Reed Business Information (RBI™)

EDN®

VOICE OF THE ENGINEER

OCT 8

Issue 19/2009

www.edn.com

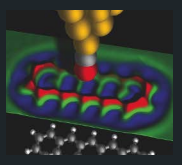

DESIGNING AN ACCESSIBLE BOARD

Page 24

COMPARING COMPARATORS: MEASURE SIGNALS, GET RESULTS

Page 15

IP QUALITY LIES BEYOND COMPLIANCE TESTING

Page 31

DEEP PACKET INSPECTION OPTIMIZES MOBILE APPLICATIONS

Page 37

APEX PRECISION POWER

Product Innovation from Cirrus Logic

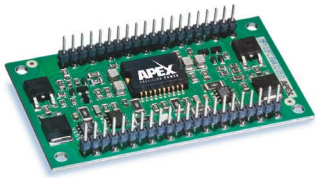

FC MODULAR 42-PIN DIP

Open Frame Product Technology
(actual footprint 65.1mm X 42.5mm)

DP 12-PIN POWER SIP

(actual footprint 30.99mm X 20.17mm)

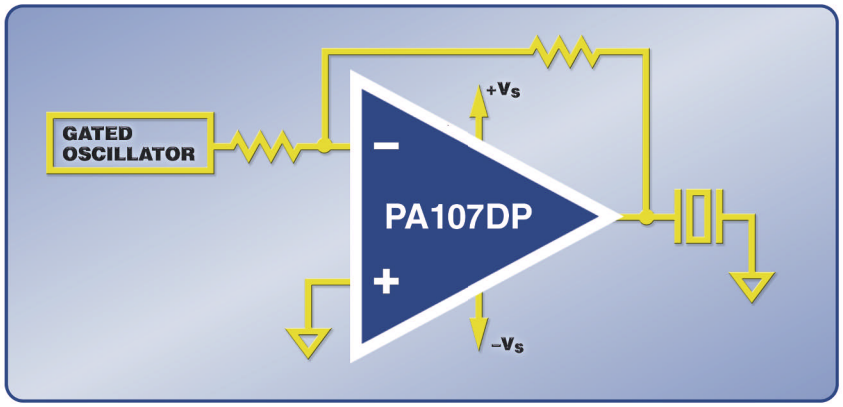

Small Package Drives Big Speed For Piezo Transducers In Ultrasonic Medical Applications

The PA107DP and MP103FC are the newest additions to the Apex Precision Power® family of high speed, high voltage power amplifiers from Cirrus Logic. The PA107DP is housed in a very small Power SIP measuring less than two inches square. The device targets medical ultrasonic and imaging applications by providing up to 3000 V/µs on voltage supplies up to 200 V. For applications requiring lower speeds, but multiple drivers, the MP103FC is a dual channel amplifier with a high power bandwidth of 230 kHz, or a 180 V/µs slew rate. The open frame form factor of the MP103FC is ideal for high speed assembly and provides a low per unit cost in comparison with many in-house discrete designs.

Model	Slew Rate	Output Current	Supply Voltage Operation
PA107DP	3000 V/µs	1.5 A continuous 5 A Peak	40 V to 200 V Dual Supply
MP103FC	180 V/µs	Up To 15 A PEAK	30 V to 200 V Dual Supply

NORTH AMERICA
+1 800-625-4084

ASIA PACIFIC
+852 2376-0801

JAPAN
+81 (3) 5226-7757

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

DOWNLOAD A COPY OF THE CIRRUS LOGIC V15 APEX PRECISION POWER® PRODUCT DATA BOOK AT WWW.CIRRUS.COM/107EDN

LEARN MORE AT
www.cirrus.com

For product selection assistance or technical support with Apex Precision Power® products call 1-800-546-2739 or apex.support@cirrus.com.

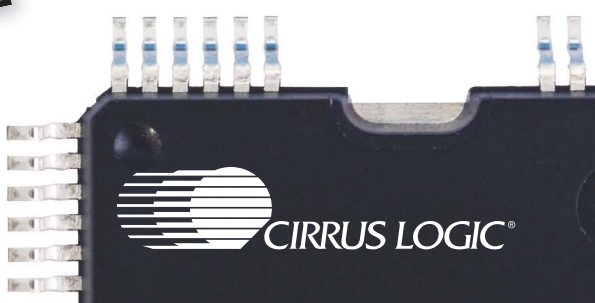

ENTER HERE.

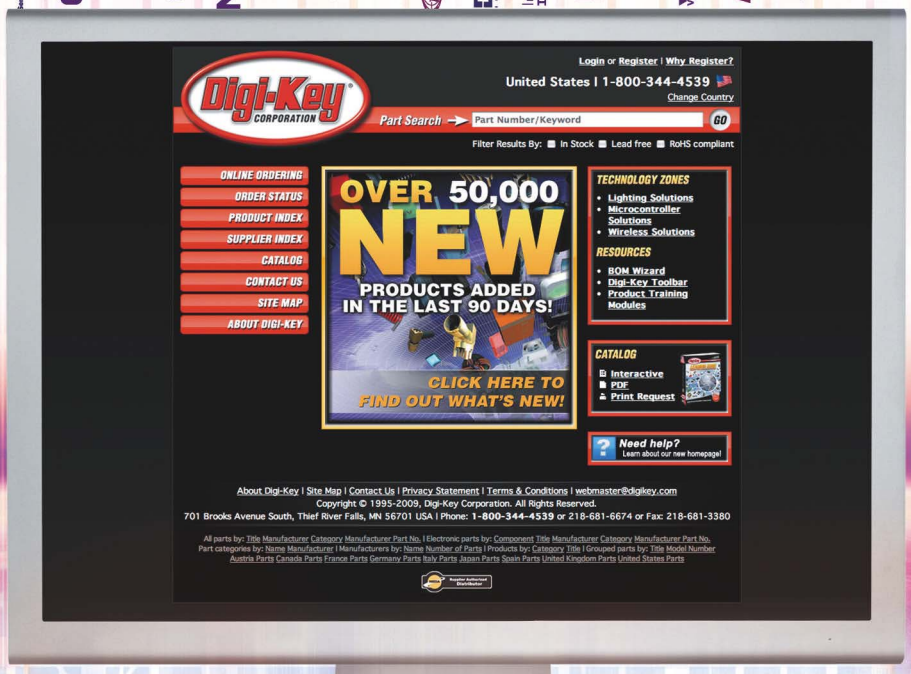

Digi-Key® CORPORATION

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

Over 400,000 products in stock from more than 400 supplier partners.*

An additional 1,000,000+ components can be sourced at digikey.com

The industry's broadest product selection available for immediate delivery

www.digikey.com

1.800.344.4539

*Digi-Key is an authorized distributor for all supplier partners. New products added daily. © 2009 Digi-Key Corporation, 701 Brooks Ave. South, Thief River Falls, MN 56701, USA

EDN 10.8.09 contents

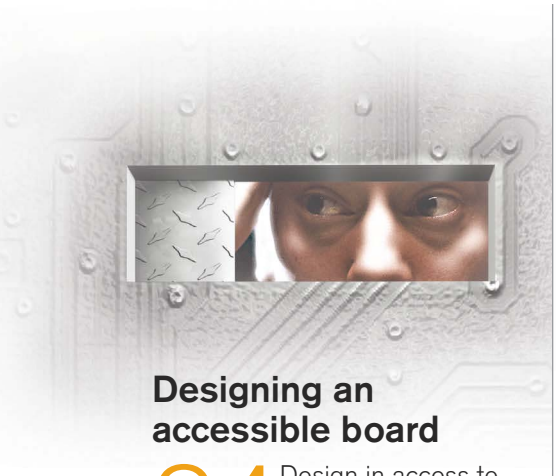

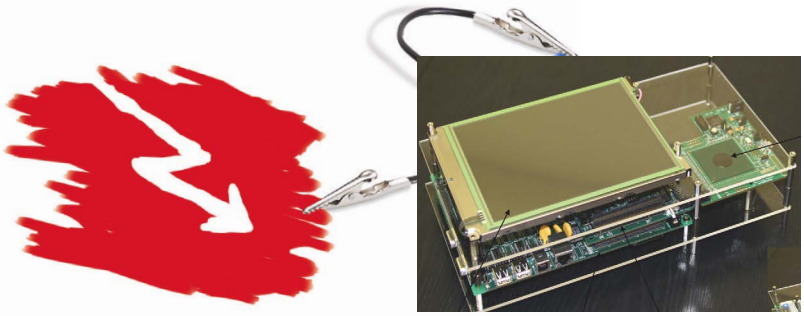

Designing an accessible board

by Ron Wilson, Executive Editor

Comparing comparators: Measure signals, get results

by Paul Rako, Technical Editor

IP quality lies beyond compliance testing

by Navraj Nandra, Synopsys Inc

Deep packet inspection optimizes mobile applications

by Mike Coward, Continuous Computing

pulse

DESIGNIDEAS

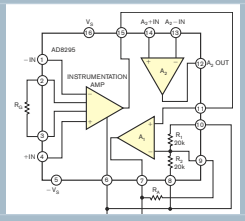

▶ Send your Design Ideas to edndesignideas@reedbusiness.com.

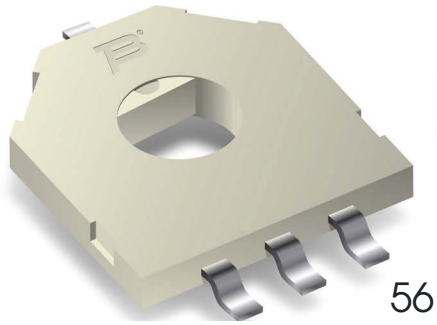

56

58

DEPARTMENTS & COLUMNS

EDN online contents www.edn.com

ONLINE ONLY

Check out these Web-exclusive articles:

Thermal-design guidelines for solid-state lighting applications using LEDs
Six critical steps are necessary to developing a successful, cost-effective thermal design for an LED application.
→www.edn.com/article/CA6695956

FPGA architectural power-saving techniques at 40 nm
As geometries shrink, FPGAs must begin to employ design-specific power-management techniques to save power while meeting timing.
→www.edn.com/article/CA6698462

READERS' CHOICE

Here is a selection of recent articles receiving high traffic on www.edn.com.

Take advantage of open-source hardware
→www.edn.com/article/CA6676116

Missing pulse detects position or produces a delay
→www.edn.com/article/CA6694947

Is there gold in green technology? A microelectronics perspective
→www.edn.com/article/CA6697713

Solar market to recover … eventually
→www.edn.com/article/CA6696791

EDN® (ISSN#0012-7515), (GST#123397457) is published biweekly, 24 times per year, by Reed Business Information, 8878 Barrons Blvd, Highlands Ranch, CO 80129-2345. Reed Business Information, a division of Reed Elsevier Inc, is located at 360 Park Avenue South, New York, NY 10010. John Poulin, Chief Executive Officer; Jeff DeBalko, President, Business Media. Periodicals postage paid at Littleton, CO 80126 and additional mailing offices. Circulation records are maintained at Reed Business Information, 8878 S Barrons Blvd, Highlands Ranch, CO 80129-2345. Telephone (303) 470-4445. POSTMASTER: Send address changes to EDN®, PO Box 7500, Highlands Ranch, CO 80163-7500. EDN® copyright 2009 by Reed Elsevier Inc. Rates for nonqualified subscriptions, including all issues: US, $179.99 one year; Canada, $229.99 one year (includes 7% GST, GST#123397457); Mexico, $229.99 one year; air expedited, $399.99 one year. Except for special issues where price changes are indicated, single copies are available for $10 US and $15 foreign. Publications Agreement No. 40685520. Return undeliverable Canadian addresses to: RCS International, Box 697 STN A, Windsor Ontario N9A 6N4. E-mail: Subsmail@ReedBusiness.com. Please address all subscription mail to EDN®, 8878 S Barrons Blvd, Highlands Ranch, CO 80129-2345. EDN® is a registered trademark of Reed Elsevier Properties Inc, used under license. A Reed Business Information Publication/Volume 54, Number 19 (Printed in USA).

BY RICK NELSON, EDITOR-IN-CHIEF

Keep your hands off my Internet content

The US Congress should promptly pass HR3458—the Internet Freedom Preservation Act of 2009. Despite the whining of the badly misnamed Hands off the Internet organization, passage of the act is necessary to ensure a level playing field for content providers and consumers, no matter how big or how small. Net neutrality is critical to ensuring that consumers, not deep-pocketed content providers in secret deals with service providers, determine what content they want to access.

Net-neutrality opponents—the marketers and financiers concocting schemes to charge a premium for faster content delivery—need to get out of the way and let their engineers deliver the bandwidth necessary to ensure that everyone's content gets through without discrimination. It shouldn't take a presidential address to a joint session of Congress to get Congress to tell net-neutrality opponents to "keep your hands off my Internet content."

As a recent editorial (**Reference 1**) points out, "On the Internet today, a Web site run by a solo blogger can load as quickly as any corporate home page. Internet-service providers, including leading cable and phone companies, want to be able to change that [situation] so they can give priority to businesses that pay, or make deals with, them." That scenario should not happen. As the editorial notes, without net neutrality, businesses could slow down or block Web content from competitors or content advocating for political or social causes. That issue is not a hypothetical one, the editorial notes.

The arguments of Hands off the Internet are disingenuous at best. The organization cites the development of high-speed Internet-access systems including cable wire, DSL (digital-subscriber line), and wireless, saying the developments resulted from "competition in its purest form." OK, but so what? They should apply that developmental expertise to building sufficient bandwidth that they need not ration it. And HR3458 makes absolutely no mention of transmission media, so what's the point of the example?

What does the bill say, specifically? In a nutshell, providers may not discriminate against anyone's ability to transmit or receive content, may not impose any charge beyond user end charges, may not prevent a user from attaching any device to the network as long as the device does not harm the network, must offer Internet access to any person, and may not prioritize one provider's traffic over another's.

One issue that net-neutrality opponents harp on is telemedicine. Here is what the bill says: "Reasonable Network Management—Nothing … shall be construed to prohibit an Internet-access provider from engaging in reasonable network management consistent with the policies and duties of nondiscrimination and openness set forth in this Act … a network-management practice is a reasonable practice only if it furthers a critically important interest, is narrowly tailored to further that interest, and is the means of furthering that interest that is the least restrictive, least discriminatory, and least constricting of consumer choice available." Does anyone reasonable think that telemedicine—at least for those who have the health insurance that might actually pay for it—isn't a "critically important interest"?

What's most appalling about net neutrality's opponents is their "can't-do" attitude. They can't possibly provide all the bandwidth consumers want, so they'll have to limit traffic—except that of their "friends." In fact, In-Stat (www.instat.com) recently reported that in the US download speeds are improving and consumers are generally satisfied with the speed of their current broadband connections.

What's happening here is that the marketers, financiers, and lawyers within the companies that oppose net neutrality want to do for the Internet what "financial engineers" did for the economy. It's time for Congress to pass HR3458 and tell opponents to "keep your hands off consumers' Internet-content preferences." Then, let real engineers step up and deliver the necessary bandwidth.

Visit my blog at http://bit.ly/sGJ6r to comment and to find links to source material. **EDN**

REFERENCE

1 "Access and the Internet," *The New York Times*, Aug 29, 2009, www.nytimes.com/2009/08/29/opinion/29sat3.html.

Contact me at rnelson@reedbusiness.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

VICE PRESIDENT, INTERACTIVE MEDIA AND SALES MANAGEMENT
Karthik Krishnan
1-646-746-7580;
karthik.krishnan@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
Newsletter/Editorial Coordinator
1-781-734-8443; fax: 1-303-265-3279
Adam Odoardi
Prepress Manager
1-781-734-8325; fax: 1-303-265-3042

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
The Quadrant, Sutton, Surrey SM2 5AS
+44 118 935 1650; fax: +44 118 935 1670;
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@idg-rbi.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

MEET THE GUY that ELIMINATED HIS TEAM'S MANUFACTURING VARIABILITY ISSUES.

FIX YOUR MANUFACTURING VARIABILITY PROBLEMS AND YOUR REPUTATION WILL PRECEDE YOU. If you're designing chips for high functionality, high speed and lower power consumption at the most advanced process nodes, you've got variability issues. We have the solution that will increase your yield, performance and prestige by a wide margin. | Get more information at mentor.com/solutions/manufacturing-variability.

Mentor Graphics®

THE EDA TECHNOLOGY LEADER

pulse

INNOVATIONS & INNOVATORS

Octal power-converter-control chip uses PMBus

Linear Technology Corp's new LTC2978 monitoring-and-control chip supervises eight switching power converters. You communicate with the IC over the PMBus (power-management bus). The controlled supplies have output voltages of 0 to 6V. The device has an onboard EEPROM that stores settings when the chip is off. Eight 10-bit DACs use the power converter's trim pin or resistor-feedback network to sequence, adjust, and margin the eight target converter chips. The device measures the target converter's output voltage by multiplexing a 15-bit delta-sigma ADC across the eight converters. The ADC can also measure the part's input voltage. A carefully trimmed onboard 1.232V reference provides a ±0.25% overall accurate measurement, including ADC errors. The device operates from a 4.5 to 12V power supply and contains 2.5 and 3.3V linear regulators. It also has an internal temperature sensor. You can cascade multiple devices if you need to control more than eight converters.

The LTC2978 allows you to bring eight analog- or digital-power-converter chips under PMBus control. It can perform as a watchdog timer, a sequencer, and a supervisor to monitor both overvoltage and undervoltage events. You can also use the part to perform margining of the target power supplies during manufacturing to ensure that your system works properly under a predetermined range of converter supply voltages. The chip logs faults to the internal EEPROM. The chip's interface and command set comply with the PMBus specification and operate as fast as 400 kHz. You can explore all of the product's features with LTpowerPlay, a free software package that can communicate over the PMBus. The demoboard kit includes a dongle to convert your computer's USB port to an SMBus (system-management bus).

The LTC2978 comes in a 9×9-mm, 64-pin QFN package. The LTC2978CUP version operates at a 0 to +70°C junction temperature, and the LTC2978IUP version operates in the −40 to +85°C range. The commercial-temperature version sells for $10.95 (1000).

—by Paul Rako

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

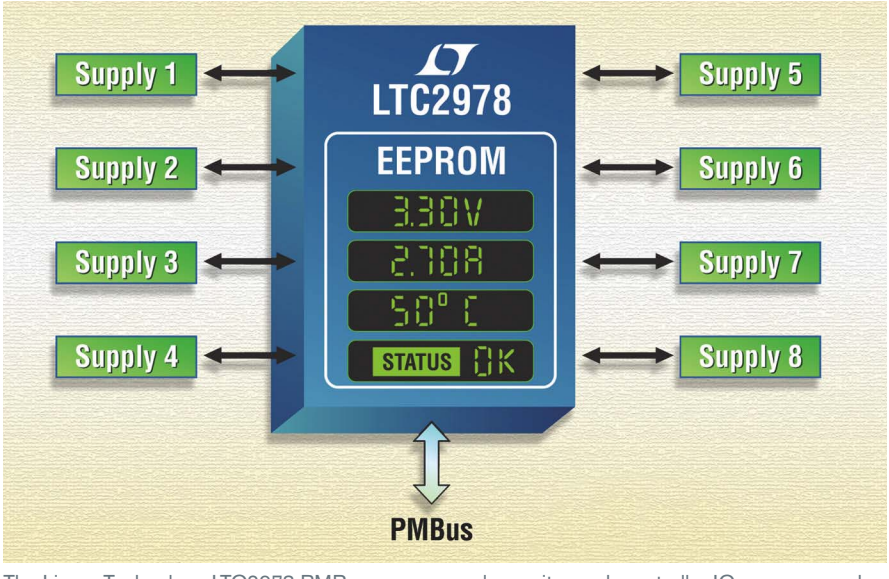

The Linear Technology LTC2978 PMBus power-supply-monitor and -controller IC manages and supervises eight regulators using a trim pin or feedback resistors.

FEEDBACK LOOP

"It seems appropriate to the idea of replacing the old whistling kettle with a software-driven alternative. Who can forget Lauren Bacall's immortal quote? 'You know how to whistle, don't you, Steve? You just put your lips together—and blow.'"

—Engineer Gary Fisher, in *EDN*'s Feedback Loop, at www.edn.com/article/CA6674029. Add your comments.

ADC modules have onboard FPGAs for speedy data analysis

Agilent's Acqiris product group recently announced three high-speed, high-accuracy, wide-dynamic-range, modular digitizers targeting medical imaging, scientific instrumentation, nondestructive testing, storage testing, LIDAR (light detection and ranging), validation of mixed-signal semiconductors in automated test equipment, and advanced research in high-energy physics, nuclear physics, and astrophysics.

The four-channel, 12-bit SVM3500 digitizer allows channel interleaving that enables acquisition at rates to 2G samples/sec. The module focuses on applications in radar, EW (electronic warfare), and SI (synthesized instrumentation). The input-amplifier bandwidth is more than twice the minimum for accurate reproduction of signals that you can digitize without aliasing at 2G samples/sec.

The unit adds an analog-digitizer mezzanine board, which includes four 12-bit, 500M-sample/sec ADCs, to the manufacturer's U1083A VME/VXS base module. Also on the mezzanine board is the MCK clock-distribution circuit, which distributes throughout the board very-low-added-jitter in-phase, quadrature, or out-of-phase versions of an external clock. These clock signals allow interleaved acquisition at rates to 2G samples/sec. In the interleaved mode, an FIR (finite-impulse-response) filter in an on-mezzanine Xilinx (www.xilinx.com) Virtex-5 FPGA corrects for static offset and gain errors and frequency-dependent gain and delay mismatch among the ADCs. This FPGA also includes an externally accessible look-up table. The resulting analog performance maintains more than 10 effective bits over a 10-MHz to 1-GHz band. Features include an SFDR (spurious-free dynamic range) to 80 dBc at 500M samples/sec and 70 dBc at 2G samples/sec, THD (total harmonic distortion) of −80 dBc at 500M samples/sec and −79 dBc at 2G samples/sec, and SNR (signal-to-noise ratio) greater than 62 dBc.

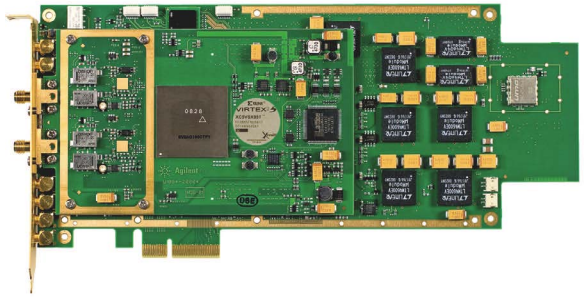

The U1084A PCIe module outputs data at 650 Mbytes/sec.

The base card provides high-performance, real-time data processing by means of the large SX55 and FX100 Xilinx Virtex-4 FPGAs. The board supports eight 3.125-Gbps serial links on the VXS backplane. Two optical links on the front panel support data transfers at rates to 3.125 Gbps, making the aggregate data bandwidth more than 3.5G bytes/sec. The VME64x interface supports the 2eSST (two-edge-source-synchronous-transfer) protocol. The optional FDK (firmware-development kit) helps you develop applications for the SX55 and FX100 FPGAs. The FDK includes a set of cores that interface to the hardware, a base design to provide ready-to-use designs, and a test-bench environment.

The dual-channel, high-speed, 8-bit, four-lane U1084A PCIe (Peripheral Component Interconnect Express) digitizer acquires as many as 4G samples/sec. This digitizer with an onboard FPGA offers a 1.5-GHz bandwidth and incorporates a 15-psec TTI (trigger-time interpolator) for accurate timing measurement. The FPGA-based peak-detection and -analysis firmware provides real-time signal processing and allows real-time acquisition and peak detection at data rates to 4G samples/sec.

Firmware options allow the data-converter card to perform user-defined postprocessing tasks, which you can easily upload into the FPGA under program control. These tasks can redefine the way in which the card acquires and processes data, making the system flexible and easy to reconfigure. The signal-peak-detection firmware option enables the board to create a peak histogram of successive acquisitions, with each histogram bin containing peak counts or summed peak amplitudes. The unit can record as many as 60M peaks/sec and exhibits trigger jitter of just 250 psec. With the high-speed PCIe bus, the onboard processing maximizes data and measurement throughput. The digitizer outputs data at rates to 650 Mbytes/sec. Software drivers for Windows and Linux enable you to easily integrate the U1084A into measurement systems or use it to replace other high-speed Acqiris data converters.

Single-unit prices begin at $26,500 for the U1084A, $46,000 for the SVM3500 (U1083A-005), and $42,000 for the SVM4800 (U1083A-007). For more, go to www.edn.com/article/CA6697248.

—by Dan Strassberg

▷**Agilent Technologies**, www.agilent.com/find/digitizers, www.agilent.com/find/u1084a, www.agilent.com/find/wideband_Rx_Tx.

DILBERT By Scott Adams

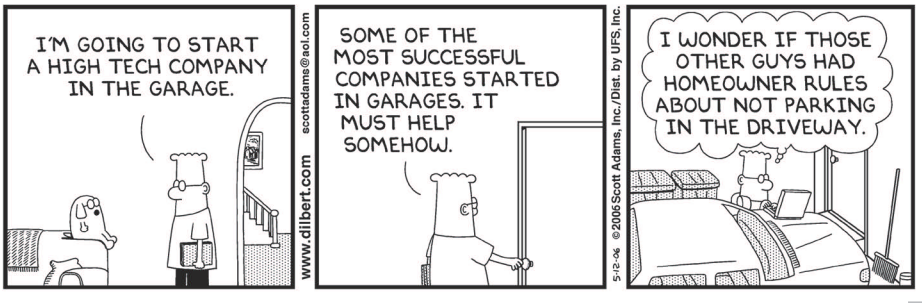

Precision JFET op amp operates with 36V power

Analog Devices' new ADA4627 JFET operational amplifier has 5-pA maximum bias current at 25°C and 500-pA maximum current at −40 to +85°C. The 19-MHz-bandwidth, A-grade part has less than 300-μV maximum offset voltage and 2-μV/°C maximum offset drift. The B-grade part has 200 μV of maximum offset and 1-μV/°C drift. The device achieves an 82V/μsec slew rate and has an open-loop gain of 120 dB. The output drives loads as low as 600Ω. Analog Devices built the device with the dielectrically isolated iPolar process technology, which enables 5.5-nV/√Hz voltage noise when operating on ±18V rails. Voltage noise at 10 Hz is 16.5 nV/√Hz.

Because JFETs are buried devices, they are not subject to noise that surface defects create. The drawback is that diffusion rather than lithography controls JFET specs, complicating the manufacture of precision devices with low offset voltage. The combination of low noise, low bias current, and high precision makes the ADA4627 suitable for circuits that require accurate signal acquisition and fast settling in optical, communications, and high-speed data-acquisition systems; professional audio and test equipment; and medical instrumentation. The device can also act as an output buffer in a DAC or as an amplifier for high-source-impedance sensors and photodiodes.

The ADA4627 comes in eight-pin SOIC packages and 3×3-mm, eight-pin LFCSPs. It operates over a temperature range of −40 to +125°C. The $6.75 (1000), A-grade part is inferior to the B-grade part, which costs $10.75 (1000).

—by Paul Rako

▷**Analog Devices**, www.analog.com.

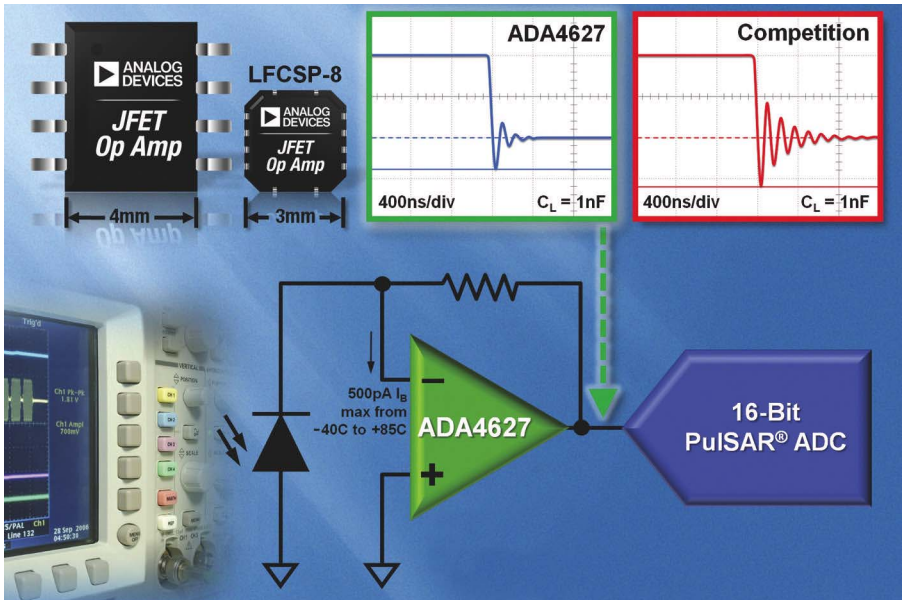

The ADA4627 JFET op amp features low offset, low drift, and low input bias current.

