accept it (e).

DISPLAY CONTENDERS

Migrate 3-D from the movie theater to the home theater, and your implementation options radically expand. Your customers could, of course, mimic a theater configuration using a single-projector or multiprojector arrangement employing DLP (digital-light projection), LCD, or LCOS (liquid-crystal-on-silicon) technology, but such setups are largely restricted to videophiles. Because modern LCD and plasma direct-view televisions switch and refresh fast enough for stutter-free playback, you can alternatively employ active-shutter glasses, timed through an infrared, RF (radio-frequency), or wired connection to the display, to synchronize with a sequentially displayed right- and left-eye-intended version of each frame. Alternatively, some LCDs leverage passive polarization, in which a polarizer filter lies directly atop each consecutive display line. This filter tailors that line's image for one of the viewer's eyes. The chief downside of the passive approach is that it halves the effective vertical resolution that each eye perceives. However, you can usually disable the polarization effect for use when viewing 2-D content or 3-D content that the display has dynamically converted to 2-D for viewing without using glasses.

Any glasses-based 3-D technology, however, has several notable shortcomings. Chief among them is the sizing concern; any mismatch between the glasses' and viewer's facial dimensions will lead to discomfort or worse. There's also the potential for breakage or misplacement and, therefore, the need for replacement of glasses, a concern to consumers but perhaps a tempting opportunity for supplier profit, particularly with the comparatively expensive active-shutter approach. Also, LCD glasses' need for periodic recharging can lead to frustration when potential customers must reschedule movie night because of the glasses' drained batteries or failure in the middle of a movie. Finally, there's the compatibility worry, a likely scenario in the early days of any new technology. Isn't it reasonable to assume that consumers will be reluctant to invest in glasses that they can't use at the homes of friends, family members, and neighbors or a 3-D technology that an upstart alternative might render obsolete?

Figure 4 The 3-D gaming kits tether glasses to computers using wired or wireless infrared or RF connectivity.

Taking a no-glasses tack at solving the problem, a number of manufacturers have developed autostereoscopic displays. These displays incorporate lenticular lenses, parallax barriers, or other mechanisms to create depth perception from a flat projection surface. To succeed to a reasonable degree, however, autostereoscopy requires that the viewer be rigidly positioned in a "sweet spot" throughout the presentation. Even under ideal circumstances, autostereoscopy doesn't create a compelling end result. I've auditioned many autostereoscopic displays over the years, and I've never walked away even remotely impressed. Fortunately, users can switch autostereoscopic displays into 2-D mode to view conventional content. Like with other

specialty-display types, such as large-screen OLEDs (organic light-emitting diodes), it's nonetheless difficult to envision that autostereoscopy can achieve sufficient early-adopter sales to appreciably reduce costs and prices for the masses.

DISTRIBUTION CHALLENGES

How can you transport discrete right- and left-eye versions of each video frame's information through a wired or wireless "pipe" that was originally bandwidth-tuned for single 2-D frame transport? In short, you can't. This pragmatic reality means that trade-offs will be necessary to accomplish a 3-D presentation. Two bandwidth-slimming possibilities are to lower the per-pixel color depth or the playback frame rate. Take HDMI (high-definition multimedia interface), for example (**Reference 9**). Now-pervasive HDMI Version 1.3 has 340 MHz of bandwidth—more than double the single-link bandwidth of its HDMI Version 1.2 predecessor. This bandwidth speedup translates to 10.2-Gbps TMDS (transition-minimized-differential-signaling) bandwidth, or 8.16-Gbps video bandwidth, which is more meaningful to the application. This bit rate is adequate for passing a 3-D variant of a full-frame 720p30, 1080i60, or even 1080p60 video presentation at 24-bit-per-pixel color (**Table 1**). Higher-color-depth, higher-frame-rate, or higher-resolution 3-D video clips, however, could find HDMI 1.3 or HDMI Version 1.4, which has the same speed as its predecessor, lacking from a bandwidth standpoint. Such a demanding data payload might come, for example, from a game running on a computer or an arcade console (**Reference 10**).

Lower-bandwidth connections, such as broadband-Internet WAN (wide-area network), wired and wireless LAN (local-area network), and ATSC-broadcast beacons, have even more challenged 3-D capabilities, despite the fact that low-bit-rate, lossy compression algorithms usually find use as their video codecs (**Reference 11**). As such, per-frame resolution reductions are often necessary so that two frames' worth of information, corresponding to the right- and

GLASSES ALTERNATIVES

Anaglyph, active-LCD-shutter, and passive-polarizer 3-D glasses may be the most prevalent options available today, but they're not the only candidates. Keep an eye on Dolby 3-D, which blends the attributes of the anaglyph and passive-polarizer approaches. Like both of these traditional approaches, the Dolby 3-D glasses are relatively lightweight and inexpensive.

Like anaglyph, Dolby 3-D relies on fine-grained-color-spectrum segmentation to differentiate between right- and left-eye-intended images. Unlike anaglyph and like passive polarization, however, each eye receives a full-spectrum presentation. And unlike passive polarization, a theater retrofit to install an expensive silvered screen is not necessary.

Chromatek's ChromaDepth, a modern descendant of anaglyph, leverages a microprism-based film that attaches to the glasses' lenses to alter perceived objects' 3-D depth based on their color. Its major shortcoming is the fixed colors that objects can display: red objects in front and blue in back, with varying depth between them matching the visible light spectrum.

Implementers of 3-D technology may also consider the Pulfrich effect, a human-visual-system oddity in which eyes perceive side-by-side motion as having depth when accompanied by an eye-to-eye-synchronization lag. Glasses manufacturers commonly implement this effect by placing a darker lens over one eye, so when something moves from left to right, it will look as if it's moving back or forward—in 3-D. Coca-Cola used this technique in a commercial that aired at halftime of Super Bowl XXIII, and reportedly distributed 40 million pairs of Pulfrich glasses before the broadcast.

Figure 5 Panasonic plans to produce a 3-D camera targeting budding professional videographers (a); Fujifilm's hardware has humbler aspirations and a lower price tag (b).

2-D CONVERSIONS

Hollywood has enthusiastic plans for rapidly transitioning upcoming movie releases to 3-D-optional or even 3-D-only formats. Such material will inherently be 3-D-ready when it comes to the home through various physical media and streamed distribution channels. For 3-D to gain rapid acceptance by consumers, however, some robust means must exist for converting a meaningful percentage of legacy 2-D content to the new format. This scenario is analogous to how, under the tenure of media mogul Ted Turner, the industry "colorized" legacy black-and-white films. For consumer-graphics-heavy content, the process is relatively straightforward: Re-render the geometry data for two eyes' worth of perspective rather than one. The makers of the 3-D version of *The Polar Express* used this approach (Reference A).

Conversely, with traditionally captured images, the 2-D-to-3-D conversion is far more involved, and the results are far less predictable in their realism. With the 3-D remake of Tim Burton's *The Nightmare Before Christmas*, for example, which the director created using stop-motion 2-D image capture of miniature physical models, the studio did a frame-by-frame digitization of the entire film (Reference B). The studio then used heavy-duty computer and graphics-processor horsepower to morph each frame into its opposite-eye counterpart. Consider that Toshiba showcased a simplified variant of the same process at January's Consumer Electronics Show in real time in the company's TVs to transform 2-D content into a pseudo-3-D variant. You'll appreciate the capabilities of the embedded Cell CPU that executed the algorithms.

REFERENCES

A Dipert, Brian, "3-D: nifty," *EDN*, Feb 22, 2006, www.edn.com/blog/400000040/post/1550002755.html.

B Dipert, Brian, "3-D stop motion: well-deserved promotion," *EDN*, Oct 31, 2007, www.edn.com/blog/400000040/post/590016659.html.

left-eye data, can squeeze into the transport space that one frame's worth of data formerly used. The resolution-reduction techniques take different approaches to optimizing the trade-off between resultant image quality and processing complexity (**Figure 3**).

The side-by-side scheme requires that the video processor in the display or projector subsequently horizontally expand each eye's image to fill the full frame size, whereas the "over-and-under," "above-and-below," or "top-and-bottom" approach requires subsequent vertical scaling. Adding support for the over-and-under approach is the impetus for the HDMI Version 1.4a specification, which is now publicly available. Alternatively, the technology can scatter the right and left eye's data across the source frame in a line-by-line or checkerboard pattern. Keep in mind that further alternation may be necessary to tailor the material to the target display technology once the video data traverses the interconnection to the destination playback device.

Rather than rapidly presenting the right and left eye's information in frame sequence and synchronizing it to active-LCD-shutter glasses, as projectors, plasma displays, and some LCD TVs do, other LCD TVs "stripe" the two eyes' data within a single display frame. Even if full-frame 3-D playback is possible over the transmission channel, the destination device may be unable to accept it without a firmware upgrade. The 1920×2205-pixel frame size of the full over-and-under 3-D implementation, for example, is incompatible with the EDID (extended-display-identification data) in almost all currently installed displays, as well as that in intermediary A/V (audio/video) receivers and HDMI switchboxes (**Reference 12**).

CAPTURE AND STORAGE

The same data-payload constraints that may hamper transmission channels in 2-D-to-3-D conversion also potentially have an impact on the storage devices that archive the video information. Storing the left and right eyes' per-frame information in full frame would require more than double the capacity of 2-D technology. The Blu-ray disc Association last December announced its support for the 3-D video format. Ironically, 3-D may finally provide sufficient justification for Blu-ray's increased capacity over its DVD predecessor because a multilayer DVD in combination with an advanced video codec, such as H.264 or VC-1, provides sufficient capacity to hold a full-length Hollywood feature film in a 2-D, high-definition format (**Reference 13**).

Speaking of film, how do moviemakers create 3-D presentations? With full computer-graphics animation sequences, the process is relatively straightforward, involving rendering distinct versions of each frame's geometry data from the right and left eyes' perspectives (see **sidebar** "2-D conversions"). The algorithms can even sometimes run in real time on a modern graphics processor, as the glasses-plus-board retail kits employing AMD/ATI and Nvidia chips demonstrate (**Figure 4**).

More traditional video-image capture requires a dual-lens setup, often with dual sensors and dual storage devices. Panasonic showed such a videocamera for professional videographers at January's CES. The $21,000 AG-3DA1, which the company had previously unveiled at the April 2009 NAB (National Association of Broadcasters) show, will become available for sale this fall. Nvidia's press briefing at the same CES showcased Fujifilm's more moderately priced and featured FinePix Real 3-D W13-D videocamera (**Figure 5**). It's probably no surprise that large consumer-electronics companies, such as dominant Blu-ray backer Sony, have plans for 3-D-supportive still cameras and videocameras, and CES demonstrations from even entry-level imaging manufacturers suggest that the technology will rapidly and pervasively make its way into the marketplace. Some degree of initial incompatibility between capture sources, playback destinations, and intermediary devices is inevitable until industry and de facto standards take hold, but the long-term future for 3-D looks realistic. EDN

REFERENCES

1 Dipert, Brian, "Expanding options bring surround sound to the forefront … and the back … and the ceiling … and the floor …" *EDN*, Jan 10, 2002, pg 34, www.edn.com/article/CA189468.

2 Dipert, Brian, "Master of some: direct-view-display technology," *EDN*, March 3, 2005, pg 38, www.edn.com/article/CA505067.

3 Dipert, Brian, "CES 2009: Blu-ray's continued struggles and the ramping ascendancy of online," *EDN*, Jan 9, 2009, www.edn.com/article/CA6627924.

4 Dipert, Brian, "Transporting high-def video broadcasts: Are wireless networks up to the task?" *EDN*, Aug 20, 2009, pg 24, www.edn.com/article/CA6676165.

5 Dipert, Brian, "Display enhancements accept no compromises," *EDN*, Dec 7, 2000, pg 47, www.edn.com/article/CA56164.

6 Dipert, Brian, "Thin air: ATSC reception isn't always easy," *EDN*, May 14, 2009, pg 20, www.edn.com/article/CA6656302.

7 Cheng, Jacqui, "Theater chain pushes back at Disney over early DVD release," *Ars Technica*, Feb 25, 2010, www.arstechnica.com/media/news/2010/02/theater-chain-pushes-back-on-disney-over-early-dvd-release.ars.

8 Dipert, Brian, "Balancing in three dimensions," *EDN*, April 27, 2000, pg 54, www.edn.com/article/CA82388.

9 Dipert, Brian, "Connecting systems to displays with DVI, HDMI, DisplayPort: What we got here is failure to communicate," *EDN*, Jan 4, 2007, pg 46, www.edn.com/article/CA6402885.

10 Dipert, Brian, "Got game? Living-room consoles grapple for consumers' eyes, wallets," *EDN*, Dec 12, 2005, pg 51, www.edn.com/article/CA6290451.

11 Dipert, Brian, "Video characterization creates hands-on headaches," *EDN*, July 25, 2002, pg 53, www.edn.com/article/CA233723.

12 Fremer, Michael, Editor, "3D HDTV and HDMI explained," HD Guru, Feb 22, 2010, www.hdguru.com/3D-hdtv-and-hdmi-explained/1336.

13 Dipert, Brian, "Advanced video codecs: the key to continued red laser-based optical storage dominance?" *EDN*, June 17, 2005, www.edn.com/blog/400000040/post/1190000919.html.

See the "Coming Soon" posts at www.edn.com/briansbrain for supplemental information.

FOR MORE INFORMATION

Advanced Micro Devices
www.amd.com

Chromatek
www.chromatek.com

Coca-Cola
www.coca-cola.com

Dolby Laboratories
www.dolby.com

Fujifilm
www.fujifilm.com

Nvidia
www.nvidia.com

Panasonic
www.panasonic.com

Samsung
www.samsung.com

Sony
www.sony.com

Toshiba
www.toshiba.com

Walt Disney Co
www.disney.com

You can reach
Senior Technical Editor
Brian Dipert
at 1-916-760-0159,
bdipert@edn.com,
and www.bdipert.com.

COMPLEX MODULATION COMES TO OPTICAL FIBER

FIBER IS RUNNING OUT OF BANDWIDTH JUST AS DIAL-UP LINES DID YEARS AGO. COMPLEX MODULATION AGAIN SOLVES THE PROBLEM.

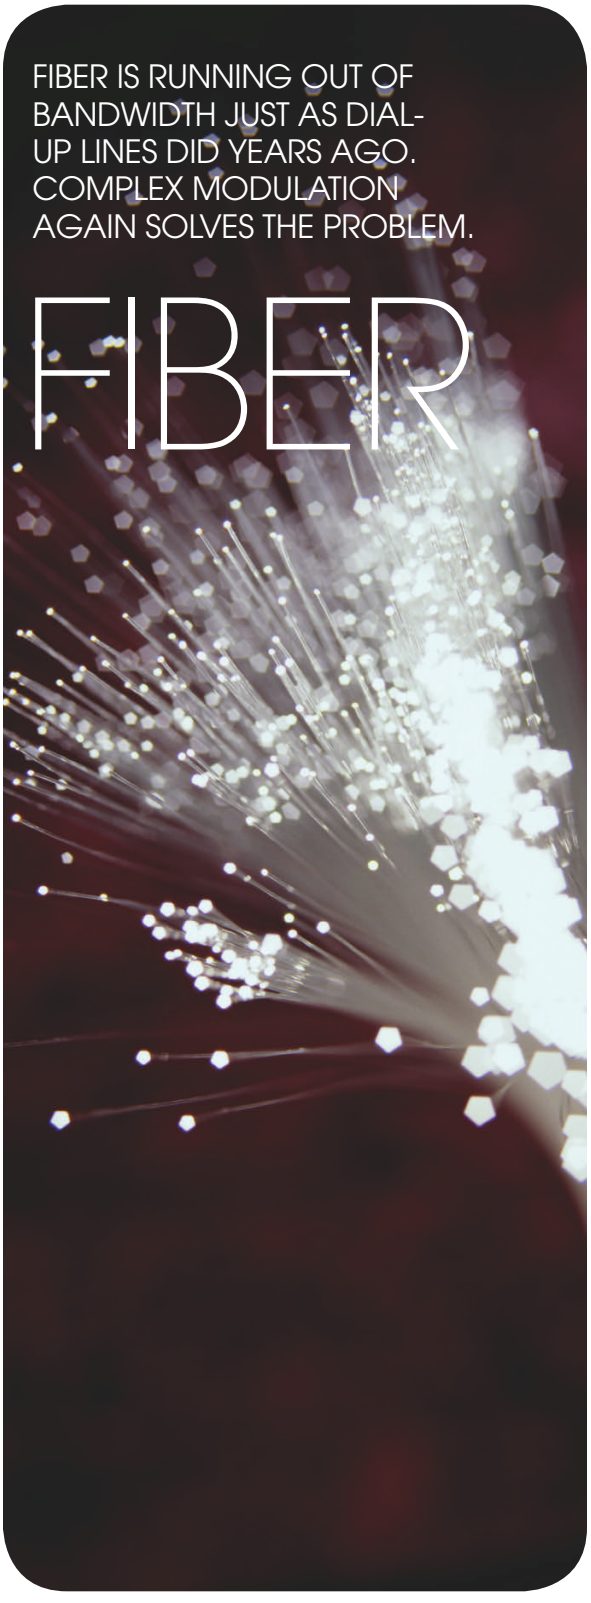

BY MARTIN ROWE • *TEST & MEASUREMENT WORLD*

The demand for greater data throughput seems endless, and it is accelerating faster than many people expected, creating bottlenecks in fiber-optics networks. Digital transmissions of 100 Gbps, which companies are just now introducing, should alleviate some of these bottlenecks. A year ago, most of the work surrounding 100-Gbps links started with 10 10-Gbps lanes over short distances. Since then, Verizon has deployed the first long-haul 100-Gbps link using four 25-Gbps lanes (**Reference 1**). With it comes complex modulation that optical communications have never before used.

