Reed Business Information (RBI™)

EDN®

VOICE OF THE ENGINEER

FEB 4

Issue 3/2010

www.edn.com

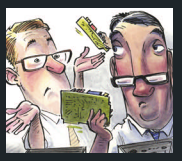

WIRELESS TO GO

Page 26

x86 PROCESSORS: CONTINUED INNOVATION IS A WELCOME CONTRADICTION

Page 19

Tools accurately simulate noise in **MIXED-SIGNAL ASICs**

Page 33

DEBUGGING AND ANALYSIS WITH **SystemVerilog test bench**

Page 38

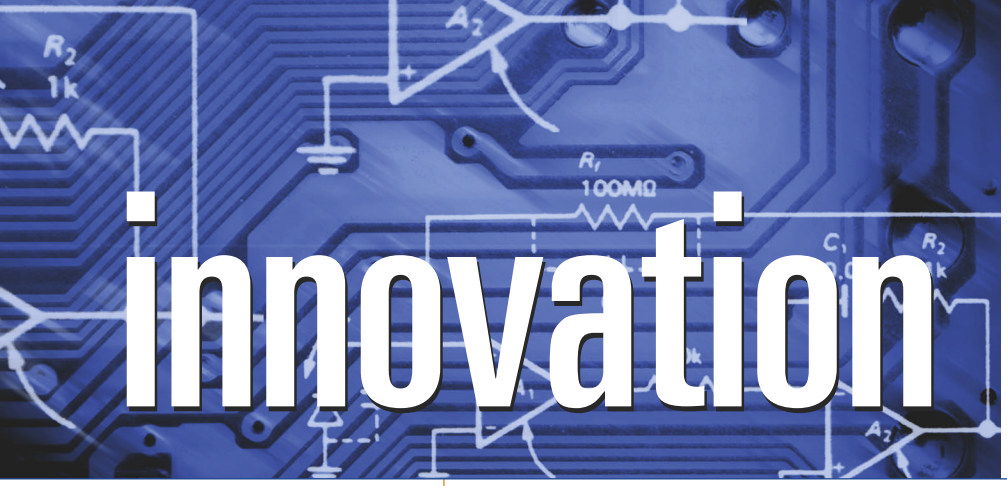

High-Precision Family of Power Meter ICs

CS546X FAMILY DELIVERS UNSURPASSED PERFORMANCE, FEATURES FOR ALL METROLOGY CALCULATIONS

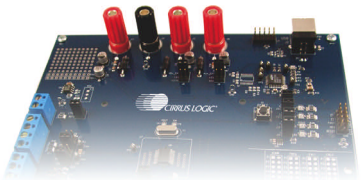

TO JUMP START YOUR DESIGN, RECEIVE A 50% DISCOUNT ON POWER METERING EVALUATION BOARDS!

DISCOUNT APPLIES TO THE FOLLOWING BOARDS: CDB5461AU, CDB5463U, CDB5464U, CDB5466U, CD5467U
HURRY! THIS IS A LIMITED TIME OFFER!

Visit **www.cirrus.com/546Xedn**

Whether it's single- or three-phase, residential or industrial, metering or monitoring, Cirrus Logic has a family of energy measurement products that are ideal for your needs. Cirrus Logic's CS546X energy management family delivers unsurpassed performance in data accuracy and are fully featured to deliver all required metrology calculations. These products are designed to accurately measure instantaneous current and voltage and calculates VRMS, IRMS, instantaneous power and average power (active/reactive/apparent) for high-performance energy measurement applications.

FEATURES

- On-chip functions:
  - Energy data linearity ±0.1% of reading over 1000:1 dynamic range
  - On-chip energy calculations
  - 24-bit fourth-order Delta Sigma A/D converters for precise measurements
  - On-chip 2.5 V reference (tempco 25 ppm/°C typ)
  - On-chip energy direction indicator
  - 24-pin SSOP package

Part	Channels	SPI Slave Interface	Real (Active) Power	IRMS VRMS	Apparent Power (S)	Reactive Power (Q)	Fundamental Power	Temp Sensor	Package
CS5461A	2	√	√	√	√			√	24 SSOP
CS5463	2	√	√	√	√	√	√	√	24 SSOP
CS5464	3	√	√	√	√	√			28 SSOP
CS5466	2		√						24 SSOP
CS5467	4	√	√	√	√	√		√	28 SSOP

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. We make it easier for you.

HOW MANY SUPPLIERS DOES IT TAKE TO MAKE A MASTERPIECE?

Digi-Key® CORPORATION

Authorized distributor for more than 440 of the world's most trusted manufacturers.*

THE INDUSTRY'S BROADEST PRODUCT SELECTION AVAILABLE FOR IMMEDIATE DELIVERY.

1.800.344.4539
WWW.DIGIKEY.COM

*Digi-Key is an authorized distributor for all supplier partners. New products added daily.

NEDA Supplier Authorized Distributors

EDN 2.4.10 contents

Wireless to go

by Rick Nelson, Editor-in-Chief

x86 processors: Continued innovation is a welcome contradiction

by Brian Dipert, Senior Technical Editor

Tools accurately simulate noise in mixed-signal ASICs

by Thierry Masson, Laurent Monge, and A Glascott-Jones, e2v

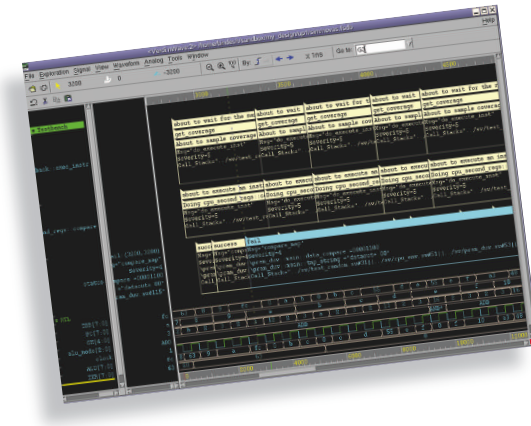

Debugging and analysis with SystemVerilog test bench

by Bindesh Patel and Rex Chen, SpringSoft

pulse

DESIGNIDEAS

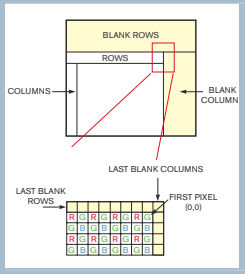

DEPARTMENTS & COLUMNS

EDN online contents www.edn.com

ONLINE ONLY

Check out this Web-exclusive article:

Determine your IC's transient thermal behavior to prevent overheating
Use Spice to determine the thermal time constants of the IC package.
→www.edn.com/article/CA6713633

INNOVATION AWARDS

Mark your calendar. *EDN* will announce the Innovation Award finalists online and in print Feb 18. No campaign speeches, attack ads, or empty promises—just recognition for worthy engineers and innovative products.
→www.edn.com/innovation

FOLLOW THE DOTTED LINES

When visiting EDN.com, you may run across special links that have a dotted underline. Unlike other sites that use such links to trick you into viewing advertising, we're using them to lead you to our Hot Topics pages, where you'll find copious links to information from not only *EDN* but also the entire electronics-industry Web. Follow a dotted link, or visit the Hot Topics home page.
→www.edn.com/hottopics

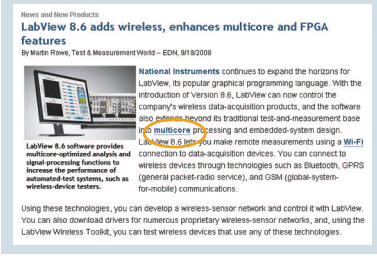
News and New Products

LabView 8.6 adds wireless, enhances multicore and FPGA features

By Martin Rowe, Test & Measurement World -- EDN, 9/18/2008

LabView 8.6 software provides multicore-optimized analysis and signal-processing functions to increase the performance of automated-test systems, such as wireless-device testers.

National Instruments continues to expand the horizons for LabView, its popular graphical programming language. With the introduction of Version 8.6, LabView can now control the company's wireless data-acquisition products, and the software also extends beyond its traditional test-and-measurement base into multicore processing and embedded-system design. LabView 8.6 lets you make remote measurements using a Wi-Fi connection to data-acquisition devices. You can connect to wireless devices through technologies such as Bluetooth, GPRS (general packet-radio service), and GSM (global-system-for-mobile) communications.

Using these technologies, you can develop a wireless-sensor network and control it with LabView. You can also download drivers for numerous proprietary wireless-sensor networks, and, using the LabView Wireless Toolkit, you can test wireless devices that use any of these technologies.

EDN® (ISSN#0012-7515), (GST#123397457) is published semimonthly, 24 times per year, by Reed Business Information, 8878 Barrons Blvd, Highlands Ranch, CO 80129-2345. Reed Business Information, a division of Reed Elsevier Inc, is located at 360 Park Avenue South, New York, NY 10010. John Poulin, CEO/CFO, RBI-US; Jeff DeBalko, President, Business Media; Jane Volland, Vice President, Finance. Periodicals postage paid at Littleton, CO 80126 and additional mailing offices. Circulation records are maintained at Reed Business Information, 8878 S Barrons Blvd, Highlands Ranch, CO 80129-2345. Telephone (303) 470-4445. POSTMASTER: Send address changes to EDN®, PO Box 7500, Highlands Ranch, CO 80163-7500. EDN® copyright 2010 by Reed Elsevier Inc. Rates for nonqualified subscriptions, including all issues: US, $179.99 one year; Canada, $229.99 one year (includes 7% GST, GST#123397457); Mexico, $229.99 one year; air expedited, $399.99 one year. Except for special issues where price changes are indicated, single copies are available for $10 US and $15 foreign. Publications Agreement No. 40685520. Return undeliverable Canadian addresses to: RCS International, Box 697 STN A, Windsor Ontario N9A 6N4. E-mail: Subsmail@ReedBusiness.com. Please address all subscription mail to EDN®, 8878 S Barrons Blvd, Highlands Ranch, CO 80129-2345. EDN® is a registered trademark of Reed Elsevier Properties Inc, used under license. A Reed Business Information Publication/Volume 55, Number 3 (Printed in USA).

BY RON WILSON, EXECUTIVE EDITOR

Can C beat RTL?

With the appearance of higher speeds and more DSP macrocells in low-cost FPGAs, more and more design teams are seeing the configurable chips not as glue but as a way to accelerate the inner loops of numerical algorithms, either in conjunction with or in place of the traditional DSP chip. There's a problem, however. You code for a DSP chip in C, and you implement it using a conventional software tool chain with familiar software debugging tools. You configure an FPGA starting in Verilog or VHDL (very-high-speed-integrated-circuit-hardware-description language)—superficially similar to C but in practice profoundly different—and you implement it using a hardware design flow. The two approaches require different skills.

Enter ESL (electronic-system-level) tools. An ESL synthesis tool lets you write your code in C, automatically synthesize RTL (register-transfer-level) logic from the C, and then feed the RTL into your FPGA flow. In reality, such tools meet with skepticism because people suspect them of poor quality of results, unreliability, and other vices. Is that assessment fair, though? BDTI (Berkeley Design Technology Inc) wanted to find out.

The company last month released the first results of its certification program for high-level synthesis tools. The first evaluation covers AutoESL's AutoPilot and Synfora's Pico. The bottom line in BDTI's findings was that both tools produced results in a reasonable amount of time and that both performed much better than software on a DSP chip. The tools were comparable in density and performance with hand-coded RTL. The fine print reveals a wealth of information below that level, however.

It is no longer prudent for design teams working with computationally oriented cores to ignore ESL synthesis tools.

First, BDTI uses a method that represents, as always, a compromise between realism and practicality. BDTI's initial benchmark is a fully functional optical-flow design comprising a three-ring binder and a DVD, which in turn contain a text description of the algorithm, the algorithm in about 600 lines of ANSI C, and a Xilinx reference design. BDTI turns the kit over to the ESL vendor, which tunes the C code for the tool and produces a design. BDTI engineers then independently repeat the process. The optical-flow core attempts to achieve maximum throughput using all the resources available in the Spartan IIIA FPGA.

Unsurprisingly, both ESL vendors produced designs with about 40 times the throughput of the best BDTI engineers could do on a Texas Instruments DM6437 DSP chip. Surprisingly, in a separate test with a smaller design, results from the ESL flows required essentially the same die area as a hand-coded RTL kernel.

The amount of work to do the FPGA design, from C to programming file, was similar to the work to program the DSP, according to Jeff Bier, BDTI president. Significant differences between the two tasks emerged, however. Optimizing the C code for one of the ESL tools caused the code to balloon from the original BDTI-supplied 559 lines to 1604 lines of C. Bier says that the work in the optimization was somewhat less than optimizing the code for the DSP chip. "It turned out that the DSP had a serious memory bottleneck that we had to code around," he explains. The synthesis tool then generated more than 38,000 lines of Verilog from the optimized C.

BDTI engineers, experienced DSP programmers, could handle the entire flow for the TI chip. A huge pile of Verilog and a stack of Xilinx tools stumped them, though. They ended up calling in an RTL-logic expert to shepherd the RTL through the Xilinx tool chain, debug it, and produce the configured FPGA.

Clearly, it is no longer prudent for design teams working with computationally oriented cores to ignore ESL synthesis tools. The analogy to the days when RTL synthesis was just beginning to displace schematic capture and Karnaugh maps or, for that matter, when programmers began to write embedded software in C instead of assembler is irresistible. Stand by for change. EDN

Contact me at ronald.wilson@reed business.com.

EDN

PRESIDENT, BUSINESS MEDIA, REED BUSINESS INFORMATION
Jeff DeBalko, jeff.debalko@reedbusiness.com
1-646-746-6573

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

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

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

EXECUTIVE EDITOR
Ron Wilson, 1-510-744-1263;
ronald.wilson@reedbusiness.com

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

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

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

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

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

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

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

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

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

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

CONSULTING EDITOR
Jim Williams, Staff Scientist,
Linear Technology

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

COLUMNISTS
Howard Johnson, PhD, Signal Consulting
Bonnie Baker, Texas Instruments
Pallab Chatterjee, SiliconMap

PRODUCTION
Dorothy Buchholz,
Group Production Director
1-781-734-8329
Joshua S Levin-Epstein,
Production Manager
1-781-734-8333; fax: 1-781-734-8096

EDN EUROPE
Graham Prophet, Editor, Reed Publishing
+44 118 935 1650;
gprophet@reedbusiness.com

EDN ASIA
Luke Rattigan, Chief Executive Officer
luke.rattigan@rbi-asia.com
Kirtimaya Varma, Editor-in-Chief
kirti.varma@rbi-asia.com

EDN CHINA
William Zhang,
Publisher and Editorial Director
wmzhang@rbichina.com.cn
Jeff Lu, Executive Editor
jefflu@rbichina.com.cn

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

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

Reed Electronics Group

THE EMBEDDED DEVELOPER

Microsoft

Windows Embedded

BUILT WITH WINDOWS® 7 TECHNOLOGIES

VOL. 5

A MALFUNCTION IN THE SYSTEM COULD COST THE PLANT MILLIONS...

THEY'RE COUNTING ON ME TO DELIVER.

THE DEVICE HAS TO WORK PERFECTLY. TO THE MICROSECOND.

AND HAVE THE CONNECTIVITY TO TRACK PERFORMANCE IN REAL TIME.

WINDOWS® EMBEDDED OFFERS A HIGHLY RELIABLE PLATFORM, WITH THE LEVEL OF PERFORMANCE YOU NEED TO HELP DELIVER CONNECTED DEVICES THAT STAND OUT.

WHICH WINDOWS EMBEDDED PLATFORM CAN HELP YOU DELIVER STANDOUT DEVICES?
FIND OUT AT WINDOWSEMBEDDED.COM/DEVICESTORIES

NI pursues robot revolution

National Instruments is moving aggressively to support the "robot revolution" with its announcement of LabView Robotics 2009, which provides a standard development platform for designing robotic and autonomous control systems. With the new release, the company hopes to capitalize on what it forecasts is a burgeoning market for robots, including the fact that one-third of US military vehicles must be autonomous by 2015, representing a $52 billion market. Further, the educational-robotics market will reach $1.96 billion by 2014. And, by 2012, service robots for professional use will increase 78%, service robots for household use will increase 109%, and robots for entertainment and leisure will increase 239%.

With LabView for Robots 2009, NI aims to overcome the challenges that currently limit robot usage in our day-to-day lives. Those challenges span software design involving modeling, simulation, and algorithm design; embedded-system design involving analog and digital I/O, protocols, motion control, concurrency, and determinism; and connectivity dealing with actuators, sensors, and motors. The company positions LabView as a key tool in meeting that challenge and has received support for that position from David Barrett, PhD, director of SCOPE (Senior Capstone Program in Engineering) at the Franklin W Olin College of Engineering (Needham, MA) and former vice president of engineering at iRobot. Barrett, who delivered a keynote address at last summer's NIWeek, says that the robotics industry requires an industrial-grade, hardened, richly supported software-development system, noting that LabView may well fit that description.

The LabView for Robots 2009 release takes aim at making sure LabView fits Barrett's description, delivering an extensive robotics library with connectivity to standard robotic sensors and actuators, foundational algorithms for intelligent operations and perception, and motion functions for robots and autonomous vehicles. The release ties together LabView's Real-Time, FPGA, Vision, Control Design and Simulation, SoftMotion, Statechart, and Mathscript RT tools as well as the PID (proportional-integral-derivative) Toolkit. It adds robotics IP (intellectual property) for sensing, connectivity, protocols, path planning, obstacle avoidance, and steering. Target hardware platforms include NI's CompactRIO, Single-Board RIO, and a new LabView robotics starter kit.

NI is collaborating with Cogmation on system simulation, Energid on kinematics, Hokuyo on LIDAR (laser-imaging-detection-and-radar) sensors, iRobot on hardware integration, Microsoft on system simulation, MobileRobots on hardware integration, MaxonMotors on smart-motor connectivity, Pitsco on OEM and academic starter kits, Skilligent on vision software, TORC Technologies on the JAUS (Joint Architecture for Unmanned Systems) protocol, and Velodyne on LIDAR sensors.

The base price for a robotics starter kit, including a 180-day evaluation of the LabView robotics-software bundle, is $1999, or $1599 for academic customers.**—by Rick Nelson**

▷**National Instruments**, www.ni.com.

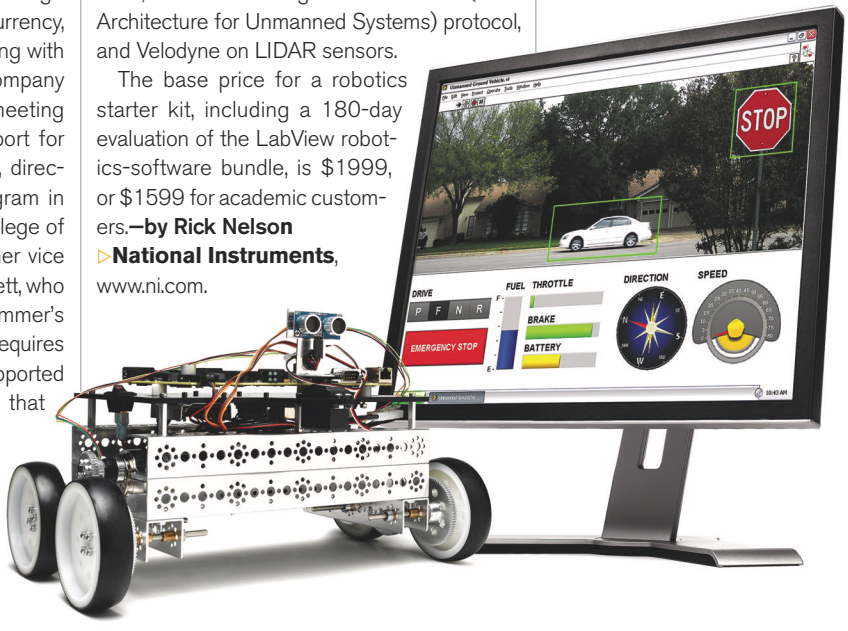

NI offers a robotics starter kit that includes the DaNI robot and a robotics-software bundle.

FEEDBACK LOOP

"Been there, done that!! And dn does it feel good when you kick its behind."**

—Engineer Jeff Droll, in *EDN*'s Feedback Loop, at www.edn.com/article/CA6625453. Add your comments.

3-D medical printer to print body parts

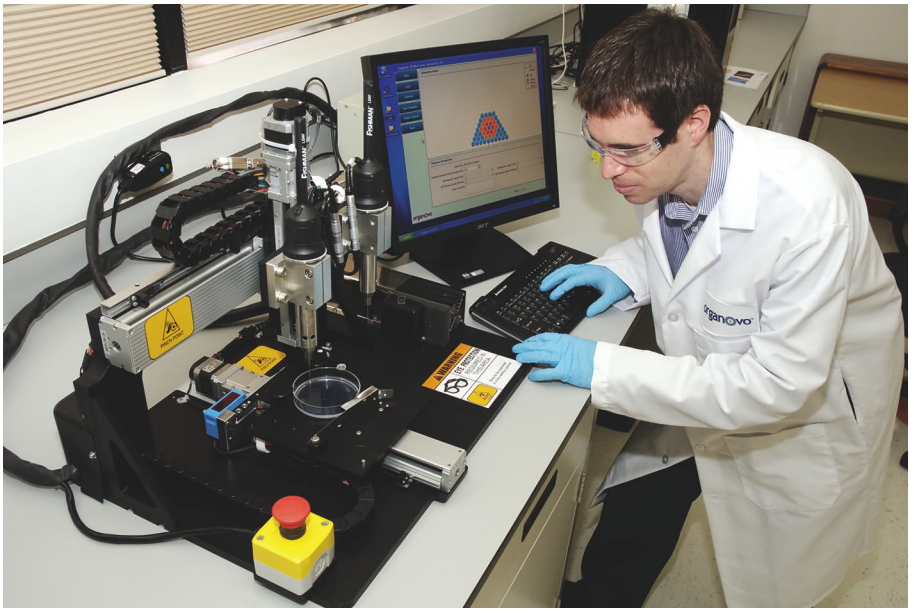

Researchers hope to custom-build tissues and organs for transplant, using the patient's own cells and a 3-D medical printer (courtesy Invetech).

At the December Printed Electronics USA conference (www.idtechex.com/printed electronicsusa09), vendors showcased printable batteries, OLED (organic light-emitting-diode) displays, and diabetes-treatment and drug dispensers. One of the most notable technologies on display comes from biosciences company Organovo, which just took delivery of the 3-D medical printer that it designed with Invetech, a design and contract-manufacturing company. The technology "prints" artificial blood vessels for transplant with the goal of creating an arterial graft for use in coronary-bypass surgery. The long-term goal, according to an article in *Information Week*, is to solve problems in medical therapy that other approaches cannot solve (**Reference 1**).

The printer allows researchers to precisely place cells into a 3-D pattern. The technology works by using a robot to lay down cells in precise positions in three dimensions, accurate to within 20 microns. It's similar to the way a laser printer prints by putting solid particles in place. The 3-D medical printer puts objects onto 2-D layers, one on top of the other. The particles comprise stem cells, formed into tiny spheres and cylinders.

"The stem cells are available for research purposes from companies including Life Technologies and Invitrogen," according to the article. Cells, such as bone marrow or fatty adipose tissues, will come from the patient, from which researchers can harvest stem cells. Because the cells come from the patient, there's no risk of rejection. The arterial segments are 5 to 20 cm long, with an interior diameter of 0.5 to 5 mm. Researchers can build arteries with larger interior diameters using Teflon or Dacron, but the smaller diameters clot with synthetic materials."

Like Organovo's equipment, other 3-D printers typically use plaster, cornstarch, or resins to create objects and most often find use in rapid prototyping for the footwear, jewelry, industrial-design, architecture, automotive, aerospace, dental, and medical industries. Organovo is not alone in finding medical applications for 3-D printing; researchers at the University of Tokyo Hospital (www.h.u-tokyo.ac.jp) and venture company Next 21 are using 3-D printers to create artificial bones for reconstructive surgery.

Oganovo is building blood vessels at first, including an arterial structure for use in coronary-bypass surgery. That procedure involves creating an object with endothelium cells on the inside, smooth muscle in the middle, and an exterior layer of fibroblasts, which are similar to skin cells. Printed blood vessels should be in clinical trials in three to five years, according to the article.**—by Margery Conner**

▷**Invetech**, www.invetech.com.

▷**Organovo**, www.organovo.com.

REFERENCE

Wagner, Mitch, "3D Printer Builds Artificial Blood Vessels," *Information Week*, Dec 23, 2009, www.informationweek.com/news/healthcare/patient/showArticle.jhtml;jsessionid=VVBNRPNXYDM2LQE1GHOSKH4ATMY32JVN?articleID=222003031.

DILBERT By Scott Adams

3.3V mixer downconverts RF with good linearity

Linear Technology Corp's new LTC5541 downconverting mixer operates at frequencies of 1.3 to 2.3 GHz, and the company plans to later introduce pin-compatible parts that will operate at 0.6 to 4 GHz. The device operates from a 3.3V supply and consumes 630 mW of power during operation and 500 μA in shutdown mode.

You can power the IF (intermediate-frequency) amplifier with 5V to improve 1-dB gain compression, and using a local oscillator with a frequency of 1.4 to 2 GHz improves performance. The mixer can switch between two local-oscillator inputs, which function with a drive level of 0 dBm (decibels referenced to milliwatts); both stay at 50Ω impedance during shutdown so that PLL (phase-locked-loop) circuits remain locked. The part achieves an IIP3 (input third-order-intercept point) of 26 dBm, a noise figure of 9.6 dB at 1.95 GHz, and a conversion gain of 7.8 dB at 1.95 GHz.

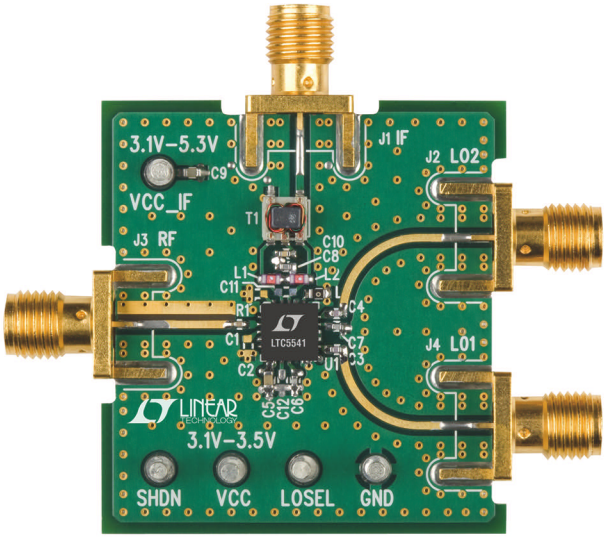

An evaluation board for the LTC5541 allows you to assess its performance in your system.