10.08.01 HIGH HOT/COLD FACTOR BOOSTS HB-LED EFFICACY

Philips' new Luxeon Rebel HB LEDs (high-brightness light-emitting diodes) have a 0.93 hot/cold factor–the ratio of junction temperatures at 100 and 25°C. The new devices thus beat typical specs for production LEDs, whose hot/cold factors range from 0.8 to 0.85. These factors are important in applications such as light-fixture luminaires, which spend most of their time operating at temperatures of 80 to 110°C–much higher than the 25°C junction temperature that most HB-LED spec sheets give as the light output. As its junction temperature increases, an LED's light output and efficacy decrease.

—by Margery Conner

▶**Philips Lumileds, www.philipslumileds.com.**

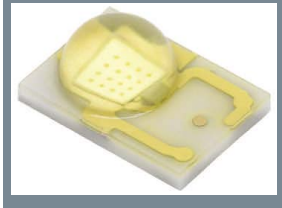

The hot/cold factor–the ratio of junction temperatures at 100 and 25°C–is 0.93 for Philips' new Luxeon Rebel LED.

DSP engine sports hybrid-SIMD/VLIW operation with ITU compatibility

Tensilica's 16-bit dual-MAC (multiply/accumulate) ConnX D2 DSP engine integrates with the company's Xtensa LX processor core. The DSP core supports modeless switching between SIMD (single-instruction/multiple-data) and VLIW (very-long-instruction-word) formats. The VLIW format supports two pipeline executions, in which one may be an SIMD operation, with register moves and operations such as autoincrement loads. The DSP engine provides 1-to-1 mapping support for ITU-T (www.itu.int) reference-code intrinsics, as well as 1-to-1 mapping to most intrinsics for Texas Instruments' (www.ti.com) C6x family. The ConnX D2 engine primarily targets the telecom infrastructure and VOIP (voice-over-Internet Protocol) applications.

The ConnX D2 DSP-engine option adds dual 16-bit MAC units and an eight-entry, 40-bit register file to the base architecture of the Xtensa LX DPU (data-plane-processing unit). It supports 16-, 32-, and 40-bit integer and fixed-point, 16-bit complex, and 8- and 16-bit vector data types, as well as seven addressing modes. For more on this product, go to www.edn.com/article/CA6685704.**—by Robert Cravotta**

▷**Tensilica**, www.tensilica.com.

Cadence links FPGA-pin allocation to PCB-layout tools

Designers who use today's large FPGAs on their PCBs (printed-circuit boards) face an increasing problem: handling the pinout and board tracking around the packages of these programmable devices. The FPGAs have large pinouts; they offer considerable flexibility in allocation of pins to internal logic functions but also have complex rules that you must follow when doing so.

FPGA designers typically accomplish this task with minimal knowledge of the connectivity of those pins to other packages on the PCB. Other engineers develop the overall circuit functions, connecting the FPGA to processors, memory, and other packages. The job of creating the PCB layout falls to yet other engineers, who must find escape routes for all the signal groups that emerge from the FPGA and route them to other packages. Those engineers must also account for features such as wide memory buses and fast signal lines.

To address these problems, Cadence has incorporated into its Allegro and OrCAD products software from Taray (www.tarayinc.com). The Taray software provides automated assistance in the FPGA-pin-allocation step, yielding a correct-by-construction process. The process, automated placement-aware FPGA-pin-I/O-assignment synthesis, gives the software knowledge of the pin-allocation rules for Xilinx (www.xilinx.com) and Altera (www.altera.com) FPGAs and of the connectivity of the logic function in the FPGAs to other packages.

Large FPGAs contain so much logic that it is difficult to represent their functions in an understandable form on one diagram, so engineers view different functions on different pages and lose sight of the bigger picture. The Taray tool tracks that information; it also allows you to make a generic placement of packages on a "canvas" when performing initial PCB design, and it yields a global view of connectivity and connection density. It further avoids manual errors in the pin-allocation step.

The offering is applicable whether you are designing an FPGA-based board as a final product or building an ASIC prototype in which the logic of the target device is divided among many complex FPGAs. Using the Taray/Allegro tool greatly simplifies such designs because the process that fragments the logic onto the FPGAs has no knowledge of placement or layout. With the new software, you can automate pin assignment across all of the FPGAs in one step.

The ability to quickly view an optimized placement for any given logic architecture also permits you to explore cost-versus-performance trade-offs at the board-design level. Under the Cadence PCB-tool brands, the software is available for the Version 16.2 release as OrCAD FPGA System Planner or Allegro FPGA System Planner L, XL, and GXL tiers, and it integrates tightly with OrCAD Capture, OrCAD PCB Designer, Allegro Design Entry HDL (hardware-design language), and Allegro PCB Design products. Taray also continues to sell its tool offering, 7Circuits, as a stand-alone product.

—by Graham Prophet

▷**Cadence Design Systems**, www.cadence.com.

FUJITSU LAUNCHES USB 3.0-TO-SATA-BRIDGE CHIP

Fujitsu Microelectronics recently announced a 3.0-to-SATA (serial-advanced-technology-attachment)-bridge chip. The company intends the device to act as a connection between a USB (Universal Serial Bus) 3.0 cable and a SATA external-storage device. It incorporates an AES (Advanced Encryption Standard) engine–critical to many external-media applications–and Fujitsu claims that the chip can support 300-Mbyte/sec throughput with flow-through encryption/decryption.

The device includes USB 3.0 Revision 1.0 and 2.0 PHY (physical)- and link-layer blocks and has 3-Gbps SATA Gen 2i PHY and link layers. Structurally, the chip includes an internal bus for data movement, the AES engine, and an unspecified 32-bit microprocessor core with attendant SRAM. The CPU's job is primarily system-control functions within the chip, such as initialization, dispatching tasks to the AES engine, and handling the intervention-required bits of the two link-layer protocols. The RAM is for code and data storage for the CPU rather than for buffering, according to Davy Yoshida, Fujitsu's director of business development. The chip also includes an SPI (serial-peripheral-interface) port to an external serial-flash chip for boot-loading of code and some general-purpose I/O pins for status LEDs and other functions.

The chip requires both 1.2 and 3.3V supplies and typically consumes 550 mW. That requirement means that a bus-powered interface with a local regulator to step the USB supply pin down to the required voltages would eat up most of the allowed 900 mA during operation without passing any power on to the mass-storage device. So interfaces using the chip will require an external power supply, probably one that the storage device furnishes. Further, according to Fujitsu, the chip's USB I/O pins do not tolerate high voltage, so they may require external protection circuitry to comply with the letter of the 3.0 specification about tolerating shorts to the power pin.

Fujitsu is looking at applications in both automotive infotainment and consumer electronics. So don't be surprised to see the company's PHY surface soon in some additional bridge chips, and perhaps in 40-nm form, in forthcoming consumer and automotive SOCs (systems on chips). For more on this development, go to www.edn.com/091008pa.–by Ron Wilson

▶**Fujitsu, www.fujitsu.com.**

10.08.09

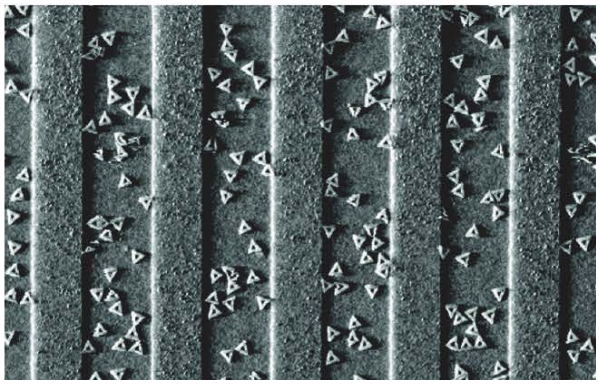

Low concentrations of triangular DNA origami bind to wide lines on a lithographically patterned surface (courtesy IBM).

RESEARCH UPDATE

BY SUZANNE DEFFREE

DNA scaffolding targets sub-22-nm lithography

Scientists at IBM Research and the California Institute of Technology (Pasadena, CA) have discovered an advancement that could be a major step forward in the semiconductor industry's move to sub-22-nm lithography. IBM researchers and collaborator Paul WK Rothemund of Caltech are discussing a development in combining lithographic patterning with self-assembly, a method of arranging "DNA-origami" structures on surfaces compatible with today's semiconductor-manufacturing equipment.

IBM says that its approach of using DNA molecules as scaffolding may provide a way to reach sub-22-nm lithography. IBM describes the scaffolding as a procedure in which millions of carbon nanotubes could self-assemble into precise patterns by sticking to DNA molecules.

The positioned DNA nanostructures can serve as scaffolds, or miniature PCB (printed-circuit boards), for the assembly of components, such as carbon nanotubes, nanowires, and nanoparticles, at dimensions significantly smaller than possible with conventional semiconductor-fabrication techniques.

The techniques for preparing DNA origami, which scientists at Caltech developed, cause single DNA molecules to self-assemble in a solution using a reaction between a long strand of viral DNA and a mixture of short, synthetic oligonucleotide strands. These short segments act as staples, folding the viral DNA into the desired 2-D shape through complementary base-pair binding. Scientists can modify the short staples to provide attachment sites for nanoscale components at resolutions as small as 6 nm. In this way, scientists can prepare DNA nanostructures, such as squares, triangles, and stars, with dimensions of 100 to 150 nm per edge and a thickness the width of the DNA double helix.

IBM uses traditional semiconductor techniques to etch out patterns and uses either electron-beam or optical lithography to create arrays of binding sites of the proper size and shape to match those of individual origami structures.

The researchers say that the discovery of the template material and deposition conditions were the keys to affording high selectivity so that origami binds only to the patterns of sticky patches and nowhere else (**Reference 1**).

REFERENCE

1 Kershner, Ryan J, Luisa D Bozano, Christine M Micheel, Albert M Hung, Ann R Fornof, Jennifer N Cha, Charles T Rettner, Marco Bersani, Jane Frommer, Paul WK Rothemund, and Gregory M Wallraff, "Placement and orientation of individual DNA shapes on lithographically patterned surfaces," *Nature Nanotechnology*, September 2009, www.nature.com/nnano/journal/v4/n9/pdf/nnano.2009.220.pdf.

▷**IBM Research**, www.research.ibm.com.

▷**California Institute of Technology**, www.caltech.edu.

IBM OBTAINS IMAGE OF MOLECULE'S CHEMICAL STRUCTURE

10.08.09

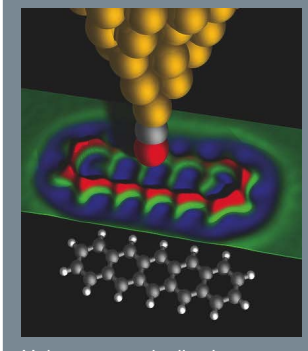

Using an atomically sharp metal tip terminated with a carbon-monoxide molecule, IBM scientists measured the short-range regime of forces that allowed them to obtain an image of the inner structure of a molecule. The colored surface represents experimental data (courtesy IBM Research–Zurich).

Using a noncontact AFM (atomic-force microscope), IBM scientists have obtained an image of the chemical structure inside a molecule. The results push the exploration of using molecules and atoms at the smallest scale.

According to IBM, understanding the charge distribution at the atomic scale is "essential for building smaller, faster, and more energy-efficient computing components than today's processors and memory devices." Such components could one day contribute to a "smarter planet" by helping instrument and interconnect the physical world. IBM Research–Zurich scientists Gerhard Meyer, Leo Gross, Fabian Mohn, and Nikolaj Moll worked with Peter Liljeroth of Utrecht University (www.uu/nl) using an AFM in an ultrahigh vacuum at −268°C (−451°F) to image the chemical structure of oblong pentacene molecules comprising 22 1.4-in.-long carbon atoms and 14 1.4-in.-long hydrogen atoms. With the AFM, the researchers looked through an electron cloud to see the atomic backbone of a molecule. For more on this research, go to www.edn.com/article/CA6687056.

▶**IBM Research–Zurich, www.zurich.ibm.com.**

XILINX

RISE ABOVE

New Xilinx Targeted Design Platforms give designers a boost to take innovation to a much higher level. These comprehensive platforms combine Virtex®-6 or Spartan®-6 devices, the ISE Design Suite 11, development hardware, IP, reference designs, documentation, and service and support—so you're way ahead of the competition before you even start. Plus, everything works together seamlessly, so design teams can focus on product differentiation, add more value, and make the whole project more successful. Learn more now at **www.xilinx.com/6**.

BY HOWARD JOHNSON, PhD

Real signals

Which step in **Figure 1** best approximates the digital signals you work with every day? The piecewise-linear step starts with a jerk, mindlessly follows a perfectly uniform ramp, and smacks hard into its upper limit. Real signals don't do that. The smooth-looking curve is a Gaussian step. It is the time integral of a Gaussian bell-shaped curve. The Gaussian step has a smooth precursor, a smooth tail, and a fast, monotonic rise in the middle. It looks more like a real digital signal, and for a good reason.

In the field of linear-system analysis, the Central Limit Theorem states that the step response of any system whose performance is limited by a large number of similar bandlimiting effects tends to become Gaussian as the number of effects approaches infinity. That theorem applies to digital devices because a typical digital driver comprises many performance-limiting stages cascaded in series, all having similar bandwidths.

An I/O driver uses multiple stages to quickly convert nanoamps of current from within your silicon into milliamps of current on the PCB (printed-circuit board). A single-stage FET amplifier cannot do that job. If you make a single-stage FET gate big enough to switch PCB-level currents, there is insufficient current available at the silicon level to quickly charge that gate. A better approach breaks down the circuit into a series of multiple stages cascaded in series. Each stage is exponentially larger than the one before it. Gate-design experts spend a lot of time choosing the number of stages within each driver and carefully crafting each stage to achieve maximum performance.

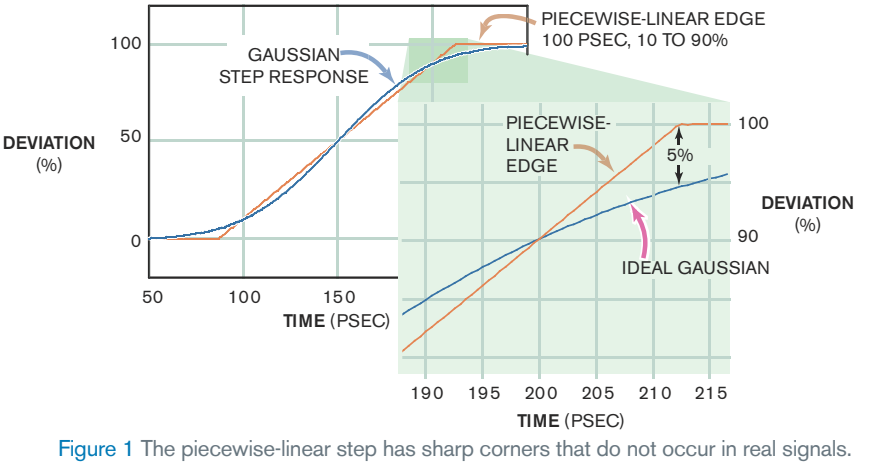

Figure 1 The piecewise-linear step has sharp corners that do not occur in real signals.

That's where the Central Limit Theorem comes into play. Imagine that you are designing an I/O driver with 10 stages. At dc, the circuit works perfectly. As you go up in frequency, various parasitic effects come into play, limiting the bandwidth. Each stage suffers perhaps 20 parasitic effects, making 200 total effects you must manage to complete the design. The system behaves like a cascade of 200 tiny low-pass-filter elements connected in series. The poorest-performing elements limit the bandwidth of the whole system regardless of how well the other elements work. If you wish to raise the overall performance, attack the worst parts first.

Successively identify the lowest-hanging effects and improve them one at a time until you have pushed every part of the system up to a uniform performance ceiling beyond which it becomes increasingly difficult to post significant gains. At that point, stop fiddling and immediately put your product into production. This process creates a system comprising many performance-limiting stages cascaded in series, all having similar bandwidths. For that reason, the step response looks Gaussian.

You can see that the maximum time-domain deviation between a Gaussian step and a piecewise-linear step amounts to only 5% (**Figure 1**). That amount doesn't sound like much, but those sharp corners in the piecewise-linear curve can induce frequency-domain deviations as large as 20 dB, a significant source of error.EDN

Howard Johnson, PhD, of Signal Consulting, frequently conducts technical workshops for digital engineers at Oxford University and other sites worldwide. Visit his Web site at www.sigcon.com or e-mail him at howie03@sigcon.com.

Go to www.edn.com/091008hj and click on Feedback Loop to post a comment on this column.

www.edn.com/signalintegrity

More measurements. More data. Less cost.

With up to 560 channels of data logging at a maximum scan rate of 1000 channels/sec, an internal 6 ½ digit DMM with 22-bit resolution, and 0.03% accuracy, it's no mystery why Agilent's data acquisition systems are among the top selling in the world. The 34970A and 34980A are both at home with either data acquisition or functional test, with a broad selection of plug-in modules. Now with Agilent Data Logger Pro software, you'll be able to collect data, check limits, perform extensive branching, and view and store the data without authoring the programming yourself. It's an ideal expansion at a significant savings.

Agilent 34970A switch/measure unit

Agilent 34980A switch/measure unit

www.agilent.com/find/daqswitch

Agilent Authorized Distributor

METRICTEST™
6,000 instruments. One source.

866-436-0887 www.metrictest.com/agilent

Agilent Technologies

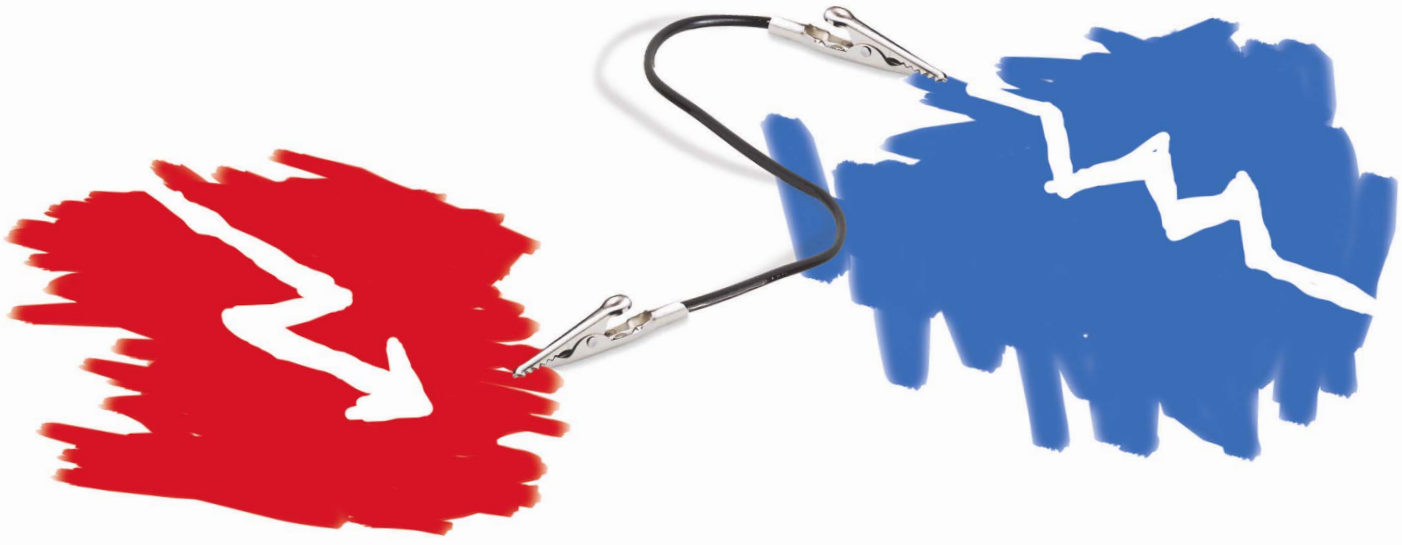

SIMPLE IN CONCEPT, VOLTAGE COMPARATORS HAVE MYRIAD SPECIFICATIONS THAT COMPLICATE THEIR APPLICATION.

COMPARING COMPARATORS: MEASURE SIGNALS, GET RESULTS

BY PAUL RAKO • TECHNICAL EDITOR

Although humble in concept, today's voltage comparators perform a simple basic task: comparing two voltages to determine which is larger. To accomplish this task, they accept two analog signals and produce a binary signal at the output. In this regard, a comparator resembles a 1-bit ADC. The basic function of a comparator comes in handy in applications requiring a comparison between a voltage and a stable reference. Such applications include level translation, radar, clock-recovery circuits, wheel sensors on antilock braking systems, precipitation gauges, and headphone-jack detection on handheld products. You can learn about hundreds of other potential applications by reading manufacturers' data sheets and application notes (**references 1** through **4**).

Despite their utility, however, comparators have countless specifications that you must be aware of (**references 5** and **6**), and no single device will make everyone happy. "For 80% of the people, a handful of comparators will fulfill their needs," says Brendan Whelan, design-section leader at Linear Technology. "For the other 20%, no two of them want the same comparator."

Further complicating the selection and application of these devices, some engineers believe that comparators are just op amps running in open-loop configuration—that is, without negative feedback. Because an operational amplifier has a well-balanced difference input and a very high gain, some op amps can serve as comparators in some functions (**references 7**, **8**, and **9**). However, in practice, dedicated comparators have several advantages over op amps. For example, dedicated voltage comparators are generally faster than general-purpose op amps that you are using as comparators. A comparator may also contain additional features, such as an accurate internal voltage reference, adjustable hysteresis, and a clock-gated input.

Another reason to choose a comparator over an op amp is that op amps stay within their linear range, whereas comparators run in open-loop mode and switch to a high or a low output. When you use an op amp as a comparator, you must first ensure that no internal clamps lie between the input pins. These clamps prevent you from pulling the input pins more than a diode drop apart. You may be able to overcome these problems by

putting series resistors in the op-amp inputs, but that approach raises the input source impedance. Also be aware that op amps may come out of saturation slowly (**Figure 1**), and, when you drive them to the rails, op amps' quiescent currents may reach excessive levels. And forget about using a handy, "free," leftover op amp from a quad-device package: One part in the quad will swing hard between the rails, almost certainly interfering and causing noise problems with the other three amplifiers in the package. An amplifier's Spice model may also not properly represent a comparator's operation during saturation.

On the other hand, engineers can legitimately use an op amp for a comparator function when their design must discriminate between small voltages of, say, 10 to 200 µV. In these cases, an op amp amplifies the input signal to a comparator. In that way, you give the comparator a low-impedance input that provides enough overdrive—the voltage margin above the nominal switching point—to properly switch the comparator. Jim Williams, staff scientist at Linear Technology, has developed several such circuits (**references 10** and **11**). According to Williams, using an op amp in front of a comparator can also work well. "Take as much gain in the preamp as you can, and let it do the work," he says.

You can also use two comparators to make a window comparator, which indicates whether the input signal is between two levels, or to ensure that the charging voltage of a lithium-ion charger stays within bounds. Further, you can use comparators with feedback to make free-running oscillators. Because comparators commonly use a reference voltage to set the trip level, hundreds of available parts combine a reference and a comparator.

AT A GLANCE

- Although it has only three pins plus power, a comparator is as tricky to apply as an op amp.
- Propagation delay and toggle rate both express the speed of a comparator.
- Using an op amp as a comparator can get you into trouble.
- Understand all the specs and charts in the data sheet to ensure that your design will work.
- To compare microvolt differences, you must use a preamp in front of your comparator.

HOW IT WORKS

The operation of a comparator is straightforward. It has a positive pin and a negative pin. When the voltage on the positive pin is higher, the output of the comparator "asserts," or drives, a signal. With an open-collector output, the comparator's output pin is the collector of a transistor or the drain of a FET. With a push-pull output, the comparator has a "totem-pole" output—that is, a complementary NPN/PNP stage—such as the one you find in operational amplifiers. An open-collector output is useful when the load and the comparator each use a different power supply. This approach allows you to implement, for example, a solenoid operating from 12V, even though your comparator may be operating from only 3.3V. Another use of open-collector outputs is to minimize quiescent current when the output is off. No base current flows in an N-type output transistor, whereas some base current always flows in one of the two output transistors in a totem-pole stage.

Open-collector outputs have a couple of drawbacks, however. For example, they require the use of external pullup resistors. The resistors must perform this pullup during high-impedance periods, so the comparator switches more quickly when its output is low than when it turns off and the pullup resistor brings the output high. Thus, using an open-collector output is unsuitable whenever you need a symmetrical waveform, such as with a clock-recovery circuit. If your circuit requires no level-shifting, you should instead select a push-pull output in a part such as Advanced Linear Devices' ALD2321APC, which can supply a 24-mA output drive and uses 90 µA of quiescent current.

Fast comparators may also have a latched output, allowing you to keep the output in a known state so that you can satisfy a setup-and-hold time to the digital input that it feeds. Once the digital section has read the comparator output,

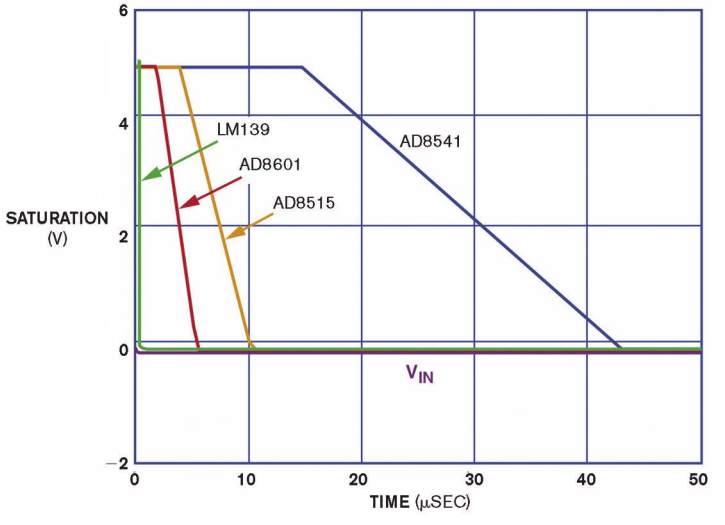

Figure 1 An op amp can take much longer to come out of saturation than a comparator (courtesy Analog Devices).

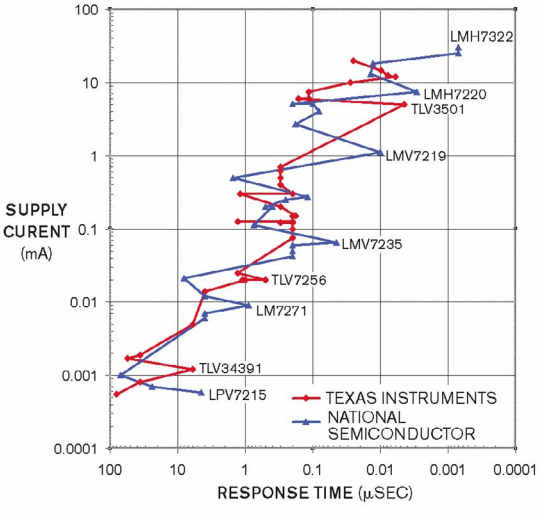

Figure 2 The faster a comparator switches, the more power it consumes.