The new modulation schemes are necessary for handling long-distance transmissions. Short-haul communications, the so-called client side within campuses and local metropolitan areas, don't need complex modulation because their distances are short enough to accommodate the higher speeds (**Figure 1**). On the client side, with distances as great as 40 km, 100-Gbps links can use four 25-Gbps lanes. IEEE 802.3ba defines these data links (**Reference 2**). Because short-haul 100-Gbps links use four wavelengths on a single fiber or even 10 10-Gbps fibers over the shortest distances, more fiber may be necessary to increase over the current 10-Gbps speed. Installing additional fiber over the short distances between buildings on a campus isn't expensive.

This is not the case for long-haul transmissions—the "line side" of networks for which service providers need transmissions of hundreds of kilometers. Adding fiber to compensate for additional lanes is just too expensive. "Carriers need to squeeze 100-Gbps throughput rates into their existing fiber plants, many of which were designed for 10 Gbps and some of which were designed for 2.5-Gbps fiber links," says Pavel Zivny, a product engineer at Tektronix (www.tek.com).

Simply squeezing a 100-Gbps NRZ (nonreturn-to-zero) stream into existing fiber is impractical. Current DWDM (dense-wavelength-division-multiplexing) fibers use 50-GHz spacing between channels. Although that

channel spacing is sufficient for 10-Gbps data streams using NRZ modulation, it is too narrow for 100-Gbps NRZ streams. "You can't put 100-Gbps streams right on the carrier," says Mike Schnecker, business-development manager at LeCroy (www.lecroy.com), because, for a 100-Gbps NRZ signal, each bit is just 10 psec wide.

"Because of crosstalk between adjacent channels, 100-Gbps data streams can't be used in DWDM systems," says Hiroshi Goto, an optical-product specialist at Anritsu (www.anritsu.com). "PMD [polarization-mode dispersion] and CD [chromatic dispersion] prevent that [scenario]. There's too much distortion. The pulses distort and overlap."

To work around the problem, the OIF (Optical Internetworking Forum, www.oiforum.com) has recommended using complex modulation to squeeze more bits per second per hertz from existing fiber. The OIF-proposed modulation uses QPSK (quadrature-phase-shift keying) and two polarizations to achieve 100-Gbps throughput on a single wavelength. QPSK is common in digital RF communications, but it's new to fiber-optics communications.

AT A GLANCE

- Carriers must squeeze 100 Gbps into their fiber plants, which were designed for 10- or even 2.5-Gbps fiber links.
- DWDM (dense-wavelength-division-multiplexing) systems can't use 100-Gbps data streams because of crosstalk between adjacent channels.
- QPSK (quadrature-phase-shift-keying) signals are more susceptible to noise and nonlinear phase distortion than NRZ (nonreturn-to-zero) signals.

A 100-Gbps link consists of two 50-Gbps streams in two polarizations—TE (transverse electric) and TM (transverse magnetic)—that propagate in two orthogonal polarization planes. Each 50-Gbps stream consists of 25G symbols/sec. QPSK modulation packs 2 bits into one symbol. Because the QPSK signal travels in two polarizations, it is called either DP-QPSK (dual-polarization QPSK) or PM-QPSK (polarization-mode QPSK); the terms are interchangeable, and both are commonly used. This article uses DP-QPSK when referring to the two polarizations and QPSK when referring to one polarization.

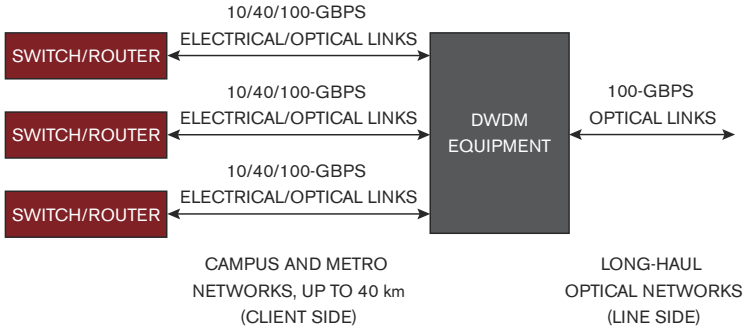

Figure 1 Carriers use the line side of a data link for long-haul transmissions between cities. Client-side transmissions link campuses and local metropolitan areas.

COMPLEX MODULATION

Figure 2 illustrates the modulation process. A single 100-Gbps bit stream splits into TE and TM polarizations. That step produces two carriers at the same frequency. Each carrier then undergoes I/Q (in-phase/quadrature) modulation, resulting in two 25G-symbol/sec streams. The total is 100 Gbps, but the actual data rate is somewhat higher (see **sidebar** "What's in a G?"). The polarization splitter in **Figure 2** appears before the QPSK modulators. Some transceiver designs may place the I/Q modulators first and then split the modulated signals into two polarizations.

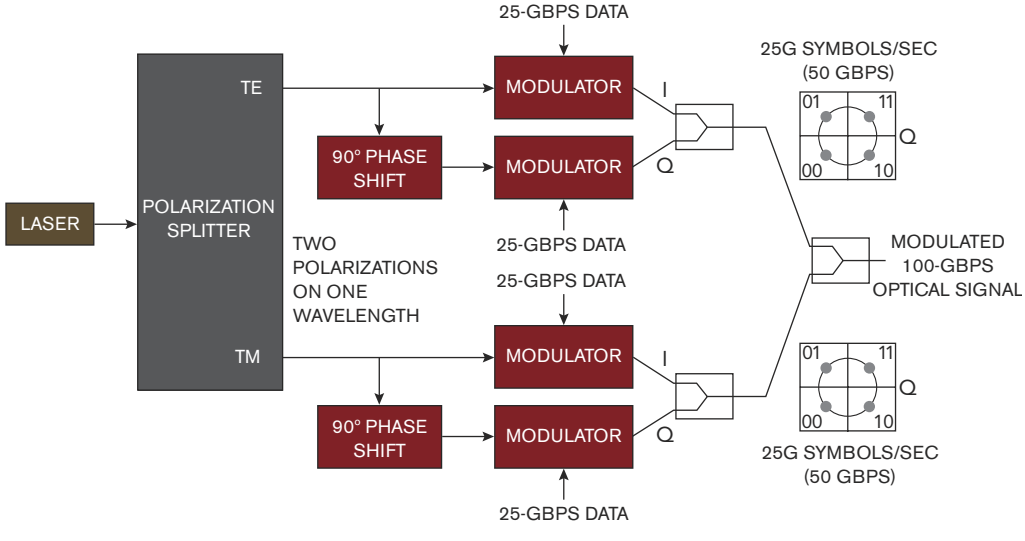

Figure 2 A 100-Gbps transmitter splits a laser into two polarizations and then modulates four 25-Gbps data streams onto a single fiber at a single wavelength.

QPSK modulation places 2 bits per symbol by phase-shifting a carrier of light in response to incoming bit pairs (00, 01, 10, 11). Each symbol represents 2 bits. A receiver demodulates each symbol into its 2 bits and produces a 50-Gbps digital data stream. In addition, bits undergo precoding before modulation and decoding after modulation (**Reference 3**). A receiver then produces four 25-Gbps electrical

signals after it demodulates and decodes the incoming DP-QPSK signal.

QPSK signals carry twice the number of bits per symbol that NRZ signals carry. Thus, the two modulations produce signals that degrade differently as they pass through fiber. Peter Andrekson, director of EXFO Sweden (www.exfo.com), explains that QPSK signals are more susceptible to noise and nonlinear phase distortion than NRZ signals. "Because of the higher noise susceptibility, QPSK-modulated signals will require higher power than NRZ signals," he says.

QPSK signals have an important advantage over NRZ signals, though. They're less susceptible to bit errors from chromatic dispersion and group delay at the same bit rate. That's because one UI (unit interval) of a 100-Gbps data stream is 10 psec wide. Because line-side transmissions use four 25-Gbps lanes, each symbol is 40 psec wide, which results in a lower bandwidth.

The 40-psec-wide symbol of a 25G-symbol/sec stream is shorter and requires more bandwidth than a 10-Gbps, 100-psec-wide NRZ signal. Thus, the 25G-symbol/sec signal is more susceptible to errors from dispersion than a 10-Gbps NRZ signal, but it's less susceptible to degradation than a 100-Gbps NRZ signal. "There is a trade-off between complexity and SNR [signal-to-noise ratio] versus dispersion tolerance and hardware bandwidth at a given bit rate," explains Andrekson.

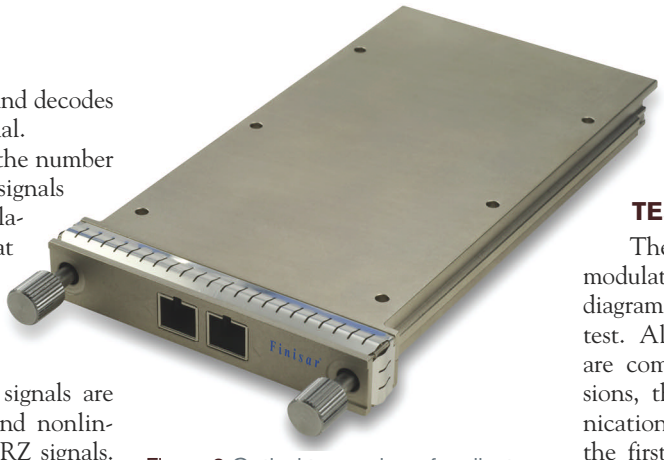

Figure 3 Optical transceivers for client-side transmission are based on the CFP multisource agreement for size and electrical connections to a line card (courtesy Finisar).

The DP-QPSK technology is so new that no transceiver modules exist for the line side. Chris Cole, senior member of the technical staff at Finisar (www.finisar.com), explains that line-side transceiver modules are larger than client-side modules (**Figure 3**), which a multisource agreement currently defines (**Reference 4**). Cole notes that designers can even implement line-side transceivers as line cards rather than as modules.

TEST WILL CHANGE, TOO

The shift from NRZ to DP-QPSK modulation brings the constellation diagram to the forefront of fiber-optics test. Although constellation diagrams are common in RF wireless transmissions, they're new to optical communications. Constellation diagrams are the first measurement you make on a QPSK transmission. Constellation diagrams provide information about the transmitted signal's integrity. Dispersion and nonlinearities can cause signal degradation, resulting in distortion. **Figure 4** shows constellation diagrams for both polarizations in a DP-QPSK signal. The constellation's points are clearly visible in **Figure 4**, but they can become indistinguishable in the presence of too much distortion.

The two lower-right traces in **Figure 4** show the QPSK-modulated signal's magnitude (upper trace) and phase (lower trace). Note the apparent discontinuities on the phase-angle diagram. They result from phase shifts due to the encoding of bit pairs in the QPSK modulation.

For testing the optical DP-QPSK signal, you can use an optical-modulation or an optical-signal analyzer. These instruments produce constellation diagrams, decode them into electrical data streams, and display them as eye diagrams. Agilent Technologies (www.agilent.com), Anritsu, EXFO, and Optametra (www.optametra.com) serve this market, and Optametra's product employs a Tektronix oscilloscope.

"There's no test specification for the 100-Gbps long-haul optical waveform, so test-equipment makers must talk to the optical-module makers to find out what they need to measure," says Finisar's Cole. "Each company will have different needs." Cole also notes that test equipment must support 28G- and 32G-symbol/sec signals. "There are DP-QPSK test systems that run at 22G symbols/sec for 40-Gbps links, but new equipment will need to run at 28G and 32G symbols/sec to support 100-Gbps links."

Testing the receiver side of optical transceivers is even more up in the air

WHAT'S IN A G?

The terms "100G," "40G," and "25G" refer to the data throughput of an optical link. Because of formatting and FEC (forward-error correction), actual data rates are higher than the numbers indicate.

For example, the data rate for a 100-GbE (Gigabit Ethernet) transmission is actually 103.125 Gbps, but the data throughput is 25 Gbps. Therefore, each 25-Gbps lane actually carries 25.78125 Gbps (26 Gbps) for the client side. So, if a test-equipment manufacturer claims that its products have 26-Gbps speed, it means that the product covers the 25.78125-Gbps data rate. A 27.739-Gbps (28 Gbps) data rate is also under consideration for Ethernet client-side networks.

For line-side networks, long-haul transmissions, links need additional FEC. The line rate for 100-Gbps links with 7% FEC is about 112 Gbps, which translates to about 28 Gbps on each lane. According to the Optical Internetworking Forum's "100G Ultra Long Haul DWDM Framework Document," the exact rate has not yet been specified (Reference A). These transmissions can also use a higher FEC overhead of 20%, which increases the bit rate to about 32 Gbps.

REFERENCE

A "100G Ultra Long Haul DWDM Framework Document," Optical Internetworking Forum, www.oiforum.com/public/documents/OIF-FD-100G-DWDM-01.0.pdf.

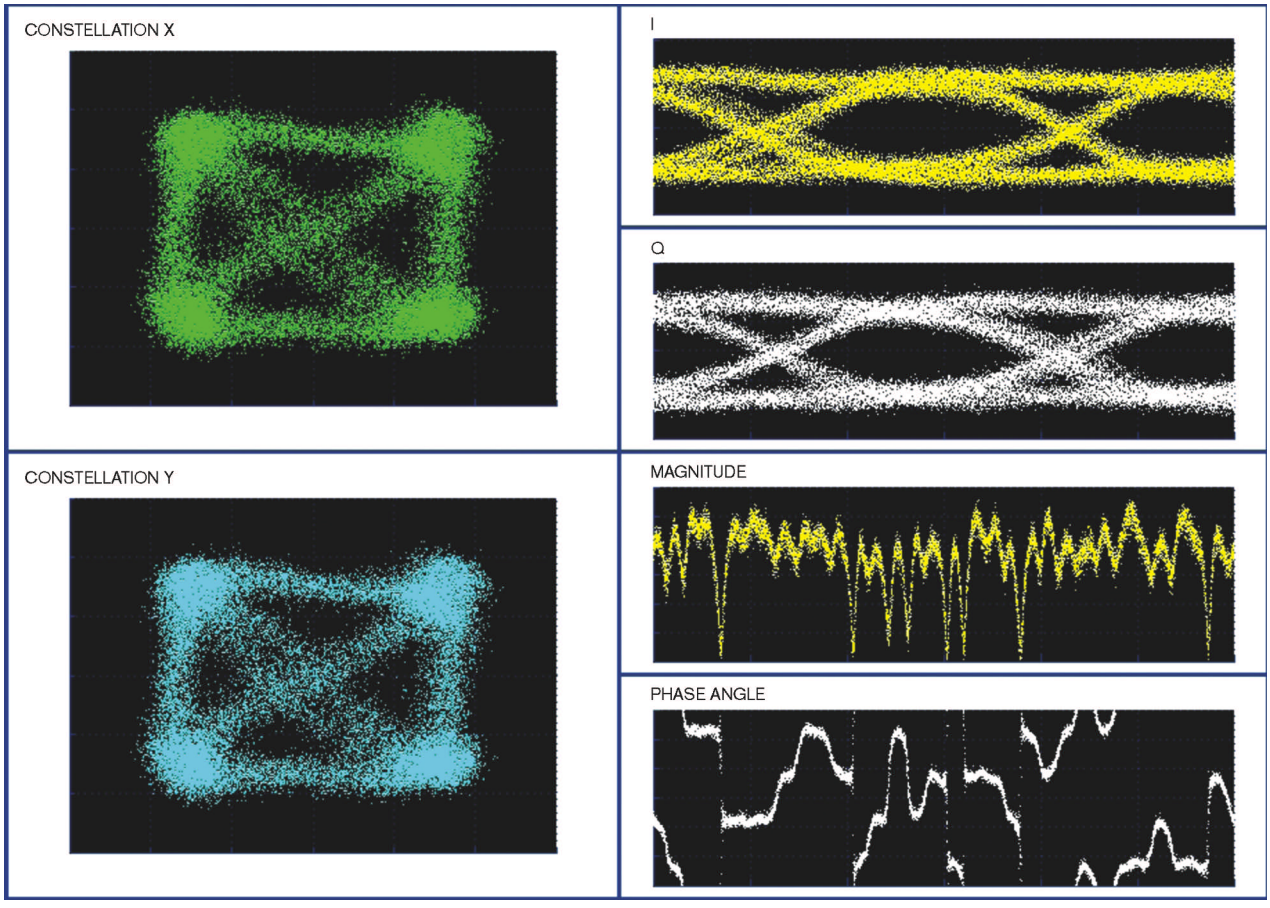

Figure 4 Constellation diagrams will become a mainstream tool for analyzing DP-QPSK modulated signals (courtesy EXFO).

because specifications do not yet exist for stressed-receiver testing. Cole says that test equipment must be able to generate DP-QPSK signals, and that requirement can introduce controlled impairments, such as chromatic dispersion and polarization-mode dispersion. These impairments cause the TE and TM carriers to rotate as they pass through fiber. The impairments must produce stressed eye patterns after demodulating and decoding so that engineers can measure the signals once they're in electrical form.

Figure 4 also shows the two eye diagrams (upper right) representing two 25-Gbps lanes from one polarization. "You'll have to look at eye-mask margins, jitter, and extinction ratio; that's the same as for 10-Gbps links," says Cole.

Engineers now use oscilloscopes and BER (bit-error-rate) testers to analyze eye diagrams. Some engineers use high-bandwidth oscilloscopes to capture DP-QPSK signals. "Because of the modulation, signals at the receiver look like noise," says LeCroy's Schnecker. "Signals are no longer repetitive, and thus you need a real-time oscilloscope." Zivny of Tektronix has also worked with engineers using real-time oscilloscopes on DP-QPSK signals. A four-channel oscilloscope lets you see all four decoded, demodulated data streams with high timebase correlation.

