Reed Business Information (RBI™)

EDN

VOICE OF THE ENGINEER

JUNE 11

Issue 11/2009

www.edn.com

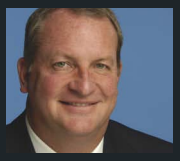

TEE UP YOUR MULTIPROCESSING OPTIONS

Page 24

USB 3.0: A SIMPLE IDEA FULL OF CHALLENGES

Page 16

RF ENGINEERS AUTOMATE TESTS

Page 20

TROUBLESHOOTING A TRANSACTION-LEVEL MODEL

Page 31

DESIGNING PORTABILITY INTO SILICON IP

Page 37

innovation

APEX™ PRECISION POWER

Product Innovation from Cirrus Logic

64-PIN QFP, PACKAGE STYLE HQ
JEDEC MO-188

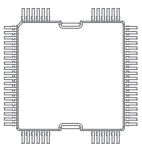

(actual footprint)

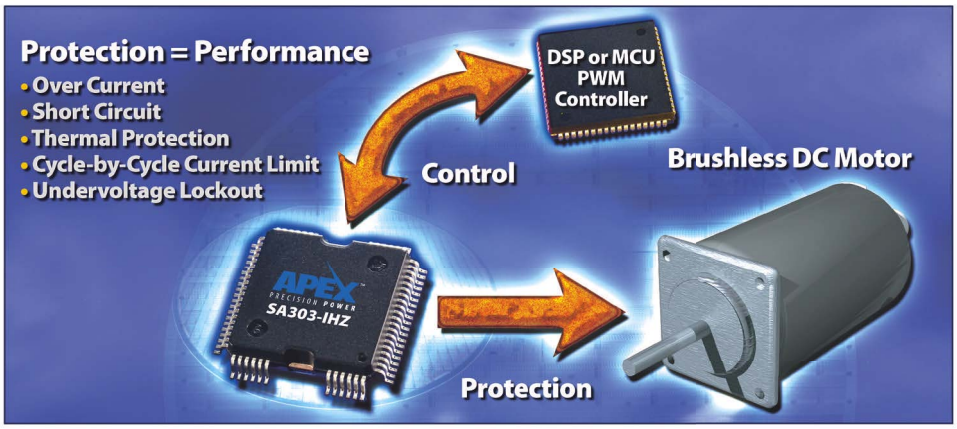

3 A, Three-Phase Motor Driver IC Uses Protection Features to Deliver Performance Gains

The SA303-IHZ is a pulse width modulation (PWM) motor driver IC designed to protect brushless motors from operational hazards such as over current, under voltage and over temperature. This newest addition to the Apex Precision Power™ product family can generate performance gains in the overall motor control circuitry by eliminating motor shutdown and re-starts. For example, should an over-temperature or short circuit occur, the SA303-IHZ will not shutdown the outputs but instead allows the system to continue running while it signals the external controller to take corrective action. In most cases, the processor can correct the condition by utilizing software to remedy the external interrupt.

APPLICATIONS

- Motor Drives – Industrial Controls
  - Factory Automation
  - Robotics
- Motor Drives – Office Equipment
  - Copiers, Fax Machines
  - Vending Machines
- Motor Drives – Aerospace, Military
  - Positioning Control
  - Aircraft Seating

Model	Motor Interface	Output Current	Supply Voltage Operation	Production Volume Pricing 10K Pieces USD*
SA303-IHZ	Brushless DC Motor	3 A continuous 10 A Peak	10 V to 60 V Single Supply	$5.86
SA306-IHZ	Brushless DC Motor	5 A continuous 17 A PEAK	< 9 V to 60 V Single Supply	$9.90
SA306A-FHZ	Brushless DC Motor	8 A continuous 17 A PEAK	< 9 V to 60 V Single Supply	$12.85
SA53-IHZ	Brush DC Motor	3 A continuous 10 A Peak	10 V to 60 V Single Supply	$4.79
SA57-IHZ	Brush DC Motor	5 A continuous 17 A PEAK	< 9 V to 60 V Single Supply	$7.15
SA57A-FHZ	Brush DC Motor	8 A continuous 17 A PEAK	< 9 V to 60 V Single Supply	$9.05

* per unit pricing for production estimating only; actual per unit cost through distribution may vary

REQUEST A FREE SAMPLE!
Request a free sample unit of the SA303-IHZ
By visiting www.cirrus.com/sa303edn

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

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/sa303edn

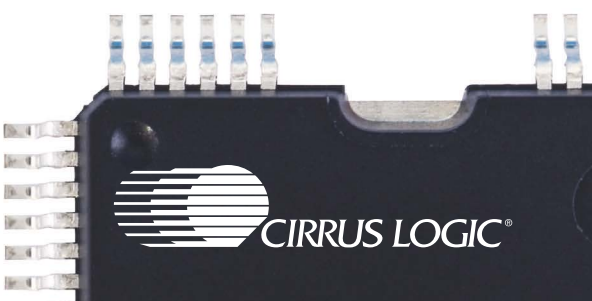

Digi-Key® CORPORATION

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

Over 400,000 products in stock from more than 400 supplier partners.*

The industry's broadest product selection available for immediate delivery

www.digikey.com

1.800.344.4539

An additional 1,000,000+ components can be sourced at digikey.com

*Digi-Key is an authorized distributor for all supplier partners. New products added daily. © 2009 Digi-Key Corporation, 701 Brooks Ave. South, Thief River Falls, MN 56701, USA

“How can I tell if a power supply is reliable?”

There’s an indicator on the front.

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

Take a brief quiz and save up to 10%.
And enter to win an iPod touch.
www.agilent.com/find/poweronquiz

Agilent Authorized Distributor

METRICTEST™
6,000 instruments. One source.

866-436-0887 www.metrictest.com/agilent

Limited time offer. See Web site for full promotion details.
Discount offer applies to 144 power supplies.

Agilent Technologies

EDN contents

6.11.09

Tee up your multiprocessing options

USB 3.0: a simple idea full of challenges

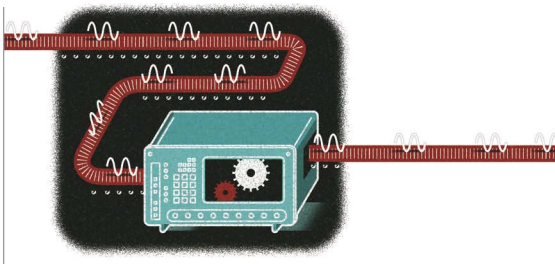

RF engineers automate tests

Troubleshooting a transaction-level model

Designing portability into silicon IP

pulse

DESIGNIDEAS

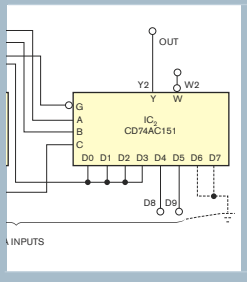

Get your $49 starter kit today.

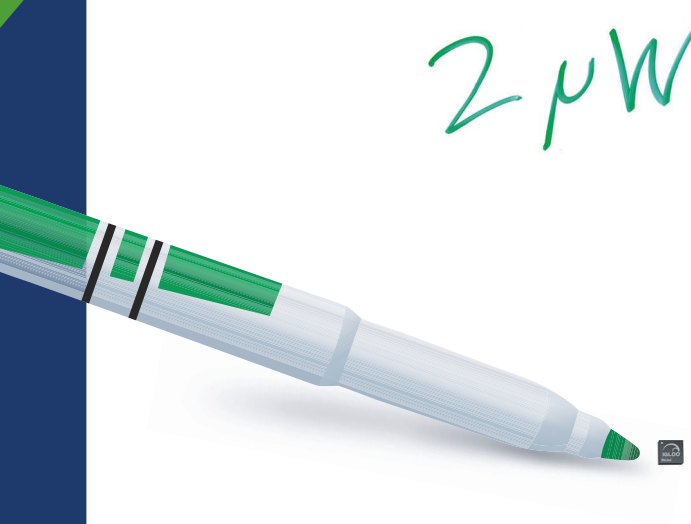

nanoPower

nanoSize

nanoPrice

Actel™ FPGAs help you do more with less.

Actel IGLOO® nano and ProASIC®3 nano low-power FPGAs bring reprogrammability to high-volume applications at the lowest cost. Available in 3x3 mm packages, with prices starting at only 49¢ (and 50 variants under $1.00), Actel gives you the industry's widest selection of low-power, low- cost FPGAs. With power consumption as low as 2 µW, IGLOO nano and ProASIC3 nano FPGAs are perfect for your volume-driven, power-sensitive applications. Now, more than ever, power matters.

Watch webcasts, read white papers, and get a starter kit at actel.com/nanoFPGAs

Actel® POWER MATTERS

contents 6.11.09

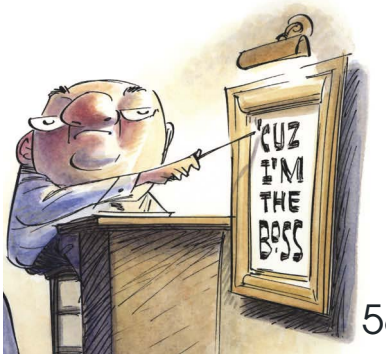

DEPARTMENTS & COLUMNS

EDN online contents www.edn.com

ONLINE ONLY

Check out these Web-exclusive articles:

For 10-Gbps interconnects, the RJ-45 once again will dominate
The industry will adopt 10-GBase-T as the primary 10-Gbps interconnect for data-center switches and servers for many of the same reasons that RJ-45 and twisted-pair cabling dominate at gigabit speeds.
→www.edn.com/article/CA6656523

Electronic-ballast circuits enhance resonant-mode power supplies
This article explains the basic functions of a typical electronic-ballast circuit, highlights the similarities between the applications, and describes a typical resonant-mode power supply using a standard electronic-ballast-control IC.
→www.edn.com/article/CA6658506

SEE NAKED ICs
The IC Insider, an exclusive EDN.com series contributed by Chipworks, one of the leading IC reverse-engineering houses, provides detailed functional and circuit analysis of significant new ICs. The IC Insider shows you how the latest ICs really work, and why—insight you'll never get from the data sheet.
→www.edn.com/icinsider

HOT TOPICS
Check out *EDN*'s Hot Topics pages. Hot Topics pages deliver continuously updated, subject-specific links from not only *EDN* but also the entire electronics-industry Web.
→www.edn.com/hottopics

EDN® (ISSN#0012-7515), (GST#123397457) is published biweekly, 24 times per year, by Reed Business Information, 8878 Barrons Blvd, Highlands Ranch, CO 80129-2345. Reed Business Information, a division of Reed Elsevier Inc, is located at 360 Park Avenue South, New York, NY 10010. Tad Smith, Chief Executive Officer; Mark Finkelstein, President, Boston Division. 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 2009 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 54, Number 11 (Printed in USA).

EDN

PRESIDENT, BOSTON DIVISION, REED BUSINESS INFORMATION
Mark Finkelstein, mark.finkelstein@reedbusiness.com; 1-781-734-8431

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

VICE PRESIDENT, EDITORIAL DIRECTOR
Karen Field
1-781-734-8188;
kfield@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
Newsletter/Editorial Coordinator
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
The Quadrant, Sutton, Surrey SM2 5AS
+44 118 935 1650;
fax: +44 118 935 1670;
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@idg-rbi.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

BY RON WILSON, EXECUTIVE EDITOR

A modest proposal for IP

The US patent system needs fixing to the point that there are debates over just whose interests we should fix first. But no one seems any longer interested in the original idea: that patents protect the inventor—not a corporation—and that, by protecting the inventor, they encourage publication, accelerating the pace of innovation. So here's a suggestion: We should prohibit the assignment, the sale, or any other transfer of patent ownership, except for inheritance.

This change would not fix everything. But it would resolve a number of problems that corrode the system, and correct one problem that threatens to blow up the whole show.

First, consider the inventors. Today, granting patents to the people who did the inventing is fiction. Unless the inventor is self-employed and not under contract, an employer gets the patent. This appropriation is poisonous in a number of ways.

Engineers find ways to work around the system. They avoid documentation. If an idea starts to pan out, the engineer quits, launches a new company, and then files patents. That move is risky, but it is a better deal than getting your name on a bronze plaque in somebody else's lobby. Most of the time, the industry and the company that funded the work get nothing, and the inventor gets two years of 16-hour days at reduced salary. And that scenario is what happens if everyone avoids litigation.

Now consider employers. The system vacuums the best R&D employees from organizations just as they are about to do their most valuable work. Any way you look at it, that move destroys equity. Further, owning patents tends to force companies to focus on generating revenue from their patents—and away from serving their customers.

Let inventors own their patents and allow them to transfer limited rights under the well-developed body of licensing law.

Suppose patents were not transferable, however. Inventors, instead of assigning patents to their employers, would license them in case-by-case negotiations based on the employment agreement. With licensing instead of assignment, the goals of the inventor and employer would align: Both want to maximize the profits from the patent, and the best way to do that is usually cooperation. If the company wants the patent, it will help defend it. If the company isn't interested, the inventor can find other licensees in the free market.

Now think about venture investors. Today, when a venture company dies, patents are often its best surviving assets. But separated from the engineers who did the work, the patents have only a fraction of their real worth. If the patents remain with the engineers, the asset would be not just a possible future cash flow. The asset would comprise the patents; all the other IP (intellectual property) associated with applying them, including collected data, prototypes, and so forth; and the actual people who did the work. This asset is a far more valuable—and tangible—one that a buyer could exploit at once.

Finally, there is an even more serious problem that the industry must address. Some investors now see patents as financial instruments: contracts that entitle the owner to an—admittedly uncertain—future cash flow. Investment banks can buy and pool patents and then create financial derivatives based on these pools. Financial technicians can model those future cash flows as stochastic processes, applying the math used for collateralization of subprime mortgages. We all know the next chapter in that story. The new demand for patents undermines the whole concept of IP in the United States.

So let inventors own their patents and allow them to transfer limited rights under the well-developed body of licensing law. It's fair to engineers. It's best for corporations and investors. And it could prevent a catastrophe that today threatens the entire system. **EDN**

Contact me at ronald.wilson@reedbusiness.com.

EDITED BY FRAN GRANVILLE

pulse

INNOVATIONS & INNOVATORS

Low-power audio codec incorporates Class D-loudspeaker amplifier

National Semiconductor recently announced the low-power LM49352 audio codec (compressor/decompressor) for handheld consumer electronics. The part includes an audio codec with ground-referenced headphone amplifiers, Class D-loudspeaker outputs, and an audio DSP. The loudspeaker delivers 93% efficiency at 970 mW from a 4.2V supply. The headphone amplifier offers 18-mW MP3 playback to extend battery life. Applications include smartphones, full-featured phones, portable gaming devices, and portable GPS (global-positioning-system) devices.

The device integrates a charge pump to increase power to the loudspeakers. It also includes a DAC, which achieves 103-dB SNR (signal-to-noise ratio), and an ADC, which provides 97-dB SNR. The part can play back audio with 96-kHz sampling rates and can record audio at 48 kHz. It also incorporates dual-stereo, five-band parametric equalizers; ALC (average-level control); and limiter and compressor functions on both the DAC- and the ADC-signal paths. The codec also features differential-microphone inputs to minimize noise, crosstalk, and RF susceptibility. The ADC has a wind-noise filter that can operate in voice- or high-fidelity-recording modes.

The low-power LM49352 audio codec for handheld consumer electronics has Class D-amplifier outputs.

The LM49352 comes in a 36-bump micro SMD package, sells for $7.55 (1000), and is available for sampling, along with evaluation boards.**—by Paul Rako**

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

FEEDBACK LOOP

"Thanks ... from an old guy just doing FPGAs for the first time (after decades of microcontrollers). I'm sure you just saved me a week somewhere."

—Engineer and *EDN* reader Ernie Murphy, in *EDN*'s Feedback Loop, at www.edn.com/article/CA6648806. Add your comments.

Two-terminal current source supplies 200 mA

Linear Technology Corp's new LT3092 two-terminal current source can supply 0.5 to 200 mA of current at a 1% initial accuracy. The device operates at 1.5 to 40V and exhibits less-than-1-mV typical load regulation and less-than-0.001%/V typical line regulation. The device needs no bypass capacitance on the input or the output in most ranges of operation, even when you use it at 0.5V. The output is short-circuit-protected and can withstand reverse voltage, and a thermal-shutdown circuit prevents overheating. You can configure the devices in parallel for greater current, or you can stack them in series for greater voltage compliance.

The LT3092 comes in 3×3-mm, eight-pin DFN; eight-pin SOT-23; and three-pin SOT-223 packages. The LT3092MP version operates in the −55 to +125°C junction-temperature range, and the LT3092E and LT3092I versions operate in the −40 to +125°C range. Prices range from $1.65 to $4.73 (1000). For more on this product, go to www.edn.com/090611pa.**—by Paul Rako**

▷**Linear Technology Corp**, www.linear.com.

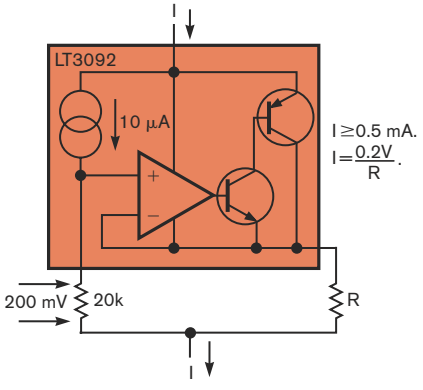

The two-terminal Linear Tech LT3092 current source works at 1.5 to 40V and can output 0.5 to 200 mA.

Agilent, ETS-Lindgren address A-GPS test

Agilent Technologies Inc has introduced GS-9000 A-GPS (assisted-global-positioning-system) test systems employing the company's 8960 wireless-communications-test set's new A-GPS-assistance-data-messaging-test capabilities. Meanwhile, ETS-Lindgren also announced an integrated test capability employing its AMS-8000 series of antenna-measurement systems and A-GPS test equipment from Agilent and other instrument vendors.

A-GPS, an enhanced position-location method, uses positioning-assistance data it obtains from base stations. With assistance data, a mobile device can within seconds determine and report back its exact location to a network. In contrast, unassisted-GPS techniques require minutes to accomplish this task. As service providers move to implement A-GPS in their networks and as vendors introduce A-GPS-capable devices, these devices require testing to demonstrate that A-GPS operation poses no interference to cellular services.

Agilent's GS-9000 A-GPS design-verification-test systems include both hardware and software for testing a device's A-GPS capabilities in a conducted environment. The hardware includes an 8960 test set and an E4438C vector-signal generator to emulate GPS satellites. The GS-9000 enables A-GPS validation, TIS (total-isotropic-sensitivity) testing, and A-GPS preconformance testing for mobile devices; the system also provides fading and SUPL (secure-user-plane-location) support. A GS-9000 Lite version omits SUPL and fading capabilities.

> A-GPS uses positioning-assistance data to determine and report back its exact location.

Within the GS-9000, the 8960 test set acts as a base station with an A-GPS message pipe for sending A-GPS-assistance-data messaging to and from the wireless device or phone under test. The phone acknowledges receipt of the data and performs location calculations, returning its location information to the 8960 test set through the appropriate UMTS (universal mobile-telecommunications system) or CDMA (code-division-multiple-access) 2000 technology-dependent protocol messaging. The system supports TIS testing in accordance with Version 3.0 of CTIA's (Cellular Telecommunications Industry Association's) "Test Plan for Mobile Station Over the Air Performance" by antenna-measurement-system providers, such as ETS-Lindgren.

For its part, ETS-Lindgren announced an expansion of the company's AMS-8000 series of antenna-measurement systems, which include a fully anechoic RF-test chamber with DUT (device-under-test) positioning equipment, antennas, instrumentation, and test-automation software. Version 1.07 of its EMQuest EMQ-100 antenna-measurement software adds support for A-GPS and other RSS (received-signal-strength)-based measurement methods. The ETS-Lindgren system supports the Agilent A-GPS test equipment and Spirent Communications' (www.spirent.com) ULTS (UMTS-location-test system).

—by Rick Nelson

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

▷**ETS-Lindgren**, www.ets-lindgren.com.

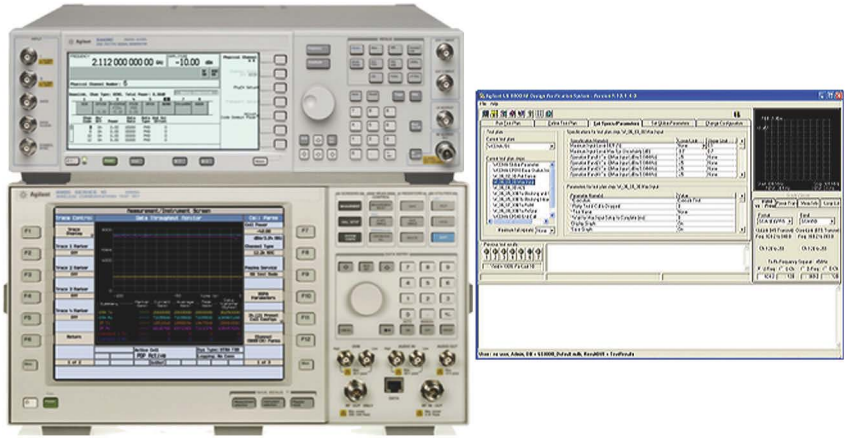

Agilent's basic mobile-device A-GPS test setup includes the Agilent 8960 wireless-communications-test set, the E4438C ESG vector-signal generator, and PC-control software.

DILBERT By Scott Adams

Kits support design of field-oriented-control motor drives

Infineon has expanded its range of application kits that support design of energy-efficient motor drives employing the company's 8- and 16-bit microcontrollers. The focus of the kits is ease of use, according to Jürgen Hoika, senior product-marketing manager of the company's industrial- and multimarket-microcontroller group. The kits include comprehensive tools and scalable reference designs.

The company based the new kits on members of the 8-bit XC800 and 16-bit XE166 families; they allow you to implement techniques such as FOC (field-oriented control) and power-factor control to obtain quiet and efficient designs. You need specific architectures to execute these techniques, Hoika says, but you can use an 8-bit device—when it has dedicated additional processing with a vector-computation block. A 16-bit controller, with a suitable architecture, runs two motors with both FOC and power-factor control.

> These kits are not just reference designs but ready-to-use, complete applications.

Infineon based the dual-motor-drive kit on the XC878 and XE166, and it supports FOC and power-factor control with a scalable architecture from 8 to 16 bits and a power board driving as much as 8A at 110 to 230V. The FOC-motor-drive kit supports design of FOC for the XC878 and XE166, with—in this case—a power board that drives as much as 7.5A at 23 to 56V. Code generation for the FOC aspect of your design is automatic. The 12V BLDC (brushless-dc)-drive kit uses the XC866 to drive motors with as much as 20A at 9 to 18V. All the kits contain a tool chain, including a compiler and a debugger; you can design with a library of preconfigured hardware and software blocks; a USB (Universal Serial Bus)-to-JTAG (Joint Test Action Group) and CAN (controller-area-network) bridge allows digital isolated real-time monitoring.

All of these kits work with Dave Drive, Infineon's autocode generator for motor drives. They are not just reference designs but ready-to-use, complete applications. For example, the dual-motor-drive kit contains two drive cards: One employs the XC878 with its vector computer, twin PWM (pulse-width-modulation) units, and an ADC with 1.5-µsec sample time; the other employs the XE164, which hosts a MAC (multiply/accumulate) unit, three PWM units, and two independent fast ADCs.

This device operates two motors and a PFC (power-factor-correction) block in parallel. The power board provides an inverter with 900 to 1800W, a second inverter with 100 to 200W, and a boost converter for power-factor control. In addition, you can run your application with control of induction motors for quick evaluation.

The 12V-BLDC kit with the XC866 costs €249 (approximately $330), the FOC-motor-drive-application kit sells for €399 (approximately $530), and the dual-motor-drive-application kit sells for €499 (approximately $665).

—by Graham Prophet

▷**Infineon**, www.infineon.com.

The dual-motor-drive-application kit supports FOC and power-factor control with a scalable architecture from 8 to 16 bits and a power board driving as much as 8A at 110 to 230V.

06.11.09

STANDARDIZATION BOOSTS SOFTWARE REUSE IN MICRO-CONTROLLERS

ARM licensees build ARM's CMSIS (Cortex-microcontroller-software-interface standard), a vendor-independent hardware-abstraction layer for processors, with Cortex-M processor IP (intellectual property). ARM intends to define a scalable interface standard across all Cortex-M series-processor vendors, enabling easier code sharing and reuse. Consistent and simple software interfaces to the processor–for silicon vendors and middleware providers–will lead to an easier learning curve for novice microcontroller developers and reduce product-development times.

In ARM's scenario, using 8- or 16-bit parts has no price benefit, and the automatic choice is the 32-bit single-chip microcontroller with ample memory, ensuring well-structured software layers. Most microcontroller projects reuse software components–but only in-house and only on a single device or device family. The CMSIS approach includes the creation of standard drivers for well-defined components that operate free of conflicts with middleware and operating systems from multiple vendors.

For more on this topic, go to www.edn.com/090611pb.

—by Graham Prophet

▶**ARM, www.arm.com.**

Op amp has offset calibration to 200 µV

Microchip Technology's new MCP651/2/5 operational amplifier targets use in microphone preamplifiers, optical-detector circuitry, digital scales, industrial instrumentation, H-bridge drivers, bar-code scanners, transmission-line drivers, and medical equipment for patient-monitoring and ultrasound functions.

The device has a maximum supply voltage of 5.5V and a minimum supply voltage of 2.5V. Input-noise voltage density is 7nV/$\sqrt{\text{Hz}}$ with a flicker-noise corner of 30 kHz. The device has a 50-MHz-gain-bandwidth product and can output 95 mA of current. The op amp has rail-to-rail outputs and an input structure that allows input-common-mode voltages 300 mV below the negative-supply voltage. It requires 1.3V of common-mode head room below the positive-supply voltage. The output-slew rate is 30V/µsec, and the supply current is 6 mA. The typical input-bias current is 6 pA, and the maximum is 5000 pA at 125°C.

The device comes in a 10-pin DFN package with an extra pin for the calibration function. When you pull the pin low, the part enters a calibration routine using an onboard DAC and memory and delay counters. The pin initiates an offset-voltage correction within 200 msec. Another version comes in a standard eight-pin-package pinout without calibration pins, so it performs the correction only on power-on and within 4 msec.

The device has an operating-temperature range of −40 to +125°C. Samples are available now, and the vendor is accepting volume orders. The single-amp MCP651, dual eight-pin MCP652, and dual 10-pin MCP655 cost $1.21, $1.49, and $1.58 (10,000), respectively.**—by Paul Rako**

▷**Microchip Technology**, www.microchip.com.

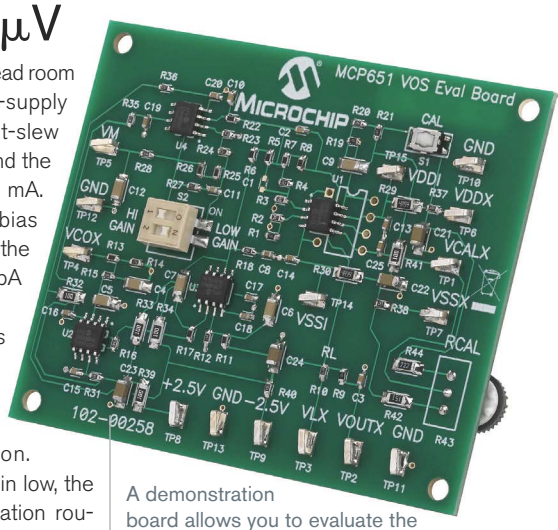

A demonstration board allows you to evaluate the capabilities of the MCP652 op amp.

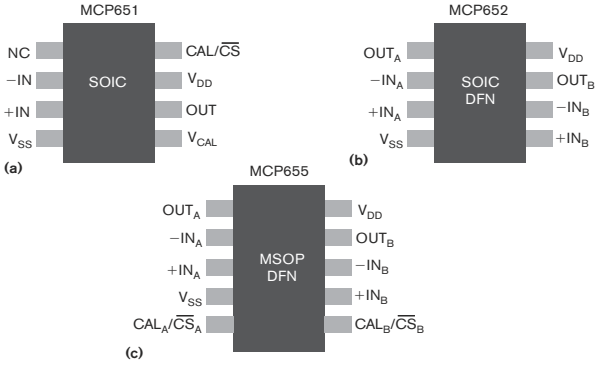

The MCP65x offset-calibrating-op-amp series is available in single eight-pin (a), dual eight-pin (b), and dual 10-pin versions (c).

