UBM Electronics

EDN®

VOICE OF THE ENGINEER

JAN 20

Issue 2/2011

www.edn.com

BATTERY-STACK-MONITOR ICs SCRUTINIZE THE CELLS

Page 36

HARDWARE-BASED VIRTUALIZATION EASES DESIGN WITH MULTICORE PROCESSORS

Page 28

MANAGING SIGNAL INTEGRITY IN TOMORROW'S HIGH-SPEED FLASH-MEMORY-SYSTEM DESIGNS

Page 44

YOU NEED IT, WE HAVE IT

Digi-Key®
CORPORATION

Over 515,000 Products in Stock
More Than 45,000 NEW Products Added in the Last 90 Days
440+ Suppliers

The industry's broadest product selection available for immediate delivery

www.digikey.com
1.800.344.4539

NEDA Supplier Authorized Distributors


We're giving power a
digital upgrade.
Find out how our new Digital EXL Core™ technology
is redefining Power Factor Correction ICs.
Cirrus Logic's CS1500 and CS1600 are the industry's first Digital Power Factor Correction (PFC) ICs to surpass ordinary analog solutions in performance and price. Enabled by Cirrus Logic's EXL Core technology, these digital solutions intelligently solve traditional power management challenges with newly patented and patent pending designs that dramatically reduce the need for bulky, high-priced components and complex circuitry, lowering overall system cost and simplifying designs. Cirrus Logic now makes it possible for more energy-efficient power supplies in digital televisions, notebook adapters, PC power supplies and lighting ballast applications through advanced digital technology that creates smarter, greener, and more energy-efficient products.
Experience Digital PFC. Register for free samples at www.cirrus.com/ednpfc
EXL CORE
CIRRUS LOGIC®
© 2010 Cirrus Logic, Inc. All rights reserved. Cirrus Logic, Cirrus, the Cirrus Logic logo designs, EXL Core, and the EXL Core logo design are trademarks of Cirrus Logic, Inc. EDN06242010

New! *LPD3015 and LPD4012 miniature coupled inductors. Perfect for LED drivers in portable equipment.*

Pick the best coupled inductor for your SEPIC converter in less time than it takes to read this ad.

(and have free samples on your desk tomorrow)

With over 135 coupled inductors to pick from, choosing the right one might seem difficult.

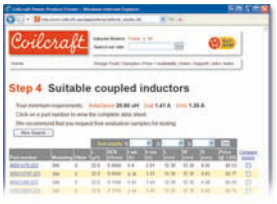

But our new on-line tools make it incredibly accurate and easy.

Enter data about your SEPIC converter and you'll get detailed specifications on all the coupled inductors that meet your needs. A few more clicks gives you a detailed analysis of core and winding losses. The whole process can literally take less than a minute.

When you're done, be sure to ask for free evaluation samples. We'll ship them that same day.

Start your coupled inductor search at **coilcraft.com/sepic**

ORDER DIRECT
BUY.COILCRAFT.COM
800-322-2645
OVERNIGHT DELIVERY ORDER BY 5
RoHS COMPLIANT

Coilcraft®

www.coilcraft.com 800/322-2645

EDN 1.20.11 contents

Battery-stack-monitor ICs scrutinize the cells

by Paul Rako, Technical Editor

Hardware-based virtualization eases design with multicore processors

by Satish Sathe, AppliedMicro

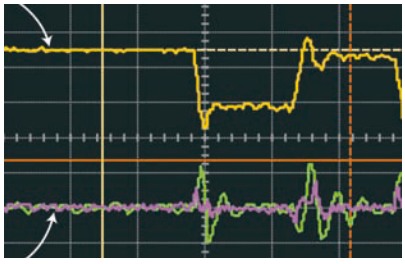

Managing signal integrity in tomorrow's high-speed flash-memory-system designs

by Perry Keller, Agilent Technologies

pulse

COVER: DASHBOARD & LANDSCAPE: NARVIKK/ISTOCKPHOTO.COM; BATTERY ICON: VADYM TYNENKO/ISTOCKPHOTO.COM

DESIGNIDEAS

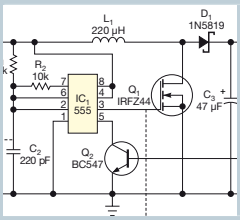

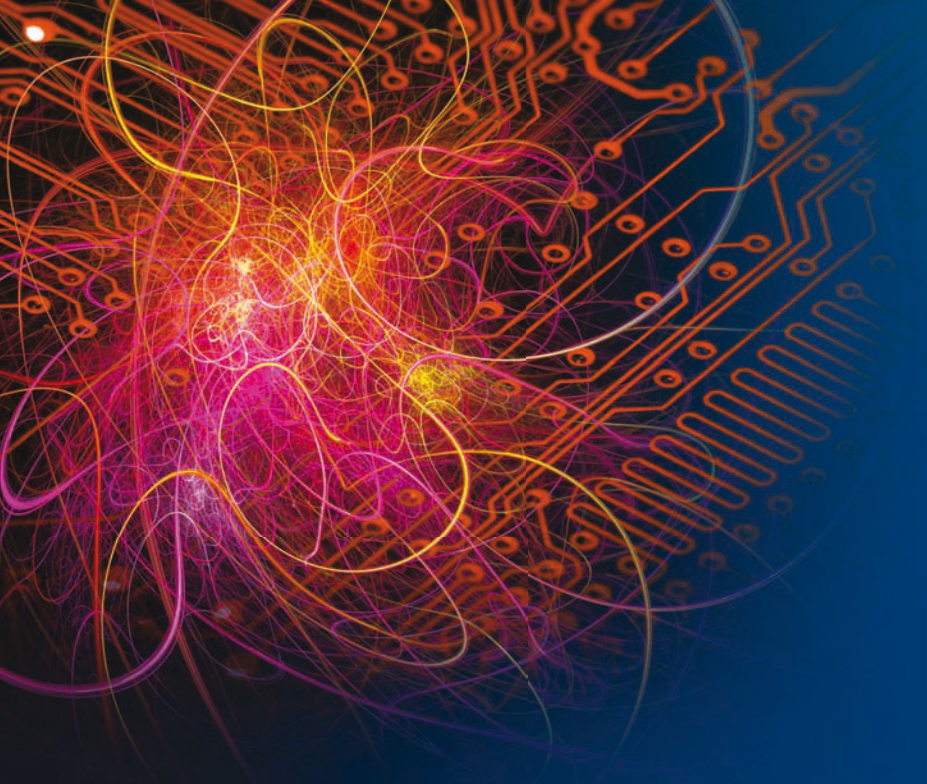

mouser.com
Semiconductors and electronic components for design engineers.

We've got what design engineers need. **The future.**

You'll find the newest products for your newest designs.
Get what's next at Mouser.

Scan Here

mouser.com

MOUSER ELECTRONICS®
a tti company

contents 1.20.11

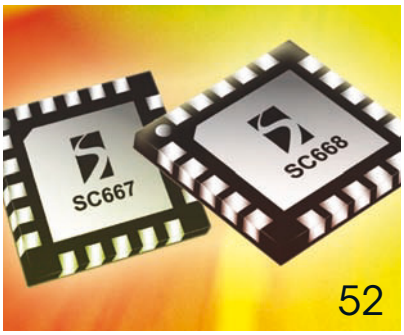

52

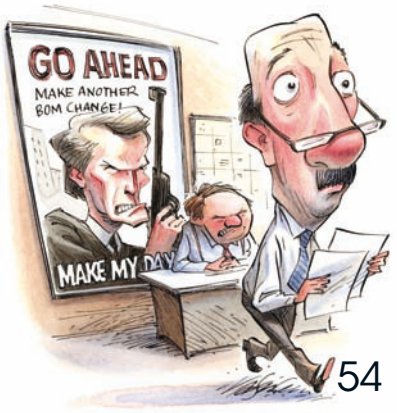

54

DEPARTMENTS & COLUMNS

EDN online contents www.edn.com

ONLINE ONLY

Check out these Web-exclusive articles:

Critical false-path analysis through sensitization methods
Static and dynamic false-path determination is essential for accurate application of false paths in an SOC design.
→www.edn.com/110120toca

Ensuring form, fit, and function for advanced interactive digital displays
The environment in which end-user equipment will be deployed has a critical influence on hardware selection, including the display front panel, which must provide a carefully optimized blend of optical, mechanical, and electrical properties.
→www.edn.com/110120tocb

READERS' CHOICE 2010
Take a look at 20 of the most-clicked-on articles and blog posts from across EDN.com in 2010.
→www.edn.com/110120tocc

PRYING EYES

In *EDN*'s Prying Eyes, we peer inside an end-user consumer gadget, a reference design, or any other interesting electronics-enabled thing we can get a good look at. Unlike your average bill-of-materials tear-down, Prying Eyes aims to illuminate the tough decisions the engineers responsible for the design had to make.
→www.edn.com/pryingeyes

What's happening in the electronics world?

The EDN.com News Center provides real-time breaking news and analysis on the global electronics industry, including coverage of semiconductors, lawsuits, IC design, microprocessor units, consumer electronics, analog, and business trends.
→www.edn.com/110120tocd

EDN® (ISSN#0012-7515) is published semimonthly, 24 times per year, by UBM Electronics, 11444 W. Olympic Blvd., Los Angeles, CA 90064-1549; 310/445-4200; FAX 310/445-4299. Periodicals postage paid at Los Angeles, CA, and at additional mailing offices. SUBSCRIPTIONS—Free to qualified subscribers as defined on the subscription card. Rates for nonqualified subscriptions, including all issues: US, $150 one year; $250 two years; $300 three years. Except for special issues where price changes are indicated, single copies are available for $10 US and $15 foreign. For telephone inquiries regarding subscriptions, call 763/746-2792. E-mail: EDN@kmpsgroup.com. CHANGE OF ADDRESS—Notices should be sent promptly to PO Box 47461, Plymouth, MN 55447. Please provide old mailing label as well as new address. Allow two months for change. NOTICE—Every precaution is taken to ensure accuracy of content; however, the publisher cannot accept responsibility for the correctness of the information supplied or advertised or for any opinion expressed herein. POSTMASTER—Send address changes to EDN, PO Box 47461, Plymouth, MN 55447. Canada Post: Publications Mail Agreement 40612608. Return undeliverable Canadian addresses to Pitney Bowes Inc, PO Box 25542, London, ON N6C 6B2.

Volume 56, Number 2 (Printed in USA).

Is Crosstalk Degrading Your 10G Multi-Channel IC Design?

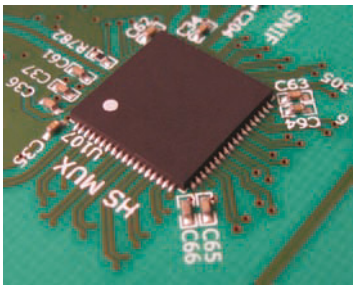

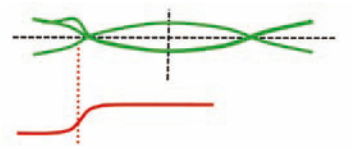

Testing with fixed delay aggressors can result in induced interference outside of the critical receiver sampling time window.

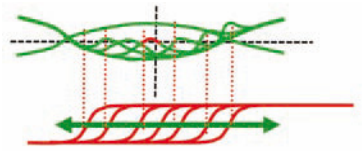

Slewed delay assures finding all crosstalk problems.

20 x 4 data switch...Running at 10 Gb/s...

first silicon due in 3 weeks...only one metal layer mask spin planned. You must find all the problems. You have lined up the test technicians, the lab space, the thermal testing station...

How are you testing crosstalk susceptibility?

If you are only driving aggressor signals on the adjacent pair of unselected ports, you are at risk of missing crosstalk induced bit errors. Yet driving all the inactive ports with single channel BERTs or pattern generators is prohibitively expensive and requires more than a rack full of equipment.

The Centellax PCB12500

parallel channel BERT is a compact, affordable alternative. Select pattern generators or error detectors for any of the five remote heads. Drive the deselected ports with either a clock or PRBS aggressor signal.

The PCB12500's exclusive slewed phase delay independently sweeps each aggressor's transition over a full UI. This creates the "perfect storm" of timing conditions that assures you will find low probability crosstalk susceptibility – those events that testing with fixed delay aggressors will miss.

To find out more about affordable solutions for signal integrity testing visit **www.centellax.com**

With your first silicon arriving in three weeks you had better call now.

CENTELLAX
SIGNAL INTEGRITY MEASUREMENT

Affordability without compromise

See us at DesignCon 2011, Booth #121

BY RON WILSON, EDITORIAL DIRECTOR

Changing times, steady goals

You may have heard, if you follow such things, that UBM (United Business Media) has purchased Canon Communications, the publisher of *EDN*. UBM publishes *EE Times* and *ESD* magazine, among others. This acquisition is part of continuing reorganization in business-to-business publishing in response to a whole range of shifts worldwide. The acquisition set off all manner of rumors in the advertising and public-relations worlds: UBM would close *EDN*, would fold it into the *EE Times* and *Designline* Web sites, or would redesign it to become the industry's first engineer-to-engineer animated comic book.

Well, sorry; as usual, the truth is more prosaic.

To begin with, *EDN* isn't going away. It is still here, and UBM is still committed to both the magazine and its network of Web communities. UBM still believes that working designers deserve technical articles by and for engineers, standing apart from the streams of news and vendor-produced marketing pieces that clamor for your attention. It will continue to be, as its tag line says, the voice of the engineer.

So, what will change? For one thing, I am back at *EDN* after a brief absence. Rick Nelson, after somehow simultaneously managing to guide two major industry journals through some difficult transitions, has returned to being editor-in-chief of just *Test & Measurement World*—rather than both that magazine and *EDN*.

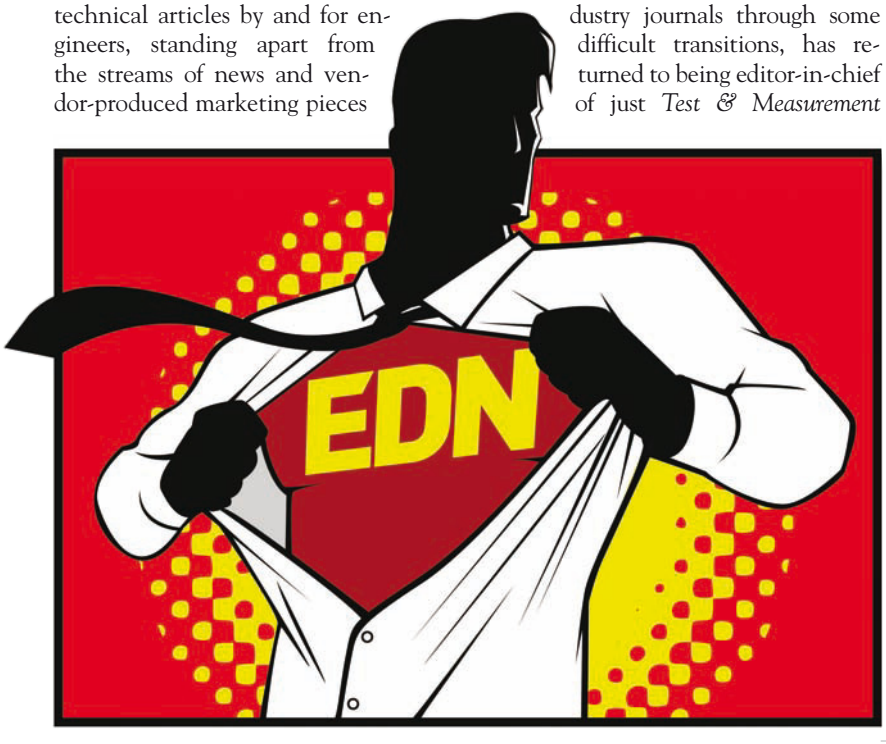

YOGY IKHWANTO/ISTOCKPHOTO.COM

Internal changes have also occurred, although these changes will more slowly become evident. The editorial teams of *EDN*, *Test & Measurement World*, *Design News*, *ESD*/Embedded.com, *EE Times Designlines*, *EELife*, and *EE Times* now all report to a senior vice president for content—and an engineer—Karen Field. For those of you as concerned as we are that advertiser influence should not intrude into editorial decisions, this change is a powerful reassurance.

The new organization also involves more expertise. We will continue to honor the distinction between engineer-written content and content from technically skilled nonengineers. We now have more voices from more technical specializations to call upon, however. We believe you will soon be reading the benefits.

We intend to keep another important distinction, as well. *EDN* has always been about what we might call foundational technologies: analog design; transforming, distributing, and delivering power; digital design; the use of processors and FPGAs; and the process of creating ICs. The material we publish comprises either staff-written or solicited and staff-edited, nonpromotional contributed articles. Some of the *Designlines* have a more application-oriented flavor, focusing on how vendors are approaching design problems in specific application areas. Especially in the print magazine, however, *EDN* will continue to focus on engineers writing about the foundations.

That's it in the proverbial nutshell. We will continue to bring you the best engineering articles from the most authoritative sources, be they working engineers, analysts, or our own staff. And we invite your feedback: This is a closed-loop system, and we need that error signal.

Thank you for being our readers, and welcome to an exciting 2011.EDN

Contact me at ron.wilson@ubm.com.

ANSYS® 13.0

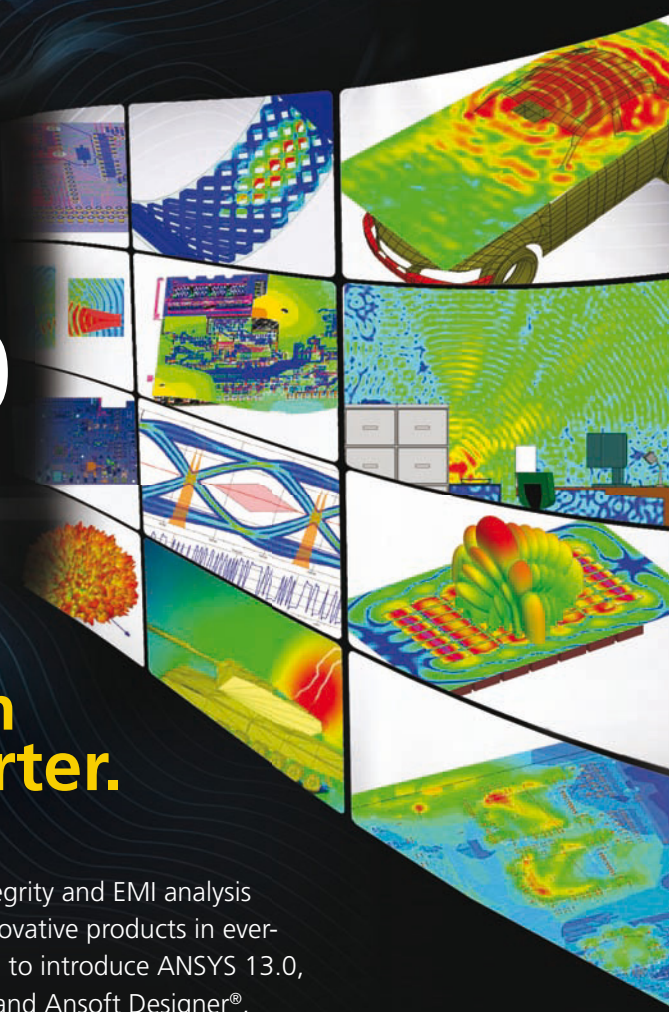

Smart engineering **simulation** is now even **smarter.**

In today's world of smart electronics, signal integrity, power integrity and EMI analysis are mission critical. Engineers must deliver both reliable and innovative products in ever-shrinking development cycles. For this reason, ANSYS is pleased to introduce ANSYS 13.0, our latest release which includes tools such as HFSS™, SIwave™ and Ansoft Designer®, allowing engineers to identify and correct potential design or performance problems before they reach the marketplace.

ANSYS 13.0 can help your design team achieve an even shorter total time to solution with several new and advanced features, including:

- Greater accuracy and solution fidelity for very large electromagnetic structures through the use of new solvers
- Higher productivity through automated 3D electromagnetic model creation, resulting in more time to design and innovate
- Increased computational efficiency, through the use of HPC, resulting in a dramatic speedup of wide-frequency band or parametric simulations

For speed, reliability and accuracy, more product development leaders worldwide trust ANSYS for simulation solutions to help create smarter products in less time.

ANSYS 13.0 **Advanced Fidelity.** **Higher Productivity.** **Leading Performance.**

ANSYS®

For more information, visit www.ansys.com/electronics13
or call us at 1.866.267.9724

Ansoft is now part of the ANSYS platform of simulation solutions.

EDN

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

EDITORIAL DIRECTOR
Ron Wilson
1-415-947-6317;
ron.wilson@ubm.com

MANAGING EDITOR
Amy Norcross
Contributed technical articles
1-781-869-7971;
amy.norcross@ubm.com

MANAGING EDITOR—NEWS
Suzanne Deffree
Electronic Business, Distribution
1-631-266-3433;
suzanne.deffree@ubm.com

SENIOR TECHNICAL EDITOR
Brian Dipert
Consumer Electronics, Multimedia, PCs, Mass Storage
1-916-548-1225;
brian.dipert@ubm.com

TECHNICAL EDITOR
Margery Conner
Power Sources, Components, Green Engineering
1-805-461-8242;
margery.conner@ubm.com

TECHNICAL EDITOR
Mike Demler
EDA, IC Design and Application
1-408-384-8336;
michael.demler@ubm.com

TECHNICAL EDITOR
Paul Rako
Analog, RF, PCB Design
1-408-745-1994;
paul.rako@ubm.com

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

SENIOR ASSOCIATE EDITOR
Frances T Granville, 1-781-869-7969;
frances.granville@ubm.com

ASSOCIATE EDITOR
Jessica MacNeil, 1-781-869-7983;
jessica.macneil@ubm.com

CONSULTING EDITOR
Jim Williams,
Staff Scientist, Linear Technology
edn.editor@ubm.com

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

COLUMNISTS
Howard Johnson, PhD, Signal Consulting
Bonnie Baker, Texas Instruments
Pallab Chatterjee, SiliconMap
Kevin C Craig, PhD, Marquette University

LEAD ART DIRECTOR
Marco Aguilera

ASSOCIATE ART DIRECTOR
Tim Burns

PRODUCTION
Michael Ciardiello,
Director of Premedia Technologies
Jeff Tade,
Production Director
Brian Wu,
Publications Production Manager
Jeff Polman, Derric Treece,
Senior Production Artists
William Baughman, Lucia Corona
Ricardo Esparza,
Production Artists

EDN EUROPE
Graham Prophet,
Editor, Reed Publishing
gprophet@reedbusiness.fr

EDN ASIA
Wai-Chun Chen,
Group Publisher, Asia
waichun.chen@ubm.com
Kirtimaya Varma,
Editor-in-Chief
kirti.varma@ubm.com

EDN CHINA
William Zhang,
Publisher and Editorial Director
william.zhang@ubm.com
Jeff Lu,
Executive Editor
jeff.lu@ubm.com

EDN JAPAN
Katsuya Watanabe,
Publisher
katsuya.watanabe@ubm.com
Ken Amemoto,
Editor-in-Chief
ken.amemoto@ubm.com

EXECUTIVE OFFICERS
Paul Miller,
Chief Executive Officer
Fred Gysi,
Chief Financial Officer
Mike Deering,
Chief Technology Officer
Stephen Corrick,
Vice President/Executive Director
Kevin O'Keefe,
Senior Vice President, Events Division
Roger Burg,
Vice President, Operations
Jason Brown,
Vice President, E-media

UBM Electronics

EDN, 33 Hayden Avenue, Lexington, MA 02421. **www.edn.com. Subscription inquiries:** 1-763-746-2792; EDN@kmpsgroup.com. **Address changes:** Send notice promptly to PO Box 47461, Plymouth, MN 55447. Please provide an old mailing label as well as your new address. Allow two months for the change.

UBM Electronics, 11444 W. Olympic Blvd., Los Angeles, CA 90064-1549; 1-310-445-4200; fax: 1-310-445-4299.

pulse

INNOVATIONS & INNOVATORS

Development kit tackles multi-dc/dc control strings for white and RGB LEDs

As LED-lighting designs become more complex, combining multiple strings of white or RGB (red/green/blue) LEDs, their control and drive circuits must also become more powerful. Texas Instruments' newest LED-development kit shows how you can use the low-cost but powerful Piccolo microcontroller to perform multi-dc/dc control for white and color LEDs. The kit includes separate dc/dc power stages to control two RGB- and two white-LED strings or as many as eight white-LED strings. The kit's C2000 Piccolo microcontroller has the bandwidth and processing power to control all of the power stages, as well as other functions, such as color mixing, PWM (pulse-width-modulation) dimming, thermal sensing and correction, ambient-lighting compensation, and LED-fault protection. TI's development kits use a control board that lets you swap in a new controller board and evaluate different microcontrollers as they become available. The kit also includes a detachable LED panel with white and RGB LEDs.

