Reed Business Information (RBI™)

EDN®

VOICE OF THE ENGINEER

MAR 4

Issue 5/2010

www.edn.com

SPECIAL SECTION

INNOVATORS Pg 45

TILT OR NEW GAME?

Page 28

PRECISELY MEASURE SETTLING TIME TO 1 PPM

Page 20

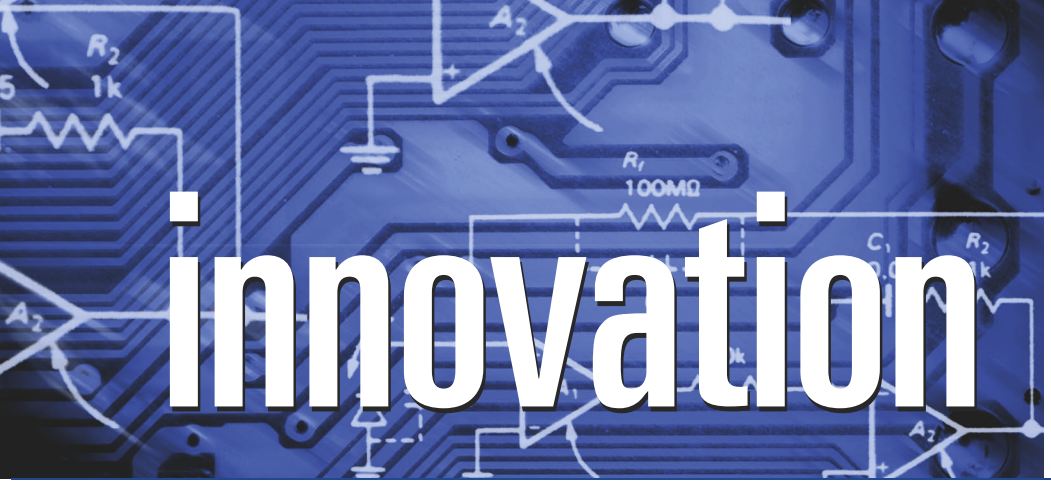

CS4353
Thermally Enhanced
4 mm x 4 mm
QFN Package
(shown at actual size)

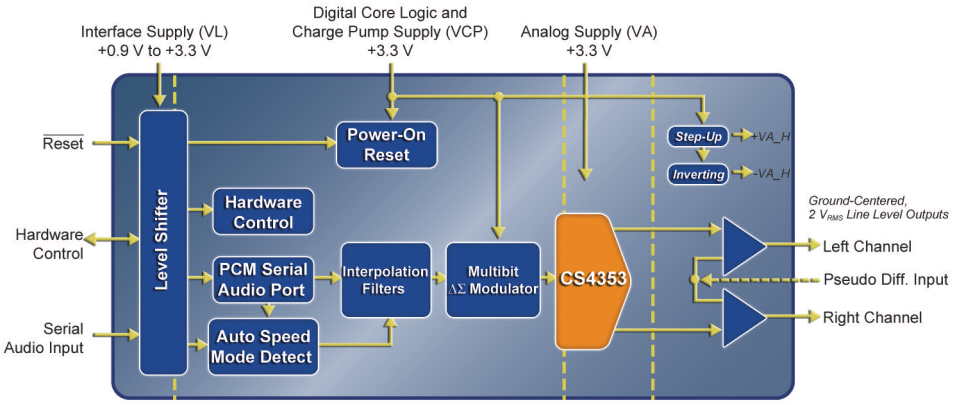

Finally! A No Compromises Digital-to-Analog Converter

CS4353 SIMPLIFIES ANALOG OUTPUT CIRCUITRY, PROVIDING SUPERIOR AUDIO PLAYBACK

Cut the clutter! When it comes to digital-to-analog converters, Cirrus Logic means business. The CS4353 is the only audio DAC to feature an on-chip 2.1 V_{RMS} line driver from a single 3.3 V power supply, allowing a 2 V_{RMS} voltage after the RC filter. The new architecture not only simplifies the design of analog output section circuitry but it also provides superior quality audio playback with the highest dynamic range and lowest pops and clicks of its kind. This is especially important in applications such as Blu-ray Disc® players, DTV and set-top boxes as they allow for the elimination of multiple system components and lower implementation cost.

CS4353 FEATURES

- 106 dB Dynamic Range, –93 dB THD+N
- Single-Ended Ground Centered Analog Architecture
- No DC-Blocking Capacitors Required
- Click-Free Stereo Output
- Selectable 1 or 2 Vrms Full-Scale Output
- High PSRR
- Small 4 mm x 4 mm QFN Package

FREE SAMPLE!

Request a free sample of the CS4353
WWW.CIRRUS.COM/4353edna

NORTH AMERICA
+1 800-625-4084

ASIA PACIFIC
+852 2376-0801

JAPAN
+81 (3) 5226-7757

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

LEARN MORE AT
www.cirrus.com

Cirrus Logic. We make it easier for you.

HOW MANY SUPPLIERS DOES IT TAKE TO MAKE A MASTERPIECE?

Digi-Key® CORPORATION

Authorized distributor for more than 440 of the world's most trusted manufacturers.*

THE INDUSTRY'S BROADEST PRODUCT SELECTION AVAILABLE FOR IMMEDIATE DELIVERY.

1.800.344.4539
WWW.DIGIKEY.COM

*Digi-Key is an authorized distributor for all supplier partners. New products added daily.

NEDA Supplier Authorized Distributors

EDN contents

3.4.10

Tilt or new game?

Precisely measure settling time to 1 ppm

INNOVATORS

pulse

DESIGNIDEAS

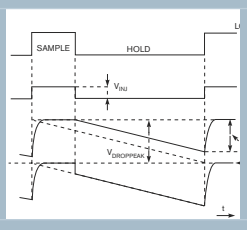

▶ Send your Design Ideas to edndesignideas@reedbusiness.com.

Microsoft

"We HAVE TO keep track of sales and distribution of over 18,000 pieces of equipment globally...

AND get the right parts, to the right place, at the right time

AND automate more processes

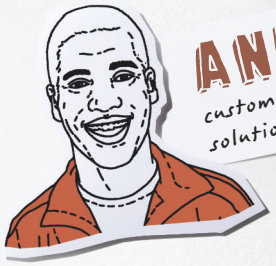

AND customize easier and cheaper solutions for improved ROI

AND reduce time spent on service calls to rectify errors

AND improve communications with our dealers

AND Microsoft Dynamics CRM lifted a heavy load for our staff by managing our relationships more efficiently."

Microsoft® Dynamics® CRM integrates information from different sources and puts it in one place, at one time, so everyone in your company sees the information needed to make better decisions faster. It's simple for your sales and support people to use, and it's ready to fit your company right away. So you can spend less time on service calls and more time building stronger relationships.

To learn more about the efficiencies Microsoft Dynamics CRM can create for your business, go to **microsoftdynamics.com/manufacturing**

Microsoft Dynamics® CRM

Because it's everybody's business

Innovative Intelligent Integration

SmartFusion™

FPGA + ARM®Cortex™-M3 + Programmable Analog

Actel® POWER MATTERS

Actel POWER MATTERS

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

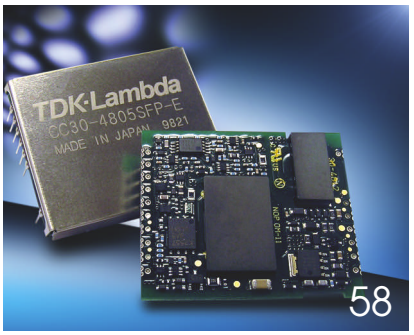

DEPARTMENTS & COLUMNS

EDN online contents www.edn.com

ONLINE ONLY

Check out this Web-exclusive article:

Acceleration of program execution
Look-up tables trade between processing resources and execution time to speed up execution of computational processing.
→www.edn.com/article/CA6717542

Designing with LEDs

Lighting for the 21st century

Practical, inexpensive HB LEDs (high-brightness light-emitting diodes) are here. Now, what can we do with them, and what will be their impact on electronics and consumers? Read more online.
→**www.edn.com/leds**

INNOVATION
VOTE NOW!

Tune in to www.edn.com/innovation20 to see a list of the innovative engineers and products our editors selected as finalists in the 20th annual *EDN* Innovation Awards and then use the online ballot to cast your vote for the winners in each category. Don't delay! Voting ends March 19.
→**www.edn.com/innovation20**

EDN® (ISSN#0012-7515), (GST#123397457) is published semimonthly, 24 times per year, by Reed Business Information, 8878 Barrons Blvd, Highlands Ranch, CO 80129-2345. Reed Business Information, a division of Reed Elsevier Inc, is located at 360 Park Avenue South, New York, NY 10010. John Poulin, CEO/CFO, RBI-US; Jeff DeBalko, President, Business Media; Jane Volland, Vice President, Finance. Periodicals postage paid at Littleton, CO 80126 and additional mailing offices. Circulation records are maintained at Reed Business Information, 8878 S Barrons Blvd, Highlands Ranch, CO 80129-2345. Telephone (303) 470-4445. POSTMASTER: Send address changes to EDN®, PO Box 7500, Highlands Ranch, CO 80163-7500. EDN® copyright 2010 by Reed Elsevier Inc. Rates for nonqualified subscriptions, including all issues: US, $179.99 one year; Canada, $229.99 one year (includes 7% GST, GST#123397457); Mexico, $229.99 one year; air expedited, $399.99 one year. Except for special issues where price changes are indicated, single copies are available for $10 US and $15 foreign. Publications Agreement No. 40685520. Return undeliverable Canadian addresses to: RCS International, Box 697 STN A, Windsor Ontario N9A 6N4. E-mail: Subsmail@ReedBusiness.com. Please address all subscription mail to EDN®, 8878 S Barrons Blvd, Highlands Ranch, CO 80129-2345. EDN® is a registered trademark of Reed Elsevier Properties Inc, used under license. A Reed Business Information Publication/Volume 55, Number 5 (Printed in USA).

BY RON WILSON, EXECUTIVE EDITOR

Toyota, drive by wire, and our failure to learn from experience

I see from the morning news that Toyota's adventure into the world of embedded software is going badly. The company's second attempt to find a quick fix for unintended acceleration in its conventionally powered vehicles is barely under way, and evidence is already emerging that the underlying problem is likely in the engine controller, not in the pedal's mechanical assembly. Now we hear from Japan that the Prius, Toyota's golden child, has a problem with its brake-by-wire control system.

Decades ago, Audi accidentally introduced drive by wire with its advanced cruise control on the Audi 5000. The cars were allegedly subject to spontaneous acceleration, a problem the company blamed on operator error. At the time, researchers at another European high-end auto company claimed to have uncovered a problem in Audi's engine-control firmware and reproduced the acceleration without requiring a driver to mistake the gas pedal for the brake. The ensuing liability litigation, however, extinguished all hope of diagnosing and documenting the problem so that the rest of the real-time-software community could avoid it.

All this came to mind when I attended a panel on achieving quality closure at last month's DesignCon in Santa Clara, CA. Despite the subject of the panel—achieving quality closure—the issue of software sat like an elephant in the corner of the room, awaiting notice. One of the panelists pointed out that the most serious quality problem in IC designs now is not the quality closure on the hardware but the integrity of the firmware and software that will run on the chip. There simply is no systematic approach to ensuring the quality of an integrated hardware/software system.

Along came C, Unix, and the cult of the hobby programmer, and the entire notion of formal correctness vanished under a smokescreen of hacking.

This situation is a tragedy. Work was well under way 30 years ago on the problem of formally proving software correctness. One company had designed a completely deterministic microprocessor that made it possible to mathematically prove all of the possible trajectories of a code set. Computer scientists such as Edsger Dijkstra were making strides in a method for creating formally proven software. But along came C, Unix, and the cult of the hobby programmer, and the entire notion of formal correctness vanished under a smokescreen of hacking.

Now, after decades invested in metrics-driven verification, formal verification, and methodology management, designers find that their chips don't work as expected because the software is still being "verified" by feeding it test cases until the schedule expires. Consumers find that their cars run into these problems for the same reason, and the press blames the problem on "electronics."

Once again, as in Audi's day, it is safe to conclude that a gag order as part of a class-action settlement will screen whatever accurate diagnostic work takes place on the Toyota problems so that no one in the industry can benefit from what Toyota engineers learn. In that way, we can repeat the situation with the next generation of software-governed systems, a new set of executives can avoid blame for the tragedies, and a new set of lawyers can make their fortunes from the resulting litigation.

The only parties in this little tragedy with an interest in improving the state of the art are the engineers, whom no one will consult, and the victims, whom the lawyers will silence. It would be better for everyone if it were a principle of civil law that, when a failure inflicts damage, the vendor and independent parties must place all of the diagnostic information they find into the public domain, and the courts may not use this information to assess or assign damages. Such a notion might somewhat restrict the income opportunities of litigators, but it would unquestionably assist the engineering community in learning from its mistakes. **EDN**

Contact me at ronald.wilson@reedbusiness.com.

GLOBALFOUNDRIES

The First Truly Global Foundry

cadence® TOSHIBA SYNOPSYS

ST Mentor Graphics® VIRAGE LOGIC IBM

MAGMA Infineon ARM

QUALCOMM® AMD MEDIATEK

Congratulations to the new GLOBALFOUNDRIES,
the merged operations of GLOBALFOUNDRIES and Chartered.

To our continued mutual success
from your Global Customers and Partners

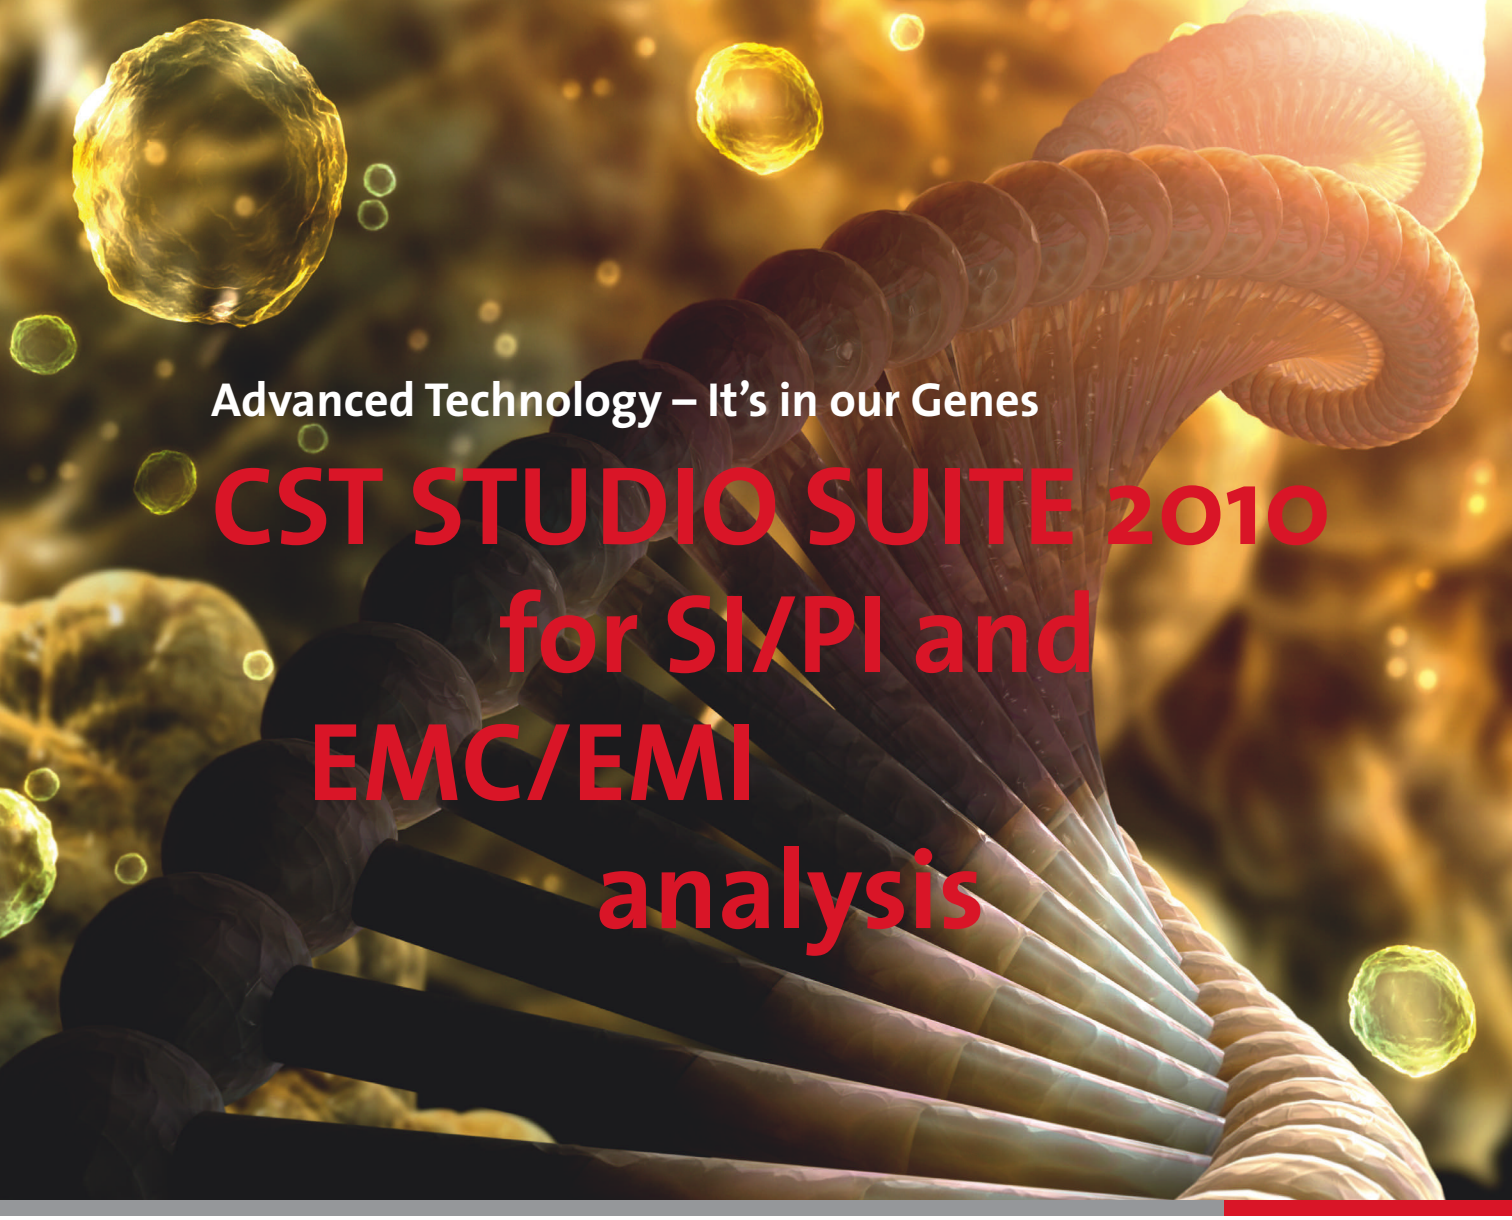

Advanced Technology – It's in our Genes

CST STUDIO SUITE 2010 for SI/PI and EMC/EMI analysis

Next generation 3D electromagnetic simulation

The best available tool for 3D EM simulation has evolved into version 2010. With a user friendly interface, easy data exchange to and from other software tools, a choice of first class solvers and world class post-processing tools, you can leverage the latest developments in 3D electromagnetics to bring designs to market faster and with lower risk.

Choose CST STUDIO SUITE 2010 – complete technology for 3D EM.

CST

CHANGING THE STANDARDS

EDN

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

CONSULTING EDITOR
Jim Williams, Staff Scientist,
Linear Technology

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

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

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

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

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

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

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

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

Reed Electronics Group

EXTREME EVOLUTION

Introducing the RX Family from Renesas.
The next-generation microcontrollers for communicating anytime, anywhere, and with anyone and anything.

- **Extremely high performance: 1.65DMIPS/MHz, up to 200MHz**
- **Extremely small code: up to 30% smaller than previous Renesas MCUs**
- **Extremely low power consumption: 30µA/MHz (typical) for the core**
- **Extremely high integration: with up to 4MB of extremely fast embedded Flash**

Stay tuned. The evolution continues.

Check the website for details.

www.america.renesas.com/ReachRX/a

Renesas Technology Corp.

Everywhere you imagine. RENESAS

pulse

INNOVATIONS & INNOVATORS

8 billion-transfer/sec PCIe 3.0 test platform integrates new digital-test console

Agilent Technologies has introduced a test platform that includes a digital-test console, which provides a complete integrated protocol analyzer and exerciser for one- to 16-lane implementations of the PCIe (Peripheral Component Interconnect Express) 3.0 specification. PCIe 3.0 is currently under development within the PCI-SIG (PCI Special Interest Group). The console offers the industry's largest capture buffer and fastest download interface.

The company designed the console to assist with and provide accurate test results during the development of PCIe 3.0 and products that support and incorporate it. The transition to PCIe 3.0 requires test engineers and validation labs to address the new specification's higher bus speeds and backward-compatibility requirements and to adapt to changes in the bus's encoding scheme and several new advanced protocol features. The console provides an ESP (equalization-snoop-probe) interface, which reliably captures data at 8 billion transfers/sec; it also incorporates an LTSSM (link-training-sequence-state-machine) exerciser for validating new encoding- and protocol-state-machine designs and a flexible GUI (graphical user interface) to aid in debugging the advanced protocol.

The test platform includes a large collection of tools for testing and validating PCIe 3.0 requirements from the electrical layer to the protocol layer. These tools include the 90000-Series Infiniium oscilloscopes, which integrate de-embedding capabilities as well as CTLE (continuous-time-linear-equalization) software that ensures the electrical compliance of future PCIe 3.0 devices; the 86100C Infiniium DCA-J (digital-communication analyzer/jitter), which provides one-button ease of use and performs precision waveform analysis as well as jitter, TDR (time-domain-reflectometer), and S (scattering)-parameter measurements; and the N4903B J-BERT (jitter/bit-error-ratio tester), a high-performance serial BERT that performs complete jitter-tolerance testing for accurate receiver characterization.

The platform also includes the N5990A test-automation software for complete automation of PCIe-compliance testing using the J-BERT, ParBERT (parallel BERT), and other Agilent instruments and the ADS (Advanced Design System), a software suite that includes a complete set of simulation tools ranging from frequency-, time-, numeric-, and physical-domain simulation to electromagnetic-field simulation for design characterization and optimization. The average system price is $100,000.

—by Dan Strassberg

▷**Agilent Technologies**, www.agilent.com/find/pci-express, www.agilent.com/find/pciexpress_backgrounder.

FEEDBACK LOOP

"SIP (system-in-package) technologies can provide the best of both worlds but with one important caveat: In the end, reliability is a function of the weakest link. So, for high-volume applications, SOC (system on chip) is a good way to go."

—Reader Dave Anderson, in *EDN*'s Feedback Loop, at www.edn.com/article/CA6715766. Add your comments.

This digital-test console is the nerve center of a platform for design and verification testing of one- to 16-lane implementations of the 8 billion-transfer/sec PCIe 3.0 bus.

Aeroflex announces LTE simulator as In-Stat calls LTE the 4G choice

To speed up real-world testing of mobile handsets for LTE (long-term-evolution) networks before network deployment, Aeroflex has introduced a one-box test system for cell-phone signal-fading simulation.

In-Stat analysts note that, although WiMax (worldwide operability for microwave access) appeared to be a competitor for 4G, the battle is now largely over, and the success of 3G, HSPA (high-speed-packet-access), and HSPA+ networks will impede LTE's deployment as mobile operators leverage their installed infrastructure (see "Getting to 4G through design and test," *EDN*, April 9, 2009, pg 38, www.edn.com/article/CA6648778). "LTE still has several glaring issues," cautions Allen Nogee, an In-Stat analyst. These issues include lack of spectrum, SNR (signal-to-noise ratio), and the lack of an established patent and royalty pool. "The shift toward 4G LTE will be gradual and protracted," he says.

Nevertheless, LTE deployments will begin this year, with North America and Asia/Pacific being the first regions to roll out. As deployment progresses, the demand is growing to meet all LTE requirements, including fading profiles that 3GPP (Third Generation Partnership Project) specifies in 36-521-1. Aeroflex's 7100 Series test platform provides fading simulation that meets or exceeds all 3GPP requirements, as well as flexibility in allocating cells and fading taps for LTE user equipment without the need for manual reconfiguration. The series features fully repeatable test scenarios with the fading simulator. These scenarios include the emulation of dynamic environments and realistic and accurate testing of MIMO (multiple-input/multiple-output) situations.

> "The simulator supports all 3GPP fading profiles, allowing users to determine whether their devices conform to 3GPP test specifications.

Fading simulators combine with noise generators to modify RF signals that the LTE-system simulator transmits and emulate degradations that real-life obstacles, such as buildings and foliage, introduce into the radio channel. For LTE developers who must profile signal fading on mobile handsets to meet 3GPP requirements before a network is available, the 7100 Series fading simulator allows engineers to perform realistic signal-fading simulations in a reliable and repeatable lab environment.

With the LTE future in mind, the 7100 Series fading simulator supports all LTE bandwidths to 20 MHz with a frequency range as high as 6 GHz. The simulator supports all 3GPP fading profiles, allowing users to determine whether their devices conform to 3GPP test specifications.

Employing Aeroflex's RF and baseband technology, the 7100 Series digital-radio test set supports both RF parametric and protocol testing for LTE-terminal devices. Targeting the R&D market—from components to handsets—the series simulates a network from the physical layer to the core-network IP infrastructure.

A typical system with two RF carriers and 2×2 MIMO sells for approximately $100,000. Customers can license and install the $10,000 fading-simulator software option and on units in the field. For more on these developments, go to www.edn.com/pa.

—by Rick Nelson

▷**Aeroflex**, www.aeroflex.com.

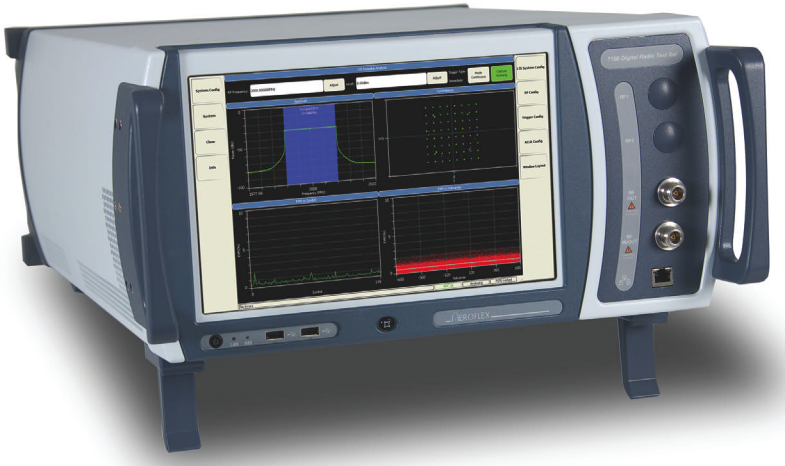

Aeroflex's 7100 Series test platform provides fading simulation that meets or exceeds all 3GPP requirements, as well as flexibility in allocating cells and fading taps for LTE user equipment.

