Reed Business Information

EDN

VOICE OF THE ENGINEER

OCT 2

Issue 20/2008

www.edn.com

VIRTUALIZATION: SILICON AND SOFTWARE SALVATION OR TECHNOLOGICAL TOWER OF BABEL?

Page 34

ON TIME, EVERY TIME: EMBEDDING REAL-TIME PERFORMANCE

Page 27

EXTENDING SPI4.2 CAPABILITIES FOR ETHERNET SERVICES

Page 47

EDN HANDS-ON PROJECT

The New Oscilloscope Experience Is Here

Quick Insight. Deep Insight.

Experience the New LeCroy Oscilloscopes.

www.insightwithconfidence.com

LeCroy

www.lecroy.com 1-800-5-LeCroy

YOUR SOURCE FOR SEMICONDUCTORS

MORE SUPPLIERS
MORE SELECTION*

Digi-Key® CORPORATION

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

Visit www.digikey.com for the industry's broadest selection of semiconductors!

Quality Electronic Components, Superior Service

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

www.digikey.com
1.800.344.4539

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

*New product added daily.

Hammer Down Your Power Consumption with picoPower™!

THE Performance Choice for Lowest-Power Microcontrollers

Performance and power consumption have always been key elements in the development of AVR® microcontrollers. Today's increasing use of battery and signal line powered applications makes power consumption criteria more important than ever. To meet the tough requirements of modern microcontrollers, Atmel® has combined more than ten years of low power research and development into picoPower technology.

picoPower enables tinyAVR®, megaAVR® and XMEGA™ microcontrollers to achieve the industry's lowest power consumption. Why be satisfied with microamps when you can have nanoamps? With Atmel MCUs today's embedded designers get systems using a mere 650 nA running a real-time clock (RTC) and only 100 nA in sleep mode. Combined with several other innovative techniques, picoPower microcontrollers help you reduce your applications power consumption without compromising system performance!

Visit our website to learn how picoPower can help you hammer down the power consumption of your next designs. PLUS, get a chance to apply for **a free AVR design kit!**

http://www.atmel.com/picopower/

Everywhere You Are®

Automotive Sensors

Circuit Protection Solutions

Magnetic Products

Microelectronic Modules

Panel Controls & Encoders

Precision Potentiometers

Resistive Products

How Will You Protect Your Circuits?

Circuit protection is a crucial part of any design decision today. No one understands that better than Bourns. We have developed the industry's broadest line of circuit protection solutions – backed by a global team of technical experts. Our customers are assured of finding the optimal solution for their application with Bourns' collective knowledge from years of circuit protection support. Beyond product depth and support, we have innovative new technologies that solve tough design challenges. Bourns' new gated thyristors provide faster and more efficient protection that reduces costs. Our symmetric gas discharge tubes regulate breakdown voltage innovatively. And Bourns' new polymer PTCs overcome voltage variability to extend product life. In the New World of Bourns, we have the technology to protect your circuits.

www.bourns.com/circuitprotection

Bourns - *the company you have relied on for more than 60 years.*

BOURNS®

Built on Trust…Based on Innovation

Green Engineering

Powered by National Instruments

MEASURE IT

Acquire environmental data from thousands of sensors

Analyze power quality and consumption

Present measured data to adhere to regulations

FIX IT

Design and model more energy efficient machines

Prototype next-generation energy technologies

Deploy advanced controllers to optimize existing equipment

For more than 30 years, National Instruments has empowered engineers and scientists to measure, diagnose, and solve some of the world's most complex challenges. Now, through the NI graphical system design platform, engineers and scientists are using modular hardware and flexible software to not only test and measure but also fix inefficient products and processes by rapidly designing, prototyping, and deploying new machines, technologies, and methods. Today, a number of the world's most pressing issues are being addressed through green engineering applications powered by NI products.

>> Download green engineering resources at **ni.com/greenengineering** **800 890 1345**

NATIONAL INSTRUMENTS™

EDN contents

10.2.08

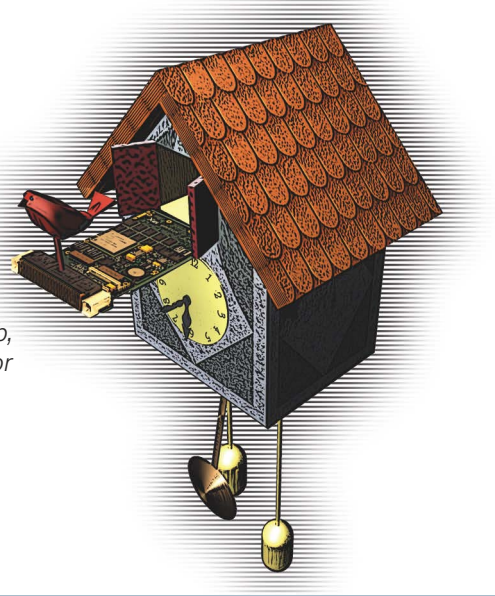

DESIGNIDEAS

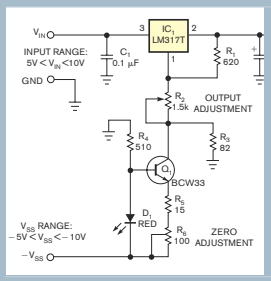

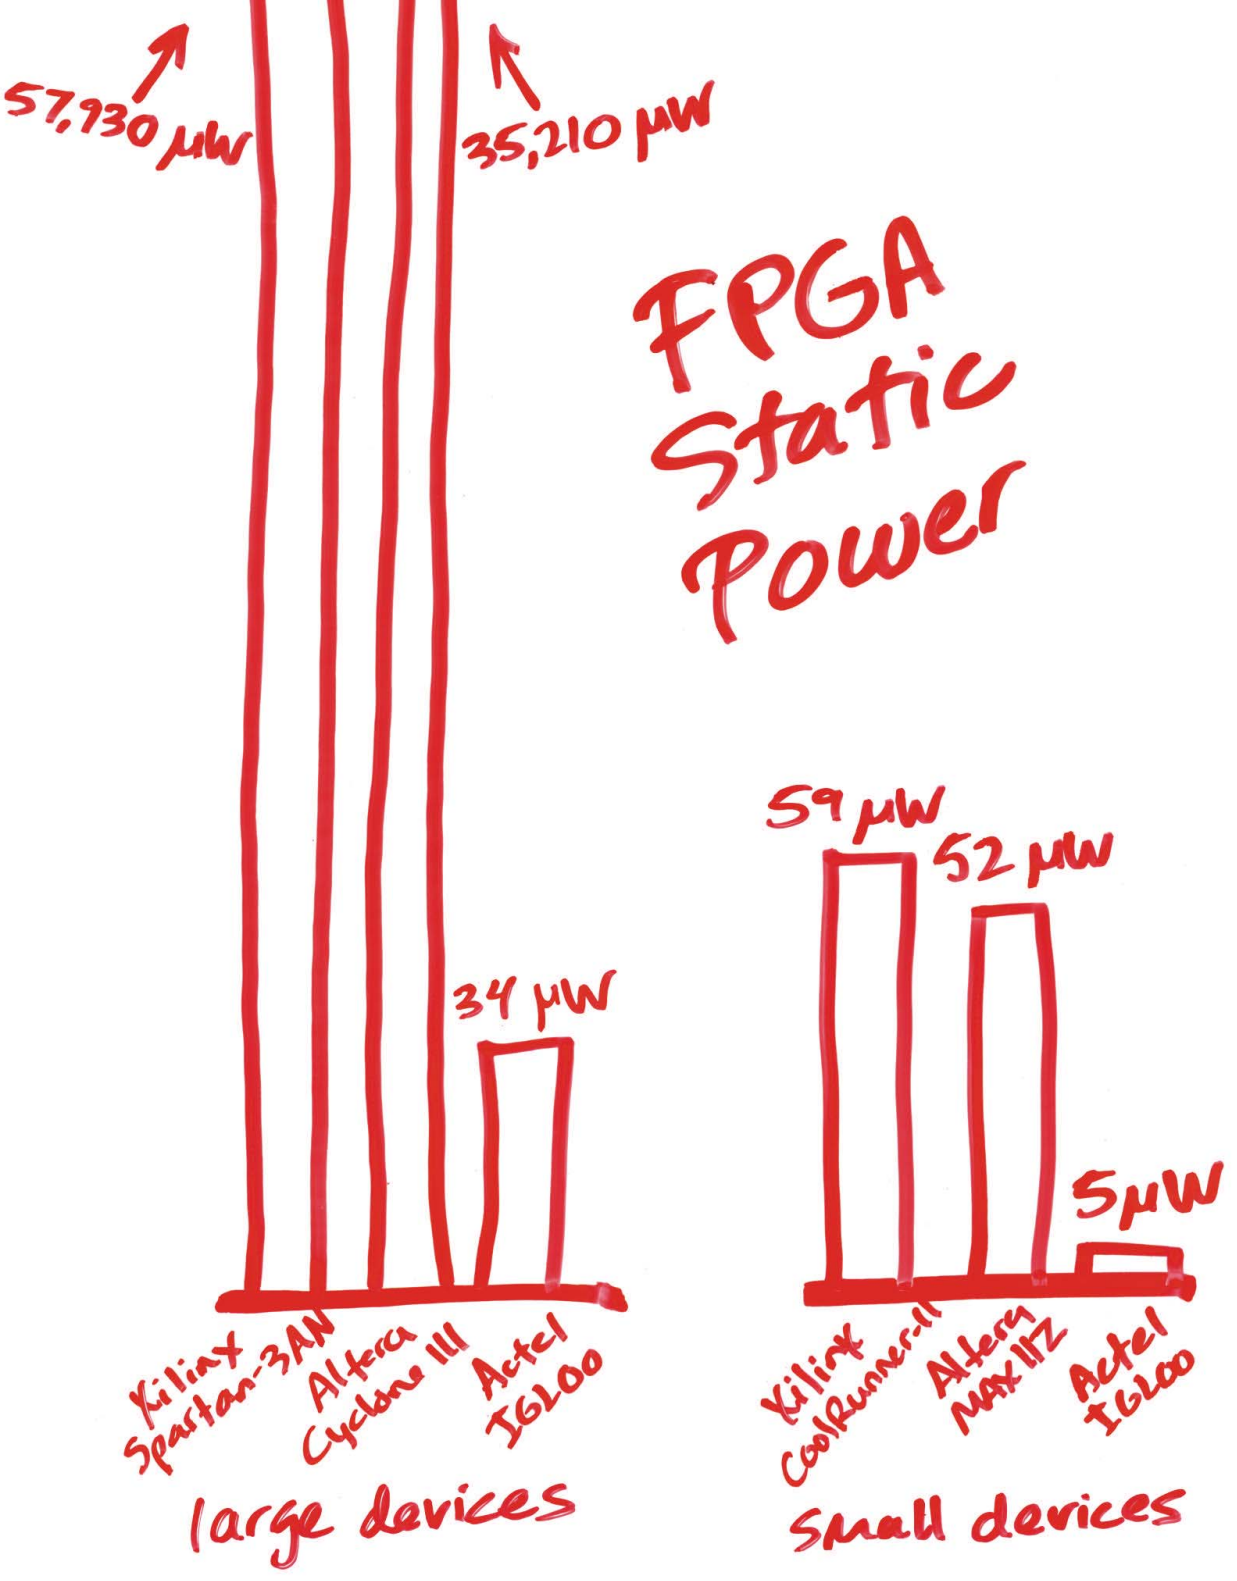

WHAT DO WE HAVE TO DO, DRAW YOU A PICTURE?

Only Actel gets you this close to zero. Any other claims of low power superiority are just that. According to their own data, Altera® and Xilinx® use between 10 and 1700 times the power of Actel IGLOO® FPGAs, depending on device and mode. Want specifics? Visit us to get the whole picture, including a video of actual measurements.

more proof and pictures at **actel.com/power**

Actel®
POWER MATTERS

contents 10.2.08

pulse

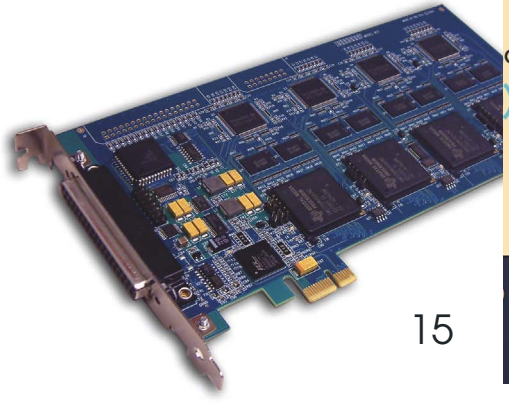

15

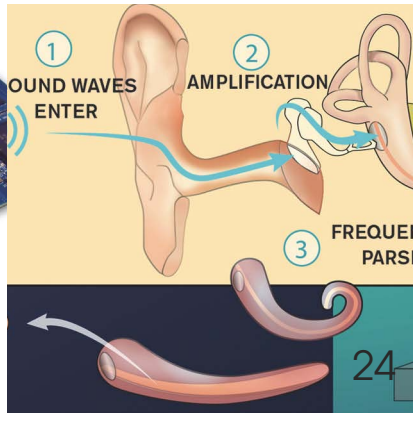

66

DEPARTMENTS & COLUMNS

PRODUCT ROUNDUP

EDN® (ISSN#0012-7515), (GST#123397457) is published biweekly, 26 times per year, by Reed Business Information, 8878 Barrons Blvd, Highlands Ranch, CO 80129-2345. Reed Business Information, a division of Reed Elsevier Inc, is located at 360 Park Avenue South, New York, NY 10010. Tad Smith, Chief Executive Officer; Mark Finkelstein, President, Boston Division. Periodicals postage paid at Littleton, CO 80126 and additional mailing offices. Circulation records are maintained at Reed Business Information, 8878 S Barrons Blvd, Highlands Ranch, CO 80129-2345. Telephone (303) 470-4445. POSTMASTER: Send address changes to EDN®, PO Box 7500, Highlands Ranch, CO 80163-7500. *Rates for nonqualified subscriptions, including all issues: US, $165 one year; Canada, $226 one year (includes 7% GST, GST#123397457); Mexico, $215 one year; air expedited, $398 one year. Except for special issues where price changes are indicated, single copies are available for $10 US and $15 foreign. Publications Agreement No. 40685520. Return undeliverable Canadian addresses to: RCS International, Box 697 STN A, Windsor Ontario N9A 6N4. E-mail: Subsmail@ReedBusiness.com. Please address all subscription mail to EDN®, 8878 S Barrons Blvd, Highlands Ranch, CO 80129-2345. EDN® is a registered trademark of Reed Elsevier Properties Inc, used under license. A Reed Business Information Publication/Volume 53, Number 20 (Printed in USA).*

The Newest Optoelectronics

EVERLIGHT PANJIT TOUCH SCREENS VISHAY LEDENGIN OPTEK KINGBRIGHT LITE-ON AVAGO LUMEX NORITAKE LAMINA SHARP TAOS NEWHAVEN DISPLAY PACIFIC SILICON SENSORS SEOUL SEMICONDUCTOR 3M TOUCH SYSTEMS

SEOUL

Acriche LEDs
www.mouser.com/seoul/a

New Products from:

AVAGO TECHNOLOGIES

PCB PolyLEDs
www.mouser.com/avagotechnologies/a

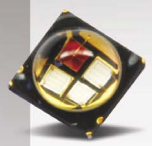

LedEngin, Inc.

10W RGB LED Emitters
www.mouser.com/ledengin/a

EVERLIGHT

XcelLed™ 1W Emitters
www.mouser.com/everlight/a

The ONLY New Catalog Every 90 Days

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

MOUSER® ELECTRONICS

a tti company

The Newest Products For Your Newest Designs

(800) 346-6873

www.mouser.com

Over A Million Products Online

EDN at edn.com

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

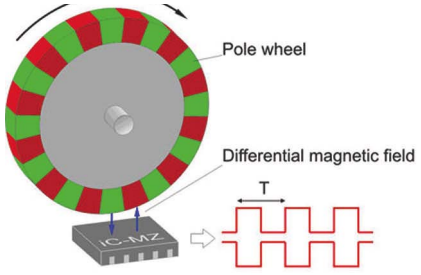

ONLINE ONLY

Check out these Web-exclusive articles:

Speed acquisition made simple
Speed acquisition is of central importance in the control and monitoring of motion processes. You can simply and efficiently tackle the task using magnetic sensors based on the Hall principle.
→www.edn.com/article/CA6594313

Challenges and design decisions for measuring multicore performance
The type of multicore processor you choose and the type of parallelism you apply in your application code will greatly affect the performance you will achieve. To ensure that you meet your goals, closely examine the benchmark scores produced by an industry-standard suite of multicore benchmarks.
→www.edn.com/article/CA6592752

The hard-drive click of death: archaic technologies and their protracted viabilities
Why are we still using rotating magnetic mass storage that's powered by a motor spinning at 4200, 5400, 7200, or 10,000 rpm, with a rapidly vacillating read/write head suspended only a few millionths of an inch above its platter mate?
→www.edn.com/081002toc1

HOT TOPICS

Check out *EDN*'s new Hot Topics pages: From WiMax to H1-B visas to automotive electronics, Hot Topics pages deliver continuously updated subject-specific links from not only *EDN* but also the entire electronics-industry Web.
→www.edn.com/hottopics

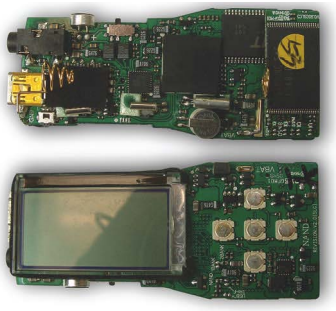

READERS' CHOICE

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

Prying apart a portable audio player
A look inside Sandisk's Sansa M250 reveals platform-minded design decisions, surprising flash-memory choices, and potential hacks.
→www.edn.com/article/CA6590195

Lithium-ion technology targets portable power
→www.edn.com/article/CA6586222

MIT researchers use plant-energy-storage system for solar-storage innovation
→www.edn.com/article/CA6583673

Handcrafted analog gets automated assist
→www.edn.com/article/CA6590188

Transimpedance-amplifier stability is key in light-sensing applications
→www.edn.com/article/CA6590187

If you can't beat patent trolls, join them
→www.edn.com/article/CA6594114

NEW ONLINE EXCLUSIVE

THE IC INSIDER

Get a peek inside the latest ICs—literally—courtesy of *EDN*'s new online feature, The IC Insider. This monthly exclusive, contributed by reverse-engineering company Chipworks, couples actual silicon imagery with circuit schematics and provides a technical discussion of some of the innovative features of a recently introduced IC. The premiere installment looks at Maxim's MAX8814ETA battery-charger chip.
→www.edn.com/icinsider

BY PAUL RAKO, TECHNICAL EDITOR

Why tout a demo board nobody can buy?

Recently, my buddy Dave was trying to make a remote-control airplane that would transmit high-definition video to the operator on the ground. Because he wanted to use the video in real time to control the plane, he needed a system with minimal latency. This requirement caused Dave to consider systems that transmit uncompressed video. Dave had been exploring a lot of technologies, and one that caught his eye was a chip set from Amimon. The chip set transmits high-definition video as a 40-MHz-wide signal in the 5.8-GHz ISM (industrial/scientific/medical) band. The company targets this chip set for use in consumer electronics because every TV and output device, including DVD players and set-top boxes, can handle uncompressed video.

Amimon's Web site lists not only the chip set, but also both a module and a demo board. Dave contacted the company regarding the price of the demo board, but the company refused to provide it. I then wrote to the company and asked the price. The company's public-relations agency responded: "There is no set pricing for the development kits; it depends on vendor needs." Dave needed a free one but realized that he couldn't qualify for a free sample because he was just at the hobby stage.

Weeks later, Dave e-mailed me. "I got an e-mail reply from Amimon," he wrote, stating that he had received the same response as the one I had received. "They asked 'What would be the cooperation model with Amimon?' Huh? There must be some supersecret marketing-indoctrination cult … that screws up people's minds so completely that they cannot speak in plain English. Put the darned development kit on Digi-Key for $800, and I'll buy it!" He went on to point out that many companies, including National Semiconductor, Linear Technology, Fairchild, Texas Instruments, and Analog Devices, all put their demo boards up for sale through distributors such as Digi-Key. "I'm dreaming of a day," he continued, "when 'call for pricing' and 'contact your local rep' will be banned from industrial suppliers' Web sites."

My friend Dave is "dreaming of a day when 'call for pricing' and 'contact your local rep' will be banned from industrial suppliers' Web sites."

I believe that all this turmoil probably derives from the fact that marketing types and engineers speak different languages. Marketing types are dying to talk to people, go to lunch, and network. Engineers, on the other hand, want to crawl into a corner with some hardware and be alone. Industrial suppliers Digi-Key, Mouser, Newark, and Jameco have made things a lot better than they were a decade ago, though, when most manufacturers insisted that you contact the factory for samples, and the big distributors only wanted to sell you a reel of 4000 parts.

The marketing types are trying to uncover the next high-volume-sales opportunity so that they can schmooze with you and beat their competitors. When engineers are producing prototypes, however, they have neither the time nor the inclination for all this socializing. The last system board I designed, a point-of-sale terminal, had 100 parts and 20 ICs. If every manufacturer had required me to call and undergo qualification just to get two bucks' worth of samples, it would have taken a month. With only two months to design and build the whole product, this approach would have been unfeasible.

How about you? Do you hate the fact that manufacturers won't give you a price and don't sell small quantities through distribution? Sound off at *EDN*'s Web site, www.edn.com/081002ed. EDN

Contact me at paul.rako@edn.com.

Go to www.edn.com/081002ed and click on Feedback Loop to post a comment on this column.

More at www.edn.com/edncomment

More Q. Less Cu

Copper
63.546

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

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

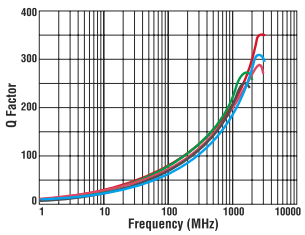

Q factors are 3X higher than standard chip inductors

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

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

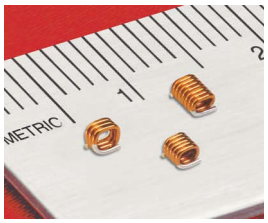

These coils are significantly smaller than existing air core inductors. We reduced the footprint by using close-wound construction and keeping the leads close to the body. The square shape cuts the height to as low as 1.5 mm and creates flat top and bottom surfaces for easy automated handling and stable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

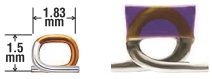

The square shape and narrow footprint reduce board space by 60-75% over conventional air core inductors.

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

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

RoHS COMPLIANT

Coilcraft®

www.coilcraft.com 800/322-2645

Power That Gives You The Best Of Both Technology Worlds

Switcher Efficiency Combined With LDO Noise And Transient Performance

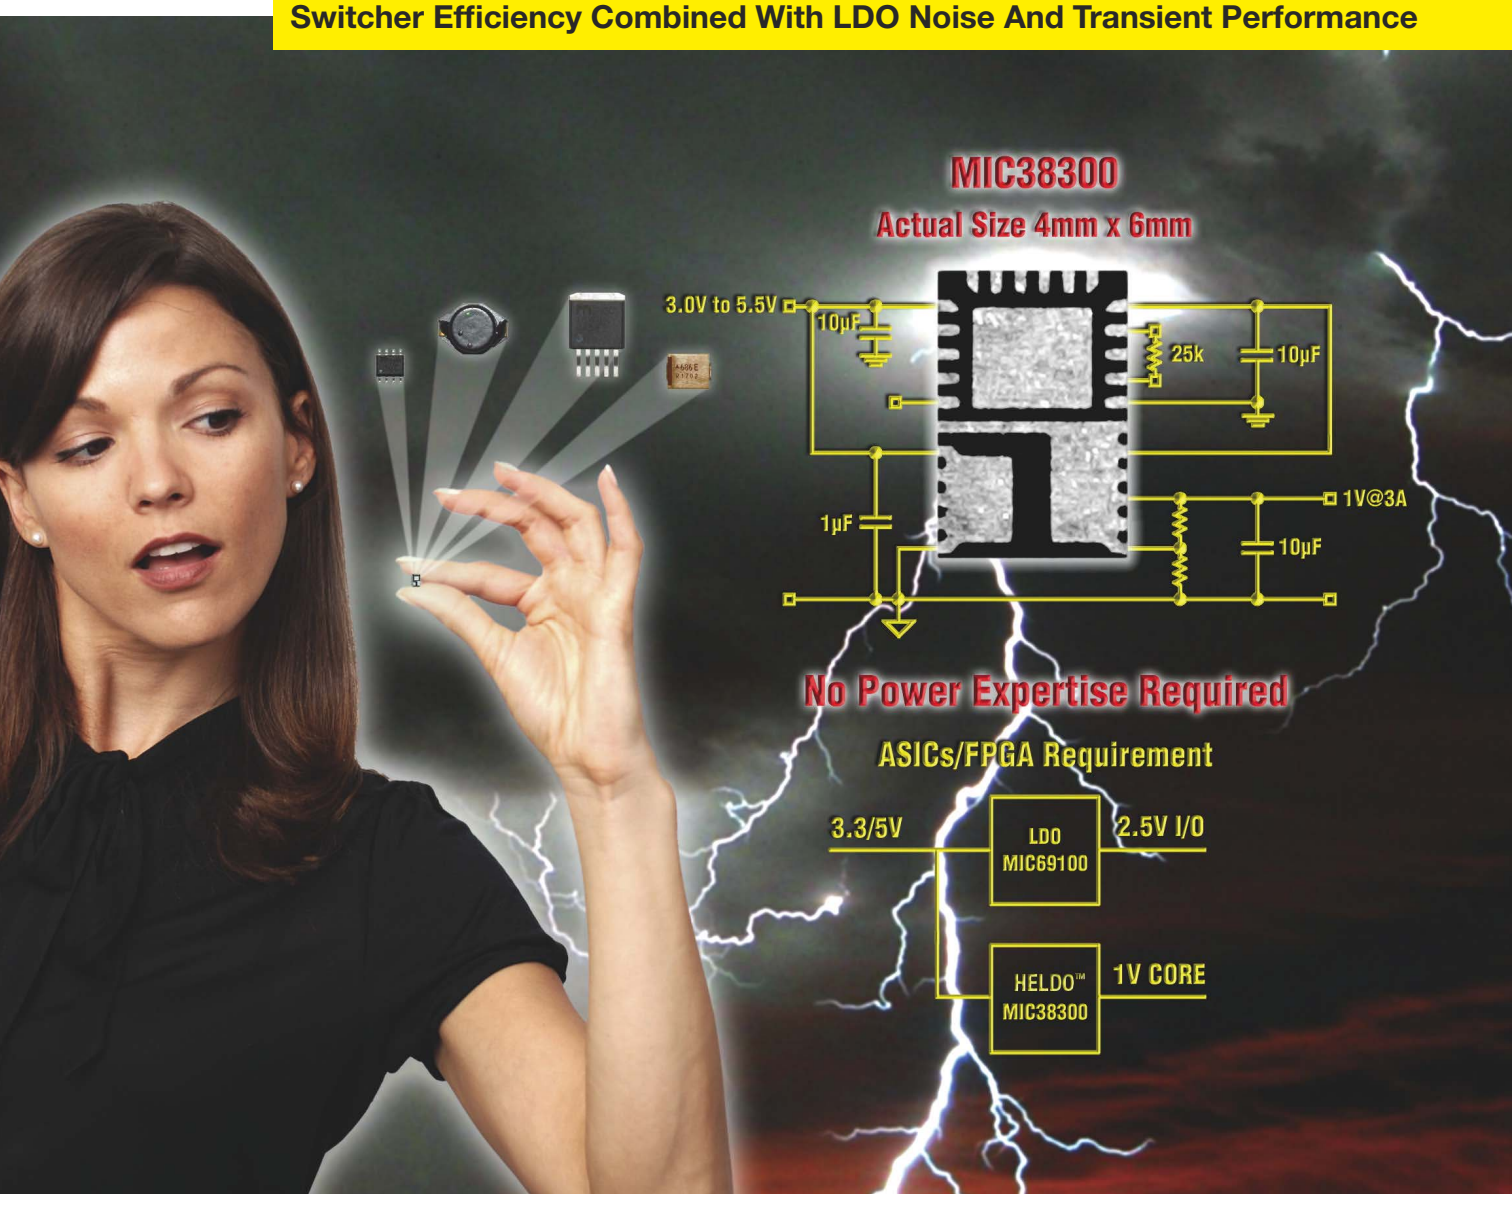

The MIC38300 is a 3A step down converter and the first device in a new generation of HELDO™ products providing the benefits of LDOs with respect to ease of use, fast transient performance, high PSRR and low noise while offering the efficiency of a switching regulator.