The mixer finds use in both the main- and the predistortion-receiver paths. Typical applications include LTE (long-term evolution) and CDMA (code-division-multiple-access) wireless base stations and WiMax (worldwide-interoperability-for-microwave-access), WCDMA (wideband-CDMA), UMTS (universal-mobile-telecommunications-system), and EDGE (enhanced-data-rates-for-global-system-for-mobile-communication evolution) systems. You can also use it for backhaul wireless service between base stations and public-safety and military-communications systems.

The LTC5541 operates at a −40 to +85°C junction temperature; comes in a 5×5-mm, 20-pin QFN package; and sells for $6.50 (1000). Production has begun.**—by Paul Rako**

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

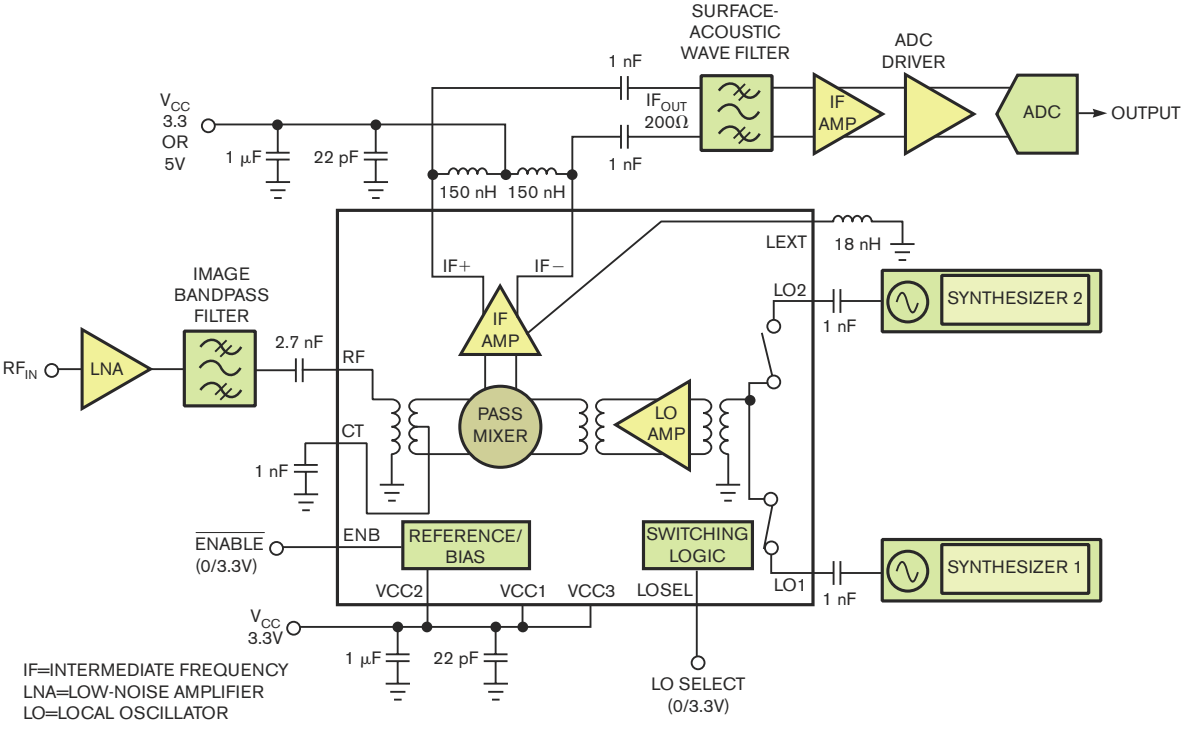

The Linear Technology LTC5541 mixer features good linearity and low noise despite the fact that it operates off 3.3V.

Affordable scope family now includes four-channel and 40-MHz models

LeCroy Corp has announced the expansion of its affordable WaveAce oscilloscope line to include four-channel models with bandwidths of 60, 100, 200, and 300 MHz and is adding a two-channel, 40-MHz-bandwidth model. The four-channel models provide waveform memory of 10,000 points/channel, or 20,000 points/channel when you interleave pairs of channels. They acquire as many as 2G samples/sec/channel with interleaving. The 40-MHz model acquires 4000 points/channel and as many as 500M samples/sec/channel.

All models offer color displays, extensive measurement capabilities, and advanced triggering to improve troubleshooting and shorten debugging time. The streamlined user interface supports 11 languages. With USB (Universal Serial Bus) host and device ports and a LAN connection, the instruments allow remote control and let you easily save data to a memory stick, PC, or printer.

The 20,000-point memory enables the full-sample-rate capture of waveforms whose length is as much as three times that of waveforms you acquire with similarly priced competitive units. With 32 built-in automated measurements, including advanced modes for skew, phase, and channel-to-channel edge timing, the scopes simplify waveform analysis. Such features as pass/fail testing, user-definable digital filters, and a waveform-sequence recorder simplify and shorten debugging.

> Buttons are backlit so you don't have to guess at the scope's operating mode.

A new remote-control command set makes the four-channel models even more versatile than their predecessors. This command set provides access to all of the controls, functions, and measurements from a remote PC through the rear-panel LAN or USB-host ports. These connections also allow you to view the waveform display and access a virtual front panel from a remote PC.

Internal storage holds as many as 20 waveforms and 20 setups. To store even more information, you can attach a USB memory stick to the front-panel USB-device port. You can access all important controls and menus from the front panel with a single button press. Pressing the associated knob resets any trace position or offset. Pressing the volts/division knob switches between fixed and variable gain; pressing the time/division knob switches among zoom modes. Buttons on the front panel that open and close menus or switch among modes are backlit so you don't have to guess at the scope's operating mode.

The 40-MHz unit sells for $695. Prices for the new four-channel models range from $1690 for 60-MHz bandwidth to $2790 for 300 MHz.

—by Dan Strassberg

▷**LeCroy**, www.lecroy.com.

Affordable, easy-to-use WaveAce scopes now include four-channel models with bandwidths of 60 to 300 MHz and a two-channel, 40-MHz unit with a US list price of $695.

BLACKFIN TARGETS LOWER BOM COSTS

02.04.10

Analog Devices' new Blackfin BF50x processors operate as fast as 400 MHz, have dual MAC (multiply/accumulate) units, and sell for as low as $4.50 (10,000). The BF504F adds 4 Mbytes of stacked, parallel flash memory for prices as low as $6.50, and the BF506F adds dual 12-bit, 2M-sample/second SAR (successive-approximation-register) ADCs and stacked flash for prices as low as $10.60 (10,000).

All three processors include dual three-phase PWMs (pulse-width modulators), SPIs (serial-peripheral interfaces), SPORTS (space-power radio-transmission systems), and UARTs (universal asynchronous transmitters/receivers) with flow control. They also feature CAN (controller-area-network) interfaces and removable storage interfaces that support SDIO (secure digital input/output), eMMC (embedded multimedia card), and CE-ATA (consumer-electronics advanced-technology attachment).

The BF504 and BF504F are available in 88-pin, 12×12-mm, 0.5-mm LFCSPs, and the BF506F comes in 120-lead, 14×14-mm, 0.4-mm LQFPs. The $199 BF50x Ez-Kit Lite and the $150 ICE (in-circuit emulator)-100B Blackfin provide support. For more, go to www.edn.com/article/CA6715679.

–by Robert Cravotta

▶Analog Devices, www.analog.com.

Tighter binning yields more consistent HB LEDs for indoor apps

Vendors of HB LEDs (high-brightness light-emitting-diodes) make white HB LEDs by dosing a blue LED with a blob of phosphor that emits white light when blue light illuminates it. If cost were no object, a manufacturing process could precisely control the LED chip/phosphor combination to get a consistent white-light color temperature, but that level of control is prohibitively expensive. However, HB-LED vendors do provide LEDs that fall within certain color ranges, or bins, in a binning process.

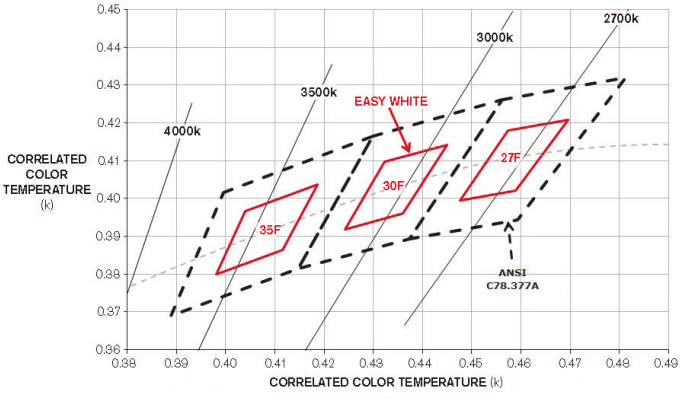

Cree's EasyWhite bins come in 3500, 3000, and 2700K color temperatures and are 75% smaller than the ANSI C78.377-standard color regions.

If your end fixture uses multiple LEDs, a common design technique for outdoor lights such as streetlights, the lighting will be consistent as the LEDs' color temperature averages out. However, some indoor applications lack the space to allow for multiple LEDs. Applications such as track lighting, spotlights, and can lights may use one HB LED, and the variation in light color temperature can be visible—as well as irritating.

Aiming at such applications, Cree recently announced EasyWhite bins in 3500, 3000, and 2700K color temperatures that are 75% smaller than standard color regions. Cree then uses EasyWhite LED chips in its multichip XLamp MC-E LEDs to obtain a more consistent white color temperature in a single LED package, which can serve as a replacement for 20 to 35W halogen light bulbs in indoor applications such as accent, track, and pendant lighting. The XLamp MC-E EasyWhite LED, at 3000K CCT (correlated color temperature), can produce as many as 560 lumens at 700 mA. EasyWhite LEDs cost about 15% more than Cree's standard-bin LEDs.

—by Margery Conner

▷**Cree**, www.cree.com.

OPEN-SILICON ADDS SILICON LOGIC ENGINEERING

SOC (system-on-chip) design now has two distinct flows. One flow creates a platform design: a completely new SOC to serve a new application. The other flow slightly modifies the platform to create a derivative design. The derivative flow exploits the platform as much as possible, sometimes just replacing one block in the physical design without changing the rest of the chip at all.

The platform team contains the application specialists, architects, verification gurus, and new-design troubleshooters. The derivative team, perhaps unfairly, can look like a group of routine implementers–that is, they can look dispensable. So the derivative teams are vulnerable to outsourcing.

That's where Open-Silicon enters the picture. The company's founders originally conceived it as a design shop for turnkey back-end-ASIC design, with management and technical customer relations in the United States and a sophisticated back-end process in India. That model was highly successful in good times, and Open-Silicon built a strong customer base among fabless start-ups. In hard times, however, start-ups aren't the best folks to have in your accounts-receivable ledger. Naveed Sherwani, company president and chief executive officer, set out to build a base among larger companies, just as the larger companies started looking to outsource their derivative-design efforts.

The result is that Open-Silicon now finds itself doing significant business in executing derivative designs. This transition required additional expertise. "To do derivative chips successfully, you need to have front-end design expertise," Sherwani says. The derivative team has to know the platform well enough to grasp the implications of the changes, which can go two or three levels into the platform.

Initially, Open-Silicon met this need by gathering a cloud of specialist design partners to assist them. Sherwani also felt that the company needed in-house expertise at the architecture level and RTL (register-transfer level). "It really helps to be involved in the original platform design or at least to have the experience to understand it," he says.

Eventually, Open-Silicon bought Silicon Logic Engineering, a full-range chip-design team. Now, relying on its newly acquired architecture, RTL, and implementation team, Open-Silicon can engage fully with a platform design, either during its development or later on, and produce derivatives to spec.

For more on this development, go to www.edn.com/100204pa.

–by Ron Wilson

▶Open-Silicon, www.open-silicon.com.

▶Silicon Logic Engineering, www.siliconlogic.com.

02.04.10

Analog Devices: enabling the designs that make a difference in people's lives.

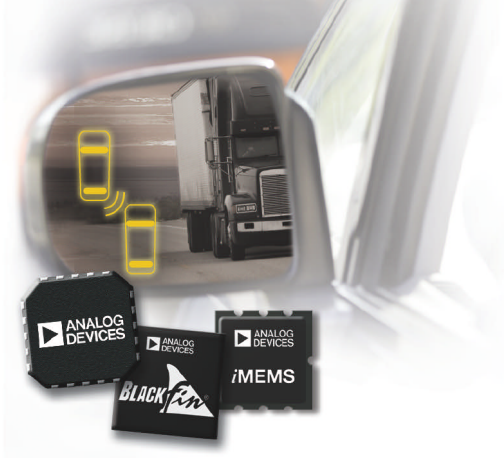

Blind spot sensors and airbag deployment are just two of many critical safety systems enabled by ADI technology.

Each generation of automobiles is safer and more responsive than the last. Enabled by advanced sensor, measurement, and monitoring technologies, today's automotive electronics gives drivers unprecedented control over their vehicles, making the roads safer for all. At ADI, years of working closely with designers of these electronics systems have yielded signal processing solutions that made adaptive cruise control, lane departure warning systems, stability control, and assisted braking systems standard equipment on many vehicles. When demanding applications require innovative signal processing, automotive engineers turn to ADI. Learn more at *www.analog.com/makeadifference*.

MAKEADIFFERENCE

www.analog.com

ANALOG DEVICES

BY HOWARD JOHNSON, PhD

Way too cool

In mid-2005, the Hilton family of hotels mandated installation of the official Hilton Family Clock Radio in 250,000 rooms. This wonderful labor-saving device was designed to rectify two long-standing problems at the hotel: the tendency of guests to foul up traditional clock-radio settings and the burden imposed on hotel staff of checking and maintaining the correct time.

Hilton's clock designers overcame the first difficulty with four buttons atop the unit. These buttons are preset to radio stations that local hotel management selects. With this design, you can wake to one of the hotel's stations, the buzzer, or nothing. No longer can pesky hotel guests leave the machine inadvertently dialed to an alternative radio station or white noise (my personal favorite wake-up sound).

Regarding the time-maintenance burden, careful studies of clock inaccuracy must have implicated those same pesky hotel guests. I can just imagine the hotel's management, red-faced and furious, demanding that guests be forever barred from messing up the otherwise perfectly good timekeeping technology provided to them. In response, the designers removed all user-accessible time-adjustment controls from the new clock.

There is no little "time" switch on the back, no pinhole for a paper clip, and no combination of buttons that can adjust the time. Only a certified hotel maintenance engineer can adjust it. The engineer must physically access the clock, remove a Phillips screw near the top on the back, pop off the lid, and depress the "time" button inside to make the adjustment. It's a slow process.

The "no user-accessible controls" strategy requires that the clock be made superaccurate; otherwise, the extreme maintenance hassle of making even occasional adjustments would invalidate the new clock's entire *raison d'être*. Toward that end, the developers equipped the clock with a precise internal time reference and a computer-controlled feature that automatically adjusts for daylight-saving time. The clock is so smart and so accurate that, in the opinion of its designers, it never requires adjustment. Initial user reviews said things like "Best clock ever made by man" and "Nice features, easy to read."

Fast-forward to Sunday, Oct 30, 2005. That day, hotel guests all across Arizona missed their flights. Arizona does not observe daylight-saving time, yet the Hilton family clocks in Arizona hotels did. Early Sunday morning, all the new clocks "fell back" precisely one hour, according to their programming, and all the guests woke an hour late. Needless to say, panic ensued. A month later, when I visited a Hilton family hotel in Tucson, AZ, the hotel building engineer still wasn't finished changing all the clocks. He said, "We'll be done in about six months." Official word from Hilton headquarters came down that the problem affected only some states, and those states would still be required to use the new clocks.

Fast-forward again to March 11, 2007. During that year, Congress changed the daylight-saving-time law in the United States to a different day than previously. Most Hilton family clocks in the United States now require major engineering maintenance four times a year.

Am I making a point? Anyone who ships product to a vast number of installations covering a wide range of application scenarios can tell you what I am talking about. Unanticipated effects always crop up in the field.

This winter, we as a nation found out that incandescent traffic lights emit plenty of infrared heat through the colored lens—enough to melt drifting snow that otherwise would accumulate, obscuring the lens. Wonderful new, energy-efficient LED traffic lights do not emit that much heat (**Reference 1**).

Bad clocks just make people late. LED traffic lights clog with snow, become indiscernible, and cause fatal traffic accidents.

Rapid adoption of large-scale societal change is a bad idea. For example, when our government talks about nationwide health care, I think we should try it for a while in one state first to see whether it works. How about, say, Massachusetts? EDN

REFERENCE

1 Saulny, Susan, "LED Signals Seen as Potential Hazard," *The New York Times*, Jan 1, 2010, www.nytimes.com/2010/01/02/us/02lights.html?src=twt&twt=nytimes.

PROVEN
25
YEAR
OPERATING
LIFE

Don't let the cost and inconvenience of hard-wired AC power send you up a tree. Instead, be as remote as you want to be with the ONLY lithium cells to offer *PROVEN* 25-year service life without recharging or replacing the battery. When it comes to wireless power solutions, Tadiran is taking innovation to extremes.

TADIRAN BATTERIES

Tadiran Batteries
2001 Marcus Ave.
Suite 125E
Lake Success,
NY 11042
1-800-537-1368
516-621-4980

www.tadiranbat.com

CONTROL ENGINEERING | design news | EDN | Test & Measurement World

MECHATRONICS IN DESIGN

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

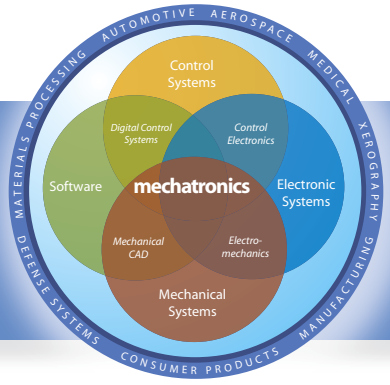

Looking beyond STEM

Engineers must lead in innovation, but STEM is only part of the solution.

As the year and decade end, I see more and more emphasis from K–12 education and private/public funding agencies on STEM (science, technology, engineering, and mathematics) activities. STEM initiatives seem to be the focal point now for addressing the innovation crisis in the United States. While I agree this emphasis is essential, I believe the focus is too narrow and exclusionary. Let me explain.

We are in competition with the entire world, and the innovation crisis we find ourselves in has been fueled by a crisis in education, not only K–12 education but university education, as well. Most students focus on facts, tests, and grades and fail to understand concepts and processes. The inability of graduating students to integrate all they have learned in the solution of a real-world problem, at any level, is a failure. The main goal of education at all levels must be to create critical-thinking problem-solvers and teach them that society's problems are complex and multidisciplinary.

Jon Jensen, associate dean at Marquette University and K–12 outreach director, makes the point that we are a nation of assessors and somewhat obsessed with comparisons, particularly with other nations. Our youth seem to lag behind other countries, most notably in the STEM areas. He says we actually do better as a country than most think, and while we do have a great educational system, we still have a long way to go.

If I am a student with no interest or particular talent in STEM activities, I feel irrelevant to solving the innovation crisis. If I am a teacher in a non-STEM area, I also feel irrelevant. Students as early as the fourth grade are segregated into a college-bound STEM track and the "other" track, the irrelevant one. Parents feel frustrated that their children are not valued for their individual abilities and passions when they do not conform to the perceived valued path as indicated by the proliferation of STEM charter schools and programs.

The STEM disciplines alone will never solve the innovation crisis. As we engineers know, they are only a part of the solution. Innovation—the process of inventing something new, desirable, useful, and sustainable—happens at the intersection of technology (is it feasible?), business (is it viable and sustainable?), human factors (is it desirable?), and complexity (is it usable?). In addition, basic science, mathematics, and engineering skills have become commodities worldwide and are available elsewhere at a fraction of the cost here. Yes, critical-thinking problem-solvers from all disciplines working together are the key to innovation. Innovation is local. You don't import it and you don't export it; you create it. It is a culture. It is a way of thinking, communicating, and doing.

STEM students and teachers, together with students and teachers from the humanities, arts, social sciences, and business, must all realize they are equal partners in solving the innovation crisis. They each play a vital role and together must be able to identify the needs of people and society, critically think and solve problems, generate human-centered ideas and rapidly prototype concepts, integrate human values and business into concepts, manage complexity, work in multidisciplinary teams, and effectively communicate results.

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

The message to our students must be that they are each vital to solving the innovation crisis, and this message must be delivered early and often and in the context of real-world problems. Students need to set high expectations for themselves, as we set high expectations for each of them. They need to discover their passion and their talents and take ownership for developing those talents knowing that in doing so they will play a vital role in transforming the world we live in. Engineers can make a vital contribution by setting a professional example and giving a real-world context to what young students study. Are you up to the challenge? We all know amazing things happen when together we attempt the seemingly impossible! EDN

Tune into the latest Mechatronics Webcast—How do you design a military robot?—which will be available on-demand at www.designnews.com beginning Feb 19, and be sure to visit the Mechatronics Zone for the latest mechatronics news, trends, technologies, and applications at www.mechatronicszone.com.

When it comes to creating your next embedded system, an important decision awaits you. Memory. Your design needs speed, reliability, performance and capacity to store the code and data your design demands. No problem. Numonyx has the broadest portfolio of parallel and serial NOR, NAND and phase change memory. And we offer extended temperature support with AEC-Q100 certification on many of our products and expanded design versatility with voltage support up to 5V. All designed to deliver a right-fit solution to help you shorten design cycles, reduce development costs and accelerate the roll-out of your next big idea. Find out how Numonyx memory matters for your next design.

NUMONYX® FORTÉ™ SERIAL FLASH MEMORY

Product Family	Density Range	Voltage/Solution
M25P (block erase)	512 k - 128 Mb	3V, single-I/O
M25PX (4KB block erase)	4 Mb - 64 Mb	3V, multi-I/O
M25PE/M45PE (page erase)	1 Mb - 16 Mb	3V

NUMONYX® AXCELL™ PARALLEL NOR FLASH MEMORY

Product Family	Density Range	Voltage/Solution
M29W/EW (JEDEC command set)	4 Mb - 2 Gb*	3V, page read†
P30/33 (Intel-based command set, sync burst)‡	64 Mb - 2 Gb*	1.8/3V core

Visit www.Numonyx.com/Embedded
for free access to the Numonyx Embedded Design Center.

numonyx™
innovative. memory. solutions.

*Includes stacked solutions †Available on specific densities ‡Other names and brands may be claimed as the property of others

Lead the handheld revolution.

With the industry's lowest power usage, and widest range of small packages, it's no wonder Actel IGLOO® FPGAs can be found in the latest handheld portable devices. Actel has been designed into smartphones, eBooks, cameras, medical devices, remote industrial monitoring devices, undersea seismic sensors, military radios, and the list goes on.

(actual size)

Actel IGLOO low-power FPGAs bring reprogrammability, design security, integration, small form factor, and live-at-power-up operation to handheld portable applications, all at a very low cost.

IGLOO®

Start designing tomorrow's handhelds today with Actel's battery-powered IGLOO Icicle Evaluation Kit.

Actel®
POWER MATTERS

Download in-depth white papers, watch short webcasts, and get the Icicle Kit at:
actel.com/handheld

CONTINUED INNOVATION IS A WELCOME CONTRADICTION

X86 PROCESSORS

BY BRIAN DIPERT • SENIOR TECHNICAL EDITOR

ALTHOUGH THE LAST SEVERAL YEARS HAVE SEEN MUCH CONSTRICTION WITHIN THE SEMICONDUCTOR INDUSTRY, DEVELOPMENT AND IMPLEMENTATION ACTIVITIES IN THE PC AREA HAVE REMAINED COMPARATIVELY BULLISH. X86 VENDORS HAVE VARIOUS PLANS FOR CONTINUING THIS TREND.

In mid-2008, when *EDN* last looked at the broad x86-processor product segment in detail, the world economy was six months into a deep recession (**Reference 1**). Roughly six months later, worldwide economic conditions remained grim (**Reference 2**). Although government officials now claim that the recession is over, key economic indicators, such as the unemployment rate, should deliver underwhelming statistics for some time to come.

You might expect, therefore, that the semiconductor industry would mirror the broader economy's malaise; this scenario has indeed occurred in many technology and product sectors. The x86 CPU business has bucked the general trend, however, as continued R&D investment and resultant new-IC output demonstrate. "Irrational exuberance," a term that former Federal Reserve Board Chairman Alan Greenspan coined in 1996, has thankfully not been evident. However, the microprocessor and core-logic-chip-set suppliers know full well that, with the rapid evolutionary pace of the PC industry, to tangibly slow down in the short term might be sufficient to ensure demise in

the long term. This insight is evident in the continuation of their prerecession momentum.

DESKTOP EXPANSION

Recent analyst reports estimate that Intel holds more than 80% of the overall x86 CPU market and an even higher percentage in key product segments, such as mobile systems and servers. The company also has a higher associated R&D budget and more employees than its key competitors. It is perhaps unsurprising then that most of the new- product roll-outs over the past several years have come from Intel. In 2008, the company launched the Core i7 CPU, its initial product employing the Nehalem microarchitecture. Intel has since substantially broadened the Nehalem line, and you must differentiate among silicon platforms and resultant products to comprehend the full scope of the roll-out.

First, recall Intel's "ticktock" product-and-process cadence, a term that company President and Chief Executive Officer Paul Otellini used in 2006. Intel aspires to reduce overall risk by making major product and process transitions at different times. In the "tick" cycle, the company shrinks a product's lithography to the next process node, making only minor feature alterations in the process. The subsequent "tock" step maintains the same process node but implements more substantial product changes.

AT A GLANCE

- Intel hopes to rapidly spread the Nehalem microarchitecture across all desktop-PC segments, whereas Advanced Micro Devices' recent moves are more modest.
- The mobile-computer market is particularly robust, and vendors' intense focus on it reflects that fact.
- Intel and its competitors disagree on whether a netbook is anything other than a notebook with neutered features.
- Buying a chip set from your CPU supplier may be not only unnecessary but also unwise.
- Servers, cell phones, and TVs round out the picture for Intel's x86 aspirations.