Engineers developing DP-QPSK transceivers use BER testers to produce the 25-Gbps data streams for each I and Q phase of a QPSK signal. They also use BER testers to measure BER on the demodulated, decoded signals. BER testers from Agilent Technologies and SyntheSys Research (bertscope.com) can measure BER at data rates as high as 28 Gbps.

Over the next few years, the industry will continue to develop 100-Gbps line-side transmissions. Test specifications will also emerge as optical-module manufacturers work with test-equipment makers and standards bodies to identify test issues and to develop test procedures and equipment. EDN

REFERENCES

1 "Verizon Deploys Commercial 100G Ultra-Long-Haul Optical System on Portion of Its Core European Network," Verizon, Dec 14, 2009, newscenter.verizon.com/press-releases/verizon/2009/verizon-deploys-commercial.html.
2 Di Minico, Chris, "802.3ba Cu specifications," Jan 2008, IEEE, grouper.ieee.org/groups/802/3/ba/public/jan08/diminico_01_0108.pdf.
3 "Trends and Issues in Ultra-High-Speed Transmission Technologies: For the implementation of 100 GbE/40 GbE and long haul transmission by optical modulation," Technical Note, Anritsu, 2009, www.eu.anritsu.com/files/MP1800A_40_100GbE_EE1200.pdf.
4 "CFP Multi-Source Agreement Draft 1.0," March 23, 2009, www.cfp-msa.org/Documents/CFP-MSA-DRAFT-rev-1-0.pdf.

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

BY BERNARD HYLAND • MAXIM INTEGRATED PRODUCTS

Measure propagation delays using time-domain reflectometry

TDR BEATS ACTIVE PROBES FOR HIGH-SPEED DELAY MEASUREMENTS.

You normally use TDR (time-domain reflectometry) to measure impedance change along a signal path (**Reference 1**). It is also a valuable tool for measuring propagation delays. The TDR technique is applicable to any high-speed circuit. You can use the propagation delay of a high-speed pin-electronics IC in ATE (automatic test equipment) to perform these measurements. These ICs contain high-speed drivers, active loads, and window comparators that operate in excess of 1 Gbps.

To perform TDR, you propagate a fast edge down a signal path and observe the reflection. The reflection shows the impedance along that signal path, as well as the delay that each change in the impedance imposes (**Figure 1**). In this case, the T_{DLY} is the delay of the PCB (printed-circuit board) run you are measuring, and Z_O is the impedance of the PCB run. Using TDR eliminates direct probing of the circuit, which is a difficult procedure because it entails placing probes on the device pins. These probes become part of the high-speed signal path and distort the signal you want to measure. Even a high-impedance active probe can load your circuit.

Rather than using active probes, you can use the TDR measurement capabilities of a Tektronix (www.tek.com) TDS8000-series oscilloscope with a model 80E04 TDR sampling module (**Figure 2**). The sampling heads have a 20-GHz bandwidth. You also use an Agilent/HP (www.agilent.com) 8082A pulse generator. You can use an evaluation board from the IC manufacturer for the DUT (device under test, **Figure 3**). The high-speed inputs to this board are DATA1 and NDATA1.

Using this setup, you make several measurements. You measure the delay due to the SMA (subminiature Type A) connectors J14 and J13 and the PCB runs under the heat sink. You measure the delay from the output of the IC through SMA connector J18. You measure the delay in the test cable connecting the DUT1 output to the oscilloscope. You measure the total delay from the DATA1 and NDATA1 inputs to the DUT1 output and through the cable to the oscilloscope. The data from these measurements let you calculate the signal delay through the IC.

Because TDR responses can be confusing, you should model the input delays using a Spice simulator (**Figure 4**). You compare the simulation with actual measurements and model the DATA1 and NDATA1 PCB runs as 6-in. lengths with 65Ω impedances. These traces are intended to be 50Ω runs, but TDR measurements show them to be 63Ω. You terminate the NDATA1 output to ground. Because DATA1 and NDATA1 are symmetrical, with identical lengths to the pins of the IC, you need to measure only the DATA1 PCB run. You also model a 12-in. cable from the generator, although that model

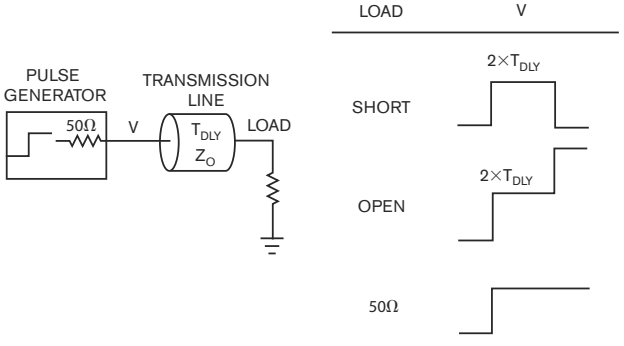

Figure 1 TDR measurements employ the reflection coefficient $\rho=(V_{REFLECTED}/V_{INCIDENT})$. The characteristic impedance, Z_O, is equal to $\rho(1+\rho)/(1-\rho)$.

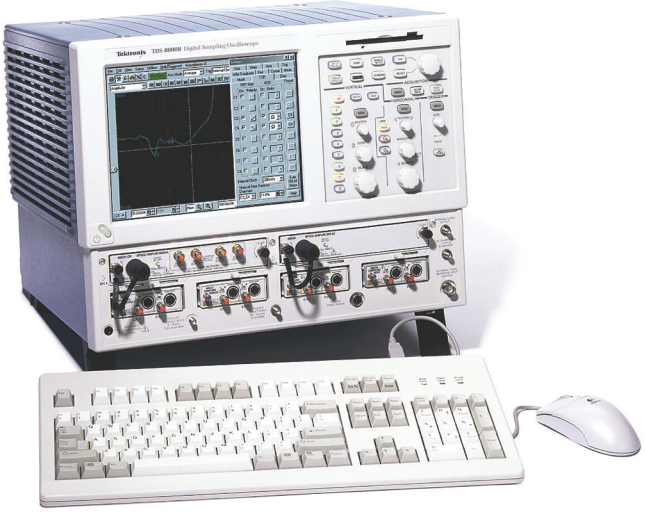

Figure 2 Rather than using active probes, you can use the TDR measurement capabilities of a Tektronix TDS8000-series oscilloscope with a model 80E04 TDR sampling module.

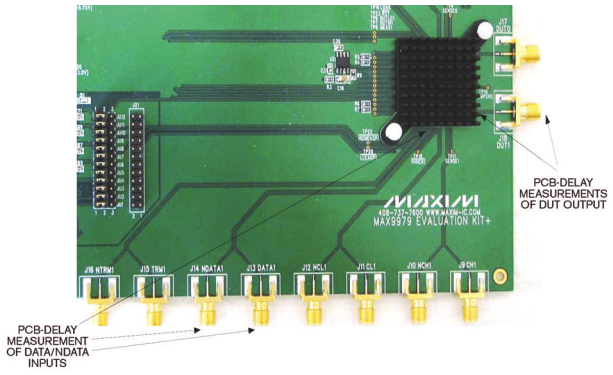

Figure 3 You can use an evaluation board from the IC manufacturer for the DUT.

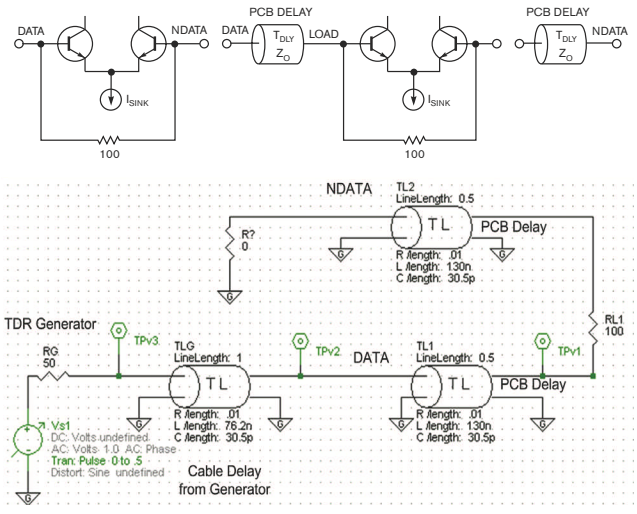

Figure 4 Because TDR responses can be confusing, you should model the input delays using a Spice simulator.

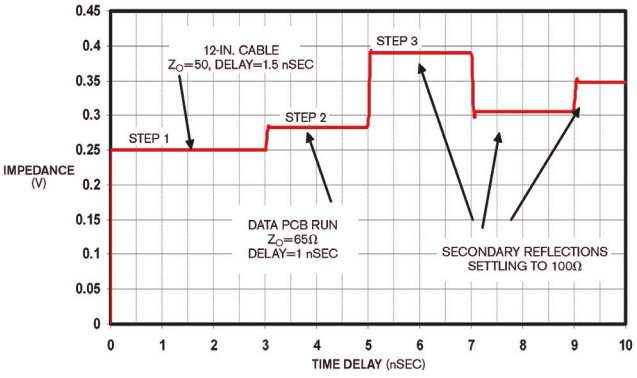

Figure 5 You solve the Spice simulation for the voltage at test point TPv3.

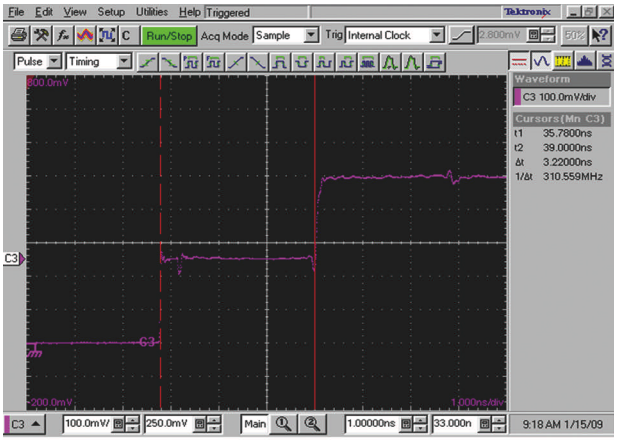

Figure 6 You measure the delay of the 2-in. SMA cable that attaches the DUT1 node to the oscilloscope's vertical input.

is not necessary for the actual propagation-delay measurement. You solve the Spice simulation for the voltage at test point TPv3 (**Figure 5**).

The simulation input signal is a step function with a 0.5V amplitude. This amplitude emulates the TDR signal from the oscilloscope. You can read the time delays for various elements in the model directly from the horizontal axis.

The part of the waveform labeled Step 1 represents the 12-in. cable from the pulse generator. The simulated delay time is about 3 nsec—twice the actual delay of 1.5 nsec. The part of the waveform labeled Step 2 represents the delay for the DATA1 PCB run. The simulation shows a delay of approximately 2 nsec—twice the actual PCB delay of 1 nsec. The other delays represent reflections of the pulse through the DATA1 PCB run.

The impedance of these various elements is proportional to voltage, as the Y axis indicates. The X axis represents signal reflections due to the single-input step signal directly in time. You can compare this simulation with the ideal version of this signal (**Figure 1**).

Use the following procedure to measure the propagation delay through the IC. First, measure the delay of the 2-in. SMA cable that attaches the DUT1 node to the oscilloscope's vertical input (**Figure 6**). Connect the 2-in. SMA-to-SMA cable to one input of the TDR module, leaving the other end open. You make the measurement using the TDR pulldown menu. Note that the waveform looks like the "open" example in **Figure 1**. Because the 804-psec delay is twice the delay of the cable, the equivalent "cable length" is 402 psec. Also note that the second waveform step is exactly halfway between the top and bottom steps. Recalling the TDR basics, this fact indicates that the impedance of this 2-in. cable is truly 50Ω.

You next measure the delay and impedance of the PCB run associated with the DATA1 input signal (**Figure 7**). You should verify the accuracy of the model by comparing this waveform with the simulation of **Figure 5**. You set the cursors to measure the impedance of the trace. The first waveform step is 50Ω, representing the cable from the oscilloscope. The second cursor shows an impedance of 97.8Ω, representing the value of the IC's internal 100Ω resistor that connects across

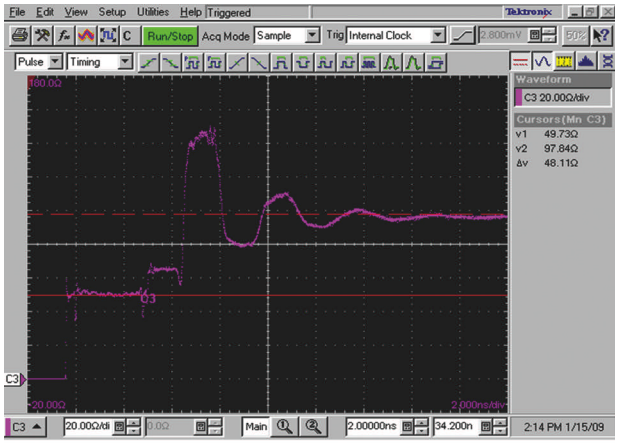

Figure 7 You measure the delay and impedance of the PCB run associated with the DATA1 input signal.

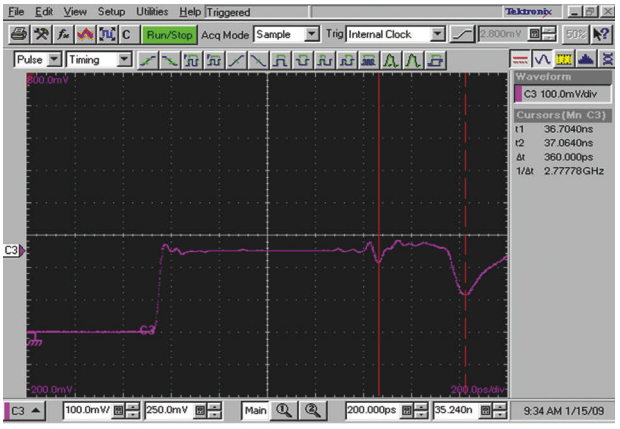

Figure 9 You measure delay and impedance for the PCB run that connects to the DUT1's output signal.

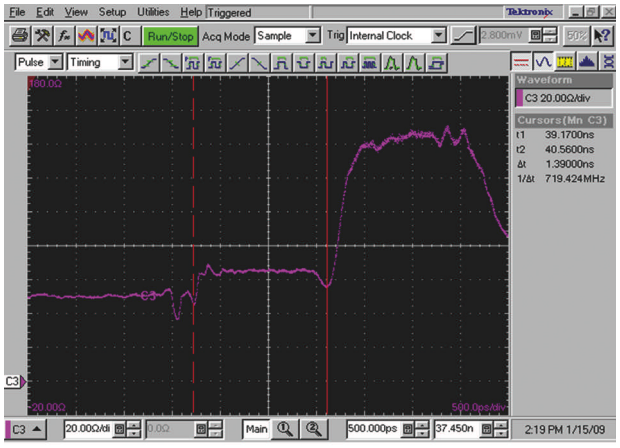

Figure 8 You can expand the waveform in Figure 7 to allow measurement of delay.

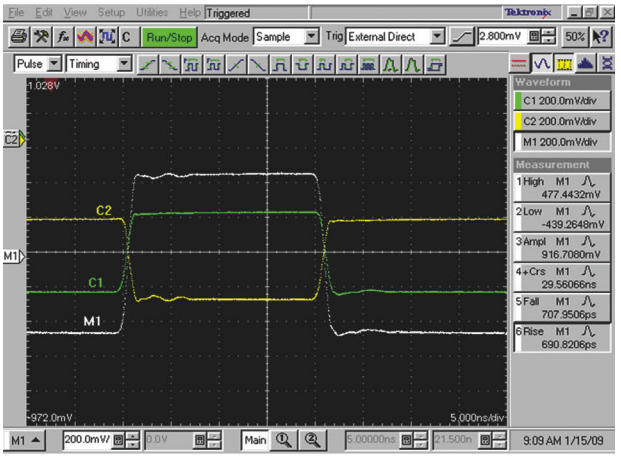

Figure 10 C1 and C2 represent the complementary-PECL DATA1 and NDATA1 signals.

DATA1 and NDATA1, which RL1 in **Figure 4** represents. The impedance of the second waveform step measures 63Ω, meaning that the PCB runs for DATA1 and NDATA1 were not designed to be 50Ω, as you would expect. The 150Ω level for the third reflection represents the sum of delays for the 50Ω cable and the 100Ω resistor.

To make this measurement, connect one end of the 12-in. SMA cable to the oscilloscope and the other end to the DATA1 SMA input connector on the evaluation board. You should ground the NDATA1 SMA connector with an SMA ground, as **Figure 4** shows. The SMA cable should be as short as possible, but its length is irrelevant to the propagation-delay measurement.

You need not apply power to the evaluation board. The measurement was made with the IC soldered onto the board and no power applied. Some users prefer to make this measurement without soldering the device. Disconnecting the IC simulates an open condition, as **Figure 1** shows, and produces a cleaner three-step signal. The delay-time measurements are the same in either configuration.

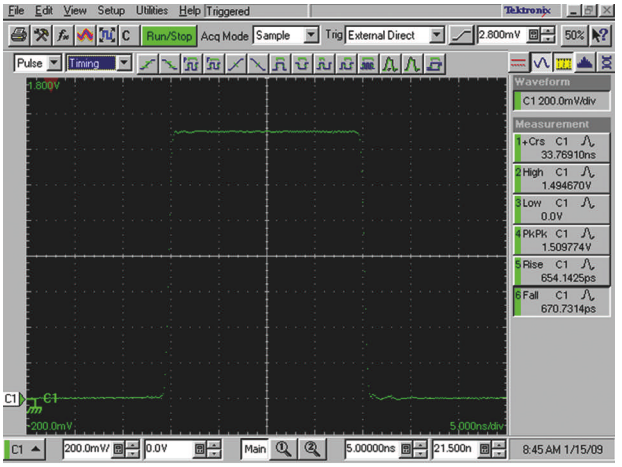

Figure 11 The IC generates a 3V signal into the 50Ω load of the oscilloscope.

