Reed Business Information (RBI™)

EDN®

VOICE OF THE ENGINEER

FEB 19

Issue 4/2009

www.edn.com

MAKING OSCILLATOR SELECTION CRYSTAL CLEAR

page 28

VERIFYING FPGA DESIGNS: SIMULATE, EMULATE, OR HOPE FOR THE BEST?

Page 22

SUPPRESSING NONSTATIONARY NOISE IN MOBILE HANDSETS

Page 38

PCIe CLOCK DISTRIBUTION IN EMBEDDED SYSTEMS

Page 42

Put the burden of custom SoC and COT solutions on our shoulders.

Benefit from an integrated approach to SoC design

Choosing between a traditional ASIC or COT strategy can be tough. Fortunately, Toshiba makes your choice easier. We successfully implement ASIC, COT and RFCMOS product flows for customers who choose to move away from pure-play foundry models that entail high-risk IP and a high cost of ownership. With Toshiba, you get access to our more than 25 years of experience as an ASIC integrated device manufacturer, the engineering talent of the world's third largest[1] semiconductor company and reliable, proven manufacturing. We invest billions annually to stay at the forefront of process technology—from 90/65/40/32/22 nm. Join with Toshiba to achieve successful custom SoC and COT solutions.

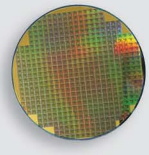

Benefits

- Use Toshiba's proven IP at no cost, including PCI express, USB 2.0, DDR1/2/3, serial I/Os up to 8.5 Gbps
- Reduce costs with our reliable infrastructure, mask-making capabilities, die/package solutions and embedded memories
- Use your custom macros/differentiating IP to minimize competitive threats
- Validate design and silicon and package integrity before design start with Toshiba "Virtual Prototyping"

ChipQuest2 design challenge

Test your skill and win a Noise Cancellation Headset, Wireless Internet Hot Spot Finder, or high-powered flexible LED light (*while supplies last*). Enter at SoCworld.toshiba.com.

TOSHIBA

Leading Innovation »»

[1] Worldwide Semiconductor Vendors by Revenue, Gartner.

YOUR SOURCE

HUNDREDS OF THOUSANDS OF PARTS
HUNDREDS OF SUPPLIERS
WWW.DIGIKEY.COM

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

Digi-Key® CORPORATION

17 YEARS IN A ROW!*

#1 Breadth of Product Line
#1 Availability of Product
#1 On-Time Delivery
#1 Responsiveness
#1 Overall Performance

Digi-Key Corporation purchases all product directly from its original manufacturer.

*Source: *EE Times* Distribution Study, August 2008

Quality Electronic Components, Superior Service

www.digikey.com
1.800.344.4539

701 Brooks Ave. South • Thief River Falls, MN 56701 • USA

The Choice for Test Management

"At Averna, we used NI TestStand, NI LabVIEW, and PXI to create an RF testing solution including AM/FM, RDS, SIRIUS, XM, DAB, and IBOC, reducing total system costs by more than 50 percent."

– Richard Maltais, Vice President, Sales and Marketing, Averna

NI TestStand, the Industry Standard in Test Management Software

Chosen by the top 10 global manufacturing companies, NI TestStand is a ready-to-run environment for developing, managing, and automating test sequences. NI TestStand works with any programming language and facilitates smoother integration with enterprise systems, getting you to first test faster.

NEW *NI TestStand 4.1*

- *Ready-to-run environment increases efficiencies and test throughput*
- *Easy sequence flow, reporting, database logging, and connectivity to enterprise systems*
- *Built-in parallel and batch execution and synchronization*

>> View demos and download free software at **ni.com/teststand** **800 891 8841**

NATIONAL INSTRUMENTS

EDN contents

2.19.09

DESIGNIDEAS

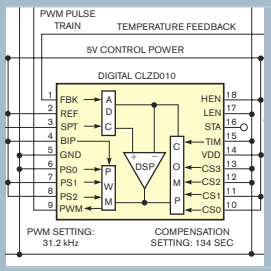

Experience a Faster Way to Search for Components!

Jameco Electronics' new catalog and enhanced Jameco.com website are two tools that are designed to work together to give electronic professionals faster access to the hottest components in the industry.

Color coded references throughout the catalog assist you in analyzing a wide range of brand choices (from franchise sourced to house brands to factory overruns) offering you more pricing options than you'll see anywhere else. Web codes allow you to quickly jump from catalog to website to view additional specifications, application notes, photos and more products. You'll find exactly what you're looking for without having to wade through hundreds of thousands of products on a complicated website or wielding a twenty pound catalog.

With a flip of the page or a click of the mouse, you have all the tools you need at your fingertips. Reach for Jameco first and order your catalog today.

Order your web-enabled catalog today!

JAMECO® ELECTRONICS

Call 1-800-831-4242 or Visit www.Jameco.com/Speed for the ride of your life!

contents 2.19.09

pulse

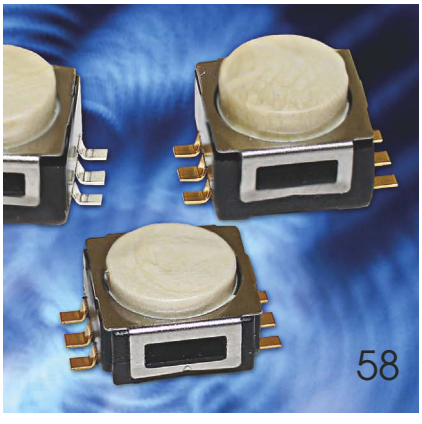

DEPARTMENTS & COLUMNS

PRODUCT ROUNDUP

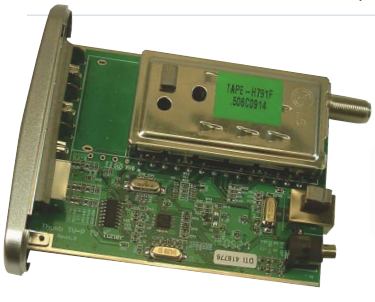

EDN® (ISSN#0012-7515), (GST#123397457) is published biweekly, 24 times per year, by Reed Business Information, 8878 Barrons Blvd, Highlands Ranch, CO 80129-2345. Reed Business Information, a division of Reed Elsevier Inc, is located at 360 Park Avenue South, New York, NY 10010. Tad Smith, Chief Executive Officer; Mark Finkelstein, President, Boston Division. Periodicals postage paid at Littleton, CO 80126 and additional mailing offices. Circulation records are maintained at Reed Business Information, 8878 S Barrons Blvd, Highlands Ranch, CO 80129-2345. Telephone (303) 470-4445. POSTMASTER: Send address changes to EDN®, PO Box 7500, Highlands Ranch, CO 80163-7500. EDN® copyright 2009 by Reed Elsevier Inc. Rates for nonqualified subscriptions, including all issues: US, $179.99 one year; Canada, $229.99 one year (includes 7% GST, GST#123397457); Mexico, $229.99 one year; air expedited, $399.99 one year. Except for special issues where price changes are indicated, single copies are available for $10 US and $15 foreign. Publications Agreement No. 40685520. Return undeliverable Canadian addresses to: RCS International, Box 697 STN A, Windsor Ontario N9A 6N4. E-mail: Subsmail@ReedBusiness.com. Please address all subscription mail to EDN®, 8878 S Barrons Blvd, Highlands Ranch, CO 80129-2345. EDN® is a registered trademark of Reed Elsevier Properties Inc, used under license. A Reed Business Information Publication/Volume 54, Number 4 (Printed in USA).

The Newest Microcontrollers

ARM®
Cortex™
8051
PSoC®
Encore!®
Z8®
Flexis™
AVR®
MSP430
HCS08
Coldfire®
PIC®
Z80®
ST7
HC12
ST6
HC08
HCS12
ST9

New Products from:

MICROCHIP
PIC32 Microcontrollers
www.mouser.com/microchip/a

NXP founded by Philips
LPC2478/2470 ARM7™ Microcontrollers
www.mouser.com/nxp/a

Texas Instruments
Authorized Distributor
MSP430 Ultra-Low Power Microcontrollers
www.mouser.com/ti-msp430

The ONLY New Catalog Every 90 Days

Experience Mouser's time-to-market advantage with no minimums and same-day shipping of the newest products from more than 366 leading suppliers.

MOUSER® ELECTRONICS
a tti company

The Newest Products For Your Newest Designs

www.mouser.com
Over A Million Products Online

(800) 346-6873

EDN at edn.com

EXPANDED ENGINEERING COVERAGE PLUS DAILY NEWS, BLOGS, VIDEO, AND MORE.

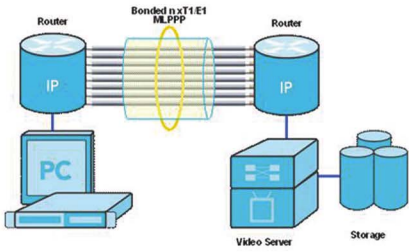

ONLINE ONLY

Check out these Web-exclusive articles:

EoPDH as an alternative for delivering rich-media applications
This study compares the performance of media-delivery techniques to show how EoPDH (Ethernet-over-plesiochronous-digital-hierarchy) technology can deliver reliable, resilient transport services for streaming-video applications.
→www.edn.com/article/CA6624633

Jasper, OneSpin seek broader audience for formal-verification tools
Not just for verification experts anymore: New tools try to attract designers at early stages of the design process.
→www.edn.com/article/CA6631000

Forging new ground using existing COMs concepts
Although they are highly popular embedded-computing components, COMs (computers on modules) lacked the continuity needed to make these boards appeal to the various industries for which they were designed.
→www.edn.com/article/CA6631704

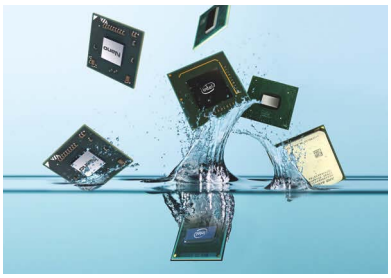

READERS' CHOICE

A selection of recent articles receiving high traffic on www.edn.com.

Evaluating value-oriented x86 CPUs: the price of falling prices
Economical microprocessors may enable you to easily hit your next design's bill-of-materials cost target, but will they also allow you to accomplish your performance and power-consumption objectives?
→www.edn.com/article/CA6625435

Outsourcing in a troubled economy
→www.edn.com/article/CA6628835

A bad-capacitor story ends happily
→www.edn.com/article/CA6625453

MIT researchers claim 50% increase in solar-cell output
→www.edn.com/article/CA6625447

Supply chains and the "weakest-link" problem
→www.edn.com/20090219toca

Intel invests $23 million in three Indian companies
→www.edn.com/article/CA6629515

INNOVATION VOTE NOW!

Tune in to www.edn.com/innovation to see the innovative engineers and products our editors selected as finalists in this year's *EDN* Innovation Awards. After perusing the candidates, help us honor electronic innovation by using the easy electronic ballot to make your voice heard.

FEED YOUR NEED

Belly up to the *EDN* Product fEEd, our ongoing buffet of new components, ICs, tools, and more.
→www.edn.com/productfeed

BY RICK NELSON, EDITOR-IN-CHIEF

Technical and political suggestions for Detroit

Do the Detroit automakers deserve a bailout, or should they be left to go the way of the dinosaurs (**Reference 1**)? Perhaps there are other options. In a recent article, Sebastian Thrun, a professor of computer science at Stanford University, and Anthony Levandowski, now a product manager at Google, discuss "four technologies that could be put on the road in the next several years." They describe these technologies, some of which they say they invest in, as win-win situations, which alone or in combination can increase energy efficiency, improve safety, cut pollution, and make commuting more convenient (**Reference 2**).

Those technologies, they write, include Wi-Fi-like dedicated short-range communication, which enables cars to form ad hoc networks with one another and with roadway infrastructure; automatically controlled X-by-wire capabilities that would allow robotic cars to drive in tight formation under computer control (**Reference 3**); solar/electric-hybrid technology, in which solar panels on one garage provide power for 10,000 miles of driving per year; and smartphone-like online-communications capabilities that will help drivers find parking spaces, make restaurant reservations, and purchase songs they hear on the car radio.

On the eve of Obama's inauguration, auto-industry journalist Steve Parker detailed his own ideas for saving Detroit (**Reference 4**). Parker's recommendations are a political wish list of what he would like the incoming administration to do. He suggests that we begin an Apollo-like program to deploy "the first national infrastructure for hydrogen- and natural-gas fuels" and build fuel-cell-based vehicles that can run on these fuels. He also suggests raising federal gasoline taxes and instituting tax rebates and other incentives for purchasers of high-mileage, clean vehicles. He'd like Detroit to become involved in the development of modern high-speed-rail and light-rail systems. And he'd like government to mandate that automakers achieve CAFE (corporate-average-fuel-economy) levels of 35 mpg—currently slated for 2020—by 2015.

Both groups of suggestions have some good ideas but no silver bullets. As for the technical solutions, technologies for robotic cars to drive in tight formation or for the deployment and use of a hydrogen infrastructure seem impractical in the short term; I wouldn't advocate extensive public financing for R&D in either area. CAFE standards have had the perverse effect of encouraging the Big Three to heavily market profitable SUVs and trucks to offset losses incurred in selling fuel-efficient autos. Tax rebates on fuel-efficient cars coupled with higher gas taxes or sales surtaxes on low-mileage cars might be a better approach, although a surtax on low-mileage vehicles would unfairly penalize individuals and families that need trucks or large cars for occasional use. (Commercial vehicles could be exempt from this surtax.) As for involving Detroit in rail projects, well, OK, but why would the Big Three be more successful in making trains than in making cars?

Solar/electric-hybrid technology shows promise. If the technology is not ready for home installation, perhaps employers and commercial establishments could offer solar-enabled parking as a perk for employees and customers. Ad hoc networking schemes and enhanced online-communications capabilities also hold promise—especially if public investment in roads and bridges could help put the electronic infrastructure in place to provide real-time information on traffic and roadway conditions and on nearby services and points of interest.EDN

REFERENCES

1 Nelson, Rick, "Evolution and Innovation, *EDN*, Jan 22, 2009, pg 8, www.edn.com/article/CA6629470.

2 Thrun, Sebastian, and Anthony Levandowski, "Four Ways for Detroit to Save Itself," *The New York Times*, Jan 3, 2009, www.edn.com/090219eda.

3 Nelson, Rick, "Balancing act," *Test & Measurement World*, March 2005, pg 30, www.tmworld.com/article/CA506541.

4 Parker, Steve, "Five Things Mr. Obama Must Say About Detroit–and Fast," *The Huffington Post*, Jan 19, 2009, www.edn.com/090219edb.

Contact me at rnelson@reedbusiness.com.

Go to www.edn.com/090219ed and click on Feedback Loop to post a comment on this column.

www.edn.com/edncomment

29

More Q. Less Cu

Copper
63.546

These tiny new air core inductors have the highest Q and current handling in the smallest footprint.

Coilcraft's new SQ air core inductors have unmatched Q factors: most are above 200 in the 1-2 GHz range! That's 3 times higher than comparably sized 0805 chip coils.

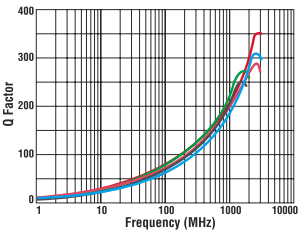

Q factors are 3X higher than standard chip inductors

And with their extremely low DCR, they can handle 4 to 8 times more current: up to 4.4 Arms.

SQ air core inductors are perfect for your LC filter and RF impedance matching applications. They come in 15 values ranging from 6 to 27.3 nH, all with 5% tolerance.

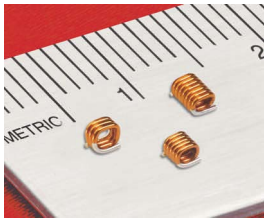

These coils are significantly smaller than existing air core inductors. We reduced the footprint by using close-wound construction and keeping the leads close to the body. The square shape cuts the height to as low as 1.5 mm and creates flat top and bottom surfaces for easy automated handling and stable mounting.

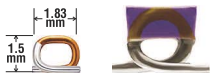

The square shape and narrow footprint reduce board space by 60-75% over conventional air core inductors.

See how the ultra-high Q and current handling of Coilcraft's new SQ air core inductors can maximize the performance of your next design. For complete specifications and free evaluation samples, visit **www.coilcraft.com/sq**

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

RoHS COMPLIANT

Coilcraft®

www.coilcraft.com 800/322-2645

More measurements. More data. Less cost.

With up to 560 channels of data logging at a maximum scan rate of 1000 channels/sec, an internal 6 ½ digit DMM with 22-bit resolution, and 0.03% accuracy, it's no mystery why Agilent's data acquisition systems are among the top selling in the world. The 34970A and 34980A are both at home with either data acquisition or functional test, with a broad selection of plug-in modules.

Agilent 34970A switch/measure unit

Agilent 34980A switch/measure unit

Now with Agilent Data Logger Pro software, you'll be able to collect data, check limits, perform extensive branching, and view and store the data without authoring the programming yourself. It's an ideal expansion at a significant savings.

www.agilent.com/find/daqswitch

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

Agilent Technologies

EDN

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

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

ASSOCIATE PUBLISHER, *EDN* WORLDWIDE
Judy Hayes, 1-408-421-3799; judy.hayes@reedbusiness.com

VICE PRESIDENT, EDITORIAL DIRECTOR
Karen Field, 1-781-734-8188; kfield@reedbusiness.com

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

EXECUTIVE EDITOR
Ron Wilson, 1-408-345-4427; ronald.wilson@reedbusiness.com

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

EDITOR-IN-CHIEF, EDN.COM
Matthew Miller
1-781-734-8446; fax: 1-303-265-3017; mdmiller@reedbusiness.com

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

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

EMBEDDED SYSTEMS
Warren Webb, Technical Editor
1-858-513-3713; fax: 1-858-486-3646; wwebb@edn.com

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

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

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

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

DESIGN IDEAS EDITOR
Martin Rowe
edndesignideas@reedbusiness.com

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

ASSOCIATE EDITOR
Maura Hadro Butler, 1-617-276-6523; mbutler@reedbusiness.com

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

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

COLUMNISTS
Howard Johnson, PhD; Bonnie Baker; Joshua Israelsohn; Pallab Chatterjee

PRODUCTION
Dorothy Buchholz, Group Production Director
1-781-734-8329
Kelly Jones, Production Manager
1-781-734-8328; fax: 1-303-265-3164

EDN EUROPE
Graham Prophet, Editor, Reed Publishing
The Quadrant, Sutton, Surrey SM2 5AS
+44 118 935 1650; fax: +44 118 935 1670; gprophet@reedbusiness.com

EDN ASIA
Susie Newham, Managing Director
susie.newham@rbi-asia.com
Kirtimaya Varma, Editor-in-Chief
kirti.varma@rbi-asia.com

EDN CHINA
William Zhang, Publisher and Editorial Director
wmzhang@idg-rbi.com.cn
John Mu, Executive Editor
johnmu@idg-rbi.com.cn

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

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

Reed Electronics Group

EDITED BY FRAN GRANVILLE

pulse

INNOVATIONS & INNOVATORS

Controller targets mobile robots

Robot technology is appealing to embedded-system designers not only because of the mechanical-control aspect but also because of the potential it has to shield humans from dangerous and even life-threatening situations. Targeting the high end of this technology, Roboteq recently introduced an intelligent, dual-channel dc-motor controller that can directly drive as much as 120A on each channel at voltages as high as 60V. The AX2860 should appeal to designers of mobile-robotic vehicles, including AGVs (automatic guided vehicles), underwater ROVs (remotely operated vehicles), and mobile robots for exploration, hazardous-material handling, and military and surveillance applications.

The intelligent, dual-channel AX2860 dc-motor controller can directly drive as much as 120A on each channel at voltages as high as 60V.

The controller accepts commands from standard radios for simple remote-controlled-robot applications, analog joysticks, or RS-232 interfaces. Using the serial port, you can use the AX2860 to design autonomous or semi-autonomous robots by connecting it to single-board computers, wireless modems, or Wi-Fi adapters. You can operate the controller's two channels independently or combine them to set the direction and rotation of a vehicle by coordinating the motors on each side for tanklike steering. You can operate the motors in open- or closed-loop speed mode. The AX2860 includes inputs for two quadrature encoders operating as fast as 250 kHz and four limit switches for precise speed and traveled-distance measurement. The AX2860 is available now at $720 in single quantities and comes with cable and PC-based configuration software.**—by Warren Webb**

▷**Roboteq**, www.roboteq.com.

FEEDBACK LOOP

"If all the dinosaurs had gone extinct, we would not have birds. A lesson for world corporations: Most of the largest species in history have outgrown their environment. Smaller species are much more adaptable and quicker to adopt changes."

—Engineer and *EDN* reader Charles Warner, in *EDN*'s Feedback Loop, at www.edn.com/article/CA6629470. Add your comments.

Companion processor supports video editing on the phone

Semiconductor start-up Movidia has announced the MA1110 multimedia processor for mobile phones. The MA1110 operates as an attached processor to the baseband or to another applications processor in medium- to high-end phone designs to enable high-performance in-phone-video postproduction in real time on low power budgets.

The concept supports social-networking applications on phones in which users capture, edit, and publish video without using a PC. The chip includes interfaces to dual cameras with resolution as high as 12M pixels, the host processors, displays, and external memory when necessary; 1.4 Mbytes of memory is on-chip. The chip has a host RISC processor and extensive graphics-manipulation IP (intellectual property) hard-coded into silicon. It performs complex video-editing tasks, such as real-time image stabilization, super-resolution zoom, slow motion, and color matching.

The MA1110 also supports the high resolutions of all standard audio and video codecs and APIs (application-programming interfaces) to ensure full compatibility with a range of content, and it can perform transcoding between the standards.

Movidia has first silicon now and expects to offer production-quality silicon for sampling by midyear. The chip consumes approximately 200 mW when performing intensive graphics operations, the company predicts.**—by Graham Prophet**

▷**Movidia**, www.movidia.com.

Electric-car development spurs investment in US battery manufacturing

LG Chem will provide the lithium-ion cells for the battery pack (foreground) of the Chevy Volt PHEV (rear), and the pack itself will be assembled in the United States.

At last month's Detroit Auto Show, General Motors announced that LG Chem, a Korean lithium-ion-battery-cell manufacturer, will be providing the cells for the Chevy Volt PHEV (plug-in-hybrid-electric vehicle). GM made much of the fact that it will keep assembly of the battery *pack* in the United States. However, the guts of a battery pack—in weight, cost, and intellectual property—are in the cells, of which the United States has apparently fallen off the map in being a volume producer.

Contrast these start-up ventures with BYD's F3DM PHEV, which it unveiled a few weeks ago in China and planned to begin selling in the United States in 2011, priced at around $20,000. BYD got its start making after-market lithium-ion batteries for cell phones, and got into the EV-battery market on the strength of its founder's belief that the future of transportation lies in EVs. The company has more than 20 years' experience developing and manufacturing lithium-ion batteries.

Innovation over the long haul in a technology often goes hand in hand with volume production. The United States may be unable to retake and maintain a lead in battery development unless it also keeps the manufacturing here.

But maybe there's hope: In a briefing to a US Senate committee, Kleiner Perkins' venture capitalist John Doerr claimed that a US company "somewhere in the Midwest" is manufacturing a stable, durable lithium-ion battery with higher effective-storage capacity. "The result is [that] electric vehicles will be able to travel twice as far and, eventually, three times as far, to over 100 miles before recharging," he says. Kleiner Perkins is investing in the company, but would give no details, such as the company's name or what the actual battery specs are.

In addition, researchers at the University of Michigan have spun out Sakti3, a Michigan-based battery start-up. The venture-capital community has invested $2 million in Sakti3, and the state of Michigan has put in $3 million, with $2.4 million in tax credits. Again, no word on when we can expect to see the batteries.

—by Margery Conner

▷**Chevy Volt**, www.chevrolet.com/electriccar.

▷**BYD Auto**, www.byd.com.

02.19.09

INSTRUMENT PERFORMS JITTER-TOLERANCE TESTS

Agilent Technologies' new N4903B J-BERT (bit-error-ratio tester) enables you to accurately characterize and test compliance of next-generation devices that support multigigabit-per-second serial-bus standards, such as PCIe (peripheral-component-interconnect-express) 2.0, USB (universal serial bus) 3, QPI (QuickPath interconnect), Hypertransport 3, and FB-DIMM (fully buffered dual-inline-memory module) 2. According to Agilent, the instrument's high characterization accuracy enhances the robustness of the design under test.

QPI, Hypertransport, and FB-DIMM 2 operate with forwarded clocks that run at half the data rate, confronting design teams with additional test challenges when characterizing receivers under real-world stress conditions.

Prices start at $139,000 for the 7-Gbps version, and $179,000 for the 12.5-Gbps version. For an expanded description of this product, go to www.edn.com/article/CA6633282.

—by Dan Strassberg

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

The B version of J-BERT tests and verifies the operation of both forwarded- and embedded-clocked ICs under a variety of signal impairments at data rates as high as 12.5 Gbps.