The late-2008 introduction of the Core i7 CPU, a high-end manifestation of the Bloomfield silicon platform, reflected a tock on Intel's 45-nm process (**Figure 1**). Its innovations versus earlier Core-microarchitecture-based products include three channels' worth of integrated memory controllers; the QuickPath Interconnect, an AMD (Advanced Micro Devices) HyperTransport-reminiscent core-to-core and CPU-to-CPU mesh; the resurrection of hyperthreading core virtualization; a trilevel integrated-cache architecture; and Turbo Mode, which automatically boosts the clock speed of active CPU cores in situations in which some cores are not in use, maintaining the chip's thermal envelope and boosting overall performance.

Late last year, Intel launched Lynnfield, its second 45-nm, Nehalem-microarchitecture-based silicon tock offering, this time for mainstream-desktop-CPU platforms. Lynnfield discards one of Bloomfield's DRAM channels, instead integrating 16 lanes' worth of PCIe (Peripheral Component Interconnect Express) 2.0 I/O that previously resided in

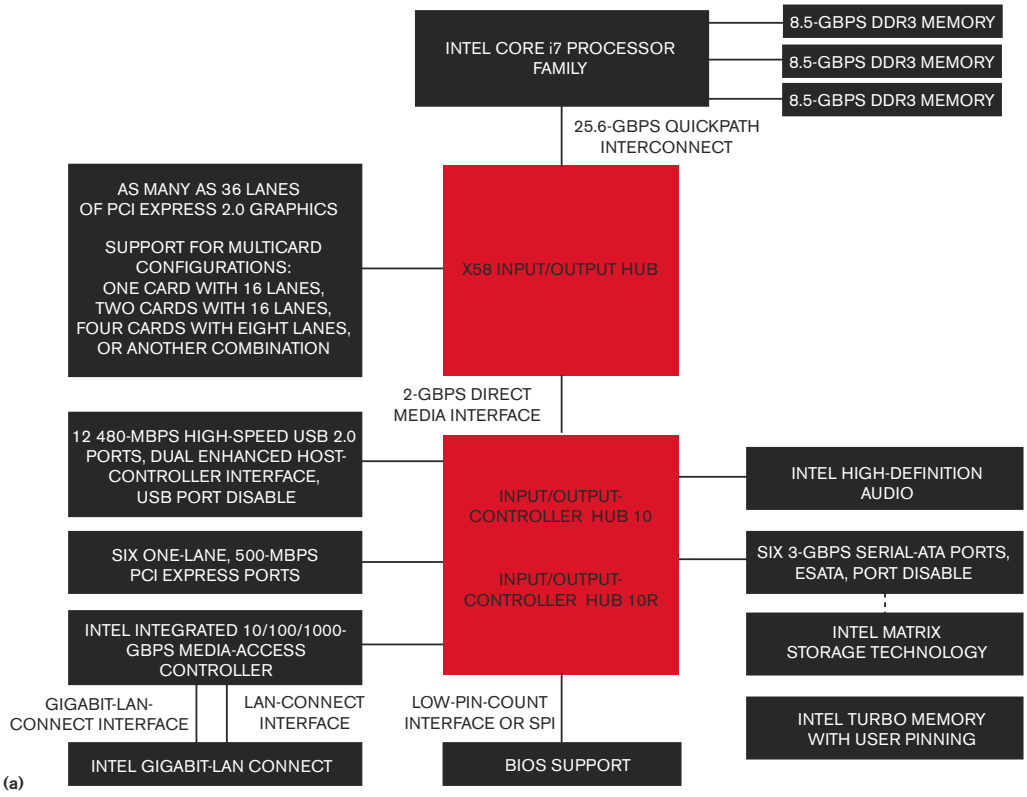

Figure 1 Intel's first three Nehalem-based CPUs, 45-nm-implemented Bloomfield (a) and Lynnfield (b) and 32-nm-fabricated Clarkdale (c), chronologically reflect a common microarchitecture's migration over time from high- through low-end system configurations.

a separate north-bridge core-logic IC. The interface to the remainder of the external core logic is also the 2-Gbyte/sec Direct Media Interface bus; Bloomfield in contrast uses the more-than-10-times-faster QuickPath bus. Lynnfield translated into both Core i7 and Core i5 product proliferations; Core i5 disabled hyperthreading support.

Intel recently executed its next tick by migrating to a 32-nm process commensurate with the unveiling of the low-end Clarkdale CPU. Clarkdale comes in both Core i5 and Core i3 product flavors; some variants disable not only hyperthreading but also turbo-mode support. Instead of relying exclusively on standalone or core-logic-embedded external graphics-accelerator circuitry, Clarkdale places graphics-processing capabilities alongside the CPU in a multidie, single-package arrangement. Clarkdale also migrates the system-memory controller off the CPU and onto the companion die, potentially degrading performance at the trade-off of lower CPU silicon cost. The technology treadmill rolls on; Intel dem-

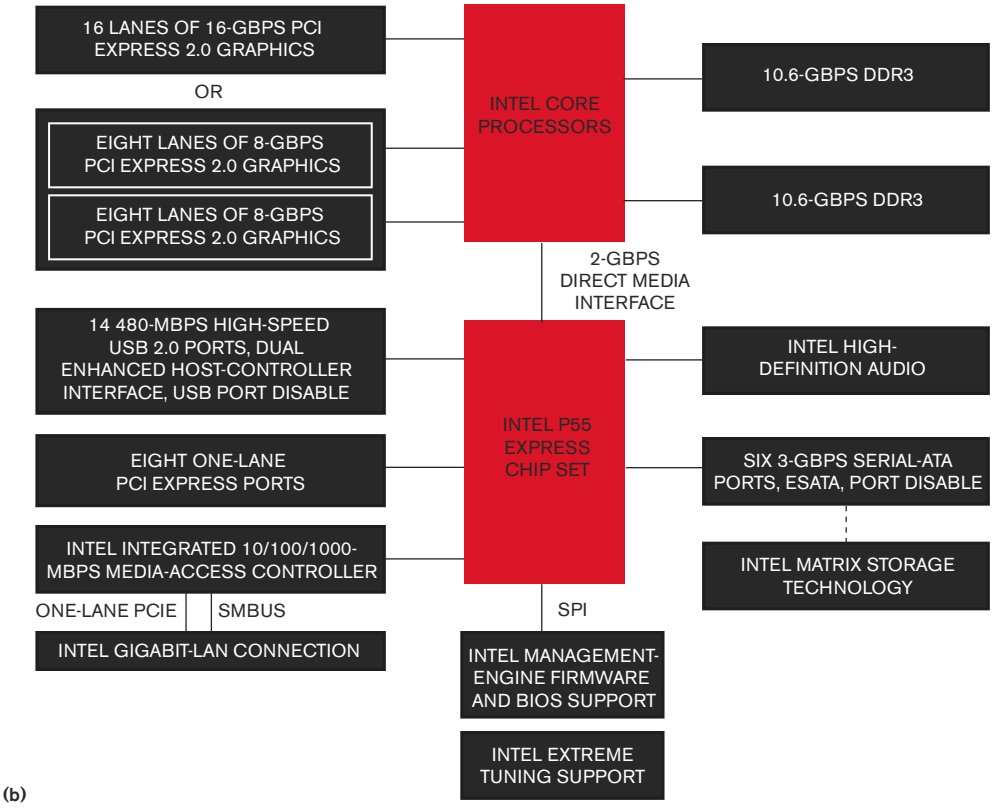

(b)

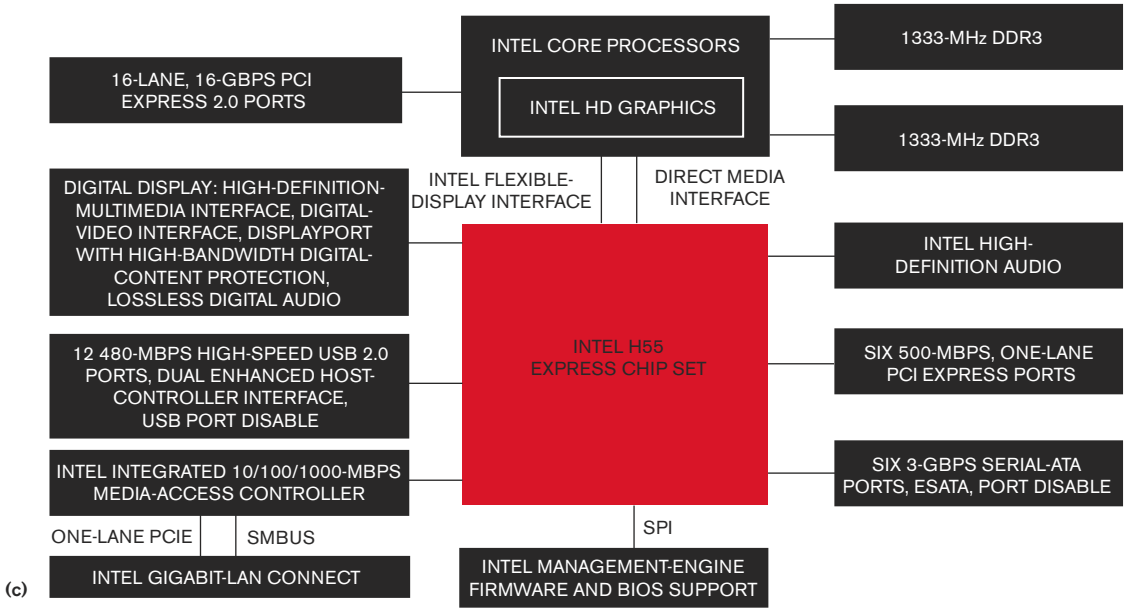

(c)

onstrated its next tock, the 32-nm Sandy Bridge microarchitecture, in functional-silicon form at last fall's Developer Forum. The company showed off a 22-nm test wafer at the same venue.

Intel's primary x86 CPU competitor, AMD, has also over the past several years been fairly busy in this product segment, befitting the company's comparatively strong market-share position in desktop PCs versus other markets. After a problem-plagued roll-out of the 65-nm Phenom CPU in late 2007, complete with poor yields and L3-cache-controller bugs, the subsequent 45-nm Phenom II transition beginning in January 2009 has seemingly gone more smoothly. Phenom II boosts the L3-cache allotment over that of its predecessor and tweaks other minor architectural details. As with Phenom, AMD sells Phenom II not only in full quad-core form but also in triple-core X3 and dual-core X2 variants. These product spins also contain four cores on each die, but AMD disables one or two of them, presumably because they don't pass functional or speed testing. The company is attempting to maximize usable product yield from each manufactured silicon wafer, but this strategy is costly and less than ideal.

Therefore, for mainstream-desktop-PC applications, AMD last summer rolled out the 45-nm Athlon II product line. Whereas the company derived the original Athlon families from the K8 and earlier microarchitectures, Athlon II evolves from the same K10 core foundation as Phenom and Phenom II. However, as a die-slimming move, AMD offers Athlon II only in dual-core form; the company also deleted Phenom II's L3-cache array, compensating somewhat by boosting the L2-cache allotment. AMD also still sells single- and dual-core, 65-nm Athlon products. More generally, the company has attempted to counterbalance a comparative dearth of fundamental product proliferations through a blizzard of clock-speed and operating-voltage spins, striving to satisfy a plethora of product-performance-versus-price expectations with comparatively few silicon foundations.

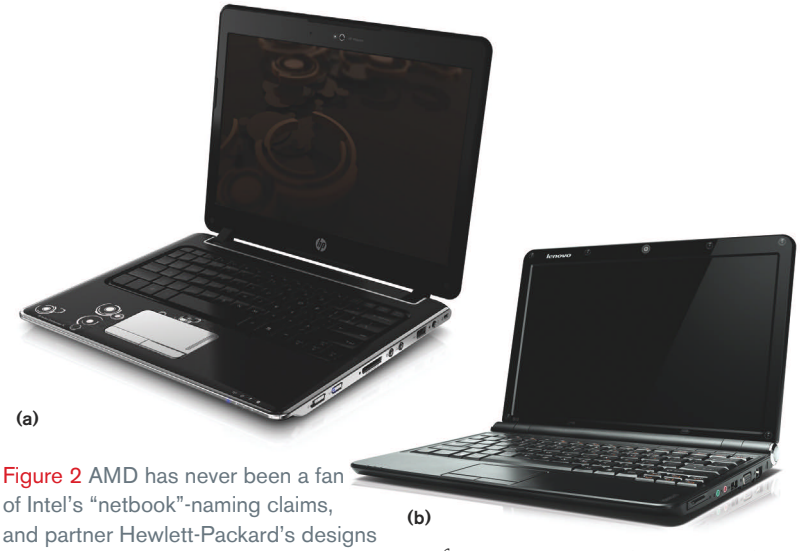

Figure 2 AMD has never been a fan of Intel's "netbook"-naming claims, and partner Hewlett-Packard's designs reflect an alternative fiscally sensitive approach (a). Via's Nano CPU is one of three fundamental options for Lenovo's IdeaPad S12; the other two variants pair an Atom CPU with either an Intel or an Nvidia Ion chip set (b).

MOBILE MIGRATION

As anyone who has recently visited computer retailers or scanned their advertisements realizes, desktop-computer sales are at best stagnant, whereas mobile-computer sales are exploding. Intel reportedly controls approximately 90% of this market, so it's no surprise that the company has focused significant attention in this area. Most Intel-based notebook computers still use so-called CULV (consumer-ultralow-voltage) CPUs employing the previous generation's Core microarchitecture, although the company last fall began shipping mobile Core i7 CPUs it derived from Clarksfield, a notebook-intended variant of the Nehalem-based Lynnfield. More recently, Intel unveiled Arrandale, a mobile version of Clarkdale. The current modern differentiation between desktop and mobile CPUs is nebulous at best. It reflects operating-voltage and clock-frequency variations but little if any underlying circuitry differences because power consumption is a critical consideration in all market segments (see **sidebar** "Server spins").

To date, AMD has not attempted to aggressively shoehorn Phenom, Phenom II, or Athlon II CPUs into notebook computers, relying instead on power-consumption-optimized earlier-generation Athlon products. The company prefers to promote and sell them in processor-plus-chip-set form. Take, for example, the Neo platform, which AMD unveiled at the January 2009 Consumer Electronics Show along with system partner Hewlett-Packard (**Figure 2**). AMD,

SERVER SPINS

Desktop and mobile versions of CPU platforms share similar features, and similar commonality exists between desktop and workstation-and-server variants. The workstation-plus-server adaptation typically offers multiprocessor connectivity, undergoes additional testing, and supports memory technologies that are amenable to high system capacities.

For example, the server-tuned version of Intel's Bloomfield is Gainestown–that is, Nehalem EP–the Xeon 5500 series. Intel plans to introduce Jasper Forest, a server variant of the Lynnfield desktop-PC and the Clarksfield mobile-system CPUs. Also on the company road map is Gulftown, a six-core shrinkage of the eight-core, 45-nm-based Beckton–that is, Nehalem EX–to the 32-nm-process node and available in both desktop and dual-processor server versions.

Advanced Micro Devices exhibits similar desktop/server commonality. For example, the 65-nm Phenom desktop and Barcelona server CPUs, respectively, migrated to the 45-nm Phenom II and Shanghai. AMD also has a desktop-PC-targeted variant of its six-core, 45-nm Istanbul server processor on its 2010 road map.

by virtue of its 2006 acquisition of ATI Technologies, now touts formidable stand-alone and core-logic-integrated graphics capabilities. And AMD isn't shy about pointing out its claimed advantages over Intel's alternatives as they relate to such topics as Windows Vista and Windows 7 UI (user-interface) enhancements, mainstream 3-D gaming scenarios, and increasingly important hardware-assisted decoding of leading-edge video codecs.

Via Technologies, meanwhile, has unveiled three variants of its latest-generation, 64-bit, single-core Nano CPU. The 1000 and 2000 series of products are identical from the silicon and microcode standpoints; the only differentiation is with respect to the marketing moniker. The newer 3000 series, on the other hand, represents a more involved circuit redesign, adding support for SSE4 (streaming single-instruction/multiple-data extension version 4) and with claimed higher performance and lower power consumption. The long-promised dual-core variant of the Nano family remains absent; the latest company road maps forecast its arrival to occur no earlier than mid-2010. Nonetheless, Via has secured several design wins with Tier 1 notebook-PC OEMs, such as Lenovo and Samsung. Via also continues to sell its earlier-generation C7 processors.

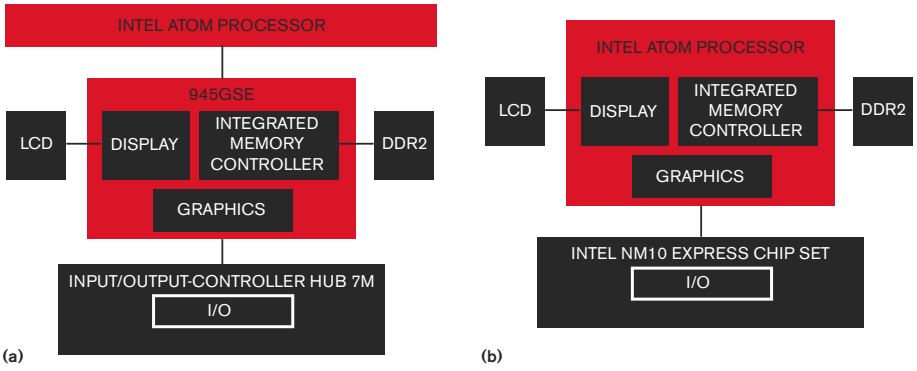

Figure 3 Selective integration within the CPU is the hallmark of the Diamondville-to-Pine Trail Atom evolution of 2008 (a) versus today (b), although Intel may be the key beneficiary from a competitive isolation standpoint.

Figure 4 Nvidia's Ion companion chip has sufficient video-decoding capabilities to transform the Atom CPU into a passable home-theater-PC foundation (courtesy Acer).

FISCAL EVOLUTION

The term "netbook" represents Intel's attempt to separate the systems from notebooks (see **sidebar** "Handsets and other CE hopes"). Netbooks contain low-cost Atom microprocessors, whereas notebooks include higher-end, higher-profit, and higher-revenue-garnering Intel CPU alternatives. A third category, the smartbook, includes 3G (third-generation) cellular-data subsystems and often offers carrier subsidization. Processor differences aside, Intel and partner Microsoft have undertaken other netbook-versus-notebook differentiation steps. These steps include restricting to 2 Gbytes the maximum system memory that the associate core logic supports, encouraging their customers to further decrease the installed system memory to 1 Gbyte, and constraining the screen sizes and resolutions that netbooks can implement. Historically, Microsoft's motivation has been to minimize the number of less lucrative Windows XP licenses that it sold into netbooks versus those for the more fiscally attractive Windows Vista in notebooks. Nowadays, its intention is to steer as many licensees as possible away from the netbook-tailored Starter Edition of Windows 7 and toward more expensive variants of the OS.

Intel in December formally unveiled the latest Atom CPU, Pineview—or Pine Trail in CPU-plus-chip-set-platform lingo—which exemplifies this segmentation trend (**Figure 3**). Pineview integrates the graphics core and memory controller that the earlier-generation Diamondville device implemented on a separate north-bridge core-logic IC. Pineview's approach has battery-life advantages, especially when you consider that the earlier 90-nm-process-based 945GC and 130-nm ICH7 consumed far more power than did the 45-nm-derived Atom N- and Z-series CPUs.

However, Intel also implemented the Pineview Atom CPU on a 45-nm process, which translates to an increase in transistor count over precursor Atom CPU generations. To minimize die size, Pineview's graphics core has fewer features than that of the 945GC. This seeming step backward may make you wonder why else Intel might have made this integration move. Performance is a likely reason; although the CPU core itself remains the same, the tight coupling to the graphics and memory subsystems is preferable to the earlier device's separation through the front-side bus. Also recall that the graphics core employs system memory as its frame buffer. This time around, the device supports 667- and 800-MHz DDR2 SDRAM versus only 533-MHz memory on the 945GC. Unfortunately, Pineview's video outputs have maximum resolutions of only 1366×768 pixels from digital ports and a maximum of either 2048×1536 or 1400×1050 pixels with an analog VGA connection.

CHIP-SET COMPETITION

Pineview's tight coupling between the CPU and the GPU (graphics-processing unit) also conveniently—at least

HANDSETS AND OTHER CE HOPES

Although Intel's aspirations to seriously combat ARM and MIPS in the CE (consumer-electronics) market have to date largely gone unrealized, the company strives onward. Coming later this year is Moorestown, a handheld-device-focused platform follow-on to the first-generation Atom Silverthorne (Figure A). As with Pine Trail for netbooks, Moorestown's CPU remains on the 45-nm-process node but is more integrated than its precursor. Also, at last fall's Intel Developer Forum, the company unveiled the Atom CE4100, a highly integrated descendant of the Pentium M-based CE3100 initial offering, targeting use in set-top boxes and network-enabled TVs.

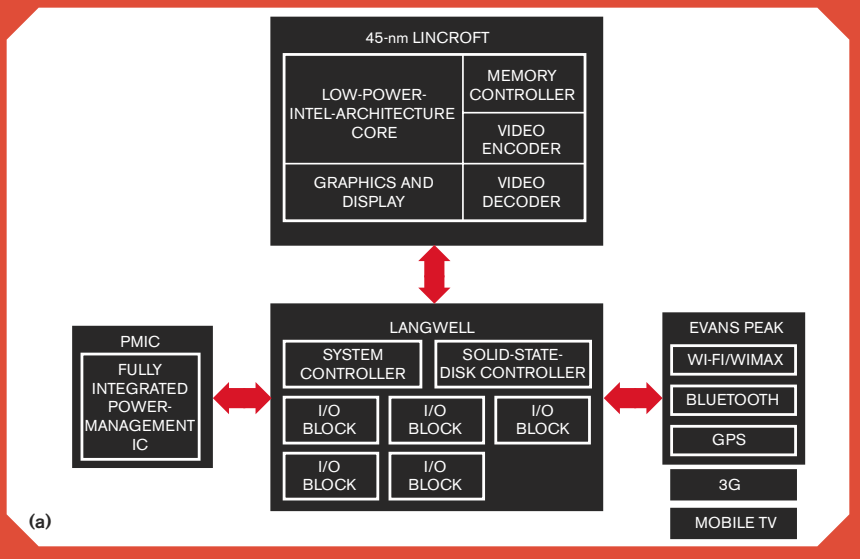

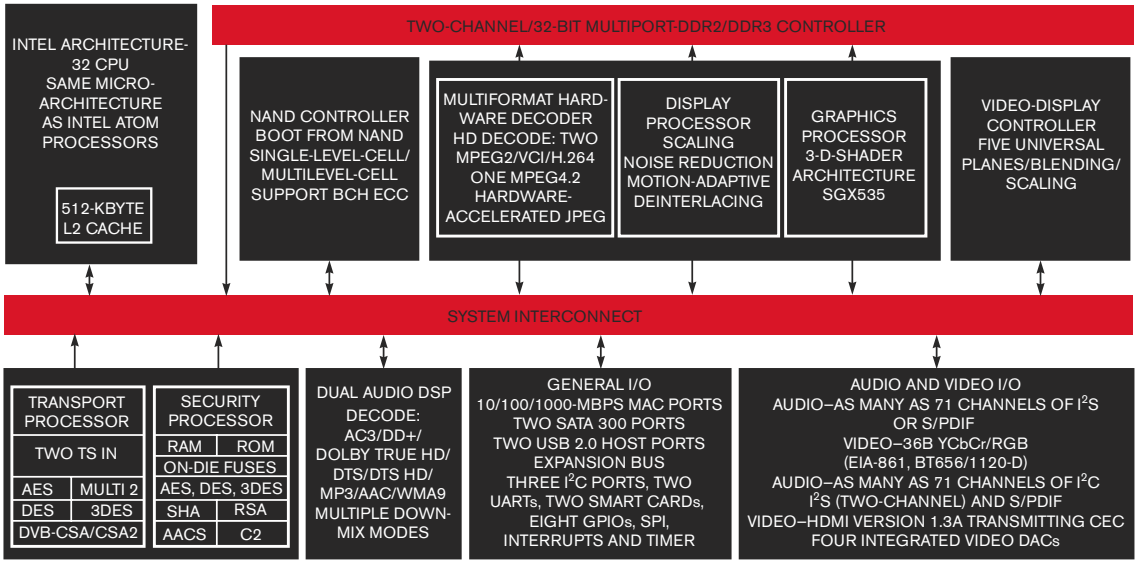

AAC: ADVANCED AUDIO CODING
AACS: ADVANCED ACCESS-CONTENT SYSTEM
AC3: AUDIO COMPRESSION 3
AES: ADVANCED ENCRYPTION STANDARD
BCH: BINARY-CODED HEXADECIMAL
C2: SECURITY
CEC: CONSUMER-ELECTRONICS CONTROL
DAC: DIGITAL-TO-ANALOG CONVERTER
DES: DATA ENCRYPTION STANDARD
DTS: DIGITAL THEATER SYSTEM
DVB-CSA: DIGITAL-VIDEO BROADCASTING COMMON SCRAMBLING ALGORITHM/COMMON SCRAMBLING ALGORITHM 2
ECC: ERROR-CORRECTING CODE
GPIO: GENERAL-PURPOSE INPUT/OUTPUT
HD: HIGH-DEFINITION
HDMI: HIGH-DEFINITION MULTIMEDIA INTERFACE
I2C: INTER-INTEGRATED-CIRCUIT SOUND
JPEG: JOINT PHOTOGRAPHIC EXPERTS GROUP
L2: LEVEL 2
MP3: MOVING PICTURES EXPERTS GROUP LAYER 3 AUDIO
MPEG: MOTION PICTURES EXPERTS GROUP
MULTI 2: TWO MULTIPROCESSORS
RGB: RED/GREEN/BLUE
RSA: RIVEST/SHAMIR/ADLEMAN
SATA: SERIAL ADVANCED TECHNOLOGY ATTACHMENT
SHA: SECURE HASHING ALGORITHM
S/PDIF: SONY/PHILIPS DIGITAL INTERFACE
SPI: SERIAL-PERIPHERAL INTERFACE
3-D: THREE DIMENSIONAL
3DES: TRIPLE DATA ENCRYPTION STANDARD
TWO TS IN: TWO INPUT TRANSPORT PROCESSORS
UART: UNIVERSAL ASYNCHRONOUS RECEIVER/TRANSMITTER
VCI: VIRTUAL-CHANNEL IDENTIFIER
WMA9: WINDOWS MEDIA AUDIO 9
YCBCR: LUMA COMPONENT/BLUE-DIFFERENCE CHROMA COMPONENT/ RED-DIFFERENCE CHROMA COMPONENT

(b)

Figure A Intel hopes its Moorestown two-chip set will lead to some serious cell-phone design wins (a), whereas the Atom CE4100 targets the living room (b).

for Intel—shuts out third-party graphics products, which previously connected to the Atom processor through that same front-side bus. Instead, third-party graphics chips must tether to Pineview through the NM10 companion chip's PCIe lanes, several bus "hops" away from the processor. System memory, which such a GPU might prefer to also employ as its frame buffer, is also several hops away, resulting in either poor performance or the necessity for a costly dedicated frame buffer that directly connects to the third-party GPU.