Other features include eight digitally controlled dc/dc power stages, allowing developers to adjust current levels to control color and brightness of individual strings; two SEPIC (single-ended-primary-inductance-converter) stages that support white LEDs; and six boost power stages that support color mixing for RGB-LED strings by providing individual current control of the red, green, and blue color components. The Piccolo F28027 control card provides 60 MHz of 32-bit performance; eight PWM channels; and a 13-channel, 12-bit ADC to support robust and flexible LED designs.

The device's GUI (graphical user interface) simplifies evaluation by providing simplified control to adjust power-stage current levels for experimentation with brightness and color mixing. Software-configurable maximum current allows for simple field upgrades and offers flexibility to adjust power stages in software as opposed to time-intensive redesign of system hardware. The devices also feature onboard isolated XDS100 USB (Universal Serial Bus) JTAG (Joint Test Action Group) emulation, which simplifies debugging and programming and reduces system cost.

Free, easy-to-use ControlSuite software includes examples for closed-loop average current control of dc/dc power LED-driver stages and detailed lab documentation of software structure and performance. The kit sells for $499. **—by Margery Conner**

▷**Texas Instruments**, www.ti.com.

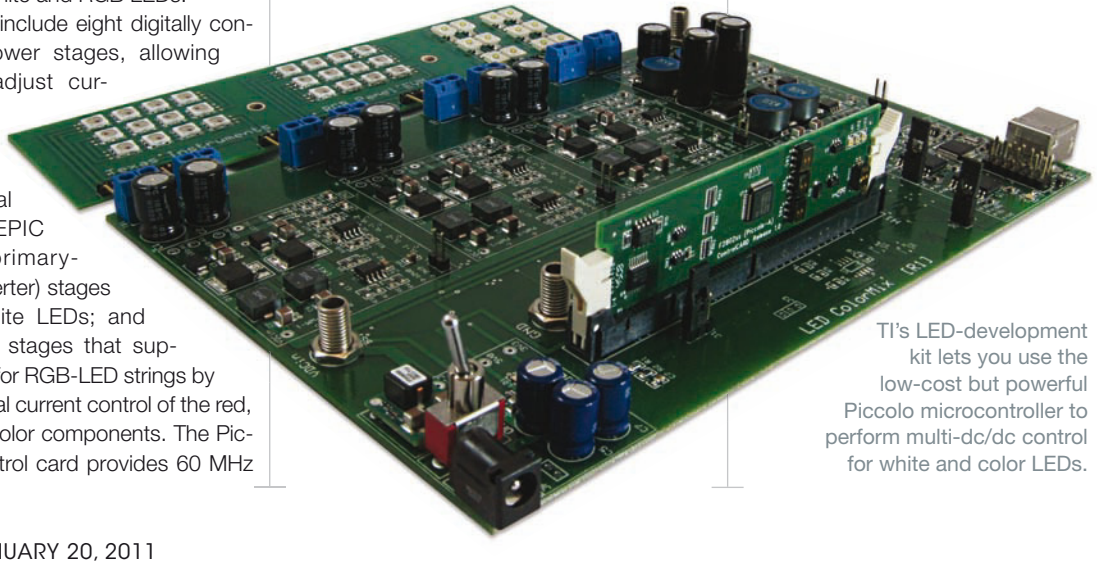

TI's LED-development kit lets you use the low-cost but powerful Piccolo microcontroller to perform multi-dc/dc control for white and color LEDs.

TALKBACK

"While a leap of imagination is still required, reading a mind may not be that far off ... for good or for ill. Who protects us all from misuse? Big Brother? Obama? I think not!!!"

—Researcher "HG," in *EDN*'s Talkback Section, at http://bit.ly/ekrbxz. Add your comments.

Move Your Automated Test Beyond the Box

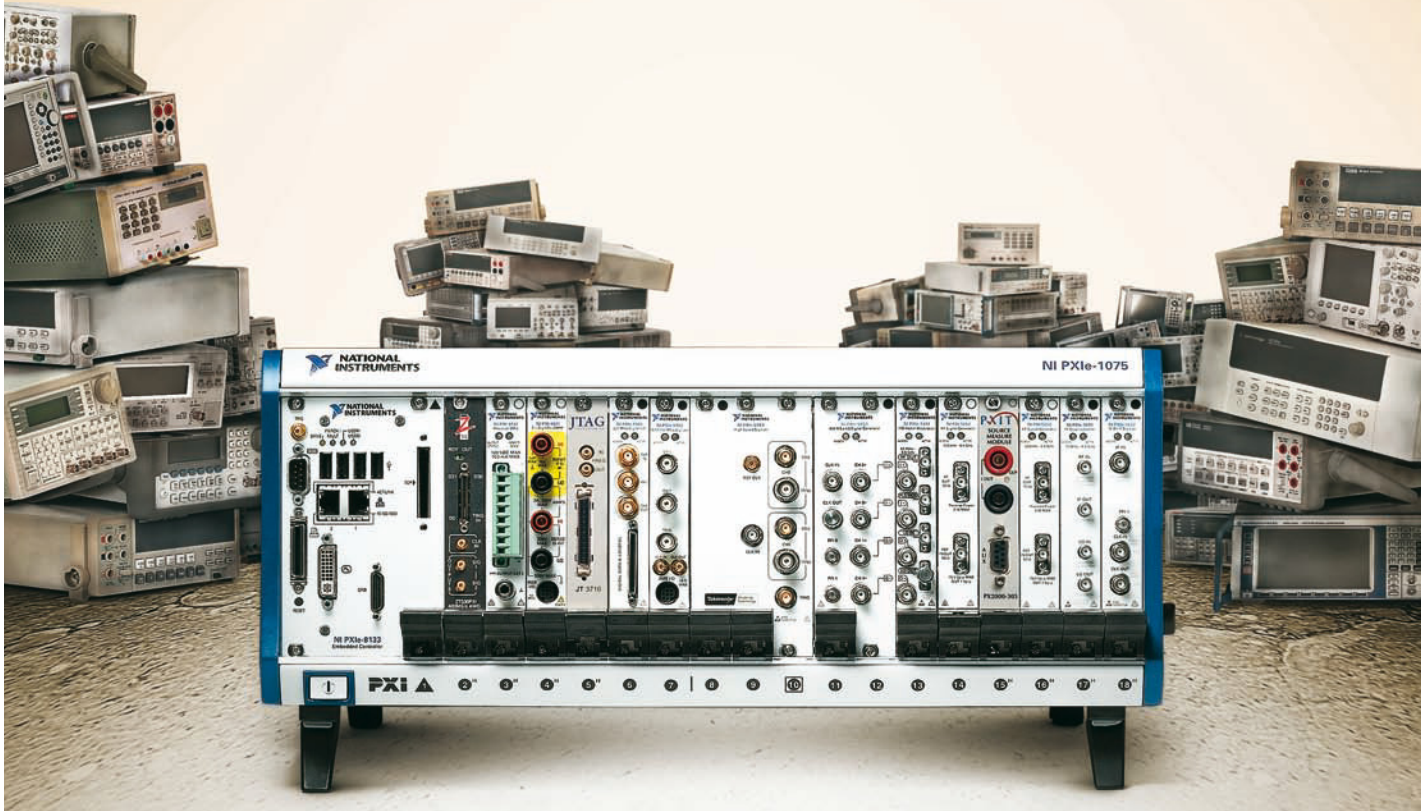

Engineers around the world are making the software-defined PXI platform the cornerstone of their test system architectures. With more than 1,500 modular instruments available from more than 70 vendors, PXI delivers the functionality and flexibility you need to build a better test system while reducing cost and size.

PRODUCT PLATFORM

PXI modular instrumentation

NI LabVIEW graphical software

NI TestStand software

>> Learn how PXI can help you at **ni.com/beyond**

800 891 8841

NATIONAL INSTRUMENTS™

Ambarella augments image-processing foundation with increasing integration

Ambarella last month unveiled its latest-generation image-processing platform, iOne, which includes a dual-core ARM Cortex-A9 system-processing nexus. Ambarella reports that its customers are increasingly embedding full-blown operating-system builds—specifically, Google's Android, due to its open-source and no-licensing-fee attributes—in their system designs. As such, iOne's CPU integration negates the need to provide separate IC-system-processor facilities.

Ambarella also equipped iOne with an Imagination Technologies (www.imgtec.com) PowerVR SGX540 graphics core, along with incremental and improved system-interface ports versus those in the previous-generation A5S and A7.

Notably, iOne for the first time embeds a SATA (serial-advanced-technology-attachment) transceiver along with a GbE (gigabit-Ethernet) MAC (media-access controller), whose bandwidth potential will come in handy in high-definition networked-camera applications. The product also includes a 32-bit DDR3 SDRAM interface, which combines high bandwidth potential with low bus pin count and finer-grained aggregate memory requirements versus a 64-bit bus alternative.

The iOne image-processing platform delivers an A5S++-class hybrid-camera function.

Like its precursors, iOne includes an ARM11 core to coordinate the activities of the proprietary-architecture still-image and video-DSP subsystems. However, Ambarella built iOne on the same 45-nm process as its predecessors: the A5S, which it introduced at the 2010 CES (Consumer Electronics Show), and A7, which it rolled out in September. The company fabricated the initial A5, which it unveiled during the 2009 CES, on a 65-nm lithography. To avoid a reticle-busting die size with requisite cost and yield negative impacts, Ambarella outfitted iOne with a subset of the A7's DSP resources, albeit still labeling the device as delivering an "A5S++-class hybrid-camera function."

Ambarella's iOne is currently available for sampling to lead customers, with production capability slated for early in the second quarter of this year. It costs $25 (volume quantities). For those of you hoping to catch the 3-D wave; the chip embeds dual image-sensor interfaces and therefore integrates the left-plus-right-eye, dual-stream-combining function that the company's stand-alone S3D Full HD camera preprocessing previously handled. **—by Brian Dipert**

▷**Ambarella**, www.ambarella.com.

Chinese company increases OLED lifetime

OLEDs (organic light-emitting diodes) have several advantages over traditional LEDs: OLEDs can be less expensive to manufacture; can serve as rugged, colorful displays for handheld devices, monitors, and TVs; and can potentially serve as large-surface-illumination/lighting devices. Imagine wallpaper that also illuminates a room.

Despite those features, OLEDs have relatively short lifetimes. For example, the Osram (www.osram.com) Orbeos OLED lasts for approximately 500 hours, compared with LED lifetimes of 20,000 hours and beyond.

Targeting this problem, Chinese OLED company Visionox and Tsinghua University (www.tsinghua.edu) have teamed up to develop a technology to prolong the lifetime of OLED panels. According to *China Daily*, the company has developed a structure that delivers emission through a composite light-emitting layer that provides a 20-fold increase in the operational life of OLEDs. Visionox claims that current commercially produced high-definition screens have a 450- to 1000-cdm luminance, whereas the new technology can provide a lifetime of as much as 100,000 hours at a brightness of 1000 cdm.

> The technology provides an OLED lifetime of 100,000 hours.

—by Margery Conner

▷**Visionox**, www.visionox.com.

DILBERT By Scott Adams

RAQ's

Rarely Asked Questions

Strange stories from the call logs of Analog Devices

Op Amp Dedicated Feedback Pin Helps Achieve New Levels of Performance.

or

This feedback pin is dedicated to all those op amps out there tonight!

Q: I've noticed a new pinout on some of your high-speed op amps. Why the change after all these years?

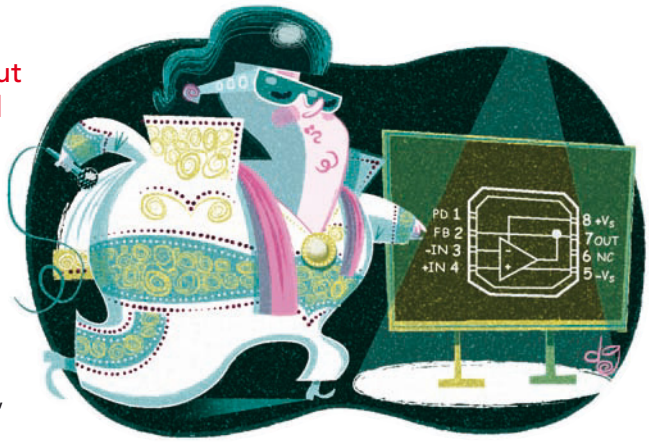

A: You're right. A few years ago, Analog Devices introduced a new pinout on some of our high-speed, high-performance amplifiers. This new pinout features a "dedicated feedback" pin, as Elvis is pointing out. Although the traditional SOIC pinout, which has been around for many years, performed admirably in the past, it imposes limitations as speeds continue to increase.

In a traditional SOIC pinout, pin 2 is the inverting input, pin 3 is the noninverting input, pin 4 is the negative supply, pin 6 is the output, pin 7 is the positive supply, and pins 1, 5, and 8 are typically no connects. But therein lies a problem: mutual inductance between pins 3 and 4 degrades the amplifier's second-harmonic distortion. To remedy this we rotated the pinout counter-clockwise by one pin on our LFCSP packages, breaking the coupling between the noninverting input and the negative supply. This change can improve second-harmonic distortion by up to 14 dB—a 500% improvement!

On some of our SOICs, we also provide a dedicated feedback pin, on pin 1. While this does not improve second-harmonic distortion, it greatly simplifies the circuit layout and reduces parasitic effects that can be detrimental to high-speed applications. In high-speed circuits, layout plays a large role in circuit performance. With a traditional op amp pinout, getting the feedback signal to the inverting input requires it to be routed around or under the amplifier, with both options adding parasitics. With the dedicated feedback pinout, the feedback pin is right beside the inverting input, so all that is required to make the connection between the two pins is a resistor or a trace. The dedicated feedback pinout enables a very compact layout that requires less board area. It also reduces parasitics, which improves high-speed performance, and provides a more streamlined layout, which improves signal routing.

It's clear that a small, but innovative change like Analog Devices' dedicated feedback pin can make a big difference in the performance of your next high-speed amplifier design.

To Learn More About High-Speed Amplifiers
http://dn.hotims.com/34920-101

Contributing Writer John Ardizzoni is a Senior Application Engineer at Analog Devices in the High Speed Linear group. John joined Analog Devices in 2002, he received his BSEE from Merrimack College in N. Andover, MA and has over 30 years experience in the electronics industry.

Have a question involving a perplexing or unusual analog problem? Submit your question to: www.analog.com/askjohn

For Analog Devices' Technical Support, Call 800-AnalogD

EDA tool accelerates place and route for SOCs

Magma Design Automation claims that its new Talus Vortex FX IC-implementation tool is the first to employ distributed computing for acceleration of placement and routing in SOCs (systems on chips). The company has also announced Talus 1.2, which integrates routing, timing analysis, and parasitic extraction to triple flat-block capacity to 3 million gates from the 1 million-gate capacity of Talus 1.1.

According to Bob Smith, vice president of product marketing at Magma, Talus Vortex FX's Distributed Smart Sync technology provides a threefold performance boost onto the improvements in standalone multithreaded capability of Talus 1.2 on a single server. On its own, the new Talus 1.2 system enables engineers to implement 1 million cells per day and perform crosstalk avoidance to detect and correct crosstalk violations during optimization and implementation, advanced on-chip variation to ensure tight timing correlation throughout the flow, and MMMC (multimode/multicorner) analysis to manage multiple timing scenarios.

Talus 1.2 is currently in use for 28-nm designs at five large semiconductor companies, according to Smith, and the company has silicon-proved the product in tapeouts at the 40-nm node.

The Talus MX timing-and-extraction engines borrow technology from Magma's Tekton timing analyzer and QCP sign-off extractor. It also features the Talus MX Router with enhanced global, tracking, and detailed routing capabilities.

The Magma Hydra design-partitioning and floorplanning tool provides the design input to Talus Vortex FX, which Magma licenses separately, to manage the distribution of parallel Talus place-and-route jobs across typically as many as 10 to 12 servers. To leverage the advantages of distributed computing, you must ensure that each server has an individual license available for the Talus 1.2 product. A typical hardware configuration uses servers with four to eight processor cores and 64 Gbytes of memory. During execution, Vortex FX synchronizes and aggregates the results of each Talus process to assemble the complete SOC.—**by Mike Demler**

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

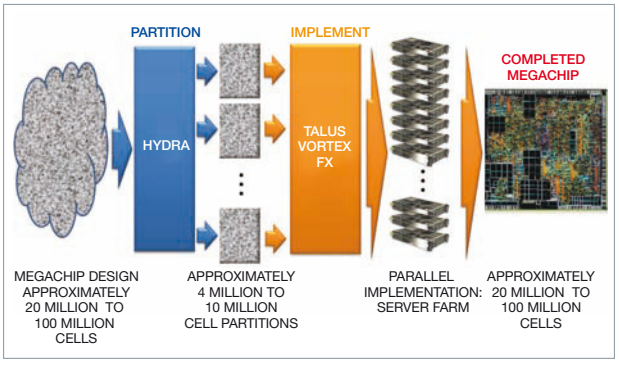

The Talus Vortex FX IC-implementation tool employs distributed computing for acceleration of placement and routing in SOCs.

400-Gbyte MLC solid-state drives target enterprises

Samsung Electronics Co Ltd is offering 100-, 200-, and 400-Gbyte MLC (multilevel-cell) solid-state drives for sampling. The drives target use as primary storage in enterprise systems. The devices use 30-nm-class MLC NAND-flash chips with a toggle DDR interface and a controller that uses a 3-Gbps SATA (serial-advanced-technology-attachment) interface. Samsung claims that the drives can process random read commands at 43,000 IOPS (inputs/outputs per second) and random writes at 11,000 IOPS. The company compares that performance to a 15,000-rpm hard-disk drive, which has an IOPS rate of 350, amounting to a 120-fold gain in random-IOPS-read performance and a 30-fold gain in random-IOPS-write performance.

The solid-state drives have a 150-times-higher IOPS/watt rate compared with 15,000-rpm hard-disk drives, making them able to process 150 times more data at the same power consumption, the company claims. The drives also feature an end-to-end data-protection function with an advanced data-encryption algorithm for drive reliability and security.

The solid-state-drive market is increasing at a solid pace, according to market-research-company projections. Gartner Inc (www.gartner.com) estimates that shipments of solid-state drives for servers and enterprise-storage systems will increase to 6.3 million units in 2014 from 324,000 units in 2009. Revenue for the solid-state-drive enterprise market should grow more than seven times from $485 million to $3.6 billion during the same period.

"As more and more server makers are adopting solid-state drives for use in eco-friendly platforms that consume less electrical power, the need for high-density drives in the server market is growing rapidly," says Byungse So, senior vice president for the memory-product planning and application-engineering team at Samsung Electronics. "We are now expanding our lineup to include high-density solid-state drives using MLC NAND-flash memory."

With the new drives, Samsung widens its range of densities to include 2.5-in., 50-, 60-, 100-, and 120-Gbyte drives using SLC NAND-flash memory and 2.5-in., 100-, 200-, and 400-Gbyte drives using MLC NAND-flash memory. The company also has 3.5-in., 100- and 200-Gbyte SLC drives. Samsung will begin mass-producing the MLC-based enterprise drives this month.—by Suzanne Deffree

▶Samsung Electronics, www.samsung.com.

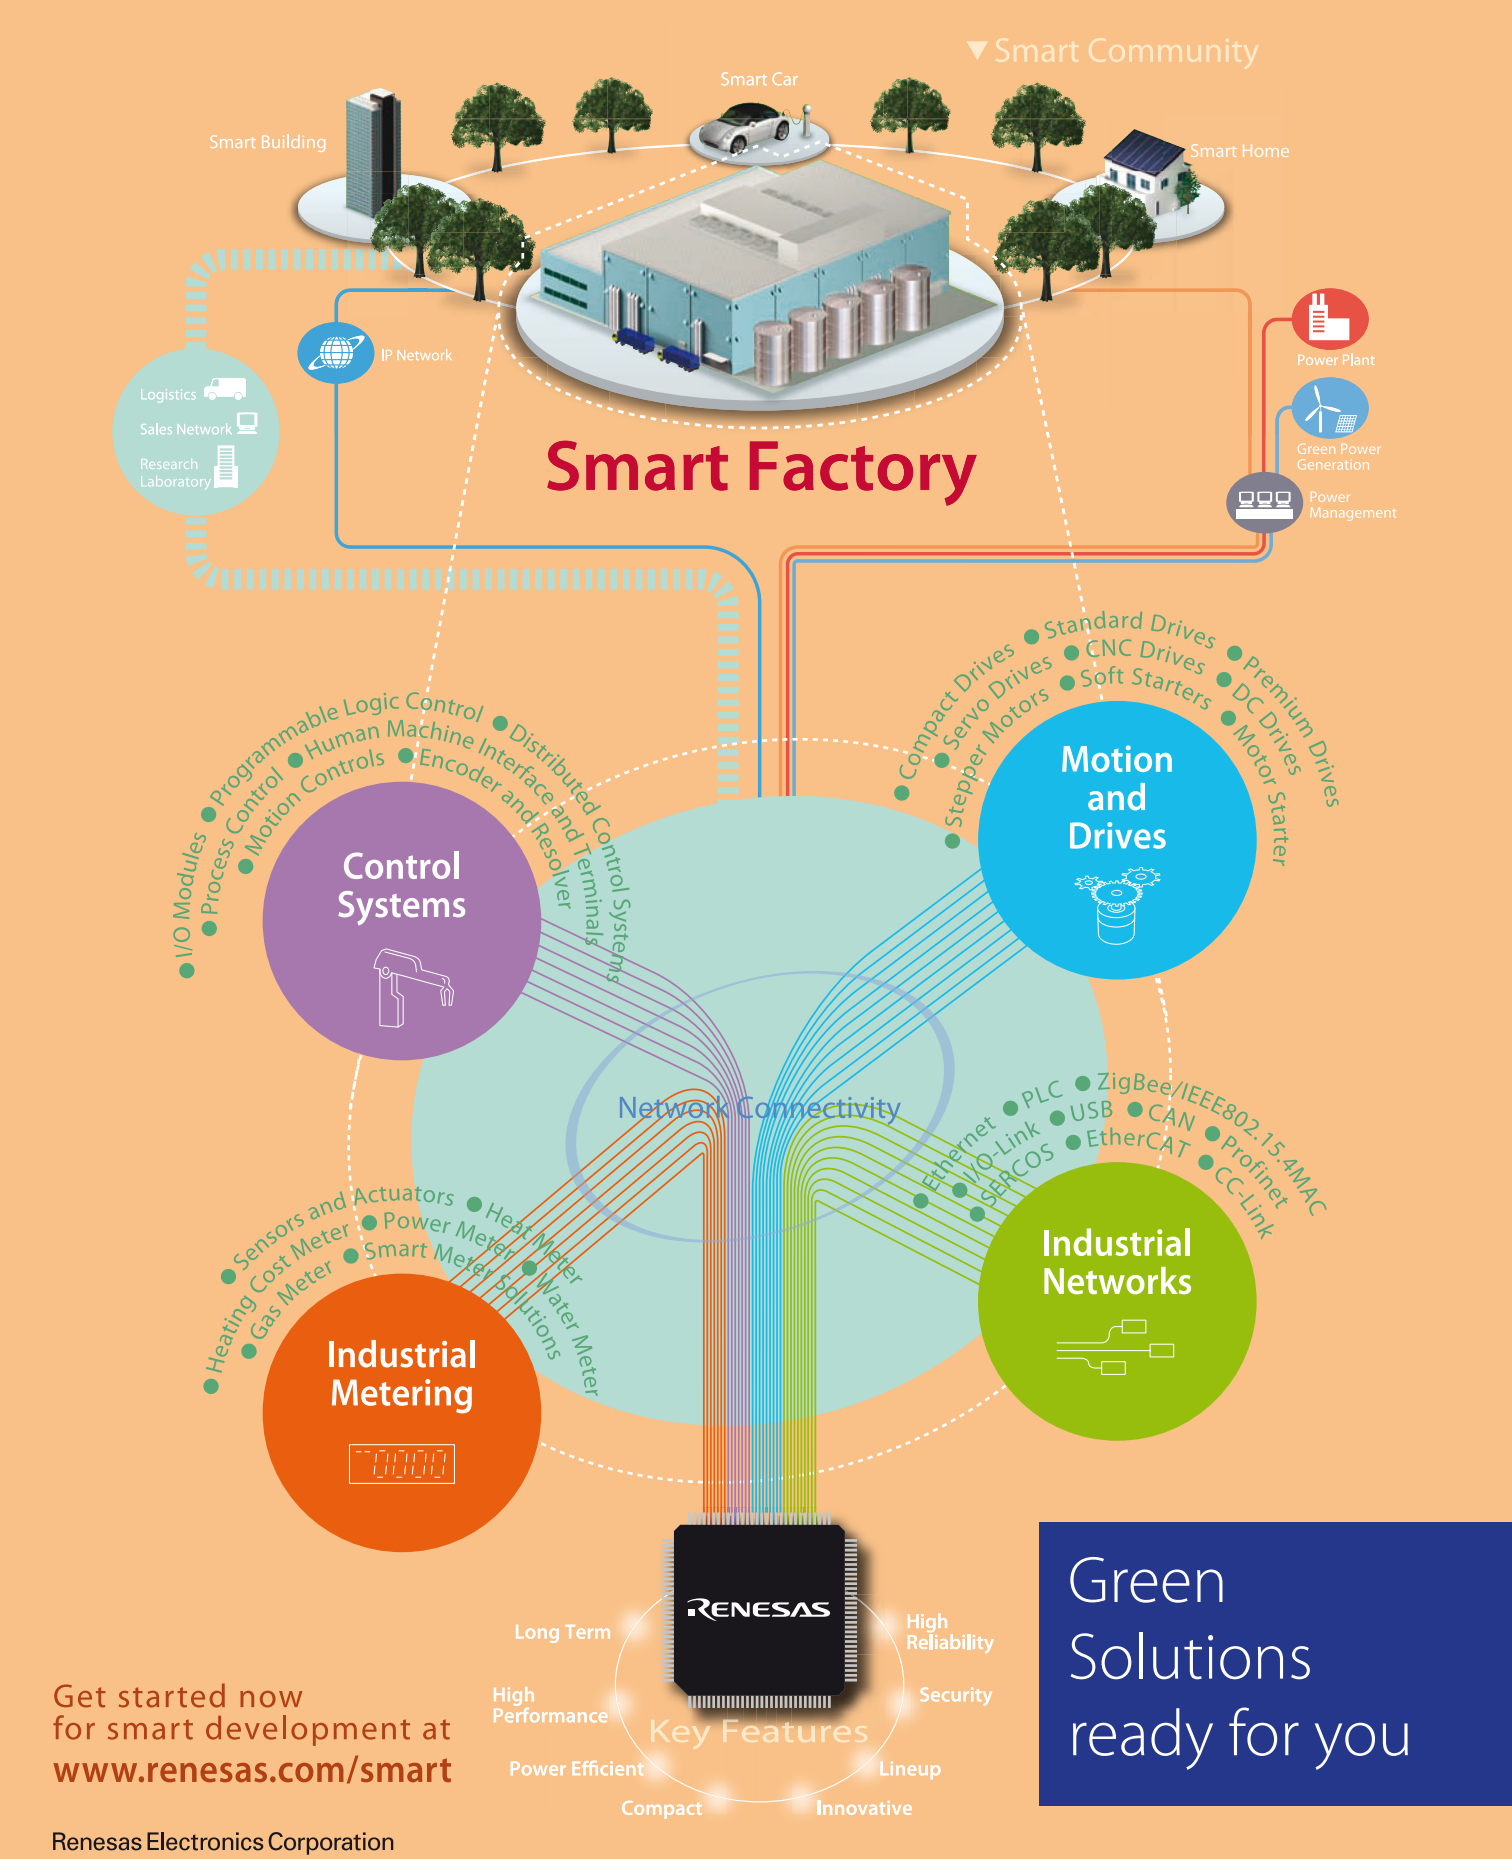
▼Smart Community
Smart Car
Smart Building
Smart Home
IP Network
Logistics
Sales Network
Research Laboratory
Power Plant
Green Power Generation
Power Management
Smart Factory
Programmable Logic Control
Human Machine Interface and Terminals
Distributed Control Systems
Encoder and Resolver
Motion Controls
Process Control
I/O Modules
Control Systems
Compact Drives
Standard Drives
Servo Drives
CNC Drives
Premium Drives
Stepper Motors
Soft Starters
DC Drives
Motor Starter
Motion and Drives
Network Connectivity
Sensors and Actuators
Heat Meter
Power Meter
Water Meter
Heating Cost Meter
Smart Meter Solutions
Gas Meter
Industrial Metering
PLC
ZigBee/IEEE802.15.4MAC
USB
CAN
Ethernet
Profinet
I/O-Link
EtherCAT
CC-Link
SERCOS
Industrial Networks
RENESAS
Long Term
High Reliability
High Performance
Security
Key Features
Power Efficient
Lineup
Compact
Innovative
Green Solutions ready for you
Get started now for smart development at
www.renesas.com/smart
Renesas Electronics Corporation
RENESAS