As output voltages move lower, the output noise and transient response of a switching regulator become an increasing challenge for designers. By combining a switcher whose output is slaved to the input of a high performance LDO, high efficiency is achieved with a clean low-noise output.

For more information, contact your local Micrel sales representative or visit us at: www.micrel.com/ad/mic38300.

The Good Stuff:

- 2.2A Continuous operating current
- Input voltage range: 3.0V to 5.5V
- Adjustable output voltage down to 1.0V
- Output noise less than 5mV
- Ultra fast transient performance
- Unique Switcher plus LDO architecture
- Fully integrated MOSFET switches
- Micro-power shutdown
- Easy upgrade from LDO as power dissipation becomes an issue
- Thermal shutdown and current limit protection
- 4mm × 6mm × 0.9mm MLF ® package

MICREL®
Innovation Through Technology™
www.micrel.com

EDN

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

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

ASSOCIATE PUBLISHER, *EDN* WORLDWIDE
Judy Hayes, 1-408-345-4437; judy.hayes@reedbusiness.com

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

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

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

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

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

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

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

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

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

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

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

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

SEMICONDUCTOR MANUFACTURING AND DESIGN
Ann Steffora Mutschler, Senior Editor
1-408-345-4436; ann.mutschler@reedbusiness.com

DESIGN IDEAS EDITOR
Martin Rowe
edndesignideas@reedbusiness.com

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

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

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

Adam Odoardi, Prepress Manager
1-781-734-8325; fax: 1-303-265-3042

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

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

PRODUCTION
Dorothy Buchholz, Group Production Director
1-781-734-8329
Kelly Jones, Production Manager
1-781-734-8328; fax: 1-303-265-3164
Linda Lepordo, Production Manager
1-781-734-8332; fax: 1-303-265-3015

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

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

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

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

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

Reed Electronics Group

See for Yourself at
www.xilinx.com/spartan

LOWEST TOTAL COST... PERIOD.

GET UP TO 50% LOWER COST

- Integrated features and only two power rails minimize need for external components
- Save up to 70% in logic cell resources with dedicated DSP blocks
- Run cool with 11mW static power, 0µW in hibernate mode

In the highly competitive high-volume market, cost is king. Our latest Extended Spartan®-3A FPGAs deliver the integrated features, low static power, complex computation and embedded processing capabilities you need to achieve the absolute lowest total cost. Period.

Combine these advantages with the industry's largest selection of IP cores, reference designs and I/O standards and you have the most complete low-cost programmable solution available for your next high-volume design.

Visit us at **www.xilinx.com** to download our free ISE® WebPACK™ design tools and start saving money today.

XILINX®

EDITED BY FRAN GRANVILLE

pulse

INNOVATIONS & INNOVATORS

Video card captures 480 frames/sec

Extending its line of OEM video-capture cards, Sensoray recently announced the model 817 PCIe (peripheral-component-interconnect-express) frame-grabber card, which captures 16 channels of compressed JPEG or uncompressed bit maps at speeds as high as 480 frames/sec. Supporting one, four, eight, or 16 PCIe slots, the board allows the user to set independent capture parameters for each channel. An internal 16×4 analog cross-point-video switch routes any combination of four input channels to external video monitors. Users can turn each of the four video outputs on or off, allowing the outputs of multiple cards to drive one monitor.

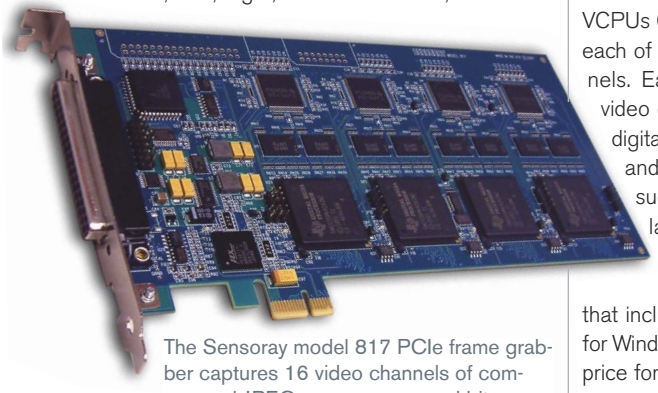

The Sensoray model 817 PCIe frame grabber captures 16 video channels of compressed JPEG or uncompressed bit maps at 480 frames/sec.

The model 817 contains four identical VCPUs (video-capture and -processing units), each of which handles four input-video channels. Each VCPU employs a four-channel video decoder to convert analog video into digital and a DSP to capture digitized video and to handle various processing tasks, such as frame decimation, caption overlay, JPEG compression, and status reporting. Sensoray provides a software-development kit for the module that includes drivers and sample applications for Windows and Linux operating systems. The price for the model 817 starts at $705 (OEM quantities).**—by Warren Webb**

▷**Sensoray**, www.sensoray.com.

FEEDBACK LOOP

"I've found that reproducing the failure is half of the problem, and, usually, ... the logistics of implementing the fix are the other half."

—Reader Ron Bauerle, in *EDN*'s Feedback Loop, at www.edn.com/article/CA6578156. Add your comments.

Book helps you resolve BGA-layout issues

As a circuit designer who also often lays out quite complex boards, I found the temptation to purchase Charles Pfeil's new book—*BGA Breakouts & Routing: Effective Design Methods for Very Large BGAs*—irresistible, and I wasn't disappointed. The approximately 160 pages and more than 100 figures within this quality paperback carry a wealth of information that far exceeds what I've amassed during my attempts over the past seven years that I've been designing with BGA (ball-grid-array) packages. And, at approximately $31, this invaluable information comes cheap.

Although Pfeil is engineering director at Mentor's system-design division and the colorful, instructive screenshots that adorn virtually every page originate from his company's products, this is no self-serving, Mentor-driven publicity exercise. Rather, it is a usable resource that tackles PCB (printed-circuit-board) designers' problems as BGA packages become ever denser.

The first two chapters explain BGA-specific concepts, terminology, and packages. In-depth discussions of HDI (high-density-interconnect)-layer stackups, fan-out patterns, and layer-biased breakouts follow. The author concludes with a chapter that explores alternative approaches to routing a 1760-ball device on a 0.8-mm pitch. This exercise gives great insight about how well different approaches may work for you, even with the smaller BGA packages that most designers use.**—by David Marsh**

▷**Mentor Graphics**, www.mentor.com/go/bga.

Analyzer/exerciser validates 5-Gbps SuperSpeed USB 3.0

Although USB (Universal Serial Bus) 3.0 with its 5-Gbps SuperSpeed technology probably won't reach the market in volume until mid-2010, manufacturers are now developing computers and peripheral devices employing the new bus. To work on these products and their all-important supporting software—especially drivers—engineers need tools for exercising developmental devices and software, analyzing bus performance, ensuring compatibility with earlier versions of USB, and debugging flaws. According to LeCroy Corp, the needed tools weren't commercially available until its August introduction at the Intel Developer Forum (www.intel.com/idf) in San Francisco of the Voyager validation system, which LeCroy describes as the first protocol analyzer/exerciser for USB 3.0.

Using custom front-end circuits, which the company developed for its 5-Gbps PCIe (peripheral-component-interconnect-express) analyzer/exerciser, the instrument performs simultaneous protocol capture of USB 2.0 and 3.0 signaling. An integrated exerciser option enables the sixth-generation unit to completely test and validate the performance of USB devices, systems, and software.

Engineers often refer to USB 3.0's SuperSpeed mode as more than 10 times as fast as USB 2.0's high-speed mode. That characterization may or may not be valid. SuperSpeed operation uses 8-bit/10-bit coding, which reserves 20% of the transmitted bits for error correction, thus allowing the 5-Gbps data stream to transmit a maximum of 4 Gbps of real data. The USB 2.0 high-speed mode, which does not incorporate error correction, can transmit bursts of data at 480 Mbps, or 12% of 4 Gbps. In sustained operation, however, the maximum high-speed-mode data rate is closer to 300 Mbps. Because the SuperSpeed protocol may prove more efficient than the high-speed protocol, USB 3.0 may more than compensate for its error-correction overhead.

Another feature of USB 3.0 is adaptive equalization, which optimizes a link's output waveshape in response to training sequences that run automatically when you establish a new connection. Unlike buses that mainly find use in backplanes, USB is primarily a cabled bus. Because such buses use cables of unpredictable length and because waveform distortion depends strongly on path length, adaptive equalization is essential to achieving the advertised data rate.

LeCroy regards the situation with USB 3.0 connectors as potentially still fluid. Whereas manufacturers have designed and prototyped connector types that are close to receiving approval from the USB Implementers Forum (www.usb.org), modified or added types are still possible. Therefore, the company designed the Voyager system to make it easy for users to change the bus connectors. Users can substitute SMA connectors for USB connectors if the "real" connectors are temporarily unobtainable.

USB 3.0 will be backward-compatible with USB 2.0 but in a different way from the way that USB 2.0 is backward-compatible with USB 1.1. USB 2.0 cables mate and work with USB 1.1 connectors. USB 3.0 devices will have connectors that mate with both USB 2.0 and USB 3.0 cables, but, if you want to use a USB 3.0 device in a system based on USB 2.0, you will need a USB 2.0 cable, which may be an optional accessory to the USB 3.0 device. According to LeCroy, testing of USB hubs will make it nearly essential that USB 3.0 analyzer/exercisers support simultaneous USB 2.0 and 3.0 data capture.

LeCroy says that fiber-optic implementations are likely to emerge in the future but that initial implementations will use copper cables. US prices for the Voyager system begin in the $20,000 to $30,000 range.**—by Dan Strassberg**

▷**LeCroy Corp,** www.lecroy.com.

LeCroy calls the Voyager validation system the first protocol analyzer/exerciser for 5-Gbps USB 3.0. The unit supports simultaneous USB 2.0 and 3.0 data capture, which is necessary for testing USB 3.0 hubs.

DILBERT By Scott Adams

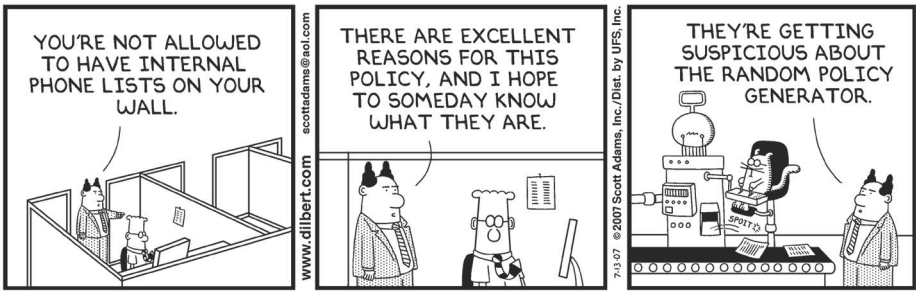

Capacitive Touch Sensing In a Flexible, Single-chip Solution

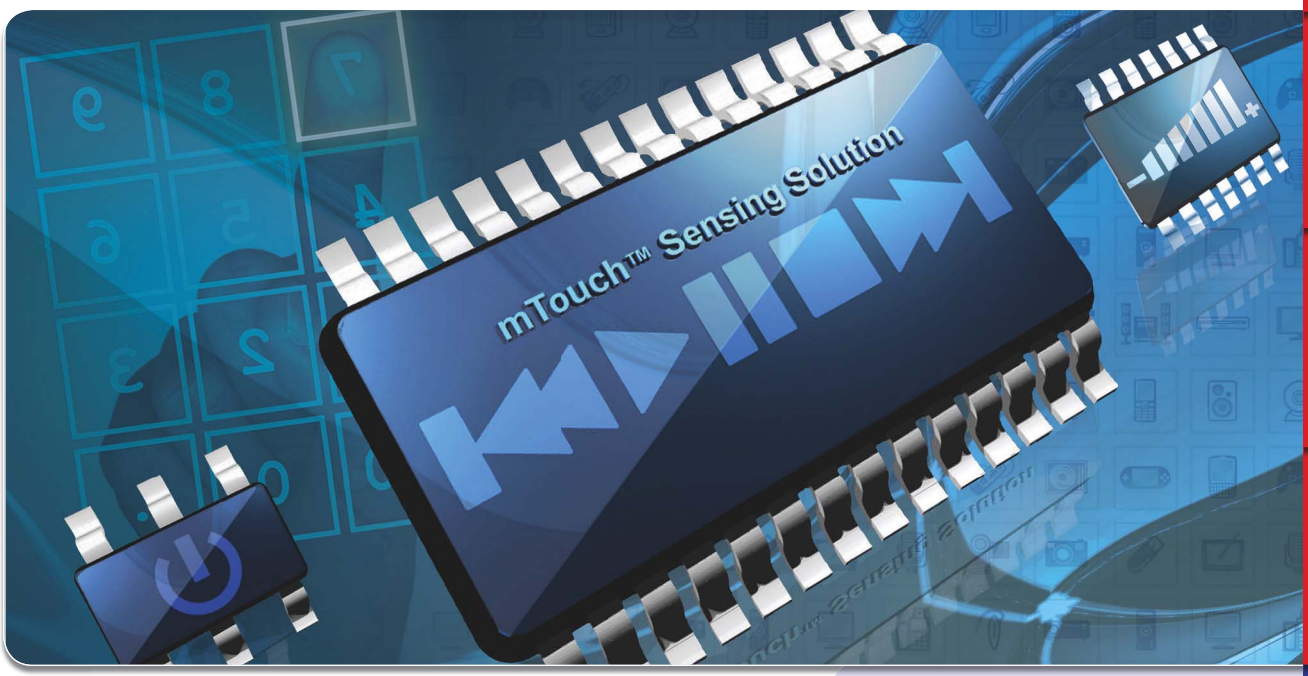

Capacitive touch interfaces provide an excellent way to add low-cost, reliable and stylish buttons into your design. Microchip Technology's mTouch™ Sensing Solution includes comprehensive development kits and a free diagnostic tool to make implementation easy and fast. Our free source code can be seamlessly integrated with your existing firmware on a single PIC® microcontroller – eliminating the need for additional controllers.

THE mTouch SENSING SOLUTION FEATURES:

- FREE license libraries and source code
- A FREE diagnostic tool
- Integration with 8- and 16-bit PIC microcontrollers
- Easy expansion, with support from 6 to 100 pins
- Low-power operation

GET STARTED IN 3 EASY STEPS

1. Visit the mTouch Sensing Solutions design center at **www.microchip.com/mTouch**
2. Download FREE libraries and source code
3. For a limited time, save 20% off a variety of touch sensing development tools when you purchase from **www.microchipDIRECT.com** and use coupon code **mTouch4U**.

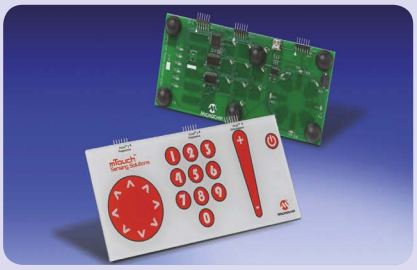

MICROCHIP

Automotive EEPROMs use two cells per bit for ruggedness, reliability

With the aim of ensuring reliable operation and long life in the demanding automotive environment, Rohm has introduced a series of EEPROMs that employ a double-cell structure. The BR25Hxx0 devices, which are SPI (serial-peripheral-interface)-bus memories, withstand the voltage surges, static discharge, heat, and vibration that they may encounter in automotive ECUs (engine-control units). Features include guaranteed operation at 125°C, with 20-year data retention at 85°C, 6-kV ESD (electrostatic-discharge) resistance, gold-pad/gold-wire connections, and a double-reset function for high reliability. The devices' high-redundancy circuits and a choice of process technology allow the company to specify a life of 1 million rewrites at 85°C and 300,000 rewrites at 125°C.

EEPROMs are increasingly finding use in critical automotive systems to record, for example, status information. As most users know, the devices also have a wear-out mechanism; charge is transferred by charge tunneling across an oxide barrier. Eventually, on a lifetime-related but somewhat-random basis, the oxide in a cell may break down. Rohm addresses this breakdown by fabricating two identical but separate memory cells for each bit and connecting them in an OR configuration, so that the combined cell will continue to function after a failure.

To protect against error conditions during power-up or voltage dips, the EEPROMs integrate a double-protection circuit comprising a POR (power-on-reset) block that resets during start-up of the power source and a low-voltage write-error-correction circuit that resets the internal circuit and prohibits write operations in low-voltage conditions.

Rohm produces variants with Microwire, I²C (inter-integrated-circuit), and SPI interfaces; the 5-MHz SPI is the fastest. The 14 devices in the family come in 6.2×5×1.71-mm SOP-8 and 6×4.9×1.75-mm SOP-J8 packages and in capacities of 1, 2, 4, 8, 16, and 32 kbits. The memories support high-speed-write and page-write modes.

—by Graham Prophet

▷**Rohm Electronics**, www.rohm.com.

EDN**BLOG**

ANABLOG

Henry Ott EMC seminar, San Francisco, Oct 15 to 17

I just got a flyer from Henry Ott about a three-day seminar on electromagnetic-compatibility engineering (www.hottconsultants.com/emccourse_3-day.html). Cost is $1375. Henry is a former Bell Labs researcher that really understands signal integrity, EMC, and RFI issues in your design. He has done great work in figuring out layout issues on PCBs.

I have blogged before about Henry's assertion that you should not cut up ground planes. I wholeheartedly agree; it usually causes more problems than it solves. The whole cutting-up-the-plane suggestion came from application engineers working at semiconductor companies, who never had to get a whole system working. The ADC guys would plop an ADC on an evaluation board and warn you to separate the analog

and digital grounds. Easy to say, but what do you do if you have multiple converters? How do you handle multiple power supplies?

Henry says that you should keep one plane, and he prefers to call it a reference plane rather than a ground plane since most designs have a chassis common, not an earth ground. Then, he says, you should use component placement and routing discipline in order to keep the digital signals out of the analog circuits. He has presented some great work on how far away traces have to be for them to not interfere. It makes me love strip line buried between planes, rather than microstrip that can radiate out 12 times its width.**—by Paul Rako**

▶www.edn.com/anablog.

▶For the full post, go to www.edn.com/081002b1.

By every measure, Analog ICs mean better data in less time.

Our Latest Data Acquisition Innovations

Programmable Gain Instrumentation Amplifiers
AD8250, AD8251, and AD8253
PGIAs offer unmatched speed and precision; require 75% less power and are up to 90% smaller than competing in-amps.

High Precision Dual Amplifier
OP2177
60 μV maximum offset, 0.7 μV/°C maximum drift, 8 $\mu V/\sqrt{Hz}$ typical low noise, and low power consumption for unmatched performance.

3-Channel, Low Noise, Low Power, 16-Bit and 24-Bit Σ-Δ ADCs
AD7792 and AD7793
Complete AFEs for high precision measurement; on-chip amplifier allows direct interface of small amplitude signals to the ADC.

16-Bit, 10 MSPS, 92 dB SNR, 130 mW Power PulSAR® ADC
AD7626
Breakthrough performance enables faster, more precise data acquisition system designs.

14-Bit, 125 MSPS, 85 dBc SFDR, Low Power ADC
AD9246
Combines excellent SFDR over a wide input bandwidth with low aperture jitter to accurately sample high frequency signals.

High Speed Digital Isolators
ADuM140x
Compared to optocouplers, ADI's 4-channel digital isolators offer faster data rates, greater reliability with isolation ratings up to 2.5 kV, and up to 70% savings in board space.

See why more engineers want ADI experience inside their data acquisition designs.

Engineers working on data acquisition (DAQ) applications depend on Analog Devices to meet the demands of their designs and the expectations of their customers. We've taken our industry-leading data converters and amplifiers, added newly developed instrumentation components—such as high speed digital isolators and low jitter clocks—and optimized them for DAQ applications. The result is unmatched component functionality, reliability, and compatibility with proven specifications to shorten the design process and yield readings you can trust.

To learn more about what Analog Devices can do for your own world-class designs, order samples, or download data sheets, visit *www.analog.com/daq-ad1*.

analog is everywhere.™

www.analog.com/daq-ad1

ANALOG DEVICES

VOICES
CriticalBlue's David Stewart

David Stewart is founder and chief executive officer of CriticalBlue (Edinburgh, Scotland) and has more than 20 years' experience in the EDA and semiconductor industries, 10 of which he spent at Cadence Design Systems, where he was a founder and business-development director of the SOC (system-on-chip)-design facility at the Alba Campus in Scotland. This initiative attracted worldwide interest, and the design center grew to more than 200 people in its first 18 months. For an expanded version of the following interview, go to www.edn.com/081002p1.

What are the challenges that your customers are now facing?

A We straddle the hardware/software-development and -design boundaries, so we see customers on both sides. In a sense, there's one challenge right there, which has been identified for a number of years as a key issue in that hardware and software have typically been designed separately, and the two disciplines don't communicate well with each other. I don't think that problem has been solved yet, although it's better.

At a company level, what I see people challenged with is how to manage the necessary investment in building silicon platforms with getting the most out of those platforms. The biggest implication of that [challenge] is that these silicon platforms need to be a lot more programmable than they were before, so the increasing use of processors [is a concern], of course, but on top of that is layered the issue of power consumption, implying a smaller number of processing elements working together to deliver the same performance at lower power consumption.

How well has the industry dealt with this problem?

A I don't think we have dealt with it. In many cases, particularly with respect to the multicore programming, we have expected some new panacea to suddenly appear, and there is a group of people who have been waiting for that to happen, and there are other people that are getting on with doing things.

What history has taught me is that engineers tend to grow in an evolutionary way; they don't tend to suddenly throw something away and start at the beginning. So, a brand-new approach is an interesting idea, but I don't see anybody with much appetite for implementing anything like that at this point because, specifically, when the markets are tough and people are being cautious … people are much more likely to stick with what they know and try to evolve it than to throw everything away and make a huge bet on something that's brand new.

But will a new design be required in the future?

A In an ideal world, that might be an interesting solution, but in terms of practically getting something that people might adopt, I struggle to find examples of where that's been done in the past. There aren't a lot of examples of brand-new things coming in and being adopted widespread across the industry; that's usually not what happens.

Are the multicore and hardware/software problems the same issues?

A They are different issues. The issue of hardware/software design is when you are building a base platform and you are trying to decide what to put on it, how to configure it, which pieces to put in hardware and which to put in software, and so on, and there are a lot of challenges with that [issue].

Then, you've got the people who are actually using the platforms—the end customers, if you like—who plan to program them and use them. And they are not designing hardware; they are purely interested in getting the best performance out of the platform they've been given and needing to learn about multicore-software programming in order to do that. There is a connection, but they are two different challenges.

What does CriticalBlue focus on in the multicore and hardware/software areas?

A In the hardware/software area, we developed a technology to allow the direct migration of software functionality into a hardware coprocessor. So, in other words, we

developed a methodology that allows you to use software to design the hardware, which is something that doesn't really exist at this point.

In the multicore space, we've been doing a lot of work to help people analyze software that they have and figure out how to redeploy it on multicore architectures—for example, a single, standard RISC core and then multiple coprocessors. So, you might look at analyzing some software running on an ARM processor and what you need to do to that software to be able to put it onto multiple coprocessors as well as the ARM.

What is CriticalBlue's approach to that issue?

A We're focused very much on pragmatic solutions that help people work within the environments and the software codes that they already have, and it harks back to the point about whether we'll ever see a shiny new language in which everything can be captured and expressed. We may do, but we've certainly been working on the principle that we're going to stick with the existing languages and that people have a lot of existing software that they wish to be able to redeploy into these multicore architectures, and they need help with how to do that.**—Interview conducted and edited by Ann Steffora Mutschler**

Signal processing ICs to help redefine what automobiles can be.

Automotive functions enhanced by our innovative ICs

Hybrid Electric Vehicle Systems
Efficiency means economy. Increased vehicle range from high voltage batteries, electric motors, and power conversion is achieved through application specific Analog Devices ICs: resolver-to-digital transmission controllers, *i*Coupler® digital isolators, and other precision linear ICs.

Safety Systems
Protection provides confidence. Airbag and vehicle stability systems with added functionality are made possible with *i*MEMS® gyroscopes and accelerometers. Customized radar and lidar designs for adaptive cruise control, lane departure warning, and parking assist are enabled by ADI's converters, amplifiers, comparators, and PLL synthesizers.

Infotainment Systems
Enjoy the ride. High fidelity decoding, equalization, and filtering in audio systems are provided by SHARC® and SigmaDSP® processors, codecs, and amplifiers. Video amplifiers and encoders/decoders enhance displays and monitors. Blackfin® processors optimize software-defined radio, hands free communications, head unit control, and iPod® connectivity.

Economical. Safe. Fun.

By offering the industry's most complete range of signal processing technologies and expertise, Analog Devices provides the functional building blocks engineers need to design safer, more enjoyable, and more environmentally friendly automobiles. We deliver ICs with industry-best functionality for concert hall quality audio, improved fuel economy, reduced emissions, advanced driver safety, comfort, and convenience.

ADI has earned a reputation for dependable IC performance and reliability in demanding applications—such as data acquisition, motor control, wireless infrastructure, and professional audio—and is applying this expertise to optimize ICs for automotive design.