Pineview's integration, reminiscent of that in the Clarkdale and Arrandale CPUs, brings to the forefront a long-standing feud between Intel and Nvidia (**Reference 3**). Nvidia's Ion chip, an alternative to Intel's two-chip 945GC-plus-ICH7 approach, muddles the distinction that Intel strives to preserve between netbooks and notebooks. Ion uses the same IC that Apple's latest Intel-based MacBook laptops and entry-level Mac mini desktop computers use. It has secured a few key design wins in portable systems alongside Intel's first-generation Atom CPUs, as well as in a host of so-called nettops for home-theater-PC and other niche applications (**Figure 4**).

Ion's fundamental appeal, aside from its single-chip function integration, is that it has notably better graphics- and video-processing capabilities than Intel's alternative. Ion offers some degree of hardware-accelerated decoding for advanced video-compression algorithms, such as H.264, which Adobe Flash is adopting, and VC-1, which Netflix's Watch Instantly online service employs. Blu-ray discs also use both H.264 and VC-1; Ion handles most of the processing itself instead of burdening the CPU with the task. Nvidia has long publicly complained, however, that the Atom pricing strategy is competitively unfair in that Intel charges significantly more for the stand-alone CPU than when it sells it with the 945GC and ICH7. Indicating the potential reality behind Nvidia's claim, systems based on Ion tend to be significantly more expensive than their "generic" Atom counterparts. It's unclear, however, whether this differentiation reflects bill-of-materials costs or PC OEMs' desires to use potential customers' perception of the superiority of an Ion-inclusive system to boost its price.

> ⊞ Go to www.edn.com/100204df to post a comment on this article.
>
> ⊞ See the "x86 Processors" posts at www.edn.com/briansbrain for supplemental information on this article's topics.

Speaking of pricing, even if Nvidia figures out how to make a cost-effective offering with compelling performance with its upcoming Ion 2 chip for Pineview, the product's success isn't assured. Pineview alone has sufficient graphics horsepower to run Windows 7's rich GUI (graphical user interface), for example. You may thus conclude that video is the primary motivation to append Ion 2 to Pineview. An even more cost-effective silicon alternative exists, however, in Broadcom's Crystal HD (high-definition) video-decoder IC. The company initially unveiled the device several years ago, but it has only now hit its stride in conjunction with the ascendancy of the Atom CPU. Intel prominently showcases Crystal HD in its Pineview promotional materials. And OEMs sell Crystal HD-augmented systems for only $30 or so more than their Atom-only counterparts, suggesting that Broadcom is aggressively pricing the part to carve out a market and fend off the Nvidia alternative. To buffer itself against Atom-only business dependence, Nvidia also plans to develop Ion variants for use with Via's Nano CPU series.

CONSUMER PROTECTION

Ion is only one of several areas of disagreement between Intel and Nvidia. Nvidia also has yet to obtain a license for either the QuickPath Interconnect or Direct Media Interface buses, meaning that the company cannot legally sell chip sets for Nehalem-generation Intel microprocessors. More generally, Nvidia is struggling from a business standpoint, with its DirectX 11-supportive graphics chips notably delayed and with AMD's microprocessor customers also tending to select AMD/ATI chip sets and discrete GPUs.

The disputes between the two companies are among several motivations that the US FTC (Federal Trade Commission) mentioned when in mid-December it filed a lawsuit against Intel for alleged anticompetitive practices. This latest legal setback is only one in a series that Intel has suffered in the past year. The European Union also last May fined Intel $1.45 billion, which the company agreed to pay, although it vowed to appeal the ruling. In November, New York State filed an antitrust lawsuit against Intel, and Intel in the same month paid AMD $1.25 billion to settle long-standing litigation between the two companies.

Nvidia may be closely following the evolution of the FTC action for more reasons than concern about chip sets and graphics. An unconfirmed long-standing rumor suggests that the company has for some time been internally developing an x86-microprocessor architecture. Both design-team acquisitions and official—albeit obscure—company comments bolster this rumor. However, Nvidia has not yet obtained permission to sell chips employing the architecture. If the x86-by-Nvidia rumor is true and if the FTC's legal aspirations are successful, one possible remedy would be for the FTC to require that Intel grant Nvidia an x86 license. EDN

REFERENCES

1. Dipert, Brian, "Embedded x86: keystone of your non-PC design?" *EDN*, May 29, 2008, pg 38, www.edn.com/article/CA6562580.
2. Dipert, Brian, "Evaluating value-oriented x86 CPUs: the price of falling prices," *EDN*, Jan 8, 2009, pg 30, www.edn.com/article/CA6625435.
3. Dipert, Brian, "CES 2009: the price of falling prices and a path to Intel's potential demise," *EDN*, Jan 8, 2009, www.edn.com/article/CA6627620.

FOR MORE INFORMATION

Acer
www.acer-group.com

Advanced Micro Devices
www.amd.com

Broadcom
www.broadcom.com

Hewlett-Packard
www.hp.com

Intel
www.intel.com

Lenovo
www.lenovo.com

Microsoft
www.microsoft.com

MIPS Technologies
www.mips.com

Nvidia
www.nvidia.com

Samsung
www.samsung.com

Via Technologies
www.viatech.com

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

BY RICK NELSON • EDITOR-IN-CHIEF

Wireless to GO

As wireless technology becomes increasingly pervasive, vendors are offering a variety of hardware and software technologies to help you add communications functions to your products. Offerings support communications standards, including Bluetooth, Wi-Fi, GPS (global-positioning system), LTE (long-term evolution), and WiMax (worldwide interoperability for microwave access). Choosing the chips and associated software necessary to add wireless connectivity—and proving that your chosen implementation works and meets relevant certification requirements—can be challenging, however. Even if you put together a working design, you might not succeed in the market if you haven't optimized performance, power consumption, cost, and size. And what's optimum today might not be optimum as communications standards and your customers' needs evolve, so you'll want to choose a hardware and software implementation that lets you adapt without starting from scratch with each new generation (see **sidebar** "Wireless at a glance").

Each of your successive generations will also probably need to remain compatible with earlier versions. "A general trend in the industry is that we are adding new standards and protocols but not subtracting the old ones as quickly," says Fanny Mlinarsky, founder and president of consulting company OctoScope, using the cellular industry as an example. "They linger, and we have to be backward-compatible. So as the whole cellular industry is reinventing itself with broadband wireless 3G [third generation] and 4G [fourth generation], it has to carry that baggage of GSM [global system for mobile communications] in Europe and CDMA [code-division multiple access] here. Even WCDMA [wideband CDMA], which is 3G, is barely making it to market now, only to be replaced soon by LTE."

For some applications, you may be able to choose a single chip to implement your wireless links. Marvell at last month's CES (Consumer Electronics Show) touted its Avastar family of wireless connectivity devices for always-on consumer-electronics devices, such as handsets, portable media players, e-readers, printers, digital cameras, netbooks, digital TVs, set-top boxes, DVD players, gaming consoles, and even thermostats. The company can snap together its IP (intellectual-property) cores to create devices for target markets. It is now sampling devices for Wi-Fi; Wi-Fi and Bluetooth; Bluetooth and FM; and Wi-Fi, Bluetooth, and FM. Marvell expects to introduce four more devices in the family during the first half of this year.

For developers of embedded systems who need to add wireless functions to a microcontroller-based design, Microchip Technology offers one-stop shopping. The microcontroller maker took aim last month at enhancing its support of embedded wireless applications with its acquisition of ZeroG Wireless, a privately held fabless semiconductor developer of Wi-Fi-certified transceivers and FCC-certified modules. Embedded-system designers can use ZeroG devices to add the ubiquitous Wi-Fi networking protocol to any of Microchip's 8-, 16-, or 32-bit PIC microcontrollers through an SPI (serial-peripheral interface).

Microchip supports other wireless standards, as well, including ZigBee. The company announced in December that it has received certification for its new ZigBee-RF4CE-compliant platform, a development kit for ZigBee that includes the XLP PIC controller, an MRF24J40 transceiver for 802.15.4 ZigBee networks, and an FCC-certified module. The kit comes with a certified protocol stack for ZigBee applications.

The XLP PIC at the center of the platform integrates several peripherals for capacitive touch sensing, USB interfaces, and analog I/O. Microchip incorporates software, including the mTouch sensing software, in the kit to use the peripherals. Other tools for the kit include the MPLab integrated development environment, MPLab Real ICE system, MPLab ICD 3, Pickit 3 debugger/programmer, and C compilers.

You'll also find single-chip implementations for cell phones. For example, Infineon offers its single-chip X-Gold 116 GSM/GPRS (general-packet-radio-service) implementation, which the company fabricated in 65-nm CMOS, for messaging phones. That chip integrates the GSM baseband, RF transceiver,

VENDORS ARE SERVING UP A VARIETY OF HARDWARE AND SOFTWARE FLAVORS TO ADD COMMUNICATIONS CAPABILITY TO FIXED AND MOBILE DEVICES.

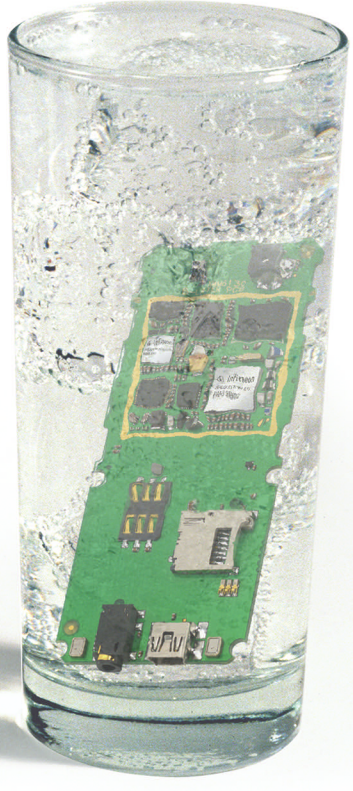

power-management functions, SRAM, and FM-radio functions in an 8×8-mm package. Infineon offers the device as part of its XMM 1160 messaging-phone platform, which incorporates about 50 components on a four-layer printed-circuit board and comes with user-interface, multimedia-framework, media-player, and Java software. The company also offers the XMM 6130 platform, which targets the entry-level Internet-browsing-phone market, supporting 3G HSDPA (high-speed downlink-packet access). The platform integrates an ARM11-based microcontroller, baseband digital and analog features, and power-management functions, and it offers dedicated interfaces for a camera, a display, a USB port, and memory cards.

Single-chip GPS implementations are available, as well. Atheros Communications last month announced the latest member of its family of GPS products, the third-generation, single-chip AR1520 GPS receiver and companion Atheros FYX 1.0 software suite. Atheros based the AR1520 on the new Atheros FYX location core. It delivers greater navigation accuracy, faster location fixes, enhanced receiver sensitivity, and lower power consumption than the company's previous products. These features make the AR1520 suitable for mobile consumer products, such as PNDs (personal navigation devices), netbooks, smartbooks, portable gaming devices, media players, and smartphones.

ST-Ericsson, meanwhile, is taking a technology-platform approach to meet consumer needs for ubiquitous connectivity. According to Thierry Tingaud, vice president and general manager for 2G, EDGE, TD-SCDMA, and connectivity products at the company, the mobile device has become the convergence point for consumers for multiple wireless technologies, including GPS and Wi-Fi, as well as cellular operation. ST-Ericsson offers platforms for smartphones, feature phones, and entry-level phones as well as platforms for connected devices, such as wireless modems in laptops.

The company announced in December that it is working with Nokia as part of a long-term partnership to develop TD-SCDMA technology, which ST-Ericsson and Nokia are deploying in China. For its part of the partnership, Nokia will use ST-Ericsson as a key supplier of chip-set platforms for its Symbian-based TD-SCDMA-device portfolio. The companies say that the partnership will enhance the leadership position they both have in China's mobile-system market. Nokia in October launched the 6788, its first Symbian-based TD-SCDMA phone, and it announced the creation of a dedicated TD-SCDMA R&D team in Beijing. ST-Ericsson, through its Chinese subsidiary T3G Technology, has been developing TD-SCDMA technology and devices for more than six years.

AT A GLANCE

- We are adding new communications standards and protocols but not subtracting the old ones as quickly.
- It's not just about making chips; it's also about gluing everything together with software.
- An effective way to support multiple communications standards is to employ a technology platform or reference design.
- People underestimate how much it costs to put wireless connectivity into a product.
- Designers can implement baseband functions in FPGAs and ASICs that they once had to implement in software.

SEAMLESS CONNECTIVITY

Broadcom is doing exciting work to combine the various technologies for integration into platforms ranging from PCs to mobile handsets and other consumer-electronic devices, according to Craig Ochikubo, vice president and general manager of wireless personal-area networking in the company's wireless-connectivity business, in which he is responsible for Bluetooth devices. A critical focus, Ochikubo says, is on the software drivers and applications necessary for providing seamless connectivity. "It's no longer about just making chips," he adds. "It's really about gluing everything together with the software and combining the hardware elements in very complementary ways."

Ochikubo elaborates on the software challenges. "We are able to put a Bluetooth chip into a PC and a similar chip into a handset, but the software associated with these applications is completely different." Even within the PC market, Broadcom provides different software packages for Microsoft Windows 7 and legacy XP and Vista systems. Serving the handset market requires yet more packages for mobile operating systems, for which the company offers software support for functions such as audio streaming and remote control that the common mobile operating systems don't support.

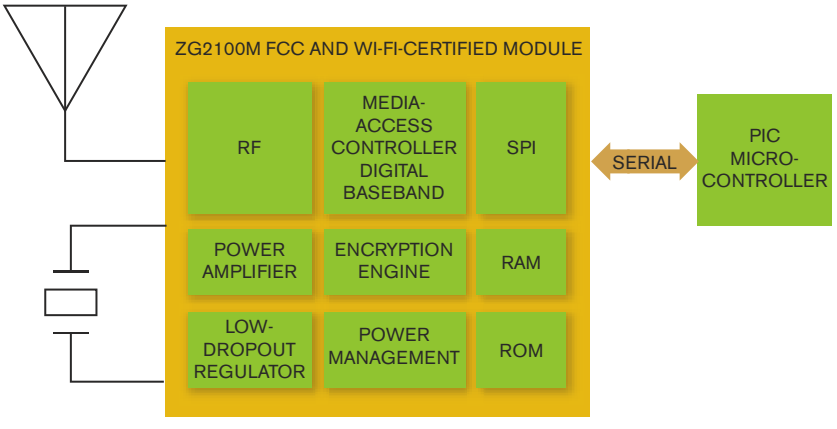

PIC microcontrollers from Microchip combine with Wi-Fi modules from ZeroG, which Microchip recently acquired, to serve embedded-system applications. The ZeroG parts include memory, media-access-control digital-baseband functions, an encryption engine, power management, low-dropout-regulator functions, and power-amplifier and other RF functions. An SPI links the Wi-Fi device with the microcontroller.

"A customer will come to us and say, 'We want to make a phone using the nascent Android operating system, but we want you to implement the most advanced features available in the market,'" says Ochikubo. "That [request] drives us to constantly push the envelope of what we offer."

The software Broadcom provides varies widely from the software that the operating system or end customer provides, he notes. In the PC market, Broadcom provides software all the way up to the user interface, offering, for example, a Bluetooth experience above and beyond what Windows can offer. Similarly, Broadcom provides a user interface for Windows Mobile that differs greatly from Microsoft's offering. However, for feature phones, Samsung, LG, and other vendors want to maintain their own look and feel in the user interface. For those customers, Broadcom provides a set of APIs (application-programming interfaces) that let customers quickly implement functions such as stereo-audio streaming and data synchronization.

Broadcom is also focusing on Bluetooth 3.0, which allows Bluetooth-centric designers to use the 802.11 physical layer to provide Wi-Fi-speed data transfers in a Bluetooth environment. Bluetooth 3.0 supports bulk synchronization of music libraries between PCs and music players or phones, supports wireless transfer of photos to printers, and sends video files from cameras or phones to computers or televisions. An alternative for Wi-Fi-centric designers, says Ochikubo, is Wi-Fi Direct, which enables Wi-Fi devices to connect and share data without joining a traditional home, office, or hot-spot network. Whatever approach Broadcom's customers choose to take, he says, he sees Broadcom's recently announced InConcert Maestro software platform as making the operation simple and transparent for the end user.

In addition to focusing on the higher speeds that Bluetooth 3.0's Wi-Fi physical layer affords, the Bluetooth Special Interest Group is also focusing on low-power applications with its "Bluetooth low-energy," or Bluetooth 4.0, specification. Bluetooth low energy will address markets such as health care, sports and fitness, security, and home entertainment.

ST-Ericsson announced in December its CG2900 device, which incorporates Bluetooth low-energy technology. With the CG2900, consumers can use their mobile phones to collect and display information from Bluetooth low-energy sensors and to link these devices to the Internet for applications such as remote health-care and fitness monitoring. Using the CG2900, the company reports, mobile phones can become the hubs of a varied ecosystem of coin-cell-powered devices, including watches, medical and sports monitors, home and proximity sensors, battery-level monitors, and temperature and pressure indicators. The phones can also act as controllers for games, gadgets, and home devices.

The XMM 6130 platform targets the entry-level Internet-browsing-phone market, supporting 3G HSDPA. It integrates an ARM11-based microcontroller, baseband digital and analog features, and power-management functions (courtesy Infineon).

MULTIPLE STANDARDS

An effective way to support multiple communications standards within a product is to employ a technology platform or reference design. Qualcomm's Snapdragon platform, for example, supports 3G mobile broadband, Wi-Fi, Bluetooth, and GPS. The second-generation QSD8650A, which the company fabricated in 45-nm technology versus 65 nm for first-generation devices, integrates a custom-designed, 1.3-GHz ARM Version 7-based microprocessor—versus 1 GHz for first-generation parts—with a 166-MHz internal bus; a 600-MHz DSP; a stand-alone, power-efficient 2-D graphics accelerator; and an enhanced 3-D graphics core, all in the same 15×15-mm package size of the first-generation version. The second-generation version offers 1080-pixel—versus 720-pixel—high-definition video recording and playback; improved Adobe Flash performance; a wider range of multimode UMTS (universal-mobile-telecommunications-system) and CDMA 3G-mobile-broadband connectivity options; and support for Wi-Fi, Bluetooth, GPS, a 12M-pixel camera, and mobile broadcast TV. The second-generation chip set achieves better than 2800-Dhrystone-MIPS performance and uses just 350 mW versus 2200 Dhrystone MIPS at 500 mW for the first-generation version. Standby power is less than 10 mW.

Companies including Acer, HTC, Sony Ericsson, and Toshiba have already introduced more than a half-doz-

WIRELESS AT A GLANCE

Bluetooth (www.bluetooth.org): wireless standard serving cable-replacement and personal-area-network applications.
GPS (global-positioning system, www.gps.gov): US government's space-based, real-time radio-navigation system.
LTE (long-term evolution): the 3GPP (Third Generation Partnership Project, www.3gpp.org/LTE) strategy for the 3GPP's cellular technology.
WiMax (www.wimaxforum.org): worldwide interoperability for microwave access, a technology based on the IEEE 802.16d and 802.16e standards for fixed and mobile communications, respectively; also known as wireless MAN (metropolitan-area network).
Wi-Fi (www.wi-fi.org): wireless local-area-network technology built on the IEEE 802.11x family of standards.
ZigBee (www.zigbee.org): a technology building on the IEEE 802.15.4 standard that supports low-power remote-monitoring applications.

en first-generation Snapdragon-powered phones, according to Qualcomm. The company expects numerous other Snapdragon smartphones and smart books to reach the market this year. At last month's CES, Lenovo announced its Lenovo Skylight, an ARM-based-processor smart-book device that the company based on Qualcomm's Snapdragon platform. Skylight connects with AT&T 3G mobile-broadband service in the United States. Also at CES, Qualcomm also said that it is working with Hewlett-Packard to develop a Google Android-based smart-book device that incorporates the Snapdragon chip set.

Lenovo at CES introduced the Skylight smart book, which incorporates Qualcomm's Snapdragon chip-set platform.

CHOOSING CHIPS

There are alternatives to buying chips and designing them in yourself, says Clay Melugin, senior marketing manager for RF at Infineon Technologies. "A lot of people underestimate what it costs to put wireless into a product," he explains. "They ask how much a chip set costs." Melugin emphasizes, however, that the initial chip-set cost represents only a portion of the costs of getting a wireless product to market. If you choose a chip set, you must pay a licensing fee for a reference design and then spend your own time and effort on product development and the certification process. "Modules are a very affordable approach if you want to get a product to market fast and avoid the high costs of embedding wireless—especially cellular—into your product," Melugin says. "If you have volume above 20,000 per year for a single product, then you can consider embedding Wi-Fi and Bluetooth, but cellular requires volumes above 200,000 per year to make the effort worthwhile." Whatever your volumes, he adds, you need to compare the basic factors of cost, size, power drain, and performance in choosing your chips or module.

If you have the volumes and want to embed a chip set, Melugin recommends that you look for a platform that includes the baseband, RF, protocol stack, and a reference design that testing and certification have proved, all the way through carrier-network testing. In that way, you avoid a lot of trouble in getting an embedded design to market. "We see many companies that want to embed cellular into their products," he says. "But, in general, they underestimate the cost to take an embedded design through the certification process." If you have the volumes and the resources to invest in the chip-selection process, you should look at the future road map of prospective chip providers, covering their upcoming releases for the new standards. "You don't want to do business with a one-hit wonder," Melugin says, adding that you will also want to ensure that your vendor will support software reuse from one platform generation to the next. "You never want to have to rewrite the code if at all possible when you take a chip set over a generational leap," he cautions. "Look for a supplier who stabilizes this process by holding universal interfaces to the chip set so you can easily port to new chip sets and platforms as they are released."

SOFTWARE-DEFINED RADIO

Infineon is developing an SDR (software-defined-radio) architecture for satellite and terrestrial communications. "The market need for SDR is here," says Melugin, and its evolving role will be interesting. "SDR essentially makes it possible to change the standard under which the radio operates. It's easy from the software side but very much of a challenge on the RF side." SDR's flexibility and size will drive designers to the technology, but they will face trade-offs in cost and power consumption. Melugin says he sees SDR approaches serving target-market needs, rather than offering a ubiquitous approach for "future proofing" a product design from evolutionary technology changes. Smaller-geometry designs that use less power will make SDR more attractive in the future, he adds.

Sigmatix, an 18-month-old start-up focusing on LTE and WiMax 4G technologies, is actively pursuing an SDR approach for cellular applications. According to Dave Kelf, president and chief executive officer, the SDR concept has been around for 20 years or so and has found use in some applications, although it has never reached the performance levels necessary to serve mainstream companies in the cellular industry. Kelf expects that situation to change as SDR designers begin to take advantage of all the parallelism and concurrency available in modern processors. Sigmatix plans to take advantage of modern processor capabilities with a multimode vector radio, which takes advantage of the vector-processing capability of processors having wide SIMD (single-instruction/multiple-data) and VLIW (very-long-instruction-word) architectures. "We can get much greater performance out of these processors," he says. Sigmatix can code 4G baseband functions on them and see a 30-fold improvement in performance or a concomitant reduction in power for the same performance. The company expects to make a formal announcement about its product offerings during the next quarter.

Kelf sees application opportunities for such an approach in handsets that combine LTE, GSM, CDMA, and Wi-Fi standards. "If you look at a typical handset, you might need five or six baseband

chips for LTE, WiMax, GSM, CDMA, and Wi-Fi," he says. "Each chip might cost from $4 or $5 up to $15 or $16 for HSDPA. And the first WiMax chips coming out cost as much as $68, and that's way too much." He adds that all these standards have similar processing needs and that cell-phone vendors are all looking at ways of using software so they can have one chip to process all the standards.

Kelf notes that cellular base stations already use a lot of software processing, but power consumption is high. He sees opportunities for low-power SDR in the emergence of femtocells, which can accommodate 300 or so cell phones within a 100-foot radius versus the 300 phones within a several-mile radius for a traditional base station. "Femtocells are a real market-changing technology," he says. Occasional-access devices also offer opportunities for SDR technology. Unlike a cell phone or a computer, such devices, including power meters and e-books, need only intermittent access to the cellular network. A power meter, for example, might report home-power usage every three months, whereas a GPS-navigation device might download traffic information at the beginning of a trip. An e-book needs cellular access only when its owner purchases a new book. All these devices have processors, which you can press into service to run the SDR algorithms. That feature can eliminate the neeed for a $10 or $20 baseband chip in a device whose total bill of materials may be only $40 to $50.

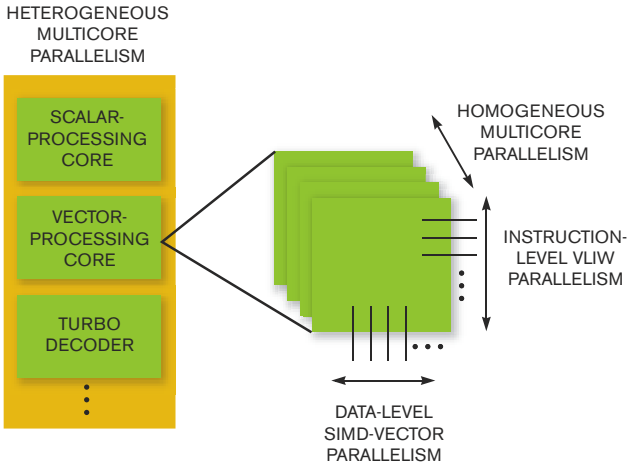

Parallelism using vector processing can speed the processing necessary to implement 4G baseband functions in SDRs (courtesy Sigmatix and OctoScope).