DILBERT By Scott Adams

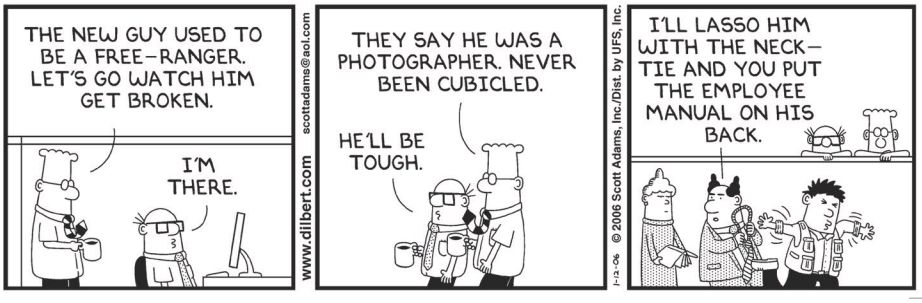

PCIe Jammer provides inline-error injection

Agilent Technologies has introduced the N5323A Jammer inline-error-injection tool for PCIe (peripheral-component-interconnect-express) protocol testing. This unit allows you to shorten test cycles and improve time to market for peripheral and host devices. PCIe has become the interconnect technology of choice in high-performance applications, including servers, storage devices, peripherals, and graphics adapters. Designers of PCIe systems and devices face demands for more reliable systems that can interoperate with the rapidly increasing number of PCIe devices.

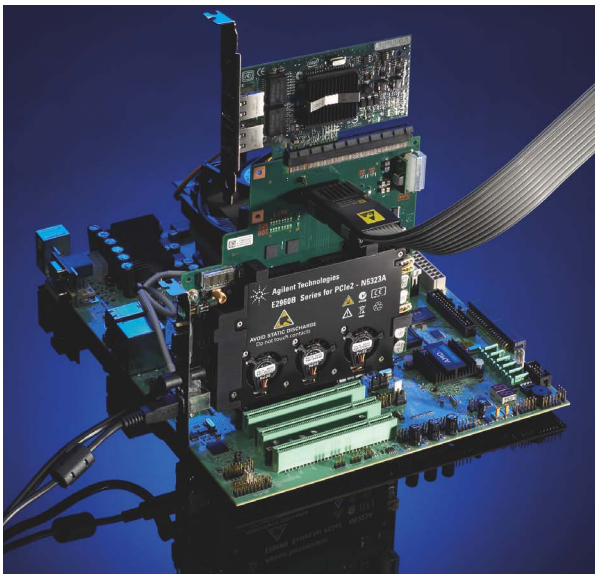

The PCIe Jammer injects errors to perform disruptive tests on live PCIe systems regardless of the operating system or application software.

Agilent designed the N5323A to address these challenges. The half-size PCIe plug-in unit measures 6.6 in. long and 7.1 in. high. It sits transparently between the host and a peripheral adapter and can modify PCIe data transfers in real time to create disruptive scenarios that increase test coverage. The unit performs these disruptive tests in live systems regardless of the operating system or application software.

You can program almost any error-recovery test case, including correctable and uncorrectable PCIe errors. With this approach, you can improve error handling and avoid costly last-minute product changes or recalls. Because of its transparency, setting up the unit is fast and easy. You need write no scripts; a graphical sequencer drives the error injections.

To seamlessly move from testing to debugging, the Jammer also works with the manufacturer's PCIe E2960B protocol analyzer. The Jammer generates errors and can trigger the analyzer to capture and display the responses. This arrangement allows you to view traffic both before and after the error injection to ensure that your application invoked the appropriate recovery mechanism. Versions of the Jammer support PCIe implementations that have one, four or fewer, and eight or fewer lanes, as well as data rates of 2.5G and 5G transfers/sec. Prices start at $27,500.

—by Dan Strassberg

▷**Agilent Technologies**, www.agilent.com/find/pciejammer.

EDN BLOG

PRACTICAL CHIP DESIGN

AMD and Virage Logic try a new model for third-party IP

An innocuous-looking announcement on the wires last month from AMD (www.amd.com) and Virage Logic (www.viragelogic.com) didn't sound like much. But it may be the sound of the curtain coming up on a whole new business model for the third-party IP (intellectual property) business—and none too soon.

Virage and AMD announced that the two companies would collaborate on the development of silicon IP, including some specific interface cores now and unspecified new cores in the future. The cores currently being productized are a Generation 2 PCIe (peripheral-component-interconnect-express) interface, an HDMI (high-definition-multimedia-interface)/DisplayPort interface, and a MIPI (mobile-industry-processor-industry) core. All three are to be available in both 65- and 40-nm versions.

So far, it doesn't sound like a big deal. But … it might be the beginning of a big deal. The graphics division of AMD, the former ATI, developed the three interface cores for internal use in standard bulk CMOS, not for the processor division's more esoteric SOI (silicon-on-insulator) process. And AMD has proved all three not only in silicon but in systems at 40 nm. Virage intends to pick up the cores as is, prepare them for commercialization as for-sale, third-party IP, and then market them.

Virage, one of the leaders in logic-cell libraries, embedded-RAM compilers, and, more recently, interface IP, has supplied the AMD designers for years. Neither the former ATI nor any other part of AMD has a history in the IP business. So, why does the IP vendor now want to license IP from its customer?

The cynic might suggest part of the answer is that AMD needs all the revenue sources it can find right now and that it would commercialize its employee lunchroom if it could figure out a business model. But Brani Buric, Virage's executive vice president of marketing and sales, says that … part of the equation is time to market. AMD engages with its foundries on a new process node at about the same time that Virage begins developing logic libraries—very early. So … AMD will have complex interface cores … up and working much earlier than would normally be feasible.

—by Ron Wilson

▶www.edn.com/practicalchipdesign.

▶For the full post, go to www.edn.com/090219pa.

TAPE-OUT COMES A LOT FASTER
WITH CALIBRE nmDRC.

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

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

Mentor Graphics®

THE EDA TECHNOLOGY LEADER

VOICES
Magma Design Automation's Rajeev Madhavan

The evolving EDA business, the depressed economy, the state of the semiconductor industry, and the future of education in the United States were all on the mind of Rajeev Madhavan, chairman and chief executive officer of Magma Design Automation, when he spoke recently with *EDN*. The following is an excerpt from that interview. For the full discussion, got to www.edn.com/090219pb.

What's your impression of the state of the semiconductor industry?

A October and November were the worst months I have ever seen in my life. People deploying new technology just stopped halfway through their projects.

What's going to happen next?

A Companies are doing some analysis to determine what they are to become. Right now, semiconductor companies are shedding what they think is bad for them, and many are moving away from large digital SOCs [systems on chips] and toward mixed signal.

Why is that?

A It's just a reflection of the fact that just putting a digital SOC together by slapping a lot of IP [intellectual property] into design and coming up with a chip that is sufficiently differentiated from the competition is very difficult.

Could you comment on Magma's recent evolution?

A Magma came out with a physical-synthesis tool in 2005, and we went through a period where we had a lot of litigation expenses. But we came out of that where we actually had a new tool, Talus. By June 2008, we finished the transition in technology, and, as of November, pretty much every customer had migrated over to the Talus platform. But what has helped us through the transition period is our analog technology, which has kicked into very high gear.

What are the challenges in the mixed-signal-EDA market?

A If you look at the history of mixed-signal and analog tools, there has really been only one player, which is Cadence. I participated in the development of the Cadence tools in 1991, and very little has changed. The flow is very manual. There was never an automated way of transferring the knowledge of a great circuit-design engineer to the layout.

What has Magma done to address the analog-design challenges?

A Rather than competing head-on with Cadence, we provide automated capabilities in our Titan tool. It's extremely fast and allows high-caliber circuit designers to capture a design, verify the design, and then in an automated fashion transfer it to a given process.

In fall 2007, you sat down with an *Electronic Business* editor and some other executives to discuss education [see www.edn.com/article/CA6487765]. In that interview, you said that, in the United States, too few engineers are graduating and that math and science have to be elevated in the middle- and high-school levels. Have you seen any improvement over the past year?

A Some of the problems of finding enough engineering graduates to hire have been dampened by the recession, but the reality is that, as far as education in the United States goes, we've got a lot of work to do.

What are the prospects for making those improvements?

A I am very hopeful about what the new administration is going to do. At least [the administration is] cognizant of the problem that exists—that there is very little emphasis on getting kids to go into technology areas. It's time we changed that and got back into leadership in every area of education—from kindergarten to college-engineering programs.

What's the solution—spend more money?

A Money is one aspect. In addition, the metrics of how we measure student performance are somewhat arbitrary. And our teachers don't foster competitiveness. We want to accommodate the sensitivities of our students, but, when our students graduate, they'll be in a competitive world. I'm not an educator, but I think we can introduce that competition much earlier in the system without having a negative impact on the kids.

—Interview conducted and edited by Rick Nelson

Optimum process control means efficiency, reliability, and Analog ICs.

Our newest optimized mixed-signal solutions

Ultralow Noise, 24-Bit Σ-Δ ADC with PGA: AD7190
Expand the operating speed and precision of sensitive measurement instruments with the industry's best combination of data rate and noise-free resolution.

16-Bit User-Programmable Current and Voltage Output DACs: AD5422 Family
Single platform simplifies system design by offering improved performance and a high level of control functionality in a compact package.

36 V, Micropower, Dual, Single-Supply Amplifier: ADA4091-2
Eases front-end design by offering integrated input protection, high accuracy, reduced power, and very small package. Developed through advancements in Analog Devices' proprietary *i*Polar™ process technology.

Precision Instrumentation Amplifier: AD8221
Affordable and easy to use. Industry-leading performance in a small package for high channel density applications.

Current/Voltage Output Drivers: AD5750 and AD5751
Increase efficiency by offering best-in-class performance with the industry's highest level of accuracy and diagnostic capability.

Low Power, High Precision Analog Microcontrollers: ADuC7060 Family
Combine 24-bit Σ-Δ ADCs with 16- and 32-bit ARM7TDMI® MCUs and Flash/EE memory for unequaled multichannel data acquisition and processing performance.

Dependable systems, advanced manufacturing, secure environments: Analog Devices

Instrumentation and process control equipment manufacturers are helping their customers to achieve increased line efficiency, productivity, and safety by leveraging advancements in analog technology. For more than 40 years, designers of industrial equipment have relied on Analog Devices' innovations to distinguish their products and to control and protect their customers' valuable equipment assets. And in today's cost- and safety-sensitive industrial world the speed, precision, and integration of ADI process control ICs enable whole new levels of capabilities—and possibilities.

Experience what Analog Devices can do for your designs. Learn more about our newest application-tuned ICs and support tools by visiting *www.analog.com/industrial-ad*.

analog is everywhere.™

www.analog.com/industrial-ad

ANALOG DEVICES

BY BONNIE BAKER

Throw those bits away

A high-quality load cell may have a 2-mV/V output-transfer function, meaning that for each volt of excitation you get ±2 mV of the full-scale output signal. With an excitation of 4.096V and a full sensor deflection, the maximum output is ±8.192 mV. In a 12-bit application, half of full-scale might represent 0 to 250 lbs for a bathroom scale. If you want 0.25-lb resolution, you need 1000 points of measurement output. To look at something that is 1/1000th of the full-scale range, you must distinguish a change of 8.192 µV of the sensor output. You can achieve this distinction by keeping the peak-to-peak sensor noise less than 8.192 µV for 99.999% of the time, using a crest factor of 4.4 (**Reference 1**). With this definition, the least-significant bit at the sensor is 8.192 µV, or 931 nV rms.

The load-cell bridge has an excitation voltage of 4.096V (**Figure 1**). The INA326 instrumentation amplifier follows the load cell with a gain of 250V/V. The system's full-scale voltage, 250×±8.192 mV, produces a ±2.048V full-scale signal. The 12-bit ADS7822 digitizes the analog signal.

This 12-bit converter system must have an analog filter. The lowpass OPA333 (www.ti.com/opa333) analog filter's primary function is to remove the high-frequency signal components at the input of the ADC (**Reference 2**). Because the load cell in the circuit operates near dc, you limit the bandwidth to 10 Hz. The components in **Figure 1** cost less than $6.

Now, look at load-cell measurement with a 24-bit system. You can simply put the load-cell signal through a first-order lowpass filter and into the delta-sigma ADC (**Figure 2**). The first-order filter in this circuit eliminates high-frequency noise around the converter's sampling frequency (**Reference 3**). The sensor provides the resistor for the filter RC pair.

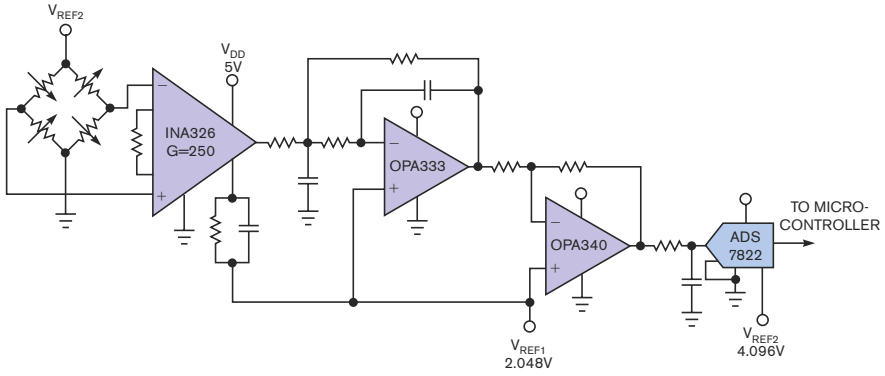

Figure 1 A 12-bit load-cell system achieves accuracy to 0.25 lb.

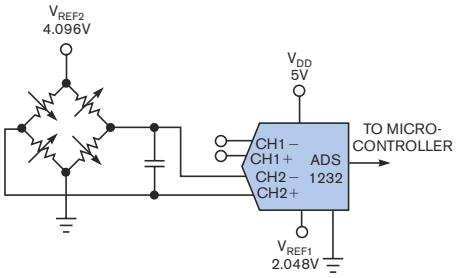

Figure 2 A 24-bit load-cell system has greater-than-0.25-lb accuracy.

Look at the errors for the 24-bit delta-sigma system in **Figure 2**. The ADS1232 (www.ti.com/ads1232) produces 3.7 µV p-p of noise, with a crest factor of 4.4. This figure is much lower than the sensor's least-significant bit. Additionally, the full-scale range of the converter is 4.096V, whereas the sensor's full-scale output range is ±8.192 mV. As you can guess, you will "throw away" most of the output bits of the delta-sigma converter. The components in **Figure 2** cost less than $4.

You may find that the 12-bit converter system ends up costing you more money, real estate, and headaches than the alternative 24-bit system. EDN

REFERENCES

1 Baker, Bonnie, "RMS and peak-to-peak noise trade-off," *EDN*, May 15, 2008, pg 26, www.edn.com/article/CA6558465.

2 Baker, Bonnie, "What's in your SAR-ADC application?" *EDN*, Dec 15, 2008, pg 22, www.edn.com/article/CA6622866.

3 Baker, Bonnie, "Analog filter eases delta-sigma-converter design," *EDN*, June 12, 2008, pg 24, www.edn.com/article/CA6566532.

Bonnie Baker is a senior applications engineer at Texas Instruments and author of A Baker's Dozen: Real Analog Solutions for Digital Designers. *You can reach her at bonnie@ti.com.*

Why is Agilent the fastest growing oscilloscope manufacturer?

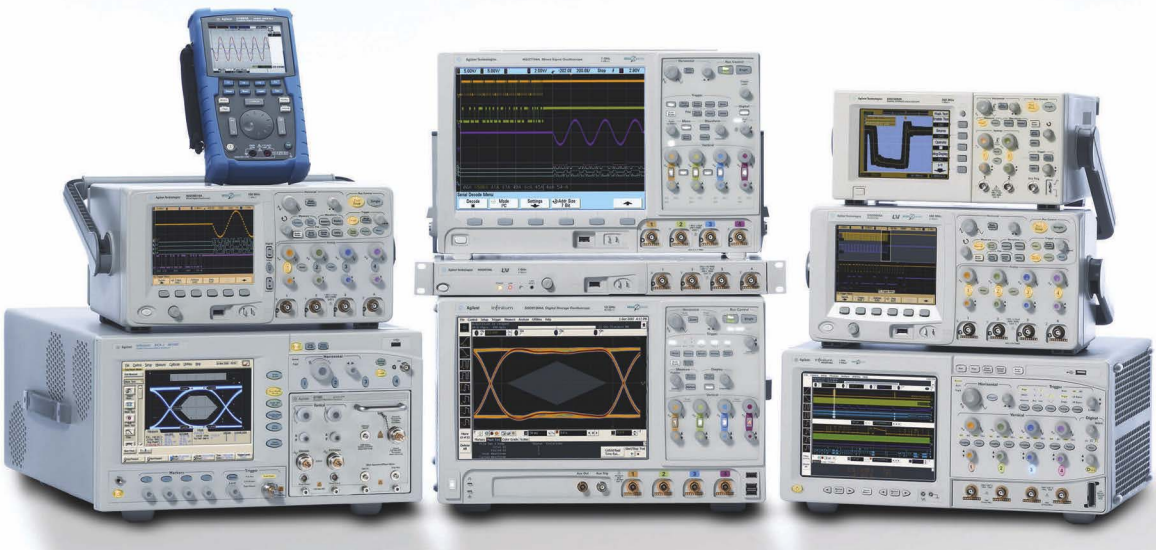

Because we listen to you.

To build our scopes, Agilent carefully examines the challenges you face. Then we deliver products that solve your problems in imaginative ways. Like the multi-chip module that enables Infiniium's industry-leading signal integrity. And the ASIC that underlies InfiniiVision's patented Mega*Zoom* deep memory giving you the industry's best signal visibility. You'll find innovations like these in each of our scopes — that's why more and more engineers are choosing Agilent over other scope brands.*

Agilent 20 MHz to >90 GHz real-time and sampling scopes

- Handhelds, portables, rack-mounts and benchtops
- 40+ application-specific software packages for exceptional insight
- Innovations that satisfy your toughest demands

See why more and more engineers choose Agilent.
Download our catalog
www.agilent.com/find/scopecatalog

Elektra 08 European Electronics Industry Awards WINNER

2008 design news GOLDEN MOUSETRAP Finalist

Best of Test 2008 Test & Measurement World Honorable Mention

 *Prime Data August 2008 CAGR T&M Report

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

Agilent Technologies

PRY FURTHER AT EDN.COM

Go to www.edn.com/090219pry for an expanded version of this article.

A tale of two digital-video converters

SOFTWARE VERSUS HARDWARE COMPRESSION DEFINES THE DIFFERENCE.

These digital-video converters offer similar functions; the less expensive unit even offers a TV tuner absent from its pricier peer. Plextor's ConvertX PX-TV100U cost $8.99 after rebate, $38.99 before rebate, and $99 original MSRP (manufacturer's suggested retail price); the PX-M402U was $29.99 after rebate, $99.99 before rebate, and $159 original MSRP. The big difference? A software- versus hardware-compression emphasis.

Plextor's PX-TV100U relies on the CPU horsepower of a USB2-tethered PC to handle the hefty audio-plus-video lossy-compression duties. The PX-M402U harnesses its integrated WIS Technologies G07007SB streaming media encoder for these meaty tasks.

The PX-M402U design offers flexibility in both primary PCB (printed-circuit-board) unpopulated sites and sufficient enclosure space for a supplemental PCB-based TV tuner that would transform the unit into a PX-TV402U. Similarly, although the PX-TV100U includes an LG Electronics "can" tuner, Plextor also sold the PX-AV100U, a tunerless variant of this design.

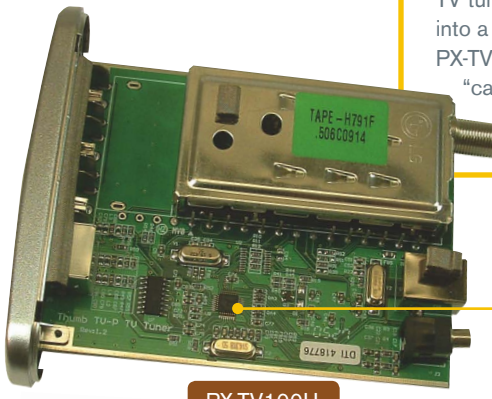

PX-TV100U

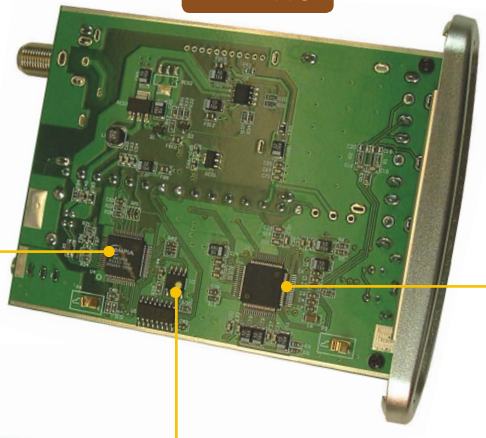

Raw audio and video conveyance to a PC is the responsibility of the PV-TV100U's Empia Technology EM2860 USB2 media-capture processor. The USB2 transport of compressed multimedia data from the PX-M402U falls to Cypress Semiconductor's CY7C68013 embedded microcontroller.

Both devices need to translate analog multimedia information into digital equivalents. With the PX-TV100U, audio and video conversion are the domains of Micronas' MSP 3425G sound processor and Texas Instruments' 5150AM1 video decoder, respectively. The PX-M402U, conversely, relies on AKM Semiconductor's 5355 dual-channel, 16-bit audio ADC and NXP Semiconductor's SAA7115 video decoder. Whereas the SAA7115 comprehends NTSC (National Television System Committee), PAL (phase-alternating-line), and SECAM (Séquentiel Couleur Avec Mémoire) video standards, the video decoder in the PX-M401U (the PX-M402U's predecessor) supported only NTSC.

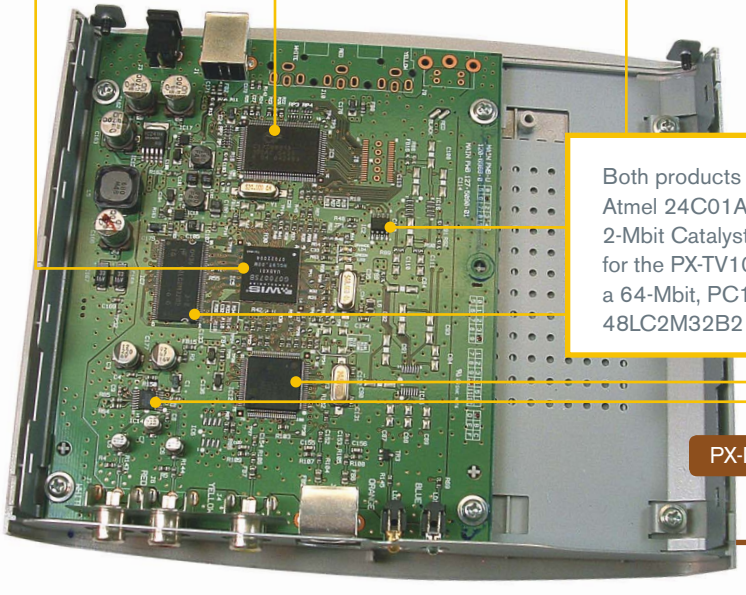

Both products include serial EEPROMs: a 1-Mbit Atmel 24C01A device in the PX-M402U and a 2-Mbit Catalyst Semiconductor 24WC02 memory for the PX-TV100U. The PX-M402U also embeds a 64-Mbit, PC100 SDRAM, Micron Technology's 48LC2M32B2 with a 32-bit interface.

PX-M402U

POWERED WITH ELECTRICITY, GAS,
AND AUTOMATICALLY-GENERATED CODE.

THAT'S MODEL-BASED DESIGN.

To create a two-mode hybrid powertrain, engineers at GM used models to continuously verify their design, test prototypes, and automatically generate the embedded code.
The result: a breakthrough HEV, delivered on time.
To learn more, visit mathworks.com/mbd

MATLAB® & SIMULINK®

The MathWorks
Accelerating the pace of engineering and science

Once upon a time, you verified a logic design for an FPGA by compiling it, loading it, and pushing the reset button on your evaluation board. But, as FPGAs have become larger, this "blow-and-go" verification style, as Xilinx's director of software-product marketing, Hitesh Patel, terms it, has become counterproductive. The odds of creating a multimillion-gate design so close to perfection that you could debug it from the pins on the package are vanishingly small. So, design teams have begun to employ software-based simulation of the design, much as ASIC teams have done for years.

VERIFYING FPGA DESIGNS: SIMULATE, EMULATE, OR HOPE FOR THE BEST?

SIMULATION IS A FACT OF LIFE FOR MANY FPGA USERS TODAY. BUT WHEN IS IT TIME TO STOP SIMULATING AND JUST DROP THE DESIGN INTO THE CHIP?