DILBERT By Scott Adams

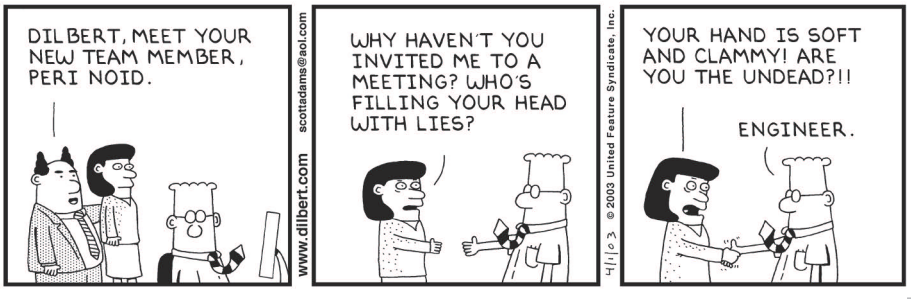

All of the Savings

None of the Loss

HV Bus Converters Enable AC-to-Load Efficiency Increase of 8.7%

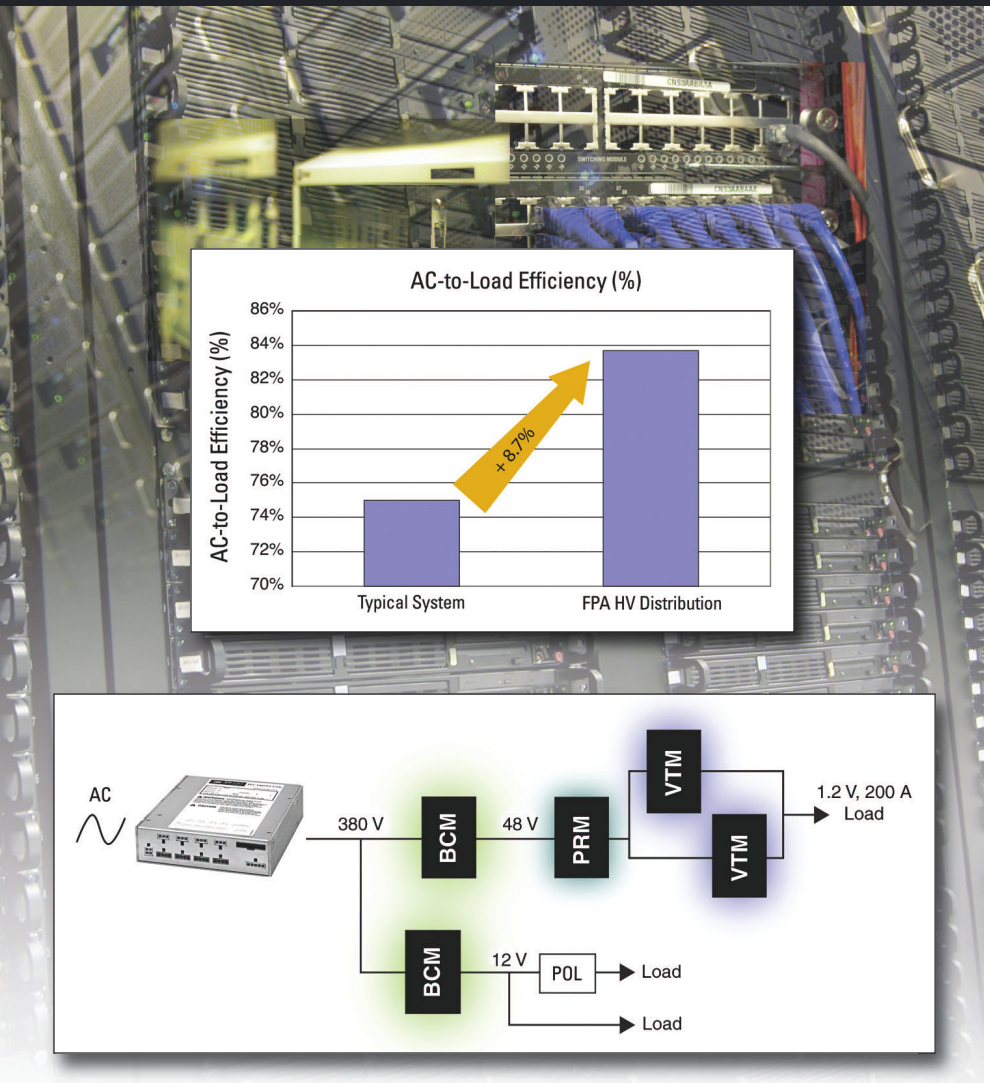

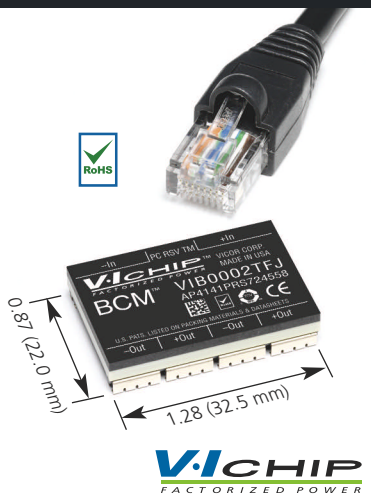

V·I CHIP
FACTORIZED POWER

- **High power density: Up to 1120 W/in³**
- **380 V to 1.2 V, 200 A conversion in 4.5 in²**
- **Low noise ZVS – ZCS MHz switching**
- **>95% efficiency**

For high power computing and telecom systems, HV BCM™ bus converters are high efficiency, low noise, high power density system building blocks that enable the highest possible distribution efficiency.

As part of a Factorized Power Architecture solution, the HV BCMs efficiently power sub 1.0 V processor and memory arrays while providing 12 V for disk drives and legacy loads.

To order call
1-800-735-6200 or email
custserv@vicorpower.com

BCM Model No.	Input Voltage (V)		Output Voltage (V)		Output Power (W)	Output Current (A)
	Nom	Range	Nom	Range		
VIB0002TFJ	380	360 – 400	48	45.00 – 50.00	325	7.0
B384F120T30	380	360 - 400	12	11.25 – 12.50	300	25.0
VIB0003TFJ	350	330 – 365	44	41.25 – 45.60	325	7.7
VIB0001TFJ	350	330 - 365	11	10.30 – 11.40	300	28.0

Products in mass production and available from stock

800-735-6200 **vicorpower.com/hvbcm/edn** VICOR

LED packs 24 chips into tiny package, reduces the need for binning

With Cree's announcement of its XLamp MPL EasyWhite LED, the company moves another step forward in its quest to eliminate "binning" from the vocabulary of LED-lighting-application designers. The XLamp MPL EasyWhite can provide as much as 1500 lumens at 250 mA. This best-case number represents a pulsed rather than a constantly on drive current. With proper drive electronics and heat sinking, however, the LED can deliver equivalent light output to—and consume 78% less power than—that of a traditional 3000K, 75W BR-30 halogen light bulb. This performance meets the Energy Star-defined efficacy-per-lumen requirements for integral LED lamps (www1.eere.energy.gov/buildings/ssl/energy_star_criteria.html).

The EasyWhite series frees light-application designers from the need to match the color output of multiple LEDs in each light-bulb fixture. Customers like to see a uniform color from lights: Imagine a string of track lights using the familiar PAR-38 bulbs—some a warm white and others a cool white. This arrangement would be distracting and unattractive.

To address that problem, LED manufacturers traditionally sorted the chips in color bins. Light manufacturers would mix and match from these bins to get a uniform color, but this approach requires a large inventory of LEDs to meet the production-line needs for a variety of bins. The idea of binning is also confusing. Traditional light sources require only two specifications: color temperature and light output. Cree is attempting to simplify the specification and reduce inventory by eliminating or at least reducing the need for binning.

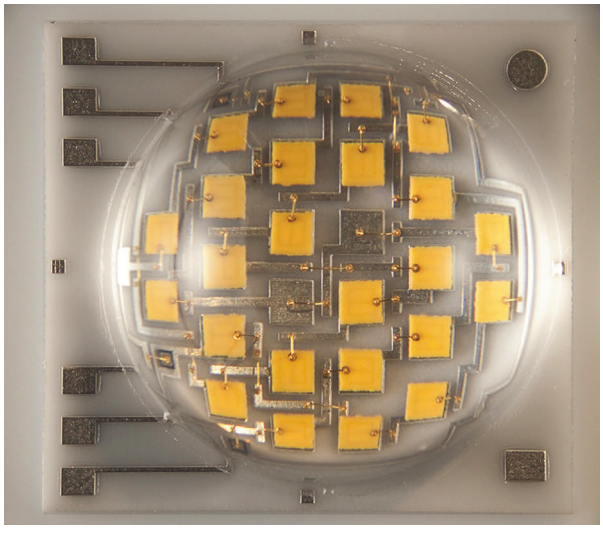

The top-side contacts for the XLamp MPL free the LED from a PCB for mounting. You can mount the MPL directly onto an aluminum heat sink.

The MPL also eliminates pixelation, the appearance of several tiny dots of light as the source, by densely packaging 24 LED chips into a 12×13-mm package that Cree believes is 72% smaller than the next-smallest alternative. In contrast, many manufacturers of traditional multichip LEDs commonly pack only four chips into a package. The XLamp MPL sells for less than $20 each (high volumes).

How does the MPL series, the first member of the EasyWhite family, fit in with the MC-E series of LEDs, which the company announced in late 2009? The MC-E has only four chips in its package. The devices can produce as many as 560 lumens at 700 mA and act as replacements for 20 to 35W halogen light bulbs in indoor-lighting applications, such as accent, track, and pendant lighting.

—by Margery Conner

▷**Cree**, www.cree.com.

Correction

In our Feb 4, 2010, cover story, a sentence on pg 31, repeated in the "At a glance" sidebar on pg 28, reads, "Designers can now implement baseband functions in FPGAs and ASICs that they once had to implement in software." It should have read "Designers can now implement baseband functions in software that they once had to implement in FPGAs and ASICs." For the full article, see "Wireless to go," *EDN*, Feb 4, 2010, pg 26, www.edn.com/article/CA6716491.

XILINX APPLIES AGILENT TOOLS TO LTE VERIFICATION

03.04.10

Agilent Technologies has announced that Xilinx has employed Agilent EEsof SystemVue software plus Agilent test equipment, including the Agilent N5106A PXB baseband generator and channel emulator, to verify 3GPP (Third Generation Partnership Project) LTE (long-term-evolution) hardware blocks using test vectors from a credible algorithmic reference. The Xilinx wireless engineering team needed to develop production-quality LTE uplink LogiCore IP (intellectual property) for Xilinx eNodeB customers. Xilinx had to verify all LTE algorithms and hardware-IP cores against published standards, but its engineers found that available test-vector-creation methods either lacked the fine parameterization and flexibility necessary for testing LTE PHY (physical-layer) algorithms and blocks or lagged behind updates to the standard.

Xilinx used Agilent SystemVue and the W1910 LTE baseband verification library to create highly specific LTE test vectors, which Xilinx downloaded to the Agilent N5106A PXB baseband generator and channel emulator, adding fading and control to the real-time test environment. For more, go to www.edn.com/pb.—by Rick Nelson

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

▶**Xilinx, www.xilinx.com.**

Bringing the Power to Light™ with High Performance LED Drivers

Wide Input Voltage, Low EMI, High Power LED Drivers. Easy to Use, Flexibility by Design.

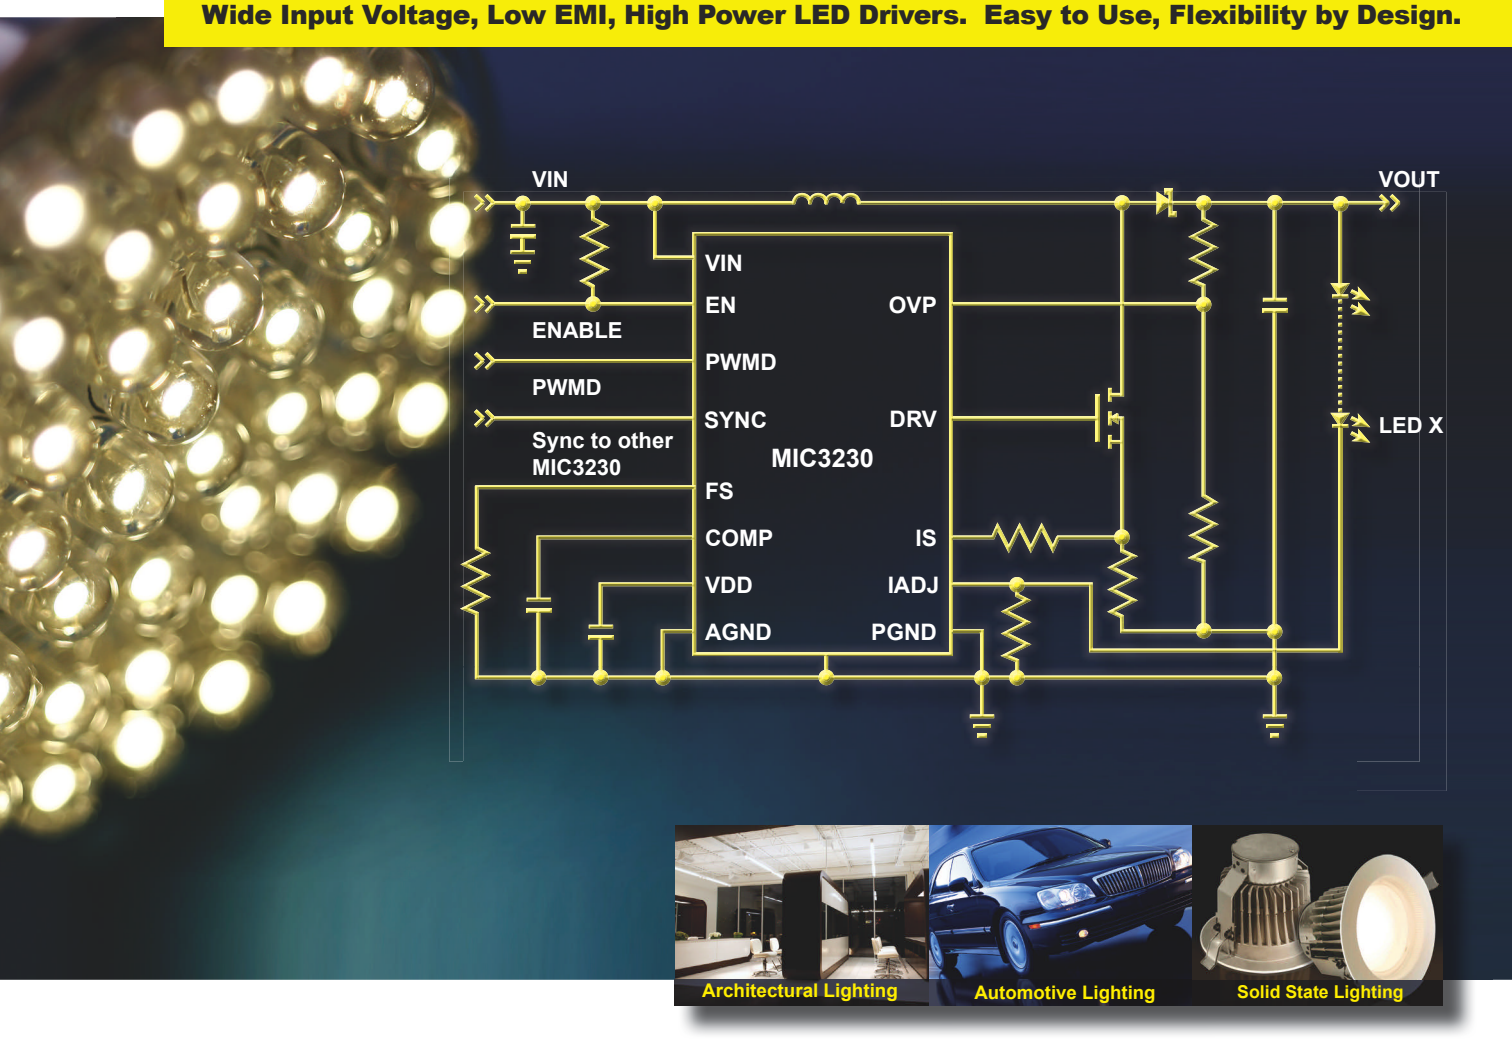

Architectural Lighting

Automotive Lighting

Solid State Lighting

The rapid growth of LED lighting applications in commercial, industrial and automotive markets has led to a wide variety of technical requirements for LED drivers. Micrel is meeting these challenges with its family of high performance LED drivers.

For example, the new 45V input, 70W, 3% accuracy constant current boost controllers like MIC3230/1/2 delivers the performance and flexibility needed. In addition to the PWM dimming and synchronization capability, these LED drivers also offer dither feature for EMI reduction up to 10dB thus providing an unparalleled performance for your solid state lighting applications.

For more information, contact your local Micrel sales representative or visit Micrel at: www.micrel.com/mic3230.

Ideal for use in:

- Solid State Lighting
- Street Lighting
- General Illumination
- Architectural Lighting
- Constant Current Power Supplies

	MIC3230	MIC3231	MIC3232
Input Voltage	6V to 45V	6V to 45V	6V to 45V
Synchronization	Yes	No	No
Dither	No	Yes	No
Frequency Range	Adjustable 100kHz to 1MHz	Adjustable 100kHz to 1MHz	Fixed 400kHz
Package	16-Pin TSSOP 12-Pin MLF®	16-Pin TSSOP 12-Pin MLF®	10-Pin MSOP

MICREL®
Innovation Through Technology®
www.micrel.com

BY HOWARD JOHNSON, PhD

Trace scaling

Imagine a microstrip trace (left side of **Figure 1**). Now scale up that trace, producing a new trace exactly k times larger in width, thickness, and height. It may surprise you to discover that this process of cross-sectional scaling does not change the per-unit-length values of trace capacitance, C, or trace inductance, L. The formulas for those quantities involve only ratios among width, thickness, and height. The absolute values don't matter. A six-mile-wide trace sitting atop a proportionately huge slab of FR (fire-retardant)-4 material enjoys the same exact capacitance per inch as its smaller cousin.

I mention this rule in case you encounter a PCB (printed-circuit-board) trace with dimensions smaller than the minimum values that your signal-integrity simulation tool allows. That scenario can easily happen in multichip-module work involving aluminum traces on silicon substrates. If that situation happens, use the rule of cross-sectional scaling to fake out your simulation tool using physical dimensions large enough to work with the tool but electrically identical to your original circuit.

The rule of cross-sectional scaling preserves the per-unit-length values of capacitance and inductance, but what about other parameters necessary for lossy-line calculations? You can handle them, too. The G parameter, conductance, describes the dielectric properties of the material surrounding your conductors. Specify the same dielectric material in your expanded circuit as the original, and you get the identical dc leakage and dielectric-loss performance. That part is easy. The R parameter represents the trace's resistance, and you must adjust it. **Figure 1** illustrates the correct procedure. The trace on the left comprises a conductor with a resistivity of ρ ohm-meters. The trace on the right scales that resistivity by a factor of k^2.

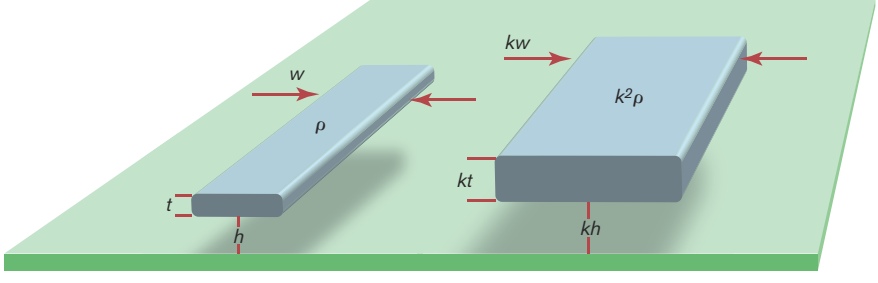

Figure 1 These microstrip traces share identical transmission parameters.

You can check the resistance-scaling procedure by thinking about the formula for the resistance of a long, thin copper bar. That formula equals the resistivity of copper, times the length of the bar, divided by the cross-sectional area of the bar. If you change the bar height and width by k, its cross-sectional area changes by k^2, decreasing the resistance by k^2. If you then multiply the resistivity of copper by k^2, the total resistance returns to its original value.

The resistivity-scaling procedure also properly handles skin-effect resistance. To understand this concept, you must remember that skin depth equals $\sqrt{2\rho/(2\pi f)\mu}$, where f is the frequency of operation, and μ is the magnetic permeability of the dielectric material. You must also know that what matters in a lossy-line problem is not the absolute value of skin depth but its value as a proportion of trace thickness. The skin-depth formula shows that scaling ρ by k^2 scales the skin depth by k, which renders the ratio of skin depth to trace thickness (also scaled by k) the same as in the original trace.

I do not know whether your simulator will allow you to make the required changes in resistivity. If it does, you can build a new layer stack with the same actual per-unit-length values of resistance, inductance, conductance, and capacitance as your original configuration, only at a physical scale acceptable to your simulator. Your simulator won't know the difference. EDN

Howard Johnson, PhD, of Signal Consulting, frequently conducts technical workshops for digital engineers at Oxford University and other sites worldwide. Visit his Web site at www.sigcon.com or e-mail him at howie03@sigcon.com.

Look to New Digital Optocouplers for Low Power and High Noise Rejection

Introduction

For many years optocouplers have reduced system noise, provided high voltage insulation, and made systems reliable and safe. Although these basic system requirements remain, Avago has developed new isolation technology and architectures that reduce power consumption substantially and increase design flexibility without compromising high voltage protection and noise rejection.

Figure 1. ACPL-x6xL power vs. first generation optocouplers.

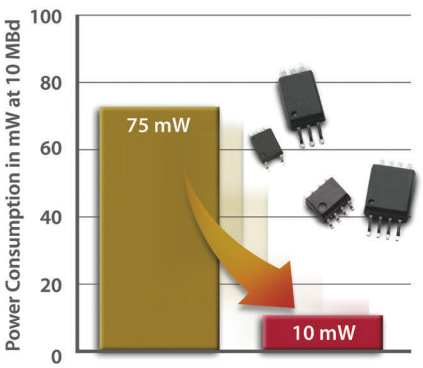

Avago's new ACPL-x6xL ultra-low power, 10 MBd CMOS digital optocouplers have the performance needed in modern communication interfaces that use SPI and I^2C serial interfaces, long distance RS-485 multi-point communications networks, CAN bus interfaces, microprocessor and microcontroller interfaces, and A/D and D/A converters. Table 1 shows the overall ACPL-x6xL family characteristics. Performance over temperature, –40 to 105°C, low LED drive current, continuous working voltages of up to 1,140 Vpeak and transient overvoltage protection of up to 8,000 Vpeak are ideal for industrial, medical, test and measurement and computer peripheral applications.

Key ACPL-x6xL Series Features

- Low current LED input allows direct drive from microprocessor CMOS outputs
- 35 kV/µs dynamic and static common mode rejection - no compromise on noise immunity!
- Controlled output slew rate over a wide range of load conditions
- Easy configuration for inverting and non-inverting operation
- Optical isolation technology is certified for safe insulation at continuous working voltages from 560 Vpeak to 1140 Vpeak and with transient voltages of 6 kVpeak / 8 kVpeak

Table 1. Key ACPL-x6xL family features.

Parameter	Specification
Propagation Delay	80 ns (max)
Skew between any two parts	30 ns (max)
Pulse width distortion	30 ns (max)
LED forward current	1.6 mA
Icc supply current	1.3 mA (max)
Common Mode Noise Rejection	35 kV/µs at 1000 V
Supply voltage	2.7 to 5.5V
Temperature range	–40 to 105°C
Continuous working voltage	560 Vpeak / 1140 Vpeak
Transient overvoltage	6 kVpeak / 8 kVpeak

Architecture

The low LED current needed to switch the new optocouplers series makes it possible to drive the input LEDs directly from most microprocessor outputs. To simplify your circuit design, all timing parameters are specified with fixed input resistor configurations, rather than current, for 3.3 V and 5 V signal levels across the temperature range. The push-pull CMOS outputs eliminate external pull-up resistors.

LED Inputs are Ideal in Noisy Environments

Common mode noise can be a significant problem in data communication applications, especially in industrial environments where electrical motors, sensors and programmable logic controllers are connected together. In such systems, isolators reduce noise levels and enhance signal performance. All isolators, regardless of technology, have a parasitic capacitive coupling between the two isolated sides of the component. A noise transient occurring on the output side may cause an unwanted voltage rise on the input side. This could result in false-triggering of the input or even latch-up of high impedance logic inputs. Optical isolators with LED inputs are ideal in environments with high levels of common mode noise.

The attenuation of the light signal through the internal, optically transparent insulation material is low, so the distance through isolator can be kept high in the new ACPL-x6xL family. A direct consequence of the large separation distance is a low parasitic capacitance. In other words, the unwanted coupling between the two sides is kept to a minimum.

The "split resistor" input LED drive approach balances the impedance across the LED input. A common mode noise voltage rise on the LED is symmetrical and therefore cannot switch the LED on. In addition, the ACPL-x6xL family LED inputs have a relatively high 70 pF input capacitance. The series connected LED and current

limiting resistor form a low pass filter that reduces noise transients further. An internal Faraday shield also minimizes the effects of common-mode noise.

Figure 2. Split input resistors give balanced impedance for higher CMR and lower noise.

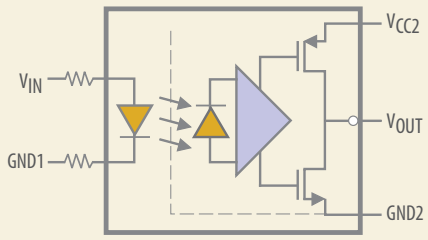

New Specification: Dynamic Common Mode Rejection

For the new ACPL-x6xL product family, Avago has an additional definition for common mode noise rejection that reflects how the device operates in an application. First generation test setups for common mode noise rejection are based on static performance with the input tied directly to Vcc or ground. Our new Dynamic Common Mode Rejection specification is based on common mode performance during data transmission. The static and Dynamic Common Mode Noise Rejection of the ACPL-x6xL family is 35 kV/µs.

Design Flexibility: Inverting and Non-inverting Operation

The output stage is inverting but it is easy to configure ACPL-x6xL devices for either inverting or non-inverting operation. In Figure 2, change Vin to Vcc and GND1 to Vin to obtain a non-inverting configuration.