PARALLEL PROGRAMMING

OctoScope's Mlinarsky elaborates on the need for the flexibility that SDR affords. The challenge is to design devices requiring pervasive connectivity that can figure out what spectrum is available and renting—rather than buying—that spectrum in the white spaces, for example, for short-term use (**Reference 1**). "Programmability is becoming important, and that's where Sigmatix is positioning itself," says Mlinarsky, who has consulted for Sigmatix, as well as companies including Atheros and Intel. Intel has a history of trying to turn analog and RF problems into software problems that its processors can solve (**Reference 2**).

Designers can now implement baseband functions in FPGAs or ASICs that they once had to implement in software. With the introduction of 3G, WiMax, and LTE, designers started quickly catching on that a lot of baseband functions share some common structures for different protocols, for coding and decoding, for performing FFTs (fast Fourier transforms), IFFTs (inverse FFTs), and a slew of other functions. Such functions embody inherent parallelism and can run on parallel-processor architectures. "The issue is that humans are not very good at doing parallel programming because you have to keep hundreds or thousands of pieces of information in your head," Mlinarsky says. "Humans are not good at that, and compilers are much worse. Today, we have a world where Moore's Law doesn't permit further scaling of single processing engines sufficiently to get more computational capacity, so we are starting to see multicore PCs. But most applications don't take advantage of multicore processors because nobody knows how to program them. They sit there and warm up our laps."

The problem lies not in the computational hardware but in the algorithms. "The processing [for SDR] must be parallel due to the nature of baseband algorithms and because we are running out of Moore's Law capacity," she says, for further performance enhancements for single-core processors. "We can put a lot of cores in there cheap," Mlinarsky adds. "But how do we most efficiently program them in a way the human mind can deal with? You don't want to use superhuman programmers; you can't find enough of them. You need something between a compiler and the hardware to make the SDR architecture work." Sigmatix and, no doubt, many of the traditional wireless-chip suppliers will want to solve that problem. EDN

FOR MORE INFORMATION

Acer
www.acer.com

ARM
www.arm.com

AT&T
www.att.com

Atheros Communications
www.atheros.com

Bluetooth Special Interest Group
www.bluetooth.com/Bluetooth/SIG

Broadcom
www.broadcom.com

Google Android
www.android.com

Hewlett-Packard
www.hp.com

HTC
www.htc.com

Infineon
www.infineon.com

Intel
www.intel.com

Lenovo
www.lenovo.com

LG
www.lge.com

Marvell
www.marvell.com

Microchip
www.microchip.com

Microsoft
www.microsoft.com

Nokia
www.nokia.com

OctoScope
www.octoscope.com

Qualcomm
www.qualcomm.com

Samsung
www.samsung.com

Sigmatix
www.sigmatix.com

Sony Ericsson Mobile Communications
www.sonyericsson.com

ST-Ericsson
www.stericsson.com

T3G Technology
www.t3gt.com

Toshiba
www.toshiba.com

ZeroG Wireless
www.zerogwireless.com

REFERENCES

1 Dipert, Brian, "White spaces: ready for development permits or off-limits?" *EDN*, Aug 7, 2008, pg 36, www.edn.com/article/CA6582851.

2 Nelson, Rick, "Analog–that computes," *Test & Measurement World*, June 11, 2008, www.tmworld.com/article/321853-Analog_that_computes.php.

EXTREME EVOLUTION

Introducing the RX Family from Renesas.
The next-generation microcontrollers for communicating anytime, anywhere, and with anyone and anything.

- **Extremely high performance: 1.65DMIPS/MHz, up to 200MHz**
- **Extremely small code: up to 30% smaller than previous Renesas MCUs**
- **Extremely low power consumption: 30µA/MHz (typical) for the core**
- **Extremely high integration: with up to 4MB of extremely fast embedded Flash**

Stay tuned. The evolution continues.

Check the website for details.

www.america.renesas.com/ReachRX/a

Renesas Technology Corp.

Everywhere you imagine. RENESAS

BY THIERRY MASSON, LAURENT MONGE, AND A GLASCOTT-JONES • E2V

Tools accurately simulate noise in mixed-signal ASICs

IN APPLICATIONS REQUIRING HIGH PRECISION, GETTING NOISE ESTIMATES RIGHT MEANS UNDERSTANDING BOTH THE APPLICATION AND THE TOOLS.

In some capacitive-sensor applications, the ability to obtain noise levels as low as 10 aF (atofarad) rms in capacitive-sensor-interface ASICs is critical to high performance. To achieve such low noise levels, a mixed-signal-IC developer must have extensive knowledge of sensor characteristics, as well as an ability to accurately model the sensor-IC interface. Developers must use a top-down method, modeling each subblock at a high level to meet overall system-performance requirements. Developers use simulation tools for this task, and the selection of those tools depends on the level in the hierarchy, ensuring an efficient design process.

To achieve low noise levels, develop-

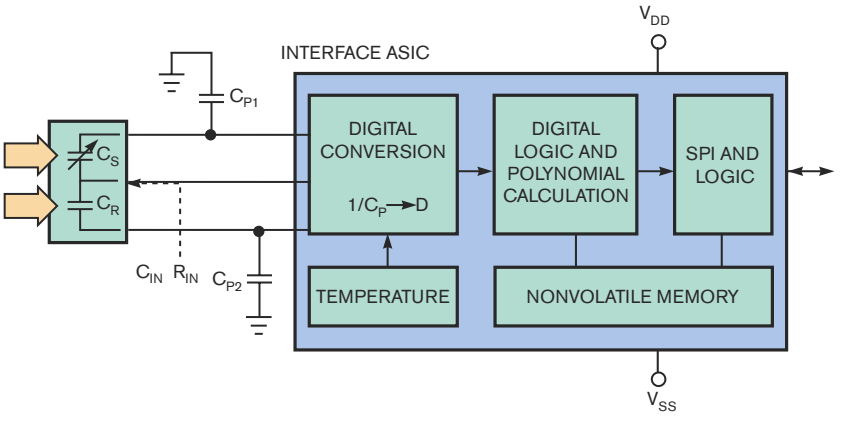

Figure 1 An architecture for a capacitive-sensor-interface IC performs a measurement of the sensor capacitor, C_S, and the related reference capacitor, C_R.

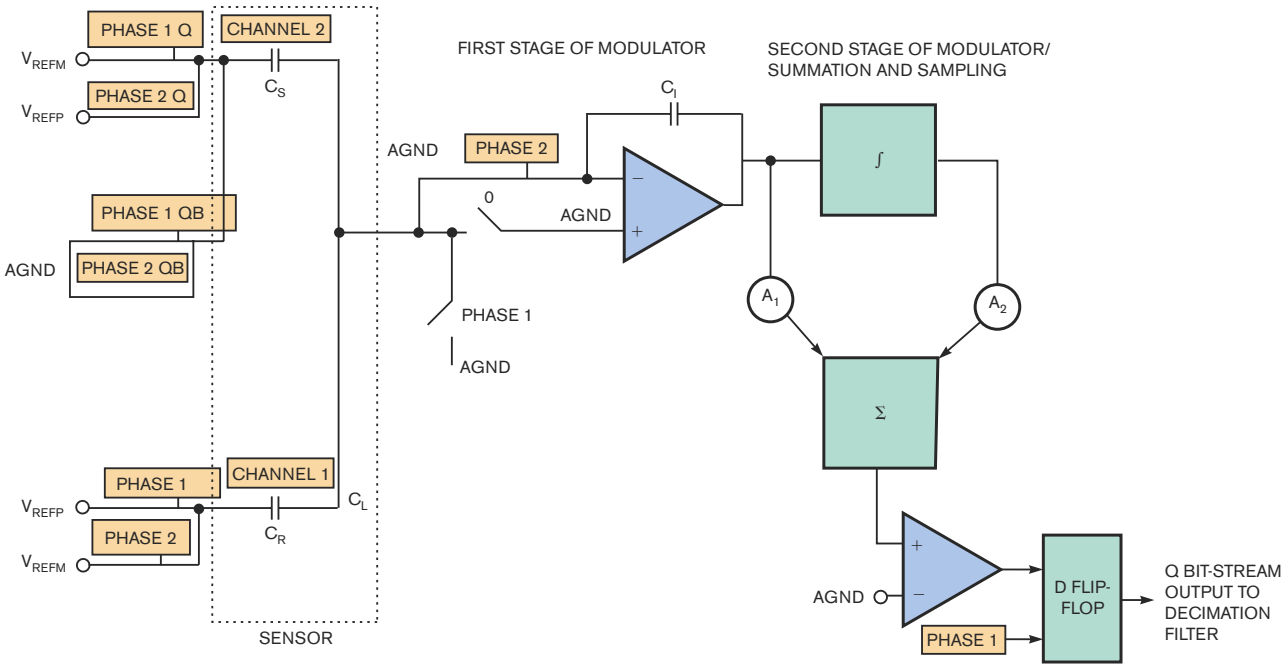

Figure 2 The input interface employs a switched-capacitor technique that merges the external capacitor to sense with the first stage of the delta-sigma modulator.

ers must vigorously examine and extensively verify the models the foundry supplies. Developers must also have precise knowledge of the simulation tool and its limitations. A knowledgeable mixed-signal-ASIC-development group also uses silicon-validated IP (intellectual property). A large portfolio of high-resolution sigma-delta converters; low-offset, low-noise amplifiers; and stable voltage references provides a basis for future designs and offers confidence that ASICs will meet required performance.

SENSOR-INTERFACE ASICs

An architecture for a capacitive-sensor-interface IC performs a measurement of the sensor capacitor, C_S, and the related reference capacitor, C_R, and can produce data for several input configurations, such as C_R/C_S, C_S/C_R, or, for differential capacitive sensors, $\Delta C/\Sigma C$ (**Figure 1**). This ability to produce different data can be useful because the effect of stray capacitance decreases for some configurations and improves linearity for others. Designers measure the external capacitor using a second-order, 16-bit delta-sigma ADC. The input interface employs a switched-capacitor technique that merges the external capacitor to sense with the first stage of the delta-sigma modulator (**Figure 2**).

A quick examination of this circuit shows that the digital output of the modulator is an image of the ratio C_R/C_S. This type of interface is most appropriate for capacitive MEMS (microelectromechanical-system) sensors, which often produce only a small capacitance change over their full range.

For example, a MEMS capacitive accelerometer produces only a 200-fF full-scale change in capacitance for a 1g acceleration. To achieve system requirements, these tiny input levels often demand that the noise level of the interface be on the order of 10^{-18}F. Simulation is essential at the start of the design to ensure that developers can meet these goals. One main challenge of this approach is predicting the final, obtainable resolution. Note that targeted figures are in the range of atofarads root mean square. A second major difficulty is sizing the various elements in the design.

NOISE SOURCES

The output noise of the measurement chain determines the final resolution. The sources of output noise come from both the quantization of the ADC and the thermal noise of electronics, including the sensor. The quantization noise depends on the delta-sigma-transfer-function coefficients and on the amount of post-filtering that the decimation filter provides. The oversampling-rate-ratio coefficient controls this second contribution. Depending on resolution, the oversampling-rate-ratio coefficient can be fixed at 128 to 1024. An additional difficulty that may arise early in the design is that some transfer-function coefficients are not fixed but depend on the sensing element itself.

Thermal noise comes predominantly from inside the first modulator stage. The main sources are switches and the operational amplifier in the first integrator. Because the delta-sigma converter uses sampled-data techniques, the converter

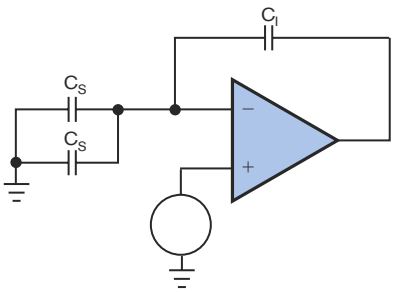

Figure 3 An op amp has several thermal-noise sources.

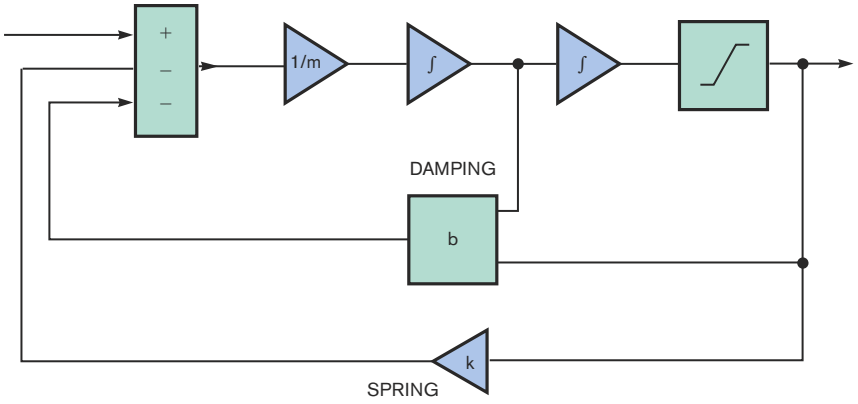

Figure 4 It is useful to perform a system-level simulation with Matlab using the parameters of the second-order modulator and the sensor.

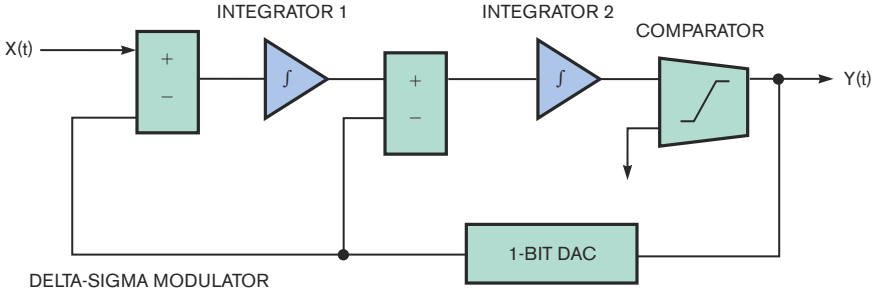

Figure 5 The delta-sigma modulator is a basic second-order design.

LISTING 1 VISUAL BASIC NOISE FUNCTION

```
Function gaussrand() As Double
  Dim a As Double
  a=Rnd()-Rnd()+Rnd()-Rnd()+Rnd()-Rnd()+
  Rnd()-Rnd()+Rnd()-Rnd()+Rnd()-Rnd()
  gaussrand=a
  End Function
```

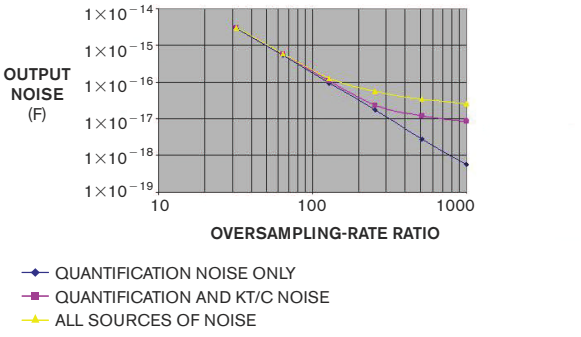

Figure 6 It is important to model all three noise sources in the capacitive interface.

aliases the sources of noise around the sampling frequency of the system. This aliasing in effect transforms the thermal noise of the resistance of switches to classic kT/C noise, where k is Boltzmann's constant, T is the temperature, and C is the capacitance. The same effect aliases the op amp's wideband thermal noise around the sampling frequency, but some gain depends on the sensing element (**Figure 3**).

SIMULATING SENSOR INTERACTION

To evaluate the system's ultimate performance, it is useful to undertake a system-level simulation with The MathWorks' (www.mathworks.com) Matlab using the parameters of the second-order modulator and the sensor. The example in **Figure 4** uses a MEMS accelerometer. The model can also incorporate nonideal behavior, such as sampling jitter and operational-amplifier noise. This top-level approach ensures the most appropriate architecture for the system and allows some upfront experimentation with the basic parameters, modulator order, and sampling rate.

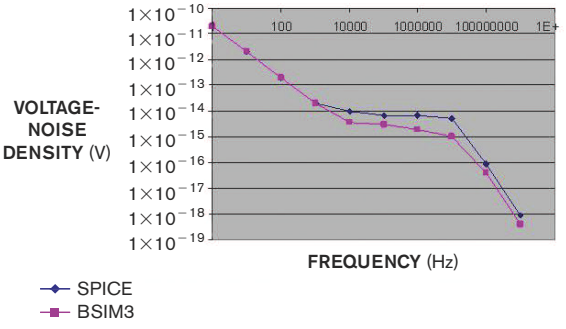

Figure 8 Different simulations give different results on the test bench.

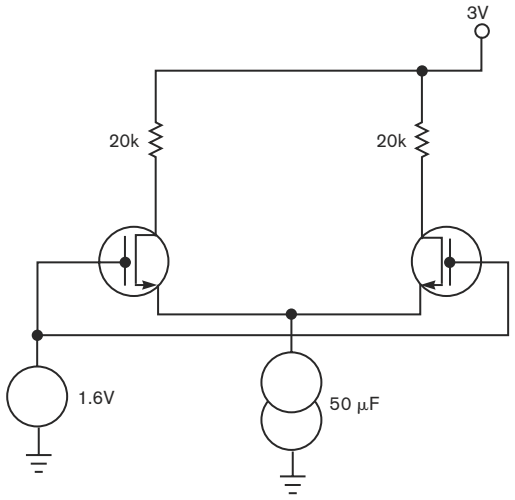

Figure 7 Use a test bench to get a reality check.

You can use any of several approaches to evaluating output noise and final resolution. The easiest to understand is to perform manual calculations. However, developers must cross-check this approach with simulations to validate the system. Developers can use popular tools, such as Matlab, or simpler programming languages, such as C++ or Visual Basic, in this system-noise simulation. Visual Basic inside Microsoft (www.microsoft.com) Excel works well as system-noise-simulation tool because it is available on most computers and allows for easy postprocessing of simulation results.

The simplest method is to simulate the delta-sigma system with difference equations (**Reference 1** and **Figure 5**). For instance, you can model a sampled-data integrator with one simple equation, $V_{OUT}(NT)=V_{OUT}((N-1)T)+V_{IN}(NT)$, where V_{OUT} is the output voltage, V_{IN} is the input voltage, N is a variable, and T is the period of the sampling frequency.

Designers add three levels of detail to the simulation model to perform noise simulation in this system: quantization of the ADC's noise, the sensor's and reference capacitor's kT/C noise, and the op amp's thermal noise. You inject thermal Gaussian noise using the sum of 12 uniform random numbers, which you create in the RND() function of the programming language. This process is simply an application of the central-

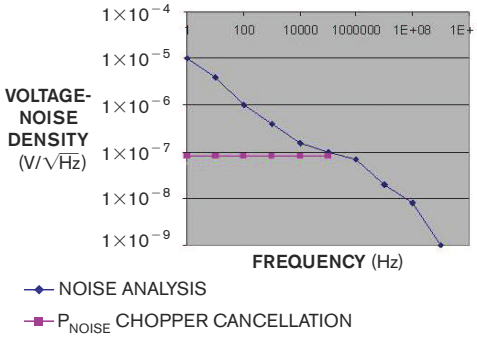

Figure 9 The chopper has a significant effect on noise simulation.

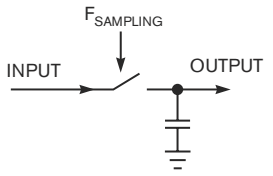

Figure 10 The model for sampling a capacitor is simple.

limit theorem, which states that the sum of many independent random variables tends to be distributed according to one of a small set of stable distributions. The Visual Basic function for producing Gaussian random numbers is simple (**Listing 1**).

Figure 6 shows noise simulation using different oversampling ratios with the three components of noise. When high resolution is necessary, the main source of noise is thermal noise, primarily from the operational amplifier. Using this tool, designers can determine the oversampling-rate-ratio coefficient, transfer-function delta-sigma coefficients, reference-capacitor size, and maximum total amount of input-amplifier noise that are mandatory to reach a given noise floor or resolution. Using these obtained values, designers can also derive other key parameters, such as sampling frequency, power consumption of the modulator, and the op amp's bandwidth and integrated input noise. Once the design team freezes the architecture, the transistor-level design can begin with a top-down approach. Noise simulation now moves to the transistor level.

TRANSISTOR-LEVEL SIMULATION

When starting to design an ASIC using a new process, it is important to check for well-modeled noise at the transistor level. Developers must be able to design low-noise structures with low power consumption. Therefore, they commonly use MOSFETs in a weak-inversion operating region, in which the gain of the transistor reaches a maximum value and the noise is low with a minimum amount of biasing current. However, the Cadence (www.cadence.com) Spectre simulator with standard Bsim3 (Berkeley short-channel-insulated-gate-field-effect-transistor 3) models often underestimates noise in this region. Therefore, you can use a small test bench to compare simulation results with those obtained using manual computation (**Figure 7**).

You can use differential pairs with two MOSFETs biased in weak inversion with large width and length and small biasing current. For example, each transistor has a current noise of 25 μA and a load of 20 kΩ. When bias is in this weak-inversion region, the MOSFETs' current-noise density is due only to shot noise. Manual computation of the total output-noise density yields $(VN0)^2=4qI_CR^2=6.4\times10-15V^2/\sqrt{Hz}$, where VN0 is voltage-noise density, q is charge, I_C is the transistor current, and R is resistance.

Using the standard parameters of the given process with Bsim3.2 models for thermal noise, the Spectre simulator gives an output noise of only $3.3\times10^{-15}V^2/\sqrt{Hz}$. By switching to Spice thermal noise, however, you can recover the hand-computed result (**Figure 8**). You must then check the parameter *noimod* to determine whether the manufacturer or the foundry partner has modified the Spice model and carefully check the flicker-noise modeling of the transistor in the process. This verification is necessary to ensure that flicker-related parameters are present in the parameter set of each transistor. After you have verified that the parameters are there, you should benchmark the flicker-noise results against results from other processes to check whether they are realistic.

CIRCUIT SIMULATION

Cadence tools, including Spectre, are appropriate for the design of analog- and mixed-signal ASICs to perform electrical and noise simulation. You can use classic noise analysis to evaluate, for instance, the op amp's integrated noise of the operational noise. To improve the amplifier's low-frequency noise, you can add chopper cancellation to the input stage to remove both flicker noise and offset. Classic noise analysis cannot predict the final output noise in chopper-stabilized designs, however. Therefore, designers use the periodic-noise analysis of SpectreRF to predict the noise at low frequencies (**Figure 9**). Using this simulation, you can establish that integrated noise is 30 μV over a 1-Hz to 1-kHz bandwidth without

TABLE 1 SIMULATION RESULTS

Sensor capacitance (pF)	Reference capacitance (pF)	Simulated sensor-capacitance noise (aF)	Measured sensor-capacitance noise (aF)
4.7	32	69	48
4.7	16	38	36
4.7	8	25	32

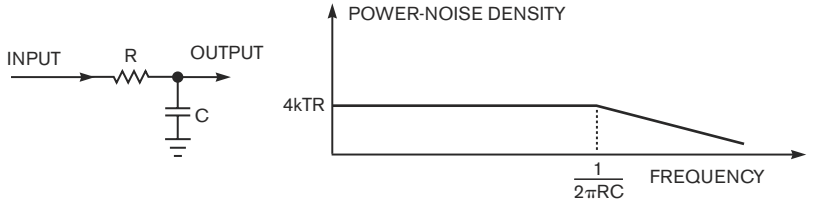

Figure 11 With the switch closed, the model is even simpler.

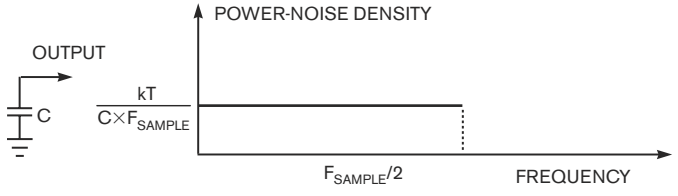

Figure 12 Opening the switch simplifies the model.

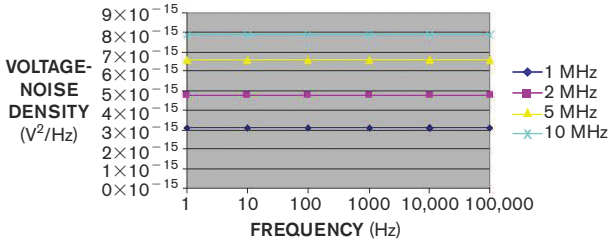

Figure 13 Bandwidth has a huge impact on simulated noise performance.

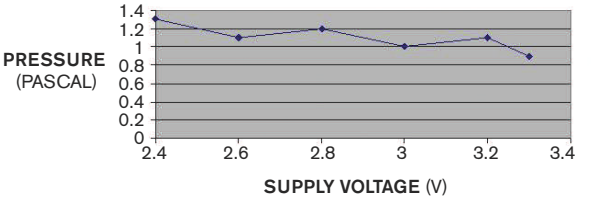

Figure 14 The pressure resolution of the sensor system is nearly 1 Pascal over the full operating-voltage range.

the chopper. With the chopper enabled, periodic-noise analysis shows a reduction to 2.3 μV. This analysis is efficient and is an important tool for securing final results.

You can also simulate the aliasing effect due to sampling using the SpectreRF option with periodic-noise analysis. To obtain solid results, however, you must compare the system's bandwidth with its sampling frequency. For instance, consider the basic effect of sampling an input signal in a capacitor (**Figure 10**). When the switch is closed, this small system is equivalent to the circuit in **Figure 11**, where R is the on-resistance of the switch and C is the sampling capacitor. The system is a lowpass filter with a cutoff frequency of ½π/RC. The total integrated noise is the well-known kT/C noise.

You can illustrate the fact that one-half of the integrated noise is below the cutoff frequency and the other half is above (**Figure 12**). After sampling, the noise is equally distributed into dc to $f_{SAMPLE}/2$ bandwidth with a density of $2\ kT/C/f_{SAMPLE}$. For instance, a resistor with a value of 10 kΩ and a capacitor with a value of 10 pF combine with a sampling frequency of 100 kHz to yield a cutoff frequency of 1.4 MHz. Total output noise, VN^2, is $(kT/C)=4\times10^{-10}=(20\ \mu V)^2$, and noise density, $VN0^2$, is $[kT/(C\times F_{SAMPLE}/2)]=8\times10^{-15}=(90\ nV/\sqrt{Hz^2}]$. This result means that accurate periodic-noise simulation of this small system must take into account at least 10 times the bandwidth of the system, even if the sampling frequency is comparatively low. In this case, you would need to use power-noise analysis over a 20-MHz bandwidth.