EDNBLOG

BRIAN'S BRAIN

CES 2009: the price of falling prices

Given the multiyear time span that elapses between when a new microprocessor architecture is proposed and when the first products based on that architecture appear, I honestly doubt that Intel (www.intel.com) anticipated the economic crisis now gripping the globe. So, while I can't give the company credit for prescient forecasting, I'm still struck by the lucky timing of its Atom unveiling.

The netbook and net-top phenomenon represents the first time I can recall since the initial unveiling of the notebook PC when CPU horsepower has taken a substantive step backward in the face of some other more compelling factor. In the desktop-to-laptop transition, consumers consciously chose to lessen their CPU demands in exchange for factors such as lower system weight and longer battery life. And in the modern notebook-to-netbook and desktop-to-net-top era, system price reductions are the fundamental motivation.

Intel and partner Microsoft (www.microsoft.com) are doing their damndest to keep the notebook- and netbook-PC segments distinct, in the hope of preserving the size and profitability of the notebook segment fed by Intel's higher-end CPU families. Take the Atom N270 CPU found in the MSI Wind U100 netbook that I tested in my study, for example. Intel's own Web site advocates pairing that particular Atom processor flavor with the Mobile Intel 945GSE Express chip set with 82801GBM I/O controller hub (ICH7M). That core-logic chip set has an archaic integrated graphics core, and it also supports only up to 2 Gbytes of system memory. OEMs also report that Intel insists that they use the N270 only with systems containing 10-in.-diagonal maximum LCD sizes.

The ascendance of the CPU-deficient netbook and net-top symbolizes a substantial change in the rules of the game, thereby providing perhaps the first (and last?) tangible opportunity for vendors to redefine the fundamental parameters of the computer.

The netbook phenomenon has opened consumers' eyes to the reality that they've long been buying far more CPU muscle than most of them require. Therefore, I suspect that the era of robust sales of competitively isolated, high-performance Intel microprocessors is largely drawing to a close.**—by Brian Dipert**

▶www.edn.com/briansbrain.

▶For the full post, go to www.edn.com/090611pd.

Texas Instruments announces dual 14-bit, 250M-sample/sec ADC

Texas Instruments has introduced the dual 14-bit, 250M-sample/sec ADS62P49 ADC, which achieves 73-dBFS (decibels relative to full-scale) SNR (signal-to-noise ratio) and 85-dBc (decibels referenced to the carrier) SFDR (spurious-free dynamic range) at an input frequency of 60 MHz. Chuck Sanna, product-marketing engineer for high-speed products at TI, calls the device the industry's fastest dual 14-bit ADC, adding that it targets communications applications, defense imaging systems, and wideband and portable test-and-measurement equipment.

Power consumption is only 625 mW per channel, reducing thermal footprint and increasing system efficiency in high-density, multiantenna-base-station receivers and software-defined radios. Programmable gain and other user-selectable settings maximize design flexibility. Internal gain adjustment of as much as 6 dB in 1-dB steps allows customers to optimize SNR, SFDR, and input swing based on their applications' needs. For instance, designers can maximize SNR to enhance linearization effectiveness in DPD (digital-predistortion) applications, or they can increase SFDR and reduce input drive to improve small-signal analysis in defense and radio-receiver applications.

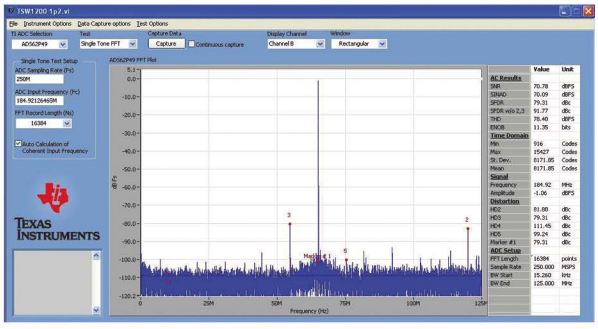

At an input frequency of 185 MHz, the ADS62P49 achieves a 79-dBc SFDR, which extends to 91 dBc if you exclude the second- and third-order harmonics.

The device forms a part of a complete signal chain; complementary products include the DAC5682Z and DAC5688 DACs, the THS4509 amplifier, the TRF3703 and TRF3761 RF components, the GC5325 DPD-transmitter processors, the CDCE72010 clock-distribution circuit, and the TMS320C6727B DSP. A comprehensive evaluation-tool suite speeds time to market. Customers can leverage TI's TSW1200 digital-capture tool for rapid evaluation of LVDS (low-voltage-differential-signaling)-output ADCs with as much as 16-bit resolution and 500M-sample/sec rates to enable prototyping of complex systems and to further speed development time.

"Communications, defense, and test-design engineers are constantly challenged to create signal- and data-acquisition receivers with increasingly wide signal bandwidths," says Art George, senior vice president of TI's high-performance-analog-business unit. "The ADS62P49 delivers high-performance [and] compact, power-efficient designs and enables rapid deployment of systems, software-defined radios, and spectrum analyzers."

The ADS62P49 is available in a 9×9-mm QFN package for $144.75 (1000). The ADS62P49EVM evaluation module has onboard sockets for a voltage-controlled crystal oscillator, a crystal filter, and TI's newest high-performance clock-distribution and jitter-cleaning chip, the CDCE72010, which allows the module to function as a system-level evaluation kit and an ADC evaluation board.

—by Rick Nelson

▷**Texas Instruments**, www.ti.com/ads62p49-pr.

PASSIVE SCOPE PROBES OFFER BANDWIDTHS TO 1.5 GHz FOR HIGH-SPEED-DIGITAL APPLICATIONS

Agilent Technologies has introduced the N2870A family of miniature passive oscilloscope probes and accessories with bandwidths from dc to 1.5 GHz. The probes' 2.5-mm-diameter heads cover only one-fourth the surface area of earlier probes' 5-mm-diameter heads. The devices also offer low input capacitance and a variety of fine-pitch probe-tip accessories.

The probes and accessories allow you to probe densely populated IC components and surface-mount devices. The sharp, spring-loaded probe tip helps you keep the probe from slipping off the device under test. Insulating IC caps center the small probe tip on the IC lead and keep it from shorting to adjacent leads. Optional probe-tip accessories provide specialized capabilities for demanding applications.

The probes are available with attenuation ratios of 1-, 10-, 20-, and 100-to-1 and probe bandwidths of dc to 35, 200, 350, and 500 MHz and 1.5 GHz. The company's Infiniium 9000 series oscilloscopes come with one 500-MHz, 10-to-1-attenuation-ratio N2873A passive probe per channel. These probes have an input impedance whose resistive component is 500Ω. Prices for N2870A series probes begin at $300.

—by Dan Strassberg

▶**Agilent Technologies, www.agilent.com/find/N2870A.**

The N2870A family of passive scope probes and accessories targets frequencies as high as 1.5 GHz. The 2.5-mm-diameter head and the large number of accessories suit the probes for use on tiny components and crowded printed-circuit boards.

06.11.09

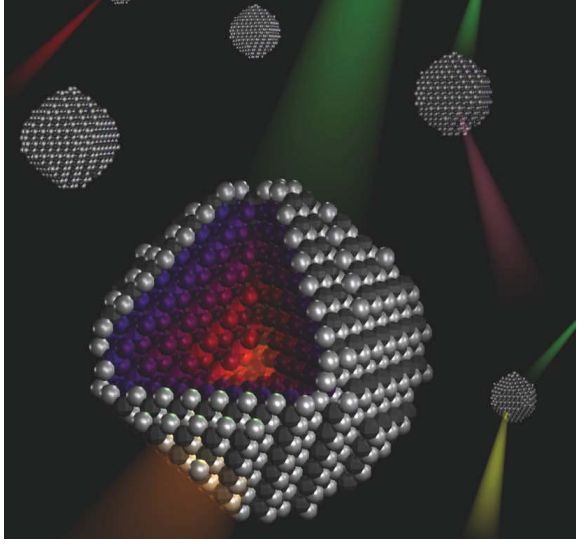

This nonblinking nanocrystal holds promise for eliminating losses due to blinking in LEDs, which show up as heat (courtesy Ted Pawlicki, University of Rochester).

RESEARCH UPDATE

EDITED BY RON WILSON

New "nonblinking" nanocrystals promise increased efficiency in LED lighting

HB LEDs (high-brightness light-emitting diodes) lose about 17% of their energy due to quantum losses, such as imperfections in the crystalline structure, or to a "blinking" effect. To address these problems, the University of Rochester with help from Kodak (www.kodak.com) and the US Naval Research Lab (www.nrl.navy.mil) has developed a nanocrystal that holds promise for eliminating losses due to blinking in LEDs, which show up as heat, a bugbear of solid-state-lighting designers.

"Many molecules, as well as crystals just a billionth of a meter in size, can absorb or radiate photons," says Todd Krauss, professor of chemistry at the University of Rochester and lead author on the study. "But they also experience random periods when they absorb a photon, but, instead of the photon radiating away, its energy is transformed into heat. These 'dark' periods alternate with periods when the molecule can radiate normally, leading to the appearance of [their] turning on and off, or blinking. A nanocrystal that has just absorbed the energy from a photon has two choices to rid itself of the excess energy: emission of light or [emission] of heat. If the nanocrystal emits that energy as heat, you've essentially lost that energy."

Krauss and fellow researcher Alexander Efros from the Naval Research Laboratory conclude that nanocrystals normally have a core of one semiconductor material wrapped in a protective shell of another, with a sharp boundary dividing the two (**Reference 1**). "The new nanocrystal, however, has a continuous gradient from a core of cadmium and selenium to a shell of zinc and selenium. That gradient squelches the processes that prevent photons from radiating, and the result is a stream of emitted photons as steady as the stream of absorbed photons," they state. Krauss refers to the new nanocrystal as being a possible successor to cheap, printable OLEDs (organic LEDs), which are just beginning to find use in displays despite significant drawbacks, such as short lifetimes.**—by Margery Conner**

▷**University of Rochester**, www.rochester.edu.

REFERENCE

Wang, Xiaoyong, Xiaofan Ren, Keith Kahen, Megan A Hahn, Manju Rajeswaran, Sara Maccagnano-Zacher, John Silcox, George E Cragg, Alexander L Efros, and Todd D Krauss, "Non-blinking semiconductor nanocrystals," *Nature* Magazine, May 10, 2009, www.nature.com/nature/journal/vaop/ncurrent/full/nature08072.html.

06.11.09

IMEC DEVELOPS DESIGN FOR DEEP-BRAIN-STIMULATION IMPLANT

IMEC (Interuniversity Microelectronics Center) revealed a design strategy for brain-stimulation implants at the recent DATE (Design, Automation, and Test in Europe) conference in Nice, France. IMEC researchers claim to have created a prototype multielectrode stimulation-and-recording probe for deep-brain stimulation, which could hold promise for medical treatment of Parkinson's disease. "To have a more precise stimulation and recording, we need electrodes that are as small as individual brain cells," says Wolfgang Eberle, senior scientist and project manager at IMEC's bioelectronics research group. "Such small electrodes can be made with semiconductor-process technology, appropriate design tools, and advanced electronic-signal processing."

A deep-brain probe using microelectronics processes could stimulate and measure local areas in the brain, allowing more precise therapy.

The IMEC group focuses on the brain-computer interface between biological and electronic systems. The group designs processes, components, and prototype microsystems for neuroelectronic interfaces that can bidirectionally exchange information with neurons through electrical mechanisms. The strategy allows for the development of advanced brain implants comprising multiple electrodes.

Multiphysics-modeling company Comsol supplied the finite-element modeling of the electrical-field distribution around the brain probe using multiphysics-simulation software. Comsol tools made it possible to investigate the mechanical and structural properties of the probe during surgical insertion and the effects of temperature. By adapting penetration depth and field asymmetry, steering the electricity around the probe results in precise stimulation. For an expanded write-up on this research, go to www.edn.com/090611pc.—by Gail Flower

▶**Interuniversity Microelectronics Center, www.imec.be.**

BY HOWARD JOHNSON, PhD

Endpoint distortion

Figure 1a and **Figure 1b** depict two common transmission-line scenarios: series termination and end termination, respectively. Both drivers are fast with negligible series output resistance. The capacitive loads represent the input capacitances of the CMOS receivers. In the series-terminated case, a step edge from the driver proceeds to the right, interacts with the load, and reflects back toward the driver. The capacitive load may create a strange-looking reflection, and it may distort the appearance of the received signal, but whatever bounces off the load returns to the driver termination and dies, never to be seen again. As a result, the receiver sees one step edge, possibly distorted but with no lingering aftereffects. Of all the things that could go wrong with a circuit, this distortion is not too bad.

Figure 1 illustrates an equivalent-endpoint circuit that defines the nature of instantaneous signal distortion at the receiver. The equivalent circuit comprises two components: a series resistance equal to the line impedance, and a shunt capacitance representing the input capacitance of the receiver (**Reference 1**).

The RC lowpass filter thus formed disperses the input rise time. It also delays the signal's time of arrival by an amount equal to the group delay of the filter, in this case Z_0C_{IN}.

If your native signal rise or fall time is much faster than Z_0C_{IN}, then the filter slows the signal edge to the point where it mimics the filter step response—a nice, clean rising edge with a 10 to 63% rise time of Z_0C_{IN}. If, on the other hand, your native signal rise or fall time is slower than Z_0C_{IN}, the filter has little effect. If you look closely, however, you will see that the filter delays the time of arrival of the signal's midpoint by the amount Z_0C_{IN}.

In **Figure 1b**, a parallel combination of the transmission line and the end termination feeds the capacitive load. If the end termination is properly set equal to Z_0, the parallel combination must be less than Z_0, so the load in this case responds more quickly, causing less distortion than in the series-terminated case.

If the PCB (printed-circuit-board) trace delay is longer than the signal's rise or fall time, the effective impedance at the end of the trace, which you measure on a scale of time commensurate with one rising or falling edge, equals simply Z_0. In that case, the parallel impedance driving the capacitor is $(½)Z_0$, and the time constant associated with the RC filtering effect is $(½)Z_0C_{IN}$, half that of the series-terminated case. On the other hand, if the PCB trace delay is shorter than the signal rise time, the effective impedance at the line's end decreases as it becomes a mixture of the driver and the line impedances, even further reducing the time constant.

The good news is that end terminations respond quickly. The bad news is that the capacitive load degrades the performance of the end termination. Upon receipt of each signal edge, the degraded termination reflects a short pulse back toward the driver. The driver has no termination, so it reflects the pulse a second time. The end of the line thus receives an initial distorted signal edge followed one round trip later by a small reflected pulse.

When signal timing fidelity is of utmost importance, the end-terminated architecture provides the least rise-time dispersion and variation in timing in response to variable load capacitance. In exchange for its improved short-term response, the end termination suffers the possibility that round-trip reflections will interfere with subsequent bits. EDN

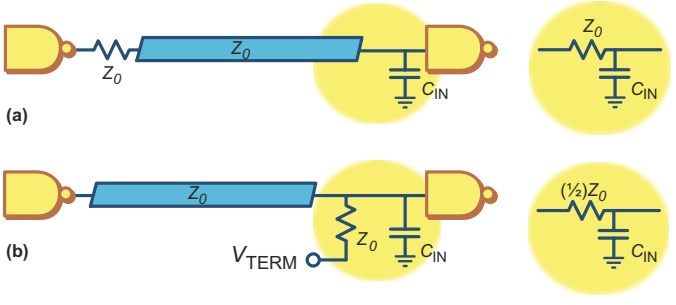

Figure 1 All PCB nets suffer endpoint distortion. Two common transmission-line scenarios are series termination (a) and end termination (b).

REFERENCE

1 Johnson, Howard, "Driving-point impedance," *EDN*, May 14, 2009, pg 12, www.edn.com/article/CA6656309.

TAPE-OUT COMES A LOT FASTER
WITH CALIBRE nmDRC.

INTEGRATED SYSTEM DESIGN + DESIGN FOR MANUFACTURING + ELECTRONIC SYSTEM LEVEL DESIGN + FUNCTIONAL VERIFICATION

Calibre® nmDRC | There's nothing like having an advantage when you're racing to market. That's exactly what you get with Calibre nmDRC. Widely recognized as the world's most popular physical verification solution, Calibre's hyperscaling architecture produces the fastest run times available. To accelerate things even more, Calibre nmDRC adds incremental verification and a dynamic results-viewing and debugging environment. Designers can check, fix and re-verify DRC violations in parallel, dramatically reducing total cycle time. Start out and stay ahead of the pack. Go to **mentor.com/go/calibre_nmdrc** or call us at **800.547.3000.**

Mentor Graphics®

THE EDA TECHNOLOGY LEADER

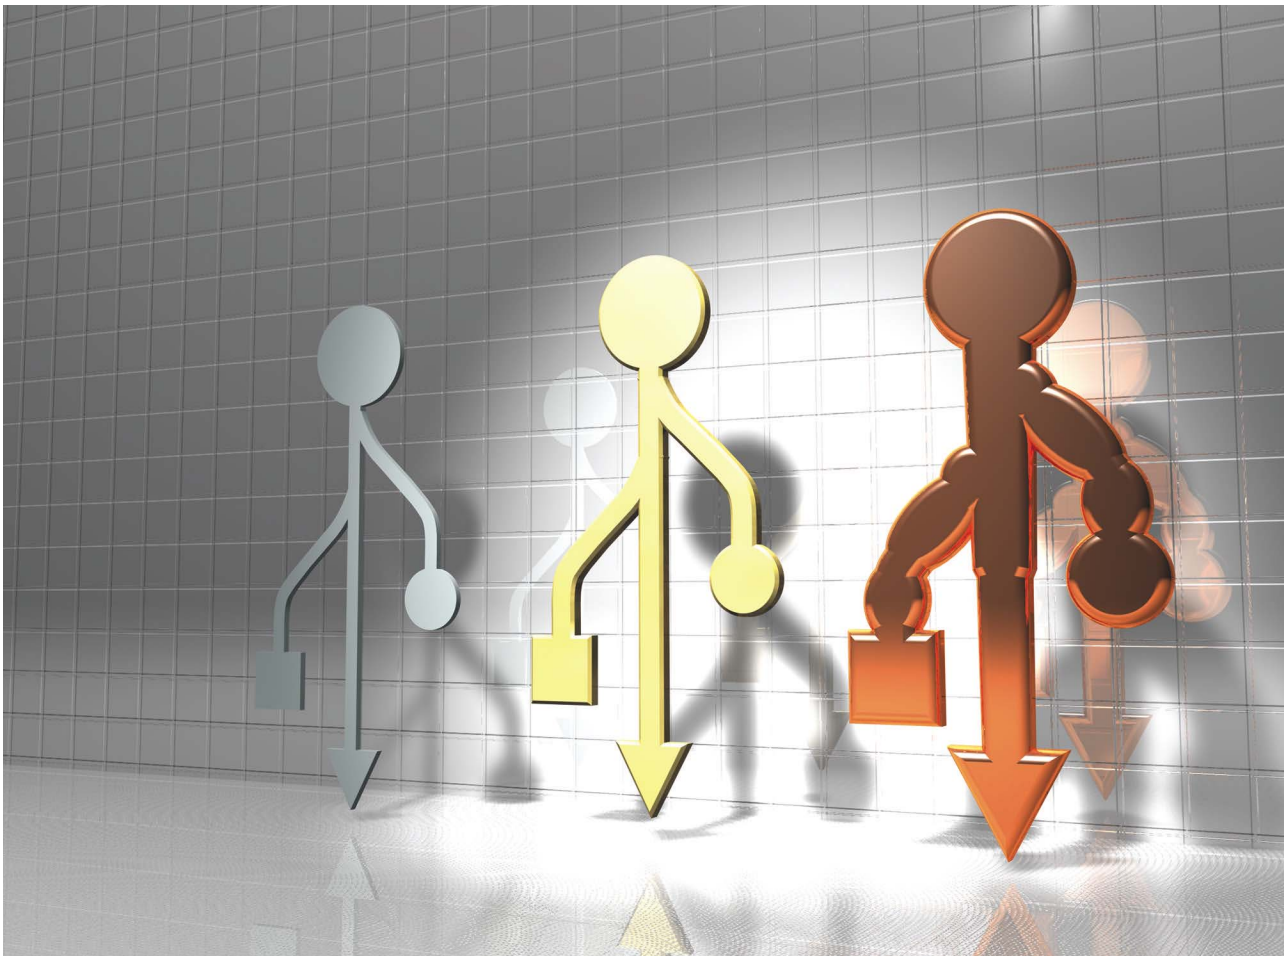

USB 3.0: A SIMPLE IDEA FULL OF CHALLENGES

COMBINING 5 GBPS WITH THE CONVENIENCE OF USB SOUNDS LIKE A SURE WIN, BUT MANY ISSUES ARE HIDING BEHIND THE PREMISE.

BY RON WILSON • EXECUTIVE EDITOR

Super-speed USB (Universal Serial Bus) 3.0 sounds like a great idea. Just start with widely used, fast, and bulletproof USB 2.0 and graft in the PHY (physical-layer) interface from another common and reliable standard, PCIe (peripheral-component-interconnect express) Generation 2. Put two differential pairs into the USB connector to carry the high-speed serial signals from the Generation 2 PHY, and you have a rugged, flexible, inexpensive interface that can operate at 5 Mbps over consumer-priced cables and connectors with interfaces cheap enough to drop into a flash drive.

The idea promises to unleash new ways of using PCs with mobile devices and with storage. With application-level throughput approaching 400 Mbytes/sec and the ability to simply plug anything from a flash drive to 3m of USB cable into a host, users could link PCs and netbooks, quickly dump the contents of huge flash drives, or easily transfer HD (high-definition) video between devices. They could even create their own external storage networks (**Figure 1**). This promise of speed and flexibility, however, carries the seeds of a difficult challenge for chip, board, and system designers.

"The concept of simply using a PCIe Gen 2 PHY in a USB controller is appealing," observes Mike Pennell, vice president of engineering at fabless-semiconductor supplier SMSC (Standard Microsystems Corp). "But you have to remember that, although the USB 3.0 PHY is based on the PCIe Gen 2 PHY, the media over which the two controllers must send data are very different. USB lives in a much more challenging world than [does] PCI."

Navraj Nandra, director of product marketing at Synopsys, explains further: "The similarity between PCIe Gen 2 and USB 3.0 stops at the speed," he says. "They both run at 5 Gbps. PCIe has to successfully work over only 20 inches of carefully designed PCB [printed-circuit board]."

Pennell says that the connection between a USB 3.0 host controller and a device controller could be far more complex than just using FR (fire-retardant)-4 material. The connection would include several inches of PCB, a connector pair, a short pigtail running to the host's enclosure, another connector pair at the back of the enclosure, the 3m cable, another connector pair on the back of the target device, another pigtail, and another PCB. All together, those components make up a highly attenuating channel full of opportunities for strong reflections, and it is highly variable.

People don't often discuss the variability issue in USB, but it becomes hugely important at speeds of 5 Gbps, according to Nandra. "Even at USB 2.0, some certified USB cables are better than others," he warns. "Even wonderful, $50 cables can degrade over time. We have already at the relatively low speeds of USB 2.0 seen poor cable performance impact the performance of the PHY. And, at 3.0, the problem will be much worse."

Cable experts agree that variability is an issue. Peter Smyth, chief executive officer of active-cable vendor RedMere Technology, points out that you can make a USB 3.0 cable good enough to meet the specifications. Doing so also requires tight manufacturing controls, which cost a lot of money, however. And, even with the best controls and within a production run, cables can vary significantly. Then there is the problem of those pigtails, which no one seems to notice much. "That pigtail that runs from the PCB to the back of the box is typically just awful," says Smyth. System designers may find that they have used up most of their eye opening just getting to the back of their own boxes. And, because cables that consistently meet the 3.0 specification will be expensive to manufacture, the door will be open for counterfeit cables. So chip and system designers both must assume the worst—mediocre board design, poor pigtails, cheap connectors, counterfeit cable, and even a wire-bond package for the PHY—when planning their approaches to the new standard.

Unfortunately, successfully pumping 5 Gbps through a messy channel isn't the only problem. The channel is also highly variable. One minute, a port could have a thumb drive plugged into it, and, the next minute, someone could plug 10 feet of cable into a cage full of disk drives. So the PHY must be flexible. Just to make matters more interesting, USB 3.0 will start out expensive, but, by 2011, it will probably be standard in netbook computers and handheld consumer products, such as cameras, media players, and flash drives. The interface must have a migration path to becoming low cost, meaning that it will require an advanced process node. And—because a flash drive, to name one technology, draws its power from the USB cable itself—the interface must be low enough in power to allow cable-powered operation. The USB 3.0 standard will allow a device to draw as much as 900 mA during operation, but it must draw no more than 150 mA before configuration. That limitation itself demands a well-studied power-management strategy at the chip level, which the system designer must understand in detail. All of these issues represent significant departures from the requirements that spawned PCIe 2.0.

AT A GLANCE

- USB (Universal Serial Bus) 3.0 brings 5-Gbps speed to USB by blending USB with a PHY (physical-layer) interface it derives from PCIe (peripheral-component-interconnect express) Generation 2.
- Combining USB with PCIe Generation 2 presents significant design challenges.
- System designers may have difficulty distinguishing a great PHY from a poor one.

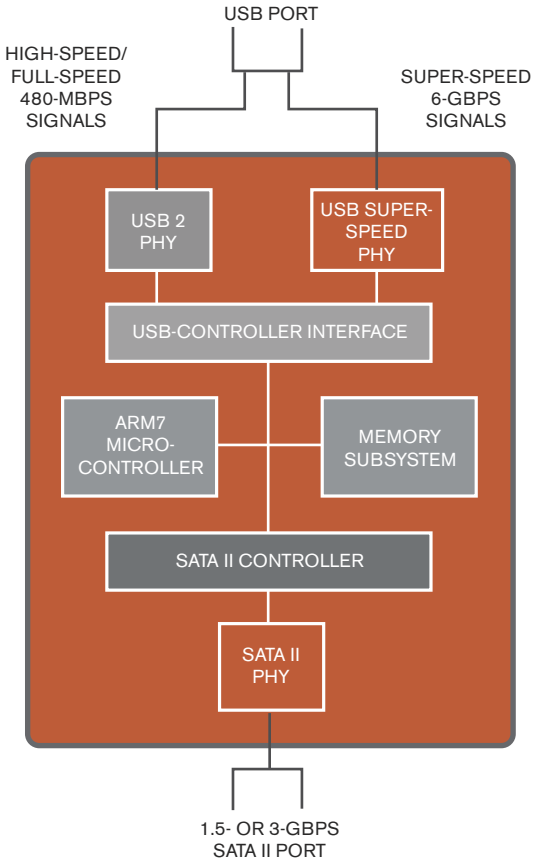

Figure 1 As USB 3.0 interfaces appear in mobile PCs, adapters such as this TI USB-to-SATA bridge will also show up.

CREATING A PHY

Much of the responsibility for getting a data stream through the travails of a USB 3.0 channel will fall on the transmitter pre-emphasis circuit and, especially, the receiver-equalization circuit. Interestingly,

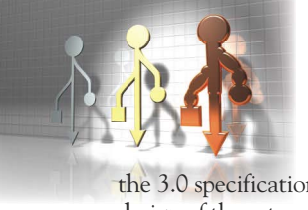

the 3.0 specification appears to leave the design of these two blocks up to the chip-design team. "One of the big differences between PCIe Gen 2 and USB 3.0 is that the PCI document specifies an eye-diagram template at the receiver input," says Synopsys' Nandra. "You have to get the signal there at a specified quality. But, in USB 3.0, the eyes can be so closed by the time the signal gets through the cable that there is no opening to put a template into. So, instead, the 3.0 document specifies an eye-diagram template at the output of the equalizer, not at the input to the receiver."