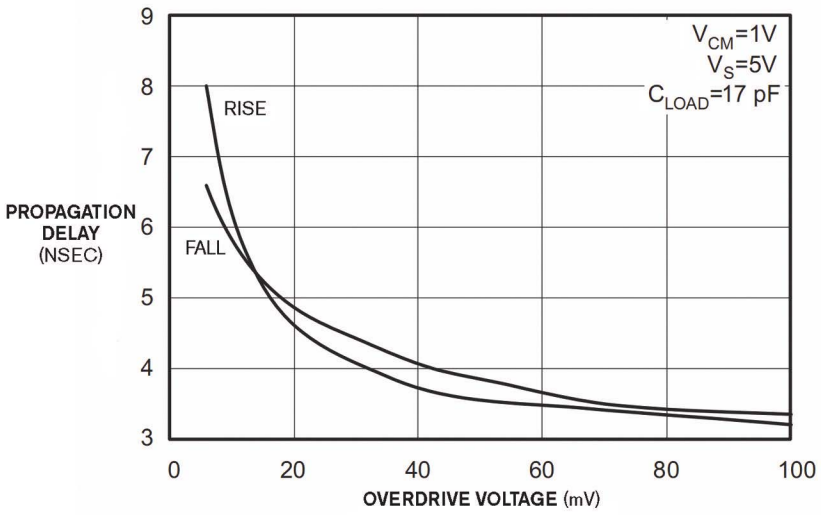

Figure 3 "Dispersion" describes the improvement in propagation delay as you increase the input overdrive (courtesy Texas Instruments).

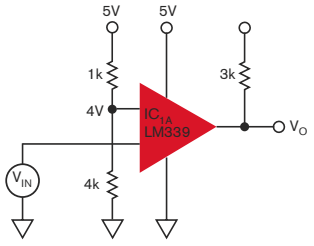

Figure 4 This circuit doesn't work because the comparator has insufficient head room. The inputs are too close to the positive rail.

you release the latch pin, and the output tracks the input. Fast comparators also may feature ECL (emitter-coupled-logic) levels of −5 to 0V. PECL (positive-emitter-coupled-logic) outputs have the same voltage swing but operate at 0 to 5V. RSPECL (reduced-swing-PECL) outputs are also available. Some fast comparators feature LVDS (low-voltage-differential-signaling) outputs with two output pins that shift 300 mV in a complementary fashion around a common-mode voltage of 1.2V. You can run these outputs directly into the LVDS input pins of FPGAs (field-programmable gate arrays) and other digital circuits.

Once you have established the output type, your next likely consideration is speed. Manufacturers generally describe a comparator as either low power or high speed. They typically build the low-power parts with CMOS processes and the fast parts with bipolar devices, illustrating the fundamental trade-off: fast, accurate parts with high power consumption versus slow parts with low-power supply currents (**Figure 2**). Another trade-off is gain versus high speed. A low-power comparator may take 70 µsec to switch and use less than 1 µA of supply current. A fast comparator with 150-psec response time, such as Analog Devices' ADCMP572, uses 44 mA. Some units stand out in the speed-versus-power trade-off. For example, National Semiconductor's LMV7219 has a 7-nsec propagation delay and uses 1.1 mA; it has relatively low gain, however. In general, an N-type device has higher electron mobility, so it switches from high to low more quickly than it switches from low to high.

A comparator consumes much more than its quiescent power when switching at its maximum toggle rate. In a quiescent state, the current is low. When you push the comparator to operate faster, you must be able to charge the capacitance, which requires current. In dynamic mode, the current increases with the speed of operation. Another factor in power consumption is the load on the chip. Capacitance also presents itself as a load to a switching circuit, and you must account for that capacitance as well as the resistive components of the load. Many parts have shutdown pins that reduce the current consumption to less than 1 µA.

As with all things analog, propagation-delay claims are true only under strictly defined conditions because how far apart you drive the input pins directly affects propagation delay. The greater the overdrive, the faster the part is. "Dispersion" is the range of propagation-delay values a device exhibits under varying degrees of overdrive (**Figure 3**). "Dispersion is a critical spec in ATE [automated-test-equipment] systems in which you are trying to measure the propagation delay of a fast logic chip," says Mike Maida, a technologist at National Semiconductor. The relationship between overdrive and speed is the reason that some engineers are loath to rate a comparator's speed as a function of toggle rate. First, you must define the output levels that qualify as a valid transition; an output level of 10 to 90% of maximum is typical. The toggle rate also implies the requirement for a hard overdrive to get the propagation delay to be as short as possible. "Propagation delay is often not a good indicator of toggle rate," says Linear Technology's Whelan. The company offers the LT1719 and LT1715 comparators, both with 4-nsec propagation delays and toggle rates of 70 and 150 MHz, respectively.

Another parameter to consider in comparator selection is noise. Manufacturers often omit noise specifications for comparators, however, instead relying on random jitter to measure noise. "In addition to just the noise signal through the gain of the device, the input's aperture errors and the output's rise and fall times can influence the jitter," says Brian Carey, senior design engineer at Analog Devices. "A clock-driver part ... is just a lower-gain comparator that's optimized for noise." National Semiconductor's Maida notes that a designer can use larger input transistors in a CMOS part to reduce flicker noise, but that approach increases the input capacitance.

Once you have selected an output type and satisfied speed and power requirements, your next concern should be the voltage rating of the comparator. Manufacturers once made slow, low-power devices in CMOS processes, but that approach meant using a 5V power supply; legacy bipolar parts, meanwhile, would work with ±15V supplies. Today, CMOS and BiCMOS analog processes often can achieve power-supply voltages of 12V or more. "In the past, people tended to use bipolar supplies for the really fast comparators because

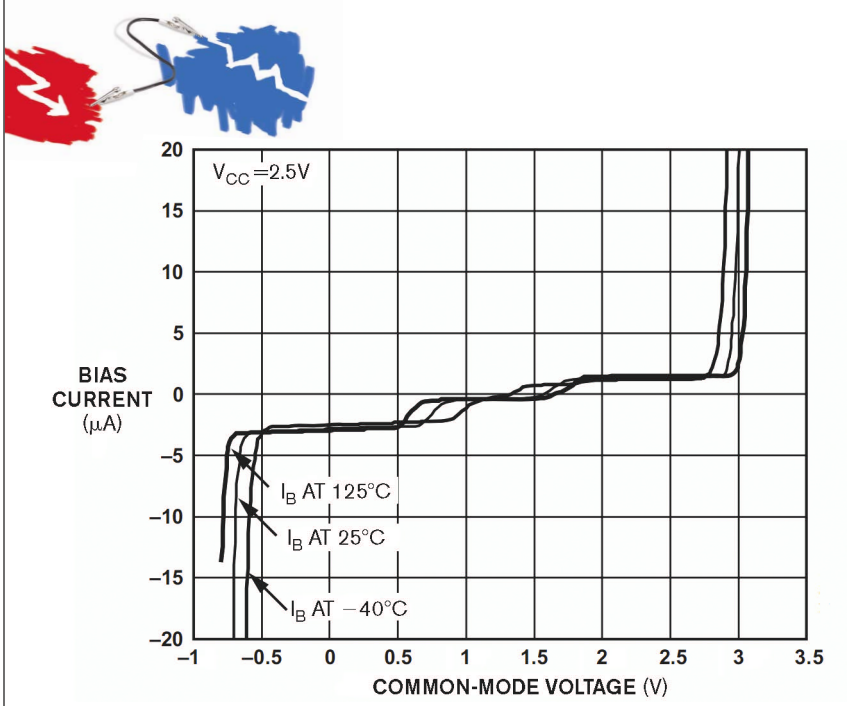

Figure 5 A rail-to-rail input part has two input stages. With bipolar parts, the input-bias current reverses direction as the input common-mode voltage sweeps through the input range (courtesy Analog Devices).

those [devices] used NPN input stages, and they could not extend the common-mode-input range to down to ground," says Maida. National Semiconductor implements fast, vertical PNP transistors in the input stage of the LMH7322 so that it can use ground as the negative rail and still allow the inputs to swing 200 mV below ground. Bipolar processes have advantages in supply voltage. Linear Technology's LT1716, for example, has a 44V input and uses only 35 μA of power. Many of the company's high-speed comparators have not only 0 to 5V but also ±5V input capability.

VERIFY THE EXACT PINOUTS OF BOTH THE OLD PART AND ITS REPLACEMENT, EVEN IF THEY COME IN THE SAME TYPE OF PACKAGE.

A factor that relates to power-supply range is the permitted common-mode voltage of a comparator's input pins. Many engineers use a legacy LM339 timer from National Semiconductor. However, its manufacturer never intended the part to work with the inputs near the top power-supply rail (**Figure 4**). Some parts allow you to drag the outputs above or below the power-supply voltage range, but others invert the output if you drag either input pin below the negative-power-supply rail (**Reference 12**). A rail-to-rail input-stage comparator, such as Analog Devices' ADCMP60x or STMicroelectronics' TS3021, extends the input-common-mode range. These devices have a dual input stage, with N-type transistors or FETs in parallel with a P-type input stage. The P-type stage works at input voltages close to ground or the negative rail, and the N-type stage works when the inputs swing close to the positive rail. IC designers usually engineer the devices to switch between stages 1 or 2V below the positive rail. Some architectures minimize the offset voltage, and the most pronounced effect occurs when the input-bias current changes from positive to negative as you sweep thought the common-mode range of a rail-to-rail part, such as Analog Devices' ADCMP600 (**Figure 5**).

Another important spec for comparators is the input-bias current—the amount of current that flows into or out of the input pins as the part operates. CMOS products have low input-bias currents, representing the mismatch in leakages in the input pin's ESD (electrostatic-discharge) structures. This input-bias current doubles for every 10°C of

temperature rise. The bias currents of fast comparators can be substantial but are not usually problematic because you tend to drive these high-speed comparators with low-impedance circuits. The input-bias current of bipolar parts changes depending on the relationship of the two inputs. In comparators, a 60-mV difference in the base voltage of a differential-input pair yields a 10-times-higher difference in the collector currents of that pair and in the input-bias currents. Thus, you may have one input pin that is sourcing or sinking at twice the rated input-bias current and the other pin with almost no input-bias current, depending on which pin has the higher voltage.

Although designers often overlook it, packaging may be the most critical spec for comparators. Legacy parts have standard pinouts for single and dual comparators that pertain to DIPs (dual inline packages) and SOIC (small-outline-integrated-circuit) packages. You may need one of the newer small packages, such as a SOT-23 (small-outline-transistor) or an SC-70. If you are replacing a legacy part, verify the exact pinouts of both the part and its replacement, even if they come in the same type of package. Other small packages include solder-bump units, or CSPs (chip-scale packages). These packages are as small as the die itself. Maxim fit the MAX9060 comparator into a four-pin CSP by tying one of the input pins to an internal rail. Some companies don't use CSPs, however, because they can't achieve the low defect rates of other packages. Manufacturers can test CSPs while sorting the wafers but not after they make the solder bumps. Manufacturers can also package parts without bond wires by mounting a solder-bump die to a lead frame. This approach yields parts smaller than 2×2 mm and still provides for die protection and product testability.

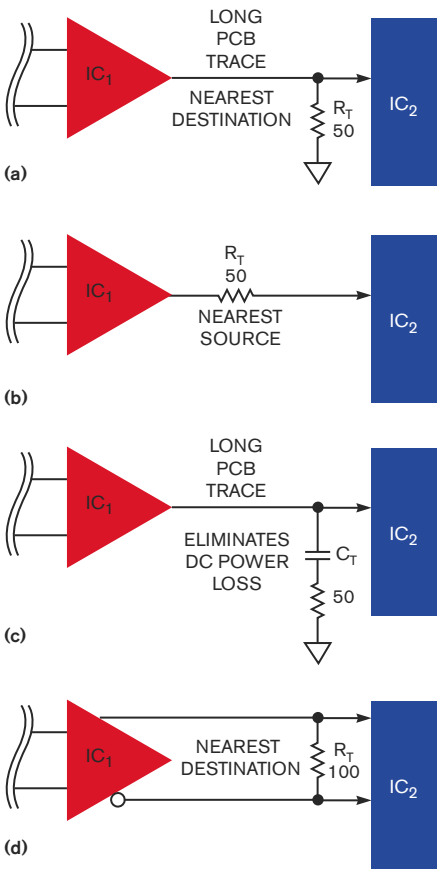

Figure 6 You should terminate fast comparator outputs to prevent ringing and reflections. You can use shunt (a), series (b), ac (c), or differential (d) termination.

PITFALLS AND PROBLEMS

All analog circuits have pitfalls, and the comparator is no exception. Application experts report that the two most common design problems are common-mode range and oscillation. To understand common-mode range—the area in which you can operate the input pins—you must understand the input structure of the comparator you are using. You must take care that your device's input range will not exceed its supply range. To ensure that this scenario does not occur, you should limit or clamp the inputs. You could place a Schottky diode on the LM339's input to ensure that its input cannot go low enough to invert the output. The ESD-protection diodes inside a device clamp the inputs to ensure that the pin cannot go more than 0.6V beyond a power rail. A current of 1 mA is safe, but 10 mA is reaching the

Go to www.edn.com/091008df and click on Feedback Loop to post a comment on this article.

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

upper limit, depending on the duty cycle. ESD-protection diodes on the input pins may limit the voltage, but problems can still occur.

The offset voltage and propagation delay of rail-to-rail devices change as the input levels move from the N-type stage to the P-type stage. In general, this transition is not a problem because you typically set up one of the input pins to a fixed dc level that determines the input stage that will switch the outputs. One important exception is in PWM (pulse-width-modulated) circuits, in which you feed one pin of the comparator a triangle wave and the other pin a waveform that the PWM represents. In this case, the inputs of the comparator sweep through the entire common-mode range.

Oscillation is the next major headache you may face. The outputs of comparators are high-speed signals no matter how slowly their inputs change. "When the guillotine goes down, it's going to shake the floor," says Paul Grohe, an application engineer at National Semiconductor. In other words, you can expect significant disruption of the power and ground in your circuit if you fail to decouple all your comparators, even the slow ones. Because of these power disruptions, Grohe warns against using a voltage divider on the power rail as a reference for the comparator.

"You have to bypass things really well with micropower parts," says Tim Regan, an application-engineering manager at Linear Technology. "The power-supply rejection is not as good as you might think because you have all these high-impedance nodes inside the part." Fast comparators are even more sensitive to bypassing and board layout (**Reference 13**). You should maintain a ground plane under the part and ensure that stray capacitances are bringing positive feedback to the inputs to make the part switch solidly rather than create oscillations.

The fundamental way of ensuring clean transitions is to introduce hysteretic resistors into the comparator circuit (**Reference 14**). These resistors return a bit of positive feedback to swamp out noise and crosstalk once the comparator begins to switch. Without hysteresis, 1 mV of ground bounce can send a part into oscillation, says Brian Hamilton, a design-section leader at Linear Technology. Many fast comparators have built-in hysteresis, and some have a pin that lets you select an adjustable hysteresis level.

Another problem with comparators is high source impedance, which makes your circuit prone to oscillation and subject to crosstalk and stray capacitance. Bob Gonzalez, an applications engineer with On Semiconductor, warns against placing 10-kΩ resistors in series on the inputs because this approach increases impedance into the inputs. The devices have internal parasitic capacitance; if the input impedance is high, that capacitance becomes prominent and adds phase shift to the circuit, leading to oscillations. You may want to use an op-amp buffer or a simple emitter-follower transistor circuit in front of the comparator to minimize the source impedance. You may also be able to overcome the problem by adding hysteresis.

Proper termination of outputs and maintaining an appropriate temperature are also critical factors in avoiding problems with comparators. Hooking a device's output to a long transition line may cause reflections from the end of the line unless you provide for a termination. You might consider an RC-termination network if you do not want to waste dc power. You could also look at using series termination that allows a reflection from the destination but then absorbs it in a series resistor (**Figure 6**).

The performance of a comparator changes over temperature. IC designers have made great strides in this area, though, so many parts are available that meet specifications over −40 to +125°C. Still, comparators are prone to oscillate at low temperatures, and high temperatures lower the device's base-to-emitter voltages and cause other performance differences. It is essential to evaluate your circuit at all the temperature extremes it will face in service.

FOR MORE INFORMATION

Advanced Linear Devices
www.aldinc.com

Analog Devices
www.analog.com

Linear Technology Corp
www.linear.com

Maxim Integrated Products
www.maxim-ic.com

National Semiconductor
www.national.com

On Semiconductor
www.onsemi.com

Rohm Semiconductor
www.rohm.com

STMicroelectronics
www.st.com

Texas Instruments
www.ti.com

"The biggest thing is propagation-delay change with temperature," says National Semiconductor's Maida. "It tends to be faster cold and slower hot. The common-mode range also shifts a little over temperature. In general, gain tends to get worse going hot, and it's more of a design challenge for head room."

Another concern is whether to simulate your circuits using Spice models. Most company representatives admit that older Spice models are far less reliable than new ones. Texas Instruments is committed to making good models for all its comparators and offers the free Tina-TI model, which allows you to cut and paste all schematics and waveforms as metafiles into Word or PowerPoint. Remember that, when you are dealing with fast comparators, the PCB (printed-circuit board) is an important component in the design; your layout may create stray capacitances and crosstalk, and these effects will overshadow any Spice simulation that does not model these second-order effects. The extreme speed of Analog Devices' new comparators, for example, causes problems for Spice. "We do not have Spice models for our newer parts because, as you go higher in performance, it becomes more difficult to get a reasonable model," says James Frame, a marketing manager at Analog Devices, which is considering developing models. The company will release them only if they are sufficiently accurate, however, so as to not mislead its customers, he says.

Over the years, the processes for manufacturing comparators have improved. Advanced CMOS processes have low power consumption and operate at more than 5V. The fast parts can take advantage of vertical PNP transistors, and the fastest have the benefit of SiGe (silicon germanium, **Reference 15**). "If you get to mix and match processes, there are obvious parts of a comparator that make sense for you to do in different processes," says Linear Technology's Hamilton.

Analog Devices uses a SiGe process in its high-speed comparators. "SiGe has a better speed-power product but also better gain; you just can't get high gain out of CMOS; we've tried," notes Carey. SiGe also has a greater voltage range. "You can't give people a 1.8V-input-range comparator; everyone wants to put in at least 2 or 3V along with a wide common-mode range on these parts," he says. Analog Devices also uses dielectric isolation to make some of the fastest comparators available. Dielectric isolation in the company's XF3 process provides low parasitic capacitance and low leakage currents.

With thousands of comparator parts and even more application circuits, you may feel overwhelmed. Armed with the basics and the subtleties, however, you can sift through all the specifications to find the comparator that provides the optimum trade-off among all your requirements. Whether you are trying to detect a pushbutton on a handheld product or sensing the trigger level in a gigahertz-frequency input to an oscilloscope, a comparator exists to fill the bill. Just heed Gordon Holton, strategic-marketing manager at Texas Instruments, when he warns you not to be too cheap. He notes some customers buy the lowest-cost comparator only to find they need the rail-to-rail input of a better part. The manufacturers' Web sites will help you

winnow the parts, and their application-engineering staffs will keep you out of trouble. By embracing the unique combination of analog and digital characteristics in a comparator, you can make sure you get the most from your circuits. EDN

REFERENCES

1. "Specifying Selected Op Amps and Comparators," Linear Brief 26, National Semiconductor, October 1973, www.national.com/ms/LB/LB-26.pdf.
2. "LM139/LM239/LM339: A Quad of Independently Functioning Comparators," Application Note 74, National Semiconductor, January 1973, www.national.com/an/AN/AN-74.pdf.
3. "Applications of the LM392 Comparator Op Amp IC," Application Note 286, National Semiconductor, September 1981, www.national.com/an/AN/AN-286.pdf.
4. Williams, Jim, "A Seven-Nanosecond Comparator for Single Supply Operation," Application Note 72, Linear Technology, May 1998, http://cds.linear.com/docs/Application%20Note/an72f.pdf.
5. "Comparators," MT-083 Tutorial, Analog Devices Inc, 2009, www.analog.com/static/imported-files/tutorials/MT-083.pdf.
6. Mancini, Ron, "Designing with comparators," *EDN*, March 29, 2001, pg 56, www.edn.com/article/CA69100.
7. "Using Op Amps As Comparators," MT-084 Tutorial, Analog Devices Inc, 2009, www.analog.com/static/imported-files/tutorials/MT-084.pdf.
8. Moghimi, Reza, "Amplifiers as Comparators?" *Analog Dialogue*, April 2003, www.analog.com/library/analogdialogue/archives/37-04/comparator.pdf.
9. "Microvolt Comparator," Linear Brief 32, National Semiconductor, June 1976, www.national.com/ms/LB/LB-32.pdf.
10. Williams, Jim, "Application Considerations and Circuits for a New Chopper-Stabilized Op Amp," Application Note 9, Linear Technology, http://cds.linear.com/docs/Application%20Note/an9.pdf.
11. Williams, Jim, "A Seven-Nanosecond Comparator for Single Supply Operation," Application Note 72, Figure 65, Linear Technology, May 1998, http://cds.linear.com/docs/Application%20Note/an72f.pdf.
12. "LM139/LM239/LM339/LM2901/LM3302 Low Power Low Offset Voltage Quad Comparators," Note 3, National Semiconductor, March 2004, www.national.com/ds/LM/LM139.pdf.
13. Williams, Jim, "High Speed Comparator Techniques," Application Note 13, Linear Technology, April 1985, http://cds.linear.com/docs/Application%20Note/an13.pdf.
14. Mancini, Ron, "Adding hysteresis to comparators," *EDN*, May 3, 2001, pg 22, www.edn.com/article/CA84881.
15. Rako, Paul, "Silicon germanium: fast, quiet, and powerful," *EDN*, Sept 18, 2008, pg 27, www.edn.com/article/CA6594090.

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

Boldly go where no system has gone before....

EPIA-P710 Pico-ITXe™ Mainboard

- Revolutionary Pico-ITXe™ Embedded Form Factor (100mm x 72mm) with PCIe, USB, & LPC Interfaces
- VIA C7 1.0GHz Processor Driving SUMIT™ I/O
- Supports 10/100/1000 Gigabit Ethernet
- ACCES Pico-I/O™ Boards (72mm x 60mm) for Stacking and System Design Flexibility

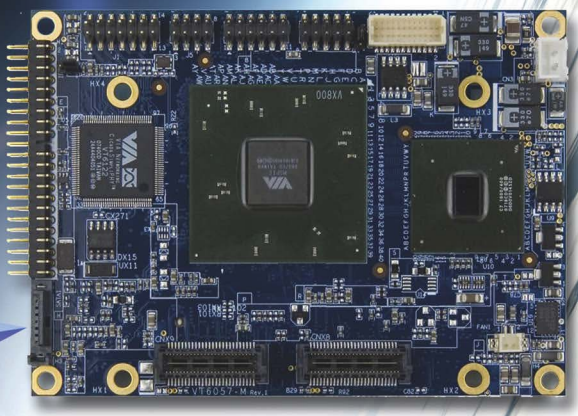

Pico-ITXe

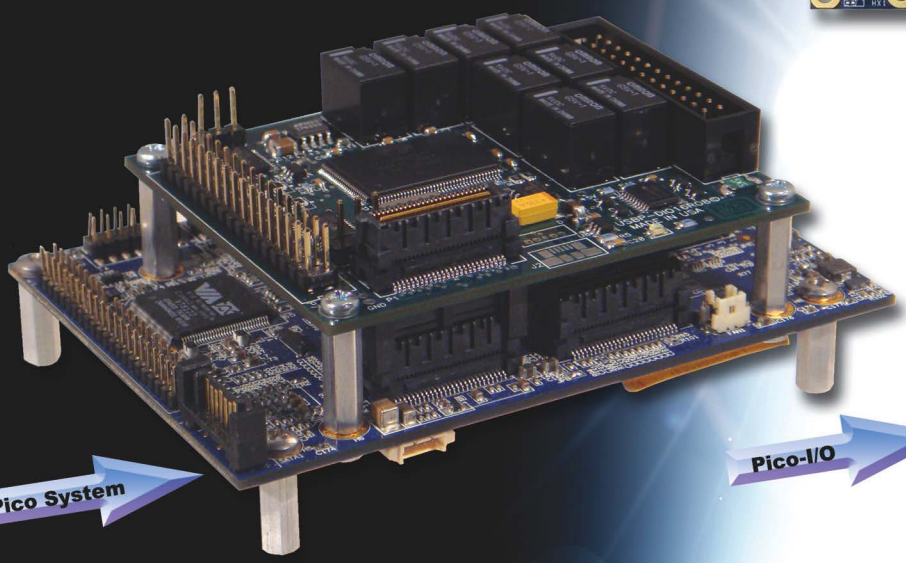

Pico System

Pico-I/O

Digital IIO, 8 Form C Relays and 16 TTL/LVTTL IIO Lines

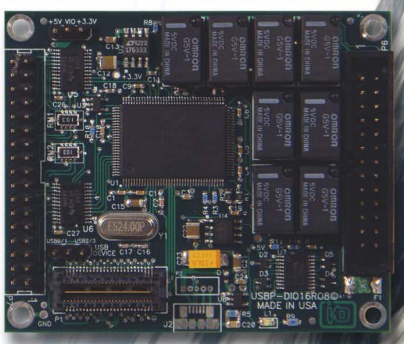

VIA Technologies' tiny Pico-ITXe™ SBC combines with ACCES I/O Products' Pico-I/O™ boards for the latest small size OEM embedded system standard.

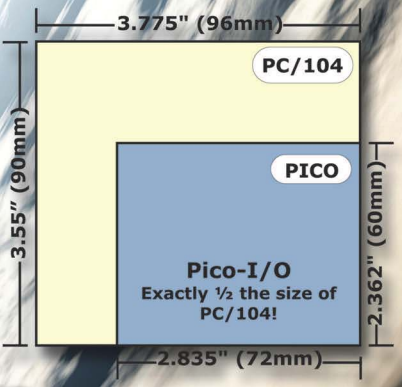

SUMIT™ standard provides a multitude of high speed and low-cost stacking I/O interfaces. Serial and low pin count I/O allows for smaller interface connectors for design and space savings. Push the limits for mobile low power systems -

Explore the Possibilities!

Saving Space, The Final Frontier

VIA embedded in your success

To find more information on EPIA mainboards, please visit www.VIAembedded.com or call 510 683 3300. Come visit us at 940 Mission Court Fremont CA 94539

ACCES I/O PRODUCTS, INC.

The source for all your I/O needs

To learn more about our Embedded Pico-I/O™ boards visit www.accesio.com or call 800 326 1649. Come visit us at 10623 Roselle Street San Diego CA 92121

Sumit™

Pico-I/O™

Pico-ITXe™

sffig

One of the responsibilities of board-level designers is to ensure that verification and failure-analysis engineers have adequate access to signals without resorting to drills, bed-of-nails testers, or focused ion beams. This requirement used to be simple when pads were farther apart, chips were less complex, and signals were more robust. Today, however, design for debugging is a major concern for creators of board-level products. Some signals—especially in high-speed serial interfaces—are virtually unprobable. Some are inaccessible in the fine-pitch tangle beneath high-pin-count chips. SOCs (systems on chips) lock within themselves some vital information. And some nets that used to require only occasional glances—power and clock networks, for example—have transformed under the pressure of advanced power-management techniques into signal nets that require functional verification. Today, design for debugging requires a plan and a systematic design approach from the beginning of the project.

DESIGNING AN ACCESSIBLE BOARD

DESIGN IN ACCESS TO VERIFICATION AND DEBUGGING DURING—NOT AFTER—THE DEVELOPMENT OF A BOARD-LEVEL PRODUCT.

BY RON WILSON
EXECUTIVE
EDITOR

This undertaking is not just another version of design for test. Test engineers want the quickest possible answer to one question: Can I ship this board? Verification and failure-analysis engineers have a different problem. They must be able to put the board through a sequence of states while observing its behavior, and, if the behavior is wrong, they must trace the problem back to its source. This requirement is far more demanding.

Reference designs, because they must be both accessible to the customer and nearly production-ready, present excellent studies for understanding how designers approach this problem. Video-processing vendor Stretch Inc, for example, is acutely aware of the challenge—and opportunity—of reference-design boards. "The reference design is more like an end product," says Ashish Thanawala, director of systems engineering at the company. "We share it with our customers, who want to take it to market as quickly as possible, and so they care about the size of the board. For that reason, we don't put in test points."

On the other hand, access to the reference design can be so well-planned that, even in its streamlined, near-production-ready form, it can be an adequate vehicle even for silicon debugging. "In this generation, we pretty much got away with using the reference design as a bring-up board," says Stretch's chief executive officer, Craig Lytle. "We are talking about whether to attempt that [approach] again with our next-generation design."

AT A GLANCE

- Modern boards sometimes don't give verification and debugging engineers the access they need.
- Designers must early and systematically address the problem.
- Different kinds of circuits require different access approaches.
- Software is the key weapon in the verification and debugging battle.