Long Life: LED Light Output over Time and Temperature

An optocoupler's life time is a strong function of input LED quality. Avago produces high reliability LEDs for optocouplers at its in-house facility. The infrared, AlGaAs LED used in the ACPL-x6xL family provides excellent stability over both temperature and time. As shown in Figure 3, light output power (LOP) is essentially unchanged after 100,000 hours of operation. Graphs showing performance over temperature are in the product data sheet.

Figure 3. LED light output power, LOP, stability over time

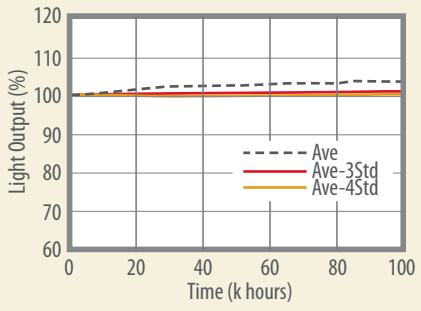

For isolators, fundamental trade-offs exist between power consumption, noise rejection and switching speed. The new ACPL-x6xL optocouplers series has been developed to achieve best-in-class noise rejection and power consumption without compromising safety specifications.

Design Example: Isolated SPI Interface Benefits from Slew-rate Controlled Outputs

In many applications, such as SPI and I^2C serial interfaces, stable switching parameters over time and temperature are far more important than the actual device switching speed.

The maximum data rate for a serial interface is limited by the skew / synchronization between signals on the data and clock lines rather than the optocoupler's absolute propagation delay. Increasing the absolute switching speed of an isolator decreases noise rejection. Rather than reducing propagation delay, Avago focused on new design features that reduced skew and enhanced signal quality. One such feature is slew-rate controlled outputs.

Differences in line capacitance can lead to differences in rise and fall time between two channels. In the case of an SPI interface, the clock and signal lines may have different fan-out capability. The output on the ACPL-x6xL family is able to precisely control the output rise and fall times over a wide range of load capacitances. The propagation delay difference (skew) between any two channels in the ACPL-x6xL family is specified at 30 ns maximum over temperature.

Supply Voltage Range and Glitch-Free Outputs

Besides being power efficient, the new ACPL-x6xL is able to deliver stable switching performance across a 2.7 V to 5.5 V supply voltage range plus the outputs are glitch-free during power up and power down. A circuit similar to under-voltage lockout ensures an output is at a known state during power up or power loss.

Summary

Avago's new generation of 10 MBd digital optocouplers provide the power efficiency modern designs need without sacrificing high voltage insulation or noise isolation performance. Outputs are slew-rate controlled for fast operation and stability over varying loads. All devices feature glitch-free outputs during power up and power down.

As Table 2 shows various package options are available. All devices are RoHS-6 compliant and meet all international safety requirements.

Table 2. Ultra low power 10 MBd digital optocouplers

Part No.	Package	Chnl	Isolation Voltage (Vrms/min)	Working Voltage (Vpeak)
ACPL-061L	SO-8	1	3750	560
ACPL-064L	SO-8	2	3750	560
ACPL-M61L	SO-5	1	3750	560
ACPL-W61L	SSO-6	1	5000	1140
ACPL-K63L	SSO-8	2	5000	1140

An evaluation aboard is available to all qualified designers by contacting your local Avago distributor or representative.

Contact us for your design needs at: **www.avagoresponsecenter.com/401**

AVAGO TECHNOLOGIES

BY PALLAB CHATTERJEE, CONTRIBUTING TECHNICAL EDITOR

Nanotechnology makes strides in quantum dots, medicine

Nanotechnology is an emerging area with contributions from both established companies and start-ups. One participant in the industry, eight-year-old Nanosys Inc (www.nanosysinc.com), is working on commercializing materials made with nanoparticles and nanotechnology devices. The company recently announced its Quantum-Rail lighting system for LG mobile-phone applications. This announcement represents the first major consumer-product application of the company's quantum-dot-phosphor nanoscale-coating technology. You can apply the coating to LEDs to help clarify the color spectrum and reduce power consumption in backlit-LED applications. The technology enables mobile devices to display a truer color even in daylight. You can also use the technology in large LED arrays, such as those in LCD televisions. In those applications, the coating technology reduces power consumption, improves color and clarity, and screens the low-end "tail" light from the LEDs that causes some TV sets to interfere with IR (infrared) remote units on set-top-box or other audio/video equipment.

The Nanosys announcement heralds the arrival of nanotechnology products from other companies that will soon be entering the market and shows that the typical eight- to 10-year gestation period for breakthrough technologies to reach commercialization is now reaching an end. For example, nanomedicine is now emerging as a major topic of investigation. To help solidify the topics in this area and to determine the best direction for commercialization, the ASME (American Society of Mechanical Engineers) held the First Global Congress on NEMB (nanoengineering for medicine and biology, www.asmeconferences.org/NEMB2010), a three-day event that took place last month in Houston.

The typical eight- to 10-year gestation period for breakthrough technologies to reach commercialization is now reaching an end.

Nanomedicine is an interesting field because a lot of the research and product development is coming from the traditional DSM (deep-submicron)-IC industry. Giving one of the plenary talks at the event was Mauro Ferrari, PhD, director of nanomedicine and deputy chairman of the department of biomedical engineering at the University of Texas Health Science Center. The NEMB event also showcased materials, characterization, modeling, circuit, and NEMS (nano-electromechanical-system)-design presentations.

In materials development, nanotechnology parallels the strides researchers are making in subwavelength-electronics-semiconductor and semiconductor-process development. Presenters at the event discussed spin-coating techniques, carbon and magnetic nanotubes, and shadow-edge lithography for the creation of nanochannels. These nanochannel structures are finding use in molecular microfluidic applications and in waveguides for monolithic optical circuits. Researchers on the circuit side of nanomedicine are examining ultra-high-sensitivity traditional circuits in which the sense levels are molecular rather than macroscopic. Developers presented polysilicon-nanowire FET chemical sensors; piezoresistive cantilevers for explosive vapor, which has applications in security and health care; nonlinear-resonance NEMS and MEMS (microelectromechanical-system) biosensors; and wireless point-of-care patient-monitoring systems.

To supplement the circuit design, presenters discussed a new development in modeling and the need for a basic understanding of the nanoscale phenomenon. For example, they covered electron-transportation behavior in single-crystal silicon nanowires, the design of nanotip sensors, and modeling and simulation of the motion of bio-inspired swimming microrobots. Semiconductor-processing techniques and physics-based modeling for the optical and medical sectors are strong extensions of the traditional electronic-semiconductor business. The large growth and R&D in these sectors will be around for several decades. Silicon and carbon are still holding strong as the engineering materials of choice. **EDN**

Pallab Chatterjee is vice chairman of the IEEE San Francisco Bay Area Nanotechnology Council. You can reach him at pallabc@siliconmap.net.

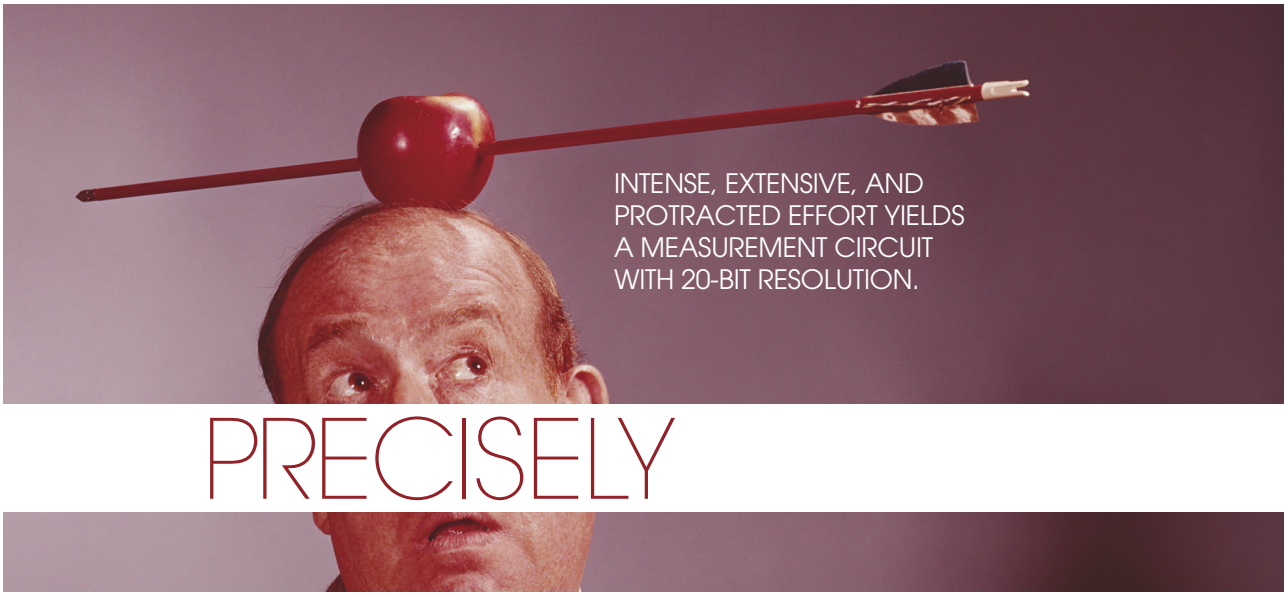

PRECISELY

measure settling time to 1 ppm

BY JIM WILLIAMS • LINEAR TECHNOLOGY

Instrumentation, function generators, inertial navigation systems, ATE (automatic test equipment), medical apparatus, and other precision applications now require 18-bit converters. The resolutions of these converters are so precise that measuring various performance parameters is difficult. DACs' dc specifications are relatively easy to verify, but their ac specifications require more sophisticated approaches to produce reliable information. A DAC's settling time is the elapsed time from an input code application until the output remains within a specified error band around the final value. The settling time of a DAC and its output amplifier is difficult to determine at 18-bit, or 4-ppm, resolution. To measure an 18-bit DAC, you must use measurement techniques with 20-bit, or 1-ppm, resolution for settling times as short as 265 nsec. Manufacturers usually specify DACs' settling times for a full-scale 10V transition.

Measuring anything at any speed to 20-bit, or 1-ppm, resolution is difficult (see **sidebar** "Components for 18-bit digital-to-analog conversion"). Dynamic measurement to 20-bit resolution is particularly challenging. Reliable 1-ppm DAC-settling-time measurement constitutes a difficult problem requiring exceptional care in approach and experimental technique. You can use an oscilloscope to accurately display DAC-settling-time information for a 10V step with 1-ppm, or 10-μV, resolution within 265 nsec. The approach permits observation of small amplitude information at the excursion limits of large waveforms without overdriving the oscilloscope.

DAC settling time has three distinct components: delay time, slew time, and ring time (**Figure 1**). The delay time is small and is almost entirely due to propagation delay through the DAC and the output amplifier. During this interval, no output movement occurs. During slew time, the output amplifier moves at its highest possible speed toward the final value. Ring time defines the region in which the amplifier recovers from slewing and ceases movement within some defined error band. A trade-off between slew and ring time normally exists. Fast-slewing amplifiers generally have extended ring times, complicating your amplifier choice and frequency compensation. Additionally, the architecture of fast amplifiers usually dictates trade-offs that degrade dc error terms (see **sidebar** "DAC-amplifier compensation").

DAC SETTLING TIME

Engineers have previously measured DAC settling time with the false-sum-node technique (**Figure 2**). The resistors and DAC form a bridge-type network. With ideal components, the DAC output steps to the negative reference voltage when the DAC inputs move to all ones. During slew, the diodes, which limit the voltage excursion, bind the settling node. When settling occurs, the oscilloscope probe's voltage should be 0V. The resistor divider's attenuation means that the probe's output will be one-half of the DAC's settled voltage.

In theory, this circuit allows you to

observe settling to small amplitudes. In practice, you cannot rely on the circuit to produce useful measurements. The oscilloscope connection presents problems. As probe capacitance rises, ac loading of the resistor junction influences the settling waveforms. A 10-pF probe alleviates this problem, but its 10× attenuation sacrifices gain. Probes with 1× attenuation are not suitable because of their excessive input capacitance. An active 1× FET (field-effect-transistor) probe works but still presents a more significant issue: oscilloscope overdrive.

The clamp diodes at the settling node reduce swing during the amplifier's slewing, which prevents the circuit from overdriving the oscilloscope. Unfortunately, oscilloscope-overdrive-recovery characteristics vary widely, and manufacturers do not usually specify them. The 400-mV drop across the Schottky diodes means that the oscilloscope may see an unacceptable overload.

AT A GLANCE

- Measuring anything to 1 ppm (part per million) is difficult, but measuring ac phenomena, such as settling time, is even more difficult.
- Use a variable-transconductance amplifier instead of a diode bridge.
- Trim out delays and parasitic effects when possible.
- Be sure to verify your measurements.

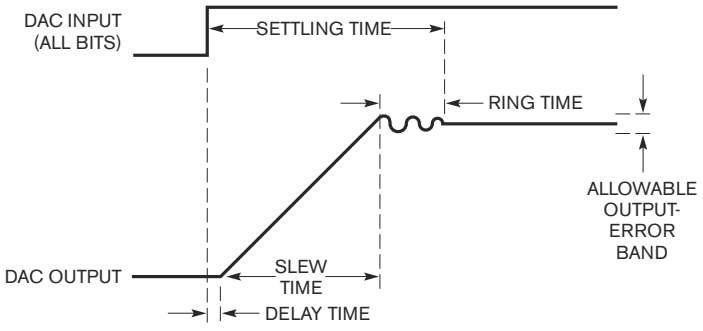

Figure 1 DAC settling time components include delay, slew, and ring time. Fast amplifiers reduce slew time, although longer ring time usually results. Delay time is normally a small term.

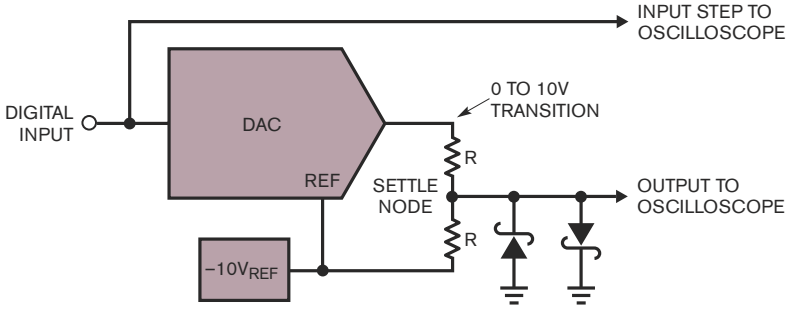

Figure 2 A popular summing scheme for DAC-settling-time measurement provides misleading results. An 18-bit measurement causes a more-than-800-times increase in oscilloscope overdrive.

An oscilloscope typically undergoes a 2× overdrive when you set it to 50 mV/division while measuring to a 10-bit resolution. A 10-bit resolution—10 mV at the DAC output—results in 5 mV at the oscilloscope. The desired 5-mV baseline is just discernible. At 12-bit or higher resolution, the measurement becomes hopeless. If you increase the oscilloscope's gain, you also increase overdrive-induced errors. At 18 bits of resolution, there is no chance of measurement integrity.

Measuring 18-bit settling time requires a high-gain oscilloscope that is also immune to overdrive. You can address the gain issue by using an external wideband preamplifier that accurately amplifies the diode-clamped settling node. Only classic sampling oscilloscopes, which vendors no longer manufacture, offer inherent overdrive immunity. However, you can construct a circuit that uses sampling techniques to avoid the overload problem. You can

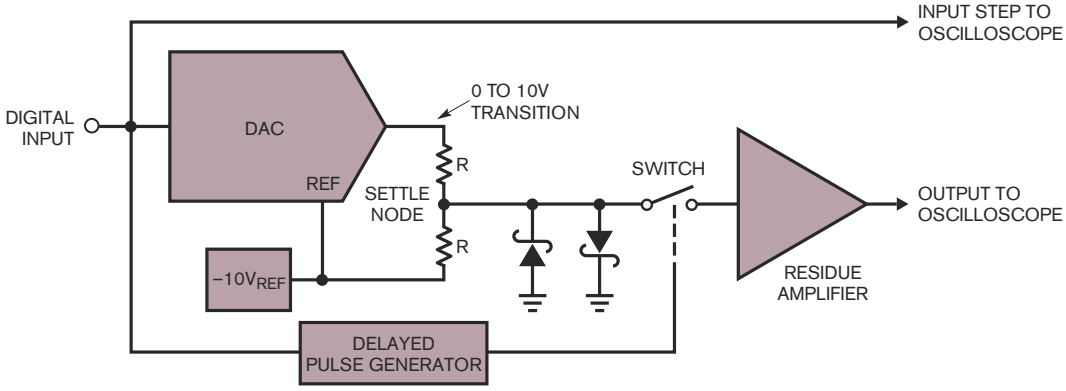

Figure 3 This arrangement eliminates oscilloscope overdrive.

also endow the circuit with features for measuring 20-bit DAC settling time.

To measure 20-bit DAC settling time, you connect the settling point to the preamplified oscilloscope with a switch (**Figure 3**). You control the switch state with a delayed pulse generator, which you trigger with the same pulse that controls the DAC. You arrange the delayed pulse generator's timing so that the switch does not close until settling is nearly complete. In this way, you sample the incoming waveform in both time and amplitude. The oscilloscope never experiences overdrive, and no off-screen activity occurs.

The sampling switch has stringent requirements. It must faithfully pass signal-path information without introducing alien components, including those from the switch-command channel. Conventional choices for the sampling switch are JFETs (junction FETs), MOSFETs (metal-oxide-semiconductor FETs), and diode bridges (**Figure 4**).

The parasitic gate-to-channel capacitances of FETs result in a large feedthrough into the signal path from the gate drive. This feedthrough is many times larger than the signal you are observing, and it obviates the switch's purpose. The small parasitic capacitance and the symmetrical differential structure of a diode bridge tend to cancel the drive signal and result in low feedthrough. You must perform dc and ac trims and use complex drive and support circuitry with a diode bridge. You can reliably measure DAC settling time to 16-bit resolution. Beyond 16 bits, residual feedthrough becomes objectionable, and you must use another approach.

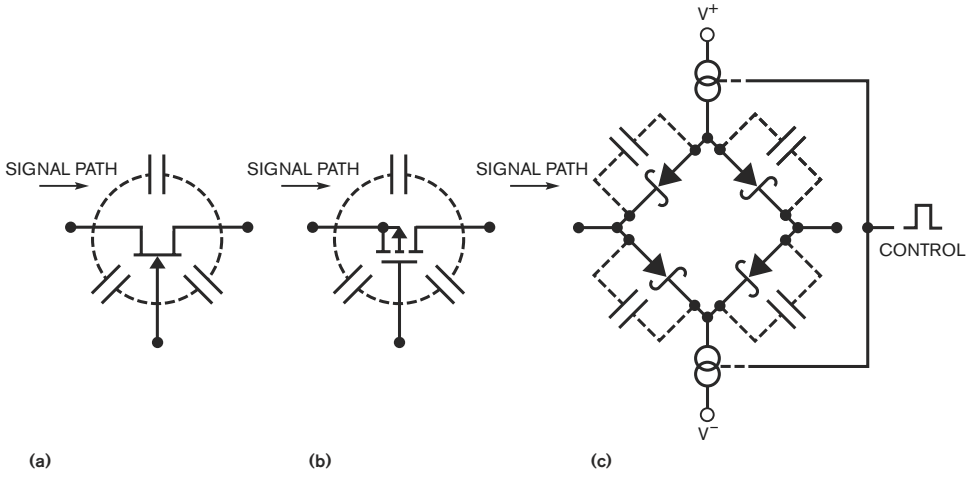

Figure 4 Conventional choices for sampling switches include JFETs (a), MOSFETs (b), and diode bridges (c).

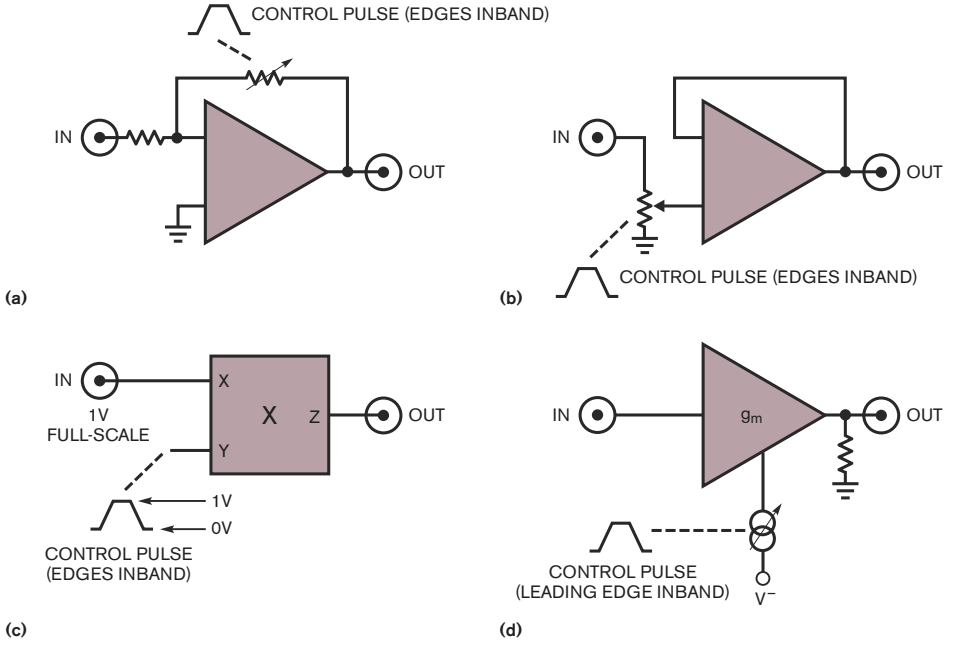

Figure 5 Some low-feedthrough electronic-switch equivalents (a and b) are difficult to implement. You must optimize one practical scheme for low feedthrough on rising and falling control-pulse edges (c). In another, inherent bandwidth reduction minimizes falling-edge feedthrough (d).

You can construct a low-feedthrough, high-resolution switch using wideband active components. With this approach, you can maintain the switch's control channel inband because its transition rate is within the circuit's bandpass frequency. The circuit's wide bandwidth means that you can always control the switch's transition. You greatly reduce feedthrough by having no out-of-band response.

Some candidates for low-feedthrough electronic switches are theoretically possible but cumbersome to implement (**Figure 5a** and **b**). You must optimize others for low feedthrough on the rising and falling control-pulse edges because of the multiplier's unrestricted wideband response (**Figure 5c**). You can minimize falling-edge feedthrough by using the collapse of an amplifier's transconductance when the control pulse goes low (**Figure 5d**). This approach allows you to minimize the feedthrough of the control pulse's rising edge without regard to falling-edge effects. This feature provides a significant advantage for an electronic switch.

An electronic-circuit sampling switch uses a transconductance amplifier (**Figure 6**). The switch dynamics are exceptionally pure because the wideband control and signal paths faithfully track a 1000-to-1 transconductance

change. You can carefully optimize the switched current source for the lowest feedthrough on the rising control edge without regard to the falling-edge characteristics. The transconductance collapses on the falling edge, which ensures low feedthrough for that condition. This technique prevents oscilloscope overdrive.

The detailed design switches signals over a ±30-mV range with peak control-channel feedthrough of only millivolts and settling times of less than 40 nsec (**Figure 7**). The unity-gain circuit approximates switch action by varying A_{1A}'s transconductance. The amplifier and its transconductance-control channel are wideband components, permitting them to faithfully track rapid variations in a transconductance setting. This characteristic means that the amplifier is always in control, affording clean response and rapid settling to the input's value.

Amplifier A_{1A} is the wideband transconductance amplifier. You set its voltage gain by the current magnitude into its I_{SET} (current-set) terminal and its output-resistor load. Amplifier A_{1B} unloads A_{1A}'s output. This amplifier provides a gain of two. When driving a back-terminated 50Ω cable, however, its effective gain is unity at the cable's receiving end. The back termination establishes a 50Ω signal-path environment. The switch-control input controls current source Q_1, which sets A_{1A}'s transconductance and gain.

Q_1 is gated off when you set the switch-control input at zero. The 10-MΩ resistor supplies 1.5 µA into A_{1A}'s I_{SET} pin, which results in a voltage gain of nearly zero. This action blocks the input signal. When you set the switch-control input high, Q_1 turns on, sourcing 1.5 mA into the I_{SET} pin, in turn forcing maximum transconductance and causing the amplifier to assume unity gain and pass the input signal. You use trims for zero and gain to ensure accurate input-signal replication at the circuit's output. You purify turn-on switching with a 50-pF variable capacitor and use a 10-kΩ thermistor with a 3300-ppm/°C temperature coefficient at Q_1. This approach compensates A_{1A}'s transconductance temperature dependence to minimize gain drift.

When you set the aberration capacitance to 35 pF, the circuit cleanly switches a 10-mV input signal (**Figure 8**). When the control input (Trace A) is low, no output (Trace B) occurs. When the control input goes high, the output reproduces the input with feedthrough settling in 20 nsec. You don't see turn-off feedthrough due to the transconductance reduction of 1000-to-1 and the attendant 25-fold decrease in the frequency bandwidth. You set the sweep speed to 10 nsec/division to examine the 0V-settling details (**Figure 9**). The output (Trace B) settles to within 1 mV within 40 nsec after the switch control (Trace A) goes high. The aberration capacitance damps the peak feedthrough excursion of 5 mV. If you set this capacitance to 0 pF, feedthrough increases to 20 mV (**Figure 10**). Settling time to within 1 mV remains at 40 nsec. Setting this capacitance to 35 pF minimizes the feedthrough amplitude but yields seven-times-higher rise time than is possible with a setting of 0 pF (**Figure 11**).

The small dc and ac error of a transconductance switch accommodates the application's requirements. The low feedthrough becomes irrelevant because the DAC's ring-time interval buries the small time and amplitude error, pointing the way toward practical 1-ppm DAC-settling-time measurements.

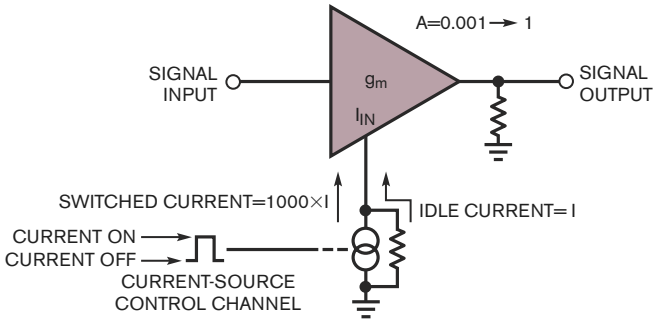

Figure 6 This transconductance amplifier switch has minimal control-channel feedthrough.

