UBM Electronics

EDN®

VOICE OF THE ENGINEER

SEPTEMBER 8
Issue 17/2011
www.edn.com

Tales from the Cube: Coffee-break mistake Pg 62

Engineering the next generation Pg 10

Semiconductors develop an EUV ecosystem Pg 24

Design Ideas Pg 49

Supply Chain Pg 58

DSOs:

THE BANNER SPECS TELL ONLY PART OF THE STORY

Page 34

CHIP-AND-PACKAGE **CO-DESIGN** RELIEVES PRESSURE ON COMPLEX DESIGNS

Page 27

The Next Generation of Signal Integrity Testing

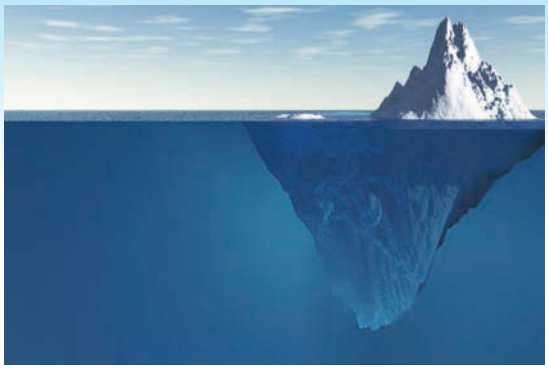

Can you see what dangers are lurking in your multi-lane design?

You already know your scope is inadequate for detecting low-probability events, but did you know that a single-channel BER tester isn't enough to find errors caused by crosstalk?

How are you testing crosstalk susceptibility?

Our parallel channel BERT provides multiple aggressors with a phase Delay Sweep feature that independently sweeps the aggressor's transitions over one, two, or four UI. Each aggressor is swept at a different rate. This creates the perfect storm – the guaranteed worst-case timing conditions that ensure you are thoroughly testing your DUT's susceptibility to crosstalk.

Don't risk missing a design flaw. Load all channels. Sweep your aggressors.

Ideal For:

- CFP
- QSFP
- CXP
- Active Optical Cables

The most compact solution available

YouTube

Scan the QR code or visit goo.gl/P6dNy to see a product video

CENTELLAX

SIGNAL INTEGRITY MEASUREMENT

Affordability without compromise

SEE FOR YOURSELF:

www.centellax.com/next_gen

EDN contents

9.8.11

DSOs: The banner specs tell only part of the story

34 To make an intelligent choice of a scope, you must know its banner specs—bandwidth, sampling rate, and memory depth. Before you sign a purchase order, though, you'd be wise to find out much more.

by Dan Strassberg, Contributing Technical Editor

Chip-and-package co-design relieves pressure on complex designs

27 Co-design accomplishes more with less by performing more design tasks in a shorter design cycle.

by Darvin Edwards, Texas Instruments

Designing reliable capacitive touch interfaces

43 Understand how to enhance the hardware and software reliability of this increasingly popular user interface for appliances and instruments.

by Sachin Gupta, Prakhar Goyal, and Kurian Polachan, Cypress Semiconductor Corp

COVER: DAVID NICASTRO/PENCIL IMAGE: ISTOCKPHOTO.COM

DESIGNIDEAS

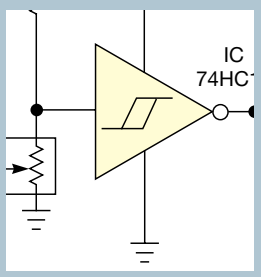

49 Polynomial rotation accelerates CRC calculations

50 Simple circuit measures optocoupler's response time

52 Circuit provides visual verification of IR pulses

54 An LED's intrinsic capacitance works in a 650-mV LRC circuit

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

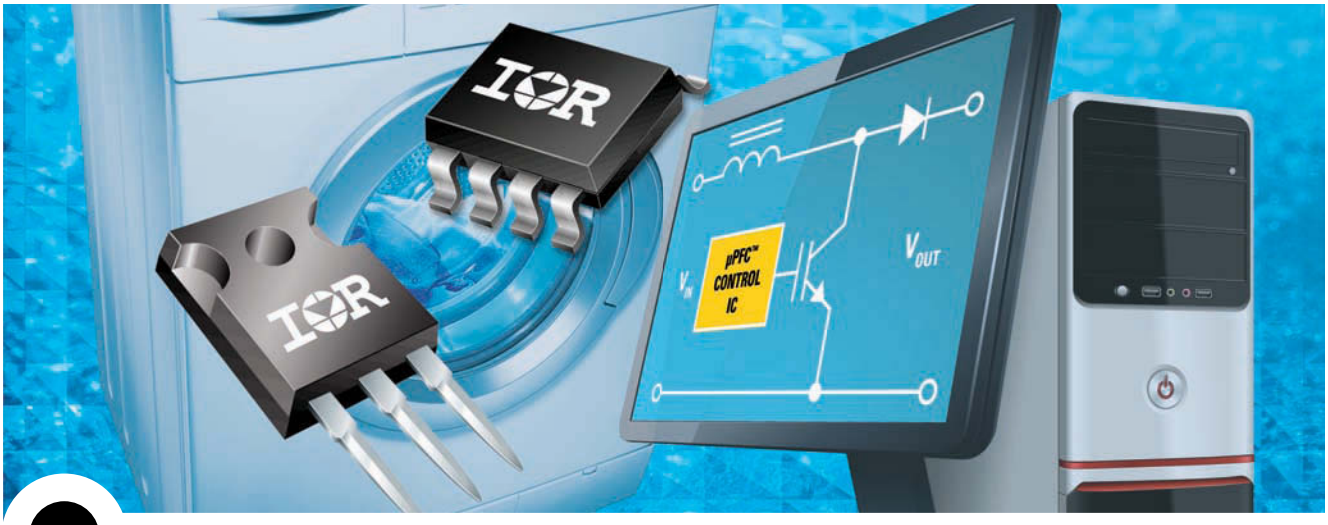

Delivering Real Power

Compact, Powerful CCM PFC from 75W to 4kW+

PFC Control ICs

Part Number	Description
IR1152	Fixed 66KHz switching frequency with brownout protection and dual OVP protection
IR1153	Fixed 22KHz switching frequency with brownout protection and programmable OVP protection
IR1155	Programmable switching frequency and programmable OVP protection

600 V PFC IGBTs for High Power Systems

Part Number	Circuit	I_c @ 100C (A)	$V_{CE(on)}$ (max) (V)	Package
IRGB20B60PD1	Co-Pack	22	2.35	TO-220AB
IRGP20B60PD		22	2.35	TO-247AC
IRGP35B60PD		35	2.15	TO-247AC
IRGP50B60PD1		45	2.35	TO-247AC
IRG4(B/IB)C20W	Discrete	6.5	2.6	TO-220AB; TO-220 FullPak
IRG4(B/IB/P)C30W		12	2.7	TO-220AB; TO-220 FullPak; TO-247AC
IRG4(B/P)C40W		20	2.5	TO-262; TO-220AB; TO-247AC
IRG4PC50W		27	2.3	TO-247AC
IRGP4069		35	1.85	TO-247AC
IRGP4063		48	2.14	TO-247AC
IRGP4066		90	2.1	TO-247AC

The μPFC family of controller ICs radically alters traditional thinking about PFC solutions. The IR115x uses a "One-Cycle Control integrator with reset" technique to deliver the high performance of Continuous Conduction Mode (CCM) PFC with the simplicity and low component count of Discontinuous Current Mode (DCM).

Features

- Small, easy, powerful solution
- Fast time to market
- Enables compliance with energy standards (1W, Blue Angel, Energy Star)
- No AC line voltage sense required
- 0.999 power factor
- Switching frequency of 22kHz, 66kHz or programmable value available
- Average current mode control
- Cycle by cycle peak current limit system protection

for more information call 1.800.981.8699 or visit us at www.irf.com

International
IOR Rectifier
THE POWER MANAGEMENT LEADER

contents 9.8.11

pulse

Dilbert 16

15 Scope, spectrum-analyzer combination provides time-correlated signals

16 Orca shows a full dual-mode Bluetooth core in CMOS

18 Synchronous, forward-converter chip set provides active clamp, deters transformer saturation

18 True-green laser module targets use in smartphone picoprojectors

20 Bridgelux plans to bring economies of GaN-on-Si to HB LEDs in 24 months

20 Watch out for well-made—but counterfeit—chips

22 New architecture promises better battery

22 Source code for µC/OS-III is now available

DEPARTMENTS & COLUMNS

9 **EDN online:** Join the conversation; Content; Engineering Community

10 **EDN.comment:** Engineering the next generation

24 **Inside Nanotechnology:** Semiconductors develop an EUV ecosystem

58 **Supply Chain:** Tool is in the works to standardize conflict-minerals tracking; Indian semiconductor consumption to grow 15.5%

60 **Product Roundup:** Switches and relays; Enclosures

62 **Tales from the Cube:** Coffee-break mistake

EDN® (ISSN# 0012-7515) is published semimonthly by UBM Electronics, 600 Community Drive, Manhasset, NY 11030-3825. Periodicals postage paid at Manhasset, NY, and at additional mailing offices. SUBSCRIPTIONS—Free to qualified subscribers as defined on the subscription card. Rates for nonqualified subscriptions, including all issues: US, $150 one year; $250 two years; $300 three years. Except for special issues where price changes are indicated, single copies are available for $10 US and $15 foreign. For telephone inquiries regarding subscriptions, call 847-559-7597. E-mail: edn@omeda.com. CHANGE OF ADDRESS—Notices should be sent promptly to PO Box 47461, Plymouth, MN 55447. Please provide old mailing label as well as new address. Allow two months for change. NOTICE—Every precaution is taken to ensure accuracy of content; however, the publishers cannot accept responsibility for the correctness of the information supplied or advertised or for any opinion expressed herein. POSTMASTER—Send address changes to EDN, PO Box 47461, Plymouth, MN 55447. CANADA POST: Publications Mail Agreement 40612608. Return undeliverable Canadian addresses to BleuChip International, PO Box 25542, London, ON N6C 6B2.

Volume 56, Number 17 (Printed in USA).

When the innovators at KUKA decided to create the foundation for the automation of tomorrow, they decided to work with Wind River. Building upon the proven reliability and breakthrough performance of our VxWorks and multi-core platforms, the KUKA team developed their revolutionary KR C4 controller—a safer, more powerful, more flexible system that allows all integrated controllers to share a common database and infrastructure. And by working with Wind River, they did it all ahead of schedule, under budget, and with far less risk.

Please visit **http://www.windriver.com/customers** to learn more about how Wind River customers have transformed their devices.

WIND RIVER

INNOVATORS START HERE.

EDN online

JOIN THE CONVERSATION

Comments, thoughts, and opinions shared by *EDN*'s community

In response to Paul Rako's Prying Eyes tear-down, "The Tektronix 1101 oscilloscope-probe power supply," at http://bit.ly/noHonC, David Lockman commented:

"To paraphrase Santayana, those who cannot learn from history are doomed to repeat it. This is an opportunity to learn how to solve design problems, and perhaps save yourself a little time and effort when you are called upon to solve a similar problem."

In response to "iFixit tears down a smart meter to evaluate electro-smog," posted in Margery Conner's PowerSource blog at http://bit.ly/EDN_ElectroSmog, Davidiode commented:

"I love this stuff. These tear-downs are fascinating. Now and historically, the scientific/technical have always been held hostage by myths and fuzzy thinking, worried about what they don't understand. The Al Gores of the world always get more traction than the Bob Peases."

In response to "Front-end, back-end, and the places in between," posted in Ron Wilson's Practical Chip Design blog at http://bit.ly/EDN_FrontBack, William Ketel commented:

"Ron certainly makes an accurate assessment of the problem. I offer here a suggestion which many will scream at as heresy, since it would reduce the options for "product differentiation," and all the marketing wonks will throw up. But how about holding the hardware constant and just changing the software? One large designed die, and more creative code? Saving a few million on a new hardware spin should be an option worth something, especially if the production yields are already high and stable. Just a suggestion."

EDN *invites all of its readers to constructively and creatively comment on our content. You'll find the opportunity to do so at the bottom of each article and blog post. To review current comment threads on EDN.com visit this page: http://bit.ly/EDN_Talkback.*

CONTENT

Can't-miss content on EDN.com

WHEN KIDS REALLY HAD FUN WITH SCIENCE

Were you a kid who almost blew up the garage or nearly burned down the house on your way to becoming an engineer? *EDN*'s Paul Rako wonders if today's supersafe, mom-approved science and tech toys for kids hinder learning by taking the fun out of experimenting. He and Analog Devices fellow Barrie Gilbert also share some of their dangerous but fun experiments from childhood in this blog post. Check it out and share your own memories in the post's comment field.

http://bit.ly/nEVdvl

ENGINEERING COMMUNITY

Opportunities to get involved and show your smarts

In *EDN*'s Tales From The Cube, engineers relate their most vexing design challenges—and how they conquered them. What's your Tale? Tell us about a memorable experience you've had successfully solving an engineering problem at work or in your spare time, and you could win a Tektronix scope! **Visit this page for contest information: http://bit.ly/Tales_EDN.**

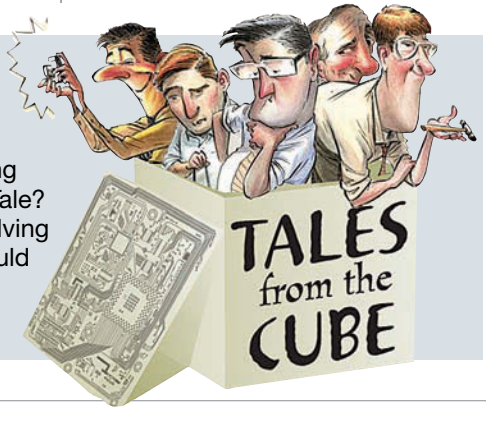

BY SUZANNE DEFFREE, MANAGING EDITOR, ONLINE

Engineering the next generation

The US Department of Commerce in July released a report on STEM (science/technology/engineering/mathematics) employment that stated some interesting facts (**Reference 1**). For example, the department projects that STEM occupations will grow by 17% from 2008 to 2018, compared with 9.8% growth for non-STEM occupations. STEM workers also command higher wages, earning 26% more than their non-STEM counterparts, and STEM-degree holders enjoy higher earnings, regardless of whether they work in STEM or non-STEM occupations.

All of these points are valuable when it comes to US competitiveness, individual earning power, and maintaining employment in what has in recent years proved to be a less-than-reliable career environment. It was the last point, however, that made me take notice of this report. Based on that point, *EDN* ran a blog post in August that started some good conversations on how to encourage kids toward STEM study and on the next generation of engineers.

One of those conversations was with Aaron Goldstein (**photo**), an Arizona State University senior. Some lucky company will surely snatch him up when he receives his bachelor's degree next May. Goldstein's story starts like that of many other engineers: He played with Legos as a child, he liked to tinker, and he had a dad who sparked a curiosity in him to ask "why" and supported him along his way to a STEM education. His father even pointed out the blog post to him.

Goldstein is growing up to be a leader. In addition to interning with the National Aeronautics and Space Administration and aerospace companies, he rallied his fellow classmates and professors to start ASU's Sun Devil Satellite Laboratory. Through this group,

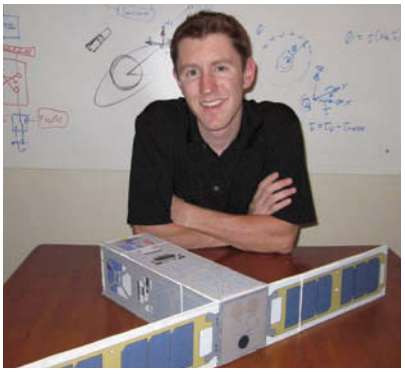

he and a handful of fellow ambitious students started out to build a functioning satellite. They have contacted NASA and are now waiting for a launch date, with expectations for early 2013.

Goldstein spent his college years on real-world projects and gained real-world experience while networking with real engineers, not simply practicing theory and academics. He speaks highly of his school but recognizes that what is taught at many universities is not engineering; it is science and mathematics on paper. Real-world application of a STEM education and what it means to be an engineer once that degree is in hand are more valuable than the degree itself in many ways, and the only way to get that experience is by working with other more experienced engineers.

Note that the headline on this editorial is not "Engineering: the next generation," which with the colon would grammatically mean that we were discussing the incoming work force of young, less-experienced engineers. The headline, instead, uses "engineering" as a verb. We have to build the next generation of STEM professionals. These are the people who will one day sit where you sit, face the design challenges that you face, and innovate in ways and fields that we may not even imagine yet—but they won't get there without your help.

Understanding that, *EDN* and its sister publications will be hosting a meeting at ESC (Embedded Systems Conference) Boston in late September. We're looking for the slide-rule crowd of experienced engineers who would like to connect with the next generation, spending an hour or so offering real-world experiences and histories. Perhaps you'll choose to work for the long term with these kids—maybe become a mentor to some of them; that's up to you.

If you are new to engineering, so new that you've never touched a slide rule, we welcome you to come and pick the brains of those who have. There's a lot of knowledge for engineers to pass on from decades of hands-on work that you won't get in a textbook.

This networking event will be the first of many efforts we will make to help engineer the next generation. You'll find details on the meeting in the original STEM salary blog that started this project (http://bit.ly/STEM_EDN). Also, be sure to sign up for ESC, the East Coast's leading embedded-systems event (http://bit.ly/ESC_EDN). Connect with your peers and spend a few minutes engineering the next generation. EDN

REFERENCE

1 Langdon, David; George McKittrick; David Beede; Beethika Khan; and Mark Doms, "STEM: Good Jobs Now and for the Future," US Department of Commerce, Economics and Statistics Administration, July 2011, http://bit.ly/ojvdjE.

Contact me at suzanne.deffree@ubm.com.

Our team will make the difference for your business.

With 42 years of success behind us, we're building upon our proven experience to implement new ways of doing business and introducing unprecedented strategies for hitting your goals.
We are injecting a new level of commitment and an elevated sense of urgency to perform better than ever for your business today.

Call us today and experience the difference.

FUTURE ELECTRONICS

www.FutureElectronics.com **1-800-FUTURE-1**

Q Series®

Proven. Flexible. Diverse.

Integral ground/power plane rated up to 9.5 amps.

Q Strip® - The industry standard for low profile board stacking on 0,5mm, 0,635mm and 0,8mm pitch.

Q Rate® - Smaller footprints and lower profile right angle without sacrificing performance.

Q2™ - Rugged contact system provides increased contact wipe.

Elevated, right angle, edge mount, power and shielding options.

QSERIES® HIGH SPEED INTERCONNECTS

www.samtec.com

SUDDEN SERVICE

EDN

ASSOCIATE PUBLISHER, *EDN* WORLDWIDE
Bob Dumas,
1-516-562-5742;
bob.dumas@ubm.com

DIRECTOR OF CONTENT, *EDN* AND DESIGNLINES
Patrick Mannion,
1-631-543-0445;
patrick.mannion@ubm.com

EDITOR-AT-LARGE
Ron Wilson
EDA Tools, Chip Design, Services, Silicon
1-415-947-6317;
ron.wilson@ubm.com

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

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

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

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

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

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

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

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

VICE PRESIDENT/DESIGN DIRECTOR
Gene Fedele

CREATIVE DIRECTOR
David Nicastro

ART DIRECTOR
Giulia Fini-Gulotta

PRODUCTION
Adeline Cannone, Production Manager
Laura Alvino, Production Artist
Yoshihide Hohokabe, Production Artist
Diane Malone, Production Artist

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

EDN ASIA
Grace Wu, Associate Publisher
grace.wu@ednasia.com
Vivek Nanda, Executive Editor
vnanda@globalsources.com

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

EDN JAPAN
Masaya Ishida, Publisher
mishida@mx.itmedia.co.jp
Makoto Nishisaka, Editor
mnishisa@mx.itmedia.co.jp

UBM ELECTRONICS MANAGEMENT TEAM
Paul Miller,
Chief Executive Officer, UBM Electronics and UBM Canon (Publishing)
Brent Pearson, Chief Information Officer
David Blaza, Senior Vice President
Karen Field,
Senior Vice President, Content
Jean-Marie Enjuto,
Vice President, Finance
Barbara Couchois,
Vice President, Partner Services and Operations
Felicia Hamerman,
Vice President, Marketing
Amandeep Sandhu,
Director of Audience Engagement and Analytics

UBM Electronics

For a complete list of editorial contacts, see http://ubmelectronics.com/editorial-contacts

Choose a bench-top SMU that's clearly more efficient.

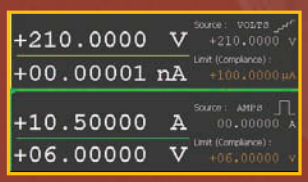

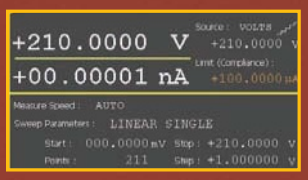

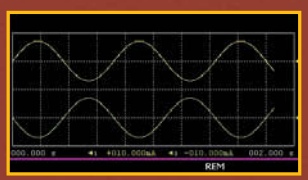

The superior GUI on the new Agilent bench-top source measurement unit (SMU) is the first in the industry to display your measurement results graphically. It also provides best-in-class output range and source/measurement resolution. For the testing speed and efficiency you need to deliver smart technologies to market faster, the choice is clear.

That's thinking ahead. That's Agilent.

Agilent B2900A Series Precision SMU	
Max output range	210 V, 3.03 A (DC) / 10.5 A (Pulsed)
Min resolution	10 fA / 100 nV
Front panel GUI	4 viewing modes on 4.3" color LCD

Learn how to make more efficient measurements by viewing a demo video or downloading application notes at:
www.agilent.com/find/AgilentSMU

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

Agilent Technologies

Scope, spectrum-analyzer combination provides time-correlated signals

To address its belief that more than 60% of oscilloscope users also use a spectrum analyzer to troubleshoot embedded-system designs with integrated wireless functionality, Tektronix Inc has developed the MDO (mixed-domain oscilloscope). The MDO4000 enables work in both the time and the frequency domain to capture time-correlated analog, digital, and RF signals for a complete system view.

With the MDO4000, the RF-input frequency range extends to 6 GHz and provides a capture bandwidth of ≥1 GHz at all center frequencies, 100 times wider than that of typical spectrum analyzers. Users can see up to four decoded serial buses, parallel buses, or both at one time on the same display. The time correlation between domains enables accurate timing measurements for understanding delays and latencies between command and control events in a design and changes in the RF spectrum.

The MDO4000 also allows designers to view the RF spectrum of a signal at any time within a long acquisition to see how the spectrum changes over time or with device state. RF time-domain traces show how the amplitude, frequency, or phase of the RF input signal changes relative to time, which eases the characterization of frequency-hop transitions, settling times, and RF event timing relative to other system components and activities.