The first waveform step corresponds to cable delay, which is not of interest. The second waveform step represents the delay of the DATA1 PCB run (**Figure 8**). The DATA1 PCB delay is 0.695 nsec—half the second-step measurement of 1.39 nsec. This result is larger than the model predicts, but the model estimate is only for comparison purposes.

THE USE OF BEST PRACTICES IN HIGH-SPEED MEASUREMENTS IS ALWAYS A GOOD IDEA.

You make measurements between dips in the signal. These dips indicate the presence of distributed capacitance, which the SMA connector on the board and the DATA1 pin of the IC create. As a consequence, you measure between the dips to ensure that the measurement includes SMA and pin delays. The inductance of the SMA connection to the board creates a waveform bump. You take the measurement before this bump to ensure that you capture the full board delay.

Next, you measure delay and impedance for the PCB run that connects to the DUT1's output signal (**Figure 9**). Use the same setup as that in **figures 7** and **8**. You connect a 2-in. SMA cable between the TDS8000 80E04 module and the DUT1 SMA connector on the MAX9979 evaluation kit. The first waveform step represents the 2-in. cable. The TDR's signal amplitude is 0.5V. The 250-mV amplitude indicates that the cable impedance is 50Ω.

You make the DUT1 delay measurement between the two dips, just as with the DATA1 measurement. The level between the dips is 50Ω, indicating that the short PCB-metal run to DUT1 is close to ideal. It's difficult to measure the PCB delay for DUT1 because its impedance appears the same as that of the cable. If the IC were not soldered to the board, you would see the three-step signal, indicating an open. You can still measure this delay with the IC soldered in place. An examination of the capacitive dips reveals one dip corresponding to the SMA connector, soldered to the board, and one dip corresponding to the IC's DUT1 pin. You should look for an inductive bump corresponding to the SMA connector. Ensure that this bump is between the two capacitive dips. The delay is 360 psec. You halve this value to obtain the actual DUT1 PCB delay of 180 psec.

The impedance of the DATA1 run is 63Ω. The DUT1 node has an impedance of 50Ω. Ideally, these impedances should be the same, meaning that the metal on the DATA1 inputs is narrower than that of the DUT1 output.

You should set up the differential signal generator with two identical SMA cables. You then measure the baseline delay on the oscilloscope. C1 and C2 represent the complementary-PECL (positive-emitter-coupled-logic) DATA1 and NDATA1 signals (**Figure 10**) with an amplitude of 450 mV. You feed the signals to the inputs of the oscilloscope from the external generator.

Waveform M1 is a mathematical calculation of the differential signal between C1 and C2. Its amplitude is 900 mV, and its rise and fall times are each 700 psec. These characteristics indicate that you have a acquired a valid set of data. You then trigger the scope and measure one of M1's zero-crossing points as 29.56 nsec. Power up the IC and measure the same crossing point as it is delayed through the evaluation board. This delay includes the delay of the two input cables. These delays cancel out because you use the same cables to measure signal delay through the PCB. Keep these cables short, but their delays are not important for the propagation-delay measurement.

Using the two cables from the setup, connect the DATA1 and NDATA1 signals to the board's DATA1 and NDATA1 inputs. The oscilloscope setup and the trigger are the same as they were previously. Refer to the evaluation board's documentation to set the IC for signal amplitudes of 0 to 3V. The output of the board terminates in 50Ω because that is the input impedance of the oscilloscope. The 50Ω load halves the output signal, producing an amplitude of 0 to 1.5V (**Figure 11**). You can verify that the rise and fall times are well within the IC's specifications, meaning that clean and valid DATA1/NDATA1 signals are driving a clean and valid output. You can measure the zero-crossing point as 33.77 nsec.

You can now calculate the IC's propagation delay. The total delay through the evaluation board is 33.77 nsec−29.56 nsec=4.21 nsec. You then subtract the PCB-trace delay of DATA1 from the total delay to get 3.515 nsec. Subtract the DUT1 PCB-run delay of 0.18 nsec, which yields 3.335 nsec. Finally, subtract the delay of the 2-in. cable to the oscilloscope for a result of 2.933 nsec. This result correlates well with the IC data sheet's specification of 2.9 nsec for a typical delay.

⊞ Go to www.edn.com/ms4357 and click on Feedback Loop to post a comment on this article.

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

Using TDR for measuring propagation delays has several advantages. It gives accurate measurements; requires no active probes, thus avoiding the inaccuracies they introduce; and is a simple technique that suffices for most propagation measurements. The measurement also lets you check for correct impedances on connectors and PCB runs. The TDR signals reveal excess capacitance and inductance in the signal path should you need to redesign the board. You can verify the TDR measurement with simple modeling and simulation tools.

As signal speeds rise, the errors and mistakes of timing measurements can cause incorrect planning decisions, faulty device selections, and bad system design. The use of best practices in high-speed measurements is always a good idea. EDN

REFERENCE

1 Rako, Paul, "TDR: taking the pulse of signal integrity," *EDN*, Sept 3, 2007, pg 48, www.edn.com/article/CA6470825.

AUTHOR'S BIOGRAPHY

Bernard Hyland joined Maxim Integrated Products in 2001 as an IC-design engineer. Before joining Maxim, he worked as an IC designer at Delco Electronics, Micro Systems Engineering, Bipolar Integrated Technologies, and Tektronix. Hyland earned a bachelor's degree in electrical engineering in signal communications from Purdue University (West Lafayette, IN) in 1985.

BY SANJEEV GUPTA • AGILENT TECHNOLOGIES EESOF DIVISION

Continuous-time equalizers improve high-speed serial links

YOU CAN DESIGN AN EFFECTIVE SERIAL-CHANNEL EQUALIZER IN A FEW MINUTES.

Impedance mismatches on PCBs (printed-circuit boards) and backplanes can cause reflections that impair high-speed chip-to-chip connections. These reflections also create ISI (intersymbol interference), an echo of a previous bit superimposed on the current bit, which causes your system to become unreliable. To overcome the interference, you can employ various signal-processing techniques, including de-emphasis, in which you predistort the signal before transmission. You use an analog filter to make a CTE (continuous-time equalizer) in the transmission path in just a few minutes.

You can best characterize a high-speed digital channel with S (scattering) parameters. S21, an insertion-loss plot, represents the signal attenuation by the channel as a function of frequency. S11, the return-loss plot, represents the reflected signal due to impedance mismatch or discontinuities in the signal path. When a channel path is long, the high-frequency part of the signal suffers severe attenuation (**Figure 1**).

High-frequency roll-off manifests itself as a rise- and fall-time degradation in the time domain. This degradation limits the data rate you can achieve with the interconnection. You can mitigate this impairment using de-emphasis of the low frequencies at the transmitter or equalization at the receiver.

Reflections also cause echoes in the time domain that appear as phase distortion in the frequency domain, and ISI causes increased jitter and a pattern-dependent variation in the signal amplitude. As a result, the eye diagram remains more closed, making detection at the receiver more prone to error. A CTE mitigates both reflections and ISI.

In preparation for a CTE design, plot and smooth the backplane response with a log-frequency scale (**Figure 2**). The insertion-loss plot looks noisy at frequencies greater than 2 GHz. You should perform moving averaging on the insertion-loss plot to remove unwanted noise from the measured data. You can decide on the number of points for the moving average based on the nature of data and to get a smoother response.

You can compare the original insertion-loss response with the moving average (**Figure 3**). The two plots look similar in the passband, but the moving average provides a much smoother response at higher frequency. At 6.25 Gbps, the eye diagram provides insufficient opening for an error-free detection at the receiver (**Figure 4**). You can design a pole-zero CTE to achieve optimum performance from a Tyco backplane running at 6.25 Gbps.

Some people refer to the complex frequency plane as the Laplace, or S plane. A simple pole-zero transfer function in the complex frequency plane represents a CTE:

$$H(s) = A\frac{(s - z_1)(s - z_2)...}{(s - p_1)(s - p_1)...}$$

The CTE representation is more complicated in cases such as the HDMI (high-definition multimedia interface).

The PCIe (Peripheral Component Interconnect Express) 3.0 and USB (Universal Serial Bus) 3.0 standards specify the nominal location of the poles and zeros. The exact optimum positions depend on the channel. Fine-tuning the location of pole and zero increases your timing margin for a given interconnect. You should create a mathematical representation to gain more control of the parameters that you can define.

You can represent the equalizer in any math package that supports complex-number mathematics. You can even use Excel for this task if you can figure out a workaround for the software's lack of support for complex-number math. You can also use models within Agilent's ADS (Advanced Design System), such as VCVS (voltage-controlled voltage source)_PZR, VCVS, and the channel-simulator receiver model, and then write the equations in the ADS data display (**Figure 5**).

You represent the CTE equalizer with a handful of equations on the data-display page. Instead of manually entering these equations, you can use the data-display template in ADS as

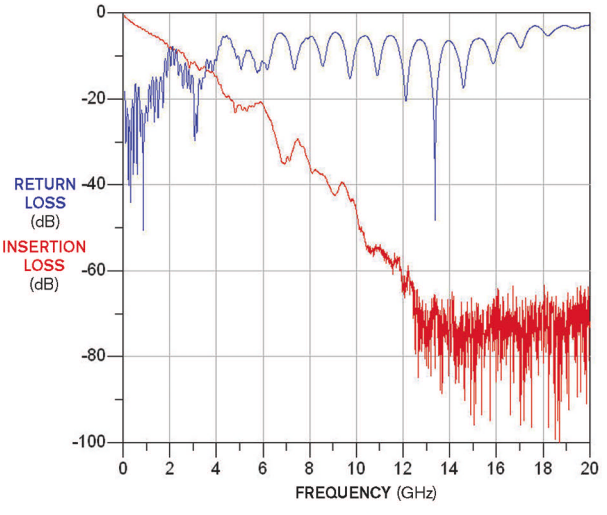

Figure 1 When a channel path is long, the high-frequency part of the signal suffers severe attenuation.

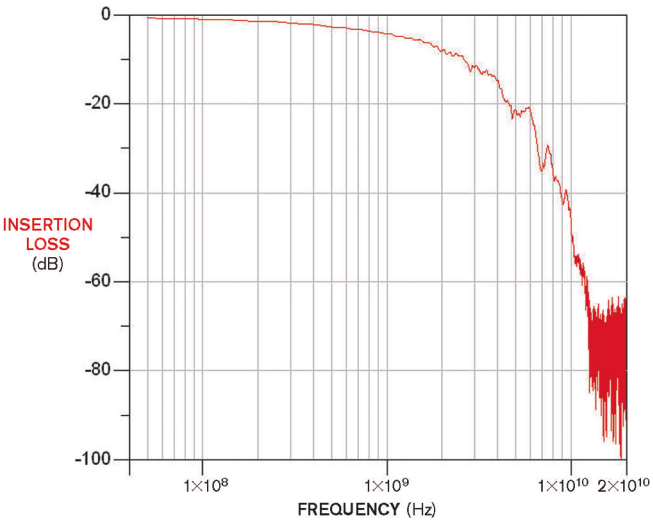

Figure 2 In preparation for a CTE design, plot and smooth the backplane response with a log-frequency scale.

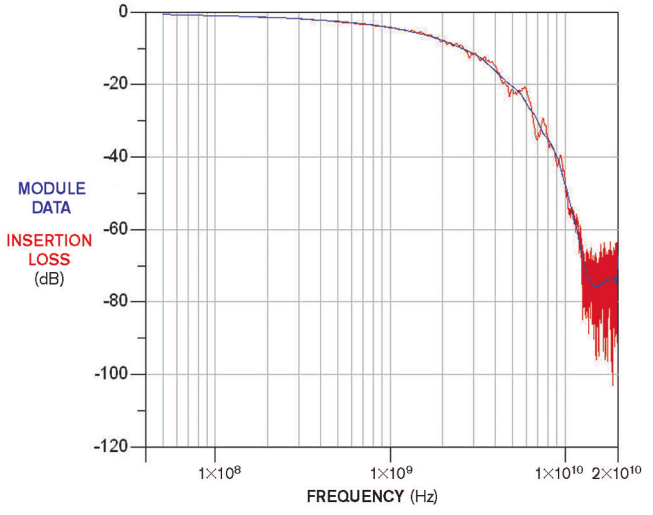

Figure 3 Apply a moving averaging to the insertion loss.

part of the PCIe DesignGuide. The plot of the transfer-function equation provides the equalizer response in decibels. Plotting the transfer function shows the effect of the pole and zero locations on the equalizer response.

A simple zero location in a transfer function introduces 20 dB of gain per decade; a pole introduces 20 dB of loss per decade. You can visually determine the approximate values of poles and zeros from the frequency-domain response if they have separate frequencies. Draw tangential lines on the parts of the frequency-response curve that have different slopes. The intercepts provide the pole and zero values (**Figure 6**).

Because a zero in the transfer function introduces gain in the system, you need more poles than zeros to make a system stable. Optimize the locations of poles and zeros to compensate for the channel insertion loss. To design an optimum equalizer, plot the inverse of the pole and zero equalizer response over channel insertion loss with a log-frequency scale.

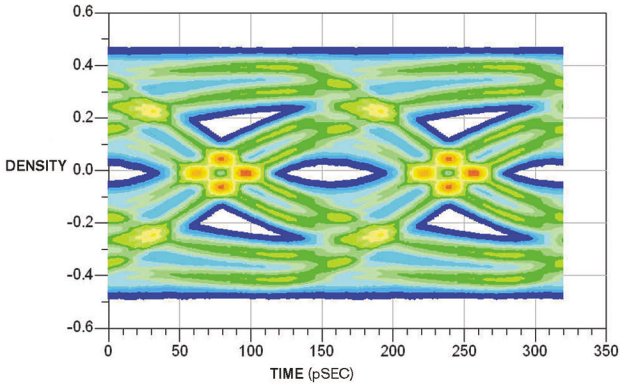

PERMUTE (HEIGHT)	PERMUTE (WIDTH)
0.036	6×10^{-11}

Figure 4 At 6.25 Gbps, the eye diagram provides insufficient opening for an error-free detection at the receiver.

```
Eqn Omega=[0::1e6::10e10]
Eqn zero1=20*log(abs(2*pi*Zero_freq+j*2*pi*Omega))
Eqn pole1=-20*log(abs(2*pi*Pole1_freq+j*2*pi*Omega))
Eqn pole2=-20*log(abs(2*pi*Pole2_freq+j*2*pi*Omega))
Eqn Gain=20*log(Gain_constant*2*pi*Pole2_freq*Pole1_freq/Zero_freq)
Eqn Transfer_function=Gain+zero1+pole1+pole2
```

Figure 5 You can also use models within Agilent's ADS (Advanced Design System), such as VCVS_PZR, VCVS, and the channel-simulator receiver model, and then write the equations in the ADS data display.

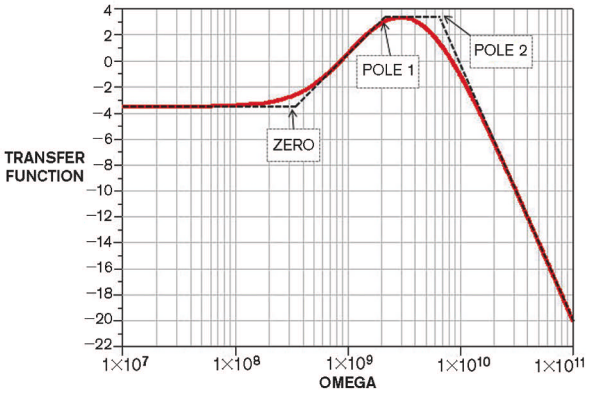

Figure 6 You can perform a graphical construct to show the approximate frequency values of the poles and zeros.

In this case, the pole and zero locations link to a GUI (graphical-user-interface) slider—the three marker positions, each on linear frequency scales. You use markers 1, 2, and 3 to change the location of poles and zeros until you get the same frequency response as the insertion loss of the channel over the maximum frequency bandwidth. A good match between the two responses provides the optimum design.

In the plot, the blue curve represents the insertion of the

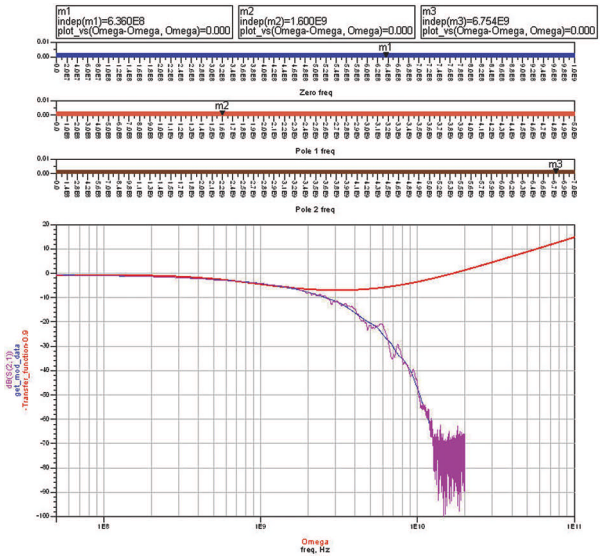

Figure 7 The blue curve represents the insertion of the Tyco backplane, and the magenta curve shows the inverted response of the pole-zero equalizer in decibels.