MEASUREMENT METHOD

You can combine this electronic switch with the DAC-amplifier summing method (**Figure 12**). You split the delayed pulse generator into independently variable delay and pulse-

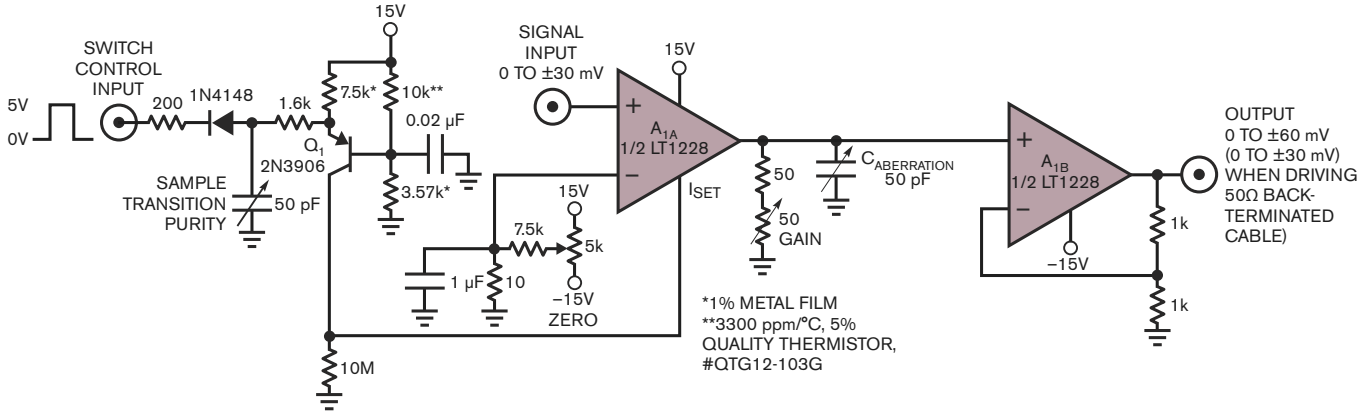

Figure 7 A low-level, 100-MHz switch has minimal control-channel feedthrough.

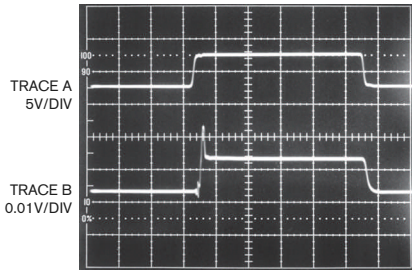

Figure 8 The control input (Trace A) dictates a switch output's (Trace B) representation of a 0.01V-dc input. The control-channel feedthrough settles in 20 nsec. Turn-off feedthrough is undetectable due to deceased signal-channel transconductance and bandwidth. The aberration capacitance is 35 pF for this test.

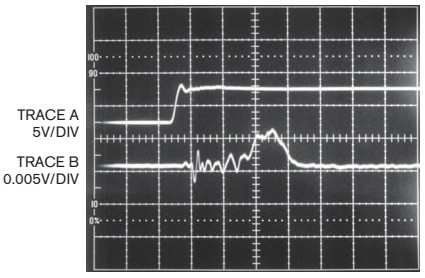

Figure 9 For a 0V signal input, the output (Trace B) peaks at 0.005V before settling to 0.001V 40 nsec after the switch-control command (Trace A). Aberration capacitance is 35 pF for this test.

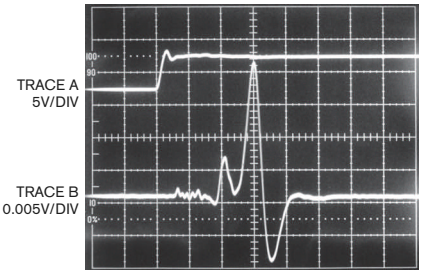

Figure 10 The aberration capacitance is 0 pF; otherwise, conditions are identical to those in Figure 9. The feedthrough-related peaking increases to 0.02V, and the 0.001V settling time remains at 40 nsec.

generator blocks. You run the input step to the oscilloscope through a section that compensates for settling-time-path propagation delays. This path includes the settling node, the amplifier, and sample-gate delays. You drive the transconductance sampling switch from a nonsaturating residue amplifier, feeding the oscilloscope. Placing the sampling switch after the residue amplifier further minimizes sample-command feedthrough.

The 20-bit DAC-settling-time-measurement circuitry uses the input pulse to simultaneously switch all the DAC bits (**Figure 13**). You must use a modified circuit and a Microchip (www.microchip.com) PIC microcontroller for serial DACs (see **sidebar** "Settling-time measurement of serial DACs"). You also route this pulse to the oscilloscope through the delay-compensation network, which comprises CMOS inverters and an adjustable RC network. It compensates the oscilloscope's input step signal for the 44-nsec delay through the circuit-measurement path (see **sidebar** "Delay and circuit-trimming procedures"). You compare the DAC-amplifier output with the 10V reference using precision 10-kΩ summing resistors. The reference IC also furnishes the DAC reference, which makes the measurement ratiometric. A_1, which also has voltage gain, unloads the clamped settling node. Amplifier A_2 has a clamped output and provides a total amplification of 40 referred to the summing node. A_2's output feeds the sampling-switch circuit. The potentiometer adjusts the circuit's gain.

You arrange the clamping and gain of A_1 and A_2 so that saturation never occurs. The amplifier is always in its active region. The input pulse triggers the 74HC123 dual one-shot. You control the pulse delay of the one-shot with a 20-kΩ potentiometer. If you set the delay appropriately, the oscilloscope sees no input signal until settling is nearly complete, thereby eliminating overdrive. You adjust the sampling switch's on-time so that you can observe all remaining settling activity.

When the sample gate goes high, switching is clean, and you can easily observe the last millivolt of ringing time (**Figure 14**). The amplifier settles nicely to a final value. When the sample gate goes low, the transconductance switch goes off, and you can discern no feedthrough. There is never any off-screen activity, and the oscilloscope is never in overdrive.

To achieve this level of performance, you must trim the circuit (see **sidebar** "Delay and circuit-trimming procedures"). If you do a poor sample-interval-zero adjustment, the circuit shifts the output baseline (**Figure 15**). A proper adjustment yields a continual baseline (**Figure 16**). You can see the barely visible sample-command feedthrough at Trace B's leading edge. Signal peaking goes to 350 μV, and settling is 50 nsec before your trimming aberrations and transition purity (**Figure 17**). By trimming these functions, you increase delay to 70 nsec from 20 nsec, but the signal peaks at only 50 μV

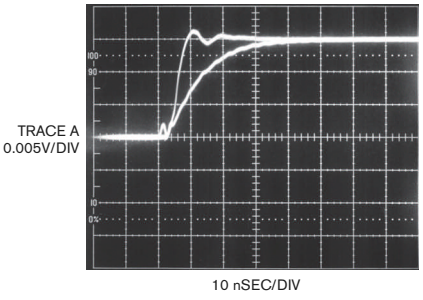

Figure 11 The signal-channel rise times for 0 pF (left trace) and 35 pF (right trace) are 3.5 and 25 nsec, respectively. The switch-control input is high for this measurement.

(**Figure 18**). Signal settling remains at 50 nsec. At the sample-command turn-off, the 1000-to-1 transconductance drop ensures a clean transition independent of the turn-on optimizing trims (**Figure 19**).

You can graph the minimum measurable settling time for a given resolution (**Figure 20**). Sample-command-path delays and sample-gate switching-residue profiles impose speed limitations on the circuit. You can measure settling time below 160 nsec at 16-bit resolution. For smaller signals, the sample gate's switching-residue profile dictates an increased minimum measurable settling time of 265 nsec at 20 bits. The DAC/amplifier, summing resistors, and residue amplifier/sampling switch impose circuit-noise limitations. Resolution beyond 15 ppm requires you to use filtering or noise-averaging techniques. You can drive the 74HC123's B2 input with a phase-advanced version of the

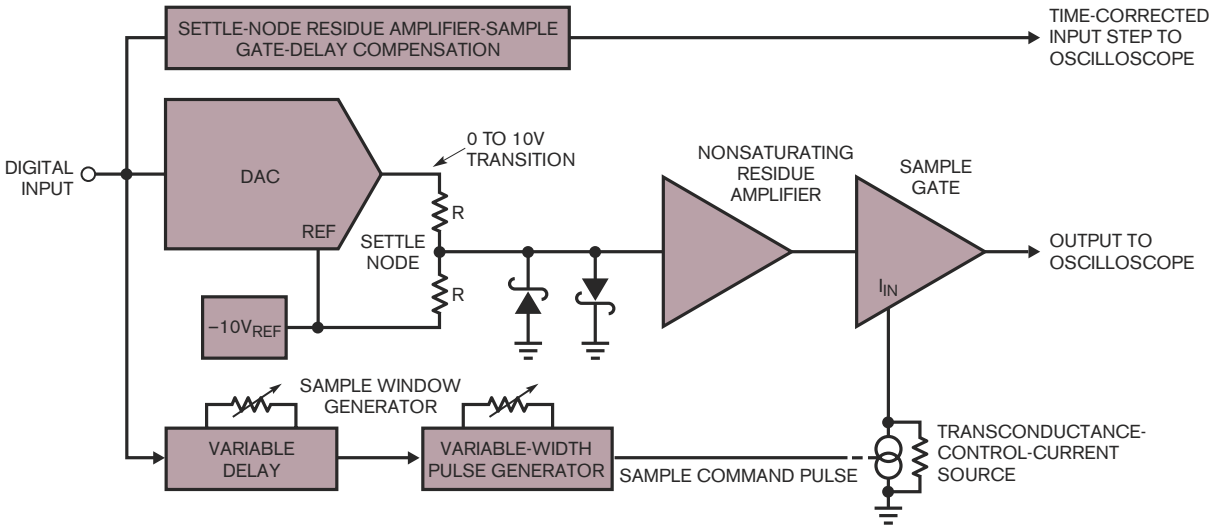

Figure 12 You can combine the electronic switch with the DAC-amplifier summing method.

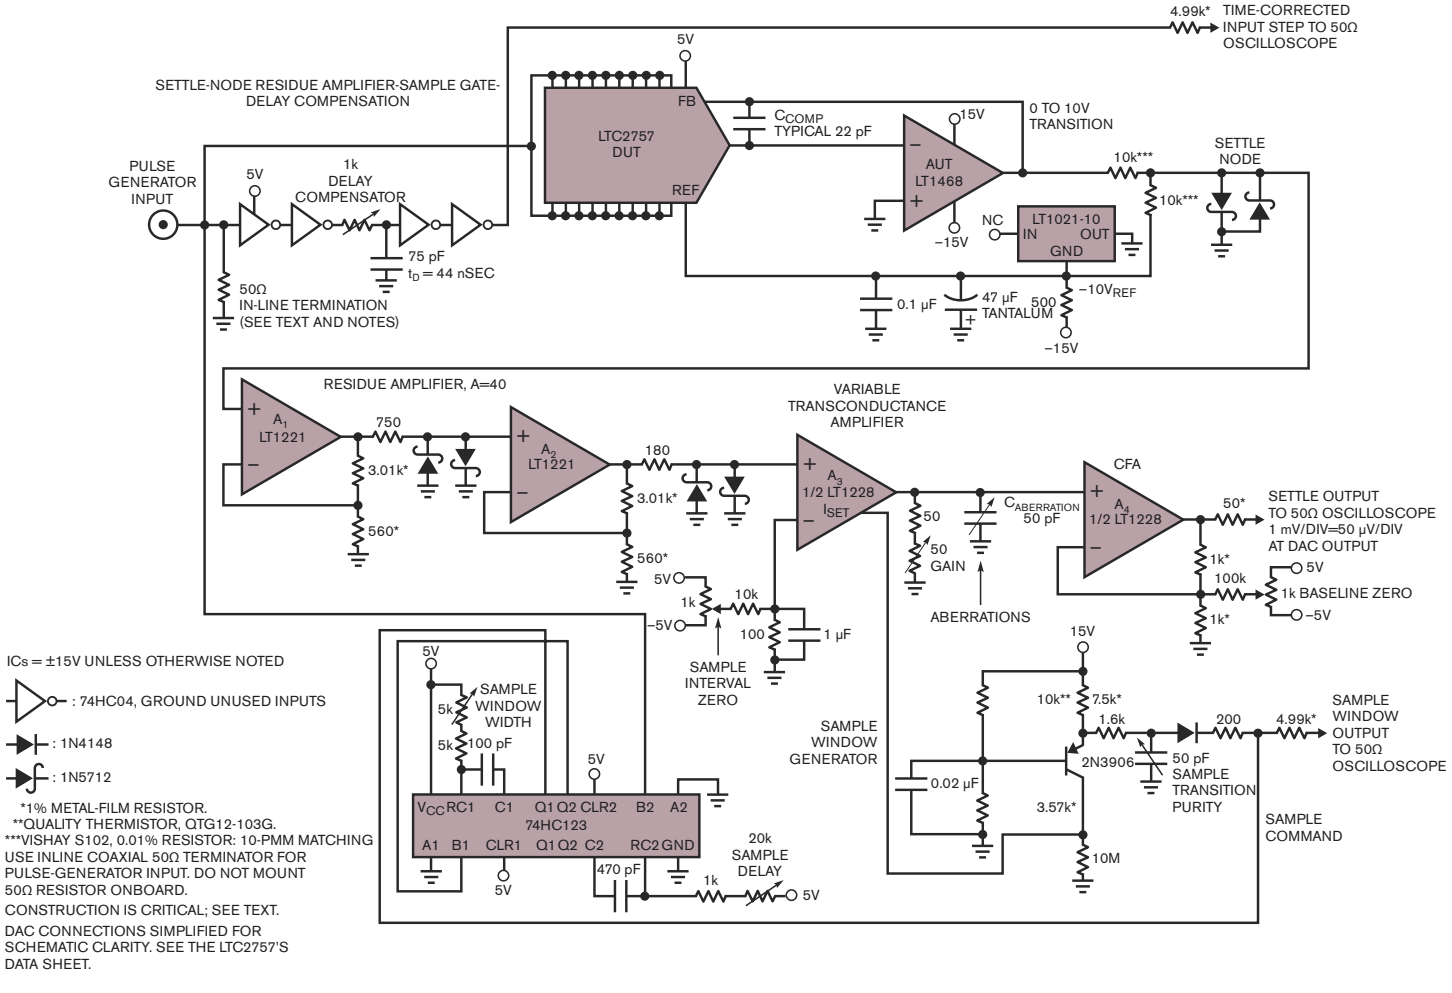

Figure 13 A detailed DAC-settling-time-measurement circuit closely follows that of Figure 12. Optimum performance requires attention to layout.

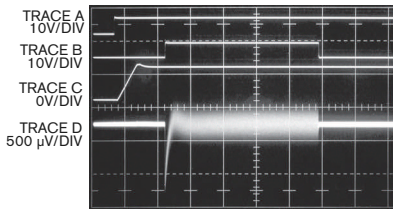

Figure 14 The settling-time-circuit waveforms include a time-corrected input pulse (Trace A), a sample command (Trace B), a DAC output (Trace C), and a settling-time output (Trace D).

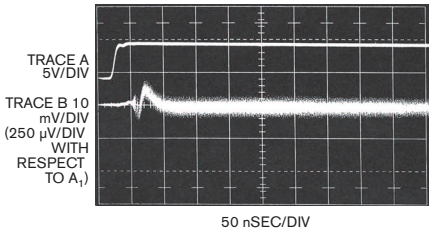

Figure 17 The output response (Trace B) to a sample command's turn-on (Trace A) before you trim out aberrations and transition purity has a signal delay of 20 nsec. The aberrations peak at 350 µV and settle in 50 nsec. You ground amplifier A_1's positive input using a 5-kΩ resistor for this figure and figures 15 and 16.

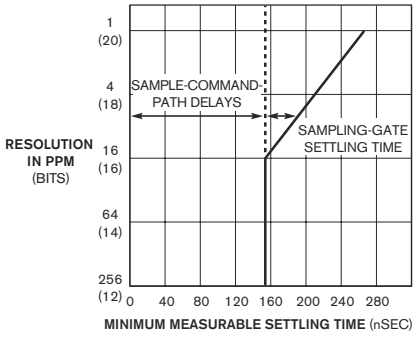

Figure 20 The sample-command-path delays and sample-gate-setting profile impose the minimum measurable settling-time-versus-resolution limits. Achieving a resolution beyond 15 ppm requires filtering or noise averaging.

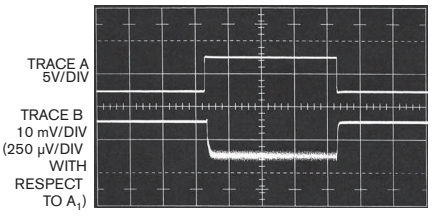

Figure 15 A poor sample-interval zero adjustment causes a shifted output baseline (Trace B) during Trace A's sample interval.

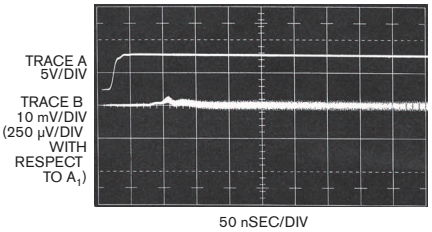

Figure 18 Trimming improves the output response (Trace B) to a sample command's turn-on (Trace A). The delay increases to 70 nsec, but aberrations peak at only 50 µV, and the circuit settles in 50 nsec.

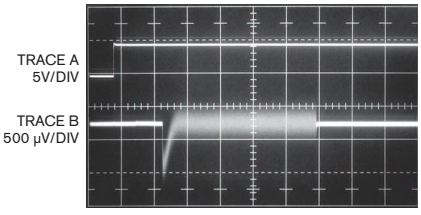

Figure 21 The 0 to 10V DAC with an unfiltered bandpass settles (Trace B) to 16 bits, or 15 ppm, in less than 2 µsec after Trace A's time-corrected input step. The sample gate's feedthrough is well-controlled, indicating that higher resolution is possible without overdriving the oscilloscope. Noise is the limitation for this measurement.

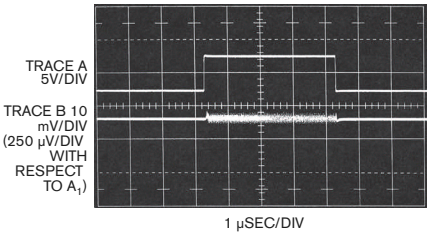

Figure 16 A trimmed sample interval's zero adjustment has no output baseline deviation (Trace B) during the sample interval (Trace A). The sample command's feedthrough is just visible at Trace B's leading edge.

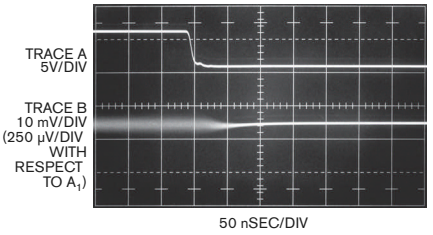

Figure 19 The output response (Trace B) to the sample command's turn-off (Trace A) shows a 1000-to-1 transconductance drop, ensuring a clean transition, independent of the trim state.

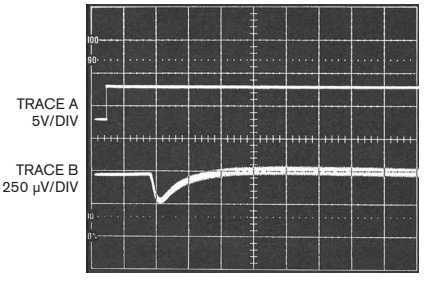

Figure 22 Reducing the measurement bandwidth to 250 kHz attenuates noise. The settling to 18 bits, or 4 ppm, requires approximately 5 µsec. Filtering permits increased resolution, although the indicated settling time increases.

pulse-generator input to largely eliminate sample-command-path-delay-induced error. This approach considerably improves minimum measurable settling time. You must use careful construction techniques and proven diligent measurements (see **sidebar** "Settling-time-circuit performance verification").

CIRCUIT APPLICATION

It is good practice to "walk" the sampling window backward in time from the settled region to the last 100 µV or so of amplifier movement so that you can observe the cessation of ringing. The sampling-based approach provides this capability and is a powerful measurement tool. Slower amplifiers may require extended delay, sampling-window times, or both, necessitating larger capacitor values in 74HC123 one-shot timing networks. Noise is the fundamental measurement limit with an unfiltered bandpass (**Figure 21**). The DAC (Trace B) settles to 16 bits, 1.7 µsec after Trace A's time-corrected input step. The DAC amplifier's compensation capacitor affects settling

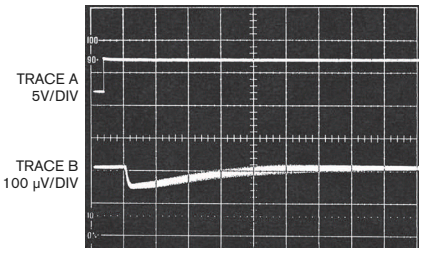

Figure 23 You can discern 19-bit, or 2-ppm, settling 9 µsec after the input command with a 50-kHz bandwidth.

time (see **sidebar** "DAC-amplifier compensation"). Sample-gate feedthrough is undetectable, indicating that higher resolution is possible without overdriving the oscilloscope. Reducing the measurement bandwidth to 250 kHz attenuates noise (**Figure 22**). The 18-bit, or 4-ppm, settling requires approximately 5 µsec. The reduced bandwidth permits higher resolution, although the indicated settling time is likely pessimistic due to the filter's lag. Decreasing bandwidth to 50 kHz permits 19-bit, or 2-ppm, resolution with an indicated settling in 9 µsec (**Figure 23**). The same filter that permits high resolution also creates a longer settling time.

You can use noise-averaging techniques to measure settling time to 20 bits, or 1 ppm (10 µV), without the time penalty of a bandlimiting filter (see **Figure 24** in the Web version of this article at www.edn.com/100304df). You can adjust the DAC amplifier for overdamped, underdamped, and optimum responses. Using averaging eliminates noise, permitting you to determine settling time due to DAC dynamics. In this case, settling time ranges from 4 to 6 µsec, and fractional LSB (least-significant-bit) tailing is evident. This measurement determines DAC settling time due solely to dynamics that a step input initiates. For this reason, you can consider averaged results academic. Noise limits the measurement certainty at any given instant to 100 µV. It is reasonable to maintain that this noise means that the DAC never settles inside this limit. The averaged measurement defines settling time after you remove noise limitations.EDN

For all sidebars related to this article, as well as a list of references, go to www.edn.com/100304df.

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

AUTHOR'S BIOGRAPHY

Jim Williams is a staff scientist at Linear Technology Corp, where he specializes in analog circuit and instrumentation design. He has served in similar capacities at National Semiconductor, Arthur D Little, and the Instrumentation Laboratory at the Massachusetts Institute of Technology (Cambridge, MA). A former student at Wayne State University (Detroit), Williams enjoys sports cars, art, collecting antique scientific instruments, sculpture, and restoring old Tektronix oscilloscopes.


EPCOT

TILT or NEW GAME?

FPGAs' EVOLVING ROLES IN NETWORKING HARDWARE MAY PREDICT THE FUTURE OF THE NEARLY UBIQUITOUS ICs.

BY RON WILSON • EXECUTIVE EDITOR

A long and intimate relationship exists between FPGAs (field-programmable gate arrays) and the networking-equipment industry, dating back to the days of the dot-com bubble. In those heady times, network-hardware vendors were under intense pressure to get new equipment out the door. With relatively slow transceivers and cabling holding back wire speeds outside the network cores, a features race developed among the vendors. The first one out the door with a new set of features claimed most of the pending orders, almost irrespective of price. If switch and router costs were to become ridiculous, the equipment vendors could just pass the cost along to networking start-ups, whose venture capitalists were always ready with another $20 million check, or to service providers desperate not to fall behind.

This climate was perfect for FPGAs because they can support rapid change in design requirements and have moderate operating speeds, little sensitivity to power consumption, and almost no sensitivity to price. FPGA vendors began to sell their largest parts—chips with prices of more than $1000 that had previously found use only a few at a time for prototyping or one-off projects—into the moderate-volume production networking boxes.

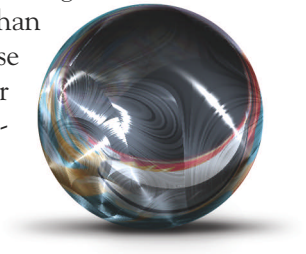

The impact on the FPGA industry was dramatic. The leading companies grew into billion-dollar businesses, creating the cash flow to fund much greater internal research and development. Altera and Xilinx moved from being conservative adopters of new process technology to being among the first few fabless companies at each new process node, which somewhat narrowed the performance gap between their chips and ASICs.

The way switch and router designers used the chips also changed. FPGAs had traditionally been convenient ways to implement glue logic. In the pressurized atmosphere of the Internet bubble, however, designers began to implement more complex functions, too—digital-signal-processing algorithms, framers, and mappers, for example. This trend, in turn, drove the FPGA vendors to reflect the new applications' demands. They began seeding their logic fabric with powerful DSP engines, designing PHY (physical-layer) and MAC (media-access-control) hardware for high-speed serial I/Os onto the chips and enlarging on-chip memory structures to support greater data throughput. They also worked to make their design interfaces more powerful, recognizing that customers weren't just designing glue anymore.

But the dot-com bubble popped. Demand for network hardware plummeted and took FPGA orders with it. The FPGA-engineering work continued, however, even as the vendors scrambled to diversify their markets beyond networking. The result was that high-end FPGAs continued to include powerful features for switch and router applications, even though these features were often going into embedded-computing and control applications.

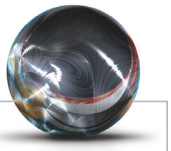

AT A GLANCE

- FPGAs (field-programmable gate arrays) and networking have had a long, symbiotic relationship.
- The pressures that catapulted FPGAs into prominence in networking are reforming.
- It is now feasible to implement most 40-Gbps line-card functions into one FPGA.
- You will soon see networking applications profoundly change the nature of FPGAs.

HERE WE GO AGAIN

Now fast-forward to 2009 and a remarkably similar networking market. At every level—from the outskirts of the access network to the concentrators through which MAN (metropolitan-area-network) carriers route their traffic into the core—a scramble is under way to get more bandwidth and to support new payloads, such as video, and to move away from legacy networks to GbE (gigabit Ethernet).

No one calls the scramble a bubble this time. Once again, though, switch and router vendors are rushing to introduce new generations of equipment. This time, speed is an issue: Network operators are clamoring for 40 Gbps now and 100 Gbps as soon as possible. Complexity is spiraling, as CE (carrier-Ethernet) networks encapsulate and carry data from the old synchronous-network protocols that had maintenance and fail-soft features for carriers. New demands, such as traffic management and security, require more detailed packet inspection, classification, and processing. As a result, equipment designers are turning to a new generation of networking ICs and developing new SOCs (systems on chips) in advanced processes. Some are also looking again at FPGAs.