In addition to the standard RF power-level trigger, an optional module, MDO4TRIG, allows additional trigger types to use the RF power level as a source, enabling customers to further isolate an RF event of interest. Users can trigger on a specific pulse width, look for a time-out event or runt, or even include the RF input in a logic pattern defined along with the analog and digital channels.

"It fundamentally changes what's involved in debugging designs with RF where there is a need to correlate events in the frequency domain with the time-domain phenomena that caused them," says Roy Siegel, general manager, oscilloscopes, at Tektronix. "Just as the MSO [mixed-signal oscilloscope] is the standard for embedded-design test, we expect the MDO will become the new standard for designs that increasingly include RF capabilities."

The MDO4000 enables debugging integration of common wireless modules below 6 GHz—WLAN, Bluetooth, and Zigbee, for example; debugging of common "homegrown" amplitude-, frequency-, or phase-modulated wireless communications; and wideband analysis of dual-band transceivers. US suggested list prices start at $19,900. **—by Colin Holland**

▷**Tektronix**, www.tek.com.

TALKBACK

"Isn't it amazing how the folks in purchasing know so much more about parts than engineers know, and how they can improve a design by using cheaper parts?"

—Engineer William Ketel, in *EDN*'s Talkback section, at http://bit.ly/qT1yrh. Add your comments.

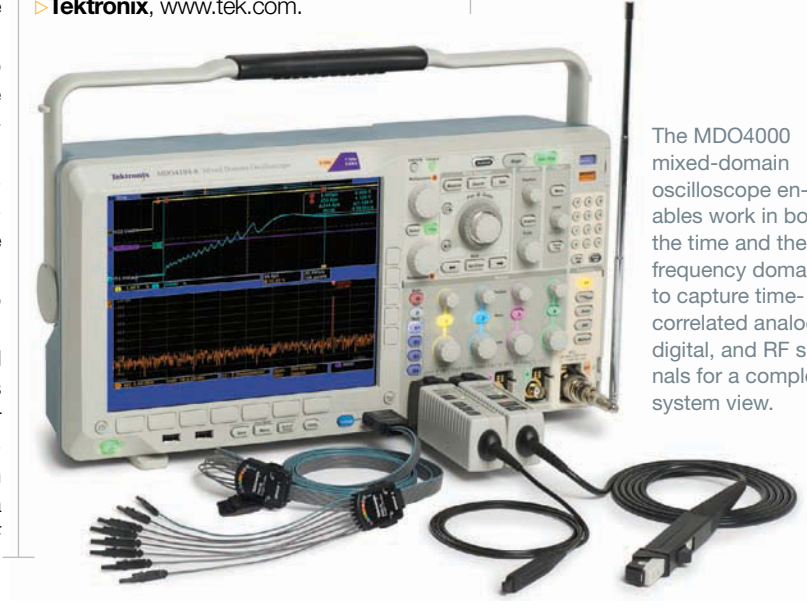

The MDO4000 mixed-domain oscilloscope enables work in both the time and the frequency domain to capture time-correlated analog, digital, and RF signals for a complete system view.

Orca shows a full dual-mode Bluetooth core in CMOS

The market trajectory of Bluetooth shows the value of perseverance. From a premature launch resulting in something remarkably similar to abject failure, Bluetooth became the unquestioned standard for people who wanted to wear a low-gain microwave oven next to their skulls; recharge the thing every night; and frequently say, "Sorry, what was that?"

From such humble success, the standard has gone on to become a plausible $75 replacement for a meter of six-conductor cable. And with the addition of the EDR (extended-data-rate) and low-energy modes, Bluetooth could become, almost by default, a major factor in all sorts of near-field communications, including the machine-to-machine kind.

As Bluetooth spreads into applications requiring significant levels of integration, it begins to present a problem for SOC (system-on-chip) designers. A Bluetooth transceiver is a relatively sophisticated 2.4-GHz spread-spectrum radio—not the sort of thing you license at RTL (register-transfer level) and hand off to a subcontractor to synthesize and integrate. Only a few complete Bluetooth-radio cores are on the market. Most of the IP (intellectual property) out there is in the form of building blocks for design teams who know their way around the inside of a digital radio.

This situation makes IP developer Orca Systems' recent announcement intriguing. The company has released a complete, turnkey Bluetooth radio—both the RF transceiver and the modem—supporting the original, EDR, and low-energy modes. Supporting all three modes in a compact core necessarily means reusing most of the blocks in the signal paths.

That necessity in turn means that power management must exploit every opportunity the standard affords for clock gating; otherwise, the low-energy mode would consume about the same power as the original mode. Even so, the design may not be as energy-efficient as a low-energy-only core. For an SOC that must work in multiple modes, however, some savings are better than none.

The block is tiny: 2.7 mm^2 in Fujitsu's (www.fujitsu.com) 90-nm CMOS, including both the transceiver and the modem. Better yet, unlike most RF-CMOS designs that are nearly process-specific, it should scale well into finer geometries. Orca delivers the core as a single hard macro, including both transceiver and modem.

All of these features are possible, according to Joe Thome, Orca's vice president for business development, because of the digital-radio technology that makes up Orca's crown jewels. The company's foundation is a digital open-loop frequency synthesizer that provides sufficient spectral purity for a Bluetooth transmitter but with low frequency-shift latency because of the open-loop architecture. Basically, it doesn't have to search and lock the way a PLL (phase-locked loop) does; it simply generates whatever frequency you program it to generate.

> Orca has so far produced the block in 130-, 90-, and 65-nm geometries, and it has scaled as predicted.

Orca designers directly modulate the synthesizer at the Bluetooth 1600-Hz hop rate, providing a fully digital signal source. Going a step further, Orca has developed a digital power amplifier to drive the antenna. This move is not unprecedented: Class D amplifiers at 2.4 GHz have appeared in the literature, notably from the University of California—Berkeley (www.berkeley.edu) at the 2008 Custom Integrated Circuits Conference. But it is hard to find any reference to a similar commercial product.

The designers did not stop there. Orca has also developed a mostly digital receiver chain. The design starts with a conventional analog low-noise amplifier, but, instead of analog filtering and mixing, the Orca design directly samples the RF, turning the continuous-time radio signal into a very-high-frequency sampled-data stream. The core does the filtering and mixing functions in the sampled-data domain, apparently somehow using digital circuitry to do so, and then passes the output to an ADC for use by the baseband processor.

These approaches provide several advantages for SOC designers, Thome says. The first is the size of the core. Without large matched RF transistors and passive components, the design is small. It is also scalable. The low-noise amp is a true analog device; otherwise, however, the radio area should scale almost like a purely digital block. Orca has so far produced the block in 130-, 90-, and 65-nm geometries, and it has scaled as predicted, Thome adds.

The DSP-RF architecture uses fewer inductors than would a conventional design. All of this digital talk suggests a standard digital process. A white paper from Fujitsu, which appears to describe the development of Orca's sample chip, says that the team used Fujitsu's 90-nm, low-power RF CMOS, with a special RF process-design kit. The core, the prototype IC, and an evaluation board are now available.

—by Ron Wilson

▷**Orca Systems**, www.orcasystems.com.

DILBERT By Scott Adams

More engineers choose ADI data converters over the next 8 suppliers combined.

More products. More performance. More design support. More education.

With so many complex design problems relying on data converters to define system performance, and so many data converter suppliers to choose from, you might think it's remarkable that so many engineers agree on their choice: Analog Devices. As the world's #1 data converter supplier, ADI is your source for 1500+ converter products, all thoroughly and conservatively specified, and backed with more tools, training, and design support than any other supplier—by far. See why more engineers turn to ADI for answers to their toughest challenges. Explore products, customer stories, and design resources at *www.analog.com/ConverterLeader.*

Leading Converter Innovations

- **AD9467: 16-Bit, 250 MSPS, IF Sampling ADC**
  Provides outstanding SFDR performance over wide bandwidths; multiple on-chip functions and adjustability features for ease of use.
- **AD5624R: Quad, 12-Bit SPI *nano*DAC®**
  Offers 12-bit accuracy (1 LSB INL) and on-chip 1.25 V or 2.5 V 5 ppm/°C reference, all in a 10-lead MSOP and 3 mm × 3 mm 10-lead LFCSP.
- **AD7626: 16-Bit, 10 MSPS PulSAR® ADC**
  Delivers an unequaled combination of effective resolution and speed, without compromising ac or dc performance, while reducing size and power.
- **AD9122: Dual, 16-bit, 1200 MSPS, TxDAC+® Transmit DAC**
  Enables multicarrier generation up to the Nyquist frequency and includes features optimized for direct conversion transmit applications.

2010 Worldwide Data Converter Market Share

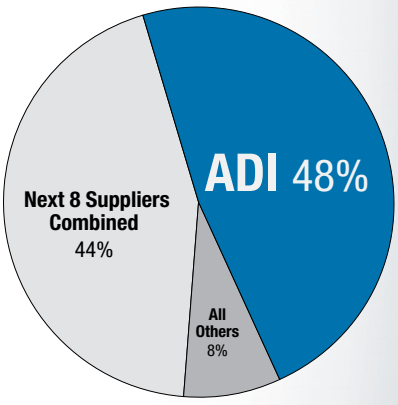

Source: Databeans, Inc.

Circuits from the Lab™
Reference Circuits
Lower your design risk and speed time to market with ADI lab tested reference circuits. Use a QR reader on your smartphone to watch video.

analog.com

ANALOG DEVICES

Synchronous, forward-converter chip set provides active clamp, deters transformer saturation

Many dc/dc converters use transformers for isolation, potentially causing a problem: Following large, sudden changes in output load or input voltage, the loop closure may be too slow to prevent transformer saturation and resultant converter failure. To overcome this problem, Linear Technology Corp has introduced a pair of ICs that together implement a synchronous, forward-converter function with active-clamp reset and a flux-limiting capability to prevent such saturation under all conditions.

The direct-flux-limit LTC3765 primary-side intelligent IC works with the companion LTC3766 to create a self-starting, isolated, forward-converter topology. This type of forward-converter design suits use in 12, 24, and 48V nominal-input voltages that are common in telecom, data communications, and industrial applications.

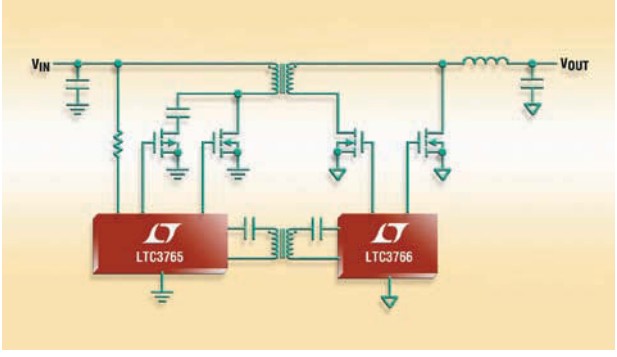

The LTC3765 and LTC3766 ICs together implement a synchronous, forward-converter function with active-clamp reset and a flux-limiting capability to prevent saturation under all conditions.

After start-up, the LTC3765 receives timing signals and bias power from the secondary-side LTC3766 controller through a tiny pulse transformer. Secondary-side control puts the intelligence close to the load, ensuring reliable control of the output voltage and current, providing fast transient response and eliminating the need for an optocoupler for isolation.

The LTC3765/LTC3766 contains control circuitry to implement an active-clamp transformer-reset technique, enabling efficiencies as high as 96% and greater power densities than those of conventional catch-winding or resonant-reset techniques. The design allows for adjustable delays for the high-current gate drivers for the main switch, the active-clamp switch, and the synchronous switches to achieve maximum efficiency. Other features include a fast and accurate average-current limit, fixed-frequency adjustable operation from 75 to 500 kHz, clean start-up into prebiased loads, multiphase operation for high-power designs, overtemperature protection, and true remote-output-voltage sense.

The LTC3765 is available in a thermally enhanced MSOP-16 package, and prices start at $1.55 (1000). The LTC3766 is available in 4×5-mm QFN-28 and SSOP-28 packages, and prices start at $2.60 (1000).

—by Bill Schweber

▷**Linear Technology Corp**, www.linear.com/product/LTC3765 and www.linear.com/product/LTC3766.

True-green laser module targets use in smartphone picoprojectors

Finnish start-up EpiCrystals Inc has announced what it claims is the smallest and most efficient green-light source for picoprojectors. EpiCrystals hopes that RGB (red/green/blue) laser technology will revolutionize smartphones, enabling them to project photos, video presentations, and movies on almost any surface with high-definition-level quality and resolution.

The EpiCrystals device uses frequency doubling to emit green light. According to Tomi Jouhti, vice president of business development, the laser comprises a GaAs (gallium-arsenide)-based, 1064-nm infrared laser chip and a PPLN (periodically poled lithium-nibate)-waveguide-type nonlinear crystal. The chip integrates a 45° folding mirror and lens, which the company builds at the wafer level, for efficient coupling. The emitted light is a bright, 532-nm color, which falls in the "true-green" spectral range of 515 to 535 nm.

> The laser has a spectrum width of 0.5 to 1 nm to reduce speckle, more than 60 mW of optical output power, and greater-than-100-MHz modulation speed.

The laser operates in passively Q-switched mode using simple construction and a three-contact laser and uses wideband QPM (quasiphase matching) and wideband SBG (surface Bragg grating) in the crystal, which locks the central frequency of the infrared laser. The laser also has a spectrum width of 0.5 to 1 nm to reduce speckle, has more than 60 mW of optical output power, and offers greater-than-100-MHz modulation speed. The TO56 module has a volume of approximately 0.2 cc and is compatible with MEMS (microelectromechanical-system), LCOS (liquid-crystal-on-silicon), and DLP (digital-light-processing) equipment. The green-laser module will be available in the first quarter of next year for $45 (10,000).

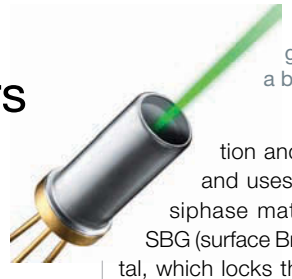

The EpiCrystals true-green laser module has a bright, true-green color.

—by Margery Conner

▷**EpiCrystals**, www.epicrystals.com.

100+ NEW CLOCK GENERATORS AND BUFFERS

LOW POWER, SMALL FOOTPRINT DEVICES FOR COST-SENSITIVE APPLICATIONS

Silicon Labs' broad portfolio of industry-leading clock generators, XOs, VCXOs, jitter attenuating clocks and clock buffers sets a new standard for flexibility, performance and lead times.

- Most frequency flexible timing solutions in the industry; ideal for multi-rate applications
- Available in industry-standard, drop-in compatible packages
- Based on patented DSPLL® and MultiSynth technologies, these low jitter products improve system performance, reduce BOM, minimize board space and simplify system design
- Helps reduce electromagnetic interference (EMI) and radio frequency interference (RFI) with a wide range of signal integrity adjustments such as rise/fall times, output impedance, drive strength and spread spectrum clocking
- All custom devices are available without NRE fees in <2-week lead times

PRODUCT LINE	WIRELESS INFRA-STRUCTURE	OTN	COMMS	BROADCAST VIDEO	EMBEDDED	SERVERS	CONSUMER
Jitter Attenuating Clocks	●	●	●	●			
Differential Clocks	●	●	●	●	●	●	
LVCMOS Clocks			●		●	●	●
PCI Express Clocks			●	●	●	●	●
x86 Embedded Clocks			●		●	●	●
Clock Buffers	●	●	●	●	●	●	
Oscillators	●	●	●	●	●	●	●

DIFFERENTIAL CLOCK GENERATORS

- Generates any frequency on any output up to 710 MHz with low phase jitter (<1 ps RMS)
- Any format on any output: LVPECL, LVDS, HCSL, CMOS, HSTL, SSTL

LVCMOS CLOCK GENERATORS

- Factory customized and I^2C programmable
- Generates multiple frequencies with equivalent frequency synthesis capability of up to 8 PLLs

PCI EXPRESS GEN 1/2/3 CLOCKS

- Up to 9 HCSL and up to 6 LVCMOS outputs
- Fully integrated terminator resistors on each output
- 35% lower power than competing solutions

EMBEDDED x86 CLOCKS

- Extensive family of Intel-compliant x86 clocks for embedded computing, communications, enterprise & industrial applications

EMI REDUCTION CLOCKS

- Spread percent from 0 to 5.0%
- Down or center spread options
- Spread modulation frequency 16 – 128 kHz

FANOUT & ZERO DELAY CLOCK BUFFERS

- Low power (2.7 mA typ), small size 10-TDFN (1.4 mm x 2.0 mm x 0.75 mm), very low phase noise (-168 dBc @ 1 MHz)

*

Find the perfect part for your application
at **www.silabs.com/parametric-search**

*Requires QR code reader; download at reader.kaywa.com. Data charges may apply.

SILICON LABS

Bridgelux plans to bring economies of GaN-on-Si to HB LEDs in 24 months

Bridgelux Inc recently announced its success in demonstrating GaN (gallium-nitride)-on-Si (silicon)-wafer LEDs with performance levels comparable to state-of-the-art GaN-on-sapphire or GaN-on-SiC (silicon-carbide)-based HB (high-brightness) LEDs. Cool-white LEDs show efficiencies as high as 160 lumens/W at a CCT (correlated color temperature) of 4350K, and warm-white LEDs deliver 125 lumens/W at a color temperature of 2940K and a CRI (color-rendering index) of 80. These numbers represent results of lab demonstrations at pulsed currents. The company claims that the lumen-per-watt values are the highest yet for GaN-on-Si LEDs.

Basing LEDs on less-expensive 8-in. silicon wafers will enable significant price drops in LEDs that will ultimately ripple through to the price of LED-based light bulbs. A thermal mismatch between GaN and Si can cause the epitaxial films to crack or the wafers to bow during or after manufacturing. The company claims, however, that its proprietary buffer-layer process produces crack-free, flat wafers. Bridgelux plans to deliver GaN-on-Si devices to the commercial market within 24 months.

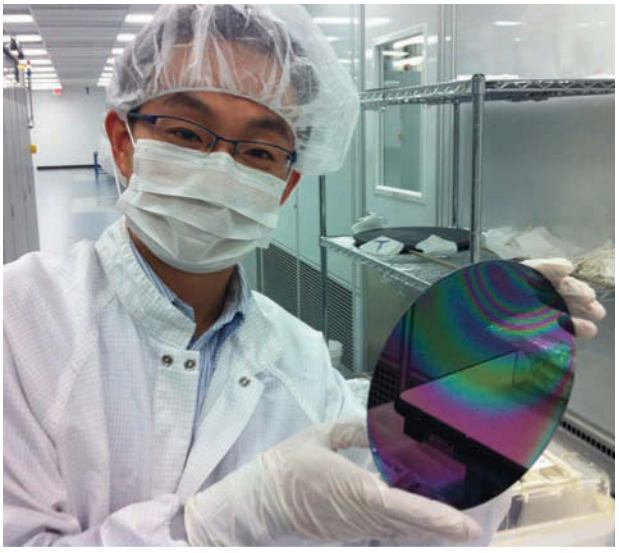

Bridgelux has demonstrated GaN-on-Si HB LEDs with performance comparable with that of LEDs employing SiC or ion-sapphire wafers, an advance that can ultimately result in lower-priced LED bulbs.

"Roughly speaking, LED chips [the emitter itself, unpackaged] cost about 25 cents/chip," says Brad Bullington, vice president of strategy at Bridgelux. "Over the next 24 months, … this technology will give us a path to sub 10 cents on launch [in 2013] and the mid- to low single digits thereafter." Bullington adds that the major barrier to adoption of LED lighting is its upfront cost and that this technology move will bring $5 billion to $7 billion in additional revenue to LED lighting in the 2013 to 2015 time frame.

—by Margery Conner

▷**Bridgelux**, www.bridgelux.com.

Watch out for well-made—but counterfeit—chips

An eye-opening blog post reveals logic-analyzer developer Saleae's experience with counterfeit parts (Reference 1). Saleae co-founder Joe Garrison first became aware that his company had a problem with counterfeits when an unusually large number of boards started failing the functional test. Garrison noticed that the USB (Universal Serial Bus) chip was running hot—not just on the failures but also on the good boards. Because the company uses Asia as a source for these boards, Garrison immediately wondered whether they were counterfeits. Aren't counterfeit parts always duds, though?

"There was little or no information online about this part, but I did uncover some random snippet that said the part was flagged for possible counterfeit activity just two months prior—not a good sign," Garrison says.

Searching further, he learned that a huge number of counterfeit parts are so-called reclaimed parts. To "reclaim" these parts, counterfeiters strip discarded electronics of any high-value components and then resell them as new in fake packaging. This process involves cleaning and relabeling the parts—in some cases with new part or batch numbers. "Apparently, this [behavior] is a staggeringly large illicit business," Garrison explains (Reference 2).

The parts may be from similar functional parts or may be different. Saleae verifies that the USB chip was a reclaimed instance of an ancient version of a part. The report also says that nonfranchised distributors are supplying almost all of the parts. Buyer beware.—by Margery Conner

▶Saleae, www.saleae.com.

REFERENCES

1 Garrison, Joe, "Counterfeit parts are a big headache," *Joe Garrison 2011 Web Edition*, Aug 11, 2011, http://bit.ly/pw32HE.

2 Hughitt, Brian, "Counterfeit Electronic Parts," Trilateral Safety and Mission Assurance Conference, National Aeronautics and Space Administration/European Space Agency/ Japan Aerospace Exploration Agency, April 2008, http://1.usa.gov/nDQSCo.