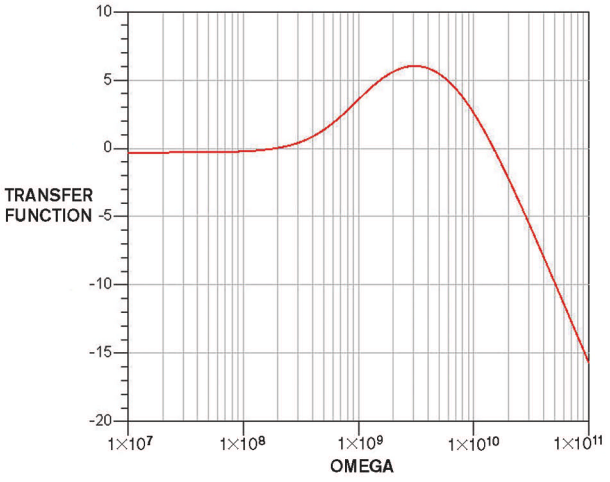

Figure 8 For this Tyco backplane, the optimized design provides the zero location at 636 MHz, and the first and second poles are at 1.622 and 6.754 GHz, respectively.

Tyco backplane, and the magenta curve shows the inverted response of the pole-zero equalizer in decibels (**Figure 7**). You should adjust the sliders until the two curves show good agreement up to the highest frequency of interest, or 2 GHz in this case. The combined insertion loss will be flat to 2 GHz, a large improvement over a nonequalized system, which has an insertion-loss slope starting at 200 MHz.

You could try to write an objective function and use machine optimization, but this manual method provides a visual way of exploring the design and finding out what is possible. For this Tyco backplane, the optimized design provides the zero location at 636 MHz, and the first and second poles are at 1.622 and 6.754 GHz, respectively (**Figure 8**).

ADS integrates the equalizer for channel analysis using a VCVS_PZR source (**Figure 9**). The channel receiver component also has a CTE component. You plot the eye diagram before and after you insert the equalizer, using the channel simulator in statistical mode. Using this equalizer, the eye diagram shows dramatic improvement in performance (**Figure 10**). The overall design of this behavioral equalizer takes less than 15 minutes.

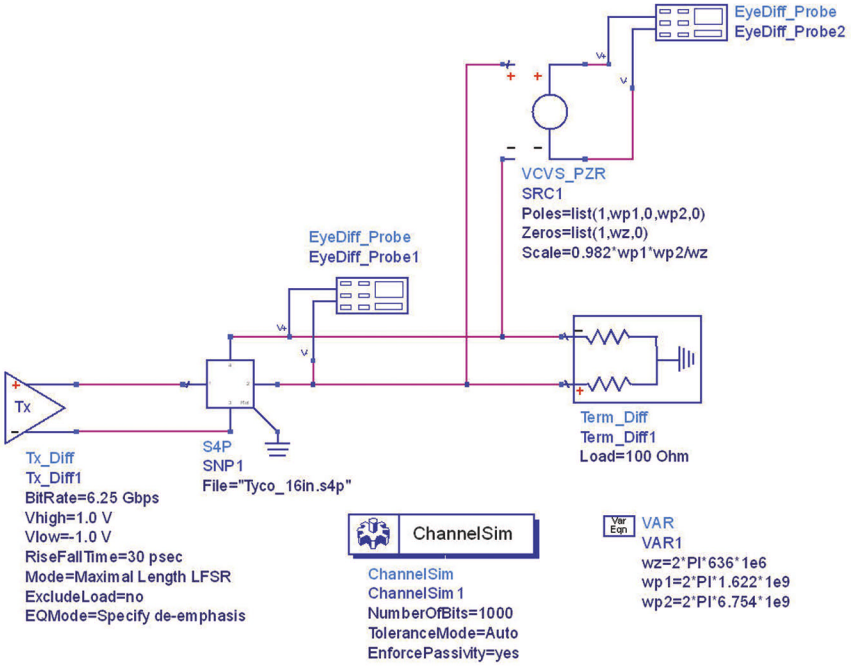

Figure 9 You use the channel-simulation schematic with the equalizer to optimize the filter.

In addition to determining the optimum locations of the poles and zeros of a CTE, you can simulate off-the-shelf CTE models within ADS. An off-the-shelf equalizer device comes with tabular data representing the gain- and phase-transfer characteristics as functions of frequency and boost levels. When you use an EDA tool to simulate with these models, you must pay attention to data formatting, data access, and the electrical representation of the device using this data.

Many semiconductor vendors provide CTE data in the CSV (comma-separated-values) file format. It lists frequency as the first column and the real and imaginary components of the transfer function with various boost settings in successive columns. EDA tools cannot use the CSV data, so you have to reformat it. Touchstone, CITI (common instrumentation transfer and interchange), or mdif (measure-

ment-data-interchange format) are acceptable formats for representing a device. Because the standard Touchstone file format supports only one boost setting within a given file and requires a full S-matrix description, it has a limited ability to represent a CTE behavioral model. The CITI file format can easily handle multidimensional data. The Touchstone mdif provides the flexibility to support arbitrary multidimensional data in which blocks of data in a file represent different boost settings. The data is organized in those blocks as a function of frequency, similar to a standard Touchstone file.

High-speed digital designers must use the frequency-domain data to predict eye-diagram performance. While running a time-domain simulation, you must search and interpolate the data with respect to the frequency and boost settings. Simulation environments such as ADS can read all these formats and, using the data-access component, search and interpolate the data as a function of frequency and boost settings. You can access data in MA (linear-magnitude-and-phase-angle), RI (real-and-imaginary), or dBAngle (decibel-magnitude-and-phase-angle) format and use this format with an electrical behavior model representing a CTE.

You can use electrical components such as the equation-based linear SNP (S-parameter/number-of-ports) component, the VCVS, and the system-amplifier model to represent a CTE device, depending on the type of data the device vendor provides. Touchstone SNP components are load-sensitive. You must use correct terminations to accurately replicate the device behavior over various boost settings. If device load conditions are unavailable, you should use a VCVS to represent a CTE device.

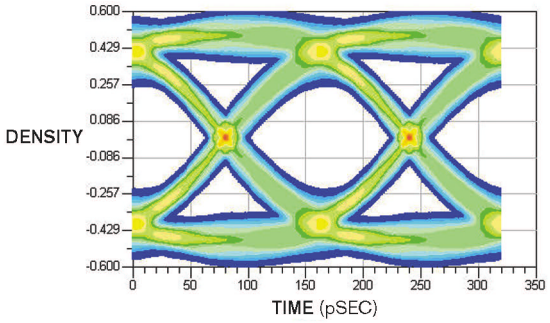

EYEDIFF PROBE 2 (HEIGHT)	EYEDIFF PROBE 2 (WIDTH)
0.427	1.216×10^{-10}

Figure 10 Using this equalizer, the eye diagram shows dramatic improvement in performance.

Once you convert the data to the required format, you can use the data-access component with an equation-based SNP component or with a VCVS device. You can then use the CTE model in a time-domain environment to predict the eye-diagram and BER (bit-error-rate) performance of the equalizer. You can now do in minutes what used to take days. **EDN**

AUTHOR'S BIOGRAPHY

Sanjeev Gupta is a signal-integrity-applications manager at Agilent EEsof. You can reach him at sanjeev_gupta@agilent.com.

±1LSB 18-Bit DAC

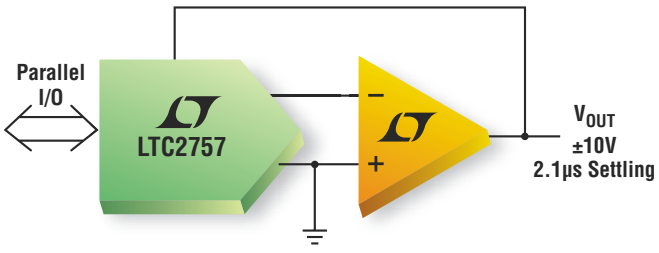

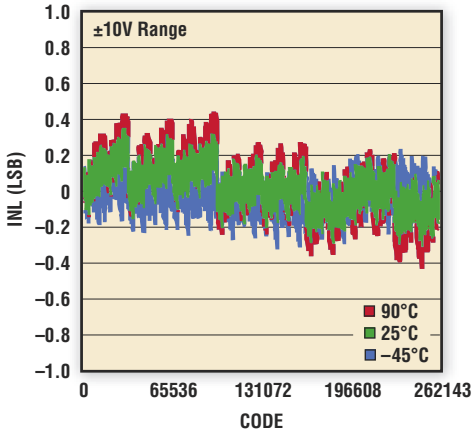

Guaranteed ±1LSB INL & DNL Over Temperature

The LTC®2757 is the first 18-bit digital-to-analog converter to offer precision DC specifications of ±1LSB INL (max) and ±1LSB DNL (max). The LTC2757 delivers 18-bit accuracy over four bipolar and two unipolar output ranges for precision adjustment in open loop applications. The device has excellent AC specifications, including full-scale settling time of only 2.1µs and low 1.4nV•s glitch impulse.

Features

- 2.7V to 5.5V Single Supply Operation
- Program or Pin-Strap Six Output Ranges: 0V to 5V, 0V to 10V, –2.5V to 7.5V, ±2.5V, ±5V, ±10V
- Low Glitch Impulse 1.4nV•s (3V), 3nV•s (5V)
- 18-Bit Settling Time: 2.1µs
- Low 1µA Maximum Supply Current
- Voltage-Controlled Offset & Gain Trims
- Parallel Interface with Readback on All Registers
- 48-Pin, 7mm x 7mm LQFP Package

Settling Full-Scale Step

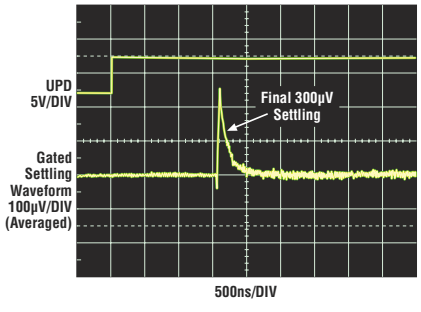

t_{SETTLE} = 1.8µs to 0.0004% (18 Bits)

Info & Free Samples

www.linear.com/2757

1-800-4-LINEAR

Free Industrial Signal Chain Brochure

www.linear.com/indsolutions

LINEAR TECHNOLOGY

EDITED BY MARTIN ROWE
AND FRAN GRANVILLE

designideas

READERS SOLVE DESIGN PROBLEMS

Efficient LED power supply has battery backup

Zhihong Yu, Juno Lighting Group, Des Plaines, IL

LEDs find wide use in emergency lighting because of their high efficiency and control simplicity. The circuit in **Figure 1** provides a highly efficient and reliable design for emergency LED lighting at 3 to 6W. The circuit's input is 12V ac, which the full-wave bridge rectifies and one or two capacitors filter into dc. The battery (not shown) is a 12V lead-acid type. IC_1 compares the battery voltage to the supply voltage. When the rectified voltage drops below the battery voltage, the battery takes over to provide LED power.

The circuit has some small switching losses, which should be acceptable as long as IC_2, a 12V PB137 battery-charging circuit from STMicroelectronics (www.st.com), keeps the battery from draining. If this switch-over is unacceptable, add a 470-µF electrolytic capacitor to filter the input voltage to maintain a certain level above the battery voltage. Note that adding this capacitor lowers the power factor.

To get 12V ac, you can use an electronic transformer. These transformers

DIs Inside

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

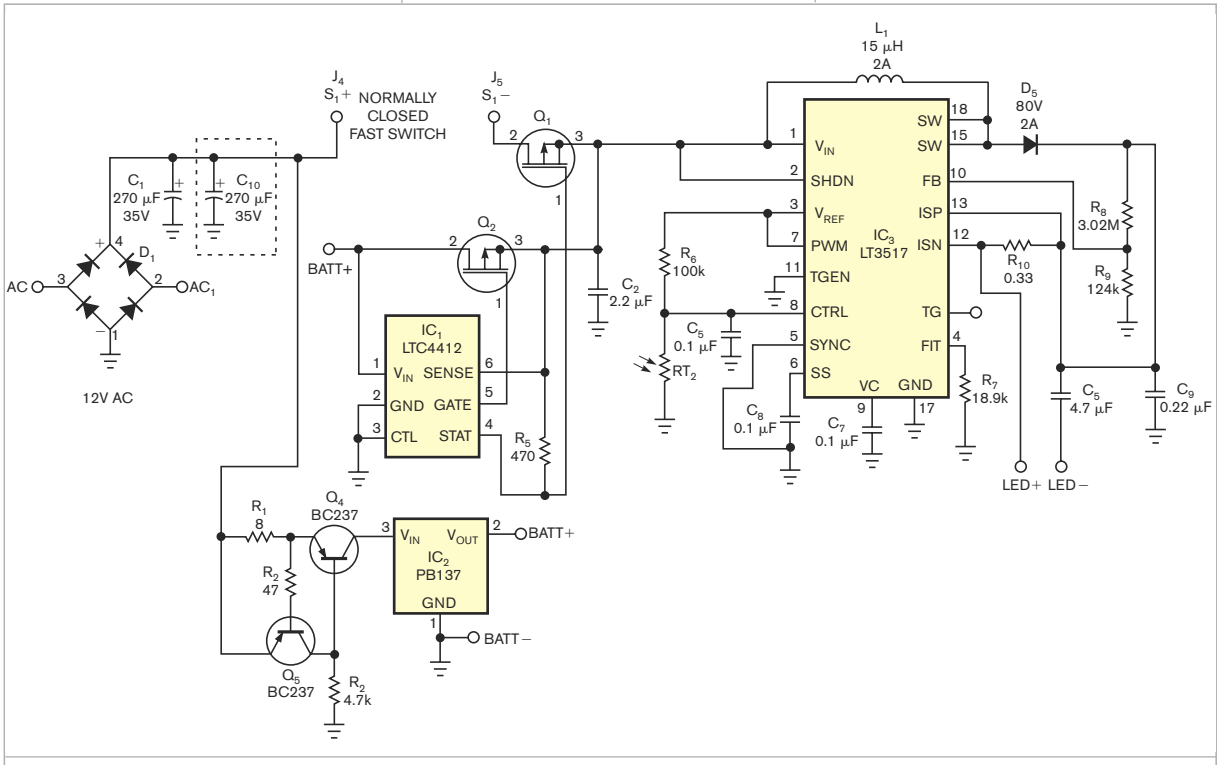

Figure 1 This circuit converts ac voltage to dc voltage, charges a battery, and drives LEDs from the ac source or the battery.

provide 12V at a higher frequency, so a 10-µF capacitor can hold the voltage high as well as provide a high power factor.

IC_1, a Linear Technology (www.linear.com) LTC4412, controls two external PFETs that create a near-ideal diode function for switching between ac and battery output. The PFETs' voltage drop is only about 20 mV compared with a normal 0.7V diode-voltage drop. Pin 5 is low when ac power is off, so you can use this pin to turn on a warning LED through another PFET. IC_2 has an internal current limit of 1.5A. Resistor R_1 limits IC_2's input; when the current reaches a certain level, Q_4 turns off the charging circuit. This IC does not require reverse-diode protection.

IC_3, an LT3517 LED driver from Linear Technology, acts as an inverting buck-boost converter because the input can range from 8 to 17V for rectified ac. R_{10} sets up the LEDs' current. Because the voltage drop from each of the three LEDs varies from 3 to 4V, the IC's output voltage can be higher or lower than its input voltage if all 300-mA LEDs connect in series.

By connecting a resistor divider, including a photocell, to the analog-dimming pin, Pin 8, you can achieve some dimming, which results in some power savings at higher ambient light. You can use IC_1's Pin 5 to turn on a transistor or an optoisolator to pull IC_3's control-pin voltage lower if you need to dim the LED when ac power is out. Resistor R_7 programs IC_1 to operate at 1 MHz. The circuit's efficiency is 82% when you power it directly from the ac power supply and about 70% from an electronic transformer.

With a few minor changes to the circuit, you can add LEDs. For example, you can use Linear Technology's LT3518, which is a pin-to-pin-compatible version of the LT3517 but with a higher switching-current limit. You may need to adjust the feedback-resistor pair R_8 and R_9 for higher output voltage. You may also need more input-filtering capacitance to hold up the voltage.

Tests show that the circuit can power as many as six LED in series. **Figure 2** shows the circuit in operation. EDN

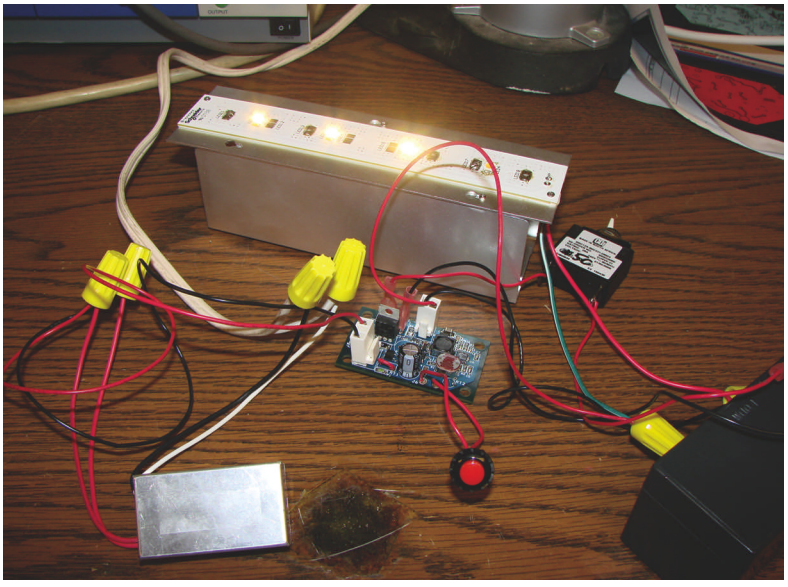

Figure 2 LEDs provide enough light for emergency lighting.

Single IC forms precision triangular-wave generator

Akshay Bhat, Maxim Integrated Products Inc, Sunnyvale, CA

The linearity of triangular waveforms makes the triangular-wave generator useful in sweep circuits and test equipment. For example, switched-mode power supplies and induction motor-control circuits often include a triangular-wave oscillator as part of their PWM (pulse-width-modulation) circuit.