To learn more about how Analog Devices' signal processing ICs are redefining the future of automobiles, visit *www.analog.com/auto-ad1*.

www.analog.com/auto-ad1

ANALOG DEVICES

BY BONNIE BAKER

Transimpedance-amplifier-noise issues

How much noise is too much noise in a photodiode-preamplifier circuit? You can derive the noise performance of a transimpedance amplifier (**Figure 1a**) with calculations or by using a Spice simulation (**Reference 1**). When calculating the noise performance of the circuit, consider six regions in the frequency spectrum (**Figure 1b**) and add each region with a root-sum-square equation or the following **equation (Reference 2)**:

$$V_{OUT}(NOISE_{RMS}) = \sqrt{e_1^2 + e_2^2 + e_3^2 + e_4^2 + e_5^2 + e_{RF}^2}.$$

The first five regions are equal to the multiple of the areas under the closed-loop-gain and amplifier-noise-density curves. The area under the noise-density curve in the e_1, flicker-noise (1/f), region is $V_{1/f:fB-fA}=A_N\sqrt{\ln(f_B/f_A)}$, where A_N is the amplifier's input-noise-density at 1 Hz and f_B is the corner frequency where the flicker noise tapers off. For many CMOS or FET amplifiers, the flicker-noise region usually ranges from dc to 100 or 1000 Hz. A calculation proves that the contribution to noise in this low-frequency region is relatively low:

$$e_1 = (1+R_F/R_{PD}) \times A_N \times \sqrt{\ln(f_B/f_A)},$$

where R_F is the feedback resistor and R_{PD} is the device's parallel resistance.

In the e_2 region, multiply the broadband noise of the amplifier, the closed-loop dc-noise gain $(1+R_F/R_{PD})$, and the square root of the region's bandwidth. Again, the contributed noise in this region is usually relatively low because of its location in the lower frequency range.

$$e_2 = (1+R_F/R_{PD}) \times e_N \times \sqrt{f_P - f_Z}.$$

Calculate the noise contribution and the e_3 region in the same manner with $f_P=1/[2\pi(R_{PD}||R_F)(C_{PD}+C_{CM}+C_{DIFF}+C_F+C_{RF})]$ and $f_Z=1/[2\pi(R_F)(C_F+C_{RF})]$.

$$e_3 = (1+R_F/R_{PD}) \times e_N \times (1\,Hz/f_Z) \times \sqrt{f_P/3 - f_Z/3},$$

where C_{PD} is the device's capacitance and C_{DIFF} is the differential amplifier's capacitance.

The noise in regions e_4 and e_5 uses the higher-frequency gain of the closed-loop-gain curve with the value of C_1 being the parallel combination of the input capacitors, or $[C_{P-R1}||2C_{CM}||C_{DIFF}]$, and C_2 is the parallel combination of C_F and C_{RF}.

$$e_4 = (1+C_1/C_2) \times e_N \times \sqrt{f_{AOL} - f_P}.$$

$$e_5 = (1+C_1/C_2) \times e_N \times \sqrt{\pi \times (f_U - f_{AOL})/2}.$$

The sixth part of the noise equation, e_6, represents the noise contribution of the feedback resistor. The amplifier does not gain the contribution of noise from the feedback resistor:

$$e_6 = \sqrt{4 \times K \times T \times R_F \times (BW)},$$

where K is Boltzmann's constant, which is 1.38×10^{-23}; T is temperature in Kelvin; R_F is the feedback resistor in ohms; and BW is the bandwidth of interest.

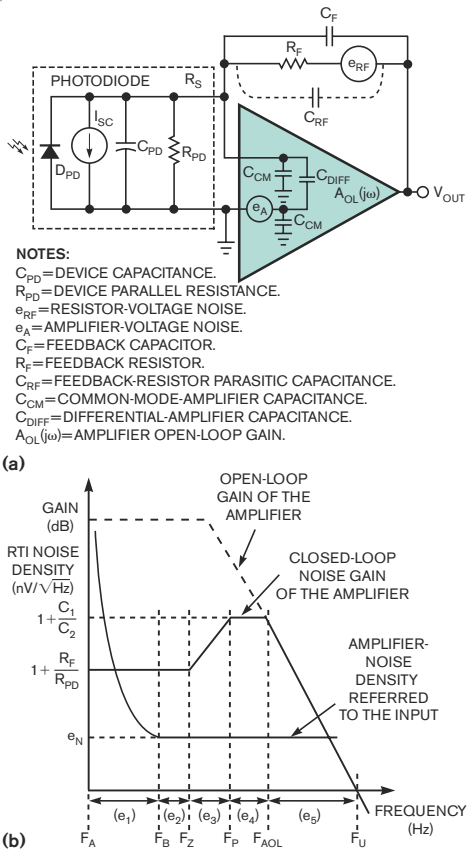

Figure 1 A typical transimpedance photo-sensing circuit (a) has five regions of overall noise response (b).

When asking how much noise is too much noise in this photodiode-preamp circuit, consider that a 12-bit system operating with a 5V input range has an LSB of 1.22 mV. The LSB for a 16-bit system with the same input-voltage range is 76.29 µV. Both LSBs are peak-to-peak numbers, and the values in this column are root-mean-square values (**Reference 3**).EDN

Bonnie Baker is a senior applications engineer at Texas Instruments. She can be reached at bonnie@ti.com.

For a list of the references cited in this column, go to www.edn.com/081002bb.

Design with the best: High Performance Analog

ANALOG DEVICES · AVAGO TECHNOLOGIES · FAIRCHILD SEMICONDUCTOR · intersil
LINEAR TECHNOLOGY · MAXIM · MICROCHIP · National Semiconductor
ON Semiconductor · TEXAS INSTRUMENTS

"On-time delivery with no minimum fees.
That's important.

Online access to all my current and past orders.
That's important, too.

The best analog brands.
That's critical."

LINEAR TECHNOLOGY LTM8021
DC/DC Switchmode µModule Regulator.
Tiny package creates space by mounting on back of board.

Design with the best brands

You'll find the best components and the best information in one site. Visit www.newark.com/instrumentation to make the most of the newest advances in high performance processing, high performance analog, power supplies, imaging, connectors, discretes and passives.

Find more at www.newark.com and 1.800.4.NEWARK

Supplier Authorized Distributor

Newark™
DESIGN WITH THE BEST

A Premier Farnell Company

PRY FURTHER AT EDN.COM

- Read an expanded version of this article at www.edn.com/081002pry.
- Go to www.edn.com/pryingeyes for past Prying Eyes write-ups.

Artificial cochlea: an example of structural processing

The cochlea is the portion of the inner ear that senses sound vibrations and converts them into electrical signals that the auditory system can interpret. The cochlea is an example of active cellular mechanical forcing and structural processing in which the shape and physical composition of the sensory organ work in combination to accomplish a complex transform. In this case, the cochlea separates a sensed, complex sound wave into its basic frequency components.

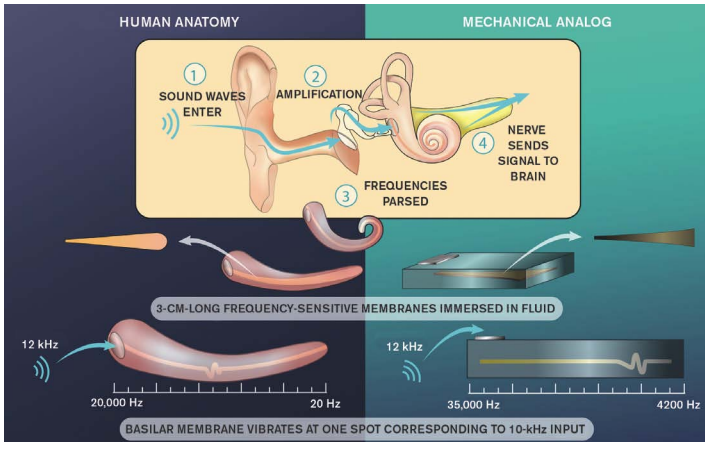

The cochlea is a curled structure, which is filled with fluid that moves in response to the vibrations coming through the oval window from the middle ear. Along the length of the cochlea are thousands of hairs that are set in motion in the liquid at the resonance points of the incoming sound wave. Each hair, based on its position within the cochlea, essentially isolates and detects a frequency band. This task is much like performing an FFT (fast Fourier transform) on a sound wave received on a conventional microphone but without performing the digital computations of the FFT (courtesy Zina Deretsky, National Science Foundation).

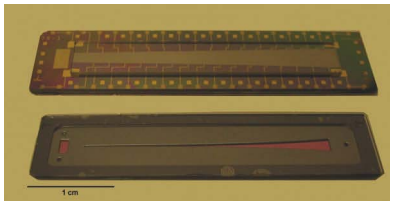

Current state-of-the-art technology for an artificial cochlea operates in a similar fashion except that, unlike the tightly curled cochlea, the MEMS (microelectromechanical-system)-based cochlea stretches out in a linear structure. The 3-cm-long device comprises an acoustic input port at the narrow end of a tapered strip. Where the strip is narrow, the sense material is stiff and vibrates in response to high-frequency compression waves in the fluid that the strip is immersed in. Additionally, as the strip widens, the material is more compliant, vibrates more easily, and absorbs the energy of lower-frequency waves (courtesy Karl Grosh, University of Michigan, and Robert White, Tufts University).

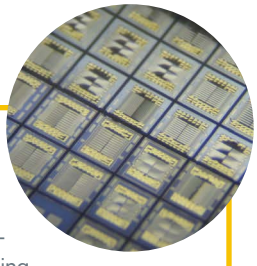

Earlier versions of the micromechanical cochlea used a capacitive-sensing approach, but the development team is working with piezoelectric sensing that takes better advantage of lithography and etching techniques. Although the structure of the cochlea passively separates a sound wave into its components' resonance points, the act of sensing the resonance points and converting them into a usable signal is an active process. The current implementation of the artificial cochlea relies on a cantilever-beam structure in which each beam is supported on only one end and the free end of the beam can sense motion of the micromechanical cochlea. This image shows sets of cantilevers that could be used for this purpose and integrated directly into the fabrication of the artificial cochlea. The pictured devices were fabricated to make many thousands of piezoelectric MEMS microphones. This technology was co-opted for the purpose of a cochlear sensor (courtesy Karl Grosh, University of Michigan).

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

There’s an indicator on the front.

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

For three measurement tips and the entire Agilent Power Products Brochure go to www.agilent.com/find/powertips

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

Agilent Technologies

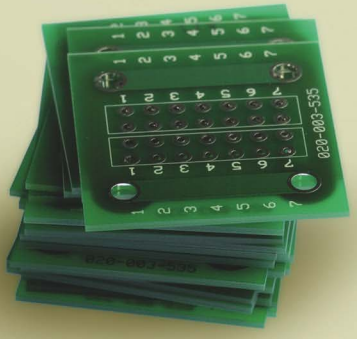

From quote to delivery, we've pioneered online PCB ordering.

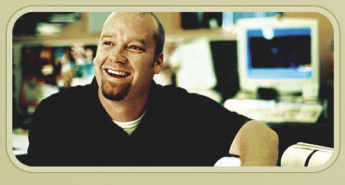

"Sunstone has done a great job with all of our orders. We rely on your easy online ordering system, quality boards, and fast lead times."

- Sunstone customer feedback

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

Visit us at
www.sunstone.com

13626 S. Freeman Road Mulino, OR 97042 | 503.829.9108

SUNSTONE CIRCUITS

HIGH-SPEED GRAPHICS, USER INTERFACES, AND NETWORKS REPRESENT THE NORM IN EMBEDDED-SYSTEM DESIGNS, AND THESE PERFORMANCE ISSUES DICTATE THE USE OF MULTITASKING FIRMWARE.

ON TIME, EVERY TIME: EMBEDDING REAL-TIME PERFORMANCE

BY WARREN WEBB • TECHNICAL EDITOR

As a new generation of computer-savvy users and gamers becomes your primary embedded-device customer, responsiveness and real-time performance become critical evaluation issues. Systems must manage priorities and tasks to provide nearly instantaneous response to users, external events, and the network. To meet these requirements, commercial vendors and the open-source community offer dozens of ready-to-run or customizable software packages ranging from minimal-resource kernels to full-featured RTOSs (real-time operating systems). Before you settle on an approach, however, it is important to analyze the inner workings of a candidate package to make the right match for your application. Scheduling algorithms, task synchronization, resource requirements, memory allocation, latencies, add-on modules, certifications, security, multicore options, and development-tool support are a few of the real-time functions and features that can make or break your design.

Real-time systems automatically execute software routines or tasks in response to external or timing events. Most embedded RTOSs are pre-emptive by design so that a task can suspend any lower-priority routines and gain control of the processor until the higher-priority task completes or until an even higher-priority task pre-empts the previous task. Unlike so-called soft real-time platforms, which simply list an average length of time to start the routine, most critical embedded operating systems must be deterministic so that they guarantee that tasks start within a precise length of time after an external event. The basic architecture of a typical multitasking RTOS for embedded devices includes a program interface, the kernel, device drivers, and optional service modules.

The kernel is the core of the real-time software and includes an interrupt processor, a scheduler, resource-sharing flags, and memory management. One of the kernel's primary functions is to handle interrupts that external or internal events cause. As interrupts take place, the processor transfers control to a service routine that logs in the event, sends a message to the scheduler, and returns to the active code. The scheduler establishes the execution order of each task based on its priority and saves execution information for each interrupt. In addition to priority scheduling, a real-time kernel also provides flags for task synchronization. For example, if several tasks want to use a data resource at once, a flag or semaphore locks the resource to a single task until the transfer is complete.

The kernel design is also important in minimizing the latencies that can degrade the performance of a real-time system. Interrupt latency is the worst-case delay between an external event, such as a switch closure, and the first instruction of the interrupt routine. If you enable the processor interrupts, the hardware delay is only nanoseconds long, but it can vary from processor to processor. The processor needs to complete only the current instruction before jumping to the requested interrupt location. If the processor has only one interrupt line, the time to poll inputs determines which interrupt routine to call as part of the interrupt latency. Another latency that the kernel contributes is the time it takes to switch between tasks. Each task has a program counter, a data-area pointer, register data, and other state information that the processor must save for the current task and restore for the pre-empting task.

AT A GLANCE

- Critical embedded operating systems must be deterministic to guarantee that tasks start within a precise length of time after an external or timing event.
- The kernel is the core of a real-time operating system and includes an interrupt processor, a scheduler, resource-sharing flags, and memory management.
- Add-on modules, such as security, safety, networking, and a GUI (graphical user interface), must work concurrently without disrupting RTOS (real-time-operating-system) performance.
- An RTOS-aware software-development-tool chain simplifies firmware debugging and real-time-performance analysis.

RTOS designers usually include a variety of optional modules and drivers outside the kernel to attract more users and to ensure that the combined package meets the advertised performance specifications. For example, almost every RTOS includes communications protocols, such as TCP/IP (Transmission Control Protocol/Internet Protocol), and most of the major RTOS vendors also offer GUI (graphical-user-interface) routines. Users can add or delete these modules depending on the application. Green Hills Software offers the Integrity RTOS in multiple configurations targeting embedded-system applications, including aerospace, automotive, medical, secure networking, wireless, and software-defined radio. These preconfigured packages integrate the most common modules and drivers into a series of off-the-shelf platforms. Product-development licenses for the Integrity RTOS start at $15,000 for a single-user enterprise license, with no royalty fees for runtime deployment.

If your application requires serious data processing or distributed processors, you should investigate an RTOS that targets use with multiple processors. You can spread tasks across several processors to gain a significant performance boost; however, all tasks must be in constant communication to maintain deterministic behavior. For example, the high-performance Enea Embedded Technology OSE (operating-system environment) targets use with high-availability, high-reliability distributed systems, such as those in telecommu-

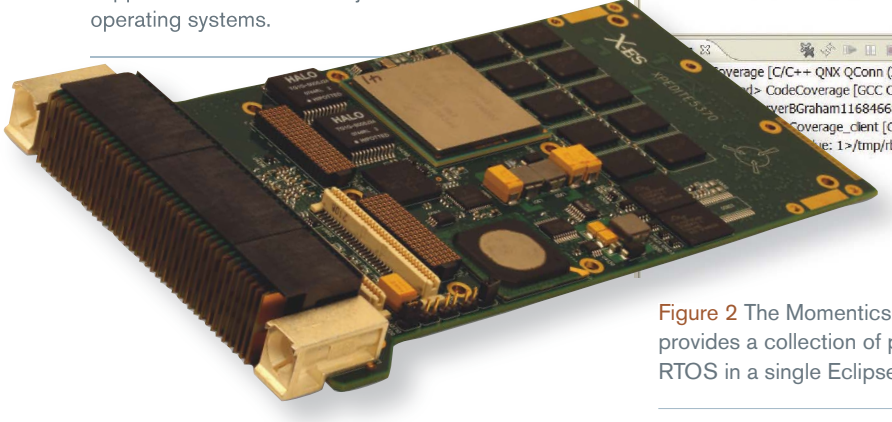

Figure 1 The XPedite5370 single-board computer from Extreme Engineering Solutions provides designers with board-support software for the major real-time operating systems.

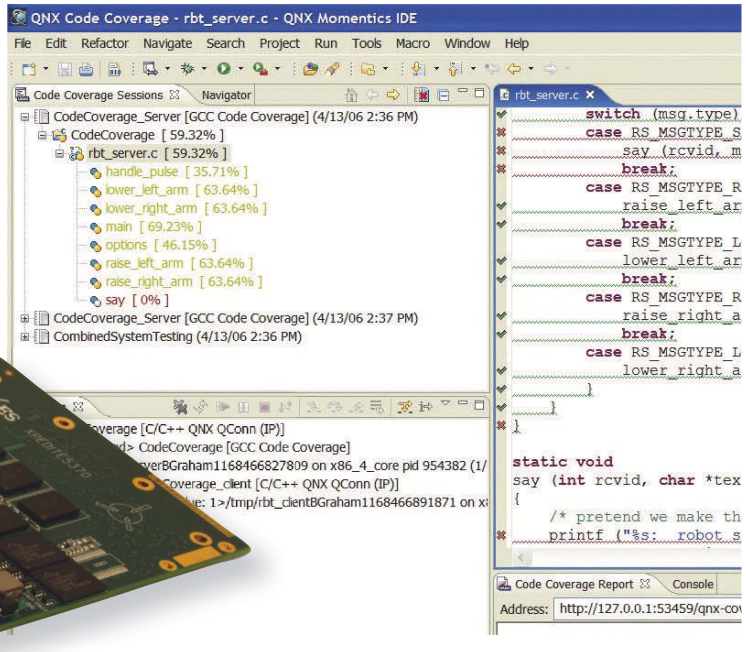

Figure 2 The Momentics development suite from QNX Software Systems provides a collection of productivity-and-analysis tools for the Neutrino RTOS in a single Eclipse-based IDE.

See the bigger picture...

with feature-rich tools for debugging mixed signal designs.

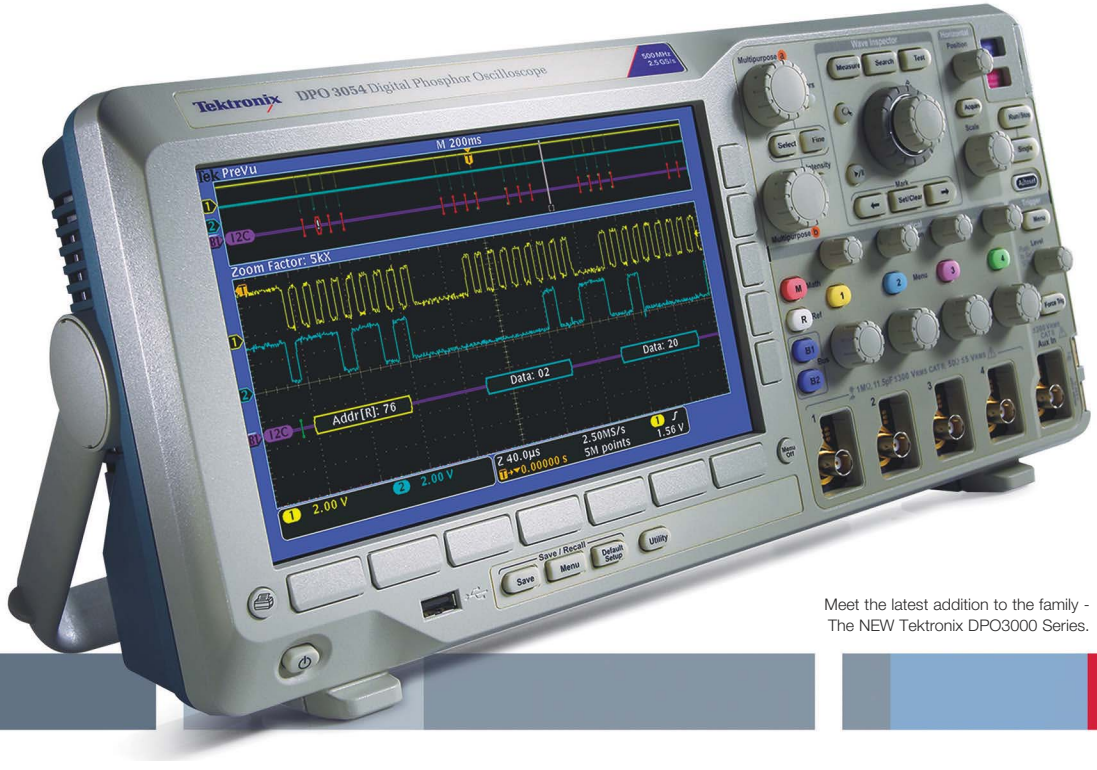

Meet the latest addition to the family -
The NEW Tektronix DPO3000 Series.

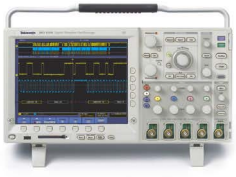

MSO4000 Series Mixed Signal Oscilloscopes Performance Specifications	
Models	MSO4032, MSO4034, MSO4054, MSO4104
Bandwidth	1 GHz, 500 MHz, 350 MHz models
Channels	2 or 4 analog and 16 digital
Record Length	10 M points on all analog and digital channels
Sample Rate	Up to 5 GS/s (analog) / Up to 16.5 GS/s (digital)
Display	10.4 in. (264 mm) XGA
Serial Bus Trigger and Decode	I²C, SPI, RS-232/422/485/UART, CAN, LIN, FlexRay

DPO4000 Series Oscilloscopes Performance Specifications	
Models	DPO4104, DPO4054, DPO4034
Bandwidth	1 GHz, 500 MHz, 350 MHz models
Channels	4 analog
Record Length	10 M points on all channels
Sample Rate	Up to 5 GS/s on all channels
Display	10.4 in. (264 mm) XGA
Serial Bus Trigger and Decode	I²C, SPI, RS-232/422/485/UART, CAN, LIN, FlexRay

NEW! DPO3000 Series Oscilloscopes Performance Specifications	
Models	DPO3054, DPO3052, DPO3034, DPO3032, DPO3014, DPO3012
Bandwidth	500 MHz, 300 MHz, 100 MHz models
Channels	2 or 4 analog
Record Length	5 M points on all channels
Sample Rate	2.5 GS/s on all channels
Display	9 in. (229 mm) WVGA widescreen
Serial Bus Trigger and Decode	I²C, SPI, RS-232/422/485/UART, CAN, LIN

See more and solve faster with Tektronix DPO and MSO Series oscilloscopes. This feature-rich family of oscilloscopes, with bandwidths up to 1 GHz, have the power and functionality to help make your work easier. With Wave Inspector® you can easily search and navigate through uncompromised long records; and large, high-resolution displays let you clearly see all on-screen data and catch glitches you might otherwise have missed. Combine this with integrated serial trigger, decode and search, and you will be debugging mixed signal designs faster and easier than ever before.

See more.
Try the virtual product demo at:
www.tektronix.com/bigpicture

Tektronix®
Enabling Innovation

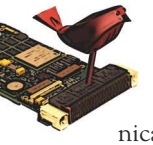

nications networks. A recently updated Linx feature adds priority messaging to maintain interprocess communications in systems that become congested due to a fault condition. Prices for the Enea OSE start at $8000 for a single developer's seat, and Linx is available for prices starting at $5000 for a multiuser license.

One way to eliminate surprises with real-time hardware integration is to base your embedded design on COTS (commercial-off-the-shelf) modules that include preconfigured support software. For example, Extreme Engineering Solutions recently introduced the XPedite5370 single-board computer featuring board-support packages for the Wind River VxWorks, QNX Neutrino, and Green Hills Integrity RTOSs (**Figure 1**). Extreme Engineering based the 3U, VITA (VMEbus International Trade Association) 46-compliant module on Freescale Semiconductor's MPC8572E PowerQUICC (quad-integrated-communications-controller) III dual-core processor, and the device supports high-speed fabric interconnections over PCIe

For more on free software for embedded design, go to www.edn.com/article/CA6582852.

Read another article about embedded Linux at www.edn.com/article/CA450620.

Go to www.edn.com/081002df and click on Feedback Loop to post a comment on this article.

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

THE SMARTER WAY TO DESIGN A SMARTPHONE

By Todd Brian, Mentor Graphics

Today, manufacturers of electronic devices face challenges that they never have before. Take the example of a mobile phone. Beyond its primary function of making a call, it must also seamlessly manage documents, such as e-mail, files, and video; have an easy-to-use GUI (graphical user interface); and be able to access and exchange data over wireless or cellular networks. All of these different yet interrelated technologies must combine to offer a homogeneous user experience. And, to satisfy demanding consumers, the latest model must reach the market within an 18-month window and at mass-market cost. How can designers balance this trade-off of features, performance, and cost?

Enter the operating-system software. Whether inside a sleek mobile phone or another multimedia device, the operating system must retain the characteristics of an embedded operating system: efficiency, scalability, determinism, and speed. Unlike the RTOS (real-time operating system) of yesterday, however, today's embedded operating-system platform must provide a complete foundation for cost-efficient innovation at the device level. Look under the hood of any mobile phone, and you will find that it relies on a set of basic technologies, such as connectivity, file storage, graphics, and all the drivers for the peripheral hardware. Those functions require a lot of software. From a device manufacturer's standpoint, this software is costly to develop and maintain in-house, yet it does not differentiate a device in the market.

A better approach is to start with a full-featured operating system containing the latest in networking and connectivity, file-storage, and USB (Universal Serial Bus) capabilities. This operating-system platform must also include the latest technology in graphics/multimedia and support industry APIs (application-programming interfaces), such as the OpenMax specification from the Khronos Group. For devices that do not comply with OpenMax, the multimedia framework can still exploit codec support and provide a common interface to the application layer. OpenMax also delivers hardware integration of other graphical and multimedia capabilities and XML (Extensible Markup Language)-driven menus. Furthermore, a cost-efficient RTOS platform must enable application-software development, whether by in-house teams or by independent software providers. Integral to the platform is a robust development platform, including an IDE (integrated development environment), a compiler, a debugger, and a profiler and simulation environment. Chip-set vendors or device manufacturers can extend and package this tool suite as a device-specific SDK (software-development kit). The next-generation operating system must work with various embedded applications. Much like a custom-tailored suit, every line of code on the device is there explicitly due to the design of the device, and manufacturers can easily extend and customize those lines so that the RTOS environment is an exact fit.

