Reed Business Information (RBI™)

EDN®

VOICE OF THE ENGINEER

MAY 28
Issue 10/2009
www.edn.com

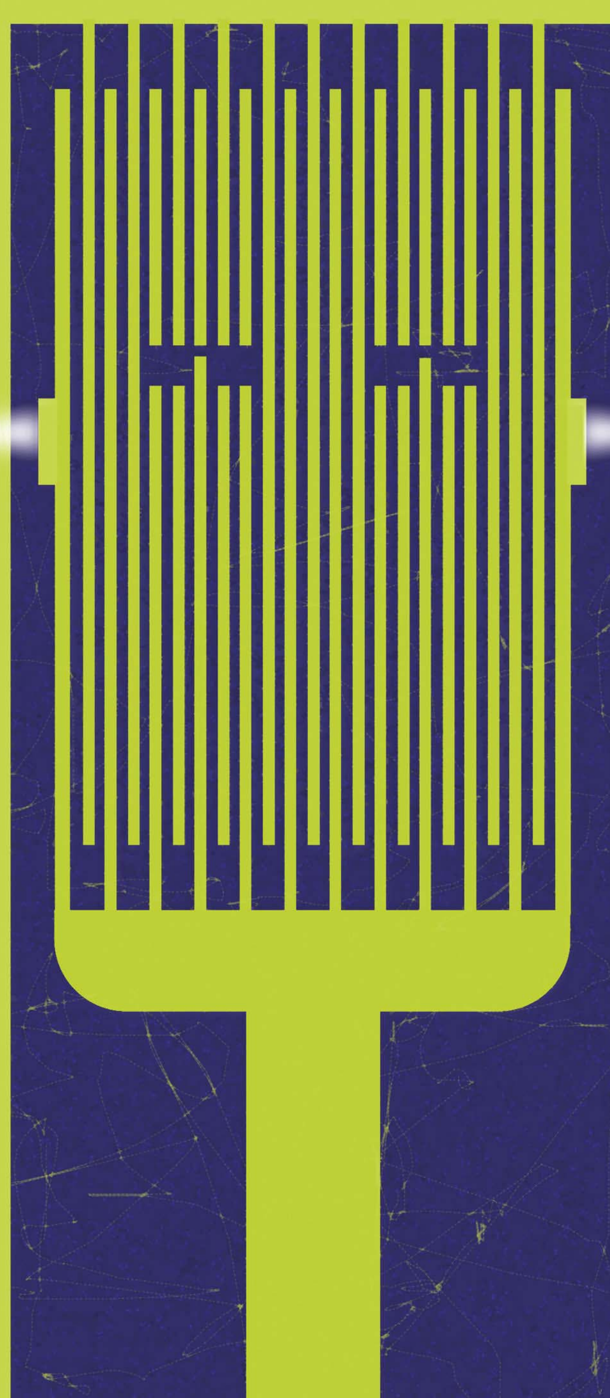

STRESS OUT; STRAIN GAUGES IN

Page 28

LEAKAGE-POWER ANALYSIS ENABLES ATTACKS ON CRYPTOGRAPHIC DEVICES

Page 23

EVOLVING TO DDR3 TECHNOLOGY

Page 37


PICO
PROVEN Critical Component Integrity
DC-DC Converters
2V to 10,000 VDC Output
2 to 5000 VDC Outputs
Ultra Miniature
Surface Mount and Thru Hole
Single and Dual Isolated Outputs
Military Upgrades Available
AV/AV/SM/AVR Series
Military Components
-55⁰ to +85⁰C Operating Temp
Military Environmental
Screening Available
3.3 to 500 VDC Outputs
M/MV Series
Isolated - Regulated
100 to 10,000 VDC Output
Proportional Control Up to 10 Watts
VV Series
Military Applications
-40⁰ and -55⁰ to +85⁰C
Operating Temperatures
Wide Input Range
3.3 to 350 VDC outputs
in 1/2 Brick and Ful Brick
Military Upgrades Available
LF/LM/FM Series
Programmable to
6000 VDC Output
HVP Series
HIGH POWER-Wide Input Range
2 to 350 VDC Output
Isolated to 300 Watts
in 1/2 and Full Brick
LP/HP/XP Series
Wide Input Range
8 to 60 VDC Input
2 to 100 VDC Output
Isolated-Regulated
OR/IR/JR/KR Series
2-20 Watts
36 to 170 VDC Inputs
Terminal Strips - Thru Hole
3.3 to 48 VDC Outputs
Single and Dual Output
LV/HV Series
Also Ac-DC single and 3 Phase.
Power Factor Corrected. 3 Watts to
2000 Watts Models.
5 to 500 VDC Output
to 50 Watts
Wide Input Range
Isolated
Regulated
QP Series
PICO
Electronics, Inc.
143 Sparks Ave., Pelham, NY 10803
www.picoelectronics.com
Complete Listing of Entire Product Line