09.08.11

New! IGBT reference book
Order now under www.infineon.com/igbt-book

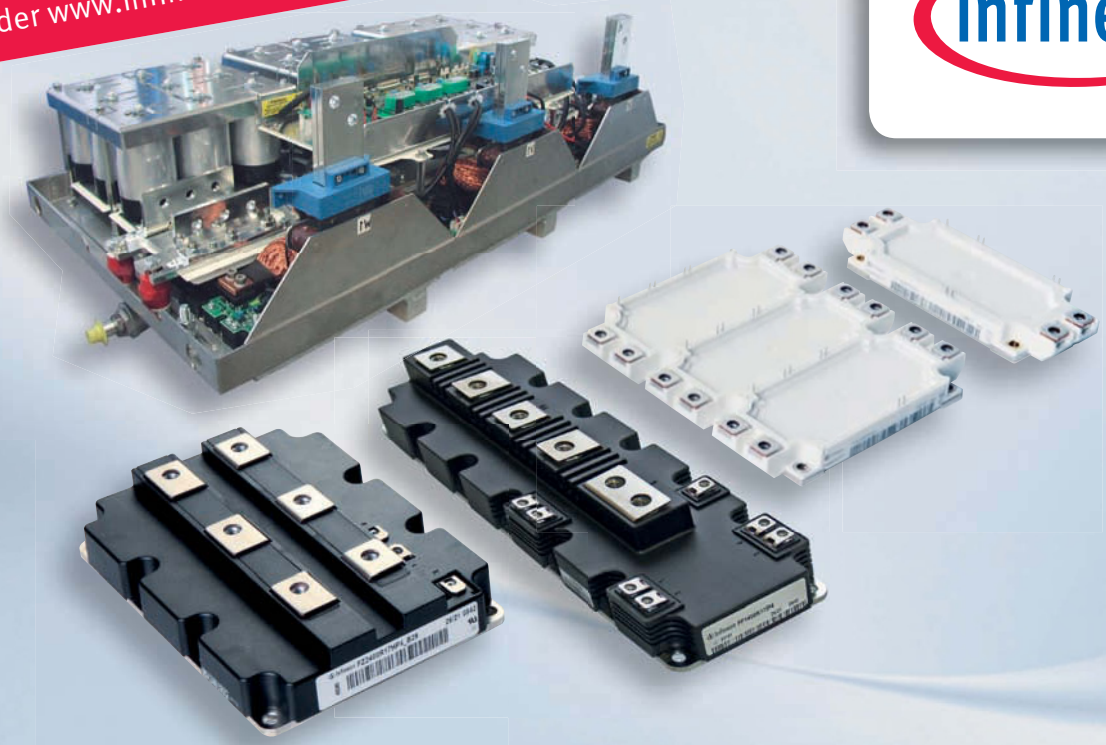

Solutions for windpower systems

Energy-efficient components for high system reliability

The Infineon product portfolio provides components for the highest energy efficiency in windmill power converter and pitch control solutions.

Our Power Modules with newest 1200V/1700V trench fieldstop IGBT4 and Emitter Controlled diode chip technology offer best in class power density solutions in conjunction with extended lifetime. The modules feature low on state losses, optimized soft switching behavior and a wide operation temperature range up to 150°C maximum junction operation temperature. The newly introduced stack assembly ModSTACK™ HD leads to more than 50 % higher power density at same footprint.

The following benefits are provided to our customers:

- Extended module utilization by 150°C maximum junction operation temperature
- Highest power density
- Supreme power cycling and thermal cycling capability

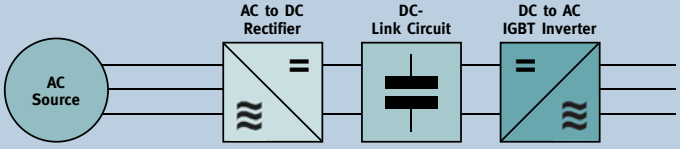

[www.infineon.com/highpower]

New architecture promises better battery

A change in architecture is promising to close the gap between semiconductor technology and battery technology, which has traditionally lagged behind semiconductors due to its dependence on unchangeable chemical reactions. Instead of storing charge in a main battery—then doling it out to individual devices on demand—a new breed of hybrid capacitor/battery is storing just enough energy for an adjacent device for its exclusive use. Ioxus Inc claims to be solving the battery problem by defining a new distributed-energy architecture.

Users are employing the company's hybrid ultracapacitors for all types of applications that were challenging for traditional battery architectures, according to Chad Hall, co-founder and vice president of Ioxus. The device satisfies the needs of short-term or backup power, without the problems of traditional batteries, he claims.

Applications for the technology range from simple to complex. For instance, it allows users to charge a simple flashlight in just 20 seconds using a hybrid ultracapacitor; the flashlight can then continue to operate for as long as two hours, according to Hall. A complex regenerative-braking system on an automobile can instead use a hybrid ultracapacitor that charges whenever you brake and stop. It then simply restarts the car when you hit the gas pedal—eliminating all the pollution that stop-and-start traffic causes.

Automotive applications allow designers to economically distribute hybrid capacitors around a vehicle, storing short-term energy as necessary for powering LEDs, onboard computers, power windows, power door locks, and security systems. In the event of a total failure—or even removal—of the main car battery, all hybrid-ultracapacitor-powered systems still work. Ioxus estimates that electric vehicles using distributed hybrid ultracapacitors rather than a centralized battery can cut 20 to 30 lbs from their weight.

A few other manufacturers claim to have developed similar hybrid ultracapacitors. For instance, Evans Capacitor Co (www.evanscap.com) has a higher-voltage, lower-energy hybrid capacitor, but it is more akin to an ultracapacitor alone than Ioxus' hybrid ultracapacitors. And JM Energy Corp (www.jmenergy.co.jp) offers a lithium-ion capacitor that is more like a lithium-ion battery than Ioxus' hybrid ultracapacitor. Ioxus claims to have intellectual property that makes its approach better than Evans' and safer than JM Energy's.

The ultracapacitors combine the architecture of a capacitor—charge accumulating on plates separated by an insulator—with the chemical storage of a lithium-ion battery. Using the same chemistry as a lithium-ion battery, the device distributes those chemicals onto the surface of a solid electrode, rather than embedding them into a porous electrode, requiring the slow process of intercalation to charge and discharge them.

As a result, users can charge and discharge the devices at nearly any rate, allowing them to absorb and discharge vast or minuscule currents, depending on an application's requirements. However, they can neither store as much charge volume nor store it for as long as a traditional battery, but the applications' distributed architecture compensates for these shortcomings, according to Ioxus.

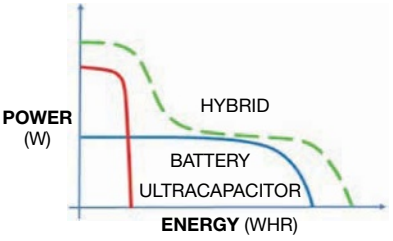

Ioxus hybrid ultracapacitors with built-in lithium-ion batteries enable distributed storage architectures that work better than either technology alone.

Hybrid ultracapacitors store less charge than do lithium-ion batteries, but they store about 100% more charge than an ultracapacitor alone and provide almost unlimited charge/discharge cycling. You can charge and discharge most lithium-ion batteries only a few hundred times, and even special long-term versions have only a few thousand cycles. You can charge and discharge hybrid ultracapacitors, on the other hand, more times than the lifetime of the product in which you are using them—more than 20,000 cycles in Ioxus tests.

—by R Colin Johnson

▷**Ioxus**, www.ioxus.com.

Source code for µC/OS-III is now available

Micrium is making available for evaluation the source code of its latest real-time operating system, µC/OS-III. The move comes as the private company claims it is seeing double-digit growth in sales and calls for evaluation boards using a variety of processors. Micrium in 2009 released its third-generation RTOS with a handful of evaluation boards and a book describing the software's capabilities. The company in 1999 released the previous version of the RTOS, which is available as source code for evaluation.

The latest version of the book on µC/OS-III, now available as a free download, also includes a section with example code for four medical-electronics systems, including a pulse oximeter and a blood-glucose meter. The company aims to expand its efforts to provide technical documentation for other vertical markets. Universities and researchers can freely download and use the company's RTOSs. Commercial users are subject to licensing fees.

Jean J Labrosse, founder and chief executive of Micrium, claims that the company is the only commercial RTOS provider that offers source code for evaluation use. OEMs that roll their own RTOS still represent the largest slice of the market, with those using free, open-source RTOSs coming in second, says Mike Phipps, director of sales, marketing, and business development for the company. However, Micrium has been gaining share since Intel (www.intel.com) purchased Wind River, Phipps claims.

—by Rick Merritt

▷**Micrium**, www.micrium.com.

SET YOUR DESIGNERS FREE WITH CALIBRE® AUTOMATIC WAIVERS | It keeps track of waivers and eliminates them from SoC verification results, saving hours of wasted effort by you and your ecosystem partners. Now your designers can do more productive things. Take a test drive at: http://go.mentor.com/autowaivers.

Mentor Graphics®
THE EDA TECHNOLOGY LEADER

BY PALLAB CHATTERJEE, CONTRIBUTING TECHNICAL EDITOR

Semiconductors develop an EUV ecosystem

The need to drive to smaller geometries is not slowing down, despite the rising cost of building wafer fabs. Lithography has been a major element in this push toward new sub-20-nm process geometries. The shift from 193-nm, DUV (deep-ultraviolet) lithography to 11- to 14-nm, EUV (extreme-ultraviolet) lithography requires a full ecosystem for the materials and methods, along with an imaging module. The ecosystem includes a supply of masks, photoresist, and both mask- and wafer-metrology tools and methods.

The system works by inspecting blanks and associating defects with the "dark-field" patterns of the blanks. The geometry of the defects in question is incompatible with standard optical-inspection techniques. As a result, vendors are using AIMS (aerial-image-message-system), ABI (actinic-blank-inspection), and advanced-PMI (patterned-mask-inspection) tools for defect review. In addition to using these inspection tools, manufacturers must create repair methods. For standard masks, the normal repair technique is to use a 30-kV FIB (focused-ion beam) to fix the patterns on the mask. For EUV masks, due to the wavelength used, underlayer and multilayer damage can occur.

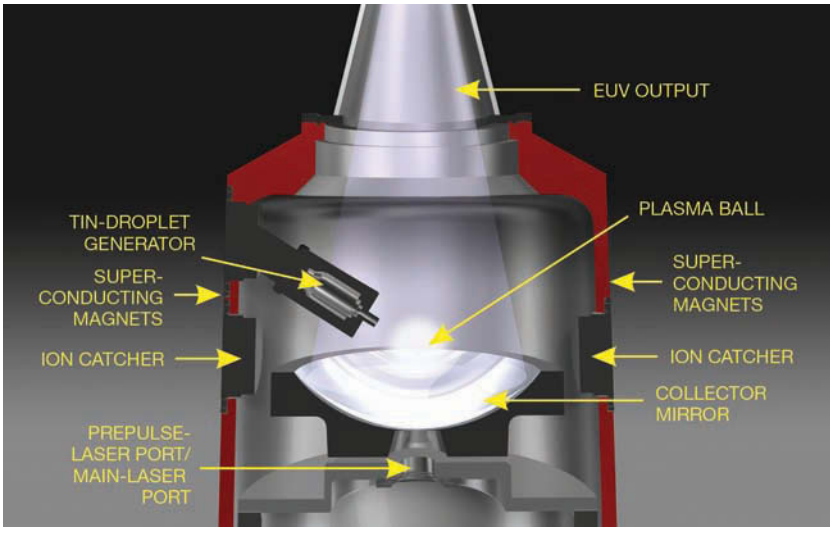

Figure 1 Gigaphoton's EUV source comprises a liquefied-tin source that produces precision 20-micron droplets at a 100-kHz rate. The system then releases and energizes these droplets with a solid-state laser prepulse and hits the resulting energized particles with a high-power laser source, which, 99% of the time, ionizes and releases the EUV pulse. A magnetic system assists a second pulse to enhance the ionization.

Several viable options are available to repair this damage. For example, manufacturers could use a 10-kV FIB, but this method has some selectivity and reflectance issues. They could also try nanomachining with a diamond tip. This method causes little substrate damage but results in a lot of debris and works only on certain patterns. Alternatively, they could attempt e-beam repair, which offers high selectivity but has issues with the high corrosiveness of the xenon difluoride, which clears the debris from the repair. Another major issue with repair and inspection is the throughput, which is currently a factor of 10 behind DUV approaches for similar tasks.

Manufacturers are making progress on the patterning side, however. Gigaphoton Inc (www.gigaphoton.com) now has a 100W EUV source that, in an imaging system, can support a throughput of 50 wafers per hour. The company is developing a 250W source that will support 150 wafers per hour. The EUV source comprises a liquefied-tin source that produces precision 20-micron droplets at a 100-kHz rate. The system then releases and energizes these droplets with a solid-state laser prepulse and hits the resulting energized particles with a high-power laser source, which, 99% of the time, ionizes and releases the EUV pulse. A magnetic system assists a second pulse to enhance the ionization (**Figure 1**). This environment yields a debris deposition rate of 1.2-nm/million pulses, and the chamber-cleaning rate is based on a debris accumulation of 4.4-nm/million pulses.

Another issue is the protection of the masks. To protect them, manufacturers must use exotic techniques, such as a pattern of buffer/absorber/buffer and antireflective caps. Techniques for reticle management must extend to reticle shipping, particle management from the shipping pod, storage, and protection from ESD (electrostatic discharge). EDN

Pallab Chatterjee is on the IEEE Nanotechnology Council.

Introducing ZYNQ, the new element in processing.

Finally, the processor comes together with the FPGA in a fully extensible processing platform called Zynq.™ More intuitive to program in the way you already know. Fully customizable to your requirements. Faster to implement and get to market. As a software engineer, if you know ARM® Cortex,™ you already know Zynq. And if you know Xilinx, you already know this is innovation you can count on.

Visit us at www.xilinx.com

XILINX ZYNQ™

Finding the right solution to your company's materials' needs has never been easier, because now the companies of Brush Engineered Materials have unified under one name, Materion. It allows us to give you greater access to all of the resources and solutions we offer. So why not find out what the new Materion can do for you? Chances are, we can make a big difference to your technologies and your bottom line. Find out more at materion.com.

MATERION

Brush Engineered Materials, Williams Advanced Materials, Brush Wellman, Academy Corporation, Barr Associates, CERAC, Technical Materials, Inc. Techni-Met, Thin Film Technologies, and CERAC are all now MATERION.

CHIP-AND-PACKAGE

CO-DESIGN RELIEVES PRESSURE ON COMPLEX DESIGNS

BY DARVIN EDWARDS • TEXAS INSTRUMENTS

The vise is closing down on design departments. Manufacturers want more capabilities in their products than ever before. It's imperative for manufacturers to remain competitive. Marketing, meanwhile, wants to cram that added performance into less space. Consider all of the handheld devices that are prevalent today. The pressure doesn't stop there, though. Design cycles keep getting shorter and shorter. A long design cycle increases the risk that the manufacturer might miss a market opportunity, and doing so can be catastrophic. Design must do more, in less space and with less time.

CO-DESIGN ACCOMPLISHES **MORE WITH LESS** BY PERFORMING MORE DESIGN TASKS IN A SHORTER DESIGN CYCLE.

IMAGE: ISTOCKPHOTO

Enter co-design. Manufacturers should be constantly examining all aspects of design to verify that they are meeting all of the performance targets for final products. With each successive performance-target verification and fine-tuning, the overall value of the product to the user increases, and the risk of a failed product launch diminishes. The goal is to accomplish more with a shorter design cycle by concurrently executing more design tasks. Of course, one of the ultimate goals is that first prototypes will function to spec, avoiding the dreaded design re-spin. Failed products that require design re-spins are costly on many fronts for both semiconductor manufacturers and their customers.

Designers have for some time employed the notion of co-design—or its predecessor concept, concurrent engineering. Managers assign teams portions of a design to execute in parallel with the rest of the design. What has changed is that co-design is now appearing in every aspect of a design project—from inception; to planning; to integration in the final system; and even to the use-profile assumptions, such as how long a user might play a game or use a smartphone. An area receiving increasing attention is chip-and-package interaction, which analyzes characteristics of the package with the performance of the chip. Designers place more emphasis on chip-and-package interaction as both chips and packages become more complex and as 3-D-packaging technologies more intimately link the interactions among package technologies and the architecture of the system in the package.

CHIP/PACKAGE CO-DESIGN

Chip-and-package co-design concurrently initiates package design with custom die and even system-level design. With this method, a team can quickly bring a new product to market and achieve the performance characteristics required to be competitive. In addition to the normal challenges that revolve around optimizing the performance characteristics of the eventual device, certain foundational challenges are inherent in every chip-and-package co-design effort. These challenges involve people working together and software tools that can share databases to quickly evaluate design characteristics.

AT A GLANCE

- Today's schedules require co-design of all facets of the system.
- Simulation and accurate modeling are vital to co-design efforts.
- An emerging subset of co-design is die/package/board co-design.
- Electrical, thermal, and mechanical issues are all important.

At the beginning of a co-design project, all of the various design disciplines must come together. A chip-and-package co-design effort includes experts on all the elements in the silicon-production process, including transistor-design libraries, I/O (input/output) cells, packaging and materials, manufacturing and assembly, and chip test. Getting such a large group together may not involve a physical meeting of everyone on the team, but it must involve a merging of intentions for everyone. For the project to succeed, all of the key players from across all design areas must share the common goal of creating an optimized end product.

At the onset of the chip-and-package co-design project, representatives of each discipline must understand the chief needs of every other discipline. For example, the electrical-layout engineers must be aware of any thermal limitations in the package that could require a change in the placement of logic or I/O blocks to reduce operating temperatures. In addition, circuit designers must create any analog circuits that are sensitive to package-stress gradients with specific package limitations in mind.

Moreover, packages must supply the required number of I/Os for the chip layout, and manufacturability concerns may dictate specific rules so that the product meets the requirements of high-volume production. In the end, the team must be cognizant of the target cost of the product to avoid choosing too-expensive options. Some concerns of team members may evoke more passionate responses than others, but eventually the entire team must pull together around the common goal, which will come down to delivering a compelling and engaging experience for the end user of the product.

Management's commitment to the project's common goal and to co-design as part of achieving this goal is essential. Conflicts among groups and even individuals who make up the co-design team may occur. Management should encourage constructive debate with the objective of arriving at the most effective approach. The leaders of the effort must occasionally guide the direction of the team when it cannot reach unanimity among conflicting positions.

THE COMMON DENOMINATOR

One of the basic characteristics of every co-design project is the necessity for design teams to concurrently complete their work on a part of the overall design while other teams are developing other aspects of the design. For example, one team may be developing a chip layout while another team is selecting and designing a package construction.

Selection of system parameters, such as the PCB (printed-circuit-board) area

WITH EACH PERFORMANCE-TARGET VERIFICATION AND FINE-TUNING, THE VALUE OF THE PRODUCT TO THE USER INCREASES, AND THE RISK OF A FAILED PRODUCT LAUNCH DIMINISHES.

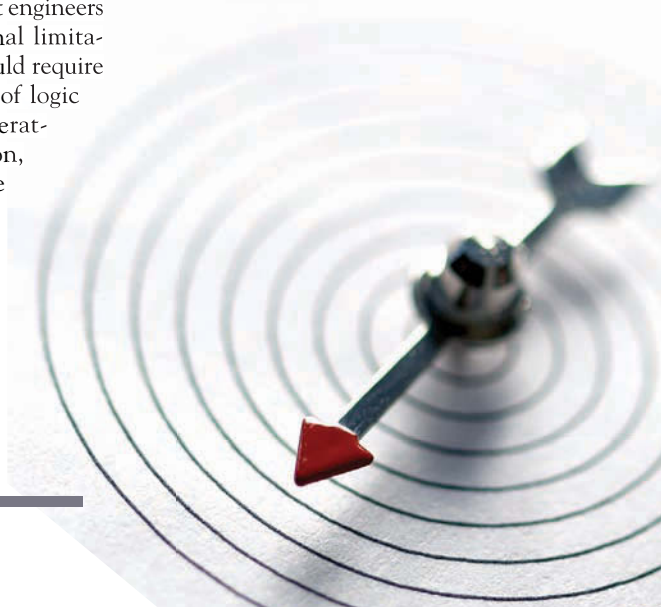

Is your scope lying to you?

Find out.

Rohde-Schwarz-scopes.com

If you're not using a new Rohde & Schwarz digital oscilloscope,
your scope may not be giving you honest answers.

Try the industry's most accurate and reliable digital oscilloscope.

888-837-8772
www.rohde-schwarz-scopes.com

ROHDE&SCHWARZ

that is available for a device, may drive packaging decisions and, in turn, affect die layout. Software simulations might sometimes identify significant power usage in a portion of the circuit, and these simulations could drive the need for thermal enhancements in the package. When tight schedules eliminate any chance of prototyping the design, designers must quickly simulate accurate component connections if the team wants to have multiple options.

Simulation tools generate a model of the chip and its package to analyze how various design options will affect the parametric performance of the device. For example, a simulated model of the device might analyze how a change in the package affects the electrical noise on the power-and-ground plane of the device or how a change in the package might alter the temperature on part of the device. Modeling the effects of packaging options on a device's I/Os is also critical. Not so many years ago, a complex device might involve only several hundred I/Os. Now, the number of I/Os on a complex SOC (system-on-chip) device could easily reach 2000 or 3000, all of which couple together. To end up with the best possible device, the modeling tool must be able to analyze how a change in a package-design parameter, such as a power plane, would electrically affect all of these I/Os.

This type of scenario can be daunting for software-simulation tools. During a co-design project, the speed and abilities of the tools environment are critical. The time window for considering various package innovations is usually narrow. A faster simulation tool produces more models over a shorter period, meaning that the design team can consider a greater

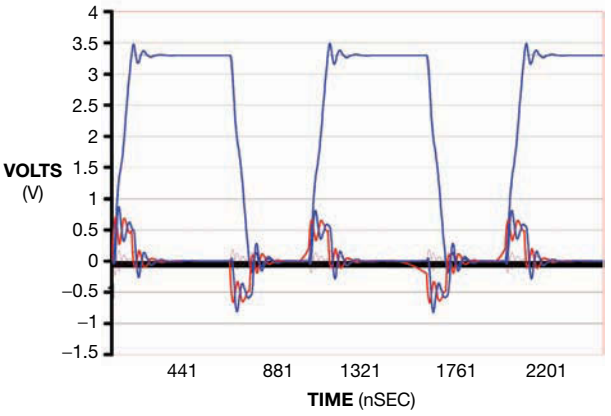

Figure 1 A typical electrical simulation of a signal injected into a package pin shows the distortion of the signal by the package "circuit" and the noise it induces on neighboring leads.

number of options to find the best approach. For example, if the software tools require a week to generate a complete model of the device and the team faces a schedule milestone three weeks away, the team can consider only relatively few alternatives without disrupting the project's overall time line.

In a perfect world, computational resources would not limit simulation. The team could characterize and understand the many interactions among the design parameters. Simulation would explore multiple routing options, material choices, reliability margins, cost minimizations, manufacturability studies, die-size minimizations, and power margins. Development-software systems would automatically adjust the chip, package, and system design until all parameters were optimal.

This scenario is the long-term dream for co-design. In reality, though, current software and hardware systems have limitations. As a result, designers use their expertise to input their best options. A finite number of simulations then point out the areas in which additional optimization may be necessary. Designers change the designs and validate the alterations.

MAKING THE MODEL

Engineers and technicians need time to describe a device, its package, and the system to the various analysis tools, slowing down modeling and affecting the effectiveness of a chip-and-package co-design project. To generate a simulated model of a device, the software tool must have a detailed description or a language-based model of the device. Changes to the device design should be simple to input because this approach eases optimization. Unfortunately, this scenario occurs infrequently.