BY RON WILSON • EXECUTIVE EDITOR

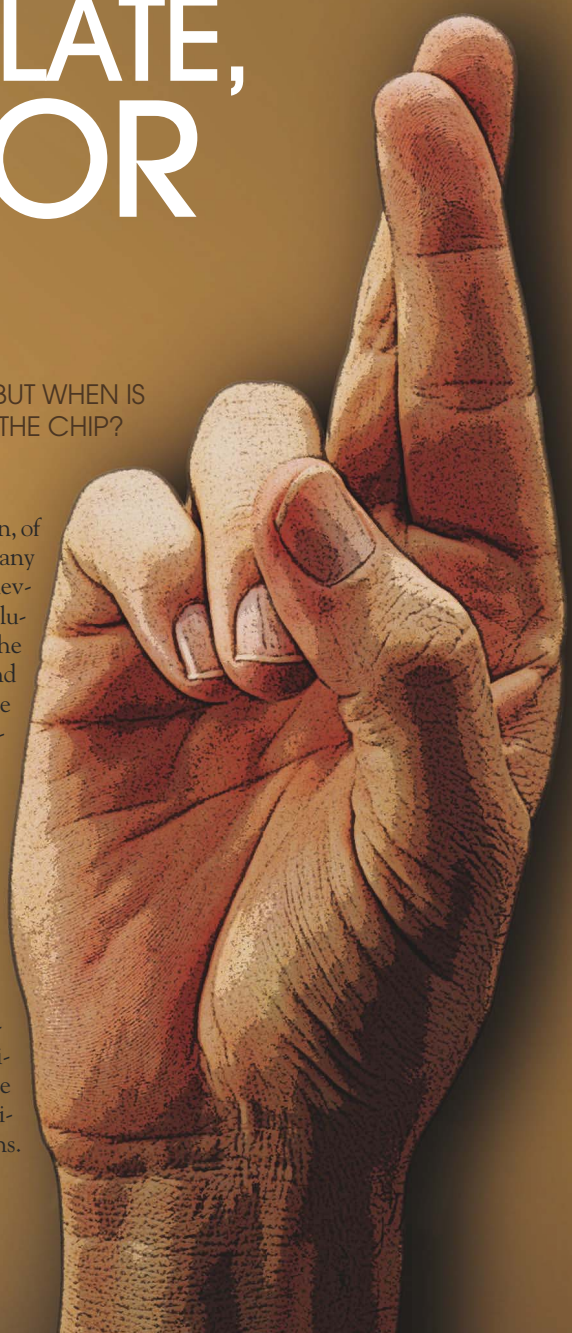

But this approach raises a series of important questions: Should the role of simulation in an FPGA design be the same as it is in an ASIC design? Should the verification team still, at some point, just put the design in the target FPGA and start testing it at speed? If so, when is that point? To find out what design teams are doing today, we asked some of the people who work most closely with FPGA users. And, for reference, we asked a few ASIC design teams who use FPGA prototypes in their verification processes for their views, as well.

PROS AND CONS

Most people begin discussing the question of verification flow by assessing the relative strengths and weaknesses of simulation versus in-FPGA verification. At the risk of boring the experienced, this article will follow the same formula.

The huge advantage of simulation, of course, is access. You can observe any signal in an RTL (register-transfer-level) design down to clock-cycle resolution. You can control the state of the design to whatever degree you find worthwhile. Your knowledge of the RTL and your skill with the simulation environment are the only limits to observability and controllability. You can work interactively on limited areas of a design, or you can set up grand experiments that may run for days. And the relatively speedy setup of simulation runs makes it possible to quickly try lots of things.

Another advantage of simulation is that most simulation environments today are friendly to the use of OVL (Open Verification Library) or SystemVerilog assertions.

Often, there is a straightforward way to transport these assertions into the simulation environment. As assertion-based verification becomes more common, this issue becomes important. In addition, the simulation environment allows you to keep your provisions for stimulating and observing the design separate from the design itself. This consideration might seem secondary, but it can become important in preserving the integrity of the design through intensive verification work.

But simulation is slow. "If you are doing 2 [million]- or 3 million-gate blocks, simulation is great," says Lauro Rizzatti, vice president of marketing at hardware-emulation vendor Eve. "But, at the multiple-block level, simulation can slow down until it is no longer useful."

The complexity of the design is not the only limitation. Phil Simpson, senior manager of technical marketing at Altera, points out that, if the design inherently requires huge amounts of data for validation, simulation may become impractical even at the block level. He offers the example of a video codec, which has so much internal state that a bug may appear only in the middle of a 15-minute video clip. But simulating the compression and decompression of 15 minutes of high-definition video could take most of a career.

AT A GLANCE

- Large-FPGA designs require an ASIC-like design flow.
- Blending simulation and FPGA-based emulation in a verification flow is necessary.
- There are no established guidelines for blending simulation.
- A quick survey suggests a consensus on a verification flow for advanced-FPGA designs.

THE IN-CIRCUIT ARGUMENT

The strengths and weaknesses of in-FPGA verification precisely complement those of simulation. To begin with the obvious, the FPGA is fast. Often, you can run the design at full speed, although, in some cases, this approach may mean more fussing with timing closure than you would want to do early in the design. And, unlike simulation, the FPGA doesn't necessarily slow down as you integrate more blocks into the design. So, it becomes possible to test the entire design rather than individual blocks and to run tests with large real-life data sets instead of surgically crafted test cases.

Because of the speed and the fact that the FPGA has the actual I/O cells that the design requires, you can also test the design in-system: either in an FPGA-development board that you have lashed into the target system or in the target PCB (printed-circuit board) if it is ready to go. Such testing eliminates the lingering uncertainty about whether the test cases really reflect the operating environment of the design. Also, testing the design in its actual board can uncover I/O-related issues—electrical problems, signal-integrity issues, or incompatibilities in high-speed serial protocols, for example—that would be virtually undetectable in any other way. And in-system test creates a software-test platform as a side benefit.

These benefits all pertain to system-level verification. But Altera's Simpson points out that there are some interesting advantages for in-chip debugging of blocks, as well. "Once you get a block into the FPGA, you can use an embedded-processor core, such as Nios, to assist in the debug process," Simpson observes. "The core can move test data on and off the chip, for instance, and it can sequence the tests. In this way, you can test a block in isolation before the circuitry around it is ready.

"In our own IP [intellectual-property]-development group, we have written transactors to run on a Nios core to generate pseudorandom tests," Simpson continues. "I don't know that this practice is common among customers yet, but it can be quite valuable."

With all the advantages of FPGAs, you might wonder what's wrong with just dropping a newly coded core into the FPGA, coding up a test fixture around it, and starting to test. The answer to that question lies in the FPGA's disadvantages.

FPGA WEAKNESSES

The first and most obvious problem is visibility. In principle, every logic element in the FPGA is visible through the chips' debugging interface. But—rather surprisingly, given the power the internal debugging ports offer—vendors estimate that only about half of FPGA users actually synthesize the debugging interfaces into their designs and use them for verification. Xilinx's Patel thinks that num-

ber may be growing as FPGAs get larger.

So, for the most part, if you want to observe a signal in your design, you have to route it out to a pin so that you can put a logic analyzer on it. Given the nature of logic analyzers, you may have to route out a number of other signals, such as internal clocks, as well. This approach means extra work, and it may mean having to reduce the clock frequency on the FPGA if you are trying to observe a fast signal that's nowhere near the I/O blocks. Hence, some managers say that it is vital to include your requirements for observability of FPGA signals in your original verification plan.

The added design work to get access to the signal is one disadvantage. But another problem with stimulating and observing internal nodes in the chip is that it requires you to change the design, rebuild, and resynthesize, risking a clean separation between the design and the testbench. Without careful separation of debugging code from design code and fanatical version control, you can lose track of these changes, risking the equivalent of leaving surgical tools inside the patient.

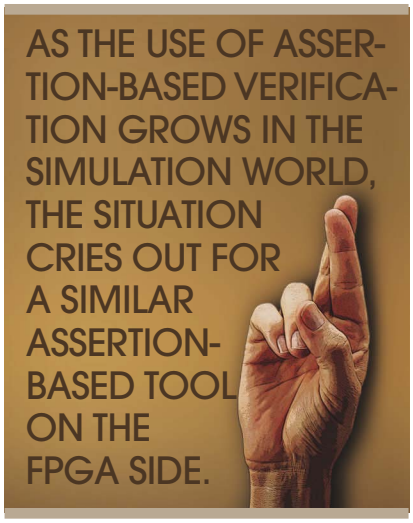

Then, there is the disadvantage of setup time. Synthesis times for large designs are not trivial, and the time to insert instrumentation, rebuild, resynthesize, and remap the design can become a factor in whether to perform a particular experiment. Incremental-synthesis tools can help here, but a 20 million-gate design can mean an overnight build-and-synthesis process.

Finally, there is the problem of transporting the testbench from the simulation environment to the FPGA environment. Stimulating a block now requires circuitry instead of a set of simulation commands. Observing a node requires circuitry and physical instruments, not just commands. And no one seems to have developed a way of systematically moving assertions from the simulation environment to the FPGA, despite the growing acceptance of assertion-based verification. "There is no solution today

SHINING SOME LIGHT ON THE COVERAGE GAPS

Everyone is in favor of the speed of in-FPGA emulation. But the difficulty of setting up, controlling, and observing experiments in an FPGA often forces laborious and time-consuming tests back into the simulation environment. In an ideal world, someone would put together a verification platform that combined the execution speed of FPGAs with the easy setup and excellent access of simulation. Not surprisingly, some vendors have targeted this ideal.

The first efforts, dating back to early in the ASIC era, were "big-iron" logic-emulation systems. These systems are, in effect, specialized mainframe computers in which either custom microprocessors or custom programmable-logic devices simulate or emulate, respectively, the operation of the logic. A representative of such systems would be the Cadence Palladium. The systems offer many times the execution speed of simulation, with what the vendors would argue is at least equal access to the design under test. But they are limited in their capacity to not much more than the size block that you conveniently simulate–unless, that is, you have an impressive capital budget. These systems are major capital investments and therefore beyond the range of most design teams targeting FPGAs for the final design.

In recent years, a number of systems have entered the market–from companies such as Eve–that perform logic emulation in lean environments using commercial FPGAs. Such systems vary in personality from being mini-mainframe-emulation systems to being basically FPGA-evaluation boards with supporting debugging software. In all cases, the attempt is to provide an FPGA-execution environment with less logic overhead in the design than would be the case with a big-iron emulation system. Because of their lower overhead, the FPGA-based systems can often run one to a few orders of magnitude faster than the mainframe-emulation systems. In general, the faster they run, the less convenience of simulation they can preserve. But they run into limitations when a design, including the debugging overhead, is too large for a single FPGA. Partitioning the design is complex and often involves multiplexing signals between FPGAs, which slows everything down.

These systems do offer the software support necessary to move testbenches and data back and forth between the FPGA system and the simulation environment. Eve, for one, reports being at work on a way to import assertions into its environment, as well.

An interesting variation on this theme is the GateRocket system, which the company positions as either a simulation accelerator or an in-circuit emulator. As a simulator accelerator, the system attempts to just slip into the user's simulation environment, accelerating simulation of time-consuming pieces of the RTL (register-transfer-level) logic without disturbing the features of the environment. Assuming the 90/10 rule–that 90% of the simulation time goes into 10% of the code–this ability allows verification engineers to continue using the simulation environment further into the verification flow than would be practical without acceleration. GateRocket claims to support what it calls a synthesizable subset of assertions.

for automatically moving assertions to the FPGA, but we are getting more requests for this capability," Simpson says.

Another weakness here is coverage metrics. Although simulation environments are developing sophisticated tools to measure verification coverage and fuse measurements from different kinds of tools, the notion of coverage barely exists in the FPGA world, and there are no established tools for measuring test coverage of a design and reporting that data back to a central coverage-closure system.

WATCH THE ASIC TEAMS

So there, in a nutshell, are the advantages and disadvantages of each approach. Given that information, how do experienced ASIC design teams—who often employ FPGAs during their own verification flows—balance simulation and FPGA-based testing?

One answer comes from video-processor vendor Ambarella. Executive Vice President Didier LeGall says, "Mostly, we do not use FPGA emulation at all. Our experience has been that, for emulation to work, you need to have very mature RTL. But, by that stage in the flow, [the process of] getting the design into an FPGA and setting up a testbench is a lot of work with little return."

The application may condition LeGall's view. Ambarella SOCs (systems on chips) process high-definition video and 10M-pixel still images at high frame rates and require fast internal clocks and complex algorithms. But LeGall follows his comments on FPGA emulation with an interesting point about the objectives of the whole verification process. "The secret of first-time-working ICs is not perfect verification," LeGall says. "It is software": That is, know where the risk areas are in your design, and plan for software workarounds from the beginning, not as afterthoughts. This strategy does reduce the value of much of the information that verification engineers can glean from extensive FPGA-based testing.

Another view comes from LSI Corp's storage-components group. Bill Wuertz, vice president and general manager of the group, describes how that team does SCSI (small-computer-system-interface) and SAS (serial-attached-SCSI) controllers.

Wuertz says that LSI uses a nearly parallel process, with one verification team working in simulation with one set of objectives while another team works on FPGAs with a different set of objectives. "Early in the design, we create a step we call trial RTL," Wuertz says. "This is the first point where we feel the RTL is basically functionally correct and that the blocks are connected to each other. At this stage, the verification splits into two tracks. The simulation team compiles the design for their tools and goes to work on the individual blocks. A separate team, the systems-engineering group, synthesizes the RTL for an internally developed FPGA board—we are on our fifth generation of the board design now—and begins exhaustive stress testing at the systems level."

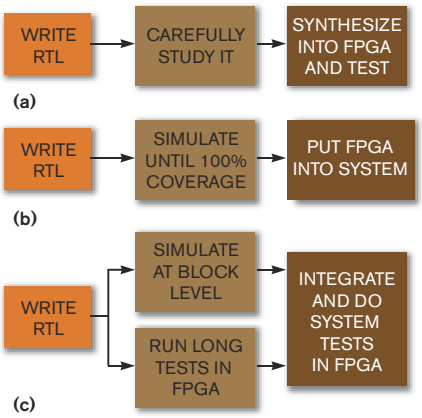

Figure 1 Three approaches to blending simulation with in-circuit debugging show the variety of verification flows teams are using today: traditional FPGA user (a), traditional ASIC designer (b), and the emerging blended approach (c).

As Wuertz describes it, the two groups have different purposes. The simulation team is trying to ensure that the circuitry is correct. The systems team is mostly unconcerned with the circuitry but is verifying that the chip works in the incredibly varied and complex environments of storage networks. Wuertz says that the FPGA prototype may run multiday tests connected to a room full of disk and tape drives. "These tests have evolved over 20 years," he says. "We've learned that it can take very long tests with a mix of different disk drives and tape units to generate just the odd timing alignment that will break the design."

LSI has worked out its own internal tools for linking the two environments. These tools allow the system team to

capture a trace leading up to a failure and convert that data to a stimulus file for the simulation team, for example. Conversely, the simulation team can alert the system team to any risk areas it sees in the design. Establishing communications between the two verification teams in their different environments has been crucial to LSI's two-track approach. The two exchange data throughout the process, and, in the end, the design manager requires closure from both teams.

A CONSENSUS METHODOLOGY

From discussions with FPGA vendors and users, we can see a consensus on a verification flow that blends simulation with emulation (**Figure 1**). Such a flow begins with block-level simulation of the design—not the exhaustive, strive-for-perfection simulation of ASIC heritage, but more of a reality check. The objective is to verify that the block is functioning, that it is doing more or less the right thing at its pins, and that it can meet timing on an FPGA in a lab environment.

At this point, many teams move a version of the block to the FPGA and begin more exhaustive in-circuit testing. This case holds particularly true if the block, such as a video codec, requires long streams of high-speed data to verify correct functioning or if the block includes high-speed-I/O functions. In other cases, simulation work continues on the blocks until all are ready for integration.

Consensus suggests that, when the team begins to integrate blocks—the trial-system build—the FPGA really comes into its own. Here, the design may simply be too large for fast simulation. Or, with the known-working blocks, it may be more productive to troubleshoot integration issues on the FPGA than on the simulator.

But the consensus also suggests that the move from simulation to emulation is not a single irreversible step. Simulation work continues during system emulation, just as it runs in parallel with software development. And most teams use the FPGA emulation to capture and isolate bugs and then pass them back to the simulation team for diagnosis. It is just too painful to do detailed diagnosis on the FPGA.

That summary describes what appears to be happening today, and it points out several serious weak spots in the method. First, it is difficult to move testbench data back and forth between environments. There seems to be no automatic mapping from simulation directives that create a test to FPGA structures that implement the same test. Second, embedded-RISC cores, available from all major FPGA vendors, appear to be a vastly underused resource—able to manage data and control tests but isolated from the simulation testbench. In principle, the simulation team could move its testbench into C code for the embedded core, rather than into RTL for the FPGA. Third, there is no simple path from the data a team collects in an FPGA experiment back to the simulation bench. And, finally, as the use of assertion-based verification grows in the simulation world, the situation cries out for a similar assertion-based tool on the FPGA side.

One indication that these issues are valid is that vendors selling FPGA-based emulation systems address each of them (see **sidebar** "Shining same light on the coverage gaps"). Examples include systems from Eve; simulation accelerators, such as GateRocket; and "big-iron" emulation boxes, such as Cadence's Palladium. Whether this infrastructure will evolve for the ad hoc board-level emulations typical of the FPGA-verification world or whether it will remain the differentiation of big-ticket simulation accelerators and emulation systems remains to be seen. EDN

FOR MORE INFORMATION

Altera
www.altera.com

Ambarella
www.ambarella.com

Cadence
www.cadence.com

Eve
www.eve-team.com

GateRocket
www.gaterocket.com

LSI
www.lsi.com

Xilinx
www.xilinx.com

Go to www.edn.com/090219df and click on Feedback Loop to post a comment on this article.

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

You can reach Executive Editor Ron Wilson at 1-408-345-4427 and ronald.wilson@reedbusiness.com.

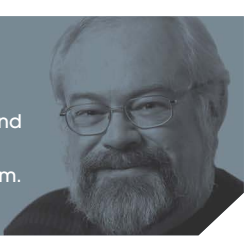

Intersil Micropower Op Amps

2 < 1

Intersil's pinPOINT™ Precision Analog Product Line now has two single-supply op amps that consume less than 1µA.

Intersil's ISL28194, with a typical supply current of just 330nA, **consumes less than half the power of the closest competitor**. The ISL28195 boasts a gain bandwidth of 10kHz on just 1µA.

Ultra-Small Packages

6 Ld SOT-23 | 6 Ld µTDFN

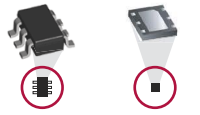

Actual Size (2.6x2.8mm) | Actual Size (1.6x1.6mm)

Other Members of Intersil's Nanopower Family

Device	I_{SS}/Amp (typ)	Function
ISL28194	330nA	Op Amp
ISL28195	1µA	Op Amp
ISL28196	2.5µA	Comparator
ISL28197	800nA	Comparator
ISL60002	400nA	Voltage Reference
X60003	500nA	Voltage Reference

Extended Battery Life:

1. Consumes less current during operation.
2. Consumes only 2nA when amplifier is disabled.

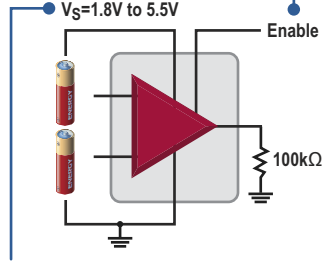

Rail-to-Rail Input/Output

Allows full dynamic range and maximum battery usage.

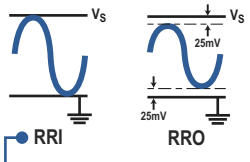

3. Designed for single-supply operation from 1.8V to 5.5V and ideal for low-power precision applications that run on 1.5V alkaline batteries.
4. True rail-to-rail input simplifes system design and supports maximum signal levels.

ISL28194

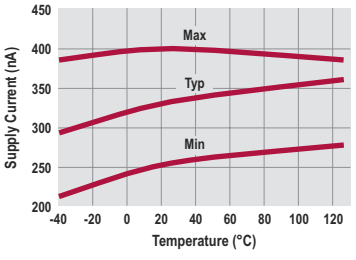

ibuy
intersil.com/ibuy

intersil®

intersil.com/pinpoint

the EVOLUTION of ANALOG™

MAKING OSCILLATOR SELECTION CRYSTAL CLEAR

BY PAUL RAKO • TECHNICAL EDITOR

SILICON AND MEMS OSCILLATORS ARE JOINING VENERABLE QUARTZ CRYSTALS AND CERAMIC RESONATORS IN THE HIGHLY DIVERSE OSCILLATOR MARKET. YOU DON'T NEED A CRYSTAL BALL TO SELECT THE RIGHT DEVICE FOR YOUR APPLICATION, BUT SOME PERTINENT FACTS WILL HELP.

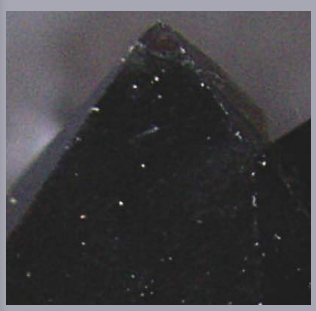

Oscillators are as ubiquitous—and, some might argue, as important—as power supplies in electronics systems, finding use in anything that needs a timing signal, from digital watches to TVs and PCs. Because of their important role in timing for electronics, their failure can bring down an entire system. For example, investigators in 1972 traced the cause of a train crash in Fremont, CA, to a faulty crystal oscillator on a control board. An inappropriate value for the oscillator's tank capacitor overdrove the crystal, causing the part to jump into an overtone frequency. As a result, the train sped up rather than slowed down as it approached a station, and the resultant crash caused many injuries. Because of problems such as this one, many engineers have stopped using raw crystals to make their own oscillators. Instead, they buy off-the-shelf components whose packages contain the amplifiers, tank capacitors, and other parts.

All digital devices require clock sources, such as silicon and MEMS (microelectromechanical-system) oscillators, quartz crystals, or ceramic resonators. Telecommunications and servers, for example, might need a dozen clocks on a single PCB (printed-circuit board). Designers have implemented traditional clock sources with quartz-crystal resonators, but MEMS and pure-silicon resonators are gaining a foothold in this highly diverse market. In addition, less accurate resonators employ ceramic materials, such as lead-zirconium titanate. The application drives the suitability of a technology. For example, if you need a clock source with better than 1-ppb (part-per-billion) accuracy, you must abandon MEMS devices and instead use atomic-resonance devices, such as rubidium- or cesium-clock sources. These devices have 1-ppt (part-per-trillion) accuracy. A GPS (global-positioning-system) satellite, for example, needs this accuracy to maintain synchronicity with the rest of the system (**Figure 1**).