A SYSTEMATIC APPROACH

A successful compromise between manufacturing-ready compactness and high accessibility requires a lot of work, most of it in the early stages of the design. "This level of design requires an understanding of what the end customer needs the board to do," says Ken Havens, tool-marketing manager for ARM North America. "Often, you start out by exploring the features of the system SOC and understanding how the final system will use those features."

As is often the case, the first step in a complicated problem may be to partition it into many smaller complicated problems. One way to do this partitioning is to divide the board into regions of control and visibility. Categorize each group of nets on the board into one of the following categories: power grids, clock trees, analog signal or control nets, digital nets that you can observe and control through a processor, digital nets accessible to an FPGA, and digital nets not in either of these last two categories. You can now work out an access strategy for each group of nets based on which category it is in.

Debugging engineers' questions about

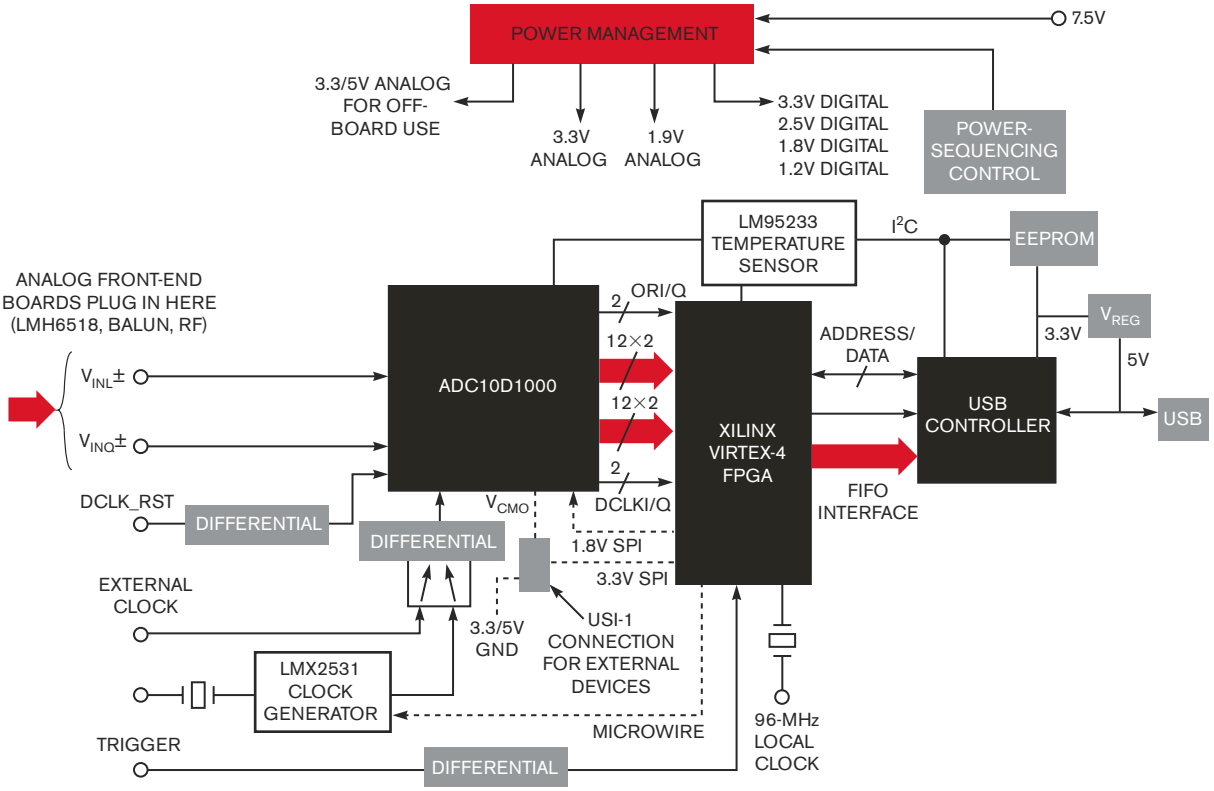

Figure 1 A reference-design board may have multiple supply voltages that must sequence properly on start-up.

power and clock signals used to be simple: Were they there, and were they within spec? The engineers needed an exposed trace, preferably near the largest load, for each supply voltage; a series resistor for current measurements; and an exposed trace on each clock. Nobody distributed high-speed clocks on a board, so they didn't have to worry about skew. Still, even on simple boards, engineers couldn't just probe around mindlessly on a power grid. "On any board, ground is a low-frequency noise source," warns Richard Fellner, a test engineer at IDT (Integrated Device Technology). Before you attach a probe and its ground somewhere, you need to think through what you might be coupling into the ground plane.

And you cannot just casually probe any old clock signal. "There are some clocks you pretty obviously can't mess with," observes Robbie Shergill, director of applications at National Semiconductor. "For instance, you can't probe clocks going into high-speed data converters. If you are going to need to see that clock, you will have to provide a replica."

With the advent of clock gating, dynamic loads, and power gating—all tools of power management—everything has become more complex. In effect, aggressive power management has turned clock and power nets into analog signal nets. Engineers may need to verify, for instance, that the clock to a chip is arriving only when the clock gate is supposed to be conducting. They also need to verify that the transitions between clock-gating modes don't cause glitches. Alternatively, they may need to verify that a power-gating sequence brings down the supplies to an SOC in the right sequence and on the right ramps. Such requirements, in turn, imply a need for greater access to clocks and power connections at the point of use.

Almost everyone faces one instance of this issue in the form of power-up sequencing. With perhaps a dozen supply voltages on a board, designers must provide enough access to verify that the supplies and clocks are all coming up in the right sequence (**Figure 1**). It sounds simple, but getting this much access can be a serious problem if, for example, the traces between a local regulator and the core-voltage pins on an FPGA are all buried—as they likely will be if you tell the layout person only to keep the regulator right next to the FPGA.

The following principle will keep showing up: Understand the functions the board must perform and plan for the measurements necessary to verify them. "You must have the test defined, including the signal characteristics and the instrumentation," says IDT's Fellner. "That information will determine what kind of connection you have to provide on the board."

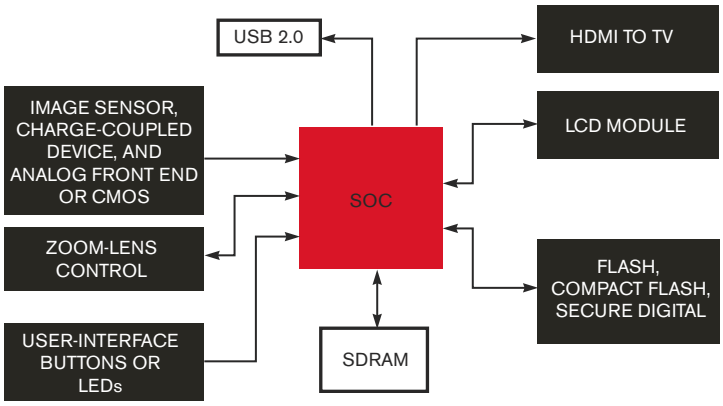

Figure 2 An SOC at the heart of a board may also be the best place to observe and control the outlying digital blocks.

UNDERSTAND THE FUNCTIONS THE BOARD MUST PERFORM AND PLAN FOR THE MEASUREMENTS NECESSARY TO VERIFY THEM.

ANALOG PATHS

In principle, planning for access to analog signal paths and control loops is similar to planning for power and clock grids. You must understand the functions the path performs, identify nodes the engineers will need to control and observe the path, and understand and provide for the measurements these activities require. But this simple statement can imply a lot of circuit analysis and an understanding of how your verification team will approach the path. And it means understanding measurement techniques that could include anything from high-voltage probes to RF network analyzers. "Many times, the choice of signals is pretty obvious," says National's Shergill. "It's a signal path. Which aspects of the signal do you need to observe—at which points?" It's always helpful to have access to the pins, Shergill adds, but the need to have the reference design as close as possible to manufacturing-ready means that you sometimes are just not going to get that level of access. "Fortunately, there's often other access to the same signal somewhere else," Shergill says. Sometimes, physical accessibility isn't the only issue. Some nodes are simply too delicate to probe. "We've had designs in which just the leakage in a chip is enough to disturb a high-impedance node," he warns. You don't want to be sticking a scope probe into this area.

In such cases, it may be necessary to replicate the original signal at a test point, even at the expense of adding components. "We almost always end up putting additional components on the board to improve access to some signals," Fellner says. One advantage, he points out, is that if you are buffering a signal anyway, you can engineer the test point to be where you want it for debugging, not where it happens to have landed in the board layout. Fellner emphasizes the phrase "engineer access." Don't just drop in test points, he says. "Sometimes you can do a quick calculation to understand what you need. Sometimes you need to run a simulation."

Chan Lee, vice president of large-

scale integration at Ambarella, says that designing accessibility is something you learn over time. "We mostly use back-of-the-envelope calculations to design test access," Lee explains. "There are too many variables to do an accurate simulation." Clearly, you must take into account the signal, the electrical environment at the node, the kinds of measurements debugging and verification engineers will have to make, and the equipment they will have to use.

Observing analog paths can be tricky, but controlling them can be even more challenging. If the signal originates off-board, you can use a signal generator with an appropriate matching network to drive the signal path. If the path begins at an onboard source, such as a thermal sensor in a critical chip, the problem may require more thought and, possibly, an analog switch to multiplex in an external reference. If the path begins at a DAC, it may be wise to provide an alternative way of jamming digital data into the converter.

HIGH-SPEED I/O

The problem with high-speed serial I/O is the difficulty of probing the signals at all. "We typically have several different image-sensor interfaces on a board, including MIPI [mobile-industry-processor interface] and LVDS [low-voltage-differential signaling]," says Lee. "Some of these [interfaces] you can probe with a scope on a test point, and some you can't."

As speeds increase, the problem becomes critical. "Some of these buses are now differential signals at 3 GHz or more," Fellner points out. "You can't probe them anywhere. If you need access, you have to design in a differential repeater." Even if you can capture the signals, though, they are difficult to interpret. "These signals are nondeterministic, so conventional automatic-test equipment has difficulties with them," he says. You need a matching receiver to recover the data.

Working with these high-speed serial links is similar to working with other analog signal paths that originate in digital circuitry. It is a huge advantage to be able to control the digital ends of the link from software. "Typically, we put in more ways to control high-speed I/Os from software on the host," Lee says.

FPGAs HAVE ONE ADVANTAGE CPUs LACK: YOU CAN ALTER BOTH THEIR HARDWARE AND THEIR SOFTWARE.

Software control usually includes a mechanism for forcing data patterns onto the bus and perhaps a loopback fixture so that the interface can check itself. Increasingly, however, these software techniques also exploit the fact that many high-speed interfaces offer programmable electrical characteristics. Even the simpler fast outputs may have programmable slew rate and drive strength. And more sophisticated transceivers have programmable pre-emphasis and equalization, sometimes with automated training sequences. These features exist so that the transceivers can adapt to the board. But you can also use this programmability to characterize the connection the transceiver sees for your own information or simply to retune the transceiver after you have attached a probe to the link.

THE DIGITAL NETS

The largest category of nets on most boards is circuits that you can control and observe from an external host processor, a local processor on the board, or an FPGA. This category is in the realm of software-assisted debugging and verification. "There is a significant role for software," says Stretch's Thanawala. "If you can detect problems in software, you don't need to probe the board. So our chip-development team tries to anticipate what the board-verification people will need."

This development is no surprise to ARM's Havens. ARM has watched as the evolution of software kernels, in-circuit emulators, and external debugging tools for the ARM cores has accelerated over the years and spread—from being simply about debugging software, to controlling and observing everything in an ARM-based SOC, to controlling and observing chips outside the SOC, as well. The tools were there, and verification and debugging engineers were quick to exploit them, often routing signals back into an SOC from other parts of the board to bring more of the design under the CPU's sphere of control.

Empowering the software to control digital blocks may require no additional hardware effort. Software-accessible registers usually control the digital hardware peripheral to a CPU. You must mentally walk through the functions the verification and debugging engineers will examine to ensure that the registers can exercise all the operations the verification team will need to watch. Observability is another matter. Thanawala suggests that a peripheral or an interface design should include an error register so that the driver software can detect and at least begin to diagnose problems. In addition, you might need to make the control state and internal registers of the device visible to the software, even if that task means routing them back into the general-purpose I/Os of the chip containing the CPU.

In this way, a processor—whether a stand-alone microprocessor, a microcontroller, or a core in an SOC—can be the heart of the access strategy for a board (**Figure 2**). But there is an even more powerful alternative: an FPGA. "You can make an FPGA the center of your design for accessibility," says Brent Przybus, director of platform marketing at Xilinx, a company that should know: It has not only counseled customers on using FPGAs as debugging controllers but also developed its own reference designs for its application-platform strategy. In part, the principle for using FPGAs as accessibility gateways is similar to the notion for CPUs: The chip by its nature sits at the junction of many key signals on the board and contains important state. This situation is likely to be true whether the FPGA serves as an SOC at the heart of the board, as a coprocessor,

Go to www.edn.com/091008cs and click on Feedback Loop to post a comment on this article.

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

or as an interface. FPGAs have one advantage CPUs lack, however: You can alter both their hardware and their software. That flexibility gives you access to internal nodes in logic functions. For example, Przybus points out, implementing a DDR3 DRAM interface in an FPGA is one of the few ways to get enough access to troubleshoot that fast memory interface. With this approach, you not only see internal nodes in the controller but also can force data patterns through the interface, fire off transactions, and bring out a copy of the data that is traversing the virtually unprobable serial links.

You can insert—or the verification team can reprogram the FPGA to insert—a virtual logic analyzer into the FPGA fabric, so an external multiprocessor can trigger, trace, and control nodes in the chip. "It makes sense to route signals from other parts of the board into the FPGA instead of out to a test header," Przybus says. If you don't feel like designing your own logic analyzer, your FPGA vendor has one available—in Xilinx's case, ChipScope. The next step in sophistication is to put the controlling CPU itself into the FPGA as a core. Both proprietary fast RISC cores and industry-standard microcontroller cores are now small enough to slip almost unnoticed into a design. The CPU core in the FPGA allows you to write both the control code for the logic analyzer and the verification or debugging routines for other parts of the board and to run the code without an external host computer.

These ideas are obviously valuable if you already have an FPGA in your design. You can simply reprogram it for verification and debugging mode, and hardware additions to the board should be minimal. Some designers are finding that it's worthwhile to move to a higher-pin-count FPGA to route more signals to or from the chip, extending the FPGA's reach to more of the design. Some even argue that this capability justifies using a low-cost FPGA in a design in which a standard product or an ASIC could do the job.

In design for accessibility, it is vital to understand the verification and debugging engineers' needs. Having a systematic approach and soliciting the views of other experienced designers can be a big help. It may be possible to prevent the oft-repeated question from evaluation, verification, and failure-analysis engineers: "How the heck am I supposed to get to that point?" **EDN**

FOR MORE INFORMATION

Ambarella
www.ambarella.com

ARM
www.arm.com

Integrated Device Technology
www.idt.com

Stretch Inc
www.stretchinc.com

Xilinx
www.xilinx.com

You can reach Executive Editor Ron Wilson at 1-510-744-1263 and ronald.wilson@reedbusiness.com.

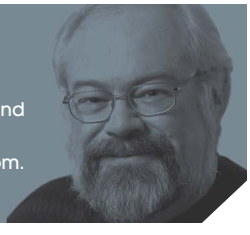

“I need a faster, more accurate audio analyzer.”

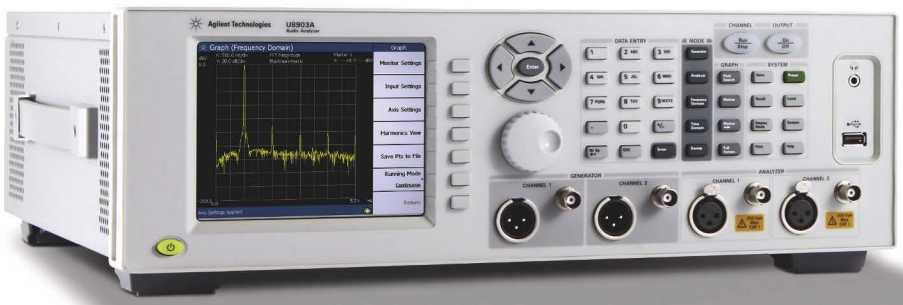

We knew you’d say that.

The new Agilent U8903A audio analyzer is four times more accurate and over seven times faster than its popular HP 8903B ancestor. In fact, it is among the fastest and most accurate currently in its class. It’s also a more fully capable audio analyzer, letting you make crosstalk measurements and phase measurements, FFT analysis, and graph sweep. And it will soon be upgradable to eight channels. You can’t find that anywhere else. Go online to learn more about this remarkable new audio analyzer.

Save $2,400
when you trade in your old HP 8903B for a new Agilent U8903A audio analyzer*
www.agilent.com/find/U8903Apromo

Trade in offer good for a limited time.

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

Agilent Technologies

BY NAVRAJ NANDRA • SYNOPSYS INC

IP quality lies beyond compliance testing

OF COURSE YOU WANT YOUR STANDARD-INTERFACE IP TO PASS COMPLIANCE TESTING. BUT THAT ACCOMPLISHMENT IS JUST THE BEGINNING. COMPLETE QUALITY ASSURANCE FOR IP CORES HAS FAR MORE CHALLENGES.

Compliance testing can reveal a great deal about whether an IP (intellectual-property) core functions according to the standard specifications, but does it really answer the most important question in the IP business, which is, Will passing compliance testing be enough to ensure that the IP will perform as expected in your chip? Long-term experience in providing IP to chip designers says that the answer to this question is "not quite."

Compliance testing alone cannot guarantee the quality or performance of IP for a number of reasons. This article, using the PCIe (Peripheral Component Interconnect Express) interface as an example, discusses the verification steps that engineers use in ensuring that the IP is of the highest quality. This process includes ensuring functional correctness, testing across corners, validating in silicon, passing compliance, and selecting an IP with advanced, built-in capabilities. All of these steps require a great deal of effort, specifically from the IP vendor.

THE VALUE OF COMPLIANCE TESTING

Given the amount of standards-based IP on SOCs (systems on chips), the ability to comply with standard specifications has become crucial. IP that complies with standards dominates the market. These standards can be de facto, such as microprocessor architectures, including ARM (www.arm.com); interconnect standards, such as USB (Universal Serial Bus), PCIe, DDR (double-data rate), SATA (serial-advanced-technology attachment), and Ethernet; or standard IC functions, such as standard cells, memories, and I/Os. The ability to use standards is the cornerstone of the IP industry (**Figure 1**). Standards-based IP enables designers to use plug-and-play techniques to take full advantage of the millions of transistors that advanced process-technology nodes provide.

Compliance testing has thus become vitally important for determining whether the standards-based IP functions as the specifications define and to confirm its interoperability with other devices. To take its standards-based IP through a compliance program, the vendor must have a silicon implementation of the IP in an appropriate test system. This test system should create a series of tests based on the specifications, including electrical specifications, such as a standard eye diagram (**Figure 2**).

Unfortunately, passing compliance and interoperability testing does not guarantee that the IP will work in your chip. The reasons range from basic functional issues to fabrication variations that the testing cannot verify. For example, compliance tests cannot check all of the functions the specification defines; they check only limited operating scenarios, and, specifically for IP, the tests do not include all configurations. Thus, compliance testing may confirm that a selected single-lane PCIe implementation passes the compliance and interoperability tests and works according to portions of the specification as defined by the standards orgainzations, but this test doesn't confirm that different configuration options built into the IP work.

Even when a version of the IP meets the standard, the tests do not capture its degree of robustness relative to the specification. This factor is especially important for mixed-signal IP—particularly for today's high-speed SERDES (serializer/deserializer) PHY (physical)-layer IP. For example, a PCIe PHY-IP core can pass the compliance tests by barely falling within the standard eye diagram and can have such a weak signal that it is less likely to work in your design than a PHY-IP core that provides plenty of margin relative to the specification.

That margin is necessary to accommodate the wide range of design-in, packaging, and board-level factors that tend to degrade signals. The ideal environment of compliance tests does

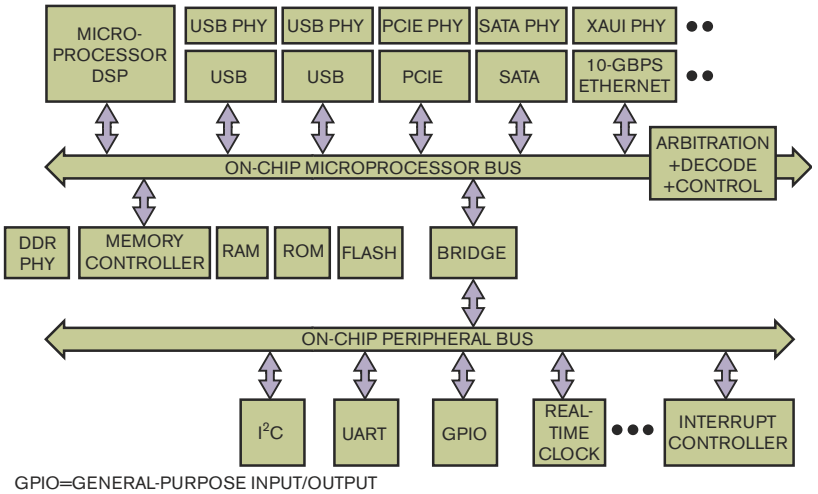

Figure 1 Standards-based IP may make up a significant portion of an SOC design.

not take these degradations into account. Moreover, the silicon for compliance tests may represent a beneficial combination of process corners, whereas your chips may tend toward corner combinations that can lead to problems. Manufacturing yield and reliability issues may result in the creation of many chips that you cannot use.

To address issues that may arise during its implementation into the overall system environment, the IP must go through much more than mere compliance testing. Ensuring high-quality IP involves a verification methodology that should consist of comprehensive simulation techniques, silicon validation, and proven interoperability. All of these factors should be a part of the IP vendor's development process.

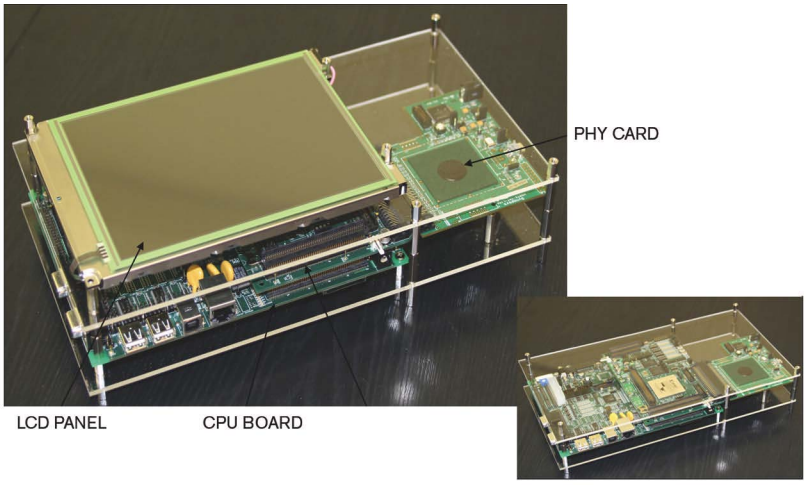

Figure 2 An appropriate test fixture is necessary for IP-compliance testing.

START WITH GOOD IP DESIGN

Experienced designers often say that you must design quality into a product from the beginning. So it is worth taking a look at how a vendor develops its IP and how experience says that the IP should be developed. Using PCIe as an example, a front-end-development flow may include many verification steps, starting with a verification plan early in the design (**Figure 3**). Rather than a separate subflow, verification is an integral part of the flow that influences the IP design at every step.

The design practices the flow uses are also critically important, and the best practices of design for reuse should be readily available. **Reference 1** details one resource, which lists best practices in reuse and SOC design based on the collective insights of Alcatel (www.alcatel-lucent.com), ARM, Atmel (www.atmel.com), Cadence (www.cadence.com), IBM (www.ibm.com), Infineon (www.infineon.com), LSI (www.lsi.com), Mentor Graphics (www.mentor.com), Philips (www.philips.com), STMicroelectronics (www.st.com), Synopsys (www.synopsys.com), Texas Instruments (www.ti.com), and several universities. The book encourages IP developers to use these best practices in their designs so that the IP is more likely to perform as they expect.

DIGITAL-IP VERIFICATION

Verification and validation are interesting to examine for both synthesizable digital IP and mixed-signal PHY IP. A good

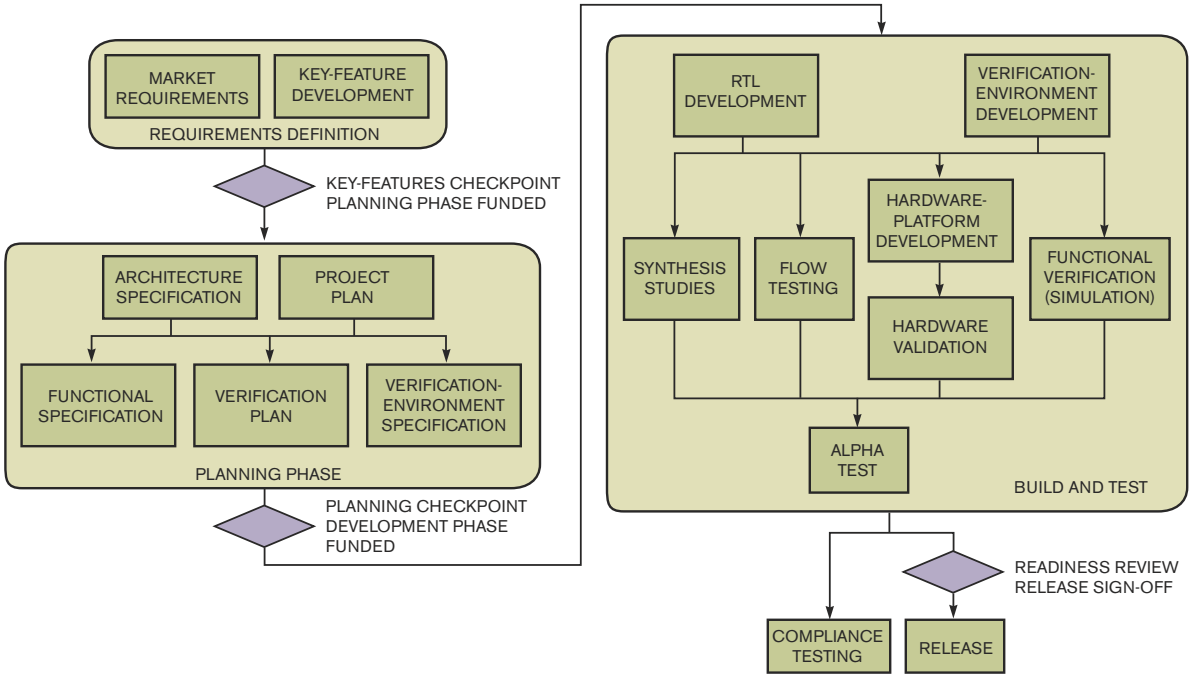

Figure 3 Verification may be complex, and it may start early in the design flow.

process for digital IP should begin with the generation of a functional-coverage check list (**Figure 4**). Verification engineers derive the check list from an interpretation of the protocol, design, and system-interface specifications, and this check list, in turn, drives the verification process.

Ensuring functional correctness typically requires a combination of techniques, such as simulation, writing assertions, emulation, and hardware prototyping. During simulation, verification engineers should conduct a series of random and directed tests to ensure maximum code coverage. The directed tests target functions within the blocks, including compliance-coverage tests, but often do not hit all of the functions of the device. Using constrained random verification techniques and functional-coverage points the verification team derives from the specifications, the team can methodically drive the IP in random simulations until it is fully verified, at almost all the corner conditions. With this approach, you can trace the verification-coverage points back to the functional-coverage check list and the primary specifications, and you can generate reports of progress against the specifications. The functional-coverage check list and report are key indicators of functional correctness, and they help ensure that the IP will work in your design.