One of the greatest differences between PCIe Gen 2 PHYs and USB 3.0 PHYs will be in the receiver-equalization circuit. Many designers expect the quality of this block to be a major differentiator in the market, for PHY-IP (intellectual-property), chips, and the systems that use them. The equalizer must be both powerful in its action and adaptive. Otherwise, a PHY would be unable to handle the range of channel conditions that USB can throw at it. As an adaptive equalizer, this circuit will require a training sequence to lock onto. Yet, the equalizer must be low in power and compact to meet the needs of consumer-product applications. Those challenges are formidable.

Also, although PCIe Gen 2 is a fully synchronous interface, Nandra points out, USB 3.0 requires that the PHY use a spread-spectrum clock in a way that makes the transmitter and receiver essentially asynchronous. The receiver CDR (clock- and data-recovery) circuit must recover the transmitting clock without having access to the spread-spectrum signal that modulated it. "That requires a very elastic CDR," observes SMSC's Pennell.

Such issues will make previous experience with PCIe Gen 2 a real asset for chip designers. "If a design team has really understood PCIe Gen 2, then they can adapt about 90% of their work to USB 3.0," says Scott Kim, manager of business development at Texas Instruments. Yes, the equalizer will need beefing up, but much of the circuitry will remain the same. For example, getting 5-Gbps performance from the PHY requires a careful trade-off between deep pipelining to meet throughput requirements and limited latency to meet bus timing.

Kim also emphasizes the importance of experience with power management. "The first step is to run all the circuits as slowly as possible to conserve energy," Kim says. "For instance, you have to keep the receiver listening for the LFPS (low-frequency-periodic-signaling) traffic so that you know when to wake up. But that [requirement] doesn't mean you have to run the whole receiver at full speed. We've figured out a way to filter

THE FIRST STEP IS TO RUN ALL THE CIRCUITS AS SLOWLY AS POSSIBLE TO CONSERVE ENERGY.

WHY DO WE NEED A 5-GBPS USB?

Why would most USB (Universal Serial Bus) users today even want so much speed in a USB link? After all, USB is supposed to provide flexibility and convenience for using mice, printers, flash drives, and the like. Consider flash drives, however. Capacity for the largest thumb drives has soared into the gigabytes, so transferring a significant portion of the contents of one of these drives through USB 2.0 can be a painful task. Meanwhile, marketers point to the supposed emergence of the home-media PC, with its enormous data files–think entire HD (high-definition) movies–and network-attached storage. Having a 5-Gbps version of USB would allow most home-media users to simply plug their external drives into their USB ports and not worry about the mysteries of SATA (serial-advanced-technology-attachment) or storage-area networks.

Also, forecasters expect a huge growth over the next few years in netbook computers and other mobile computing, Web-accessing, and media-gulping devices. Today, about the only option for creating a high-speed link between such a mobile device and your home PC is to set up a home network, complete with the hardware hassles and administration problems. A single USB cable would be a much simpler approach. And companies such as PLX Technology are working on protocols for using SuperSpeed USB as a virtual network between computers.

Jimmy Chou, director of marketing for storage and USB products at PLX Technology, forecasts that consumer PCs will be the early adopters of USB 3.0, with discrete PHY (physical-layer) chips going into notebooks early next year. A lot of notebook vendors have such designs under way now, he suggests. Meanwhile, other sources suggest that makers of handheld devices that trade in large data files, such as videocameras and movie players, may be on a similar time line.

A big milestone will arrive when Intel, in about 18 months, releases a south-bridge chip with integrated 3.0. With that announcement, momentum should grow until 3.0 has become universal by mid-2011, according to forecasts. Predictably, the fly in the ointment may be software. Chou says that the host-side driver for 3.0 is about an order of magnitude more complex than the device-side driver. Device vendors will also have to supply some code to the host to execute with the resident host driver. Microsoft does not plan to release any host driver for 3.0 before the Windows 7 service-pack 1, according to industry talk. So, it may turn out that, for all the added hardware complexity that USB 3.0 is imposing on chip vendors, the real delay may be in just getting the drivers right.

LFPS while running at very low power." Kim also cites TI's extensive low-power-design method, which now routinely includes clock gating and power gating. Such techniques allow designers to run the transmitter-pre-emphasis circuit at varying power levels, depending on the needs of the channel. Similarly, Synopsys 3.0 IP will reduce receiver-equalizer power to just the level the circuit needs for the required equalization.

All of these challenges represent a substantial investment, in both design and testing, according to Jimmy Chou, director of marketing for storage and USB products at PLX Technology. PLX has the benefit of an established presence in both USB 2.0 and PCIe Gen 2, but Chou says that the engineering investment was still substantial. In addition to the usual costs of chip design, he says, the company has spent about a year and a half in the lab testing and validating its PHYs, especially the equalizer algorithms.

LEGACY TROUBLES

One of the serious implementation challenges with USB 3.0 turns out to be the requirement that the PHY operate simultaneously in both legacy USB 2.0 mode and 3.0 mode. Most design teams have looked at the problem and concluded that the best approach is to simply drop a commodity USB 2.0 PHY-IP block into the design next to the newly designed 3.0 hardware. TI, for example, takes this approach, and most IP vendors will probably also do so. The cell exists in their libraries, so why not just use it? According to Synopsys' Nandra, however, there are good reasons not to. At the layout level, you can combine the digital portions of the two PHYs and save some real estate. With cleverness, you can also reduce the pin count of the block and the frontage on the perimeter of the die. More important, according to Nandra, is the fact that a lot of signals must pass between the 2.0 and 3.0 PHY devices during operation. It may be unwise to expose these internal signals to the chip designers who don't know the details of USB 3.0 operation.

Nandra also worries about crosstalk when both 2.0 and 3.0 modes are in simultaneous operation. He points out that, if you are using a continuous-time linear equalizer, turning up the gain on the equalizer won't help because doing so amplifies not just the signal, but also the crosstalk. So you may have to implement active crosstalk-suppression circuitry for simultaneous operation.

With complex channels varying in characteristics, power and cost constraints, legacy compatibility issues, and some lingering questions about the robustness of the standard, USB 3.0 is not a walk in the park. Given the range of applications (see **sidebar** "Why do we need a 5-Gbps USB?") and the momentum behind the standard, however, IP and chip vendors will support the effort. They will present system designers with a variety of choices with different price and power points.

It may prove exceedingly complex for system vendors to determine exactly what kind of performance they are getting for their dollar and their milliwatt, however. Extensive testing with board layouts, different lengths and qualities of cabling, and varying connector quality may be the only way to tell the really top-quality PHY, which can deliver consistently high data rates and low bit-error rates, from the bargain-basement PHY, which will work well only under ideal conditions. **EDN**

Go to www.edn.com/090611df and click on Feedback Loop to post a comment on this article.

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

FOR MORE INFORMATION

Intel Corp
www.intel.com

Microsoft
www.microsoft.com

PLX Technology
www.plxtech.com

RedMere Technology
www.redmere.com

Standard Microsystems Corp
www.smsc.com

Synopsys
www.synopsys.com

Texas Instruments
www.ti.com

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

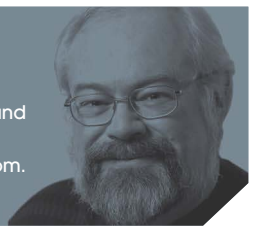

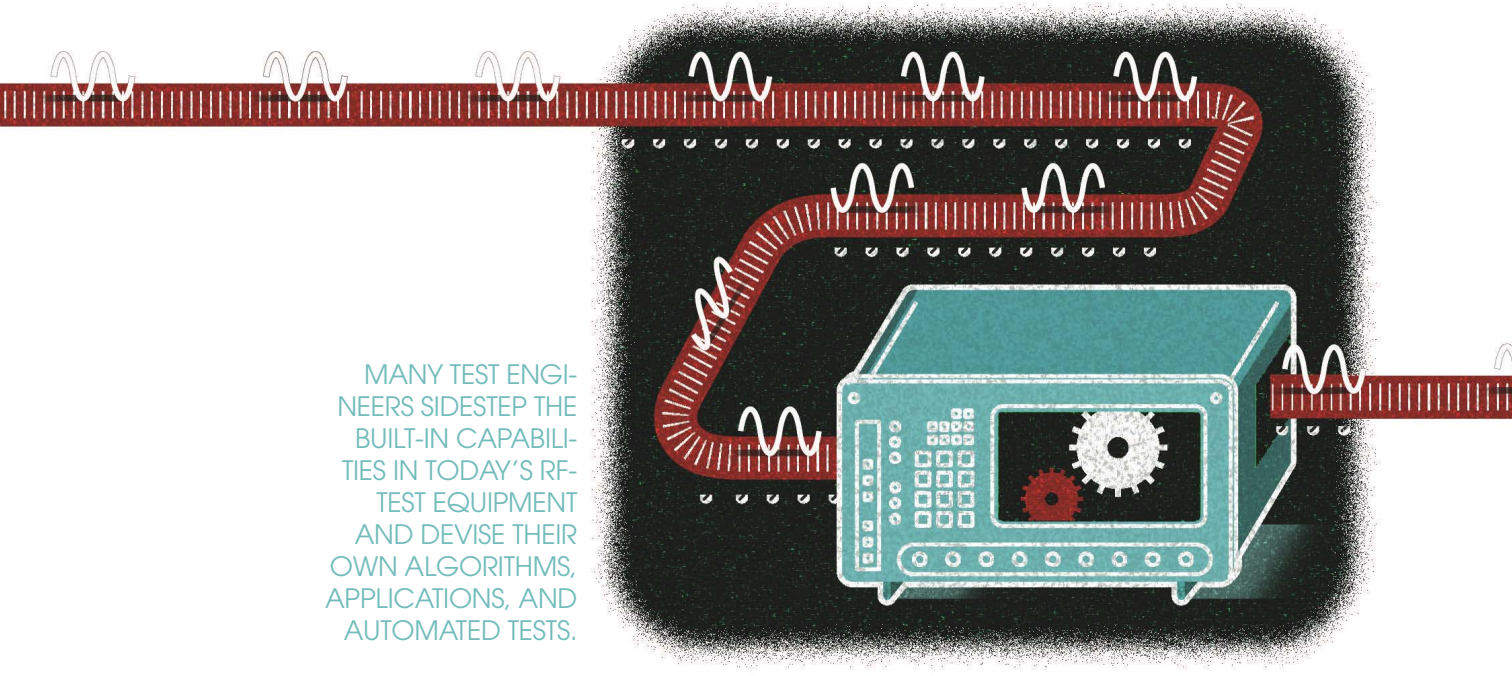

MANY TEST ENGINEERS SIDESTEP THE BUILT-IN CAPABILITIES IN TODAY'S RF-TEST EQUIPMENT AND DEVISE THEIR OWN ALGORITHMS, APPLICATIONS, AND AUTOMATED TESTS.

RF ENGINEERS AUTOMATE TESTS

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

Testing of RF devices, such as amplifiers and RFICs (radio-frequency integrated circuits), can be tedious work. Such devices work over a wide range of frequencies and power levels, and they must meet specifications over temperature and power-supply ranges. Testing for all of those conditions can generate loads of data. Fortunately, automation can cut test time and help you make sense of all that information. You may have a spectrum analyzer, a network analyzer, or a power meter with features that can improve testing, but you may be unable to use the instrument if you must maintain compatibility with older models. If you're testing leading-edge RF products, your test equipment may lack the necessary dedicated features, and you'll have to develop your own.

Bill Drago, a test engineer at L-3 Communications' Narda Microwave East (www.nardamicrowave.com/east), supports production of RF amplifiers, downconverters, upconverters, and transceivers that operate in the C through Ka bands. These products often remain in production for years, which is the reason that Drago is unable to take advantage of features in newer test equipment that can automate many of the measurements he needs to make. He developed his tests before such automation was available, and he needs all of his test instruments to continue to follow identical procedures. Thus, he has written software that performs automated measurements, such as amplifier gain, 1-dB compression, IMD (intermodulation distortion), return loss, spurious noise, and noise figure. According to Drago, spurious-noise testing is important. "Downconverters and upconverters mix an input signal with a local oscillator," he says. "The converter's local oscillator must be tuned to customer specifications within a certain range and step size. The converter needs a frequency synthesizer that's programmable with specified steps over its frequency range. The frequency synthesizer can't introduce any spurs into the converter, so we must test for that."

To measure spurious noise, Drago uses an Agilent Technologies (www.agilent.com) spectrum analyzer to perform a frequency sweep through an approximately 1-GHz band around the carrier. He usually breaks that sweep into a number of smaller sweeps, each comprising 601 frequency points. Each step might be a few kilohertz wide. If Drago were to use one large sweep, its step size and the instrument's resolution bandwidth would be too wide and might miss a spur. He adjusts sweep span and resolution bandwidth so that spurs don't fall between the points in a sweep. Using a number of

smaller sweeps is also faster than a single sweep for the bandwidth that Drago needs. He notes that a single sweep could take as long as an hour, depending on resolution bandwidth and frequency. Furthermore, smaller sweeps can reveal failures sooner than waiting for a large sweep to complete.

Although some of Drago's spectrum analyzers have built-in test applications for measuring spurious noise, he doesn't use them, because not all of his spectrum analyzers have that function. If he were to use that feature, he might not have a replacement instrument for the production line should that instrument fail. Instead, he wrote his own applications, and, by keeping the test applications outside the instrument, he can use any spectrum analyzer that's available.

Drago has written several other test programs, including one that measures an amplifier's 1-dB compression point. He builds this test into some of his VNAs (vector-network analyzers). The algorithm uses a binary-search process, similar to the type that SAR (successive-approximation-register) ADCs use. He starts with an input signal from an Agilent RF-signal generator that's the highest possible value for the amplifier under test. He then measures output power with an Agilent RF-power meter. If the output signal is compressed by more than 1 dB, he cuts the input signal in half and then increases it or decreases it by half of that value until he finds the 1-dB compression point (**Figure 1**).

AT A GLANCE

- Testing over temperature and power-supply ranges can generate lots of information, and automation can help you make sense of it.
- Engineers may or may not use an instrument's built-in functions, depending on the needs of both their customers and the design teams within their companies.
- GPIB (general-purpose-interface-bus) communications are good for sending a series of short commands, whereas Ethernet works with the logic analyzer to collect data on digital-baseband signals.
- With loop-back testing, a manufacturer can wirelessly test a Bluetooth product as it moves along a production line.

FINDING HARMONICS

Michael Ford is a test engineer at Comtech PST (www.comtechpst.com), a manufacturer of RF amplifiers that operate at 500 MHz to 6 GHz at power levels of 100W to 10 kW. Ford's typical test station contains an RF-signal generator, a spectrum analyzer, a network analyzer, a power meter, and RF switches (**Figure 2**). A USB (Universal Serial Bus) digital-I/O module from Measurement Computing (www.measurementcomputing.com) controls the RF switches. The amplifiers mount on an environmental plate that changes their temperature. The test station measures gain, output power, harmonic distortion, IMD, efficiency, spurious noise, and harmonics.

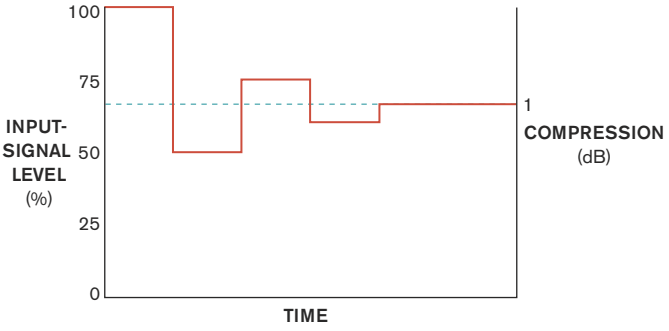

Figure 1 A binary-search algorithm finds an amplifier's 1-dB compression point (courtesy L-3 Communications).

On occasion, Ford uses an instrument's built-in functions. For example, he might measure spurious noise with an Agilent spectrum analyzer's built-in application. But he also writes his own applications to make those measurements if his instrument doesn't have that feature. Ford supports both engineering and production. When making measurements for engineering evaluations, he uses built-in functions, such as spurious noise and harmonic analysis. For production, Ford always uses his own software routines. "We write our own routines using a modular format because it lets us use equipment from multiple vendors, such as Agilent or Rohde & Schwarz [www.rohde-schwarz.com]. We use the same test routines and just change instrument-command libraries." For example, his routine for spurious-noise measurements works with spectrum analyzers from either manufacturer.

Ford's routines for measuring harmonics of a carrier frequency use variables for parameters such as center frequency, span, resolution bandwidth, and video bandwidth. After receiving those parameters, the routine runs sweeps at multiples of the carrier frequency to find the power of its harmonics. The results go to a spreadsheet for analysis.

CELLULAR NETWORKS

Although Drago and Ford must support products that remain in production for years, engineers developing tests for RFICs face different problems and have different reasons for not always using an instrument's automation features. Joe Flynn is a staff engineer at fabless-semiconductor company Sequoia Communications (www.sequoiacommunications.com). He evaluates RFICs that support several cellular-wireless standards (**Reference 1**). He has developed several test stations for evaluating the devices. "We characterize transceivers for gain, noise figure, IMD, cross-modulation, and EVM [error-vector magnitude]," he says. Flynn also develops his own automation tools, but measurement speed is the critical feature in choosing equipment. Because RFICs transmit and receive modulated signals, Flynn's test stations include spectrum analyzers for characterizing frequency content, and they include modulation analyzers for characterizing modulated content. **Figure 3** shows a system that tests the receivers in Sequoia's ICs. The receiver test bench lets Flynn evaluate how a receiver performs in the presence of undesired blocking signals, such as those from simulated cellular base stations, other cell phones, and broadcast-radio stations.

As part of the blocking-signal test, Flynn must measure SNR (signal-to-noise ratio) over a frequency range of 100 kHz to 12.7 GHz in 200-kHz steps. That

task takes approximately 60,000 measurements per channel, and the RFIC has 1300 channels over seven frequency bands. An SNR test generates loads of measurement data, and it's just one of many tests that Flynn must run on a preproduction lot of parts. To help analyze the data, Flynn developed a data converter that produces data plots. The tool lets him use production-test-analysis tools to view engineering data. For example, he might want to see the distribution of parameters, such as gain, return loss, and current consumption across the 100 parts in a preproduction run. "When we look at the data, we get a feel for how production parts will behave," he says.

Sequoia also developed an in-house Visual Basic Web-based tool that manipulates the bench-characterization data. The tool uses .netCharting (www.dotnetcharting.com) software to create some 1200 data plots on the parts. "When we have a test review, the tool lets us find any out-of-spec measurements," he says. The tool lets him select data parameters to plot and refine the data by selecting certain test conditions. For example, it lets him look at receiver gain and noise figure versus temperature or power-supply voltage.

Because he must make so many measurements, speed is key for Flynn when he selects a spectrum analyzer and develops a test method. To cut measurement time, he uses two Agilent MXA spectrum analyzers, one each for a receiver's in-phase and quadrature channels. The instruments are frequency- and phase-locked, which synchronizes the measurements. He then optimizes resolution bandwidth and sweep time to minimize test time: "When I started, a single SNR measurement on one channel at one blocker frequency took about a second. Now, I make each measurement in 18 msec." Flynn also learned how to minimize test time by using both LAN (local-area-network) and GPIB (general-purpose-interface-bus) communications. Most instruments in the test system use GPIB because, according to Flynn, "You can't beat GPIB when sending a series of short commands. The overhead needed to use Ethernet is apparent only when transferring large blocks of data." Thus, he uses Ethernet for the logic analyzer, which collects data on digital-baseband signals. He also uses a dedicated GbE (gigabit-Ethernet)-LAN card to avoid packet collisions between the test equipment and the corporate network. He uses three GPIB cards to cover all of the test equipment.

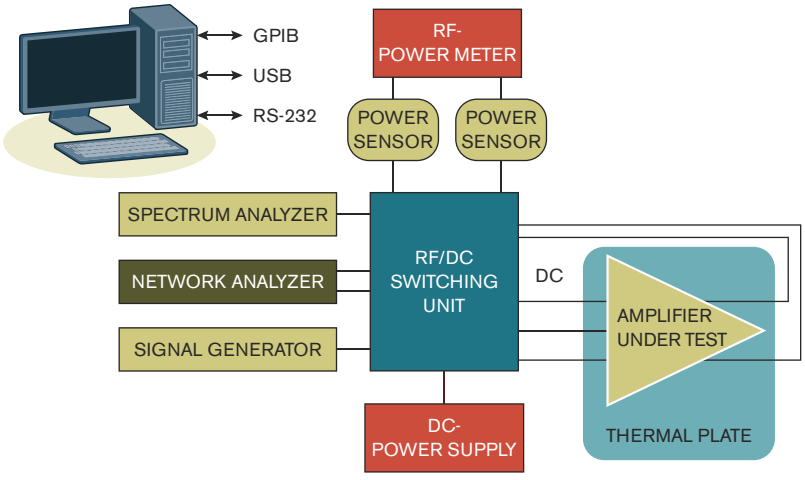

Figure 2 A typical automated tester for RF amplifiers uses a power meter, a spectrum analyzer, a network analyzer, and a signal generator (courtesy Comtech PST).

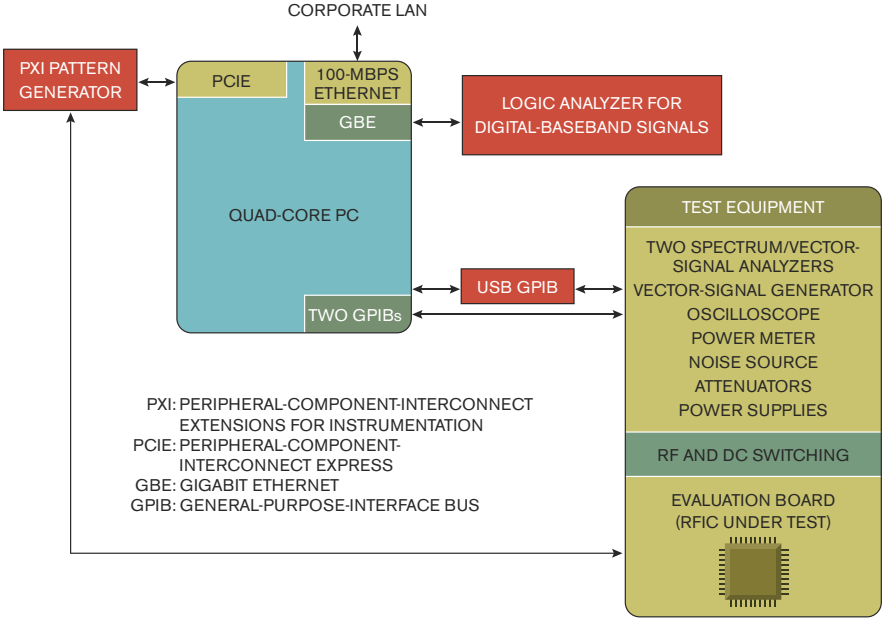

Figure 3 An RFIC-test station for testing receivers uses two GPIB cards and a LAN to communicate with instruments (courtesy Sequoia Communications).

MIX AND MATCH

CSR (www.csr.com) is a fabless-semiconductor company that develops wireless-communication RFICs for PANs (personal-area networks), such as Bluetooth and Wi-Fi. James Blackwell, who heads the company's applications-engineering group, helps customers evaluate CSR's RFICs and develop products based on the company's devices. CSR engineers use a mixture of in-house and purchased test equipment, usually starting with in-house testers. "Because we're often on the leading edge," says Blackwell, "we have to develop our own test suites until the test-equipment companies catch up." One example is a Blue-

tooth tester for performing loop-back tests. Blackwell notes that the Bluetooth specification defines a loop-back test in which the tester wirelessly controls the DUT (device under test). With loop-back testing, a manufacturer can wirelessly test a Bluetooth product as it moves along a production line.

When CSR's engineers several years ago developed a Bluetooth RFIC, engineers had to build their own RF-test system using RF-signal generators, spectrum analyzers, and vector-signal analyzers. They used The MathWorks' (www.mathworks.com) Matlab scripts to control the instruments, process the data, and produce test reports. In a loop-back test, the tester commands the DUT to produce a 2.405-GHz tone with a specified modulation signal, for example. The tester measures as many as 20 parameters, such as peak and average transmitted RF power. It also performs receiver tests, such as sensitivity and bit-error rate.

Some Bluetooth devices transmit at controlled power levels, so the DUT must operate in steps in its power table while the tester measures differences in power. It also measures frequency tolerance and drift. A test of the frequency response of the DUT's modulation filter uses 10101010 and 11110000 bit patterns. Over time, test-equipment makers developed Bluetooth testers, and CSR was able to use them. Today, the company's engineers use Bluetooth testers from Rohde & Schwarz, Agilent, and Anritsu (www.anritsu.com). CSR has all three instruments, so the engineers always have one that their customers use. That requirement is crucial when an engineer needs to reproduce a customer's test.

CSR engineers didn't immediately switch to a dedicated Bluetooth tester. "We work with the test-equipment manufacturers to develop test applications for their equipment," Blackwell explains. "But sometimes we must wait for a second or third generation of a tester before we can use it. Even after we adopt a commercially available tester, we may still use our own test suites for certain tests." Blackwell notes that dedicated Bluetooth testers may perform some tests faster or more accurately than CSR's in-house testers, but the company's engineers still use the in-house tester when they feel that it performs the tests better than a dedicated tester can.

TEST COMPANIES RESPOND

Despite the best efforts of test-equipment manufacturers, some engineers still often find that they need to develop their own test algorithms. Test-equipment makers point out that some tests require application software in the instrument to perform a test because you need real-time results. **EDN**

REFERENCE

1 Rowe, Martin, "The RFIC evaluator," *Test & Measurement World*, March 2009, pg 9, www.tmworld.com/article/CA6639373.

A version of this article appeared in the April 2009 issue of *Test & Measurement World*.