The humble ceramic resonator lies at the opposite end of the accuracy spectrum. You measure the accuracy of these devices as percentages because the parts-per-million measurement yields an unwieldy, large number. The typical initial accuracy of a ceramic resonator ranges from 0.5 to 0.1%, and drift due to aging or temperature

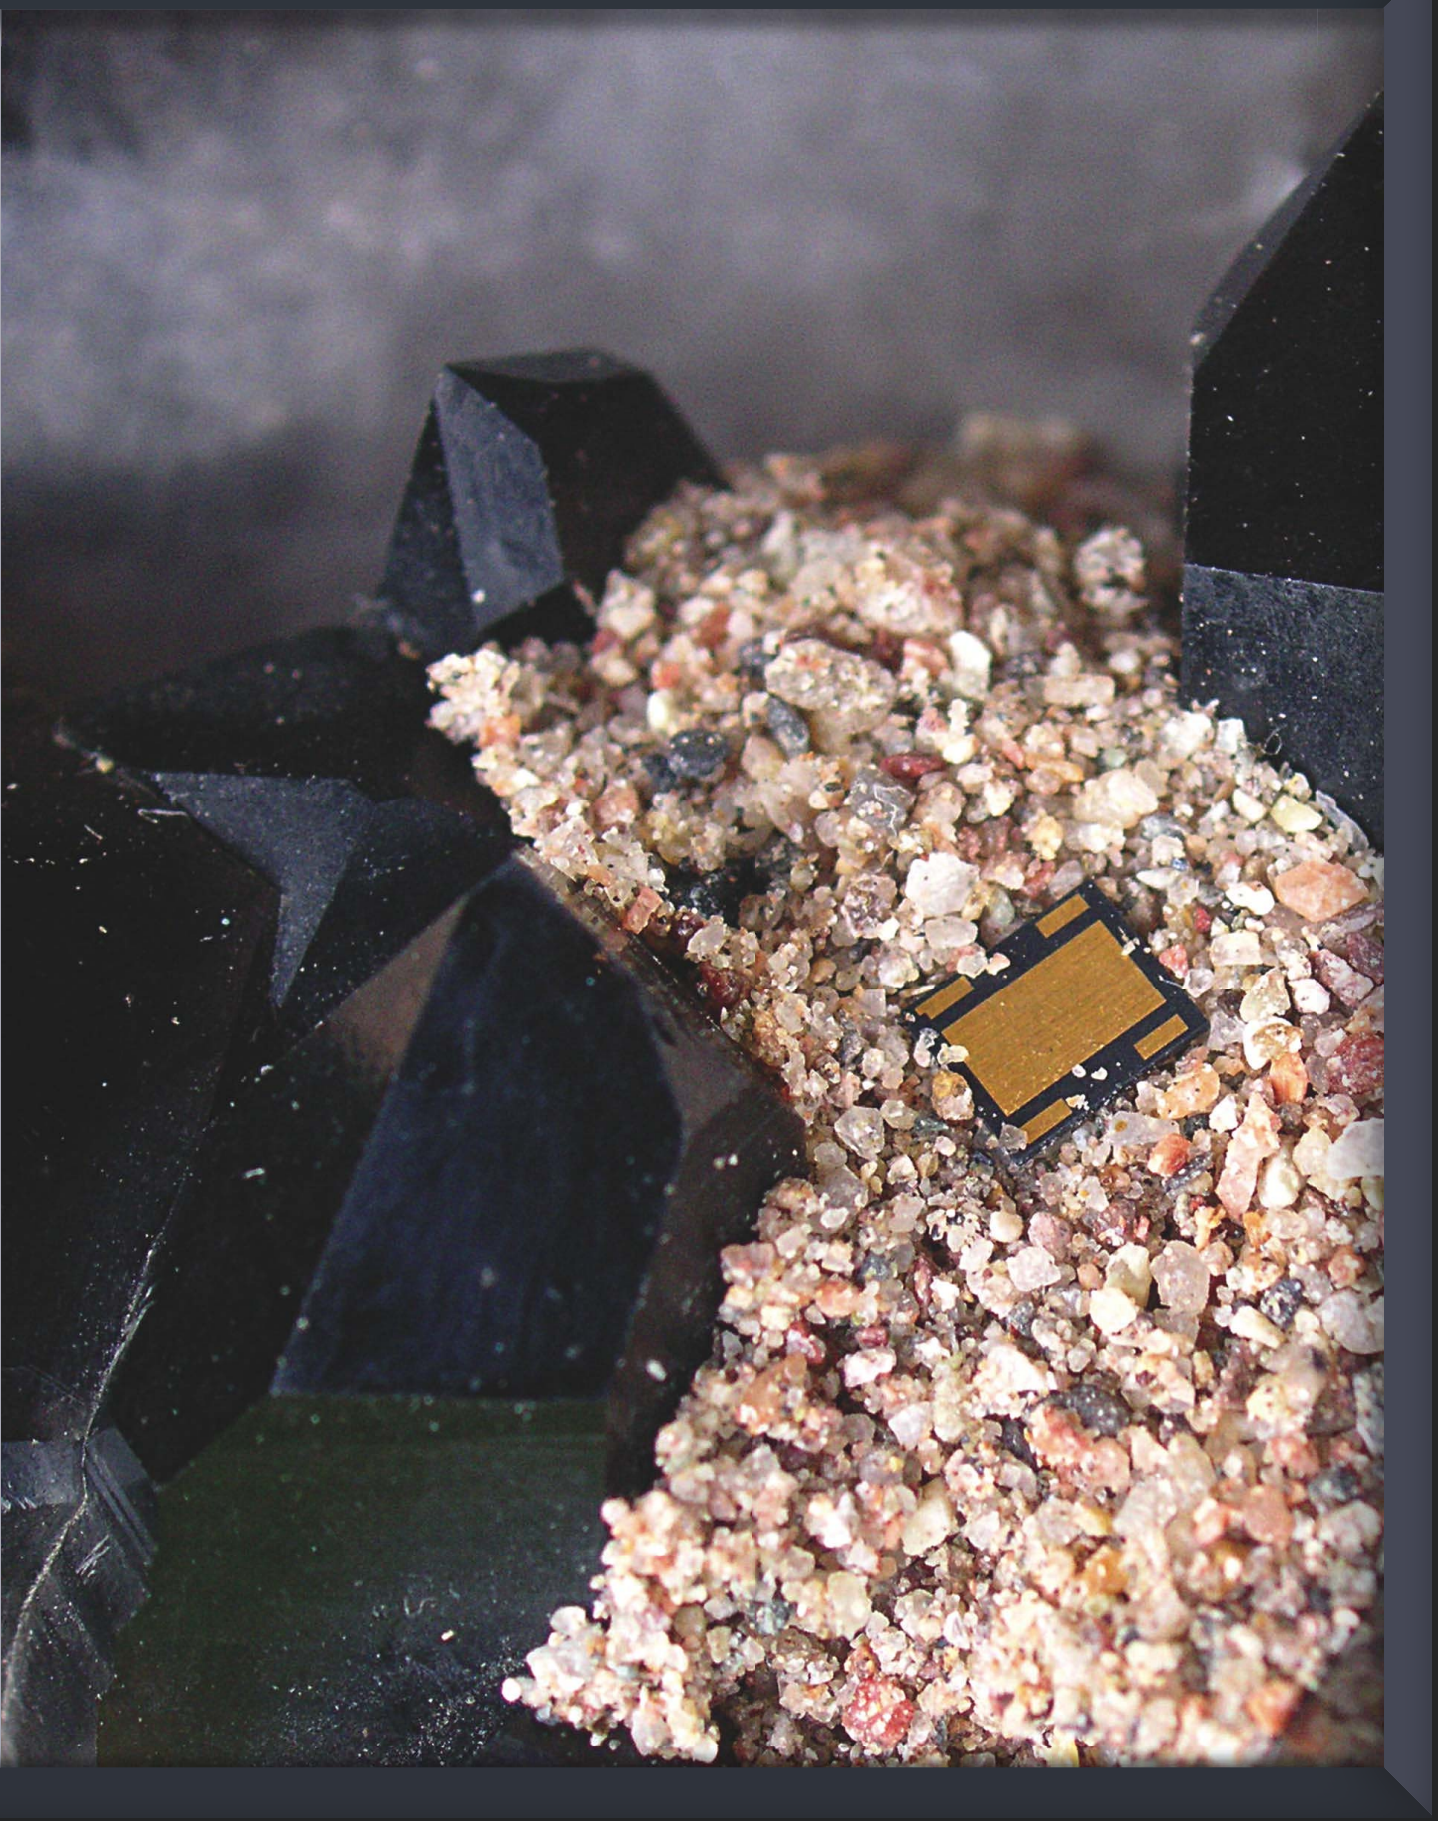

changes can alter this range. As a result, an inexpensive ceramic oscillator can have a tolerance of only ±1.1%, whereas higher-end automotive and commercial products can have accuracies of ±0.25 and ±0.3%, respectively. These tighter-tolerance ceramic resonators target use in USB (universal-serial-bus) 2.0 circuits in commercial and automotive CAN (controller-area-network)-bus applications that operate at −40 to +125°C. Low-cost ceramic resonators, at frequencies of 200 kHz to almost 1 GHz, work well in embedded systems in which timing is not critical. Ceramic devices offer faster start-up and are often smaller than their quartz counterparts. They are also more tolerant of shock and vibration. Ceramic resonators are available from such manufacturers as Murata, Oscilent, AVX, TDK, and Panasonic.

For digital systems using UARTs (universal synchronous/asynchronous transceivers), you should do an error-budget analysis to ensure that the baud rate you derive from the resonator frequency stays in spec. If you use the UART only during code development, you might be able to switch in a ceramic resonator in production and save money.

Note that some silicon oscillators use RC (resistance/capacitance) or LC (inductance/capacitance) tanks and no ceramic or quartz. These oscillators also have a broad range of accuracy commensurate with their price. Companies such as STMicrosystems manufacture oscillators that have all the advantages of ceramic resonators and can be even smaller and less expensive. "The main point regarding silicon oscillators is their robustness versus the fragility of a crystal," says Louis Grantham, product-marketing engineer at the company. "Plus, the manufacturability of crystals is more difficult than that of ICs."

AT A GLANCE

- Ceramic resonators have accuracy of 1 to 0.1%, compared with 1 to 100 ppm (parts per million) for quartz oscillators and 1.5 to 100 ppm for silicon devices.
- Silicon and MEMS (microelectromechanical-system) oscillators are more resistant to shock and fit into smaller packages than can ceramic devices.
- Quartz oscillators take longer to start but typically use less power than other types.
- Power consumption of any oscillator type depends on output loading.

IT BEGAN WITH QUARTZ

A crystal oscillator uses the mechanical resonance of a vibrating crystal of piezoelectric material to create an electrical signal with a precise frequency. This frequency commonly keeps track of time, as in quartz wrist watches; provides a stable clock signal for digital integrated circuits; and stabilizes frequencies for radio transmitters and receivers. Engineers have been using these crystals to establish radio frequencies since the 1920s when AM Nicholson at Bell Telephone Laboratories and Professor WG Cady of Wesleyan University, working with Rochelle salt as the crystal, observed the reaction of a resonant piezoid on a driving circuit (**Reference 1**). However, researchers did not develop methods for high-volume manufacturing until World War II (**Reference 2**). If you cut the resonator elements in a quartz crystal at the correct angle in relation to the crystal matrix, you can eliminate effects due to temperature. Some cut crystals have a zero temperature coefficient, whereas the LC cut finds use as a thermometer (**Figure 2**).

Just because you can make a quartz crystal from a mineral, do not assume that a crystal oscillator is a low-tech device (**Reference 3**). Manufacturers of today's quartz crystals grow the crystals in large reactor ovens, or autoclaves, at high temperatures and pressures of 30,000 psi (pounds per square inch) or more (**Reference 4** and **Figure 3**). It can take months to grow the quartz crystals in an autoclave, and any seismic activity or the slightest degradation or loss of electric power to the heaters would ruin the entire lot. For that reason, NDK, a Japanese company that has made crystals for decades, now has autoclaves in Belvidere, IL. The company based its decision to locate its newest facilities in the Midwest on the reliability of the region's electric-power grid and low incidence of earthquakes.

"We put mined quartz into a huge vessel that we adapted from battleship-cannon technology," says Craig Taylor, general manager of business and applications development at the company. "We then put seed quartz in baskets above the mined quartz. By adding [a sodium-carbonate or sodium-hydroxide] electrolyte and applying great pressure and temperature, all the natural quartz dissolves and migrates upward. It attaches itself to the seed quartz, and all the dirt and impurities are left in the bottom of the vessel."

Amplification and buffering turn a crystal into an XO (crystal oscillator).

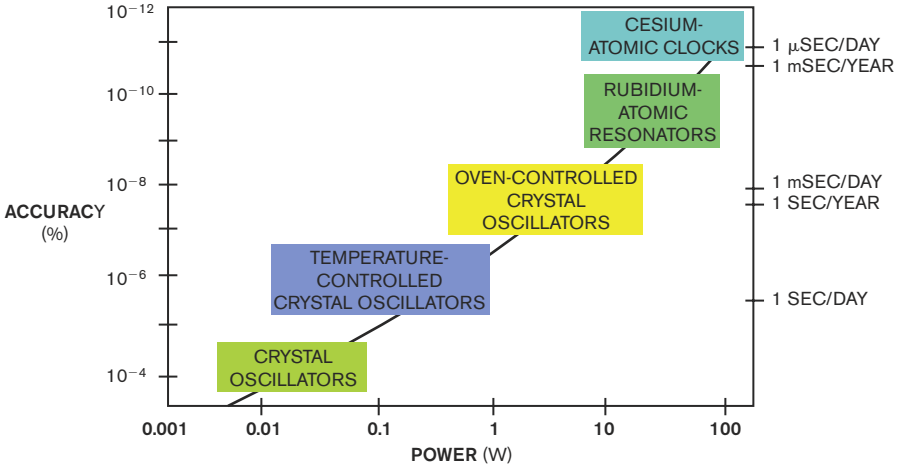

Figure 1 Oscillators trade off power versus accuracy, so selection depends on the needs of the application (courtesy John R Vig, US Army Communications-Electronics Command).

Optical Sensing Made Easy

Melexis' New SensorEyeC™ Series Simplifies Design In Virtually Any Optical Sensing Application

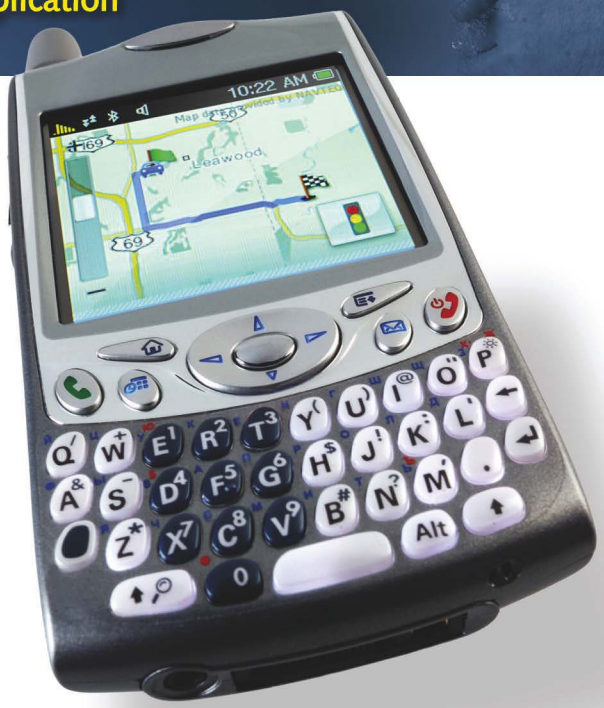

Features and Benefits:

- Integrated Sensor with Switch Output or Linear Output (voltage or frequency)
- Low Temperature Coefficient
- High Linearity
- Solder Reflow 260°C

Applications:

- LCD Backlight Dimming
- LED Lighting Control
- Proximity Sensing
- Printer/Copier Sheet Detection

Integration Makes Life Easier

Melexis introduces three new **Low-Cost SensorEyeC™ Optical Sensor ICs**, each with integrated amplifier, output stage and temperature compensation included in one chip.

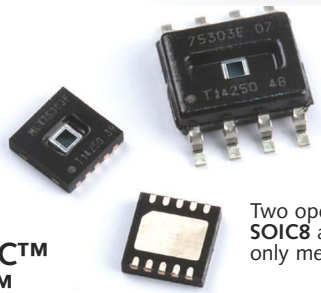

The Melexis Sensor EyeC™ Family uses unique and patented open-cavity package technology which combines high robustness (including solder reflow of 260°C) with an unobstructed optical path.

Two open-cavity packages are available: **SOIC8** and a small size **DFN** package that only measures 3mm x 3mm x 0.65mm.

MLX75303 Optical Switch SensorEyeC™
MLX75304 Light-to-Frequency SensorEyeC™
MLX75305 Light-to-Voltage SensorEyeC™

Small things make a big difference.

For more information call
(603) 223-2362
or visit our website:
www.melexis.com

Melexis
Microelectronic Integrated Systems

Adding temperature-compensation circuitry yields a TCXO (temperature-controlled crystal oscillator) with accuracy of 1 ppm (part per million), and putting the entire oscillator into a temperature-controlled-oven package yields an OCXO (oven-controlled crystal oscillator) with accuracy in the part-per-billion range. A 30-MHz oscillator with 1-ppm accuracy will have an error of only 30 Hz over time and temperature. Only rubidium and cesium-atomic standards are more accurate, in no small part because atomic resonance is independent of temperature. Some companies are also offering PXCOs (programmable crystal oscillators), which allow you to write to digital registers in the chip to adjust the frequency.

Adding a PLL (phase-lock loop) to a crystal allows it to emit a higher frequency at lower cost than a quartz crystal, according to Nancy Zhang, product-marketing manager at Pericom. According to Kay Annamalai, senior marketing director at the company, a third-overtone crystal allows frequency of only 150 MHz. For needs beyond that frequency, designers often add a PLL. He describes a proprietary Pericom technique that multiplies the frequency without using a PLL. This method offers the same reduction in crystal cost but also improves jitter performance. The company's XP technology avoids the use of a PLL but allows frequencies greater than 150 MHz.

A PLL can also improve performance, according to CS Lam, a director at Epson Electronics America. Lam notes that the company has achieved less-than-10-ppm accuracy using fractional-PLL circuitry. He also points out that the first PLL-based crystal oscillator with less than 1-psec-rms phase jitter at 12 kHz to 20 MHz appeared in 2004 (**Reference 5**).

Adding a PLL also allows you to electronically vary the frequency of operation to help comply with FCC (Federal Communications Commission) and CE (Conformité Européenne) radiation standards. When the PLL varies the clock frequency, the high-amplitude spike of EMR (electromagnetic radiation) or EMI (electromagnetic interference) spreads the radiation over a frequency band. Note that this technique does not reduce the amount of radiation; it just sweeps over a band such that the energy-measurement instruments give a lower reading. The EMI leaves the measuring bandwidth of the spectrum analyzer, lowering the measurement readings and helping your product pass the test for compliance.

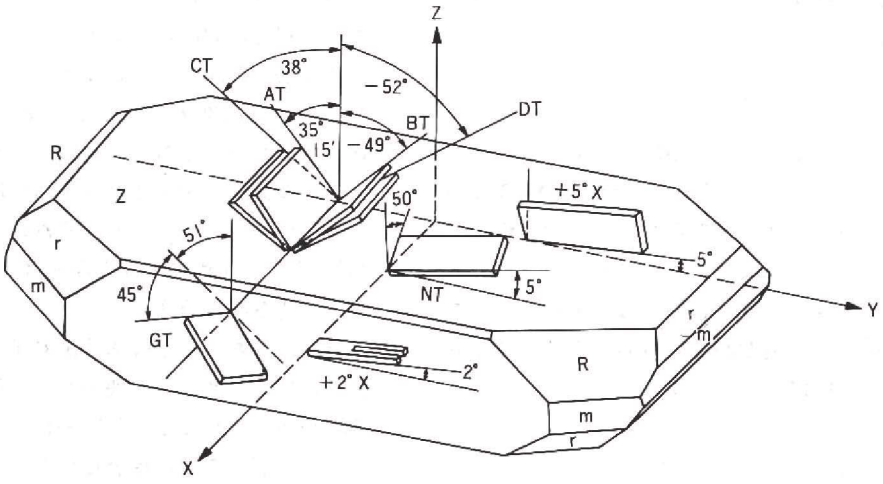

Figure 2 You can take many cuts from a quartz crystal. Watch producers employ small tuning-fork crystals to the right of the GT cut (courtesy Quartz Crystal Industry Association of Japan).

Spread-spectrum clocking also plays a role in oscillator selection. It has two broad applications: power supplies and system clocks in computers and telecommunications. Power supplies can use oscillators that vary by as much as 10%, spreading the energy over a wide band and dramatically reducing the measurement results. These devices, which employ ring oscillators or LC tanks, do not require quartzlike accuracy. The PLL portion of the oscillator circuitry uses the output of the silicon oscillator to create a spread-spectrum clock. Just as with other pure-silicon oscillators, the parts are more resistant to shock and have faster start-up. Because an LC tank or a ring oscillator has a much lower Q (quality) factor than any quartz crystal or MEMS resonator, you might expect the silicon oscillator to take more energy to keep going. However, it takes only microwatts to maintain the oscillation because the power consumption in oscillators depends on the process and architecture of the PLL and temperature-compensation circuitry.

The other application of spread-spectrum clocking is digital systems that have

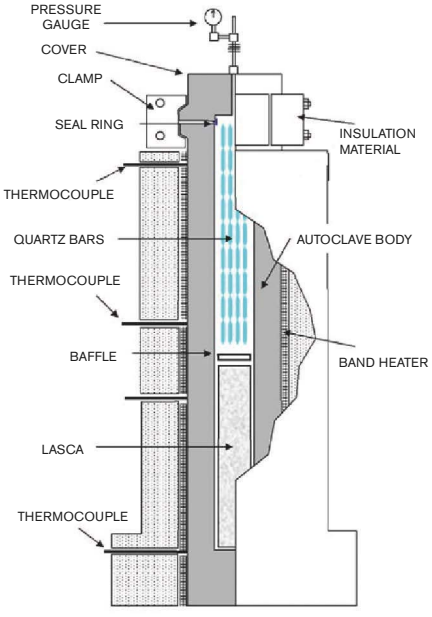

Figure 3 Quartz crystals grow to nearly perfect purity in an autoclave. A thick iron casing, which some manufacturers adapt from cannon barrels, contains the tremendous pressure and temperature of the molten system. An electrolyte melts the raw quartz, or lasca, and then deposits seed bars on the quartz at the top of the autoclave. It takes months for manufacturers to grow crystal bars, which they then cut and grind for electronics use (courtesy Quartz Crystal Industry Association of Japan).

Microsoft

it's everybody's business to get answers as FAST as yesterday.

Does your enterprise software help give you the insight to keep up with the pace of demand?

ask for *people_ready* enterprise high tech solutions

microsoft.com/hightech

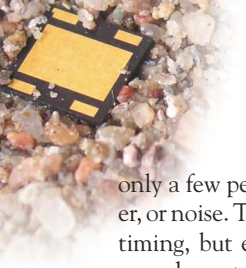

only a few percentage points of dither, or noise. They must maintain tight timing, but even a small amount of spread-spectrum clocking can allow a system board to pass FCC testing. Pericom's Annamalai notes that the spread-spectrum clocking works especially well in memory subsystems. "Memories tend to be high-speed, so you want to spread that single spectrum," he says. The company uses the Hershey's Kiss spreading profile, which takes its name from the popular candy, whose shape it mimics.

Lexmark discovered and patented this profile. To understand this response, imagine that a sinusoidal frequency is modulating the operating frequency of the system clock; the average time that the oscillator spends at the end frequencies will be greater than the time it spends between the ends. In other words, the clock lingers at the outside bounds of the frequency excursion, where the sinusoidal modulation is slowly changing direction. This change gives rise to the "bat-ears" frequency-domain profile (**Figure 4**). By using the Hershey's Kiss waveform, manufacturers can eliminate bat ears and allow your system to pass FCC testing.

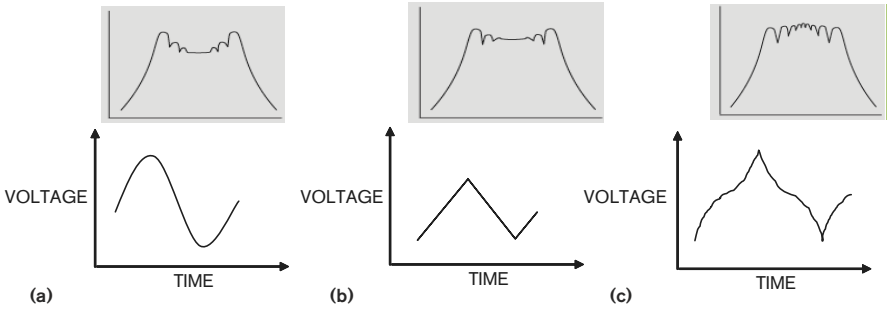

Figure 4 Spread-spectrum oscillators have changing frequencies to reduce electromagnetic radiation. Circuitry that modulates the clock with a sine wave produces a frequency-domain plot that exhibits "bat ears" at the sides of the plateau (a). A triangle-wave modulation creates a frequency-domain plot that still has discernible amplitude overshoots at the edges of the plateau (b). With the Hershey's Kiss modulation (c), the plateau in the frequency response is optimally flat (courtesy Pericom Semiconductor).

Pericom uses quartz, a high-Q source with low jitter for system clocking. By melding this crystal with a high-performance, low-jitter PLL, the company provides a swept-spectrum oscillator that consumes minimal power and combines the benefits of both quartz and silicon.

Power consumption is another factor to consider when selecting an oscillator. Start-up Mobius Microsystems offers a pure-silicon oscillator that provides near-quartz accuracy, fast start-up, and high shock resistance. However, the company achieves these features by running the silicon tank at a high frequency and then dividing the frequency down, resulting in higher power consumption than that of quartz-based devices. Silicon-process and -design techniques are improving quickly, though, so silicon oscillators should rapidly improve in almost every specification.

Another company making advances in silicon technology, Silicon Laboratories, makes both pure-silicon oscillators and devices that adapt a PLL to a quartz crystal (**Reference 6**). These parts offer accuracy matching that of low-end crystal oscillators. "The Holy Grail is to eliminate the [need for a] mechanical resonator," says Mike Petrowski, director of marketing for timing products at the company. "If you eliminate that mechanical resonator, you improve reliability, simplify your manufacturing flow, and [make the device] easier to mass-produce." Petrowski maintains that silicon oscillators do not consume excess power because they achieve their accuracy with temperature compensation, not a PLL that divides down a higher frequency.

Note that *silicon oscillator* can mean a lot of things—from a cheap part that replaces a ceramic resonator to a device with quality matching that of quartz. Always evaluate the power consumption to ensure that the technology you use is appropriate for your application. Be aware of subtleties, such as the fact that a quartz crystal or a MEMS oscillator draws more current for the several milliseconds that it takes to start up. This excess current draw may present a problem in micropower applications or in applications that require the part to continually start and stop.

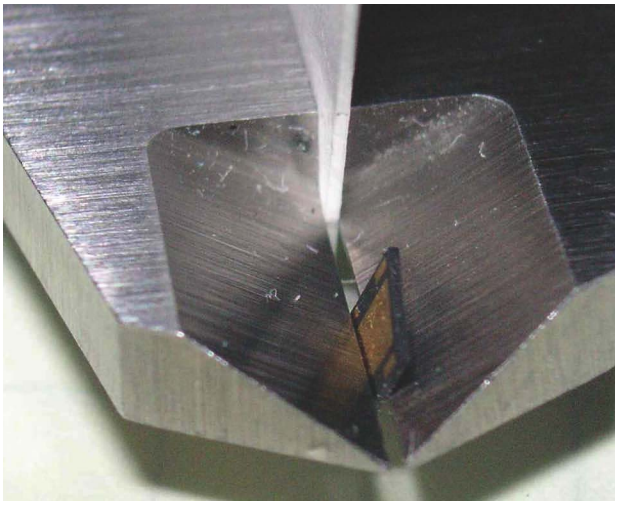

Figure 5 Electronics continually get smaller, and oscillators are no exceptions. Silicon and MEMS oscillators can fit into tiny packages, including this 0.2-mm-thick oscillator that is thinner than a business card (courtesy SiTime).