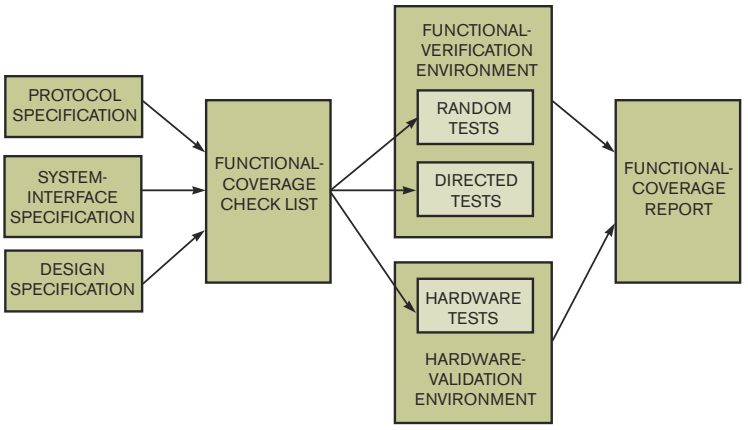

Figure 4 A functional-coverage check list informs the rest of the verification flow.

Using directed and random verification testing, Synopsys routinely runs 35 billion to 40 billion cycles of simulations per day on IP-core regressions. This work requires the company to have approximately 500 CPUs running regressions all the time. That process goes a long way toward ensuring correct functioning of the IP. It also ensures that customers, no matter what tool chain they use, will get the same results the IP provider does. Consequently, the full validation process includes thorough flow-testing of the IP with individual EDA tools, as well as testing with major simulators from multiple vendors.

Even with billions of cycles of simulation, the design still needs to work in real-world applications. For example, even though each PCIe interface originates from the same specification, some areas of the specification are not completely defined, or designers can interpret them differently. This problem is generally what leads to interoperability issues. An IP provider must test the IP by placing it into hardware, usually through an FPGA-prototyping board. Synopsys generally uses the same FPGA prototype for both hardware verification and compliance testing, so the FPGA implementation enables additional testing and running of interoperability tests, such as "plugfests," with other PCIe systems and available test equipment.

Unfortunately, it is difficult for an IP vendor to include every possible configuration of a flexible IP product into the FPGA prototype. So it is best to ask the IP provider not only what configurations it has tested in FPGAs but also what configurations its customers have placed into silicon and taken into volume production. Another customer may have already taken the configuration you need through compliance testing.

MIXED-SIGNAL-IP VERIFICATION

Although synthesizable digital IP offers plenty of verification challenges, challenges also arise when verifying analog/mixed-signal PHY IP. A look at a PCIe PHY at process nodes of 65 nm or smaller shows some fascinating trends, many of which involve physical-design dependencies.

One trend is the occurrence of variations due to changes in device size. Predicting these variations requires Monte Carlo simulations of threshold voltage and drain-to-source saturation current mismatches, and these simulations, in turn, demand large amounts of computation time and many tool licenses. Many other variations, including systematic variations due to physical effects, such as STI (shallow-trench-isolation) stress, WPE (well-proximity effect), and contact stress, also occur. The IP may also suffer from time-dependent variations due to NBTI (negative-bias temperature instability), HCI (hot-carrier injection), and EM (electromigra-

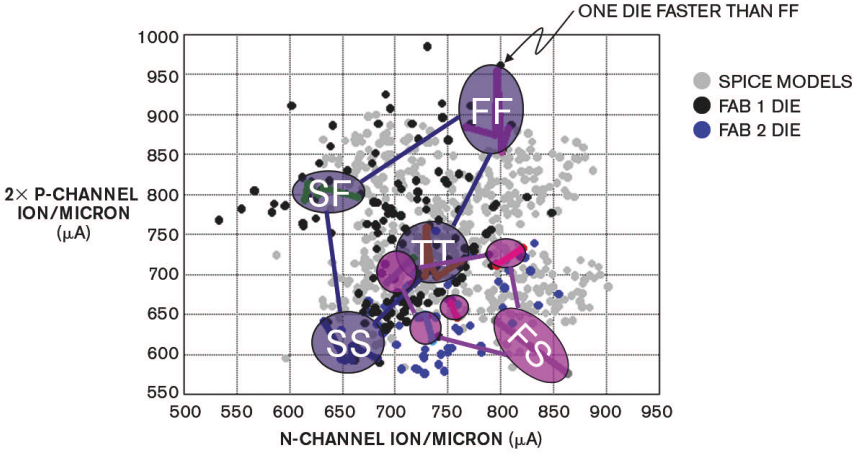

Figure 5 An on-chip process monitor can allow you to pick dice representative of each process corner.

tion). In the face of all of these variations, simulations that use prelayout parasitic-extraction parameters are becoming poor predictors of performance, simply because performance increasingly depends on physical factors that are subject to large layout-dependent variations. Postlayout parasitic-extraction effort is therefore increasing.

All of these variations in effect create process corners. But just as Monte Carlo simulations absorb computing time, simulating mixed-signal IP across process corners is an excellent way to consume huge amounts of computational resources. The simulation to verify the DesignWare PHY for PCIe, for example, covers 14 basic analog blocks. On average, the simulations require 18 test benches per block for a total of 250 test benches. Vendors must run each test bench across process, voltage, and temperature corners, including IR (current/resistance)-drop effects. Each test bench also includes independent corners for thin- and thick-gate devices, polysilicon resistors, and capacitance, as well as other variables. When you add up all of these considerations, each test bench runs across an average of 90 corners. This part of the PCIe PHY-layer verification process thus requires approximately 23,000 simulations.

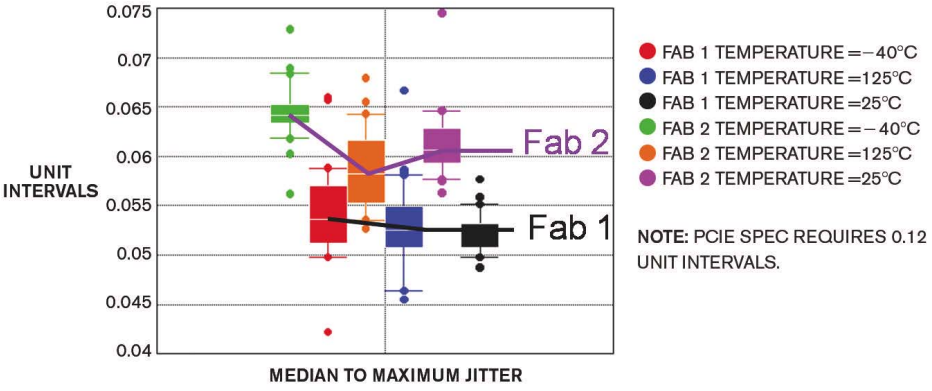

Figure 6 A process monitor can help visualize the range of jitter performance for dice.

Manufacturers may also do Monte Carlo simulation in addition to traditional corner analysis to assess yield and design centering. Altogether, the typical total simulation time for a DesignWare PHY layer for PCIe is approximately 3.25 CPU years. Following simulation is reliability analysis. Some companies analyze EM for both power and signal nets in postlayout parasitic-extraction simulation.

BUILT-IN INSURANCE

Because the performance of a high-speed PHY-layer IP core relies on a complex mix of difficult-to-predict factors, success may require more than just exhaustive analysis. You may need some new techniques for mixed-signal IP. For example, you can use an on-chip process monitor and an on-chip sampling scope in each die. The scope shows the on-chip signal from the PHY-layer chips, so you can determine its performance independently of packaging and board effects. The on-chip process monitor allows you to find out what corner combination you have on a die—information that is otherwise difficult to obtain. If you compare data from Spice models with on-chip monitor data for chips from two fabs, you can see where a die lies in the process (**Figure 5**). When you have this information, you can select dice representative of each process corner. You can predict yields and see where the center of the process is relative to Spice.

This ability enables more meaningful jitter analysis (**Figure 6**). Because jitter limits yield, chips must be better than the specification defines. If all your dice end up at the slow/slow corner, will your design still be OK? Without the on-chip monitor, it's virtually impossible to tell.

Even before good Spice models are available, you can use the information from the on-chip monitor to take better advantage of programmability in the analog blocks. In a PHY-layer chip, for instance, you may be able to adjust the dc level,the crossover point, and the rise and fall characteristic to make the block faster or slower. Pre-emphasis is useful, but only if you can implement it without causing ringing. Knowing each die's corner combination from the on-chip process monitor enables you to accurately and efficiently set the programmable PHY-IP characteristics.

PLAYING SAFE WITH IP

You are highly likely to achieve working silicon using IP that was designed and verified using the best design practices and methods. In addition, incorporating advanced functions, such as the on-chip abilities for mixed-signal IP, provides a clear means of differentiating IP from various vendors in the market. Although passing compliance certification is a requirement, it is not sufficient to guarantee IP quality. It is also mandatory to examine a vendor's verification/validation and design practices. For digital IP, a functional verification plan and check list are keys. For analog- and mixed-signal IP, a development method for robust designs and margins, with a clear path to high yields, is essential. EDN

REFERENCE

1 Keating, Michael, and Pierre Bricaud, *Reuse Methodology Manual for System-on-a-Chip Designs,* Third Edition, Kluwer Academic Publishers, 2002.

AUTHOR'S BIOGRAPHY

Navraj Nandra joined Synopsys in February 2005 as director of product marketing for the mixed-signal products that include SERDES, USB, and DDR2. He has worked in the semiconductor industry since the mid-'80s as an analog/mixed-signal-IC designer for Philips Semiconductors, austriamicrosystems, and EM-Marin. Before joining Synopsys, Nandra was director of application engineering at Barcelona Design. He holds a master's degree in microelectronics, with a focus on analog-IC design, from Brunel University (London) and a postgraduate diploma in process technology from Middlesex University (London). For information on Synopsys' DesignWare IP, visit www.synopsys.com/designware.

3MHz Buck Regulator Delivers 1A Load in a Tiny 2x2mm Package

Integrated FETs, Tiny Footprint and Minimal External Components

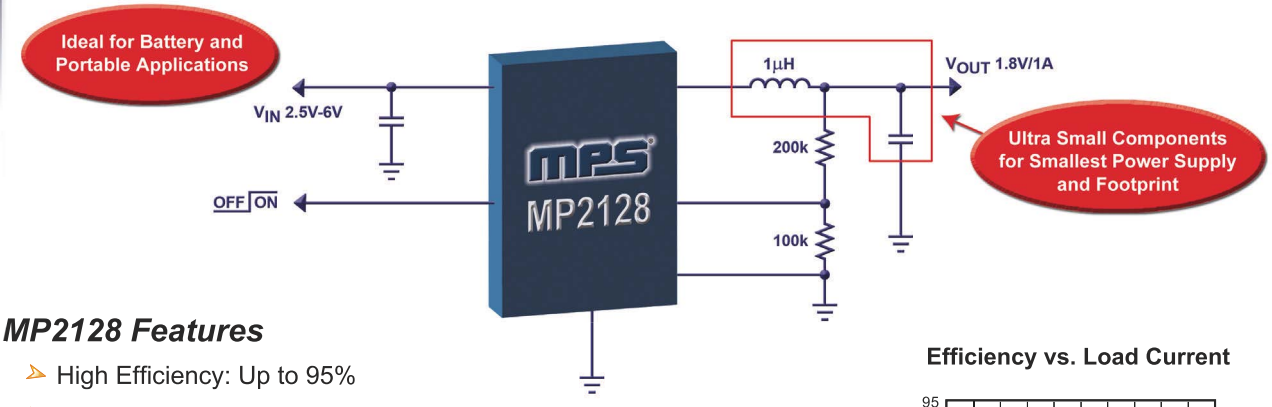

MP2128 Features

- High Efficiency: Up to 95%
- 100% Duty Cycle, Ideal for Battery Applications
- 3MHz Constant Switching Frequency for Smaller Inductor
- 1A Continuous Load Current
- 2.5V to 6V Input Voltage Range
- <1µA Shutdown Current
- Internal Soft-Start and Compensation
- Space Saving 5-Pin SOT23 Package and QFN8 2mm×2mm Package

Efficiency vs. Load Current

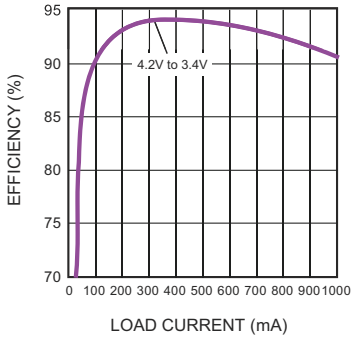

New Buck with Integrated FETs and Optional LDO Output

Part Numbers	V_{IN} Range	Load Current	Switching Frequency	Maximum Duty Cycle	Integrated LDO	Package
MP2128	2.5V-6V	1A	3MHz	100%	No	SOT23 & QFN 2x2mm
MP2116	2.5V-6V	1.5A	1.25MHz	100%	Yes	QFN 3x3mm
MP2126	2.5V-6V	2.5A	1.25MHz	100%	Yes	SOIC

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

ISO9001:2001 Certified

LEAD FREE RoHS COMPLIANT

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

AVNET® Memec

The path of least resistance

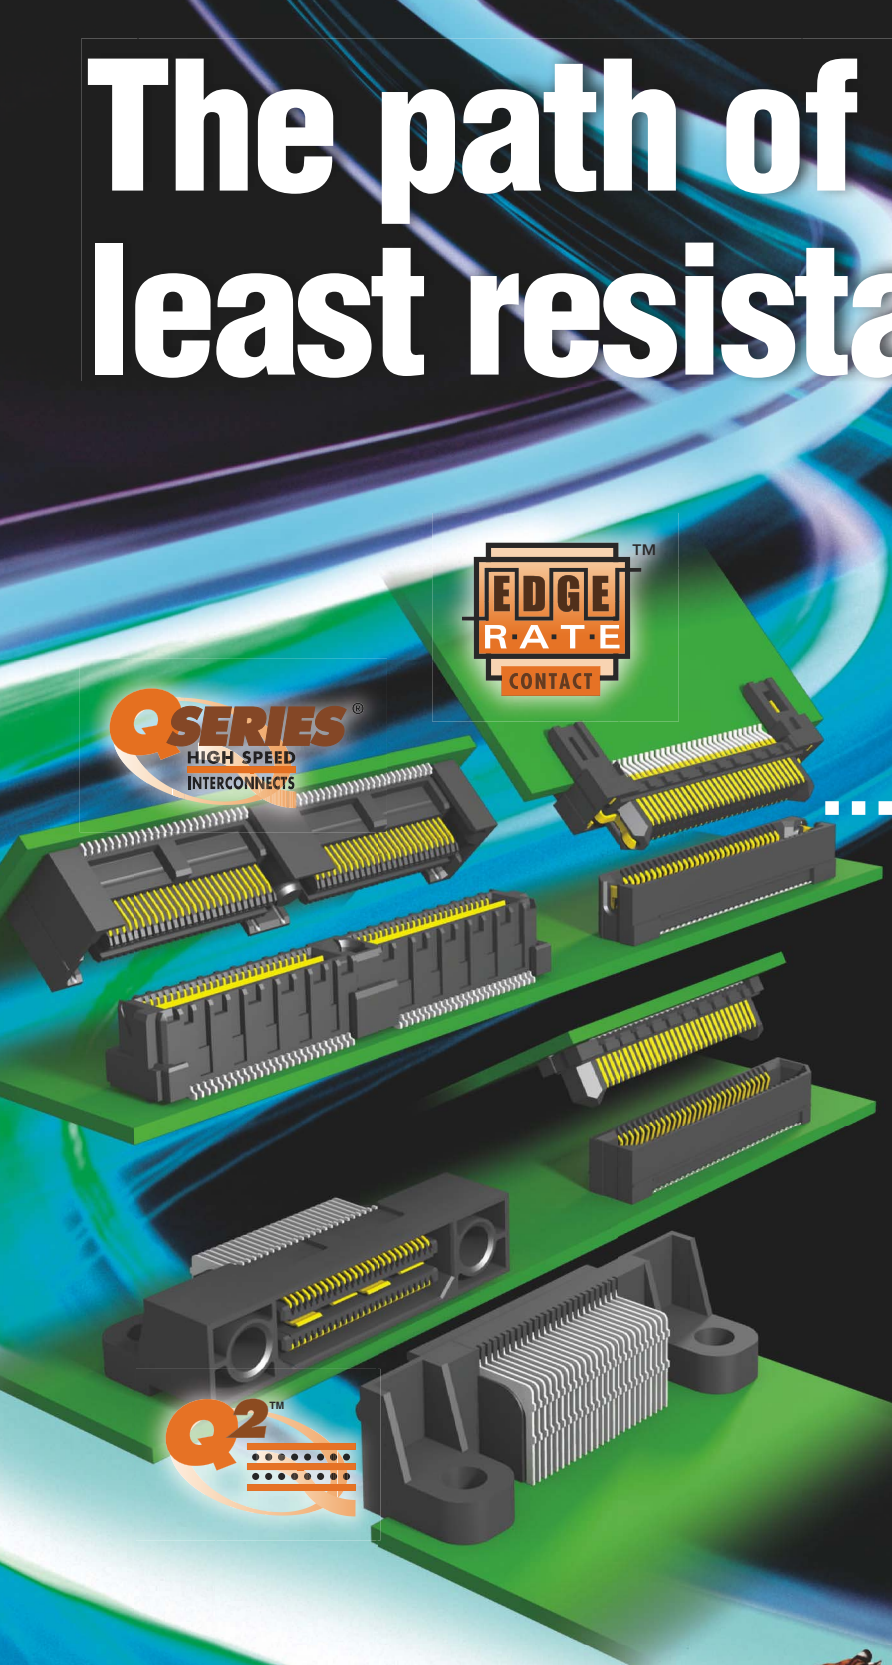

…for high speed board-to-board.

Q Series®

- Three pitches (.8mm, .635mm, .5mm), 5mm to 30mm stack heights
- Rated at 9 GHz single ended
- Integral power/ground plane

Q2™ Series

- Rugged design on .635mm pitch
- Rated at 9 GHz single ended

Edge Rate™ Series

- Contact system on .8mm pitch is robust when zippered during unmating
- Rated at 9.5 GHz in differential pairs

All products available in mezzanine, perpendicular and coplanar applications

samtec

TRANSMISSION LINE SOLUTIONS

www.samtec.com/tls

Deep packet inspection optimizes mobile applications

MOBILE DATA PROVIDES A NOTABLE GROWTH OPPORTUNITY FOR CELLULAR OPERATORS, BUT ADOPTION WILL BE SLOW UNLESS NETWORKS CAN PROVIDE GOOD SERVICE TO ALL SUBSCRIBERS. DEEP PACKET INSPECTION ENABLES OPERATORS TO UNDERSTAND THE APPLICATIONS THEIR SUBSCRIBERS ARE USING AND TO BUILD APPROPRIATE NETWORKS.

The communication industry widely recognizes mobile broadband as the growth engine for wireless carriers. As cell-phone penetration reaches saturation and voice revenue plateaus, mobile-data revenue is one segment that continues to grow. Mobile data is available through 3G (third-generation) networks, which are increasingly exhibiting the strain of the increased traffic. Most operators have announced plans to move to LTE (long-term-evolution) networks, which promise greater-than-100-Mbps data rates, a fivefold increase over 3G HSPA+ (high-speed-packet-access) networks. A quick look at customer-traffic patterns makes it clear that even network upgrades from 3G to LTE won't deliver sufficient data rates to guarantee good service to all users. However, additional technologies such as DPI (deep packet inspection) can ensure that prioritized and managed traffic optimizes the user experience.

Last year finally saw the explosive growth in wireless-data traffic and revenues that industry analysts have for years forecast. Wireless-data services are now a $130 billion global market, with $40 billion in North America. Research company IDC's (www.idc.com) projections for the growth in wireless-data revenues through 2012 are particularly welcome to wireless carriers, given the plateau in voice revenues in 2008 (**Figure 1**).

Network operators can use one of several techniques for adding capacity and meeting the growing demand for mobile broadband. First, they can increase spectral efficiency. LTE promises a fourfold improvement—in bits per seconds per hertz—over its WCDMA (wideband-code-division-multiple-access) predecessor. However, data rates are growing too fast for improvements in spectral efficiency to make more than a dent in the problem. It will take many years from the original release of LTE for full network rollouts and adoption to occur. Cisco (www.cisco.com) expects mobile-data traffic to increase 66-fold between 2008 and 2013 (**Reference 1**).

Another technique is to use more spectrum; wider channels mean more bandwidth. However, spectrum is an expensive resource, and most operators have little available spectrum to spare. Alternatively, operators can use cell-splitting techniques; adding smaller cell sites reduces subscriber density in a given cell site. By shrinking the cells, the number of subscribers per cell decreases, and the average bandwidth per subscriber accordingly increases. This approach is both the easiest and the most expensive option because radio access is the priciest portion of the operator network, with costs directly proportional to the number of cell sites.

Another alternative is to more efficiently allocate user bandwidth by employing new DPI technologies. DPI manages data networks and optimizes data traffic. It reaches beyond the IP (Internet Protocol) headers and examines the packet con-

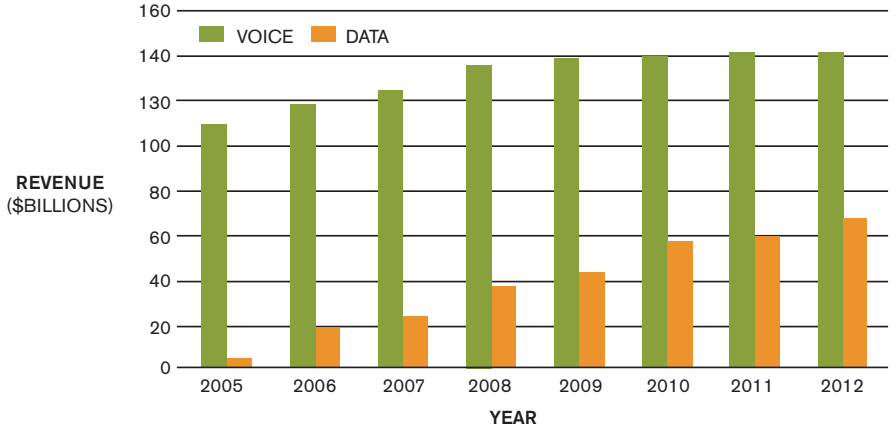

Figure 1 IDC's mobile-operator analysis reveals the disproportionate increase in wireless-data usage and revenues compared with traditional voice services.

tents. Although you can use this technique to look for keywords or other content that most consumers feel violates privacy norms, the most common use of DPI is simply to determine the application of the packet: Is it an e-mail, Web, video, or P2P (peer-to-peer) packet? The shortcomings of other approaches increase the attractiveness of more efficiently using the overall bandwidth and ensure that applications and subscribers fairly share bandwidth. DPI enables the implementation and enforcement of these policies.

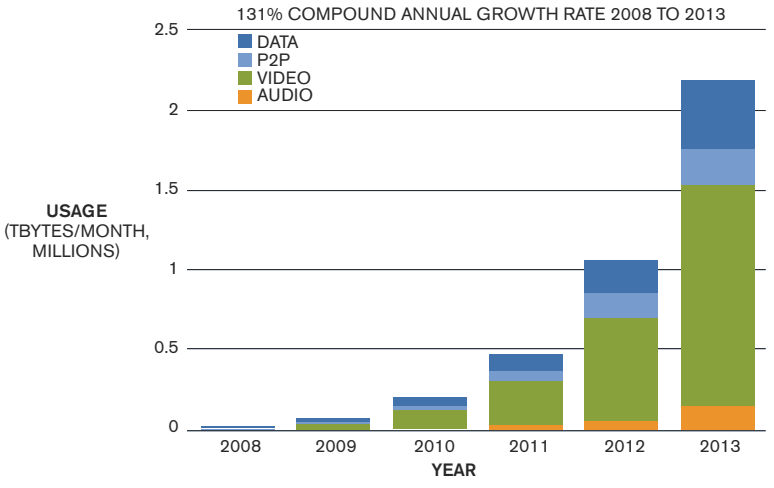

Figure 2 Cisco's forecast indicates that video services will see usage growth in the coming years.

THE REAL THREAT

Media coverage of mobile broadband tends to suggest that smartphones, such as the iPhone and BlackBerry, are driving the bulk of the congestion on the network. However, most of the data growth actually comes from laptop computers with data cards and USB (Universal Serial Bus) dongles. Operators that want to encourage mobile-data growth have been promoting the idea of mobile broadband, and some are even subsidizing networks with built-in cellular modems. Although a smartphone generates as much data traffic as 30 regular phones, laptops with data cards generate 15 times more traffic than smartphones, or as much traffic as 450 regular phones (**Reference 2**). Laptops are more conducive to Internet browsing and have larger, higher-resolution screens that demand higher-quality content. Laptop users are also more likely to run P2P applications that can consume huge swaths of bandwidth.

SHARED DATA CHANNEL

The developers of both 3G and LTE networks employed the common concept of a single data channel that all subscribers in a given cell share. They chose this architecture because they assumed that users would employ mobile data primarily for "bursty" activities, such as Web surfing and e-mail. In these activities, a shared data channel means that users get high bandwidth for downloads but also that users can employ the channel while other users read e-mail or scroll through Web pages. The shared data channel performs poorly when it encounters large sustained transfers. These transfers fill up the shared channel and lead to dropped packets and long latencies for every subscriber in the cell. With the growth in mobile data and broad usage of wireless-data cards, video and P2P activities, which fit the profile of large sustained transfers, are becoming more common. Video and P2P traffic currently accounts for 60% of all data and should grow to 74% by 2013, according to Cisco (**Figure 2**).

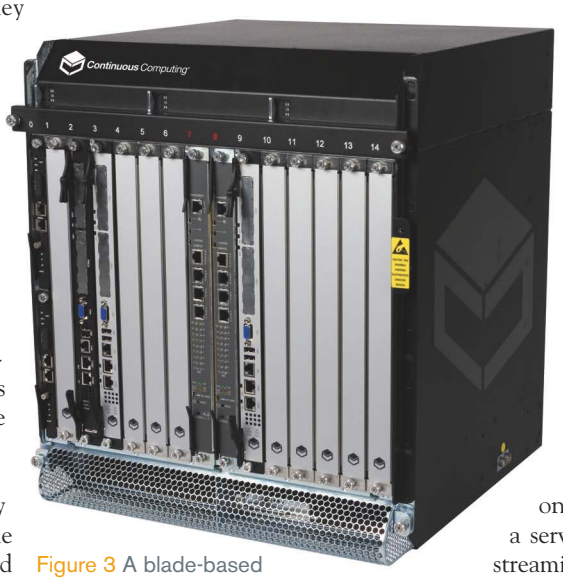

Figure 3 A blade-based appliance is amenable to scalable DPI capabilities (courtesy Continuous Computing).

In the early days of the Internet, IP headers clearly marked applications, but a combination of numbering limitations and corporate firewalls has encouraged application developers to mask the type of traffic. Most traffic today looks like Web traffic if you examine only the headers. Once operators determine the application and the subscriber, they can apply a range of policies to ease network contention, implement new data-service plans, or block traffic that violates the operator's terms of service. DPI allows operators to offer tiered data-service plans that they base on a range of criteria. Most mobile operators now offer one data plan and cap data usage at approximately 5 Gbytes/month. One simplistic alternative approach is to offer data plans with size caps of 5, 10, 25, and 100 Gbytes/month, for example. Operators can also offer plans that cap the mobile-broadband speed at levels such as 128 kbps, 256 kbps, and 1 Mbps.

By using DPI, operators can build plans showing an understanding of how customers use their service. You might be able to optimize some sample plans for Web surfing and e-mail sessions of approximately 64 kbps but with a tight bandwidth cap on any P2P traffic. Others might offer a service that allows YouTube-style video streaming at approximately 250 kbps but with limits on high-definition video streaming of 4 Mbps or higher. Another service might attract gamers, offering low latency

for gaming packets. Corporations might choose a premium package, which offers traffic priority in any cell site for e-mail, CRM (customer-relationship management), and other corporate applications. A service for P2P users would offer unlimited bandwidth during off-peak hours but tightly cap the P2P bandwidth during peak-usage times.

> Go to www.edn.com/ms4343 and click on Feedback Loop to post a comment on this article.
>
> For more technical articles, go to www.edn.com/features.

TECHNICAL CHALLENGES

DPI technology is simple in concept but complex in practice. At a conceptual level, looking at a packet to determine the application and subscriber and then taking action on that identification sounds easy. The complexity in DPI arises from network line rates and rapidly evolving applications. The packet rates in carrier networks are staggering. A single 10-GbE (gigabit-Ethernet) channel can support 30 million packets/sec with minimum-sized packets. Even with a more realistic traffic profile and packet sizes of 200 bytes, that 10-GbE channel has 10 million packets/sec.