As the collapse of the dot-com bubble stopped the intense pressure for new features, networking-hardware developers no longer needed the fast time to market that FPGAs offer, and they could no longer pass along to their customers the cost of a boardful of $1500 chips. But FPGAs didn't disappear from routers. The chips could still be a good source of the high-speed transceivers the routers needed for interfaces, especially as moderate-speed transceivers began to appear in smaller, lower-cost FPGA families.

The chips are also useful for their flexibility in some difficult situations. "We use FPGAs only in critical time-to-market situations or to implement functions for which the requirements are still changing," explains Martin Skagen, senior director and chief architect at Brocade. "Once the requirements are stable, we tend to lock down the design in an ASIC."

The FPGA vendors recognize this reality. "Today, the established box makers have ASIC designs," says Arun Iyengar, senior communications-business-unit manager at Altera. "They use FPGAs for modifying the function of the ASICs or for glue." That state of affairs is what's occuring now, but it may not be in the future. The new pressures for bandwidth and traffic management to cope with TV over IP (Internet Protocol), for example, are bearing down on the carriers. As a result, CE equipment is starting to experience accelerating, bubble-like change.

FPGAs IN CE

"As speeds increase, we are seeing a transition in CE," says Morteza Ghodrat, director of CE technology at Vitesse

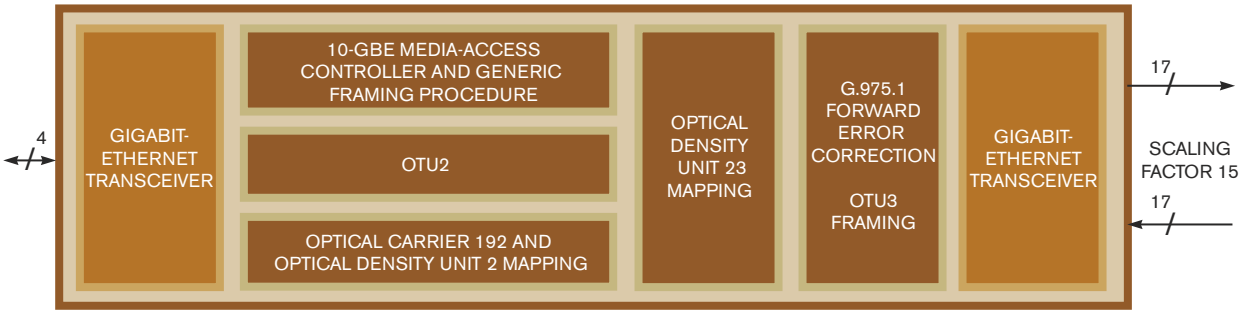

Figure 1 A single FPGA can handle a full set of packet-processing tasks, as in this example from Xilinx.

Semiconductor. "Packet processing is shifting from network processors to FPGAs as architects try to cobble together enough components to support CE functions at increasing wire speeds. This [scenario] is happening now, but it may not be a viable approach at 40 Gbps. The future is integrating functions into SOCs and then using smaller FPGAs as backplane or link drivers or to offload carrier-proprietary functions, such as Layer 3 services or DPI [deep-packet inspection]."

What SOC vendor Vitesse sees as a transitional form, however, Altera and Xilinx see as the leading edge of a permanent transition. "What's changing people's thinking is that wire-speed DPI just isn't possible for an NPU [network-processing unit] operating faster than 10 Gbps," says Iyengar. "So architects of 40-Gbps systems look to something like an [Altera] Arria-sized FPGA to offload the inspection." The next logical step is to ask what more you could do with more FPGA.

The answer, Iyengar says, is a great deal more. "Today's 40-nm parts are fast and dense enough to bring in four lanes at 40 Gbps and handle packet processing and inspection at wire speed. So a large FPGA can handle the whole packet-processing task, not just one inspection stage," he adds (**Figure 1**). This approach requires the FPGA to support both an external TCAM (ternary-content-addressable-memory) interface and a significant amount of DRAM. "You can use an algorithmic search engine and DRAM in place of the TCAM," Iyengar says. "But, generally, an FPGA doing DPI, QOS [quality-of-service] functions, and policing is going to need about four banks of DDR-3 DRAM." Ironically, the high-speed transceivers that made the FPGAs such good glue are now gluing memory and other engines to the FPGAs.

Gilles Garcia, business group director at Xilinx, agrees. In the short term, FPGAs may enter the architecture to increase differentiation, but, in the long run, "FPGAs can make ASSPs [application-specific standard products] irrelevant," Garcia says. "All 40- and 100-Gbps designs are using FPGAs in one way or another." To some extent, Garcia admits, the FPGAs are there because off-the-shelf parts aren't ready. Trials of 100-Gbps designs are going on, and some vendors are developing 160-Gbps devices. Only one commercial framer/mapper chip is available for that kind of speed, however, he claims. And if you put together the needs of a 100-Gbps product line, including packet processing, traffic management, QOS, and security, he says (**Figure 2**), "SOCs are just nonstarters."

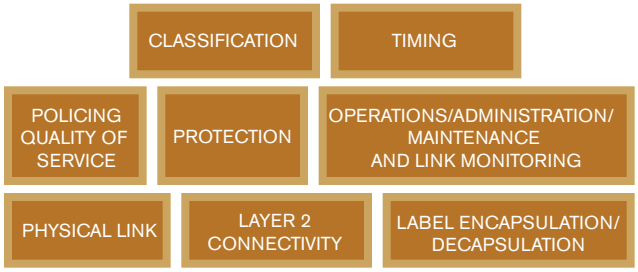

Figure 2 A single FPGA can handle a wide variety of carrier-Ethernet functions.

Analyst Bob Wheeler of the Linley Group doesn't see NPUs disappearing just yet. Although Wintegra and LSI are focusing on 10-Gbps or slower applications in the access network, EZchip, Xelerated, and Broadcom are all in the game for MANs as they move to 100 Gbps. The NPU and SOC vendors may have longer chip-development cycles than do FPGA users, but they also have enormous advantages in speed, density, and power.

TODAY'S 40-nm PARTS ARE FAST AND DENSE ENOUGH TO BRING IN FOUR LANES AT 40 GBPS AND HANDLE PACKET PROCESSING AND INSPECTION AT WIRE SPEED.

A vendor that wishes to remain anonymous agrees with the FPGA suppliers' view. This company uses FPGAs to do all the per-packet heavy lifting in an advanced Layer 3 flow-control system. The system manages to use a little of everything: ASSP switches on the front end, an NPU for setting up and taking down flows, and the FPGAs for inspecting individual packets. This arrangement allows a 1U box to handle 40 Gbps. The company doesn't see the FPGAs as interim solutions. By the time 100-Gbps links are common, the company expects to have access to FPGAs fast enough to do the work.

This vendor uses FPGAs much as Vitesse's Ghodrat describes: as one component in a heterogeneous pipeline. But other design teams are moving in the direction Xilinx's Garcia suggests: toward eliminating ASSPs and NPUs from their product lines. One example is Danish intellectual-property vendor TPack. Lars Pedersen, chief technology officer at the company, says that his organization is focusing on the rise of the OTN (optical-transport network) at the MAN level as well as the expansion of fast Ethernet. "We see both 40- and 100-Gbps Ethernet coming this year, as well as new MPLS (multiprotocol-label-switching) and OTN standards with more emphasis on packet transport," he says. Pedersen sees the same issues all across the network, from access aggregation to MAN routing. As packet-based services, such as video, dominate, carriers start wanting all packets delivered in the order and with the timing with which they arrived. This level of service is far beyond conventional Ethernet switching, and this QOS must be in every node. "Operators want the same functionality on every box—from a 5-Gbps pizza box to a 5-Tbps rack," he says.

TPack's approach is to create the functions for a line card in RTL (register-transfer-level) logic and then synthesize the design into FPGAs for customers. The latest 40-nm Altera parts, Pedersen claims, allow TPack to put a full 40-Gbps Ethernet switch into one FPGA, with all the carrier-class features operators are requesting. Pedersen says that the 5-Gbps pizza-box implementation fits into one $50 FPGA. The plans don't stop there. TPack last year began designing a 100-Gbps, two-FPGA Ethernet device, which he expects to release in the second quarter of this year. One FPGA, the packet processor, connects through a 100-Gbps Ether-

TO STAY IN THE NETWORKING MARKET, AN FPGA VENDOR MUST ARRIVE EARLY WITH FAST-ENOUGH TRANSCEIVERS AND FAST-ENOUGH FABRIC.

net CAUI (100-Gbps attachment-unit-interface) port to an optical module. The chip implements Ethernet/MPLS-TP (transport-profile) switching and tunneling; protection and timing facilities; and an Interlaken interface, a standard protocol for packet transfers between components in communications systems. The second FPGA, a packet manager, provides per-flow queuing and policing. This chip also has an Interlaken interface, and both FPGAs have private DRAM connections. Together, the two FPGAs will provide essentially the same level of CE wire-speed functions as planned 100-GbE ASSPs will.

TPack officials see FPGAs taking all the functions of the line card, but Altera's Iyengar sees them going even further. "Some architects are looking at an FPGA with 48 transceivers and asking why they couldn't use it to build switch fabric," he says. "Such a design would be easily upgradable to new speeds and might make a lot of sense. But it is a new area for us."

BRACKETING THE FUTURE

Clearly, it is possible to implement an entire line card's functions in FPGAs. For markets in transition—emerging 100-Gbps CE or new-generation OTN, for example—the flexibility of having everything in FPGAs may prove essential. Still, Brocade's Skagen points out that, as standards mature, FPGAs quickly lose their luster. Being able to put all the functions for CE switching into FPGAs is one thing, but wanting to do so is another. "I can't agree that FPGAs are approaching SOC capabilities," Skagen says. "For complex functions, a cell-based design will be 15 or 16 times denser than an FPGA. The speed will be much higher, as well. FPGAs are simply not fast enough to sit in our pipeline. And ASICs are an order of magnitude more energy-efficient. Keeping the power down on the chips makes life a lot simpler for everyone."

The role for FPGAs, then, remains tightly circumscribed. On one boundary, the chips are helpful to network-equipment manufacturers only when they are available with fast-enough transceivers and logic fabric to keep up with wire speeds. On the other boundary, FPGAs become rapidly less competitive as soon as standards gel enough to implement with firmware-programmed ASICs or ASSPs. To stay in the networking market, then, an FPGA vendor must arrive early with fast-enough transceivers and fast-enough fabric. And the vendor must quickly drive these features down market to smaller, less expensive devices before alternative ways of implementing the line card start to appear. Only with less expensive devices can the FPGAs retreat into their customary special-function role instead of having vendors sweep them from their product lines.

This scenario may help explain why Altera last month began discussing its 28-nm generation—long before the products' availability. The company is anticipating the next step in networking hardware beyond 100 GbE and weighing the strengths and weaknesses of its 28-nm-process capabilities against the coming task. "We see transmission going to 100 Gbps and then right on to 400 Gbps," says Luanne Schirrmeister, senior director of component products at the company. "Today, it takes over 350,000 logic elements to implement the front-end block for 100 GbE: That includes MACs and the Interlaken interface," she says. "To move this [speed] up to 400 Gbps is just impractical in 40-nm FPGAs. But even moving to 28 nm by itself doesn't solve the problem." The transceiver isn't the issue. Schirrmeister says that Altera's 28-nm chips will offer enough 28-Gbps transceivers to support 400-Gbps ports. The problems are the speed, density, and power of the programmable-logic fabric. There will not be that much more logic density or less power and little increase in speed in moving from 40 to 28 nm.

So instead of counting on scaling, Altera will offer Embedded HardCopy.

Altera and its customers will be able to implement certain blocks using the company's HardCopy metal-programmed ASIC capability: a step roughly halfway between implementing in programmable-logic fabric and doing a full cell-based ASIC implementation. Blocks that designers put into HardCopy will thus be denser, faster, and lower in power than in programmable logic but not as much so as cell-based portions of the FPGA, such as the DSP blocks. Altera will place these HardCopy blocks inside the FPGA chip for optimum routing to the other resources. The result will be an application-directed FPGA, with certain functional blocks hard-embedded in an otherwise field-programmable chip.

This move is necessary to meet the networking industry's needs. It will also influence the way FPGAs relate to the rest of the electronics industry. Application-directed parts will move rapidly down the price curve, becoming attractive to users in other applications but with algorithmically similar problems. And the Embedded HardCopy design flow will be available to all customers. Once again, as with the long-ago appearance of DSP blocks and high-speed transceivers, features necessary to the networking industry will become available and affordable to everyone else. The wise designer will figure out how to apply them. **EDN**

FOR MORE INFORMATION

Altera
www.altera.com

Broadcom
www.broadcom.com

Brocade
www.brocade.com

EZchip
www.ezchip.com

Linley Group
www.linleygroup.com

LSI Corp
www.lsi.com

TPack
www.tpack.com

Vitesse Semiconductor
www.vitesse.com

Wintegra
www.wintegra.com

Xelerated
www.xelerated.com

Xilinx
www.xilinx.com

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

EDITED BY MARTIN ROWE
AND FRAN GRANVILLE

designideas

READERS SOLVE DESIGN PROBLEMS

Circuit lets you test sample-and-hold amplifiers

Marián Štofka, Slovak University of Technology, Bratislava, Slovakia

Sample-and-hold amplifiers sample an analog voltage and hold it until an ADC can digitize it. A perfect sampling circuit holds a voltage until digitizing is complete. Thus, the amplifier's output is identical to its input. Real sample-and-hold amplifiers, however, can gain or lose voltage, producing an error. Offset voltages in amplifiers cause a static additive error. Further, there occurs a specific additive error, the so-called voltage pedestal, which originates within the transition from the sample state to the hold state because of a parasitic charge transfer to the hold capacitor.

A sample-and-hold amplifier uses an analog switch to connect a signal to a holding capacitor. When the switch closes, thus having low resistance, the capacitor charges to the sampled input voltage. During the hold time, when the switch has high resistance, the sampling capacitor holds the voltage until the ADC digitizes it. During the transition from low to high switch resistance, a parasitic charge injection, mainly from the gate of the switch to the hold capacitor, continues to charge the capacitor until the switch's control voltage reaches a steady logic level. The injected charge produces an error voltage at the capacitor. Additional errors may occur during the hold time. Leakage and bias currents in the amplifier combine with tens of picoamps of leakage current in the switch and capacitor to cause the capacitor to charge or discharge during hold time.

By applying a logic-control signal with a duty cycle of D and 1−D, you can measure a mean output voltage difference, $[\Delta V_{OUT}]=[V_{OUT}-V_{IN}]$, which the following **equations** show. $[\Delta V_{OUT}]=V_{STAT}+(1-D)V_{INJ}+½(1-D)^2V_{DROPPEAK}$, and $[\Delta V_{\overline{OUT}}]=V_{STAT}+DV_{INJ}+½D^2V_{DROPPEAK}$, where ΔV_{OUT} and $\Delta V_{\overline{OUT}}$ are the output-voltage differences for D and 1−D, respectively; V_{STAT} is the steady output-voltage difference for a selected value of the reference input voltage, D is the duty cycle, V_{INJ} is the voltage pedestal, and $V_{DROPPEAK}$ is the peak voltage drop. **Figure 1** shows how the voltages in the **equations** change over time. If you apply a complementary control waveform with a duty cycle of 25%, you can measure another dc component of the sample-and-hold amplifier's output voltage. Finally, when the sampling switch is continuously on, you can measure the V_{STAT} voltage, which is a real dc voltage. V_{OUT} and $V_{\overline{OUT}}$ contain a waveform superimposed onto a selected value of the reference voltage. Thus, you should measure the mean values of these voltages using a series resistor with a value of, say, 10 kΩ.

Multiplying the voltage pedestal, a simple rectangular waveform, by the duty cycle yields the av-

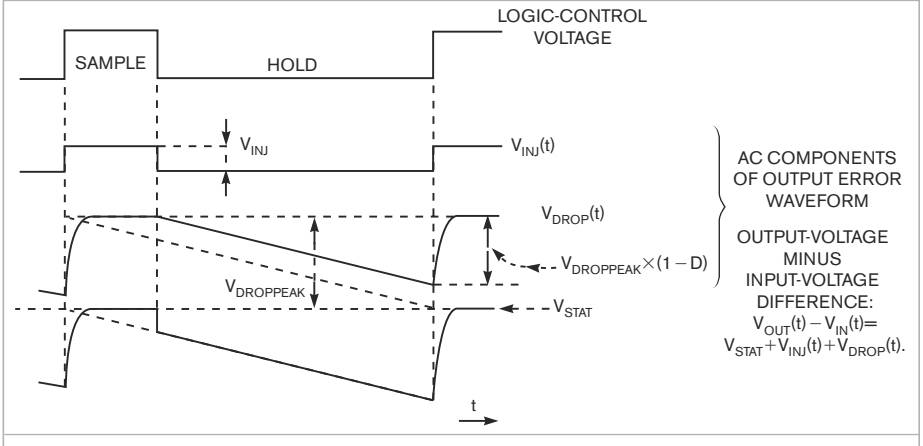

Figure 1 A sample-and-hold amplifier's holding capacitor experiences a voltage drop due to leakage and bias currents plus a voltage step, which results in a difference between the amplifier's output and input voltages.

DIs Inside

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

LINEAR TECHNOLOGY

DESIGN NOTES

Precision, Matched, Baseband Filter ICs Outperform Discrete Implementations – Design Note 476

Philip Karantzalis

Introduction

In digital communication systems, baseband signals must be band-limited in the transmitter or the receiver. Although the bulk of baseband signal shaping and analysis is accomplished using digital signal processing (DSP), analog filtering is used in a number of places along the signal chain. For instance analog filters reduce the imaging of a digital-to-analog converter (DAC), filter out the high frequency noise of an RF demodulator or reduce the aliasing inputs of an analog-to-digital converter (ADC).

Typically, 3G communication systems (CDMA, GSM, UMTS or WiMax) feature a baseband channel bandwidth of 1.25MHz to over 20MHz. In this frequency range, discrete analog filters—those constructed with high speed op amps, resistors and capacitors—are sensitive to PCB layout parasitics, component tolerances and mismatches. The pitfalls of using discrete components can be avoided by using integrated, pin-configurable, precision analog filter ICs, such as the LTC6601-1/-2 and the LTC6605-7/-10/-14. The LTC6601-x is a single 2nd order lowpass filter and the LTC6605-x is a dual, matched filter.

The LTC6601-x Lowpass Filter

Figure 1 shows a block diagram of an LTC6601-x and the 2nd order function it implements. High frequency filters are easily implemented using the LTC6601-x, which integrates a low noise (1.5nV$\sqrt{Hz}$) wideband (600MHz), fully differential amplifier with precision resistors and capacitors. The standard deviation of the on-chip resistors and capacitors is ±0.25% and their matching is ±0.1% (in a differential amplifier, the common mode rejection at high frequencies depends on tight matching of the signal paths). Furthermore, the gain bandwidth (GBW) of the LTC6601-x amplifier is trimmed to ±5%. Note that this level of precision cannot be achieved with discrete analog filters in any practical manufacturing process.

In many LTC6601-x filter implementations, no external components are required. For instance, lowpass filters can be produced by simply hardwiring the input pins for a variety of filter gain and cutoff frequencies from 5MHz to 27MHz.

The product and ratio of the on-chip resistors and capacitors determine the f_0 and Q values of the 2nd order filter function. The f_0 and Q pair sets the filter's cutoff frequency and passband gain peak. The value of the feedback resistor determines the range of the f_0 frequencies. The 400Ω feedback resistors can be shunted by input resistors to increase the f_0 range.

Using an LTC6601-x with input RC or LC filters, 3rd, 4th and 5th order lowpass filters can be implemented (refer to the LTC6601-1 or -2 data sheets for higher order filter options).

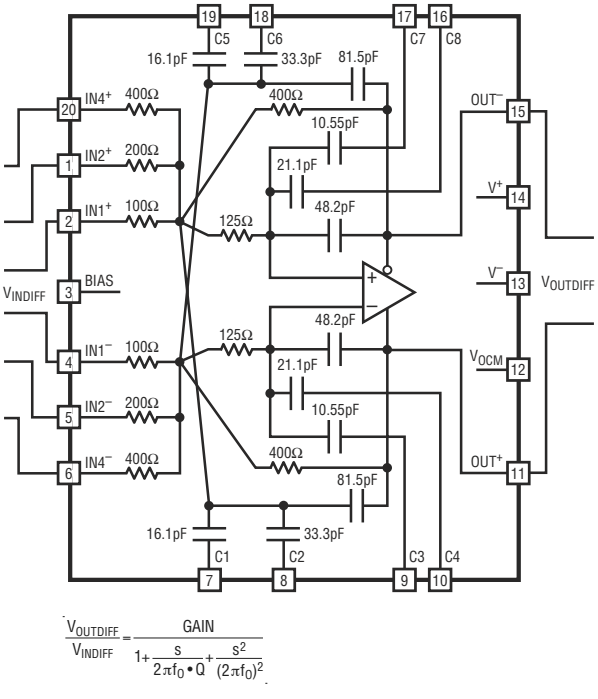

Figure 1. Block Diagram of the LTC6601-x and the 2nd Order Transfer Function it Implements.

The LTC6601 family is offered in two options that trade off distortion and noise. Use the LTC6601-1 for low noise and the LTC6601-2 for low distortion at low power.

The LTC6605-x, Dual, Matched, Lowpass Filter

Typically, quadrature down conversion is used in a direct conversion or zero-IF receiver. In a quadrature demodulator, the RF signal is split into two paths and mixed with the LO (local oscillator) to produce an I (in phase) and a Q (90° phase or quadrature) signal. In a direct conversion receiver, high image rejection depends on very tight gain and phase matching. High image rejection maximizes the signal-to-noise ratio (SNR) and minimizes the bit error rate.

The LTC6605-x contains two LTC6601 ICs configured and tested as a dual, matched, 2nd order, lowpass filter. There are three LTC6605 versions: the LTC6605-7 with an adjustable $f_{(-3dB)}$ range of 6.5MHz to 10MHz, the LTC6605-10 with an adjustable $f_{(-3dB)}$ range of 9.7MHz to 14MHz and the LTC6605-14 with an adjustable $f_{(-3dB)}$ range of 12.4MHz to 20MHz (the $f_{(-3dB)}$ or $f_{(-1dB)}$ frequency is set by an external resistor).

The maximum gain and phase matching error of an LTC6605-x is ±0.35dB and ±1.2° respectively (equivalent to –32dB of image rejection). This level of gain and phase matching of an LTC6605-x IC is impractical using discrete resistors and capacitors.

The LTC6605-7 or -10 input pins can be hard wired for gains of 1, 4 or 5; the LTC6605-14 for gains of 1, 2 or 3. Either can produce noninteger gains using external input resistors with a slight reduction in gain and phase matching.

Accurate gain and phase matching, when combined with very low noise and distortion, allows for a high dynamic range differential circuit. Figure 2 shows an LTC6605-7 driving a dual LTC2205 16-bit ADC and Figure 3 shows the FFT plot of the ADC output. The LTC6605-7 is configured for a 12dB gain and a 7MHz, –1dB frequency (suitable application for WiMax).

As with an LTC6601-x, an LTC6605-x can be used with input RC or LC filters, to implement 3rd, 4th and 5th order lowpass filters. The LTC6605-7 is available in a compact 6mm × 3mm, 22-pin leadless DFN package.

The power consumption for an LTC6601-x and for each LTC6605-x 2nd order section, is set by a three-state BIAS pin, allowing a choice between shutdown (I_S = 350µA), medium power (I_S = 16mA) or full power (I_S = 33mA).

Conclusion

The LTC6601-x fully differential, lowpass filter with precision on-chip resistors and capacitors can be configured by hardwiring pins to implement 2nd order filters in a frequency range of 5MHz to 27MHz.

The LTC6601-x is insensitive to PCB parasitics and is a higher performance circuit than a discrete analog filter. The LTC6605-x is a dual, matched, 2nd order lowpass filter for driving dual ADCs in a high performance direct conversion receiver. Using RC or LC filters with an LTC6601-x or LTC6605-x, 3rd, 4th or 5th order lowpass filters can be implemented.

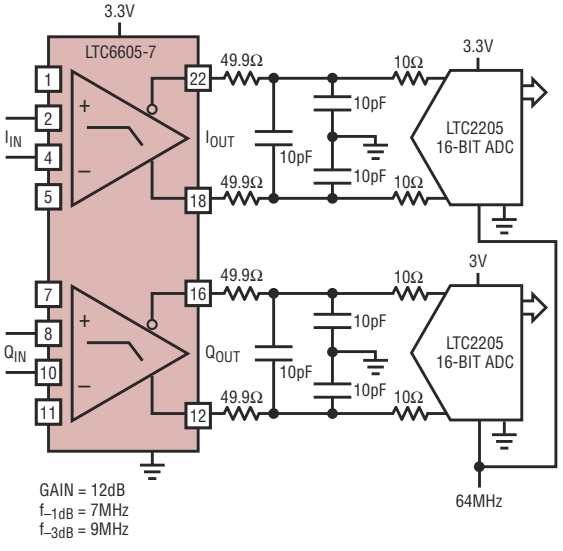

Figure 2. An LTC6605-7 Driving Two LTC2205 ADCs

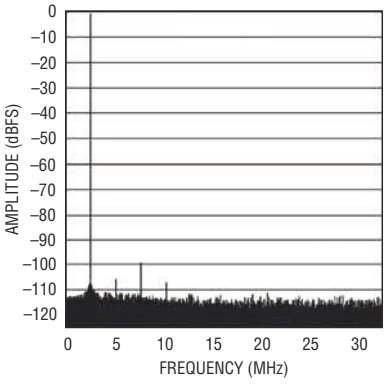

Figure 3. The FFT Plot of the LTC2205's I or Q Output Channels of Figure 2 (75dB SNR, –98dB SFDR, 64K-Point FFT, 64Msps, f_{IN} = 2.5MHz, –1dBFS).

Data Sheet Download

www.linear.com

For applications help,
call (408) 432-1900, Ext. 3761

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

dn476f LT/TP 0310 116K • PRINTED IN THE USA

erage value. In contrast, the voltage-drop waveform appears as a sawtooth. Its mean rises as one-half of the duty cycle squared. The peak-voltage-drop value denotes a hypothetical voltage drop at the end of a whole period, T, of the SAMPLE/$\overline{\text{HOLD}}$ logic-control waveform.

You can use the previous **equations** to find the values of the voltage pedestal and the peak voltage drop. A 75% duty cycle is a convenient value. The following **equations** are valid for this duty cycle: $V_{INJ}=6[\Delta V_{OUT}]-2/3[\Delta V_{\overline{OUT}}]-16/3V_{STAT}$, and $V_{DROPPEAK}=16[-[\Delta V_{OUT}]+1/3[\Delta V_{\overline{OUT}}]+2/3V_{STAT}]$. You must find the optimal repetition rate, f_{REP}, of the logic-control signal. As the optimal repetition rate increases, the difference in output voltage from the input is almost purely due to dc voltage offset plus the voltage pedestal: $(V_{\overline{OUT}}-V_{STAT})/(V_{OUT}-V_{STAT})\simeq 3$. The following **equation** finds the maximum value for the optimal repetition rate: $f_{REP}\leq(0.01/4)\times 1/(t_{ON}-t_{OFF})$, where t_{ON} and t_{OFF} are the on and off times, respectively. This **equation** ensures that the difference in values between the turn-on and turn-off times of the sample-and-hold amplifier's internal analog switch won't affect the accuracy of the precision 25 and 75% duty cycles by more than 1%.