TEE UP YOUR

MULTIPROCESSING OPTIONS

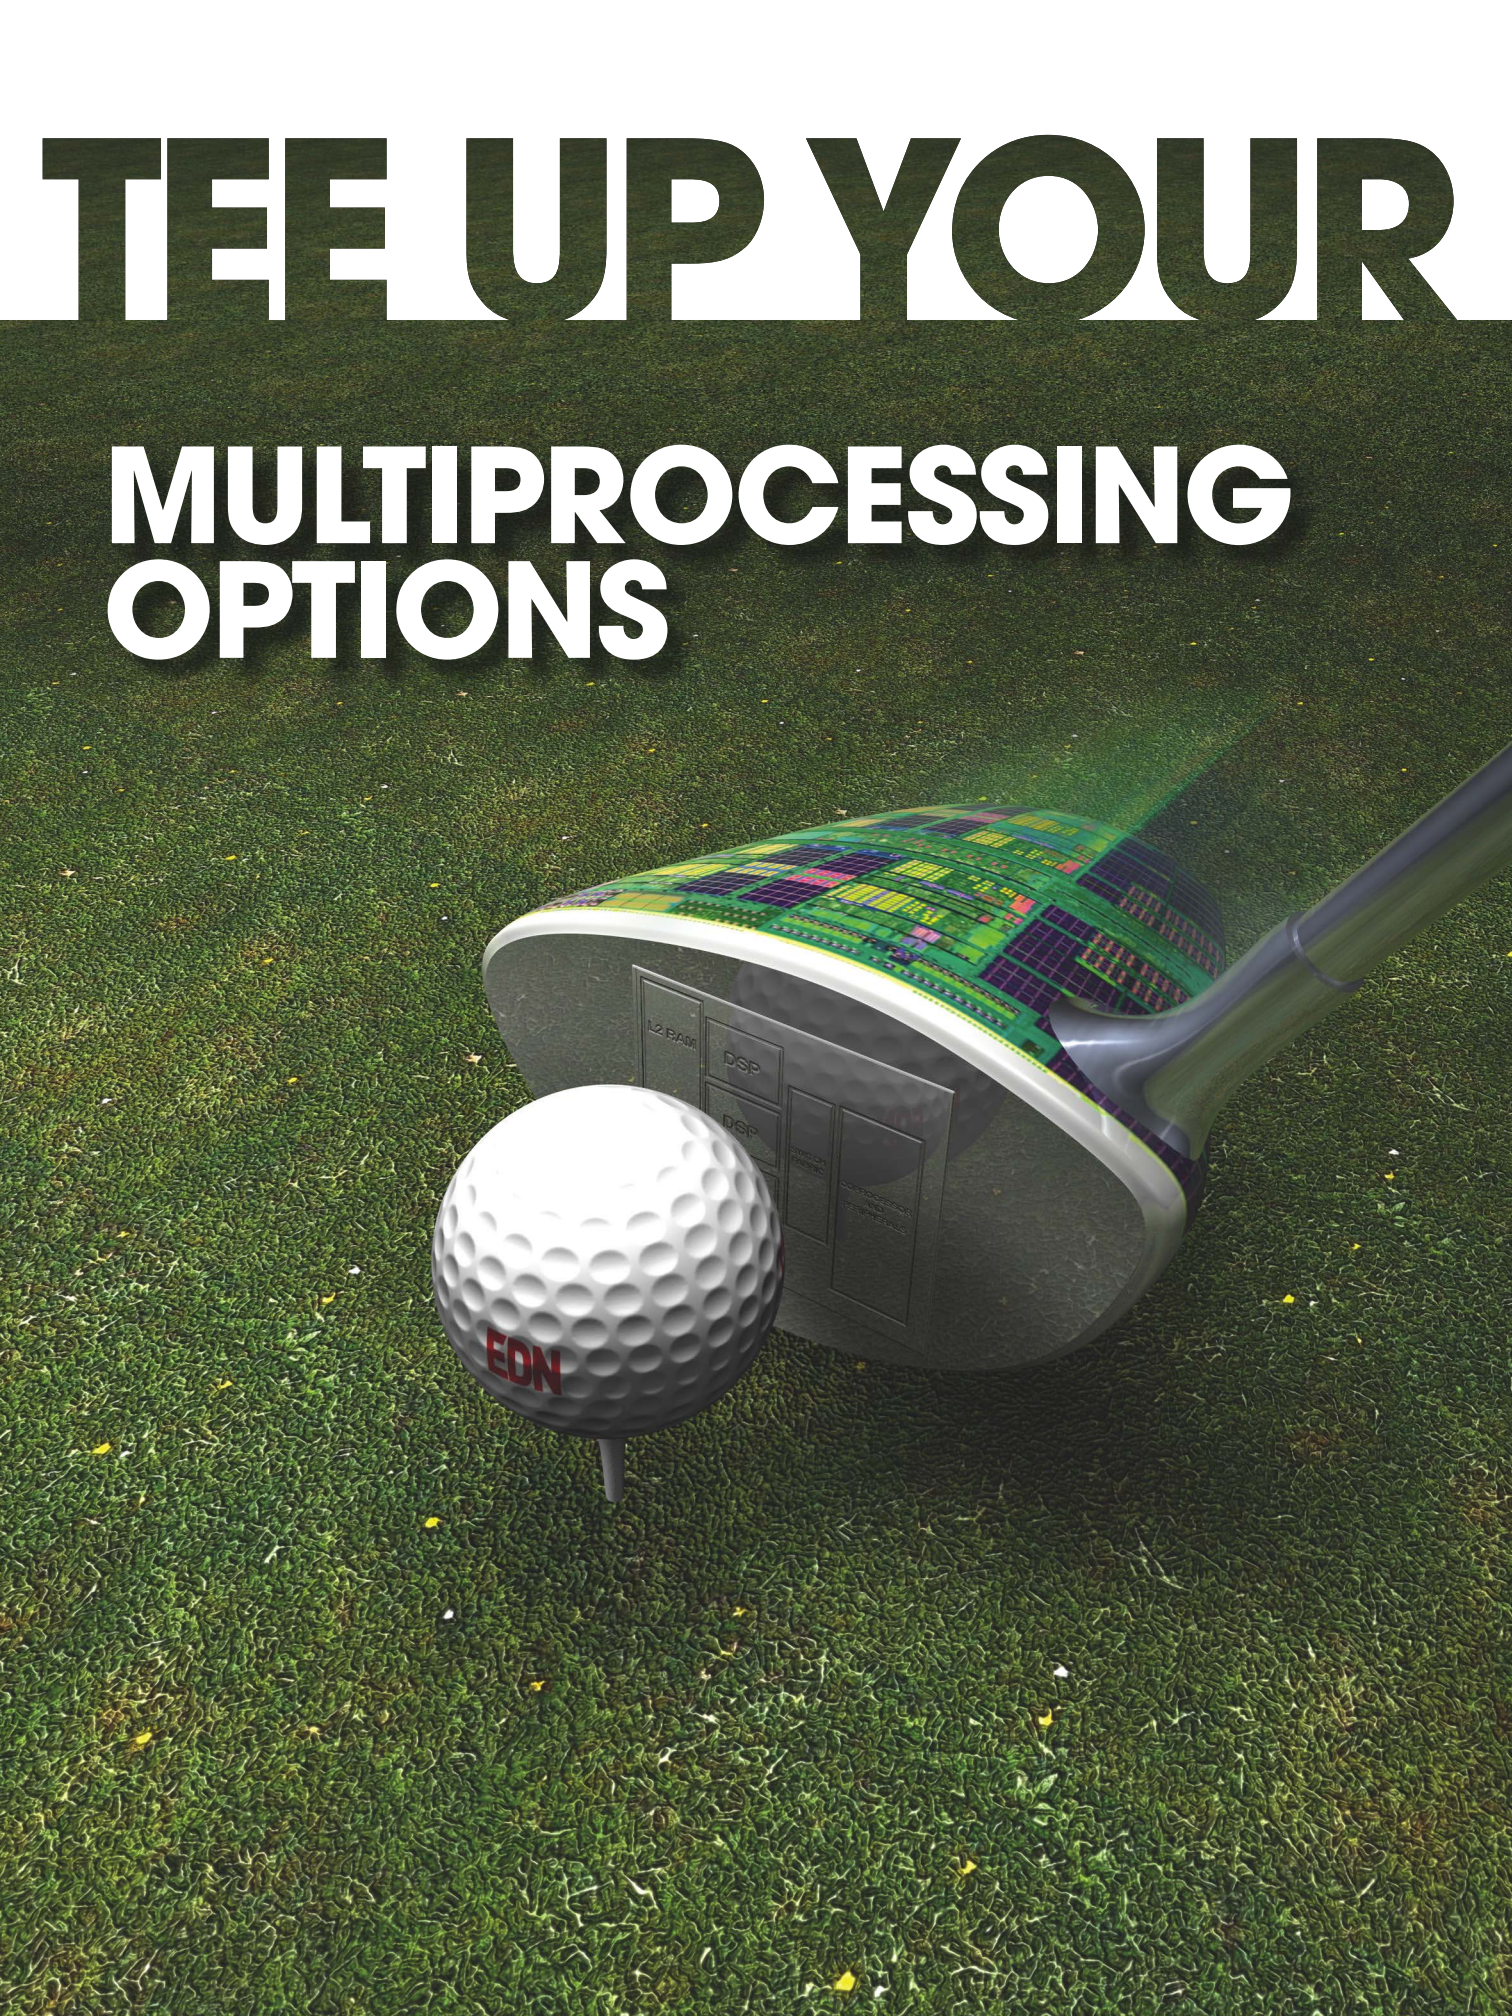

A GOOD PROCESSOR TAXONOMY HELPS YOU UNDERSTAND PROCESSING ARCHITECTURES' SWEET SPOTS. YOU SHOULD PROBABLY LOOK AT MULTICORE OPTIONS IN THE SAME WAY.

Vendors categorize single-processor architectures by the types of applications they target. For example, microcontrollers, DSPs, network processors, and DSCs (digital-signal controllers) employ different implementations to meet the needs of their target domains. Classifying processors at this top level enables designers to quickly identify candidate processors without delving into the implementation details of every processor. Differences in implementation details within a processor category become differentiators with a much shorter list of candidates as a result. Categorizing processors in this way allows different versions of each type of processor, such as low-power DSPs and microcontrollers, to use the same implementation techniques without letting that implementation feature confuse what the target is. This issue is one of the points of confusion affecting a contemporary multicore taxonomy.

Another point of confusion is that the term *multicore* has a different meaning depending on who is using it. One use of the term describes multiple processors within a system. Another use constrains the multiple cores to one package, and yet another usage further constrains the multiple cores to one die. People also use the term to refer to identical cores or heterogeneous cores in each of these contexts. These usages fail to converge on one meaning, even though these differences, in many cases, are basically transparent to the software-development effort. This article uses the loosest interpretation of *multicore* because a chip-level implementation today was a board-level implementation yesterday, and that difference does not change how you build the software. Less transparent to the software-development effort, however, is the application domain; each type of processing, most of which the single-core taxonomy captures, requires different analysis and programming knowledge, and different tools and libraries serve them.

Some multicore processors place their primary focus on performance by stating how many cores they support, how much bandwidth they can handle, and whether they support SMP (symmetrical-multiprocessing) and AMP (asymmetrical-multiprocessing) configurations. As relevant, they also disclose information about coherency mechanisms, shared-data performance, multithreading support, operating-system support, and the ability to balance workloads between cores. You may have to infer the target application and the favored trade-offs from the information the vendor offers.

In contrast, some multicore offerings, similar to single-core offerings, place their primary focus on identifying their target application, which is often digital media, networking, wireless infrastructure, and servers. They then get into the implementation details, such as coherency or processor interconnects, to differentiate them from other similar multicore offerings. However, multiprocessing is not limited to this short list of applications. Embedded-system designers have for decades been quietly and successfully implementing multicore designs. To understand how a more-inclusive multicore taxonomy might look, you should explore a taxonomy for single-core architectures and how you might extend

it to incorporate multicore offerings.

Each of the single-core processor architectures available best fits into a sweet spot of processing complexity (**Figure 1**). Although each type of processor can perform other types of processing, they are best at the type of processing that the target needs the most. The proposed primary characteristics for mapping each processing sweet spot are computational load and number of states, or contexts. Both of these characteristics are measures of complexity. Computational load can indicate the peak magnitude, total amount, or sustained amount of processing performance the system needs within a system cycle. The number of states can indicate internal system states, number of system inputs and outputs, or a level of possible states or contexts that the system must support. The reason for proposing loose definitions for these characteristics is to rein in complexity and accommodate the range of processing scenarios that include managing scalar data or control flows, managing aggregate data or control flows, and managing many channels of scalar or aggregate data or control flows. Each of these characteristics is implementation-independent and permits an orthogonal differentiation for each type of processor. Each of these processor types evolved over time at different rates, and each trades one or more measures of performance to maximize one or more other measures of performance.

AT A GLANCE

- Processing architectures choose optimization preferences that target a processing sweet spot.
- Each type of processing has its own development ecosystem.
- Focusing on functional capabilities enables stakeholders to converge their efforts to evolve their type of processing.
- Multicore discussions should expand beyond digital media, networking, and server applications.

SWEET SPOTS

Microcontrollers are specialized processors that offer a cost and power-efficiency advantage at the expense of flexibility and processing performance. They provide a cost advantage by incorporating memories and peripherals in the same package. They provide a power-consumption advantage partly because they support lower clock rates and partly because they implement only the minimum set of circuitry to perform control processing. If a design needs the flexibility of a larger or a smaller memory, different peripheral set, or higher clock rate, the design must swap out the processor for another. To accommodate this change, processor vendors offer families of devices with many variations of the same microcontroller that include differing amounts of memory, peripheral sets, and supported clock rates so that developers can realize the best cost and power efficiency without losing flexibility.

The processing sweet spot for microcontrollers is systems that must respond—often deterministically—to external real-world events, such as for motor control, with tight latency requirements. Microcontrollers are capable of rapid, frequent, and prioritized context switching; they usually can also handle many different contexts. A common differentiation for handling context switching includes employing specialized hardware-interrupt processing that can sense and react to external events and change the microcontroller's internal context within a few clock cycles. Alternatively, the system can avoid context-switching overhead by using peripheral-to-peripheral communication coprocessors that are independent of the processor core. Some microcontrollers differentiate for deterministic operation by not using caches or pipelines, so as to avoid the uncertainty from stalls during cache misses or pipeline flushes.

Microcontroller suppliers differentiate their devices with features such as vertically appropriate peripheral sets, on-chip memories, a range of clock rates, packaging options, power-management options, deterministic operation, and development support. Because projects that include microcontrollers are often cost-sensitive, it is important to include only the necessary peripherals and amount of on-chip memory to complete the project. To support higher clock rates, microcontroller suppliers can employ many techniques, such as wide data buses, to mask the access latency to on-chip flash memory. Some of these approaches focus on burst/peak performance versus sustained performance before encountering a wait state. Power management is common in microcontrollers with many ways to implement sleep and low-power modes that have varying granularity for turning off parts of the processor's resources. Each of these differentiations is an example of implementation details that do not

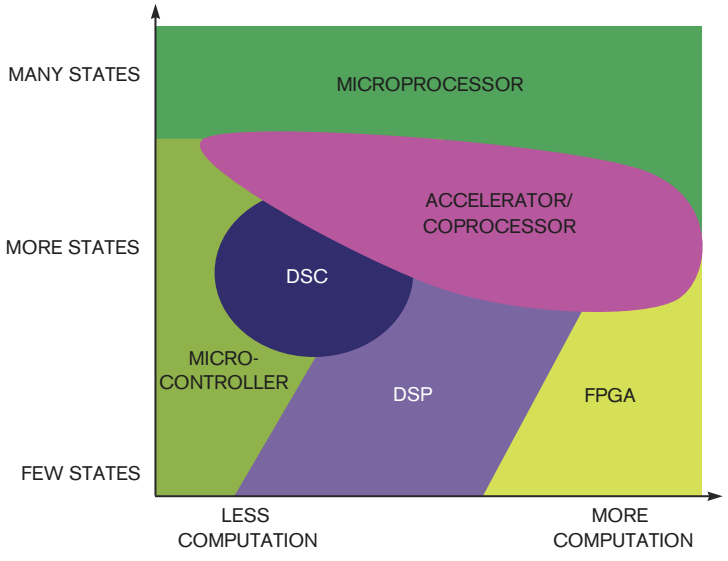

Figure 1 This taxonomy mapping highlights processing sweet spots of mainstream processing architectures and how embedded developers can combine them to complement each other and cost-effectively cover the range of an application's processing requirements.

change the primary function of the microcontroller but enable it to compete with other microcontrollers.

Software development with microcontrollers often involves open- or closed-loop control algorithms as well as filtering of real-world inputs. Microcontrollers are common in board- and system-level multicore embedded designs in which they find use in distributed or periodic system monitoring, distributed control, management of simple user interfaces, and even low-power-supervisory functions to power up and down more expensive parts of the system. They are also present in some multicore chips alongside a DSP as an efficient way to complement the strengths and weaknesses of the DSP architecture.

DSPs are similar to microcontrollers in that they offer advantages in cost and power efficiency at the expense of flexibility and processing performance. However, they substantially differ from microcontrollers because they sacrifice efficiency in handling context switching to maximize their performance of continuous and repetitive calculations, such as in most signal-processing tasks. They also do not integrate many peripherals because they are not ideally structured to handle the context switching that handling many peripherals could require. The most common integrated peripheral is an ADC that the DSP uses to collect a stream of real-world data to process. DSPs often support fixed-point operations to optimize the performance and energy efficiency for fractional mathematics, and they use special address generators to optimize algorithms that work on arrays, circular buffers, or even bit-reversed values. They employ multiple buses and memory structures so that they can do simultaneous memory operations to support continuous single-cycle MAC (multiply/accumulate) operations. They may employ specialized registers to minimize memory accesses and to enable zero-overhead looping.

The processing sweet spot for DSPs is systems that process continuous and sustained streams of data, especially if there is extensive computational processing on that data stream. The upper limit of this sweet spot overlaps with FPGAs, except for highly computationally intensive algorithms that include a high level of state or decision complexity. VOIP (voice over Internet Protocol) is an algorithm that would work better on a DSP than as an FPGA-only implementation because the domain-specific knowledge for compressing data involves many control states and decisions. Software development with DSPs usually requires an understanding of signal-processing algorithms. Many application- and domain-specific analysis and development tools and libraries are available to assist developers with their DSP applications.

DSP suppliers can differentiate their devices by including application-specific hardware accelerators or by implementing a SIMD (single-instruction/multiple-data) or VLIW (very-long-instruction-word) architecture that supports the simultaneous operation of multiple execution units. Contemporary DSPs use orthogonal instruction sets so that software developers can use compiled C, rather than assembly, to program most of the code. Programming signal-processing code requires different experience and skills from those for developing control or application code. An important differentiator for DSPs is access to signal-processing libraries that either help jump-start a project or enable an application- or control-software developer to use the DSP as a coprocessor without understanding how to program those algorithms.

DSPs commonly find use in embedded-multicore designs; designers often pair them with a complementary microcontroller or microprocessor for control- and application-level support. They also often combine them with hardware accelerators or FPGAs to perform heavy lifting for intensive computational processing. Another common multicore configuration involves a cluster or array of DSP cores that work together on multiple streams of data or wide streams of data, such as for image processing. The arrays of DSPs can involve dozens of cores that connect at the board level. A lot of experimentation with multicore-DSP configurations has occurred over the years. The commercial failure—versus obsolescence—of many of those attempts serves as a reminder that chip-level multicore offerings have to balance between technological advantage and serving a set of applications that generate enough volume to support the device.

DSCs are new to this taxonomy; most of the market in 2007 formally acknowledged them as separate types of processor. Before then, market participants usually referred to them as "hybrid," or "unified," processors because they combined features of DSPs and microcontrollers in a single processor. The processing sweet spot for DSCs focuses on applications that must handle a fair amount of not only context switching, but also signal processing. These devices can alleviate the need to use two heterogeneous cores—a microcontroller and a DSP, for example—in a design. This approach can lower the system BOM (bill-of-materials) cost, and it can lower the development-tools cost because software developers can use the same development tools to program the system. The software developer still must understand the details of programming control or signal-processing algorithms, however.

Because DSCs are relatively new, designers do not have a rich base of experience in implementing them in multicore designs. An early motivation for developing DSCs was to replace multicore designs involving a microcontroller and a DSP with a single device executing a single instruction thread that included both signal-processing and control code. Dual DSCs are now available as single chips. One rationale for using multiple DSCs in a design is that you could normally use one core as a control processor and the other as a signal processor. When one function heavily outweighs the other function, however, the system can more easily perform load balancing between the cores because all of the cores

can support both types of processing.

FPGAs provide a programmable platform that can leverage arbitrarily wide signal-processing algorithms acting as hardware-acceleration blocks. This feature gives them an advantage over DSPs when the signal-processing algorithm is sufficiently wide enough that it can efficiently use more than the available processing units in the DSP. FPGA signal-processing blocks work well when they involve few decision states and large amounts of processing per data point. As the number of decision states increases in an algorithm, it can make sense to place one or more processor cores in or next to the FPGA fabric. FPGA signal-processing implementations function well as coprocessors or accelerators next to a DSP, a microprocessor, or both.

Another processing option is to employ specialized hardware blocks or processors that accelerate software functions. These blocks or processors are usually not stand-alone, instead typically connecting to another processor, such as a DSP or a microprocessor. Some processor vendors integrate these specialized hardware blocks as coprocessors into the same device and even into the instruction-pipeline flow to differentiate their devices. Besides the quick and energy-efficient performance of some tasks, hardware accelerators help offload computational load from the main processor, thus freeing that processing capacity for some other differentiating function. This taxonomy currently lumps network, graphic, and video processors in with the hardware-accelerator and coprocessor category. These types of accelerators are usually not stand-alone systems, and, when they are stand-alone, they exist in a multicore structure, and the proposed multicore categories encompass them. Please contribute to the companion blog post for this article at www.edn.com/blog/1890000189.html to expand on these architectures or add any others.

Microprocessors round out the taxonomy. They employ general-purpose architectures that enable them to perform well enough across the range of processing tasks. They do not handle context switching as quickly or deterministically as microcontrollers do because they usually employ software within the interrupt handling. They do not handle looped processing as quickly or efficiently as DSPs do because they usually neither include zero-overhead looping nor employ the bus and memory structures to keep a single-cycle MAC busy every cycle. Microprocessors are ideal when the processing requirements that the system must support are not well-known, such as when the system supports user-loaded applications.

Microprocessors generally support large memory-address spaces and rely on large on-chip caches to compensate for time penalties from off-chip memory accesses. They can execute complex operating systems and support a range of legacy code. Microprocessors are also appropriate for "quick and dirty" prototyping and proof-of-concept exploration when cost and energy efficiency are less important than a short development cycle.

The sweet spot for microprocessors is to support systems that exhibit significant uncertainty in processing behavior and loads, such as when executing disparate processing threads over the same processor resources. A microprocessor's flexibility and operating-system abstractions of the hardware peripherals make them ideal for the development of sophisticated user interfaces and high-level application code. Microprocessors can perform most functions—at higher cost and energy consumption—that the other specialized processor architectures can perform. Using a microprocessor for control or signal-processing tasks can make sense if the design needs a microprocessor anyway and there is sufficient head room to accommodate the processing; otherwise, for embedded systems, designers may migrate well-defined tasks to the appropriate specialized processor to save cost and energy consumption.

Multicore-microprocessor designs include desktops and servers. As cell phones evolve to include sophisticated graphics and user interfaces, microprocessors and graphics accelerators are supplanting the microcontrollers that sit next to the DSP that handles the signal processing in these devices. In all of these cases, each processor architecture has use cases for single-core and multicore designs, and, in many cases, the multicore-use cases involve combining processor architectures so that they complement each other.

MULTICORE

This taxonomy highlights the major single-core architectures and the distinct development ecosystems that exist for each type of processing architecture. Each ecosystem includes its own silicon offerings, domain-specific analysis and development tools, and libraries that abstract some of the complexity for those developers who specialize in the other types of processing. Perhaps looking at multicore offerings in a similar fashion will help our industry address a problem that still seems intractable. A number of companies have come and gone that offered their vision of a multicore architecture. Were they off-base, or were they just missing a single component that could have made all the difference for success?

Multicore terminology is not standard; this article proposes a couple of multicore-configuration names. The terminology instead attempts to propose functional-level descriptions for multicore architectures. A developer should be able to incorporate one or more of these types of multicore architectures in the same design, much as many embedded designs incorporate more than one type of processor architecture.

Channel-multicore architectures target the parallel nature of some applications, such as infrastructure-networking equipment, which apply the same type of processing across many input sources at once. The processing for each input is relatively independent of the processing of the other inputs. These systems implement multiple copies of the processing engine so that each copy can apply the same processing across multiple

channels of data. The system contains duplicate processing resources so that they are local to each processing engine; this approach provides lower latency and better processing performance than do architectures that access the same resources through a shared-access mechanism. The nature of the processing engine is application-specific, and it is not limited to a single function. Software development for these systems focuses on optimizing the domain-specific processing of the inputs so that the system minimizes how much it must manage program resources.

Aggregate-multicore architectures target processing tasks, such as video or image processing, that are impractical or less practical to perform in one instruction thread. The processing engine consists of multiple homogeneous or heterogeneous processing elements that perform separate parts of the same overall task. The architecture aggregates the results from the separate processing elements in some fashion to complete the processing. This type of multicore system has existed for decades as high-end signal processing or simulation systems. These types of systems are of increasing importance for high-performance processing as processor clock rates have stalled and are no longer racing to higher rates. Much research is going into how tools can help developers extract parallelism in their software. No widely used tools are currently available that automate parallelizing general code. Some domain-specific analysis tools for signal-processing-algorithm exploration assist domain experts in determining how to structure their algorithms, but they do not completely automate the parallelization of algorithms and software to take advantage of these types of systems.

Multidomain multicore architectures target those applications that encompass multiple software domains, such as cell phones, automobiles, and many consumer products. This area is the bread and butter of many embedded-multicore designs that developers have been doing for decades. Some coupling and communication may occur between the cores in the system. Each processing element in the system not only makes optimization choices that align with the function it performs, but also works in tandem with and complements the other processing elements in the system. These types of systems have existed as proprietary board-level designs. Single-chip, multidomain multicore devices are available as vertically targeted devices for applications such as mobile devices with rich video support. Software developers are aware of each type of processing engine they are targeting, and they use the appropriate set of analysis and development tools to perform their tasks. There are challenges in debugging these types of systems because each processor has different on-chip resources and may not provide visibility to all of those resources in the same way from chip to chip.

⊞ For related blog posts about embedded processing, go to www.edn.com/blog/1890000189.html.

⊞ For a related article about processing options, go to www.edn.com/article/CA6298267.

Feedback-multicore architectures target those applications, such as autonomous systems, that employ feedback within and between subsystems or different domains of the system to perform their functions. These systems may employ a central command unit or aggregate decisions across distributed decision makers. The feedback loop may be as simple as an actuator sense-and-control loop, or it may involve larger-scope feedback that can assist the system with trending, predicting, planning, and learning. These emerging designs currently exist mostly in the academic, aerospace, and military domains, but they are set to spill into the industrial and consumer areas. A candidate industrial system is a smart building that manages systemwide resources across distributed decision makers. Consumer candidates include automated active functions in automobiles as well as household appliances, such as vacuum cleaners. Software developers need an integrated and correlated aggregate of analysis and development tools that cover filters, signal processing, control algorithms, and closed-loop simulations.

CONVERGENCE

This taxonomy for single-core architectures has successfully enabled the industry to converge the vocabulary and key requirements for each application domain between the relevant chip architects, board-level-system designers, software developers, and development-tool providers. Each processor type has its own ecosystem with its own set of vendors for chip-, board-, and software-level-design support. Even for those hardware and software companies that offer products and services across multiple types of processor, different teams support each processor market.

The multicore-system market could benefit from a convergence of vocabulary and requirements that employs functional capabilities rather than raw implementation and performance details. Rather than a single general-purpose discussion on SMP versus AMP, there should be at least four and at least one for each of the multicore categories. You could say the same for core topologies, coherency mechanisms, messaging, multithreading, operating-system support, virtualization, hypervisors, code parallelizing, and balancing workloads. By defining and establishing the optimization preferences employing required functional capabilities, chip architects, board-level-system designers, and software developers can meaningfully discuss how to allocate technical approaches among all of the stakeholders rather than an ad hoc fitting together of what each group builds. **EDN**

FOR MORE INFORMATION

Altera
www.altera.com
AMCC
www.amcc.com
ARM
www.arm.com
Atmel
www.atmel.com
Freescale
www.freescale.com
Intel
www.intel.com
Microchip
www.microchip.com
Microsoft
www.microsoft.com
MIPS
www.mips.com
National Instruments
www.ni.com
QNX
www.qnx.com
Radisys
www.radisys.com
Tensilica
www.tensilica.com
Texas Instruments
www.ti.com
Xilinx
www.xilinx.com

Maximize Multicore

Increase Your Performance with NI LabVIEW

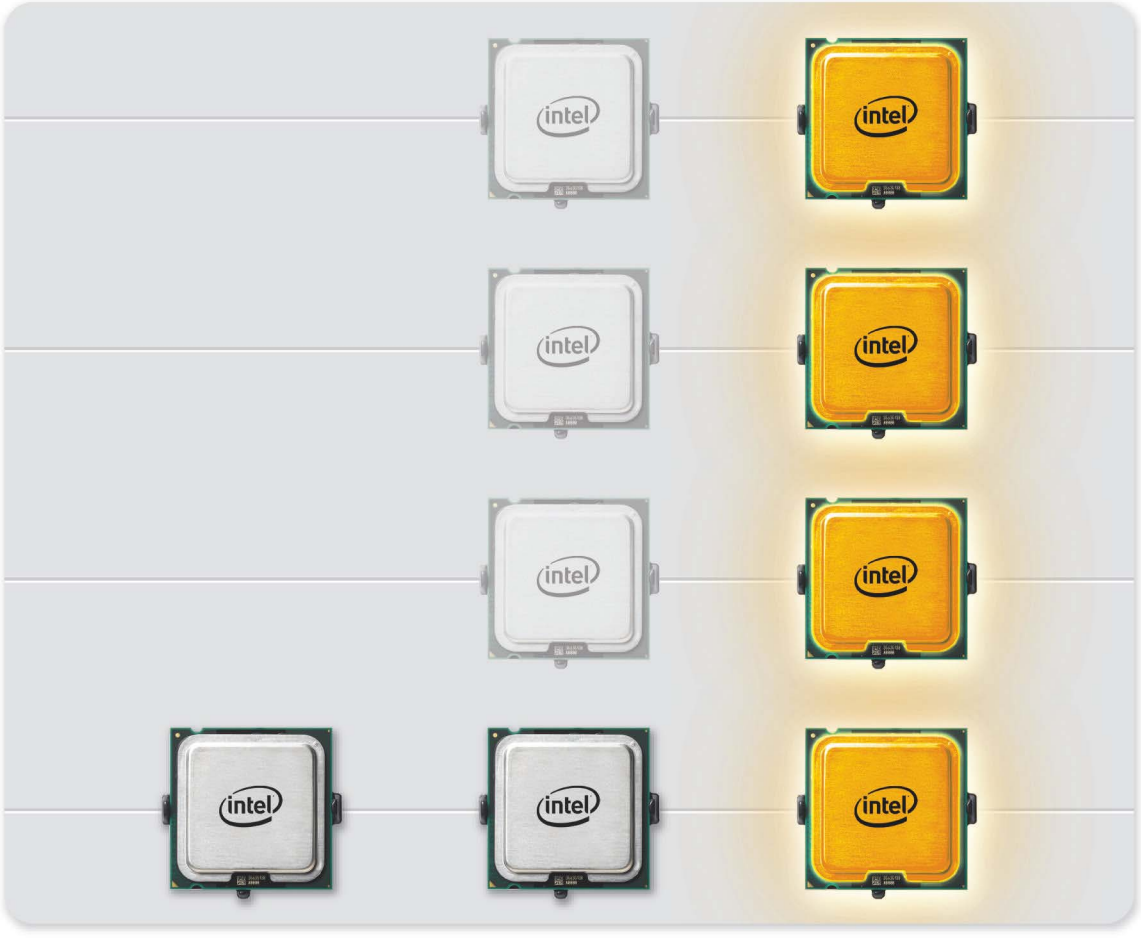

Single-Core Application | Sequential C Program on Multicore | **LabVIEW on Multicore**

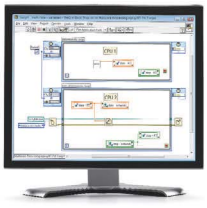

Using LabVIEW — an inherently parallel graphical programming language — you can take advantage of the latest advances in multicore systems, including real-time symmetric multiprocessing and FPGA-based embedded systems. LabVIEW simplifies visualizing and programming parallel applications for test, control, and embedded design. Using LabVIEW, customers are realizing a near-linear performance increase on their multicore systems.