Average consumers may be relatively unaware of the operating system inside their mobile phones. Obscure as it may be, however, it is this next-generation RTOS platform that enables better economics of mobile phones and other electronic devices and enables device manufacturers to quickly deliver the latest features to their customers.

AUTHOR'S BIOGRAPHY

Todd Brian is product-marketing manager at the embedded-systems division of Mentor Graphics. You can reach him at todd_brian@mentor.com.

The power of the industry's broadest product portfolio.

ON Semiconductor now has the industry's broadest selection of power efficient solutions—your single source for everything from ASICs and ASSPs to standard ICs and discrete components.

When it comes to power efficient solutions, ON Semiconductor is the one name to know. In addition to providing a wide variety of discrete components and standard integrated circuits, ON Semiconductor now also offers specialized solutions like custom ASICs, ASSPs, and power efficient GreenPoint™ reference designs. No matter what your application is—ranging from medical to automotive and industrial to power supplies, and everything in between—ON Semiconductor is the only resource you can be sure has a solution that's complete, effective, and efficient. Now that's powerful.

Find your next power solution at www.onsemi.com/thepower

ON Semiconductor®

(peripheral-component-interconnect express), Serial Rapid I/O, and GbE (gigabit Ethernet). Prices for the XPedite5370 start at less than $8000, and volume discounts are available.

Certifications are another way to ensure that an RTOS meets some or all of your performance requirements. For example, Wind River Systems claims that its VxWorks operating system was the first to receive certification for conformance to the POSIX (portable-operating-system-interface) PSE52 real-time-controller 1003.13-2003 product standard. POSIX comprises a family of related IEEE standards that define the API (application-programming interface) and critical performance areas. In addition to ensuring source-code portability, PSE52 certification confirms that the RTOS delivers the predictable response times that time-critical applications require.

The development-tool chain is another big issue in the selection of an RTOS. You will spend most of your software-development and -debugging efforts interacting with the IDE (integrated development environment) to gain quick access to the editor, compiler, linker, downloader, and runtime tools. The open-source Eclipse IDE has become a popular and easy-to-use standard interface for embedded-software-tool vendors. The platform allows users to easily create custom Eclipse configurations because everything other than a small runtime kernel is a plug-in. For example, the Momentics tool suite from QNX Software Systems provides a set of productivity-and-analysis tools for the Neutrino RTOS, all integrated into a single Eclipse-based IDE (**Figure 2**).

If you decide on a commercial-software package, you must also determine whether you need or want to purchase the source code for the vendor's RTOS. Some vendors automatically supply the source code when you purchase their software, whereas others charge extra. With the source code, you have the option to tweak the vendor's software to remove every line of unused code and reclaim vital memory space. The source code can also help you understand those subtle bugs in your application code. However, vendors that supply only object code claim that, if you change the source code, you have created a unique and unsupportable RTOS. You can always purchase a copy of the source code to document your system or to deliver to your customer.

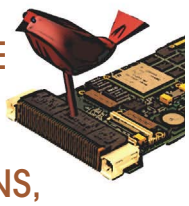

REAL-TIME SYSTEMS ARE CRITICAL IN MANY APPLICATIONS, AND A MISSED DEADLINE CAN RESULT IN CATASTROPHIC LOSS OF LIFE OR PROPERTY.

Embedded devices come in all sizes and capabilities, but one thing is for sure: As the software complexity grows, the need for real-time, deterministic performance becomes vital (see **sidebar** "The smarter way to design a smartphone"). In fact, real-time systems are critical in many applications, and a missed deadline can result in catastrophic loss of life or property. As an embedded-system designer, you can expect your customers to ask for higher performance and increased complexity as the technology bubble expands. Fortunately, a huge group of software vendors and open-source volunteers are already working on the next generation of real-time firmware to simplify your next embedded-software application. EDN

FOR MORE INFORMATION

Enea Embedded Technology
www.enea.com

Extreme Engineering Solutions
www.xes-inc.com

Freescale Semiconductor
www.freescale.com

Green Hills Software
www.ghs.com

Khronos Group
www.khronos.org

Mentor Graphics
www.mentor.com

QNX Software Systems
www.qnx.com

VITA (VMEbus International Trade Association)
www.vita.com

Wind River Systems
www.wrs.com

Highest Speed, µPower ULDO Delivers 48dB (1MHz) PSRR!

800mA Load Current & Superior Transient Response

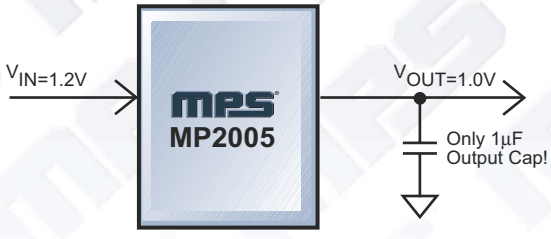

Superior Load Response

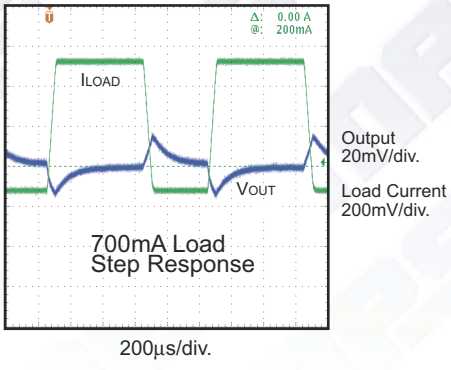

- *High PSRR: 65dB (1kHz), 48dB (1MHz)*
- *Stable with 1µF output capacitor*
- *100µA Supply Current*
- *Ultra-low dropout voltage: <90mV*
- *Wide input voltage range: 1.0V to 5.5V*
- *Adjustable VOUT range: 0.5V to 4.0V*

PSRR vs. Frequency

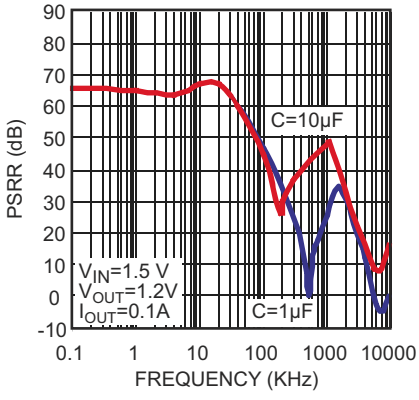

MPS Precision LDO Family

Part	Dropout	Max IOUT	VIN Range	Min VOUT	Adj or Fixed	Package
MP2000	250mV (100mA)	150mA	1.35V - 6.0V	0.5V	Adj	TSOT23-5
MP8801	150mV (150mA)	150mA	2.7V - 6.5V	1.25V	Adj/Fixed	TSOT23-5
MP8802	230mV (250mA)	250mA	2.7V - 6.5V	1.25V	Adj/Fixed	TSOT23-5
MP2002	290mV (500mA)	500mA	1.35V - 6.5V	0.5V	Adj	QFN8 (2x3)
MP89046	360mV (600mA)	600mA	2.5V - 6.5V	0.5V	Adj	QFN8 (2x3)
NEW MP2005	<90mV (800mA)	800mA	1.0V - 5.5V	0.5V	Adj	QFN8 (2x3)

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

ISO9001:2001 Certified

LEAD FREE RoHS COMPLIANT

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

AVNET Memec

VIRTUALIZATION:

BY BRIAN DIPERT • SENIOR TECHNICAL EDITOR

STABLE, ROBUST CODE SPEAKS ONE LANGUAGE; NEW CPUs SPEAK ANOTHER. IS A SOFTWARE REWRITE NECESSARY TO RESOLVE THE SEEMING CONTRADICTION, OR CAN VIRTUALIZATION TEMPORARILY—OR EVEN PERMANENTLY—EASE THE TRANSLATION?

EDN HANDS-ON PROJECT

SILICON AND SOFTWARE
SALVATION OR TECHNOLOGICAL TOWER OF BABEL?

Imagine that you've amassed a large library of mature software for a CPU architecture or a system containing it but that the silicon manufacturer abruptly goes out of business, and no viable second source exists. Alternatively, imagine that you want to move your next-generation platform design from one microprocessor or embedded-controller family to another for performance, power-consumption, price, or other reasons but that you lack the schedule, manpower, tools, and education budget to port your code base to the new chip. How can you keep your development plans moving forward?

Hardware virtualization provides a possible approach. It's equally attractive to silicon suppliers hungry for new applications that will exploit Moore's Law-fueled incremental-IC capabilities. You can potentially apply the concept to any design because all processor architectures will sooner or later exhibit the same sorts of clock-speed boosts, on-chip-core-count increases, and dedicated-function-hardware advancements that are most evident today in the PC market. The bulk of the industry's current attention on virtualization, however, focuses on computers, thereby making such virtualization equally appropriate for x86-CPU-based embedded systems. This focus reflects the huge market that exists for laptop, desktop, workstation, and server-hardware applications compared with other applications.

Virtualization for Apple's Macintosh computer lines predated the company's shift from PowerPC to x86 microprocessors. However, the concept became notably more attractive after the CPU transition for several key reasons (**Reference 1**). For one thing, the resultant underlying instruction-set compatibility between the host operating system and the emulated software, such as Linux or Windows, along with the commonality of other system-building blocks, such as core-logic chip sets, graphics processors, and mass-storage devices, meant that virtualization was speedier and functions were more robust than they were previously. The resultant performance and power-consumption improvements, along with price reductions, also made Apple hardware more attractive than before to consumers who might potentially switch to the Macintosh but still need to run a few Windows applications.

After enduring a glitch-filled, nearly two-year-long attempt at natively running Windows XP Professional on a first-generation MacBook using Apple's Boot Camp partitioning software and driver suite, I decided a few months ago to instead try using Windows in a virtualized fashion (**Figure 1**). I'd previously experimented with Parallels' Desktop for Mac, the initial product in the virtualization-on-OS-X category, but I'm now using Fusion from long-time virtualization pioneer VMware, which is now an independent subsidiary of EMC. Unlike Parallels' Workstation, which, at press time, could access only one CPU core, VMware Fusion Version 1 can tap into two cores' worth of resources.

It's nice to know I've got the added processing muscle available if I need it, although to date I largely haven't. Fusion touts not only a robust array of features, but also surprising speed—even on one CPU core. It also uses less battery power in this configuration than it does in the two-core alternative approach. As such, I've qualified and quantified my impressions in the hope that they'll be of interest as you evaluate virtualization for your upcoming designs as well as for your potential personal use. The specifics of various virtualization approaches differ, so there's no guarantee that your results will mirror mine, especially if you use a dissimilar host operating system or alternative host and virtu-

alized-CPU architectures. Nonetheless, after perusing the pages to come, I think you'll agree with me that virtualization holds tremendous promise for bridging the hardware-versus-software divide.

TEST-BED STATS

The hardware heart of my Apple MacBook is a dual-core, 2-GHz, Yonah-generation CPU. Although Intel refers to it as a Core microprocessor, this marketing moniker doesn't make it a Core microarchitecture product; the company based it on the earlier Pentium M microarchitecture. As such, it's in effect two single-core Pentium M CPUs, each with a dedicated 1-Mbyte L2 cache, on one die, communicating with each other as well as with the remainder of the system over the shared front-side bus. Contrast this arrangement with the follow-on 65-nm Merom and 45-nm Penryn CPUs, in which all on-die cores share a unified L2 cache, and with the upcoming 45-nm Nehalem products, which have core-specific L2 caches but a common L3 cache. My Yonah processor, a T2500, implements Intel Virtualization Technology's hardware hooks, which are useful, albeit not completely necessary, as VMware's industry presence many years before Intel unveiled VT attests.

The system-hardware suite that Windows XP and my benchmarking program of recurring choice, SiSoftware's Sandra (system analyzer, diagnostic, and reporting assistant) Lite XII.SP2c, report when operating virtually under VMware Fusion, differs in several key areas from the Boot Camp-derived native system's building-block counterparts (**Table 1**). For one thing, the benchmarking program incorrectly reports that the virtualized CPU's front-side bus is running at 4 MHz because the virtualized-core-logic chip set is Intel's archaic 440BX. This chip set interfaces to EDO (extended-data-out) asynchronous DRAM. Further, the virtualized graphics and audio subsystems have fewer and less robust features than their real-life counterparts. What's more, the reported amount of onboard RAM cache associated with the hard-disk and optical drives, when operating virtualized, is substantially less than the actual amount of RAM cache. In addition, VMware Fusion Version 1 doesn't support IEEE-1394 virtualization, although it does implement robust USB 2 support. Finally, Fusion leverages NAT (network-address translation), bridged, and host-only pairings to the host operating system's LAN, WAN, and Bluetooth PAN (personal-area-network) connections, so its sole virtualized-networking subsystem is a 1-GbE (gigabit-Ethernet) transceiver.

The virtualized system's DRAM allocation begs for additional explanation. The MacBook supports only as much as 2 Gbytes of main memory; note, too, that Yonah CPUs aren't 64-bit-capable. VMware Fusion recommended that I allocate 512 Mbytes of memory to the virtualized OS; I overrode the default settings and assigned 1 Gbyte to the virtual machine to minimize paging-induced Windows-performance degradation. However, if I were simultaneously running memory-intensive applications on the OS X host operating system, I might not have bumped up the virtual-memory allotment. In that case, I also probably would have selected the "optimize-for-Mac-OS-application-performance" option instead of the "optimize-for-virtual-machine-disk-performance" option.

AT A GLANCE

- Virtualization promises to cross the chasm between seemingly incompatible code and CPUs.
- Minor hiccups can't disguise the notable functional compatibility accomplishment that VMware Fusion delivers with virtualized Windows XP on OS X.
- CPU, memory, and networking speed suffer little from the virtual-machine-manager intermediary.
- A constantly and audibly running system fan implies higher power consumption.
- Virtualization approaches span a spectrum of options.

TABLE 1 SYSTEM-BUILDING BLOCKS

	BOOT CAMP	VMWARE FUSION V1
CPU	Intel Core Duo T2500 (Yonah, 2-GHz core, 667-MHz front-side bus, dual-core, 2x1-Mbyte L2 cache)	Intel Core Duo T2500 (Yonah, 2-GHz core, 4-MHz front-side bus, dual-core, 2x1 Mbyte L2 cache)
Core-logic chip set	Intel Mobile 945 Express (Napa)	Intel 440BX
DRAM	2-Gbyte DDR2-667 (Kingston Technology)	1-Gbyte EDO DRAM
Graphics processor	Intel GMA 950	VMware SVGA II
Hard-disk drive	Seagate ST9160821AS (160-Gbyte total with 60 Gbytes allocated to the NTFS partition, 5400 rpm, SATA/150 with native-command queing, 8-Mbyte cache)	VMware virtual-IDE hard drive (20 Gbytes, ATA33, 32-kbyte cache)
Optical drive	Matsushita UJ-857 (ATAPI66, 2-Mbyte cache)	Matsushita UJ-857 (ATAPI33, 32-kbyte cache)
Audio	Sigmatel 82801G high-definition audio	Creative Sound Blaster AudioPCI 64V, AudioPCI 128
USB	Intel 82801G USB universal host controller	VMware 82371AB/EB/MB PIIX4/E/M USB controller, VMware Abstract USB2 enhanced-host-controller-interface controller
IEEE-1394 (FireWire)	Agere Systems FW322/323 IEEE-1394 OHCI FireWire controller	NA
Ethernet	Marvell Yukon 88E8053 PCIe GbE controller	AMD PCnet family PCI GbE adapter
Wi-Fi	Broadcom 802.11n network adapter	NA

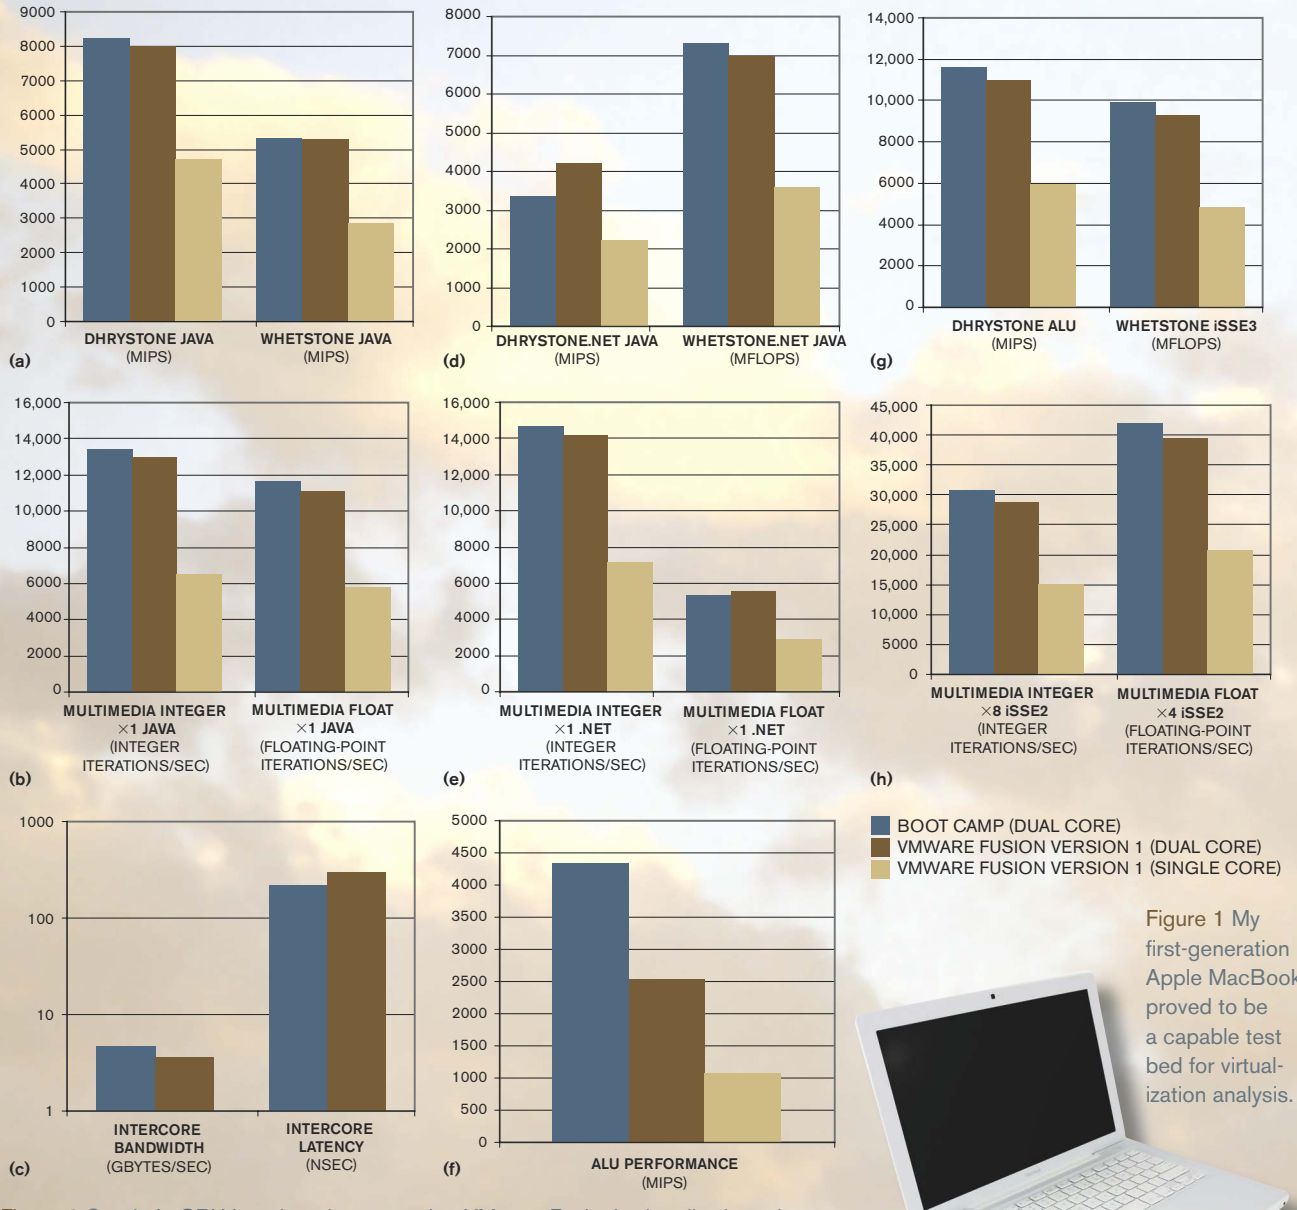

Figure 2 Sandra's CPU benchmarks attested to VMware Fusion's virtualization talents: Java Arithmetic (a), Java Multimedia (b), Multicore Efficiency (c), .NET Arithmetic (d), .NET Multimedia (e), Power-Management Efficiency with ALU power/performance of 30 Hz (f), Processor Arithmetic (g), and Processor Multimedia (h).

Figure 1 My first-generation Apple MacBook proved to be a capable test bed for virtualization analysis.

I went along with the 20-Gbyte-maximum virtual-hard-disk-drive size that Fusion recommended because, at the time, I had limited free space available on the 160-Gbyte hard-disk drive. I was happily surprised to see that this virtual-partition size is plenty for Windows XP, Office 2000, and miscellaneous other programs, though I'll need to store bit-intensive data files, such as still images, video clips, and music tracks, elsewhere. And, if I ever want to beef up this peak-capacity setting, I'll be unable to directly do so; VMware officials say I'll instead need to create a new virtual partition of the desired dimensions and then mirror my Windows image to it.

USAGE IMPRESSIONS

In migrating to virtualized Windows XP from a natively running predecessor, I expected to experience a substantial decrease in perceived performance, along with numerous functional hiccups. Happily, neither forecast came to pass, although the virtualization intermediary somewhat impacted processing-intensive applications and battery life. I should also point out that, so far at least, I'm not running any 3-D-graphics-based applications, and I therefore haven't yet enabled Fusion Version 1's experimental Direct3D "v9" feature. Fusion Version 2, which was still in beta testing at press time, touts improved graphics-API (application-programming-interface)-virtualization capabilities, hardware-accelerated video-decoding support, and a more general decrease in the virtual-machine manager's consumption of CPU and other system resources.

Every USB peripheral I've attempted to use with virtualized Windows XP has worked without a hitch. For example, I synchronized several Microsoft porta-

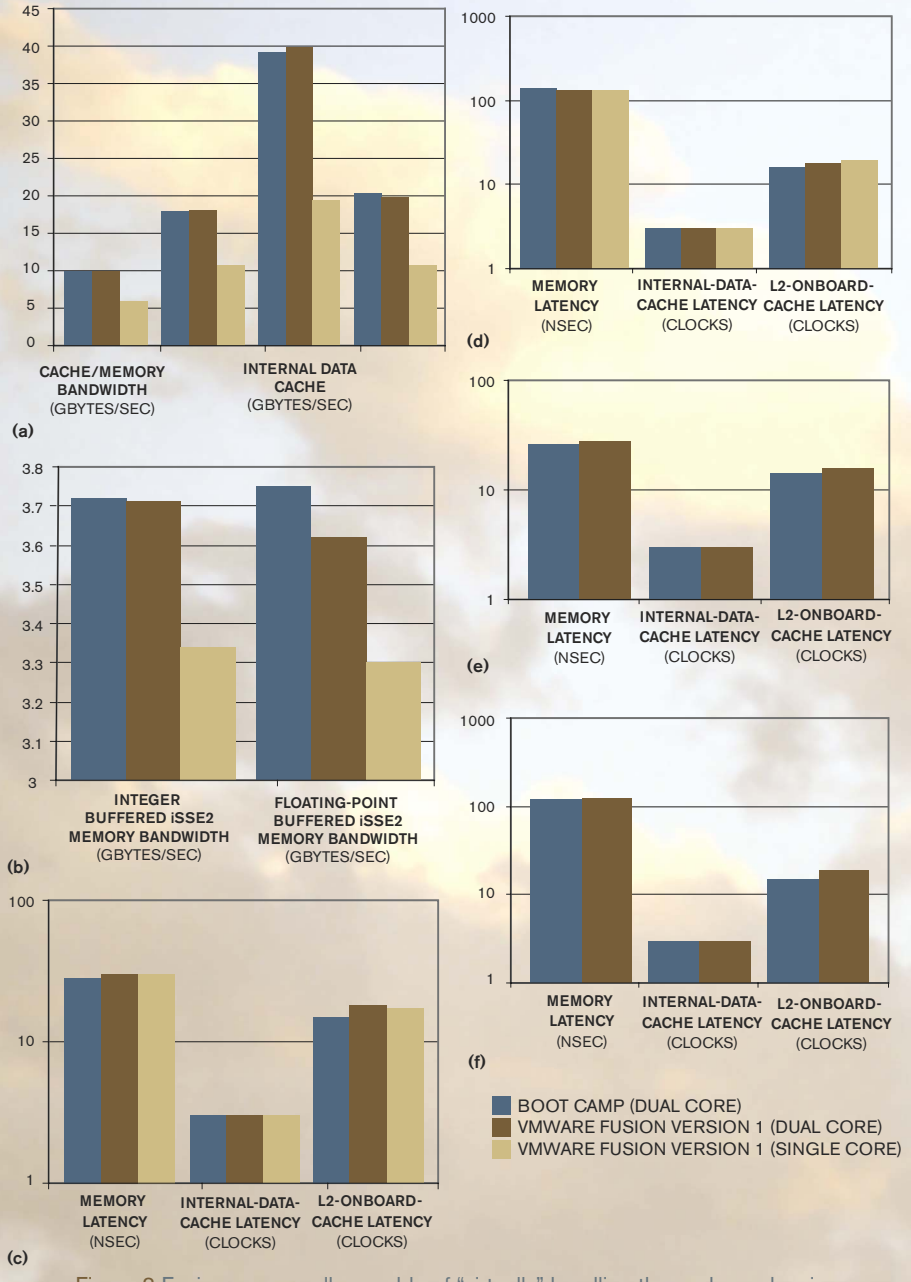

Figure 3 Fusion was equally capable of "virtually" handling the cache and main memory at nearly real-time speeds: Cache and Memory (a), Memory Bandwidth (b), Linear-Memory Latency for CPU 1 (c), Random-Memory Latency for CPU 1 (d), Linear-Memory Latency for CPU 2 (e), and Random-Memory Latency for CPU 2 (f).

ble music players and a T-Mobile Dash Windows Mobile Smartphone using Microsoft's Zune software and ActiveSync utility, respectively. (Note that these tasks sometimes don't work even with *real* USB transceivers!) I established network connections in OS X using Category 5-cable-wired, Wi-Fi wireless, and a Bluetooth wireless-PAN tether to my acting-as-a-cellular-modem cell phone. After I made any of these connections, Fusion's virtual-network adapter consistently tapped into it, too. Perhaps the most significant stumble I've encountered so far in my Fusion experimentation is that I couldn't "see" other network resources assigned to my Windows work group until I switched the virtualized-networking mode from its default NAT setting to the bridged alternative. Getting the MacBook's built-in webcam working within Windows wasn't intuitive, either, though I eventually succeeded by extracting the necessary drivers from the Boot Camp Version 2 suite on an OS 10.5 CD. Windows-resident Bluetooth and IR (infrared) support currently requires jumping through similar hoops; VMware promises more straightforward support for all three peripherals in Fusion Version 2.