Microcontrollers operate as fast as 120 MHz

STMicroelectronics has announced the Cortex-M3 F-2 series of microcontrollers, which add more than 30 pin- and software-compatible parts to the STM32 family.

The company has also revealed its plans for the ARM Cortex-M4 and -M0 and plans to make products available for sampling starting in midyear and for production by the end of this year.

The company manufactures the STM32 F-2 in a 90-nm CMOS process, and it operates at clock speeds as high as 120 MHz. The devices include on-chip voltage regulators to convert external supplies of 1.65 to 3.6V for internal 1.2V core operation. The product portfolio offers a range of 64- to 176-pin packages, including LQFP64, LQFP100, LQFP144, and UFGBA176, along with the WLCSP64, which measures less than 4×4 mm.

The STM205 includes one USB (Universal Serial Bus) port and 128 kbytes to 1 Mbyte of flash memory, and the STM207 adds a second USB port and an Ethernet port. It also includes an 8- to 14-bit parallel camera interface that supports data rates as high as 48 Mbytes/sec. The STM215 and STM217 include 512 kbytes of flash memory and add a cryptography/hash processor to implement 3DES (Triple Data Encryption Standard), AES (Advanced Encryption Standard) 256, and SHA (secure hash algorithm) 1 encryption. Three 2M-sample/sec ADCs are available, with an on-chip temperature sensor and a two-channel, 12-bit DAC.

The company's ART (adaptive-real-time) memory accelerator reduces program-branch-execution delay by saving as many as 64 branches in a cache memory. The 128-bit-wide flash-memory interface enables simultaneous fetching of multiple 8-bit instructions and provides a 128-bit-wide prefetch queue for data.

> The adaptive-real-time memory accelerator reduces program-branch-execution delay by saving as many as 64 branches in a cache memory.

The memory resources include 4 kbytes of battery-backed SRAM and 528 bytes of OTP (one-time-programmable) memory for storage of critical user data, such as Ethernet MAC (media-access-control) addresses or cryptographic keys.

An evaluation board, the STM3220G-Eval, is available as a vendor-neutral platform and comes with no IDE (integrated device electronics) or JTAG (Joint Test Action Group) cable, supporting a choice of development environments.

The STM Cortex-M3 F-2 series is now available for sampling for OEMs and will be available for production during the first quarter of this year. Prices for the F-2 series start at $3.18 (10,000) for the ST-205 with 128 kbytes of flash memory.

—by Mike Demler

▷**STMicroelectronics**, www.st.com/stm32.

The Cortex-M3 F-2 series of microcontrollers adds more than 30 pin- and software-compatible parts to the STM32 family.

GARTNER PREDICTS INDIAN PC SHIPMENTS TO GROW 25% IN 2011

PC demand is growing beyond the largest cities in India and will see shipments rise nearly 25%, according to a recent forecast from Gartner Inc. The market-research company has estimated that Indian PC shipments will total 13.2 million units in 2011, a 24.7% increase from 2010 levels. "There is an increased aspiration to own a PC in India," says Vishal Tripathi, principal research analyst at Gartner. "The younger generation is contributing to the large-scale adoption of PCs in India in a significant way. Besides reading and surfing the Internet, the usage of PCs for watching movies, listening to music, and gaming is playing a pivotal role in driving demand."

Gartner notes that vendors are seeing an increasing demand from smaller cities in India in which PC penetration has grown considerably. The company attributes that growth to rising income and declining PC prices.

The research company predicts that mobile PCs will outnumber desktops in India by the fourth quarter of this year. Still, the total number of desktop shipments should be higher than those of mobile PCs. In 2011, the desktop-PC market will grow 5%, totaling 7.2 million units, and mobile-PC shipments will grow 61% with 5.9 million units. "The Indian PC market is far from reaching maturity, unlike the other developed countries," says Tripathi. "India contributes significantly to the overall growth of PC shipments worldwide."—by Suzanne Deffree

▶Gartner, www.gartner.com.

Integrated Device Technology – Leader in PCI Express® Solutions

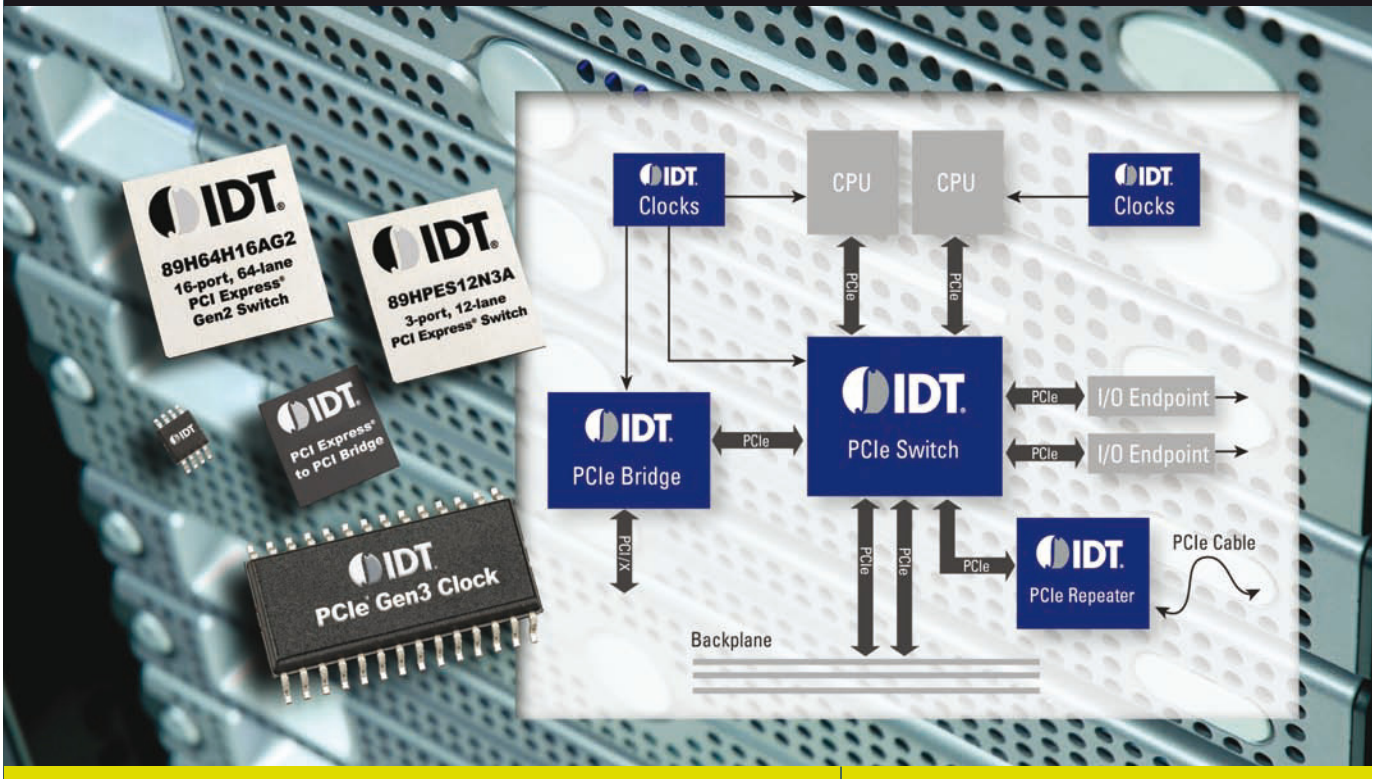

PCI Express Switches
Industry's most comprehensive family of high-performance, scalable PCIe switching solutions
- Gen 1 and Gen 2 Portfolio
- Up to 64 lane and 24 port devices
- Highly flexible port configurations
- Unprecedented 8 non-transparent bridging (NTB) functions to enable multi-root applications

PCI Express Repeaters
5Gbps Gen 2 PCIe Signal Repeaters
- Four and eight channel signal repeaters
- Compensates for cable and PCB trace attenuation and ISI jitter
- Programmable receiver equalization to 30dB and de-emphasis to -8.5dB
- Leading edge power minimization in active and shutdown modes
- No external bias resistors or reference clocks required
- Channel mux mode, demux mode, 1 to 2 channels multicast, and Z-switch function mode

PCI Express Bridges
High performance PCIe bridging to legacy PCI and PCI-X protocols
- Ultra low power version for consumer applications
- Forward mode buffer optimization
- Only PCI Express bridges with Short Term Caching for significant improvement in PCI Read performance
- Pin compatibility with competitive offering for dual sourcing solution

PCI Express Clocking Solutions
Industry's broadest offering of PCIe Gen1, Gen2 and Gen3 clock generation and buffering solutions
- World's First PCI Express Gen3 Family of Timing Devices
- Clock Generators and Synthesizers
- Spread Spectrum Clock Generators
- Fanout Buffers and Muxes
- Jitter Attenuators
- Frequency Translators
- Zero Delay Buffers

IDT®
Integrated Device Technology

The Analog and Digital Company™

For more information visit us online at www.IDT.com

VOICES

Samsung Electro-Mechanics' Jong-Woo Park: targeting smart-grid, electric-vehicle, and biotechnology applications

Samsung Electro-Mechanics has grown continuously through constant investment and R&D since its establishment in 1973 to become one of the leading electronics-parts manufacturers, according to Chief Executive Officer Jong-Woo Park. He says that the company has continued to break quarterly sales records since last year, becoming the fifth-largest electronics-parts maker. Sales reached approximately 5.5 trillion Korean won on a consolidated basis in 2009 and approximately 3.5 trillion Korean won in the first half of 2010. *EDN* recently conducted an e-mail interview with Park, a portion of which follows. For the complete version, go to www.edn.com/110120pa.

What product lines does Samsung Electro-Mechanics offer?

A Samsung now produces essential parts and components for mobile phones, TVs, computers, and other advanced IT products. For mobile phones, Samsung makes mobile-phone substrates, camera modules, precision motors, and MLCCs [multilayer ceramic capacitors]. In addition, Samsung makes components for display devices, including power modules, network modules, digital tuners, and MLCCs, and the company makes components for computers, including network modules, semiconductor substrates, and MLCCs.

What is the significance of the MLCCs?

A MLCCs and PCBs [printed-circuit boards] have greatly contributed to corporate growth in recent years. MLCCs are essential parts for all kinds of electronic products, including mobile phones, PCs, and TV sets. The device controls the flow of current in electric circuits.

How have PCBs helped Samsung's business?

A In the PCB-for-semiconductor segment, we have maintained the top share of the world market since 2007. This [achievement] has been made possible by our launching new models, such as one with the world's thinnest substrate, at the right time and securing stable product yields. The semiconductor substrate is an auxiliary part that acts as a bridge between the semiconductor and the motherboard. Semiconductors, which have high density, cannot be directly attached to the main PCB in the way that condensers, resistors, and other passive parts can. Thus, semiconductor substrates are essential for transferring semiconductor signals to the motherboard.

Could you comment on your supply-chain-management capabilities?

A Since I was appointed chief executive officer in January 2009, we have built a companywide SCM [supply-chain-management] system to prevent inefficiency and realize smart management. SCM is a companywide innovation program designed to ensure operational predictability by seamlessly connecting all parts, from customers to suppliers. The SCM solution guarantees just-in-time production and reduces resource waste, raising profitability.

Of course, even though we build an SCM system, it is of no use if we do not use it properly. Top management must be the first to understand the system and apply it in everyday operations. Then, the rest of the work force must all follow the rules and processes that have been laid out to achieve the expected results. This [approach] will also allow us to work smarter.

What are your plans for the future?

A We will join the ranks of the top-tier players by 2015 by improving the way individual employees do their jobs and realizing even greater changes within the company. Samsung is now expanding the weight of value-added products in the current IT-based product portfolio. At the same time, internal competitiveness is being strengthened and new technology developed to advance into new business areas.

To this end, we plan to grow as a top-tier maker of MLCCs, substrates, and power supplies—all core business areas. At the same time, we will take advantage of our core technologies in materials, optics, and radio-frequency devices to cultivate new growth opportunities in the smart grid, EV (electric-vehicle) power-supply, EV traction-motor, biotechnology, and health-care fields.

Can you elaborate on those applications?

A In the smart-grid sector, Samsung is now involved in power conversion—photovoltaic inverters and smart meters for remote measuring—and module services. The company is also participating in various test-bed projects organized by the Korean government. We also plan to expand into drive motors, electric-power-control systems, quick chargers, and parts for automotive electrical systems by leveraging our technologies related to power supplies and motors.

As for biotechnology, we are now promoting cell chips for toxicity monitoring and drug-releasing systems by applying MEMS [microelectromechanical-system] technology obtained through our ink-jet-printer business. The future strategy of Samsung is to expand into energy, environment protection, EV parts, and biotechnology by taking advantage of convergence and modularization.

—interview conducted and edited by Rick Nelson

The World's Standard.

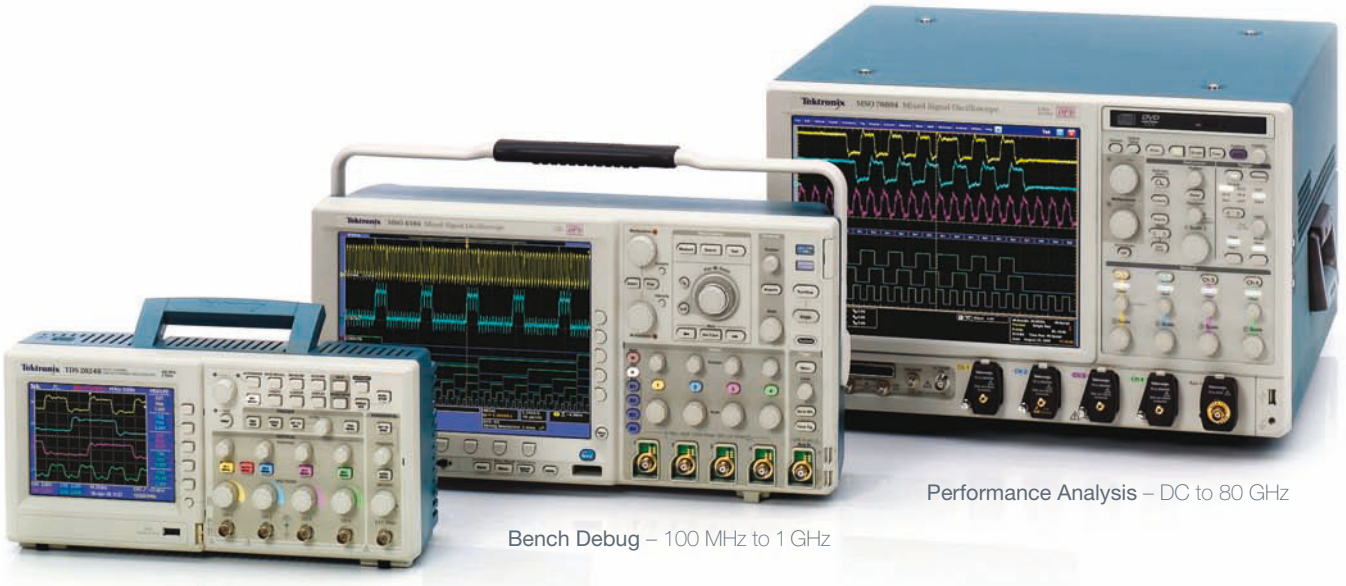

Performance Analysis – DC to 80 GHz

Bench Debug – 100 MHz to 1 GHz

Basic Visualization – 40 to 500 MHz

Oscilloscopes from Basic to Bench and Beyond.

Whatever your need, there is a Tektronix oscilloscope to fill it. Tektronix offers the world's broadest portfolio of oscilloscopes to debug and test tomorrow's designs, today. Start with Tektronix Basic oscilloscopes: fast, familiar and affordable; they're what you know and trust. Our Bench oscilloscopes provide next-level productivity, with the feature-rich tools you need to debug today's complex mixed signal designs. And for those who demand fast, flexible, in-depth PHY layer analysis, our Performance oscilloscopes offer the industry's best signal fidelity, verification and characterization capabilities to help you shorten your design cycles.

No wonder **8 of 10 engineers worldwide trust Tektronix** to help them bring advanced designs to market on time and on budget. Wherever you're going, we'll get you there.

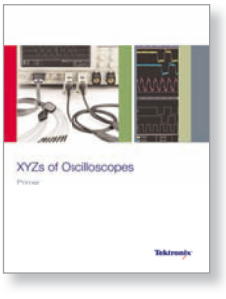

See why Tektronix is the World's Standard in Oscilloscopes and download your free Oscilloscope Primer.

www.tektronix.com/worldstandard

Tektronix®

BY BONNIE BAKER

Remote photo sensing

Photodiodes transform a basic physical occurrence, light, into an electrical form, current. Design engineers methodically convert the photodetector's current to a usable voltage, which makes the manipulation of the photodiode signal manageable. There are many ways to approach the photo-sensing-circuit problem, but one issue came to mind as I read the comments in the Talkback section for a recent column (**Reference 1**). A reader requested a circuit that reduces the noise impact of a photodiode with a large parasitic capacitance.

A classic photo-sensing system circuit has a photodiode, an operational amplifier, and a feedback resistor/capacitor pair at the front end. Recalling a circuit from another column, in this circuit, the photodiode, amplifier, and feedback-capacitance elements limit the bandwidth of the circuit (**Reference 2**). Yet another column details the stability of this circuit (**Reference 3**).

When sensing with a photodiode with a large parasitic capacitance or from a remote site, the input of the amplifier has a large capacitance across its input. The result of this added capacitance increases the circuit's noise gain unless you increase the amplifier's feedback capacitor. If the feedback capacitor, C_F, increases, the bandwidth of the circuit decreases.

To fix this problem, you can use a bootstrap circuit (**Figure 1**). Photodiodes with a relatively low diode capacitance do not benefit from this circuit. A unity-gain buffer, A_2, removes the cable's capacitance and the photodiode's parasitic capacitance from the input of the transimpedance amplifier, A_1.

When designing this circuit, you'll find that the type of amplifier you select for A_2 is somewhat easy. The only important performance specifications are low input capacitance, low noise, a wider bandwidth than A_1, and low output impedance.

In this design, A_2's input capacitance is the only capacitance that plays a role in the ac-transfer function of the transimpedance system. The input capacitance of the buffer replaces the summation of the input capacitance of A_1, the cable capacitance, and the parasitic capacitance of the photodiode. A good rule of thumb is to have $C_{A2}<<(C_{A1}+C_{CA}+C_{PD})$, where C_{A1} and C_{A2} represent the sum of their input differential and common-mode capacitance.

With this design, you exchange one noise problem, A_1, with another, A_2. The unity-gain buffer removes the noise effect from A_1. A good approach is to make the noise of A_2 less than or equal to that of A_1.

The difference between the input signal and the output signal in this system falls across the cable/diode capacitance. You can keep this difference low by selecting A_2 with wider bandwidth than A_1 and keeping A_2's output impedance low. A_2's gain roll-off introduces an upper limit for the bandwidth improvement, making A_2's bandwidth much greater than that of A_1. This circuit requires stability optimization as you balance C_F and the input capacitance of A_2 (**references 4** and **5**).EDN

A_2's input capacitance is the only capacitance that plays a role in the ac-transfer function.

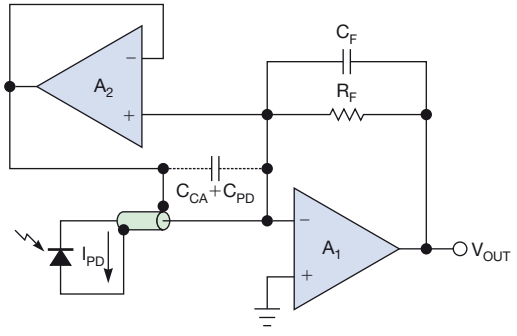

C_{PD}: PHOTODIODE PARASITIC CAPACITANCE
C_{CA}: CABLE CAPACITANCE

Figure 1 You can use bootstrapping to eliminate diode capacitance and cable capacitance from the transimpedance-design problem.