>> Learn to maximize your multicore processor at **ni.com/multicore** **866 337 5041**

NATIONAL INSTRUMENTS

BY ANDREA KROLL, PhD • JEDA TECHNOLOGIES

Troubleshooting a transaction-level model

SEEMINGLY MINOR VIOLATIONS OF THE TLM 2.0 STANDARD CAN TURN A SYSTEM-LEVEL MODEL INTO AN AGENT OF EVIL IN YOUR DESIGN FLOW.

The goal of OSCI (Open SystemC Initiative) TLM (transaction-level modeling) 2.0 is to enable high-level component models to simply plug into and play with each other in a system model. The standard identifies modeling styles, several interfaces, a generic payload, and more than 150 rules to define the expected behavior in the simulation. For the system model to work correctly, designers must comply with the standard. Finding errors in compliance can be a major problem, however. Nevertheless, approaches exist for troubleshooting system models to find noncompliant areas.

TLM 2.0's 150 rules specify the expected behavior of a transaction and restrictions to avoid system malfunctions. The OSCI TLM working group defined TLM 2.0 specifications with respect to initiator and target pairs. A typical design has several initiators and targets, which connect through interconnects (**Figure 1**). The interconnect serves as both a target and an initiator.

The communication path is the path between a communication's starting point and its ending point in a system. A communication path typically contains more than one TLM 2.0 initiator/target pair. In a simple example, there are always two initiator/target pairs in the full communication path. Real systems contain interconnect hierarchies; multiple interconnects that connect bridges; and multiple subsystems, multiple cache hierarchies, or both.

Most of the TLM 2.0 rules apply to each initiator/target pair, not the full communication path. This requirement ensures interoperability between any two components in a system. It is easy to validate that models obey these rules at the starting point of a communication. However, it is more difficult to validate that the models obey the rules in the system because this type of validation forces designers to observe multiple communication paths in parallel. If the models obey all the rules, TLM 2.0 guarantees the interoperability and, thus, the communication between components. The following example illustrates what happens when a model violates a rule.

NONCOMPLIANT COMPONENTS

Violations of rules can create significant interoperability problems in a system. Interoperability issues are often not immediately visible. They might result in system-behavior malfunctions later in a simulation. The problems often occur in a different location in the system or trigger an unexpected corner-case situation. So interoperability issues require constant monitoring of the communication path between the initiator/target pairs and a comparison of the actual with the expected behavior.

One such violation is a generic-payload-access violation. Several rules limit access to the generic payload to specific components in a model. For efficiency reasons, the generic-payload objects do not go through the high-level system model. Rather, the model uses references from each component in the communication path to one copy of the generic payload to mimic the transmission of data through the system. This technique also keeps the communication between components flexible.

With this reference mechanism, the designer can explore multiple sequential and parallel transmission schemes without changing the model. The disadvantage is that all components in a communication path have access to all the information in the generic payload. Implementing a mechanism in the generic payload that identifies which component accesses which data field is too

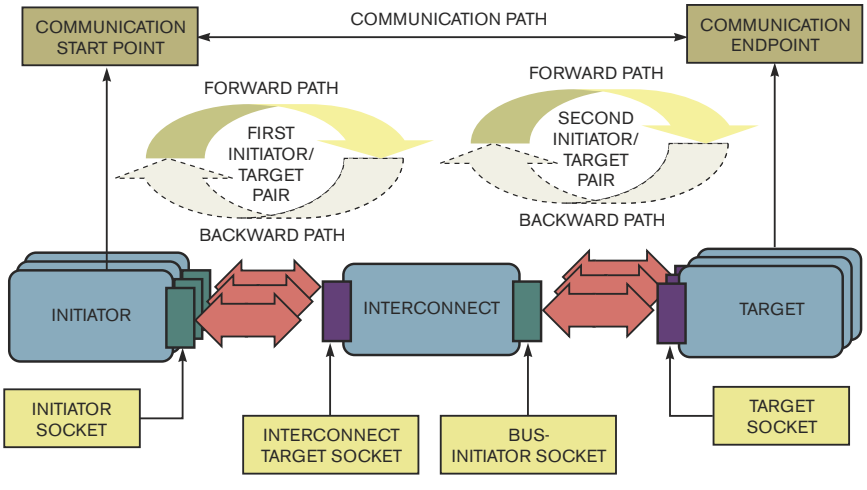

Figure 1 TLM 2.0 provides a strict framework to model the interactions between blocks in a high-level system model.

expensive in simulation performance. Still, some access rules are necessary to guarantee interoperability and avoid a malfunctioning system.

For example, one rule specifies that only the interconnect models can modify address values. Doing so easily mimics virtual-to-real-address translation in a system. In many cases, the interconnect uses address fields in an address-translation table. For that reason, another rule prohibits targets from modifying the address attribute of the generic payload. A target that modifies the address field causes the interconnect's address-translation table to fail. The model either cannot complete the response or sends the response to the wrong initiator.

A second example is that only the targets can enable the attribute DMI (direct-memory interface). So another rule states that the initiator must initialize the DMI-allowed attribute with the value "false." Only targets can change this attribute. Only the target knows whether it can support the DMI. An initiator setting "DMI enable" to "true" with a target that does not support DMI creates side effects in system behavior that are difficult to observe and to debug.

PROTOCOL-STATUS VIOLATIONS

OSCI TLM 2.0 provides a set of rules specifying the usage and the transition of protocol-status information. These rules capture the expected behavior of the protocol. In this way, every user has the same view of the states and state transitions the model allows. TLM 2.0 also restricts the number of allowed "in-flight" transactions between an initiator/target pair before the initiator can issue the next transaction. This restriction protects the targets from memory or register overwrites. The rules do not apply to the full communication path; they apply only to each initiator/target pair. In this way, designers can easily model features such as out-of-order execution and parallel-thread processing without making assumptions about the communication partner.

TLM 2.0 specifies two-phase communication for LT (loosely timed) and four-phase communication for AT (approximately timed) models. It has some restricted ability to pre-empt a communication early by sending a TLM-complete return value. TLM 2.0 predefines the order of the phases and explicitly specifies the allowed transitions between phases. In this way, every component has the same view of all possible states and need not be able to handle any other state transition. This requirement allows the user to implement efficiently tuned models for high-speed simulation. Unfortunately, it also means that malfunctions in models can go undetected.

OSCI TLM 2.0 specifies that you can issue only one transaction with a begin-request command to one target. The target must first respond with an end-request command before the next begin-request command can occur. This rule protects the target from information overwrite by ensuring that only one transaction can be in flight. The target must actively issue an end request before the initiator can start the next transaction.

FOR A USER OUTSIDE THE MODEL-DEVELOPMENT GROUP, THE IMPACT OF INTEROPERABILITY ISSUES INCREASES TO THREE TO FIVE DAYS OR EVEN TO WEEKS.

Ensuring these interoperability restrictions from the starting point of a communication seems trivial, but, because TLM 2.0 models share interconnect systems, multiple transactions can be in flight in the interconnect model. The interconnect model must now ensure that all of its connections of initiator sockets to targets obey the rule. For some systems, this requirement means that 10 to 100 communication paths work in parallel, so 10 to 100 transactions can be in flight at once.

If an initiator accidentally issues more than one begin-request command to one target, depending on the implementation of the target, the second begin-request command will overwrite information from the first transaction before it can process the command. This scenario breaks the communication of the first transaction. The target would respond with only one end-request command, instead of the required two. The initiator of the first transaction waits for the end-request command before it can send the next begin-request command, so a deadlock arises.

When an initiator issues a phase multiple times, so that it issues an additional transport call, the sequence can push the traffic generator out of synchronization. Especially for protocol-status violations, the problem is typically not visible immediately in the simulation but later causes malfunctions in the system in the simulation. Tracing the malfunction back to the original protocol violation is time-consuming.

In general, all of these examples show that rule violations cause interoperability problems. These problems, in turn, create major system-behavior failures. Those system failures may be immediately visible in the form of system crashes, obvious wrong behavior, or deadlock situations. Some others might be visible only to system experts who could spot inconsistencies with the expected behavior or performance. Those interoperability issues are so difficult to find, but designers can use debugging and troubleshooting approaches if they cannot automatically check a model for TLM 2.0 compliance.

FINDING COMPLIANCE VIOLATIONS

All major electronic-system houses and semiconductor companies that are implementing ESL (electronic-system-level) models have an average project delay of three to five hours per issue due to interoperability problems in their models. They confirm that this number is conservative and applies to a highly skilled model developer. The assumption is that the developer can trace the issue to a component in a system or has found it during unit testing.

For a user outside the model-development group, the impact of interoperability issues increases to three to five days or even to weeks. The further the user is from the model developer, the larger the time increase is. For a typical project, these delays lead to 40 to 60% of project time for debugging. To understand why it takes so long to identify interoperability problems, you must understand debugging and identify the steps

necessary for troubleshooting interoperability problems.

Debugging ESL models employs a multilayer approach involving the application-, software-, component-, and interface-debugging layers. This approach is one way that virtual-platform developers cope with system complexity. Typical virtual platforms contain at least 20 to 40 hardware- and software-component models, all with their own debugging infrastructure. To be able to quickly identify problems, a user traces key parameters in the system. If those key pieces of information are inconsistent with expectations, something is wrong, and debugging starts.

IDENTIFYING AN INTERFACE-SYNCHRONIZATION PROBLEM REQUIRES DETAILED INTERACTIVE SOURCE-CODE DEBUGGING.

At the application-debugging level, the user looks at application-specific data. This data might be a sequence of video frames, an audio stream, or a wireless trace. This approach allows narrowing down the area in which a problem occurs.

For example, in debugging a video frame, you may look at several video frames. You might identify that some frames are intact, but several are broken. Some lines in a frame are still intact, but the picture fragments shift randomly throughout the frame. This problem is a synchronization issue with the display of the frame. The problem can occur within one frame or between frames. Lines of a picture are recognizable, and some picture segments are intact. This pattern means that the problem probably occurs after decoding and before displaying the frame.

After identifying a frame-synchronization issue, the next step is to look at the software. One of the processors is responsible for synchronizing frame processing. The next step is to identify the piece of software code that is responsible for triggering frame processing. At the software-debugging level, the designer might use software debuggers such as GDB (Gnu debugger). The user sets breakpoints in the source code for frame processing and observes values at register or memory locations. In most debuggers, the user can set "watch points" to specific values. A watch point on the register value triggers the next frame's processing and stops the simulation at the frame-synchronization point.

These advanced tools help to filter out other behavior and focus on frame synchronization. It still takes time to identify the source code and set the breakpoints and watch points in the correct location and then identify the register values responsible for triggering the frame processing in the hardware accelerators.

In this example, frame synchronization seems to work. The user has set the register values, and frame processing in the hardware component starts. You have thus determined that the software application did not cause the problem and can move to the next debugging step by looking into the hardware models.

Multiple hardware components, including those for encoding, decoding, and frame-buffer processing, are responsible for video processing. Most of the components move pixel information into and out of various local and global memories in parallel. Multiple components share some memories.

To identify the location of problems, developers use monitor components to observe data arrays in memory. In addition, they trace read and write values from processing elements to memory and vice versa. Most ESL-debugging tools are SystemC- and TLM-2.0-aware. This feature allows the tools to automatically and conveniently display transactions. The ability to set breakpoints and watch points on TLM 2.0 transactions significantly eases debugging. The user still must know the expected values of the transactions throughout the platform, however, to be able to pinpoint issues in data communication.

The user traces transactions from the point at which the data stream is in synchronization with the location in the system where it gets out of synchronization. This example identifies that the transactions are correct until they reach the frame buffer.

The example identifies that the start addresses of the frames are randomly distributed when the frame processing starts. You can expect multiple start addresses because multiple frames are in the buffer. After consulting the documentation, you can identify the expected values for the start addresses of frames, but they won't match what you saw during debugging.

Now, trace back to where the address is written. The generated addresses are correct. Thus, you can pinpoint the frame buffer as the location of the problem. More accurately, the communication between the processor and the frame buffer is causing the problem. You thus should more closely examine the TLM 2.0 socket interface of those components.

The interface-debugging step is the most detailed one. It requires detailed knowledge about both the synchronization mechanism in the system and its implementation in the model. At the transaction level, multiple ways exist to mimic the behavior of one interface protocol. TLM 2.0 significantly reduces the number of ways of implementing an interface protocol, however. Nevertheless, identifying an interface-synchronization problem requires detailed interactive source-code debugging. This task is especially cumbersome if you have incorrectly implemented the TLM 2.0 interface.

In the case of video, the memory is a shared resource that connects through an interconnect to the processor. Multiple applications can concurrently transmit traffic to this interconnect, meaning, in some cases, hundreds of transactions from various applications are in flight at the same thing.

You must filter out only those transactions that are relevant for your video application. You also must predict and confirm the expected values as well as the expected time each value should occur. For TLM 2.0, this requirement means that, for each transmitted value, you must trace two or four transaction calls for LT and AT, respectively, to ensure that they transmit the correct value

⊞ Go to www.edn.com/ms4321 and click on Feedback Loop to post a comment on this article.

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

at the correct time. It may take hours of staring at transaction traces and comparing expected traces with the simulated traces.

INTEROPERABILITY ISSUES

Eventually, an interoperability problem arises. For example, the processor did not respect the TLM 2.0 rule in some corner-case situations. It updated the register, which triggered the hardware accelerator to start processing the next frame for display, without waiting for the end-request command for the previous frame. The assumption in the system is that data processing in hardware is faster than in software. The model does not implement detailed interrupt processing, which is also not relevant for the initial analysis. So, instead of modeling the full interrupt mechanism, the model uses an end-request command to mimic the interrupt to indicate the end of processing in the display accelerator. This action notifies the processor that it can send the next transaction to trigger processing for the next frame.

The process to wait for the end-request command and send the next transaction is broken for some corner-case situations. The frame view shows that multiple frames had started, but, because the processor triggered another frame too early, it finished transmitting only parts of each frame before the next frame started.

In a project that is supposed to quickly create some initial analysis results, debugging can take several days, so the results may not be available in time. Automatic compliance checking would have instantly detected this problem and made the results available in a reasonable amount of time.

Although the effect of TLM 2.0 protocol violations in high-level models may seem minimal in passing, failing to catch such a problem has deeper implications. Erroneous models with TLM 2.0 violations alter the behavior of the system and cause system malfunctions or biased performance results. Undetected errors can propagate into the RTL (register-transfer level). Especially in the case of biased performance results and corner-case user scenarios, it is nearly impossible to detect these problems at RTL because it is too detailed and simulations are too slow to run the system-level tests. You might be able to identify the problems during FPGA testing, but it is too late in the development cycle and causes project delays. In some cases, you might not detect the problem at all, and that problem would result in misbehaving silicon.EDN

AUTHOR'S BIOGRAPHY

Andrea Kroll is a vice president of marketing and business development at Jeda Technologies, where she has worked since early 2008. Previously, Kroll worked at Synopsys as an application specialist for CoCentic System Studio and CoWare Processor Designer. In 2005, she took over product marketing for CoWare Processor Designer. By feeding back technical requirements from customers to engineering, she helped double the company's processor-designer business year on year for three years. Kroll has a doctorate in electrical engineering with a concentration in high-level-model validation from the University of Technology (Aachen, Germany).

120 Watts

1/32 Brick Footprint

48 V : 12 V Bus Converter for Server, ATE, Telecom, and Industrial Control

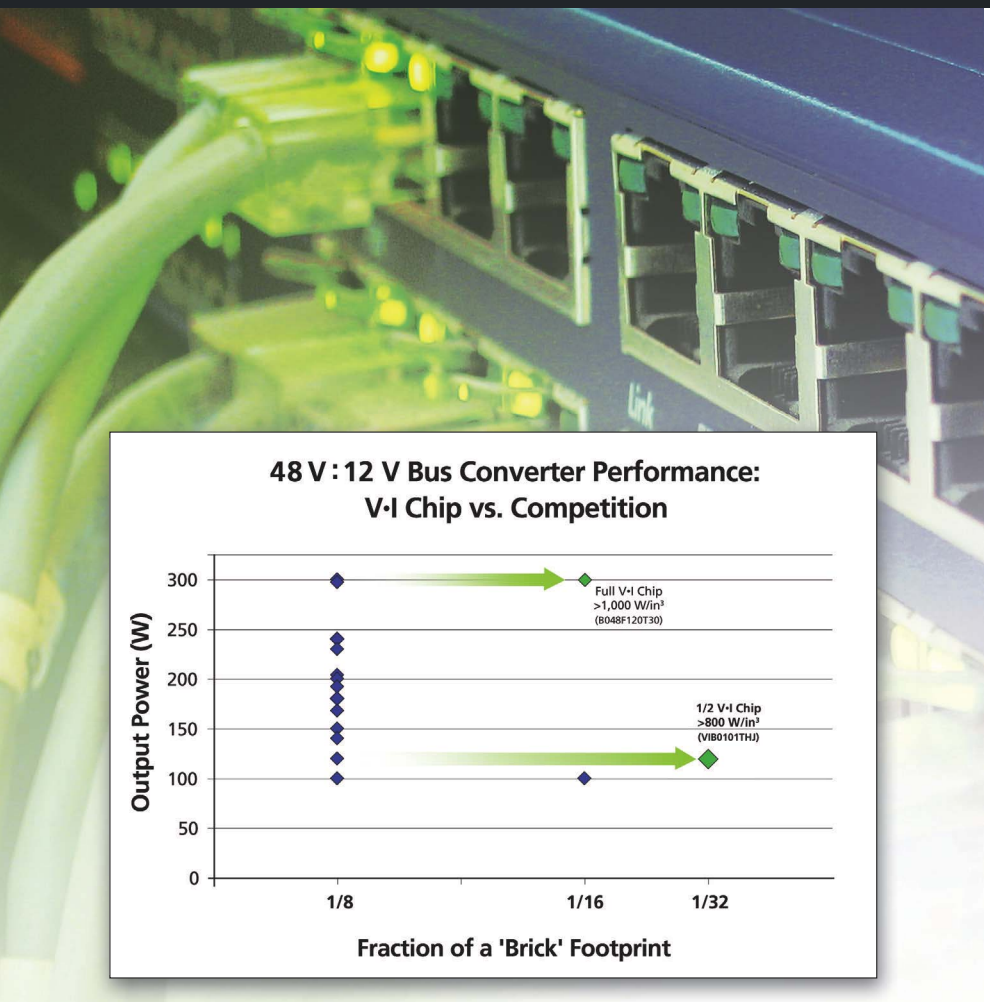

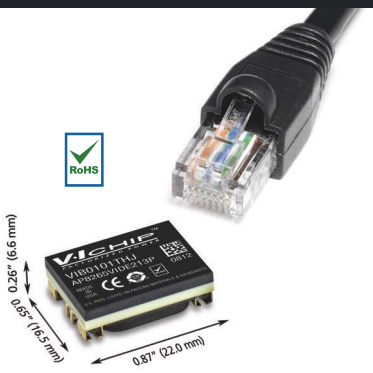

V·I CHIP
FACTORIZED POWER

For designers working IBA applications in server, ATE, telecom, industrial control, etc., VIB0101THJ offers layout flexibility, reduced power converter footprint, and easy design for fast time-to-market.

- **4x smaller than previous best bus converters**
- **Simple design for fast time-to-market**
- **High efficiency**
- **Powers standard 12 V POL converters**
- **16x capacitance reduction feature**

BCM Model Number	Package	V_{IN} Range (V)	V_{OUT} Range (V)	Output Power (W)	Efficiency (%)
VIB0101THJ	Half	38 – 55 Vdc	12.0	120 W	94.5
B048F120T30	Full	38 – 55 Vdc	12.0	300 W	95.1
B048F096T24	Full	38 – 55 Vdc	12.0	240 W	96.2

Call 1-800-735-6200
(US & Canada only) to order.

vicorpower.com/vichip

Products in mass production and available from stock

800-735-6200 **vicorpower.com/bcm-half/edn** VICOR

16V, High-Efficiency 2A & 4A Synchronous Buck Converters

4A Continuous Load in a 3mm x 4mm Footprint

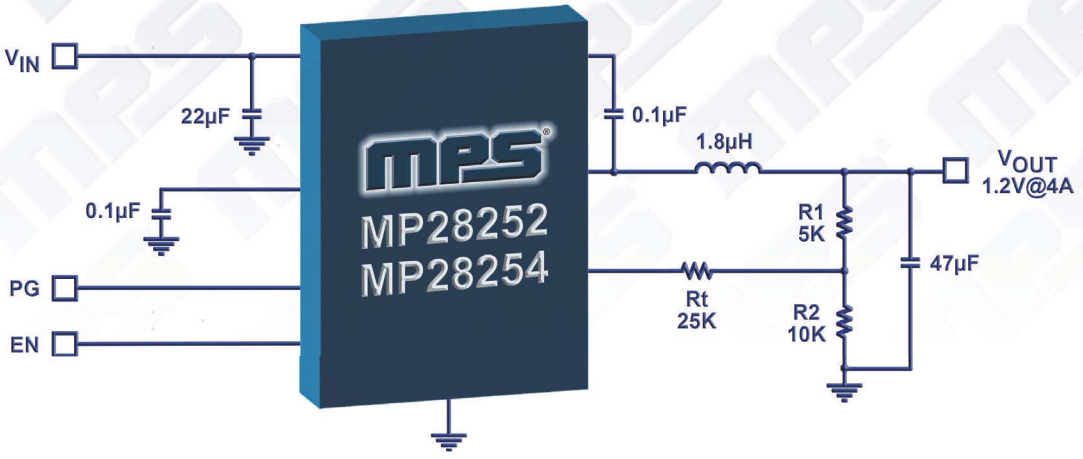

MP28252/54 Features

- *Wide 4.5V to 16V Operating Input Range*
- *Load Current Options:*
  - *MP28252: 2A Continuous*
  - *MP28254: 4A Continuous*
- *Power Good Indicator (PG)*
- *Low Rds(on) Internal Power MOSFETs*
- *Proprietary Switching Loss Reduction Technique*
- *Fixed 500kHz Switching Frequency*
- *Sync from 300kHz to 2MHz External Clock*
- *Internal Compensation and 4ms Soft-Start*
- *Integrated Bootstrap Diode*
- *Available in a 3mm x 4mm 14-Pin QFN Package*

Efficiency vs. Output Current

V_{IN} = 12V, V_{OUT} = 1.2V

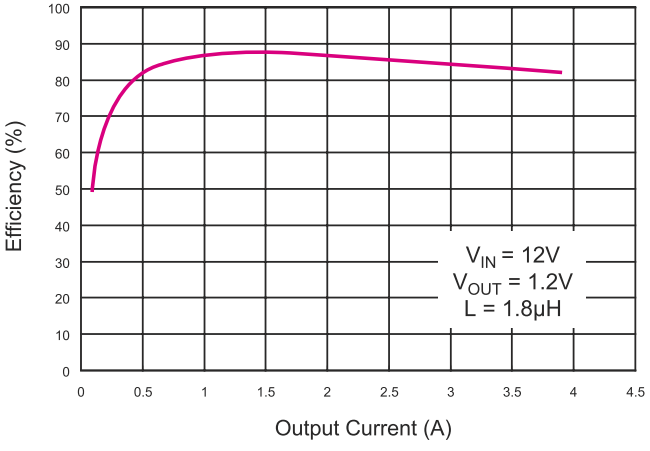

DC to DC Converters CCFL / LED Drivers Class D Audio Amplifiers Linear ICs

ISO9001:2001 Certified

LEAD FREE RoHS COMPLIANT

Monolithic Power Systems®
The Future of Analog IC Technology®
www.monolithicpower.com

AVNET Memec

BY JOSEPH PUN • GENNUM CORP, SNOWBUSH IP GROUP

Designing portability into silicon IP

DESIGN FOUNDRY PORTABILITY INTO IP CORES, RATHER THAN APPLYING IT AFTER THE FACT.

Designing IP (intellectual-property) cores that you can port to various foundry protocols or process geometries is paramount in today's highly diverse standards-based-IP technologies. This technological advantage allows IP-core vendors to gain access to new customers and markets, enabling chip designers to be first to market with new products that ultimately consume less power and area. It is also imperative for IP-core vendors to offer products that can target multiple industry-standard protocols, such as PCIe (peripheral-component-interconnect express), SATA (serial-advanced-technology attachment), and GbE (gigabit Ethernet), which are the building blocks for switches, hard drives, chip-to-chip communications, optical communications, and networking.

For successful SERDES (serializer/deserializer)-IP development, you must design SERDES cores that you can easily and quickly port from foundry to foundry and into the latest and smallest-geometry processes. This approach ensures a comprehensive portfolio of IP, which results in a wide pool of potential deployments in customer applications. Customers all have their own reasons for wanting to go with one foundry and process technology rather than another. Going with the latest process technology can afford lower power dissipation and increases the number of transistors customers can fit onto a die, decreasing costs and increasing profit margins.

To design portable SERDES IP, analog- and digital-system designers must build a range of programmability into their circuits to make them robust over a range of process variations, and they must be able to rapidly verify circuits using various I/O-supply voltages. As with all development efforts, trade-offs arise between designing portability into IP and optimizing IP for a process. Designing SERDES IP with wide programmability allows for porting efficiencies and costs IP developers increased footprint for the IP on the chip.

Ideally, a design team can rapidly port a design without much redesign or modification of the physical design or layout. Verification, however, is always necessary because each foundry has its own set of model libraries, and designers must place and route all of the digital synthesized logic using each foundry's specifications. Even the same process technology can vary widely in performance among foundries. For example, one foundry's 65-nm, general-purpose process can differ from another foundry's 65-nm, general-purpose process in process and device parameters, as well as the typical variances in physical geometries and spacing rules. Therefore, efficient development of IP cores always takes into account minimizing effort and time to address these differences.