To illustrate the importance of this point, examine the simulation results of noise density of the sampling capacitor with 1-, 2-, 5-, and 20-MHz bandwidths of analysis (**Figure 13**). Using the time-domain mode of analysis for the simulation yields a theoretical noise density of $8\times10^{-15}\ nV^2/\sqrt{Hz}$. Manual calculations find noise analysis with a bandwidth of more than 20 MHz. Hence, to obtain reliable results, the bandwidth analysis should be 10 times the system's cutoff frequency. You must always use manual computations first and then compare the results of those computations with simulation results. With this method, you can have confidence in the noise-simulation results of the modulator input stage.

Using this approach, the simulated noise matches well with actual silicon measurements, which show a high level of performance (**Table 1**). Using the capacitive-sense circuit with a MEMS pressure sensor would give the system extremely fine pressure resolution (**Figure 14**). Note that resolution of 1 Pascal is equivalent to the change in pressure of a 10-cm increase in altitude. The combination of manual computation and simulation can accurately predict the behavior of sensing-interface ASICs. You can perform high-level modeling with difference equations in Visual Basic or C++ high-level languages and perform electrical-noise simulation at the transistor level with an RF simulator—even if the frequency of interest is near zero. Manual computation skills and experience in the application and technology are also key factors for development success. New design techniques and architectures mean that noise performance of better than 10 aF rms is now possible. EDN

REFERENCE

1 Nosworthy, Steven R, Richard Schreier, and Gabor C Temes, *Delta Sigma Data Converters*, ISBN: 0780310454, John Wiley & Sons Inc, Oct 14, 1996.

AUTHORS' BIOGRAPHIES

Thierry Masson is an analog- and mixed-signal expert designer at e2v (Grenoble, France), where he has worked for 23 years. His responsibilities include design of the company's mixed-signal ASICs and custom ICs for sensor-interface applications in automotive, industrial, and medical markets. Masson has a degree in engineering from École Nationale Supérieure de Télécommunications (Brest, France) and a doctorate from Grenoble University.

Laurent Monge is a design and development manager at e2v (Grenoble, France), where he has worked for one year. His responsibilities include the design of custom ICs for sensor-interface applications in automotive, industrial, and medical markets. He manages the development team that designs mixed-signal ASICs. Monge has a master's degree in electronics.

Andrew Glascott-Jones is an application engineer at e2v (Grenoble, France), where he has worked for one year. His responsibilities include the design of custom ICs for sensor-interface applications in automotive, industrial, and medical markets. He is responsible for the development kits that e2v offers customers to support them during the ASIC-development phase. Glascott-Jones has a bachelor's degree in physics and a master's degree in electronics.

BY BINDESH PATEL AND REX CHEN • SPRINGSOFT

Debugging and analysis with SystemVerilog test bench

SOME MINOR ADDITIONS TO OVM LIBRARIES CAN OPEN A NEW WORLD OF VISIBILITY FOR DESIGN VERIFIERS.

The verification component of SystemVerilog has dominated the rapid adoption of the language. The new verification syntax in the language allows for dramatic productivity gains in the verification cycle, which itself is now the major portion of the entire design and verification cycle. However, the debugging process of such environments brings about additional complexities. First, unlike with RTL (register-transfer-level) logic, no callback-driven "dumpvars"-type, dump-everything schemes exist for the object-oriented, software-like SystemVerilog verification constructs. Using these capabilities, engineers have access to the entire history of events that occur during a simulation. Once available, this data can drive basic capabilities, such as waveforms and overlay of the trace data onto familiar source code and schematic views, and automation, such as multicycle tracing to the root cause with one command or mouse click.

Second, it is unclear what "access to all the data" would mean in the first place for SystemVerilog object-oriented test-bench code, which creates and destroys objects at runtime. Would the code record local variables or object-creation time? More important, how would this information be useful to users even if it were available? It is clear that the dump-all paradigm that is applicable to RTL logic is not meaningful for test-bench code for these reasons.

Early adopters of SystemVerilog test-bench syntax used crude logging of useful information to text files using the built-in $display or sprintf constructs. The users then inspected the resulting text files during debugging and manually correlated the text-file data with the RTL-dumped data to get a picture of what happened. In the extreme case, users invoked the simulation in interactive mode because this mode can potentially allow for full introspection of objects. This approach was less than ideal, however, because interactively running a simulation raises obvious performance issues, is inefficient, and is not friendly to today's simulation-farm, batch-based environments.

More recent advances have leveraged the notion of logging by recording the user-specified information into a proper debugging database, such as SpringSoft's FSDB (Fast Signal Database) format. Once users load such a specialized database, they can employ an automated debugging system that allows them to analyze the resulting data in a waveform or a table view that synchronizes with the rest of the system (**Figure 1**).

With this structured logging approach, a user can specify the data to record using a special system task, $fsdbLog (**Figure 2**). This task can record types of data and class local variables, thereby providing maximum flexibility. This structured logging approach comprises a message label—"success" in the example—and other optional attributes or properties of that message, such as a severity value. This breadth gives a more flexible capability to users than the previous string-only text-based approach. Users can debug and analyze using the result-

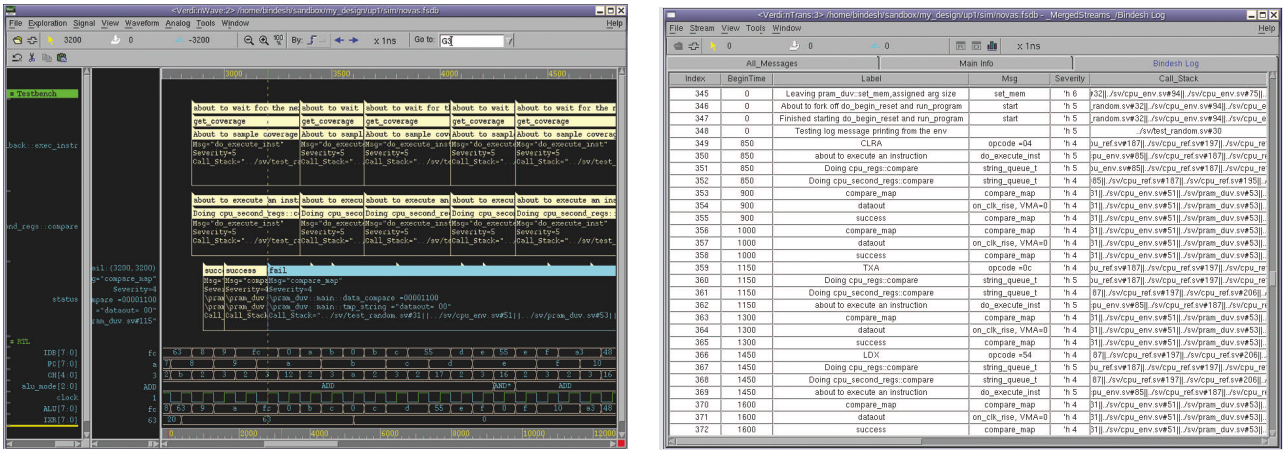

Figure 1 Capturing test-bench messages at runtime allows visualization and analysis of SystemVerilog test-bench messages alongside design waveforms.

ing messages within the same environment as the design RTL or gate-level dumped-signal data, thereby significantly boosting productivity. The availability of such a logging task that dumps data into a proper debugging database is a key foundation for further refinements in the flow and eventual automation.

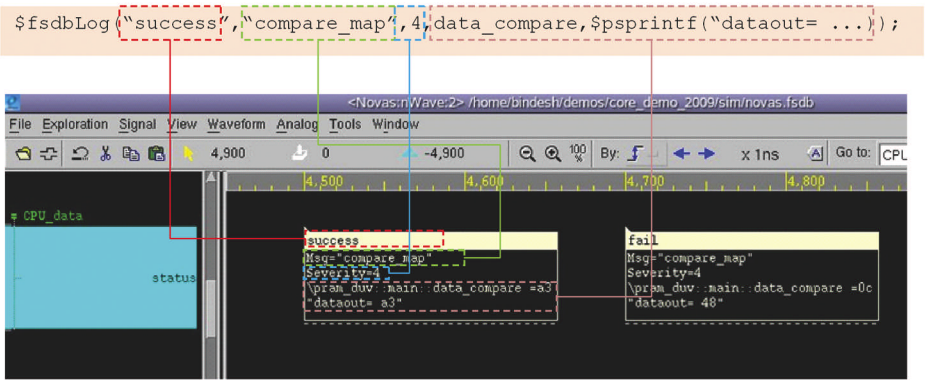

Figure 2 Message logging permits users to visualize the structure of captured data during runs.

LOGGING SCHEMES

Someone must be responsible for the instrumentation of the described logging tasks. The recording happens whenever a user encounters any such task during the execution of the test-bench program during simulation. The brute-force method is to replace or augment all uses of the built-in $display construct with a corresponding call to the logging task, $fsdbLog. Even this primitive use model leads to a more productive user experience than before. A further refinement would be to bury the logging code into the underlying OVM (Open Verification Methodology) or VMM (Verification Methodology Manual) libraries, which are foundation base-class libraries that most SystemVerilog verification environments use. **Listing 1** shows how you can add the logging task in **boldface** to the `vmm_error macro that you use to log error messages in environments employing the VMM library.

You could also extend the newer vmm_log_catcher and vmm_log_format classes with the logging instrumentation. Similarly, you could extend the OVM reporting functions (**Listing 2**) for the ovm_report_info function. Once users have performed this underlying coding, those who use these library macros or functions to record information will also be able to record into a proper debugging database with all the resulting benefits.

AUTOMATION

Using the basic $fsdbLog task to record messages into a proper database as a foundation, the next logical step is the complete automation of transaction recording in OVM- and VMM-based test benches. To achieve this goal, you can leverage the well-defined structure of OVM and VMM to watch transactions flowing up and down the test bench and to automatically record this activity.

Like VMM, OVM promotes a transaction-based approach to writing test benches. In this context, a transaction is a higher-level abstraction of data. The basic components of an OVM-based transaction-level verification environment are a stimulus generator or sequencer to create transaction-level traffic to the DUT (device under test); a driver to convert these transactions to signal-level stimuli at the DUT's interface; a monitor to recognize signal-level activity on the DUT's interface and convert it into transactions; and an analysis component, such as a coverage collector or scoreboard, to analyze the transactions (**Figure 3**).

VMM-based environments are similar. Regardless of the underlying library, the most interesting and useful data for engineers are the transactions between the sequencer and the driver and those between the monitor and the analyzers. You could apply these techniques to record the transactions of interest. The ultimate goal, however, would be to automatically record all the transactions between the sequencer and the driver. You can use the structure of OVM test benches to achieve this goal.

The OVM and its library require that the sequencer and the driver communicate with each other in specific ways. For example, the sequencer sends requests to the driver using the

LISTING 1 LOGGING ERROR MESSAGES

```
`define vmm_error(log, msg) \
do \
    if (log.start_msg(vmm_log::FAILURE_
TYP, vmm_log::ERROR_SEV)) begin \
        void'(log.text(msg)); \
        $fsdbLog("vmm_error_m", msg, 1,
    $psprintf("%m")); \
        log.end_msg(); \
    end \
while (0)
```

LISTING 2 EXTENDING OVM LOGGING FUNCTIONS

```
function void ovm_report_info(string id,
      string message,
      int verbosity = OVM_MEDIUM,
      string filename = "",
      int line = 0);
   string stream_name;
   ovm_top.ovm_report_info(id, message, verbosity,
   filename, line);
   $sformat(stream_name,"%s:%d",filename,line);
   $fsdbLog(id,message,verbosity,stream_name);
endfunction
```

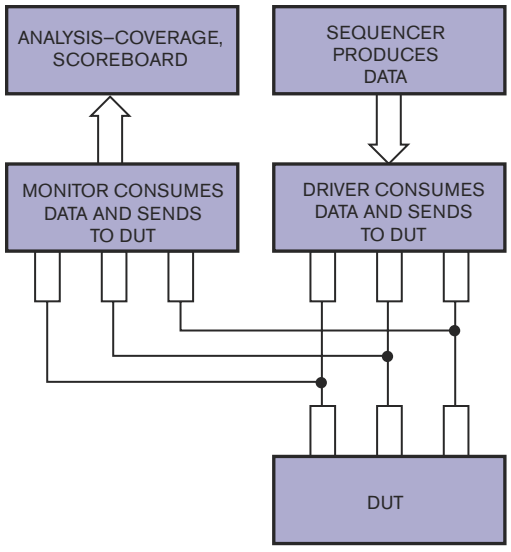

Figure 3 An OVM environment includes four basic elements.

ovm_sequence::send_request() function. This scenario is true even when you use the `ovm_do and `ovm_do_with macros to generate the sequence data because these macros call the same send_request() function. The driver itself sends the response to the sequencer using the ovm_seq_item_pull_port::put() or ovm_seq_item_put_port::item_done() function. It then follows that you can embed your logging function—sps_generic_log(), for example—in these OVM-standard functions and in the process be able to log all transactions into your debugging database. You can provide a modified OVM library with these changes for end users. This modified library can include the logging function sps_generic_log() in the ovm_sequence_item class. As a starting point, assume that you simply include a basic call to the $fsdbLog task in the ovm_sequence_item::sps_generic_log() to record a label, a string message, and a severity for each transaction (**Figure 4**).

You can provide hooks to map actual labels, strings, and severities pertaining to an environment using virtual functions that the user can then easily code. From a user's perspective, this additional coding is trivial and involves overriding a function, as **Listing 3** shows in **boldface**. The end result will be meaningful, intuitive transactions that the debugging system will capture and display (**Figure 5**). Even without these user modifications, the system would still record all the traffic between the sequencer and the driver, albeit with nonintuitive labels.

ADVANCED VISUALIZATION AND ANALYSIS

You can also instrument the modified OVM library to automatically log information connecting a sequencer that makes a request and a driver that sends back the corresponding response. This approach allows you to associate each request transaction with a response transaction. Although this information is useful to the user even in the waveform and table views, the availability of such complete information can inspire other views that can more intuitively show the testbench activity.

LISTING 3 ADDING VIRTUAL FUNCTIONS

```
class simple_item extends ovm_sequence_item;

          rand int unsigned addr;
              constraint c1 { addr < 16'h2000; }
          rand int unsigned data;
              constraint c2 { data < 16'h1000; }
          rand int unsigned flag;
              constraint c3 { flag < 3'b100; }

          `ovm_object_utils_begin(simple_item)
              `ovm_field_int(addr, OVM_ALL_ON)
              `ovm_field_int(data, OVM_ALL_ON)
              `ovm_field_int(flag, OVM_ALL_ON)
              `ovm_object_utils_end

          // new - constructor
     function new (string name = "simple_item");
              super.new(name);
          endfunction : new

          function string_queue_t sps_define_log-
      ging_info();
          string_queue_t string_queue;
          string s;

          if (flag==RESPONSE)
              $sformat(s, "Response");
          else if (flag==BUS_WRITE)
              $sformat(s, "BUS WRITE");
          else if (flag==BUS_READ)
              $sformat(s, "BUS READ");
          else
              $sformat(s, "Simple Item");
          string_queue.push_back(s); // Label
          $sformat(s, "");
          string_queue.push_back(s); // Message
          $sformat(s, "0");
          string_queue.push_back(s); // Severity
              return string_queue;
          endfunction

      endclass : simple_item
```

One avenue for pursuing this inspiration is the UML (Unified Modeling Language), which provides intuitive graphical representations of systems, including software systems. The UML sequence diagram shows components and the interactions between these components in their temporal, sequential order. Although UML's primary use has been to document a program or system, you can use a sequence diagram to represent the dynamic behavior of a SystemVerilog test-bench program. The components in this case would be the OVM or VMM, and the interactions would be transactions passing between these components.

You now have practical schemes for extracting all the transactions and even for getting the additional details about which components are providing the transactions and which components are receiving them. Capturing all this rich data in the debugging database during simulation can drive the sequence

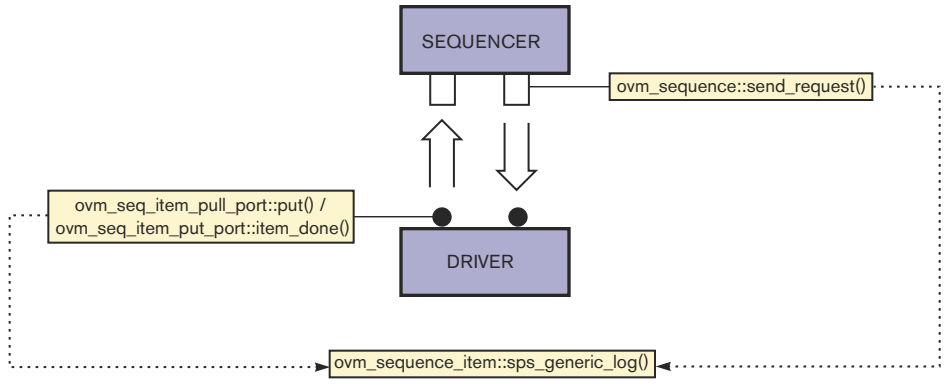

Figure 4 This OVM library automatically records all transactions between the sequencer and the driver.

diagram, thereby providing an additional, more intuitive and appropriate view of the test-bench behavior. A sequence-diagram view representing test-bench behavior working alongside and in synchronization with traditional hardware-behavior views, such as waveforms, can provide an ideal system for engineers to debug, understand, and analyze their entire design and verification environment. A prototype sequence-diagram view shows test-bench activity for a sample XBus transaction (**Figure 6**).

In a sequence view, the time axis is vertical, increasing downward. The components appear horizontally at the top, and the interactions or transactions between them appear as horizontal arrows at the respective simulation times. You can synchronize this view with other views, such as waveforms, so that the selection of a transaction causes the time cursor to move to the same time in the waveform. Other useful features can include tracing forward and backward. In the sample, tracing the "write" from the upper_sequencer causes the diagram to highlight the "write" and corresponding "response" between the sequencer and driver.

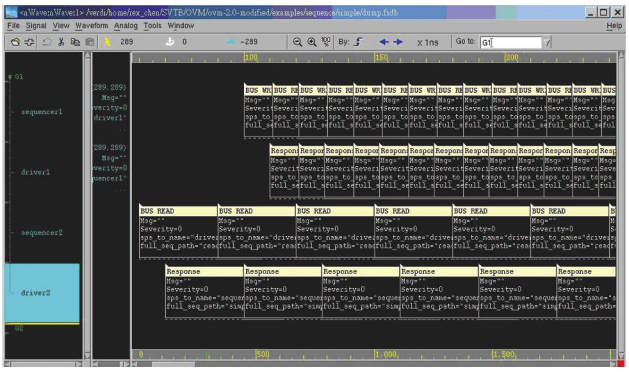

Figure 5 This sample waveform view shows transactions between the sequencer and the driver for an OVM test bench.

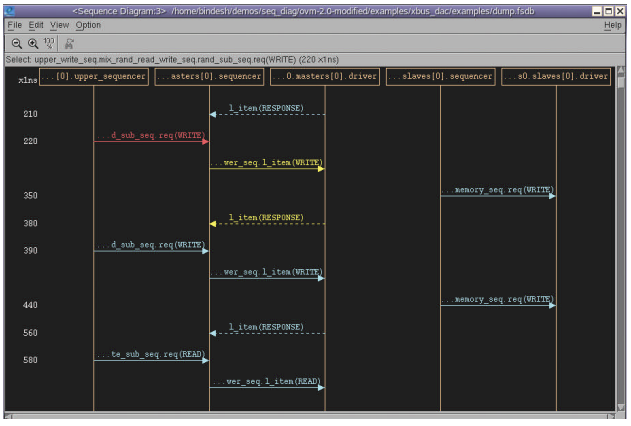

Figure 6 A prototype sequence-diagram view shows test-bench activity for a sample XBus transaction.

Using the foundation of a flexible logging function that records messages into a debugging database, you can instrument a modified OVM library to automatically log transaction-level activity between a sequencer that makes a request and a driver that sends back the response. This approach allows the debugging system to record each request transaction and associate it with its appropriate response transaction. Although this information is useful to the user even in the context of the waveform and table views, the easy availability of such complete data allows you to introduce a sequence-diagram view that can more intuitively show the test-bench activity.

Future work will focus on refining the use model for automatic logging so that you can eliminate the need to provide a modified library. This advance would remove dependencies in keeping in lockstep with OVM or VMM library revisions. Another focus area is to explore other UML representations, such as the component diagram for design and verification. EDN

AUTHORS' BIOGRAPHIES

Bindesh Patel is a technical manager for the verification group of SpringSoft USA (San Jose, CA), where he has worked for eight years. In his current position, he provides direction on future capabilities and products by engaging with customers, partners, and industry organizations, such as Accellera. He also touts the benefits of new methods, especially those that pertain to debugging. Patel has a bachelor's degree from the University of California—Santa Cruz. His personal interests include proving that baseball is a version of cricket for those who want it easy, along with hiking, camping, traveling, reading, and playing volleyball. You can reach him at bindesh_patel@springsoft.com.

Rex Chen is a technical manager of R&D at SpringSoft Inc (Hsinchu, Taiwan), where he has worked for nine years. In his current position, he provides research in verification and debugging methods and technologies. Chen has a master's degree in computer science from National Tsing Hua University (Hsinchu, Taiwan). His personal interests include photography. You can reach him at rex_chen@springsoft.com.

Buck. Boost. Both.

Actual Size

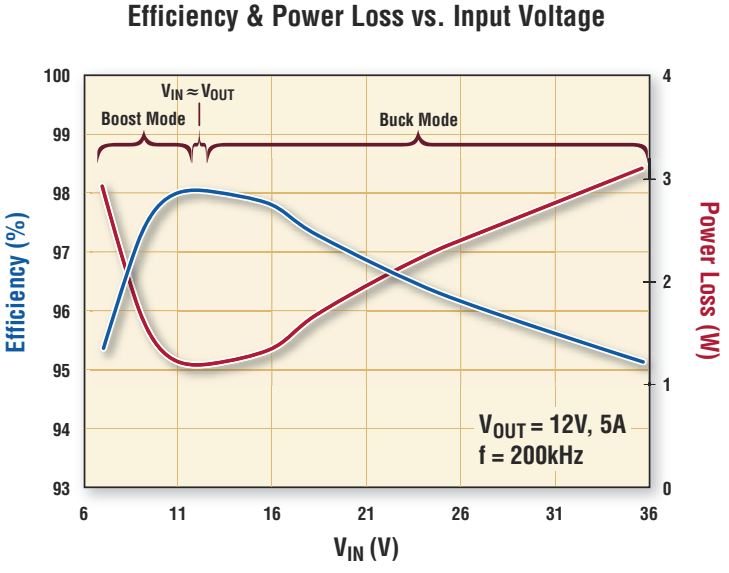

Simplify High Power Conversion with Buck-Boost DC/DC µModule Regulators

These new DC/DC µModule® regulators do three things well: buck, boost and buck-boost. The LTM®4605, LTM4607 and LTM4609 are pin compatible with identical layout and cover voltages up to $36V_{IN}$ and $34V_{OUT}$. Simply parallel them to increase the output power. They ease board layout by housing a DC/DC controller IC, four synchronous MOSFETs and compensation circuitry in a compact 15mm x 15mm x 2.8mm surface mount package.

Pin-Compatible DC/DC Buck-Boost µModule Regulators

Part Number	LTM4609	LTM4607	LTM4605
V_{IN}	4.5V to 36V	4.5V to 36V	4.5V to 20V
V_{OUT}	0.8V to 34V	0.8V to 24V	0.8V to 16V
I_{OUT}*	3A to 7A		
Output Current Sharing	Up to Four Devices for More Output Power		
Synchronous Current Mode Architecture	Yes		
PLL	200kHz to 400kHz		
Pin-Compatible Package	15mm x 15mm x 2.8mm 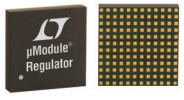		

*Can be adjusted depending on external inductor. See data sheets.

Info & Free Samples

www.linear.com/BuckBoost

1-800-4-LINEAR

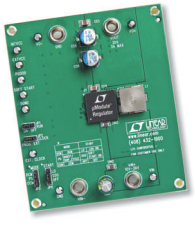

Purchase Demo Board

LINEAR TECHNOLOGY

EDITED BY MARTIN ROWE
AND FRAN GRANVILLE

designideas

READERS SOLVE DESIGN PROBLEMS

Spice simulators provide behavioral sources to mere mortals

William H Morong, Paoli, PA

Spice simulators include three types of tools for voltage and current sources: independent, dependent, and behavioral. Independent sources have two terminals and provide a specified number of volts or amperes you enter as a constant, just as a bench supply generates a set voltage or current. These simulators, like signal generators, also provide waveforms.

Figure 1a shows two independent sources. Voltage source V_1 provides 5V dc, and V_2 provides a 1V-peak sine wave. Dependent sources have at least four terminals. A first terminal pair is a voltage or current output. Another terminal pair is an input that attaches in shunt or in series with two circuit nodes. The output responds to the voltage across or current through the input nodes. Some dependent sources have multiple input-terminal pairs. The output responds to the inputs according to some linear mathematical rule. For example, a voltage-dependent voltage source set to the constant 100, such as E_1 in **Figure 1a**, is an ideal voltage amplifier with a gain of 100.

Behavioral, or arbitrary, sources are the least-used but most powerful of these sources. They have only an output-terminal pair, but they are more powerful than their simpler counterparts. They can implement a set of mathematical functions that roughly correspond to those available on scientific calculators. Their outputs can be independent or dependent. In the simplest case, a behavioral voltage source set to V=5 produces 5V, just as a power supply does.