Installing Windows XP Professional under Fusion was a breeze, courtesy of the bundled virtual-machine-assistant utility. The utility first polled my system and recommended a setting for the maximum-virtual-hard-drive size. It also allowed me the option of entering the Windows product key, along with my desired user-account login and password. It then prompted me to insert the Windows-installation CD. Fusion took over from there, including installing drivers for VMware-virtualized subsystems along with the VMware Tools add-in. Because virtualized Windows is the primary OS that my machine uses, I've tweaked some of the OS X-keyboard settings to make them more Microsoft-like, and I've also installed an open-source program called AutoHotKey to give me a dedicated delete key.

One of the many headaches I experienced with Boot Camp was its unreliable power management; transitions into and out of standby mode weren't rock-solid—sometimes with disastrous consequences. Conversely, Fusion and the virtualized Windows running under it act like any other Mac OS X application, thereby enabling leverage of Apple's ironclad power management. To wit, while playing music in Windows, I shut the system bezel, putting the MacBook in standby mode, and the tune continued without a hitch when I subsequently reawakened the system. System crashes inevitably occur, however. To protect myself from their aftereffects, I could take a "snapshot" of the virtual machine at any time for subsequent restoration on an as-needed basis. And backing up the entire virtual-machine image or, for that matter, migrating it to another VMware-inclusive system, was as simple as a one-file copy.

I could copy and paste text and other information between OS X and Windows using either operating system's "clipboard" because Fusion links them, and I could similarly swap files either by dragging and dropping them between OS desktops or through the shared-folders feature, which neatly translates between HFS+ (hierarchical file system plus) and NTFS (New Technology File System). Virtual-machine-display options include windowed; full-screen; and

Cadence OrCAD Capture CIS with Digi-Key Integration

Design with confidence in your part selection

Reliable

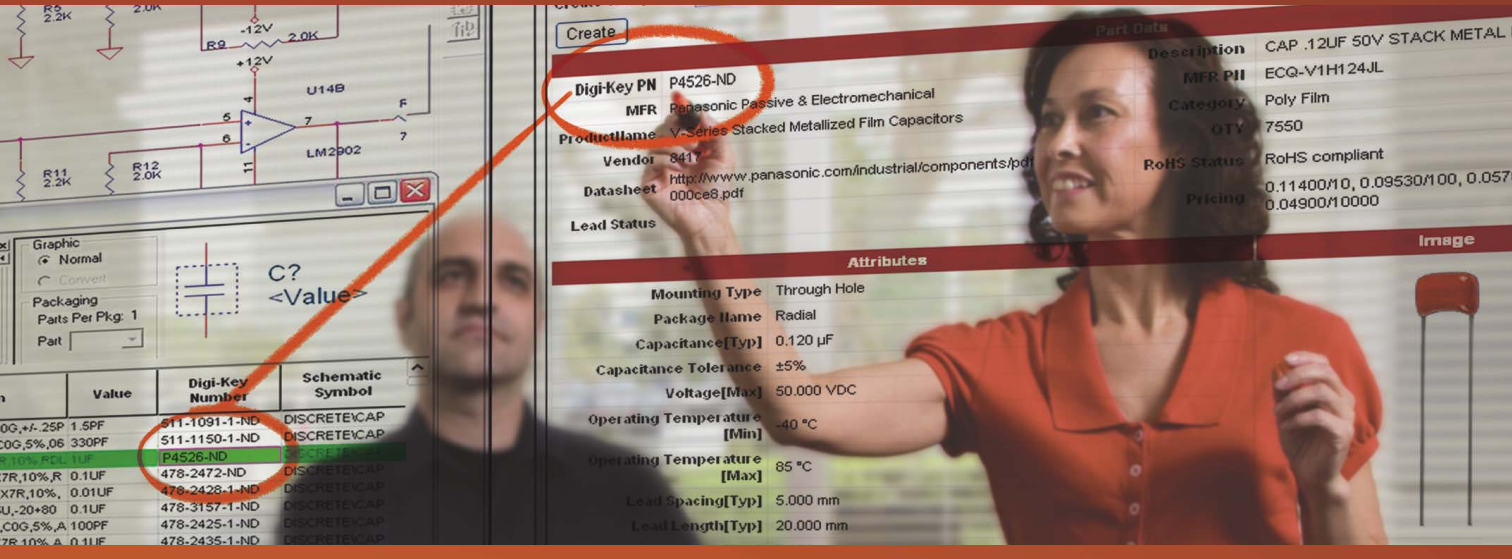

PCB

SI

Custom IC

Analog

Digital

Libraries

Mechanical

PLM

Training

An enterprise solution for capturing schematic content

Cadence® OrCAD® Capture Component Information System (CIS) with Digi-Key® integration provides a wealth of content at the engineer's fingertips. With Digi-Key, an OrCAD Capture CIS user has access to over 1,000,000 parts in the Digi-Key database with device parameters, RoHS compliance status, cost, quantity on hand, and mechanical dimensions.

Streamline the part introduction process

Often times a temporary part has to go through an extensive internal review process by purchasing, manufacturing, and documentation groups before it can be approved for use. Every new part introduced requires significant administrative overhead to verify and add to the system. With the Cadence OrCAD Capture CIS and Digi-Key integration, you can reduce administrative overhead because your part information is accurate, shows current part availability, and has correct part numbers.

The power behind Cadence OrCAD Capture CIS with Digi-Key integration

The Component Information Portal (CIP) is a web based interface that accesses the parts database behind OrCAD Capture CIS. CIP provides engineers and administrators with a more effective way to update and access the CIS database. With enterprise integration, the CIS database will hold data from external systems like ERP, MRP, PLM, and PDM.

Experience the capabilities of the integration with an on-demand demo!

Witness how you can access part information for over one million components from more than 380 manufacturers by integrating OrCAD Capture CIS with Digi-Key. Visit EMA, a Cadence Channel Partner, online at www.ema-eda.com/Digi-Key, or call us at 800.813.7288

Digi-Key .com

cadence™ CHANNEL PARTNER

EMA Design Automation™ ema-eda.com

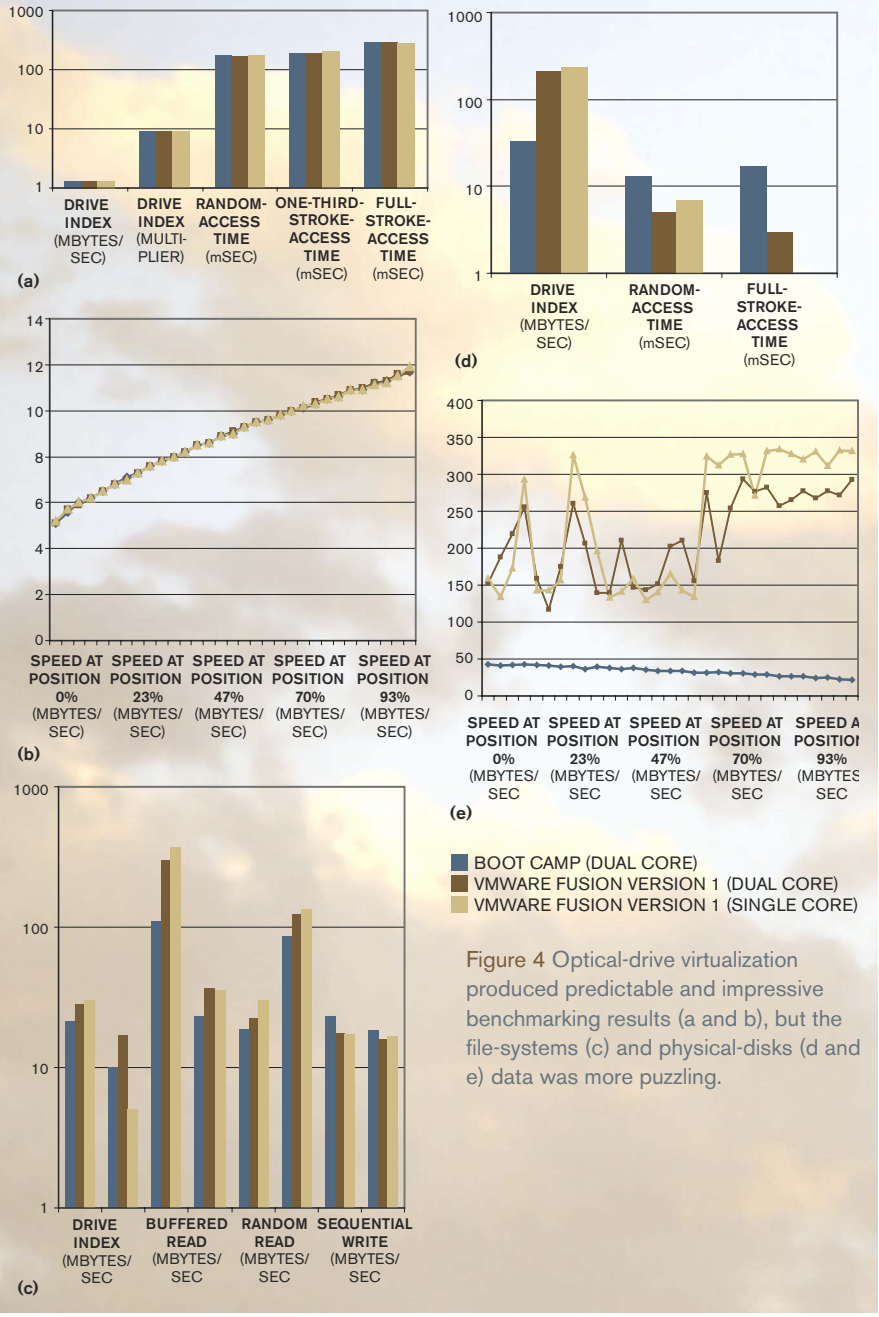

Figure 4 Optical-drive virtualization produced predictable and impressive benchmarking results (a and b), but the file-systems (c) and physical-disks (d and e) data was more puzzling.

Unity, which merges the Windows desktop with OS X. With Fusion Version 1, I could neither, for example, assign all HTTP (HyperText Transfer Protocol) links to go to Firefox in the host OS X nor assign all mail-to links to go to Outlook in the virtualized Windows. Parallels Desktop does have those features, and VMware plans to include them, along with inter-OS file-to-application assignments, in Fusion Version 2.

To quantify my earlier comment about Fusion's surprising speed even on one CPU core, I ran Sandra's various microprocessor-centric benchmarks on Windows XP native with Boot Camp and with both CPU cores enabled—that is, I didn't implement the /onecpu flag in BOOT.INI. I also ran the benchmarks on Windows XP virtualized with Fusion, running on dual-core-enabled—that is, I didn't disable a core using the CHUD (computer-hardware-understanding-development) kit—OS X, and with Fusion's two-virtual-processor setting enabled. I then ran the benchmarks on virtualized Windows XP using Fusion, again running on dual-core-enabled OS X but this time with the one-virtual-processor default setting enabled.

The results clearly showcase the tangible performance benefit of microprocessor-instruction-set compatibility between the host machine and the virtualized operating system (**Figure 2**). With some of the more processing-intensive benchmarks, dual-core virtualized Windows XP is noticeably, albeit not significantly, slower than its native counterpart. This result reflects the incremental CPU usage that the OS X-based virtual-machine manager incurred, and single-core virtualized Windows XP was proportionally slower still. Other benchmarks, which are more systemic in nature, reveal near-parity between the virtualized and the native dual-core configurations. And, in a few cases, virtualized Windows XP even came out ahead of the native alternative.

Next, take a look at how well Fusion handled virtualization of cache and main memory (**Figure 3**). Only a perpetual pessimist would fail to be impressed with the virtualized-versus-native benchmark numbers that Sandra delivered. In both these and the earlier tests, SiSoftware's utility also output scaled performance-versus-clock-rate and performance-versus-power-consumption data, but **Figure 3** doesn't show it because the utility is unreliable in a virtualized configuration. Recall that the virtualized CPU and core logic connect through a 4-GHz front-side bus and that the virtualized DRAM is of the ancient, asynchronous-EDO flavor. All of these factors greatly distort the virtualized per-megahertz- and per-watt-performance results. "The TPD [thermal-power dissipation] is estimated, not calculated, except for CPUs that report CID [CPU identification] as well as VID [voltage identification]; AMD's Phenom/Barcelona [is] the only one," says C Adrian Silasi, chief technology officer at SiSoftware. "Otherwise, TPD is based on processor model/type. It is a database look-up based on published Intel specifications adjusted for reported frequency and voltage: Power is approximately equal to the frequency times the voltage squared."

The MacBook's mass-storage subsystems beg for Sandra inspection, too, and the results in this case can at first glance be more baffling (**Figure 4**). The optical-drive-read tests are straightforward enough and attest to the robustness of Fusion's virtualization of this peripheral. However, look at how much faster in many cases the virtualized hard-disk drive is than its native counterpart. The physical-disks test bypassed operating-system-specific API calls that the file-

Optical Sensing Made Easy

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

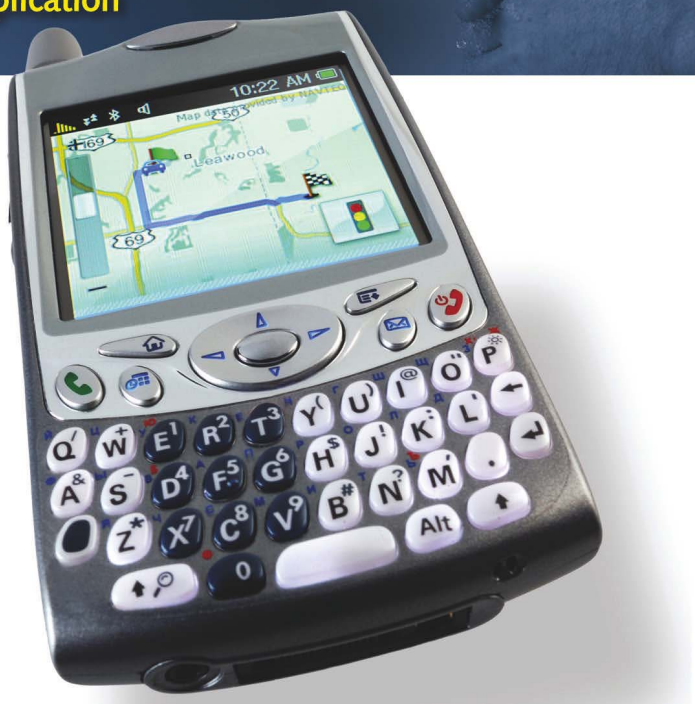

Features and Benefits:

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

Applications:

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

Integration Makes Life Easier

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

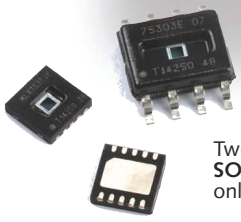

The Melexis Sensor EyeC™ Family uses unique and patented open-cavity package technology which combines high robustness (including solder reflow of 260°C) with an unobstructed optical path.

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

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

Visit Melexis Booth 1419
CONVERGENCE 2008
October 20-22, Cobo Hall, Detroit

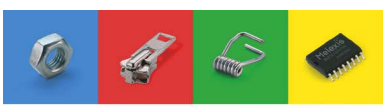

Small things make a big difference.

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

Melexis
Microelectronic Integrated Systems

systems test employed. The physical-disks test was therefore communicating directly with the physical or virtualized hardware. I had also configured the file-systems test to circumvent Windows buffering.

However, Sandra can sidestep only the system-memory-based caching schemes that it's aware of, and, as was the case with my past storage tests, I suspect that it was unsuccessful in this case (**Reference 2**). "In both benchmarks, a suitably large amount of data is read and written to flush or flood any caches that cannot be disabled," says SiSoftware's Silasi. "However, the tests disable only software caches and not hardware caches, [such as] RAID-controller caches or the hard-disk caches themselves. Most likely, VMware caches disk reads and writes in main memory, and, being 'hardware,' these caches are not disabled while being as fast as normal software caches running in system memory." Note that

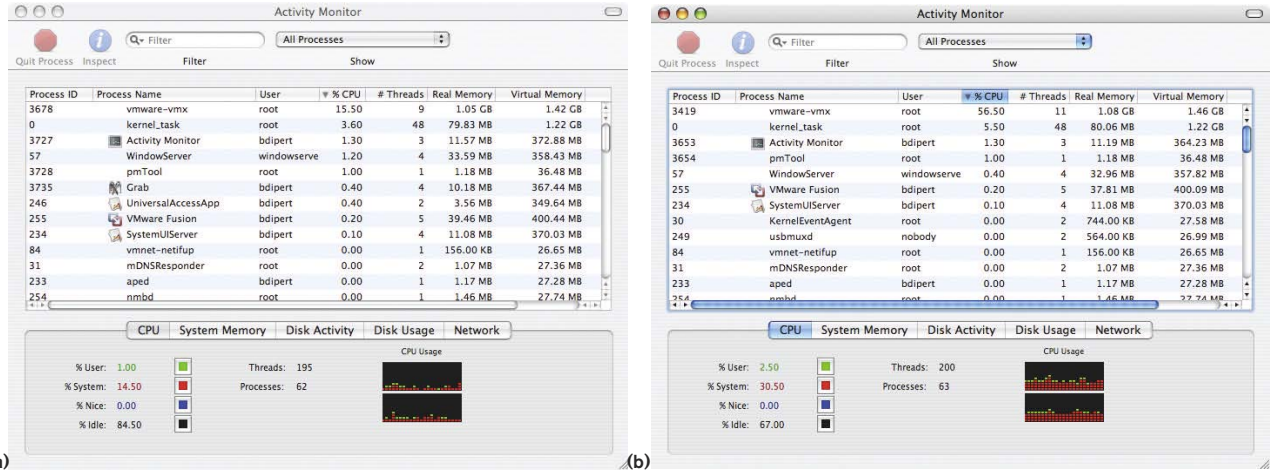

Figure 5 The virtual-machine manager consumes perceptible CPU resources in both single-core (a) and dual-core (b) virtualization configurations, even with the virtualized OS at idle.

FIRST-PASS SYSTEM SUCCESS

First-Pass System Success, [def., Engineering] The accurate prediction of the behavior of a high-performance electronic product design before committing to hardware.

At Ansoft, we are committed to delivering design and analysis software to achieve First-Pass System Success. We deliver the power you need to solve the fundamental physics accurately and quickly and to use these correct models at the circuit and system levels to precisely predict the performance of your design. By leveraging our industry leading electromagnetic field simulators dynamically linked to our next-generation circuit and system tools, Ansoft has taken high-performance electronic design to a new level of accuracy, speed and precision. We deliver the capabilities you need to balance high-speed, low power, high density and process variability to get your product to market faster, at lower-cost and with higher performance.

Contact us, we will be happy to show you how to achieve **First-Pass System Success**.

ANSOFT.COM

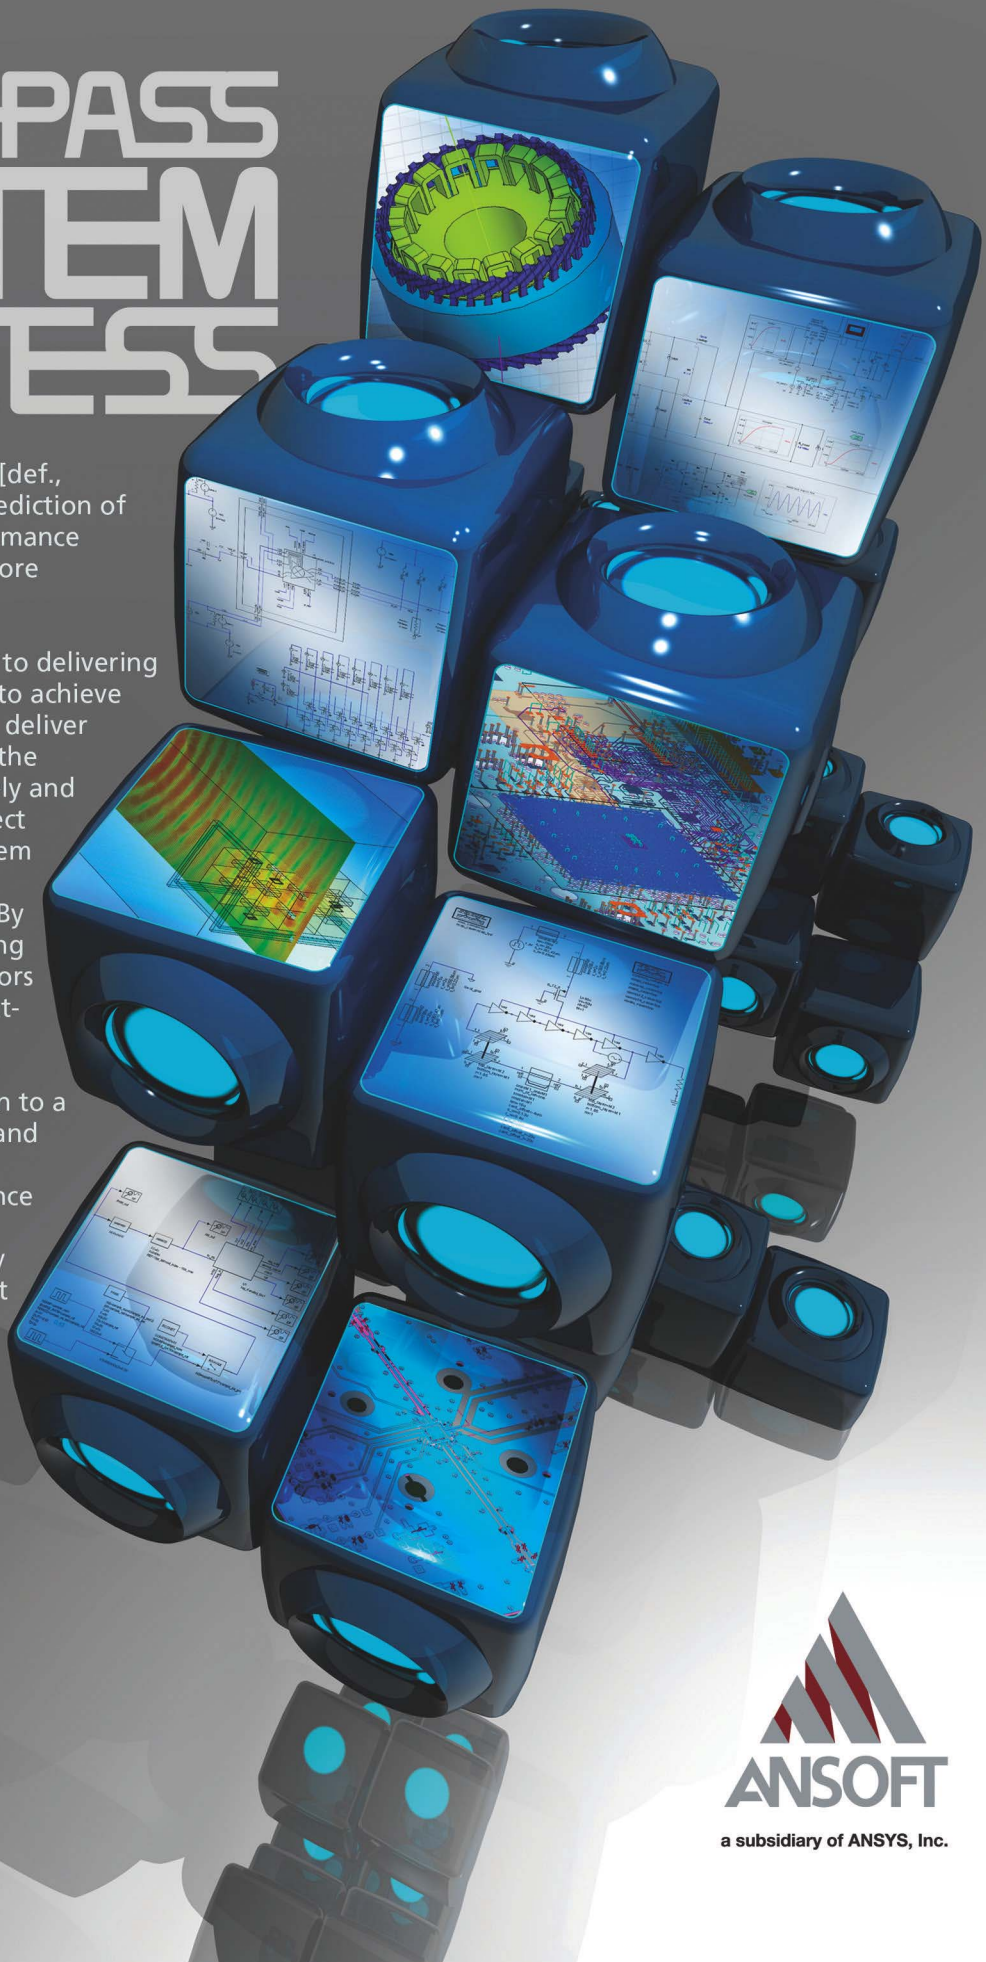

ANSOFT

a subsidiary of ANSYS, Inc.

I didn't run the physical-disks write tests because they require a blank drive.

The virtualized drive's improved performance versus that of a "real" drive is also likely in part a function of VMware's and other virtualizations' unique approaches to file storage. An NTFS, HFS+, or FAT (file-allocation-table) partition marks versions of files as old when you update or delete them, and other data eventually fills the space these obsolete files took up on the drive. Virtualized drives, on the other hand, employ databaselike linked-list structures. Such approaches deliver comparatively fast accesses. However, as is also the case with databases, virtualized drives require periodic compaction to cull the no-longer-current file entries. VMware calls this function shrinking, and it operates in cooperation with the virtualized Windows OS through the VMware Tools interface. During shrinking, the virtualized operating system is inaccessible; the function is fast, however.

As for testing virtualized-versus-native networking performance, my evaluation using Sandra's various benchmarking utilities was erratic, and I didn't trust the results I got. So, I instead chose a more elementary approach, a broadband-speed test, that produced nearly identical host-versus-virtualized results to within a reasonable run-to-run margin of impermanence. "Pings" to various LAN peripherals and WAN servers also produced identical results regardless of whether they came from the host OS X or the virtualized Windows XP.

POWER PENALTIES

When using Windows XP and its applications under Fusion, my MacBook's system fan runs more regularly and more robustly than when I use OS X alone. I decided to search for the source of the incremental heat generation, and the pursuit didn't take long to bear fruit (**Figure 5**). Note that Activity Monitor still assumes the existence of a single-core CPU in the system; the 56.5% figure for vmware-vmx on dual-core virtualized Windows XP, for example, translates to an overall 28.25% CPU burden. Note, too, that, even with Fusion in single-core mode, the CPU and EFI (extensible-firmware-interface) code controlling it still spread the vmware-vmx load across both cores when available. But keep in mind that I captured these Activity Monitor screenshots with virtualized Windows XP at an idle state—that is, with negligible CPU usage, according to Windows' Task Manager. When virtualized Windows XP is in normal use, vmware-vmx's system burden is substantially bigger.