ANALOG DESIGN AND PROGRAMMABILITY

Designing analog circuits with appropriate programmability in mind can vastly reduce the effort of porting from foundry to foundry or porting into new process technologies within the same foundry. SERDES IP typically includes a CDR (clock- and data-recovery) block in the receiving path, as well as a transmitter PLL (phase-locked loop), which generates the transmit high-speed clock. The CDR is responsible for recovering a clock from a serial binary-data stream that samples the high-speed data. The data and clock then go to a serial-to-parallel converter to deserialize the serial stream into bytes and words, which the converter then sends on to the core digital logic. The transmitter PLL is responsible for multiplying a reference-clock frequency into a high-speed bit clock that a parallel-to-serial converter then uses to serialize and transmit data words. The CDR and transmitter PLL contain VCOs (voltage-controlled oscillators) that provide a range of output frequencies. The chip designer tunes this circuit for the specified process-technology port.

Typically, digital control of the voltage inputs into the VCO tunes the circuit. For example, if a process technology operates at a slower rate than the chosen benchmark specifies, design-

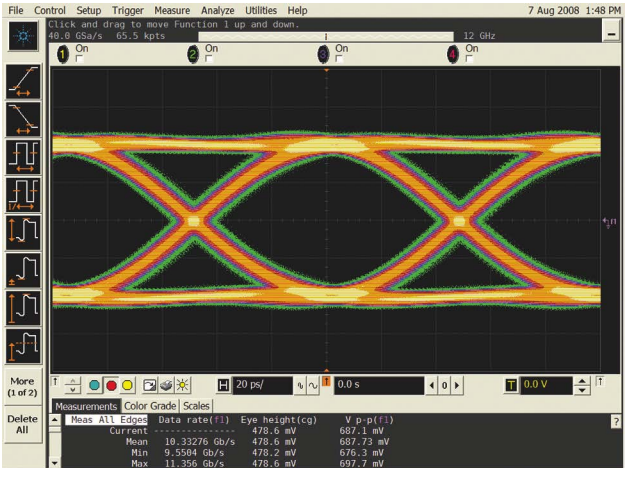

Figure 1 An open-eye diagram is one indication that a port to a new process has been successful.

ers would tune the VCOs to a higher setting to provide the desired frequency of operation. This programmability allows the IP vendor to design and lay out CDRs and transmitter PLLs in a way that provides enough tuning range to achieve the correct output frequency for various process technologies. The IP designer must determine how much tuning range and resolution to build in to target all desired foundries and process technologies for a design.

Another parameter that typically varies among processes is the sheet resistance of polysilicon resistors. A polysilicon resistor with a given geometry may be more or less resistive than a polysilicon resistor of the same geometry in another process. To overcome these resistance mismatches between process technologies, you may need to allow for programmable resistances. Designers need resistors that they can switch into and out of a circuit. Although not the most area-efficient approach, this feature allows the designer to control the resistance without modifying the layout. An example of this usage is in transmitting/receiving-termination circuits.

When designing SERDES IP to enable quick porting to smaller process geometries, it is important to take note of the I/O-supply voltages that each process technology supports. For instance, some 65-nm processes support 2.5 and 3.3V I/O transistors, whereas others may support only 1.8V devices. To further complicate matters, overdriven, 2.5V devices may have a wider range of operation. IP-circuit designers should design the SERDES to accept a range of I/O voltages and lay out the circuits such that the IP's physical designer can easily swap in different I/O transistors. This step usually involves having the IP's physical designer develop a floorplan with room for the larger I/O devices in the layout and for the IP-verification engineers to verify operation of the circuits with all the devices. With careful design, this step can be as simple as swapping layers in the physical-design software based on the customer's preference. Another possibility is to use the overdriven I/O devices with extended widths, which can support a range of supply voltages larger than 10% that vendors typically specify for I/O-supply variations of SERDES. A designer must be mindful not only of multiple-foundry design and portability into new process geometries, but also of the process-technology options that are available, such as different I/O devices.

DIGITAL CALIBRATION OF ANALOG CIRCUITS

An effective technique for using programmability is automatic digital calibration. In this process, a designer puts a digital-calibration loop around an analog-circuit block. This technique is useful not only for minimizing design changes when porting, but also for extending a SERDES' ability to minimize effects due to process, voltage, and temperature variations. Automatic calibration and a wide enough tuning range on key analog circuits aid in getting better porting results without requiring extensive verification for different process technologies and options. The digital-calibration circuitry should take up minimal area and should become smaller than the analog block as process geometries get smaller. Plus, the benefits can be substantial. The calibration helps compensate for differences in performance among each foundry's technology and among newer technologies. For example, the calibration loop can automatically tune the VCO's voltage inputs so that the VCO oscillates at the correct frequency even if the performance of a new foundry's process technology differs greatly from that of the process currently in use. Programmability and calibration are important tools for ensuring first-pass silicon success when porting to new technologies.

Using this design approach to achieve first-pass silicon success in a port to a different process technology, one company designed a test chip for a new process technology. The company had recently completed a test chip for a foundry's general-purpose process, and the goal was to quickly port to the equivalent geometry's lower-power node. The designers placed a wide tuning range on the VCOs, added programmability to their transmitter/receiver-termination resistors, and designed for multiple-I/O transistor devices. The team had time to verify only the major blocks, such as the master bias, and little time for physical-design changes. The team relied on the use of built-in programmability and calibration.

The design would need few changes to meet the specifications of the lower-power process technology. A few months later, the team received both the lower-power and the general-purpose test chips and began characterizing both. While testing, the team verified the calibrated values of the lower-power and the general-purpose test chips and observed similar performance. The end result was that the lower-power test chip performed as successfully as the general-purpose test chip, with minimal time and effort in porting from the general-purpose to the lower-power chips.

AUTOMATED VERIFICATION

The other missing link of designing portability into IP is the use of specially designed and optimized CAD (computer-aided-design) tools that streamline porting. The goal is to design IP that designers can quickly and easily port, and the CAD tools you use play a big part in achieving this goal. You can devise a customized verification-suite setup for rapid verification. Such a suite would comprise custom PERL (Practical Extraction and Report Language) scripts and Microsoft (www.microsoft.com) Office macros that automatically generate testbenches, run simulations, and extract formatted results. In such a setup, the highly automated verification process relieves designers of the task of extracting data so that they can focus on designing the circuits and analyzing the pertinent results. An automated verification suite that all the designers use also makes it easier to share data and for designers to easily and quickly begin new projects without long learning curves. Some companies also offer specification templates for each major analog-circuit block, and these templates allow quick, repeatable, and consistent verification and produce formatted results for easy review.

PHYSICAL DESIGN FOR PORTING

Supporting a range of I/O-device options inevitably requires allocating a large enough area to support the largest devices a user might drop into the circuit. Similar ideas apply to supporting multiple foundries; the layout must work for different sets of design rules. It may not always be feasible to target designs for all foundries and processes of interest. You may have to take a look at the design rules of each foundry and make a

layout compatible with only a subset of those rules. Supporting too large a set of foundries inevitably increases the size of your layout, resulting in wasted area.

Minimum device geometries and spacing requirements for each foundry can vary widely, so having a custom set of design rules that meets requirements for multiple foundries is a good way to do the physical design. When taping out designs to customers, however, you must ensure that you use the target foundry's design rules rather than your custom set. The most time-consuming form of physical design is manual analog layout, so, for efficient design porting, minimize custom design in favor of digital-synthesized logic that you can place and route. Although some tools can partially automate analog physical design, you should avoid their use for high-speed SERDES applications. IP designers should investigate any opportunity to replace a traditional analog-circuit block with a digital implementation. This extensive use of digital circuits makes the design easier to port to new geometries and allows the design to more easily reap the benefits of the area savings than would an analog implementation. The design must still meet all required specifications, including area, power, performance, and eye-diagram templates (**Figure 1**).

Providing IP solutions in multiple foundries and in multiple process-technology nodes allows IP providers to widen their customer base. To target as many technologies as possible, the IP designer must optimize design, verification, and physical design with porting in mind. This approach means designing analog blocks with a wide tunable range and placing a digital-calibration loop around these blocks. Adding calibration loops around wide-tuning analog-circuit blocks enables the circuit to calibrate process variations and improve yield. Designers should choose digital implementations of analog-circuit blocks wherever feasible. EDN

AUTHOR'S BIOGRAPHY

Joseph Pun is an applications engineer at Gennum Corp's Snowbush IP Group. He obtained a bachelor's degree in electrical engineering from the University of Toronto (ON, Canada).

125Msps ADC at 125mW

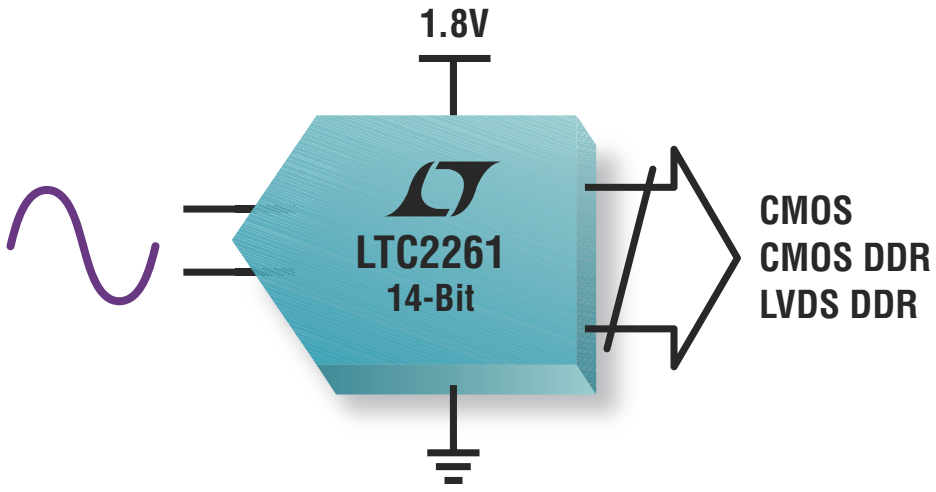

One-Third the Power of Comparable High Speed ADCs

Linear's newest high-speed ADC family achieves 1/3rd the power consumption of existing solutions without compromising AC performance. Operating from a low 1.8V supply, the 14-bit, 125Msps LTC®2261 dissipates 125mW while maintaining 73.4dB SNR and 85dBc SFDR. Digital outputs can be configured as DDR CMOS, DDR LVDS or standard CMOS for minimizing FPGA pin count. Setting a new low power industry standard enables a wide variety of applications to "cut the cord" and move into the portable world.

Features

- 73.4dB SNR, 85dB SFDR
- Single 1.8V Supply
- CMOS, DDR CMOS or DDR LVDS Digital Outputs
- Optional Data Output Randomizer
- Shutdown (0.5mW) and Nap (9mW) Modes
- 40-Pin (6mm x 6mm) QFN Package

Lowest Power ADC Family

Part Number	Speed	Power	SNR
LTC2261-14	125Msps	125mW	73.4dB
LTC2260-14	105Msps	106mW	73.4dB
LTC2259-14	80Msps	89mW	73.4dB
LTC2261-12	125Msps	125mW	70.8dB
LTC2260-12	105Msps	106mW	70.8dB
LTC2259-12	80Msps	89mW	70.8dB

Info & Free Samples

www.linear.com/2261

1-800-4-LINEAR

Free High Speed ADC Brochure

www.linear.com/adcsolutions

LINEAR TECHNOLOGY

EDITED BY MARTIN ROWE
AND FRAN GRANVILLE

designideas

READERS SOLVE DESIGN PROBLEMS

Fast 10-line-to-one-line data selector/multiplexer comprises only two ICs

Marián Štofka, Slovak University of Technology, Bratislava, Slovakia

When dealing with logic operations over BCD (binary-coded-decimal) numbers, you often need a 10-line-to-one-line data selector/multiplexer. In the past, you could use the famous 16-line-to-one-line 74150 multiplexer IC. Nowadays, however, when you look at the Web sites of the renowned semiconductor houses for the 150 and similar 16-to-one multiplexers, such as the 250, the 850, or the 851, you find that vendors have labeled them obsolete or no longer available. On the other hand, the eight-line-to-one-line multiplexers not only have survived but also are parts of advanced logic families, such as HC (high-speed CMOS) and AC (advanced CMOS).

The circuit in **Figure 1**, a 10-line-to-one-line data selector/multiplexer, comprises two eight-to-one multiplexers, IC_1 and IC_2. The A, B, and C bits of the address input of IC_1 connect to corresponding address bits—A, B, C, and D—of the main address input. The eight data inputs, D0 to D7, of the circuit are identical to the equally denoted data inputs of IC_1.

Whenever the main address is a binary-coded eight or nine, when A, B, C, and D=eight, the data input, D4 of IC_2, is active. When A, B, C, and D=nine, D5 of IC_2 is active. This shift in addressing of IC_2's data inputs is due to the IC's modified addressing: Address bit C connects to the MSB (most-significant bit) D of the main address input. The A and B are common to IC_1 and IC_2, respectively. To unite their outputs without using any additional logic, you must connect the noninverting output, Y, of IC_1 to data inputs D0 through D3 of IC_2. The eight lowest values, zero through seven, of the address always activate a signal of D0 through D3 in

DIs Inside

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

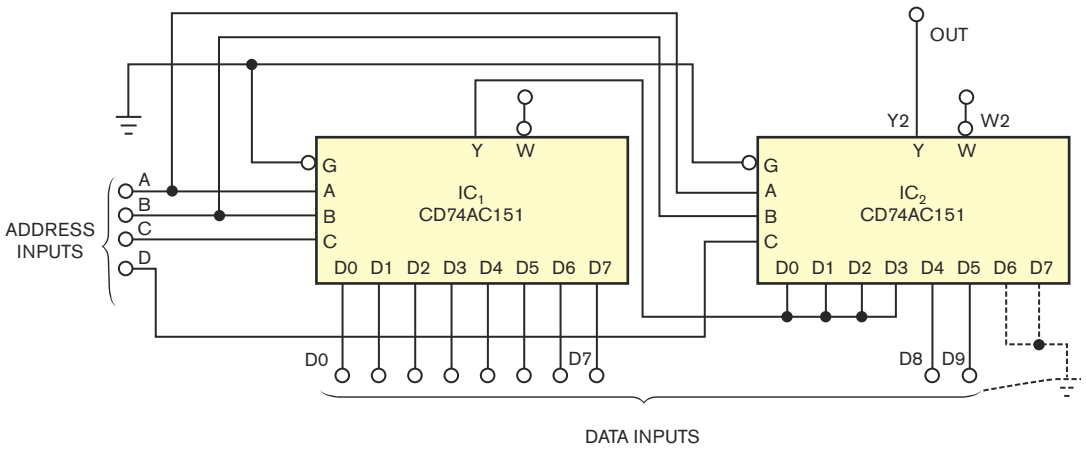

Figure 1 The maximum worst-case propagation delay of this 10-line-to-one-line data selector/multiplexer is 27 nsec, whereas the typical value is only 6.8 nsec. The circuit can also serve as a 12-to-one multiplexer.

IC_2. The output signal of IC_1 passes through one of these data inputs to the main output, Y2. If necessary, you can also use the W2 inverting output. Although the propagation delay from D0 through D7 to Y2 output is twice that from D8 and D9 to Y2, it is still less than 2×13.5 nsec=27 nsec for the CD74AC151 with a 5V supply. The typical delay is only 6.8 nsec.

Note that you can also use the circuit as a 12-line-to-one-line data selector/multiplexer by using the remaining data inputs, D6 and D7 of IC_2, which are idle in this circuit. In such a case, you attribute another notation of D10 to the D6 input of IC_2, and D11 holds for D7 one. Simultaneously, you must code the A, B, C, and D address in duodecimal code and, eventually, hexadecimal code, instead of BCD. EDN

Implement a simple digital-serial NRZ data-recovery algorithm in an FPGA

Jef Thoné and Bob Puers, ESAT-MICAS,
Katholieke Universiteit Leuven, Leuven, Belgium

Serial-data links embed clocks in their data streams, and those clocks must be recovered at the receiver end. This Design Idea describes a data/clock-recovery algorithm for an NRZ (nonreturn-to-zero), 1.5-Mbps data stream in a Xilinx (www.xilinx.com) Spartan XC3S200 FPGA. The algorithm employs a modified data-recovery application note (**Reference 1**). The application note uses the DCM (digital-clock manager) on the Xilinx Spartan and Virtex models, but this application uses a simplified algorithm that compares the data edges, if any, with internally generated clock edges, dynamically changing the data-input-to-data-output delay. The simplified algorithm allows integration in smaller CPLDs or FPGAs that lack a DCM (**Figure 1**).

The algorithm uses a 3-bit, free-running counter to generate the output clock, an 8-bit shift register to sample the serial data, seven XOR ports for edge detection, a 7-to-1 multiplexer with decoding for multiplexing the right-shift-register bit to the output, and some buffering registers. The algorithm runs at eight times the serial-data-stream speed, without a known phase relationship between both. It clocks the data into the shift register, which implies that, after eight clock cycles, the shift register will contain a rising edge; a falling edge; or, when the input data remains the same, no edge. The multiplexer does not take into account cases in which the shift register contains no edges or more than one edge.

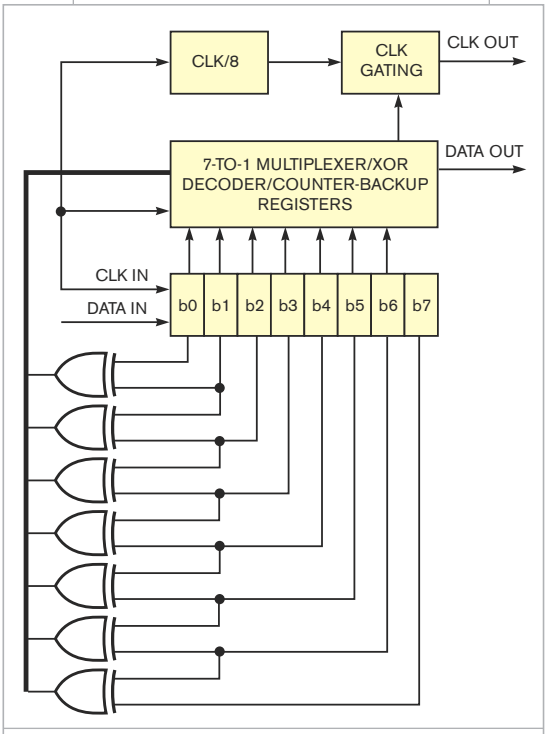

Figure 1 A clock-recovery circuit in an FPGA recovers data in a 1.5-Mbps data stream.

The edge location is checked in the shift register using the XOR-port array, which compares shift-register bit 0 with bit 1, bit 1 with bit 2, and so on. Depending on the output of the XOR array, showing where the edge occurs, a certain bit of the shift register multiplexes to the output. This action ensures that the output clock always toggles around the middle of the output-data bits.

When there are slight differences in clock speed and serial-input-data speed—for example, in the case of clock jitter or clock tolerances—the data-input phase continuously changes with regard to the output-clock phase as the algorithm tries to track the input-data phase. In this case, the multiplexer has an overflow, which happens when shift-register bit 7 multiplexes to the output, the next bit is

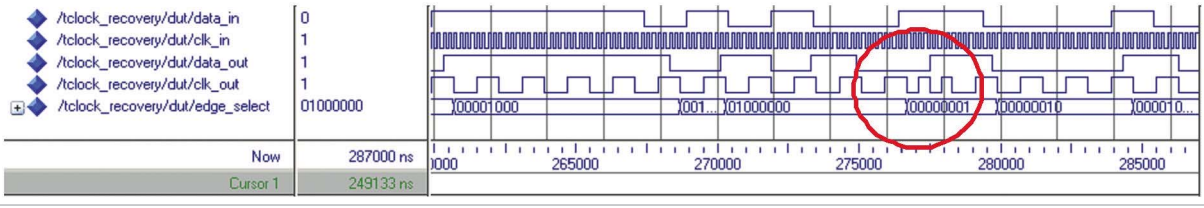

Figure 2 Doubling a clock output prevents a backward phase jump.

shift-register bit 1, or vice versa.

If bit 7 is output first—that is, the signal edge_select is 0100 0000—and the next selected bit is bit 1, with an edge_select of 0000 0001, a sudden phase jump in output data occurs. This phase jump is −360°×7/8, or −315°. Because the next input-data bit already had shifted in completely in the shift register, you need to employ a double-output clock once, so that the register doesn't miss a data bit (circled area in **Figure 2**).

When bit 1 is output, with an edge_select of 0000 0001, and the multiplexer jumps to bit 7, with an edge select of 0100 0000, a sudden phase jump in output data of 360°×7/8, or 315°, occurs. Because the shift-register data bit 7 is a delayed version of the last clocked bit, b1, the output clock must be stalled for one cycle. Otherwise, one bit too many will clock at the output (circled area in **Figure 3**). You can solve the overflow-phase jumps by gating the output clock using combinatorial logic. EDN

REFERENCE

1 Sawyer, Nick, "Data Recovery," Xilinx Application Note XAPP224, Version 2.5, July 11, 2005, www.xilinx.com/support/documentation/application_notes/xapp224.pdf.

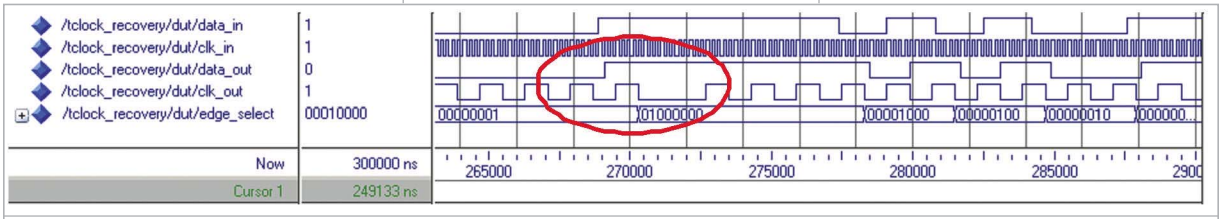

Figure 3 Stalling the output clock prevents a forward phase jump.

LED strobe has independent delay and duration

Michael C Page, Chelmsford, MA

The circuit in **Figure 1** is not complex, but it saved the day in an application involving visual inspection of the spray pattern of fuel injectors for quality and consistency. In this application, xenon strobe lights did not work because they take up too much space, and the light they emit is too intense. With a bank of six injectors with isolation panels, the reflection off a person's shirt or the wall behind him would interfere with the visual inspection. So the application instead used white HB LEDs (high-brightness light-emitting diodes) on "gooseneck"-type stands for adjustability in the chamber sections. Although the applica-

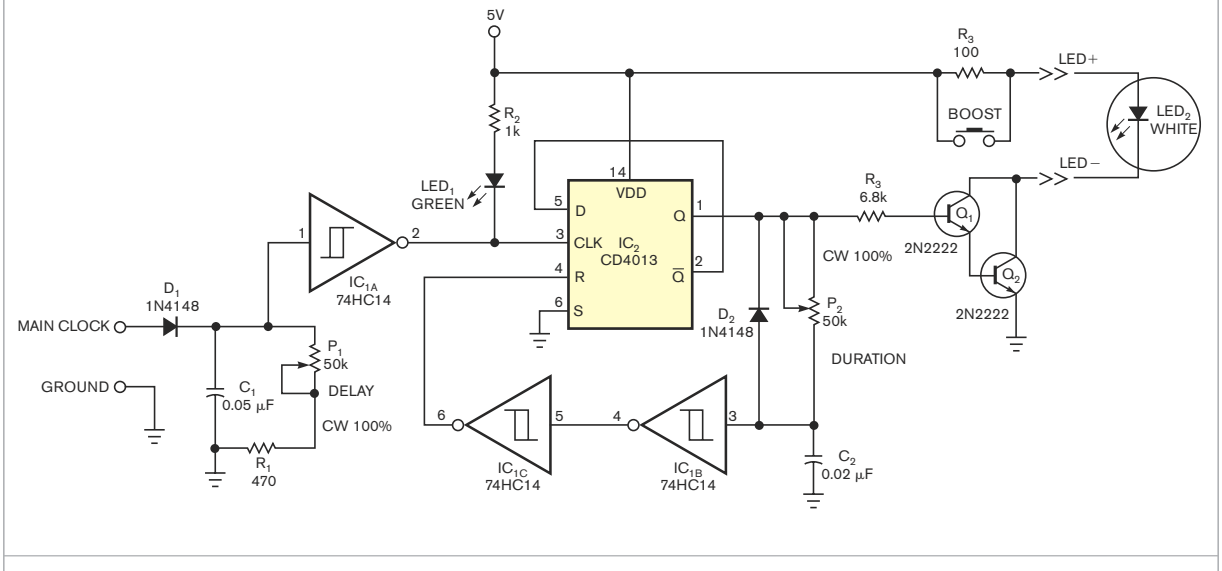

Figure 1 This circuit employs HB LEDs for a visual-inspection application.

LINEAR TECHNOLOGY

DESIGN NOTES

Hot Swap Controller, MOSFET and Sense Resistor Are Integrated in a 5mm × 3mm DFN for Accurate Current Limit and Load Current Monitoring in Tight Spaces – Design Note 466

Vladimir Ostrerov

Introduction

In general, a Hot Swap™ controller provides two important functions to boards that can be plugged and unplugged from a live backplane:

- It limits the potentially destructive inrush current when a board is plugged in.
- It acts as a circuit breaker, with the maximum current and maximum time at that current factored into the breaker function.

Two of the components required to implement these functions, the power MOSFET and sense resistor, tend to dominate the board real estate taken by the Hot Swap circuit. The LTC®4217 saves space by combining these two components with a Hot Swap controller in a 16-pin 5mm × 3mm DFN package (or 20-lead TSSOP). This 2A integrated Hot Swap controller fits easily onto boards operating in the voltage range from 2.9V to 26.5V. A dedicated 12V version, LTC4217-12, is also available, which contains preset 12V specific thresholds. Figure 1 shows how little space is required for a complete Hot Swap circuit.

LTC4217 Features

Figure 2 shows a simplified block diagram of the LTC4217. The controller provides inrush current control and a 5% accurate 2A current limit with foldback. For soft-start, an internal current source charges the gate of the N-channel MOSFET with 300V/s slew rate. Lower soft-start output voltage slew rates can be set by adding an external gate capacitor.

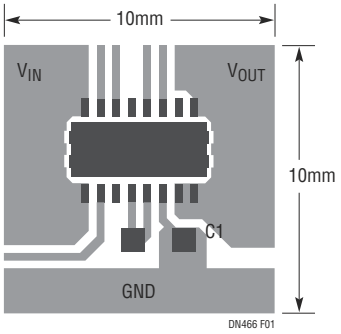

Figure 1. Tiny Integrated Controller Package Results in a Small Footprint

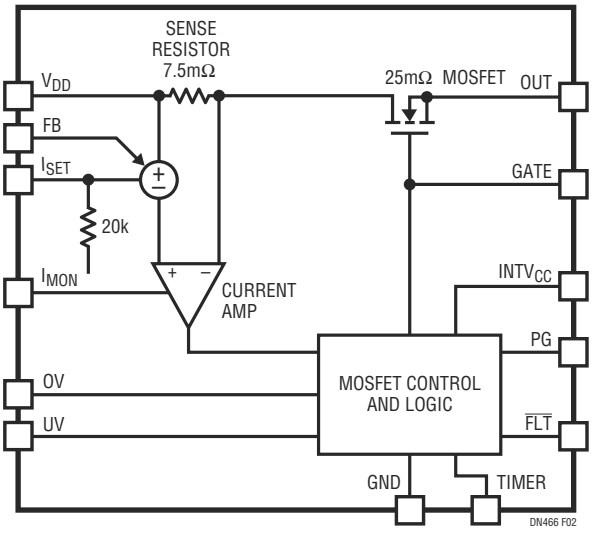

Figure 2. Simplified Block Diagram of the LTC4217

Integrated MOSFET and Sense Resistor

The LTC4217 integrates a 25mΩ MOSFET and 7.5mΩ current sense resistor. The default value of the active current limit is 2A, which can be adjusted to a lower value by adding an external resistor. Using an external analog switch for the connection of this resistor allows a start-up current to be larger than the maximum load current in steady state operation.