A behavioral source can also respond to a designated voltage or current somewhere in the simulated circuit. In **Figure 1b**, B_1 responds to the voltage at the input node. Applying

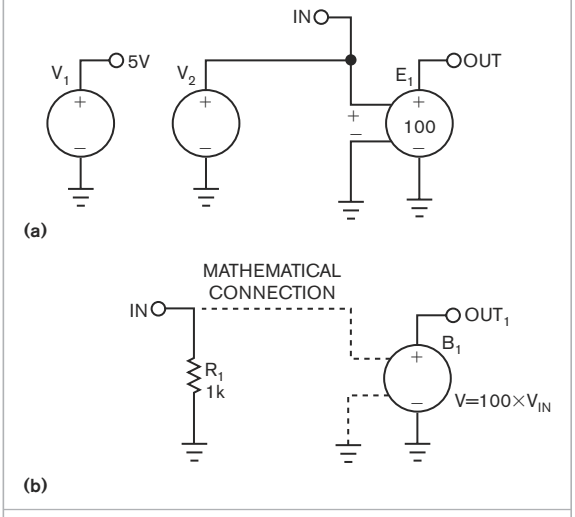

Figure 1 Independent voltage source V_1 provides 5V dc, and V_2 provides a 1V-peak sine wave (a). Behavioral source B_1 responds to the voltage at the input node (b).

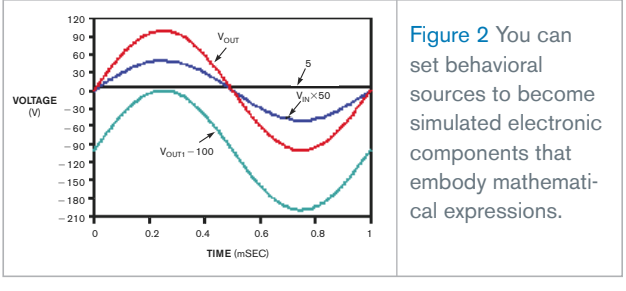

Figure 2 You can set behavioral sources to become simulated electronic components that embody mathematical expressions.

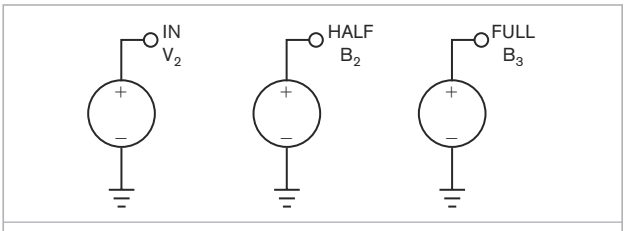

Figure 3 Behavioral source B_2 amplifies and half-wave-rectifies the voltage at the input node. Behavioral source B_3 also amplifies and full-wave-rectifies the same voltage.

the **equation** beside B_1 sets the behavioral source to $V=100\times V_{IN}$ and amplifies the voltage across R_1 by 100V.

In this case, the behavioral source behaves as the dependent source when the input node and ground drive it, but the behavioral source's input connection is purely mathematical. **Figure 2** shows several possible plots for output signals. You can set behavioral sources to become simulated electronic components that embody mathematical expressions. In the preceding example, if you set the behavioral source to $V=uramp(100\times V_{IN})$, the amplified output is ideally half-wave-rectified. If you substitute the abs function for uramp, you get full-wave rectification.

In **Figure 3**, behavioral source B_2 amplifies and half-wave-rectifies the voltage at the input node. Behavioral source B_3 also amplifies and full-wave-rectifies the same voltage. **Figure 4** shows the results. Engineers commonly generate ABMs (analog behavioral models) for Spice simulations to represent entire blocks of analog functions. Sometimes you need simple models when none are available, but you can use behavioral models to solve those problems. For example, you can set $I=uramp(V_{FB}-1.25)$ as a behavioral current source to draw 1 mA/mV at the FB (feedback) node that exceeds 1.25V (**Figure 5**). The behavioral source looks like an idealized three-terminal shunt regulator (**Figure 6**). If you need more refinement, you can set the feedback node to act as an RC filter or another function.

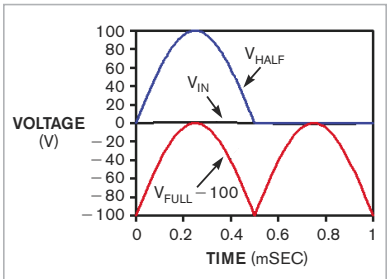

Figure 4 Behavioral source B_2's amplification and half-wave-rectification of the voltage at the input node and B_3's amplification and full-wave rectification of the same voltage yield these results.

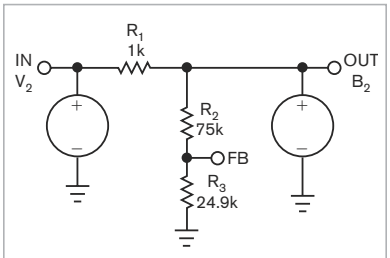

Figure 6 The behavioral source looks like an idealized three-terminal shunt regulator.

Because behavioral sources can embody mathematical functions, you can use them to test the mathematical viability of a design before simulating the circuit. Long mathematical expressions can be difficult to read, so you should break the idea into small blocks and set a behavioral source to implement each block. Just because your idea works in math, though, doesn't mean that its behavioral model is a practical circuit. In Spice, it is easy to make a model of a current source with 1 MV of near-perfect compliance, sometimes without intending to do so, but most practical engineers would rather defer the realization of such a circuit.

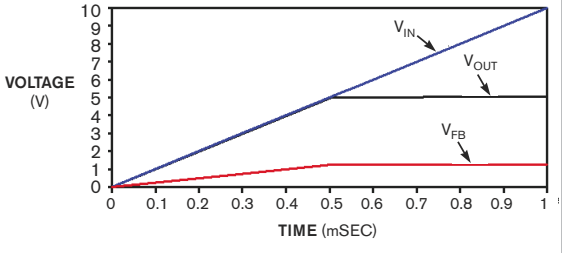

Figure 5 You can set $I=uramp(V_{FB}-1.25)$ as a behavioral current source to draw 1 mA/mV at the FB (feedback) node that exceeds 1.25V.

The caution to be realistic applies to Spice in general and to behavioral sources in particular. When you draw two equal resistors in your Spice schematic, they will be perfectly matched in simulation. The resistors in your stockroom are not that good. EDN

Multidecade BCD DAC uses resistors of only six values

Marián Štofka, Slovak University of Technology, Bratislava, Slovakia

A previous Design Idea uses a three-decade BCD (binary-coded-decimal) DAC to precisely set the output current of a current source (**Reference 1**). The circuit acts as a code-to-conductivity converter. The values of resistors of this DAC are staggered by powers of two within any of the decades, and the values of resistors at corresponding bits of the decades are staggered by powers of 10. Thus, the circuit needs 12 values of resistors, ranging from 125Ω to 100 kΩ.

In comparison, the circuit in this Design Idea enables you to construct a BCD DAC using only six resistor values, regardless of the number of decades. Moreover, these six resistor values vary within a relatively narrow 1:8 range. The voltage-output DAC operates ratiometrically. That is, if the temperature coefficients of the resistors are approximately the same—and you can assume they will be within this narrow range of values—then the variation of resistance with temperature has almost no detrimental effect on accuracy. This situation is not true, however, for code-to-conductivity DACs, in which the temperature coefficient of the resistors directly influences the temperature coefficient of the DAC.

LINEAR TECHNOLOGY

DESIGN NOTES

µPower Buck Regulator Safely Rides Through Automotive Load Dumps – Design Note 475

Juan G. Aranda

Introduction

When a buck regulator is used in applications where output short-circuit and overload conditions occur, the duty cycle required to keep the inductor current under control can be lower than the regulator's minimum duty cycle when operating at full speed. Therefore, to effectively protect a traditional regulator under such conditions, its switching frequency must be reduced to a speed that can safely handle the maximum expected input voltage.

In some cases, frequency foldback can help reduce the effective duty cycle by reducing the switching frequency as the output voltage falls out of regulation. However, this technique might not provide enough protection if the *folded* frequency is not low enough. In the end, the duty cycle problem imposes a limit on the maximum switching frequency at which the regulator can be safely operated, especially in automotive applications where the input supply can see positive voltages several times higher than the normal 12V operating voltage.

The LT®3682 is a new 1A buck regulator that overcomes the duty cycle limitation by monitoring the current through the external catch diode, and delaying the generation of new switch pulses if this current exceeds a defined value. The LT3682 safely accommodates output shorts and overload conditions up to its maximum adjustable operating frequency of 2.2MHz regardless of input voltage. This added level of protection allows automotive systems designers to take advantage of the maximum switching frequency of the regulator without concern for transients on the input supply.

The LT3682 accepts input voltages from 3.6V to 36V and transients up to 60V. Additional features such as soft-start, power good flag, frequency foldback and thermal shutdown are all included in the thermally enhanced 12-lead 3mm × 3mm DFN package.

Minimum Switch On-Time

Under normal load conditions the internal switch current limit is regulated to meet the required peak inductor current. Due to internal delays, the power switch does not turn off immediately when the internal current limit is reached, but instead takes the minimum on-time ($t_{ON(MIN)}$) to do so. This delay allows the inductor current to continue rising to values that depend on the current slope and the value of $t_{ON(MIN)}$. During positive input voltage transients the slope of the rising inductor current is greatly increased, thus raising its peak value. A well designed buck converter must skip switch pulses to maintain regulation if the new duty cycle requirements cannot be met due to minimum on-time limitations.

Overload conditions during the input transients may aggravate the situation further. Under extreme overloads the internal current limit is ultimately clamped to its maximum value, I_{LIM}, and the output voltage falls out of regulation, reducing the negative slope of the inductor

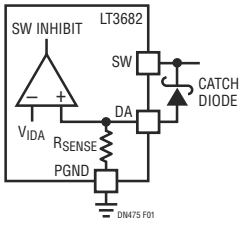

Figure 1. The LT3682 Monitors the Current Flowing Through the External Catch Diode at the DA Pin

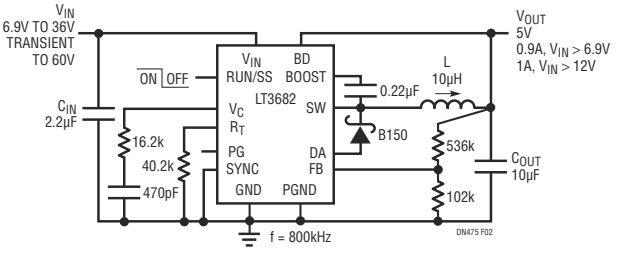

Figure 2. A Typical 800kHz Application

current accordingly. If during the switch off-time the inductor current does not return to the same or smaller values it had at the end of the previous cycle, its peak value will ratchet higher in every cycle and increase to unacceptably high levels.

Monitoring Current Via the DA Pin

The LT3682 constantly monitors the inductor current during the switch off-time by looking at the current flowing through the external catch diode via the DA (diode anode) pin (see Figure 1) and delays the generation of new switch pulses if this current does not fall below a defined threshold, I_{DA}, thus reducing the regulator's effective duty cycle.

Thus, it is now possible to use the small footprint inductors found in high frequency applications without sacrificing robustness during a number of fault conditions. Figures 2 and 5 show how to configure the LT3682 in a $5V_{OUT}$ application with a 800kHz and 1.7MHz programmed switching frequency, respectively. The resistive load at V_{OUT} is increased until the regulator hits its maximum current limit. Figures 3, 4, 6 and 7 show the DA pin current sense protection for input voltages of 12V and 36V for both applications. In all cases the lowest value of the inductor current is pinned to about 1.1A, which keeps its peak value well under control. By delaying the generation of new switch pulses the switching frequency is effectively reduced to satisfy the new duty cycle requirements introduced by the fault condition.

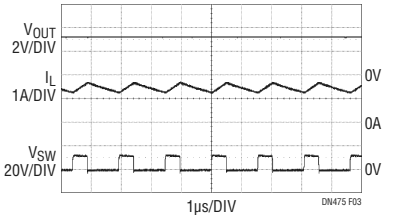

Figure 3. The 800kHz Application with V_{IN} = 12V. Overload Condition Forces V_{OUT} to Drop to About 3.2V

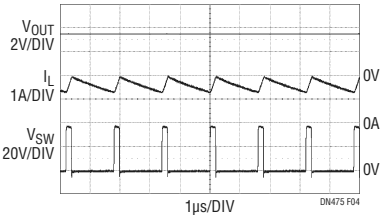

Figure 4. The 800kHz Application with V_{IN} = 36V. Overload Condition Forces V_{OUT} to Drop to About 3.5V

Conclusion

The LT3682 is a 1A monolithic buck switching regulator that accepts input voltages from 3.6V to 36V and transients up to 60V. It features an adjustable and synchronizable switching frequency from 250kHz to 2.2MHz. It also has the ability to monitor the current flowing through the external catch diode, thus providing an extra level of protection against output fault conditions over the entire operating frequency range, regardless of input voltage. These features, together with its typical 75µA no load quiescent current makes the LT3682 the right choice in high frequency automotive and battery-powered applications.

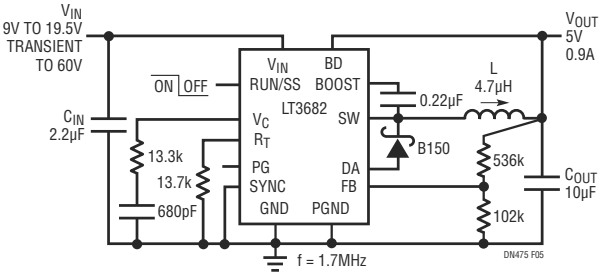

Figure 5. A Typical 1.7MHz Application

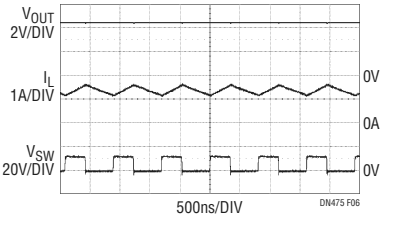

Figure 6. The 1.7MHz Application with V_{IN} = 12V. Overload Condition Forces V_{OUT} to Drop to About 4.4V

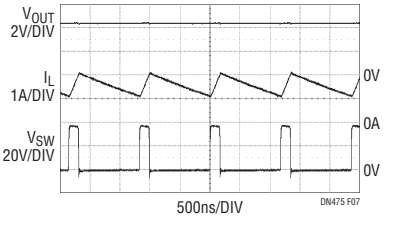

Figure 7. The 1.7MHz Application with V_{IN} = 36V. Overload Condition Forces V_{OUT} to Drop to About 4.4V

Data Sheet Download

www.linear.com

For applications help,
call (408) 432-1900, Ext. 3513

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

dn475 LT/TP 0210 116K • PRINTED IN THE USA

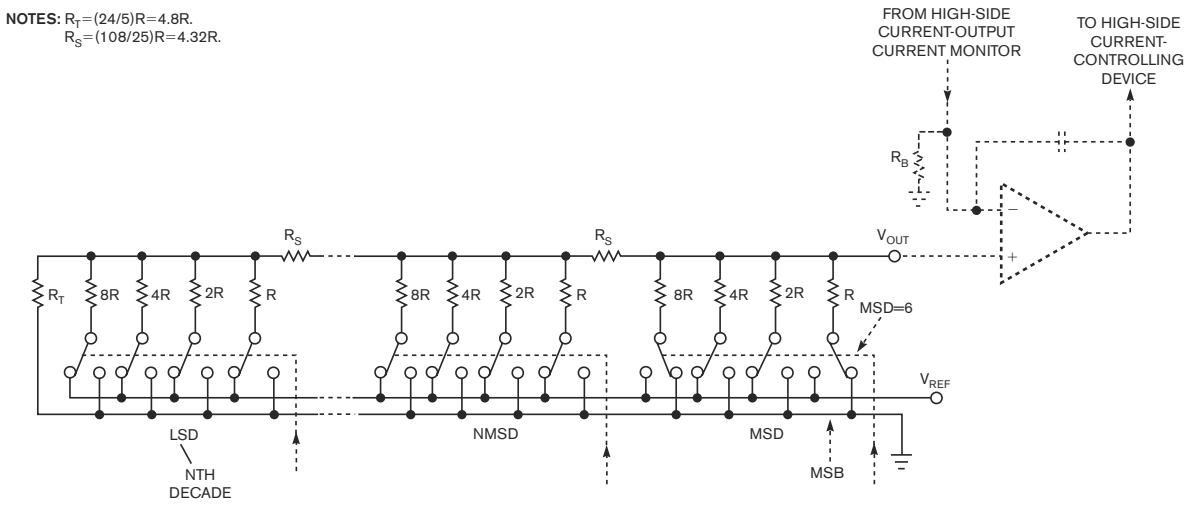

Figure 1 A voltage-output BCD DAC needs just six precision-resistor values. Resistors connect either to ground or to the reference voltage.

Figure 1 shows the voltage-output BCD DAC. The values of resistors are staggered by powers of two within each decade. The values of resistors at corresponding bits of the decades are of equal values. The switched ends of the resistors connect to either ground at logic zero or to the reference voltage at logic one. The voltage-output DAC thus has an advantage in that all the resistors' ends are at a defined potential. In a code-to-conductivity converter, on the other hand, one end of the resistor remains open at logic zero, and these open ends might act as capacitive sensors or even antennas, which could introduce additional errors. The common ends of four resistors in four bits of the MSD (most-significant decade) form the output. The common outputs of resistor quads in the less-significant decades successively connect to the main output through the series resistors, R_S, which all have the same value. Thus, R_S=(108/25)R=4.32R, where R is the value of the resistor at the MSB (most-significant bit) of any of the decades.

The common ends of bit resistors in the LSD (least-significant decade) connect to ground through terminating resistor R_T. This resistor represents the theoretically infinite number of decades having weight lower than the actual LSD, whereas these hypothetical decades are all set to zero. Thus, they contribute no voltage at the output. They do, however, influence the properties of the resistive network. R_T sets this influence and is equal to (24/5)R, or 4.8R. The full-scale output of the voltage-output BCD DAC is $3/5\times(1-10^{-N})V_{REF}$, where N is the number of decades—in this case, three.

To exploit a voltage-output BCD DAC in a single-supply programmable-current source, connect the output of the voltage-output BCD DAC to the noninverting input of an op amp, which accepts input voltages as low as 0V. The inverting input of the op amp connects to ground through resistor R_B, which has the value of $(3/5)\times(V_{REF}/10^{-2}A)$. EDN

REFERENCE

1 Diószegi, Gyula, and János Nagy, "Current-sense monitor and MOSFET boost output current," *EDN*, May 28, 2009, pg 44, www.edn.com/article/CA6659409.

Converter translates Bayer raw data to RGB format

Yu-Chieh Chen and Tai-Shan Liao,
National Applied Research Laboratories, Hsinchu, Taiwan

CMOS image sensors include the color filters of an RGB (red/green/blue) Bayer array, which lets the sensor detect colors. The image data, the output from the image sensor, is Bayer raw data (**Figure 1**). Unfortunately, most consumer-grade image-displaying devices require an RGB-image-data format with red, green, and blue in each pixel's data. Therefore, you often need a Bayer-raw-data-to-RGB converter between an image sensor and a displaying device. This Design Idea describes such a converter in Verilog HDL (hardware-description language). You can implement the code, available from the online version of this Design Idea at www.edn.com/100204dia, into a CPLD or an FPGA.

To make the design easy to understand, the RGB data is only 24 bits deep. A 320×240-pixel test-bench pattern verifies the design (**Figure 2**). The image data for red, green, and blue are 88h, 66h, and 22h, respectively. **Figure 3** shows the timing of the converter, and **Figure 4** shows the flow

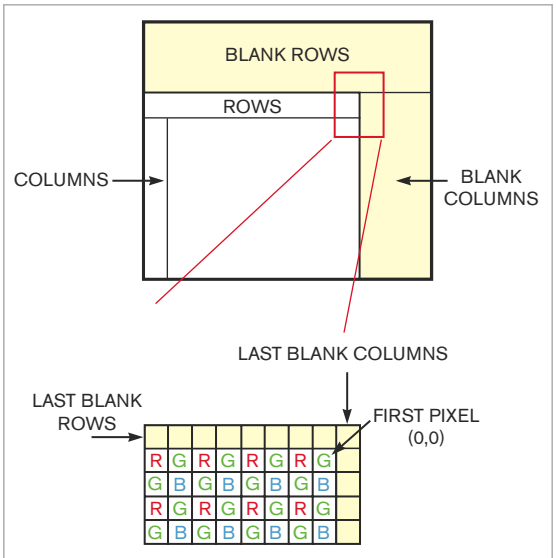

Figure 1 The Bayer raw data from an image sensor represents a 320×240-pixel array.

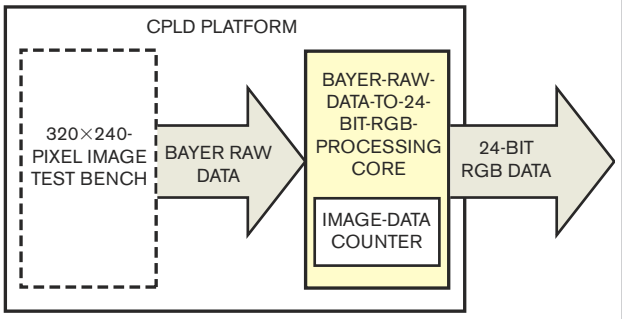

Figure 2 A CPLD processes Bayer raw data into 24-bit RGB data.

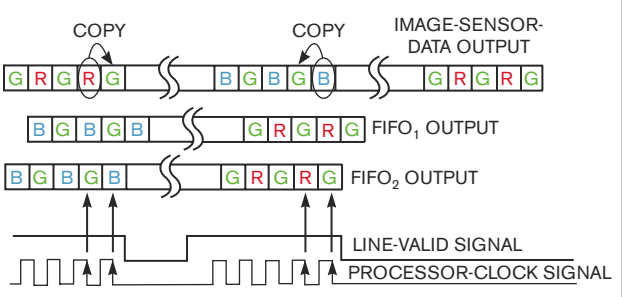

Figure 3 Rising edges of a clock latch Bayer raw data from FIFOs into RGB data.

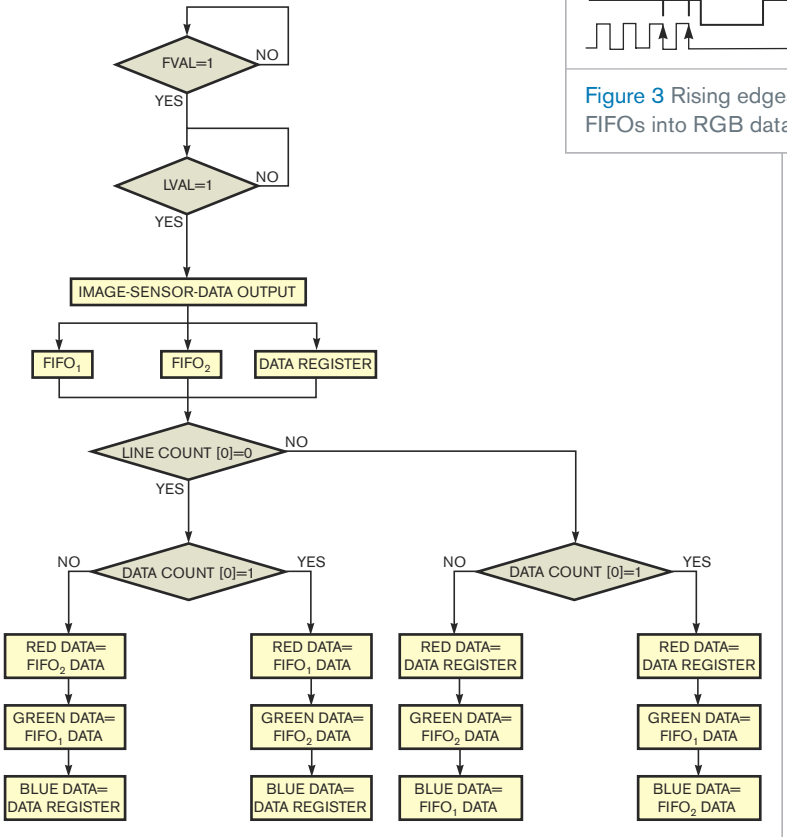

Figure 4 RGB data represents red, green, and blue for display.

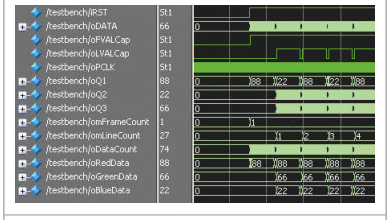

Figure 5 The values 88h, 66h, and 22h, respectively, represent red, green, and blue data.

chart. In **Figure 3**, the PCLK (processor-clock) signal's rising edge latches the 24-bit RGB data. The LVAL (line-valid) cap signal is the synchronized line-valid signal for reading the data. The process begins when the FVAL (frame-valid)-data signal goes high. When that action occurs, sensor data begins writing to $FIFO_1$ (first-in/first-out), $FIFO_2$, and the data register. After saving the data, $FIFO_1$ reads out the data. $FIFO_2$ and the data register read their data at the same time. Those readouts occur one clock cycle

Maxim's bidirectional level translators seamlessly interface between ICs

Level translate without a direction pin and save board space

1-Wire®

Memory cards

I²C

Multipurpose, high speed

SPI™

LCD displays

Low-voltage operation
- V_L = 1.2V to 5.5V, V_{CC} = 1.65V to 5.5V

Save board space compared to discretes
- Small UCSP™, µDFN, and UTQFN packages

Save processor overhead
- Bidirectional operation with no direction pin

Robust protection
- ±15kV ESD protection* (V_{CC} pins)

Save power
- Ultra-low, 1µA supply current in shutdown mode

Part	I/O Channels	V_L Supply (V)	V_{CC} Supply (V)	Data Rate (Mbps, max)	I/O V_L Shutdown State	I/O V_{CC} Shutdown State	Package (mm x mm)
MAX13046E	1	1.1 to 3.6	1.65 to 5.5	8	High-Z	High-Z	6-µDFN (1 x 1.5)
MAX13047E	2	1.1 to 3.6	1.65 to 5.5	8			10-UTQFN (1.4 x 1.8)
MAX13042E	4	1.62 to 3.2	2.2 to 3.6	100			12-UCSP (1.5 x 2.1)
MAX13030E	6	1.62 to 3.2	2.2 to 3.6	100			16-UCSP (2 x 2)
MAX3002E	8	1.2 to 5.5	1.65 to 5.5	20		6kΩ to GND	20-UCSP (2 x 2.5)
MAX13103E	16	1.2 to 5.5	1.65 to 5.5	20		High-Z	36-UCSP (3 x 3)

1-Wire is a registered trademark and UCSP is a trademark of Maxim Integrated Products, Inc.
SPI is a trademark of Motorola, Inc.
*Human Body Model.

www.maxim-ic.com/MAX13046-info

MAXIM DIRECT™
www.maxim-ic.com/shop

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

MAXIM
INNOVATION DELIVERED™

For free samples or technical support, visit our website.

after $FIFO_1$ starts reading. The next 24-bit RGB data remains the same in the data register, and it combines the data from $FIFO_1$ and $FIFO_2$, which read out at the same rising edge of the clock. The line-count signal shows whether the data is even or odd, which influences the combinational sequences of the data that reads out from $FIFO_1$, $FIFO_2$, and the data register.