REFERENCES

1. Baker, Bonnie, "A good holiday-season project," *EDN*, Dec 15, 2010, pg 18, http://bit.ly/fG59Ly.
2. Baker, Bonnie, "The eyes of the electronic world are watching," *EDN*, Aug 7, 2008, pg 24, http://bit.ly/i5lr65.
3. Baker, Bonnie, "Transimpedance-amplifier stability is key in light-sensing applications," *EDN*, Sept 4, 2008, pg 24, http://bit.ly/hG4bNW.
4. Graeme, Jerald G, *Photodiode Amplifiers: Op Amp Solutions*, McGraw-Hill, 1996, ISBN 0-07-024247-X.
5. Kurz, Dov, and Avner Cohen, "Bootstrapping reduces amplifier input capacitance," *EDN*, March 20, 1978.

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

Who makes the fastest real-time oscilloscopes?

20 MHz - 40 MHz

100 MHz - 200 MHz

100 MHz - 500 MHz

100 MHz - 1 GHz

60 MHz - 200 MHz

100 MHz - 1 GHz

The fastest-growing oscilloscope company.*

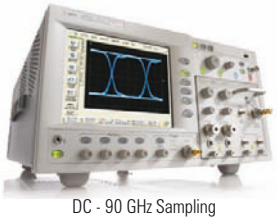

DC - 90 GHz Sampling

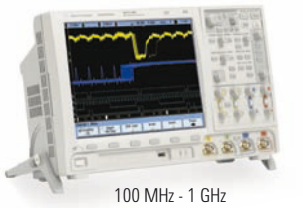

100 MHz - 1 GHz

2.5 GHz - 13 GHz

Introducing 16-32 GHz
Agilent Infiniium 90000 X-Series

600 MHz - 4 GHz

Our portfolio offers families engineered to deliver the best:

- Best measurement accuracy
- Broadest measurement capability
- Best signal visibility
- More scope than you thought you could afford

Are you using the best scope?
Take the 5-minute scope challenge and find out.
www.agilent.com/find/scopecheck

*Prime Data 2009 Market Growth Analysis.

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

Agilent Technologies

PRY FURTHER AT EDN.COM

Go to www.edn.com/pryingeyes for past Prying Eyes write-ups.

The Schauer TB10012 battery charger

This charger worked fine for 30 years. When it stopped working a couple of years ago, I pried into its interior to see what was wrong. It turns out that the failure was purely mechanical. I must have dropped the unit, which dislodged the copper PCB (printed-circuit board) from its isolated mounting and allowed it to short-circuit the case. In addition, the movement of the heat sink caused a pin to break off the output transistor that was mounted on it. I repaired the unit by pushing the heat sink back onto the insulated boss and re-soldering the wire to what remained of the transistor pin. It seemed that the charger was still broken until I remembered that it was designed to have no output until it was hooked to a battery. Once I biased up the control board with a few volts from a dead motorcycle battery, the charger worked fine, including the adjustable output voltage on the control board. An important factor in the longevity of the unit is that the designers did not use electrolytic capacitors.

The design uses puck-style diodes similar to those in automotive alternators. Manufacturers solder them directly to the copper heat sink. The two wires to the diodes are the ends of a center-tapped transformer's secondary winding.

A thermal circuit breaker limits output current in case the output transistor short-circuits. The breaker is wired into the negative side of the output.

The mounting of the control PCB allows an adjustment potentiometer to extend from the back of the case. The board regulates the output voltage and prevents output unless an attached battery is present.

The charger's main transformer features primary and secondary windings on separate bobbins, similar to the high-isolation transformers that medical equipment uses. The case includes a ground wire in accordance with UL (Underwriter Laboratories) requirements.

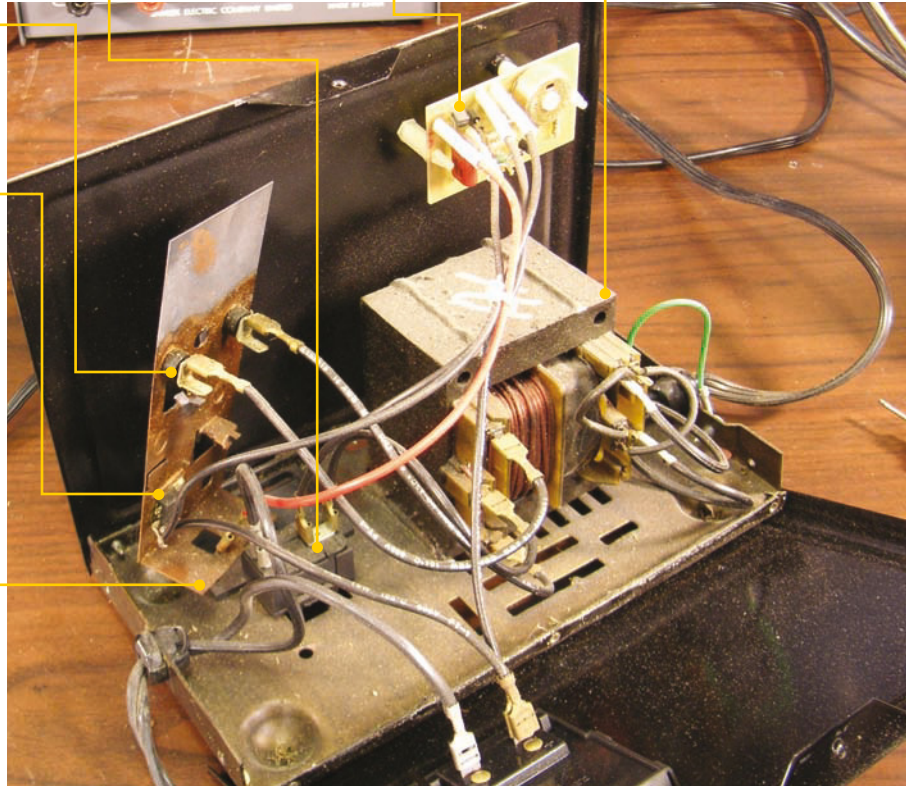

The output pass transistor has two thick wires soldered to each of its pins. The wire from the control PCB caused stress, which broke off the transistor's base pin.

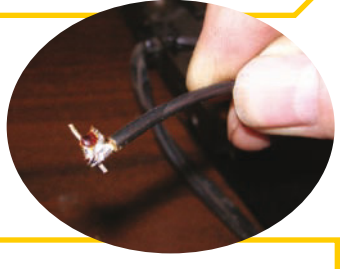

A thin copper sheet serves as a heat sink. A square hole in the bottom snaps over a plastic mounting boss in the outer case. The heat sink operates at 15V, the raw unregulated rectified output of the transformer.

Support **Across The Board.**™
From Design to Delivery™

Now, you can have it all.™

Faster and easier than ever before. Our commitment to customer service is backed by an extensive product offering combined with our supply chain and design chain services – which can swiftly be tailored to meet your exact needs. We have dedicated employees who have the experience to provide the highest level of customer service with accuracy and efficiency. All of our technical experts are factory certified on the latest technologies, providing you the expertise to move projects forward with speed and confidence.

Avnet offers the best of both worlds: extensive product and supply chain knowledge, and specialized technical skill which translates into faster time to market – and the peace of mind that comes from working with the industry's best. **Avnet is ranked Best-In-Class* for well-informed sales reps, knowledgeable application engineers and our design engineering services – proof that we consistently deliver:**

> Industry recognized product expertise
> Specialized technical skills

Ready. Set. Go to Market.™
Visit the Avnet Design Resource Center™ at:
www.em.avnet.com/drc

AVNET®
electronics marketing

Accelerating Your Success™

50 Years AVNET ON THE NEW YORK STOCK EXCHANGE NYSE

NEDA Supplier Authorized Distributor

1 800 332 8638
www.em.avnet.com

Follow us on Twitter!
www.twitter.com/avnetdesignwire

*As rated by Hearst Electronics Group: The Engineer & Supplier Interface Study, 2009.

UBM Electronics
DesignNews EDN Test & Measurement World

MECHATRONICS IN DESIGN

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

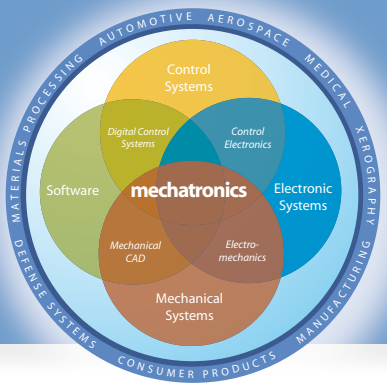

You can't just say it's reliable

Understanding the physics of failure and having a systems approach are essential to modern engineering.

"A man's got to know his limitations." This line is one of the more memorable that Clint Eastwood delivers in the movie *Magnum Force*, and it shows great wisdom. How does this line relate to engineering-system reliability, however?

What do we mean when we say that someone is reliable? Is it possible to say that a person is always reliable or just sometimes reliable, in all circumstances or in just some circumstances? You must apply the same questions to an engineering-system design because reliability cannot be an afterthought.

As you become more dependent on complex mechatronic systems, it is not sufficient to understand just how they work; you must also understand how they fail. Fault-tolerant system design, not just fault-tolerant component or subsystem design, has become paramount. Reliability is the probability that an item performs a required function under stated conditions for a stated period of time. An engineer thus must define the functions a system must perform, the boundary conditions under which the system will operate, and the time during which the system must be reliable.

As mechatronic systems become more complex, the interactions among the subsystems become more difficult to manage, and this integration affects the overall system reliability.

Tim Kerrigan, a fluid-power consulting engineer at Milwaukee School of Engineering's Fluid Power Institute, works to ensure that industrial and government systems are reliable. He believes that a physics-of-failure approach to reliability is consistent with the model-based approach of modern mechatronic-system design. It uses modeling and analysis to design reliability into a system, perform reliability assessments, and focus reliability tests in areas in which they will be most effective. The approach involves understanding and modeling the potential failure mechanisms, including fatigue, wear, and temperature; the failure sites; and the failure modes.

The failure modes of a mechatronic system include those of mechanical, electrical, computer, and control subsystems—that is, hardware and software failures. A physics-of-failure approach can improve reliability, reduce the time to deploy systems, reduce testing and costs, and increase customer satisfaction.

As mechatronic systems become more complex, the interactions among the subsystems—mechanical, electrical, computer, and control—become more difficult to manage, and this integration affects the overall system reliability. Therefore, an assessment of overall system reliability must have an adequate margin for safety. A useful analogy in this case is a feedback-control system. It provides great benefits, but a feedback-control system can become unstable if there is an imbalance between the strength of corrective action, or gain, and system dynamic, or phase, lags. You quantify model uncertainty by assuming that either gain changes or phase changes occur and that the tolerances of gain or phase uncertainty are the stability margins, gain margin, and phase margin. Real systems must have adequate stability margins. Real systems must also have adequate reliability margins.

Mechatronics can enhance the reliability and fault tolerance of a system with prognostics, diagnostics, and built-in test capabilities. The additional sensors and control elements must be reliable, and they add cost. The long-term cost of unreliability is much larger, however, than the initial design cost of reliability. In addition, designing for reliability enhances energy efficiency and sustainability. Reliability and fault tolerance are competitive advantages in the commercial market and absolute requirements in the healthcare, military, and transportation sectors. EDN

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

Introducing EDN's New Microsite on LEDs

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

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

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

OFFLOADING QUEUE AND TRAFFIC MANAGEMENT TO HARDWARE REDUCES DESIGN COMPLEXITY, IMPROVES APPLICATION EFFICIENCY, AND MAXIMIZES PERFORMANCE.

HARDWARE-BASED VIRTUALIZATION EASES DESIGN WITH MULTICORE PROCESSORS

BY SATISH SATHE • APPLIEDMICRO

Today's SOC (system-on-chip) processors integrate a diversity of cores, accelerators, and other processing elements. These heterogeneous multicore architectures provide increased computational capacity, but the resulting complexity also poses new challenges for embedded-system developers across a variety of applications, including control-plane processors, video servers, wireless base stations, and broadband gateways. Discrete cores each have full access and control of their resources. Such predictable access allows straightforward management and deterministic performance in applications with real-time constraints. In a multicore architecture, however, cores share access to resources, and potential contention complicates many design factors, such as processing latency and deterministically handling interrupts.

To provide deterministic behavior equivalent to that of single-core devices, multicore architectures have begun to implement resource-sharing and management techniques that have been proved in network communications. These architectures use established queue- and traffic-management techniques to efficiently allocate resources among multiple cores, maximize throughput, minimize response latency, and avoid unnecessary congestion.

AT A GLANCE

- The complexity of multicore processors poses new challenges, such as the risk of contention for access to shared hardware resources, for embedded-system developers.
- To manage the myriad interactions within a multicore SOC (system on chip), you must apply queue- and traffic-management techniques that have been proved in network communications.
- Virtualization of on-chip resources enables each application to treat a resource as if it were the sole owner.
- Hardware-based virtualization accelerates the software virtualization layer, enabling faster development with simplified debugging and increased system reliability.

RESOURCE VIRTUALIZATION

From an architectural standpoint, SOCs are complex systems with multiple cores that connect across a high-speed fabric to a variety of controllers and resources (**Figure 1**). In many ways, the myriad interactions within an SOC resemble a communications network with multiple sources, or cores, that interconnect to the same destinations, including memory, peripherals, and buses. Not surprisingly, bandwidth-management techniques, such as virtualization, which designers developed to improve network efficiency, have proved useful in managing traffic among multiple processor cores and shared peripherals.

Virtualization of on-chip resources enables cores to share access; this shared access is transparent to applications. Each application can treat a resource as if it were the sole owner, and a virtualization manager aggregates shared ownership—measured by the amount of allocated bandwidth. Virtualizing and sharing access to resources require both a queue manager and a traffic manager. Applications use one or more queues to buffer access to a resource. Virtualization adds events or transactions to the queue and pulls them off when the resource is available. Queues comprise a list of buffer descriptors pointing to data in a buffer, and you can implement queues in many ways, depending on the needs of the applications. The number of supported queues varies in an SOC from a few hundred to hundreds of thousands to meet the needs of various applications.

The queue manager updates the queue state—that is, the queue size, head pointer, tail pointer, and start address—and maintains fill levels and thresholds, including full, almost full, almost empty, and empty. The queue manager also provides full memory management for each queue, including allocation and deallocation of buffers from free pools and checking of access rights when an event is added to a queue (**Figure 2**). Multiple requesters may simultaneously add descriptors to one or more queues, as well as allow selection from multiple queues waiting for a service.

The manager serves as the arbitrator for available bandwidth among queues assigned to the same resource. It performs this task not only between applications sharing a resource but also among the multiple queues an application may have to enable QOS (quality of service).

Traffic management employs policing and shaping mechanisms to measure and control the amount of bandwidth assigned to a flow or a group of flows. Policing controls the rate at which the traffic manager adds events to a queue, and shaping is the rate at which the traffic manager removes events from the queue. For the most control and ability to manage queue priority, you must implement policing and shaping on a per-queue basis. The traffic manager also maps multiple queues to a single shared resource based on a predefined servicing algorithm.

By bringing queue and traffic management together, you can provide reliable, end-to-end QOS. This approach allows multiple paths to share a resource without negatively affecting bandwidth subscriptions. Fine-grained QOS supports SLAs (service-level agreements), guaranteeing minimum, average, and maximum bandwidth on a per-flow basis. Developers can implement queue levels for marking and metering traffic to prevent congestion. Early notification of congestion allows the queue manager to take corrective action through feedback to traffic sources to eliminate the unnecessary processing of packets that are

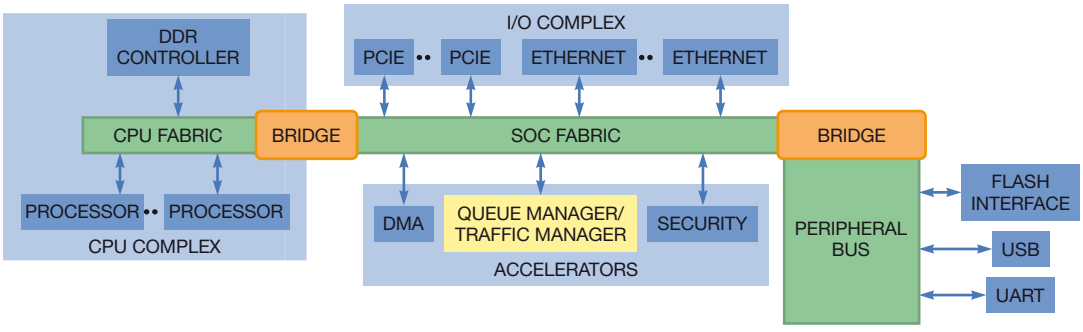

Figure 1 Next-generation SOCs are complex systems with multiple cores accessing the same shared resources. These systems resemble a network fabric and can benefit from the same queue- and traffic-management technology that communications networks use.

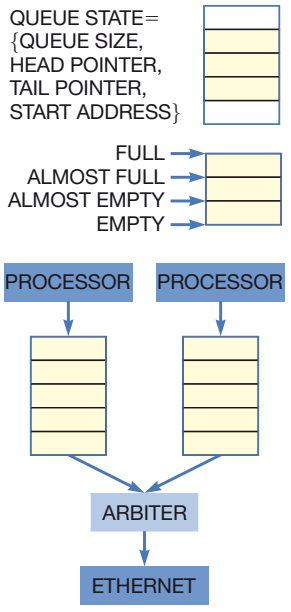

Figure 2 Hardware-based virtualization offloads queue management from application processors, including updating queue states, maintaining fill levels and thresholds, allocating and deallocating buffers, and confirming access rights when an application accesses a queue.

likely to be dropped or, ideally, to avoid congestion altogether.

For example, a queue- and traffic-management-based Ethernet driver prevents any one processor from unfairly monopolizing port bandwidth. It also guarantees bandwidth allocations and maximum-latency constraints regardless of other queue states. The driver supports a choice of arbitration schemes—strict priority or weighted round robin, for example—and facilitates reliable real-time services, such as video streaming. In the end, multiple sources can share the Ethernet port without adversely affecting bandwidth subscriptions. Tasks such as IP (Internet Protocol) forwarding become straightforward to implement robustly, and latency-sensitive applications, such as audio or video delivery, benefit from deterministic and reliable port management. In addition, when you implement the queue and traffic management in hardware, the driver can maintain end-to-end QOS with little to no software overhead.

THE VIRTUALIZATION LAYER

Early multicore SOCs, like the original network processors, left all of the work of virtualizing resources to developers. Applications, to some degree, had to recognize that they might share a resource with other applications. When an application used a shared resource, it had to do so in a way that allowed coexistence with other applications. The operating system also needed to support virtualization.

In a traditional architecture, processors manage their own access to shared resources through a software layer (**Figure 3a**). Processors must be aware of what resources are available and how often they can use them. As the number of processors increases, so does the complexity of resource sharing. One downside of software-based virtualization is that it introduces overhead to every transaction to store and later retrieve packets. Such overhead consumes processor cycles and introduces complexity to the coding process. It also places the burden of bandwidth management and meeting subscription guarantees on the virtualization software. Even when using tools to automate the creation of virtualization code, developers still must troubleshoot application interactions as they pass through the virtualization code.

The added overhead and complexity of virtualization have limited the use of multicore SOCs. Queue and traffic management, however, is a fairly deterministic process that you can implement in hardware. Developers configure queues once for an application, and the hardware mechanisms can then completely offload queue management, thus restoring substantial computational cycles back to the application processors. The ability to dynamically change allocations allows modification of the overall configuration at runtime to accommodate changing task loads.

In an architecture using a hardware-based queuing and synchronization mechanism, each processor operates independently of the others (**Figure 3b**). Through virtualization of resources, sharing becomes transparent to the applications. The mechanism allocates each processor and each task resource bandwidth, and each processor and task operates as if it were the only controller of the resource. Although the gains from implementing queue and traffic management vary from application to application, hardware-based resource virtualization and sharing significantly improve system efficiency.

A hardware-based virtualization layer removes or accelerates the software-virtualization layer. Offloading virtualization substantially increases processor efficiency. In some cases, hardware-based virtualization removes the need for software-based virtualization, other than during initial configuration. In other cases, hardware-based queue and traffic management significantly accelerates virtualization software in the datapath.

A hardware-based virtualization layer also lowers design complexity and speeds development because it eliminates the need for developers to implement and design around the virtualization layer. This approach simplifies design and speeds time to market. This hardware-based layer also increases determinism. Elimination of virtualization overhead reduces a major source of system interrupts. This elimination in turn reduces processing latency and increases system responsiveness.

Another benefit of this approach is

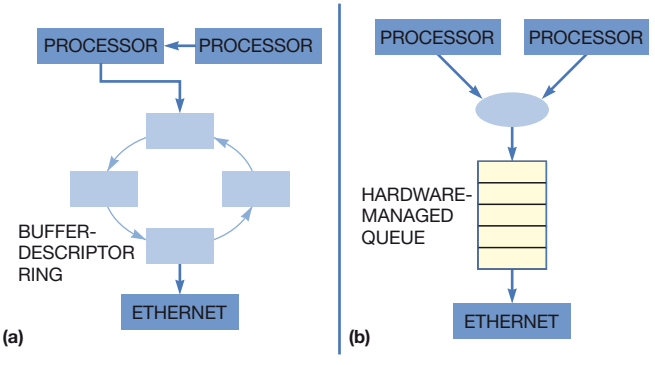

Figure 3 In a traditional architecture, processors manage access to shared resources through a software layer (a). As the number of processors increases, so do the complexity and overhead of resource sharing. Hardware-based virtualization eliminates the need for a software-based virtualization layer (b). In addition to offloading queue and traffic management, resource sharing becomes transparent to applications, and each processor can operate independently of the others.

that it simplifies debugging. Because virtualization and resource sharing are hardware functions, the virtualization layer itself is not part of the development process. However, developers still have full access to and control of queues if necessary for troubleshooting. A hardware-based virtualization layer also increases reliability because hardware-implemented queue and traffic management is not vulnerable to many of the issues that can arise with a software-based implementation. For example, if the core-handling software-based virtualization becomes compromised, the entire system is vulnerable. With a hardware-based implementation, there is no centralized control routine to compromise.

PROCESSOR OFFLOADING

The level of supported queue offloading depends on the implementation. For example, some SOCs might provide locking mechanisms but not perform all state management of queues. Ideally, developers want a flexible system that supports different configurations, is straightforward to integrate with software, and minimizes the software changes necessary to adopt the SOC. A virtualization mechanism may be efficient; if it requires significant deviation from traditional programming models, however, porting application code will increase system cost and delay time to market.

How you implement queues can also affect system performance. For example, queue location affects which processors can access those queues. Some queues must reside in memory types, be spread across multiple chips, or be tied to a resource. Dynamically allocated queues give developers the flexibility to appropriately partition queues to applications and resources. For systems using multiple multicore SOCs, the ability to manage queues over a system bus, such as PCIe (Peripheral Component Interconnect Express), enables sharing of resources not just between cores on the same SOC but also between those on different SOCs. For example, a cluster of processors can share a single forwarding database. Alternatively, a multiple-SOC system may have a single deep-packet-inspection engine that applications running on different SOCs must access. Such multichip sharing of resources allows even further virtualization of system resources.

One of the greatest design challenges in multichip architectures is partitioning tasks in a way that equally spreads resource requirements among all processors. In software-based virtualization, this process can be time-consuming and places a burden on designers, including the challenge of efficiently managing free memory pools. In addition, any change in software can result in a shift in resource requirements, requiring developers to repartition the system. Many of these issues apply to both asymmetrical and symmetrical multiprocessor architectures.

With hardware-based virtualization, most partitioning management takes place in hardware, and the operating system handles a small remainder. With this abstracted partitioning, developers can make system changes without manually repartitioning the system. This approach also offloads tasks, such as managing free memory pools, from the application and operating system.

BANDWIDTH GUARANTEES

Control of a resource also extends to limiting the maximum allocation a processor can receive to address potential processing bottlenecks on the receiver side. For example, many communications, audio/video, data-acquisition, and test-and-measurement applications have a maximum transmission data rate that the receiving processor is expecting or can handle. In these cases, even if there is more capacity available on the peripheral because other processors are not currently using their allocations, the application may not want the queue flushed at a faster rate because this flushing may overwhelm the receiving processor and result in loss of data.

Many developers take a worst-case approach to design; they make sure there is enough capacity to support worst-case loading. This approach means, however, that, under typical operating conditions, there will be underused resource capacity. A typical round-robin arbitration algorithm, for example, supports only minimum allocations. If the system can have as many as 10 requesters for a resource, each can expect to always have at least 10% of the bandwidth. However, if only one requester is active, that requester could receive 100% of the bandwidth.