DC-DC Converters • AC-DC Inductors • Transformers
and Inductors.
E Mail: info@picoelectronics.com
Call Toll Free:
800-431-1064
Fax:
914-738-8225
ISO 9001:2000
ANAB
BSI
AS 9100 REV. B


1-800-DIGI-KEY
ONE SOURCE
OVER 400,000 PRODUCTS IN STOCK FROM MORE THAN 400 SUPPLIER PARTNERS*
Digi-Key®
CORPORATION

Green Engineering

MEASURE IT – FIX IT

La Selva Biological Station developed a wireless sensor system to monitor the rain forest ecosystem.

Nucor Steel optimized its steel melting process to drastically reduce electricity consumption.

CEMS Engineering built a control system to reduce energy use of industrial air chillers by 30 percent.

Vehicle Projects created a complex control system for a zero-emission, fuel-cell-powered locomotive.

MEASURE IT

Acquire — Acquire and measure data from any sensor or signal

Analyze — Analyze and extract information with signal processing

Present — Present data with HMIs, Web interfaces, and reports

FIX IT

Design — Design optimized control algorithms and systems

Prototype — Prototype designs on ready-to-run hardware

Deploy — Deploy to the hardware platform you choose

Engineers and scientists around the world are solving some of today's most pressing environmental issues using the NI graphical system design platform to design, develop, and deploy more efficient and environmentally friendly products, technologies, and processes. Using modular hardware and flexible software, they are not only testing and measuring existing systems but also creating innovative ways to fix the problems they find.