At that speed, the system has only 100 nsec to receive and inspect the packet, determine its application, modify it if necessary, and send it. Assuming a modern, 3-GHz, single-core processor, this time frame equates to only 300 instructions' worth of execution, which is usually not enough to even receive the packet. This reality has driven the adoption of multicore, multithreaded processors for packet inspection. With 32 cores, or threads, attacking the problem, even at a more modest 1-GHz core-clock speed, that same system now must process only 300,000 packets per core/sec and a more reasonable 3200 clock cycles/packet, enough for inspection, classification, and modification.

Even with this performance, carriers are demanding systems that can today process 40 Gbps of traffic, increasing to 100 or 200 Gbps by 2011. This trend has pushed the industry toward bladed approaches because a single appliance might handle a few gigabits per second or even 10 Gbps, but not 20, 40, or 100 Gbps. A blade-based DPI system can scale to 80 Gbps of traffic handling, with each DPI blade handling 10 Gbps of traffic (**Figure 3**).

The second challenge in DPI is reliably identifying applications based solely on the traffic flow. This feature is paramount because the DPI system might be using the application type to set the priority of the packets, decide which packets to drop if congestion arises, bill a customer based on different applications, and even block certain applications. Given the potential impact of misidentification, it is critical to accurately classify as many of the applications as possible and to minimize the number of false application matches. Yet this requirement must take place in an environment in which developers create applications every day and in which different applications are popular in different parts of the world.

To make matters worse, certain applications, particularly some P2P variants, actively try to disguise their identities to thwart corporate firewalls and traffic-shaping systems. Solving this problem requires a database of application signatures as well as a team that can actively update these signatures upon detection of new applications. The DPI-system vendor can address the requirements; alternatively, third parties that specialize in traffic identities are often willing to license application signatures.

ADAPTIVE TRAFFIC SHAPING

One example of an advanced DPI application is adaptive traffic shaping. In any mobile network, peer-site cells load others at any time. Cell-site loading depends on location, the number of users, and the types of applications, and it varies over time. As an extreme example, a cell site serving a sports stadium will see huge surges in traffic during a game but is quiet the rest of the time. Similarly, a cell site covering a business park is most active during the day, whereas a cell site covering a suburb gets busier during the evening.

Traffic shaping in wireless networks is inherently more complex than in wire-line networks because of the dynamic of variable loading and capacity over time. Adaptive traffic shaping enhances wire-line traffic shapers by providing them with visibility into this additional dimension, thereby allowing operators to enact policies that employ knowledge of cell-site loading. Operators might block P2P traffic and downgrade video traffic during peak stadium hours, for example, to ensure that enough capacity remains for users to download e-mail or browse the Web. Conversely, bit-heavy services can run at full speed when the network is not busy. EDN

REFERENCES

1. "Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update," January 2009, www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.html.
2. Roberts, Mike, *Future Mobile Broadband*, Third Edition, Informa Telecoms and Media, April 2009, http://shop.informatm.com/marlin/30000001001/MARKT_EFFORT/marketingid/20001744820.

AUTHOR'S BIOGRAPHY

Mike Coward became chief technology officer at Continuous Computing in June 2006, having previously served as general manager of the company's platform-business unit. Since co-founding the company in 1998, he has held senior-leadership positions in technology and engineering. Coward specializes in system architecture and the design of highly available redundant platforms and is a prolific speaker and author. Previously, he was the lead designer for an experiment that flew in 1996 on the National Aeronautics and Space Administration's Space Shuttle Endeavour. He has a master's degree in electrical engineering from the California Institute of Technology (Pasadena, CA).

2 GHz Clock Generator

CG635...$2490 *(U.S. list)*

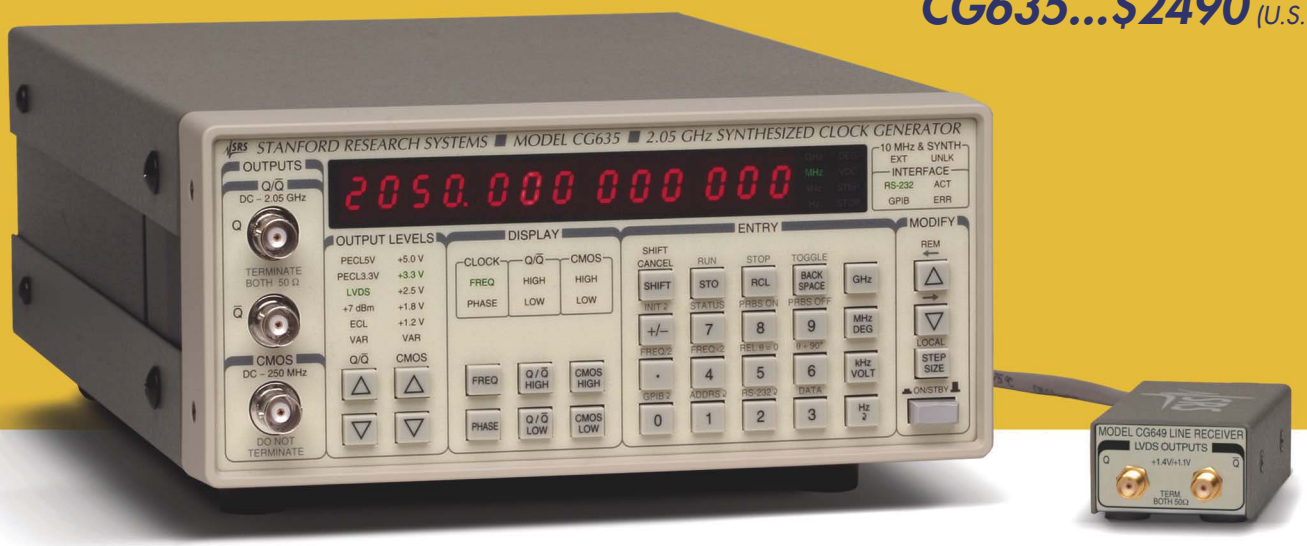

- ***Square wave clocks from DC to 2.05 GHz***
- ***Random jitter <1 ps (rms)***
- ***80 ps rise and fall times***
- ***16-digit frequency resolution***
- ***CMOS, LVDS, ECL, PECL, RS-485***
- ***Phase adjustment & time modulation***

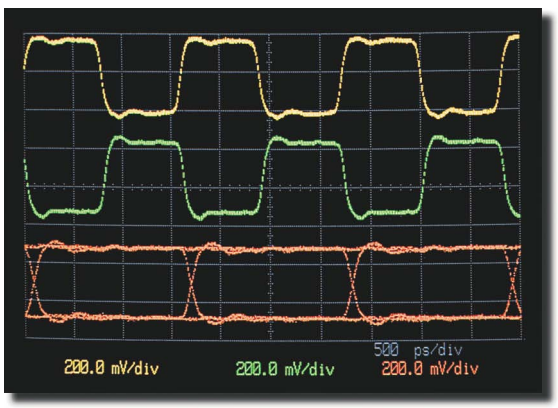

Plot shows complementary clocks and PRBS (opt. 01) outputs at 622.08 Mb/s with LVDS levels. Traces have transition times of 80 ps and jitter less than 1 ps (rms).

The CG635 generates clock signals—flawlessly. The clock signals are fast, clean and accurate, and can be set to standard logic levels.

How fast? Frequency to 2.05 GHz with rise and fall times as short as 80 ps.

How clean? Jitter is less than 1 ps and phase noise is better than –90 dBc/Hz (100 Hz offset) at 622.08 MHz.

How accurate? Using the optional rubidium timebase, aging is better than 0.0005 ppm/year, and temperature stability is better than 0.0001 ppm.

You would expect an instrument this good to be expensive, but it isn't. You no longer have to buy an rf synthesizer to generate clock signals. The CG635 does the job better—at a fraction of the cost.

Stanford Research Systems

Phone: (408) 744-9040 · Fax: (408) 744-9049 · info@thinkSRS.com · **www.thinkSRS.com**

Tiny Octal DAC with 10ppm/°C Reference

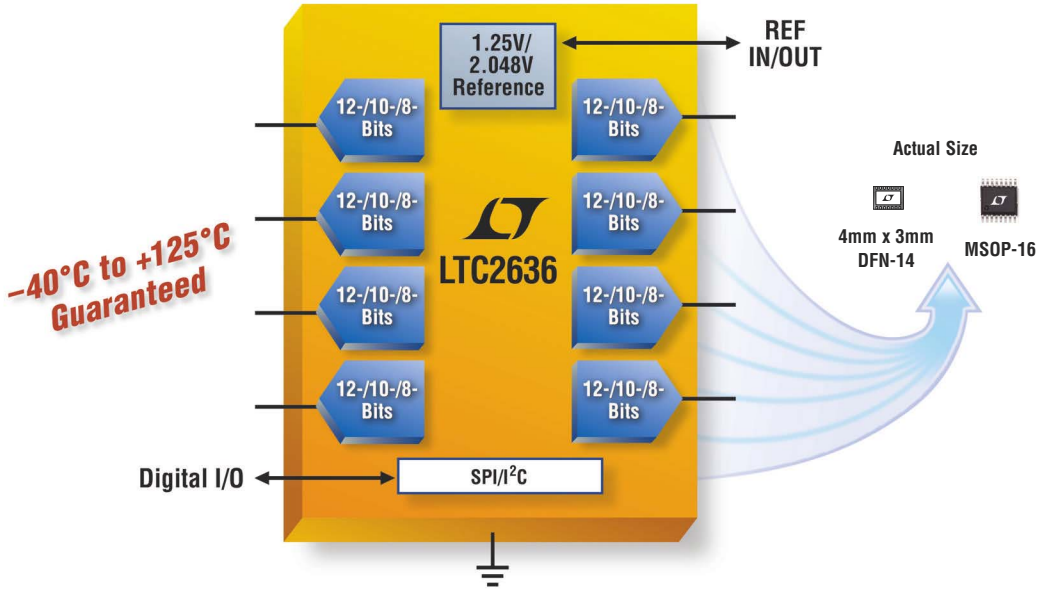

12-, 10- & 8-Bit DACs Reduce Solution Size by Half

Our LTC®2636 DAC family includes singles, duals, quads and octals with ordering options for 12-, 10- or 8-bits of resolution, internal 1.25V or 2.048V 10ppm/°C reference, SPI or I²C interface, power-on reset to zero- or mid-scale, and commercial, industrial or automotive guaranteed temperature specifications. The smallest solution available, these DACs are ideal for industrial, optical networking and automotive applications. Customize your design today.

Features

- Internal Precision Reference: 1.25V or 2.048V, 10ppm/°C
- 2.7V to 5.5V Supply Range
- Ultralow Crosstalk Between DACs (<2.4nV•s)
- –40°C to +125°C Guaranteed Temperature Range
- Power-on Reset to Zero-Scale or Mid-Scale
- 16-Pin MSOP, 4mm x 3mm DFN-14 Packages

Multichannel DAC Family

Part Number	Bits	I/O	DACs
LTC2636	12/10/8	SPI	8
LTC2637	12/10/8	I²C	8
LTC2634	12/10/8	SPI	4
LTC2635*	12/10/8	I²C	4
LTC2632*	12/10/8	SPI	2
LTC2633*	12/10/8	I²C	2
LTC2630/ LTC2640	12/10/8	SPI	1
LTC2631	12/10/8	I²C	1

*Future Products. Contact factory for availability.

Info & Free Samples

www.linear.com/2636

1-800-4-LINEAR

Free Industrial Signal Chain Brochure

www.linear.com/indsolutions

LINEAR TECHNOLOGY

EDITED BY MARTIN ROWE
AND FRAN GRANVILLE

designideas

READERS SOLVE DESIGN PROBLEMS

DIs Inside

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

Unused port adds a PWM/analog channel to a microcontroller

Vishwas Vaidya, Tata Motors Ltd, Pune, India

Low-cost, 8-bit, single-chip microcontrollers are stingy when it comes to on-chip PWM (pulse-width-modulation) resources. The use of a PWM resource often forces a designer to sacrifice a capture/compare or timer channel because the PWM channel shares the same on-chip hardware. This Design Idea describes how you can use an on-chip unused synchronous serial port to generate PWM signals and convert them to a slow-moving analog signal (**Figure 1**). Many microcontroller-based stand-alone electronic units don't use the synchronous serial port. Thus, you can use the microcontroller's baud-rate generator and parallel-to-serial-converter blocks to generate bit patterns to form a 256-bit PWM pattern. You can then filter the PWM output with an RC filter to extract an analog signal (**Reference 1**). The synchronous communication is devoid of the start and stop bits of asynchronous mode, so the bit pattern can generate long periods of high or low level.

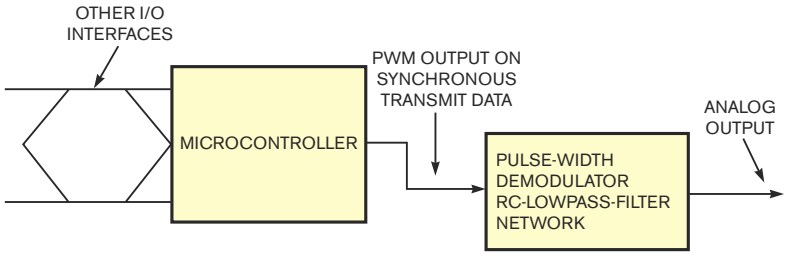

Figure 1 You can use an on-chip unused synchronous serial port to generate PWM signals and convert them to a slow-moving analog signal.

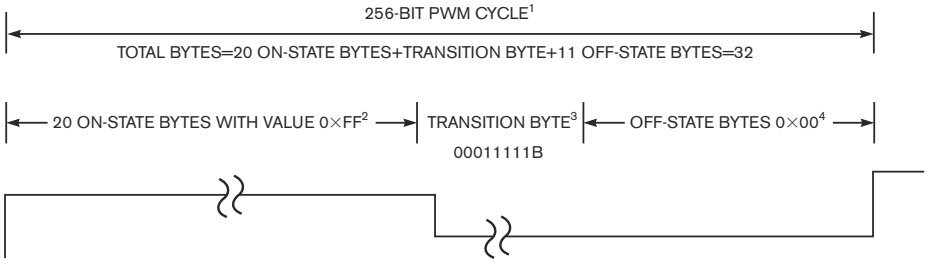

[1]RAW DATA=165; ON-STATE=5V; OFF-STATE=0V.
[2]NUMBER OF ON-STATE BYTES=165/8=20 (INTEGER DIVISION).
REMAINDER=165 – (8×20)=165 – 160=5.
[3]TRANSITION BYTE=00011111B=0×1F (NOTICE FIVE ONES FROM LSB SIDE).
[4]NUMBER OF OFF-STATE BYTES=TOTAL 32 BYTES – ONE TRANSITION BYTE – 20 ON-STATE BYTES=32 – 1 – 20=11 BYTES.
ANALOG OUTPUT AFTER LOWPASS FILTER=(165/256)×5=3.22V.

Figure 2 You can generate raw data with a decimal value of 165 using this concept.

You can generate raw data with a decimal value of 165 using this concept (**Figure 2**). A PWM-conversion cycle consists of generating 256 bits—that is, 32 bytes. The number of "on" bits corresponds to the value of the raw data to convert into PWM. Hence, for 165 bits as the raw data, 165 bits are on and 91

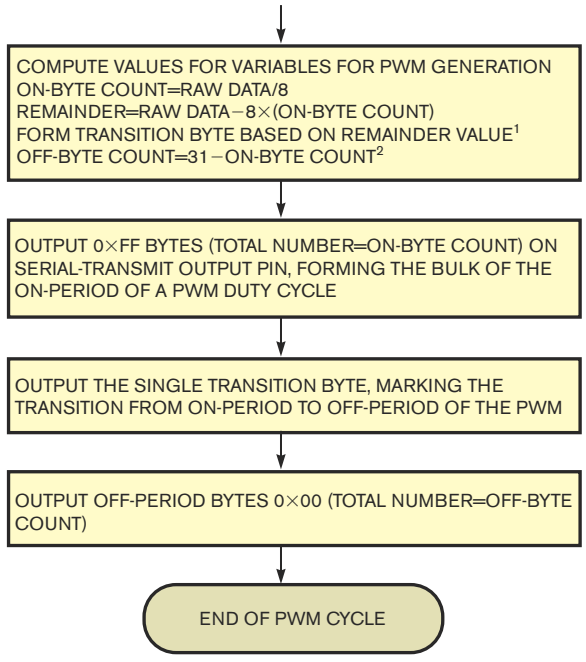

[1]TRANSITION BYTE CONSISTS OF A NUMBER OF ONE BITS STARTING FROM LSB. REMAINING BITS ARE ZERO. THE VALUE OF THE REMAINDER DECIDES THE NUMBER OF ONE BITS. IF THE REMAINDER IS THREE, THE TRANSITION BYTE IS 00000111B. FOR A REMAINDER OF FIVE, THE TRANSITION BYTE IS 00011111B. THE LSB TRANSMITS FIRST.

[2]FOR A SINGLE PWM CYCLE, TOTAL BITS=256 AND TOTAL BYTES=256/8=32. BECAUSE THE TRANSITION BYTE OCCUPIES 1 BYTE, THE OFF-BYTE COUNT=31 – ON-BYTE COUNT.

Figure 3 You can tailor the PWM frequency to your application by selecting a crystal, PLL, and baud rate.

bits are off. To generate a 165-bit on-period, the first 20 bytes—that is, 160 bits—transmit as 0×ff on-state bytes. The trick lies in judiciously composing the 21st, or transition, byte. This byte has some of its LSBs (least significant bits) as ones and the rest as zeros to form the required length of the on-period. In this case, the circuit needs five more on bits: 160+5=165. Hence, the transition byte should have a 00011111b pattern (byte=0×1f).

Figure 3 illustrates the process in flow-chart form. You can tailor the PWM frequency to your application by selecting a crystal, PLL (phase-locked loop), and baud rate. A simple RC filter can convert the PWM into a slow-moving analog value. Although this idea describes an 8-bit PWM, you can increase or decrease resolution by changing the total bits per PWM cycle. You correspondingly increase or decrease the conversion time.

Listing 1, which is available at www.edn.com/091008dia, provides a sample code for illustrating the concept. The code uses the Microchip (www.microchip.com) PIC18F4525, which has a 4-MHz crystal and 10-kHz baud rate for the synchronous serial communication, yielding 10,000/256=39.31 Hz of PWM frequency. You can filter it with a 0.1-sec RC filter, which is sufficient for slow-moving analog signals, such as speed setpoints for motion-control applications. By using a 20-MHz crystal, you can achieve synchronous serial baud rates greater than 1.5 MHz and PWM frequencies of a few kilohertz. EDN

REFERENCE

Mitchell, Mike, "Make a DAC with a microcontroller's PWM timer," *EDN*, Sept 5, 2002, pg 110, www.edn.com/article/CA240913.

Capacitance meter uses PLL for high accuracy

Jim McLucas, Broomfield, CO

An old *Electronics Designer's Casebook* described a circuit that provided capacitance measurements of 10 pF to 1 µF with 1% accuracy (**Reference 1**). A number of issues emerged with the circuit during testing, and this Design Idea describes an improved circuit. The meter circuit in **Figure 1** (pg 48) lets you measure capacitance from 10 pF to 10 µF with high accuracy. It needs no microprocessor; thus, it needs no code. Even in the 1- to 10-pF range, the circuit is accurate to about ±1 pF when reading values as low as 5 pF.

The circuit requires a high-input-impedance device to interface with the high-value resistors, R_6, R_8, R_9, and R_{10}, and a fast comparator to interface with the PLL (phase-locked loop). IC_1, an Analog Devices (www.analog.com) AD8033 op amp, does the job because of its 1000-GΩ input impedance and 1.7-pF input capacitance. It also has only 50 pA of input bias current over temperature. Its 80-MHz bandwidth and 80V/µsec slew rate are more than enough for this application. It can operate with just an 8V power supply. Unfortunately, the AD8033 is available only in surface-mount packages, which makes breadboarding somewhat tedious. IC_2, an Analog Devices ADCMP601 comparator, interfaces

DESIGN NOTES

One Device Replaces Battery Charger, Pushbutton Controller, LED Driver and Voltage Regulator ICs in Portable Electronics

Design Note 470

Marty Merchant

Introduction

The LTC®3577/LTC3577-1 integrates a number of portable device power management functions into one IC, reducing complexity, cost and board area in handheld devices. The major functions include:

- Five voltage regulators to power memory, I/O, PLL, CODEC, DSP or a touch-screen controller
- A battery charger and PowerPath™ manager
- An LED driver for backlighting an LCD display, keypad and/or buttons
- Pushbutton control for debouncing the on/off button, supply sequencing and allowing end-users to force a hard reset when the microcontroller is not responding

By combining these functions, the LTC3577/LTC3577-1 does more than just reduce the number of required ICs; it solves the problems of functional interoperability—where otherwise separate features operate together for improved end-product performance. For instance, when the power input is from USB, the limited input current is logically distributed among the power supply outputs and the battery charger.

The LTC3577/LTC3577-1 offers other important features, including PowerPath control with instant-on operation, input overvoltage protection for devices that operate in harsh environments and adjustable slew rates on the switching supplies, making it possible to reduce EMI while optimizing efficiency. The LTC3577-1 features a 4.1V battery float voltage for improved battery cycle life and additional high temperature safety margin, while the LTC3577 includes a standard 4.2V battery float voltage for maximum battery run time.

Pushbutton Control

The built in pushbutton control circuitry of the LTC3577/LTC3577-1 eliminates the need to debounce the pushbutton and includes power-up sequence functionality. A PB Status output indicates when the pushbutton is depressed, allowing the microprocessor to alter operation or begin the power-down sequence. Holding the pushbutton down for five seconds produces a hard reset. The hard reset shuts down the three bucks, the two LDOs and the LED driver, allowing the user to power down the device when the microprocessor is no longer responding.

Battery, USB, Wall and High Voltage Input Sources

The LTC3577/LTC3577-1 is designed to direct power from two power supply inputs and/or a Li-Ion/Polymer battery. The V_{BUS} input has selectable input current limit control, designed to deliver 100mA or 500mA for USB applications, or 1A for higher power applications.

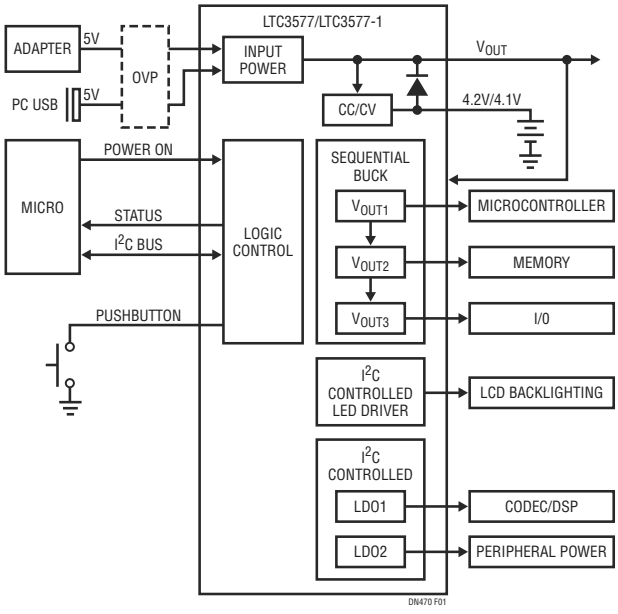

Figure 1. Portable Device Power Distribution Block Diagram Featuring the LTC3577/LTC3577-1

A high power voltage source such as a 5V supply can be connected via an externally controlled FET. The voltage control (V_C) pin can be used to regulate the output of a high voltage buck, such as the LT3480, LT3563 or LT3505 at a voltage slightly above the battery for optimal battery charger efficiency.

Figure 1 shows a system block diagram of the LTC3577/LTC3577-1. An overvoltage protection circuit enables one or both of the input supplies to be protected against high voltage surges. The LTC3577/LTC3577-1 can provide power from a 4.2V/4.1V Li-Ion/Polymer battery when no other power is available or when the V_{BUS} input current limit has been exceeded.

Battery Charger

The LTC3577/LTC3577-1 battery charger can provide a charge current up to 1.5A via V_{BUS} or wall adapter when available. The charger also has an automatic recharge and a trickle charge function. The battery charge/no-charge status, plus the NTC status can be read via the I^2C bus. Since Li-Ion/Polymer batteries quickly lose capacity when both hot and fully charged, the LTC3577/LTC3577-1 reduces the battery voltage when the battery heats up, extending battery life and improving safety.

Three Bucks, Two LDOs and a Boost/LED Driver

The LTC3577/LTC3577-1 contains five resistor-adjustable step-down regulators: two bucks, which can provide up to 500mA each, a third buck, which can provide up to 800mA, and two LDO regulators, which provide up to 150mA each and are enabled via the I^2C interface. Individual LDO supply inputs allow the regulators to be connected to low voltage buck regulator outputs to improve efficiency. All regulators are capable of low-voltage operation, adjustable down to 0.8V.

The three buck regulators are sequenced at power up (V_{OUT1}, V_{OUT2} then V_{OUT3}) via the pushbutton controller or via a static input pin. Each buck can be individually selected to run in Burst Mode® operation to optimize efficiency or pulse-skipping mode for lower output ripple at light loads. A patented switching slew rate control feature, set via the I^2C interface, allows the reduction of EMI noise in exchange for efficiency.

The LTC3577/LTC3577-1 LED boost driver can be used to drive up to 10 series white LEDs at up to 25mA or be configured as a constant voltage boost converter. As a LED driver, the current is controlled by a 6-bit, 60dB logarithmic DAC, which can be further reduced via internal PWM control. The LED current smoothly ramps up and down at one of four different rates. Overvoltage protection prevents the internal power transistor from damage if an open circuit fault occurs. Alternatively, the LED boost driver can be configured as a fixed voltage boost, providing up to 0.75W at 36V.

Many circuits require a dual polarity voltage to bias op amps or other analog devices. A simple charge pump circuit, as shown in Figure 2, can be added to the boost converter switch node to provide a dual polarity supply. Two forward diodes are used to account for the two diode voltage drops in the inverting charge pump circuit and provide the best cross-regulation. For circuits where cross-regulation is not important, or with relatively light negative loads, using a single forward diode for the boost circuit provides the best efficiency.

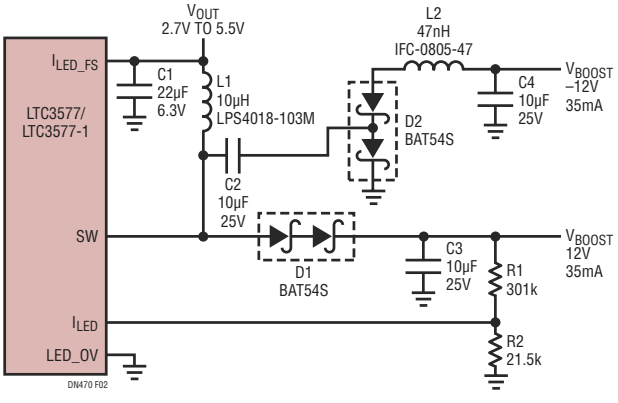

Figure 2. Dual Polarity Boost Converter

Conclusion

The high level of integration of the LTC3577/LTC3577-1 reduces the number of components, required board real estate and overall cost; and greatly simplifies design by solving a number of complex power flow logic and control problems.

Data Sheet Download

www.linear.com

For applications help,
call (978) 656-4778

with the AD8033 op amp and IC_3, a 74HC4606A PLL. The comparator has a typical propagation delay of only 4.3 nsec. It has built-in hysteresis and needs only a 5V supply. It is also available only in surface-mount packages.

The capacitance meter generates two signals; one of them lags the other by 60°. A 3-bit, self-correcting, divide-by-six twisted-ring counter comprising IC_6, IC_7, and IC_{13B} provides the lagging signal. The lagging signal connects to the COMP input of the PLL (Pin 3), and the other signal is applied to an RC circuit, which provides a 60° phase lag before it gets to the SIG input of the PLL (Pin 14). The PLL adjusts the frequency of its VCO (voltage-controlled oscillator) so that the two input signals are in phase. The resulting period of the VCO's output signal (Pin 4) is proportional to the measured capacitance.

On the low-capacitance range, signals with frequency F_O are applied to the PLL. On the high-capacitance range, the frequency is $F_O/1000$. IC_8 through IC_{10} provide the division, and S_2, IC_{4B} through IC_{4D}, IC_{5D}, IC_{5E}, and the associated components provide the high-capacitance/low-capacitance range switching. The VCO of the PLL runs at $6F_O$. The circuit divides this signal by three to provide an output with a period that's proportional to the measured capacitance. It provides the correct digits when you measure with a frequency counter that you set to measure the period. You can calculate F_O or $F_O/1000$ from $0.1505/R_XC_X$, where R_X is R_6, R_8, R_9, or R_{10}, depending on the selected range.