Virtual and transparent resource allocation means that an application does not know how much bandwidth it might receive. For applications with receiver-side bottlenecks, the ability to set a maximum allocation for a resource is important for the stability of the system. This maximum allows developers, no matter what allocation algorithm is in use, to control resource bandwidth per application, to prevent swamping the receiver-side processor, and to prevent data loss. Developers also have the option of implementing standard mechanisms, such as IEEE 802.1Qav or 802.1Qau, to manage congestion.

SYSTEM STABILITY

An application may sometimes attempt to use a resource to which it does not have access. This situation can occur because of an error in programming, when an only partially updated application is in use, or when an overwriting of code or data memory has occurred. You must prevent such an application from corrupting other applications—that is, by writing in their memory space—or negatively affecting their performance—for example, by seizing control of a shared resource. In software-based resource-sharing implementations, a corrupted application may ignore its bandwidth allocation and monopolize a shared resource. Similarly, if the processor hosting virtualization becomes corrupt, queuing mechanisms may fail and bring down the entire system.

Hardware-based queue management allows you to protect the various components of the system from each other. The most basic form of fault isolation is preventing access to memory and resource bandwidth allocated to other applications. To keep sharing of virtualized resources completely transparent to applications, the queue and traffic manager must take action only on the corrupted application. In other words, you must shield applications from both the actions of other applications and the need to accommodate the failure of another application to maintain stability. Dedicated queues, by their nature, isolate faults and prevent other processors and applications from effects. Such queues also facilitate effective error recovery; dedicated queues can completely clear with no loss of data for other applications.

A queue- and traffic-management controller can implement several levels of response to resource-access violations. The simplest response is to prevent the access and generate an alarm to the application, typically through an interrupt. This alarm tells the application that it has tried to do something it shouldn't have. A second method logs the violation for use by developers to help trou-

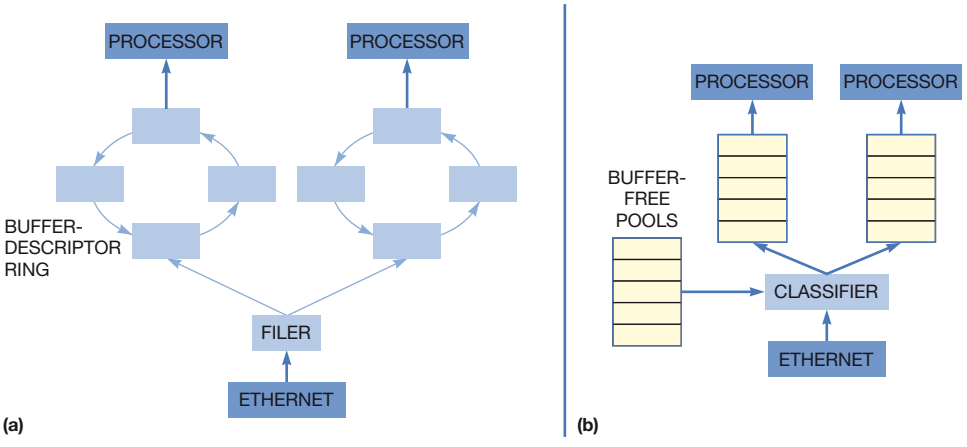

Figure 4 With software-based virtualization, bandwidth allocations between processors are often fixed in a manner that limits effective management of QOS and oversubscription (a). Hardware-implemented virtualization enables resource allocation, even among SOCs with full bandwidth management, including rate policing, traffic shaping, and queue arbitration that reflect and balance the needs of each processor and application (b). This approach enables efficient sharing of resources, such as a hard drive or a security engine, across the entire system, not just one processor.

UBM Canon

Accelerate Your Projects into the Future

AEROCON

AEROSPACE • AVIATION • DEFENSE

FEBRUARY 8–10, 2011
Anaheim Convention Center
Anaheim, California

MARCH 16–17, 2011
Orange County Convention Center
Orlando, Florida

MARCH 14–15, 2012
Fort Worth Convention Center
Fort Worth, Texas

Introducing the new conference and exhibition for the aerospace and defense industries

Featuring suppliers for the design and production of aerospace and defense products, including:

- Software
- Materials
- Sub-Assemblies
- Electronics
- Machinery
- Components
- Contract Services
- *...and much more!*

Your Project Acceleration RESOURCE

For more information on attending or exhibiting, visit:

AeroConShows.com

Promo Code: AA

17411_AS_AE11

Produced and managed by UBM Canon | canontradeshows.com

bleshoot issues in the field. The queue- and traffic-management controller also must be able to escalate its response by triggering a reset and the reinitialization of a potentially corrupted application. Ideally, a developer can create a policy that controls this response. For example, a developer could set a threshold dictating that, if an application makes three illegal accesses, it is assumed to be a corrupted application and must restart.

PARTIAL VIRTUALIZATION

When a series of transactions must take place in order, this requirement becomes a blocking instance because other requesters must wait until the transaction is complete before they can gain control of the resource. Consider a typical SATA (serial-advanced-technology-attachment) transaction in which you first configure the SATA port and it then executes a sequence of commands. Unlike an Ethernet port, in which packets are single events, the SATA port must lock to an application until the transaction is complete; otherwise, two applications may overwrite each other before either has completed its task.

Although applications cannot fully share resources supporting transactions of this nature, allocation can be partially virtualized. Applications wanting to use the resource first must make sure that the port is available and then lock the port while it is in use. Support of locking requires a thin software layer between operating systems to enable them to communicate to see which application has control of the lock. The use of hardware, however, can manage and accelerate acquisition of the lock. You must implement lock acquisition in hardware to provide a failsafe mechanism for the resource; otherwise, a locked processor could also lock the resource.

Depending on the application, a system must support shared resources that can be completely virtualized and those that require locking. An SOC could, for example, provide a SATA port that is not shared, but only one processor could use it, and sharing of the resource would have to be in software. By also supporting lockable resources, cores within the SOC can still share the resource in a manner transparent to all applications with failsafe reliability.

Ease of integration is an important aspect of multicore architectures. The ability to bring multiple processors onto a single chip requires straightforward application software to migrate; otherwise, developers might as well design a new system.

You must consider a number of factors in determining ease of migration to a virtualized architecture. For example, the architecture must support multiple operating systems because multiprocessors often use multiple operating systems across cores, depending on the applications that require support. Multicore architectures that support only one operating system force developers to use that operating system and then port all code to it. By supporting multiple operating systems, a multicore SOC simplifies code migration.

You also need to consider QOS because applications have different bandwidth needs. Latency-sensitive applications, such as video streaming, need real-time access to shared resources, whereas data-based applications, such as content downloading, can tolerate delay and take advantage of underused bandwidth. The ability to service different bandwidth requirements enables developers to bring together divergent applications under the same processor cluster.

Also consider whether the architecture includes transparent resource sharing because transparency allows developers to migrate both applications that support virtualization and applications that don't support it. Another aspect is removal of the software-virtualization layer. Although some code rewriting is necessary when migrating between SOCs, for many applications, most changes when moving to an SOC with hardware-based resource sharing involve not changing the software but rather eliminating the software-virtualization layer. Removal of this layer simplifies system design and troubleshooting and increases system efficiency. In cases in which manufacturers have licensed virtualization code, removing this layer also reduces system cost.

Another factor to consider is whether the architecture consolidates system resources. When a system employs multiple chips and operating systems, each must have its own storage resources.

By managing resources in hardware, all tasks can share access to a resource they need. For example, rather than requiring multiple drives, as a traditional architecture does, a single hard drive or Ethernet port can serve an entire system, even across SOCs. This approach results in equipment savings and lower overall system-component count.

Communication among SOCs is also important. Applications can share resources over a high-bandwidth system bus, such as PCIe, to enable sharing among SOCs. Traditional architectures allocate a fixed bandwidth to each processor in a manner that limits effective QOS management among cores and makes it difficult to oversubscribe reliably (**Figure 4a**). With hardware-based virtualization, resource allocation is flexible, even among SOCs (**Figure 4b**). Full bandwidth management is possible, with rate policing, traffic shaping, and queue arbitration that reflect and balance the needs of each processor and application. Together, this approach enables the efficient sharing of resources, such as a hard drive or a security engine, across the entire system, not just one processor.

You should also consider interprocessor communications because multiprocessor systems often must transfer significant amounts of data among processors. Queue-management mechanisms provide a simple and efficient means of accelerating communication among processors on an SOC and across multiple SOCs.

Today's next-generation SOCs are complex multiprocessor environments that must share on-chip resources without incurring additional overhead that reduces system efficiency. Queue management helps virtualize chip resources and simplifies resource sharing by providing a reliable mechanism for allocating bandwidth, isolating faults, and facilitating robust error recovery. Traffic management ensures that a system fairly shares resources in a manner that meets the differing latency and throughput needs of applications by controlling the rate at which traffic enters and leaves queues. With hardware-implemented virtualization, developers can offload queue and traffic management to improve application efficiency, maximize resource throughput, reduce latency, and increase system reliability. EDN

AUTHOR'S BIOGRAPHY

Satish Sathe is a senior systems architect at AppliedMicro, where he is responsible for SOC-architecture and -feature definition. He has more than 20 years of experience in building and designing state-of-the-art network-switching systems and controllers and has played key architecture roles at LSI, Stratacom (which Cisco Systems acquired), Acclaim Communications (which Level One and then Intel acquired), Ample Communications, and Brocade. Sathe holds a master's degree in electrical engineering from the University of Illinois—Urbana-Champaign.

BY PAUL RAKO • TECHNICAL EDITOR

BATTERY-STACK-MONITOR ICs SCRUTINIZE THE CELLS

ELECTRIC-VEHICLE AND DATA-CENTER BATTERIES NEED PRECISE MEASUREMENTS USING ROBUST CHIPS.

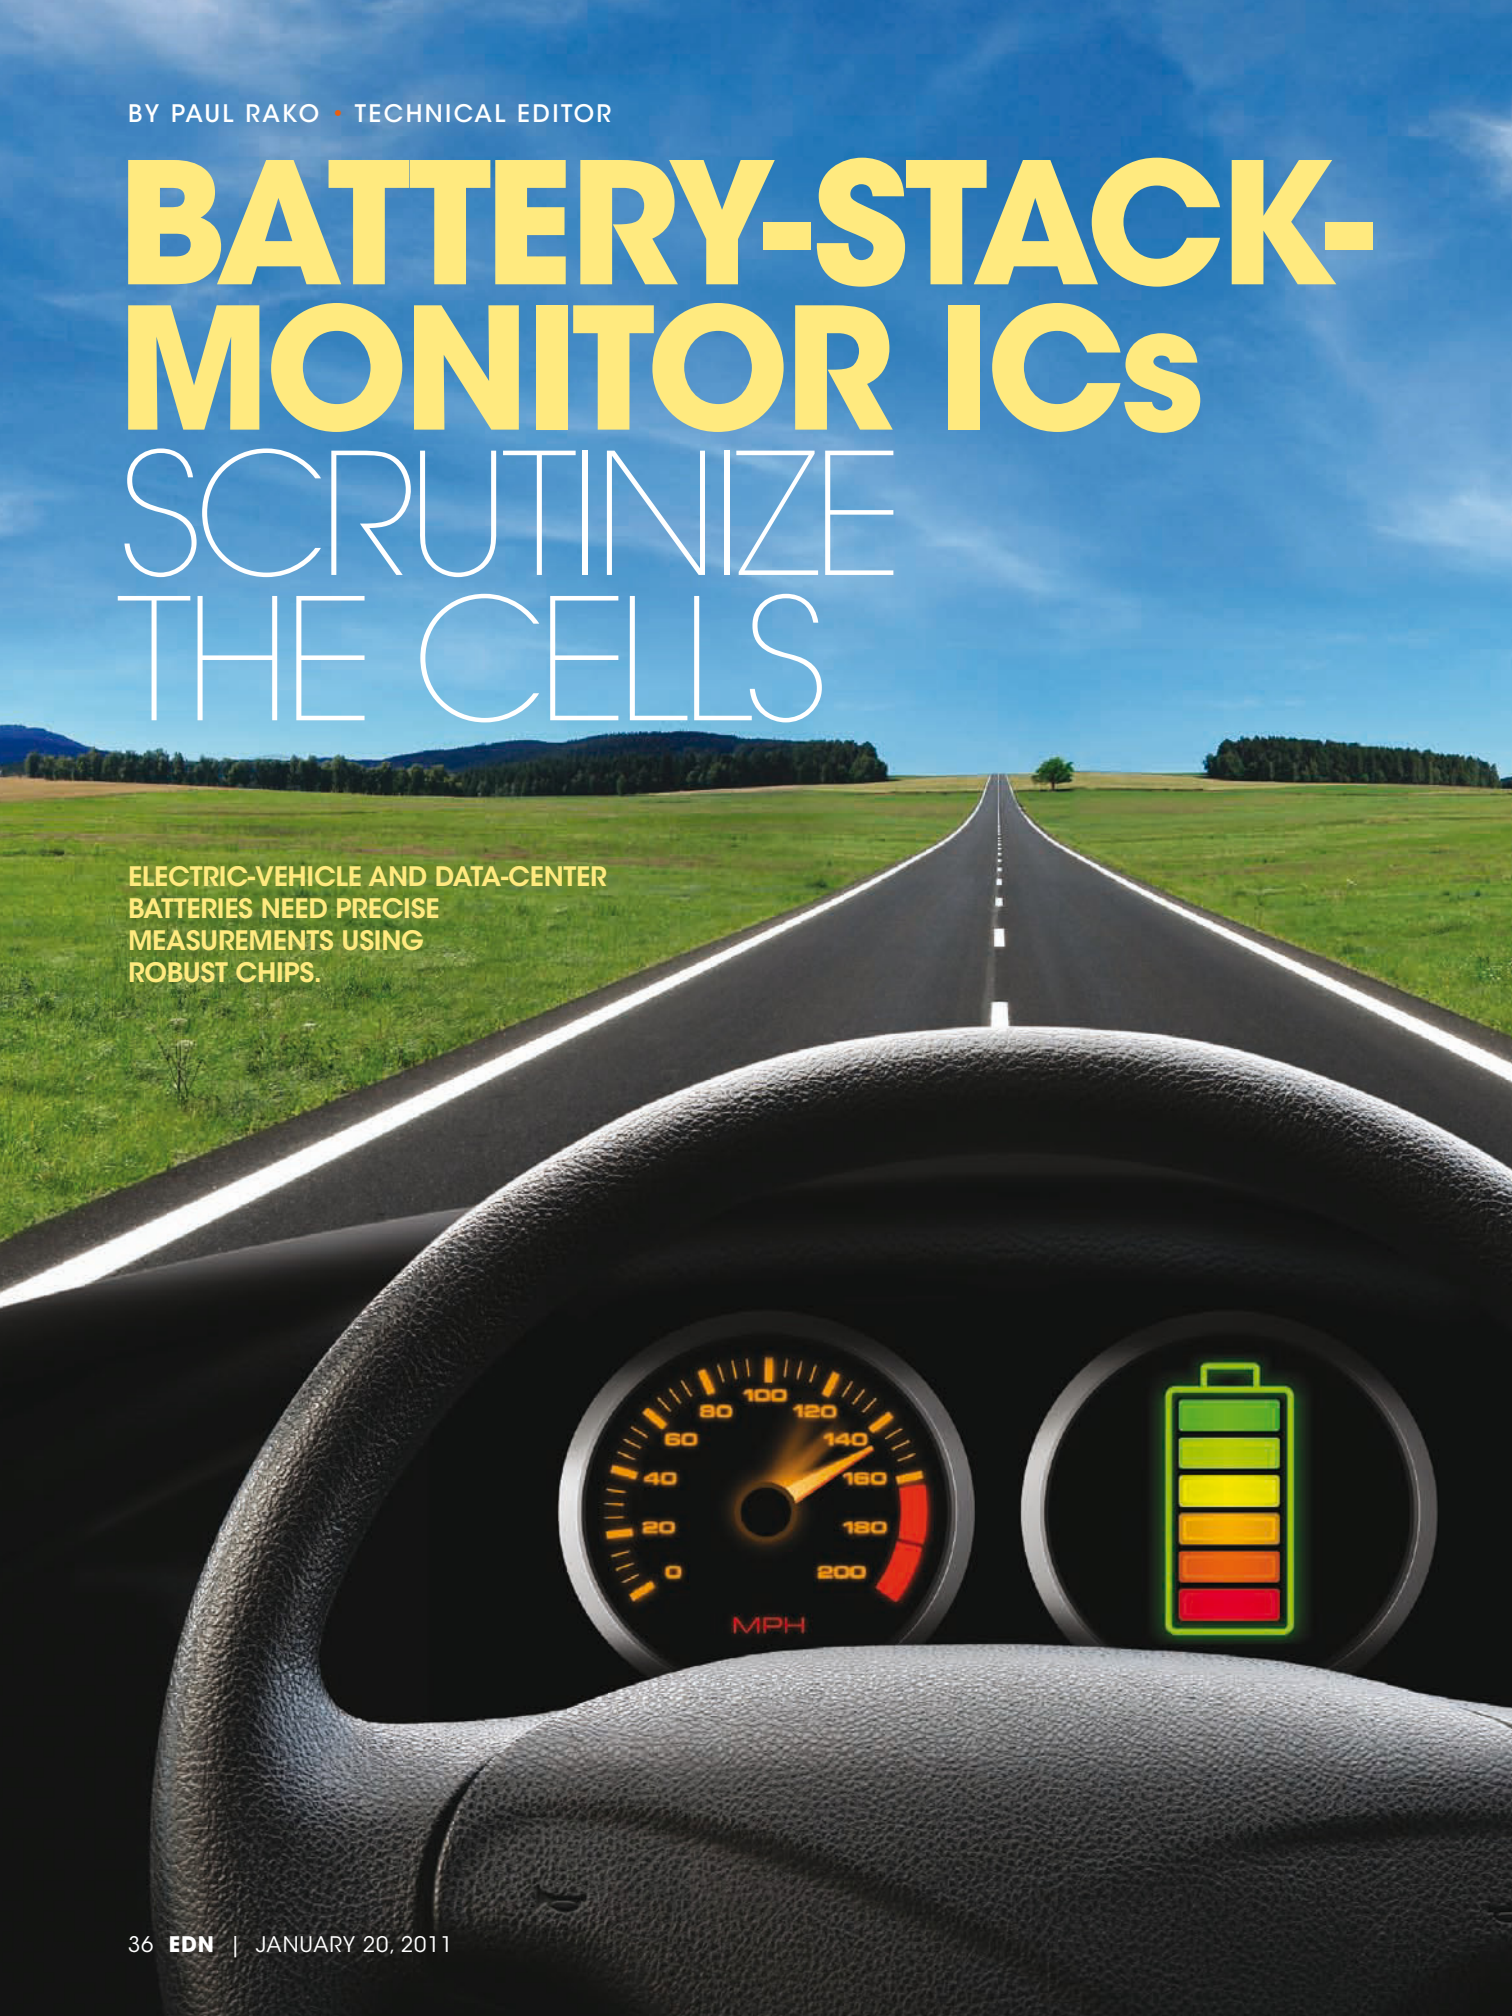

DASHBOARD & LANDSCAPE: NARVIKK/ISTOCKPHOTO.COM; BATTERY ICON: VADYM TYNENKO/ISTOCKPHOTO.COM

Battery cells, supercapacitors, and fuel cells need careful monitoring to extend range, prolong life, and ensure safety in energy-storage systems, such as those in electric and hybrid vehicles (**Figure 1**). The use of batteries in automobiles is developing along a range of applications. Micro hybrid vehicles use a conventional 12V lead-acid battery and have alternator-motor units that allow the engine to stop when you bring the vehicle to a halt. When you press the gas pedal, the engine smoothly starts and then operates conventionally. Hybrid vehicles, such as the Toyota (www.toyota.com) Prius, the Honda (www.honda.com) Insight, and the Chevy (www.chevrolet.com) Volt, have much larger batteries. These batteries produce more than 200V. Cell chemistries have traditionally been NiMH (nickel-metal hydride), but various lithium-ion chemistries are providing more energy for a given weight (**Figure 2**). Fully electric vehicles, such as the Tesla (www.tesla.com) Roadster and the Nissan (www.nissan.com) Leaf, have the largest batteries; their battery-stack voltages range from 300 to 400V.

The higher the voltage in a battery, the lower the current for a given power will be, reducing the gauge of expensive copper cabling. More important, the higher voltage allows the winding of higher-output motors. In 2004, Toyota added a boost converter to the Prius that raised the battery-stack voltage from 200 to 500V. This step allowed Toyota to redesign the propulsion motor and improve torque from 350 to 400 Nm and power from 33 to 50 kW (**Reference 1**).

Data centers also use 300V battery strings for UPS (uninterruptible-power-supply) backup power. In this application, lithium-ion batteries are replacing lead-acid batteries. Vehicles take advantage of lithium-ion's better gravimetric energy density—that is, the energy per pound or kilogram. UPS applications instead involve the volumetric energy density of lithium-ion batteries. Data-center floor space is expensive; although a lithium-ion-battery system may cost more, it takes up only one-fourth the space that a lead-acid-battery system requires. This fact often allows data centers to combine the battery and inverter systems into one room. Some data centers are considering removing the inverters and distributing dc voltage to data-server computers that can accept dc inputs.

Grid-leveling applications share the same benefits as data centers when it comes to a lithium-ion battery's size. Some grid-leveling schemes intend to use fuel cells, and high-voltage stacks of fuel cells need the same careful monitoring as electrochemical batteries. Fuel cells have special requirements; they can have either polarity on the cell during use and exhibit various failure modes. IC manufacturers are adapting their battery-stack-monitor chips to handle these negative cell voltages.

A similar problem occurs when you are monitoring stacks of supercapacitors. Users want to get all the energy from the capacitor, and doing so means discharging it to 0V. When this scenario occurs, dielectric effects can cause the capacitor to exhibit a negative voltage, often as large as –0.5V. Several IC manufacturers have ruggedized their battery-stack chips to handle these negative potentials. Supercapacitors store less energy than do batteries or fuel cells, so they are finding less use in high-energy applications (see **sidebar** "Battery characteristics").

CELL MONITORING

Auto and UPS manufacturers want to

accurately measure each cell in a battery stack. "You don't want to shut the car down for a bad cell, but you would shut it down for an overtemperature condition," says Paul Maher, a hybrid- and electric-vehicle-segment marketing manager at Analog Devices. The care of auto batteries is critical. "You expect a laptop battery to last two years, but an automotive pack should last 10 years," he adds.

The measurement must be accurate because a few millivolts can represent a large amount of charge. The measurement has a common-mode problem, in which you try to make accurate cell measurements in the presence of hundreds of volts of common-mode potential. The measurement is not a dc measurement in which you can use integrating ADCs. The battery voltage may be changing at kilohertz rates due to the chopping action of the motor-inverter circuitry. Furthermore, you need galvanic isolation in the measurement system because battery voltages are hazardous. The chips must consume little power so that they don't drain the battery. In addition to the difficulty of the measurement itself, you must communicate the measurement to various places in a vehicle or a data center.

Your first challenge with a battery-stack-monitor circuit is accuracy. A modern lithium-ion cell has a flat discharge curve. "A 5-mV measurement error represents a 10% error in the state-of-charge estimation of the cell," says Matthew Borne, power-marketing manager at Texas Instruments. You must stop discharging the batteries before you damage them, so better accuracy translates directly into greater range; 8-mV accuracy on a 4V cell translates to 0.2% precision. To deliver system accuracy of 0.2%, the voltage-reference accuracy would have to be perhaps 0.1% over time and temperature (**Reference 2**).

AT A GLANCE

- Measuring battery cells is difficult, requiring accuracy, common-mode rejection, and fast response.
- Chips make measurements in environments full of EMI (electromagnetic interference), heat, vibration, and motor noise.
- The automotive market requires high volumes and low prices.
- Automakers' safety and liability concerns require redundant, fault-tolerant designs.
- ISO (International Standards Organization) 26262 safety requirements should emerge in 2011.

Once you achieve sufficient accuracy, you face another problem: measuring a 4V cell that may be wired in series with dozens of other cells. You would need precise resistor dividers if you were to use an attenuator to measure the cell voltage (see **sidebar** "Common-mode problems"). Even thin-film resistors are not accurate enough and cannot track closely enough over temperature.

You can charge capacitors up to the cell voltage and then switch them to a chassis-referenced potential. This so-called flying-capacitor approach works but has drawbacks. For example, capacitors begin to transfer mismatched charges between cells of different potentials, according to Stephen G LaJeunesse, business manager of automotive- and industrial-battery products at Maxim Integrated Products. "They also require high-voltage switches, and these switches have losses of their own," decreasing the circuit's efficiency, he says.