The basic triangular-wave generator includes an integrator for generating the triangular-wave output and a comparator with external hysteresis, such as a Schmitt trigger, for setting the output amplitude (**Figure 1**). You can implement these components with a Maxim (www.maxim-ic.com) MAX9000 IC, which includes a high-speed operational amplifier, a 185-nsec comparator, and a precision 1.23V bandgap reference.

The integration of a constant current, which you obtain by applying constant voltage across a resistor, produces a linear ramp at the op amp's output. This output feeds a Schmitt trigger whose output feeds back to the integrator resistor. Abrupt state changes in the Schmitt trigger's output determine the peak voltages for the triangular-wave output. These changes in turn depend on the input threshold voltages you set for the Schmitt trigger.

Unfortunately for this circuit, the triangular-wave peaks must be symmetrical about the reference voltage you apply to the comparator's inverting input. To generate a triangular wave from 0.5 to 4.5V, for example, you must provide a reference voltage of (0.5V+4.5V)/2=2.5V.

It would be preferable to set this voltage range independently of the standard bandgap-reference voltage

LINEAR TECHNOLOGY

DESIGN NOTES

Driving Lessons for a Low Noise, Low Distortion, 16-Bit, 1Msps SAR ADC

Design Note 477

Guy Hoover

Introduction

Designing an ADC driving topology that delivers uncompromising performance is challenging, especially when designing around an ultralow noise SAR ADC such as the 1Msps LTC2393-16. For both single-ended and differential applications, a well thought out driving topology can fully realize the ultralow noise and low distortion performance required in your data acquisition system.

The LTC2393-16 is the first in a family of high performance SAR ADCs from Linear Technology that utilizes a fully differential architecture to achieve an excellent SNR of 94.2dB and THD of –105dB. And in order to take full advantage of the ADC performance, we present driving solutions for both single-ended and differential applications. Both topologies fully demonstrate the ultralow noise and low distortion capabilities of the LTC2393-16.

Single-Ended to Differential Converter

The circuit of Figure 1 converts a single-ended 0V to 4.096V signal to a differential ±4.096V signal. This circuit is useful for sensors that do not produce a differential signal. Resistors R1, R2 and capacitor C2 limit the input bandwidth to approximately 100kHz.

When driving a low noise, low distortion ADC such as the LTC2393-16, component choice is essential for maintaining performance. All of the resistors used in this circuit are relatively low values. This keeps the noise and settling time low. Metal film resistors are recommended to reduce distortion caused by self-heating. An NPO capacitor is used for C2 because of its low voltage coefficient, which minimizes distortion. The excellent linearity characteristics of NPO and silver mica capacitors make these good choices for low distortion applications. Finally, the LT6350 features low noise, low distortion and a fast settling time.

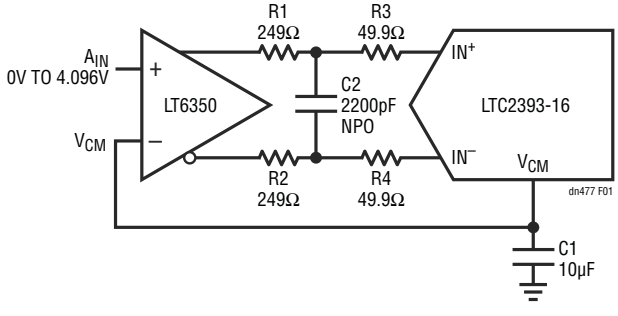

Figure 1. Single-Ended to Differential Converter

The 16k-point FFT in Figure 2 shows the performance of the LTC2393-16 in the circuit of Figure 1. The measured SNR of 94dB and THD of –103dB match closely with the typical data sheet specs for the LTC2393-16, showing that little, if any, degradation of the ADC's specifications result from inserting the single-ended to differential converter into the signal path.

Fully Differential Drive

The circuit of Figure 3 AC-couples and level shifts the sensor output to match the common mode voltage of the ADC. The lower frequency limit of this circuit is about 10kHz. The lower frequency limit can be extended by increasing the values of C3 and C4. This circuit is useful for sensors with low impedance differential outputs.

The circuit of Figure 1 could be AC-coupled in a similar manner. Simply bias A_{IN} to V_{CM} through a 1k resistor and couple the signal to A_{IN} through a 10µF capacitor.

PCB Layout

The circuits shown are quite simple in concept. However, when dealing with a high speed 16-bit ADC, PC board layout must also be considered. Always use a ground

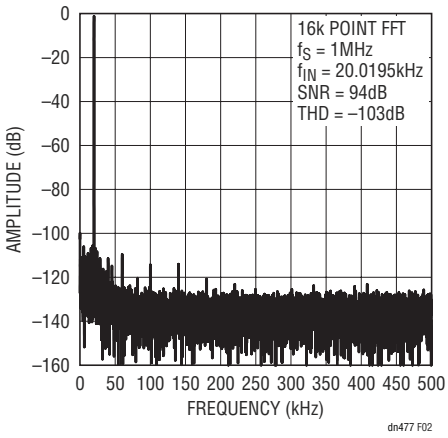

Figure 2. LTC2393-16 16k Point FFT Using Circuit of Figure 1

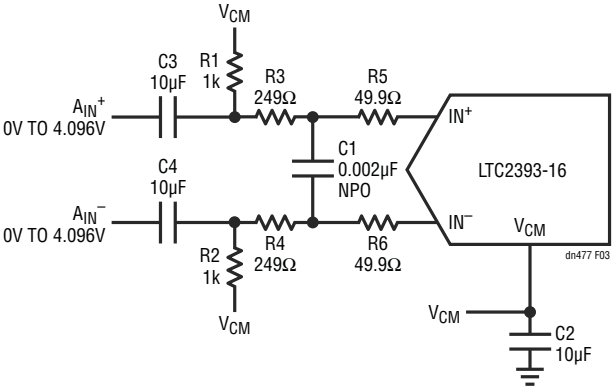

Figure 3. AC-Coupled Differential Input

plane. Keep traces as short as possible. If a long trace is required for a bias node such as V_{CM}, use additional bypass capacitors for each component attached to the node and make the trace as wide as possible. Keep bypass capacitors as close to the supply pins as possible. Each bypass capacitor should have its own low impedance return to ground. The analog input traces should be screened by ground. The layout involving the analog inputs should be as symmetrical as possible so that parasitic elements cancel each other out.

Figure 4 shows a sample layout for the LTC2393-16. Figure 4 is a composite of the top metal, ground plane and silkscreen layers. See the DC1500A Quick Start Guide available at www.linear.com for a complete LTC2393-16 layout example.

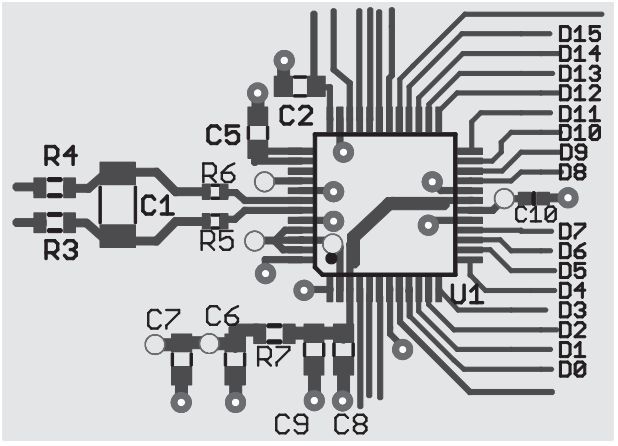

Figure 4. Sample Layout for LTC2393-16

Conclusion

The LTC2393-16 with its fully differential inputs can improve SNR by as much as 6dB over conventional differential input ADCs. This ADC is well suited for applications that require low distortion and a large dynamic range. Realizing the potential low noise, low distortion performance of the LTC2393-16 requires combining simple driver circuits with proper component selection and good layout practices.

Data Sheet Download

www.linear.com

For applications help,
call (408) 432-1900

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

dn477 LT/TP 0410 116K • PRINTED IN THE USA

available, 1.23V. You can achieve this flexibility by adding resistor R_3 to the hysteresis network in a single-IC version of the circuit (**Figure 2**). R_3 lets you set the triangular-wave peaks independently of the reference voltage.

To build the Schmitt-trigger comparator, you first select R_2. The comparator's input-bias current at C_{IN+} is less than 80 nA. To minimize the error this current causes, the current through R_2, $[(V_{REF}-V_{OUT})/R_2]$, should be at least 8 µA. R_2 requires two **equations**, corresponding to the two possible comparator-output states: $R_2=V_{REF}/I_{R2}$, and $R_2=(V_{DD}-V_{REF})/I_{R2}$.

Use the smaller of the two resulting resistor values. For example, if the supply voltage is 5V, the reference voltage is 1.23V, and the reference current is 8 µA, the two R_2 values are 471.25 and 153.75 kΩ, so this circuit uses the standard value of 154 kΩ.

Next, select R_1 and R_3. During a rising ramp, the comparator output is logic low (V_{SS}). Similarly, the comparator output is at logic high (V_{DD}) during a falling ramp. Thus, the comparator must change state according to the required peak and valley points of the triangular wave.

Two simultaneous **equations** result when you apply nodal analysis at the noninverting input of the comparator and solve for these two thresholds:

$$\frac{V_{IH}}{R_1}+\frac{V_{SS}}{R_2}=V_{REF}\left(\frac{1}{R_1}+\frac{1}{R_2}+\frac{1}{R_3}\right),$$

and

$$\frac{V_{IL}}{R_1}+\frac{V_{DD}}{R_2}=V_{REF}\left(\frac{1}{R_1}+\frac{1}{R_2}+\frac{1}{R_3}\right).$$

In this example, the voltage range of the triangular wave is 0.5 to 4.5V. You therefore substitute a value for V_{IH} of 4.5V, V_{IL} of 0.5V, V_{DD} of 5V, and V_{REF} of 1.23V into the above **equations** to obtain a value of 124 kΩ for R_1 and 66.5 kΩ for R_3.

You can now design the integrator. Considering the comparator's two possible output states, the magnitude of current flowing through R_4 is: $I_{R4}=(V_{DD}-V_{REF})/R_4$, or $I_{R4}=V_{REF}/R_4$. The op amp's maximum input-bias current is 2 nA. To minimize error, therefore, the current through R_4 must always be greater than 0.2 µA. This constraint implies that R_4's value is less than 6.12 MΩ.

The triangular-waveform frequency is:

$$f=1/\left(\frac{V_{OUT\,P-P}}{(V_{CC}-V_{REF})}(R_4C)+\frac{V_{OUT\,P-P}}{V_{REF}}(R_4C)\right).$$

For this example, the frequency is 25 kHz, the output voltage is 4V p-p, or 0.5 to 4.5V for a triangular wave, and the reference voltage is 1.23V. Solving for the resulting time constant, R_{4C}=9.27 µsec. Select a capacitance of 220 pF and a value of 42.2 kΩ for R_4.

The resulting output should match the desired frequency, provided that the op amp is not slew-limited. Because the feedback capacitor charges or discharges with a constant current, the output signal's maximum rate of

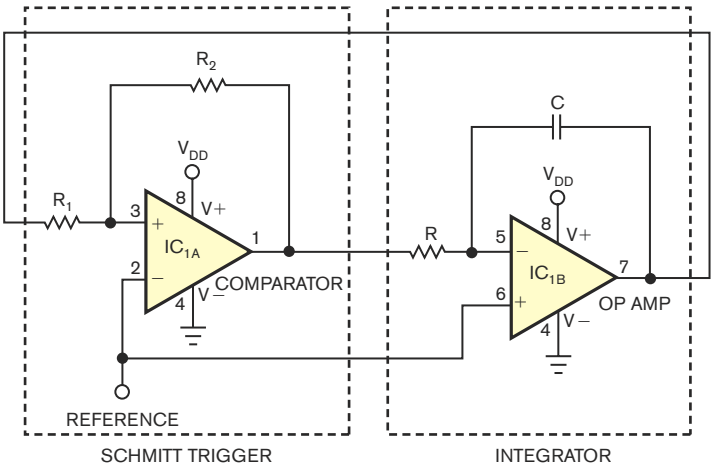

Figure 1 The basic triangular-wave generator includes an integrator for generating the triangular-wave output and a comparator with external hysteresis, such as a Schmitt trigger, for setting the output amplitude.

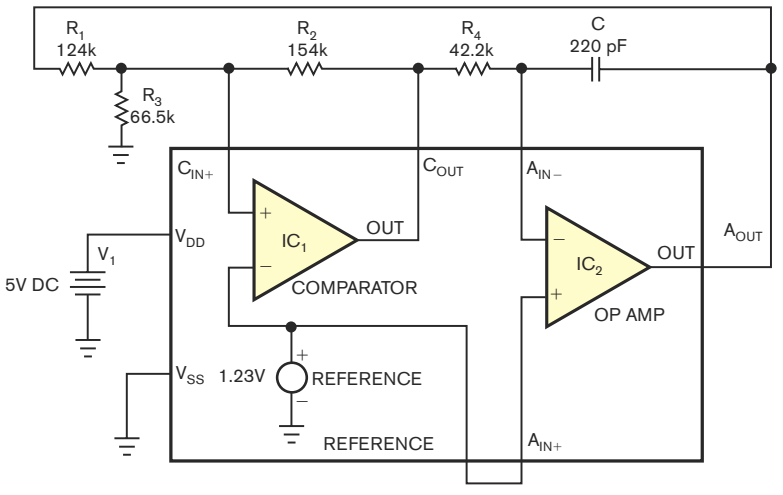

Figure 2 A triangular-wave generator employs an IC that includes an op amp, a comparator, and a bandgap reference.

change is:

$$\frac{dV_{OMAX}}{dt} = \frac{I_{R4MAX}}{C} =$$

$$\frac{V_{CC}-V_{REF}}{R_4C} = 0.406\frac{V}{\mu SEC}.$$

To provide a margin against process variations, the op amp's typical slew rate should be 40% higher than the maximum rate of change of the output signal—0.56V/µsec or greater in this case. The op amp's slew rate is 0.85V/µsec, which is therefore adequate for this 25-kHz waveform (**Figure 3**).EDN

REFERENCE

1 Terrell, David L, *Op Amps: Design, Application and Troubleshooting*, Butterworth-Heinemann, 1996.

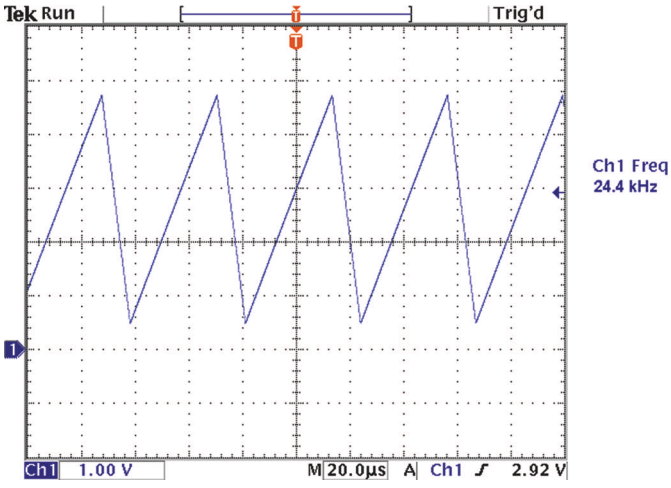

Figure 3 The op amp's slew rate is 0.85V/µsec, which is therefore adequate for this 25-kHz waveform.

Use a low-cost PWM ramp generator in switch-mode power supplies

Dwayne Reid, Edmonton, AB, Canada

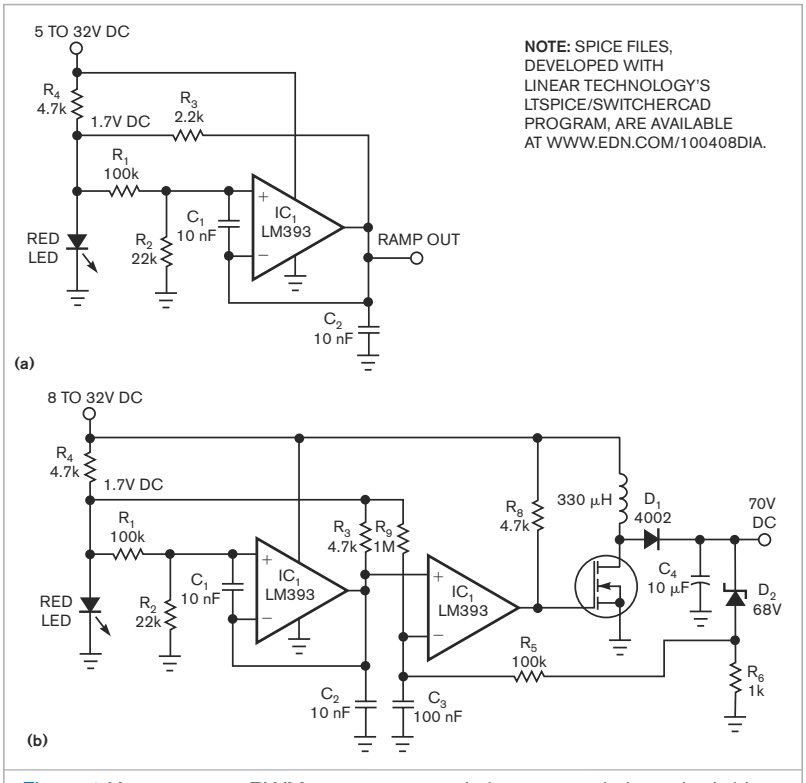

Figure 1 You can use a PWM ramp generator in low-cost switch-mode dc/dc power supplies (a). A 70V-dc upconverter employs the ramp generator (b).