Describing the device geometries and the package layout to one or more tools can be laborious and tedious. It is difficult to describe the 3-D world of a real chip, package, and system in the words of the computer-simulation tools. Designers must divide the 3-D geometry into many elements through "meshing." Mathematical functions, which relate elements to surrounding elements, describe the parameters for modeling, producing many simultaneous equations for the computer to solve. The accuracy of the solution is dependent on the quality of the subdivision of the 3-D solid and on how

well the individual elements capture the gradients of the fields to be calculated. The physics of a device design could require different types of meshes. Some may need only surface meshes, whereas others require full 3-D representations.

In most instances, models must capture the interfaces of one material to another, sometimes requiring that the mesh be continuous across the boundaries of all components or structures in the design. For example, the mesh of a circuit trace must be continuous with the mesh of the surrounding substrate's polymer to calculate interfacial stresses. In addition, whenever someone makes a change to the design, designers must remesh it.

These tasks often prove too difficult for current software tools without significant human intervention guiding them and supplying knowledge to the software to identify the order the tool should take for meshing complex geometries. Because each area of physics involves different types of meshes, experts often must manually repeat this intensive process during various steps.

This type of human intervention adds little value to the design-optimization task. The software tools should use adaptive meshing to automate this repeated meshing and check the mesh against the physics of the problem to identify where mesh refinement is necessary and to generate an accurate result. Once adaptive meshing and other automation features become standard features of all analysis tools, the time to describe a device will become dramatically shorter.

OPTIMIZED PACKAGING

Time to market is one of the primary forces that drives chip-and-package co-design projects. After all, time is money, and you can never regain missed opportunities in the market. Another good reason for chip-and-package co-design is to achieve a product that excels in the market. Poorly optimized designs don't deliver the performance that users expect, often resulting in product recalls and loss of market share.

For example, a poorly optimized chip design might consume more power than the system can efficiently handle, necessitating more expensive cooling features. A design that uses less power for the same performance will achieve more market success in a portable-system application, in which battery life is a key consideration. Users prefer a smaller portable device with a low profile and small footprint over a larger device with the same features. Moreover, if one chip runs faster than another for the same performance and cost, the faster chip will likely win market share. In general, the chip-and-package co-design team will examine, analyze, and evaluate all of the electrical, thermal, and thermomechanical characteristics of a device and its package.

ELECTRICAL ISSUES

One of the first concerns of chip and package designers is which chip-interconnection technology to use. Flip-chip interconnection can still be more expensive than a wire-bonded approach. A wire-bonded package that is compatible

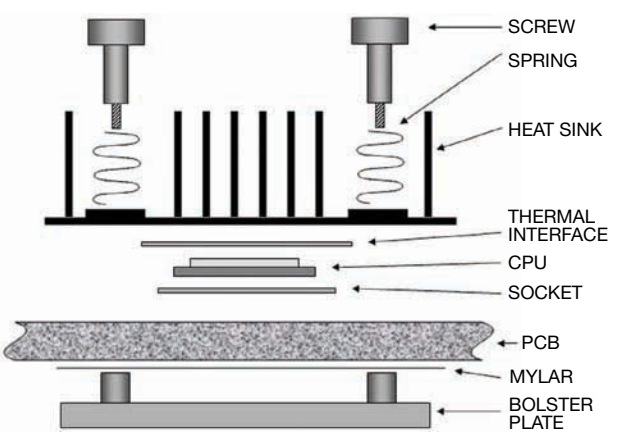

Figure 2 Co-design optimizes a heat-sink design and attachment mechanism for a high-powered microprocessor component using different materials and thicknesses.

with the device's I/O count, speed, and performance specifications can result in significant cost reductions over the lifetime of a device. To select the best alternative, teams do electrical analysis on those products whose performance requirements overlap the capabilities of wire-bond and flip-chip interconnections. The rest of the package design flows from this selection.

A primary concern is whether enough power is getting onto the device for it to function at its specified level. The device's package plays a major role in this regard because the resistance, inductance, and capacitance of the power and ground pins can distort the supply of clean power to the chip, creating power and ground bounce or noise. Chip-and-package co-designers simulate the electrical design of chips and packages with an eye toward analyzing the electrical noise on the power and ground pins and planes.

The simulation software can highlight areas of high noise, inductance, current density, or resistance for optimization. A redesign can then minimize the problems. A simulation sometimes indicates that filtering capacitance is necessary to reduce power and ground noise. These analyses should also consider the power- and ground-distribution planes in the system-level PCB, or at least apply a valid simplification of the system-level planes as a boundary condition.

Engineers should also devote much analysis to simulations of the electrical load on the device's I/Os. Timing and noise analysis will identify a poorly laid-out high-speed I/O on the substrate (**Figure 1**). For example, two single-ended lines that are too close together may induce coupling noise on both lines. Separating these lines or adding a ground between them might be necessary to reduce the signal coupling. In other instances, changing the current-return paths of an I/O to reduce or increase inductance might be necessary. Simulations can also explore manufacturability issues, such as trace-width or substrate-thickness variations, and response versus frequency. The goal is to ensure that the package imparts the least distortion to the signal.

THERMAL CONSIDERATIONS

Every power-consuming device is also a small heater that requires a means of power dissipation. Often, when a component's power consumption exceeds 1W, the thermal needs of the chip drive the package choice. An effective chip-and-package co-design team examines how to best design the device's package to conduct heat away from the device and toward the underlying PCB or to a heat sink, depending on the system application.

Co-design ensures that the package conducts heat where it needs to go, helping to minimize the temperature of hot spots on the die. With BGA (ball-grid-array) packages, for example, thermal vias can conduct heat away from the die and to the PCB. Some package configurations include heat sinks to efficiently conduct heat energy away from the die (**Figure 2**). All devices have maximum-temperature limits. Exceeding these limits causes the device to stop operating or causes damage to the device. Co-design plays a major role in ensuring that devices remain within their temperature limits.

STRESS

Thermomechanical-stress issues as they relate to chip-and-package co-design involve the mechanical interaction of the die, the package, and the system. These issues affect the reliability of not just the chip but also the end-use device. For example, DSPs for wireless base stations must be able to withstand swings from the high temperatures of a desert to the extreme cold of an Alaskan winter or a high mountaintop. These extreme conditions must not cause delamination of the various package layers or fatal fatigue of the solder balls that make the electrical connection to the PCB. The base station is only as reliable as the DSP chips that operate it, and the interaction between the chip and its package affects the DSP chips' reliability.

Thermomechanical analysis reveals whether the package's stress gradient is the best for the chip's parametric performance. This analysis is particularly important with analog chips because stress-related gradients across the device can create shifts in the parametric characteristics of the transistors or resistors, resulting in deviations in the linearity, gain, voltage offset, and other characteristics of the device.

A chip-package analysis of the mechanical robustness of a device might lead to the deployment of a different

material in the package or the inclusion of a structural element in the device. For example, process engineers work to implement new on-chip dielectrics in semiconductors because these materials reduce capacitance between the lines on a chip and enable higher switching speeds. These dielectrics are sometimes less dense and less structurally robust than other materials, however. The chip-and-package co-design team would need to simulate a new dielectric in the device to analyze whether the chip would be able to withstand a certain level of mechanical stress. Dangerously high levels of stress in the predictions might trigger a materials change or a new design rule to structurally strengthen the device.

Warping is another important concern for thermomechanical co-design engineers. Packages are laminates of multiple layers of different materials, each with its own expansion and stiffness. Passage through a temperature cycle can cause warping, which can in turn cause the package's leads or solder balls either to not contact the PCB during soldering or to put too much pressure on the joint, which could result in solder bridging. Thermomechanical analysis can predict this warping and point out the need for different materials or structures if the warping is too severe. Shadow moiré, a measurement tool, further characterizes the warping and validates the thermomechanical modeling.

CALIBRATION AND VALIDATION

In many regards, chip-and-package co-design depends on modeling and other tools to simulate either portions of a design or the entire device so that the design team can study the effects of various design options, eventually resulting in a fully optimized design. The outcome of a chip-and-package co-design effort will be successful only if the tool's outputs are accurate. A tool that returns an inaccurate result can wreak more havoc on a chip-and-package co-design project than having no result at all. The design team could take the wrong path by using faulty simulation information and not realize the mistake until the project has progressed significantly along its time lines. The team can never recoup the wasted time.

Because of these issues, experienced co-design teams constantly calibrate and validate the tools that are essential to the process. The most fundamental way of completing this task is to analyze the characteristics of devices both before and after manufacturing. Any discrepancy indicates a tool in need of recalibration.

When feasible, designers can plan test dice early in the process to place known stresses on sensitive components. These test dice can contain structures such as strain gauges, thermal sensors, leakage sensors, and variations on design rules. The test structures come early in the development process to enable calibration of the simulation tools before designers must use the tools for the driver products.

As investigators stress the test structures, they compare the devices' outputs with the results from the simulation tools. When discrepancies exist, some degree of destructive analysis further characterizes the test structures and identifies the cause of the discrepancy. When the team identifies the causes, the members adjust the tools to match the test-structure results.

CO-DESIGN'S IMPACT

The pressures on design teams will continue to grow in the foreseeable future as the electronic content in everyday devices, machinery, office equipment, consumer gadgets, medical systems, home appliances, and transportation vehicles accelerates in the years ahead. This acceleration heightens the design risks for system manufacturers and highlights the critical importance of all aspects of co-design. Extending the well-documented benefits of co-design into the realm of chip-and-package design will become increasingly imperative with so much riding on each design project. EDN

AUTHOR'S BIOGRAPHY

Darvin Edwards is a Texas Instruments fellow and manager of TI's semiconductor-package-modeling group. He is responsible for thermal, electrical, and stress analysis for a wide range of TI product families. In 1980, Edwards earned a bachelor's degree in physics from Arizona State University (Tempe, AZ). He has authored and co-authored more than 40 papers on IC packaging, has written two book chapters, and holds 17 US patents.

DSOs: THE BANNER SPECS TELL ONLY PART OF THE STORY

BY DAN STRASSBERG • CONTRIBUTING TECHNICAL EDITOR

If you are contemplating the purchase of a new scope for debugging high-performance embedded systems and you want to know the amount of time the scope will require to capture ephemeral, aberrant waveforms associated with seemingly random SUT (system-under-test) failures, you're unlikely to find the scope's banner specs enlightening. Performance in capturing anomalous waveforms depends on little-known and often poorly understood instrument characteristics. Depending on how you plan to use the scope, you might discover that the headline specifications fail to reveal much of the information you need about important aspects of the instrument's behavior (see **sidebar** "Running into a brick wall").

ILLUSTRATION: DAVID NICASTRO; CHECKLIST IMAGE: ISTOCKPHOTO

TO MAKE AN INTELLIGENT CHOICE OF A SCOPE, YOU MUST KNOW ITS BANNER SPECS—BANDWIDTH, SAMPLING RATE, AND MEMORY DEPTH. BEFORE YOU SIGN A PURCHASE ORDER, THOUGH, YOU'D BE WISE TO FIND OUT MUCH MORE.

Moreover, if you want one number that describes the scope's performance in capturing transitory waveforms, the manufacturer may be only too happy to supply one; the number you get from the field engineer, however, may prove to be irrelevant to your application. Such application-dependent parameters as waveform duration, sweep time, and sampling rate can alter aberrant-waveform-capture performance by orders of magnitude. The reason is that most real-time-capture DSOs (digital-storage oscilloscopes) spend large portions of time preparing to display the waveform data they have just captured.

Converting a series of digitized samples—that is, a record—into a recognizable and informative display is more complicated than you might think. Even though most modern high-performance real-time-capture DSOs segregate the display memory from the capture memory, the instruments cannot act on new trigger commands while they are preparing the display (see **sidebar** "A funny thing happened on the way to the screen").

Some industry participants call the interval during which a scope cannot process trigger commands sleep time, blind time, or dead time. However, during the intervals when it cannot process trigger commands, the scope is neither asleep nor blind, and it most certainly isn't dead. More optimistic and convenient terms are aware time and percentage of aware time. Although the percentage of aware time can be less than 1%, it is the number that you probably want to know and use in subsequent calculations.

A related term, waveforms per second, is the number of waveforms per second the scope can display and is the display-speed spec about which scope manufacturers are most likely to boast. Unfortunately, like blind, sleep, and dead times, the waveforms-per-second spec depends strongly on such application-dependent parameters as waveform duration, sweep time, and sampling rate (see **sidebar** "DPX, the original high-waveform/second technology).

AT A GLANCE

- A scope's performance in capturing anomalous waveforms depends on little-known and often poorly understood characteristics. The headline specifications may fail to reveal much of the information you need about important aspects of its behavior.
- In everyday use, scopes whose fast updates give rise to high awareness-time percentages provide users with a "live" feel, as opposed to the frustratingly sluggish behavior of scopes that update more slowly.
- To avoid aliasing, you must sample at more than twice the highest-signal-frequency component whose amplitude is no greater than the ADC's least-significant-bit weight.
- For years, the standard for reconstruction filters in scopes had been gaussian response. Now, some scope manufacturers provide a choice of reconstruction-filter characteristics; you can choose the characteristic that best suits the measurement.

LeCroy's four-channel, 400- and 600-MHz-bandwidth WaveRunner 6 Zi HROs are unusual not only for their 12-bit ADC resolution but also for their pivoting displays. Another LeCroy family, the WaveMaster 8 Zi-A, uses patented digital bandwidth interleaving to achieve industry-leading 45-GHz bandwidth on one channel at 120G samples/sec. Bandwidth is 30 GHz/channel at 80G samples/sec on two channels and 20 GHz/channel at 40G samples/sec on all four channels.

WAVEFORMS/SECOND

The scopes that provide the fastest update rates without requiring the use of special operating modes are units from Agilent and Rohde and Schwarz that offer –3-dB bandwidths in the 1- to 2-GHz range. Both companies' instruments top out at slightly more than 1 million waveforms/sec at 10 nsec per division but drop to approximately 1 waveform/sec at 100 msec per division. These values apply to only a few models from each company. Don't assume that the vendors' other models are equally fast or that selecting a higher-priced model automatically gets you faster waveform updates. Ask the vendor to supply a complete table of update-speed specifications for the instrument model you are thinking of purchasing.

Currently, the scopes with the highest update rates have 10 to 20% aware

Rohde and Schwarz's RTO series includes a custom ASIC that enables the capture of more than 1 million waveforms/sec.

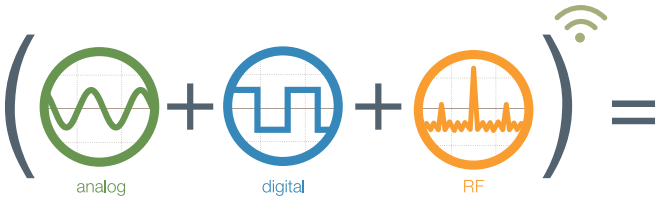

The World's First Mixed Domain Oscilloscope

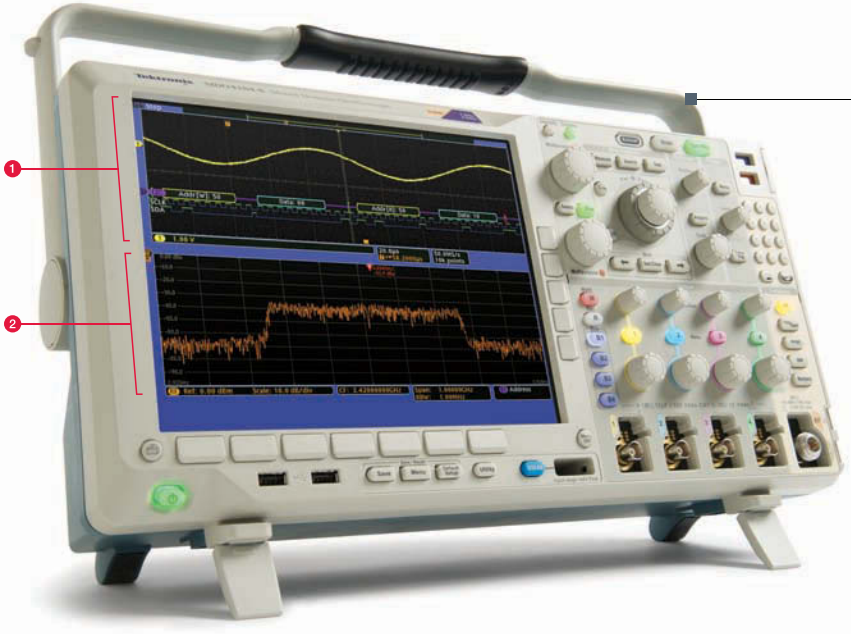

MDO4000 Mixed Domain Oscilloscope

1 Time Domain 2 Frequency Domain

- 4 analog channels
  - 500 MHz and 1 GHz bandwidth models
- 16 digital channels
- 1 RF channel
  - 50 kHz–3 GHz and 50 kHz–6 GHz frequency range models
  - Ultra-wide capture bandwidth up to 3 GHz
  - Unique RF analysis tools: automated markers, spectrogram display, RF vs. time traces, advanced RF triggers
- Parallel bus triggering and analysis, included standard
- Serial bus triggering and analysis options
- Built on the award-winning MSO4000B mixed signal oscilloscope platform

Introducing the MDO4000 Series from Tektronix, the revolutionary oscilloscope with a built-in spectrum analyzer. It's more than just a new scope—it will transform the way you test. Capture time-correlated analog, digital and RF signals for a complete system view of your device. See both time and frequency domains in one glance. View the RF spectrum at any point in time to see how it changes. Quickly and efficiently solve the most complicated design issues—with an oscilloscope as integrated as your designs. Two domains. One remarkable scope. Only from Tektronix.

Take a look at how we've transformed testing: See the scope in action, analyze the specs and learn more at **www.tektronix.com/revolutionary**.

Tektronix®

time at 10 nsec/division. According to Agilent and Rohde and Schwarz, under conditions comparable with those at which the fast-update scopes exhibit such aware-time percentages, it is common to find other scopes with times of 1% or less. Aware times of 10% are good, and times of 20% are exceptional. At present, aware times greater than approximately 20% are unheard of.

Manufacturers that have achieved the highest aware-time percentages emphasize the importance of achieving this characteristic without resorting to special operating modes. They point out that, in everyday use, scopes whose fast updates give rise to high aware-time percentages provide users with a "live" feel, as opposed to the frustratingly sluggish behavior of scopes that update more slowly.

You find the highest aware-time percentages in scopes with –3-dB bandwidths in the 1- to 2-GHz range. Although such scopes require sophisticated ASICs to achieve the fast waveform processing that enables their high percentages, wider-bandwidth scopes would need even more complex waveform-processing ASICs. At their necessarily higher sample rates, these scopes would require deeper waveform memories to capture records of the same duration as those that the lower-bandwidth

RUNNING INTO A BRICK WALL

Aliasing creates frequency components in a signal record that don't exist in the original signal. Once the aliased components have become part of the record, you can't easily distinguish them from real information; hence, you can't reliably remove them. A common misconception is that aliasing won't occur if the signal's –3-dB bandwidth is less than half the sampling rate. To avoid aliasing, you must sample at a rate at least two times the f_{LSB}, the highest signal-frequency component whose amplitude is no greater than the ADC's LSB. Many octaves can separate f_{LSB} from the –3-dB frequency, and f_{LSB} is usually higher.

To take advantage of the high stopband insertion loss, you might like the idea of a brick-wall antialiasing filter, which exhibits steep attenuation-versus-frequency characteristics in the transition band between the passband and the stopband. However, in wideband real-time-sampling applications, such as high-performance DSOs (digital-storage oscilloscopes), brick-wall antialiasing might be practical only in those situations that involve oversampling, in which the scope acquires many more than two samples in the period of the LSB's frequency. In high-performance, real-time-sampling DSOs, whose sampling frequency exceeds 20 GHz, the ratio of the sampling frequency to the –3-dB frequency ranges from approximately 2-to-1 to 3-to-1. Although real-time scopes that contain ADCs that take 100G samples/sec and more have recently become available, it isn't yet possible to construct physical ADCs that take, say, 200G samples/sec in real time. Thus, among high-performance DSOs, the only possible candidates for using brick-wall-antialiasing filters would appear to be instruments that support RETS (random equivalent-time sampling), sometimes also called random interleaved sampling (see sidebar "A funny thing happened on the way to the screen"). By sampling only repetitive signals, albeit not necessarily periodic ones, RETS can create waveform records that mimic those that would be created by an ADC that samples in real time at a rate much higher than that of any currently available device.

RETS is neither for every designer nor for every application; you must understand its possible pitfalls before you attempt to apply it. What's more, as a practical matter, brick-wall filters are DSP-based and exist only in the digital domain. The data they process must therefore be in digital form, which means that a brick-wall filter must follow—not precede—the ADC. Reconstruction filters are different; scopes can include a DSP-based brick-wall reconstruction filter, and reconstruction filters with brick-wall characteristics are common in high-performance scopes. In cases that involve sampling at the limits of ADC capabilities, however, antialiasing filters must be analog circuits.

These filters affect the waveforms you observe. If you are interested in how your SUT (system under test) handles very-fast signal edges, you should look closely at the characteristics of your scope's reconstruction filters. In other words, the step response that you observe can have more to do with the response of your scope than with the response of the SUT.

The standard for reconstruction filters in scopes for years has been gaussian response, which corresponds to a constant time delay at all passband frequencies and necessitates a gradual roll-off of the sine-wave-amplitude response as the signal frequency increases. Now, however, some scope manufacturers provide a choice of reconstruction-filter characteristics; you can choose the filter characteristic that best suits the measurement. Gaussian filters tend to slow the displayed rise times, whereas properly chosen brick-wall filters do not. Unlike gaussian filters, however, brick-wall filters introduce their own artifacts, most notably preshoot that precedes the beginning and overshoot that follows the end of signal transitions, or steps. To provide the best of both worlds, reconstruction filters can even combine constant delay at lower frequencies with a brick-wall characteristic in the transition band.

Experts argue that preshoot—a response to a stimulus that occurs before the stimulus reaches the SUT—is a purely mathematical artifact that cannot exist in real physical systems, whereas overshoot often is real, although the overshoot your scope displays may not accurately represent your SUT's actual response. The scope may introduce artifacts that can produce erroneous results in eye-diagram tests, a widely used technique for estimating the BER (bit-error ratio) of data-transmission channels.

scopes capture. As ASIC technology advances and allows higher-density chip designs, higher-performance scopes will take advantage of the improved technology, and update rates of 1 million waveforms/sec and more will appear in high-performance scopes—those that have –3-dB bandwidths of 10 GHz and higher. Manufacturers won't commit to a timetable for such advances, however. LeCroy WaveMaster 8Zi scopes, with bandwidth greater than 10 GHz, capture 1 million waveforms/ sec—but only in the sequence mode.