Adjustable Current Limit

The voltage at the I_{SET} pin determines the active current limit, which is 2A by default. This pin is driven by a 0.618mV voltage reference through a 20k resistor. An external resistor placed between the I_{SET} pin and ground forms a resistive divider with the internal 20k resistor. The divider acts to lower the voltage at the I_{SET} pin and therefore lower the current limit threshold.

The I_{SET} pin voltage increases linearly with temperature, with a slope of 3.2V/°C when no external resistor is present. This compensates for the temperature coefficient of the sense resistor—important for applications that must maintain monitoring accuracy over a wide temperature range. This also provides a convenient means to monitor the MOSFET temperature. If the die temperature exceeds 145°C, the MOSFET is turned off. It is turned on again when the temperature drops below 125°C.

Voltage and Current Monitoring

The LTC4217 protects the load from overvoltage and undervoltage conditions with a 2% accurate comparator threshold. The LTC4217 also features an adjustable current limit timer, a current monitor output and a fault output.

The adjustable current limit timer sets the time duration for current limit before the MOSFET is turned off. The current monitor produces a voltage signal scaled to the load current. The fault output is an open drain that pulls low when an overcurrent fault has occurred and the circuit breaker trips.

Typical Application

The LTC4217 application circuit shown in Figure 3 operates with a 100ms auto-retry time and a 2ms overcurrent condition when the load current reaches 2A. It also produces a voltage signal for an ADC to monitor load current.

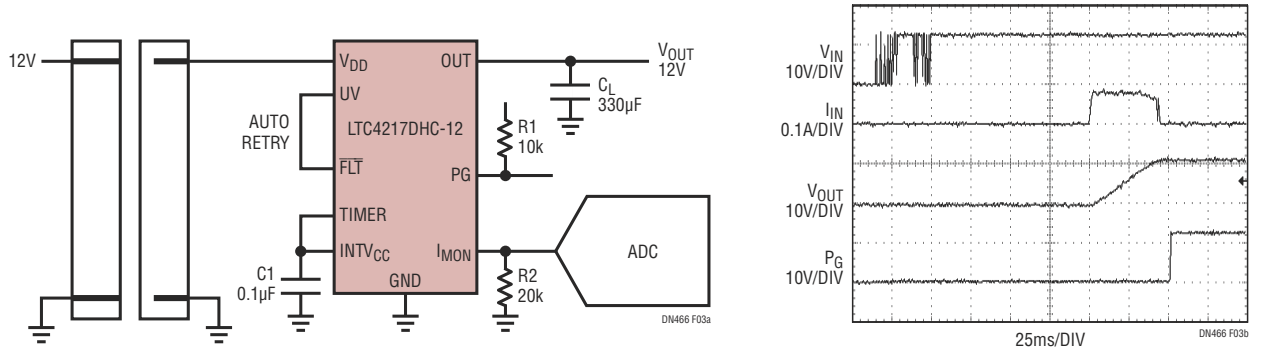

Figure 3. Typical Application with Current Monitored by an ADC

tion could have used the trusty old 555 timer, the delay and duration duty-cycle controls interact, which is an awkward situation.

The circuit in **Figure 1** shows the main-clock input; the delay and duration potentiometers, P_1 and P_2; and the HB-LED output. The circuit also includes an onboard general-purpose LED for bench testing to indicate an input signal, although, when the circuit is operating at high speeds, this LED is useless. The main-clock input is a 5V pulse of approximately 30 µsec coming from the fuel-pump index. Delay potentiometer P_1 adjusts the on-time delay of the LED from about 40 µsec to 2 msec, and duration potentiometer P_2 adjusts the LED-on, or flash, time with a range of approximately 15 µsec to 15 msec.

The circuit applies a 5V pulse, the main clock, to diode D_1 and capacitor C_1 to form a peak-hold circuit. C_1 then discharges at a rate that P_1 sets. Schmitt trigger IC_{1A} monitors C_1's voltage, and, when it reaches the low threshold of IC_{1A}, it outputs a high level to IC_2's clock input, setting the Q output high. With IC_2's Q output high, the Darlington-transistor pair comprising Q_1 and Q_2 turns on, driving the output to the HB LED low at the output, lighting the LED. At this time, capacitor C_2 charges at a rate that P_2 sets. When this voltage reaches the upper threshold of IC_{1B}, IC_{1C}'s output switches to high, resetting flip-flop IC_2's output back to low and turning off the HB LED. The circuit is now ready for another round. Diode D_2 ensures a complete discharge of capacitor C_2 for repeatability when you reset the Q output of IC_2 to low. Because IC_2 requires an active-high signal, you can omit IC_{1B} and IC_{1C}, but you should use a Schmitt trigger following an RC circuit for repeatability, especially on slow capacitor-charge/discharge times.

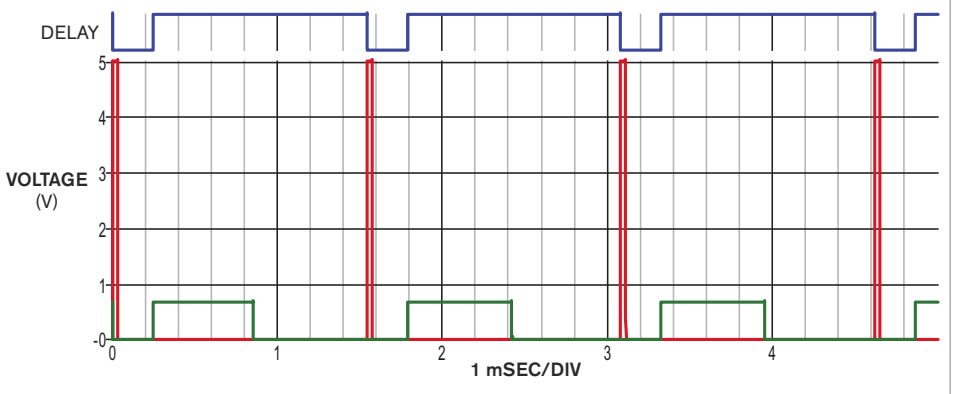

Figure 2 With a main-clock input of 650 Hz, the delay is approximately 250 µsec, with P_1 at 10%, and the duration is approximately 600 µsec, with P_2 at 75%. The top trace (blue) represents the strobe delay, the lower trace (green) represents Q_1's base duration, and the 5V trace (red) represents the main clock.

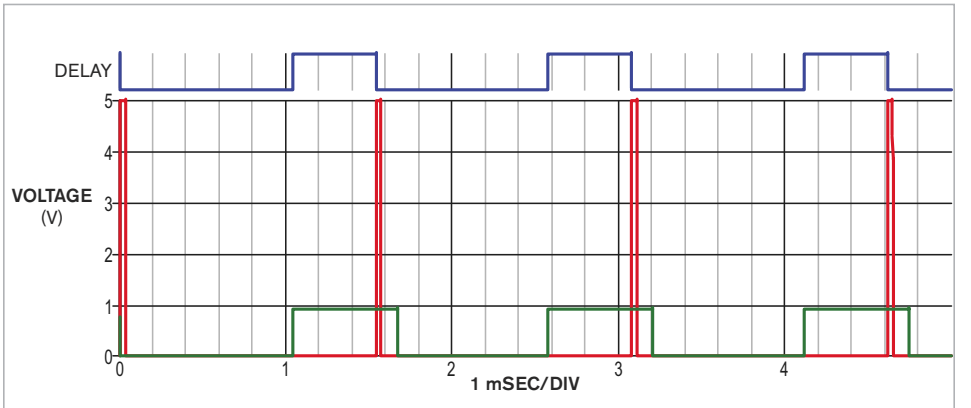

Figure 3 An adjustment change of delay occurs with the same duration as in Figure 2. The top trace (blue) represents the strobe delay, the lower trace (green) represents Q_1's base duration, and the 5V trace (red) represents the main clock.

Figure 2 shows the results of the circuit running with a main-clock input of 650 Hz and a delay of approximately 250 µsec, with P_1 at 10%, and a duration of approximately 600 µsec, with P_2 at 75%. **Figure 3** shows an adjusted change of delay with the same duration setting as in **Figure 2**. The new flash period overlaps the following fluid burst. You could, depending on the injector nozzle, see the end of one fuel burst of calibration fluid and the start of another in the chamber during the same flash period without encountering an error. The circuit also has a boost switch for a momentary intensity increase; otherwise, R_3 normally limits the current to approximately 40 mA. When you press the boost switch, the Darlington pair, two 2N2222 transistors with current of approximately 400 mA, still limits the current, but long-term use of the switch will shorten the LEDs' life. You should tailor the values of C_1, C_2, P_1, and P_2 to the application. Calculations will vary depending on the logic family you use, but, generally, $T=0.7\times R\times C$, where T is the time in seconds, R is the resistance, and C is the capacitance. EDN

Cancel sensor-wiring error with bias-current modulation

W Stephen Woodward, Chapel Hill, NC

The approximately −2-mV/°C temperature coefficient of diode junctions is a popular means of temperature measurement, especially in cryogenic applications (**Figure 1**). Diode temperature sensors are compact, stable, robust, sensitive, and inexpensive, and, unlike thermocouples, they require no reference junction. All of these benefits help explain the durable popularity of this—to use the polite term—"mature" technology.

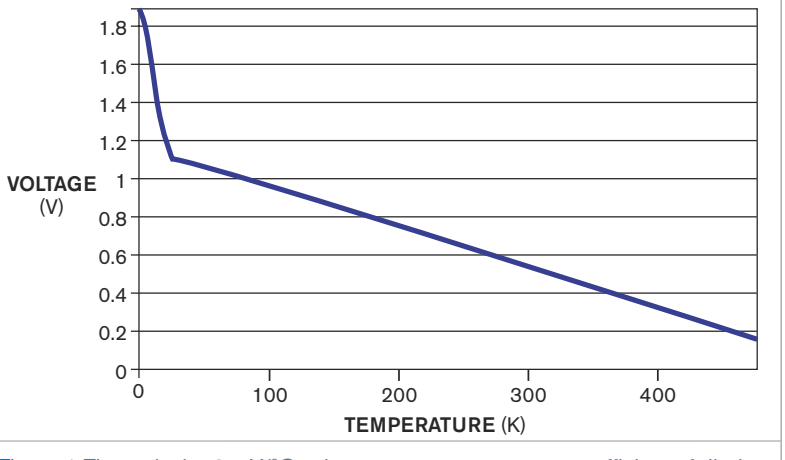

Figure 1 The typical −2-mV/°C-voltage-versus-temperature coefficient of diode sensors is large and nearly constant over a wide range of temperatures.

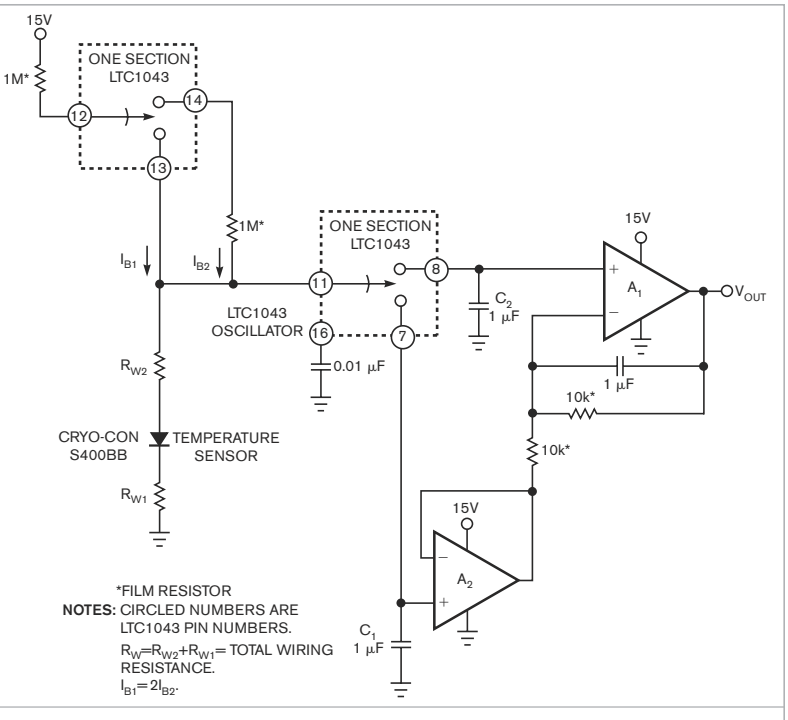

Figure 2 This circuit cancels the wiring-resistance error inherent in diode temperature sensors and requires only two conductors in the sensor cable.

A complicating factor and potential error source affecting these sensors arises from their need for bias-current excitation, however. The resulting contribution of ohmic IR (current/resistance)-voltage drop in the wiring and the connectors' resistance to the sensor's output voltage create spurious and temperature-sensitive voltage offsets. These offsets can introduce unacceptably large measurement error. This situation is especially likely when you use small and, therefore, high-resistance-gauge wire for sensor cabling, such as in cryogenic applications. In those cases, designers prefer exceptionally fine-gauge wire to minimize thermal conductivity and leakage.

The usual solution to the IR problem is to employ four-wire "Kelvin"-interconnection topologies, in which one pair of conductors carries the sensor's bias current and a separate, independent pair differentially senses the sensor's output voltage. This approach prevents corruption of the sensed voltage by IR drop in the bias pair. This traditional fix works well but complicates the wiring and doubles undesirable thermal leakage due to the extra wires, thus defeating much of the point of using fine-gauge cabling in the first place.

Figure 2 illustrates a circuit that implements a different approach. It cancels the wiring-resistance error and needs only two conductors in the sensor cable. It takes advantage of the fact that IR-voltage drop is directly proportional to current, but the sensor voltage is mostly constant. It works by alternating the magnitude of the excitation current, I_B, between two values, I_{B1} and I_{B2}, where $I_{B1}=2I_{B2}$. The ac component of the resulting signal is thus approximately I_BR_W, where R_W is the total wiring resistance plus a minor contribution from nonzero sensor impedance.

The clock for both I_{B1}/I_{B2} excitation modulation and synchronous demodulation of the resulting response is the internal oscillator of the LTC1043, which you set to approximately 500

Eliminate CAN or LIN interface with smart 18-bit bidirectional LVDS SerDes

The MAX9257/MAX9258 bidirectional LVDS serializer/deserializer (SerDes) chipset eliminates the need for an external CAN or LIN interface in automotive safety applications by integrating all control functions. This chipset simplifies system design by removing the costly system-configuration bus and associated wiring. These devices have smart features that enhance bandwidth utilization and suppress EMI.

Why do this with the typical automotive safety application?

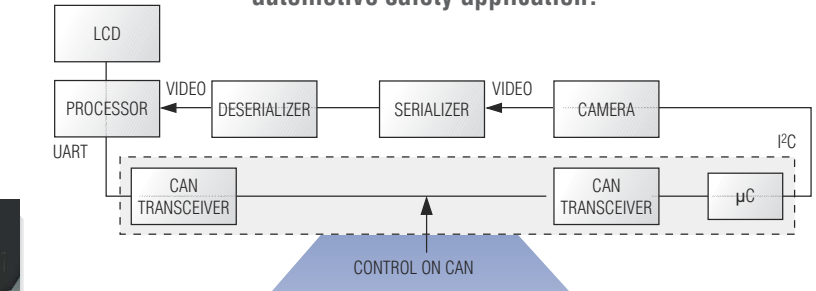

Night vision

Lane departures

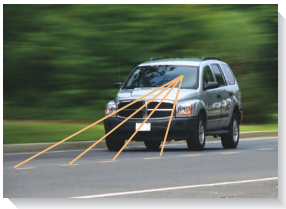

When you can use Maxim's simple solution!

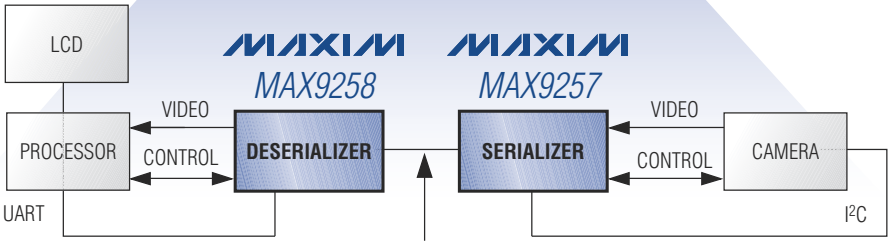

INTEGRATED CONTROL

Features

- **Programmable spread spectrum**
- **Deserializer does not require external reference clock**
- **ESD protection: ±10kV contact and ±30kV air-gap discharge**

Applications

- **Sign recognition**
- **Collision avoidance**
- **Night vision**
- **Lane departure**

www.maxim-ic.com/LVDS-SerDes

MAXIM DIRECT
www.maxim-ic.com/shop

AVNET electronics marketing
www.avnet.com

MAXIM
INNOVATION DELIVERED™

For free samples or technical support, visit our website.

Hz by connecting the external 0.01-µF capacitor to Pin 16. The resulting toggling of the excitation ballast resistance between 1M and 1M+1M=2M creates the 2-to-1 current modulation and an ac-signal component proportional to wiring resistance: I_BR_W.

The other side of the LTC1043 synchronously rectifies the I_BR_W ac component, storing the $I_{B1}R_W=V_{C1}$ phase on C_1 and the $I_{B2}R_W=V_{C2}$ phase on C_2. Op amp A_2 buffers V_{C1} and inputs it to the resistor network and A_1, which subtracts it from the average sensor signal, producing an output voltage independent of cabling-resistance offset. One downside of the technique is that, due to sensor-impedance effects on the order of 20 mV, the thermometric diode usually requires custom temperature calibration. EDN

Simple FSK modulator enables data transmission over low-speed link

Israel Schleicher, Prescott Valley, AZ

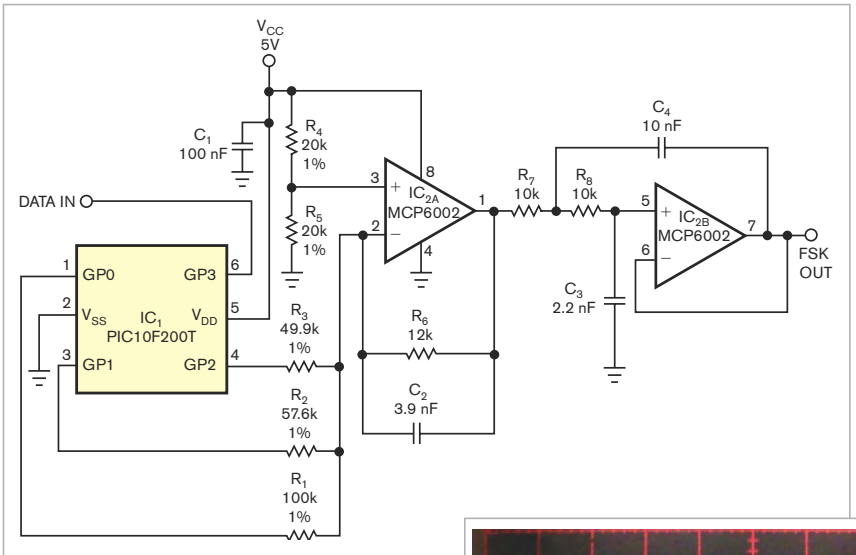

Figure 1 This microcontroller-based circuit generates Bell 202-compatible FSK modulation.

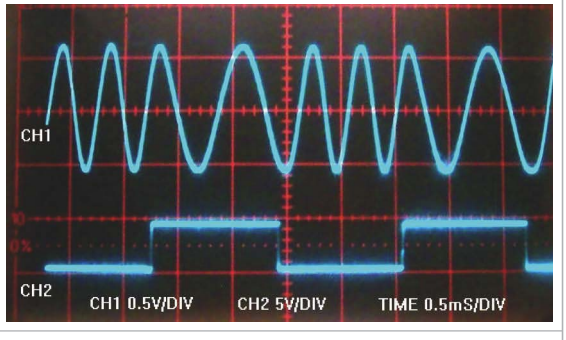

Figure 2 The FSK modulator's output changes frequency based on a digital input.

FSK (frequency-shift keying) is a type of signal modulation for transmitting digital data over an analog communication link. An FSK modulator comprises a digitally controlled sine-wave generator whose frequency shifts between two predetermined frequencies in response to the two logic levels of the digital data. The circuit in **Figure 1** generates a sine wave by continuously sampling a single sine cycle. The output of IC_{2A} is proportional to the currents through R_1, R_2, and R_3. These resistors connect together at one end to the inverting input of IC_{2A}, which is biased at $V_{CC}/2$. The outputs GP0, GP1, and GP2 of microcontroller IC_1 produce nonoverlapping pulse trains. When you set either output high or low, the others are off—that is, at high impedance. When you set an output high, the voltage across the resistor that connects to it is $V_{CC}/2$. When you set the output low, the voltage across the resistor is $-V_{CC}/2$.

Select the values of R_1, R_2, and R_3 so that the current pulses have magnitudes proportional to samples of sin 30, 60, and 90°, respectively. Setting all the outputs of IC_1 to off produces the sample of sin(0°), and no current flows through the resistors. Thus, starting with all outputs of IC_1 at off and consecutively and periodically setting GP0, GP1, and GP2 to high and then, in reverse order, setting GP1 and GP0 high again generates the positive half of a sine wave. Repeating the process but setting the outputs to low generates the negative half of the waveform.

This scheme produces a sampled sine waveform with 12 samples per cycle. In addition to the desired frequency component, f_0, this waveform contains higher-frequency components at $(12k+1)f_0$ and $(12k-1)f_0$, k=1,2,3, and so forth. The lowpass filter comprising IC_{2B}, R_7, R_8, C_3, and C_4 easily filters out these undesired components of smaller amplitude. **Listing 1**, which is available with the Web version of this Design Idea at www.edn.com/090611dia, is the assembly-program code that implements the Bell 202 FSK standard. When the control input Data In is high, the output frequency is 1200 Hz; when the control is low, the output frequency is 2200 Hz. The transition from one frequency to the other occurs in a manner that retains phase continuity. **Figure 2** shows the FSK-modulator output (CH1) in response to a modulating signal (CH2). EDN

Great things come in small packages

Industry's smallest current-sense amplifier

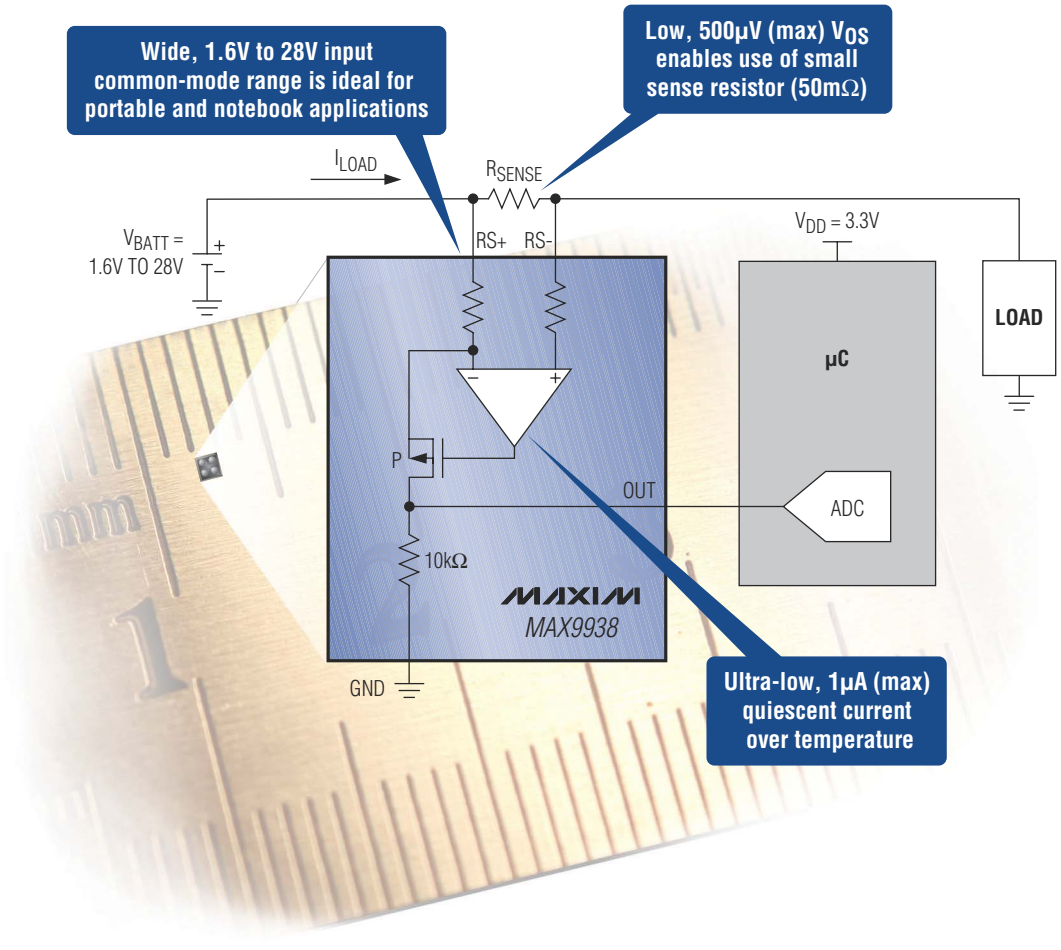

Tiny, 1mm x 1mm, 4-bump UCSP™ is 50% smaller than the competition

UCSP is a trademark of Maxim Integrated Products, Inc.

www.maxim-ic.com/MAX9938-info

MAXIM DIRECT
www.maxim-ic.com/shop

AVNET electronics marketing
www.avnet.com

MAXIM
INNOVATION DELIVERED™

For free samples or technical support, visit our website.

BY SUZANNE DEFFREE

supplychain

LINKING DESIGN AND RESOURCES

Economy opens more doors to demand creation for distributors

The ongoing downturn is causing an upturn for distribution-demand creation—using distributor-supplied suites of engineering services to encourage product training and guidance—as suppliers rely more on distributors to encourage complete designs. "What I've seen in a number of suppliers—not universally but in a number of key supplier relationships—is maybe a greater willingness to examine the way business is being done," says Mark Larson, president of components distributor Digi-Key Corp (www.digikey.com). "These changes were occurring even before the downturn. Now, everyone is standing back, and they are seeing that success is not coming as easily or, in some

cases not at all, and it forces a re-examination. It's an interesting thing, and I think it's going to make for more creative solutions."

Mike Long (**photo**), newly appointed chief executive officer of Arrow Electronics Inc (www.arrow.com), agrees that, in many ways, distributor relationships with suppliers are getting stronger. "We are looking for ways that we can increase or take on different portions of the supply-chain workload so we are not duplicating efforts as much as we used to," he says. "Suppliers would like us to continue on our journey to be more technically astute and bring their technology to more customers than we have in the past."

Long notes that companies are changing the ways they are doing designs. "If you go back 10 years, every engineer wanted to see another engineer that was around product specialization," he says. "Today, the engineers are more interested in solving a solution problem versus a part problem. Customers aren't interested in our doing a design for one part. They want the total solution to be efficient, and they want the solution to be designed with products that will be readily available for longer periods of time."

GREEN UPDATE

EPA FINDING ON GASES COULD AFFECT US IC MANUFACTURING

A recent move by the United States EPA (Environmental Protection Agency, www.epa.gov) could discourage IC manufacturing in the country. In late April, the EPA issued a proposed finding that greenhouse gases contribute to air pollution that may endanger public health or welfare.

Of the six gases identified as potential threats, hydrofluorocarbons, perfluorocarbons, and sulfur hexafluoride are used in chip manufacturing. Although the companies using them regulate the gases and although the EPA finding does not offer proposed regulation of the gases, it does open the door to possibly stringent and costly restrictions on equipment makers and semiconductor manufacturers.

The proposed endangerment finding is now in the public-comment period, which will end on June 23, 2009. Before taking any steps to reduce greenhouse gases under the Clean Air Act, the EPA must conduct "an appropriate process and consider stakeholder input." Both President Barack Obama and EPA Administrator Lisa P Jackson have repeatedly stated their preference for comprehensive legislation to address climate change and to create the framework for a clean-energy economy. You can submit written comments on the proposed finding through http://epa.gov/climatechange/endangerment.html.