Jim Williams, staff scientist at Linear Technology, has developed a novel circuit that uses small, inexpensive transformers that he interrogates for each cell's voltage (**Figure 3** and **Reference 3**). The circuit performs well, but the transformers add cost and might fail due to vibration.

Battery-stack-monitor-IC manufacturers avoid the common-mode problem by floating the chips at the cell's stack voltage. They convert the analog measurement to a digital value and then communicate those digital bits down a daisy chain of other chips. This step removes the resistive attenuators from the system and eliminates any common-mode-attenuation errors from the measurement (**Figure 4**).

Placing the measurement chips in daisy chains helps with another important requirement. It requires one galvanic isolator instead of many. Decades ago, engineers would try to transmit the analog voltage across an isolation barrier. The architecture of battery-stack-monitor chips reveals a modern trend. You measure the analog voltage and then convert that voltage to digital bits. Many ways exist for transmitting digital data across a galvanically isolated boundary (**Reference 4**). You can use optocouplers, capacitive isolators, transformer isolators, RF isolators, or even magnetostrictive isolators. If you send analog voltage levels across the boundary, you can use delta-sigma modulators, such as those from Avago.

Once you ensure an accurate measurement and solve the common-mode problem, you must make sure to meet the power requirements of the design. The battery stack itself provides power for most battery-monitor ICs, meaning that you deplete the batteries unless you

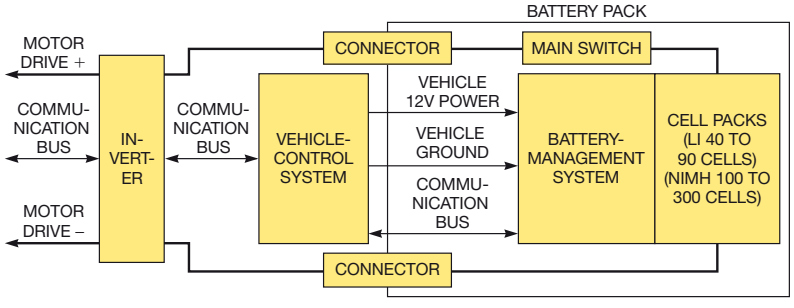

Figure 1 The battery-management system of an electric car must measure the voltage and temperature of individual cells (courtesy Maxim Integrated Products).

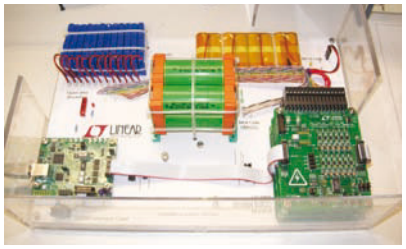

Figure 2 Electric vehicles and UPS applications might use various battery types, all of which need monitoring with battery-stack-measurement ICs (courtesy Linear Technology).

BATTERY CHARACTERISTICS

Cell manufacturers often issue glowing reports about how wonderful their batteries are, but you need to read the comments of a cell's users to get a realistic idea of a cell chemistry's limitations (Reference A). Lithium-ion cells have a flat discharge curve, so you must measure them with millivolt accuracy (Figure A). Overcharging lithium-ion cells or storing them at 100% charge can cause damage. As with most other cell chemistries, discharging lithium-ion cells to 0% capacity also causes damage (Figure B). For these reasons, the Chevy Volt's battery-management system charges the vehicle's 16-kWhr battery to 90% and stops the discharge at 25% capacity. This approach yields a usable battery capacity of 10.4 kWhr but extends the cycle life.

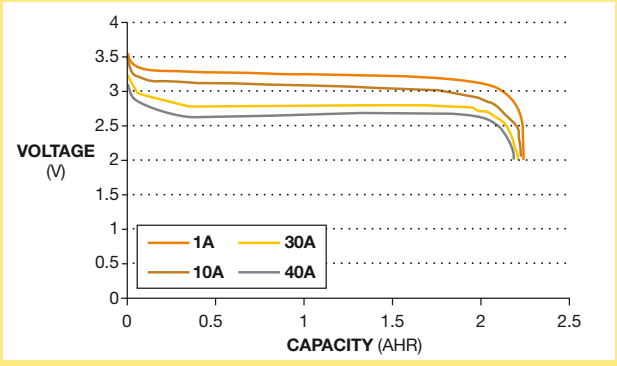

Figure A Lithium-ion cells have flat discharge curves, so they require accurate measurement to determine state of charge (courtesy A123).

Battery designers must make tradeoffs between energy density and power density (Figure C). A battery with acceptable energy density holds a lot of energy but must draw it out slowly. Batteries with acceptable power density have low internal impedance, so large pulses of power can come out without damaging the cell.

Although discharging a cell to 0% may damage the cell, it won't endanger a user. Overcharging cells, on the other hand, can endanger you. Excessive current in any cell chemistry causes the cell to heat up, boil, and even burst. Lithium-ion cells are particularity problematic in this regard. The electrolyte in some lithium cell chemistries is flammable, as is the lithium metal in the cells. Rumor has it that both Tesla and Chevrolet have lost cars to battery-pack fires during testing. Older lithium-ion chemistries using cobalt are the most prone to fire. Short-circuiting or puncturing them can deliver enough energy to the short circuit to create a conductive plasma that can cause even more short-circuiting of adjacent cells.

Newer iron-phosphate lithium-ion cells are less prone to bursting into flame; they instead sizzle and simmer. For this reason, most large automobile companies use iron-phosphate lithium batteries. The downside is that the iron-phosphate cells have 33 to 50% less gravimetric energy density, which makes the battery heavier. The major auto companies thus are cautious about any radically different propulsion system, such as those that electric cars use.

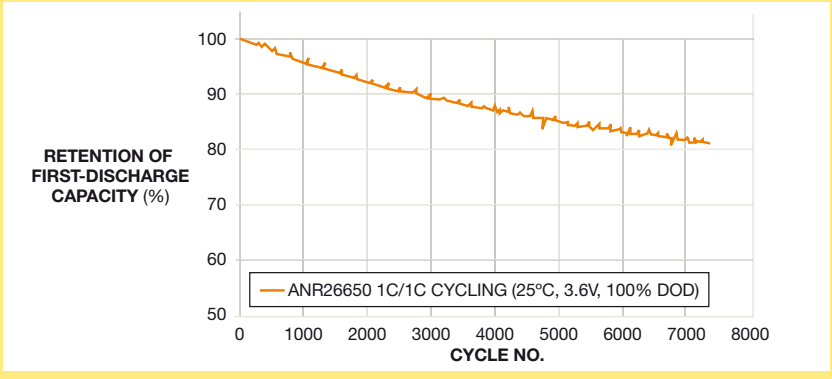

Figure B Modern cells can perform many charge-discharge cycles, but a key factor is in the definition of 100% discharge. Letting the cells discharge to low voltages immediately damages them (courtesy A123).

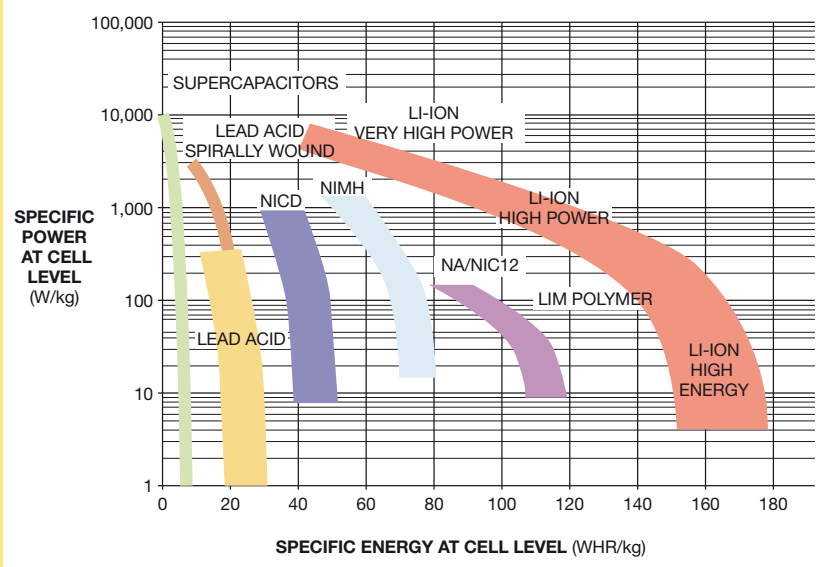

Figure C Supercapacitors have low impedance but hold little energy. High-energy lithium-ion cells hold a lot of energy but have high internal impedance (courtesy Saft).

REFERENCE

A Hoffart, Fran, "Proper Care Extends Li-Ion Battery Life," *Power Electronics Technology*, April 2008, http://bit.ly/eBY5v0.

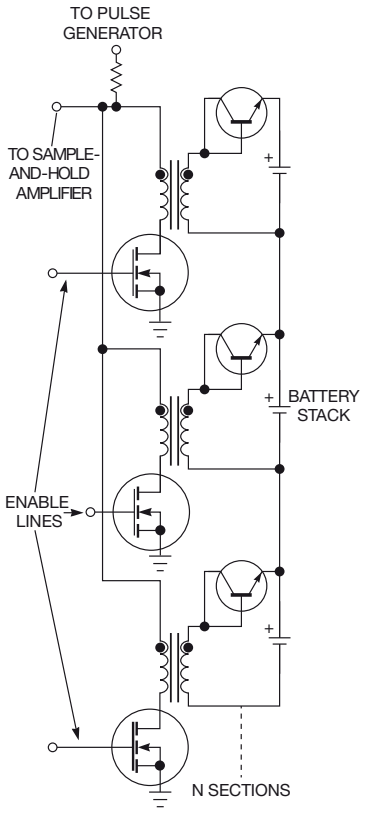

Figure 3 This circuit solves the problem of measuring common-mode voltage by transferring the cell voltage plus a diode drop across small isolation transformers.

are charging them. Just as important, each chip must have the same power consumption so that one doesn't unbalance a group of cells by drawing more power than its neighbors. You could also provide for isolated power from the car battery or an external source, as Analog Devices does on its monitor chips (**Figure 5**). Thus, the monitoring circuits do not deplete the propulsion battery.

Once you design the measurement system, you must send the data over a communication link. Some manufacturers convert a simple local serial protocol, such as SPI (serial-peripheral interface), to a high-level protocol, such as a CAN (controller-area-network) bus. Decades of use have established the reliability of CAN communications in automobiles.

These considerations of the measurement system are just the basic requirements. To meet auto manufacturers' reliability and liability concerns, you must measure each cell in the battery. To minimize the required number of measurement converters, most IC manufacturers use high-voltage, fault-protected multiplexers on their battery-stack-monitor chips, allowing the chip to accurately measure four to 12 cells and then communicate the measurements across a serial bus to the next chip in the daisy chain.

Another major reliability concern occurs when assembly workers or mechanics use hot-plug replacement to connect the measurement circuitry to the battery. This approach provides no guarantee about which cell the workers will connect first. The arbitrary connection of any cell's sense line can inject bulk current into the silicon die. Chips also must have rugged junctions, according to Maxim's LaJeunesse, who notes that the chips don't need dielectrically isolated processes. "DMOS and CMOS get the job done, but you have to know how to bias the circuit elements and how to strap the transistor wells and gates internally," he says. He advises IC designers to use guard rings to cure hot-carrier injection.

Most manufacturers provide for both external and on-chip temperature sensing, enabling system designers to measure ambient and cell temperatures. The designers can then factor those temperatures into charging and safety algorithms. Reliability concerns have also prompted auto manufacturers to ask for fully redundant measurement systems (**Figure 6** and **Reference 5**). In these setups, one set of daisy-chained chips makes the measurements, and another set of chips supervises the measurements using comparators to set minimum and maximum limits (see **figures** illustrating these effects in the Web version of this article at www.edn.com/110120cs). Manufacturers such as Intersil strive to ensure that the chip-to-chip serial communication is passive because, if one chip fails electrically, it will still pass the communications of all the other chips in the daisy-chain string. Most battery-stack-measurement chips also have bidirectional serial communication, allowing you to query the chip with your system's BMU (battery-management-unit) microcontroller to ensure that the chip is powered up and working.

In addition to this system-level redundancy, many manufacturers build redundancy and self-test into their chips (**Reference 6**). Kenneth Lenk, marketing manager of automotive products at Intersil, notes that chips include multiple voltage references. The company also incorporates hidden DACs in its chips for calibration and self-test, he adds.

Most manufacturers stress that their devices have passive communications links. These links will continue to work even if the chip fails. Manufacturers must not only build in internal redundancy but also ensure that the redundant safety chip is working, according to

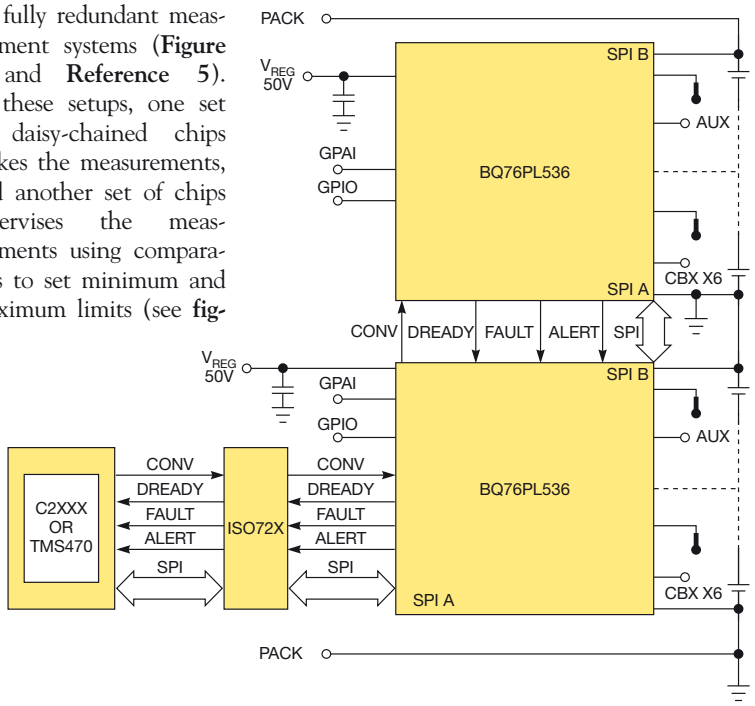

Figure 4 Most chip companies design their battery-monitor chips to operate in a daisy chain so you can use multiple chips to measure dozens of cells (courtesy Texas Instruments).

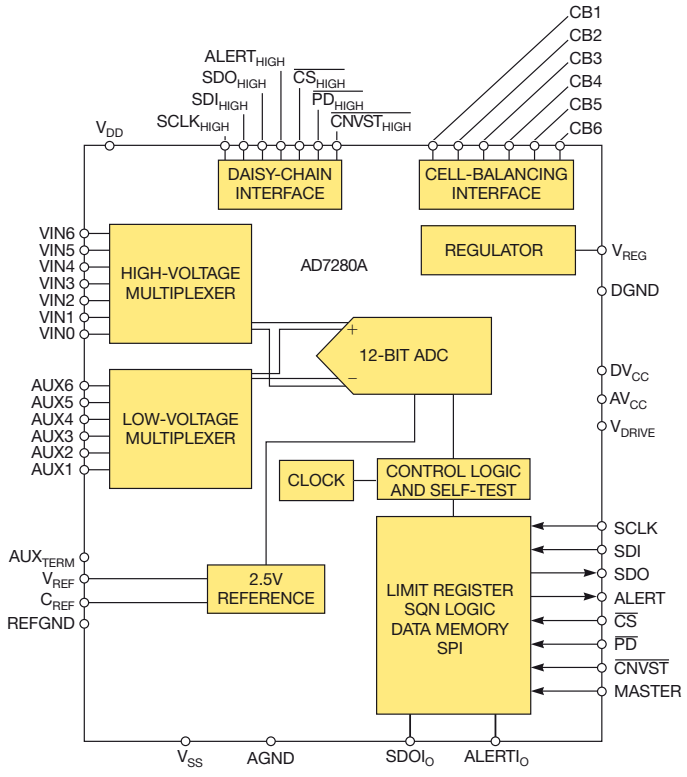

Figure 5 Analog Devices lets you power the chips from an external source or the battery you are monitoring. The company provides six auxiliary inputs for temperature or other measurements.

Sam Weinstein, a precision-amplifier-product-line manager at Analog Devices. He notes that BIST (built-in self-test) is expensive but essential to satisfying the requirements of the auto industry.

An engineering committee is working to formalize the fault protection and redundancy features of automotive-battery-stack systems into the ISO (International Organization for Standardization) 26262 standard, which the organization expects to issue this year. The developers adopted this standard from industrial-machinery standards, and it will provide comprehensive guidelines for the analog, digital, and software components in electric vehicles. Companies such as Texas Instruments, Analog Devices, STMicroelectronics, and NXP are working to provide the analog and digital hardware for these mission-critical advanced-powertrain modules.

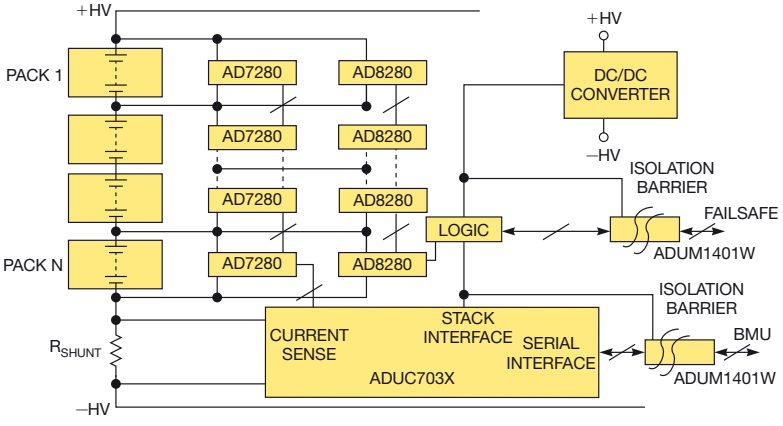

Figure 6 Auto makers require redundant measuring systems with chips that monitor other chips (courtesy Analog Devices).

COMMON-MODE PROBLEMS

You will run into the common-mode-measurement problem whenever you try to resolve a small voltage difference that sits atop a much larger voltage. In a vehicle's battery stack, for example, a 4V cell may sit atop a 400V stack. Your first instinct might be to make a resistive voltage divider to reduce the voltage to a smaller value—4V, for example.

Although formal ways exist for analyzing the accuracies of the resistors, you can also use an intuitive method. If 10 bits equals 1 mV on 1V, 10 bits would equal 4 mV on a 4V battery. For error budgets, you would make the measurement with 12-bit converters, a job that stack-monitor chips are now performing.

When you use resistive dividers, however, you reduce both the measurement voltage and the common-mode voltage. The factor of 100 in this example means that you are trying to resolve 0.00004V, or 40 µV. With a 2-bit error margin, you would need to resolve 10 µV—a more challenging problem that also brings noise concerns.

The increased accuracy requirement is not the biggest drawback of resistive dividers, however. The approach also requires extremely low-tolerance resistors. For example, consider a voltage divider that is making 4V from 400V. The lower resistor's normalized resistance would be 4Ω, the upper resistor's value would be 396Ω, and the nominal tap voltage would be 4V. If you were to increase the upper resistor's voltage by 1%, you'd get a voltage of 399.96V, rounded to 400V. The string has 404Ω resistance; 400V across that resistance yields 3.96V instead of 4V at the tap: a 40-mV error.

The LSB (least-significant bit) is 10 µV, so a 1% error in one resistor would yield 4000 times more error in the measurement than is acceptable. A factor of 1000 implies a resistor-accuracy spec of 0.001%. Factoring in the 4V yields a 0.00025% accuracy spec.

With two 400V resistor dividers, one for each side of the cell you are measuring, one resistor can be out of tolerance, meaning that you have to divide by two for a resistor-accuracy spec of 0.000125%.

Nevertheless, a resistor-tolerance spec of ±125 ppm is not feasible for a low-cost design. Temperature-coefficient problems will also crop up, so you must match the dropping resistors over temperature changes. You should instead directly measure the 4V cell and then transfer that analog voltage down to the chassis reference.

Alternatively, you could place the ADC at the cell voltage and then just transfer the bits across an isolation boundary—the approach that most analog-IC companies take in their battery-stack-monitor chips.

BATTERY ENVIRONMENTS

In addition to measuring automotive battery stacks, you must also engineer the system to survive in the harsh environment that modern vehicles experience. All of their components are subject to vibration and acceleration. Some of the greatest acceleration occurs when shipping cars by rail with chained-down suspensions. Surface-mount chips and passive parts are vibration-resistant.

Your systems must also withstand wider temperature ranges than consumer electronics require. Battery cells cannot withstand a temperature of 125°C, for example. Still, most chips operate at temperatures as high as 85°C, and Maxim's and Analog Devices' chips work at temperatures as high as 105°C. Intersil and other manufacturers provide chips that work at a temperature the auto maker specifies. Low temperatures also cause problems. Transistors' base-emitter junction voltages and transconductance rise with decreasing temperature, causing amplifier oscillation.

You must design your battery-stack-measurement systems to withstand EOS (electrical overstress). This phenomenon can occur, for example, when a mechanic disconnects the battery cable of an operating engine to determine whether the alternator is working. In this case, the alternator could place a 100V pulse into the electrical system. Although electric-vehicle-battery-stack chips might not be subject to this stress, the bus bar connecting cells can break while large currents are flowing, causing a large overshoot in the battery voltage.

EMI interferes with measurements and is one of the biggest environmental challenges in electric vehicles (**Figure 7**). All of the traces and high-impedance nodes are subject to EMI, which can ruin the cell-voltage measurement. According to Tim Regan, application manager at Linear Technology, ac ripple can show up anywhere. That ripple is caused by the inverter's chopping frequency and adds to the electrical noise from the motor.

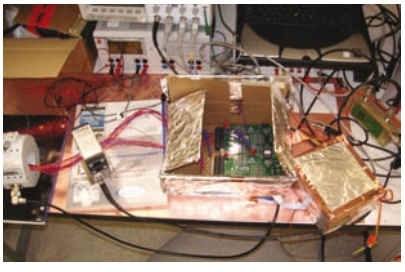

Figure 7 This lab setup injects currents into the measurement system to ensure that the chip can withstand abuse. The foil-covered boxes keep external radiation from affecting the measurements (courtesy Linear Technology).

"Basic decoupling works wonders," Regan adds. Decoupling is only a starting point, however. You must also pay attention to noise sources, PCB (printed-circuit-board) layout, and shielding.

Once you remove EMI from your measurement, you still need to consider the fact that EMI can cause a loss in the serial communications between chips in the daisy chain. The communication links between chips may have different lengths. "You might have 400 to 600V potentials across the links," says Intersil's Lenke.

Electric-vehicle designers also must deal with the issue of magnetic fields, which arise due to the large currents around the vehicle that shuttle between the charger, the battery, and the motor. It is difficult to shield magnetic fields. Doing so requires heavy metal or steel plates to keep the fields from the electronics. As with all other noise prob-

lems, it is a good idea to address this problem at its source. Keep all large currents in small loop areas.

Cost is an overriding problem in automotive-battery-stack-monitor systems. Mass-market automobiles cannot have a dozen $60 chips in the battery pack. In this regard, automotive design is even more difficult than military design because automotive companies cannot spend unlimited amounts of money to solve the vibration, temperature, and high-performance requirements of the market. Chips must feature high performance and low cost and be available in high volume and with good yields, says Erik Soule, general manager of signal-conditioning products at Linear.

Analog companies have made a considerable achievement with their battery-stack-monitor chips. The parts must be highly accurate, small, and robust because the applications involve EMI and electrical overstress. Auto makers require redundancy and fault protection. Although vendors provide evaluation boards, it is doubtful whether you can strap that board somewhere in your system and expect it to work properly. Instead, you will have to understand the measurement, noise, and interference problems in your application and then apply good design and layout techniques. With careful design and judicious shielding, you can make a monitor system that will keep a vehicle's propulsion system operating for a decade or more. EDN

FOR MORE INFORMATION

Analog Devices
www.analog.com

Avago
www.avagotech.com

Infineon
www.infineon.com

Intersil
www.intersil.com

Linear Technology
www.linear.com

NXP Semiconductors
www.nxp.com

Maxim Integrated Products
www.maxim-ic.com

STMicroelectronics
www.st.com

Texas Instruments
www.ti.com

REFERENCES

1. Hsu, JS; CW Ayers; and CL Coomer, "Report on Toyota/Prius Motor Design and Manufacturing Assessment," ORNL/TM-2004/137, Oak Ridge National Laboratory, July 2004, http://bit.ly/fzizOk.
2. Rako, Paul, "Voltage references hold steady," *EDN*, Oct 21, 2010, pg 38, http://bit.ly/eY1zpS.
3. Williams, Jim, "Novel measurement circuit eases battery-stack-cell design," *EDN*, Jan 10, 2008, pg 47, http://bit.ly/hoFWow.
4. Rako, Paul, "Draw the line: Isolation shields systems from shocking surprises," *EDN*, Sept 3, 2009, pg 22, http://bit.ly/ePO0Hf.
5. Boyle, Steven, "Lithium Ion Battery Monitoring–Redundancy Required?" *ECN*, Jan 28, 2010, http://bit.ly/eOM7a0.
6. Allen, Tony, "Monitoring systems ensure Li-ion battery safety and efficiency," *EE Times*, Dec 2, 2010, http://bit.ly/eOkGM7.