The circuit in **Figure 1** shows a PWM (pulse-width-modulated) ramp generator that you can use in low-cost switch-mode dc/dc power supplies. Its supply voltage can range from 5 to 35V dc, and you can set the output-ramp amplitude of 0.3 to 1V. You can also set a minimum off time that lets you set a maximum 50% duty cycle for magnetic components that need duty-cycle limiting.

The ramp generator (**Figure 1a**) uses one-half of an LM393 dual comparator. The other half of the comparator is available to generate the PWM portion of the converter. The ramp amplitude and frequency depend on the reference. An ordinary red LED can act as a low-cost reference. Its forward voltage of approximately 1.7V is reasonably constant over indoor temperature ranges. The ratio of R_1 to R_2 sets the ramp amplitude relative to the reference, and R_1, R_2, and C_1 set the minimum off time. R_3 and C_2 establish a time constant, which sets the period. Note that the R_1, R_2, and C_1 network also affects the period. **Table 1** shows examples of various configurations.

Figure 1b, a 70V-dc upconverter, employs the ramp generator. You can easily configure it at any output ranging from the highest input voltage to whatever the FET can handle. This

Industry's leading innovator of integrated-mixer technology

- #1 share in BTS mixers
- Industry's most extensive wireless infrastructure mixer portfolio
  - Solutions covering all 400MHz to 4000MHz wireless standards
  - 41 cores designed specifically for infrastructure applications
  - 2.5x more devices than the closest competitor
- First to offer 100% monolithic, fully integrated up-/downconverters with LO buffers and switches
- First to offer innovative "Green Mixer" operating modes with up to 50% power savings

Complete portfolio of single-/dual-channel, pin-compatible solutions for all 400MHz to 4000MHz wireless standards

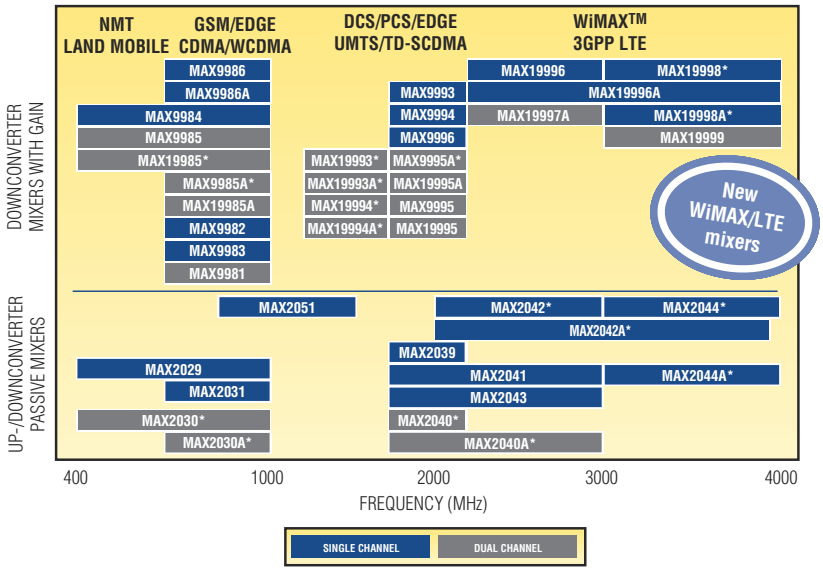

Fully integrated monolithic mixer designs with IF amps, LO buffers, switches, and baluns

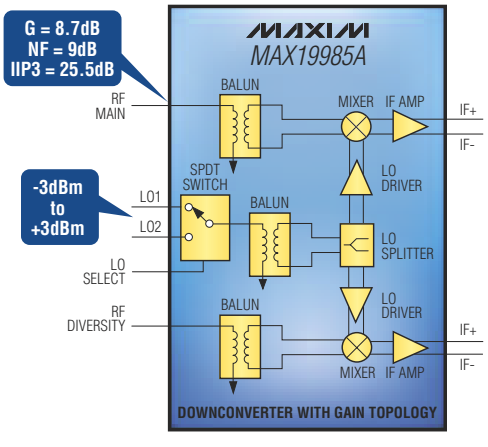

Largest portfolio of high IIP, low NF mixer cores

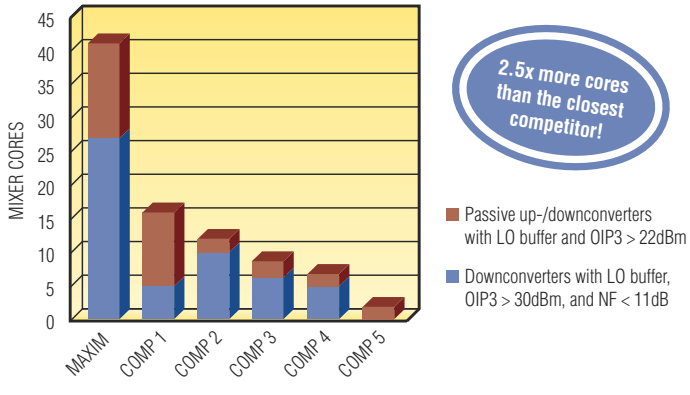

WiMAX is a trademark of the WiMAX Forum.
*Future product—contact factory for availability.

www.maxim-ic.com/wi-infra

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.

example uses a 330-µH inductor, but you can easily change that value by choosing the appropriate PWM frequency.

Note that the output FET does not turn on quickly, and it doesn't need to, but it does turn off quickly. You can enhance the turn-off speed by adding a 2N4403 PNP transistor between the output of the comparator and the pullup resistor. Connect the base to the comparator, the emitter to the FET gate, and the collector to ground. Add a 100Ω resistor from the base to the emitter.

The circuit has slow load-transient response, which you can adjust by altering the time constant that R_5 and C_3 form. Note that R_9 and R_5 form a voltage divider that ensures the lowest error voltage at the PWM comparator is above the ramp's lowest point. The converter cannot operate without R_9. EDN

TABLE 1 EXAMPLES OF CONFIGURATIONS

Frequency (Hz)	R_3 (Ω)	C_2 (nF)	R_1 (Ω)	R_2 (Ω)	C_1 (F)	Approximate duty cycle (%)
500	120k	100	1M	220k	10n	100
700	100k	100	100k	22k	100n	100
62k	4.7k	10	100k	22k	10n	95
100k	4.7k	10	100k	22k	100p	100
200k	2.2k	3.3	100k	22k	3.3n	60
200k	1.8k	10	100k	22k	100p	95
400k	2.2k	1	100k	22k	10p	95
400k	5.6k	2.2	100k	22k	470p	50

MSP430's port-interrupt-request logic helps debounce contacts

Richard Neubert, Manchester, NH

Contact debouncing requires monitoring an input and waiting for it to stop toggling or at least establish that it's definitely switching from its initial state. You can use either analog or digital filtering plus hysteresis to accomplish contact debouncing, but this approach uses a lot of resources, including parts, board space, and CPU time, when multiple inputs need to be conditioned. Alternatively, you can either detect just the first state change or sample the input at least twice as often as the contacts can bounce. Sampling at mechanical vibration frequencies must be avoided. Both methods also require time delays to ensure that the contacts have finished bouncing.

These 1-bit approaches are attractive when you must condition multiple contact inputs. The first method requires conservative delay setting to avoid resuming the edge monitoring before the last bounce, and it's unsuitable for re-

LISTING 1 INTERRUPT-FLAG CHECKER

```
#define SwitchMask  0x01          // Port 1 bit(s) connected to switches.
#define Ndebounce   3             // (# of consecutive checks w/o an edge to wait for valid state)-1
                                  // (must be nonzero)
void chk_sw(void);
static unsigned char debounced_sw;      // The debounced switch state(s)
                                  //(output)

void main()
{
        P1DIR = ~SwitchMask;      // Set unused pins to output mode, unless
                                  //switched to a peripheral.
        debounced_sw = P1IES = P1IN;    // Init so we catch the first state
                                  //change.
        P1IFG = P1IE = 0;         // Keep Port 1 interrupts disabled.
        _EINT();                  // Set General Interrupt Enable (if used).

        while (1)
        {
                .                 // Main program loop.
                chk_sw();         // Do this here if loop execution period
                                  // is fixed by a timer,
                .                 // o/w use a timer to run chk_sw(). If
                                  // chk_sw() is made a timer  ISR, then
                .                 // debounced_sw will be volatile here.
        }
}

void chk_sw()                     // Port1 interrupt request
                                  //checker, called periodically.
{
        static int db_count = Ndebounce;  // Initialize to force P1IES
                                  //update on first pass of chk_sw().

        if (P1IFG & SwitchMask)
        {
                P1IFG &= ~SwitchMask;     // Clear the switch interrupt flag(s).
                db_count = Ndebounce;     // Must see P1IFG not set for
                                  //Ndebounce+1 passes.
        }

        if (db_count)
        {
                P1IES = P1IN;     // Cock for opposite edge (repeated in case we
                --db_count;       //  miss an edge while doing it).
        }
        else                      // Finished debounce
        {
                debounced_sw = P1IES;     // Debounced switch output

                // Can put code here that you want to run only when switch
                // inputs are stable.
                // If you care about the switch inputs only in this block, you can read the debounced
                // read the debounced state directly from P1IES, no need to save it
                // in debounced_sw.
        }
}
```

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.

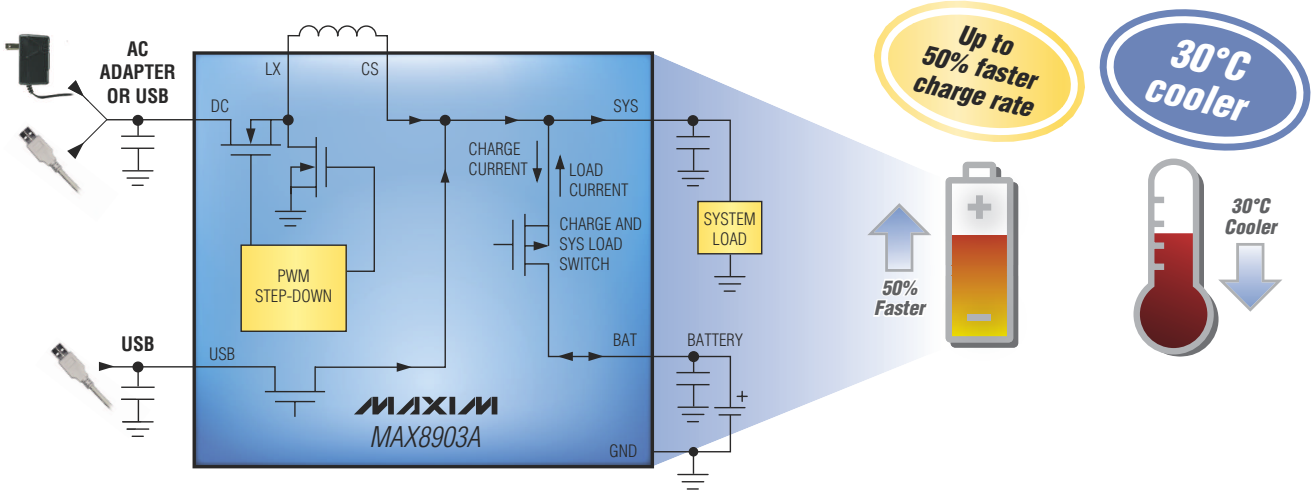

Feature	Competition	MAX8903A	MAX8903A Advantages
Input type (USB/DC)	Single	**Dual/single**	**Greater design flexibility**
Input range, absolute max	4.35V to 5.5V, 7V	**4.15V to 16V, 18V**	**No external overvoltage protection IC needed**
UVLO/OVLO	Yes/no	**Yes/yes**	**Better design robustness**
Switching 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.

jecting noise. The second method adds to the actual bouncing time only the delay necessary for bridging the longest quasistatic input state during bounce—not the longest duration of bouncing. When you implement this function in software, however, it can add substantial overhead for monitoring the input for further transitions when the system detects a transition.

The Texas Instruments (www.ti.com) MSP430-series microcontrollers have I/O ports with configurable interrupt logic for each bit. You can select the rising or falling edge of the bit as the trigger. Even with the interrupt disabled, the microcontroller can read its interrupt-request flag to determine whether an active edge has occurred. You can use this technique in place of high-rate sampling in software. You must periodically check the interrupt-request flag at a rate high enough only to keep the switch response delay to an acceptable time. Calling the routine at the frequency of a mechanical vibration isn't an issue; the interrupt logic monitors the switch between calls.

IF YOUR APPLICATION HAS NO TASK TO RUN EXCEPT IN RESPONSE TO A SWITCH, YOU CAN MAKE A VERSION IN WHICH THE CPU SLEEPS WITH THE PORT INTERRUPTS ENABLED.

You can use multiple switch inputs, provided that the switches either never change state simultaneously—for example, with a keypad—or you don't mind delaying response to a switch until all simultaneously changing switch inputs have settled. You would lose the sequence of switch operations in this case; in the debounced output, they would all change at once.

Listing 1 periodically calls function chk_sw() to check the interrupt-request flags and update the output value, debounced_sw. The time interval between calls times Ndebounce should be short enough to satisfy the required response time after the last bounce and longer than the longest time between transitions when the contacts are bouncing. In a noisy environment, making the delay too long is counterproductive because noise transients during the delay extend the delay.

If your application has no task to run except in response to a switch, it's fairly simple to make a version in which the CPU sleeps with the port interrupts enabled (P1IE=SwitchMask) and the P1IES bits set to the last input state when there is no switch input activity. A port interrupt-service routine must respond to the first input change to set P1IE=0 and set a timer to periodically call chk_sw() until chk_sw() resolves the input state. When reacting to a brief noise impulse, the CPU would wake up once for the port interrupt and Ndebounce+1 times to run chk_sw(). EDN

Power Manager II.
INTEGRATE POWER MANAGEMENT. LOWER COSTS.

Power Manager II features an innovative programmable core and precision analog sensors to reduce component count and increase reliability. Lattice programmability enables your power management design to meet changing requirements without major circuit redesign or a board re-spin.

Learn more about the cost-efficiencies of Power Manager solutions at:
latticesemi.com/powermanager

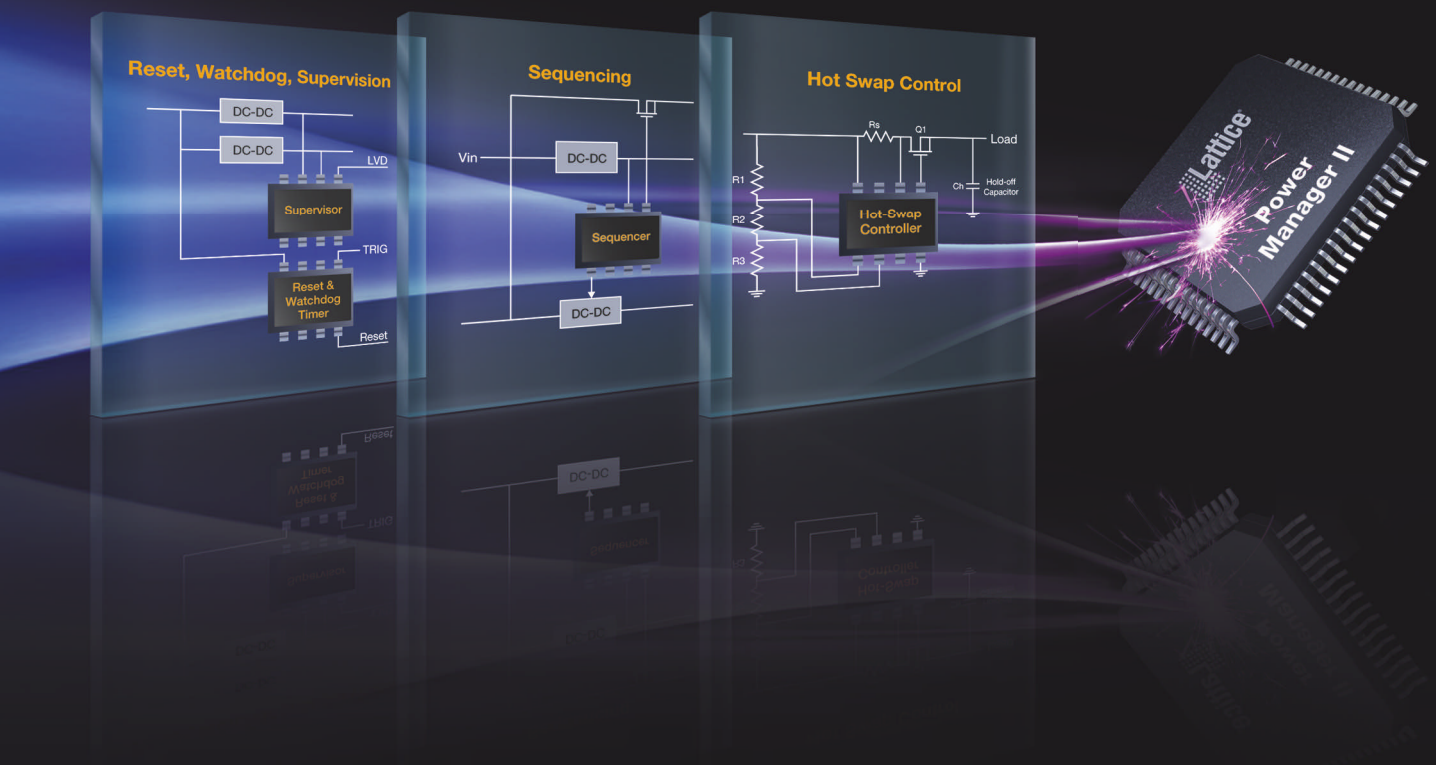

Features

- Charge pumps for hot swap MOSFET control
- PLD for flexible supply sequencing and reset generation
- 0.2% error (typ.) programmable comparators for fault detection
- Differential sensing for centralized control
- 10-bit ADC with I²C interface for system monitoring
- High volume pricing starts at <$1.00

And more...