Figure 5 is the ModelSim simulation waveform of the converter. The 24-bit RGB output-data values 88h, 66h, and 22h are red, green, and blue data, respectively. The **figure** shows the 24-bit RGB data as having red, green, and blue values of 88h, 66h, and 22h, respectively, during every line-data period. The line-data period matches the default pixel value in the image data's test-bench pattern. EDN

Drive 12 LEDs with one I/O line

Charaf Laissoub, Maisons Alfort, France

Several Design Ideas expand the I/O of a pin-limited microcontroller (**references 1** through **4**). The circuit in this Design Idea uses an alternative method (**Figure 1**). It limits external additional parts to one IC, and it drives as many as 12 LEDs in dot-bar or bar-graph mode. You can use the same technique in a dot-bar design (**Figure 2**). If you need a seven-segment LED display, you can use the circuit in **Figure 3**, which shows how to rearrange the circuit according to a classic multiplexed, four-digit common-cathode display. The prototype display uses Kingbright's (www.kingbright-led.com) SC52-11EWA high-efficiency LEDs, which emit 2000 to 5600 µcd at a forward current of 10 mA. The driver is a 12-stage NXP (www.nxp.com) 74HCT4040 binary counter or a 74HC4040 version for a lower power supply.

Listing 1, which you can download at the online version of this Design Idea at www.edn.com/100204dib, contains an assembly-language routine. It generates a precise quantity, Q, of high-frequency pulses, which deliver the number, N, that the outputs of the 74HCT4040 require. The relations are $Q=2^{N-1}$ in dot-bar mode and $Q=2^N-1$ in bar-graph mode. **List-**

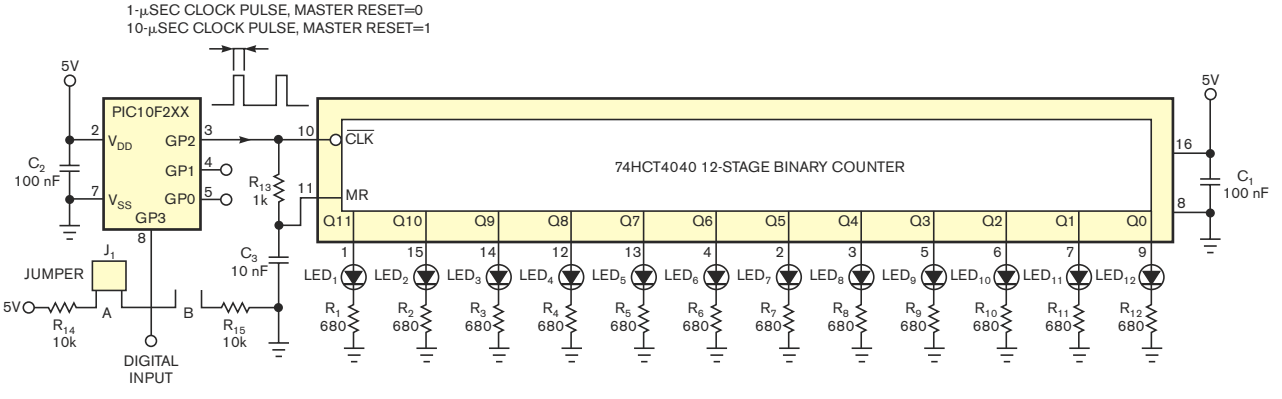

NOTE: BEFORE POWER-ON, PUT JUMPER INTO A FOR BAR-GRAPH MODE OR PUT JUMPER INTO B FOR DOT-BAR MODE.

Figure 1 A 12-LED dot-bar/bar-graph display uses a 12-stage binary counter.

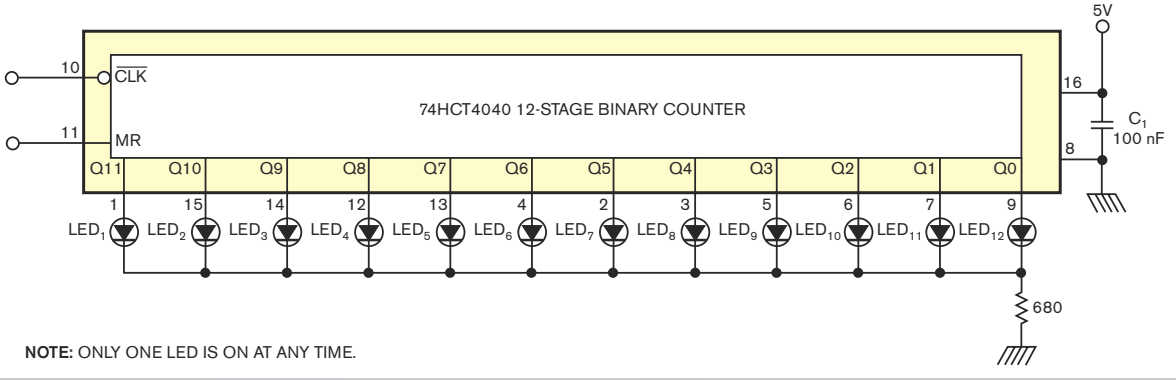

NOTE: ONLY ONE LED IS ON AT ANY TIME.

Figure 2 Using one current-limiting resistor can reduce parts count.

Extend HDMI 1.3 cable reach beyond 30 meters

Cable equalizer improves HDMI™ 1.3 Deep Color™ receiver performance

The MAX3815A automatic equalizer for HDMI cables enables high-quality digital interconnections on inexpensive cables beyond 30m. It also improves receiver margin for short cables. This equalizer compensates for high-frequency cable loss on three transition-minimized, differential signaling (TMDS) channels up to 2.25Gbps each, plus a one-tenth (0.1x) rate clock. Due to its fully automatic equalization, no user control is required.

Extends 2.25Gbps TMDS interface length

- 0 to 22m over HDMI cable, 28AWG STP
- 0 to 35m over HDMI cable, 24AWG STP

Extends 1.65Gbps TMDS interface length

- 0 to 40m over HDMI cable, 24AWG STP
- 0 to 28m over HDMI cable, 28AWG STP

HDMI and Deep Color are trademarks of HDMI Licensing, LLC.
Blu-ray Disc is a registered trademark of the Blu-Ray Disc Association.

www.maxim-ic.com/digitalvideo

MAXIM DIRECT™
www.maxim-ic.com/shop

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

MAXIM
INNOVATION DELIVERED™

For free samples or technical support, visit our website.

ing 2, which is also available at www.edn.com/100204dib, is a full example of using this routine with Microchip's (www.microchip.com) PIC10F202, a member of the PIC10F series, which is the company's most pin-limited microcontroller family.

Although the PIC's internal unique clock frequency is 4 MHz, you'll notice little flicker effect. You can reduce the flicker by using a midrange pin-limited PIC microcontroller, such as the PIC12F629, which has an internal clock frequency of 20 MHz. **Listing 3**, also available at www.edn.com/100204dib, uses a look-up table to convert the desired number into seven-segment code to replace the 12 LEDs with a four-digit display. EDN

REFERENCES

1 Anonymous, "Microcontroller provides low-cost analog-to-digital conversion, drives seven-segment displays," *EDN*, May 10, 2007, pg 80, www.edn.com/article/CA6437954.

2 Raynus, Abel, "Squeeze extra outputs from a pin-limited microcontroller," *EDN*, Aug 4, 2005, pg 96, www.edn.com/article/CA629311.

3 Jayapal, R, PhD, "Microcontroller's single I/O-port line drives a bar-graph display," *EDN*, July 6, 2006, p 90, www.edn.com/article/CA6347254.

4 Lekic, Nedjeljko, and Zoran Mijanovic, "Three microcontroller ports drive 12 LEDs," *EDN*, Dec 15, 2006, pg 67, www.edn.com/article/CA6399101.

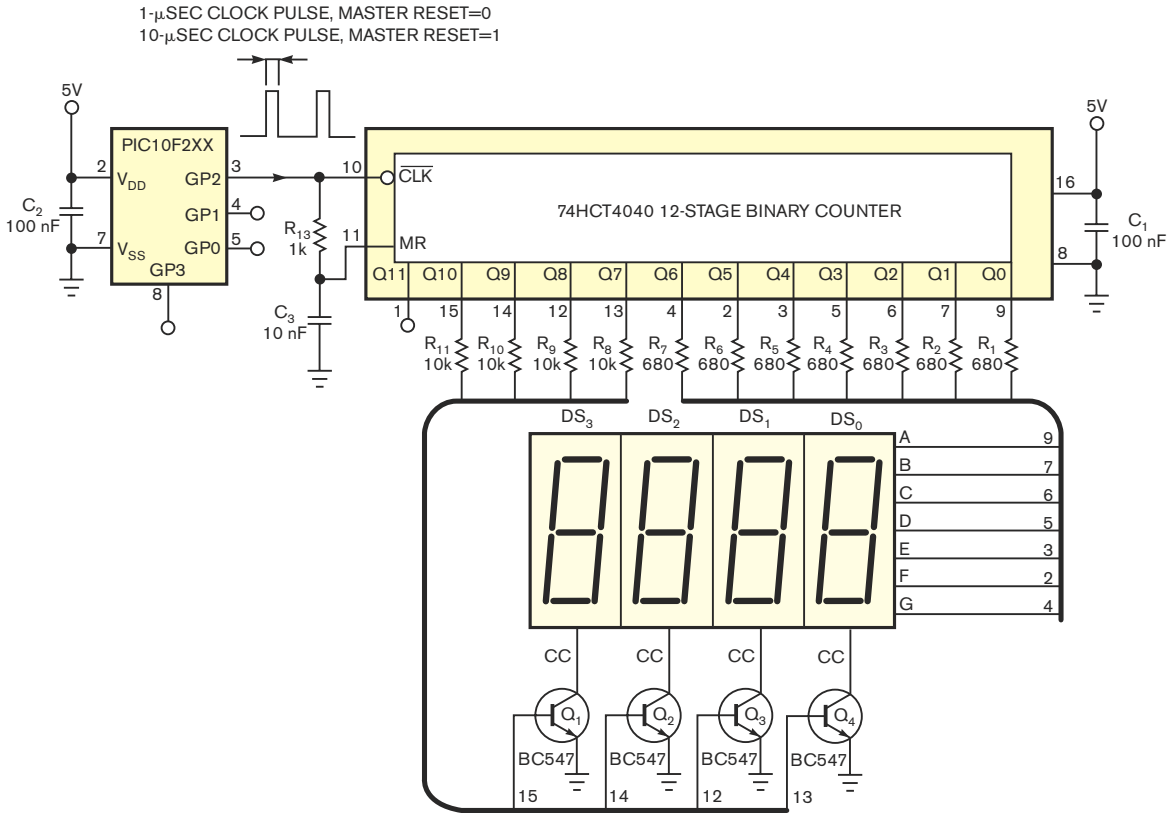

Figure 3 The same technique works with seven-segment displays.

Circuit precludes common-mode conduction

Ken Herrick, Oakland, CA

When driving an H-bridge or a similar circuit, you usually must ensure that two or more transistors are not on at the same time. Eliminating multiple transistors from turning on reduces power consumption and lowers EMI (electromagnetic interference). Crossover-delay circuits solve that problem. **Figure 1** shows a simple, two-phase design that lets you adjust the crossover delays equally by changing the value of one component with minimal phase delay.

Each Schmitt trigger inverter is driven during one half-cycle through

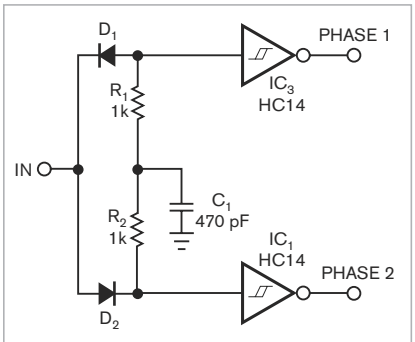

Figure 1 Each Schmitt trigger inverter is driven during one half-cycle through a diode. The RC delay occurs during the alternate half-cycle. Equal-value resistors R_1 and R_2 serve alternatively as delay elements and gate-coupling resistors.

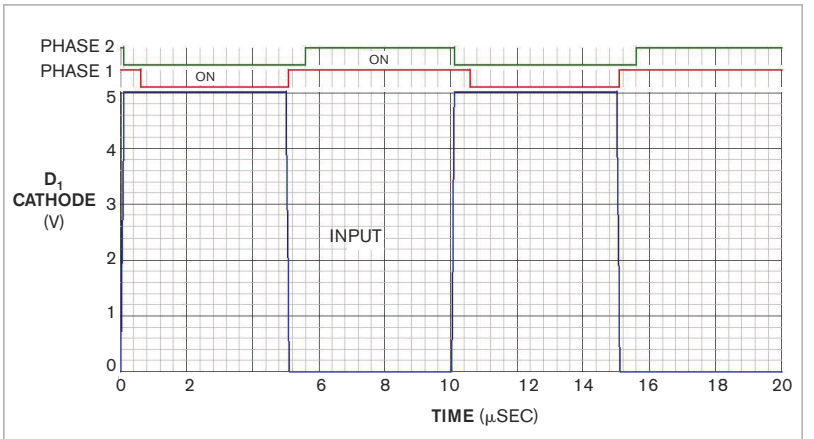

Figure 2 For the two out-of-phase half-cycles, leading edges are delayed equally with respect to the input transition, and trailing edges are coincident with the transition within about one gate delay.

a diode. The RC delay occurs during the alternate half-cycle. Equal-value resistors R_1 and R_2 serve alternatively as delay elements and gate-coupling resistors. The waveform in **Figure 2** shows the result. For the two out-of-phase half-cycles, leading edges are delayed equally with respect to the input transition, and trailing edges are coincident with the transition within about one gate delay. If you need equal polarity "on" half-cycles, insert an inverter in one of the two phase outputs. Alternatively, if biphase drivers, such as those for driving coupling transformers, will follow this circuit, merely interchange the outputs of one of those drivers to effect the inversion. EDN

BY SUZANNE DEFFREE

supplychain

LINKING DESIGN AND RESOURCES

NEDA targets counterfeit components with return guidelines

The NEDA (National Electronic Distributors Association) has released its guidelines on product returns, aiming to provide a clear-cut process to ensure the legitimacy of products customers purchase through an authorized distribution channel. According to NEDA, the primary entry of counterfeit electronic parts into the electronics supply chain is purchasing from unauthorized sources, such as brokers and unauthorized distributors. The group's goal is to strengthen the integrity of the electronics supply chain when customers purchase through authorized distribution.

Robin B Gray (**photo**), executive vice president of NEDA, describes the guidelines as a way to codify best practices and as a way to thwart the entrance of counterfeit components into the electronics supply chain. "It's something that authorized distributors have been doing for several years now," he says. "We thought it was a good opportunity to publish these [guidelines] as best practices for distributors." Customers have expressed concern about the soundness of the supply chain with respect to counterfeit products. "About the only way counterfeit products can get into the authorized supply chain is through returns," Gray adds. "We wanted to make sure that authorized distributors practice these types of policies and that those who don't [do so] consider implementing them as soon as possible."

The guideline recommendations include use of sealable packaging; verification of purchase orders, quantities, and dates; and visual inspection of parts to determine whether additional handling or processing may have occurred. "You should have some sort of verification that the product you are getting back from the customer is, in fact, the product that you sold to them," says Gray. "We want the component manufacturers to also consider their stock rotation/scrap-allowance policies with respect to their authorized supply chain." The NEDA guidelines are available at www.nedassoc.org.

GREEN UPDATE

DOE INTRODUCES LIGHTING-FACTS LABEL

Focusing on future growth in the SSL (solid-state-lighting) market, the US DOE (Department of Energy, www.energy.gov) has teamed with the NGLIA (Next Generation Lighting Industry Alliance, www.nglia.org) to create the SSL quality-advocates program and the lighting-facts label (**photo**). Some LED-luminaire manufacturers voluntarily commit to testing products and reporting performance results according to industry standards. They pledge to use the lighting-facts label to document the performance of products they manufacture, sell, distribute, or promote.

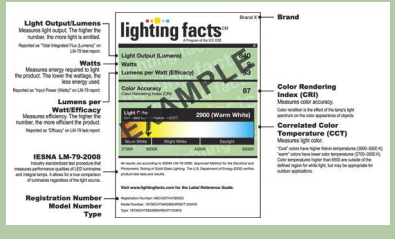

The label provides a summary of product-performance data about lumens, efficacy, watts, CCT (correlated color temperature), and CRI (color-rendering index), according to IES (Illuminating Engineering Society of North America, www.iesna.org) LM-79-2008.

The DOE reports that it may in future editions of the label consider additional metrics related to reliability, product consistency, construction, and other parameters.

OUTLOOK

MOBILE-CE-DEVICE SHIPMENTS TO INCREASE 55-FOLD

ABI Research (www.abiresearch.com) forecasts that shipments of mobile broadband-enabled CE (consumer electronics), including e-book readers, mobile digital cameras and camcorders, PMPs (personal media players), PNDs (personal navigation devices), and mobile gaming devices, will increase 55-fold between 2008 and 2014. Total shipments should reach 58 million in 2014.

"Consumers are snapping up connected PNDs and e-book readers in numbers and will continue to do so," says Jeff Orr, a senior ABI research analyst. The consumption is not coming without questions. Several of these devices have for some time featured Wi-Fi. "When you embed a cellular or mobile broadband modem in a device," says Orr, "it becomes tied to a particular operator's service billing. That [scenario] dramatically changes the device vendors' business model."

Paying each time a consumer downloads content works for the e-book-reader supply chain. "In the case of a multiplayer game, questions arise," says Orr. "Paying to download the game is straightforward, but, beyond that, what's the appropriate model?"


XILINX
ACCELERATE CHANGE.
ADVANCE INNOVATION.
The power to innovate is in your hands. Focus your energy on innovation and add differentiating value to your electronic systems. Xilinx Targeted Design Platforms integrate advanced FPGA silicon and design tools, IP cores, development boards, and reference designs—giving you the freedom to innovate without risk. Find out more at **www.xilinx.com**
© 2010 Xilinx, Inc. All rights reserved. Xilinx and the Xilinx logo are registered trademarks of Xilinx in the United States and other countries. All other trademarks are property of their respective holders.

productroundup

PASSIVES

Radial leaded inductors provide EMI/RFI protection

Suiting EMI filters and switching power supplies, the DR217-series radial-leaded inductors protect against EMI/RFI events in current-handling applications using currents as high as 8.4A. Features include a 1-µH to 150-mH inductance range and a −20 to +80°C temperature range. The DCR meets specifications of 0.008 to 520Ω over a maximum current rating of 0.045 to 8.4A. The DR217-0 through DR217-7 devices have minimum Q specifications based on test frequencies of 25.2 kHz to 7.96 MHz. The DR217 series is available in round packages with diameters of 5 to 14 mm and with 2- to 7.5-mm lead spacing. Prices range from 7.5 to 22 cents (volume quantities).

Datatronic Distribution, www.datatronics.com

Planar transformer handles as much as 22W for SMPS applications

The Sfernice PLAC 100 user-configurable planar transformer targets switch-mode-power-supply applications, including flyback, forward, push-pull, and half-bridge converters. Using the vendor's software allows adaptation of the transformer to a range of applications by setting the power-supply type and desired level of input voltage, output voltage, power, and frequency. After entering the variables, the software determines the appropriate configuration and determines which outputs should connect. The transformer has a 50- to 400-kHz frequency range, handling as much as 220W for SMPS applications. The SMD and through-hole versions operate at −55 to +125°C. The Sfernice PLAC 100 planar transformer costs $80.

Vishay Intertechnology, www.vishay.com

Fixed resistors use stripline termination

Aiming at applications requiring high-power termination from a remote location, the Model CHF-190104KxBF500L fixed-resistor series uses RF power termination. The technology uses 50Ω stripline termination, functioning at 1000W at 1-GHz frequencies. Running at 800 MHz enables a reduction in reflections at a 1.06 maximum voltage-standing wave ratio. Additional specifications include 50Ω resistance, ±5% tolerance, a thick resistive film, and use of a beryllium-oxide substrate. Targeting use in combiners and dividers in television and radio broadcast, radar and medical equipment, or other applications requiring high power, the CHF190104KxBF500L costs $43 (1000).

Bourns, www.bourns.com

INTEGRATED CIRCUITS

Sensor IC interfaces with piezoelectronic sensors

The AS1716 capacitive-sensor IC interfaces with piezoelectronic capacitive sensors and finds use in automotive knock sensors. The device provides differential inputs; differential-to-single-ended conversion; a programmable-gain stage; a two-pole, lowpass, multiple-feedback filter; and a first-order lowpass filter, cutting off the high-frequency noise components. The IC operates over a 4.5 to 5.5V supply. The analog front end aims at unbiased capacitive sensors, interfacing with sample-and-hold input stages and ADCs. The AS1716 costs 66 cents (1000).

austriamicrosystems, www.austriamicrosystems.com

Video-reconstruction filter targets high-definition media players

The six-channel ADA4424-6 video-reconstruction filter provides luma channels with four modes for detecting and concealing dc-input-offset voltages as high as 1.1V, with no dc offset at the output pins. The video filter supports S1/S2 signaling that the EIA CPR-1202 specification defines and the D-connector interface that the EIAJ CP-4120 standard defines. The filter's SD channels operate flat to 6.75 MHz and include two fifth-order Butterworth filters for SD filtering. The SD section also includes internal summation of Y/C channels for composite-video baseband-signal output. Available in a lead TSSOP-38, the ADA4424-6 costs $1.49 (1000).

Analog Devices, www.analog.com

EDN ADVERTISER INDEX

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

Swan song for swap-out strategy

DANIEL VASCONCELLOS

I started my career in electronics some 25 years ago as an engineer in the test-equipment-maintenance and -calibration department of a major company. One of the first jobs I undertook was to assist in the repair of a true-rms RF-power meter, which worked by dissipating the input signal in a precision resistor and measuring the temperature rise. These units were rather bulky by today's standards and consisted of a horizontal "backplane" with slots into which were plugged a power-supply board, an analog board, a digital board, and a user-interface board. One of these meters gave no reading after the segment test at power-on, instead displaying all zeros.

I was the new guy, so my manager teamed me up with a couple of the senior engineers. We had an identical power meter on hand—identical except that this one was working. We decided to put good boards one at a time from this unit into the defective unit and thus identify the faulty board. We carefully noted the serial numbers of the boards so we would know where they belonged and got going.

We started with the analog board, which was the most likely suspect. No luck. We soldiered on until we had swapped all the boards. Still no joy. So we put all the boards from the good unit into the bad one. Nothing. Could the backplane be the culprit? Surely not; all it had were the connectors, but we took it out of the chassis and painstakingly did continuity checks anyway in case of contact problems. None surfaced. These boards were not from some cheap consumer equipment; they were sturdily built.

We decided to put the two units side by side and compare the signals at various points. We put all the boards we had taken from the good unit back into the good unit and switched it on. It didn't work. Now it was time to start panicking; instead of one faulty unit, we had two.

We unseated and reseated the boards in their slots, mentally said a few prayers, and tried again, but there was no getting around it: Both units were now showing the same fault. The fault had somehow propagated itself. The look on the boss's face cut us to the quick. It was closing time, so we wound up and went on our way with heavy hearts.

NOW IT WAS TIME TO START PANICKING; INSTEAD OF ONE FAULTY UNIT, WE HAD TWO.

The next morning, we started from scratch, as we should have done in the first place, using circuit diagrams, an oscilloscope, a multimeter, an RF-power source—the works. When we finally located the fault, all became clear. The unit had a dual-slope ADC with the linear and digital ICs split between the analog and digital boards. The analog board got the control signals from the digital board, and the output of the final comparator went back to the digital board, in which the input gate was defective: It had developed a dead short to the power-supply voltage. When the comparator on the analog board tried to pull this input low, it, too, died. For that reason, the good analog board had also become defective.

The moral of the story? Blindly swapping boards between defective and working units is risky and may not work. EDN

Sujit Liddle is an engineer in New Delhi, India.

www.edn.com/tales

30 MHz Function/Arb Generator

- 1 μHz to 30 MHz
- 1 μHz frequency resolution
- Sine, square, ramp, triangle, noise
- 12-bit, 40 MSamples/sec arb waveforms
- 2.5, 5, or 10 MHz reference input
- Log and linear sweeps, AM, FM, PM and burst modulation
- GPIB and RS-232 (opt.)

$1595 *(U.S. List)*

The DS345 is a flexible and powerful function/arb waveform generator for a wide range of applications.

Using Direct Digital Synthesis (DDS), the DS345 can generate very high resolution waveforms. Its full suite of features includes AM, FM, PM and Burst modulation, linear and log sweeps and a 10 MHz reference input.

Dedicated buttons make front panel operation intuitive, unlike competitive models with obscure menu driven interfaces. This makes the DS345 very easy to use.

If you're looking for outstanding value in a function generator, take a close look at the DS345.

Stanford Research Systems

1290-D Reamwood Avenue • Sunnyvale, CA 94089 • E-mail: info@thinkSRS.com
Phone: (408) 744-9040 • Fax: (408) 744-9049 • **www.thinkSRS.com**

Data acquisition just got a lot easier.

More ways to control. More ways to connect.

The new Agilent 34972A Data Acquisition Switch Unit takes our best-selling Agilent 34970A to the next level. For starters, you get convenient built-in LAN and USB connectivity. Plus, you can control your data acquisition remotely via Web interface. And transfer logged data to your PC with a simple flash drive. No more expensive adapters and connectors. That's easy. That's Agilent.

***NEW* 34972A**	**34970A**
USB and LAN	GPIB and RS232
Graphical web interface Benchlink data logger software SCPI programming	Benchlink data logger software SCPI programming
3-slot LXI unit with built-in 6 ½ digit DMM	
~~$1,845~~ $1,597*	$1,597*

Hurry! For a limited time, get the 34972A at the 34970A price.
www.agilent.com/find/Agilent34972A

*Prices are in USD and are subject to change.
See participating distributors for details.

Agilent Technologies

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

Newark
800-463-9275
www.newark.com/agilent