The 74HC4046A PLL can exhibit several problems. For example, it may not start when you apply power, or it may hang with the VCO running with the VCO-input pin (Pin 9) stuck high or low. The start-up circuitry, comprising IC_{13F}, Q_4, and associated components, applies a positive voltage of approximately 2V to the VCO's input, which forces the VCO to oscillate. After the VCO starts, D_4 becomes back-biased, which disconnects the start-up circuitry from the VCO's input pin. If the VCO is running but hung with its input stuck high or low, one-shot IC_{12A} detects that it's not phase-locked by responding to pulses

THE 74HC4046A PLL MAY NOT START WHEN YOU APPLY POWER.

from Pin 1 of IC_3. The one-shot then issues a 1.5-sec pulse that causes IC_{12B} to produce a 0.5-sec pulse that causes either a positive pulse at the inhibit pin or a low pulse at the VCO's input pin, depending upon whether the PLL is low or high. After the 0.5-sec pulse ends, the pulse from IC_{12A} continues for 1 sec, giving the PLL time to lock. LED D_7 indicates phase lock. If the PLL phase locks, all is well. If it does not, the IC_{12A}/IC_{12B} one-shots continue issuing pulses. Experiments determined these methods for recovering from the anomalous states. It's possible that the circuit won't always recover, but these methods have been effective on the test unit.

The circuit applies the $6F_O$ signal, divided by three, to buffer IC_{5F}'s Pin 5. This action provides an output frequency whose period is proportional to the value of the measured capacitance. The output provides the correct digits without regard to the location of the decimal point. To determine the value of the unknown capacitance, observe the setting of S_1 and S_2.

You can calibrate the circuit by using a capacitance of a known value of approximately 1000 pF, with S_2 at the low-capacitance position and S_1 at the 100- to 1000-pF/0.1- to 1-µF position. Set R_{22} at its midposition, connect a frequency counter to Pin 6 of IC_{5F}, and set the meter to measure the period of the signal. Adjust R_{12} for a period whose digits agree with the known value of capacitance. Next, use a capacitance of approximately 100 pF and set S_1 to the 10- to 100-pF/0.01- to 0.1-µF position. Record the measured value of the capacitor. Then, using the same capacitance of approximately 100 pF, set S_1 to the 100- to 1000-pF/0.1- to 1-µF position and adjust R_{22} to get the same value as you obtained on the 10- to 100-pF/0.01- to 0.1-µF position. The R_{22}/C_{13} combination provides a small variable delay relative to the signal at Pin 14 of IC_3. This fine adjustment improves accuracy in the lower range.

Employing measurements made with the available equipment, which did not include an accurate, high-resolution capacitance meter, this meter is accurate to approximately ±2% over 100 pF to 10 µF (**Table 1**). The accuracy degrades over 10 to 100 pF because of the input capacitance of the op amp and the associated parasitic capacitance at IC_1's Pin 3. R_7 and C_6 provide some compensation at the 10- to 100-pF range for the inherent capacitance at that node. R_5 and C_5 provide compensation at the 1- to 10-pF range.

You can also measure the inherent capacitance and then subtract it from the reading on the two lower ranges. If you take this approach, omit R_5, R_7, C_5, and C_6 from the circuit. Then, with S_1 at the 1- to 10-pF range and S_2 at the low-capacitance position, you can measure the capacitance at that node with no external capacitance. The in-

(*continued on page 50*)

TABLE 1 CAPACITANCE MEASUREMENTS

Range	1 to 10 pF	10 to 100 pF	100 to 1000 pF	1000 to 10,000 pF	0.01 to 0.1 µF	0.1 to 1 µF	1 to 10 µF
Capacitance	5.24, 10.04	10.04, 23.22, 47.6, 102.68	102.68, 469.32, 1022.1	1022.1, 5226.9, 10,140	0.01014, 0.10052	0.10052, 1.034	1.034, 10.07
Measurement error (%)	−8.85, 2.89	6.37, −4.78, −3.68, −0.61	−0.86, −2.5, −0.7	−0.89, −1.28, 0	0.89, 0.88	2.27, −0.87	2.03, 1.24

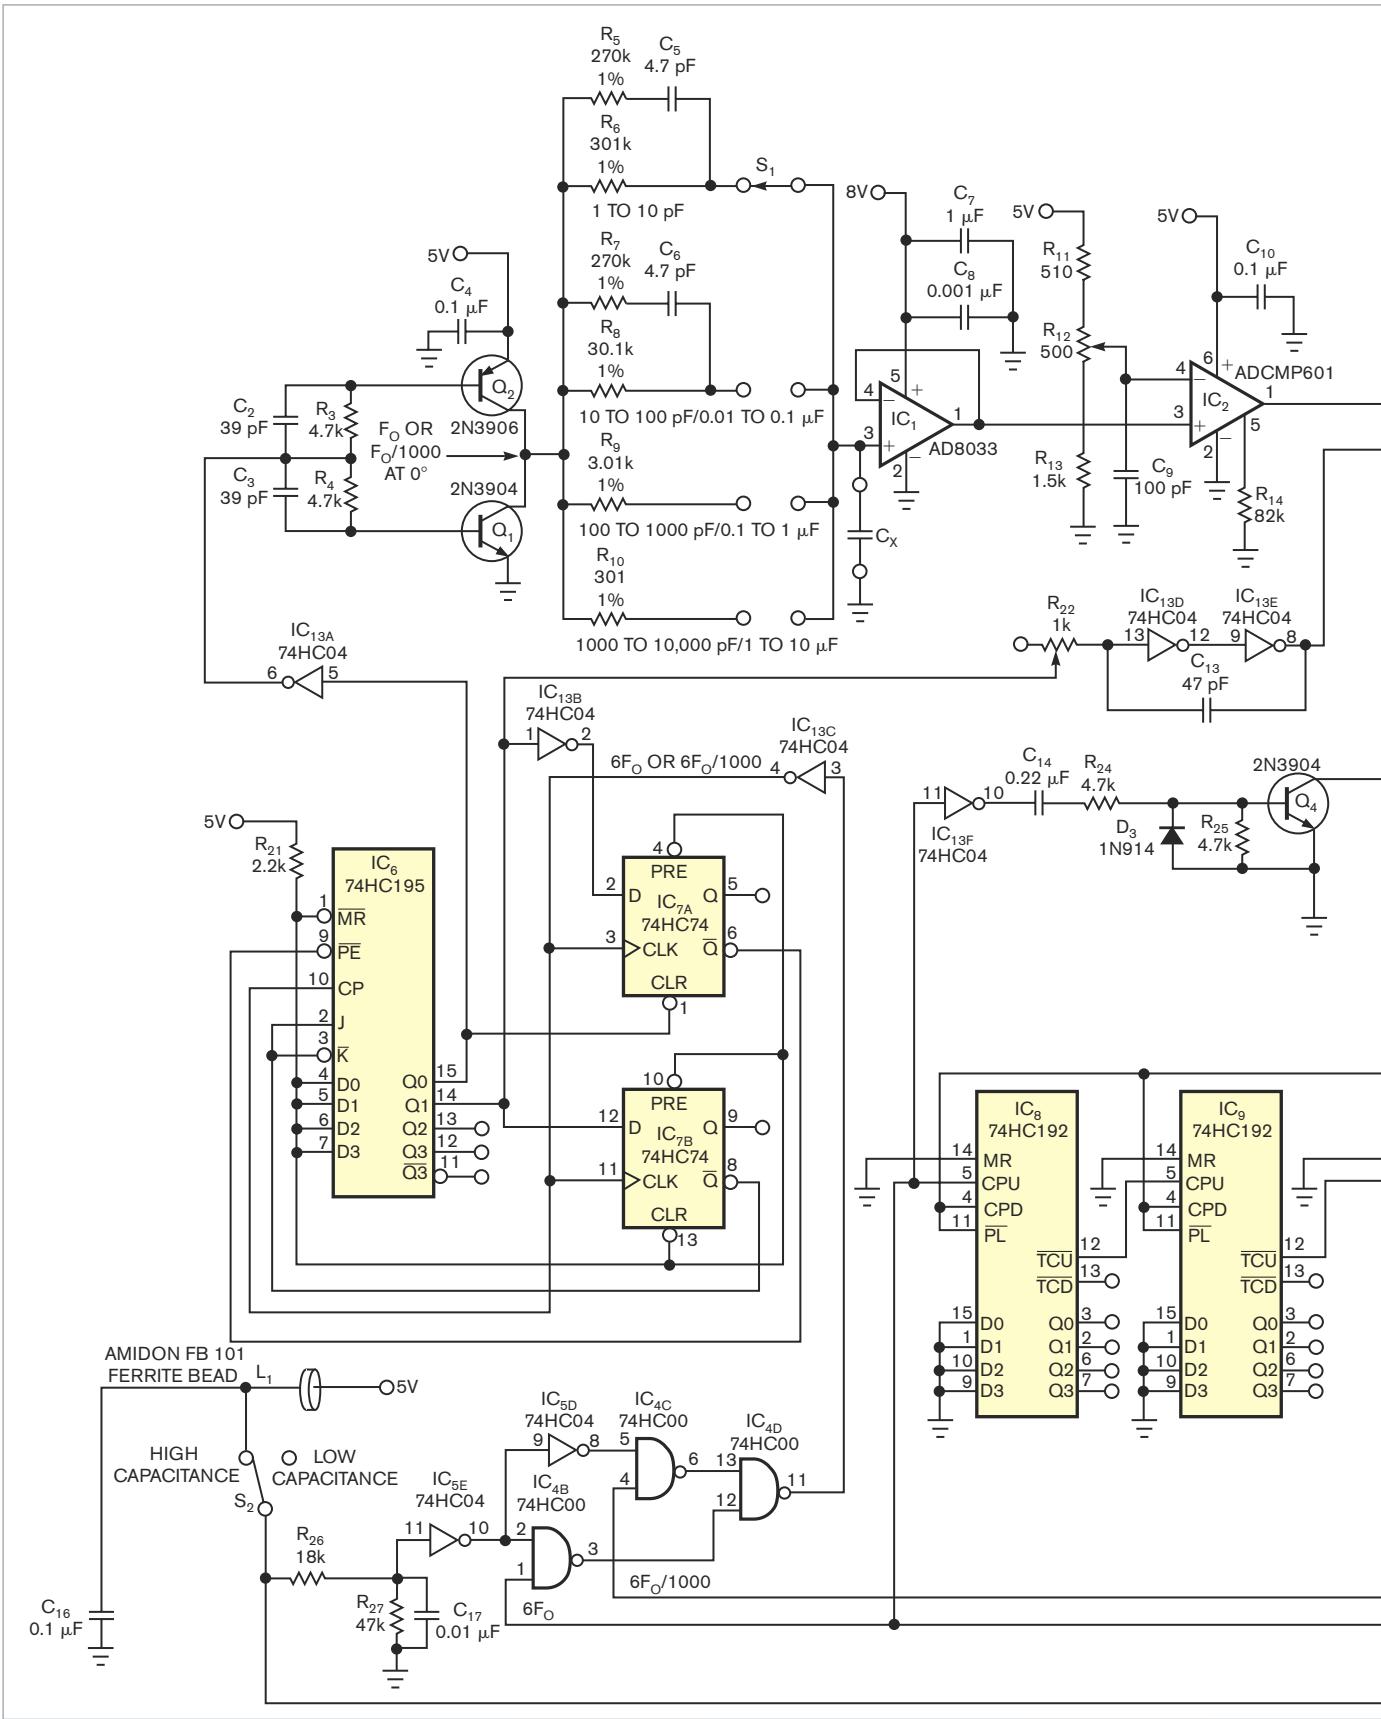

Figure 1 A capacitance meter connects to a frequency counter measuring pulse width to provide a capacitance measurement.

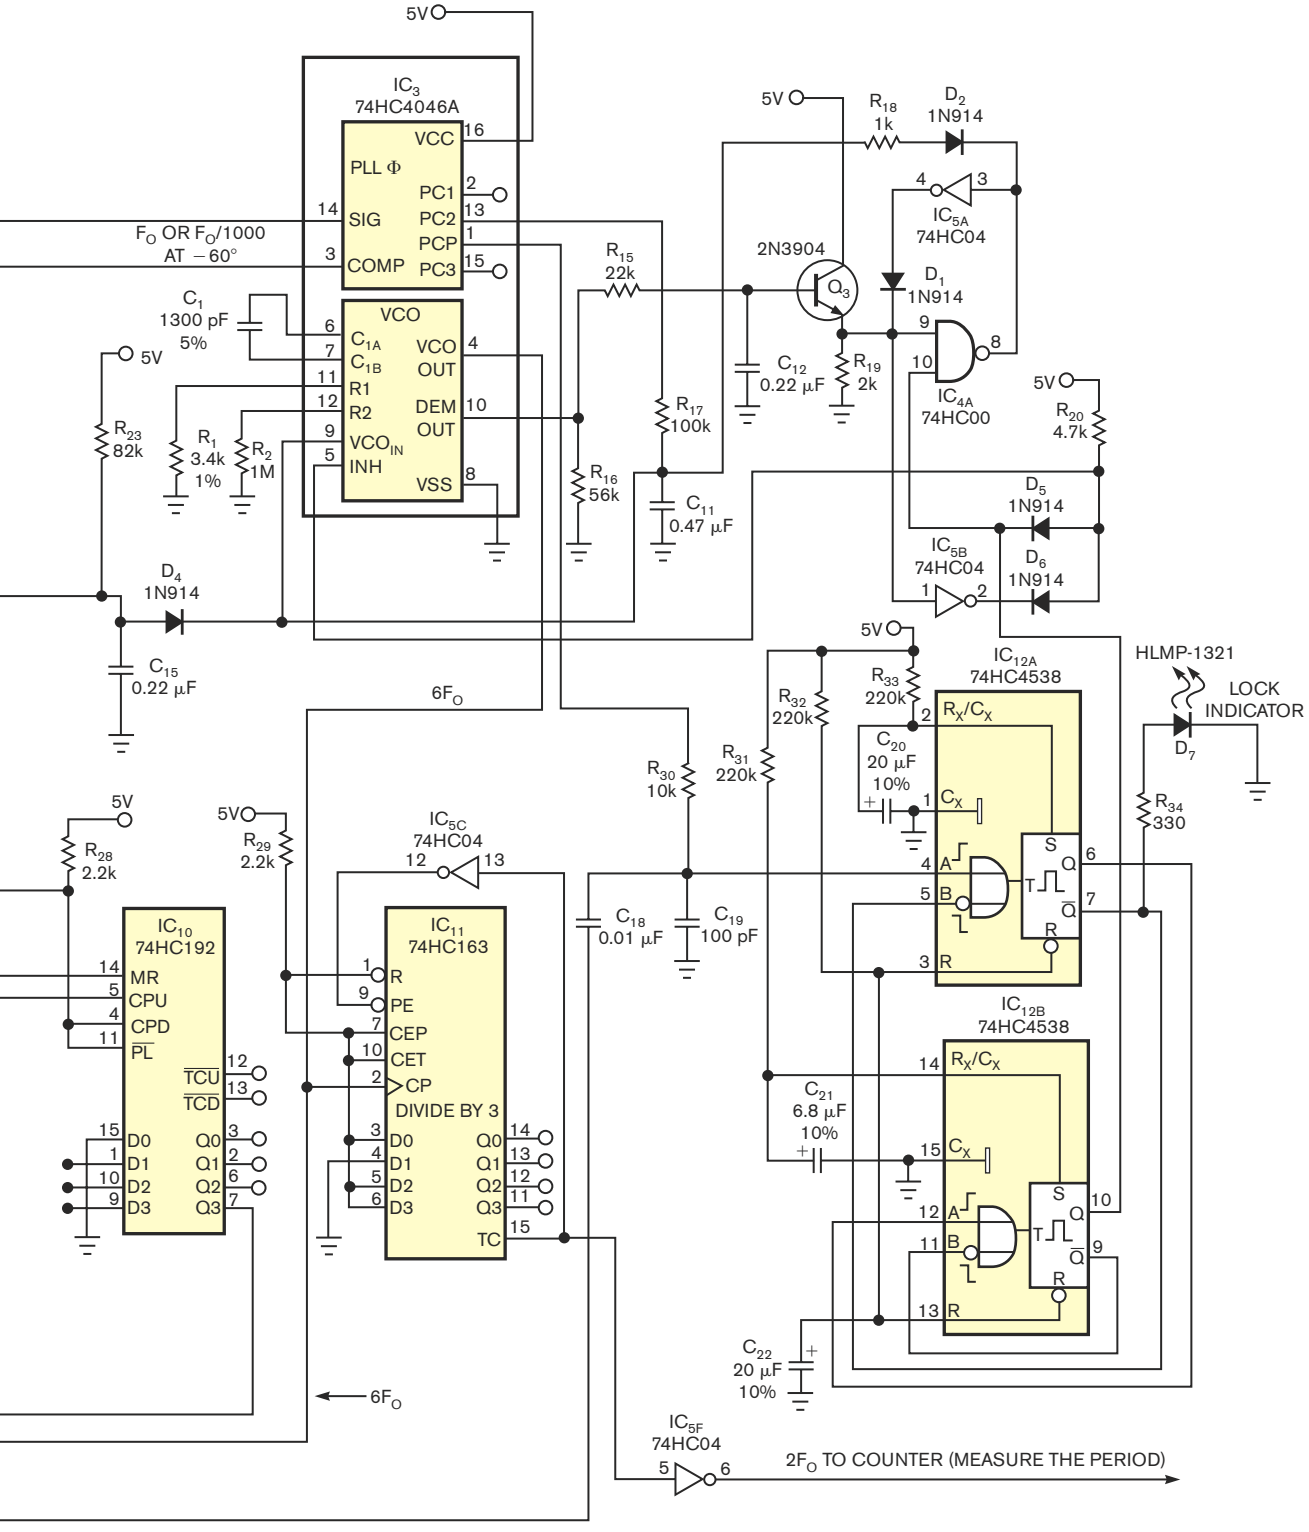
5V
IC3
74HC4046A
VCC
PLL Φ
PC1
PC2
PCP
PC3
SIG
COMP
VCO
VCO OUT
DEM OUT
VCOIN
INH
VSS
FO OR FO/1000
AT −60°
C1
1300 pF
5%
R23
82k
R1
3.4k
1%
R2
1M
R15
22k
R16
56k
R17
100k
C11
0.47 µF
D4
1N914
C15
0.22 µF
6FO
R18
1k
D2
1N914
IC5A
74HC04
2N3904
Q3
D1
1N914
C12
0.22 µF
R19
2k
IC4A
74HC00
R20
4.7k
D5
1N914
IC5B
74HC04
D6
1N914
HLMP-1321
LOCK
INDICATOR
D7
R34
330
IC12A
74HC4538
R33
220k
R32
220k
R31
220k
C20
20 µF
10%
RX/CX
CX
S
Q
T
R
R30
10k
C18
0.01 µF
C19
100 pF
IC12B
74HC4538
C21
6.8 µF
10%
C22
20 µF
10%
R28
2.2k
R29
2.2k
IC10
74HC192
MR
CPU
CPD
PL
TCU
TCD
D0
D1
D2
D3
Q0
Q1
Q2
Q3
IC11
74HC163
PE
CEP
CET
CP
DIVIDE BY 3
TC
IC5C
74HC04
IC5F
74HC04
2FO TO COUNTER (MEASURE THE PERIOD)

(*continued from page 47*)
trinsic capacitance of the test circuit is 2.8 pF. Using this correction, the values you obtain on the lowest two ranges are accurate to approximately ±2%, or ±1 pF.

You must observe capacitor polarity when measuring electrolytic capacitors. Connect the negative end of the capacitor to the grounded terminal. Also, the circuit provides no overvoltage or ESD (electrostatic-discharge) protection, so be sure to discharge the capacitors before connecting them to the capacitance meter and use an ESD wrist strap to avoid damaging the circuit. For best results, you need accurate and stable 5 and 8V power supplies. Both supplies should be accurate to ±2%. You can raise the 8V supply to 9V and relax the accuracy to 5%. If you use a 9V battery to supply the 8V, you can let the voltage drop to about 7.9V before adversely affecting the performance of the meter. You must, however, maintain the 5V supply at a constant, accurate value. Note that all of the ICs except IC_1 have 0.1-µF bypass capacitors from their 5V pins to ground. EDN

REFERENCE

1 Pyle, Ronald E, "Phase-locked loop aids in measuring capacitance," *Electronics Designer's Casebook*, No. 4, pg 32.

Resistor compensates for instrumentation-amp gain drift

Kenneth Gustafsson, Karlskoga, Sweden

Some instrumentation amplifiers use external resistors to set their gain. Unfortunately, the lack of temperature-coefficient matching between the external and the internal resistors results in a high gain drift. If, however, another on-chip resistor is available, you can use it to compensate for gain drift as a result of temperature.

As an example, Analog Devices' (www.analog.com) AD8295 has a drift of as much as −50 ppm/°C, even if you use a zero-drift gain-setting resistor. In this Design Idea, you can compensate this drift with an extra zero-drift resistor in combination with an internal chip resistor.

The gain-set **equation** from the data sheet (**Reference 1**) is

$$\text{GAIN} = 1 + \frac{49{,}400}{R_G}.$$

From this gain-set **equation**, you can assume that the chip uses two 24.7-kΩ resistors with the external gain resistor, R_G, to set the amplifier's gain. The chip has two more 20-kΩ resistors. Because all of these chip resistors are of the same magnitude, they probably will have good temperature-coefficient matching, and you can use this matching for compensation. If the amplifier resistance, R_A, and the gain resistor are zero-drift resistors (**Figure 1**), then

$$\text{GAIN} = \left[1 + \frac{49{,}400\,(1+\Delta)}{R_G}\right]\left[1 + \frac{2 \times R_A}{20{,}000\,(1+\Delta)}\right],$$

where Δ is the drift of the internal matched resistors. If

$$\frac{49{,}400}{R_G} = \frac{R_A}{10{,}000},$$

then the first-order drift of the gain cancels, and the gain splits equally between the instrumentation amplifier and A_1. Solving for R_G and R_A yields

$$R_G = \frac{49{,}400}{\sqrt{\text{GAIN}} - 1};$$

$$R_A = 10{,}000\left(\sqrt{\text{GAIN}} - 1\right).$$

For gain greater than 100, the amplifier resistance becomes greater than 90 kΩ, which can be problematic. In this case, you can use A_1 in an inverting configuration with a gain of −1

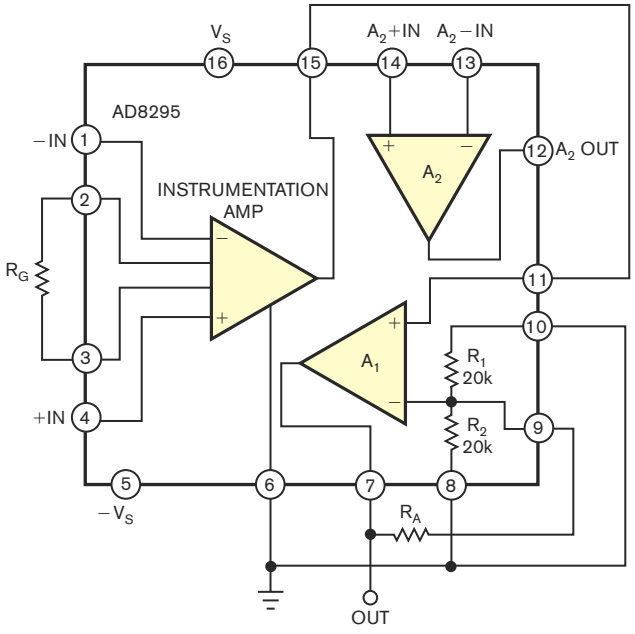

Figure 1 In this configuration, the first-order drift of the gain cancels, and the gain splits equally between the instrumentation amplifier and A_1.

High-voltage integrated pulsers and switches save space and reduce power consumption

Ideal for portable ultrasound equipment applications

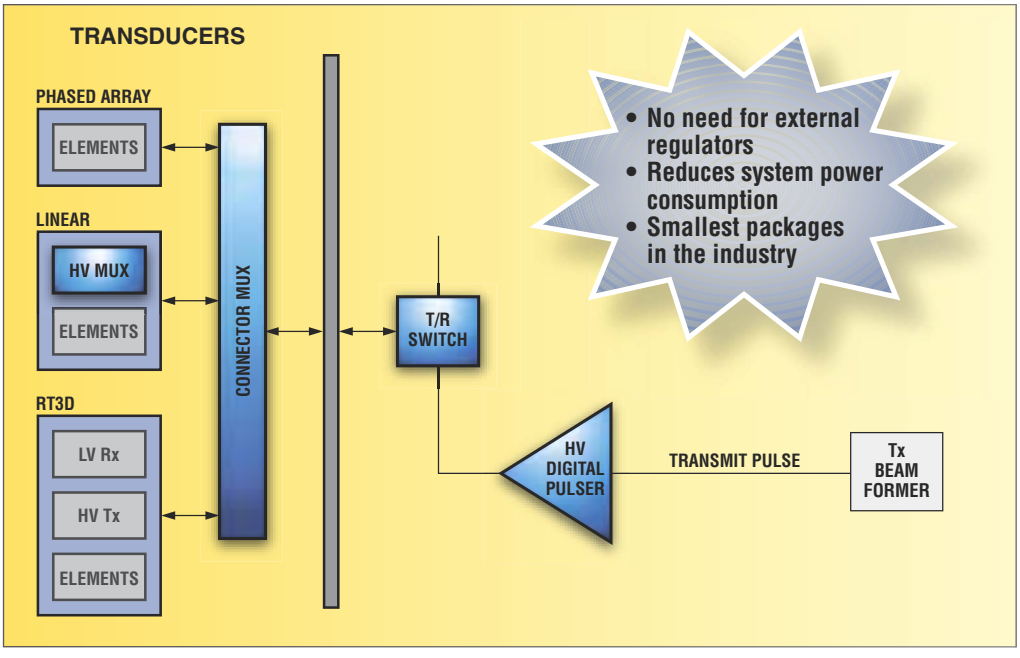

Improvements over the competition

Quad-channel, high-voltage, 2.0A digital pulsers (MAX4940/MAX4940A)

- Integrated active clamps minimize 2nd harmonic distortion and reduce power consumption by 66% vs. HDL6V5581
- Capacitive architecture eliminates need for large, costly floating regulators and complex power sequencing
- AC-coupled architecture reduces overall footprint by up to 45% vs. HDL6V5581
- Five-level pulsing with active damping (MAX4940A) enables transmission for complex systems

16-channel, high-voltage switches (MAX14800*/MAX14801*/MAX14802/MAX14803)

- Ultra-small, 7mm x 7mm CSBGA package saves board space
- Integrated overvoltage protection diodes prevent damage caused by high-voltage overshoots (MAX14802/MAX14803)
- Designed with larger, high-voltage supply range for better linearity
- 33% faster turn-on/off times provide fast frame rates for improved imaging

*Future product—contact factory for availability.

www.maxim-ic.com/Ultrasound-info

MAXIM DIRECT
www.maxim-ic.com/shop

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

MAXIM
INNOVATION DELIVERED™

For free samples or technical support, visit our website.

(**Figure 2**). With an amplifier resistance of 10 kΩ,

$$\text{GAIN} = \left[1 + \frac{49{,}400\,(1+\Delta)}{R_G}\right]$$

$$\left[\frac{2 \times R_A}{20{,}000\,(1+\Delta)}\right] = \left[\frac{1}{(1+\Delta)} + \frac{49{,}400}{R_G}\right].$$

This case sizes R_G using a value from the data-sheet formula. If the gain is 50, the internal matching and the negative drift compensate the "49" part of the gain, and the "one" part is just the drift divided by 50 in the total gain, resulting in a typical figure of −1 ppm/°C. In both cases, the resulting gain temperature coefficient can be less than 5 ppm/°C, which is 10 times better than the original outcome. EDN

REFERENCE

1 "AD8295: Precision Instrumentation Amplifier with Signal Processing Amplifiers," Analog Devices, www.analog.com/en/amplifiers-and-comparators/instrumentation-amplifiers/ad8295/products/product.html.