Benefits

- Integrate multiple functions to reduce BOM cost
- Increase fault coverage
- Reduce footprint
- Increase reliability
- Wide application coverage
- Reduce risks through programmability

And more...

Lattice® Semiconductor Corporation

latticesemi.com/powermanager

Distributors step up embedded-systems support

If there's ever been a period in the design chain during which cost and time to market haven't been major factors, it's faded into distant memory. As engineers continue to pursue the best designs in the quickest way possible, distributors are stepping up their embedded-system support to help ease some of the pressure. In mid-February, for example, Avnet Electronics Marketing (www.em.avnet.com) announced that it had formed a sales division for embedded systems. In December 2009, Arrow Electronics (www.arrow.com) expanded its relationship with embedded-products supplier Kontron GmbH (www.kontron.de). Both changes are aimed at extending the channel's reach deeper into the embedded-system market.

"The cost today for proprietary system design continues to rise, while embedded solutions offer early access to key technology and help our customers get their products to market faster—at a lower cost," says Jeff Ittel (**photo,** left), Avnet's senior vice president and head of Avnet Electronics Marketing's embedded unit. "In the past few years, we have seen increased demand for these products from our customers across multiple industries, and, as that trend continues, we want to intensify our own focus in this area," he adds.

OEMs are turning to the channel for a number of reasons, executives say. Using off-

the-shelf, or merchant, embedded boards helps reduce OEM development costs, accelerates time to market, and contributes toward product differentiation. "It's a classic make-versus-buy decision," says Andrew Femrite (**photo,** right), head of Arrow Electronics Engineering Solutions Center. "There's a strong incentive to use commercial off-the-shelf boards."

Although manufacturers design off-the-shelf embedded boards to be standard, the market involves a vast variety of boards and a lot of complexity. Different board characteristics lend themselves better to specific applications, explains Troy Smith, director of Intel Corp's (www.intel.com) Embedded and Communications Alliance. The channel's span and scope, particularly in regard to its supplier base, can reduce some of that complexity.

The channel's size and scope are also working to the benefit of suppliers. "Our embedded-systems business is one of the top growth opportunities in the corporation," Smith says.

OEMs are also using embedded designs to extend the life of their products and are looking to the channel for guidance. "Many components, particularly processors, are targeted at the PC market and have a shorter life cycle than embedded designers feel comfortable with," said Arrow's Femrite.

"Designers of embedded applications, such as kiosks or medical equipment, don't want to switch products every couple of years," says Avnet's Ittel.

Through their supplier relationships, distributors have a lot of visibility into product-development plans. Chip makers such as Intel share their product road maps as part of their standard training and support. "We recognize that many market segments require longer product life cycles, so we have our own embedded road map where we commit to a seven-year life span," says Smith.

Another competitive advantage of off-the-shelf boards is that they enable OEMs to focus on product differentiation. OEMs are distinguishing their products now through software, a highly customized part of the system, or a combination of discrete components, explains Femrite. "If we can steer them toward the right development system, the right components, or the right partnerships, they can focus more on how they differentiate their products."

Suppliers and distributors are well-positioned to take advantage of embedded-market trends, according to market research company Venture Development Corp (www.vdcresearch.com). Prior to the effects of the 2009 economic recession, VDC had forecast that the merchant embedded-board market would grow by as much as 8 to 10% through 2012. The forecast has since been revised downward but still reflects positive growth (**Table 1**). Certain vertical segments, including digital signage, infotainment, energy, and power, should outpace others.

–by Barbara Jorgensen

TABLE 1 BOARD-MARKET GROWTH

Year	Market size ($ millions)	Growth (%)
2007	4694.8	NA
2008	5087.2	8.36
2009	5250.5	3.21
2010	5275.2	0.47
2011	5597	6.1
2012	6070.5	8.46

Source: Venture Development Corp

Introducing ...

SR1 Audio Analyzer

- **−110 dB THD + N** *(1 kHz, 2 V, 22 kHz BW)*
- **± 0.008 dB flatness**
- **200 kHz system bandwidth**
- **24-bit / 192 kHz digital audio**
- **<600 ps residual jitter**
- **Cross-domain analysis**

SR1 ... $6900 *(U.S. list)*

Introducing *SR1* Dual-Domain Audio Analyzer
— AP 2700 class performance at $6900.

SR1's outstanding specifications and rich suite of measurements make it ideal for analog, digital and cross-domain audio signal analysis.

Standard measurements include Level, THD + N, Harmonic Distortion, IMD, FFT, Frequency Response, Multi-Tone, Crosstalk, Histogram, Jitter Amplitude & Spectrum, and more. Hardware options include a digital audio carrier digitizer with full-color eye diagrams and carrier spectra, multi-channel I/O switchers, and an atomic rubidium system clock.

SR1 offers uncompromised performance at a fraction of the price of the competition.

Call us or visit our web site for further details.

thinkSRS.com

SRS **Stanford Research Systems**
Ph: (408)744-9040 www.thinkSRS.com

productroundup

DISCRETE SEMICONDUCTORS

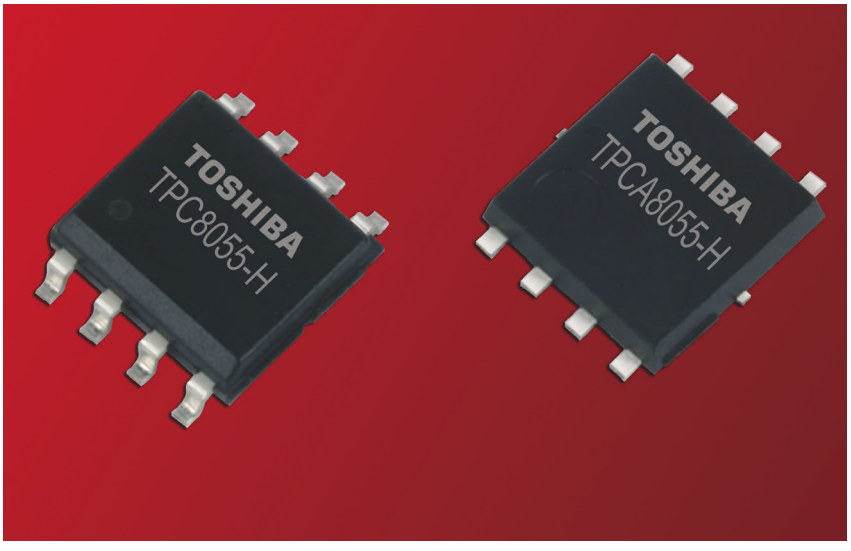

Fast-switching MOSFET aims at dc/dc conversion

The first two MOSFETs in the company's UMOS VII-H series and suiting synchronous dc/dc conversion, the TPCA8055-H and the TPC8055-H power MOSFETs use the vendor's trench process in the UMOS VII-H fast-switching series. The TPCA8055-H provides a 2.2-mΩ on-state resistance, a 10V gate-to-source voltage, and a 30V drain-to-source voltage. The TPC8055-H comes in an industry-standard 5×6×1.6-mm SOP-8 package, and the TPCA8005-H comes in a low-profile 5×6×0.95-mm Toshiba SOP Advance housing. Prices for the TPCA8055-H and the TPC8055-H power MOSFETs start at 75 cents.

Toshiba America Electronic Components, www.toshiba.com/taec

Integrated device enables 1-MHz frequencies

The SiC762CD DrMOS device integrates PWM high- and low-side N-channel MOSFETs, a full-featured MOSFET-driver IC, and a bootstrap diode. Compliant with the vendor's DrMOS specification for voltage regulators in servers and desktop computers, the device enables 1-MHz operating frequencies at 92% efficiency. Operating over a 3 to 27V input-voltage range, the device delivers 35A continuous output current. The integrated MOSFETs have output voltages of 0.8 to 2V and a nominal input voltage of 24V. An advanced gate-driver IC accepts a PWM input from the controller, converting it into the high- and low-side MOSFET gate-drive signals. The PWM input aims at use with controllers with tristate-PWM-output functions. Available in a PowerPak MLP-40 package, the SiC762CD DrMOS device costs $2.65.

Vishay Intertechnology, www.vishay.com

OptiMOS voltage-regulation MOSFET and DrMOS families have 93% efficiency

The 25V OptiMOS device suits voltage regulation in power supplies for computer servers and telecommunications and data-communications switches. The new MOSFETs are also integrated into the TDA21220 DrMOS devices that are compliant with the Intel DrMOS specification. The discrete devices come in SuperSO8, CanPak, and 3.3×3.3-mm S3O8 packages. The S3O8 package enables a six-phase converter. The multichip TDA21220 package integrates two OptiMOS transistors and a driver IC. In a SuperSO8 package, a 1-mΩ, 25V OptiMOS device costs $1.40 (2000). The TDA21220 DrMOS devices cost $2.17 (2000).

Infineon Technologies, www.infineon.com

N-channel MOSFETs have a variety of package options

Adding 12 devices to the vendor's MOSFET portfolio, the NTP641x and NTB641x 100V N-channel power MOSFETs target designs requiring voltage-overstress protection from unclamped inductive loads. Features include a 500-mJ avalanche rating, a 13-mΩ on-resistance, and a 76A current capability. Operating over a −55 to +175°C temperature range, the MOSFETs come in lead-free, ROHS-compliant TO-220, D²Pak, DPak, and

IPak packages. Prices for the devices range from 92 cents to $1.90 (10,000).
On Semiconductor, www.onsemi.com

Power-rectifier diode provides low forward-voltage drop

The STPS50U100C high-efficiency power-rectifier diode suits use in the output of power supplies for adapters, desktop PCs, servers, and TV and video products. The forward-voltage drop is the primary source of power loss in power Schottky diodes in secondary rectification. The power rectifier reduces this voltage drop using a driver to a typical figure of 0.64V at 25A and 0.38V at 5A. Additional features include 200-mA maximum leakage current, 50A maximum average forward current, and 250A surge nonrepetitive forward current. Available in TO-220 and I²Pak packages, the device costs $1.50 (1000).
STMicroelectronics, www.st.com

COMPUTERS AND PERIPHERALS

Mouse family offers high tracking precision

The $10.99 M-300 retractable mini-notebook mouse measures 74 mm and has a retractable cable. The $14.99 M-500G gaming mouse has a blue fluorescent light that shines in dark environments. The $24.99 M-900LS laser notebook mouse features 2.4-GHz digital cordless technology and delivers interference-free operation. The $29.99 M-905BT Bluetooth notebook mouse connects instantly to any Bluetooth-enabled notebook, eliminating the need to tie up a USB port. It works from a distance of 10m.
Agama, www.agamazone.com

External 3.5-in. drive boasts 2-Tbyte capacity

The recently redesigned desktop SimpleDrive offers capacities of 500 Gbytes, 1 Tybte, and 2 Tbytes and operates at 7200 rpm. The drive is available on the vendor's SimpleTech Pro Drive, SimpleTech Duo Pro Drive, and internal hard-drive kits. The SimpleTech Pro Drive has USB 2.0, FireWire 400, FireWire 800, and 3-Gbps eSATA interface options. You can stack multiple units or stand them upright. The 500-Gbyte, 1-Tbyte, and 2-Tbyte devices sell for $119.99, $169.99, and $299.99, respectively. The SimpleTech Duo Pro Drive uses RAID 0 or 1. RAID 0 targets use in home-video editing and graphics-intensive files. Interfaces include eSATA and USB 2.0. The 1-, 2-, and 4-Tbyte devices sell for $199.99, $299.99, and $499.99, respectively. The internal-hard-drive kits each feature a drive, installation instructions, a SATA cable, and mounting screws. They come in 5400- and 7200-rpm versions with a cache buffer as large as 32 Mbytes. The 2-Tbyte, 7200-rpm kit sells for $249.99.
Hitachi Global Storage Technologies, www.hitachigst.com

EDN ADVERTISER INDEX

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

HAROLD STILTNER, INGERSOLL RAND CLIMATE SOLUTIONS

Can't put a Band-Aid on a boomerang

DANIEL VASCONCELLOS

While working for a company specializing in sensors, I received a board to test—a signal-conditioning module for a strain-gauge torque transducer that bolted directly to the hub of a drive motor. The board had been designed with size and cost reduction as key directives. The schematic showed that the input amplifier stage comprised a single operational amplifier acting as a current-input amplifier. The sensor comprised four 350Ω strain gauges acting as a Whetstone bridge, equivalent to driving the inputs with Thevenin sources with 175Ω source resistance. While the board was on the bench, everything seemed fine; manually applying torque corresponded to changes in the analog output signal.

My mechanical counterpart discovered during the calibration process that the transducer couldn't pass the 60-second stability test that acquired 1 sample/sec to confirm the output was clean of noise or drift. After a couple of failed attempts, I rolled an oscilloscope over to the device and began to monitor the output. There was "bobble" all right, but it didn't look random enough; it was information—not noise. Attaching a piezoelectric speaker to the analog output revealed that we were picking up the news broadcast of a local AM-radio station.

While troubleshooting, I found a grounding problem in the calibration system; correcting the problem eliminated the radio broadcast, and I was able to explain it all away. I did not know that what I had just tossed was a boomerang.

For a display at an upcoming trade show, the sales manager wanted to use the sensor with a variable-frequency-drive motor—a common application. After powering up the drive system, I realized that I was no longer dealing with a little bobble; I was now witnessing noise that swung to 75% of the output range. The output signal was essentially useless. I couldn't explain away this problem.

Although the transducer's metal body was isolated from the strain gauges and the wiring by hundreds of gigohms, a stray capacitance of approximately 140 pF had sneaked directly from the op amp's inputs to the transducer's body and, thus, to the motor housing. I saw three major issues: saturation, implying that the coupled input spikes were exceeding the amp's common-mode range; RF rectification in the op amp's inputs; and the implication that the gain of the signals coupling down the pathway would increase with frequency.

To block high-speed transients and other RFI, I added 100-µH chokes in series with the amplifier inputs, reducing the noise by more than 40 dB. Careful grounding of the drive yielded further improvement. The Band-Aid fix worked for the trade-show booth, and the company ultimately redesigned the module with an instrumentation amplifier with input RC networks to roll off common- and normal-mode signals to within the bandwidth of interest before the signals reached the instrumentation amp's input pins. The RC networks effectively attenuated the transients to an amplitude that did not cause saturation.

Looking back, two key points pop out. The first is that the design with the lowest cost and the fewest parts is not necessarily the best. The second is that, although I "solved" the radio-station problem with an easy fix, I should have questioned why the module was sensitive to that condition in the first place. EDN

Harold Stiltner is a hardware engineer for Ingersoll Rand Climate Solutions (LaCrosse, WI).

⊞ www.edn.com/tales

THE EMBEDDED DEVELOPER

Microsoft

Windows Embedded

BUILT WITH **WINDOWS 7 TECHNOLOGIES**

Vol. 5

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

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

Innovative Intelligent Integration

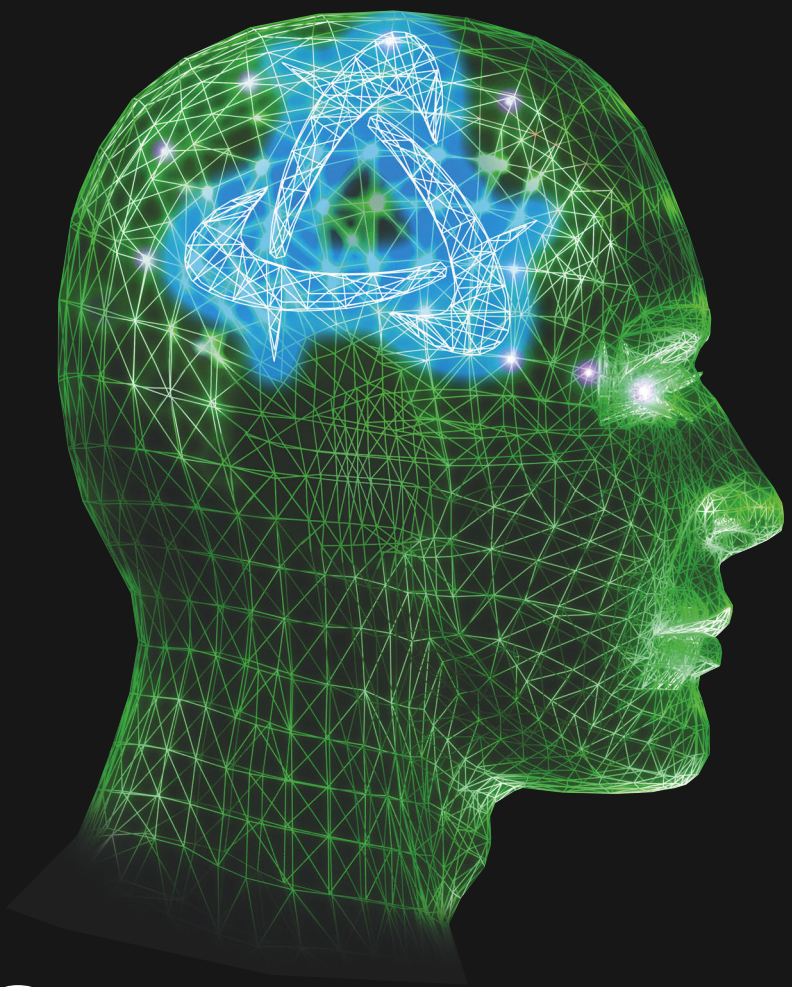

SmartFusion™

FPGA + ARM®Cortex™-M3 + Programmable Analog

Actel POWER MATTERS

Get Smart, visit: www.actel.com/smartfusion

Visit us at ESC Silicon Valley, Booth #930