A "SOMETIME" THING

Suppose that the system you are designing exhibits the following mysterious but not uncommon behavior: While performing its intended function, at intervals that you have determined occur, on average, approximately twice an hour, the system fails to do what you have designed it to do. Then, with no input from you, it performs normally again. After a period that could be as short as five minutes or as long as five hours, it misbehaves again. You try to determine whether the average interval between the instances of misbehavior has changed, and you eventually satisfy yourself that your original estimate of an average of once every 30 minutes still seems about right.

Now suppose that you want to get a look at the waveforms at key points in the system so that you can determine how those that accompany misbehavior differ from the normal waveforms at the same probe points. If your four-channel DSO's aware-time percentage is, say,

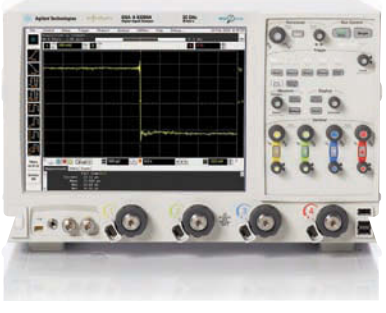

The Agilent 90000 X-Series delivers 32-GHz true analog bandwidth on two channels and 16 GHz on four. This unit samples at 80G samples/sec/channel with two channels active and 40G samples/sec/channel with four channels in use.

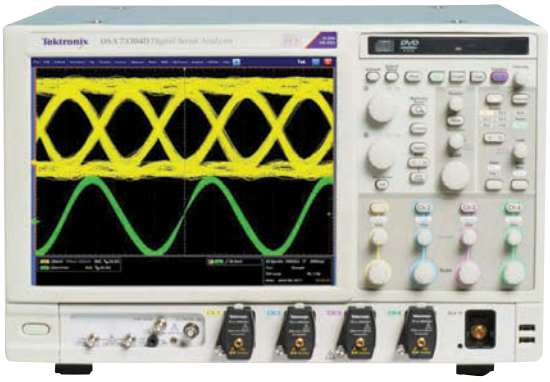

Tektronix's high-performance real-time DPO/DSA 70000D series has a bandwidth of 33 GHz and takes 50G samples/sec/channel on four channels or 100G samples/sec/channel on two channels. For fast sampling of repetitive signals, it offers RETS, which, in effect, captures 10,000G samples/sec/channel.

FOR MORE INFORMATION

Agilent Technologies
www.agilent.com

LeCroy Corp
www.lecroy.com

Rohde and Schwarz
www2.rohde-schwarz.com

Tektronix Inc
www.tektronix.com

Yokogawa Corp of America
http://tmi.yokogawa.com/products/oscilloscopes/

0.1%, it might, with luck, immediately capture the four ectopic waveforms. On the other hand, because of your scope's low aware-time percentage, you could reasonably expect to have to wait for 500 hours, or approximately three weeks of round-the-clock testing, to capture just one set of ectopic waveforms.

In contrast, if your scope has an aware time of 10%, you could expect to capture a set of ectopic waveforms in five hours. Because the failures occur at random intervals whose duration averages 30 minutes, you could have to wait considerably longer, or, with luck, you could capture your first set of ectopic waveforms in much less than five hours. The point, though, is that the scope can capture transitory events only when it is aware of trigger commands. If it is aware of triggers only a small fraction of the time, you must compensate for its inattentiveness by increasing the duration of your test.

Under certain circumstances, some DSOs that do not normally have high aware-time percentages provide a way to, in effect, increase their percentages without resorting to the use of specialized waveform-processing ASICs. Most manufacturers refer to this mode as segmented memory, although some manufacturers use names such as FastFrame or fast acquisition. Suppose that the scope has a waveform-memory depth of 100 Mbytes/channel. Segmented memory allows you to divide that memory into multiple shorter segments. You might, for example, be able to divide the memory into segments whose depth is 1 Mbyte/channel or some other convenient value. The scope can then acquire 100 sweeps per channel before it begins waveform processing. Therefore, if you are capturing signals that occur in rapid succession, the scope will be aware of triggers almost without interruption. The downside is that, at the required sampling rate,

SOME DSOs PROVIDE A WAY TO INCREASE THEIR PERCENTAGES WITHOUT RESORTING TO THE USE OF SPECIALIZED WAVEFORM-PROCESSING ASICs.

A FUNNY THING HAPPENED ON THE WAY TO THE SCREEN

The waveform record of a DSO (digital-storage oscilloscope) comprises a series of digital representations of analog samples, in which each sample corresponds to the signal voltage at a specific instant. Equal amounts of time separate each instant from those that immediately precede and follow it. This statement is true even though the scope might not have captured the samples sequentially—for example, the scope might have used RETS (random equivalent-time sampling), also called random interleaved sampling, to increase its effective sampling rate beyond its maximum real-time sampling rate (Reference A). In RETS, the scope obtains samples during multiple waveform iterations and keeps track of precisely how long after the most recent trigger event it captured each sample. The scope then arranges the samples as though it had captured them sequentially during one waveform iteration.

A little thought reveals that brief waveform records—that is, records with fewer points than the number of pixel columns in the display—must do some work to produce meaningful and easily interpreted displays. If the oscilloscope were to place the samples on the screen without processing, the individual samples would constitute a series of dots. To be recognizable as a waveform, the scope must convert these dots into a series of—ideally—curvilinear line segments, each joined to its immediate neighbors (Reference B).

At the other end of the density spectrum, the number of samples in a waveform record can greatly exceed the display's number of pixel columns. The scope can accommodate this possibility by requiring you to scroll horizontally through many screen widths to see the complete record. This arrangement is fine—and necessary—for obtaining a close-up view, but most scope users also want the big picture; they want to be able to see the entire record without scrolling.

To present long records on a relatively narrow screen without scrolling, the scope must combine multiple samples into single pixel columns. In bygone days, scopes would light all pixels in the range from the lowest to the highest value in the group of samples that were to appear in one pixel column. The result was a bright, thick trace, to which scope users objected because it failed to reveal how often the several illuminated pixels should have been lit. The solution was to make each pixel's intensity or color indicate its "popularity." Such intensity- or color-graded displays are more informative and easier to understand, but producing them forces the scope to perform rather complex math. The math takes time, and the initial result was decreased display-update rates. Ultimately, though, implementing the calculations in custom ASICs enabled faster updates despite the increased information density.

REFERENCES

A Pupalaikis, Peter J, "Random-interleaved sampling (RIS)," Lecroy Corp, http://bit.ly/qSmep2.

B "Waveform Reconstruction (Interpolation)," Agilent Technologies, http://bit.ly/pVhEP5.

Read two more sidebars—"Exclusion triggering: A worthwhile feature?" and "Oops! My scope wasn't looking"—online at **www.edn.com/110908cs**.

DPX, THE ORIGINAL HIGH-WAVEFORM/SECOND TECHNOLOGY

More than a decade ago, Tektronix introduced DPX (digital-phosphor technology), an innovative approach to increasing DSO (digital-storage-oscilloscope) display-update rates (Reference A). The company continues to offer the technology in its midrange and high-end digital scopes and has made it a cornerstone of its DSP-based RTSA (real-time-spectrum-analyzer) family.

DPX uses a custom ASIC to create an intensity- or color-graded pixel map in the scope's memory but does so in a unique manner: Newly arriving samples go immediately into the map, in which a register or a bin represents each pixel; the bin stores a number that corresponds to the number of hits the pixel has received since the most recent trigger event. When operating in the DPO (digital-phosphor-oscilloscope) mode, the scope does not preserve a sequentially ordered record of the sample values, thereby avoiding much of the overhead of generating the display and allowing update rates that, at approximately 300,000 waveforms/sec in Tek's highest-performance scope models, remain among the fastest available.

Almost from the day it debuted, however, DPX drew fire from competitors, who could not copy the technology without infringing on Tek's patents. The criticisms were not entirely unjustified: The loss of the sequentially ordered sample record precludes useful postacquisition analysis. The user must decide whether to operate the scope as a conventional scope or as a DPO and, after making that decision, cannot change it until the start of a new acquisition. DPO mode is useful mainly for short records. If you want to capture long records, you should select a different mode. Tektronix says that the DPO mode in its scopes is invaluable for identifying transitory ectopic waveforms; once you have seen, saved, and perhaps printed out a picture of what you are looking for, you can operate the scope as a more conventional DSO, trigger on the ectopic waveforms you have identified and characterized, and do all of the postacquisition analysis you want.

REFERENCE

A "See a World Others Don't: The Story Behind DPO, DPX, and Waveform Capture Rates," Tektronix Inc, 2011, http://bit.ly/oBsZ4X.

the number of samples that represents one waveform must not exceed your selected segment depth. EDN

ACKNOWLEDGMENT

Portions of this article are based on material that the author developed for a column he briefly wrote for EDN*'s Scope Expert micro site, http://scopes.edn.com.*

REFERENCES

1 Lim, Jit, "Understanding the impact of digitizer noise on oscilloscope measurements," *EDN*, April 7, 2011, pg 18, http://bit.ly/o6wZ11.

2 Nelson, Rick, "Vendors target oscilloscope sweet spot," *EDN*, Sept 9, 2010, pg 30, http://bit.ly/ng7Y7v.

3 "The Impact of Digital-Oscilloscope Blind Time on Your Measurements," Rohde and Schwarz GmbH, August 2010, http://bit.ly/ptC9KU.

4 Strassberg, Dan, "InP ICs halve jitter, double true analog bandwidth of real-time digital scopes," *EDN*, May 13, 2010, pg 14, http://bit.ly/q8W6LW.

5 Rowe, Martin, "Handheld oscilloscopes offer varied features," *EDN*, Nov 8, 2007, pg 85, http://bit.ly/rndVHd.

6 "Analog and Digital Filtering for Anti-Aliasing," Microstar Laboratories, http://bit.ly/ouB4Po.

You can reach Contributing Technical Editor Dan Strassberg at StrassbergEDN@att.net.

“I need the fastest DMM they make.”

They don’t make them any faster.

Nobody but Agilent makes a range of DMMs this fast, this accurate, or this reliable. Up to 1000 times more readings per second than the nearest competitor’s, and far easier to use, you’ll rip through tests in a fraction of the time. It’s what you’d expect from the leader in DMM technology.

DMM	Digits	DC Accuracy	Max Readings	Function/Range Changes	IO
34405A	5 ½	0.0250%	19 / sec	0.2 sec	USB
34401A	6 ½	0.0035%	1,000 / sec	.02 sec	GPIB, RS-232
34410A	6 ½	0.0030%	10,000 / sec	2.6 ms	GPIB, USB, LAN (LXI)
34411A/ L4411A	6 ½	0.0030%	50,000 / sec	2.6 ms	GPIB, USB, LAN (LXI)
34420A	7 ½	0.0030%	250 / sec	.02 sec	GPIB, RS-232
3458A	8 ½	0.0008%	100,000 / sec	3.0 ms	GPIB

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

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

Download the latest measurement brief and tips
www.agilent.com/find/fastestdmm

Agilent Technologies

BY SACHIN GUPTA, PRAKHAR GOYAL, AND KURIAN POLACHAN
CYPRESS SEMICONDUCTOR CORP

Designing reliable capacitive touch interfaces

UNDERSTAND HOW TO ENHANCE THE HARDWARE AND SOFTWARE RELIABILITY OF THIS INCREASINGLY POPULAR USER INTERFACE FOR APPLIANCES AND INSTRUMENTS.

Capacitive sensing offers an intuitive and robust interface that increases product reliability by eliminating mechanical parts in many appliances, or white goods, and instruments. Because of their experience with personal-electronics devices, many consumers are used to touch interfaces employing capacitive sensing, and they have come to expect these interfaces to be reliable and to operate accurately.

Environmental noise and other factors can affect capacitive technology, however, causing systems to be unresponsive to finger touches or to trigger false touches. Developers must fine-tune these sensors or risk a severe reduction of accuracy and reliability. By understanding how capacitive sensors work and how you can design them to be self-tuning to compensate for noise, you can build robust systems that make appliances more reliable, more cost-effective, and easier to use.

CAPACITIVE SENSING

To understand the challenges behind designing a robust user interface, it helps to first take a brief look at the technology behind a capacitive-measurement system (**Figure 1**). To sense the presence of a finger, a capacitive-sensing system must first know the sensor capacitance in the absence of a finger, or parasitic capacitance (**Figure 2a**). When a finger approaches or touches the sensor, the sensor's capacitance changes, resulting in another capacitance, the finger capacitance, in parallel to the parasitic capacitance (**Figure 2b**). The following **equation** yields the total sensor capacitance in the presence of a finger: $C_X=C_P\pm C_F$, where C_X is the total sensor capacitance, C_P is the parasitic capacitance, and C_F is the finger capacitance.

To analyze the sensor capacitance using a microcontroller, you must convert the sensor capacitance into a digital value. Several methods are available for measuring sensor capacitance. One type uses a switched-capacitor block that emulates the sensor capacitance using an equivalent resistance, a programmable current source, an external capacitor, and a precision analog comparator (**Figure 3**). The programmable current source continuously charges the external capacitor until its voltage crosses the reference voltage and the comparator's output is high. The programmable current source then disconnects, and the external capacitor discharges through the resistance until the capacitor's voltage drops below the reference voltage. The comparator's output is now low until the capacitor again charges to the reference voltage. The sen-

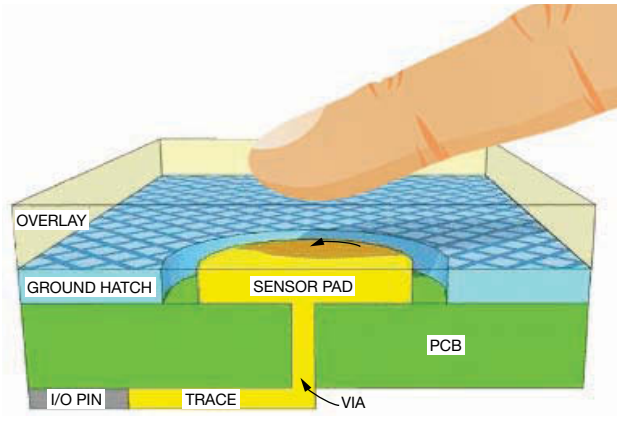

Figure 1 To understand the challenges behind designing a robust user interface, it helps to first take a brief look at the technology behind a capacitive-measurement system.

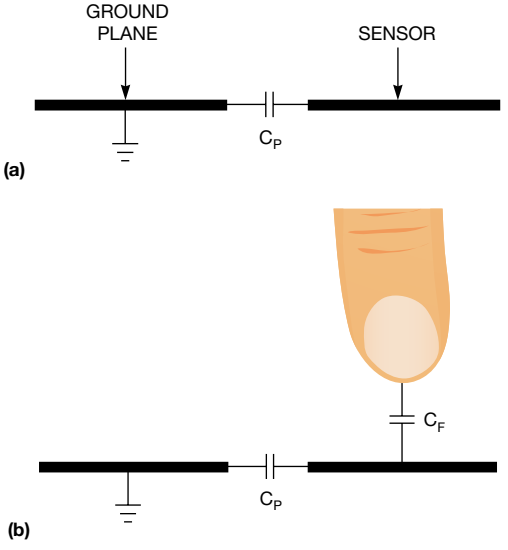

Figure 2 To sense the presence of a finger, a capacitive-sensing system must first know the sensor capacitance in the absence of a finger, or parasitic capacitance (a). When a finger approaches or touches the sensor, the sensor capacitance changes, resulting in another capacitance, the finger capacitance, in parallel to the parasitic capacitance (b).

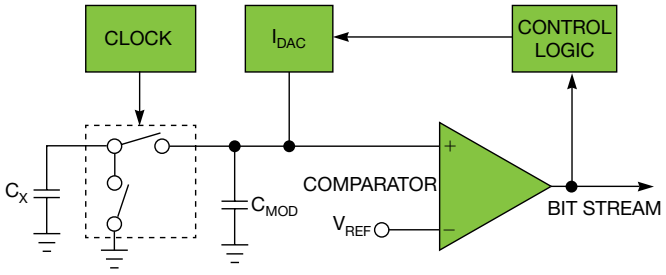

Figure 3 One method for measuring sensor capacitance uses a switched-capacitor block that emulates the sensor capacitance using an equivalent resistance, R_{EQ}; a programmable current source, I_{DAC}; an external capacitor, C_{MOD}; and a precision analog comparator.

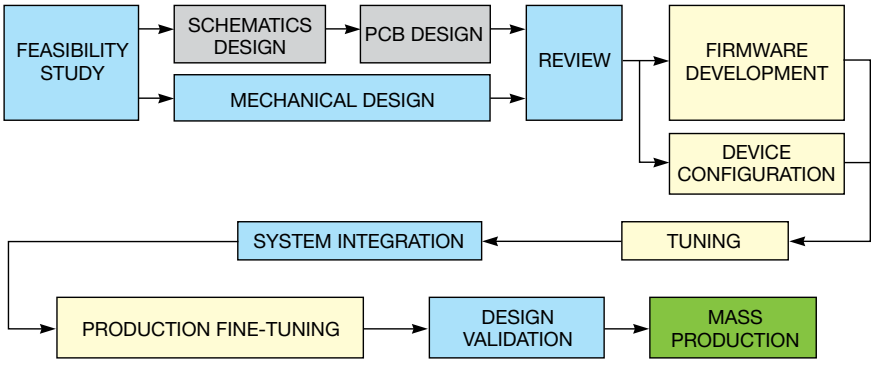

Figure 4 The design flow for a capacitive-sensor touch interface must account for real-world conditions, in which variations occur in components, environmental operating conditions, and noise.

sor's capacitance is greater in the presence of a finger, and the emulated resistance is lower, according to the following **equation**: $R_{EQ}=1/F_SC_X$, where R_{EQ} is the equivalent resistance and F_S is the switching frequency of the switched-capacitor block.

Thus, when a finger is present, the external capacitor discharges faster, and the comparator's output stays high for a shorter time, meaning that a higher capacitance value corresponds to a shorter high time for the comparator. You can feed the resulting bit stream to a counter for a fixed amount of time. This counter value, or raw count, provides an indication of the magnitude of the sensor's capacitance. The fixed amount of time for which the counter counts also determines the number of raw counts, or resolution. When you increase the resolution, the counter counts for a longer period, thus increasing the raw count. In other words, resolution is also the highest possible number of raw counts.

TUNING

Figure 4 shows the design flow for a capacitive-sensor touch interface. However, capacitive sensors must operate in the real world, in which variations in components, environmental operating conditions, and noise can affect sensor performance and reliability. Tuning is a critical process for ensuring that a sensor functions correctly and consistently. You achieve this goal by identifying and determining optimum values for a set of sensor parameters to maintain a sufficient SNR (signal-to-noise ratio) and finger threshold. In general, a 5-to-1 SNR is the minimum requirement for a robust sensor design (**Figure 5**). To avoid false triggering due to changes in capacitance resulting from atmospheric changes, a finger threshold of 65 to 80% of the signal strength ensures reliable finger detection.

Although sensor-controller manufacturers provide guidelines to aid engineers in tuning, achieving the ideal tuning parameters for the system involves iteration. For a sensor controller with a capacitive-sensing algorithm, the tuning procedure follows a particular set of steps (**Figure 6**).

Developers can implement tuning parameters either by writing code specific to the operation of the sensors in firmware, through external components, or by configuring the controller. With a firmware approach, developers have flexibility; however, whenever tuning parameters require changes, developers must also modify and update the firmware.

Alternatively, designers can simplify system-firmware development by using a fixed-function, nonprogrammable capacitive-sensor controller. In this case, designers must implement tuning parameters either by using external components on the board or by sending configuration data over a communication interface, such as an I^2C (inter-integrated circuit). With this approach, whenever tuning parameters require changing, the developer must either rework the user-interface board or update the configuration data. Developers must be aware that tuning can be time-consuming, especially if the PCB (printed-circuit board) or overlay requires changes between iterations.

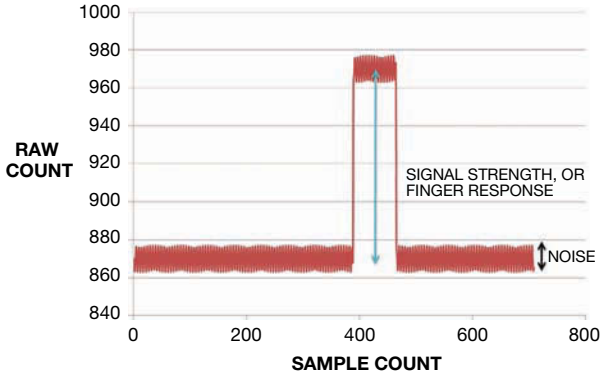

Figure 5 A 5-to-1 SNR is the minimum requirement for a robust sensor design.

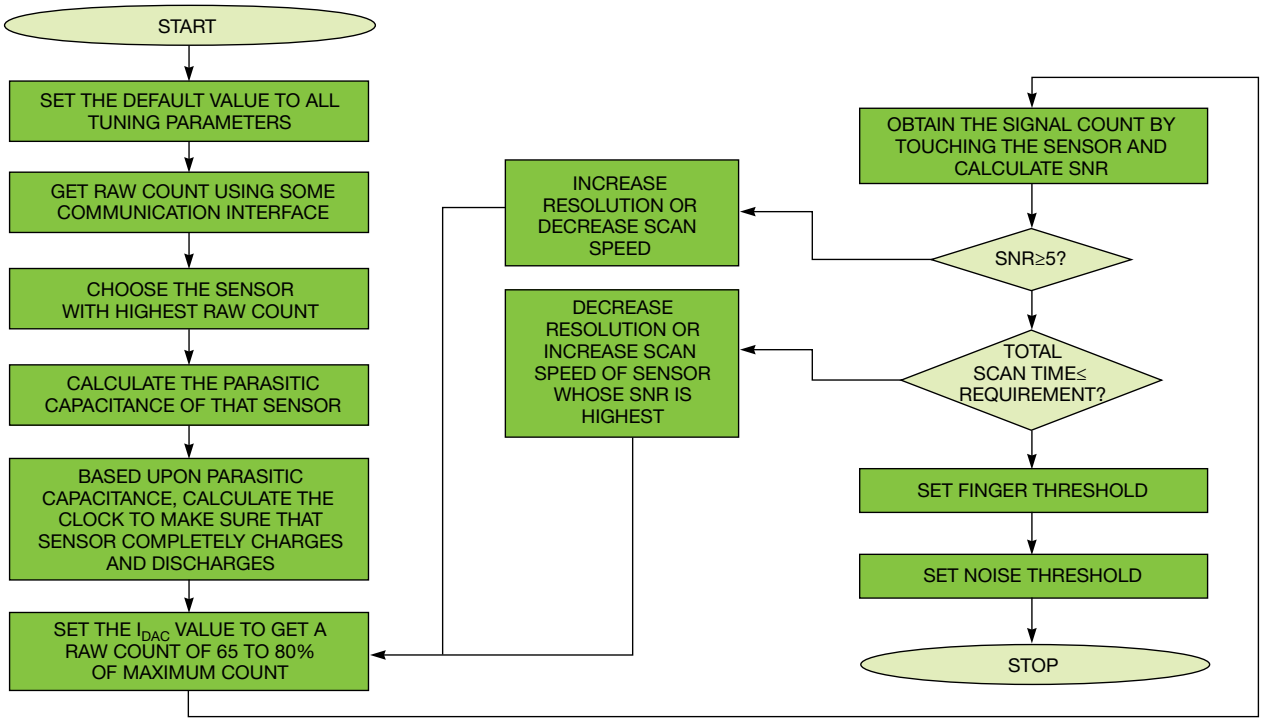

Figure 6 Although sensor-controller manufacturers provide guidelines to aid engineers in tuning, achieving the ideal tuning parameters for the system involves iteration. For a sensor controller with a capacitive-sensing algorithm, the tuning procedure follows these steps.