OUTLOOK

SYSTEM ICs SET FOR DOUBLE-DIGIT GROWTH IN 2010

The global core-silicon market, including ASICs (application-specific integrated circuits), ASSPs (application-specific standard products), and PLDs (programmable-logic devices), appears to have hit bottom and should move toward an expansion in 2010, iSuppli Corp (www.isuppli.com) reports.

According to iSuppli, the core-silicon market saw its revenue fall by nearly one-third over the six months leading up to May and is set to see its revenue rise to $19 billion in the third quarter, up 8.3% from an estimated total of $17.6 billion in the second quarter. The revenue increase will mark the first sequential rise since a year earlier, when revenue increased by 6.7% in the third quarter of 2008.

"The core-silicon market hinges on a few major applications for most of its revenue," says Jordan Selburn, principal analyst for core silicon at iSuppli. "While some areas, most notably desktop PCs and 1G/2G [first-generation/second-generation] mobile handsets, now are forecast to actually suffer a decline in unit shipments from 2008 to 2009, others, such as 3G [third-generation] wireless phones, netbook PCs, and set-top boxes, are still expected to grow, despite the crippled economy."

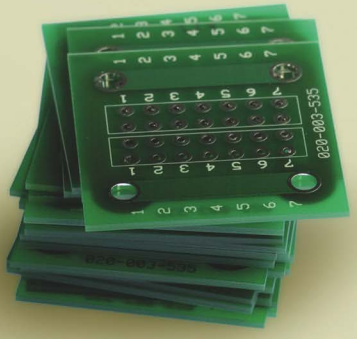

From quote to delivery, we've pioneered online PCB ordering.

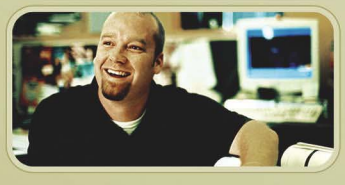

"Sunstone has done a great job with all of our orders. We rely on your easy online ordering system, quality boards, and fast lead times."

- Sunstone customer feedback

- 2-6 layer, quickturn proto-boards
- Complete customer service, 24/7, 365 days a year
- Over 30 years of manufacturing experience
- The ultimate in easy online PCB ordering

Visit us at
www.sunstone.com

13626 S. Freeman Road Mulino, OR 97042 | 503.829.9108

SUNSTONE CIRCUITS

productroundup

POWER SOURCES

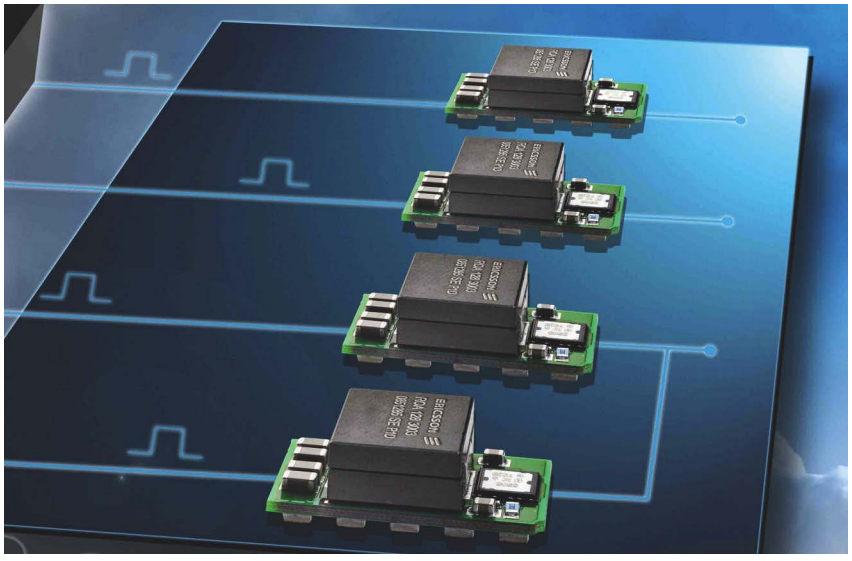

Power block suits centralized-control architectures

Aiming at use in power FPGA and other processors, the 30A ROA14283003 power block targets use in centralized-control architectures, such as Internet-switching routers, computer servers, and telecommunications switches. The device allows you to insert it as a power stage that a single- or a multiple-phase PWM controls and works with the vendor's isolated and nonisolated power modules to enable power management in information- and communication-technology applications. Using a digital interface, such as the I²C or the PMBus, the device accepts 7 to 13.2V inputs, provides a 0.8 to 3.3V output, and claims 94% efficiency at full load. Built with an 11.4-mm maximum profile, the ROA14283003 power block costs $10.05.

Ericsson Power Modules, www.ericsson.com

1050W front end has wide output-load range

Providing a maximum 86.7A with a 12V output, the SFP1050-12BG 1050W ac/dc front end exceeds Climate Savers Computing Initiative's Silver Standard efficiency requirements over output loads of 20 to 100%. Compliance with these efficiency requirements and active current sharing make the device suitable for use in redundant-system applications for datacom, server, and other distributed-power applications. Features include a 3.3V standby voltage at 3A output, I²C monitoring and control, input of 90 to 264V ac, and 91% efficiency in loads as low as 50%. The device also provides output-overvoltage, output-overcurrent, and overtemperature protection. Measuring 12×3.1×1.6 in., the SFP1050-12BG 1050W ac/dc front end costs $512 (250).

Power-One, www.power-one.com

DC-power supplies come in 2U packages

The programmable XR dc-power-supply series includes 54 models with 2-, 4-, and 6-kW power ratings; 1000V-dc voltage ratings; and 375A dc-current ratings. The series features the vendor's LXI-certified Ethernet-interface option with an embedded Web server. The interface allows control and monitoring of the power supply in any Web browser or through National Instruments' LabView with certified drivers. Accepting an ac-input-voltage range of 208 to 480V ac at 50-, 60-, and 400-Hz, the series provides a 92% power factor at rated power. The combined features function at 88% efficiencies. Available in a 2U package, the XR series costs $2795.

Magna-Power Electronics, www.magna-power.com

Point-of-load dc/dc converter claims 92% power-conversion efficiency

The 7A iBD series of nonisolated point-of-load dc/dc converters provides a Distributed Power Open Standards Alliance-compatible pinout and claims 92% high-power-conversion efficiency. The 6 to 14V input range suits the series for powering distributed power and 4-to-1 and 5-to-1 unregulated bus-converter systems. A 0.8 to 5.5V output-adjustment range, a prebiased output, and a 7A maximum output current enable the devices' use in ASIC, FPGA, and other applications requiring multiple voltages. Available in a 1×0.5×0.33-mm through-hole-mounted SIP, the iBD dc/dc converter costs $9.67 (10,000).

TDK-Lambda, www.tdk-lambda.com

INTEGRATED CIRCUITS

Class D-amplifier IC increases power output in portable stereo amplifiers

The two-channel TS4999 Class D-stereo-amplifier IC claims to improve sound quality for portable consumer equipment using 3-D-audio effects. The device features a 2.4V minimum supply voltage, a 10-nA power-saving mode, and 90% efficiency. It achieves 1% distortion and delivers as much as 2.8W per channel into 4Ω speaker loads. The TS4999 is available in an 18-bump, flip-chip lead-free package and costs $1.10 (1000).

STMicroelectronics, www.st.com

DAC features digital filtering targeting professional-audio equipment

Aiming at professional-audio applications, the WM8742 DAC suits audio/visual receivers, compact discs, DVDs, SACDs, and home-theater systems. The 24-bit, 192-kHz DAC combines a programmable digital filter and a 123-dB stereo SNR. The system includes a dithered digital-interpolation filter, fine-resolution volume control, digitally enabled de-emphasis, and a multibit sigma-delta modulator/switched-capacitor stage with differential-voltage outputs. Pin-compatible with the WM8740 and the WM8741, the WM8742 comes in an SSOP-28 package and costs $3 (1000).

Wolfson Microelectronics, www.wolfsonmicro.com

Hall-effect switch comes in lead-free, "green"-compliant package

Adding to the vendor's Micropower and Omnipolar line, the MLX90248 Hall-effect noncontact magnetic switch provides ESD protection as high as 8 kV. The MLX90248 requires 1.5V operating voltage, compared with previous generations, which required 2.5V operating voltage. Providing bounce-free switching, the device suits use in open/close detection in mobile phones and notebooks and screen-rotation detection in digital cameras and camcorders. The IC typically consumes 5 µA of power at 3V and 3 µW at 1.5V. Allowing detection of magnetic fields lower than 6 mT, the device detects the presence of north or south magnetic poles. Available in lead-free, "green"-compliant TSOT packages and 2×1.5×0.4-mm CSPs, the MLX90248 Hall-effect switches cost 16 and 18 cents, respectively.

Melexis, www.melexis.com

Digital step attenuator has neither gate lag nor phase drift

Covering a 31.5-dB attenuation range in 0.5- or 1-dB steps, the PE43502 digital step attenuator provides low insertion loss and high attenuation accuracy. The attenuator has neither gate lag nor phase drift, and it has fast settling. These features make it suitable for wireless-broadband-access applications, including TDE, WiMax, and TD-SCDMA; cellular base stations; repeaters; femtocells; and power-amplifier-distortion-canceling loops. The 50Ω devices feature a dc to 6-GHz operating-frequency range, an input third-order intercept point of 58 dBm, and a 500V HBM (human-body-model) ESD tolerance. Available in 3.3 or 5V supply-voltage options, the attenuator uses onboard CMOS logic, facilitating 2.75V control. Available in a QFN-24 lead package, the PE43502 digital step attenuator costs $2.14 (100,000).

Peregrine Semiconductor, www.psemi.com

COMPUTERS AND PERIPHERALS

Integrated platform provides processor options

Part of the vendor's IPnexus MicroTCA line, the AMP5071 integrated platform uses an Intel Core 2 Duo processor or a Freescale MPC8641D dual-core 1-GHz PowerPC processor. The platform allows for configuration and integration of AMC modules for I/O, storage, and computing functions into the system, meeting a range of IP-based communications-design criteria. The system comes with the vendor's NexusWare carrier-grade Linux OS and development environment and NexusWare Portal remote-system-management software. Prices for the AMP5071 application-ready system, configured with a processor, storage, and NexusWare, start at $4595.

Performance Technologies, www.pt.com

Hardware-port multiplier integrates a RAID controller

Featuring an integrated RAID controller, the 4X1 eSATA/USB-hub hardware-port multiplier connects as many as four SATA/SATA II hard drives to a host using a single SATA or USB port. The four connected drives allow you to set up a RAID configuration by dialing an onboard rotary switch. Available in system and enclosure versions, the 4X1 eSATA/USB-hub hardware-port multiplier costs $89.95.

Addonics Technologies, www.addonics.com

1.5-Tbyte hard drive has 3-Gbps SATA interface

The EcoGreen F2EG hard-disk family delivers a 1.5-Tbyte capacity with 500 Gbytes per disk. Suiting use as an external hard disk and in desktop PCs, the devices incorporate 16- or 32-Mbyte buffer memory; a 3-Gbps SATA interface; and 500-Gbyte, 1-Tbyte, or 1.5-Tbyte capacity. The EcoGreen F2EG 1.5-Tbyte hard-disk drive costs $149.

Samsung Electronics, www.samsung.com

EMBEDDED SYSTEMS

USB RS-232 adapter has programmable baud rate

Appearing as a standard COM port to the host computer, the SeaLink+232-DB9 single-port USB-to-RS-232 adapter enables setup and provides compatibility with legacy software. The device features a programmable baud rate and data formats with 128-byte transmitter and 384-byte receiver buffers. Compatible with all standard PC baud rates, the USB adapter supports 921.6-kbps communication. Powered by the USB port, the adapter provides status LEDs molded into the enclosure, indicating serial-data activity and connection to the host. The SeaLink adapter costs $79.

Sealevel Systems, www.sealevel.com

Rack-mount platform provides 15×GbE

The rack-mountable, modular PL-80090 platform uses a redundant power supply and provides 15×GbE, making it suitable for use in critical networking and network-security applications. Built with Intel embedded IA components, the platform supports Intel quad-core processors with 1333-MHz front-side bus and Intel Virtualization technology. The platform provides two unbuffered ECC or non-ECC DDRII-800 DIMM slots with as much as 4-Gbit memory. The storage interfaces include two 3.5-in. SATA hard-disk drives and CompactFlash. The PL-80090 costs $1739.

Win Enterprises, www.win-enterprises.com

Computer-on-module uses Intel Atom N270 processor

The low-power, long-life-cycle computer-on-module line includes Diamondville, a low-power Intel Atom N270 processor that operates at 1.5 GHz. Available in ETX and PICMG's COM Express basic and compact sizes, the Express-AT basic measures 125×95 mm, and the Express-ATC compact measures 95×95 mm. Suiting medical equipment, portable devices, and industrial controllers, the Express-AT, Express-ATC, and ETX-AT claim seven-year life cycles. Prices for the modules start at $200.

Adlink Technology, www.adlinktech.com

Development tool addresses stack-overflow issues

The StackX tool allows developers to avoid stack-overflow problems frequenting embedded systems. The development tool performs an analysis of a complete application at the executable-code level, computing the worst-case stack usage the application can experience. The StackX costs $1000 per developer for a perpetual license.

Express Logic, www.rtos.com

EDN ADVERTISER INDEX

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

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

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

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

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

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

BuyerZone
Where Smart Businesses Buy and Sell

A division of Reed Business Information.

BY STEVE LUBS • DEPARTMENT OF DEFENSE

The boss is always right, even when he's wrong

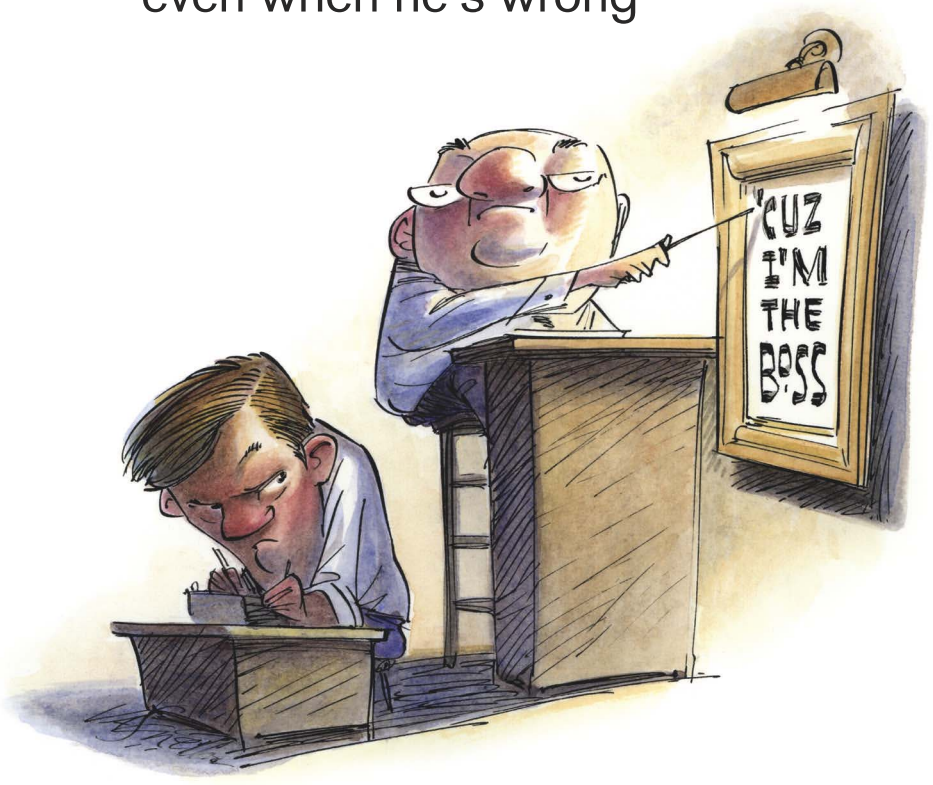

DANIEL VASCONCELLOS

In the early 1980s, I was working at a new division that built special-purpose modems. Our boxes were effective and efficient, working with well-known modulation, such as FSK (frequency-shift keying), BPSK (binary-phase-shift keying), and QPSK (quadrature-phase-shift keying). When an application required a different kind of signal, however, we ran into power and space limitations and other headaches.

My first assignment was to handle an unusual, high-data-rate signal. My boss, a 20-year analog-veteran designer, had decided to try digital processing. His protocol, however, was to use only his "tried-and-true" analog techniques in digital form—no "fancy programming" allowed. I sighed, studied previous equipment designs, and sketched several proposals. Then I hit the first of several roadblocks.

My boss and I argued about theory and my designs. When I asked questions about details in earlier boxes, he would lecture me as if I were a neophyte. After more design reviews than I can count on both hands, plus lectures, I finally had a completed system on my bench. The first test, as per my boss' protocol, was to input a square wave as data; the display output should be a single slowly rotating vector. My display output, unfortunately, included three vectors, and they weren't rotating slowly. He glanced at it and said, "Double-check your work."

So I spent the next three weeks double-checking wiring, schematics, clock signals, timing diagrams—anything that might be out of place. Finding nothing wrong, I triple-checked the schematics and replaced parts. The output remained the same. In desperation, I changed wiring and clock signals until the display showed only one vector rotating. Unfortunately, when I tried simulated data instead of a square wave, my new circuit produced only noise.

Frustrated, I begged one of my coworkers to help me. Studying the original design, he pointed out that it was working: My data-sampling clock was a noninteger multiple of the system clock, so the two clocks couldn't exactly synchronize—hence, the three rotating vectors.

My boss responded: "It doesn't work."

"Yes, it does," I countered. "It's doing what you—what we designed it to do."

His observation: "Too many errors in the output data."

"It'll work! It needs a faster sampling rate. We just multiply …"

But I was too late. He had decided that "digital" processing was a mistake and handed the project to a co-worker to do in analog form.

Two years later, Texas Instruments' first commercial DSP chip appeared. One of our new guys got to use it. He showed me how to replace several analog boards with one DSP board and some support ICs. I saw my original project and a few later ones implemented in software.

I have moved on, trying to keep an open mind. It's been difficult, with the deluge of new technology every year. Whenever I find myself slipping, I remember my old boss, spouting his tried-and-true philosophy, keeping me out of DSP and away from what is now a major part of electronic design. I hope I don't force anyone else into that scenario. EDN

Steve Lubs has been an engineer in a variety of roles at the Defense Department for 30 years and has always argued with his bosses. You can reach him at salubs1@verizon.net.

⊞ www.edn.com/tales

The path of least resistance

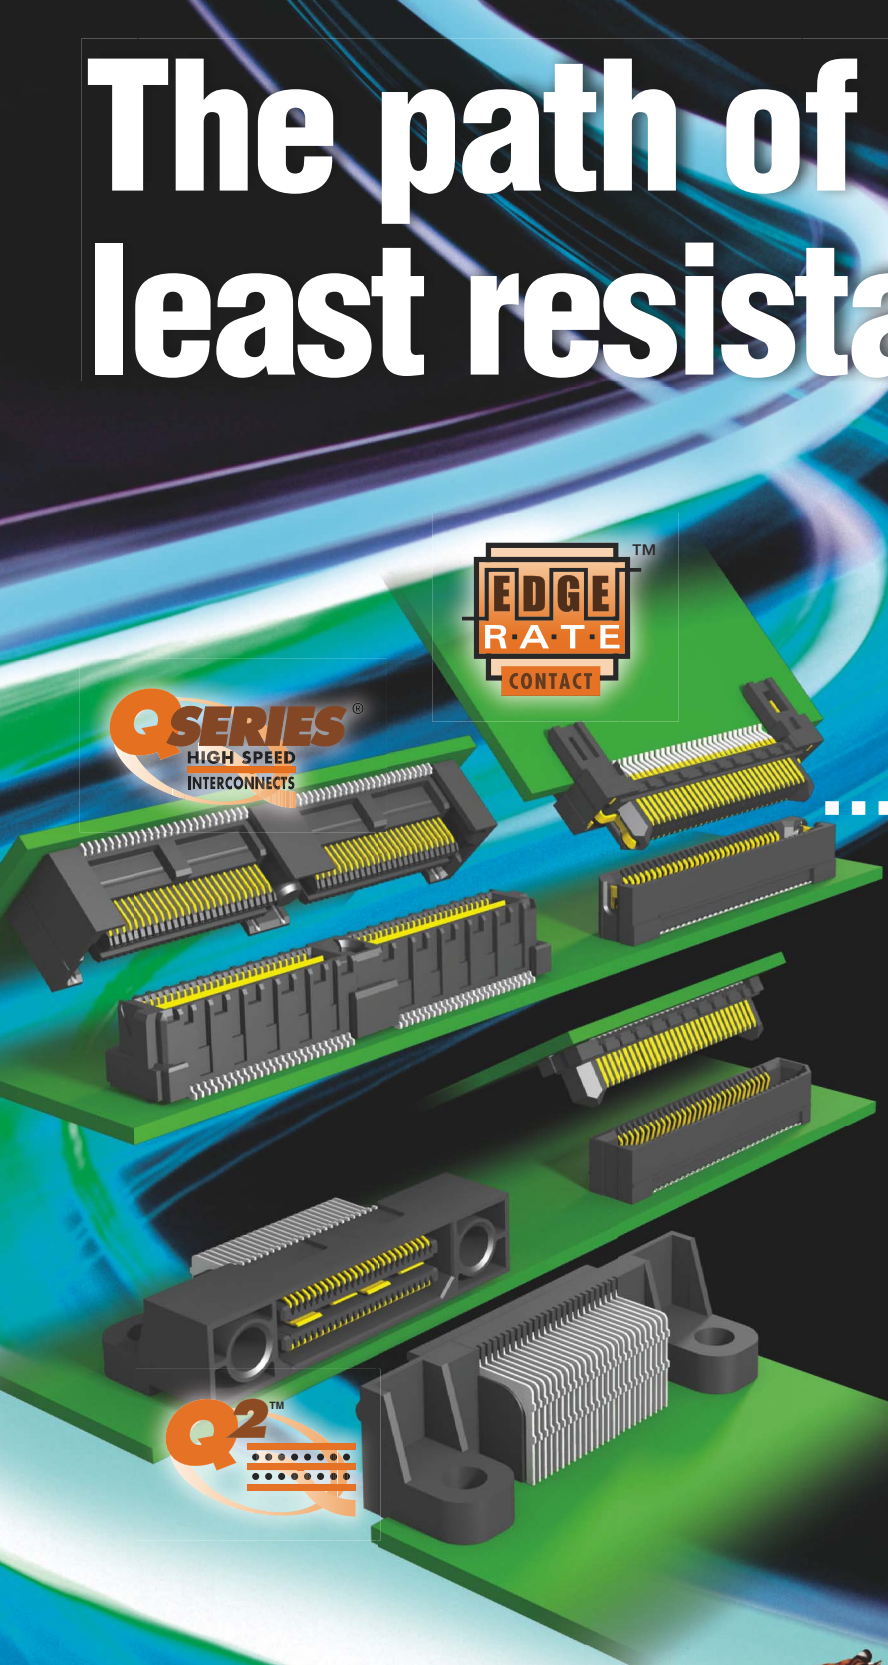

...for high speed board-to-board.

Q Series®

- Three pitches (.8mm, .635mm, .5mm), 5mm to 30mm stack heights
- Rated at 9 GHz single ended
- Integral power/ground plane

Q2™ Series

- Rugged design on .635mm pitch
- Rated at 9 GHz single ended

Edge Rate™ Series

- Contact system on .8mm pitch is robust when zippered during unmating
- Rated at 9.5 GHz in differential pairs

All products available in mezzanine, perpendicular and coplanar applications

samtec

TRANSMISSION LINE SOLUTIONS

www.samtec.com/tls

CADENCE INTRODUCES FPGA SYSTEM PLANNER

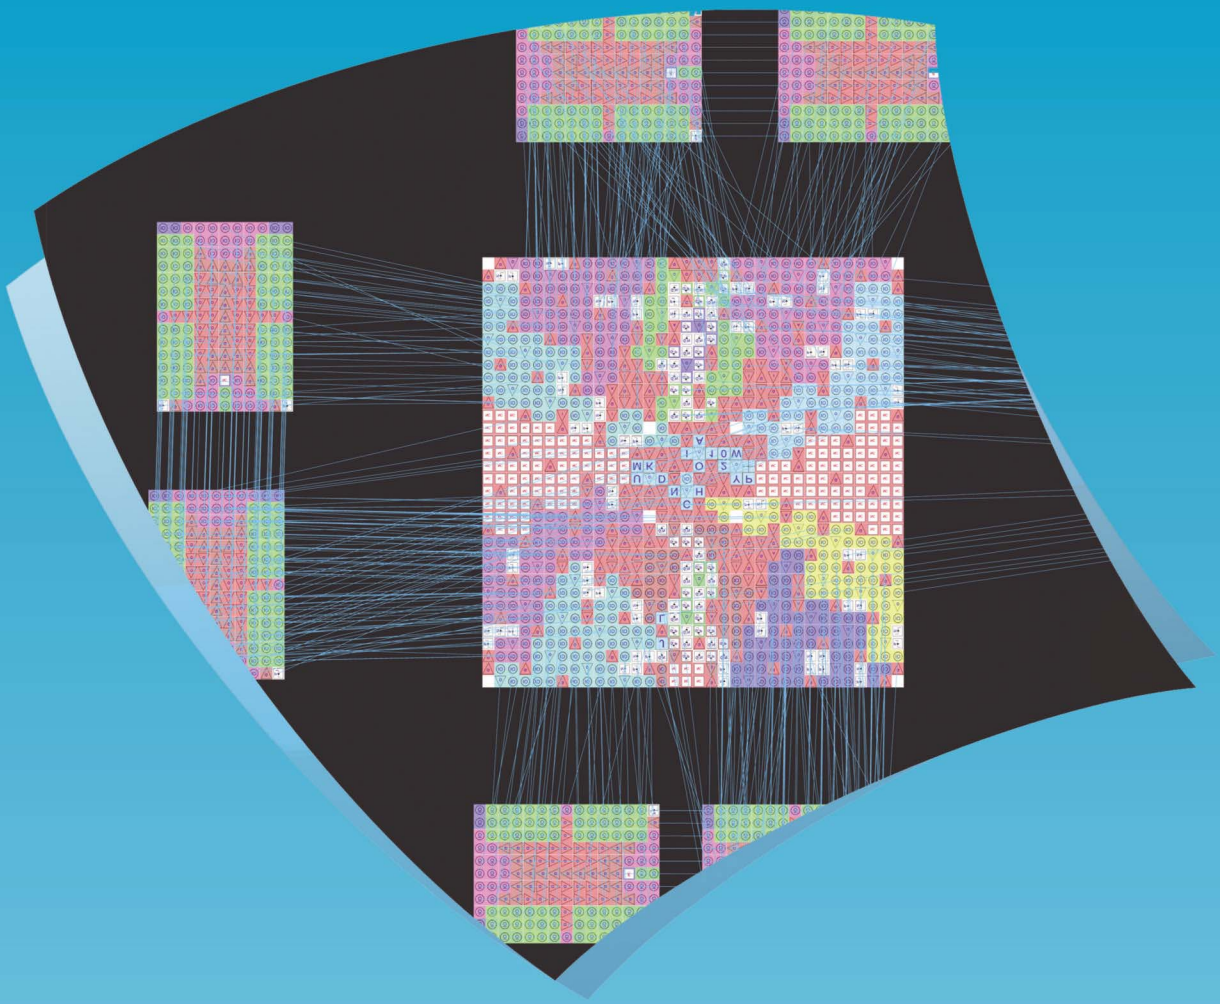

Innovative FPGA PCB Co-design Technology
The Cadence® Allegro® and OrCAD® FPGA System Planners offer optimized correct-by-construction FPGA pin assignments that minimize the number of iterations during PCB layout while reducing the number of layers required to route the FPGA. This innovative design solution is scalable, shortens design time, reduces product costs and mitigates risk.

Available from EMA Design Automation
To learn more about the Cadence FPGA System Planner visit EMA Design Automation, a Cadence Channel Partner, at www.ema-eda.com/FPGA or call us at 800.813.7288.

cadence™

EMA Design Automation™
Cadence Channel Partner