Besides accuracy and power, another key spec in oscillators is jitter, or phase noise, the cycle-by-cycle change in frequency. For example, a stable device could alternate operation at 1 MHz during one cycle with operation at 2 MHz during the next cycle, providing an average frequency of 1.5 MHz. However, this huge change in cyclic frequency would make the part useless in most applications; a switching power supply could

not work over such a broad range, and a PLL would have difficulty locking into such a high-jitter source. Any system using such an oscillator could not include ADCs or DACs because the variance in frequency would ruin the digital processing, even though the average frequency is stable. For this reason, the oscillator-design groups in many companies are in the analog sections of the company. A PLL is an analog component, and many of the specs, such as jitter, are important in analog circuitry.

Although jitter and phase noise are the time- and frequency-domain representations, respectively, of the same condition, it is easy to misrepresent jitter

AS RECENTLY AS FIVE YEARS AGO, THE PRESENCE OF A PLL IN ANY OSCILLATOR WOULD HAVE HAD A NEGATIVE IMPACT ON THE DESIGN'S JITTER PERFORMANCE.

specs, according to Doug LaPorte, design-section leader for signal-conditioning products at Linear Technology. Some companies spec jitter only over a certain frequency range, he says. These companies may produce phase-noise plots that integrate only a certain amount of that phase noise and omit other bits of the noise. Optical-communications standards, such as SONET (synchronous-optical networking), transmit, perform PLL, and then retransmit. The loop has a design bandwidth that allows the system to reject phase noise outside the loop but allow noise inside the loop. "[These manufacturers] get away with a spec of, say, 20 kHz to 10 MHz," LaPorte notes. "Beyond that [limit], they don't care."

As recently as five years ago, the presence of a PLL in any oscillator would have had a negative impact on the design's jitter performance. Silicon Labs' Petrowski says that the company used to worry about the bad reputation that older PLLs had. "When we launched our PLL-based oscillators, we were concerned there would be some negative connotations," he says. "We have a significant amount of R&D and patents rolled into these parts, and it is absolutely possible to make a low-jitter PLL, especially with finer IC geometries."

The nature of a PLL, with its analog filtering, phase detection, and VCO (voltage-controlled oscillator), gives rise to increased jitter at all points of the circuit. During the last five years, designers began to use bond wires as small inductors on ICs or to place discrete spiral inductors on IC die. Now that IC designers can use inductance as well as capacitance as reactive elements, the filters and tank circuits can all have higher Q and more poles and zeros. Maxim Integrated Products, for example, uses LC-based oscillators rather than ring oscillators in its designs. "Ring oscillators tend to have more jit-

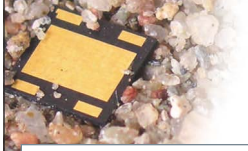

Go to www.edn.com/090219cs and click on Feedback Loop to post a comment on this article.

ter than an LC type," says Paul Nunn, business manager for precision oscillators at the company. Such companies as Pericom, Silicon Labs, SiTime, On Semiconductor, and Fox Electronics use these high-quality PLLs because they allow oscillators to have adjustable frequency and low jitter.

MEMS DEVICES ADD A TWIST

MEMS oscillators share the amplifier and perhaps the PLLs of quartz oscillators but use a small, vibrating silicon mass rather than a quartz crystal. This approach offers better MTBF (mean time between failures), shock resistance, and reliability. For example, JEDEC (Joint Electron Device Engineering Council) and HTOL (high-temperature-operating-life) testing of silicon yields a 500 million-hour MTBF, whereas quartz yields only a 10 million- to 30 million-hour figure, according to Piyush Sevalia, vice president of marketing at SiTime. And, whereas a 1-kHz vibration readily shows up in the jitter performance of a quartz oscillator, neither MEMS nor silicon oscillators are sensitive to this vibration. MEMS devices resonate at a fundamental frequency in a mode that incident vibration does not modulate. However, MEMS and quartz oscillators have slower start-ups than those of pure-silicon oscillators.

One challenge in manufacturing MEMS oscillators is keeping the vibrating silicon element atomically clean. Even a monomolecular layer of atoms on the vibrating beam can cause the part to go out of spec, and manufacturers use various methods to overcome this challenge. For example, Discera uses "getters," reactive materials for removing traces of gas to absorb any incidental gases or material over the life of a part. SiTime, on the other hand, uses a technology that Bosch first developed (**Reference 4**).

Rather than place a glass or epoxy cap over the MEMS element, SiTime creates the silicon beam in a matrix of glass, caps that matrix with polysilicon, and then dissolves the glass with hydrofluoric acid. The company then seals off the beam with a thicker layer of polysilicon. All this work occurs in an epitaxial reactor, a high-vacuum semiconductor machine that offers one of the cleanest environments on earth. This exotic processing allows SiTime to sell oscillators with quality that rivals that of quartz oscillators. The company's products incorporate both a MEMS-resonator die and a CMOS die into one package that can be smaller and thinner than a quartz oscillator (**Figure 5**).

Both Discera and SiTime oscillators are fully programmable because they integrate PLLs. Discera also offers a less-than-$500 kit that includes a handheld programmer and 200 parts; the programmer connects to the USB port of your computer. According to Gerry Beemiller, vice president of sales and marketing at Discera, the kit allows you to build a highly accurate oscillator with 1- to 150-MHz frequency. In contrast, SiTime touts fast turnaround rather than in-field programmability. Because it employs no quartz processing, the company claims that it can—within days—provide you a part that operates at any frequency.

Accurate time is always critical in a sampled-data digital system. If the somewhat-shabby timebase of a ceramic resonator or low-performance silicon oscillator doesn't fit your needs, you can choose a part from the entire spectrum of quartz technology. Add those choices to the higher-performance silicon oscillators from Silicon Labs and the MEMS oscillators from SiTime and Discera, and

FOR MORE INFORMATION

AVX
www.avx.com

Discera
www.discera.com

Epson Electronics America
www.eea.epson.com

Fox Electronics
www.foxelectronics.com

Lexmark
www.lexmark.com

Maxim
www.maxim-ic.com

Mobius Microsystems
www.mobiusmicro.com

NDK
www.ndk.com

On Semiconductor
www.onsemi.com

Oscilent
www.oscilent.com

Panasonic
www-images.panasonic.com/industrial/components/ceramic/ceramic.htm

Pericom
www.pericom.com

Silicon Labs
www.silabs.com

SiTime
www.sitime.com

STMicrosytems
www.st.com/stonline/products/families/clock_and_timing_circuits/silicon_oscillators/silicon_oscillators.htm

TDK
www.tdk.co.jp/tefe02/piezo_buzzer.htm

you can see that choosing an oscillator is crucial. Understand all the trade-offs involving accuracy, power consumption, jitter, and programmability, as well as any spread-spectrum requirements. And remember that it is always desirable to at least provide for a spread-spectrum oscillator on your power supply or system clock in case you fail FCC testing. This scenario always occurs at the worst possible time: just when you are ready to ship your product; having a sophisticated oscillator that you can substitute for the fixed one is always good insurance. After weighing all these factors and how they match the needs of your application, choosing an oscillator should become crystal clear. EDN

REFERENCES

1 Botton, Virgil E, "A History of the Quartz Crystal Industry in the USA," *Proceedings of the 35th Annual Frequency Control Symposium*, 1981, pg 3, www.ieee-uffc.org/fc_history/bottom.html.

2 Thompson, Richard J, Jr, *Crystal Clear: The Struggle for Reliable Communications Technology in World War II*, John Wiley & Sons Inc, 2007, ISBN-13: 978-0-470-04606-7, ISBN-10: 0-470-04606-6.

3 Travis, Bill, "Crystal oscillators: looking good in wireless systems," *EDN*, Nov 20, 1997, pg 39, www.edn.com/archives/1997/112097/24cs.htm.

4 Conner, Margery, "Crystal grinding: when electronics were *really* hands-on," *EDN*, April 4, 2008, www.edn.com/blog/1470000147/post/580024458.html.

5 Lam, CS, "A Review of the Recent Development of MEMS and Crystal Oscillators and Their Impacts on the Frequency Control Products Industry," 2008 IEEE International Ultrasonics Symposium, 2008.

6 Prophet, Graham, "'Crystal' oscillator comes without the crystal," *EDN*, Oct 30, 2008, pg 20, www.edn.com/article/CA6607209.

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

BY LLOYD WATTS • AUDIENCE

Suppressing nonstationary noise in mobile handsets

MOBILE CARRIERS ARE INTIMATELY AWARE OF THE ROLE THAT VOICE QUALITY PLAYS IN CUSTOMER RETENTION. ONE OF THE PRIMARY FACTORS AFFECTING VOICE QUALITY IS ENVIRONMENTAL NOISE, SO ANY MEANS OF SUPPRESSING NOISE PROVIDES A POTENTIAL DIFFERENTIATOR FOR HANDSET MANUFACTURERS.

Until recently, noise-suppression technology focused on reducing slow-changing stationary noise sources. However, nonstationary noise sources are fast-changing, and the current technology does not suppress them. As a result, subscribers cannot reliably use their handsets on busy streets, in crowded restaurants, or even at home.

Suppressing nonstationary noise brings substantial benefits to both subscribers and carriers. Users gain the freedom to speak and hear clearly wherever and whenever they want, enjoy increased privacy by being able to speak softly in noisy environments, and need not leave important conference calls. Carriers will see a reduction in customer churn, increased airtime usage, more efficient use of network bandwidth, and significant savings of capital and operational expenses.

You can readily recognize stationary noise, such as a loud fan in the background, because of its relatively constant nature, and you can effectively subtract this noise through conventional signal-processing techniques (**Figure 1**). Nonstationary noise, in contrast, involves rapid or random change, such as a person talking, background music, or keyboard typing (**Figure 2**). By the time you recognize nonstationary noise as noise, it has already passed, so it requires more sophisticated noise-suppression techniques.

MULTIPLE-MICROPHONE NOISE SUPPRESSION

Next-generation noise-suppression techniques, such as ASA (auditory scene analysis), beam forming, and BSS (blind source separation), use multiple microphones to more accurately identify, locate, and suppress noise sources than is possible with a single microphone. ASA uses psychoacoustic grouping principles to separate noise sources from the voice of interest. ASA's developers based the technology on the human auditory pathway; ASA processes noise in the same way that people listen to sound, using powerful cues such as pitch and spatial location of sound sources.

Beam forming uses multiple microphones to track a cone-shaped area of interest to locate and identify noise sources. Because the voice of interest lies within the cone, listeners can quickly differentiate as noise any sound sources outside the cone. Today's handset manufacturers recognize the trend toward multiple microphones and have begun to introduce second microphones into handset architectures.

BSS uses a linear unmixing technique to decompose the input sound mixtures into independent sources. The challenge in using BSS is that its linear unmixing technique requires as many microphones as there are noise sources, and it can suffer from convergence problems in the presence of reverberation when too many simultaneous noise sources are present.

USING MULTIPLE CUES FOR GROUPING

The human auditory system can hear voices in noisy environments because it uses all the information available in the signals arriving at the two ears. Like the human auditory system, ASA uses many methods to analyze the signals, resulting in multiple cues that can group the spectral energy into the corresponding sound sources. Some of the more important cues include pitch, spatial location, and common onset time. *Pitch* refers to the harmonics that a pitched sound source generates.

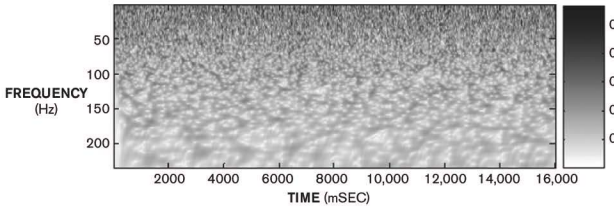

Figure 1 Stationary white noise is patterned and slow-changing.

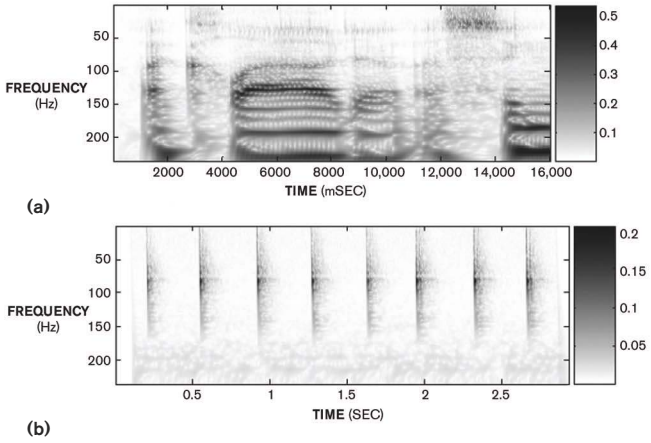

Figure 2 Speech (a) and pen-tap noise (b) are examples of nonstationary noise: rapidly changing, random, and often also containing harmonics across a wide frequency range.

These harmonics form distinct frequency patterns, so you can use them to distinguish one sound from another. Pitch is one of the primary cues, for example, in distinguishing between male and female voices. *Spatial location* refers to the location of a sound, based on its distance and direction; you can use spatial location to group sounds, thereby differentiating them from the voice of interest. *Common onset time* refers to the fact that, when two bursts of sound energy and their corresponding harmonics simultaneously occur, they are likely from the same source.

Traditional noise-suppression techniques must first converge before they can remove noise, making them ineffective in suppressing nonstationary noise sources. By using fast-acting cues to characterize sound, you can identify and remove even instantaneous events, such as a finger snap.

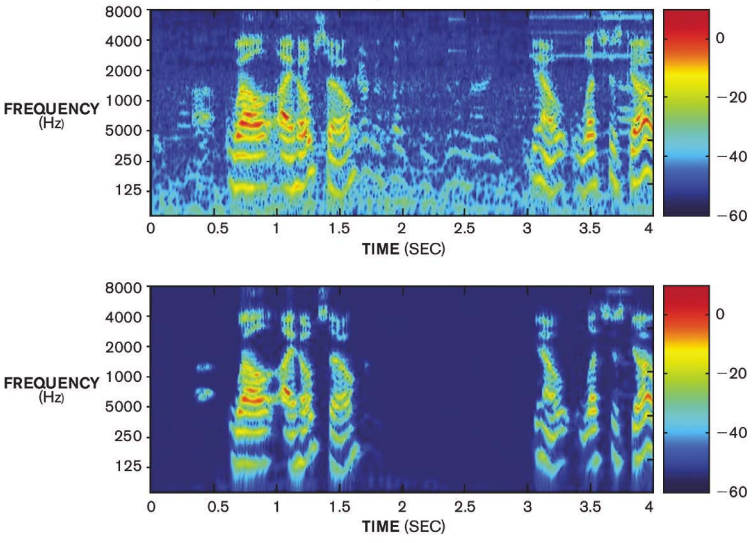

Figure 3 Next-generation techniques can dramatically improve the signal-versus-noise characteristics of captured audio.

LOGARITHMIC VERSUS LINEAR SCALES

The familiar FFT (fast Fourier transform) decomposes frequency components on a linear scale that limits spectral resolution at low frequencies; it also uses a constant frame size and frequency-independent bandwidth. In contrast, an approach such as the FCT (fast cochlea transform) mimics characteristics of the human cochlea and operates on a logarithmic frequency scale. As a result, the FCT does not limit spectral resolution. By operating continuously instead of in frames, FCTs also reduce processing latency, making them appropriate for identifying nonstationary noise sources. Additionally, FCTs operate with frequency-dependent bandwidth, so you can more precisely match the time-versus-frequency trade-off at each frequency of the human hearing range.

Certain techniques, such as beam forming, require a specialized cardioid unidirectional microphone. Cardioid microphones cost more and have tighter tolerances than do omnidirectional microphones. They also require individual calibration and matching to within 1 dB, introduce restrictions on spacing, and add as much as 12 dB of noise because of sensitivity to wind and breath. Beam forming is also limited in that any distractions in the beam of interest will incorrectly pass through as part of the voice of interest.

> Go to www.edn.com/ms4303 and click on Feedback Loop to post a comment on this article.

You traditionally remove echoes using separate echo-cancellation techniques. Such techniques can be computationally intensive because they must calculate echo reflections, and they offer poor performance in the presence of rapidly changing noise sources. Grouping cues enable you to treat echoes as simply another noise source. Instantaneous suppression becomes possible because you need neither to calculate nor to track the changes of echoes, providing echo-suppression performance to 46 dB.

NEW TESTING STANDARDS

The mobile-equipment industry continues to drive test standards to reflect higher levels of voice quality through innovations in noise suppression. To ensure the best quality for products, the recently amended ITU (International Telecommunication Union) P.835 specification provides a consistent test method for measuring and reporting voice quality with active-noise-suppression technology.

Effective suppression of both stationary and nonstationary environmental noise is essential if handset manufacturers and carriers are to keep pace with their competitors. By employing next-generation noise-suppression techniques, developers can reduce noise levels in handsets by as much as 35 dB under a range of operating conditions (**Figure 3**). EDN

AUTHOR'S BIOGRAPHY

Lloyd Watts is the founder and chief technology officer of Audience and, as such, provides ongoing guidance and impetus for the company's core technology direction as well as the vision of neuromorphic computing for voice systems. Before joining Audience, he was principal researcher at Paul Allen's Interval Research Corp, where he collaborated with leading auditory neuroscientists on his vision of a machine that could hear as people do. Before joining Interval, Watts developed ICs and software for satellite-communications systems, telephony systems, optical-character-recognition systems, and LCDs for Microtel Pacific Research, Synaptics, and Arithmos. He also invented a low-delay digital-speech-coding algorithm that was sold to Cisco in 1999. Watts is the author of three issued and seven pending patents. He holds a doctorate from the California Institute of Technology (Pasadena, CA), a master's degree in digital-speech coding from Simon Fraser University (Burnaby, BC, Canada), and a bachelor's degree in engineering physics from Queen's University (Kingston, ON, Canada).

ELECTRONIC COMPONENTS

NAND flash | Solid State Drives | removable flash storage | embedded NAND solutions | multi-chip packages

Don't fall behind. Choose Toshiba NAND technology to enhance your leading designs.

Toshiba's reliable, high-performance storage solutions enable design success by delivering optimum performance for mobile applications. As the world's most experienced producer of NAND flash, Toshiba has the commitment and capability to deliver advanced NAND solutions, including Solid State Drives (SSDs), MCPs with configurable MLC and SLC areas, and embedded NAND with MMC, NAND and SD interfaces. Rely on Toshiba for leading NAND technology and finish ahead of the competition.

Toshiba NAND Applications:
- Laptops, UMPCs, netbooks
- Mobile phones
- Digital audio/MP3 players/portable media players
- Mobile gaming and GPS devices
- Digital video camcorders
- Mobile peripheral applications
- and more

www.memory.toshiba.com

BY IAN DOBSON AND JIM HOLBROOK • INTEGRATED DEVICE TECHNOLOGY INC

PCIe clock distribution in embedded systems

AS PCIe RAPIDLY BECOMES THE I/O INTERFACE OF CHOICE FOR HIGH-PERFORMANCE SYSTEMS, CLOCK-DISTRIBUTION AND JITTER-REDUCTION REQUIREMENTS CREATE CHALLENGES FOR THE EMBEDDED-SYSTEM-DESIGN TEAM.

The PCIe (peripheral-component-interconnect express) protocol is highly desirable for communication across backplanes in embedded and other system types. However, for an embedded-system environment in which backplane connector pins are often at a premium, PCIe's preferred clock-distribution scheme—using a star configuration of point-to-point connections—is less than ideal. You can distribute a PCIe-compatible clock using a single multidrop signal and still meet the tight jitter requirements of the PCIe Generation 2 specification.

CLOCKING IN PCIe

PCIe Base Specifications 1.1 and 2.0 define three clock-distribution models for the 2.5- and 5-Gbps signaling rates (**figures 1** through **3**). The common-clock architecture is the common method for a variety of reasons. First, most of the commercially available chips supporting PCIe interfaces use only this architecture. Second, this architecture is the only one that directly supports spread-spectrum clocking, which can be important in reducing EMI (electromagnetic-interference) peaking and, hence, simplifies the task of meeting electromagnetic emissions limits for the system (**Figure 4**). Finally, this architecture is the simplest to conceptualize and design.

The most significant disadvantage of the common-clock architecture is the need to distribute the reference clock to each PCIe endpoint in the system. The clock's 100- or 125-MHz frequency and the PCIe protocol's tight jitter requirements further complicate this task. For 2.5-Gbps operation, the limit is 86-psec p-p phase jitter for a sample set of 106 samples. The 5-Gbps operational limit is 3.1-psec-rms jitter. However, to operate at 5 Gbps, a transceiver first negotiates at 2.5 Gbps and then moves up to the higher rate if both ends can do so. That is, if the system supports any 5-Gbps links, then the reference clock must meet both jitter specifications.

The separate- and data-clock architectures avoid these limitations but substantially increase the complexity of the clock-system design and don't support spread-spectrum clocking without the use of sideband signaling. The governing specifications for reference-clock jitter are PCIe Base Specifications 1.1 and 2.0, and PCIe Jitter-Modeling Revision 1.0D and PCIe Jitter and BER (bit-error-rate) Revision 1.0 detail the method for verifying jitter compliance. The electromechanical specifications provide mechanical-form-factor information, electrical-signal definition, and functional definitions. Some of these specifications, such as Card-Electromechanical Specifications 1.1 and 2.0, also provide jitter budgeting among the reference clock, transmitting PLL (phase-locked loop), receiving PLL, and media. Strictly speaking, the Card-Electromechanical Specification applies only to PC-, server ATX- (advanced-technology-extended), and ATX-based form factors. Industry groups have published additional electromechanical specifications to cover other form factors, such as Mini-Card-Electromechanical Specification 1.2 for mobile-computing platforms.

For most embedded systems, these specifications provide guidelines that designers can use in whole or in part to specify the embedded system's PCIe clock-distribution scheme. For example, many of the Card-Electromechanical documents

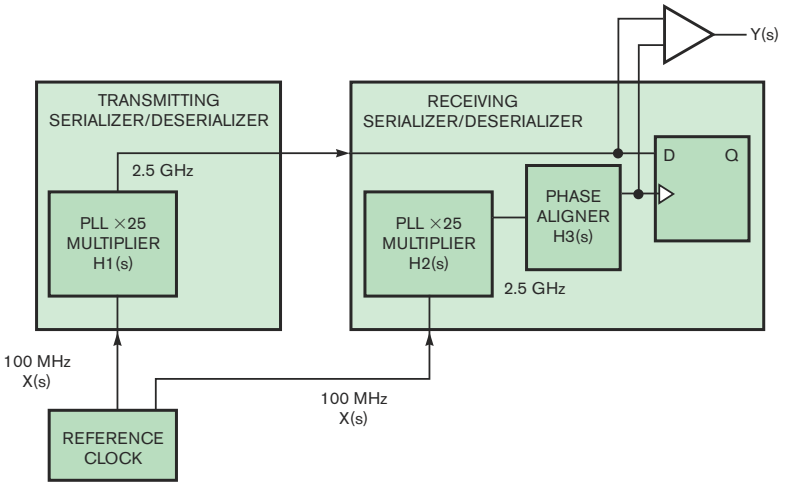

NOTES:
SYSTEM-TRANSFER FUNCTION IS H(s)=H3(s)×[H1(s)-H2(s)].
OUTPUT Y(s)=X(s)×H(s).

Figure 1 The common-clock architecture is one of three methods for clock distribution that the PCIe base specification defines.

specify the use of the HCSL (host-clock-signal-level) protocol for distributing the reference clock. However, many embedded systems use LVPECL (low-voltage-positive-emitter-coupled-logic) signaling or M-LVDS (multipoint-low-voltage-differential signaling) to achieve a greater reach, noise margin, or both on their clock-distribution network.

Many embedded systems distribute a large number of high-speed signals, including clocks, across their backplanes. To deal with the often-heavy electrical loading on those backplanes, these signals tend to have powerful drivers and, hence, high edge rates. This situation presents the danger of crosstalk and other signal-integrity problems, especially when the backplane has a lighter load than the worst-case design. Another related design challenge is that PCIe specifies reference clocks of 100 or 125 MHz, which are difficult to distribute cleanly over a long, heavily loaded backplane.

In addition to the PCIe specifications' tight jitter limits and need for a longer signal reach, the number of signals that can transit the backplane connectors and the backplane itself also constrain embedded systems. Defining the connector pinouts is one of the more critical tasks when specifying the system.

COMMON-CLOCK-DISTRIBUTION SCHEME

Due to the clock frequency and jitter constraints, most common-clock-architecture designs distribute their reference clocks using point-to-point differential-signaling pairs, one of which goes to every PCIe endpoint in the system. If your design has multiple PCIe endpoints on a single card, you can take in a reference-clock input from the backplane and provide a clock-distribution network on the card using zero-delay buffers. Even this task can be difficult to design, however, given the jitter constraints of 5-Gbps PCIe operation.

Assuming that you could design such on-card distribution schemes, they still require a point-to-point connection from the PCIe root to every card in the system. In embedded systems, this requirement adds a lot of connector pins to the root-card slots and a lot of traces with special routing requirements to the backplane. It also means that the slot that the root card plugs into has a different pinout from that of the other slots.

One approach to solving these problems is to divide the PCIe reference clock on the root card and distribute it across the backplane using a multidrop M-LVDS and then to multiply it to the desired frequency or frequencies on the destination cards. Although conceptually simple, this approach is tricky to achieve within the jitter constraints of PCIe (**Figure 5**).