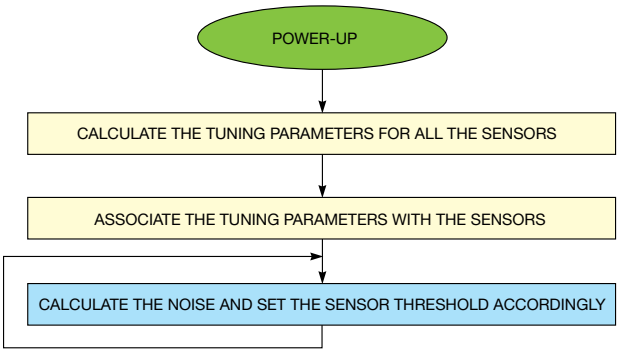

Figure 7 To cost-effectively handle design constraints and market needs, it is best to implement tuning in the appliance itself. This ideal self-tuning system performs this task.

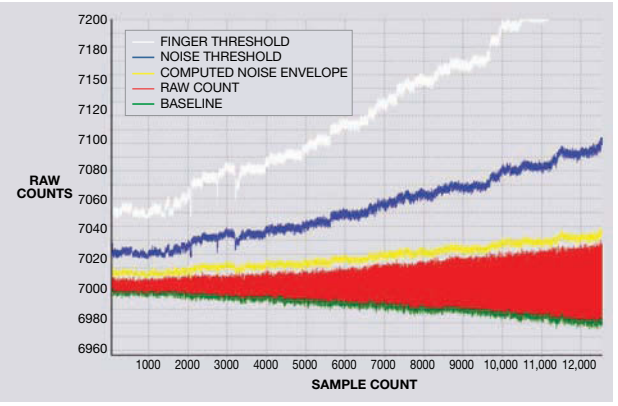

Figure 8 Adjust the threshold value for finger detection depending on detected noise in the sensor's raw counts. A self-tuning system can adjust the finger threshold depending on system noise.

PRODUCTION TUNING

Capacitive-sensor performance depends heavily on the physical properties and characteristics of the sensor board and environmental and operating conditions. For example, sensor capacitance can change due to variations in the PCB-manufacturing process, in overlay material or thickness, or in the PCB vendor. All of these variations can affect sensor performance, and the challenges do not stop there. Parasitic capacitance also varies with environmental conditions, such as noise, temperature, and humidity. Thus, a board you tune in the Swiss Alps may not work in the hot and humid climate of Hong Kong, resulting in more time and labor to retune the board. To minimize yield losses due to process variation or vendor change, tuning must take into account expected differences based on statistical analysis.

You may need to redo the board layout for various reasons, including changing a button's size, moving traces to incorporate minor changes in the schematic, and resizing to address EMC (electromagnetic-compliance) and EMI (electromagnetic-interference) issues. All of these modifications require that you retune the sensors. Moreover, tuning requires a communication protocol and a host-side application to observe and analyze the raw sensor data. Tuning also requires extra I/Os because it must take place after the final board is complete, creating potential issues for systems with pin constraints.

Tuning is a difficult job, requiring significant expertise with the chips and an understanding of capacitive-sensing effects at low signal levels. Couple these issues with the time-to-market constraints of the appliance market, and tuning can impose significant delays and increase system cost. To cost-effectively handle design constraints and market needs, it is best to implement tuning in the appliance itself. An ideal self-tuning system performs this task (**Figure 7**).

In systems with self-tuning capacitive sensing, numerous algorithms achieve a workable touch-sense system. At a basic level, appliance-implemented self-tuning is the same as manual tuning. In one-time compensation, some tasks occur once at power-up; in dynamic compensation, the appliance must continuously perform some tasks.

Self-tuning capacitive-sensing systems must calculate the best parameter settings for the sensors depending on the appliance and expected operating environment. The capacitive-sensing systems in this article use switched-capacitor theory. In contrast, a physical sensor capacitor forms a resistor by charging and discharging the sensor capacitor in consecutive cycles. Emulated resistance is proportional to the sensor capacitor's value, and you measure it using a current source with an analog-to-digital-conversion stage to compute the value of the sensor's capacitance.

Proper emulation of resistance requires the sensor capacitor to fully charge and discharge at a frequency that provides enough time for this operation. Therefore, you should adjust the switching frequency in accordance with the absolute sensor capacitance and reduce the frequency if the sensor's capacitance is higher.

Because the system converts the capacitance of the sensor to counts, the smallest change in capacitance that you can measure depends on the system's resolution. You can calculate the required resolution using the parasitic capacitance and required sensitivity.

Scan time is one of the most important tuning parameters from a system-specification point of view. However, increasing the resolution of scanning adds increased noise to the system. To compensate for this increase in noise, the scan time of sensors must stretch to integrate the noise and reduce its effect on capacitance measurements. A self-tuning algorithm must not exceed the scan time beyond system requirements. The best way to deal with scan time is to lay out the board to keep parasitic capacitance as low as possible.

Upon selecting the scanning resolution, you must automatically adjust a baseline measurement of the sensors—that is, the raw count when a finger is not present—to nearly 80% of the maximum count. This approach ensures that variations in neither environmental conditions nor silicon parameters adversely affect the sensor's measurement precision or its ability to accurately detect a finger touch.

Noise, by its nature, is a random function of time. It is not the same on power-up, a moment after power-up, or an hour after power-up. As a result, you should adjust the threshold value for finger detection depending on detected noise in the

sensor's raw counts. A self-tuning system can adjust the finger threshold depending on system noise (**Figure 8**).

SNR and scan time are the main factors determining the robustness, reliability, and efficiency of appliance-implemented self-tuning. Ensure that the SNR from self-tuned sensors is higher than the minimum requirement of 5-to-1 across the parasitic-capacitance range to preserve robustness and reliability. Scan time affects the power efficiency of the self-tuning algorithm; the more time a sensor takes to scan a sensor, the more power it consumes. Even though a longer scan may fit application requirements, the self-tuning algorithm should optimize power consumption by cutting the scan time as much as possible without compromising the SNR.

Designers often overlook one of the most important aspects—layout—during the initial design stage of a design; board layout affects the whole system's performance. Parasitic capacitance affects the amount of effort necessary for tuning, product yields, scan time, and other system characteristics. You should follow the guidelines from controller manufacturers when designing the layout to minimize the parasitic capacitance of the sensor. These guidelines can be used to improve the performance of the system through self-tuning, which helps you meet changing market needs. Cypress, for example, provides the SmartSense self-tuning-based capacitive sensor, which automatically optimizes the scan speed to be as high as possible, minimizes power consumption, and maintains an SNR of greater than 5-to-1 to avoid any false triggering.

Self-tuning controllers eliminate the overhead of repeatedly tuning capacitive-sensing-based appliances as specs and operating conditions change. In some cases, a sensor's parasitic capacitance may be high, requiring external components and some manual tuning to maintain the capacitance within a typical range.

Manual tuning can impose significant design challenges for developers implementing capacitive sensing in appliances and other systems. Tuning is sometimes necessary for different lots due to process variations; following board redesigns to accommodate changes in overlay thickness, button size, or other requirements; or due to noise and interference issues. Tuning improves performance and reliability, but manual tuning can add cost and delay the product's release. Self-tuning controllers eliminate these costs and delays, allowing developers to quickly implement reliable systems without becoming capacitive-signaling experts. EDN

ACKNOWLEDGMENT

This article originally appeared in two parts on EDN*'s sister site, Medical Designline (http://bit.ly/rhsD8E and http://bit.ly/npv5hq).*

AUTHORS' BIOGRAPHIES

Sachin Gupta is a senior applications engineer on the global-applications team at Cypress Semiconductor.

Prakhar Goyal is an application engineer at Cypress Semiconductor.

Kurian Polachan is a senior applications engineer in Cypress Semiconductor's consumer and computation division, where he focuses on CapSense applications.

FEED YOUR NEED

Online in the ProductFeed, *EDN*'s technical editors highlight notable new products, including analog and digital ICs, power components, sensors, passives, and boards. Contact the appropriate editor to submit products for consideration. Don't know who to contact? See the following link: http://bit.ly/qT7xKC.

→www.edn.com/productfeed

Pure DC

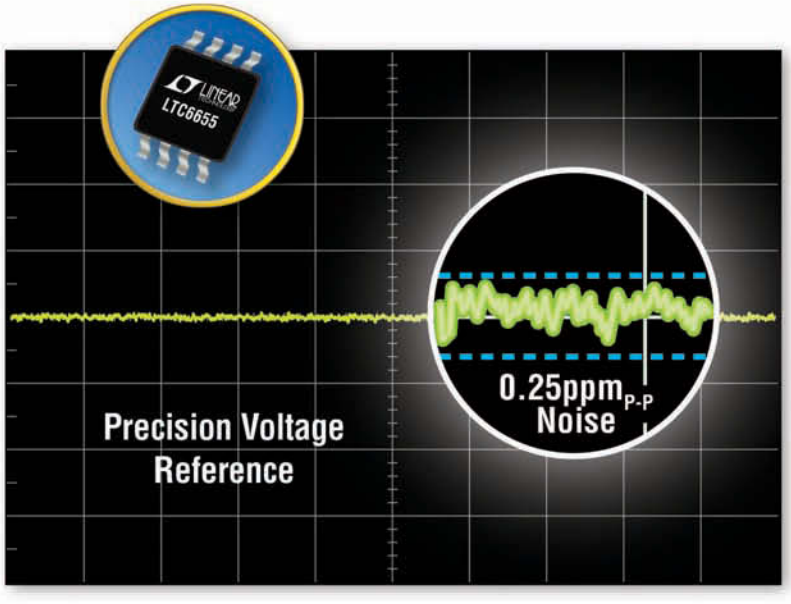

Announcing the Industry's Lowest Noise, Complete Voltage Reference Family

A lot of noise is made about precision voltage references. At Linear Technology, we're taking the quiet approach—introducing the LTC®6655 precision reference, with less than ***0.25ppm*** of low frequency noise. And if that is not enough, the LTC6655 exhibits less than 2ppm/°C of drift and an initial accuracy within ±0.025%. Such rock solid, pure DC performance ensures that the LTC6655 will be the new standard for high performance instrumentation and test equipment.

Features

- Low Noise: 0.25ppm$_{P-P}$ (0.1Hz to 10Hz)
- ±0.025% Accuracy, 2ppm/°C Drift
- Fully Specified from –40°C to 125°C, & 100% Tested at –40°C, 25°C and 125°C
- Sinks and Sources Current: ±5mA
- V_S: 500mV Dropout to 13.2V
- Load Regulation: 3ppm/mA
- Low Power Shutdown: <20µA
- 1.25V, 2.048V, 2.5V, 3V, 3.3V, 4.096V, 5V
- 8-Lead MSOP Package

Temperature Drift

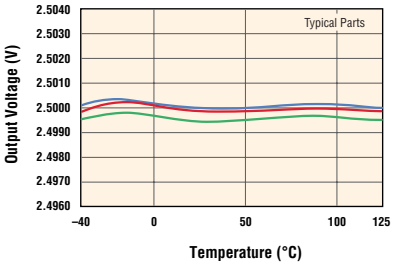

Info & Free Samples

www.linear.com/6655

1-800-4-LINEAR

www.linear.com/an124

LINEAR TECHNOLOGY

EDITED BY PAUL RAKO
AND FRAN GRANVILLE

designideas

READERS SOLVE DESIGN PROBLEMS

Polynomial rotation accelerates CRC calculations

Josef Valasek, AWOS, Pardubice, Czech Republic

Engineers commonly use two methods of CRC (cyclic-redundancy-check) calculations. One method uses shift registers, and the other uses precalculated CRC values. The first method shifts the data bit by bit, and, depending on the bit value, a polynomial divides the bits. **References 1** and **2** explain this method in detail. You can find assembly-language procedure codes for 8- and 16-bit calculations in **references 1**, **3**, and **4**.

THIS IDEA INTRODUCES A METHOD OF CRC CALCULATION THAT ROTATES A POLYNOMIAL INSTEAD OF SHIFTING DATA.

References 1, **3**, and **4** describe procedures that use just a few bytes of program memory, but they work slowly. When you study procedure codes, you'll find that shifts and cycles use more than 70% of execution time, whereas the CRC calculation uses less than 30%. **Listings 1** and **4**, which are available at www.edn.com/110908dia, show simpler and faster procedure codes for 8- and 16-bit CRC calculations, respectively, for 8051 microcontrollers.

The second method for CRC calculations uses the table of precalculated CRC values. **References 1**, **3**, and **4** contain links to 8- and 16-bit assembly-language procedure codes, respectively. The procedures work quickly, but they consume many bytes of program memory, which can cause a problem if you use a microcontroller with little program memory.

This Design Idea introduces a method of CRC calculation that rotates a polynomial instead of shifting data. You know the polynomial's value before you write the procedure's code. Therefore, you can prepare polynomial values and write them into the procedure code as constants. Online **figures**, also available at www.edn.com/110908dia, show the principle of the method.

The method exclusive-ORs the data with the old CRC value and places the result to bits 7 to 0 of a 16-bit register. Each calculation has eight steps. In each step, a polynomial value shifts one bit to the left, and the code tests one bit of the register. The results of the tested bit determine whether the code performs or skips the polynomial division. The resulting CRC value is subsequently written to bits 15 to 8 of the register. When all steps are complete, bits 15 to 8 of the register comprise a new CRC value. The lowest significant bits of the register remain unused after testing. You can use these dormant bits in the register to store result bits by modifying the method using polynomial rotation instead of a polynomial shift. **Listings 2** and **5** show this procedure for 8- and

DIs Inside

50 Simple circuit measures optocoupler's response time

52 Circuit provides visual verification of IR pulses

54 An LED's intrinsic capacitance works in a 650-mV LRC circuit

▶What are your design problems and solutions? Publish them here and receive $150! Send your Design Ideas to edndesignideas@ubm.com.

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

TABLE 1 COMPARISON OF 8-BIT CRC-CALCULATION PROCEDURES

Method	Procedure code	Code length (bytes)	Minimum/maximum execution times (µsec)
Shift-register simulation	Reference 3, calculation method	26	119/127
	Reference 1, Example 1	28	120/128
	Listing 1	15	45/53
Polynomial rotate	Listing 2	45	20/28
Look-up table	Listing 3	265	8/8

TABLE 2 COMPARISON OF 16-BIT CRC-CALCULATION PROCEDURES

Method	Procedure code	Code length (bytes)	Minimum/maximum execution times (µsec)
Shift-register simulation	Reference 4, calculation method	26	92/124
	Reference 1, Example 4	37	156/180
	Listing 4	23	77/101
Polynomial rotate	Listing 5	57	22/36
Look-up table	Reference 4, table method	550	40/40
	Reference 1, Example 5	533	16/16
	Listing 6	530	14/14

16-bit CRC calculations, respectively. These routines perform only bit-testing and polynomial-dividing instructions.

Tables 1 and **2** highlight calculation procedures for 8- and 16-bit calculations, respectively, using this method. The **tables** describe execution times for a standard 12-clock 8051 core with a 12-MHz crystal. The code lengths and execution times of procedure codes come from **Reference 1**; **references 3** and **4** calculate without saving or restoring the procedure's working registers.

You should properly set the initial CRC value before starting CRC calculation. Some communication protocols, such as the one in **Reference 2**, use an initial CRC value other than 00 Hex. **Listings 7** through **10** show the procedure codes for 8- and 16-bit CRC calculation with PIC (www.microchip.com) microcontrollers. EDN

REFERENCES

1 "Understanding and Using Cyclic Redundancy Checks with Maxim iButton Products," Maxim Integrated Products, Application Note 27, 2001, http://pdfserv.maxim-ic.com/en/an/AN27.pdf.

2 *MODBUS over Serial Line Specification and Implementation Guide*, Version 1.02, Modbus Organization, December 2006, pg 39, www.modbus.org/specs.php.

3 Wren, John C, "A CCITT-8 CRC calculator," www.8052.com/codelib.htm.

4 Wren, John C, "A CCITT-16 CRC calculator," www.8052.com/codelib.htm.

Simple circuit measures optocoupler's response time

Peter Demchenko, Vilnius, Lithuania

You can use the circuit in this Design Idea to measure the attack and release times of photoresistor-type optocouplers (**Figure 1**). Such devices often find use in audio compressors or volume-control circuits. The design uses an oscillating Schmitt trigger with the optocoupler DUT (device under test) in the feedback loop. The photoresistor and resistor R_1 form a voltage divider that controls the input of the Schmitt trigger. The optocoupler's LED connects to the trigger output. You can measure the duration of the output pulses with an oscilloscope or a digital meter. The duration of the negative output pulses is equal to the switching on-time, or attack time. The duration of the positive pulses is equal to the switching off-time, or release time. The attack and release times depend on the value of R_1; you can

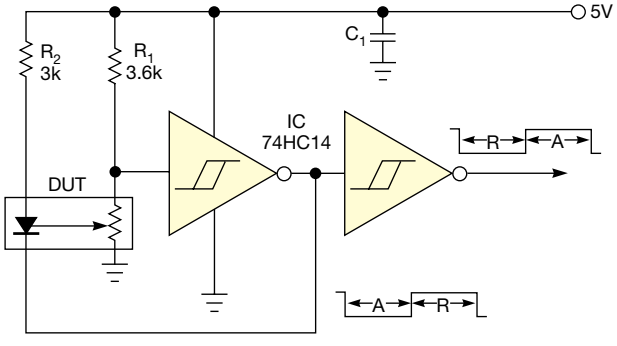

Figure 1 You can determine an optocoupler's rise and fall times by incorporating a photoresistor in the feedback loop of an oscillator circuit.

Create a Brighter Future

Maxim's LED Drivers Enable Drop-In Replacements for Offline/MR16 Lamps

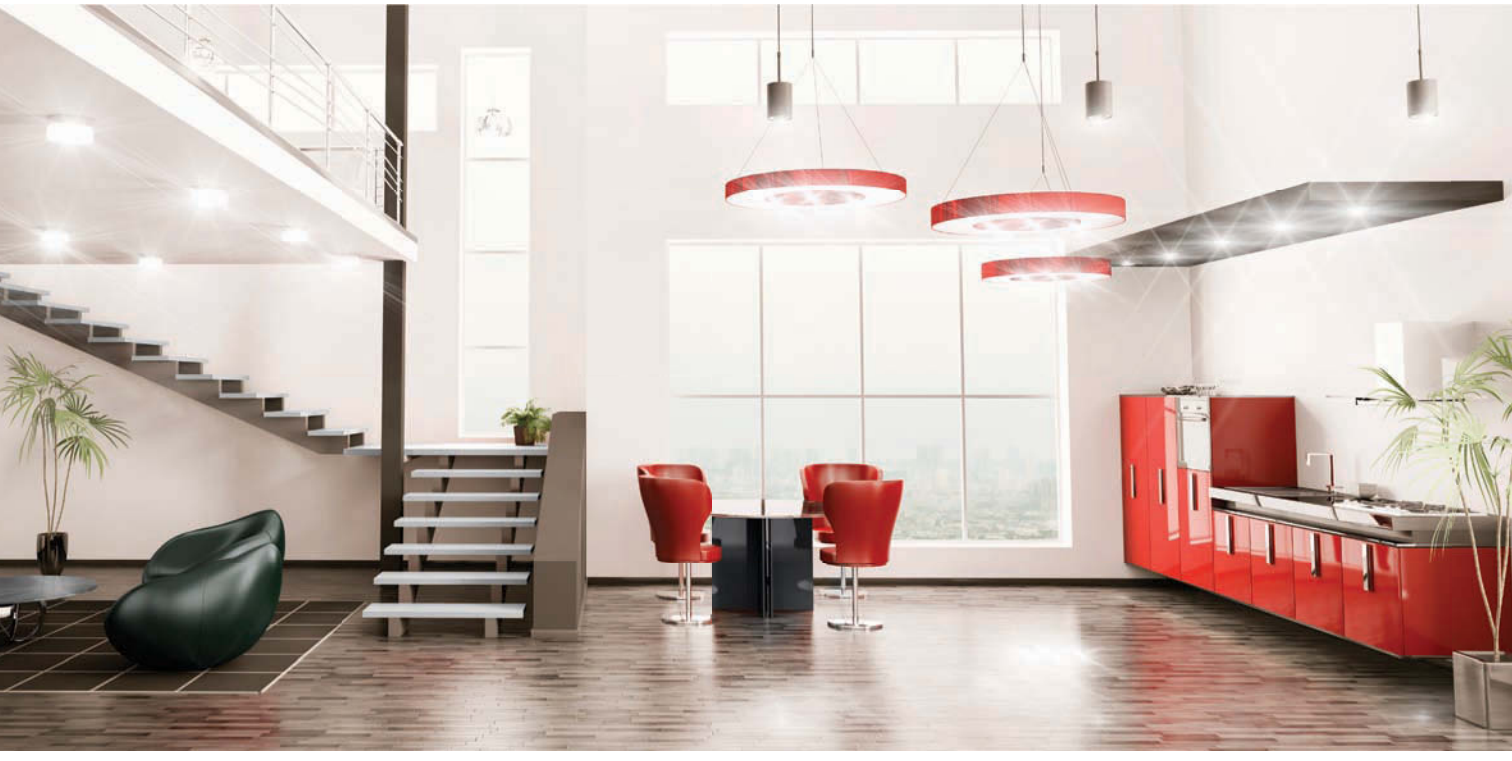

Maxim's solutions for dimmable offline (A, R, PAR, GU10, etc.) and MR16 LED retrofit lamps use a proprietary architecture to achieve flicker-free operation with dimmers and electronic transformers. This ensures compatibility with the existing electrical infrastructure, helping manufacturers overcome a critical obstacle to large-scale adoption.