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

BY PERRY KELLER • AGILENT TECHNOLOGIES

Managing signal integrity in tomorrow's high-speed flash-memory-system designs

NEXT-GENERATION FLASH-MEMORY TECHNOLOGY WILL TOUT DATA-TRANSFER RATES AS MUCH AS 10 TIMES FASTER THAN CURRENTLY AVAILABLE. HOWEVER, INCREASING DISTORTIONS IN THE DATA-CARRYING DIGITAL SIGNALS CAN CAUSE DATA-TRANSFER FAILURES, COMPLICATING THE MANAGEMENT OF SIGNAL INTEGRITY. PROPER DESIGN STRATEGIES CAN HELP YOU DELIVER RELIABLE, HIGH-PERFORMANCE PRODUCTS.

Signal integrity will matter more in next-generation flash devices, and several changes in the technology will make managing signal integrity more critical. For example, data rates in these devices will range from 400 MHz to 6 GHz. To support the faster data rates, edge rates will have to become 10 to 100 times faster. Demand for increased storage capacity will also drive the need for denser packaging and more complex interconnects. Pressure to reduce costs will force you to make trade-offs that affect signal integrity, such as using lower-cost materials or even eliminating the use of ground planes.

As rise times become shorter, a signal's high-frequency components become more pronounced. Higher-frequency signals are more sensitive to interconnect quality, so signal-integrity problems tend to proliferate. As signal frequencies increase, signal loss also increases. Therefore, high-frequency components of fast-rise-time signals experience more loss than the low-frequency components, leading to signal distortion and ISI (intersymbol interference).

SIGNAL-INTEGRITY PROBLEMS

Understanding the causes of signal-integrity problems and their remedies is critical. Signal-integrity problems include reflections and distortions, crosstalk, ground bounce, and jitter. Reflections and distortions relate to signal quality on an individual net. When a signal encounters an impedance discontinuity, it generates a reflection that becomes further distorted as it continues along the net. The reflection travels from the impedance discontinuity in two directions—toward the receiver and back to the driver. The reflections themselves react to other discontinuities, creating further reflections that distort the true signal in complex ways, generating effects such as ringing, overshoot, and slope reversal. Careful design to maintain well-controlled impedance along key traces is the best way to improve signal integrity.

Crosstalk-induced signal-integrity problems involve multiple signal nets. If you place an active net near a quiet one, capacitive and inductive coupling can cause some of the energy from the signal on the active net to couple over to the quiet side (**Figure 1**).

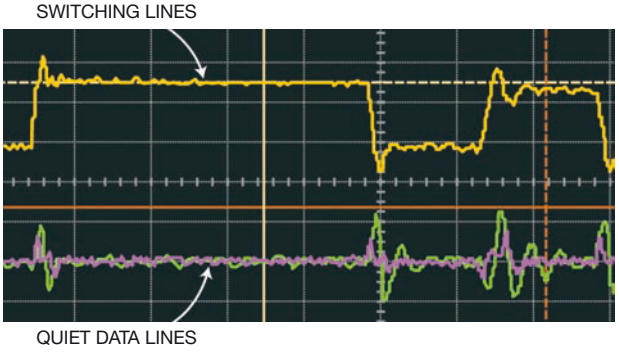

Figure 1 A fast transition on the upper trace couples into the lower trace, generating noise on the lower trace.

The quality of grounding and current-return paths in your design is among the biggest factors affecting crosstalk. In most PCB (printed-circuit-board) designs, ground planes are available as return paths. This approach is best if you can afford the extra plane. If cost pressure forces you to eliminate using a ground plane, you must use other strategies, such as placing a ground trace next to the signal or using differential instead of single-ended signaling.

Ground bounce also affects signal integrity and relates to power distribution. As with any network that has interconnects, inductance exists in power and ground networks. As the I/O signals transition from zero to one or one to zero, transient current flows in the power-distribution network. Many signals' switching at once generates large transient currents. Any inductance or resistance in the power- and ground-distribution network converts these transient currents into voltage spikes that appear as noise in other signals or even as a shift in the ground voltage. Ground planes or multiple ground or power connections reduce the impedance and therefore the SSO (simultaneous-switching-output) noise. Using lower voltage swings and protocols that minimize the number of signal transitions also helps.

Jitter issues also affect signal integrity. Reflections, crosstalk, and SSO all can contribute to jitter. In addition, ISI created

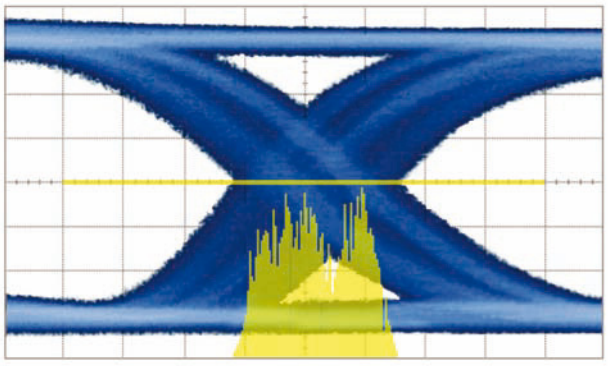

Figure 2 Knowing whether the jitter (yellow trace) is random, periodic, or correlated to some other event in the system helps you determine the best ways to address the problem.

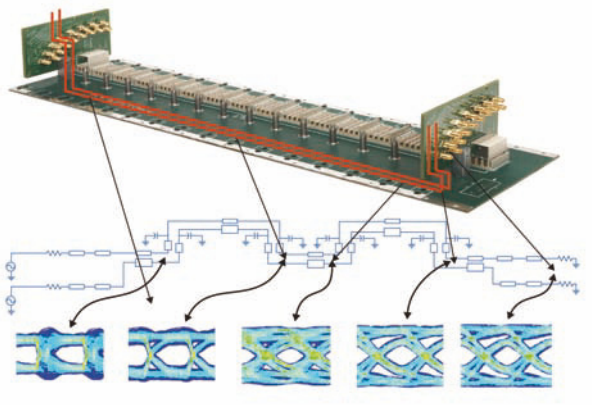

Figure 3 Physical measurements can confirm an accurate simulation model to accurately predict signal integrity at various nodes in the design.

on lossy channels, PLL (phase-locked-loop) noise, EMI (electromagnetic interference), differences in transmitting and receiving threshold voltages, and ordinary delay mismatches in internal logic can generate jitter. The strategy for managing jitter differs, depending on the cause of the jitter. Proper shielding can help with EMI-induced jitter, but it cannot fix a noisy power supply. Knowing whether the jitter is random, periodic, or correlated to some other event in the system helps you determine the best ways to address the problem (**Figure 2**).

MANAGE SIGNAL INTEGRITY IN YOUR DESIGNS

As every RF engineer knows, everything in a circuit can affect the signal. To manage signal integrity, it is critically important to first identify the parts of the design that affect signal integrity. A common approach is to start by creating a model of your design and its components and interconnects. However, a model typically is less accurate than it needs to be. You must make measurements of your circuits, compare them with your model, and adjust the model to make it consistent with your measurements. Once the model is accurate, you can use the simulator to predict which changes will improve signal integrity. Pay particular attention to vias, wire bonds, packages, PCB traces, and connectors when considering components that will affect signal integrity.

The goal of simulation and modeling is to predict the real-world behavior of your design. Engineers have traditionally used modern design and simulation environments, such as Agilent's ADS (Advanced Design System), for microwave and RF design. As digital speeds approach microwave speeds, engineers have been applying these tools to digital design, especially for evaluating signal integrity. These tools accurately simulate high-speed effects, such as distortion, mismatch, and crosstalk, in your channels.

Integration of system, circuit, and EM (electromagnetic) simulators provides accurate answers and avoids error-prone and time-consuming data transfer among point tools (**Figure 3**). The ADS model in the **figure** can assess signal integrity and predict eye diagrams at various nodes in the backplane. You can see the eye after the daughtercard, at the high-speed backplane connector, and after the backplane traces. From these eye diagrams, you can determine where signal-integrity problems caused the eye to degrade and at what point you should modify the design. An accurate model such as this one gives you insight into your design and lets you rapidly evaluate changes that will improve its performance.

PHYSICAL-MEASUREMENT TOOLS

Making physical measurements is key to assuring the accuracy of your model and validating the final performance of your design. At the speeds of next-generation flash design, it is important to analyze the data in both the time domain and the frequency domain. You can make physical measurements

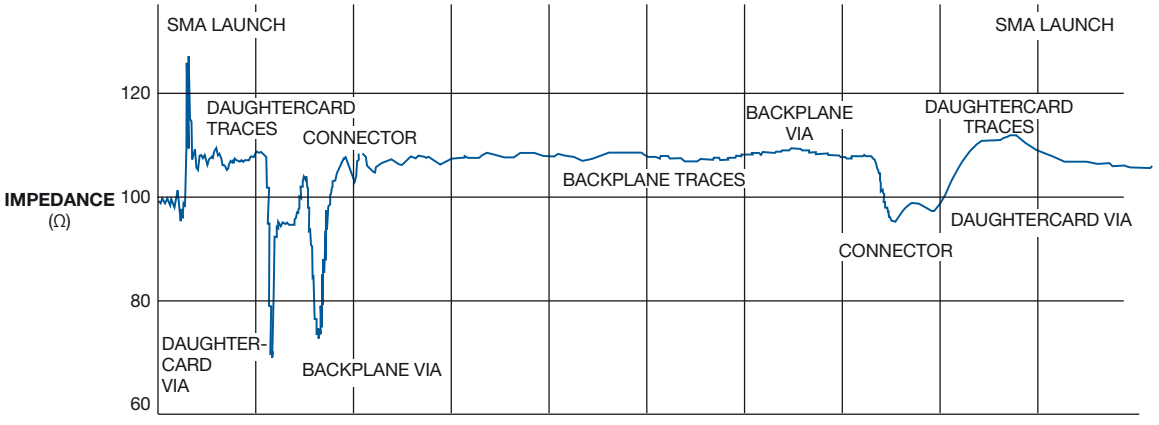

Figure 4 You can use time-domain reflectometry to measure interconnect impedance at each point in your design.

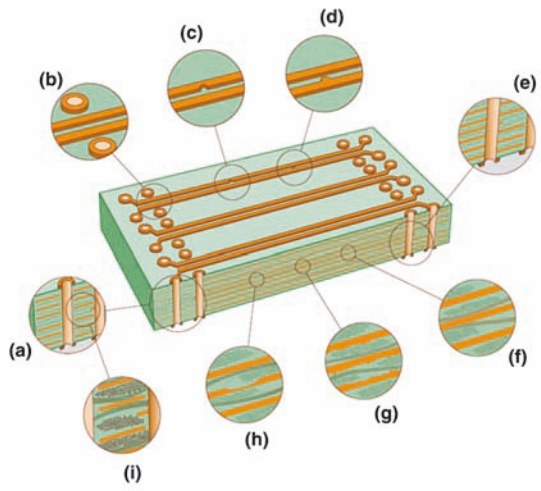

Figure 5 Many signal-integrity problems can appear on a PCB. These problems include excess capacitance of a through hole, which can affect differential-transmission-line characteristics (a). You can address this problem by using a larger antipad. Another problem is localized crosstalk between vias in close proximity to a differential transmission line (b), which you can fix by increasing trace spacing. Fix varying conductor width (c) by increasing the trace width. If conductor spacing is uneven (d), increase the separation between conductors. Via stubs can create capacitive reflections and load down transmission-line impedance (e). Consider instead using blind, buried, or back-drilled vias. If the PCB has a nonuniform dielectric (f), try using low-flow preimpregnated composite fiber. Some surface treatments can cause thickness nonuniformity (g), which you can fix by using example plating and finishing instructions. Adjust materials or processes to handle localized changes in foil thickness (h). Gaps in the lamination layers of epoxy-resin and fiberglass can cause anodic-conductive-filament shorting (i). To deal with this problem, use extra nonconductive layers or thicker layers.

with a time-domain instrument, such as a TDR (time-domain reflectometer), or a frequency-domain instrument, such as a VNA (vector network analyzer).

If you're new to measuring interconnect behavior, you may want to consider starting with a TDR, such as the Agilent 54754A differential and single-ended TDR module. It provides an intuitive waveform and good first-pass model. For greater accuracy and higher bandwidth, you'll need to learn S parameters and use a VNA, such as the Agilent N5241A. Another possibility is to use tools such as Agilent's PLTS (physical-layer test system). With the PLTS, you can work in either the frequency domain or the time domain, whichever best suits your requirements, regardless of whether you are using a TDR or a VNA (**Figure 4**).

As the signal travels through the physical structures, the TDR measures and displays impedance. Ideally, the impedance trace is a flat line, indicating no discontinuities. At the daughtercard via, you can see a large discontinuity, so you know you must modify your design to raise its impedance.

MANAGING INTERCONNECT DESIGN ON PCBs

Active circuitry, packaging, and connectors all influence signal integrity; PCBs can also influence signal integrity, often in difficult-to-detect ways. Competing demands for high speed and low cost often collide on the PCB. For example, FR4 is low cost and has relatively good performance at lower data rates. However, when data rates exceed 10 Gbps, problems increase dramatically. **Figure 5** should help you identify which elements of your design might be causing problems and suggests ways to deal with those problems.

The next generation of flash will enable next-generation performance, but next-generation speeds will require designers to pay increased attention to signal integrity. Understanding what can affect signal integrity and how to model and measure it will help you deliver reliable, high-performance products on time and on budget. EDN

AUTHOR'S BIOGRAPHY

Perry Keller is program lead for applications and standards at Agilent Technologies, where he has more than 25 years of experience in software and system engineering, high-speed-hardware and ASIC design and validation, and product marketing. Keller has a master's degree in electrical engineering from Rice University (Houston).

EDITED BY MARTIN ROWE
AND FRAN GRANVILLE

designideas

READERS SOLVE DESIGN PROBLEMS

Oscillator has voltage-controlled duty cycle

Luca Bruno, ITIS Hensemberger Monza, Lissone, Italy

The classic Colpitts oscillator circuit in **Figure 1** generates a clock signal with a fixed duty cycle. By replacing CMOS inverter gate IC_{1A} with a voltage comparator (**Figure 2**), you can obtain a more versatile and more useful clock generator. You can set not only the oscillation frequency but also the duty cycle. You must use the comparator in an inverting configuration, which introduces a 180° phase shift. That shift, along with an additional phase shift of 180° from the capacitor-input network, lets the circuit oscillate. The circuit compares the sine wave at the output of the capacitor-filter network with the reference voltage, which drives the output of the threshold comparator high and low.

The network sets the oscillation frequency as follows:

$$f_O=\frac{1}{2\pi\sqrt{LC/2}},$$

where f_O is the oscillation frequency. With a suitable choice of the values of inductors and capacitors, the circuit can oscillate at frequencies as high as 10 MHz.

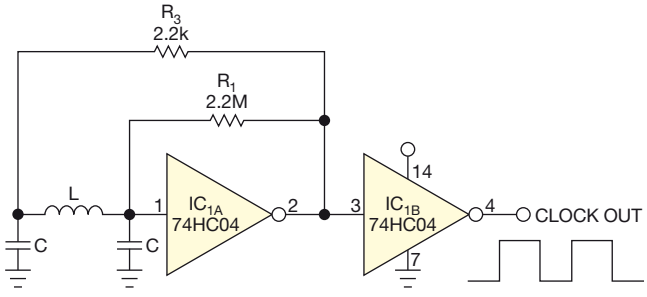

Figure 1 A typical Colpitts oscillator uses an inverter and has a fixed duty cycle.

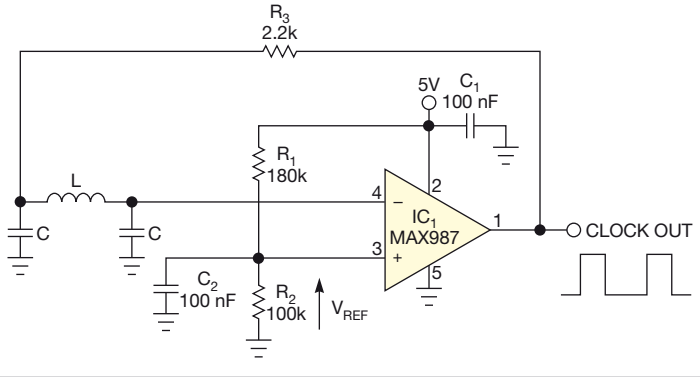

Figure 2 A comparator lets you vary the duty cycle based on resistors R_1 and R_2.

DIs Inside

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

The output clock's duty cycle depends on the reference voltage, which you can easily set through the voltage divider comprising R_1 and R_2. Unfortunately, the mathematical relationship between the reference voltage and the duty cycle is nonlinear because the sine wave at the output of the capacitor-input-filter network is not a linear function. Also, its amplitude is not constant but depends on the duty cycle of the output clock. You can easily obtain this mathematical relationship by testing the circuit with an inductance of 10 µH and a capacitance of 10 nF.

You can use any high-speed compara-

TABLE 1 OSCILLATOR DUTY CYCLE BASED ON REFERENCE VOLTAGE

Reference voltage (V)	Duty cycle (%)
0.5	15.2
1	28.3
1.5	37
2	43.5
2.5	50
3	56
3.5	62.6
4	71.5
4.5	85.4

tor with rail-to-rail inputs and outputs, such as the MAX987 from Maxim (www.maxim-ic.com), to achieve a wider input range for the reference voltage. That wider range gives you wider control of the duty cycle, although you can't reach the minimum duty cycle of 0% or the maximum duty cycle of 100%.

The propagation delay, T_{PD}, of the comparator introduces a further phase shift of value $\Delta\phi=2\pi f_O T_{PD}$, where $\Delta\phi$ is the phase shift. The capacitor-input network compensates for the phase shift, slightly reducing the oscillation frequency. For safe operation of the circuit, you should vary the reference voltage by 0.5 to 4.5V. The duty cycle varies from approximately 15 to 85% (**Table 1**). You can produce a bipolar output signal if you use a dual power supply. EDN

Generate noisy sine waves with a sound card

José M Miguel, RF-Electronics Ltd, Barcelona, Spain

Testing audio-noise-reduction circuits, PLLs (phase-locked loops), and audio-frequency filters may require a noisy sine wave, one that is summed with white noise. Using a typical computer sound card, free software, and an external amplifier circuit, you can create a noisy sine wave.

Free Generatosaur software from Wavosaur (www.wavosaur.com) turns your sound card into a low-frequency wave generator. It lets you independently choose amplitude, frequency, and waveform for the left and the right channels. The Generatosaur's user interface is a dialogue-box-style control panel (see **figure** at www.edn.com/110120dia). If you select a sine wave for the left channel and a white noise for the right channel, you then need only to use an amplifier to add the signals. **Figure 1** shows the complete circuit.

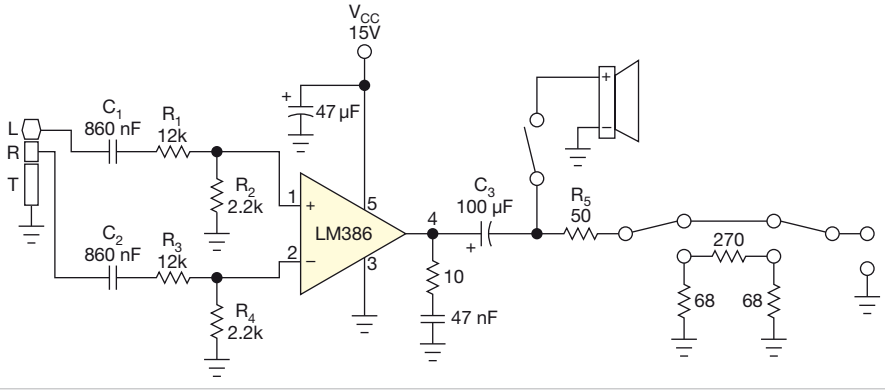

Figure 1 Generatosaur is free software that turns your sound card into a low-frequency wave generator.

The differential amplifier employs a National Semiconductor (www.national.com) LM386 audio power amplifier with a supply voltage of 15V. The output of the LM386 has a self-centered quiescent voltage that is half the power-supply voltage and that requires a blocking capacitor, C_3. Resistor R_5 sets the output impedance to 50Ω. You need the voltage dividers R_1/R_2 and R_3/R_4 because the output-voltage range for a standard sound card is 0 to 2V. Taking into account that the voltage gain of the LM386 amplifier is internally set to 20 and that its output voltage range is 7V, you need an attenuation factor, K, of 7/(2×20) in each amplifier input. The circuit also includes a selectable 20-dB attenuator that you can invoke with the two DPDT (double-pole/double-throw) switches.

Another **figure**, also available at www.edn.com/110120dia, shows a 450-Hz sine wave with a 10-dB SNR (signal-to-noise ratio). The root-mean-square noise voltage of this signal is 0.5V measured on an oscilloscope and following the tangential method. If you need to hear the generated noisy signal, connect a loudspeaker to the output of IC LM386. EDN

Decode a quadrature encoder in software

Sid Levingston, Gentec-EO, Lake Oswego, OR

Quadrature encoders work in many applications to determine displacement and direction of mechanical travel. They vary in design, but they all do the same thing: supply a set of square waves 90° out of phase. **Figure 1** shows the typical output signals.

The encoder rotates clockwise when Channel A leads Channel B. If Channel B leads Channel A, the encoder is rotating counter clockwise. By counting the pulses and the direction of rotation, you can find the position of the encoder. Although ICs can decode quadrature encoders, you can easily and less expensively have the processor decode the signal. The signals from Channel A and Channel B go through a Schmitt trigger if necessary, but many encoders and processors include this trigger internally. The signals are then applied to two I/O pins on the processor that support edge-triggered inter-

Industry's lowest power port expanders and LED drivers

Drive LEDs with no additional components in portable applications

- Controlled brightness with PWM and constant current
- Low 1.62V operating voltage
- I²C and SPI™ serial interfaces
- Low 1.5µA (max) standby current
- Small chip-scale package (2mm x 2mm)
- EV kits available

	Part	No. of Ports	Interface Type	Supply Voltage Range (V)	Sink Current (mA, per Pin)	Standby Current (µA at +85°C)	Features
New	**MAX7302**	9	I²C	1.62 to 3.6	25	2.0	4 slave IDs, PWM, blink, $\overline{RST}$
New	**MAX7306/07**	4	I²C	1.62 to 3.6	25	2.0	4 slave IDs, PWM, blink, $\overline{RST}$
	MAX7319–23	8	I²C	1.71 to 5.5	20	1.9	16 slave IDs
	MAX7324–27	16	I²C	1.71 to 5.5	20	1.9	16 slave IDs
	MAX6966/67	10	SPI	2.25 to 3.6	20	1.9	1.5% constant current, PWM
Now offered in CSP	**MAX6946**/47	10	I²C	2.25 to 3.6	20	1.5	1.5% constant current, PWM, **2mm x 2mm CSP**
	MAX7315	8	I²C	2 to 3.6	50	3.3	64 slave IDs, PWM
	MAX7316	10	I²C	2 to 3.6	50	3.3	64 slave IDs, PWM, $\overline{RST}$, INT
	MAX6964	18	I²C	2 to 3.6	50	3.3	4 slave IDs, PWM, blink
	MAX7300/01	20/28	I²C/SPI	2.5 to 5.5	10	11	16 slave IDs

SPI is a trademark of Motorola, Inc.

www.maxim-ic.com/GPIO-info

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

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

MAXIM
INNOVATION DELIVERED®

For free samples or technical support, visit our website.