This approach allows you to use an M-LVDS pair to drive or receive a PCIe-compliant reference clock. In many embedded systems, the cards operate as roots or endpoints depending on the application, the slot assignment, or both. A card that operates in only one of those modes would be simpler than the one in **Figure 5**. One card in the system would act as the root, generating a reference clock meeting the PCIe constraints from its onboard crystal. This clock would drive any onboard PCIe devices from an internal clock-distribution network. The clock would also go to a non-PLL divider circuit that would divide it from 100 or 125 MHz to the backplane frequency of 25 MHz. It would then drive the divided-down reference clock to the rest of the cards in the system. All the other cards in the system would disable the use of their onboard clock generators, tristate

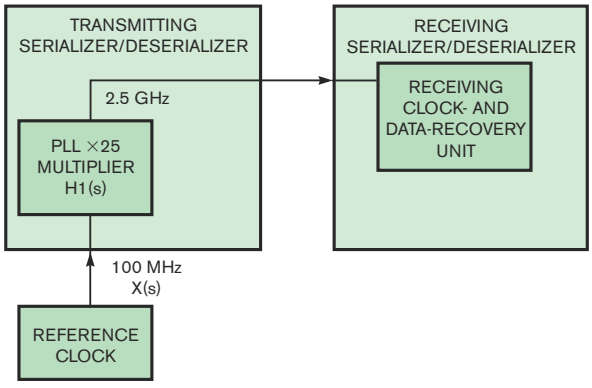

NOTES:
SYSTEM-TRANSFER FUNCTION IS H(s)=H1(s).
OUTPUT Y(s)=X(s)×H(s).

Figure 2 The PCIe data-clock architecture combines the reference clock with the transmitted data.

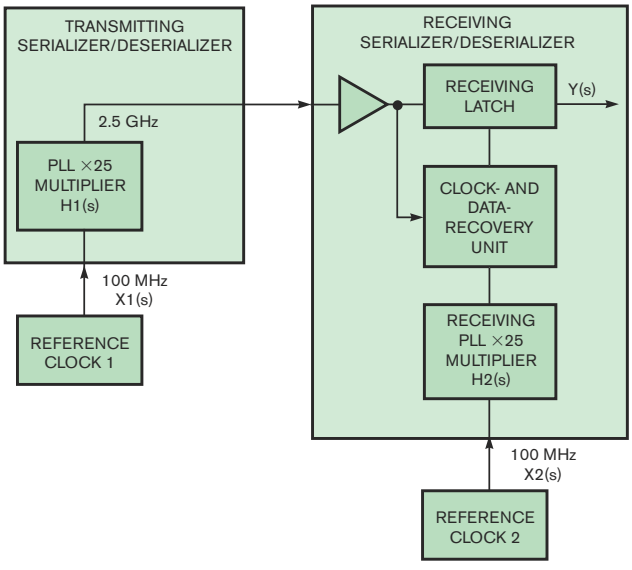

Figure 3 The PCIe separate-clock architecture employs a second reference clock in the receiver.

their drivers for the reference-clock traces, and receive the reference clock from the backplane. This clock would multiply using a PLL-based zero-delay buffer to the required onboard reference-clock frequency and then travel to the other cards. The circuitry that receives and multiplies the reference clock from the backplane would usually reside on the root card and could generate the second reference-clock frequency, if necessary. To achieve the low jitter that PCIe requires, you can incorporate jitter attenuators for the clock synthesizer and the zero-delay buffer.

One of the main challenges of a design such as this is that PLLs filter high-frequency jitter higher than their loop bandwidth but add jitter at modulation frequencies lower than their loop bandwidth. PLLs also induce tracking skew because they

do not perfectly track phase and frequency variations of the reference-clock input. For a backplane-PCIe implementation such as this one, which involves two or more cascaded PLLs for frequency generation and translation, you must take great care to minimize phase jitter and PLL-tracking skew.

PCIe-JITTER MEASUREMENT

Before diving into an analysis of the performance of this design, you must understand the process by which PCIe analyzes jitter performance. One of the overarching concerns of the PCIe Jitter Working Group was to neither overspecify nor underspecify the reference clock. To that end, the group accounted for the filtering effect of the transmitting and receiving PLLs and phase interpolator on the reference clock and the peaking effects of these PLLs.

Although the group has yet to detail many portions, the process now has four high-level steps. First, determine the accumulated phase error for each cycle. For serial-data transfer, the accumulated phase error is more important than cycle-to-cycle jitter or period jitter, which are important characteristics of parallel buses. Second, apply the DFT (discrete Fourier transform) to the accumulated phase-error data to change from time-domain to frequency-domain analysis. Then, apply the system-transfer function to the DFT of the accumulated phase-error data and perform an inverse DFT to transfer the filtered accumulated phase-error data back into the time domain.

You perform the filtering analysis of the PLL system in the complex frequency domain by setting $s=j\omega$ in the system-transfer functions. This equation works well for continuous systems, but most modern PLL implementations are not pure-analog systems because they have digital components, such as the phase detector and feedback divider; thus, Z-domain digital analysis is more accurate. However, brief studies by the PCIe Jitter Working Group showed that S-domain analysis imposes minimal error, so the group used S-domain analysis for modeling. The S-domain approximation deviates significantly from reality when the reference frequency is less than 10 times the PLL bandwidth, and designers must keep that fact in mind when selecting a PLL (**Reference 1**).

JITTER-MEASUREMENT TIPS

An improper measurement method can easily lead to jitter measurements that are twice as great or more than you would get using correct techniques. Here are a few tips:

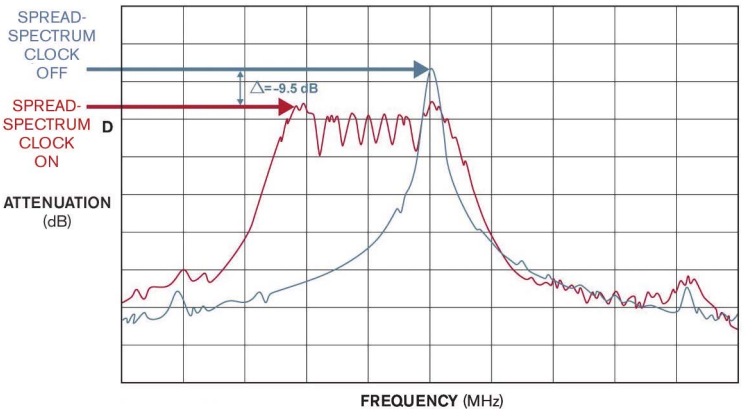

Figure 4 Spread-spectrum clocking can reduce EMI peaking to meet system-emissions limits.

Use shielded coaxial cables from the device under test to the oscilloscope, and terminate the clock to the oscilloscope input. If using high-impedance probes, use a low-capacitance probe and a ground clip rather than a wire. Use the highest possible sampling rate consistent with the required sample size. Maximize the vertical scale on the oscilloscope screen for accurate voltage measurements. Keep monitors, switching power sup-

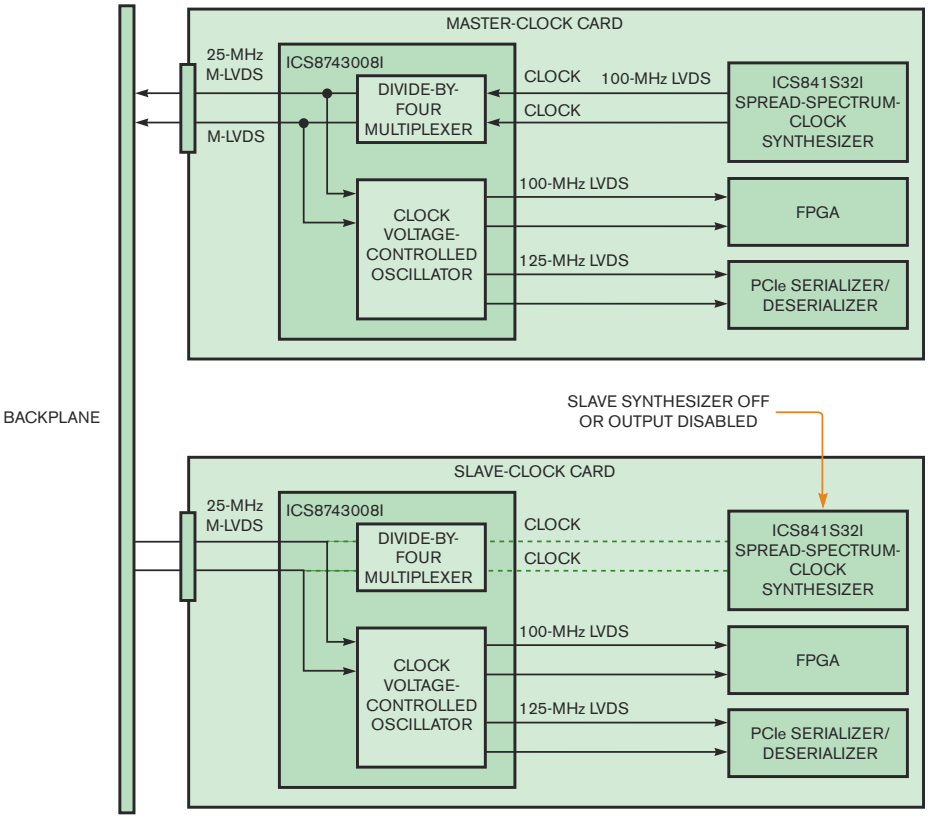

Figure 5 One technique for meeting the jitter constraints of 5-Gbps operation is to divide the reference clock, distribute it across the backplane, and restore the desired frequency at the destination.

plies, and cell phones away from the device under test. Use a linear power supply whenever feasible. When performing differential measurements, ensure that you have deskewed the two cables relative to one another.

ANALYSIS OF IDT'S SOLUTION

Engineers built a prototype of the circuit in **Figure 5** daisy-chaining an IDT (Integrated Device Technology, www.idt.com) ICS841S32I characterization board, an IDT ICS8743008I board, and a second ICS8743008I board representing the slave card. They took measurements at the output of the second ICS8743008I. They then offloaded the clock-period data from the oscilloscope and postprocessed the data with a jitter-analysis script, which performs the necessary frequency- and time-domain analysis (**Figure 6**).

The result for the 2.5-Gbps-analysis method is 18.91 psec. This result meets the PCIe peak-to-peak-phase-jitter spec of 86 psec with a factor-of-4.5 margin. For 5-Gbps operation, PCIe specifies rms phase jitter rather than peak-to-peak phase jitter. These results also exceeded specifications: 0.52-psec rms low-band jitter and 1.47-psec high-band jitter versus a 3.1-psec specification limit.

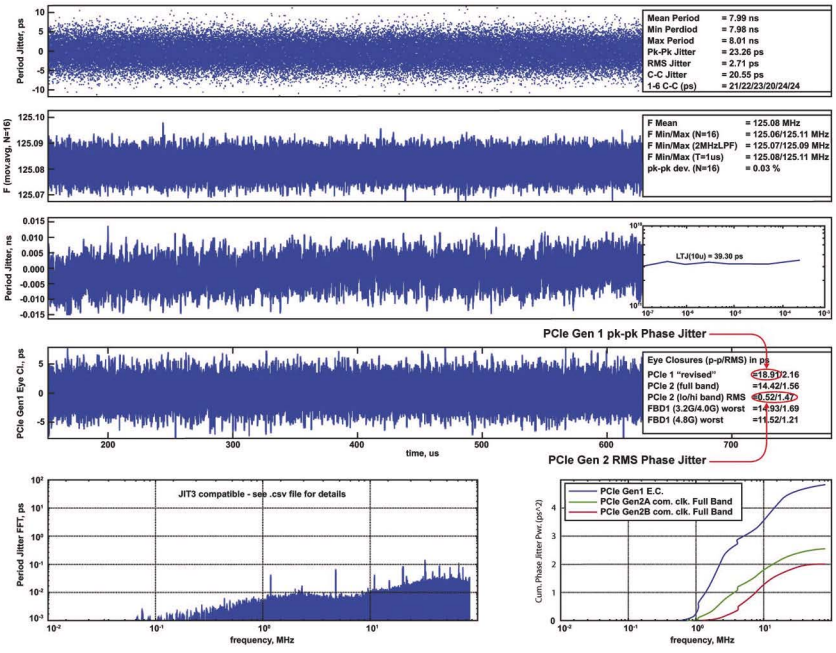

Figure 6 The results of the jitter analysis show that the clock-division strategy meets and exceeds the PCIe specifications for both 2.5- and 5-Gbps operation.

For 5-Gbps operation, PCIe specifies two transfer functions and two frequency ranges for analysis in the frequency domain. The pole frequencies for these transfer functions are 5 and 16 MHz for the first transfer function and 8 and 16 MHz for the second transfer function. The two frequency bands over which you analyze the jitter are 10 kHz to 1.5 MHz for the low band and 1.5 MHz to the Nyquist frequency for the high band. For the Nyquist frequency, you analyze as much as half of the reference-clock frequency. For example, for a 100-MHz frequency, the frequency-domain analysis would extend to 50 MHz. The script reports the worst case between the two transfer functions across each frequency-analysis band.

The originators of the PCIe standard defined it primarily for use in PC systems, but, due to its low pin count and scalable high performance, it is rapidly becoming the I/O interface of choice for components in almost all applications. The high speed of the reference clock that you must distribute, along with the option for two compliant reference-clock speeds, poses some challenges for embedded-system designers who want to use PCIe components.

One tested approach allows a system to use components supporting the 100- and 125-MHz reference-clock options and allows you to distribute this clock over an M-LVDS differential pair to all cards in the system. This approach also allows you to configure cards so that they can act as a root or an endpoint as the application dictates and can reside in any slot in the system. Furthermore, the approach lowers the operating frequency for the reference clock on the backplane, easing the routing constraints and crosstalk performance for that signal. EDN

Go to www.edn.com/ms4314 and click on Feedback Loop to post a comment on this article.

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

REFERENCE

1 "PCIe Reference Clock Requirements," Integrated Device Technology.

AUTHORS' BIOGRAPHIES

Ian Dobson is senior manager of system technology at Integrated Device Technology, where he has worked for five years. He identifies how IDT can best address customers' needs. Dobson holds a bachelor's degree in applied sciences from the University of Waterloo (Waterloo, ON, Canada). You can reach him at ian.dobson@idt.com.

Jim Holbrook is director of systems engineering in the communications division of Integrated Device Technology, where he has worked for seven years. He provides product and technology definition, strategic planning, and market analysis. Holbrook has 13 years' experience in the semiconductor industry. He holds a bachelor's degree in electrical engineering from the US Naval Academy (Annapolis, MD). He completed Navy nuclear-engineering training in 1988, followed by seven years in the US Nuclear Submarine Force. You can reach him at jim.holbrook@idt.com.

High Voltage, Precision, Battery Stack Monitor

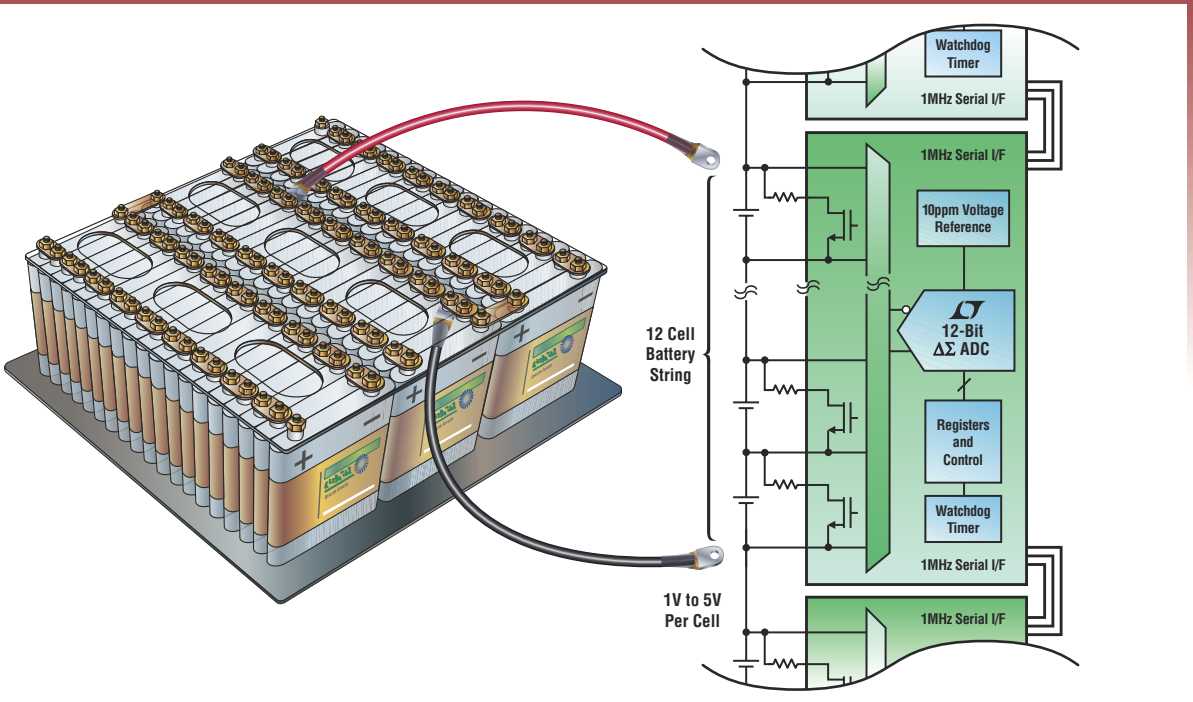

Rugged IC for Hybrid/Electric Vehicles & Battery Backup Systems

The LTC®6802 is a highly integrated multicell battery monitoring IC, capable of measuring up to 12 individual battery cell voltages. Using a novel serial data communication technique, multiple LTC6802s can be stacked in series without optocouplers or isolators, allowing precision voltage monitoring of every cell for 1000V+ systems. With high ESD, EMI and noise immunity, the LTC6802 stands up to real world conditions in a range of automotive and industrial applications.

Features

- 0.25% Maximum Total Measurement Error from –40°C to 85°C
- ΔΣ ADC with Inherent FIR Filtering
- Onboard FETs for Cell Discharge
- External Temperature Inputs
- Diagnostics and Fault Detection
- AEC-Q100
- 44-Lead SSOP Package
- Fully Specified for –40°C to 85°C

Cell Measurement Accuracy

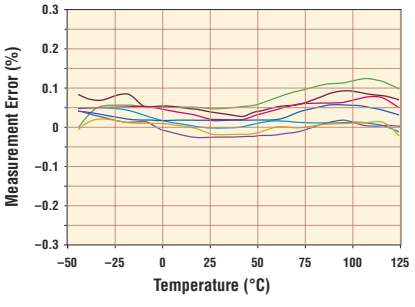

Typical devices show outstanding accuracy over extended temperature extremes.

Info & Free Samples

www.linear.com/6802

1-800-4-LINEAR

Free Automotive Electronics Solutions Brochure

www.linear.com/autosolutions

LINEAR TECHNOLOGY

EDITED BY MARTIN ROWE
AND FRAN GRANVILLE

designideas

READERS SOLVE DESIGN PROBLEMS

Digital controller compensates analog controller

David Caldwell, Flextek Electronics, Carlsbad, CA

Emerging digital ICs for power control lack basic features, such as the built-in gate drive and current limiting, that you would normally find in analog ICs. Digital-power controllers generally have only PWM (pulse-width-modulated)-logic output, and discrete gate drivers rarely include current limiting. In addition, most protected FETs work only in low-frequency, low-side applications.

The LM3485 IC from National Semiconductor (www.national.com) includes high-side gate drive with current limiting (**Reference 1**). However, the hysteretic-control scheme of this analog IC is likely to yield questionable performance in some applications due to variable switching frequency and overshoot, as well as an inability to regulate feedback below the 1.24V reference. A traditional PID (proportional-integral-differential)-control scheme can get around these limitations but adds considerable complexity.

The CLZD010 CLOZD (Caldwell-

DIs Inside

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

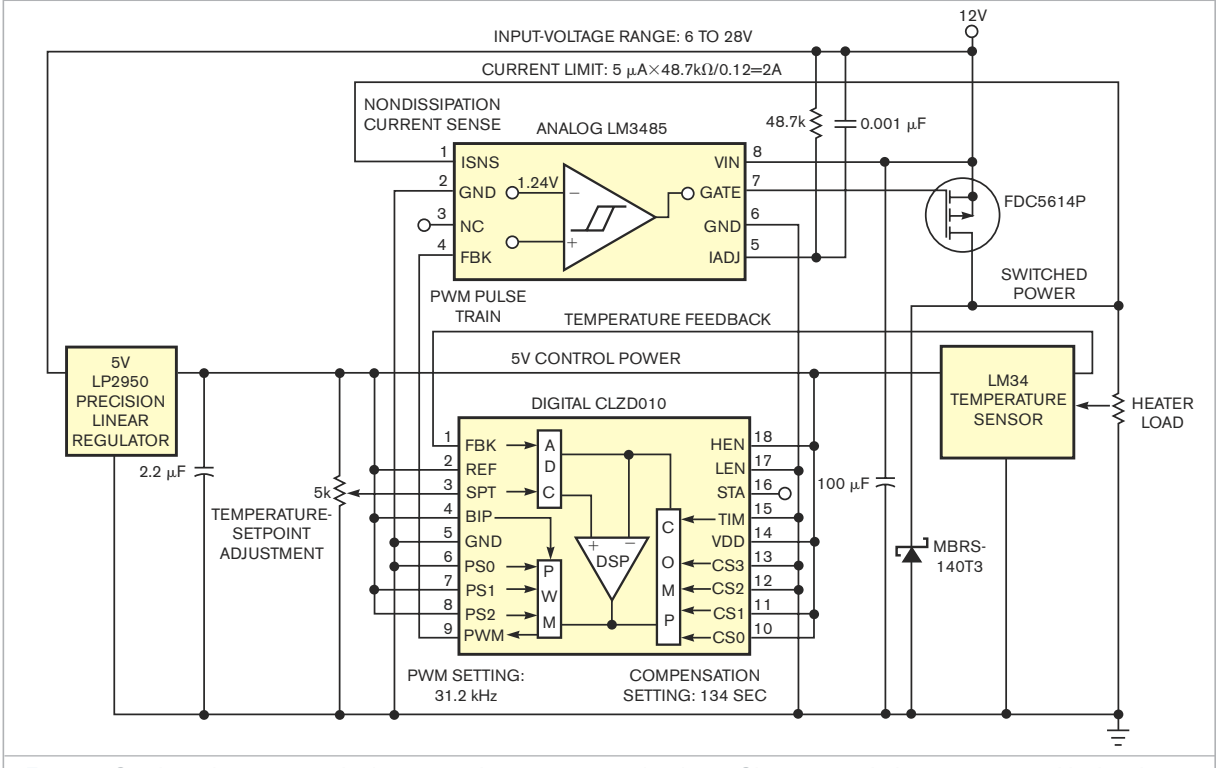

Figure 1 Combine the simple and robust closed-loop control of the digital CLZD010 with the current-limited high-side gate drive of the analog LM3485 for the best of both worlds.

LINEAR TECHNOLOGY

DESIGN NOTES

Triple Buck Regulator Features 1-Wire Dynamically Programmable Output Voltages

Design Note 459

Andy Bishop

Introduction

The LTC®3569 is a compact power solution for handheld devices. Its tiny 3mm × 3mm QFN package includes three buck regulators with individually programmable output voltages. One regulator supports load currents up to 1200mA, while the other two support currents to 600mA. Two regulators can be paralleled for increased load capability. Each current-mode regulator is internally compensated with excellent load and line regulation. A complete 2- or 3-output solution requires a minimum number of external passive components.

Three Individually Programmable Bucks

The LTC3569's three output voltages are independently programmed by simply toggling their respective enable pins. Each time an enable pin sees a falling edge, a 4-bit counter is decremented. After a time-out delay of 120µs from the last rising edge at the enable pins, the counter state is latched into the feedback reference voltage DAC. In this way, the reference voltage can be programmed from 800mV (full scale) to 425mV in 25mV steps.

Configure Parallel Power Stages for Different Loads

The LTC3569's buck regulators can be paralleled for higher load capability. By pulling the feedback pin of one of the two 600mA regulators up to the input supply voltage, that regulator's power stage is reconfigured as a slave, where switching is synchronized to its upstream master. Buck 2 can be a slave to Buck 1, or Buck 3 a slave to Buck 2. When operating in slave mode, the slave switch pin is tied in parallel with the master switch pin and the maximum output currents sum. This yields three possible combinations: three independent regulators (a 1.2A buck and two 0.6A bucks), two independent 1.2A bucks, or two independent bucks of 1.8A and 0.6A.

Power Good Indicator

The LTC3569 has a PGOOD pin to indicate when any enabled regulator output voltage has risen to within 8% of the programmed value. If any of the enabled output voltages are lower than programmed, the PGOOD pin pulls low. If all of the regulators are off, the PGOOD pin pulls low and the LTC3569 enters a low power shutdown mode with <1µA of supply current.