Translating CPU usage into battery life is difficult, inexact, and a moving target; as VMware and its competitors' virtualization capabilities improve, they'll be better able to exploit systems' various power-efficient hardware-acceleration capabilities. For example, playing a DVD within virtualized Windows currently incurs a substantial CPU burden because the requisite tasks are implemented in software. However, upcoming Fusion Version 2 promises to improve video-playback performance, presumably by tapping into the MacBook's Intel graphics core's built-in MPEG-2-decoding circuitry. For now, I'll stick to watching movies on OS X's DVD player; the same native-versus-virtualized preference, when possible, is equally valid for other demanding applications. **EDN**

REFERENCES

1 Dipert, Brian, "Mac (under) the knife: piecing together the PowerPC puzzle," *EDN*, Sept 15, 2005, pg 44, www.edn.com/article/CA6255047.

2 Dipert, Brian, "Interface overkill? Is eSATA necessary for your next system design?" *EDN*, May 10, 2007, pg 48, www.edn.com/article/CA6437950.

FOR MORE INFORMATION

Apple
www.apple.com

EMC
www.emc.com

Intel
www.intel.com

Microsoft
www.microsoft.com

Parallels
www.parallels.com

SiSoftware
www.sisoftware.net

VMware
www.vmware.com

You can reach
Senior Technical Editor
Brian Dipert
at 1-916-760-0159,
bdipert@edn.com,
and www.bdipert.com.

EMBEDDED SERIAL NOR FLASH MEMORY

Product	Erasable Sector	Speed	Density
M25P†	256 Kb - 2 Mb	50 MHz, 75 MHz	512 Kb - 128 Mb
M25PX	4 KB, 64 KB	75 MHz (Dual I/O)	8 Mb - 64 Mb
M25/M45PE	256 B, 4 KB, 64 KB	50 MHz, 75 MHz	1 Mb - 16 Mb

EMBEDDED PARALLEL NOR FLASH MEMORY

Product	Voltage	Performance	Density
J3 v.D	2.7V - 3.6V	X8, X16 Page	32 Mb - 256 Mb
P30	1.7V - 2.0V, 1.7V - 3.6V	X16 52 MHz Burst	64 Mb - 512 Mb
P33	2.3V - 3.6V	X16 52 MHz Burst	64 Mb - 512 Mb
M29W†	2.7V - 3.6V	X8, X16 Page	4 Mb - 128 Mb
M58BW†	2.7V - 3.6V, 2.5V - 3.3V (55ns)	X32 Burst	16 Mb - 32 Mb

EMBEDDED NAND FLASH MEMORY

Product	Voltage	Page Size	Density
SLC small page	1.8V/3V	512 B	128 Mb – 1 Gb
SLC large page	1.8V/3V	2 KB	1 Gb - 8 Gb
CompactFlash* card	3V/5V	2 KB	64 MB - 4 GB
eMMC*	1.8V/3V	2 KB	1 GB

*Other names and brands may be claimed as the property of others.
†Automotive temp available.

Imagine the future. Reach for it. And don't let a little thing like memory hold you back.

At Numonyx, our mission is to make embedded memory one less barrier to your next breakthrough. By making reliability our obsession. By providing more choices. By delivering expert solutions and innovations to fuel the future.

So go ahead. Dream big. Design big. Discover how far we can help you go.

Visit **www.numonyx.com/embedded** to learn how you can win a free development board.

numonyx™

EFFICIENT POWER

While the world around you goes green, CUI and V-Infinity can keep you one step ahead of the efficiency curve. CUI offers a large selection of EISA 2007, CEC Level IV, and Energy Star compliant external power supplies with RoHS compliance and UL approvals. Efficient internal power supplies from V-Infinity include compact switching regulators and open frame switching power supplies.

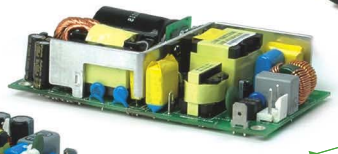

V78 series

- dc switching regulator
- pin compatible with 78XX linear regulators
- 500 and 1000 mA output current models
- efficiency up to 96%
- no need for a heatsink
- wide input range
- thermal shutdown
- low ripple and noise
- non-isolated
- short circuit protection

USB style series

- switching power supply
- 2.5 & 5 W output power
- USB type A receptacle output
- Energy Star / CEC / EISA 2007 compliant
- insulation resistance 100 M Ohm at 500 V dc
- no-load power consumption 0.5 W max.
- universal input
- UL/cUL approvals
- RoHS compliant

EMT series

- switching power supply
- 30 W output power
- interchangeable blades
- ac power cord inlet
- Energy Star / CEC / EISA 2007 compliant
- insulation resistance 50 M Ohm at 500 V dc
- no-load power consumption 0.5 W max.
- universal input
- UL/cUL, CE, FCC; TUV/GS, C-Tick approvals
- RoHS compliant

VMS-160 & VMS-365 series

- switching power supply
- **VMS-160**: 160 W output power in a 2”x 4” footprint and 18.2 W/in³ power density
- **VMS-365**: 365 W output power in a 3”x 5” footprint and 19 W/in³ power density
- single output voltages of 5, 12, 24, and 48 V dc
- 90% typical efficiency
- medical approvals
- universal input (90 – 264 V ac)
- built-in active PFC function
- 12 V auxiliary fan output

AC-AC series

- linear power supply
- 3 – 12 W output power
- 6ft. cord length - custom lengths available
- output voltage tolerance: ±5 % at rated load
- class 2 power supply
- Energy Star / CEC / EISA 2007 compliant
- no-load power consumption 0.5 W max.
- North American wall plug
- UL/cUL approvals
- RoHS compliant

CEC Level IV- The California Energy Commission has mandated requirements for power supplies used with certain types of products. The most current requirements are the same as the EISA 2007 requirements and are referred to as either “Tier 2” or “Level IV.”

Energy Star- Energy Star is a joint program of the US Environmental Protection Agency (EPA) and the US Department of Energy (DOE) aimed at preserving the environment through energy efficiency. Adapters meeting the Energy Star guidelines are up to 30% more efficient than non-compliant versions and must meet both active and no-load minimum efficiency requirements set forth by the EPA and DOE. Compliance with these requirements is voluntary.

EISA 2007- The Energy Independence and Security Act of 2007 was passed by Congress in December of 2007 and addresses minimum efficiency standards for external power supplies manufactured on July 1, 2008 and after. This law stipulates the energy efficiency criteria for adapters in active mode depending upon their power rating. The stipulated energy consumption for all adapters in no-load mode must be less than 0.5 W according to EISA 2007. Compliance with these requirements is mandatory.

ENERGY STAR PARTNER

available through Digi-Key CORPORATION

CUI INC | V-INFINITY

visit: www.cui.com/efficient *or* www.v-infinity.com/efficient **call:** 800.275.4899

BY SHAKEEL PEERA • LATTICE SEMICONDUCTOR

Extending SPI4.2 capabilities for Ethernet services

WITH THE PROLIFERATION OF INTERNET PROTOCOL-BASED SYSTEMS IN THE TELECOMMUNICATIONS MARKET, DESIGNERS ARE TURNING TO FPGAs TO CREATE INTELLIGENT ETHERNET BRIDGES AND TRAFFIC MANAGERS.

As carriers and cable providers start delivering integrated video, voice, and data over a common broadband infrastructure to their customers, OEMs are increasing their efforts to roll out IP (Internet Protocol)-based systems, including passive-optical networks, cable-modem-termination systems, DSLAMs (digital-subscriber-line-access multiplexers), multiservice switches, and other access and back-haul equipment. The underlying physical layer for this equipment is the ubiquitous Ethernet technology. The system-level components in next-generation Ethernet-service cards are framers, NPUs (network-processing units), and off-the-shelf Ethernet switches—leaving design engineers with the challenge of developing a programmable bridge between the parts, fitting the solution, and cost-effectively implementing it.

Line cards with network processors and framers often use SPI4.2 (system-packet interface, Level 4, Phase 2). Although the specification addresses the challenge of achieving fast, low-latency, point-to-point, 10-Gbit physical connectivity, it leaves the user to implement efficient buffer-management schemes as they relate to the system design. The consequences of the user's decisions affect system-bandwidth efficiency and are concerns to system vendors, which must demonstrate the suitability of an all-IP network to service providers seeking the determinism, reliability, and SLA (service-level-agreement) guarantees of SONET (synchronous-optical-networking) and SDH (synchronous-digital-hierarchy) multiplexing protocols and ATM (asynchronous-transfer-mode)-based systems. These service providers also desire the ubiquity and reduced development, capital-equipment, and operational costs of Ethernet.

SPI4.2-BASED BRIDGES

An FPGA can play a host of roles, but, in its most basic form, it is a programmable gasket or bridge between the framer and the NPU or between the NPU and the carrier-class-Ethernet switch (**Figure 1**). There are a number of must-have components in this context.

The first component is an XAUI (10-Gbit attachment-unit interface) that supports class-of-service and port-switching-information overlays. These interfaces need to run above and beyond the IEEE 802.3-mandated rates of 3.125 Gbps to maintain a 10-Gbps line rate. One example of this interface is the proprietary Broadcom HiGig+ interface, which runs at 3.75 Gbps. Conversely, if the system is aggregating multiple lanes of 1-GbE (gigabit-Ethernet) or 10/100-GbE streams, the best choice is usually multiple 1000BaseX-compliant interfaces in the form of IEEE 802.3z, a GbE interface, or SGMII (serial-gigabit media-independent interface).

The second must-have component is a fully compliant 10-GbE interface or multiple triple-speed MACs (media-access controllers) that can add, strip, and process-overlay the headers.

The third necessary component is the bridge, which is responsible for accepting packets in the ingress direction, ensuring that transfers occur to the Ethernet domain under the right conditions—for example, throttling during flow control. In the egress direction, the system presents traffic depending on the condition of the SPI4.2 status channel. The user requests Ethernet flow control based on the system criteria. The bridge also performs bus translation in both directions.

Finally, for control, a system bus monitors status, provides program registers for user-controlled options, and provides interrupt support.

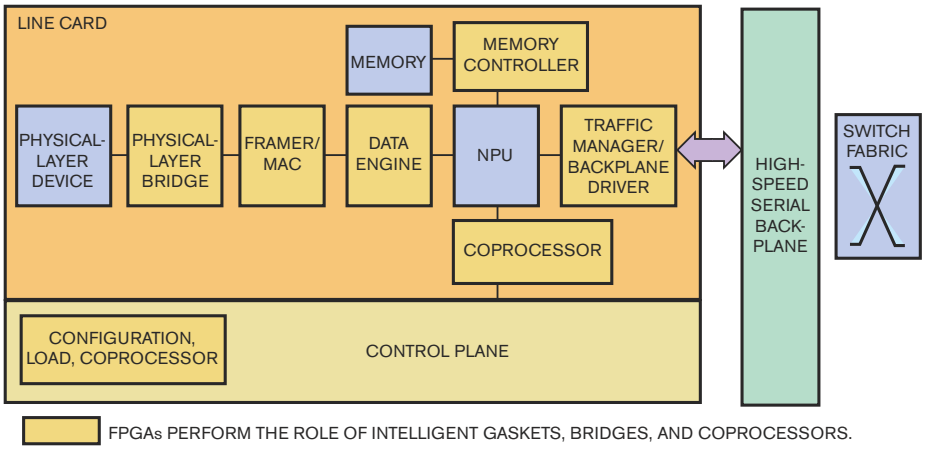

Figure 1 An FPGA can play a host of roles, but, in its most basic form, it is a programmable gasket or bridge between the framer and the NPU or between the NPU and the carrier-class-Ethernet switch.

You can find these attributes in either off-the-shelf ASSPs (application-specific standard products) or FPGAs. However, FPGAs provide several additional benefits. First, they allow the user to instantiate multiple bridges into a single monolithic device. Because carrier-Ethernet systems require more determinism and quality-of-service guarantees than the traditional IEEE 802.3 specification provides, switch vendors are taking proprietary approaches to provide flow control, channelization, interleaving, and OAM (operation/administration/maintenance). With each generation of products, switch vendors are moving away from traditional Ethernet and adding their secret sauces, making FPGAs ideal for keeping pace with burgeoning proprietary interfaces.

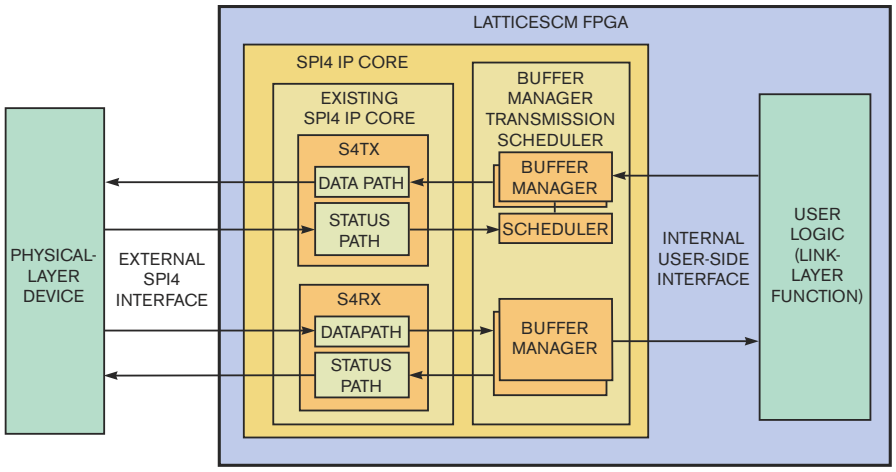

Figure 2 It is best to implement a suitable buffer manager that resides outside the basic SPI4.2 datapath and dictates transmit bandwidth commensurate with the end application.

One of the biggest advantages of programmability is that the customer can define the bandwidth-management and -provisioning capabilities that best suit the nuance of the end application. An ASSP or ASIC would have to support multiple schemes to be a truly catchall product. However, a low-cost FPGA can implement the scheme the user needs, eliminating the unnecessary power, real estate, and cost of ASSPs and ASICs.

SPI4.2 BANDWIDTH AND BUFFER MANAGEMENT

Many bridging applications have packets arriving at full 10-Gbps rates over XAUI, for example, but may need to play out at a much slower rate for a given channel on the SPI4.2 side—155 Mbps, for example. The system must also support dynamic channel reprovisioning so that you can add or subtract individual subscriber lines or bandwidth to and from a channel while a system is live without disrupting service on the other channels. It is best to implement a suitable buffer manager that resides outside the basic SPI4.2 datapath and dictates transmit bandwidth commensurate with the end application (**Figure 2**).

PHYSICAL PER-CHANNEL BUFFER

A physical-buffer approach is ideal when the application requires a few FIFO buffers and serves multiple asynchronous physical interfaces that are independent of each other—for example, when bringing multiple 2.5- or 1-Gbps interfaces into one 10-Gbps SPI4.2. Individual FIFOs work well because the number of FIFOs is small—10 in the 1-Gbps scenario—and each provides an independent interface, including independent clocks for the external interfaces. This approach is the simplest and most straightforward for solving the per-channel-buffer-design question. The architecture is N channels=N physical buffers of equal depth=N physical interfaces=N FIFO controllers, where N does not exceed 16 channels.

Additionally, this architecture would handle error and overflow conditions—usually, a packet-drop capability with FIFO-pointer readjustments—that you base on user-defined criteria. However, this architecture has several potential drawbacks that the virtual buffer addresses. With Ethernet, you must con-

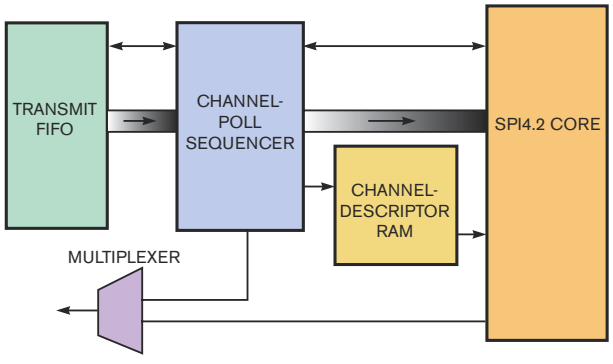

Figure 3 The scheduler could comprise a poll sequencer and a channel descriptor.

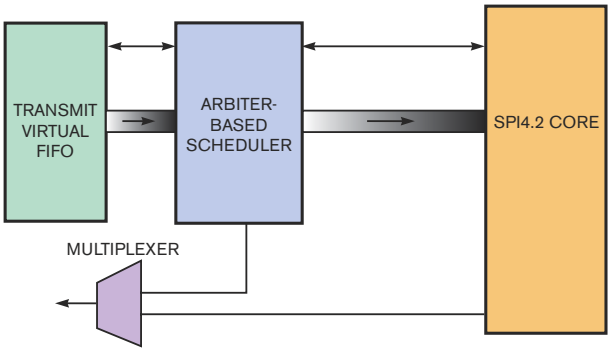

Figure 4 The objective of an arbiter-based architecture is to skip channels that have nothing to send and prioritize those channels that do.

LINEAR TECHNOLOGY

DESIGN NOTES

Current Sense Amp Inputs Work from –0.3V to 44V Independent of Supply

Design Note 451

Glen Brisebois

Introduction

Monitoring current flow in electrical and electromechanical systems is commonly used to provide feedback to improve system operation, accelerate fault detection and diagnosis, and raise efficiency. A current monitoring circuit usually involves placing a sense resistor in series with the monitored conductor and determining the voltage across the sense resistor. To minimize power loss in the sense resistor it is kept as small as possible, resulting in a small differential voltage that must be monitored on top of what may be a fairly large varying common mode voltage. The LT®6105 is an ideal current sense amplifier for this application. Just give it any reasonable supply voltage, say 3V, and its inputs can monitor small sense voltages at common modes of –0.3V to 44V and anything in between. The accuracy of the LT6105 over this range is displayed in Figure 1.

Solenoid Monitoring

The large input common mode range of the LT6105 makes it suitable for monitoring currents in quarter, half and full bridge inductive load driving applications. Figure 2 shows an example of a quarter bridge. The MOSFET pulls down on the bottom of the solenoid to increase solenoid current. It lets go to decrease current, and the solenoid current freewheels through the Schottky diode. Current measurement waveforms are shown in Figure 3. The small glitches occur due to the action of the solenoid plunger, and this provides an opportunity for mechanical system monitoring without an independent sensor or limit switch.

Figure 4 shows another solenoid driver circuit, this time with one end of the solenoid grounded and a P-channel MOSFET pulling up on the other end. In this case, the inductor current freewheels around ground, imposing a negative input common mode voltage of one Schottky

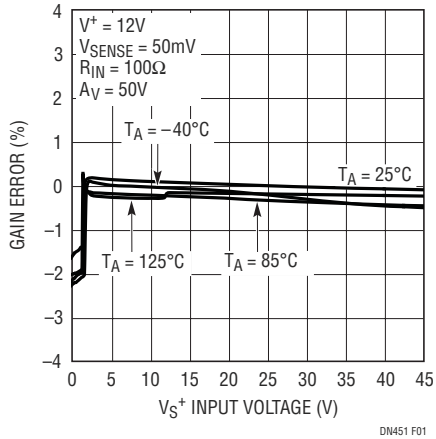

Figure 1. Gain Error vs Input Common Mode

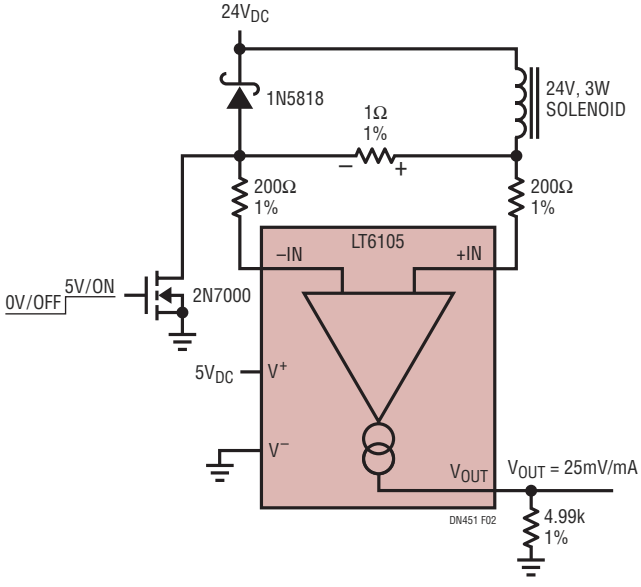

Figure 2. Solenoid is Pulled Low, Freewheels High. Input Travels from 0V to 24.3V

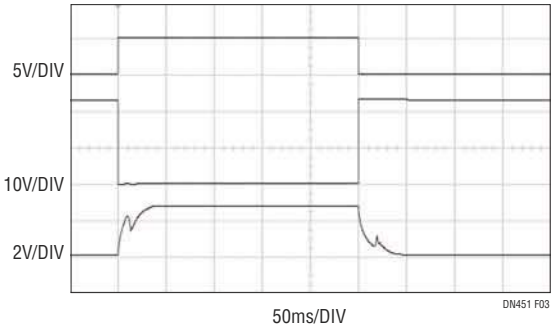

Figure 3. Solenoid Waveforms: MOSFET Gate, Solenoid Bottom and Current Sense Amp Output. Bumps in the Current Result from Plunger Travel

diode drop. This voltage may exceed the input range of the LT6105. This does not endanger the device, but it severely degrades its accuracy. In order to avoid violating the input range, pull-up resistors can be used as shown in Figure 4.

Supply Monitoring

The input common mode range of the LT6105 allows it to monitor either positive or negative supplies. Figure 5 shows one LT6105 applied as a simple positive supply monitor, and another LT6105 as a simple negative supply monitor. Note that the schematics are practically identical and both have outputs conveniently referred to ground. The only requirement for negative supply monitoring, in addition to the usual constraints of the absolute maximum ratings, is that the negative supply to the LT6105 must be at least as negative as the supply it is monitoring.

Conclusion

Current measurement is popular because it offers improved real time insight into matters of efficiency, operation and fault diagnosis. The wide input range of the LT6105 and its accuracy over that range make it easy to measure currents in a variety of applications.

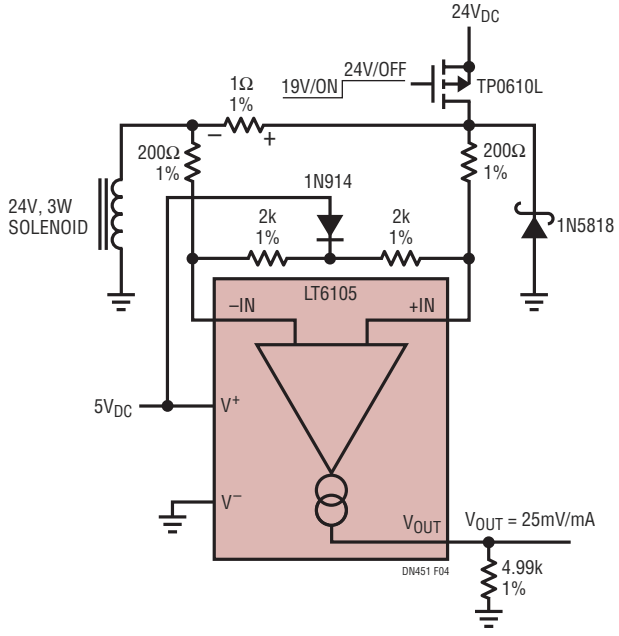

Figure 4. Solenoid is Pulled High to 24V, Freewheels Low to a Schottky Below Ground. LT6105 Inputs are Kept Within Range by 2k Pull Ups

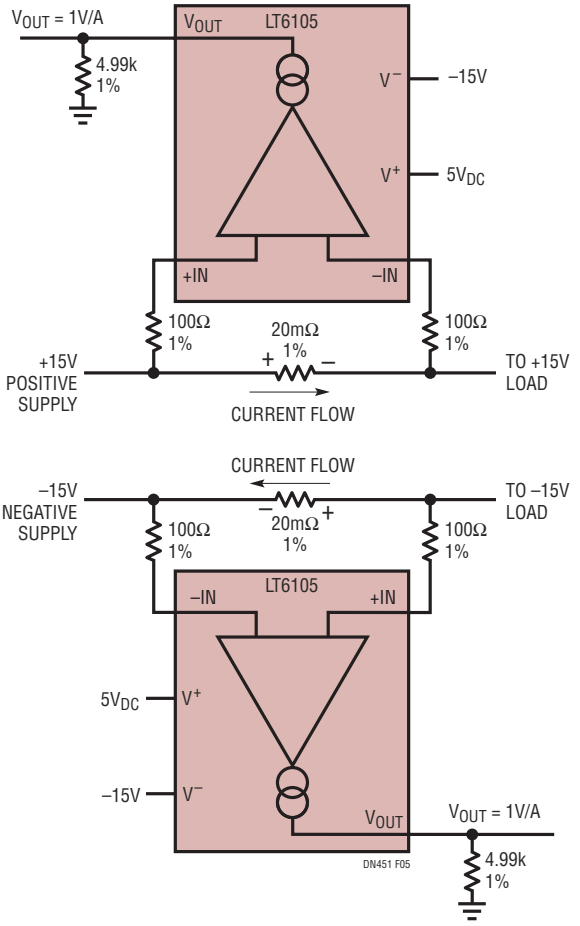

Figure 5. The LT6105 Can Monitor the Current of Either Positive or Negative Supplies, Without a Schematic Change. Just Ensure that the Current Flow is in the Correct Direction

Data Sheet Download

www.linear.com

For applications help,
call (408) 432-1900, Ext. 3755

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

dn451f LT/TP 1008 246K • PRINTED IN THE USA

sider packet sizes of 9 kbytes or greater for worst-case scenarios. If your application requires more than one or two instances of jumbo-packet storage in both the transmitting and the receiving directions for one of the physical channels, you cannot use any of the other memory from other underused channels to reduce overall memory demands. This architecture does not work well for applications in which you must add or subtract bandwidth on the fly without affecting other channels. Therefore, you should budget the worst-case amount of per-channel memory beforehand.

Because this architecture assumes multiple independent asynchronous clock domains—one for each physical channel—the system cannot share FIFO-controller intelligence across channels to optimize logic resources.

> ⊞ Go to www.edn.com/ms4293 and click on Feedback Loop to post a comment on this article.
>
> ⊞ For more feature articles, go to www.edn.com/features.

VIRTUAL PER-CHANNEL BUFFER

The virtual buffer uses a large memory structure in which one FIFO controller can manage many logical FIFO buffers. The system treats the large memory as a global resource that it can allocate dynamically for a channel through linked lists. Because the virtual buffer uses dynamic-memory allocation, it allows the engineering of peak memory-resource requirements at the system level rather than at the physical channel level, lowering overall memory needs. This approach is viable when a large number of synchronous channels are in operation—for example, when aggregating multiple SONET/SDH tributaries through a virtual-concatenation framer to an XAUI bridge.