If you evaluate the **equation** for a high-performance analog switch, such as the Analog Devices (www.analog.com) ADG1213, you get a repetition rate of 33 kHz or less. The difference due to voltage drop prevails at low-value repetition rates. In this case, the repetition rate can be the value of the frequency at which $V_{\overline{OUT}}-V_{STAT}\leq 1/10\times V_{INMAX}$, where V_{INMAX} is the maximum input-voltage range. The best way to determine the lower limit of the repetition rate is through experimentation.

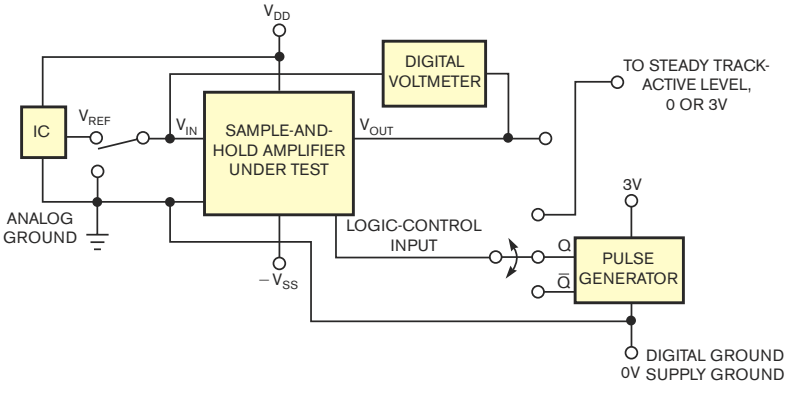

Figure 2 Use a digital voltmeter to measure the difference between the output and the input voltages of a sample-and-hold amplifier.

A tested sample-and-hold amplifier using the circuit in **Figure 2** uses a supply voltage of −1V, a drain-to-drain voltage of 5V, and a supply voltage of 3.3V for logic circuits in the pulse generator. Two sets of measurements at 25, 75, and 100% duty-cycle values for the AGD1213's internal switch control used input voltages of 0 and 2.5V. You will measure the output-voltage difference, approximately −0.0366 mV, and the pedestal voltage, approximately −0.0333 mV, at a repetition rate of 1.762 kHz. The value of the residual effective charge injection, Q_{INJ}, into the hold capacitor, $C_H=2$ nF, is $Q_{INJ}=C_H\times V_{INJ}$. The value is negative and doesn't exceed −75 fC. The following **equation** defines the difference of charge injection within the 2.5V range of input voltage: $\Delta Q_{INJ}=Q_{INJ}(2.5V)-Q_{INJ}(0V)$ and yields a value of −6.7 fC. The following **equation** determines the residual effective leakage current from the acquired values of peak voltage drop at a repetition rate of 160 Hz: $I_{LEAK}=C_H\times V_{DROPPEAK}\times f_{REP}$, where I_{LEAK} is the leakage current. A leakage current at the input voltage of 0V is approximately 17 pA, and a leakage current at the input voltage of 2.5V is approximately −17 pA. EDN

REFERENCE

1 "Low Capacitance, Low Charge Injection, ±15 V/+12 V iCMOSTM Quad SPST Switches," Analog Devices Inc, 2005, www.analog.com/en/switchesmultiplexers/analog-switches/adg1212/products/product.html.

Add hysteresis to a voltage comparator

Luca Bruno, ITIS Hensemberger Monza, Lissone, Italy

Positive feedback is a typical technique for distributing hysteresis around a comparator, provided that you have a resistive path between the comparator's output and the noninverting input. Positive feedback forms two threshold voltages that have (or assume) fixed values. In addition, they depend on the saturation values of the comparator's output stage. Plus, the load conditions affect their accuracy. The circuit in **Figure 1** provides an alternative for applications requiring a comparator with hysteresis that has precise thresholds that you can easily and independently set. The circuit includes two inverting and noninverting threshold comparators whose outputs directly drive a set/reset latch. You can use a latch with either active-low or active-high inputs.

You can generate the positive and negative threshold voltages using a

precision voltage reference that powers a resistor divider (not shown) or by driving the inputs with DACs if you need a digitally programmable comparator. The circuit's high input impedance facilitates this task. Because of its hold state, the latch nullifies the effects of frequent switching on the comparator's outputs due to noise on the input signal. The circuit thus acts as a Schmitt trigger even if there is no positive feedback. The latch introduces a propagation delay that's normally a few tens of nanoseconds and is negligible in low- to medium-speed applications. Because the latch has complementary outputs, the circuit provides a noninverting characteristic on the Q output and an inverting characteristic on the $\overline{Q}$ output (**Figure 1a** and **b**).

Some integrated latches have only the Q output. If you need an inverted output, you need only to exchange the comparator outputs with the latch inputs for both circuits; the upper comparator drives the reset input, and the lower comparator drives the set input. You can use open-collector or open-drain comparators to process bipolar or positive signals higher than the supply voltage of the latch. You can easily interface them without using clamping diodes. You must add only a pullup resistor that the logic supply powers.

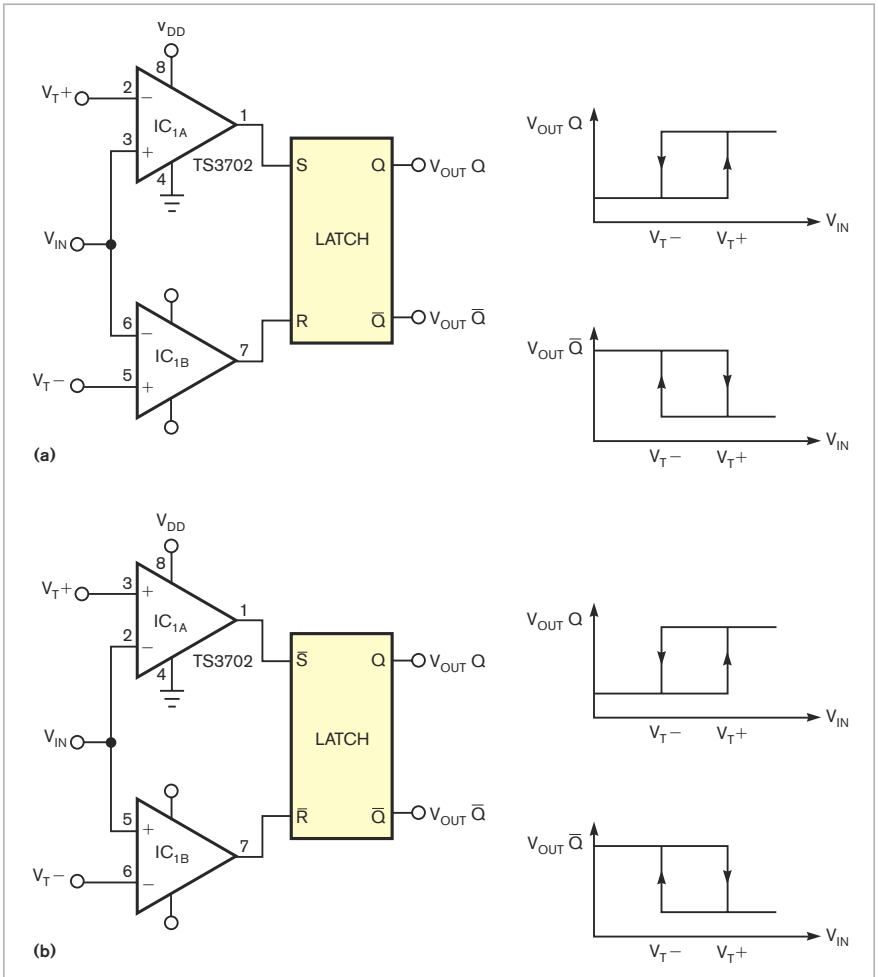

Figure 1 The set/reset latch provides hysteresis for the comparator. The circuit provides a noninverting characteristic on the Q output and an inverting characteristic on the $\overline{Q}$ output (a and b). Exchange the comparator outputs with the latch inputs for both circuits: The upper comparator drives the reset input, and the lower comparator drives the set input, and you get an inverting characteristic on the Q output and a noninverting characteristic on the $\overline{Q}$ output.

The circuit uses IC_1, an STMicroelectronics (www.st.com) dual micropower comparator with a push-pull output stage. In this case, the supply voltage must be the same as that for the latch. EDN

Broken-coil detector is simple yet robust

Juan Pablo Caram, Santiago, Chile

The circuit in this Design Idea was originally designed to detect damaged conveyor belts in the mining industry. Thin coils are embedded in the conveyor belt. If the belt suffers damage, it stretches at the affected location, causing one or more coils to break. The method for detecting the broken coil is to allow a "sensing" coil to magnetically couple with the passing coils in the belt, thus changing the total inductance of the magnetic pair. The sensing coil is part of an LC oscillator (**Figure 1**). When an intact coil passes the sensing coil, the frequency of the oscillator changes. If the conveyor belt moves at a fixed speed, the frequency of the oscillator modulates

Isolated RS485 + 1W Power

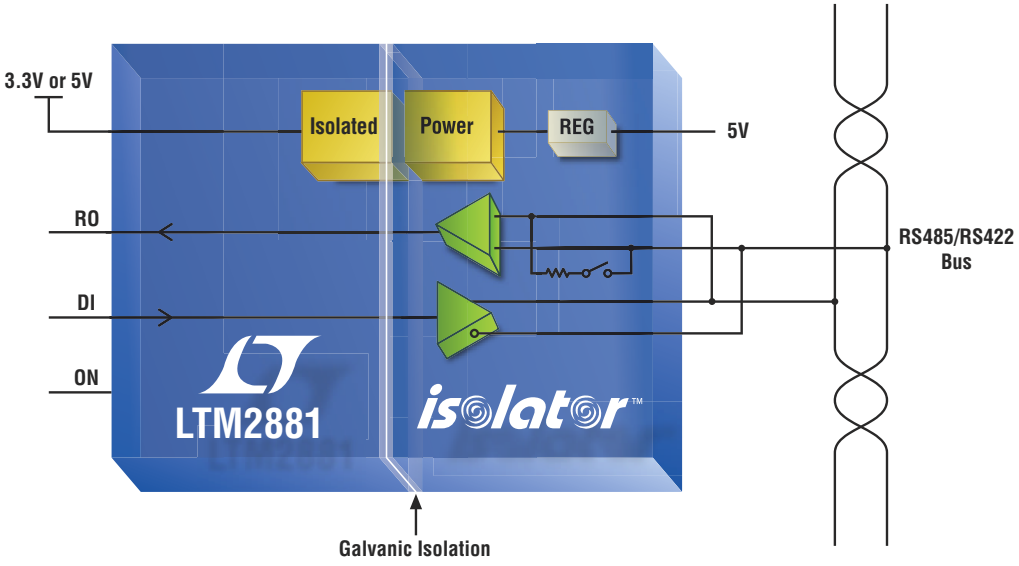

Complete 20Mbps µModule® Transceiver Includes 2500V_{RMS} Isolated Power—No External Components Required

The LTM®2881 is an isolated RS485 transceiver that guards against large ground-to-ground differentials. The LTM2881's internal inductive isolation barrier breaks ground loops by isolating the logic level interface and line transceiver. An onboard DC/DC converter provides power to the transceiver with an isolated 5V supply output for powering additional system circuitry. With 2500V_{RMS} galvanic isolation, onboard secondary power and a fully compliant RS485 transmitter and receiver, the LTM2881 requires no external components and provides a small, complete µModule solution for isolated serial data communications.

Features

- Isolator µModule Technology
- Isolated RS485/RS422 Transceiver: 2500V_{RMS}
- Integrated Isolated, 1W DC/DC Converter
  - Good Efficiency (up to 62%)
  - Low EMI
- 3.3V or 5V Input Supply Voltage (LTM2881-3/LTM2881-5)
- 20Mbps or Low EMI 250kbps Data Rate
- High ESD: ±15kV HBM
- Common Mode Transient Immunity: >30kV/µs
- Integrated Selectable 120Ω Termination
- Small Footprint, Low Profile (11.25mm x 15mm x 2.8mm) in Surface Mount LGA & BGA Packages

LTM2881 Demo Board

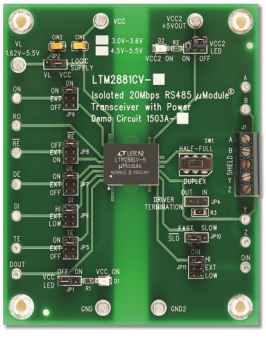

Info & Free Samples

www.linear.com/2881

1-800-4-LINEAR

Free Telecom, Datacom and Industrial Power Products Brochure

www.linear.com/48vsolutions

LINEAR TECHNOLOGY

at a fixed rate. When a broken coil passes the sensing coil, there is more time between modulations (and this is what you want to detect).

The oscillator doesn't generate a pure sine wave and is power-hungry but stable. It oscillates over a large range of LC pairs, even with a low quality factor and with almost any transistor. The amplitude is flat across a large bandwidth. The frequency is $1/(2\pi\sqrt{L_1C_1})$, where L_1 is the inductance of the sensing coil. R_2 represents its resistance. L_2 and switch S_1 represent either undamaged or broken coils in the conveyer belt. When S_1 is closed, the coils are undamaged, and when S_1 is open, the coils are broken. When the coupling between the sensing coil and a conveyor coil is perfect, it is equivalent to having the two connected in parallel, which would reduce the total inductance and increase the frequency of oscillation.

The problem now becomes how to detect different frequencies by implementing an FM demodulator or frequency-to-voltage converter. An easy way to accomplish this task is to pass the signal from the oscillator through an appropriately tuned lowpass filter. If the frequency range of the oscillator lies at the beginning of the roll-off, a higher frequency causes higher attenuation. At this point, an FM waveform has become an AM waveform, which you can easily demodulate using envelope detection.

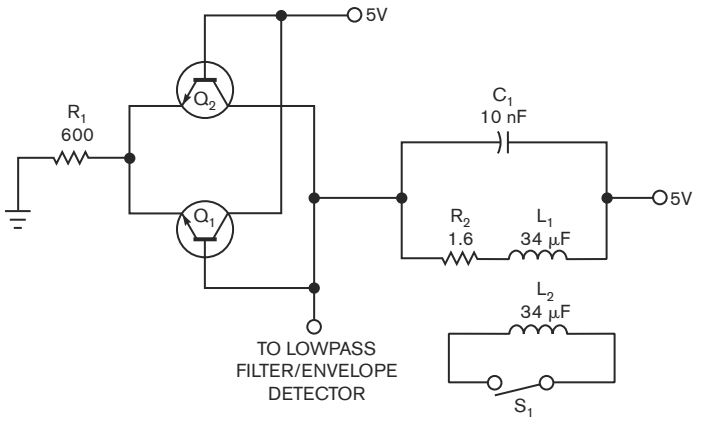

Figure 1 An oscillator generates a stable signal based on the values of L_1, L_2, and C_1.

Figure 2 shows a trivial RC lowpass filter comprising R_5 and C_2 and driving the envelope detector comprising D_1, R_6, and C_3. To find out whether the sensing coil couples to an external coil, you can check at the output of this circuit to see whether the voltage is above or below a certain threshold. You can experimentally determine the best value for the threshold. You can perform the comparison with an analog comparator or with a microcontroller after digitizing the signal. This last method would allow you to measure the time since the most recently detected undamaged coil passed. **Fig-**

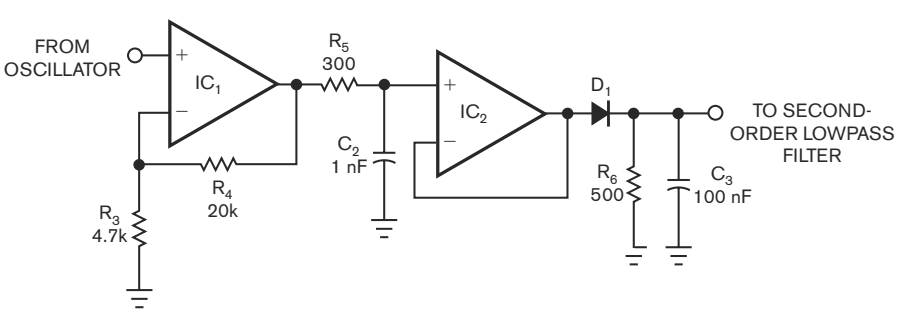

Figure 2 A lowpass filter changes the oscillator output into an envelope signal.

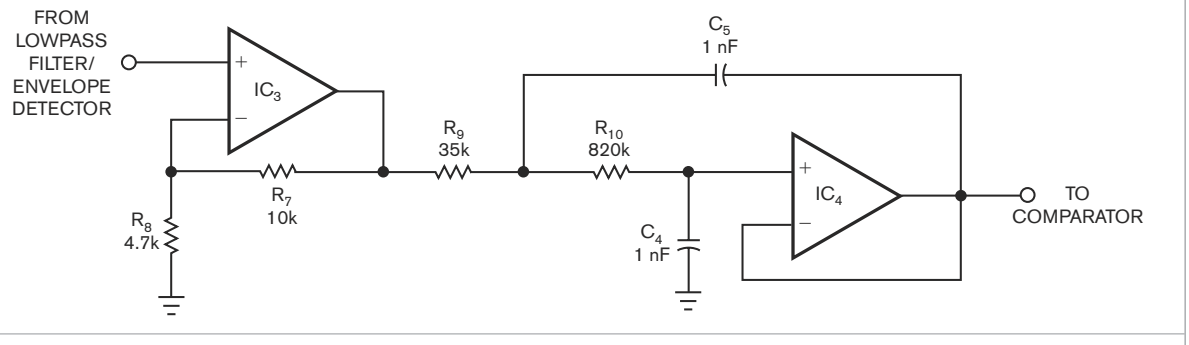

Figure 3 A second-order lowpass filter removes ripple from the envelope detector's output.

High-performance, miniature crystal oscillators

Deliver frequencies to 622.08MHz and less than $1ps_{RMS}$ jitter

The DS4-XO crystal oscillators use fundamental AT-cut crystals to enhance frequency stability over temperature. They provide higher stability (±50ppm) than SAW-based oscillators (±100ppm) over the extended temperature range, and they eliminate the spurious modes of operation associated with overtone designs. Their miniature ceramic package allows precise placement in designs, thus making them ideal for Fibre Channel, InfiniBand™, and other networking applications.

Features

- **Low jitter: < $1ps_{RMS}$ from 12kHz to 20MHz**
- **Wide frequency range: 77.76MHz to 622.08MHz**
- **Small size: 5mm x 3.2mm x 1.49mm**
- **High frequency stability: ±50ppm (-40°C to +85°C)**
- **Low voltage: 3.3V ±5%**

Applications

- **Fibre Channel:** DS4106, DS4212, DS4425
- **InfiniBand:** DS4125
- **Ethernet, 10GbE:** DS4125, DS4156, DS4312
- **SAS/SATA:** DS4150
- **SONET/SDH:** DS4155, DS4311, DS4622, DS4776
- **PCI Express®:** DS4100H

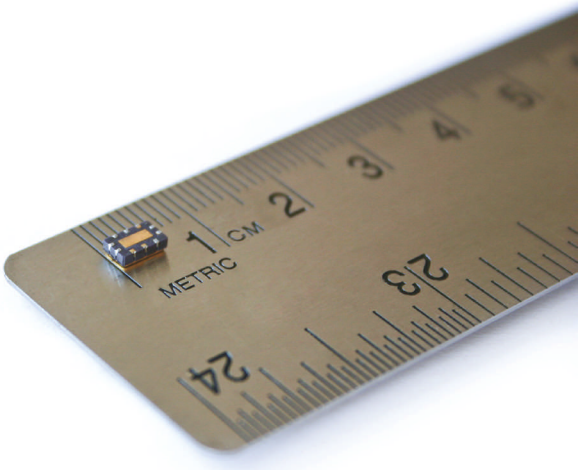

Housed in a tiny, 5mm x 3.2mm x 1.49mm LCCC ceramic package, the DS4-XO crystal oscillators offer the industry's highest performance and save critical board space.

InfiniBand is a trademark of the InfiniBand Trade Association.
PCI Express is a registered trademark of PCI-SIG Corp.

www.maxim-ic.com/DS4-XO

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.

ures 3 and **4** show a simple and reliable analog method, using the envelope signal, to detect that a bad coil has passed. The output from the envelope detector contains ripple, so lowpass filtering allows a more precise discrimination in frequency. In this example, the filter in **Figure 3** makes the ripple insignificant without deteriorating the frequency response of the system.

The output of the filter then feeds into the comparator (**Figure 4**). R_{14} sets the threshold, which should be at the midpoint between the generated voltages with and without an external coil coupled to L_1. D_2, C_6, R_{12}, and comparator IC_6 behave as a countdown timer set to a time slightly longer than the period between passing coils. C_6 quickly charges to a maximum voltage when a good coil passes and then slowly discharges. If the time between (sensed or detected) consecutive undamaged coils is below a certain maximum, the voltage across C_6 should never go below the threshold that R_{15} sets, thus keeping the output of IC_6 low and lighting LED D_3. If a damaged coil passes, the oscillator's frequency should remain the same, allowing the voltage across C_6 to drop enough to trigger the comparator, turning off the "all-good" light.

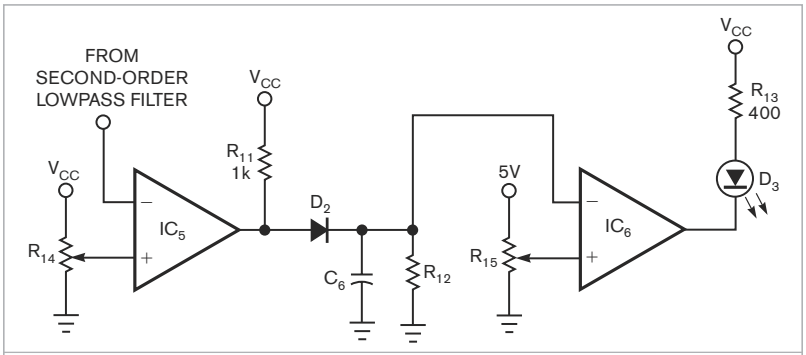

Figure 4 Comparator IC_6 monitors the voltage across C_6 and activates when the C_6 voltage exceeds a preset value from R_{15}.

In a real-world application, you should latch the output to ensure that the operator notices the alarm condition. The circuit could simply shut down the power to the conveyor belt, allowing immediate repairs and indicating where the fault occurred. The circuit uses just a few common components with large tolerances on their values. Its use of transistors, op amps, discrete linear and nonlinear components, oscillators, filters, demodulators, converters, and magnetically coupled circuits make it an excellent teaching resource. It even gives some insight into how modern proximity-card technologies, such as RFID (radio-frequency identification) work. EDN

The best **Rubidium Clock**

is also the most affordable — **PRS10...$895**

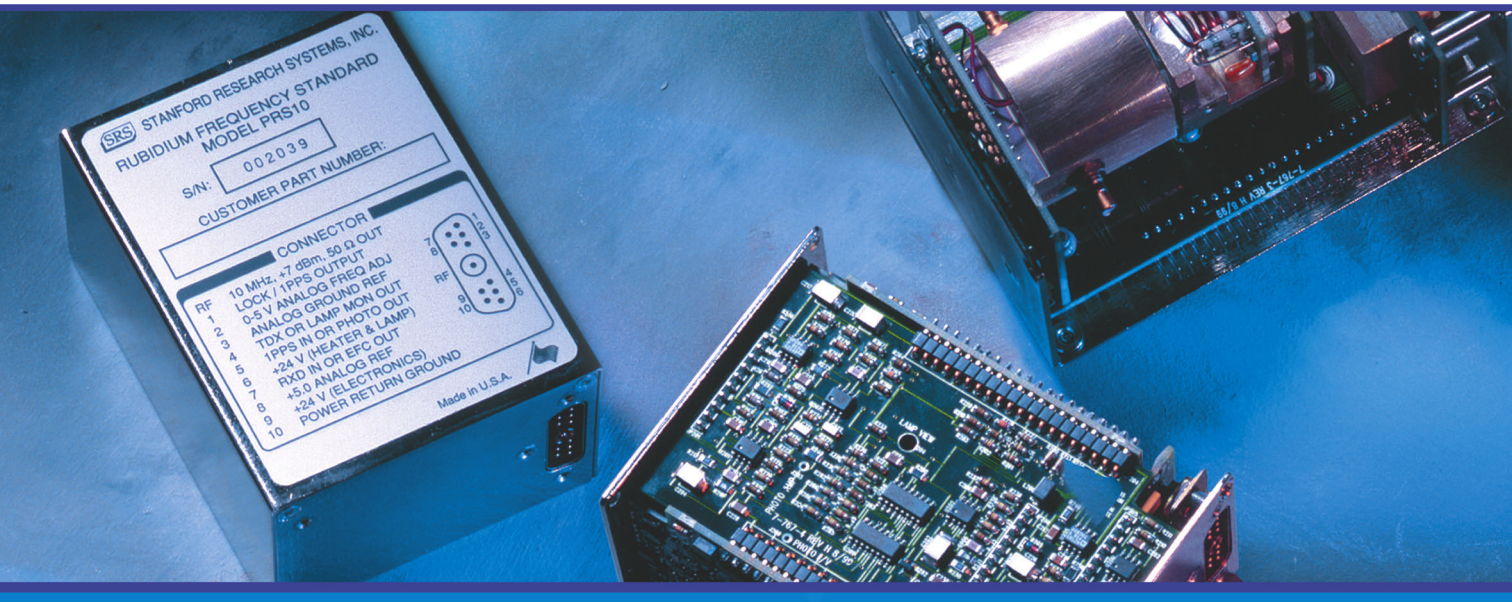

The 10 MHz PRS10 Rubidium Frequency Standard offers a unique combination of features, performance, reliability and low cost.

Its ultra-low phase noise output, low temperature coefficient and excellent aging rate place the PRS10 in a class by itself. The long-life rubidium lamp and robust design provide unequaled reliability.

A 1 pps input and output are provided for synchronizing to GPS, making the PRS10 ideal for critical timing applications.