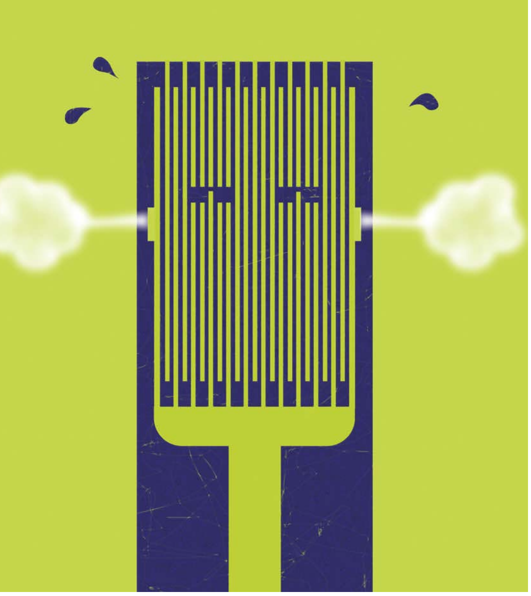

EDN 5.28.09 contents

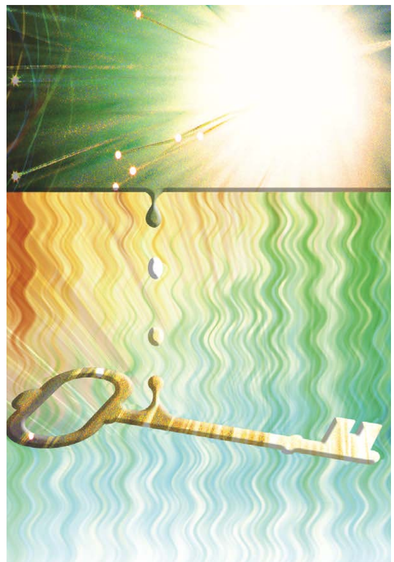

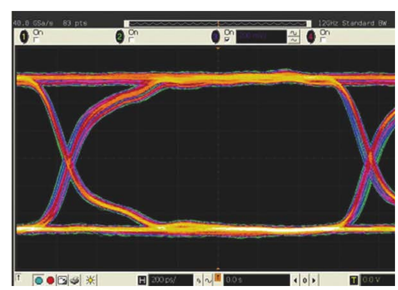

pulse

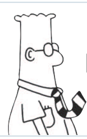

DESIGNIDEAS

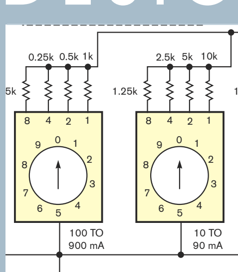

contents 5.28.09

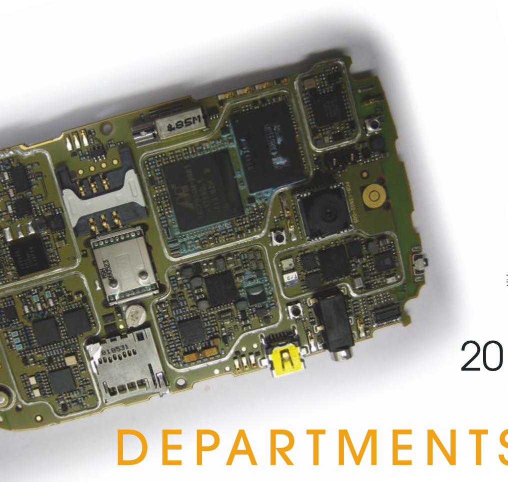

20

58

DEPARTMENTS & COLUMNS

EDN online contents www.edn.com

ONLINE ONLY

Check out these Web-exclusive articles:

A universal algorithm for implementing an infrared decoder
A microcontroller-based remote-control design provides the flexibility and ease of upgrade necessary to support a diversity of modern infrared protocols.
→www.edn.com/article/CA6653644

Innovating out of the downturn
GUEST OPINION: If history is any indicator, innovation, not the business cycle, will pull the semiconductor industry out of recession. But, this time, young, innovative companies jumping on advanced processes may play a disproportionate role.
→www.edn.com/article/CA6653667

One size fits all
GUEST OPINION: IEEE 802.11 Wi-Fi is a technology that has taken a decade to mature and fit into a market segment.
→www.edn.com/article/CA6656616

CHILLING TALES OF ENGINEERING TERROR

OK, they're not really that scary, but our Tales from the Cube articles, which appear on the back page of every issue, do relate hair-raising incidents as told by real-life engineers. Check out these recent Tales, as well as the entire archive of every heart-stopping episode—if you dare.

Conspiracy theory
→www.edn.com/article/CA6643387

A lesson in humidity? No sweat
→www.edn.com/article/CA6639818

Noisy radio link puts brakes on racing-crew communication
→www.edn.com/article/CA6629491

Complete Tales archive
→www.edn.com/tales

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

BY BRIAN DIPERT, SENIOR TECHNICAL EDITOR

802.11n: Complicated and about to become even messier

In January 2007, I bemoaned the blizzard of "draft" 802.11n gear at that year's Consumer Electronics Show, whose presence constrained the continued evolution of the then-under-development next-generation Wi-Fi standard. More than two years later, the IEEE is still slaving away on the specification, although tentative indications suggest that the approval work—albeit perhaps not the pending lawsuits—may wrap up by year-end. In a mid-2008 interview with Stephen Palm, the technical director of Broadcom's broadband-communication group, I grumbled about the fact that the draft standard allowed for access points and access-point-containing routers to optionally support either 2.4- or 5.8-GHz—but not necessarily both—ISM (industrial/scientific/medical) bands, pointing out the market confusion that might result but acknowledging the cost-driven motivations for the decision.

Unfortunately, the situation is about to get even more muddled. First, keep in mind that LAN clients also do not need to support both the 2.4- and 5.8-GHz bands. Even more bewildering is the relationship or—perhaps more accurately—lack of relationship between 802.11n-cognizant equipment's MIMO (multiple-input/multiple-output)-antenna arrays and the number of simultaneous-stream transfers they support along with the dearth of product documentation regarding these specifications. The minimum number of antennas on both sides of the link limits the number of possible simultaneous data streams. However, the radios often further limit the number of spatial streams that can carry unique data.