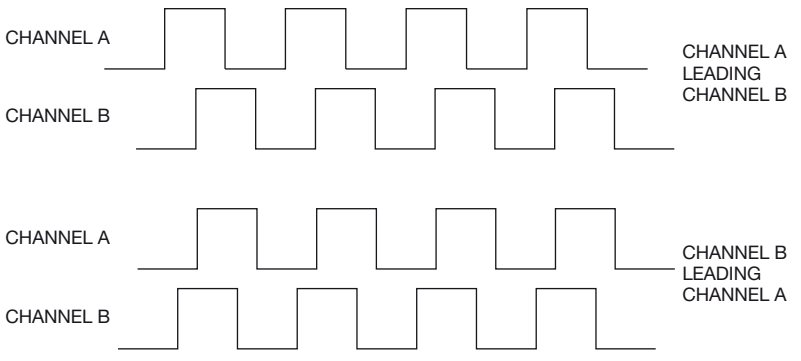

Figure 1 A quadrature encoder produces pulses that are 90° out of phase.

rupts. The code in the interrupt handler implements a standard decoder algorithm, but all algorithms typically follow these steps:

1. Set up a state table like the one in **Table 1**. The table wraps around from State 3 to State 0.
2. Initialize a counter.
3. Measure the current state of Channel A and Channel B. Find that state in the table and set a pointer to it.
4. Enable the interrupts.

In the interrupt handler, use the following steps:

1. Read the state of Channel A and Channel B.
2. If the state is the one preceding the pointer, decrement the counter.
3. If the state is the one following the pointer, increment the counter.
4. Set the pointer to the new state.
5. Clear the interrupt.

This method requires that a state table exists, that the previous state remain, and that the handler determine on each interrupt which of four states exists and then make a decision based on two possible conditions. The handler accomplishes this task with a four-case switch, in which each case has two if conditions.

Now, consider what happens in the real world. If the I/O pin generates an interrupt on a rising edge, then when the interrupt happens, that channel goes from low to high. Therefore, there's no reason to read the state of the pin that interrupted. The other channel did not interrupt because the signals are 90° out of phase. So, to determine the current state, you need only to read the state of the pin that didn't initiate the interrupt. The state of the unchanged low or high signal tells which way the encoder rotated. If it is low, then the interrupting pin is leading. If it is high, then the interrupting pin is trailing. You can use these facts to implement an efficient interrupt handler with no state table and no memory of a previous state.

The code in **Listing 1** was tested on an MSP430F processor connected to an Encoder Products (www.encoder.com) model 15T. The encoder monitored the position of a linear stage. The stage traveled 85 mm and could be tracked with resolution of 5 microns.

The define statements (highlighted in red) make the code more readable. The initPort() function (highlighted in blue) sets up the rising-edge interrupt on channels A and B. The final piece is the interrupt handler (highlighted in green). Note that it contains only six lines of code compared with the 20 or 30 lines it would take to implement the traditional method of decoding the channels. EDN

TABLE 1 LOGIC STATES OF QUADRATURE ENCODER

	Channel A	Channel B
State 0	0	0
State 1	1	0
State 2	1	1
State 3	0	1

LISTING 1 QUADRATURE-ENCODER CODE

```
#define CHA BIT0
#define CHB BIT1

initPort()
{
  P1SEL = 0;
  P1DIR = 0;
  P1OUT = 0;
  P1IES = 0;    // rising edge.
  P1IE = CHA + CHB;    // interrupt on CH A or CH B rising edge
}

#pragma vector = PORT1_VECTOR
__interrupt void port1_ISR (void)
{

  if(P1IFG & CHA )  // who interrupted? If A then A is high
   {
    (P1IN & CHB) ? gStepCount--  : gStepCount++ ;   // test B
     P1IFG &= ~CHA;    // clear interrupt
   }

  if(P1IFG & CHB ) // who interrupted? If B then B is high

   {
    (P1IN & CHA) ? gStepCount++  : gStepCount-- ;  // test A
     P1IFG &= ~CHB;     // clear  interrupt
   }
}
```

Power an LED driver using off-the-shelf components

TA Babu, Chennai, India

High-power LEDs challenge electronics engineers to design accurate and efficient, yet simple, driver circuits. Conventionally, driving high-power strings with accurate current requires dedicated switching regulators. Choosing a discrete driver circuit requires an understanding of LED lighting to make the best trade-off. This Design Idea describes a simpler and equally good way to employ the ubiquitous 555 IC.

In the converter circuit in **Figure 1**, IC_1's pins 2 and 6 connect together, which lets the device retrigger itself on each cycle. Thus, it operates as a free-running oscillator. During each cycle, capacitor C_2 charges up through timing resistor R_1 and discharges through resistor R_2. The capacitor charges up to two-thirds of the power-supply voltage, the upper comparator limit, which $0.693(R_1C_2)$ determines, and discharges itself down to one-third the power-supply voltage, the lower comparator limit, which $0.693(R_2C_2)$ determines. The total time period, T, is $0.693(R_1+R_2)C_2$.

ONCE THE VOLTAGE DROP REACHES THE BASE-EMITTER THRESHOLD OF TRANSISTOR Q_2, IT STARTS CONDUCTING.

During the on time, transistor Q_1 conducts and stores the energy in inductor L_1. When it stops conduction, the stored energy transfers to capacitor C_3 through Schottky diode D_1.

You can use the following **equations** to calculate the inductor value. The selection of an inductor depends on input voltage, output voltage, maximum current, switching frequency, and availability of standard inductor values. Once you know the inductance, you can choose the diode and the capacitor.

MOSFET Q_1 determines the duty cycle, according to the following **equation**:

$$D=1-\frac{V_{INMIN}\times\eta}{V_{OUT}},$$

where V_{INMIN} is the minimum input voltage, V_{OUT} is the desired output voltage, and η is the efficiency of the converter, estimated at 80%.

The average inductor current is

$$I_{LAVG}=\frac{I_O}{1-D},$$

where I_{LAVG} is the average inductor current and I_O is the output current.

The peak inductor current is

$$I_{LPEAK}=I_{LAVG}+\frac{\Delta I_L}{2},$$

where I_{LPEAK} is the peak inductor current and ΔI_L is the change in inductor current.

Assume that the change in inductor current is 25% over the average current. You can compute inductor L_1 as

$$L=(V_{IN}\times D)/(F_{OSC}\times\Delta I_L),$$

where F_{OSC} is the oscillator frequency. The inductor's saturation-current rating should be greater than the peak current.

To ensure constant illumination, you must monitor the current through the LED. Resistor R_3 senses the output current. Once the voltage drop across this resistor reaches the base-emitter threshold of transistor Q_2, it starts conducting, and this conduction reduces the on time of the 555 timer.

The following **equation** thus sets the LED current:

$$I_{LED}=\frac{0.6V}{R_{SENSE}},$$

where I_{LED} is the LED's current and R_{SENSE} is the sense resistance.

The minimum and maximum input and output voltages for this circuit are 10.5 and 15V, respectively. The LED string's voltage and current are 21V and 350 mA, respectively. The 6W LED driver can find numerous applications, including battery-operated portable lighting, solar-operated garden lighting, automotive lighting, bike headlights, and underwater lights. Driving high-power LED strings with standard off-the-shelf components simplifies your design without sacrificing performance. EDN

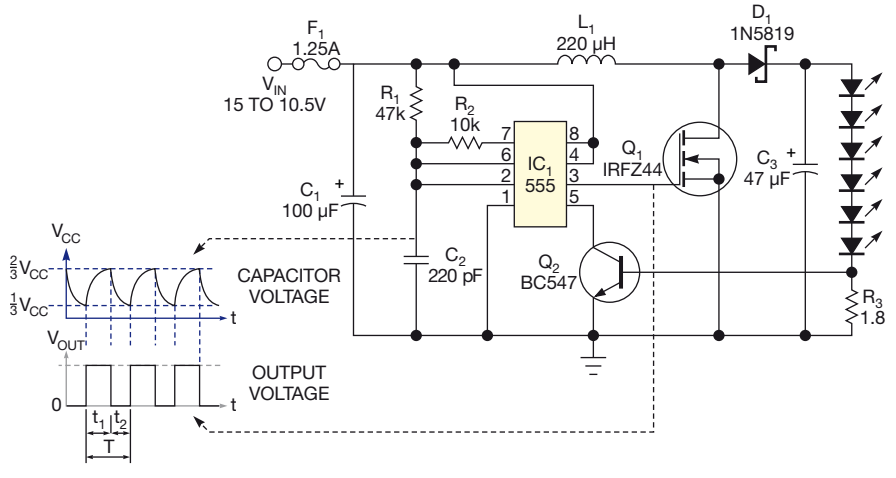

Figure 1 A free-running 555 timer provides the stimulus for driving a series of LEDs.

productroundup

OPTOELECTRONICS AND DISPLAYS

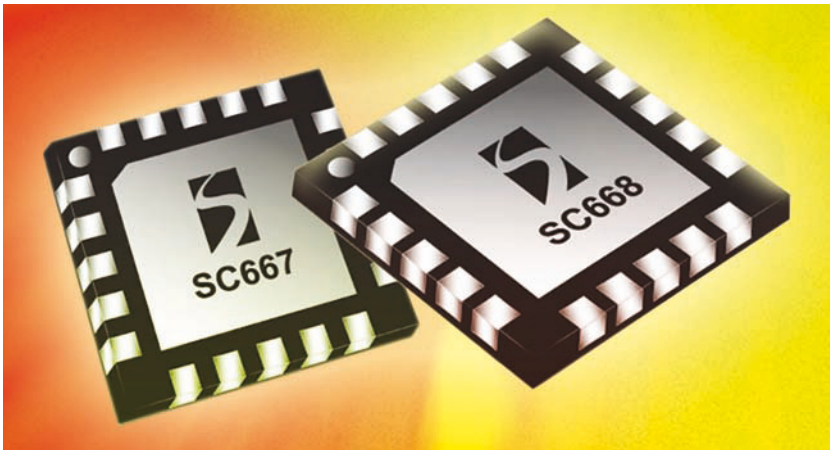

Smart current-sink LED-backlighting platform targets high-end handheld displays

The SC667 and SC668 smart current-sink LED-backlighting chips with on-chip digital-lighting effects provide the ability to incorporate fade, breathe, and blink effects without changing the firmware. The devices feature automatic dropout prevention, enabling them to reduce the total parts count and extend battery life over that of boost converters and charge pumps in high-end handheld displays. The devices eliminate the need for the capacitors and inductors that boost circuitry requires, reducing component count, board size, and system cost; eliminating switching cost; and extending operating time. The SC667 and SC668 sell for $1.15 and $1.29 (3000), respectively.

Semtech Corp, www.semtech.com

Single-die LEDs deliver 1000 lumens at 100 lumens per watt

The XLamp XM-L LEDs target use in very-high-lumen applications, such as roadway lighting. The 6500K, cool-white LEDs deliver 1000 lumens with 100-lumen/W efficiency at 3A. In a 5×5-mm footprint, the devices deliver 160 lumens/W at 350 mA and as much as 315 and 150 lumens/W at 700 mA. Prices start at approximately $4.50 (low volumes).

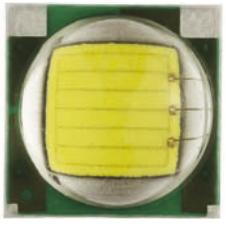

Cree Inc, www.cree.com

LED-driver IC touts 92% power factor

The R2A20134 LED-driver IC allows designers to select MOSFETs with a voltage tolerance of 300 to 500V, rather than the typical 700V. A high-voltage control and critical-conduction mode enable zero-current switching at turn-on and switching in a low-drain voltage state. Power factor is 92%. Designers can achieve a variety of control modes, including buck-boost, average-current, peak-current, and constant-input-power control, by changing the configuration of the externally connected elements. Designers can also use the IC in TRIAC-dimmer applications by applying a voltage matching the TRIAC phase to the feedback pin for constant-current control. Samples are available for $1 per unit.

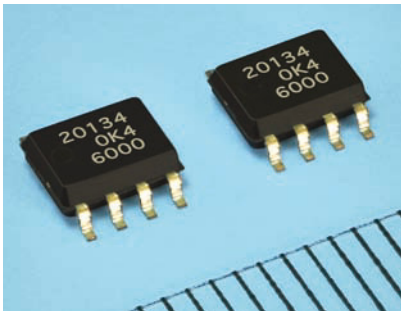

Renesas Electronics America Inc, www.renesas.com

Visible-color LEDs target military, harsh-environment applications

The OVM4x series of visible-color LEDs features hermetically sealed TO-46 packages with built-in ESD protection and is available in blue, green, red, and yellow. The devices withstand the severe levels of mechanical, environmental, and electrical stress typical in field-deployed military-electronics systems. The devices sell for $14.56 each (1000).

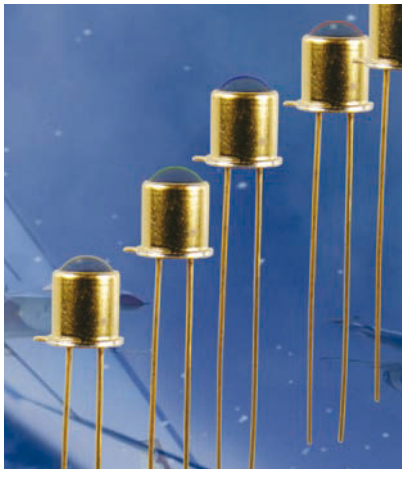

Optek Technology, www.optekinc.com

LED panel light replaces fluorescent tubes

The Pollux LED-panel-light series replaces traditional tube or grid lights, providing evenly distributed cool-white light. The devices consume 50W of power and provide 3400 lumens, targeting use in indoor settings. Other features include a beam angle of 120°, an input voltage of 100 to 240V ac±10% at 50 or 60 Hz, and a 30,000-hour life span. Measuring 600×600×13 mm, the devices have a two-year warranty and sell for $199.

GlacialLight, www.glaciallight.com

EDN ADVERTISER INDEX

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

DISCRETE SEMICONDUCTORS

Gate-driver optocoupler provides 35-kV/µsec CMR

The 3A-output-current FOD31-84 gate-driver optocoupler drives power MOSFET and IGBTs at frequencies as high as 250 kHz. The FOD3184 comprises an aluminum-gallium-arsenide LED optically coupled to a CMOS detector with PMOS and NMOS output-power transistors in the circuit's power stage. The device has a minimum CMR of 35 kV/µsec and an operating range of 15 to 30V, 3A maximum peak output current, and a guaranteed operating-temperature range of –40 to +100°C. The device comes in an eight-pin dual inline housing and is compatible with a 260°C reflow process. It sells for $1.51 (1000).

Fairchild Semiconductor, www.fairchildsemi.com

Go ahead. Make my day!

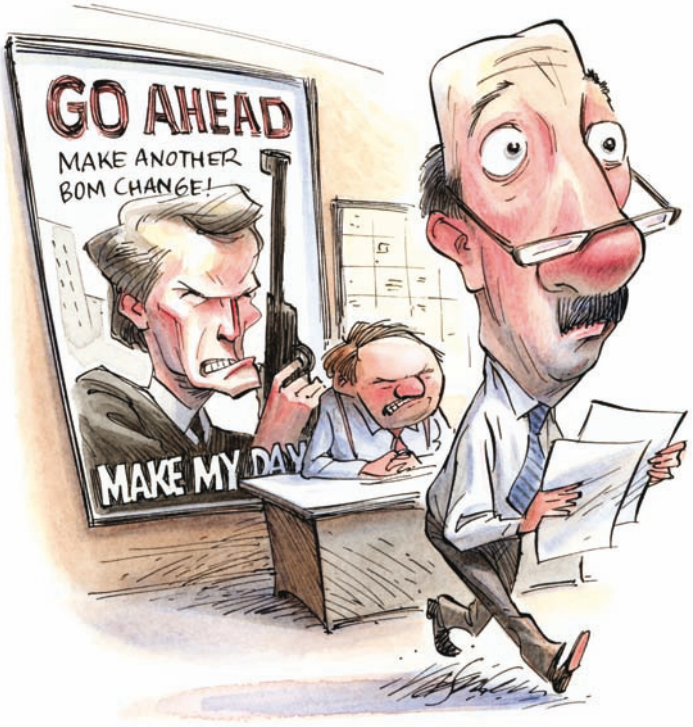

You usually understand that you have done a good design when, after launching the product on the production line, the unit goes on the test fixtures well within the data-sheet parameters and passes the final burn-in test, no BOM (bill-of-materials) changes crop up, and no malfunctions occur after assembly with components from second-source manufacturers. After six to 12 months of running the unit on the production floor with no problems, the infant-mortality phase is over for this product—at least that's how the textbooks describe it.

At my job, we had in regular production a small 50W power supply that had been operating for almost two years with no problems. One morning, some operators suddenly encountered problems with running the power supply on the test fixtures. This unit was my responsibility, so I dug out the design folder and the schematics and went unhappily toward the production floor. On my way, I met the manufacturing director, who had a poster on his office wall of Clint Eastwood, gun in hand, and the quotation "Go ahead! Make another BOM change! Make my day!"

The unit on the test station was a standard power-MOSFET switch-mode power supply. Under full load, the converter was "hiccupping" as if one of the outputs were shorted. Under no load and 10% of full current, the converter ran relatively normally, but the load regulation was not good. I checked the output capacitors, the output rectifiers, the resistive loads of the test fixture, and the PCB (printed-circuit-board) layout for possible track shorts. Everything looked normal. I hooked the scope probe onto the auxiliary winding of the transformer, which was supplying the IC controller, and found overloading: The voltage was collapsing with a full output load.

I again checked the board for a short on the secondary. This explanation was the only one I could think of for the problem. There were ±12V and 5V outputs, and I figured that they were OK. I wasn't so sure about the third output voltage, 20V. The output rectifier comprising two ultrafast parallel diodes in a TO-220 case was behaving as a diode, but I couldn't visually check the markings on the package. I had laid out the PCB, and, in my desire to minimize the current loops and to pass the EMC (electromagnetic-compatibility)-qualification test, I had placed the components so close to each other that I couldn't see the front of the rectifier case. Visual inspection was impossible for every component.

I unsoldered the rectifier to see whether the right diode had been used. When I pulled out the part, I saw another power MOSFET in the same type of TO-220 case as the diodes. As a result, the parasitic body diode of the MOSFET was acting exactly as a TO-220-encased output rectifier in a flyback converter! For that reason, I couldn't detect any problems with this output "rectifier"; I was simply measuring the body diode of a wrongly soldered power MOSFET. OK, but why was the parasitic body diode still working at a switching frequency of approximately 100 kHz? Parasitic diodes should not have high switching speeds.

A look at the production folder revealed that we had used a second-source manufacturer for this MOSFET. I went online to see the MOSFET's data specifications and found that the manufacturer had replaced this MOSFET with a FET with the same electrical parameters but with an enhanced high-speed drain-to-source diode. As a result, the output "rectifier" was able to conduct at the high switching frequency.

I made a third revision of the PCB, this time with a cutout on the board so that the output rectifier's case was visible through it. A simple visual control ensured installation of the right part. Fortunately for me, it was only a revision and not a BOM change. EDN

Todor Arsenov is an electronics engineer in Toronto, ON, Canada.

www.edn.com/tales

DANIEL VASCONCELLOS

DC-DC Converters
2V to 10,000 VDC Output

New AV Series
Isolated
100 to 5000 VDC Output
Surface Mount/Thur-Hole

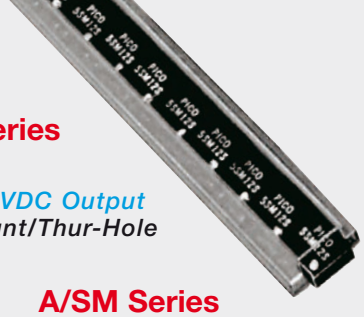

New Dual Isolated Outputs
to 250 VDC Isolated Output
Thru-Hole and Surface Mount

VV Series
Up to 10,000 VDC Output
Up to 10 Watts

A/SM Series
Surface Mount/Thur-Hole
Isolated
3.3 to 1000 VDC Output

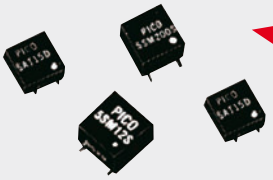

AVR Series
Regulated to IKV
Low Noise

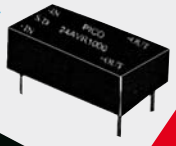

M/MV Series
Military Components
Military Environmental
Testing Available
3.3 to 500 VDC

HVP Series
Programmable to
6000 VDC Output

P/HP/XP Series
HIGH POWER
3.3 to 350 VDC
Output
Isolated to
300 Watts

LV/HV Series
36 to 170 Inputs
Terminal Strips
PC Board Mount
3.3 to 48 VDC Output

OR/IR/JR/KR Series
Wide Input Range
8 to 60 VDC
2 to 100 VDC Output
Regulated

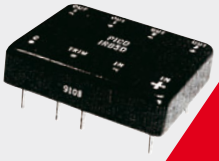

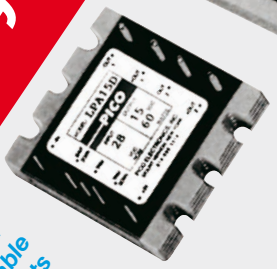

LF/LM/FM Series
-40° and -55° to +85°C
Operating Temperature

3.3 to 350 VDC outputs
in 1/2 Brick and Full Brick

LP Series
Wide Input Range
Watts to 75 Watts
2 to 100 VDC Output Full Brick

2500 Standard Models
Regulated/Isolated
Wide Input Available
1 Watt to 300 Watts

- **COTS**
- **MILITARY**
- **INDUSTRIAL**

Also AC-DC single and/or 3 Phase. Power Factor Corrected. 3 Watts to 2000 Watts Models.

PICO ***Electronics, Inc.*** 143 Sparks Ave. Pelham, NY 10803

• ***Call Toll Free*** 800-431-1064 • ***Fax:*** 914-738-8225

ISO CERTIFIED 9001:2000

SURFACE MOUNT • MILITARY UPGRADES • 2500 Standard Models • Custom Models

www.picoelectronics.com

E-mail: info@picoelectronics.com • ***Send for free 180 pg PICO Catalog***

RF-to-Bits

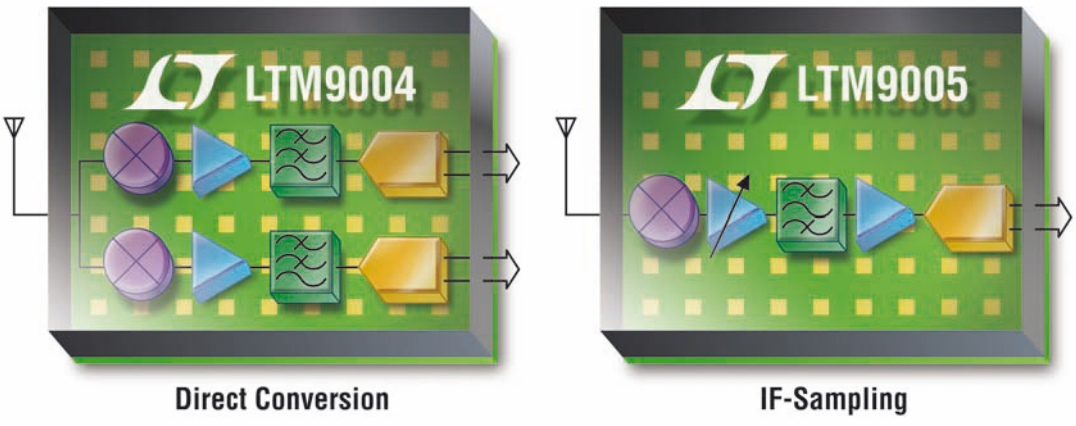

Complete µModule® Receivers Simplify Radio Design

The LTM®9004 (Direct Conversion) and LTM9005 (IF-Sampling) µModule receivers offer unprecedented integration for smaller, low power designs and dramatically faster time-to-market. Integrating the RF, signal filtering, gain stages and high-speed ADCs, these receivers eliminate time-consuming design, layout and component sourcing associated with high-performance communication systems. With no performance compromise, they enable dense, multichannel macrocells or compact remote radio heads.

Features

- Fully Integrated RF-to-Digital Receivers
- 14-Bit, 125Msps Low Power ADC
- Direct Conversion Architecture (LTM9004)
  - 800MHz to 2.7GHz RF Input Range
  - I/Q Demodulation & Dual ADC
  - 5V & 3V Supplies, 1.8W Total
- IF-Sampling Architecture (LTM9005)
  - 400MHz to 3.8GHz RF Input Range
  - 20MHz SAW Filter, 140MHz IF
  - 3.3V Supply, 1.3W Total
- 22mm x 15mm LGA Package

LTM9004 Demo Board

Info & Free Samples

www.linear.com/9004

1-800-4-LINEAR

Free Wireless Solutions Brochure

www.linear.com/wireless

LINEAR TECHNOLOGY