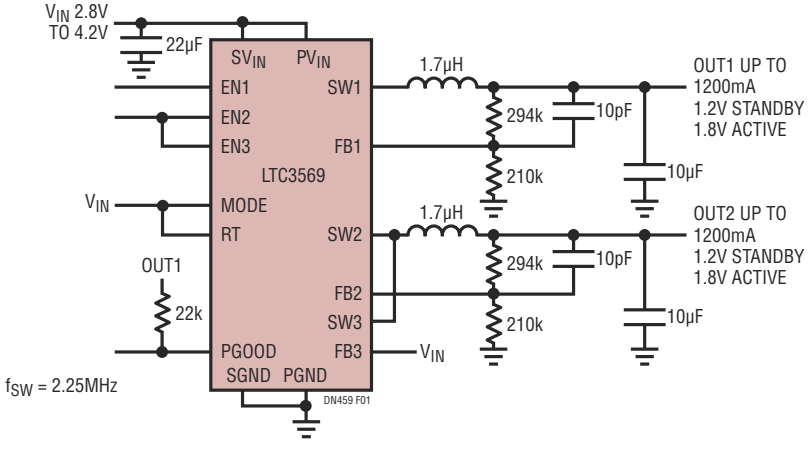

Figure 1. LTC3569 Configured As Dual 1.2A Programmable Buck Regulators

Power Saving Operating Modes

The LTC3569 offers two modes of operation (set via the MODE pin) that improve efficiency at light loads. Burst Mode® operation is the most efficient at low load currents, while pulse-skipping mode produces lower ripple currents. At startup, until the end of the soft-start ramp, pulse-skipping mode is automatically selected.

Programmable Clock Frequency

The switching frequency is fixed at 2.25MHz by pulling RT up to the input supply, or the clock can be programmed to a frequency between 1MHz and 3MHz with a timing resistor to ground. If a clock signal is applied to the MODE pin the LTC3569's clock is injection locked to the external clock as long as the frequency is greater than that programmed using the RT pin. With injection locking, the operating mode is automatically set to pulse-skipping.

2-Output, Individually Programmable 1.2A Regulators

Figure 1 shows a 2-output application where each output can be reprogrammed at any time to a standby voltage of 1.2V or an active voltage of 1.8V. Both outputs provide up to 1.2A of load current from a Li-Ion battery voltage between 2.8V and 4.2V. Burst Mode operation is selected for high efficiency at light loads. Figures 2 through 5 show independent programming of the two output voltages via toggling of the respective enable pins while supplying a constant 625mA to each load.

Conclusion

The LTC3569 is a flexible solution for powering handheld Li-Ion battery applications. The ability to adjust or disable individual output voltages on the fly provides a simple solution to support energy saving operating modes in advanced microprocessor-based designs.

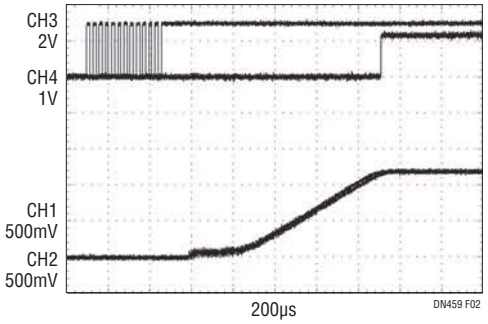

Figure 2. Soft-Start Both Bucks Into Standby. CH1 = OUT1, CH2 = OUT2, CH3 = EN1 = EN2, CH4 = PGOOD

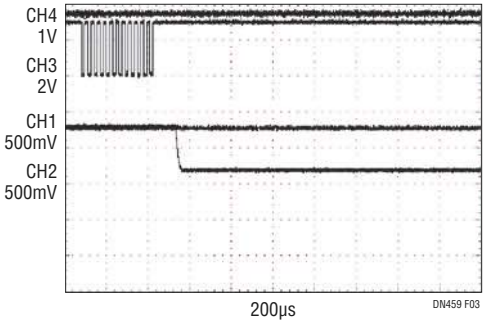

Figure 3. Reprogram Buck 2 From Active to Standby With No Cross-Talk On Buck 1 Output. CH1 = OUT1, CH2 = OUT2, CH3 = EN2, CH4 = PGOOD

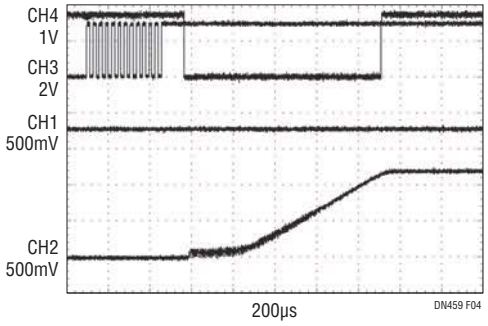

Figure 4. Buck 1 Active, Buck 2 Soft-Start to Standby. CH1 = OUT1, CH2 = OUT2, CH3 = EN2, CH4 = PGOOD, No Cross-Talk On Buck 1 Output

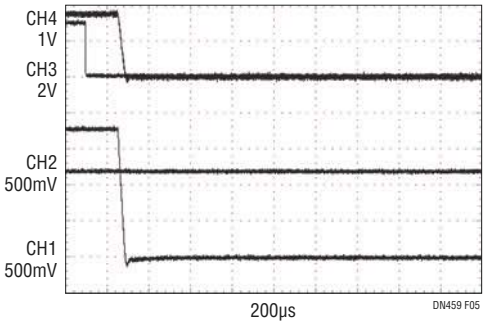

Figure 5. Buck 1 Active to Shutdown, Buck 2 Standby. CH1 = OUT1, CH2 = OUT2, CH3 = EN2, CH4 = PGOOD, Note PGOOD Falls As It Is Tied to OUT1, No Cross-Talk On Buck 2 Output

Data Sheet Download

www.linear.com

For applications help,
call (408) 432-1900, Ext. 3161

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

loop-optimization-in-Z-domain) controller-chip IC from Flextek Electronics (www.flex-tek.com) both broadens and simplifies control applications though the embedded intelligence of a digital device (**Reference 2**). A single time-domain compensator replaces the three frequency-domain PID parameters, eliminating complex stability analysis. The circuit requires no PC interface because you inspect the open-loop response and then use pin settings to configure the closed-loop compensation. However, the PWM output is only a logic-level driver.

Combine the simple and robust closed-loop control of the digital CLZD010 with the current-limited high-side gate drive of the analog LM3485 for the best of both worlds (**Figure 1**). The PWM-logic level of the digital IC overrides the hysteretic comparator of the analog IC to switch the FET. A second comparator at ISNS, Pin 1 in the LM3485, turns off the FET if the voltage across it exceeds a predetermined value during conduction to limit current.

In the thermal-response example (**Figure 2a**), the circuit takes about three minutes for the open-loop temperature to reach roughly two-thirds of its final value, so closed-loop compensation, at 134 seconds, is slightly faster in **Figure 1**. The resultant closed-loop temperature quickly nears its final value due to maximum drive; voltage then decreases to allow the temperature to settle at the setpoint without overshoot (**Figure 2b**). You can use this basic circuit combination to satisfy a broad range of applications in multiple industries. EDN

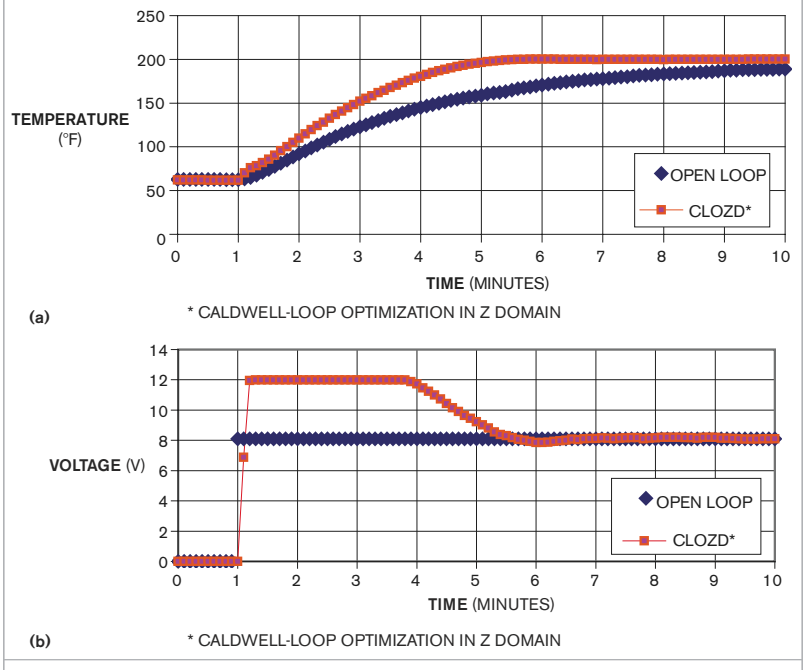

Figure 2 In the thermal-response example (a), the circuit takes about three minutes for the open-loop temperature to reach roughly two-thirds of its final value. The resultant closed-loop temperature quickly nears its final value due to maximum drive; voltage then decreases to allow the temperature to settle at the setpoint without overshoot (b).

REFERENCES

1 "LM3485 Hysteretic PFET Buck Controller," National Semiconductor, September 2004, www.national.com/ds/LM/LM3485.pdf.
2 "CLOZD Loop Controller Chip, Part CLZD010," Flextek Electronics, 2004, http://flex-tek.com/CLZD010.pdf.

Circuit provides constant-current load for testing batteries

Vladimir Rentyuk, Modul-98 Ltd, Zaporozhye, Ukraine

Suppose that you need to test a 1.5V, AA-size alkaline battery. You can apply a short circuit and measure current, or you can measure open-circuit voltage, but neither method properly tests the battery. A suitable test current of approximately 250 mA gives you a more reasonable test. You can use a 6Ω resistive load at 1.5V, which produces an output voltage of 1.46V at an ambient temperature of 25°C if the battery is in excellent condition. A poor battery might produce less than 1.2V. Given the load, the output current at 1.2V will be 200 mA instead of 250 mA. The battery will have just 80% of a full load current. Instead, you can use the circuit in **Figure 1** to produce a constant-current load.

The circuit uses a 9V battery and a voltage regulator to produce a steady power-supply voltage of 5V. From that voltage, the circuit produces a constant sink current, which is independent of the battery's output voltage, using IC_1, IC_2, and Q_3. Your choice of current depends on battery size. You calculate the sink current of this circuit as $I_{TEST} = 1/R_{19} \times [V_{CC} \times R_{18}/(R_4+R_{18})]$, where I_{TEST} is the current you are testing and V_{CC} is the voltage of resistive divider R_4 and R_{18}.The voltage across R_{19} should range from 0.3 to 0.85V for AAA and AA batteries. Transistor Q_3 should be in its active region. Resistor R_{14} limits Q_3's base current to a safe level.

A suitable choice for the operational amplifier, IC_2, is also important. You should use a single-supply op amp with a rail-to-rail input and a rail-to-rail output, such as Analog Devices' (www.

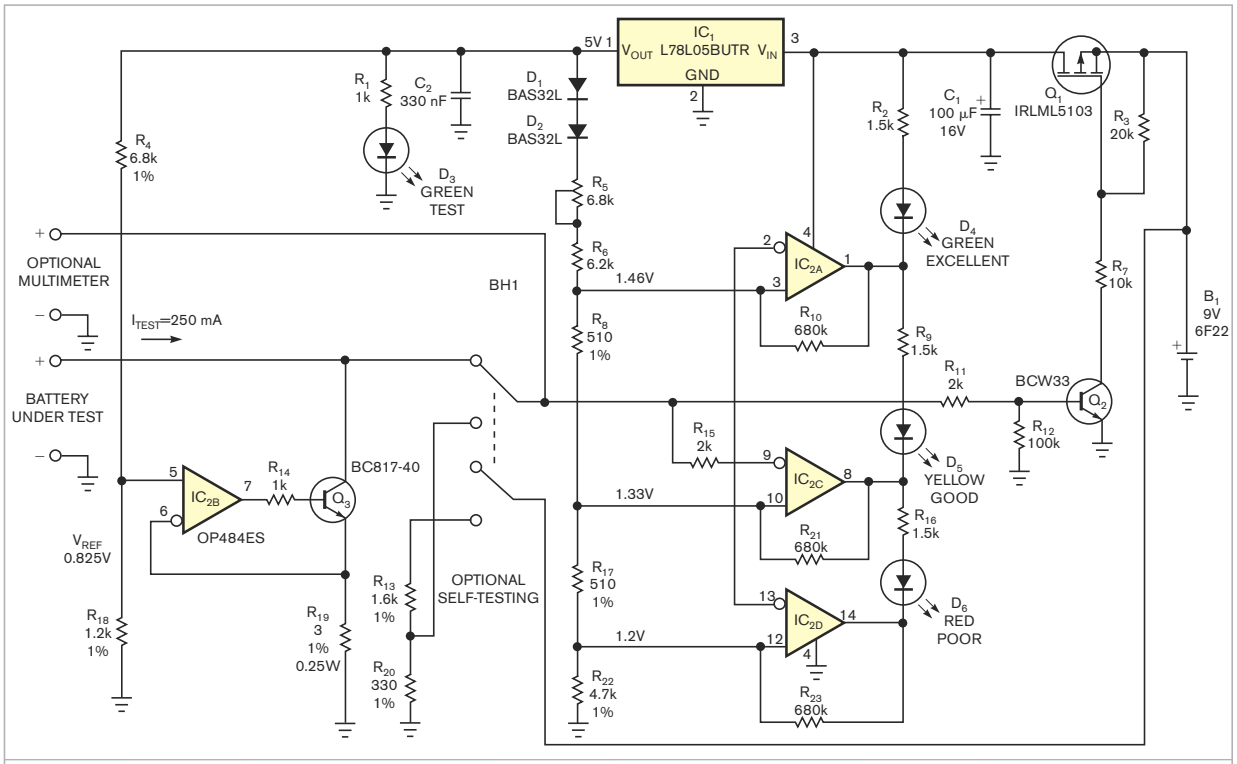

Figure 1 A tester of AA- or AAA-sized batteries uses constant-current load.

TABLE 1 VOLTAGE RANGES FOR LEDs

Condition	Battery voltage[1] (V)	D_2	D_4	D_5	D_6	D_3
Excellent	>1.46	Yes	Yes	No	No	Yes
Good	>1.33	Yes	No	Yes	No	Yes
Poor	>1.2	Yes	No	No	Yes	Yes
Bad	>1[2]	Yes	No	No	No	Yes
Unable to test	>12	No	No	No	No	No

[1]Ambient temperature is 25°C.
[2]This estimated value can be less.

analog.com) OP484ES or OP496GS.

When you connect the battery under test, Q_2 turns on, which then turns on Q_1, applying voltage from the 9V battery to the regulator. That action lights D_3, indicating that the battery under test has enough voltage to be tested. LEDs D_4, D_5, and D_6 indicate the battery's condition. **Table 1** shows the voltage ranges necessary for these LEDs to light.

Op amps IC_{2A}, IC_{2C}, and IC_{2D} work as comparators with some hysteresis for operational stability. The resistive divider comprising R_5, R_6, R_8, R_{17}, and R_{22} sets the voltage levels. Diodes D_1 and D_2 are optional but are useful when you need to operate the circuit outdoors, where temperatures vary widely. Resistor R_{15} protects the inputs of IC_{2A}, IC_{2C}, and IC_{2D}.

When you connect a battery to test, you should test it for at least 5 seconds. LED D_3 shines if the battery is in relatively normal condition. In this case, switch Q_1 applies power to the battery tester. The sink-current generator comprising IC_{2A} and Q_3 loads the battery under test, and the resistor-divider network sets the comparator voltages.

You can add an optional self-testing button for checking the 9V battery to ensure that it has enough voltage to drive the circuit. You can also connect a digital multimeter to the multimeter terminals if you need a more accurate measurement. You can use a suitable rotary switch or a variable resistor and change the value of the test current by changing the value of R_4 to test another type or size of battery. EDN

MOSFET-based, analog circuit calculates square root

Abhirup Lahiri, Netaji Subhas Institute of Technology, New Delhi, India

Square-root-calculating circuits find wide use in instrumentation and measurement systems for such tasks as calculating the rms (root-mean-square) value of an arbitrary waveform, for example. Hence, designers need an effective analog square-root calculator. Because manufacturers do much of the IC fabrication in MOS technology, a MOSFET-based, analog square-root calculator seems appropri-

First line-card timing IC to enable 1G and 10G Synchronous Ethernet

Complete single-IC solution for traditional and next-generation line cards

The DS3104-SE is a low-cost, feature-rich, line-card IC that supports all traditional SONET/SDH clocks and includes additional clock rates that enable Synchronous Ethernet architectures. Gigabit Ethernet (GbE), 10-Gigabit Ethernet (10GbE), and Fast Ethernet xMII clock rates are all supported, allowing clock synchronization at the physical layer per ITU-T Recommendations G.8261 and G.8262.

Key features

- Hitless reference switching
- Holdover on loss of all inputs
- Programmable 1Hz to 400Hz PLL bandwidth
- Frequency conversion between SONET/SDH rates and Ethernet MII/GMII/XGMII rates

TWO COMPLETELY INDEPENDENT DPLLs

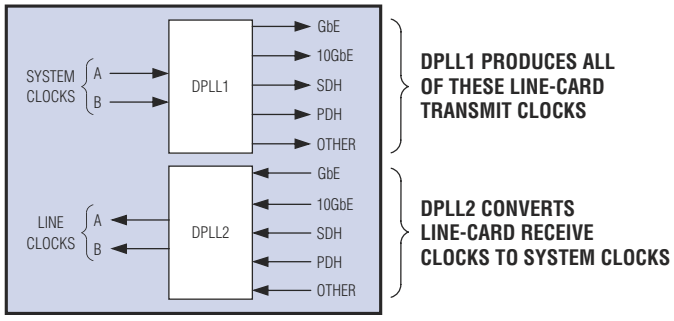

Supports More Industry-Standard Clock Rates than Any Other Single IC

- Ethernet xMII
- SONET/SDH
- PDH
- Base station
- CMTS
- Wireless, GPS, and more

Feature	DS3104-SE	Competitor A	Competitor B
Clock Inputs	**8**	8	3
Clock Outputs	**14**	6	2
Two Independent DPLLs	✓	✓	
Locks to 25MHz MII Clock	✓	✓	✓
Synthesizes 25MHz MII Clock	✓	✓	
Locks to 125MHz GMII Clock	✓		✓
Synthesizes 125MHz GMII Clock	✓		
Locks to/Synthesizes 156.25MHz XGMII Clock	✓		
Synthesizes SONET/SDH, MII, GMII, and XGMII Rates at the Same Time	✓		
Custom Frequency Synthesis (Hz)	**n x 2k to 77.76M, n x 8k to 311.04M**	n x 8k to 100M	—

MAXIM DIRECT
www.maxim-ic.com/shop

AVNET electronics marketing
www.avnet.com

MAXIM
www.maxim-ic.com

For free samples or technical support, visit our website or call 1-800-998-8800.

ate. This Design Idea describes such a circuit, which uses only MOSFETs to provide the square-root function (**Figure 1**). The design is simple and versatile and can provide the output as the square root of the difference of two voltages.

The circuit uses the nested connection of MOSFETs Q_1 and Q_2. Q_2 works in the saturation region as it is diode-connected, forcing Q_1 to work in the triode region. All other MOSFETs work in the triode region. The first part of the circuit, comprising Q_3, Q_4, Q_5, and Q_6, creating the current I_{O1}, is basically a MOS-resistive circuit. The essential **equation** governing the circuit operation is:

$$V_O = \left(\sqrt{\frac{1}{K_2}} - \sqrt{\frac{1}{K_1}}\right)\sqrt{I_{O1}},$$

where K_1 and K_2 represent the aspect ratios of transistors Q_1 and Q_2, respectively: $K_1=(\mu C_{OX}W)/2L_1$, where $I=K_1=K_2$. The MOSFETs creating the MOS-resistive circuit and hence responsible for the current creation are identical, having the same aspect ratio and threshold voltage. The current relates to inputs V_1 and V_2, as the following **equation** shows: $I_{O1}=G(V_1-V_2)$, where $G=2K(V_A-V_B)$ and represents the conductance of the MOS-resistive circuit—$k=(\mu C_{OX}W)/2L$—of the identical transistors forming the MOS-resistive circuit, and V_A and V_B are control voltages applied to the gate of the MOSFETs that are working in the triode. This approach provides the advantage of voltage controllability of the output; hence, the square-rooting function is voltage-controllable.

The following **equation** gives the output voltage:

$$V_O = 2\left(\sqrt{\frac{1}{K_1}+\frac{1}{K_2}} - \sqrt{\frac{1}{K_2}}\right)\times$$

$$\sqrt{K(V_A-V_B)(V_1-V_2)}.$$

It is evident from this equation that the output voltage, V_O, is the square root of the difference of input voltages V_1 and V_2. If you ground V_2, then the output voltage is proportional to the square root of input voltage V_1. As noted, control voltages V_A and V_B can vary the proportionality constant. Hence, you have devised a new all-MOSFET-based, voltage-controllable analog square-root calculator.

You can test the circuit using a variety of commercially available MOSFETs, such as the 2SK1228, which is available from many sources; the buffer can be a MOSFET-based op-amp buffer, such as the BUF04701 from Texas Instruments (www.ti.com). For the operation of the circuit to be in accordance with the output-voltage **equation**, the four MOSFETs you use to create the MOS-resistive circuit should be identical and should work in the triode region, for which inputs V_1 and V_2 should be less than V_A-V_{TH} and V_B-V_{TH}, respectively. The MOSFETs in the current mirror, Q_7 and Q_8, should be identical, and the diode-connected MOSFETs, Q_1 and Q_2, should be different and have different aspect ratios. You can test the circuit onboard using commercially available ICs, or you can simulate it on a computer using any standard version of Spice. The supply voltage must be in accordance with the selected components. EDN

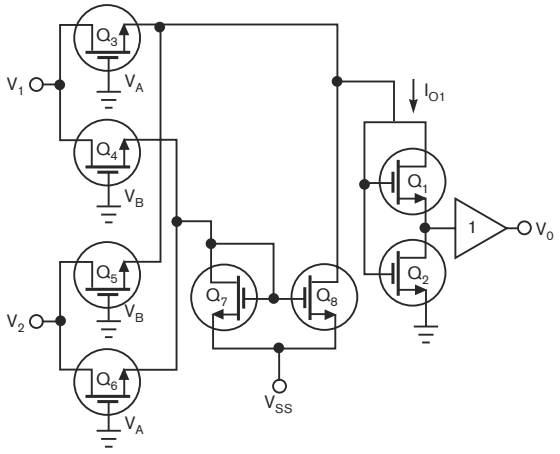

Figure 1 This circuit uses only MOSFETs to provide the square-rooting function.

"Hippasian" nonlinear VFC stretches dynamic range

W Stephen Woodward, Chapel Hill, NC

Hippasus of Metapontum was a Greek philosopher who lived approximately 500 BC. A disciple of Pythagoras, Hippasus discovered some interesting properties of square roots. This Design Idea describes a VFC (voltage-to-frequency converter) that also uses an interesting property of square roots: their ability to extend VFC dynamic range by orders of magnitude (**Figure 1**).

Linear VFCs are one of the oldest types of ADCs, and their simplicity and noise rejection preserve their popularity. However, their Achilles' heel is the direct proportionality between dynamic range and conversion time. Because the voltage resolution of linear VFC conversion is equal to the full-scale voltage reference, V_{REF}, divided by full-scale frequency, f_{FS}, multiplied by the counting interval, large dynamic range is inevitably associated with long counting intervals and slow conversion, even when clever VFC design provides for fast full-scale frequency.

For example: If you use a 3-MHz VFC-based ADC, such as Analog Devices' (www.analog.com) AD7742 with a 2.5V reference voltage in a design that requires 1-mV resolution, then

Tiny Hi-Speed USB/audio/video switches enable single connector

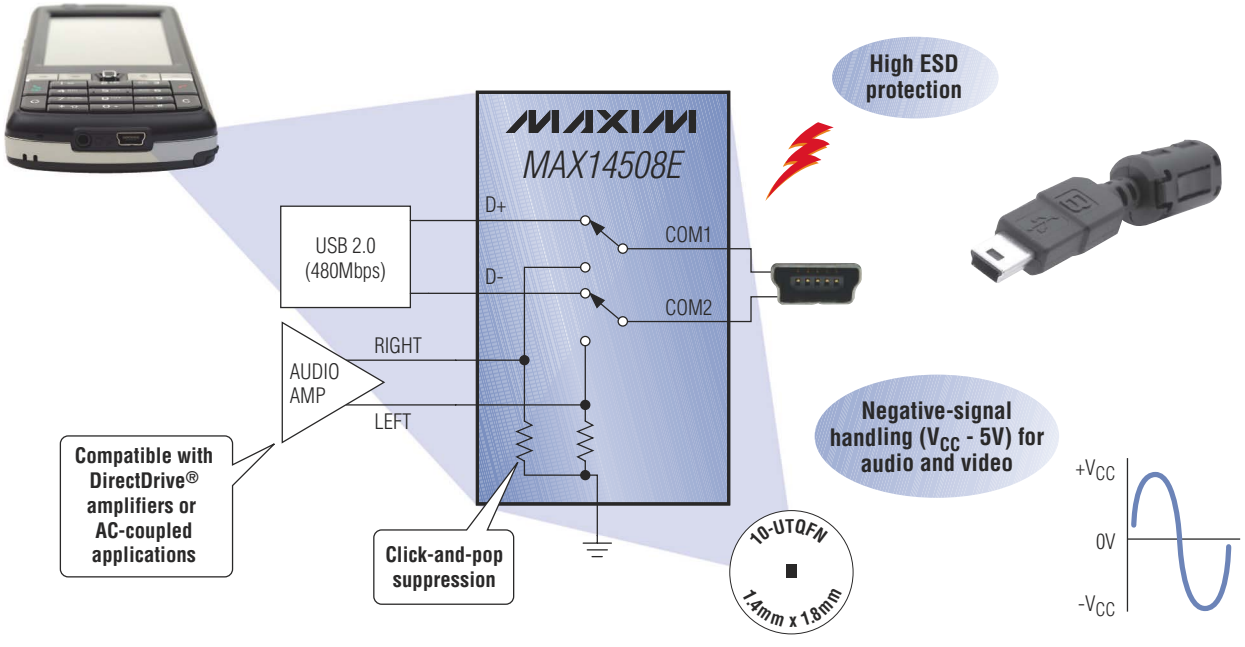

Optimize Performance

- Handle Hi-Speed USB (480Mbps) signals
- Low 0.05% THD+N (audio channel)
- High 950MHz bandwidth (video)

Save board space

- Allow interfacing to small USB connectors

Part	Configuration	Supply Voltage (V)	V_{BUS} Detection/Enable	Fault Protection	Shunt Resistors
MAX14508E	DPDT	+2.7 to +5.0	Enable	✓	✓
MAX14509E*			Enable		
MAX14509AE			V_{BUS}		✓
MAX14510E*			V_{BUS}		

EV kit available

DirectDrive is a registered trademark of Maxim Integrated Products, Inc
*Future product—contact factory for availability.