- ***20 year aging <5 ppb***
- ***20 year lamp life***
- ***Ultra-low phase noise (–130 dBc/Hz at 10 Hz)***
- ***±1 × 10^{-10} temperature stability***
- ***1 pps input and output***
- ***RS-232 interface***

Qty	***Price***
1	$1495
500	$995
1000	***$895***

SRS ***Stanford Research Systems***

Phone: (408) 744-9040 · Fax: (408) 744-9049 · info@thinkSRS.com · **www.thinkSRS.com**

Allied

Your MRO solutions provider

ALLIED ELECTRONICS®

Personal Service

Personal expertise from a dedicated Allied Sales Rep

Availability of Product

Over 115,000 products on the shelf and ready to ship

Latest Order Time

Order up to 10pm ET for same day shipping or up to 8pm ET for next day delivery

| Test & Measurement | Interconnect | Enclosures | Power | Control | Passive & Active | Assembly | Tooling |

alliedelec.com 1.800.433.5700

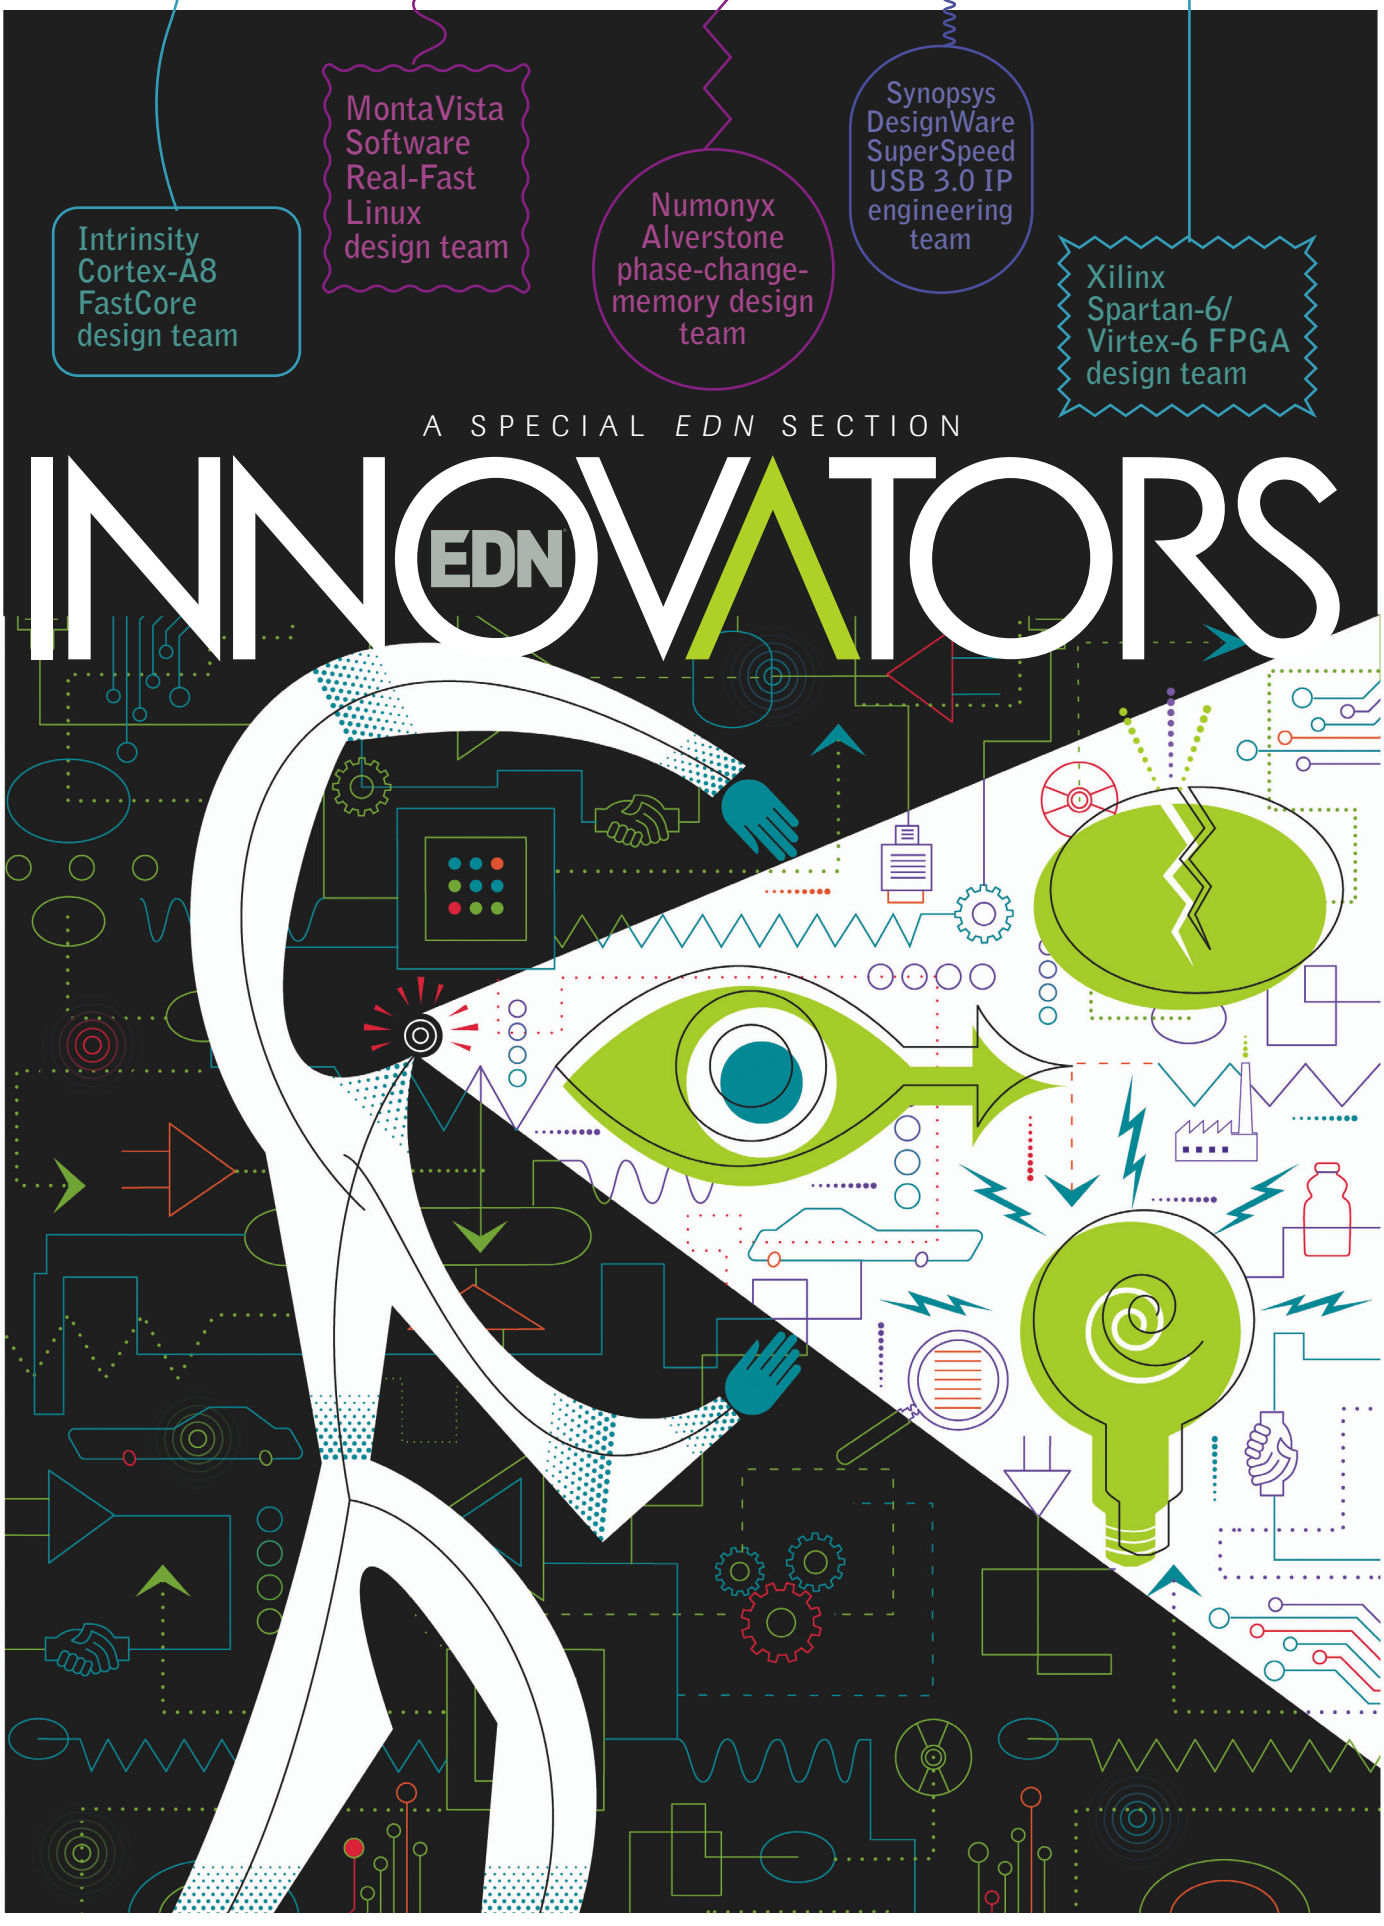
Intrinsity Cortex-A8 FastCore design team
MontaVista Software Real-Fast Linux design team
Numonyx Alverstone phase-change-memory design team
Synopsys DesignWare SuperSpeed USB 3.0 IP engineering team
Xilinx Spartan-6/ Virtex-6 FPGA design team
A SPECIAL EDN SECTION
INNOVATORS
EDN

BY RICK NELSON, EDITOR-IN-CHIEF

In a jobless recovery, innovators will have jobs

THE ECONOMY IS RECOVERING, OR SO THE ECONOMISTS TELL US. The employment picture remains grim, however, and high levels of joblessness may be a permanent condition, according to a recent item in *The Huffington Post* (**Reference 1**) by Martin Ford, a Silicon Valley entrepreneur, computer engineer, and author of *The Lights in the Tunnel: Automation, Accelerating Technology and the Economy of the Future*.

Ford begins with a look at an article in *Salon*, "Why *Dilbert* is doomed: The jobs of tomorrow are not what you'd expect" (**Reference 2**). That article predicts that the stable jobs of tomorrow will be those that cannot be automated–those that require creativity or proximity, with proximity including health care and education. Ford says, however, that neither health care nor education is immune to automation. Online education is becoming commonplace, and expert systems can perform many health-care tasks.

Ford describes an employment pyramid with a few skilled professionals and entrepreneurs at the top; the vast majority below the apex performs routine, repetitive tasks. Automation, he says, will consume the entire base of the job-skills pyramid. People in that base must move up or leave.

He doubts that most people will have the ability to move up. He adds that people with the creativity, talent, or personality traits to be successful writers, performers, and commissioned sales people exhibit a power-law income distribution, which leads to drastic income inequality.

Ford sees a future of job sharing and income supplementation for most people. He echoes Gregory Clark, a professor of economics at the University of California—Davis. "The economic problems of the future will not be about growth but about something more nettlesome: the ineluctable increase in the number of people with no marketable skills and technology's role not as the antidote to social conflict, but as its instigator," Clark has noted (**Reference 3**).

Where do engineers stand in this picture? Engineering is not immune to automation. I've written much over the past few months about the model-based design tools and hardware-in-the-loop tests that drastically reduce the amount of time-consuming low-level programming and hardware prototyping that engineers must do (**Reference 4**). To stay at the top of Ford's employment pyramid, engineers must exploit the new tools available to them to develop products that represent significant innovation and creativity. Marginal improvements won't cut it.

In this special section, we profile five engineering teams who have pushed to stay at the apex of the job-skills pyramid during 2009: the Intrinsity Cortex-A8 FastCore team, the MontaVista Software Real-Fast Linux team, the Numonyx Alverstone Phase-Change-Memory team, the Synopsys DesignWare SuperSpeed USB 3.0 IP team, and the Xilinx Virtex-6/Spartan 6 team.

EDN editors have selected these teams as finalists for Innovator of the Year in our annual Innovation Awards program. We encourage you to peruse their stories on the following pages and cast an online ballot for your favorite at www.edn.com/innovation20. More important, though, we hope their stories will serve as inspiration as you hone your innovation skills and remain at the top of Ford's employment pyramid throughout 2010 and beyond. EDN

REFERENCES

1 Ford, Martin, "A Jobless Recovery ... And A Jobless Future?" *The Huffington Post*, Feb 2, 2010, www.huffingtonpost.com/martin-ford/jobless-recoveryand-joble_b_445439.html.

2 Lind, Michael, "Why Dilbert is doomed," *Salon*, Nov 2, 2009, www.salon.com/news/opinion/feature/2009/11/02/healthcare_employment/index.html.

3 Nelson, Rick, "Robots, jobs, and war," *Test & Measurement World*, Aug 17, 2009, www.tmworld.com/blog/Taking_the_Measure/21001-Robots_jobs_and_war.php.

4 Nelson, Rick, "High-level software for embedded-system design: doing your job?" *EDN*, July 9, 2009, pg 20, www.edn.com/article/CA6668610.

Analog Devices: enabling the designs that make a difference in people's lives

A broad range of products are enabled by ADI's small, low power, low cost ICs.

RF and broadband technologies that connect even the most remote users. Whether transmitting critical medical information, reporting breaking news, or linking people with family and friends, wireless systems connect the world. And designers of these broadband systems have come to rely on Analog Devices to make a difference in their designs. Our converter, amplifier, RF, and power technologies enable the low power, high capacity, and cost-effective results demanded in the network of point-to-point communications systems, private mobile radios, and wireless infrastructure equipment. Our tradition of innovation continues, making wireless systems, and people's lives, better. Learn more at www.analog.com/makeadifference.

MAKEADIFFERENCE

www.analog.com

ANALOG DEVICES

Intrinsity Cortex-A8 FastCore team

BY ROBERT CRAVOTTA, TECHNICAL EDITOR

Domino logic boosts performance

INTRINSITY'S FASTCORE TECHNIQUES OFFER AN ALTERNATIVE TO HIGH-PERFORMANCE, LOW-POWER CORES.

Intrinsity's Fast14 NDL (1-of-N domino logic) can increase the performance of a low-cost, standard RTL (register-transfer-level) core by 40 to 50% and maintain cycle accuracy. The total design cost is approximately 90% less than the full-custom approach, and the design cycle takes less than a year.

Intrinsity's team of chip designers, using its Fast14 NDL, developed the 1-GHz, less-than-700-mW Cortex-A8 RTL FastCore, with cycle accuracy and Boolean equivalency to ARM's original Cortex-A8 golden specification. About one-fifth of the A8's functions are benefiting from the domino logic. Intrinsity's FastCore techniques offer an alternative approach to implementing high-performance, low-power cores.

The fastest processors, such as those from Intel and AMD, use fast dynamic- or domino-logic gates in a full-custom-design process. Resolving the timing and noise issues associated with dynamic logic in such designs requires large design teams, multiyear design cycles, and hundreds of millions of dollars.

Many SOC (system-on-chip) vendors license processor cores from ARM or MIPS. These cores use synthesized static-logic gates in an automated design process that provides moderate performance and quick design cycles with small teams of engineers. It is possible to increase the performance of these cores by implementing changes to the ISA (instruction-set architecture). However, this approach also requires substantial resources. ISA license fees are substantially higher than those for synthesized RTL cores. Optimizing the ISA is a multiyear process that typically requires large design teams and deep pockets. Changing the ISA destroys compatibility with the original core and requires the development of new test suites and software, which can add a year to the design cycle.

Synthesized static logic has more relaxed timing constraints than does domino logic. It allows as many as 15 to 20 gates per clock cycle, but each gate ex-

CONTINUED ON PAGE 53

The Intrinsity Cortex-A8 FastCore team includes Brent Chambers, team leader, and Pat Thomas, vice president of R&D.

MontaVista Software Real-Fast Linux team

BY PAUL RAKO, TECHNICAL EDITOR

Innovators get Linux to boot in 1 second

DOING REAL-TIME LINUX IS HARD ENOUGH. MAKING IT BOOT IN 1 SECOND TOOK REAL INNOVATION.

MontaVista Software has always been a leader in embedded-Linux commercialization. The company has developed Linux-development platforms since 1999, when founder Jim Ready pledged to bring "100% pure Linux" to the world under the GNU (GNU's not Unix) GPL (general public license). Since then, MontaVista has specialized in embedded and real-time Linux.

Its approach is not simply an RTOS (real-time operating system) that runs Linux as one of its tasks. The company has changed the Linux kernel to provide determinism and real-time performance in a real Linux operating system. Cavium Networks recently acquired the company, which just announced the release of Version 6 of its operating system.

In addition to designing real-time Linux, MontaVista has been working on the development of real-fast Linux, a Linux operating system that boots in less than 1 second. The team who worked on the project includes Alexander Kaliadin, Nikita Youshchenko, and Cedric Hombourger. Many on the team also worked on the MontaVista real-time Linux. "One of the first things we did years ago was to make the Linux scheduler pre-emptive and deterministic," says Hombourger. These fast-boot developments are not necessarily limited to real-time or an embedded Linux; however, they can get a conventional Linux distribution to boot in 1 second, as well.

"The methods we have developed are independent of whether [you use] a real-time kernel," notes Kaliadin. He says that the team first considered netbooks and mobile Linux fast-boot approaches. It could not adapt them, however, because they needed a fairly heavyweight Linux distribution, and the requirements for a 1-second boot are far more extreme than those of any netbook or notebook application. "We had to devise a set of new techniques in both the kernel and the boot loader," says Kaliadin.

A customer who needed his automotive-dashboard system to come up quickly prompted the development of a fast-booting Linux. Kaliadin remarks that the customer loved the MontaVista OS but had a fixed requirement to show data on the screen in less than a second. "Our first reaction was that it was impossible," says Kaliadin.

The 1-second-boot-time achievement came in three stages. The first stage was optimization in all the obvious places, including the boot loader. The team could eliminate some of

CONTINUED ON PAGE 53

The MontaVista team includes Alexander Kaliadin (left) and Cedric Hombourger (right).

Numonyx Alverstone Phase-Change-Memory team

BY BRIAN DIPERT, SENIOR TECHNICAL EDITOR

Phase-change memory enhances and supplements other memory technologies

PHASE-CHANGE MEMORY IS THE LATEST CANDIDATE IN THE RACE TO KNOCK FLASH MEMORY OFF ITS LEADERSHIP THRONE.

Innovative memory technologies appear with reasonable regularity, at least if you measure them by the frequency of press releases and technical-conference presentations. Invariably, however, most of them never make it out of the laboratory, or, if they do, they remain single-vendor-sourced niche products.

Innovative new memory technologies that achieve high-volume production, multiple-vendor sourcing, and consequent widespread industry adoption, on the other hand, are rarer. Three of today's dominant approaches, DRAM, SRAM, and EEPROM, all date from the 1970s and have since experienced only deliberate evolutionary advancement. Perhaps the most recent memory innovation that achieved pervasive popularity, in fact, was flash memory, which both Intel, with NOR, and Toshiba, with NAND, unveiled in the mid-1980s.

Flash memory's in-system rewritability represented a notable advancement beyond its ROM, PROM, and EPROM predecessors and was an attractive attribute in both code-execution and mass-data-storage applications. Yet flash memory is by no means a semiconductor panacea. Erasure—that is, changing a stored zero into a one—is slow and occurs only on a bulk block-by-block basis. Bit-by-bit one-to-zero writes aren't terribly fast, either. Sequential-cell-rewrite operations, or cycles, of any sort inevitably lead to transistor-oxide degradation and eventually to unrecoverable access failure. As semiconductor processes steadily scale downward in various transistor dimensions, oxide thickness and other size-related factors threaten to drive flash memory into a lithographic brick wall.

It's no surprise, therefore, that several upstart contenders have attempted to knock flash memory off its leadership throne. Neither FRAM (ferromagnetic-random-access memory) nor MRAM (magnetic RAM), however, has to date shown that it has the sustainable moxie to credibly counterbalance the formidable momentum of the entrenched flash-memory competitor. PCM (phase-change memory) is the latest candidate and, judging from the number of semiconductor suppliers pursuing its development, has a fighting chance of succeeding.

PCM employs the same kinds of chalcogenide-glass materials and leverages the same temperature-driven transitions between the material's crystalline and amorphous states as writable optical discs, such as CDs, DVDs, and Blu-ray discs.

CONTINUED ON PAGE 54

The Numonyx team includes (from left) Duane Mills, Rich Fackenthal, Greg Atwood, Roberto Gastaldi, Roberto Bez, and Fabio Pellizzer. Not shown: Ferdinando Bedeschi

Synopsys DesignWare SuperSpeed USB 3.0 IP team

BY RICK NELSON, EDITOR-IN-CHIEF

One-stop shop provides easy-to-use USB 3.0 IP

WITH DESIGNWARE, DEVELOPERS CAN QUICKLY AND SUCCESSFULLY IMPLEMENT USB 3.0 IN THEIR ASIC DESIGNS.

When looking to support its customers' adoption of the new emerging USB (Universal Serial Bus) standard, the Synopsys DesignWare SuperSpeed USB 3.0 IP (intellectual-property) engineering team wanted to ensure interoperability and ease of use by providing a single-vendor approach—from instantiation of the on-chip IP through certification and test.

The approach optimizes both area and performance. For example, it requires only one external reference-clock input and one external calibration resistor, which the high-speed and SuperSpeed operational modes share. In addition, the IP is compatible with a range of technology nodes—from the most advanced processes to older ones more suitable for low-volume or cost-sensitive applications.

Gervais Fong, product-marketing manager, cites some challenges and motivations for the team's approach. He notes that high-speed USB 2.0 runs at 480 Mbps, whereas SuperSpeed USB 3.0 runs at 5 Gbps. "The USB 3.0 specification includes the requirement to support all four USB speeds—low speed, full speed, high speed, and now the new SuperSpeed," he says. "We find from an IP perspective that providing a complete solution to customers instead of giving them individual blocks makes it easier for them to integrate and provides faster time to market at lower cost. At the system level, we provide a complete USB solution consisting of the PHY [physical layer], the controller, the verification IP, and the reference drivers. We try to be the one-stop shop to make it easy for the design engineers to quickly and successfully implement USB 3.0 in their ASIC designs."

Robert Lefferts, R&D director at the company, says that Synopsys also focuses on applying the resources and extra time and effort necessary to take the IP through certification and compliance testing. Combining certain functions expedites this process. "If we broke apart the SuperSpeed functionality from the standard USB 2 functionality, we couldn't follow through," he says. "Attaching all the functionality together into one solution was a fairly big undertaking, but it helped get a complete solution that we can verify. One of the first things we looked at was crosstalk through the connector. How much margin would we need to take into account to handle that additional crosstalk of combining

CONTINUED ON PAGE 55

The DesignWare SuperSpeed USB 3.0 IP engineering team includes (from left) Tri Nguyen, Subramaniam Aravindhan, Eric Huang, Minh Pham, John Youn, and Saleem Mohammad.

Xilinx Virtex-6/Spartan-6 FPGA team

BY RON WILSON, EXECUTIVE EDITOR

FPGA design team pushes for a global effort

ONE GLOBAL DESIGN TEAM, TWO PRODUCT LINES, AND A UNIQUE ORGANIZATION MOVE FORWARD FPGA TECHNOLOGY.

The tangible achievement of Xilinx's Virtex-6/Spartan-6 FPGA design team speaks for itself. The team produced two complete families of FPGAs reaching from parts costing a few dollars to a 2.5 billion-transistor behemoth and used two processes at two foundries to build them. Both families have a common tool set and a common pool of IP (intellectual property), and the team produced both families in a global 24-month design effort.

There is a back story, however. The effort introduced a fundamental cultural change to Xilinx's engineering community.

The heart of this change was the drive for consistency across two divergent product lines. "Customers said they could see our organizational boundaries in our products," says Victor Peng, senior vice president for programmable platforms. "We believed from the beginning that the way to get a consistent product line was to create a foundation of IP and reuse it throughout the families—not just for the design of the FPGAs but for the tools and the customer-IP libraries, as well."

The effort was not simply a matter of removing inconsistencies from the customer experience, however. Consistency was necessary because of the scale of the undertaking. For the first time, Xilinx would simultaneously be creating two FPGA families. Further, FPGAs increasingly must adopt architectural and implementation features for specific application clusters. So the team had to produce a range of features, as well. "We had to expand our market without proportionately expanding our engineering effort," explains Steve Douglass, vice president of product development.

The only feasible approach was to build a base of silicon IP in both target processes—UMC [United Microelectronics Corp] 40 nm and Samsung 45 nm—and use that IP to construct the FPGAs. "Just linear scaling of effort was not enough," Douglass says. "We had to reorganize the work."

"The foundation was the Virtex-5 fabric," says Suresh Menon, vice president of product development. "It was created to be scalable, and in fact we were able to reuse the architecture and much of the connectivity—even the serial I/O—in this Virtex-6/Spartan-6 effort."

"We made different embodiments of the same concepts for different product points," Douglass says. "For instance, the PCIe [Peripheral Component Interconnect Express] and DRAM interfaces across the product line are different implementations of the same architecture."

Success in leveraging the common IP platform depended on the organization of the

CONTINUED ON PAGE 55

The Xilinx team includes Steve Douglass and Suresh Menon, vice presidents of product development. Not shown: Victor Peng, senior vice president for programmable platforms

INTRINSITY CONTINUED FROM PAGE 48

ecutes less logic than an NDL gate. An NDL gate can execute two to five times more logic per clock cycle than a synthesized static-logic gate by taking advantage of multiple out-of-phase clocks so that the system completes more work per system clock cycle.

However, domino logic consumes more power and silicon area than static-logic gates, so you must use domino logic judiciously. Intrinsity's NDL uses less power than traditional domino logic by representing values zero through three with four wires so that only one of the transistors driving the wires is powered at a time. The system represents zero as 0001, one as 0010, two as 0100, and three as 1000.

The Fast14 technology originally targeted multigigahertz applications allowing only one NDL gate per clock phase because that was the only way to guarantee the timing requirements. Originally, developers manually applied NDL gates to large sections of a design—full blocks or units—because of the complexity of interfacing the domino logic to the synthesized static logic in the chip. The new tools automatically size the NDL transistors to achieve the required timing and drive strengths.

Intrinsity developed a simplified interface between static and dynamic logic that enables more efficient transitions between dynamic and synthesized static logic and supports an automated placement flow. The simpler static/dynamic interface allows the use of NDL in smaller subcomponents, ameliorating power and area issues and giving designers the freedom to use small amounts of it in more areas of the chip. The ability to use smaller NDL subcomponents enables the FastCore design team to selectively use NDL only in the critical paths of the core, rather than in large blocks.