You can implement the link list by storing the current state information, including descriptor locations, fill level, and other information, in a high-density RAM that loads into the FIFO controller when you select a channel. The system separately stores data memory and dynamically allocates it to the channels without regard to any per-channel constraints. When an application needs to control the amount of memory for a channel, this process can occur in the user domain of the buffer manager simply by using the per-channel status information, which indicates the number of memory segments each channel uses.

This architecture achieves the objective of a buffer design that supports a large number of synchronous channels, ensuring dynamic-memory allocation and maintaining full and constant traffic flows. The obvious drawback is that this scheme assumes fully synchronous physical interfaces. You can design a retrofit with small clock-domain-conversion FIFOs in front of the large virtual FIFO, allowing it to operate more asynchronously.

SEQUENCER-BASED SCHEDULER

You can implement transmitter scheduling through the use of the SPI4.2 calendar sequence, which allows the user to re-

ceive status updates regarding each channel's bandwidth allocation. You can then design a scheduler to schedule channels for transmission, basing its action on the direct-channel-sequence information that the user enters in the SPI4 status calendar's RAM.

From an implementation perspective, the designer needs to ensure that the SPI4.2 circuit supports random access to the status of any channel in a single clock cycle for the scheduler. The scheduler itself could comprise a poll sequencer and a channel descriptor (**Figure 3**). The poll sequencer would use channel ID to poll channel status. Transmission for the channel would begin as soon there were enough data available to satisfy far-end-receiver requirements. If the poll sequencer did not meet the criterion, the scheduler would poll the next channel for transmission. The channel descriptor, stored in RAM, contains the channel ID and burst parameters for that channel.

One potential drawback of such an architecture is the time it spends polling channels with no assurance that any of the channels have anything to send. In such a case, you could consider an arbiter that would service only relevant channels.

ARBITER-BASED SCHEDULER

The objective of an arbiter-based architecture is to skip channels that have nothing to send and prioritize those channels that do (**Figure 4**). The scheduled RAM sequencer polls every channel, which can add undesirable latency.

One approach is to assign each channel a location in a RAM table and then to add a weight for each channel. The weight would specify how much of the total SPI4.2 pipe the system could allocate to each of the channels. A conventional round-robin arbiter would then flag channels that have a larger weight and have data to send first. Traffic-management and shaping circuits commonly use this scheme, weighted round robin, to handle and police priority requests for traffic. You can effectively use the same scheme in SPI4.2 buffer management.

SUMMARY

With both the power consumption and the prices of programmable-logic devices plummeting, as well as the advent of embedded high-speed Ethernet-related interfaces, it becomes difficult to ignore the case for FPGAs in carrier-Ethernet systems. Most important, for buffer-management schemes, it is improbable that an ASSP can achieve the type of cost- and power-effective flexibility that FPGAs can offer in tailor-made, intelligent bridges, gaskets, and traffic managers. EDN

AUTHOR'S BIOGRAPHY

Shakeel Peera is director of marketing for high-performance solutions at Lattice Semiconductor (Bethlehem, PA). He previously held various technical-marketing positions in FPGAs, ASICs, and networking intellectual property at AT&T Microelectronics, Lucent Technologies, and Agere Systems. Peera holds a bachelor's degree in electrical engineering from the University of Rochester (New York). You can reach him at shakeel.peera@latticesemi.com.

POWER BY DESIGN

Custom power solutions designed to fit your specific needs

Vicor Custom Power: Small company responsiveness, large company resources

The sole focus of Vicor Custom Power is designing and manufacturing turnkey custom power systems that meet your specific needs. Vicor Custom Power maintains the flexibility of a small entrepreneurial company while taking advantage of Vicors technical and business resources to deal effectively with your most challenging power requirements. Vicor has invested in the tools and resources to offer you full service solutions from prototype to mass production with the shortest lead times and the lowest risk.

General Capabilities:

- Electrical and Mechanical Design
- Rapid Prototyping
- High Volume Production Capacity
- MIL-STDs Compliance
- Reliability / Certification Testing:

High Temperature Operational Life	Altitude
HALT *(Highly Accelerated Life Test)*	Explosive Atmosphere
Mechanical / Thermal Shock	Temperature Cycling
Vibration	Burn In
Humidity	EMI
Acceleration	Transient Immunity

Put Vicor Custom Power to work for you today, call 1-800-496-5570 to speak with a Vicor Custom Power engineer, or e-mail apps@vicorcustom.com

vicorcustom.com

VICOR CUSTOM POWER

99.91% Accurate Silicon Oscillator

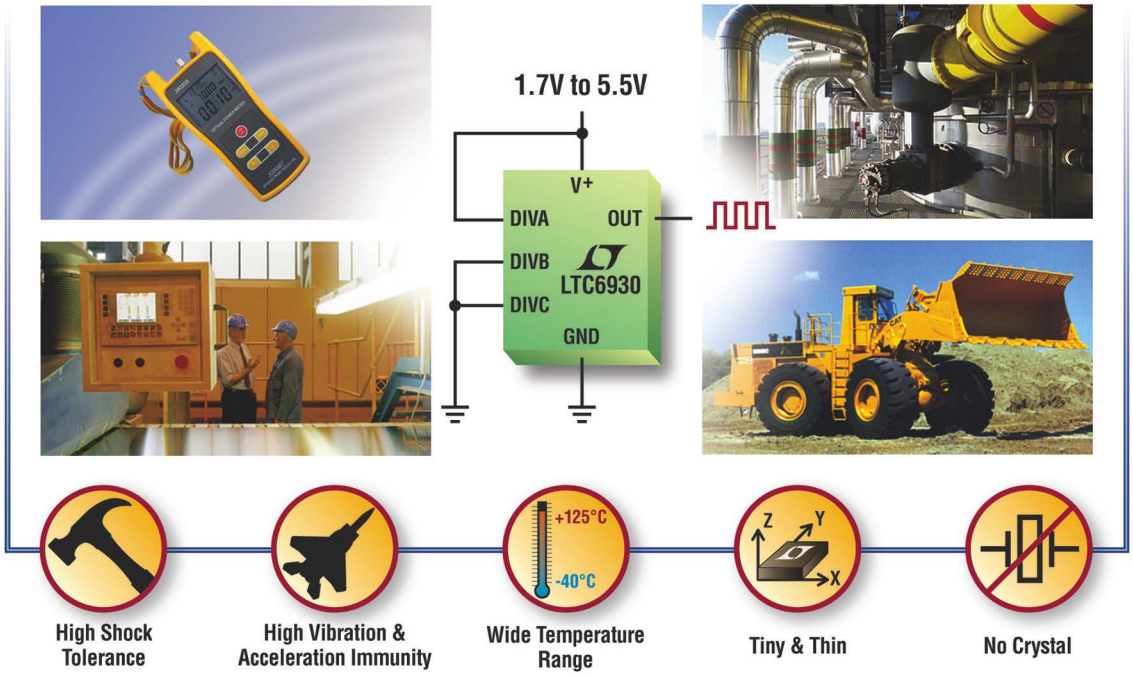

Are your systems subject to shock, vibration, acceleration or temperature extremes? Linear Technology's new LTC®6930 silicon oscillators offer solid advantages over traditional clocks. Built with standard silicon fabrication and assembly techniques, these oscillators require no crystals and no trim components. You can count on short lead times, glitch-free startup and reliable operation. The LTC6930 oscillators are the perfect choice for automotive and industrial applications.

Features

- Frequency Error 0.09% Max. at 25°C
- Frequency Error 0.45% Max. over 0 to 70°C, 0.65% Max. over -40 to 85°C
- 110µs Maximum Startup Time
- 1.7V to 5.5V Single Supply Operation
- 490µA Supply Current at 8MHz
- -40°C to 125°C Operation
- 2mm x 3mm DFN or MSOP Package
- Short Lead Times

5 Options Cover 32kHz to 8.192MHz:

LTC6930-4.19	32.768kHz to 4.194MHz
LTC6930-5.00	39.063kHz to 5.000MHz
LTC6930-7.37	57.600kHz to 7.373MHz
LTC6930-8.00	62.500kHz to 8.000MHz
LTC6930-8.19	64.000kHz to 8.192MHz

Each Oscillator Offers 8 Pin-Selectable Frequencies

Info & Free Samples

www.linear.com/6930
1-800-4-LINEAR

Free Industrial Signal Chain Brochure

www.linear.com/industrial

LINEAR TECHNOLOGY

EDITED BY MARTIN ROWE
AND FRAN GRANVILLE

designideas

READERS SOLVE DESIGN PROBLEMS

Use an LM317 as 0 to 3V adjustable regulator

Vladimir Rentyuk, Modul-98 Ltd, Zaporozhye, Ukraine

Most engineers know that they can use an inexpensive, three-terminal adjustable regulator, such as Fairchild Semiconductor's (www.fairchildsemi.com) LM317, as an adjustable regulator to only some necessary value of voltage, such as 36 or 3V. This value cannot be less than 1.25V without employing other approaches, however. The devices' inner reference voltage is 1.25V, and their output voltage accordingly cannot be less than this value without potential bias (**Reference 1**). One way to solve this problem is to use a reference-voltage source based on two diodes (**Reference 2**). Although this approach is suitable for a 1.2 to 15V or higher-voltage regulator, it is not appropriate for an extra-low-voltage fixed- or adjustable-voltage regulator. The two 1N4001 diodes it employs do not provide the needed potential bias of 1.2V, and they have additional temperature instability of approximately 2.5 mV/K (**Reference 3**). Hence, additional temperature drifting of the output voltage is approximately 100 mV; it is more than 6% for a 1.5V output voltage and 10% for a 1V output voltage if you adjust the temperature to 20°C—a typical indoor situation. You can solve these problems by using a Fairchild Semiconductor LM185 or an Analog Devices (www.analog.com) AD589 adjustable-voltage-reference IC. These devices are expensive, however, and, in this case, they require not only additional zero adjustment but also matching. These adjustments at their reference voltages are 1.215 to 1.255V and 1.2 to 1.25V for the LM185 and AD589, respectively. Note that the reference voltage of the LM317 is 1.2 to 1.3V.

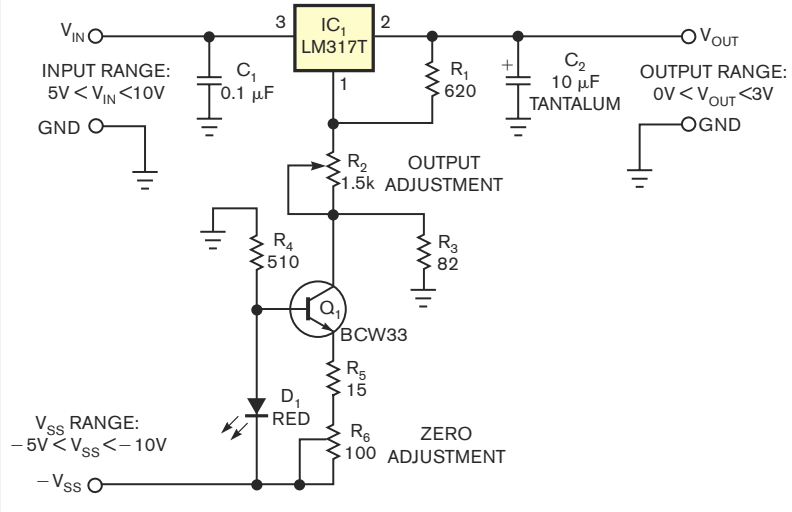

Figure 1 This circuit is an inexpensive approach using a simple 0 to 3V adjustable regulator.

Figure 1 shows an inexpensive approach using a simple 0 to 3V adjustable regulator. You implement the necessary potential bias using a simple temperature-stabilized constant-current source (**Reference 4**). You calculate this current source using the following **equation**: $I=(V_F-V_{EBO})/(R_5+R_6)$, where V_F is D_1's forward voltage of approximately 2V and V_{EBO} is Q_1's emitter-base voltage of approximately 0.68V. The current is approximately $1.32/(R_5+R_6)$. The constant-current source creates a bias voltage of approximately −1.25V on resistor R_3. You implement the zero adjustment using resistor R_6, which can change the current of the constant-current source. Resistor R_5 protects transistor Q_1. You can use D_1 as a light indicator. You can adjust the output voltage using resis-

DIs Inside

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

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

tor R_2. Calculate the output voltage as follows: $V_{OUT}=V_{REF}(1+R_2/R_1)-V_{R3}$, where V_{REF} is the reference voltage of IC_1 and V_{R3} is some compensative voltage of resistor R_3. You should establish this voltage to equal the reference voltage for its compensation. In this case, $V_{OUT}=V_{REF}(R_2/R_1)$. With R_2 having a value of 1.2 kΩ, this circuit found use as the equivalent of a typical battery with an output voltage of 1.56V for development projects. EDN

REFERENCES

1 "LM317 3-Terminal Positive Adjustable Regulator," Fairchild Semiconductor Corp, June 2005, www.fairchildsemi.com/ds/LM/LM317.pdf.
2 "LM350 3-Terminal 3A Positive Adjustable Voltage Regulator," Fairchild Semiconductor Corp, 2001, www.fairchildsemi.com/ds/LM/LM350.pdf.
3 Schenk, C, and Ulrich Tietze, *Halbleiter-Schaltungstechik*, Springer-Verlag Berlin Heidelberg, 2002, ISBN: 3540428496.
4 Rentyuk, Vladimir, "The Simple Temperature-Stabilized Constant-Current Source," *Electronics World*, November 2006.

Alarm monitors rotational speed of dc motor

Peter Demchenko, Vilnius, Lithuania

You can use the circuit in **Figure 1** to monitor the rotating speed of a dc fan motor and sound an alarm if the motor stalls. One potential application of the circuit is monitoring the CPU-fan speed in a PC in which overheating the CPU can ruin the whole system. A PC BIOS (basic input/output system) often has a limited capability for monitoring the speed of CPU or chassis fans during boot-up. Moreover, if you enable the CPU-fan-protection function of BIOS today, you can have a problem with it tomorrow: If the fan's starting acceleration slows down, the BIOS powers down the PC at the beginning of the boot sequence, not allowing you to go into BIOS settings to correct the situation. So, the manual often advises you to disable this fan function. The circuit in **Figure 1** shows how to implement continuous monitoring and sound an alarm and automatically power off the system if a fan problem occurs.

The impulses on R_1, arising from commutation in the fan's brushless motor, start up the Schmitt trigger, Q_1/Q_2, which controls transistor switch Q_3, commutating the sense pin of the fan's motherboard connector; the frequency of commutation is proportional to the rotation speed. Optionally, the output of the trigger resets the timer with two time-out periods; the expiration of the first time-out activates the alarm buzzer.

After the second time-out, transistor Q_5 powers down the PC with or without the relay switch. The relay switching is more consistent, is less prone to interference, and is preferable when the distance between this circuit and the power-switch connector on the motherboard exceeds 20 to 30 cm. You must connect the collector of Q_5 or the contacts of the relay in parallel with the power-switch button. The alarm circuit comprises Q_4 and a three-terminal piezoelectric buzzer. EDN

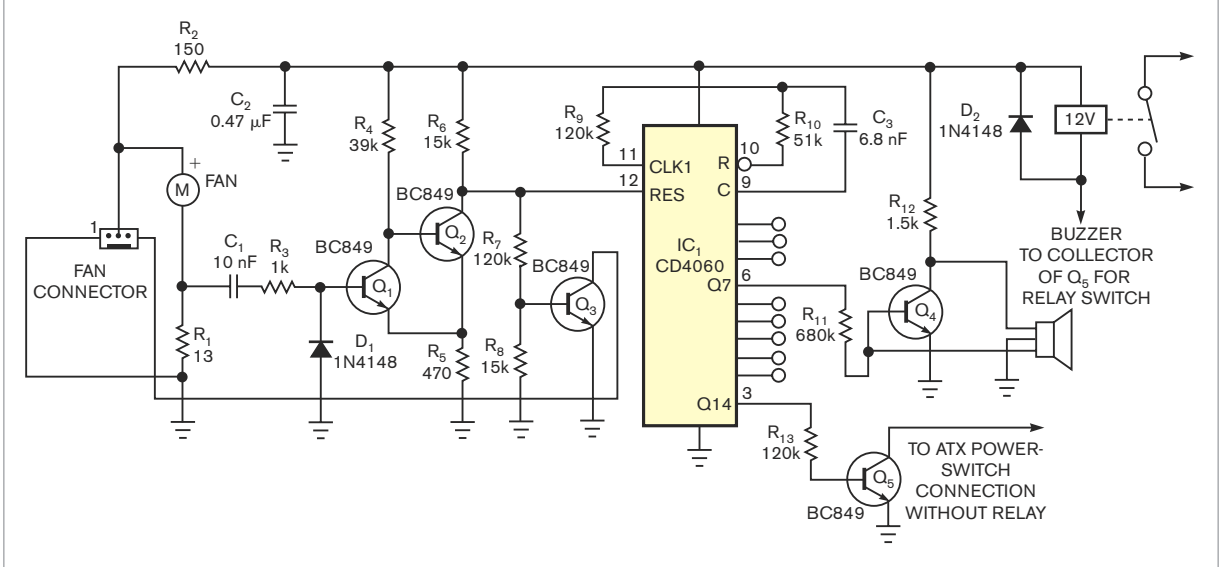

Figure 1 This circuit provides an optional audible alarm after a time-out when a brushless-dc fan motor slows down. Then, after a second time-out, the circuit powers down the PC.

16-Bit, 130Msps ADC + Driver

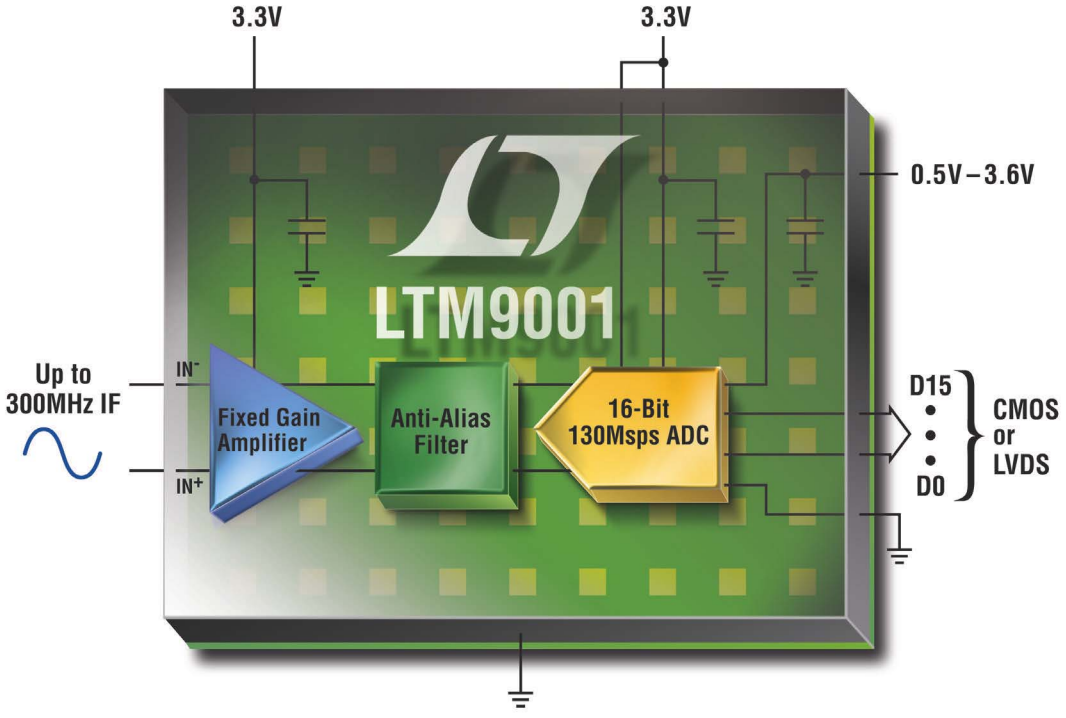

µModule™ Subsystem Dramatically Reduces Design Complexity & Board Space

The LTM®9001 is a 16-bit, 130Msps IF/baseband receiver subsystem that leverages years of applications engineering expertise to maximize high speed ADC performance. The LTM9001 alleviates the need for driver and ADC impedance matching, filtering, bypass placement and layout, eliminating long hours of troubleshooting and reducing time to market. With no external components required, the LTM9001 provides a high performance solution in less than half the board space of a discrete implementation.

Features

- Integrated 16-Bit, 130Msps ADC, Filter and Amplifier
- 72dB SNR, 82dBc SFDR @ 162.5MHz
- Up to 300MHz IF Range
- Fixed Gain: 8dB, 14dB, 20dB or 26dB
- 50Ω, 200Ω or 400Ω Input Resistance
- Integrated Bypass Capacitance, No External Components Required
- Optional Internal Dither
- Optional Data Output Randomizer
- 11.25mm x 11.25mm LGA Package

LTM9001: 16-Bit IF/Baseband Receiver Subsystem

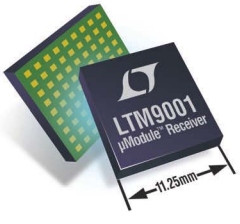

Info & Free Samples

www.linear.com/9001

1-800-4-LINEAR

Wireless Solutions Brochure

www.linear.com/wireless

LINEAR TECHNOLOGY

Add charging status to simple lithium-ion charger

Peter T Miller, Applied Inspirations, Bethlehem, CT

Like other simple, single-cell lithium-ion battery chargers, Microchip's (www.microchip.com) MCP73812 provides no means of indicating the charging status. You can remedy this situation by adding four components (**Figure 1**). Add one more LED, and you also get a charging-complete indication. This two-LED configuration has the added benefit that one of the LEDs is always on, providing an indication that the charger is powered.

While the cell is in the constant-current charging mode, 401 mA flows through the 1N4001 diode, D_1. The additional 1 mA is the supply current of the control chip. Because the 1N4001 conducts before the base-emitter junction of Q_1, it prevents Q_1 from turning on until the forward voltage across it reaches about 450 mV. Q_1 then starts to conduct and turns on D_2, a red LED that indicates charging. Because the forward-voltage drop for a green LED is typically higher than that of a red LED—2.1 versus 1.7V—the voltage across D_2 and Q_1 is less than the turn-on voltage of the green LED, D_3, and it remains off.

For the last part of the charging cycle, the controller switches to constant-voltage mode. As the cell voltage gets closer to this 4.2V terminal voltage, the current through D_1 drops, and at 15 to 40 mA, both LEDs illuminate.

Tests measured this range for several 2N3904 transistors. Testing with 2N4401s gave a lower range of 4 to 18 mA. When the current drops below about 15 mA, Q_1 turns off D_2. The voltage across D_3 now rises above its forward-voltage threshold, and the green charging-completed LED lights. EDN

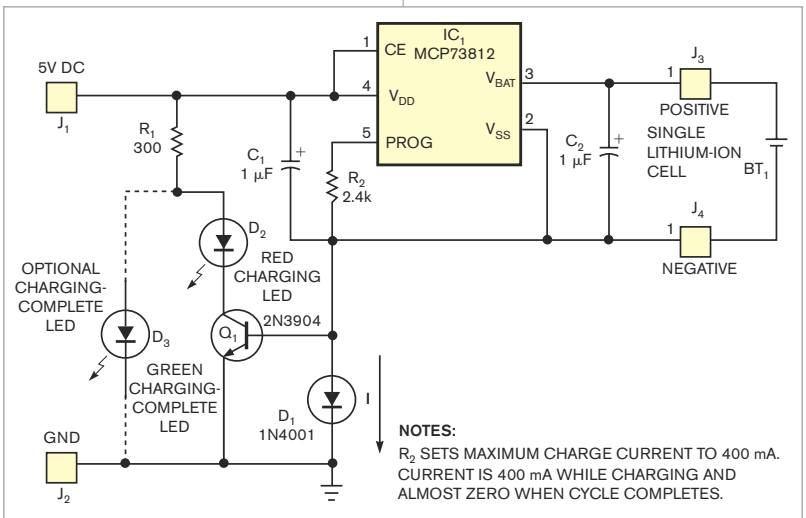

Figure 1 Adding a few components to a lithium-ion cell charger provides an indication of charging status.

555 timer drives multiple LEDs from one NiMH cell

Chuck Irwin, Hendersonville, NC

Using a CMOS 555 timer and a single NPN transistor, you can drive as many as seven LEDs using a minimal amount of voltage and power from a single NiMH (nickel-metal-hydride) AA cell. The circuit works by creating much higher-voltage pulses than the voltage for powering the circuit by pulsing a high-Q power inductor. The circuit

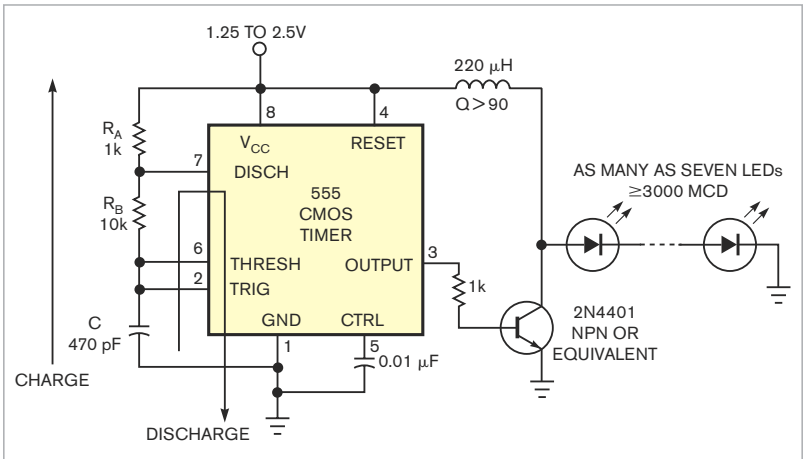

Figure 1 Using a CMOS 555 timer configured as a switching power supply, you can drive seven high-brightness LEDs from a single 1.25V cell.

Gigabit multiport LVDS crosspoint switches minimize system cost

Reduce the number of point-to-point LVDS links

Ideal for automotive applications

- LIN, I²C, or pin programmable
- -40°C to +105°C operating temperature range
- ±25kV ESD protection*
- Preemphasis improves signal integrity

Greater flexibility

- Flexible crosspoint switch
- Mux any input to any output
- Broadcast any input to all outputs
- 3:2 (MAX9132) and 3:4 (MAX9134) configurations

*ISO Air Discharge Model.

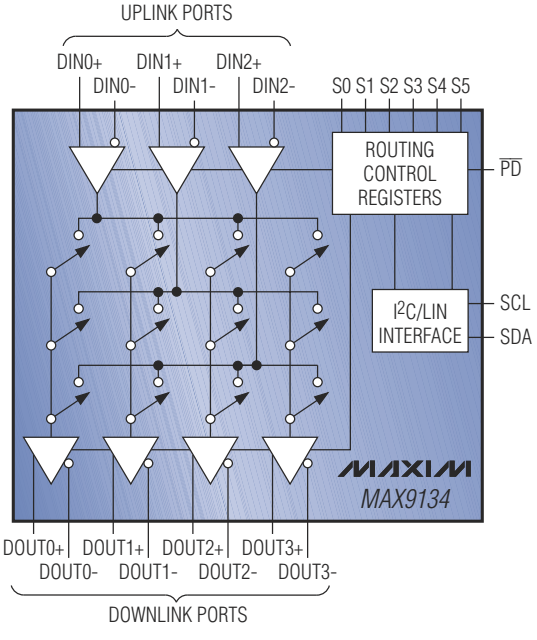

MAXIM DIRECT
www.maxim-ic.com/shop

AVNET electronics marketing
www.avnet.com

MAXIM
www.maxim-ic.com/MAX9132-info

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

creates voltage pulses of 23V using a 1.25V NiMH cell with seven connected LEDs.