Technology Benefits

- Smooth dimming with installed cut-angle dimmers from 0 to 100% brightness
- Flicker-free operation with most electronic transformers (MR16)
- Universal input (90VAC to 265VAC) allows a single lamp design to support worldwide infrastructure requirements (offline)
- Optional operation without electrolytic capacitors extends lamp lifetime

Find your LED driver solution at:
www.maxim-ic.com/LED-info

MAXIM
INNOVATION DELIVERED®

MAXIM DIRECT
www.maxim-ic.com/shop

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

For a complete list of Maxim's sales offices and franchised distributors, visit **www.maxim-ic.com/sales.**

observe both by varying the value of R_1. With the component values in **Figure 1**, the durations of the output pulses are a 0.15-msec attack time and a 2.7-msec release time.

During oscillation, the resistance of the photoresistor sweeps in from R_{P1} to R_{P2}. The circuit sweeps these photoresistor values according to R_1, the power-supply voltage, and the Schmitt-trigger thresholds, as the following **equations** show: $R_{P1}=R_1\times V_{T2}/(V_{CC}-V_{T2})$, and $R_{P2}=R_1\times V_{T1}/(V_{CC}-V_{T1})$, where V_{T1} is the positive-going threshold voltage and V_{T2} is the negative-going threshold voltage of the Schmitt trigger.

In the case of the 74HC14 logic family, you can determine the thresholds from the data sheet and your power-supply voltage, according to the following **equations**, which yield typical values: $V_{T1}=0.53\times V_{CC}$, and $V_{T2}=0.31\times V_{CC}$. Using 5V as a power-supply voltage and solving the following **equations**, you can determine the photoresistor range: $R_{P1}=0.45\times V_{R1}$, and $R_{P2}=1.13\times V_{R1}$.

This approach lets you pick a value for R_1 so that the photoresistor range is suitable for your device. You can also vary the value of resistor R_2 to observe the LED-current-to-attack-time characteristic of the DUT but not affect the release time. Note that R_2 limits the current through the LED; if its value is too large, oscillation will not occur.

Using this circuit allows you to match custom optocouplers comprising green, superbright LEDs and an MPY7P photoresistor. A recent Design Idea, although thorough, lacked data on response time (**Reference 1**). EDN

REFERENCE

1 Foit, Julius, and Jan Novák, "Photoresistor provides negative feedback to an op amp, producing a linear response," *EDN*, May 27, 2010, pg 49, http://bit.ly/oPQMfo.

Circuit provides visual verification of IR pulses

Michael J Gambuzza, General Electric Measurement and Control Solutions, Billerica, MA

You can test an IR (infrared) link with a circuit that converts an IR-generated photocurrent to an amplified current that drives a standard LED. This approach provides a visual feedback to indicate that the transmitter is working. The circuit can be enclosed in a small plastic or metal box and requires just a 9V transistor battery for operation. Diode D_1 is a basic Everlight (www.everlight.com) PD333-3C/H0/L2 or equivalent IR photodiode in a T1¾ package.

You can configure amplifier IC_{1A} as a photovoltaic amplifier. When the IR-light energy impinges on photodiode D_1, it generates a small photocurrent that tries to pull the inverting input negative. Meanwhile, the output of IC_{1A} goes positive, maintaining the virtual-ground node on Pin 2 of the amplifier at 0V. The transfer function for the circuit is $V_{OUT}=I\times R_1$. If you set the gain high, IC_{1A} goes to the power-supply rail when the circuit detects light. Analog Devices' (www.analog.com) AD823AR JFET-input amplifier directly drives the LED through a 750Ω current-limiting resistor. C_1 compensates the amplifier, preventing it from oscillating due to capacitive load from D_1 and the input parasitic capacitance.

If the output of IC_1 oscillates, you may need to increase the value of C_1. You can determine the value of C_1 by using the following **equation** for a 45° phase margin: $C_1=\sqrt{(C_D/2\pi R_1 F_C)}$, where F_C is the unity-gain-crossover frequency of IC_{1A}—typically, 16 MHz for the AD823—and C_D is D_1's 0V junction capacitance, including any parasitic capacitance on that node. Adjust R_1 for optimum gain. For testing, the remote-control transmitter window should be as close as possible to photodiode D_1 for maximum signal transfer. EDN

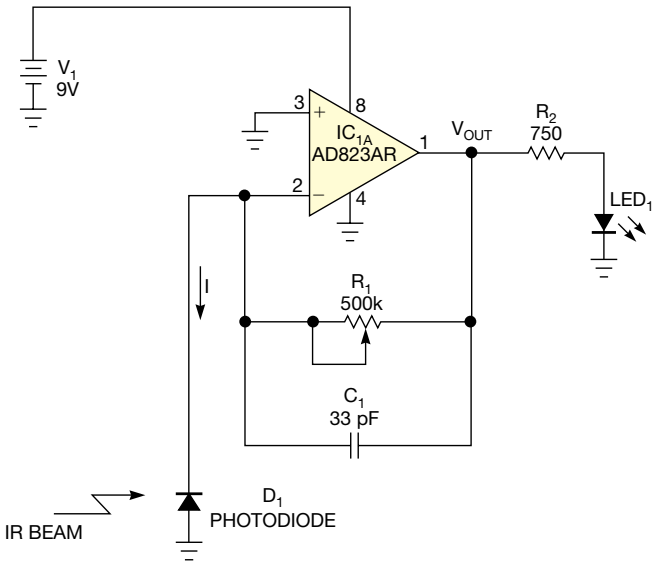

Figure 1 This circuit detects the transmitter in an IR communication link and provides a visual output from the LED.

Explore the EM simulation universe

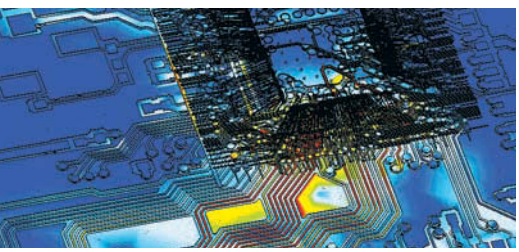

Crosstalk analysis **www.cst.com/pcb**

→ Get equipped with leading edge EM technology. CST's tools enable you to characterize, design and optimize electromagnetic devices all before going into the lab or measurement chamber. This can help save substantial costs especially for new or cutting edge products, and also reduce design risk and improve overall performance and profitability.

Involved in signal or power integrity analysis? You can read about how CST technology was used to simulate and optimize a digital multilayer PCB's performance at **www.cst.com/pcb**. If you're more interested in EMC/EMI or microwave components, we've a wide range of worked application examples live on our website at **www.cst.com/apps**.

Now even more choice for SI/PI simulation. The extensive range of tools integrated in CST STUDIO SUITE enables numerous applications to be analyzed without leaving the familiar CST design environment. This complete technology approach enables unprecedented simulation reliability and additional security through cross verification.

→ Learn more about CST STUDIO SUITE. Register for the CST Leading Technology webinars on EDA at **www.cst.com/eda-webinar**.

CST

CHANGING THE STANDARDS

An LED's intrinsic capacitance works in a 650-mV LRC circuit

Sajjad Haidar, University of British Columbia, Vancouver, BC, Canada

You can use the inherent capacitance of an LED to make a series resonant boost circuit that can create a voltage large enough to light the LED. Depending on the color of the LED, you need a voltage higher than 1.6V to turn it on. The threshold, or knee, voltage rises higher as the LED wavelength becomes shorter. All PN-junction diodes, including LEDs, have capacitance due to depletion and diffusion profiles.

You can light an LED using its capacitance in a series LRC (inductance/resistance/capacitance) resonant circuit. In such a circuit, the Q factor determines the multiple of the generator voltage that appears across LC. If you fashion a circuit with a high enough Q factor, you boost the generator voltage enough to light the LED. The Q factor of the resonant circuit is a function of the resistance, inductance, and capacitance, as the following **equation** shows: $Q=(1/R)\sqrt{(L/C)}$.

You can verify this calculation with a simple circuit using a blue LED in series with an inductor (**Figure 1**). The knee voltage of the LED is 2.45V, and the signal generator has an internal resistance of 50Ω. An inductance of 100 µH and the 50-pF capacitance of a typical LED yield a Q of 28. The amplitude of the sinusoidal signal generator is set at 650 mV p-p. You can then vary the generator's output frequency until you see the circuit's resonant point. As the circuit approaches the

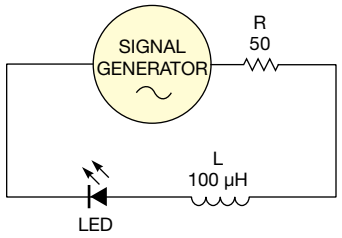

Figure 1 In this LRC circuit, the resistance is the source resistance of the sine-wave generator. The LED provides the capacitance.

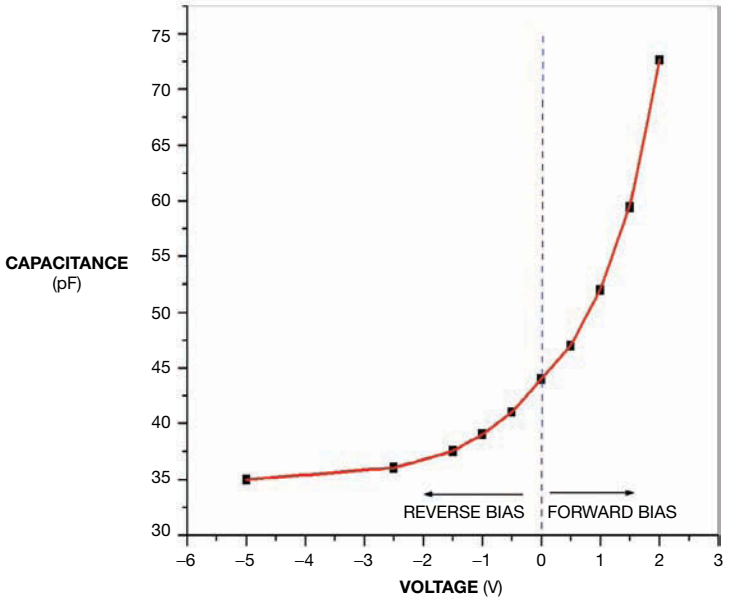

Figure 2 The junction capacitance of an LED increases as you go from a reverse bias of –5V to a forward bias of 2V.

FACTORY AUTHORIZED SAFETY STOCK™

100% GUARANTEED

We have the semiconductors you need *NOW!*

You know Rochester Electronics as the solution for EOL products, but did you know that we also have millions of current products? All product is 100% guaranteed and factory traceable for unquestionable quality. As our suppliers reach over-stock of certain products, Rochester acquires these and makes them available for times when supply is limited due to factory production lead time. This is Rochester's "**Factory Authorized Safety Stock™ [FASSt™]**." A tremendous resource to keep your production lines running without interruption. Be confident with semiconductors that you can get today. That's "**FASSt!**"

	Part Number	Product Family	SAFETY STOCK™
ANALOG DEVICES	AD5398BCPZ-REEL7	D/A Converter	226,847
	ADM5170JP-REEL	Line Driver	11,250
FAIRCHILD SEMICONDUCTOR	NC7S32P5X	OR Gate	871,379
	P6KE200A	Voltage Suppressor	111,602
freescale semiconductor	MC56F8013VFAE	Microcontroller	35,935
	MC9S08QE8CTG	Microcontroller	52,032
intel	RK80530PZ009256	Microprocessor	78,144
	SB80486DX2SC50	Microprocessor	59,193
ON Semiconductor	NB4N7132DTR2G	Interface	12,500
	NCP1651DR2G	Controller	99,600

These are just a few of the thousands of part numbers available *immediately!*

Transforming End-of-Life into Extension-of-Life®

Rochester Electronics turns End-of-Life problems into **Extension-of-Life** solutions that provide you with the authorized, guaranteed components you need for as long as you need them. Call us to find out more about our comprehensive, customizable **Extension-of-Life** products and services: continued manufacturing, re-creation, special product agreements, and inventory.

EOL™

Captain Rochester's Quest for Quality

Rochester Electronics has created Captain Rochester to personify the importance of quality and a secure and trusted supply chain in the semiconductor industry! Just like Rochester Electronics, Captain Rochester will arrive "just-in-time" to save the day when critical semiconductors are needed. Because the threat of counterfeit and substandard devices has risen significantly, now more than ever, there needs to be a secure supply chain of quality, authorized product.

Rochester Electronics

Worldwide Corporate Headquarters

16 Malcolm Hoyt Drive . Newburyport, MA 01950

phone **978.462.9332** . *email* **sales@rocelec.com** . *web* **www.rocelec.com**

Leaders on the "Trailing Edge of Technology."®

September 26-29, 2011.
Attend the East Coast's Leading Embedded Systems Event.

ESC Boston drives product discussions and faster applications, delivers training and education on the latest technologies, and helps you discover new solutions to your real-world design challenges.

NEW & EXCLUSIVE!
Available ONLY to All Access Pass Holders:

- Android™ Certificate Program
- BeagleBoard-xM
- eZ430 Chronos Watch
- Hands-On Sessions
- Build Your Own Embedded Systems

Highlights Include:

- 100+ Conference Sessions
- 14 Industry Tracks
- 20+ Sponsored Training Sessions
- 1000's of FREE giveaways
- 2 Co-Located Expo Floors
- Multiple Keynotes & Industry Addresses
- FREE Technology Tote
- Teardowns, panel discussions, studies & more

Keynote Speakers

Joerg Borchert
Vice President, Chip Card & Security ICs
Infineon Technologies North America

Charles G. Sodini
LeBel Professor EECS
MIT

Register Now for 10% off ANY package!
Expo Pass registration is FREE.
esc.eetimes.com/boston
Promo Code: AD1

Four Events. One Unique Experience.

This year, ESC Boston is part of DesignDays 2011: a distinctive electronics industry event comprised of 4 design conferences—ESC Boston, DesignCon East, DesignMED and Designing with LEDs. Attend all 4 conferences with the All Access Pass!

UBM Electronics

resonant frequency, the voltage across the LED starts to increase. The resonant point manifests itself as a small jump in voltage, rather than a smooth progression, due to a positive feedback at resonance. The positive feedback happens because the capacitance of any PN-junction device is not linear (**Figure 2**). As the circuit approaches the resonant frequency, the LED voltage increases, which also increases the LED capacitance, resulting in lower resonant frequency.

For a blue LED, the voltage waveform as the circuit approaches resonance is 1.55 MHz. The circuit settles at 1.69 MHz (**Figure 3**). The forward-biased LED is thus emitting light, clipping the positive parts of the boosted waveform. Using the same 650-mV-p-p generator amplitude on other colors of LEDs produces different resonant frequencies. You can see a similar effect with a square-wave generator because it also contains the fundamental components of the resonant frequency. EDN

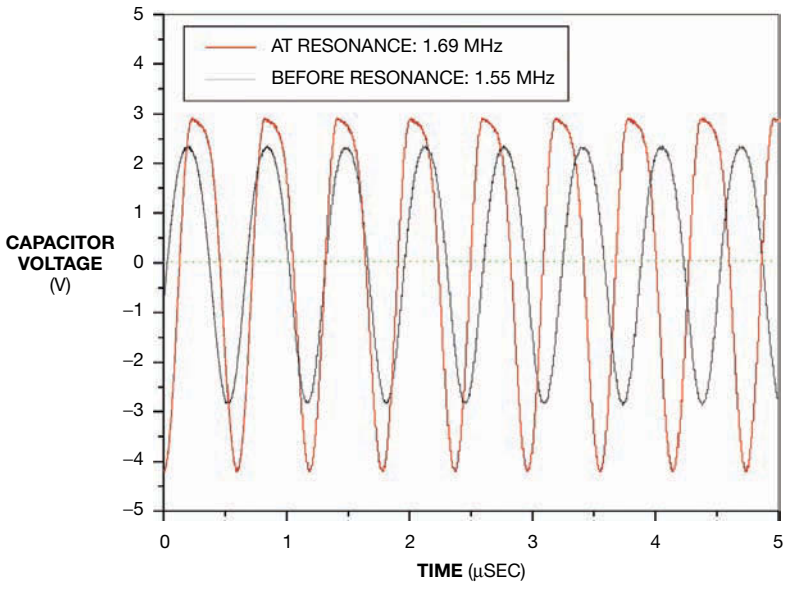

Figure 3 The effect of the LED capacitance increasing with applied voltage means the circuit will jump to a frequency as you approach the resonant point.

Tool is in the works to standardize conflict-minerals tracking

One of the most difficult challenges with any new regulatory initiative is establishing a standard for the measurement of compliance. For example, when the European Union was deploying its ROHS (restriction-of-hazardous-substances) directive—banning lead and other substances from electronics products for sale in the European Union—establishing a baseline for exempt products was a sticking point. Which products could use lead for safety reasons, and how much lead was acceptable?

> Once they gather data, companies can use the standard as a baseline to measure their compliance.

The movement to ban so-called conflict minerals from the electronics supply chain is facing the same hurdles. Regions that ignore human rights are the sources of many materials key to the electronics supply chain. Industry participants also suspect that factions within these regions are using profits from these minerals to finance civil wars.

Industry associations IPC (www.ipc.org), the EICC (Electronic Industry Citizenship Coalition, www.eicc.info), and the GeSI (Global e-Sustainability Initiative, www.gesi.org) are developing a standard that will allow businesses to demonstrate compliance with a law that discourages the use of conflict minerals. Under the Dodd-Frank Wall Street Reform and Consumer Protection Act, publicly traded companies must report the use of tin, tantalum, gold, and tungsten from conflicted regions. The standard will allow companies to share data on conflict minerals and assist in the preparation of compliance reports.

According to IPC, the data-exchange standard will build on a due-diligence communication tool that the EICC and GeSI recently released. The standard will use an XML schema in the IPC-175x family of standards. The EICC/GeSI template enables companies to track conflict-mineral-related information within their supply chains and share that data with their business partners. IPC will use the template to develop a data-exchange standard for the global electronics industry. According to an EICC press release, the EICC/GeSI tool includes questions regarding a company's conflict-free policy, engagement with its direct suppliers, and the smelters the company and its suppliers use.

The rapid development and adoption of this tool will significantly benefit the electronics industry. A standard helps avoid duplication of effort. Companies won't focus on their supply chain to the exclusion of others. Once they gather data, companies can use the standard as a baseline to measure their own compliance efforts, and the information will be available in a consistent format.

IPC will begin development of a data-exchange standard at a kickoff meeting on Sept 22, 2011, in Schaumburg, IL, during its Midwest Conference and Exhibition.

—by Barbara Jorgensen, *EBN* Community Editor

This story was originally posted by EBN: http://bit.ly/p78Xb9.

IMAGE: ISTOCK

OUTLOOK

INDIAN SEMICONDUCTOR CONSUMPTION TO GROW 15.5%

Gartner Inc (www.gartner.com) projects that India's semiconductor consumption will grow 15.5% year over year in 2011 to total $8.2 billion. "Changing demographics, increasing consumer affluence, economic growth, and favorable government policy continue to drive the electronic-equipment-manufacturing industry in India," says Ganesh Ramamoorthy, research director at Gartner. "Numerous global electronic-equipment-manufacturing companies have set up production facilities in India to take advantage of the growing domestic market and to cater to neighboring markets in the region."

The communications-electronics and data-processing segments will account for nearly 52% and about 26%, respectively, of India's semiconductor consumption in 2011. "Given the low penetration and the growing demand for key electronic equipment, … the Indian market will be able to sustain high growth rates in the coming years," Ramamoorthy says.

Gartner expects India's semiconductor consumption to grow the fastest across the globe through 2015 at a compound annual growth rate of 15.9% to reach nearly $15 billion.