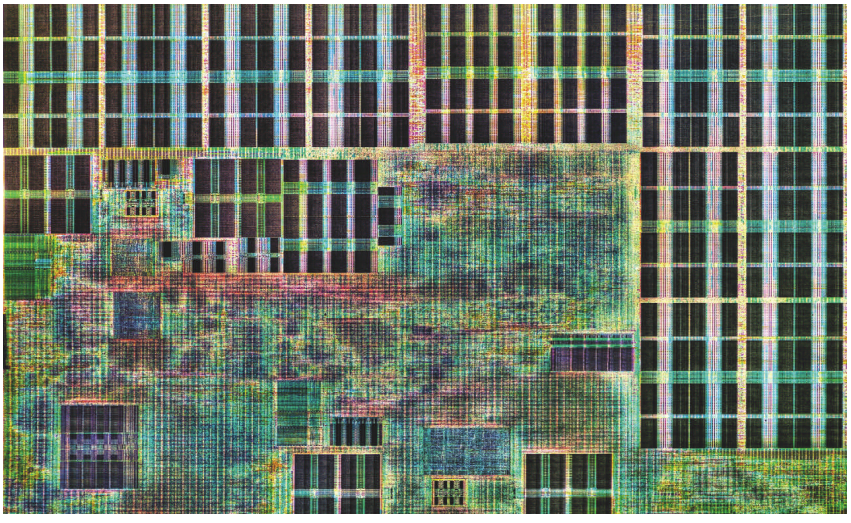

Intrinsity's engineers developed a static/dynamic interface that allows the use of NDL in smaller subcomponents, ameliorating power and area issues and giving designers the freedom to use small amounts of it in more areas of the chip.

Intrinsity engineers used their circuit-design expertise to build a semiautomated timing-analysis flow employing the PrimeTime timing-analysis tool. The tool looks at all the paths in the design to identify those with the most common features: either common start and endpoints or paths that traverse through a common set of logic. By speeding the logic areas that affect the largest numbers of critical paths, the design team could speed up a large number of critical paths with a minimal amount of NDL, minimizing any die-size or power-consumption penalty.

Intrinsity engineers determined that the lower clock-frequency targets for the ARM Cortex families of cores allow the use of four to 12 gates per clock phase, and they still meet the timing requirements. Increasing the number of NDL gates in each clock phase contributes to achieving the 1-GHz clock rate of the Cortex-A8 FastCore in a 45-nm LP (low-power) process—about 40% more speed than you can achieve in a 45-nm LP process using only synthesized static logic. EDN

MONTAVISTA CONTINUED FROM PAGE 49

the boot time in the boot loader because the hardware is the same for everything the system boots. The team also omitted many drivers the OS didn't need and minimized the OS configuration. "Even [Linux founder] Linus Torvold admits that Linux is getting pretty bloated," says Kaliadin. This first stage got the team members down to a 7-second boot time.

The second stage required an intimate knowledge of the hardware. The Linux boot loader is a serial process. The team's epiphany came when the developers realized that they could use DMA (direct-memory-access) methods to parallel many tasks in the boot process.

The DMA agents can move many boot tasks between flash memory and the processor memory and can accomplish this task in the background without processor overhead. "These days, CPUs have a pretty large cache memory, so they are capable of doing all these things in parallel," says Kaliadin. Using DMA and the processor cache saves 3 more seconds, which further reduces the optimized boot time to 4 seconds from 7.

The next logical place to reduce the boot was in the user's application, but customers fix and determine that variable. The MontaVista team then looked at the loading of the customer's applications. The developers could use the RAM disk that has been available in the Linux kernel since Revision 2.4, but Linux still cached that memory, and that process slowed things down. Since the development of the 2.6 kernel, Linux has supported loading the file system into this RAM disk. "We ditched the whole buffer scheme and just loaded the customer's application into the Linux page-cache memory," says Kaliadin.

The second part of this innovation was the developers' realization that they

didn't have to load the customer's entire application, just the parts that the initial application required to start up. That realization allowed the boot time to near 1 second. "We made some big advances but then had to find 100 milliseconds here and there," he says.

Even though they worked with a customer's automotive dashboard, the developers' work is applicable to any application that needs to have an embedded Linux, complete with a file system, boot in less than a second. Although this application is processor-specific, MontaVista says that it can help you apply the technology to any 32- or 64-bit processor. "We designed the fast-boot process to be architecture-independent," notes Hombourger. He comments that the only hardware-specific features involve the DMA commands that speed the boot loader, but it is a relatively small task to port that function to different hardware.

For those who wonder why a processor doing billions of instructions per second needs to take minutes to boot up, wonder no more.

For those who wonder why a processor doing billions of instructions per second needs to take minutes to boot up, wonder no more; MontaVista has shown that you can boot up a complete modern operating system in less than a second. EDN

NUMONYX CONTINUED FROM PAGE 50

Whereas optical storage uses a laser to precisely inject heat into the chalcogenide-glass material, PCM employs cell-specific applications of voltage and current. Also, whereas developers manipulate the disc material's optical properties to store zero and one data-sequence combinations, PCM alters and, during subsequent reads, senses variances in chalcogenide glass's electrical resistivity.

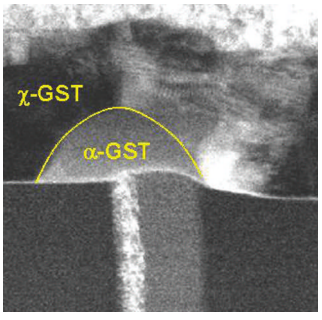

PCM employs the same chalcogenide-glass materials that writable optical discs, such as CDs, DVDs, and Blu-ray discs, use.

PCM is fully bit-alterable, with the added potential for higher write speeds than flash memory, and its wear-out mechanisms are more subtle than those of flash memory, leading to extended cycling specifications. Stored PCM data also exhibits notable stability characteristics across voltage, temperature, and chronological variances. And the technology is highly resistive to the data corruption that radiation bombardment causes.

One serious participant in the PCM market is Numonyx, a company crafted from the NOR-flash-memory development teams at Intel and STMicroelectronics. Numonyx, which Micron Technology recently announced plans to acquire, has to date achieved several notable PCM accomplishments. After many years' worth of internal development and industry investment, then-Intel and STMicroelectronics in October 2006 jointly revealed a 90-nm-fabricated, 128-Mbit PCM device, which Numonyx has now advanced to low-volume-production status. Roughly 18 months later in February 2008 at the International Solid-State Circuits Conference, the two companies unveiled an MLC (multilevel-cell) variant that stores 2 bits of information in each PCM cell. Last October, Numonyx showcased the ability to vertically stack multiple PCM arrays on a single-die sliver of silicon, and the company announced in December its intentions to productize a 1-Gbit PCM device it fabricates on an advanced 45-nm-lithography process.

Will the looming PCM steamroller flatten flash memory and, if so, when? Company Chief Technical Officer Ed Dollar cautions against "irrational exuberance," quoting a phrase that former Federal Reserve Board Chairman Alan Greenspan notably uttered in 1996, favoring a more cautious crawl-before-walk-before-run approach.

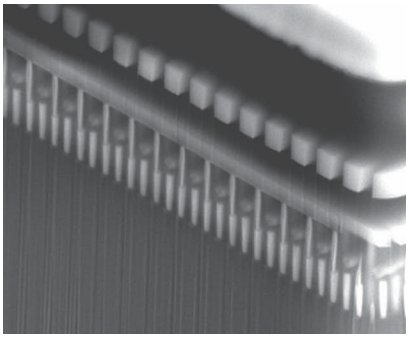
Numonyx vertically stacks multiple PCM arrays on a single-die sliver of silicon.

"Like most new technologies, PCM offers benefits to those who know where and when to apply it," he says. "To understand where PCM fits today and to appreciate its potential value, we need to evaluate its relative cost, reliability, and performance compared to incumbent technologies such as SLC [single-level-cell] and MLC NAND-flash [memory], as well as system solutions, including hard-disk drives and SSDs (solid-state drives). As we examine where PCM fits in the memory landscape, it is important to view it not as a replacement technology for other types of memory but as a supplemental technology capable of providing key benefits when the system requirements are right.

"Whether it is supplementing RAM or supplementing NAND flash, PCM can be used in just the right amounts to deliver improvements in reliability and performance in high-end applications, such as enterprise computing and e-commerce. If solid-state drives based on NAND-flash technology are currently at an early stage of adoption in many applications, PCM can be found at an even earlier stage of the adoption curve, where the most innovative developers work. Part of my job is to spend a lot of time talking with innovators to understand how they think and how we can help meet their needs."

Dollar wraps up on a high note, though. "PCM can deliver significant performance gains," he says. "Once PCM technology is utilized in high-end applications, its popularity will grow." EDN

SYNOPSYS CONTINUED FROM PAGE 51

the 480-Mbps and 5-Gbps functions?" If the team had just broken the problem in half, he says, the customer would have had to worry about the answer to that question later.

Subramaniam Aravindhan, R&D manager at the company, comments on the USB 3.0 controller: "If you have separate 2.0 and 3.0 controllers, you need to have two subsystems—two bus interfaces—and that increases integration and verification costs. One major concern is software drivers. If you have a 2.0 programming model that is different from the 3.0 [model], you need to write different software drivers, and you need to validate the two drivers. Having a combined software model for 2.0 and 3.0 created a large initial effort for us, but, from the long-term point of view, it's a clean, complete, single solution."

Lefferts cites one challenge of combining functions into one device. "One of the logistical problems for us was that the project involved design teams at seven locations." Synopsys had to orchestrate those teams, which were in the United States, Canada, India, and Armenia, to deliver the final product.

Working out the logistics had a beneficial result, explains Lefferts. Two years ago, IP teams didn't tend to overlap. For example, USB 2.0 and high-speed SERDES (serializer/deserializer) teams were relatively isolated. "As a result of trying to go after a full solution for USB 3.0, all those barriers are down," he says.

You can synchronize music in a matter of seconds with a USB 3.0 media player versus minutes with a USB 2.0 media player.

"If customers have an issue, it doesn't matter whether it's on the high-speed 480-Mbps or SuperSpeed 5-Gbps side. There are people they can call now that understand both."

Lefferts also addresses test issues. "The other feature we've implemented in the USB 3.0 solution came out of some of the work we've done on some of our other standards-based IP," he says. "A lot of our customers for PCIe [Peripheral Component Interconnect Express] started out working with PCI at much lower frequencies. They had never worked with anything over a gigabit per second, and, when they went to PCIe Generation 1 at 2.5 Gbps, there was obviously a learning curve from a signal-integrity perspective. So we build our IP with diagnostic and debug capabilities to help them isolate problems. In this case, we are working with customers that were working at 480 Mbps and now are working at 5 Gbps. They might not even own a real-time scope that's even capable of looking at the signal. What we've done on our SERDES in general is to provide functions to be able to help separate where the problems are—whether it's a signal integrity-issue or it's a functional issue. We try to make our IP not only easy to use and integrate but also easy to debug so you can isolate and work on problems very quickly."

Lefferts points out that manufacturers tend to implement other high-speed serial-I/O standards in advanced process technologies, whereas they may deploy USB 3.0, which provides a high-speed interface to some of the lowest-cost electronics devices available, on older process nodes. Consequently, the team had to address many technology nodes in a short period of time. Synopsys now has customers who have first silicon back and who will be taping out for volume production this year.

Eric Huang, product marketing manager for DesignWare USB IP at Synopsys, and Terry Moore chief executive officer of MCCI, teamed up at the Consumer Electronics Show in January to demonstrate how you can synchronize music in a matter of seconds with a USB 3.0 media player versus minutes with a USB 2.0 media player. The demonstration consisted of the Synopsys DesignWare SuperSpeed USB digital controller and MCCI SuperSpeed USB software running on an FPGA hardware platform. You can view a video of the demo at www.synopsys.com/IP/pages/Videos.aspx.EDN

XILINX CONTINUED FROM PAGE 52

design team. At the hardware level, the vice presidents created 10 Centers of Excellence around the world, each center responsible for implementing a category of blocks—programmable fabric, high-speed transceivers, or embedded memory, for example. These centers fed into product-integration teams, and three common silicon design flows bound them together, heading off any problems with tool incompatibility among centers.

But the design flows had an additional responsibility. The microarchitecture and chip implementation were moving ahead at Xilinx in parallel with the process developments at Samsung and UMC. There were no mature models,

The Virtex-6/Spartan-6 resulted from a single effort that bound together the 250 software, silicon, IP, and system engineers around the world.

rule decks, or PDKs (process-design kits). Further, Samsung was essentially new to the foundry business and was just building the infrastructure to support an external chip design team.

"We had to develop our own test vehicles to characterize these processes for our designs, and we evolved what I'd call look-ahead circuit design," says Menon. "You can't design for the process as the test chips show it to you. You have to design to the process as you believe it will have evolved by production time."

In practice, that requirement means not only an intimate relationship with the highest levels of process engineering in the foundries, but also the experience

to understand where a process variable is going to end up, independent of the foundry's stated objectives. It is also necessary to build flexibility into the circuits to adapt to possibly unanticipated changes in process parameters during the silicon design flow. "We call that [flexibility] a shock absorber," Douglass says.

The needs of FPGAs—not ordinarily on a foundry's radar screen—further complicated these goals. "We had to be very good at understanding leakage variations and single-event-upset characteristics," Menon says. "These are different issues for a large FPGA than they are for, say, a consumer SOC [system on chip]. In our world, things that are second-order effects for other people become very important."

The development of the design software and IP that customers would use to program the FPGAs ran in parallel with the silicon design effort. "From architecture definition on, the hardware and tool developments were side by side," Douglass says. With devices of this complexity, it was possible neither to design the hardware and toss it over the wall to tool developers nor to create a generic architecture on the assumption that it would meet increasingly application-specific customer needs.

> The team is looking for critical challenges that will drive the next round of IP, software, and microarchitecture development.

Xilinx's approach to this problem was Solution Teams. Each team focused on a cluster of applications, interviewing customers, hiring applications experts, and developing representative use cases. The Solution Teams became central. "Implementation really started with the use models," Douglass explains. "From the use models, the Solution Teams defined the critical functional blocks their applications would need. Then they built reference designs and ran them through the tools."

"These weren't hypothetical exercises," Peng emphasizes. "They are real, usable reference designs." Running the reference designs through the tools produced feedback for the silicon and the tool teams, directing the entire project toward the goals customers would have for their own designs with the FPGAs.

A single effort emerged that bound together the 250 software, silicon, IP, and system engineers around the world. The hardware Centers of Excellence built a platform of reusable silicon IP and integrated it into specific chips. The Tools Teams created the programming flow. The Solutions Teams validated the flow and chip models against target reference designs. All of it converged in a successful product launch.

There is also a longer-term benefit. "We know how to do this again," Peng says. The team is diving into target applications, building new use models, and looking for critical challenges that will help to drive the next round of IP, software, and microarchitecture development. EDN

EDITED BY SUZANNE DEFFREE

supplychain

LINKING DESIGN AND RESOURCES

Niche distributors specialize for market edge

Whether it's supplier support or the newest technology on the block, niche distributors have developed specialized line cards and services to attract and hold customers to prosper in a component world dominated by broad-line global distributors such as Arrow, Avnet, and Digi-Key.

Nu Horizons Electronics Corp (www.nuhorizons.com), a distributor of advanced semiconductor, display, illumination, power, and system products, defines itself as having an edge in its niche, in both technology and supply services. "We're a true extension of our supplier's sales force," says Kent Smith, vice president of global sales and marketing at Nu Horizons. "On the technology side of it, we're more knowledgeable."

> We're there when they want to do 1000 prototypes.

The company's goal is to master a small line of products well enough to fully support design engineers and to provide global supply services. "From a design standpoint, we offer our customer full support," says Smith. "For fulfillment, we can support our customers seamlessly in a global way. If they have an EMS [electronic-manufacturing-services] partner, that's easy for us."

Mouser Electronics (www.mouser.com), a catalog distributor that focuses on supporting design engineers through development, prototype, and early production runs, has also found a niche. The company provides small to midsized quantities of a wide range of the latest components. In focusing on the design and early production stage for its customers' production cycle, Mouser forgoes interest in high-volume distribution. "We deal with small production quantities," says Kevin Hess, marketing director at the company. "We're there when they want to do 1000 prototypes. A design engineer or small-product engineer can buy small amounts and broken packs from us. But when it comes to volume, we leave that to the volume distributors."

For more on this story, go to www.edn.com/article/CA6717240.

—by Rob Spiegel

OUTLOOK

LEDs HELP BOLSTER LCD-TV SALES

Global revenue from shipments of large LCD panels for use in televisions should rise to $49.2 billion in 2010, up 40% from $35.2 billion in 2009, according to iSuppli Corp (www.iSuppli.com). The projected revenue increase for this year follows a 5.2% decline from 2008 to 2009.

According to Sweta Dash, senior director of LCD research at iSuppli, rising demand, the shift to larger displays, the accelerating sales of higher-value panels, and increased manufacturing efficiency will drive revenue growth and profitability in the global LCD-TV-panel market in 2010. "These higher-value panels sport features including LED backlighting, full high-definition resolution, and faster refresh rates of 120 and 240 Hz," says Dash. "Furthermore, television brands in 2010 are planning to focus on 3-D, which will lead to faster adoption of higher-frequency panels." Increased production efficiency will make larger TVs more affordable for consumers, Dash adds, and also notes that suppliers will cut costs by using fewer components in TV panels.

ISuppli expects manufacturers this year to ship 25.4 million LCD-TV panels featuring LED backlighting—a fivefold increase from 4.8 million units in 2009.**—SD**

GREEN UPDATE

UNITED KINGDOM REMINDS INDUSTRY OF BATTERY REGULATIONS

The United Kingdom's Environment Agency is urging small businesses that make, import, or sell batteries and battery-operated equipment to ensure that those products comply with new regulations to reduce the environmental impact of batteries. Any business that in 2009 placed batteries for the first time on the market in the United Kingdom should have registered as a battery producer. The first deadline for submitting battery data for 2009 was Jan 31, 2010.

The actions are new responsibilities under the United Kingdom's Waste Batteries and Accumulators Regulations that came into force in 2009. The regulations set out how the United Kingdom collects, treats, and recycles all types of waste batteries and rechargeable batteries. The Environment Agency estimates that approximately 700 million batteries, which can contain substances that are harmful to the environment, such as cadmium, enter UK landfills each year. Meanwhile, the government agency estimates that just 3% of the 30,000 tons of portable batteries sold in the UK market annually are currently recycled. "Ensuring that portable batteries are correctly disposed of and their component parts, such as their metal casings, are reused is good news for the environment," says Bob Mead, the Environment Agency's batteries-project manager.

For more information, go to www.environment-agency.gov.uk/batteries.**—SD**

productroundup

POWER SOURCES

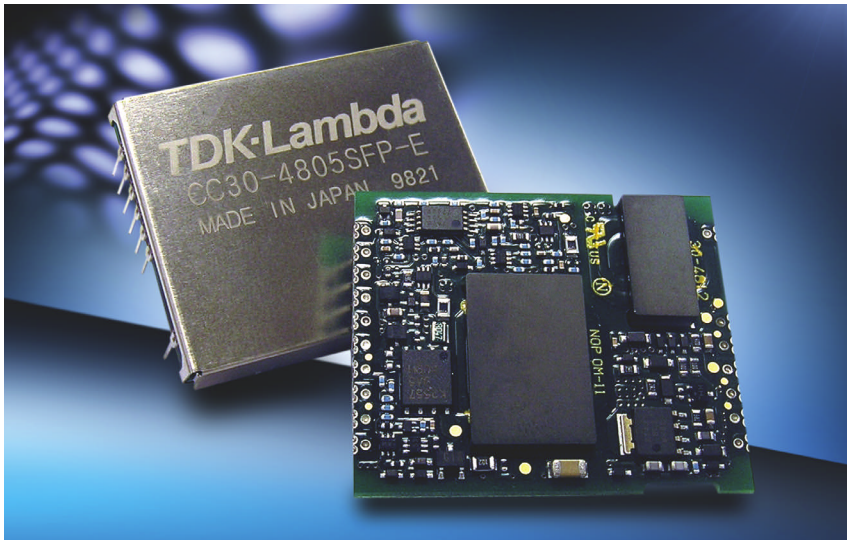

Isolated 15 and 30W converters enable parallel capability

Providing nominal 24 and 48V-dc inputs, the CC-P-E dc/dc-converter family targets use in telecom, datacom, distributed power, battery-operated devices, FPGAs, and process-control applications. Available with 3.3, 5, 12, or 15V-dc output voltages, the devices include 15 or 30W output-power ratings. The fully regulated converters have 92% typical efficiency, and designers can parallel converters for higher-power applications. The devices offer a nominal 24V-dc input operating from 18 to 36V dc or a 48V dc input operating from 36 to 76V dc. The 15W converters measure 29.6×38.4×6.8 mm, and the 30W units measure 33.5×38.4×8.3 mm. The devices come in through-hole or SMT packages with or without shielded metal casings. Operating over a −40 to +85°C temperature range, devices in the CC-P-E converter family cost $28 each (100).

TKD-Lambda, www.us.tdk-lambda.com/lp

Synchronous buck regulator provides 15 digitally programmable outputs

The SC195 500-mA synchronous buck step-down regulator provides a flexible output for low-voltage point-of-load regulation and mobile-device applications. Operating from 2.9 to 5.5V inputs, the regulator provides 15 pin-programmable output voltages from 0.8 to 3.3V, covering all the typical core and I/O voltage rails with one device. Claiming 94% efficiency, the device provides fast transient response, 3.5-MHz PWM operation with an automatic power-saving feature at light loads, 38-μA quiescent current, and a 0.1-μA typical shutdown mode. Using four digitally programmable pins to set the output voltage eliminates the need for external feedback components or individual fixed output-voltage regulators for different loads, reducing space and improving reliability. Available in a 1.5×1.5×0.6-mm MLPQ package, the SC195 500-mA step-down regulator costs 69 cents (3000).

Semtech Corp, www.semtech.com

Single-output power supplies have fan-cooling options

Suiting medical and industrial applications, the SHP650 and the MHP650 650W ac/dc power supplies provide a single output and high power density. Two fan-cooled formats enable the option of mounting the fan internally on the top or externally on the end of the unit. The load-dependent integral fan's speed minimizes noise. Using the U channel format enables designers to use their own airflow, requiring 5.5m/sec forced airflow. The device includes a 12V dc, 6W fan supply when using the U-channel format. Additional features include an 85 to 264V input range, operation at full power from 90V ac, and a −20 to +70°C temperature range with no derating below 50°C. The SHP650 and the MHP650 ac/dc power supplies cost $179 and $281, respectively.

XP Power, www.xppower.com

Switch-mode power modules suit use in POL regulators

Available in surface-mount packaging, the ISL8204M and ISL8206M switch-mode power supplies include a PWM controller, power MOSFETs, a power inductor, and associated discrete components. Pin-compatible with the vendor's 10A ISL8201M module, the ISL8204 and the ISL8206M power supplies have 4 and 6A ratings, respectively. The devices enable the construction of a high-performance POL regulator using input and output capacitors, and a resistor for programming the output voltage. Claiming 95% efficiency, the modules provide a 1 to 20V input-voltage range. A resistor sets output voltage from 0.6 to 6V with ±1% accuracy. Available in lead QFN-15 packages measuring 15×15×3.5 mm, the ISL8204M and the ISL8206M switch-mode power modules cost $8.12 and $11.34 (1000), respectively.

Intersil Corp, www.intersil.com

EDN ADVERTISER INDEX

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

EDN product mart

This advertising is for new and current products.

MARVIN MOSS, PhD

No good deed goes unpunished

DANIEL VASCONCELLOS

In 1979, I was on assignment for a contract to update the radar simulator for the C-5A flight simulator at Altus Air Force Base. I was working for Lockheed Aircraft (now Lockheed Martin Corp). Work was progressing reasonably well until it was time to integrate the radar system with the main flight-simulator part of the system. While doing so, I noticed a considerable amount of drift in the flight portion, an integral part of what sends data to the radar part. I thus had to investigate why there was so much error.

The flight system used the 24-bit-word SEL840, at the time a state-of-the-art computer from SEL (System Engineering Labs). The big computer used 64 kbytes of core memory and had discrete component boards with RTL (resistor-transistor-logic) devices (circa 1975). The entire system worked from one sample-and-hold circuit. All of the inputs were either digital, such as switches, or analog, such as control column positions. The outputs were digital-to-analog conversions. The sample-and-hold circuits would sample an analog input and then convert the held analog voltage to a digital word with an ADC.

I used an oscilloscope to observe the output of the sample-and-hold circuit to see why it was not holding the sampled voltage steady while the ADC was doing its job. The scope showed that the output had drift and unsteadiness when the sampled input's analog voltage was as steady as it could be. Careful observation revealed something else interesting. The waveform that was controlling the sampling appeared to be upside down, meaning that when it should have been sampling and holding, it was actually holding and then sampling.

The entire process happened on one plug-in card. We had replaced the card with a newer version because, after plugging and unplugging the older ones many times over the course of many years of routine maintenance, we had caused damage to them. We carefully checked the original schematics and the card for any discrepancies. Looking at the waveforms showed that the digital sampling pulse into the card fed through two simple inverter stages to clean up the signal and make sure that the waveform was clean and of the proper magnitude. Because the waveform appeared on the new card to be inverted from what it should have been, I checked the original card and saw two inverters. What I found on the new card was unbelievable! There, as clear as a bell, were three inverters cascaded together, obviously giving the wrong waveform, yielding the upside-down pulse.

The obvious change was to eliminate one of the inverters. I cut one trace and jumpered the extra inverter out of the string of inverters. Absolutely perfect sample-and-hold action was the result, and all of my needed signals were as stable as a rock. The Air Force technicians were just overjoyed at the results of my fix. I thought that they were going to award me a medal.

All was well until I got back to my home base in Marietta, GA, when my boss called me into his office and read me the riot act. The Air Force higher-ups had found out about my "unauthorized fix," and I received a severe reprimand for going outside the approved method of making recommendations and changes to the simulator. I learned that following proper procedures was more important than having a simulator that was working completely up to spec. I never made that mistake again. EDN

Since 1998, Marvin Moss, PhD, has been retired from Lockheed Martin after 32 years there on the job. He now enjoys playing at home with all of his electronic toys.

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

Data acquisition just got a lot easier.

More ways to control. More ways to connect.

The new Agilent 34972A Data Acquisition Switch Unit takes our best-selling Agilent 34970A to the next level. For starters, you get convenient built-in LAN and USB connectivity. Plus, you can control your data acquisition remotely via Web interface. And transfer logged data to your PC with a simple flash drive. No more expensive adapters and connectors. That's easy. That's Agilent.

NEW 34972A	34970A
USB and LAN	GPIB and RS232
Graphical web interface Benchlink data logger software SCPI programming	Benchlink data logger software SCPI programming
3-slot LXI unit with built-in 6 ½ digit DMM	
~~$1,845~~ $1,597*	$1,597*

Hurry! For a limited time, get the 34972A at the 34970A price.
www.agilent.com/find/Agilent34972A

*Prices are in USD and are subject to change.
See participating distributors for details.

Agilent Technologies

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

RS ELECTRONICS
888-852-7234
www.rselectronics.com