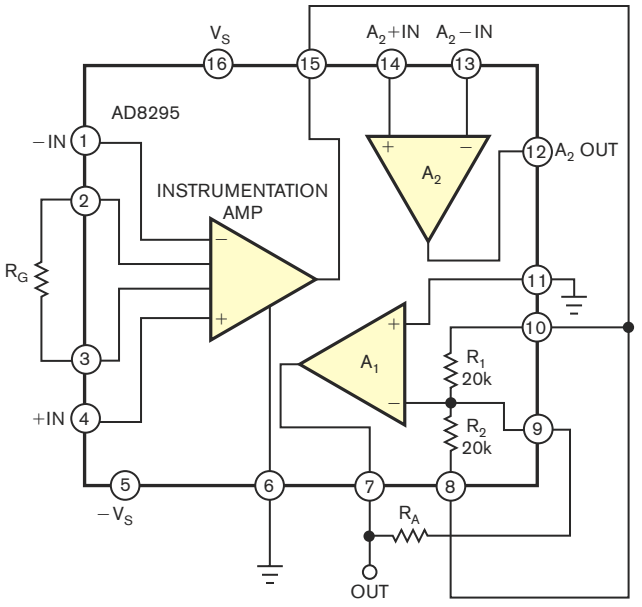

Figure 2 For gain greater than 100, the amplifier resistance becomes greater than 90 kΩ, in which case you can use A_1 in an inverting configuration with a gain of −1.

High-performance, miniature crystal oscillators

Deliver frequencies to 622.08MHz and less than $1ps_{RMS}$ jitter

The DS4-XO crystal oscillators use fundamental AT-cut crystals to enhance frequency stability over temperature. They provide higher stability (±50ppm) than SAW-based oscillators (±100ppm) over the extended temperature range, and they eliminate the spurious modes of operation associated with overtone designs. Their miniature ceramic package allows precise placement in designs, thus making them ideal for Fibre Channel, InfiniBand™, and other networking applications.

Features

- **Low jitter: < $1ps_{RMS}$ from 12kHz to 20MHz**
- **Wide frequency range: 77.76MHz to 622.08MHz**
- **Small size: 5mm x 3.2mm x 1.49mm**
- **High frequency stability: ±50ppm (-40°C to +85°C)**
- **Low voltage: 3.3V ±5%**

Applications

- **Fibre Channel:** DS4106, DS4212, DS4425
- **InfiniBand:** DS4125
- **Ethernet, 10GbE:** DS4125, DS4156, DS4312
- **SAS/SATA:** DS4150
- **SONET/SDH:** DS4155, DS4311, DS4622, DS4776
- **PCI Express®:** DS4100H

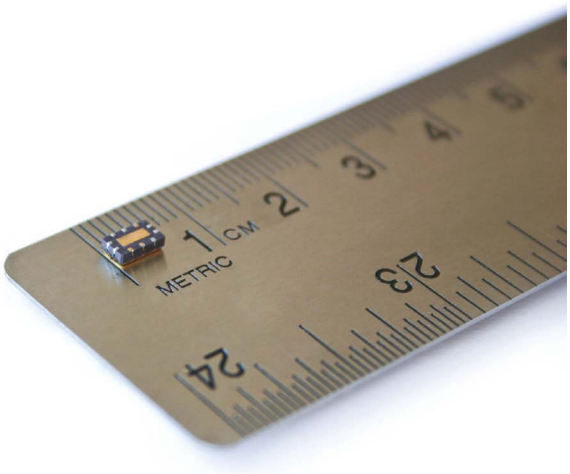

Housed in a tiny, 5mm x 3.2mm x 1.49mm LCCC ceramic package, the DS4-XO crystal oscillators offer the industry's highest performance and save critical board space.

www.maxim-ic.com/DS4-XO

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.

EDITED BY SUZANNE DEFFREE

supplychain

LINKING DESIGN AND RESOURCES

With spotlight on Asia/Pacific, don't forget Americas design

The components industry's obvious growth may be in the Asia/Pacific region, but distributors are not dismissing more modest areas of growth, including the Americas. "We consider the Americas to be an extremely viable market on a number of fronts," says Alex Iuorio (**photo**), senior vice president of supplier management for Avnet Electronics Marketing Americas (Avnet EM, www.em.avnet.com). "We … believe there is a customer base here that will figure prominently in … innovative design."

Iuorio points to the defense/aerospace, medical, security, lighting, and transportation segments as strong opportunities for the Americas. "There are certain members of these segments where the business [and] the products lend themselves more to not only an Americas design but also to potential manufacturing, specifically defense/aero and medical," he says. "We made a bet to continue to support the design efforts of the overall customer base … but specifically these verticals because that's where we think the growth is going to be. When it starts to happen … we are positioned to take advantage of it."

As evidence of Iuorio's comments, Avnet EM recently inked a deal to distribute PUI's (Projects Unlimited Inc's) line of audio-technology products in the Americas. With a military background, PUI makes indicators, speakers, and microphones that find use in medical glucose meters and infusion pumps, industrial scissor lifts and construction equipment, and consumer electronics, among other areas.

"Asia is very important," says Iuorio, but he emphasizes that a company should also focus on other regions of the world "for their growth potentials and their potentials to contribute to the overall global value chain. We're not biased about growth," he says. "We'll take it exactly where we can find it, where we can anticipate it, and where we can resource for it. It doesn't matter if that [growth is] in Asia or the West; we'll find ways to contribute."

GREEN UPDATE

EUROPEAN UNION, OTHER ORGANIZATIONS PROPOSE 15 "SUBSTANCES OF VERY HIGH CONCERN"

The ECHA (European Chemicals Agency, http://echa.europa.eu) has published proposals to identify SVHCs (substances of very high concern) under the REACH (Registration/Evaluation/Authorization of Chemicals) regulation and is welcoming comments on these 15 substances by Oct 15. The agency proposed these substances as SVHCs because they are considered carcinogenic, mutagenic, or reprotoxic, or because they are considered persistent, bioaccumulative, and toxic. For a list of the substances or to comment, go to http://echa.europa.eu/consultations/authorisation/svhc/svhc_cons_en.asp.

The European Union/European Economic Area member states and the European Commission proposed the substances to the ECHA, which will take comments into account when deciding whether the substances will appear on the candidate list from which lawmakers select substances for required authorization of use. Substances appearing on this list may eventually be subject to authorization in the European Union.

The ECHA has advised that comments should focus on the hazardous properties that qualify the chemicals as SVHCs. Parties can also provide comments and further information on the uses, exposures, and availability of safer alternative substances or techniques, although the EHCA will consider these aspects mainly at the next stage of the process.

OUTLOOK

CPE MARKET SEES SOME BRIGHT SPOTS

Although the economic downturn forced 2008 to have the slowest total global broadband CPE (customer-premises-equipment)-market growth this decade of 2.4% at 154 million units and will force 2009 to show even slower growth, researchers at In-Stat (www.instat.com) report that there are some bright spots to the segment.

According to the company, cable gateways, wireless and VOIP (voice-over-Internet Protocol) routers, and wireless and VOIP DSL (digital-subscriber-line)-CPE units continued to grow in at least the double digits. In addition, several segments of the broadband-CPE market, including cable gateways, FTTH (fiber-to-the-home) gateways, and FWB (fixed-wireless broadband), should grow considerably faster than the overall market over five years, In-Stat reports.

"Gigabit Ethernet, VOIP, the DSL Forum's TR-69, and 802.11n are all drivers for CPE upgrades and replacements," says Joyce Putscher, In-Stat analyst. "We're also seeing accelerated growth in FTTH-CPE unit shipments in 2009." According to In-Stat, VOIP-enabled DSL-CPE unit shipments saw healthy growth in 2008, and more than half of DSL-CPE unit shipments in 2009 will feature VOIP capabilities.

Accelerating the **Speed of Design.**™

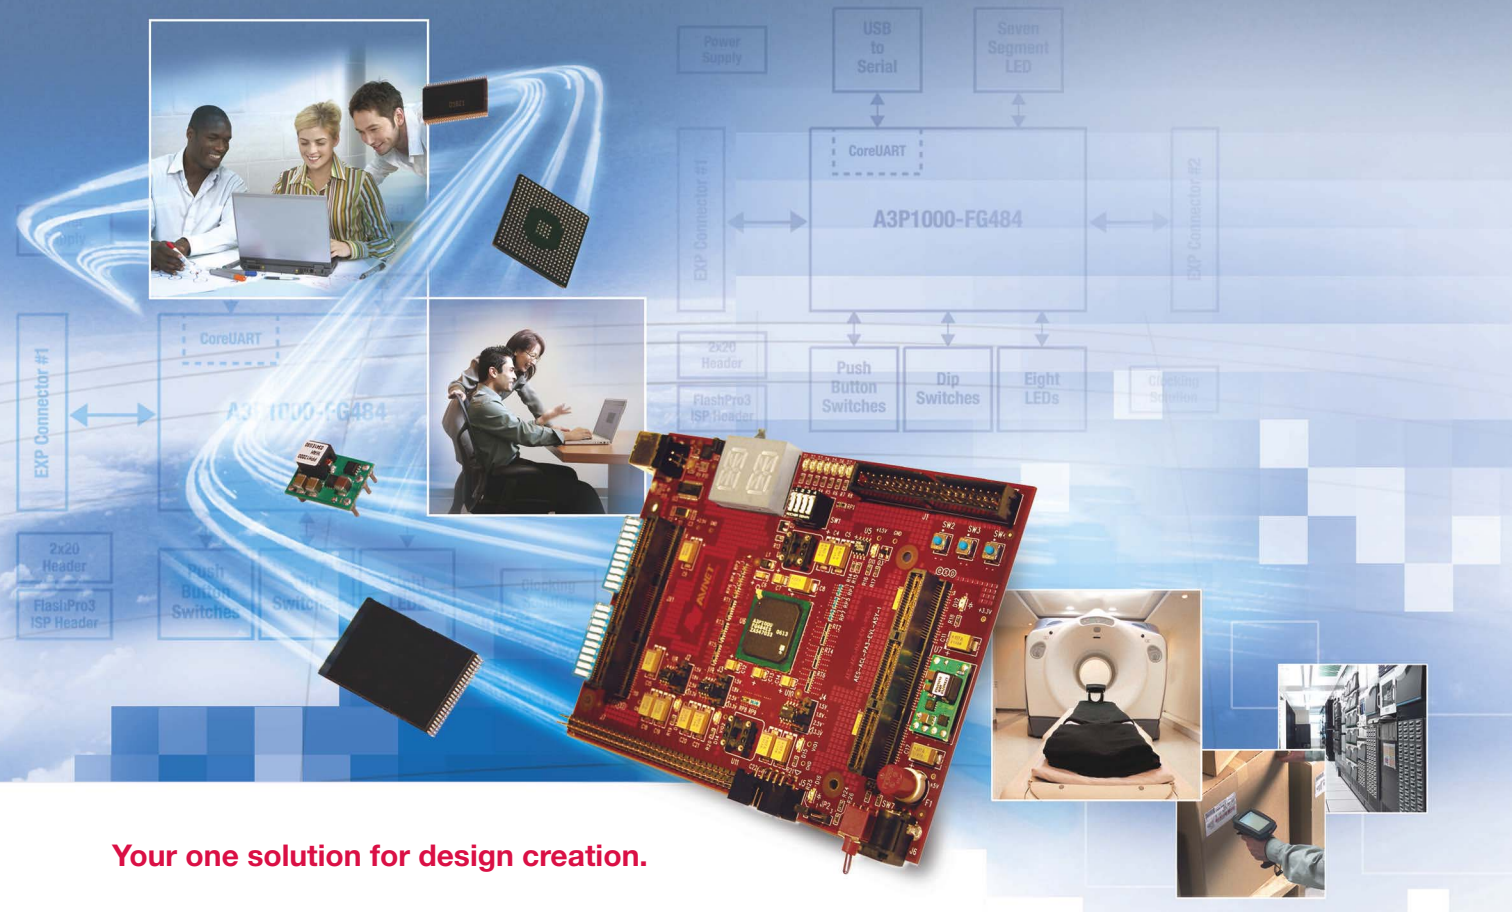

Your one solution for design creation.

Avnet Memec focuses on a select group of leading suppliers to provide original equipment manufacturers (OEMs) a high level of technical expertise — enabling the swift and accurate design of leading-edge products.

With factory trained and certified support, Avnet Memec specialists offer a level of technical depth unmatched in the industry. With easy access to our industry experts, your projects move forward with greater velocity, driven by in-depth technical expertise. That translates into faster time to market for you, with greater ease and confidence. From Design to Delivery.™ **Avnet is ranked Best-in-Class* for online seminars, design & engineering services and knowledgeable application engineers – proof that we consistently deliver:**

> Proactive engineering support
> Factory certified FAEs

Accelerate Time to Market.™
Visit the Avnet Design Resource Center™ at: www.em.avnet.com/drc

SAME-DAY SHIPPING
orders placed before 8:00P.M. central time

AVNET® Memec

Accelerating Your Success™

1 800 332 8638
www.em.avnet.com

Avnet Green Initiative

NEDA Supplier Authorized Distributor

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

productroundup

SENSORS AND TRANSDUCERS

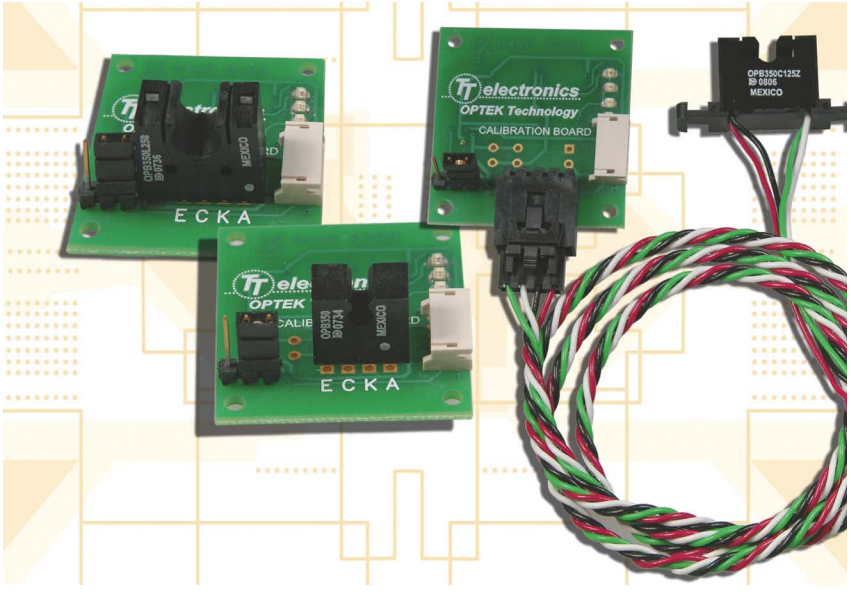

Sensor detects presence of fluid in transparent tubing

The OCB350 fluid-sensor series with self-calibration circuitry uses an LED/phototransistor for sensing change in refractive properties between full and empty tubing as it passes through the sensor housing. The housing has an opening accommodating clear tubing with 1.6-, 3.2-, and 4.8-mm outer diameters. Fluid passing through the sensor housing creates multiple output states, such as fluid present, no fluid present, and no tube present. The devices reduce the effects of changes in the sensor's optics from manufacturing variances, ambient light, temperature fluctuation, and degradation of the LED and phototransistor. An LED and phototransistor in an opaque-plastic housing enhance ambient-light rejection; clear liquid causes phototransistors to sink the maximum current, and dark liquid causes phototransistors to sink the minimum current. As bubbles pass through the tube, the signal varies between the liquid-present and no-liquid states. If no tube is present, the phototransistor sinks the current between the dark- and clear-fluid states. The fluid sensor requires an optical device and a power supply for operation. The OCB350 fluid sensor costs $13.15 (1000).

Optek Technology, www.optekinc.com

Rotary-position sensor suits 3.5- and 4-mm-diameter D shafts

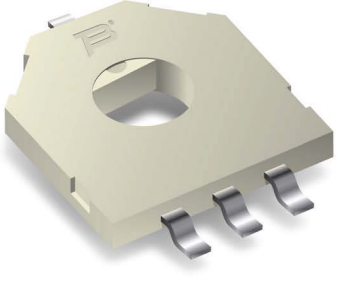

The Model 3382 rotary-position sensor includes an H-pin through-hole option and an increased standard-rotor slot accommodating 4- and 3.5-mm-diameter D shafts. The sensor provides 2.5- to 100-kΩ resistance ranges. A rotary-position sensor enables the control of variable outputs, including frequency, speed, contrast, brightness, and volume. Available with a 2.1-mm profile, the Model 3382 rotary-position sensor costs $1.12 to $1.25 (1000).

Bourns Inc, www.bourns.com

Dual-axis gyroscope improves optical image stabilization

Targeting camera phones and digital still cameras, the digital IDG-2000 dual-axis gyroscope has an integrated temperature sensor that provides temperature-compensated bias adjustment. The device is smaller and thinner than competing gyroscopes, enabling optical image stabilization in digital still cameras and high-resolution, 8M-pixel photos in camera phones. The digital dual-axis, pitch-and-roll gyroscope integrates a 16-bit ADC, allowing direct telecommunication for advanced image processors over an I^2C or SPI bus, requiring no discrete components. The device measures hand jitter over a 0.1- to 20-Hz-wide frequency range and has ±1% cross-axis sensitivity. Available in a 4×4×0.9-mm QFN plastic package, the IDG-2000 costs $1 per axis.

InvenSense Inc, www.invensense.com

COMPUTERS AND PERIPHERALS

256-bit hardware encryption enables daisy-chaining for increased security

The CipherChain AES 256-bit hardware full-disk encryption system comes in a module the size of a Compact Flash. Requiring no drivers, the device works with SATA I/II hard

COMPUTERS AND PERIPHERALS

drives, solid-state drives, and storage devices with SATA interfaces. Aiming at PCs, servers, rack-mounted systems, and data-storage equipment, the device encrypts the boot sector, partition table, and information in the drive and allows daisy-chaining for increased security. The CCM35MK1 model includes the CipherChain, a pair of Cipher keys, and a mounting bracket for a 3.5-in. drive bay; the price is $79.

Addonics Technologies, www.addonics.com

RAID controllers provide zero-maintenance cache protection

Joining the vendor's Unified Serial SATA/SAS family, the 5Z RAID-controller series provides flash-based zero-maintenance cache protection. The family includes the RAID 5445Z with four internal ports and four external ports, the RAID 5805Z with eight internal ports, and the RAID 5405Z with four internal ports. Features include a dual-core RAID-On-Chip, a 533-MHz DDR2 write cache, and an eight-lane PCIe host-interface-bus connection. The RAID 5405Z, the RAID 5805Z, and the RAID 5445Z Unified Serial SATA/SAS devices cost $785, $965, and $1045, respectively.

Adaptec, www.adaptec.com

Dual displays fold closed

The InterView system combines two rotating, superthin, high-resolution, 17-in. TFT LCDs on a desktop stand. The stand allows 180° rotation of screens on a horizontal axis, and the screens fold 90° from completely closed to a full width apart. The base features three USB ports and controls for microphones, power, and DMS connections. The InterView dual-display costs $649.99.

EVGA Corp, www.evga.com

Portable hard drive includes FireWire and USB 2.0 capabilities

The SureFire FW800/USB2 portable hard-disk drive includes a USB 2.0 port and a FireWire 800 port and requires no ac-power adapter. Measuring 81×128×15 mm, the 0.41-lb device features a 5400-rpm spindle and an 8-Mbyte cache memory. The device comes in 250-, 320-, and 500-Gbyte options and costs $109.99, $139.99, and $179.99, respectively.

Verbatim Americas LLC, www.verbatim.com

EDN ADVERTISER INDEX

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

JB GUIOT • DIGILEA SA

Keep it simple, stupid, and kiss problems goodbye

DANIEL VASCONCELLOS

Fresh out of college and after attending military school, I joined an international company that developed, built, and sold CNC (computerized-numerical-control) machine tools. It was a small company employing approximately 200 people in three countries. Our location was responsible for the integration of the CNC and all the electronics and electrical equipment into the machine tools. We also had a repair-and-support department for customers.

After two months of training on the CNC and the different machines, I first landed in the repair department. It was an interesting time to discover all the insides of such a control. This CNC used only good old TTL (transistor-transistor logic), the 74xx series, with a 74181 ALU (arithmetic-logic unit) at its core. No microprocessor! The program counter, ALU, interrupt controller, instruction decoder, TTY (Teletype) interface, memory control, numerically controlled oscillator, and axis controller were all discrete single boards in a 19-in. rack. I could follow each instruction clock pulse by clock pulse into the decoder, the ALU, and so on. It is rare today for an electronics engineer to get such an opportunity to understand in detail the processors' internals.

The programming of this control was a marvel for me, too. Although most people nowadays have a hard time with assemblers, I would have loved to have an assembler for this machine; the language was machine code! I could follow that the Nth bit of the code opened an AND gate, which made the ALU shift left instead of right. I still sometimes look at the machine codes of modern microcontrollers to see the pattern.

One day, we received a new spindle driver to fit into a series of machines. Within a few days, everything worked well except that the voltmeters for speed and torque were operating at 15V but the driver output operated at only 10V full-scale. We could not replace the indicators within the short time we had. Electronically speaking, though, it was easy stuff! We needed no high precision or stability, and we had ±15V supplies in the electrical-equipment cabinet, so the solution was obvious. I took some pieces of a PCB (printed-circuit board)—a simple 741 op amp, four resistors, and two capacitors—and, in no time, I had a 1.5× amplifier. I quickly checked it in the lab. It worked as I had expected.

I was working at the machine table, hooking up the parts of my new design, when the boss came by. Interested, he asked me what I was doing and to explain the details. When I was finished with my explanation, he opened the back of the voltmeter, cut out the 1.5-kΩ resistor inside, replaced it with a 1-kΩ resistor, and closed the indicator. "Voilà!" he said. His simple replacement of just one resistor by another outsmarted my naive attempt to use a complex—though theoretically correct—circuit!

That day I learned what may be the most important lesson in a developer's life: Always ask yourself whether you might be able to find a simpler approach. I don't know who first said "keep it simple, stupid," but ignoring this advice is the "KISS" of death! EDN

JB Guiot is an engineer living in Mulhouse, France, but working for more than 25 years in Switzerland.

Go to www.edn.com/091008tales and click on Feedback Loop to post a comment on this article.

The Newest
Renewable Energy Harvesting

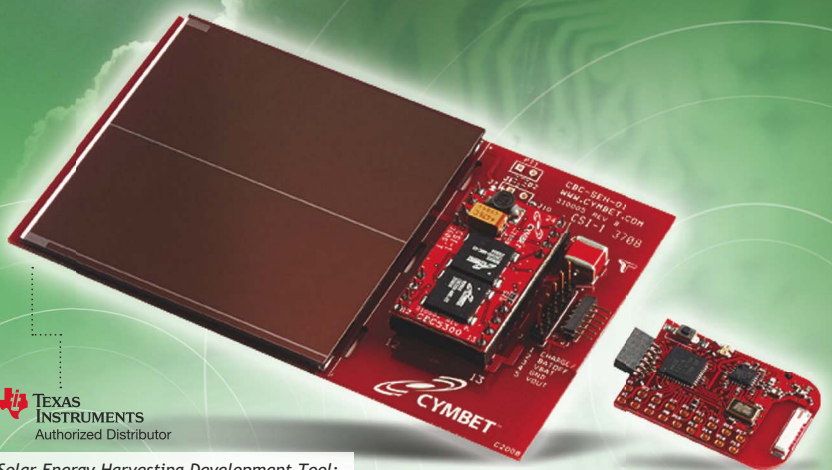

TEXAS INSTRUMENTS
Authorized Distributor

Solar Energy Harvesting Development Tool:
eZ430-RF2500-SEH
www.mouser.com/ti_ez430_rf2500_seh/

www.mouser.com
Over A Million Products Online

New Products from:

ADAPTIVENERGY
Random Vibration Joule-Thief™ Module
www.mouser.com/adaptivenergy_joule-thief

CYMBET CORPORATION
EnerChip™ EH CBC5300 Module
www.mouser.com/cymbet_cbc5300

ADVANCED LINEAR DEVICES, INC.
Energy Harvesting Modules
www.mouser.com/ald/a

The ONLY New Catalog Every 90 Days

Experience Mouser's time-to-market advantage with no minimums and same-day shipping of the newest products from more than 390 leading suppliers.

MOUSER ELECTRONICS®
a tti company

The Newest Products For Your Newest Designs

(800) 346-6873

Are you up for a challenge?

What is the missing component?

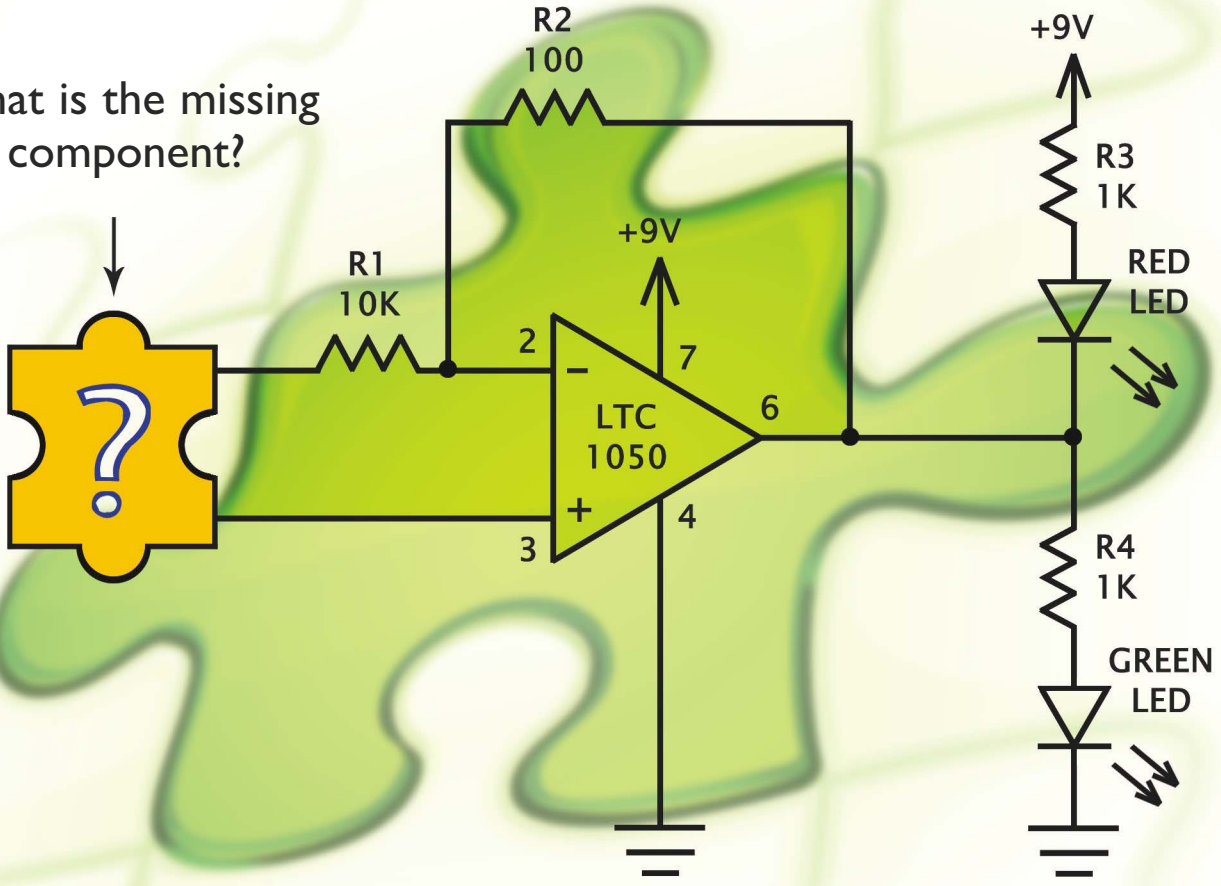

Industry guru Forrest M. Mims III has created a stumper. Video game designer Bob Wheels needed an inexpensive, counter-clockwise rotation detector for a radio-controlled car that could withstand the busy hands of a teenaged game player and endure lots of punishment. Can you figure out what's missing? Go to www.Jameco.com/unfold to see if you are correct and while you are there, sign-up for our free full color catalog.

JAMECO®
ELECTRONICS
1-800-831-4242