The circuit uses a CMOS timer because it functions on low voltages—in this case, as low as 1V. A single white LED rated at 9300 mcd maintains its brilliance down to this low voltage. The circuit works for 192 hours using a 2000-mAHr-rated NiMH cell. The output of the timer is a 4.5-µsec pulse repeating at a 222-kHz rate. Although you can use the circuit to power any LED, it works best using high-brightness, high-power LEDs rated at 3000 mcd or higher. Obviously, the higher the millicandela rating, the brighter the LED will appear.

You can connect the LEDs in parallel if their forward voltages match; otherwise, the LED with the lowest forward voltage will dim out the other LEDs. Using the parallel connection, all LEDs will glow with equal brightness if their forward voltages match. Adding LEDs does not increase the current drawn from the battery but reduces the brilliance of all of the connected LEDs.

The advantage of connecting the LEDs in series—which is possible because of the high pulse voltage they produce—is equal brilliance of all LEDs, regardless of their individual forward-voltage drops and millicandela ratings. Each additional LED decreases additional voltage and lowers the resulting current into the series string of LEDs, lowering their brilliance. Using seven LEDs with a single 1.25V cell draws a current of only 8 mA. By adding a 1.25V cell to the power input, the LEDs become so brilliant that it is difficult to look at them. With a 2.5V supply, the peak voltage pulses increase to 70V with no connected LEDs. With the LEDs connected, the output voltage peaks at 25V. Current draw at 2.5V is 20 mA. EDN

Microcontroller inputs parallel data using one pin

Rex Niven, Forty Trout Electronics, Eltham, Victoria, Australia

Inputting multiple bits of information using a single entry pin of a microcontroller without the complexity of UARTs can prove useful. Such a scheme could allow scanning of a keyboard, mode switches, or any relatively slowly changing digital data. **Reference 1** details a technique for outputting signals with a single pin. The data from switch bank S_1 first presents itself to IC_3, a 74HC165 parallel-to-serial converter from NXP Semiconductors (www.nxp.com, **Figure 1**). Loading the data into the shift register requires a pulse on the PL line (Pin 1). Line CK accomplishes this pulse by sending as output a long pulse on the microcontroller-pin line. R_2 and C_2 introduce a delay, and, once the pulse exceeds that delay, the PL line goes low, and the data loads.

After the PL signal rises, shorter pulses on the microcontroller's I/O port generate pulses at the shift register's clock input, CP, but not at the PL input. The duration of these clock pulses must be long enough to exceed delay R_1C_1 but not R_2C_2. These clock

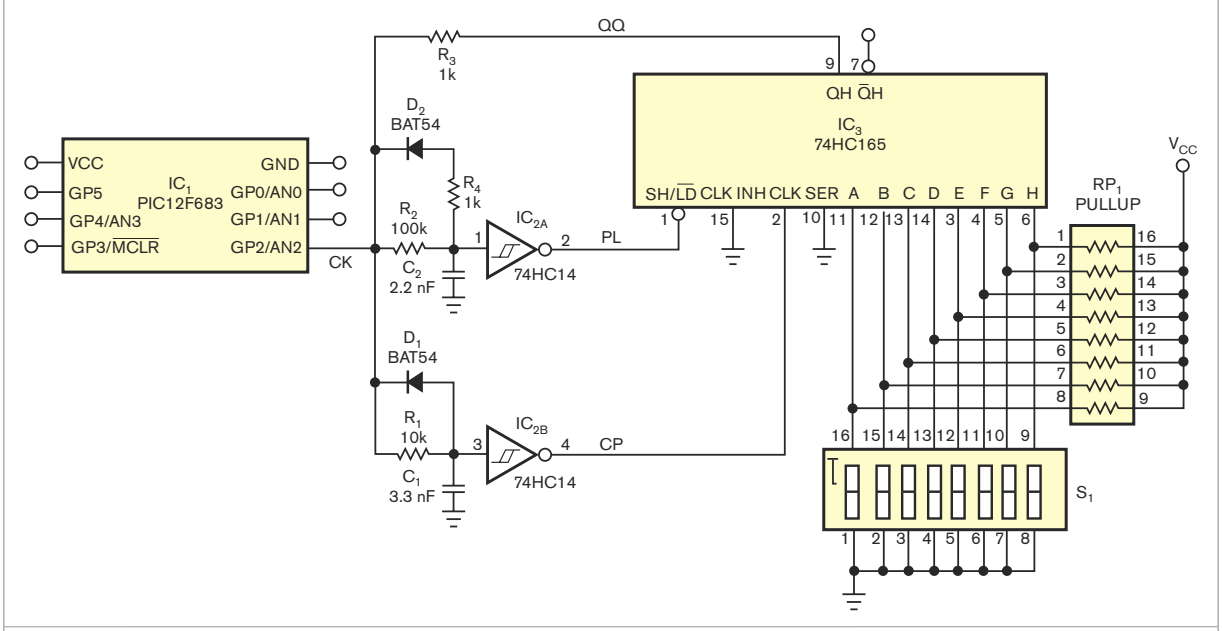

Figure 1 Careful adjustment of the RC time constants allows a microcontroller to input a serial-data stream using a single I/O pin.

Versatile Class D amp offers battery operation and lowest BOM cost

Highly integrated 2 x 15W mono/stereo amplifier simplifies multimedia/TV design

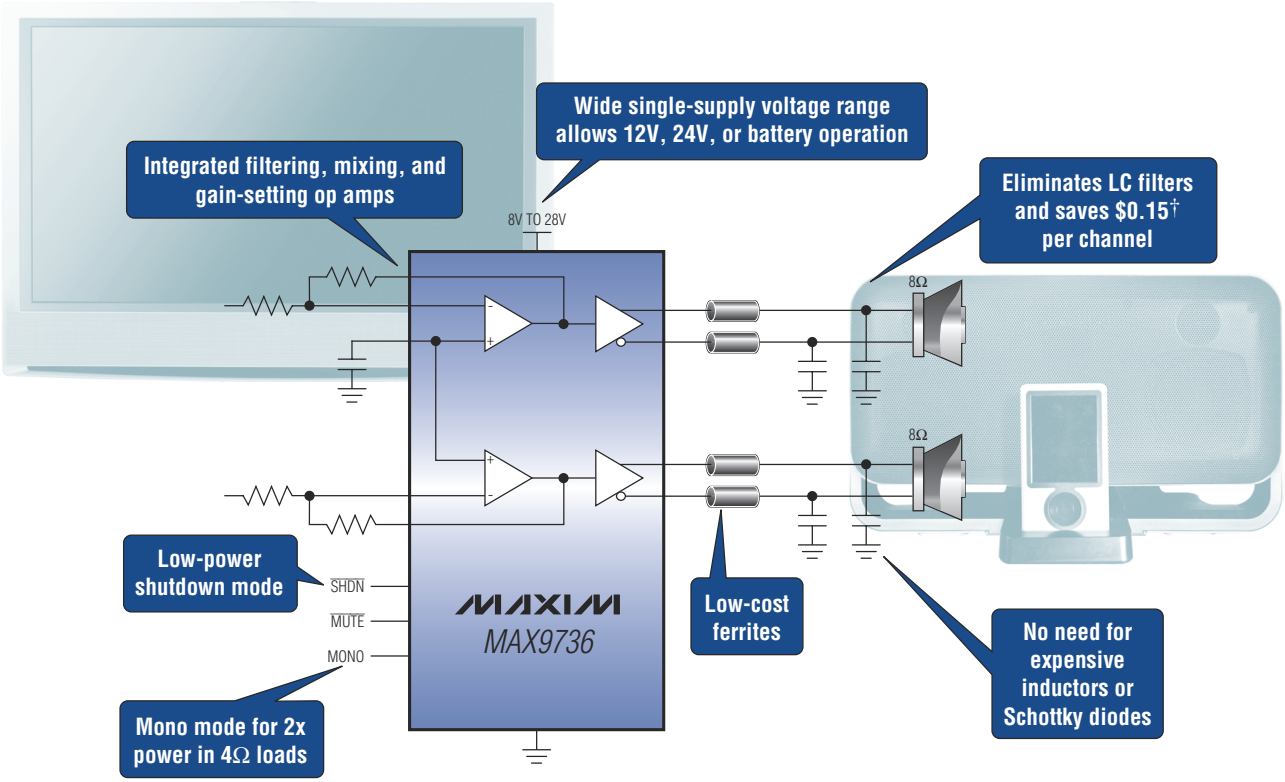

†Based on consumer-electronics volumes. Prices provided are for design guidance and are FOB USA. International prices will differ due to local duties, taxes, and exchange rates. Not all packages are offered in 1k increments, and some may require minimum order quantities.

MAXIM DIRECT
www.maxim-ic.com/shop

AVNET electronics marketing
www.avnet.com

MAXIM
www.maxim-ic.com/MAX9736-info

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

pulses shift the data so that the 8 bits appear in sequence at the shift-register output, QQ.

If the microcontroller's data direction briefly changes to input with high impedance, this shift-register data dominates because of the relative values of R_1, R_2, and R_3, with R_3 being a much lower value. The high-impedance state must exist only for a time less than the R_1C_1 time constant (**Figure 2**). The microcontroller now reads the single bit of data. The action of three differing periods generates three functions: load, clock, and data read. The time the microcontrollers need to change port direction, read the pin data, and reset the pin's direction to output determines the timing. For example, a 1-µsec microcontroller requires 10 µsec.

To avoid spurious CP pulses, this time constant must be less than $0.33R_1C_1$, so R_1C_1 could be 30 µsec and R_2C_2 could be 200 µsec. These settings would allow a complete 8-bit read in about 1 msec. To achieve faster operation, replace the RC delays with a precision retriggerable monostable multivibrator, such as NXP's 74HC123, and logic gates. You can expand the scheme with more shift registers to read dozens of signals.

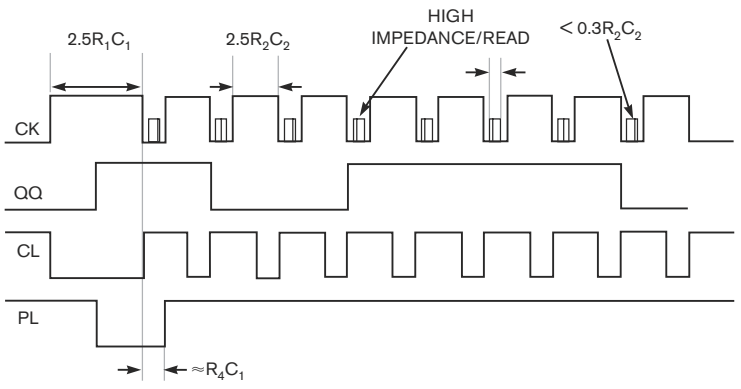

Figure 2 The high-impedance state must exist only for a time less than the R_1C_1 time constant.

Note that internal logic in the 74HC165 shift register prevents the CP signal from shifting data when LD is active. Resistor R_4 ensures the correct sequencing of LD and CP. Diodes D_1 and D_2 quickly discharge the capacitors to "reset" the delay function of R_1C_1 and R_2C_2. EDN

REFERENCE

1 Niven, Rex, "RC lowpass filter expands microcomputer's output port," *EDN*, June 21, 2007, pg 74, www.edn.com/article/CA6451248.

"Sweating the details" redefined by Toshiba.

Toshiba sweats every detail in order to deliver expert, responsive technical service/support for its electronic components. The Toshiba design centers have a team of engineers dedicated to power analysis to ensure optimum power management. Toshiba Virtual Prototyping—our exclusive system for chip/package/system co-design for Custom SoCs—delivers a power-optimized model usually in less than eight weeks. And we employ advanced low-power design techniques like Multi-Vth, Clock Gating and much more. This commitment to delivering unmatched technical expertise is our formula for customer success.

$P_{ff} =$

Which equation below is the most accurate estimate for power consumed by flip flops on a chip? Use the legend at **transform.toshiba.com** to solve the equation and win a hoody sweat shirt (while supplies last).

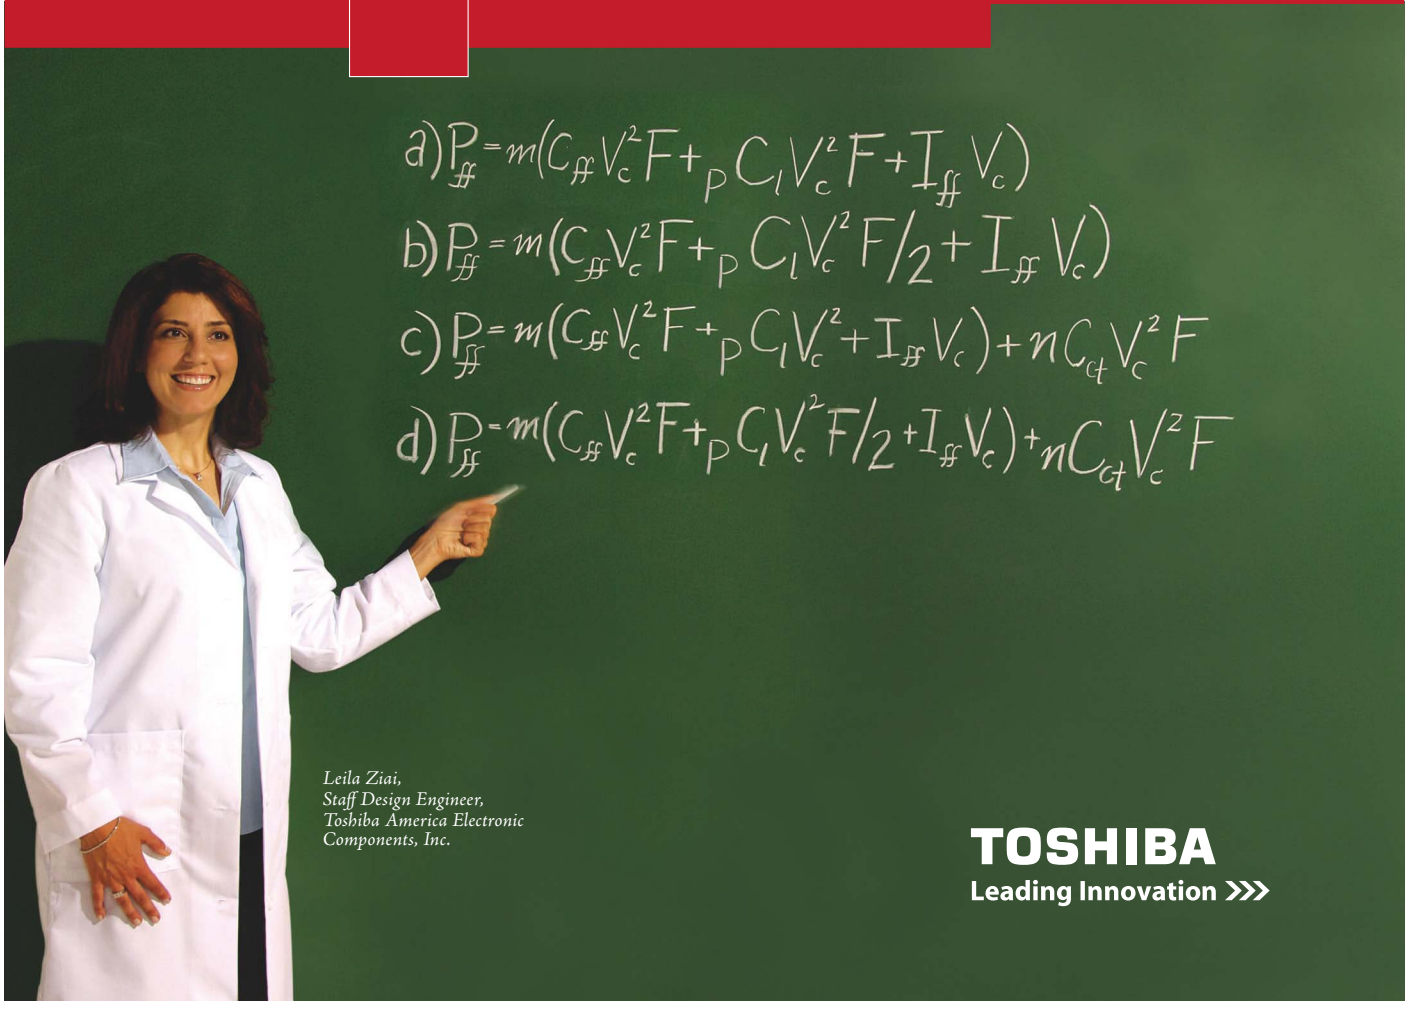

Leila Ziai,
Staff Design Engineer,
Toshiba America Electronic Components, Inc.

TOSHIBA
Leading Innovation >>>

productroundup

AMPLIFIERS, OSCILLATORS, AND MIXERS

Current-sensing amplifier incorporates precision resistors

Precision resistors in the high-side MAX9938 current-sensing amplifiers eliminate the need for external gain-setting resistors. Operating over a 1.6 to 28V input-voltage range suits the amplifier for smartphones, digital cameras, MP3 players, notebook computers, and other battery-operated devices. The product comes in 25, 50, and 100V/V voltage-gain versions, with a −40 to +85°C temperature range. The precision IC consumes 1-µA IQ and provides a ±0.5-mV maximum input-offset voltage and a 0.5% maximum gain error. Available in a 1×1×0.6-mm, four-bump UCSP package, the MAX9938 current-sensing amplifier costs 64 cents (1000).

Maxim Integrated Products, www.maxim-ic.com

Autozero-operation amplifier targets industrial and medical markets

Aiming at industrial and medical markets, the MCP6V01/2/3 autozero operational amplifiers suit battery-powered devices and instrumentation in sensor calibrators for defibrillators and computer-tomography scanners, as well as high-precision temperature sensing and dc-offset correction. The amplifiers have a 2-µV maximum input-offset voltage. Additional features include an 18 to 5.5V operating-voltage range, a 300-µA typical quiescent current, and a rail-to-rail I/O structure. Available in an SOIC-8 package, the MCP6V01 costs $1.26 (10,000).

Microchip Technology, www.microchip.com

Dual ADC driver functions with low noise

A guaranteed ±0.56-dB matched gain and a ±0.1° typical match phase allow the LTC6420-20 dual differential-ADC driver to reduce errors in multichannel systems, such as IQ-demodulation and diversity receivers. Features include 80-dB channel separa-

tion at 100 MHz, 1.8-GHz bandwidth at 3 dB, and 80-mA supply current. The driver provides a 20-dB fixed gain with −84-dBc third-order intermodulation distortion at 100-MHz input frequency. The device has a 2.2-nV/√Hz input-voltage noise and operates from a 3V supply with rail-to-rail swing. The driver operates over a −40 to +85°C temperature range. Available in a 3×4-mm QFN-20-lead plastic package, the LCT6420-20 costs $5.22 (1000).

Linear Technology, www.linear.com

Low-power comparators come in tiny packages

The MAX9060/61/62/63/64/65 family of low-power comparators operates over a 0.9 to 5.5V power-supply range. The family includes five variations of single comparators and one window comparator. Consuming a 350-nA maximum quiescent current at 5.5V, the MAX9060 and MAX9061 feature open-drain outputs, drawing the supply current from an external 0.9 to 5.5V reference voltage. The MAX9062, MAX9063, and MAX9064 offer a push-pull or an open-drain output, consuming 1.1 µA at 5.5V, and include a 0.2V internal reference. The MAX9065 3 to 4.2V window comparator features a push-pull output and consumes 1.3 µA maximum at 5.5V. All of the devices include a 0.3 to 5.5V input-common-mode range independent of the supply voltage. The devices come in lead- and halide-free, ROHS-compliant, 1×1-mm, four-bump UCSP and SOT23-5 packages. The comparator family operates over a −40 to +85°C temperature range, and prices start at 85 cents (1000).

Maxim Integrated Products, www.maxim-ic.com

Instrumentation amplifier extends battery life

The INA333 instrumentation amplifier suits portable-medical-system, handheld-instrumentation, scale, and data-acquisition applications. A switched-capacitor notch filter eliminates chopping noise and provides a 50-nV/√Hz input-voltage noise and a three-operation-amplifier architecture. Features include 75-µA quiescent current, 25-µV offset voltage, 0.1-µV/°C offset drift, and 200-pA input-bias current. The amplifier integrates special filters in series with the inputs to reduce RFI (radio-frequency interference). Available in MSOP-8 and DFN-8 packages, the INA333 instrumentation amplifier costs $1.80.

Texas Instruments, www.ti.com

EDN ADVERTISER INDEX

Company	Page
Actel Corp	6
Agilent Technologies	25
Analog Devices Inc	19, 21
Ansoft Corp	43
Atmel Corp	2
austriamicrosystems AG	62
Bokers Inc	65
Bourns Inc	3
Cirrus Logic Inc	C-4
Coilcraft	11
CUI Inc	46
Digi-Key Corp	1
EMA Design Automation	39
Express PCB	52
International Rectifier Corp	9
LeCroy Corp	C-2
Linear Technology Corp	49, 50
	54, 57
Maxim Integrated Products	59, 61
Melexis Inc	41
Micrel Semiconductor	12
Micro/sys	42

Company	Page
Microchip Technology	17
Mill Max Manufacturing Corp	13
Monolithic Power Systems	33
Mornsun Guangzhou Science & Technolgy Ltd	52
Mouser Electronics	8
National Instruments	4, 32
NewarkInOne	23
Numonyx	45
ON Semiconductor	31
Pico Electronics	44, 64
Prism Sound Ltd	51
Stanford Research Systems Inc	C-3
Sunstone Circuits Inc	26
Tektronix	29
Tern	65
That Corp	51
Toshiba America	63
Vicor	53
Xilinx Inc	14

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

BY THOMAS BLACK • DIGITAL PRODUCTS COMPANY

Good as gold

DANIEL VASCONCELLOS

During my career, I worked with a team on medical imaging. Our system had a 400-MHz ADC board behind its large VME backplane. The board employed early ECL (emitter-coupled-logic) technology and used several Siemens SDA8010 flash-ADC ICs, which were unpopular due to their $200 price tag. Along with an arsenal of other ECL parts, several large coils of RG174 coaxial cable acted as precision delay lines for board timing. The design was unpopular but worked well.

One morning, the purchasing manager mentioned that the SDA8010 had become unavailable for purchase. We shared the bad news with the engineering department, and the engineers advised us that they needed several months to design a replacement board. Our interim solution was to buy enough chips for 16 months of production. Given our modest five-systems-per-month requirements, this purchase commitment was not a serious burden.

A few months later, however, we received a large order from an overseas medical distributor. The increase in production would in a few weeks deplete our inventory of SDA8010s. The engineering department had not yet started the board redesign. To make matters worse, Siemens had dismantled the production facility that produced the chips. The only parts remaining were some raw IC wafers that the company intended to destroy.

In a production meeting, someone joked that we could make our own chips. Facing a production stoppage, I decided to turn that joke into a project. Management directed purchasing to obtain the raw wafers from the supplier. We essentially got them for the cost of shipping.

None of us had any idea how to process a silicon wafer into a production-quality IC. We began by locating a wafer-dicing facility and an IC-packaging house. Our only production spec was a working SDA8010 IC, which we used as an ad hoc blueprint. Within a couple of weeks, we had produced ICs that looked like the real deal. I had created some test software to verify performance, and we eagerly tested each of them. Unfortunately, they all suffered missing bits, horrible linearity, and other serious issues. At this point, we were days away from a full-blown production stoppage.

I borrowed a working chip from a system and one from my "bone pile" of useless parts—traditional-looking DIP ceramic parts in 24-pin packages. With the aid of a heat gun, I gingerly removed the top metal lid from one of the bad parts. I could clearly see the IC's die and bond wires. I did the same heat-gun surgery to the good part, and the only difference I could see was that our packaging vendor used bond wires that appeared to be aluminum, whereas the good parts from Siemens had gold-finished wires.

With fingers crossed, we sent off another group of dice to the packaging house with instructions to copy the gold bond-wire material that the working sample used. A few days later, the first articles arrived. To my amazement, every one of them had performance identical to that of the Siemens parts.

The wafers continued to give high yields, and the surplus of ADC chips was a welcome relief. To add to our good fortune, our system's cost decreased by several hundred dollars because our do-it-yourself chips cost us only $22 after processing. About a year later, engineering delivered a completely new ADC-board design that worked better than our old ECL monster. EDN

Thomas Black is president of Digital Products Company (Folsom, CA). Like Thomas, you can share your Tales from the Cube and receive $200. Contact edn.editor@reedbusiness.com.

www.edn.com/tales

Powerful, Versatile and Affordable...

100 kHz Dynamic Signal Analyzer

- **DC to 100 kHz frequency range**
- **100 kHz real-time bandwidth**
- **Dynamic range**
  - **90 dB (FFT)**
  - **145 dB (swept-sine)**
- **Low-distortion (–80 dBc) source**
- **Up to 32 Mbyte memory**
- **GPIB and RS-232 interfaces**

FFT analyzers starting at $4950

The SR785 Dynamic Signal Analyzer offers state-of-art performance at a fraction of the cost of competitive analyzers. Standard features include swept-sine mode, order tracking, octave analysis, curve fit/synthesis, and arbitrary waveform source.

When compared to the Agilent 35670A, the SR785 comes out on top. It offers twice the frequency range (2 ch.) and 10 times the real-time bandwidth at less than half the price of the 35670A.

The SR785 is ideal for filter design, control systems analysis, noise measurement, audio or acoustical test, and mechanical systems analysis.

Take a close look at the SR785.

Stanford Research Systems
1290-D Reamwood Ave., Sunnyvale, CA 94089 • e-mail: info@thinkSRS.com
Phone (408) 744-9040 • Fax (408) 744-9049 • **www.thinkSRS.com**

Home Theater: It's Our Specialty

Cirrus Logic knows home theater. You might say it's our specialty. Our families of audio DSPs are designed for audio applications—no general purpose processors here. Select from a variety of single or dual-core 32-bit processors that are right for you, with a comprehensive firmware library, along with the industry's widest selection of audio converters that offer the right level of performance and features.

CS49700 FEATURES

- 32-bit dual-core audio DSP
- Ideal for Blu-ray® Disc home theater audio
- Multistandard HD decoding + post processing

Cirrus Logic. We make it easier for you.

NORTH AMERICA
+1 800-625-4084

ASIA PACIFIC
+852 2376-0801

JAPAN
+81 (3) 5226-7757

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

LEARN MORE AT
www.cirrus.com

CIRRUS LOGIC®