—by Suzanne Deffree

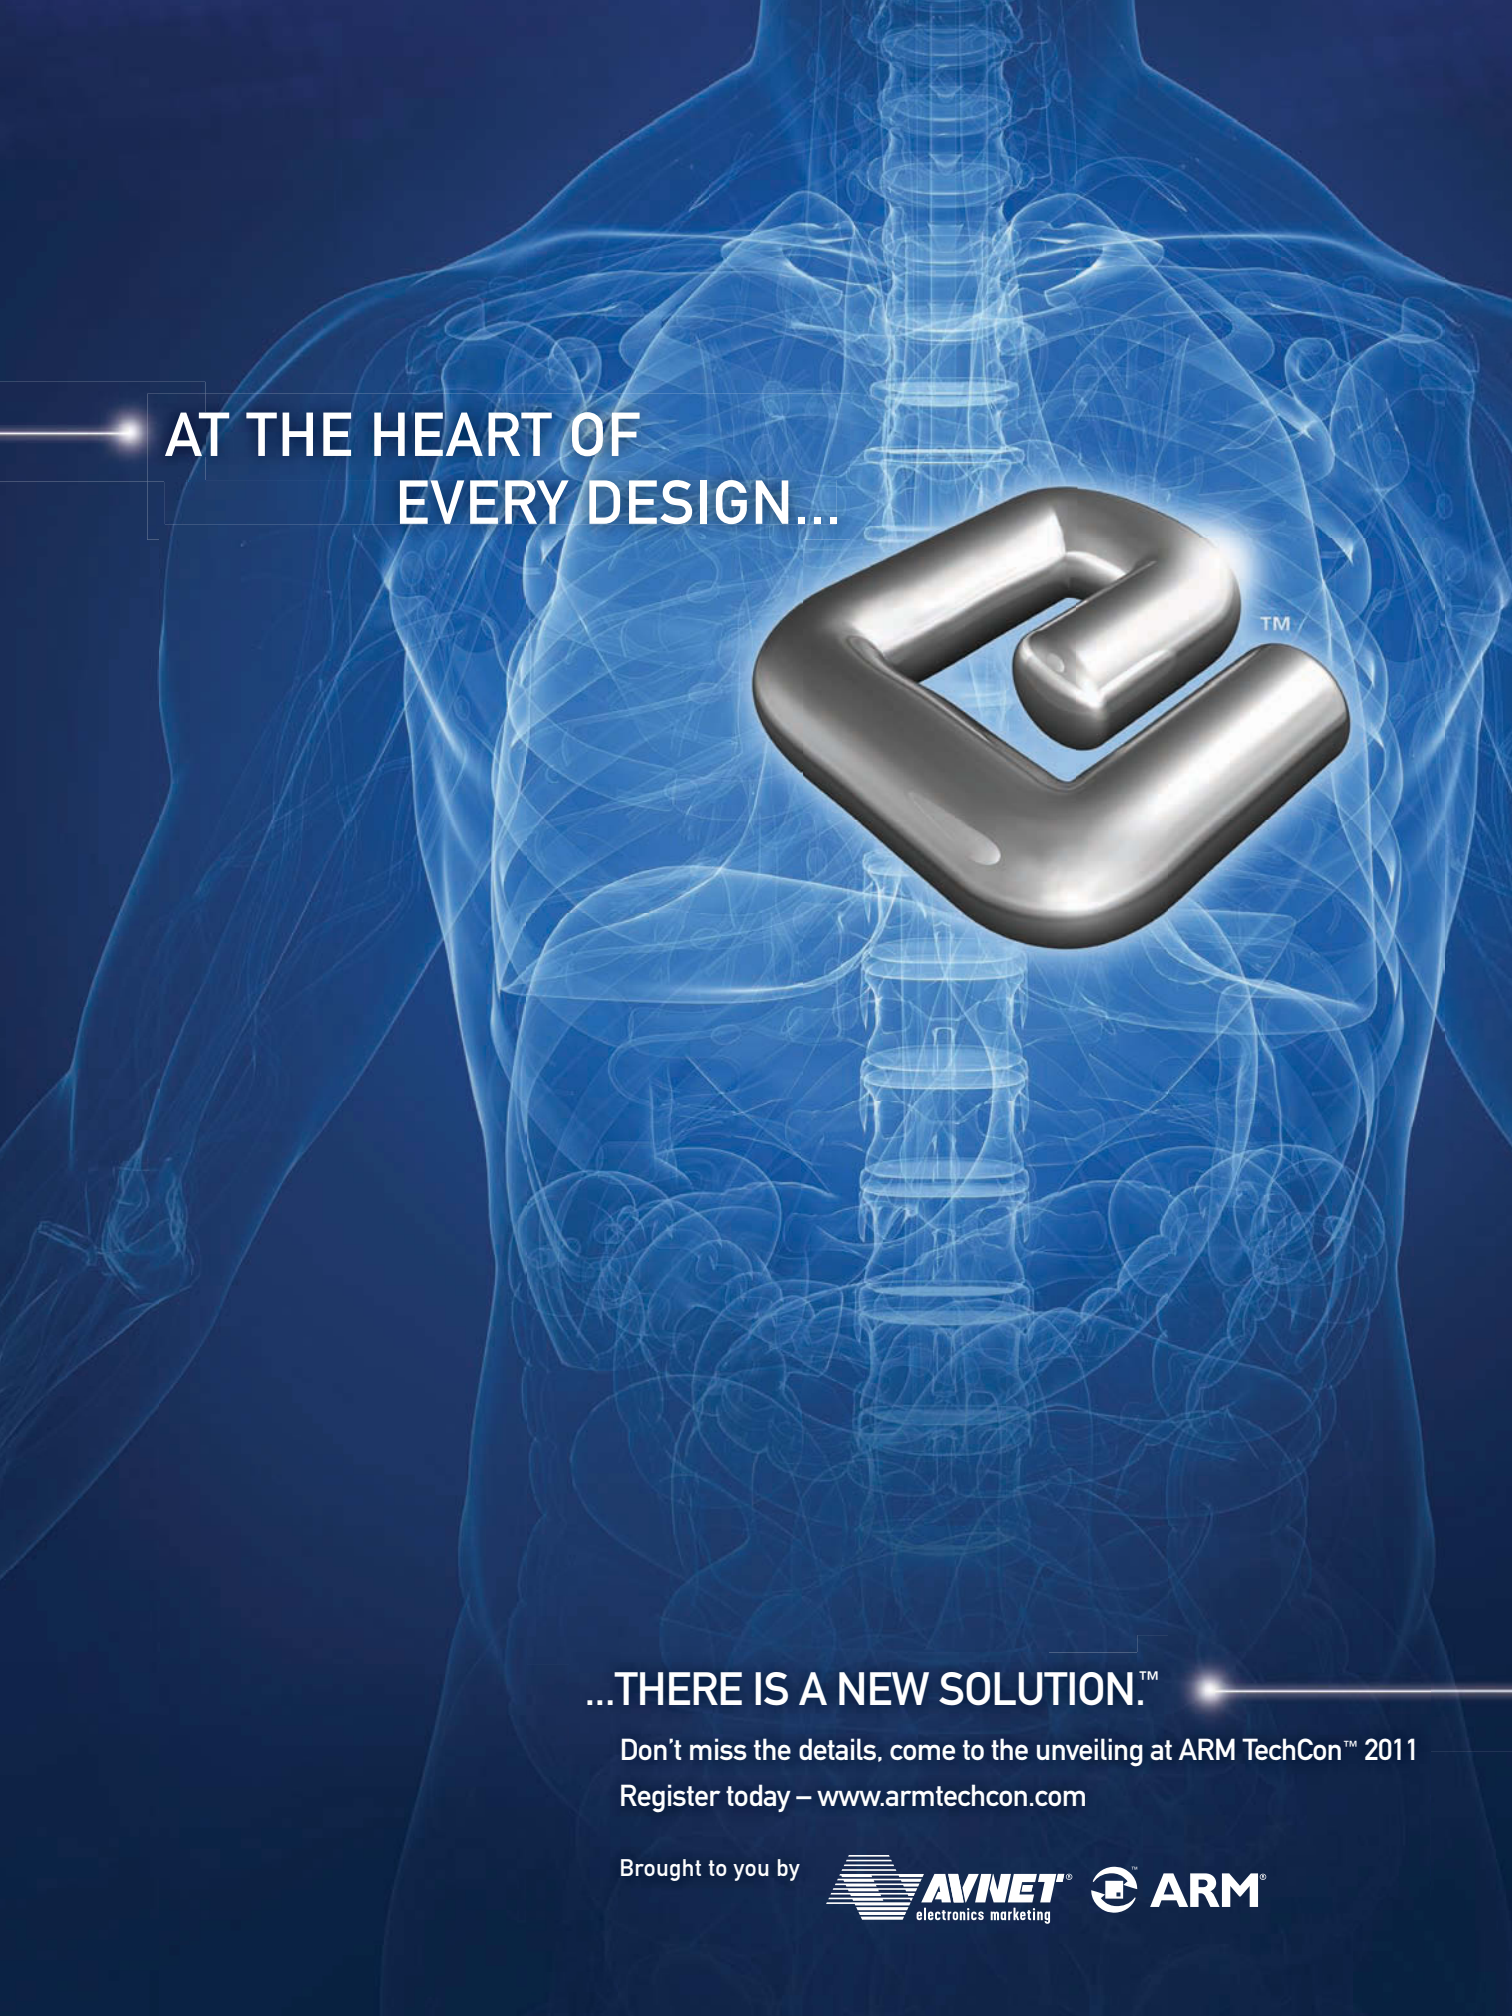
AT THE HEART OF
EVERY DESIGN...
...THERE IS A NEW SOLUTION.™
Don't miss the details, come to the unveiling at ARM TechCon™ 2011
Register today – www.armtechcon.com
Brought to you by
AVNET electronics marketing
ARM

productroundup

SWITCHES AND RELAYS

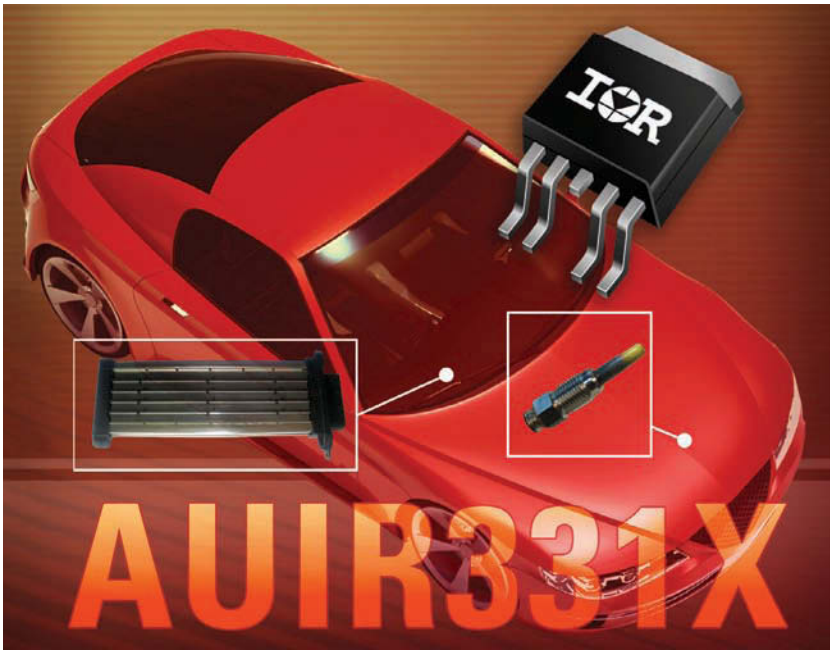

IR's AUIR331x intelligent power switch features current sensing for automotive applications

The AUIR331x family of high-side intelligent-power-switch devices features accurate current sensing and built-in protection circuits for automotive applications. The current sensing, particularly at low currents, enables precise monitoring of load current to provide additional data to the microcontroller for diagnostic applications, including open-load detection and early warning of overload and stalled-motor conditions. The family also includes a slow-switching version, which helps minimize noise in EMI-sensitive automotive applications. The AUIR331x ensures safe and reliable automotive application at continuous currents of 30A and peak currents of 90A. In the event of a reverse-battery condition, an integrated protection circuit turns on the main MOSFET switch, helping to relieve the intrinsic body diode and reducing or eliminating thermal problems. Prices range from 89 cents for the AUIR3315S to $1.20 for the AUIR3313S (100,000).
International Rectifier, www.irf.com

E-T-A's 1620 Type II mini circuit breaker targets automotive applications

The Type II version of the automotive-market-targeted 1620 mini circuit breaker meets all of the requirements of SAE J553 and is available in current ratings of 5 to 30A, with a voltage of 12V dc. The product's housings are color-coded consistently with industry-standard current ratings and comply with the form, fit, and function of other mini blade products, making it interchangeable with blade-type fuses. Integration of the 1620 into designs allows for trouble-free retrofit capability, eliminates the need for spare fuses, saves space, and reduces vehicle downtime. The device finds use in applications for the protection of onboard, 12V electrical systems in passenger cars, trucks, buses, watercraft, specialty vehicles, and extra-low-voltage wiring systems that require modified reset ability. The price is less than $2 (high volumes).

E-T-A Circuit Breakers, www.e-t-a.com

C&K's KMT series includes ultra-low-profile, double-action switch

The KMT series of ultraminiature tactile switches now includes a double-action, double-tactile version. The KMT 6 series provides a lifetime of 150,000 electrical/mechanical cycles and is sealed to IP68 certifications, making it suitable for use in handheld and portable electronic devices. The ROHS-compliant, halogen-free KMT 6 features a footprint of 3.8×3.1 mm and a height, including the actuator, of 0.76 mm. The normally open switches are constructed of lead-free silver plating. The first and second actuation forces are 1 and 2N, respectively, with a minimum difference of 0.6N between the two forces to ensure proper discrimination. The devices feature maximum power of 0.5 VA, a voltage of 20 mV to 32V dc, and current of 1 to 50 mA. Operating temperature ranges from –40 to +85°C, and prices are 0.135 cents, depending on quantity.
C&K Components, www.ck-components.com

AnalogicTech's AAT2138 dc/dc converter targets USB 3.0 peripherals

The AAT2138, a 5A-output-current, synchronous dc/dc converter, integrates a ±10% precision current-limit load switch for safe and efficient regulation of power in USB peripherals. The device supports power levels beyond the USB 3.0 standard and combines the load switch and a 95%-efficient dc/dc converter in one package. Input-voltage range is 2.7 to 5.5V, and output-voltage range is as low as 3V. The device features 85-mΩ high-side FETs and 50-mΩ low-side FETs and has a 2.8-MHz typical switching frequency. It sells for $1.19 (1000).

Advanced Analogic Technologies Inc, www.analogictech.com

EDN ADVERTISER INDEX

Company	Page
Agilent Technologies	C-3, 14, 42
Analog Devices	17
ARM	59
austriamicrosystems AG	57
Centellax	C-2
CST - Computer Simulation Technology AG	53
Digi-Key Corp	C-1, 3
Express PCB	30
Future Electronics Inc	11
Infineon Technologies AG	21
Intel	C-4
Interconnect Systems Inc	39
International Rectifier	6
Ironwood Electronics	61
Linear Technology	46A–46B, 48
Materion Corp	26
Maxim Integrated Products	51
Mentor Graphics	23
Mill-Max Manufacturing	13
Mornsun Guangzhou Science & Technology Co Ltd	47
Mouser Electronics	4
Pickering Electronics Ltd	41
Pico Electronics Inc	7, 32
Rochester Electronics	54, 55
Rohde & Schwarz	29, 31, 33
Samtec	12
Silicon Labs Inc	19
Tektronix	37
UBM *EDN*	45, 56
Wind River	8
Xilinx Inc	25

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

ENCLOSURES

OKW's Minitec enclosures target personal electronics

The Minitec line of pocket-sized enclosures targets medical-device, personal-safety, detection, monitoring, emergency-call, audio, alarm, and remote-control electronics. The family comprises eight case sizes with dimensions of 2×1.25×0.51 to 3.3×2.08×0.75 in. Each case comprises a top part, a base part, and a soft intermediate ring. The top and base parts are molded in off-white or lava-gray ABS or black, infrared, transparent PMMA. The intermediate rings are molded in soft-touch TPE in volcano gray, orange, lime green, and bright blue. The intermediate rings are available plain or with molded eyelets or strap loops for fitting the accessory key ring, wrist strap, or lanyards. The recessed top of the enclosure is designed for attaching a membrane keypad or a product label. The Edge-style product also comes with a smooth top for fitting rubber keypads or similar accessories. Location bosses are molded in the base parts for mounting PCBs and components, and the cases are assembled using a self-tapping fixing screw. Prices start at $8.

OKW Enclosures Inc, www.okwenclosures.com

NIGEL ADAMS • AEROFLEX LTD

Coffee-break mistake

A number of years ago, I worked as a senior support engineer for a large ATE (automated-test-equipment) system manufacturer. We supplied a large number of test systems to many customers worldwide, all with successful results. One company had five similar machines in its production process, and all initially seemed to be operating well, with good reliability and performance. After a few months, though, problems started to emerge. Our service engineers began to get calls to investigate system crashes. The factory was not local, and getting there required a lengthy journey. By the time the engineer reached the site, the problems had vanished, however, and the systems were working normally.

We swapped out various cards, changed power supplies, and reloaded software—all to no avail. This problem continued intermittently for weeks. One minute the customer was testing boards; the system would then suddenly and inexplicably freeze. Nothing except a reboot would fix the problem. There appeared to be no pattern to the failures. They were not operator-dependent, and they occurred when different tests and even different programs were running.

The customer was concerned over system downtime, blaming our hardware, and we were running out of ideas. Taking drastic measures, we decided to replace the ATE. We selected a demonstration system, fully reworked it so that it was in tiptop condition, and calibrated it to within an ounce of perfection. We loaded the customer's programs, and they performed perfectly under test. Upon delivery of the system, all appeared well.

The next morning, though, we got the dreaded phone call: The new system was behaving exactly the same as the old one. Meanwhile, we had returned one of the customer "rogue" systems to the workshop and could find nothing wrong with it. Reviewing the service sheets and job records revealed little in the way of fixes that I wouldn't have tried under similar circumstances. It was time for a clean sheet of paper.

So, I loaded up the car with all manner of test equipment and headed for the customer site, not knowing what to expect or with an idea of what to try to fix the problem. I decided the best first attack would be to simply sit and watch what happens. The first day revealed nothing out of the ordinary. On the second day, I volunteered to act as the operator to alleviate the tedium and to gain firsthand experience of the fault when and if it ever showed up.

On the third day, around midmorning, the system just froze. Nothing obvious was different, but I now had a clue because the fault recurred about 20 minutes later. The factory had a large work force who took tea and coffee breaks in groups, and large numbers descended on the tea and coffee machines at about the same time in an area adjoining the production floor.

I hooked up the mains analyzer to the incoming supply and discovered the problem. A couple of captured-event waveforms showed that, when all the drinks machines were operating at once, they produced cumulative patterns of noise, rising to well over the ratings for our standard industrial filtering. Our computers simply could not cope with the excessive interference.

It turned out that the drinks vendor had installed its machines soon after we had installed our systems, thus explaining the initial trouble-free period. The drinks machines also had little more than the most basic suppression circuitry, and all of them were wired to the same phase of mains supply as our systems. Changing the wiring on our systems to another phase immediately fixed the problems. EDN

Nigel Adams is a senior product specialist at Aeroflex Ltd (Stevenage, UK).

Do you have a Tale to share? Enter *EDN*'s Tell us your Tale contest, sponsored by Tektronix, and you could win a scope! For details, go to http://bit.ly/Tales_EDN.

DANIEL VASCONCELLOS

From extreme value to extreme performance.

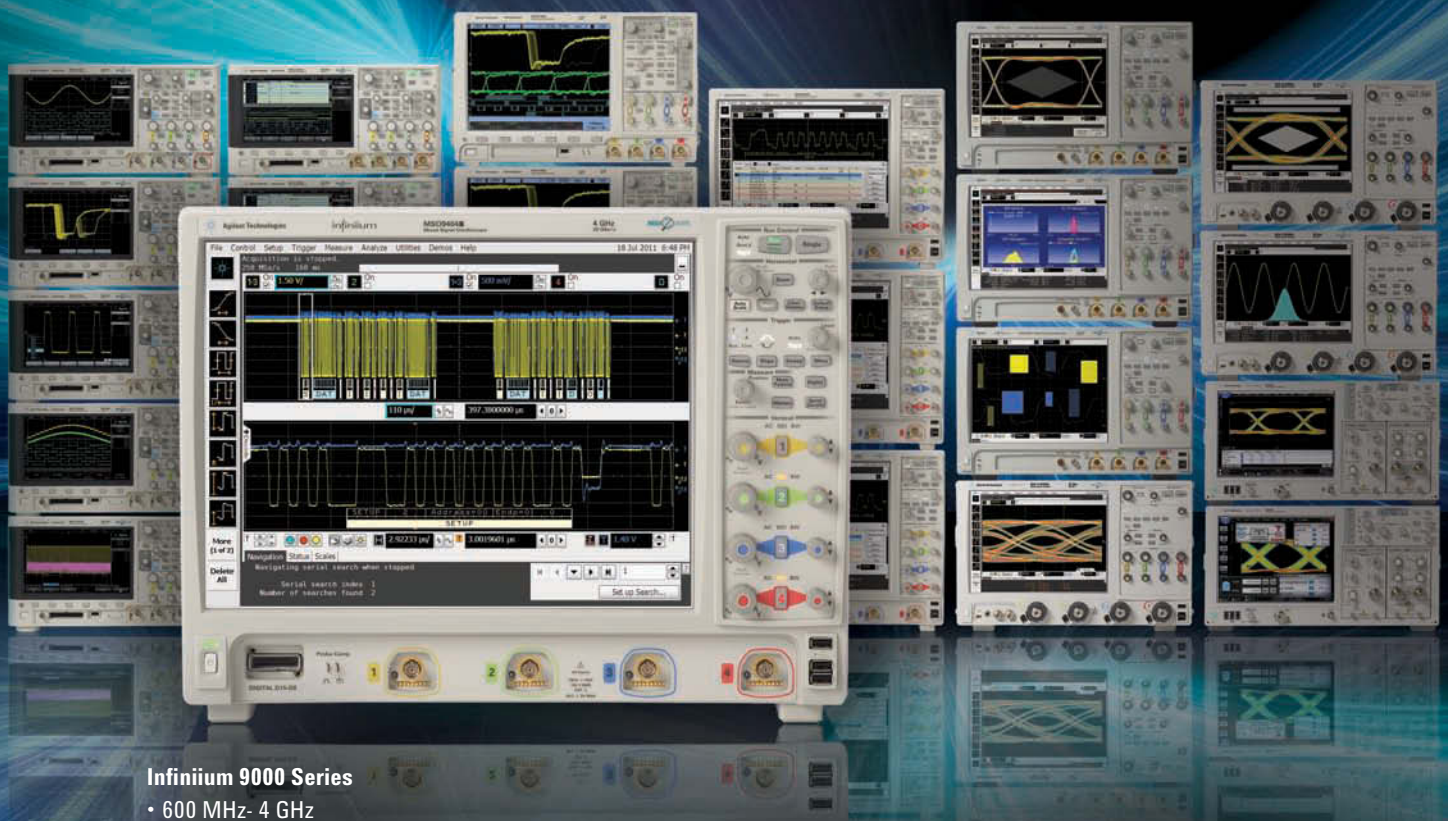

Infiniium 9000 Series

- 600 MHz- 4 GHz
- MSO and DSO models
- 20 GSa/s sample rate
- 40 Mpts memory std. (optional to 1 Gpt)
- >30 protocol, compliance and analysis options

Whether you need the fastest real-time oscilloscope on earth with the highest measurement accuracy anywhere, or something a little more basic, from 20 MHz to 90 GHz Agilent oscilloscopes outperform in every category. That's why discerning engineers have made Agilent the fastest growing oscilloscope company in the world since 1997. Like you, we're working on what's next.

That's thinking ahead. That's Agilent.

Scan code or visit http://goo.gl/NR3Xm for 9000 Series demo video

Buy the Best. Get Rewarded.
Find out how to get a complimentary 3000 X-Series scope
www.agilent.com/find/ScopeReward

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

Agilent Technologies

IDF 2011
INTEL DEVELOPER FORUM
San Francisco
September 13-15
Register Now. Change the Future.
intel.com/idf

EMBEDDED WITH INNOVATION: Virtualization

Squeeze more capability from a single platform.

Challenge >

Develop new embedded apps, while preserving investments in legacy software.

Solution >

Consolidate multiple operating systems and software stacks on a single platform.

Virtualization helps you push computing envelopes without expanding footprints. Intel® Xeon® and Intel® Core™ vPro™ processors let multiple OSs and software stacks co-exist, each in its own virtual machine. Real-time, general purpose, legacy, you name it. Run them all, without modification on a single device to boost performance, maximize security and reduce BOM cost.

+48%

Up to 48% performance boost in MMU-intensive benchmarks using Intel® processors.[2]

More powerful

Intel® vPro™ Technology

3 µsec interrupt latency attained by virtual machine monitor vendors using Intel® vPro™ technology hardware assist.[3]

Faster response

Virtualization.
Exclusive Webinar.

Learn how to optimize your embedded apps.

For platform consolidation, data protection, security and reliability, embedded virtualization is ready from prime time. Join the discussion, and see how Intel and our growing virtualization ecosystem are redefining the embedded world.

READ > WATCH > DISCUSS >>> **intel.com/embedded/innovation/virtualization**

intel™
Embedded