MAXIM DIRECT
www.maxim-ic.com/shop

AVNET electronics marketing
www.avnet.com

MAXIM
www.maxim-ic.com/MAX14508E-info

For free samples or technical support, visit our website or call 1-800-998-8800.

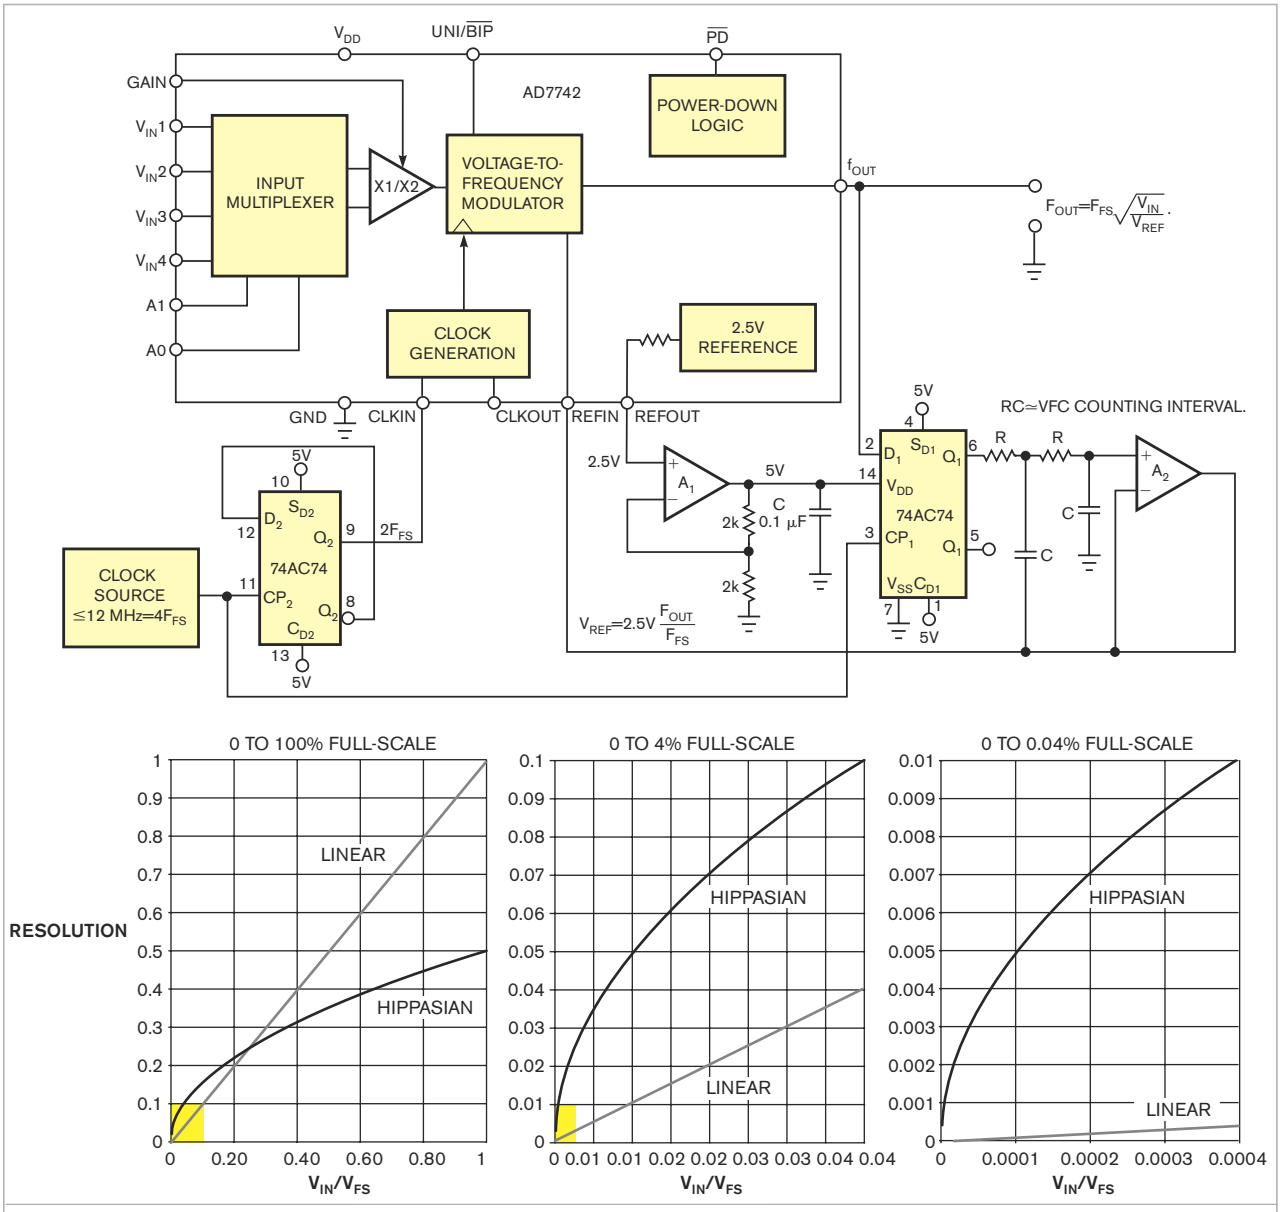

Figure 1 This nonlinear, wide-dynamic-range voltage-to-frequency converter exhibits 25-times improvement in counting time over other approaches.

you would need a minimum counting interval of 2.5/1 mV/3 MHz=2500/3 MHz=833 µsec. That counting interval yields only 1200 conversions per second, which for many applications is inconveniently slow.

The "Hippasian" VFC avoids this problem with a semiparabolic-transfer function instead of a linear one. It works by substituting V_{REF2}, which, instead of the constant V_{REF} of a linear VFC, is proportional to the output frequency. Then, $V_{REF2}=V_{REF}\times f_{OUT}/f_{FS}$, $f_{OUT}=V_{IN}\times f_{FS}/V_{REF2}=V_{IN}\times f_{FS}/(V_{REF}\times f_{OUT}/f_{FS})$, $(f_{OUT}/f_{FS})^2=V_{IN}/V_{REF}$, and $f_{OUT}=f_{FS}\times(V_{IN}/V_{REF})^{1/2}$.

Generating the dynamic, output-frequency-proportional reference voltage is the job of op amps A_1, which boosts the VFC's internal 2.5V reference to power flip-flop Q_1, and Q_1 and A_2, which compose a high-performance frequency-to-voltage converter. The accuracy of the reference voltage depends on precise 50-to-50 symmetry of the VFC's input-clock reference. Flip-flop Q_2 guarantees this symmetry.

The effect on conversion resolution of low-level signals is dramatic. To return to the example of a 2.5V full-scale, 1-mV-conversion resolution, which requires a 2500-count, 833-µsec conversion interval with a linear 3-MHz VFC, the Hippasian version needs only 100 counts and 33 µsec—a 25-fold improvement. Software linearization of the Hippasian VFC conversion is easy, requiring only one multiplication. EDN

Decoder lights the way

Jean-Bernard Guiot, Mulhouse, France

To display the status of two digital outputs, you can simply connect an LED and its resistor on each output. You must, however, interpret, or "decode," the displayed binary code. In addition, when no LED lights, users have no way of knowing whether it means that both outputs are off, that the power is off, or that a malfunction has occurred. In some applications, including industrial and medical settings, an indicator sending an ambiguous signal would be unacceptable. This Design Idea describes a simple circuit that resolves this problem by displaying four states on four LEDs (**Figure 1**). The operator need not understand binary coding, and, if no or more than one LED lights, it can mean only "no power" or "default."

The circuit works in the following way: If both inputs A and B are low, Q_1 allows current to pass through D_1 and resistor R_2 to A; only D_1 will light. Symmetrically, if both inputs A and B are high, Q_2 passes, and the current can pass from A through R_4, Q_2, and D_4; only D_4 will light. If both inputs are on different levels, only D_2 or D_3 will light. **Table 1** shows the possibilities; all other displays point to a default, such as a bad connection, a no-power condition, or a malfunction.

A totem-pole output that can sink and source the current for one LED must drive the A and B inputs. Resistors R_2, R_3, and R_4 are for applying a 12-mA LED current if the power supply is 5V. No component is critical. For example, you can use generic transistors, such as the NPN 2N3904 and the PNP 2N3906. You can even use transistors with integrated base resistors, further reducing the component count. EDN

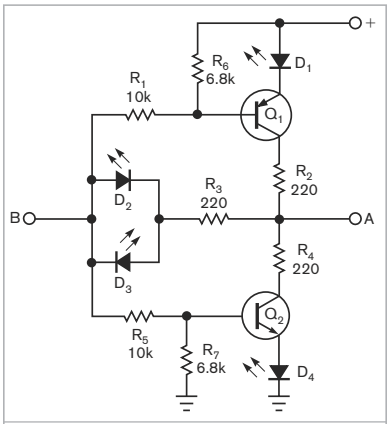

Figure 1 This simple circuit displays four states on four LEDs.

TABLE 1 LED-LIGHTING POSSIBILITIES

IN		LED			
A	B	1	2	3	4
0	0	1	0	0	0
0	1	0	1	0	0
1	0	0	0	1	0
1	1	0	0	0	1

productroundup

SWITCHES AND RELAYS

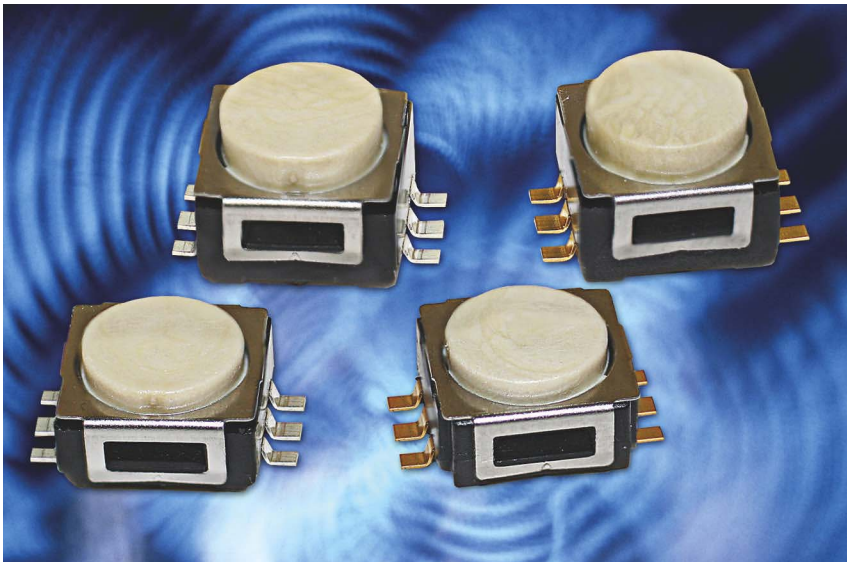

Snap-action key switches have tactile feel

The K5G and the K5A2 snap-action key switches feature silver- and gold-plated dome inserts, improving user feel and tactile feedback. Measuring 8.2×8×6 mm, the K5G features a 3.7N mean operating force for the first stroke and 6.5N for the second stroke and has 0.5- and 0.85-mm travel for the first and second strokes, respectively. The momentary-action switch includes an operational life of 30,000 actuations, targeting car-window and sunroof controls, as well as multifunction controllers. The K5A2 key switch measures 8.2×8×4.5 mm and has a 5.5N mean operating force; it is also available in a 6N version. Suiting use in automotive, computing, and networking equipment, the K5A2 has 0.5±0.1-mm electrical travel, with a 50,000-actuation operational life. The K5G and the K5A2 snap-action key switches cost 97 cents and $1.50 (1200), respectively.

C&K Components, www.ck-components.com

Triac photocouplers have high noise immunity

The TLP3782 and the TLP3783 800V zero-crossing triac-output photocouplers control current in the forward and the reverse directions. The devices provide 1500V typical impulse-noise immunity and 800V-peak off-state output voltage. The optical isolation between logic and the ac current reduces EMI (electromagnetic interference) and protects the circuit from current surges. The TLP3782 and the TLP3783 feature 10- and 5-mA trigger LED currents, respectively; a 30-mA input-forward current; and a 0.1A on-state current. Additional features include a gallium-arsenide infrared-emitting diode and a zero-crossing-turn-on photo triac. Available in DIP-6 packages, the TLP3782 and TLP3783 cost 50 and 55 cents, respectively.

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

Load switches feature low quiescent current

The FPF202x IntelliMAX load-switch series provides a 1-µA qui-

escent current. The series features a 1.6 to 5.5V operating range, addressing the voltage specification for biometric-sensor modules. The load switches come in 1×1.5-mm WL-CSP packaging with 5.5-kV ESD protection. Using lead-free terminals, the FPF202x load switches cost $1.49 (1000).

Fairchild Semiconductor, www.fairchildsemi.com

Rocker-switch series is water-resistant

Aiming at applications in fluids, weather, or other harsh conditions, the water-resistant KR rocker-switch series uses silver-alloy contacts, achieving a 20A current rating at 12V dc and 16A at 125V ac. A silicone-rubber actuator and a snap-fit mounting make the switches resistant to dust and water. Available in lighted and nonlighted versions with maintained- and momentary-action options, the devices also come in single-throw and double-throw designs. Fitting into a 13×19.8-mm, industry-standard panel cut, the KR series costs $4.

Cherry Electrical Products, www.cherrycorp.com

Power-distribution switches guard against short circuits

Suiting use in USB ports/hubs, cell phones, laptops, flat-panel televisions, and set-top boxes, the TPS2552/53 power-distribution switches enable end equipment to handle heavy capacitive loads and prevent short circuits. The devices provide a programmable 75-mA to 1.3A current-limit threshold using an external resistor with as much as 6% accuracy. The switches limit the output current to a safe level by switching into a constant-current mode when the output load exceeds the current-limit threshold, and the TPS2552-1/52-3 turns the switch off during overcurrent events. The units' enhanced accuracy allows you to use them as simple circuit-breaking or current-limiting devices for optical modules and low-current-management rails in telecommunications systems. The switches meet USB-current-limiting requirements and have a 2.5 to 6.5V input range and a 2-µsec overcurrent response. Available in SOT23-6 and 2×2-mm SON-6 packaging, the TPS2552/53 costs 75 cents (1000).

Texas Instruments, www.ti.com

Push-in terminals suit slim relays

The G2RV-SL500 SPDT (single-pole/double-throw) industrial-relay series features a push-in terminal-socket option, which allows for wiring and disconnection with minimal assistance from a flat-blade screwdriver. A relay-and-socket combination switches 6A at 250V-ac/30V-dc resistive load and provides 100,000 operations in a typical life span. Features include a transparent casing for visual assurance that the relay is making or breaking contact and a mechanical indicator window on the top surface, displaying red when operational, and a green LED on the socket, indicating that the coil is powered up. Coil voltages come in 12, 24, and 48V dc or 24, 48, 110, or 230V ac. Available in a slim, 6-mm package, the G2RV-SL500 relays with socket cost $10.18 (500).

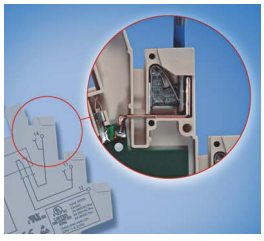

Omron Electronic Components, www.omron.com

PowerWise

Best-in-Class Dimming Performance. High Efficiency.

High Performance, TRIAC Dimmable LED Driver

National's industry-leading TRIAC dimmable offline LED driver solution is perfect for any application where an LED driver must interface to a standard TRIAC wall dimmer. National's new TRIAC dimmable LED driver delivers a wide, uniform dimming range free of flicker, best-in-class dimming performance, and high efficiency—all while maintaining ENERGY STAR® power factor requirements in a typical application.

- ✔ Online design tools
- ✔ Reference designs
- ✔ Application notes

Full Range Dimming Capability

National's TRIAC dimmable LED driver offers 100:1 full range dimming capability, going from full light to nearly imperceptible light in a continuous range without being extinguished, and maintains a constant current to large strings of LEDs driven in series off of a standard line voltage.

Easy To Use

National's TRIAC dimmable LED driver enables a direct replacement of incandescent or halogen lamp systems that are currently interfaced to a TRIAC dimmer without having to change the original infrastructure or sacrifice performance. In addition, the new TRIAC dimmable LED driver is available in WEBENCH® LED Designer to allow for easy and quick design in.

Uniform Dimming Without Flicker

National's TRIAC dimmable LED driver allows master-slave operation control in multi-chip solutions which enables a single TRIAC dimmer to control multiple strings of LEDs with smooth consistent dimming, free of flicker.

national.com/LED

National Semiconductor

EDN ADVERTISER INDEX

Company	Page
Advanced Interconnections	46
Advanced Linear Devices	37
Agilent Technologies	10
	19
Allied Electronics	C-4
Analog Devices Inc	17
austriamicrosystems AG	57
Coilcraft	9
Digi-Key Corp	1
Ellsworth Adhesives	45
ESS Technology Inc	39
Express PCB	45
International Rectifier Corp	7
Intersil	23
	25
	26
	27
Jameco Electronics	4
Keil Software	59
Linear Technology Corp	47
	49
	50
LPKF Laser & Electronics	35
MathWorks Inc	21
Maxim Integrated Products	53
	55
Melexis Inc	31
Mentor Graphics	15
Microsoft Corp	33
Mill Max Manufacturing Corp	11
Mouser Electronics	6
National Instruments	2
National Semiconductor	60
Pico Electronics	36
	58
Signal Consulting Inc	40
Toshiba America	C-2
	41
Xilinx Inc	C-3

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

EDN productmart

This advertising is for new and current products.

Reversal of fortune

Years ago, I had to design a dc/dc converter, or SMPS (switched-mode power supply), to generate isolated 48V dc from a 5V rail. A manager asked why I could not just use an off-the-shelf 48V-to-5V converter and run it in reverse. I explained about output-rectifier diodes and their one-way characteristics.

Years later, SMPS conversion efficiencies had reached 90% or more. Designers achieved this advancement by using low-on-resistance, low-forward-voltage-drop FETs as synchronously driven rectifiers in place of diodes.

A recent project for a high-current server power supply required running two SMPSs in parallel with current sharing for fail-safe redundancy. Each SMPS could support the full load should the other fail. To maintain efficiency, we used back-to-back series FETs instead of diodes at each SMPS output to combine the two outputs into a common rail. The advantage was low FET forward-voltage drop, but, unlike diodes, the FETs conduct in both directions and needed a controller to monitor forward and reverse currents and voltages and to turn these FETs on and off (**Reference 1**). Reverse current is most troublesome during start-up when the SMPS voltage and current-share control loops are stabilizing. The controller temporarily blanks its reverse-current sensing to allow for this problem and keeps the FETs turned on.

A common 48V rail powered both SMPS devices. If the 48V rail should momentarily droop due to the finite-response time to isolate a short elsewhere, series diodes would reverse-bias, and storage capacitors would keep each SMPS running during the few microseconds of droop.

The first prototype units turned out to be disasters. The SMPSs would not start up properly; if one was a slightly lower voltage than the other, it received reverse current and would continuously shut down and then restart, only to shut down again. The SMPS manufacturer said the SMPS should not shut down in the event of reverse current into its output. A colleague had seen this happen before, and he suggested measuring the input voltage to the SMPS that was shutting down.

Huh? The input is 48V, right? Wrong. The input voltage started at 48V and ramped up to 57V, at which point the SMPS detected excessive input voltage and shut down just as it was supposed to. Then, the input dropped back to 48V, and the cycle repeated. The synchronous-FET rectifiers were chopping the reverse current and causing current to flow from the SMPS input. This current flow caused the input isolation diode to reverse-bias, and the current had nowhere to go but into the input storage capacitor, where it ramped the voltage upward.

I was amazed by this revelation and remembered that meeting long ago when I told the manager that dc/dc converters do not work in reverse. It seems that I was wrong in the case of synchronous-FET rectifiers.

The solution was to include a temporary switched-FET path between the two 48V SMPS inputs to circulate reverse current during the 40-msec start-up phase. After these changes and a few others for different problems, it worked great. EDN

REFERENCE

1 "ORing MOSFET Controllers with Fastest Fault Isolation for Redundant Power Supplies," Maxim Integrated Products, 2004, http://datasheets.maxim-ic.com/en/ds/MAX8535-MAX8585.pdf.

Contact design consultant Glen Chenier at glen@teetertottertreestuff.com.

www.edn.com/tales

DANIEL VASCONCELLOS

CAN'T FIND PRODUCT QUICKLY?

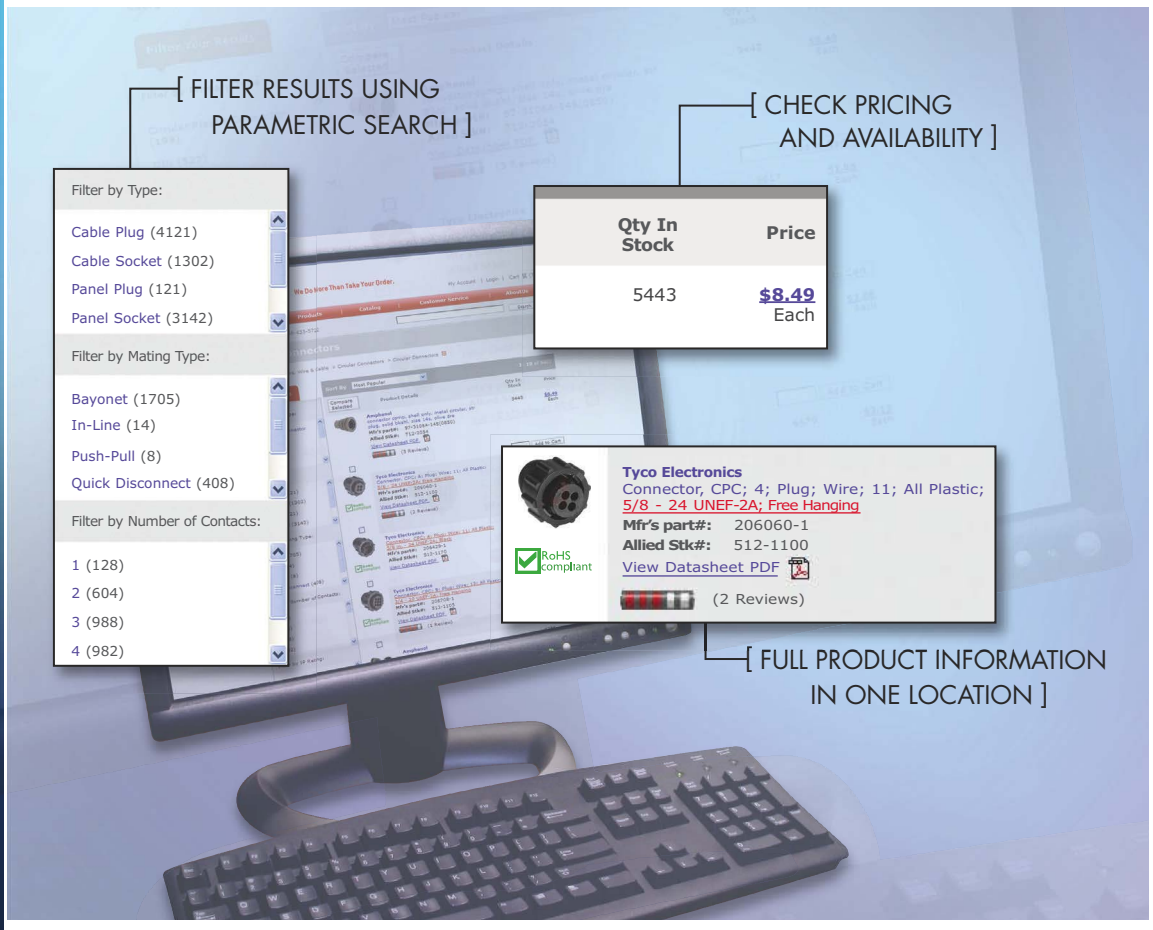

Quick and easy solutions @ alliedelec.com

alliedelec.com

We do more than take your order℠

1.800.433.5700

ALLIED ELECTRONICS®