To help identify a radio's performance, the IEEE uses the a×b:c notation, where a is the maximum number of transmitting antennas or RF chains the radio can use, b is the maximum number of receiver antennas or RF chains the radio can use, and c is the maximum number of data-spatial streams the radio can use. For example, a 2×3:2 radio can transmit on two antennas and receive on three but can send or receive only two data streams. The 802.11n specification working draft allows for configurations that can transmit and receive on as many as four antennas and can send or receive as many as four data streams. Common retail configurations of 11n devices are 2×2:2, 2×3:2, and 3×3:2. All three configurations have the same maximum throughputs and features and differ only in the amount of diversity the antenna systems provide.

The current 802.11n specification working draft, according to knowledgeable sources, requires that access points and routers containing them deliver 2×2:2 or better transfer-rate capabilities, although it allows LAN clients to drop down to 1×1:1 levels. Its aspirations differ, however, from the certification requirements of the Wi-Fi Alliance, an organization that ramped up its activities in the absence of formal specification approval, as a means of jump-starting the market transition to 802.11n. Like the IEEE, the Wi-Fi Alliance requires that 802.11n access points and routers be 2×2:2 or better. However, with respect to LAN clients, the Wi-Fi Alliance differentiates between PCs, which must receive at least two simultaneous streams but must minimally support only single-stream transmissions, and mobile devices, for which single-stream support for both basic transmitting and basic receiving operations is acceptable.

I acknowledge the cost-versus-performance background of this differentiation, but the split between the two categories is imprecise. Consider the laptop, especially the cost-sensitive netbook. Is it a PC or a mobile device? And what about the "low-power 802.11n" mode supposedly coming in the next-generation iPhone and iPod touch? Does this feature simply mean "single-stream capable," or will these devices also artificially limit the maximum per-stream bit rate as some suppliers have done in the past with 802.11b and 802.11g? Will anyone ever know the answer to this question, save for Apple and its rumored silicon supplier, Broadcom?

As 802.11n matures, and as suppliers strive to differentiate themselves, two- and three-stream-capable clients and three-stream-capable access points and routers will inevitably become more common. However, unless vendors clearly document the gear's capabilities and limitations and unless end users intuitively understand these limitations, the end result will inevitably be a combination of implementation frustration and purchase paralysis. EDN

Contact me at bdipert@edn.com.

+ For a more in-depth version of this write-up, go to www.edn.com/090528eda.

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

Buck regulators with integrated MOSFETs simplify designs up to 25A/phase

High efficiency

EFFICIENCY VS. OUTPUT CURRENT

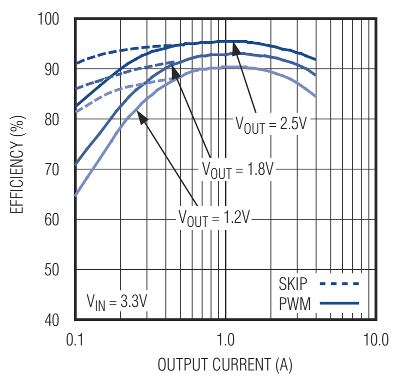

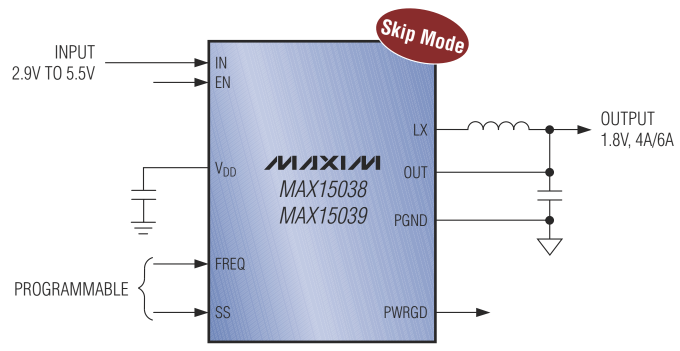

1- to 8-phase buck regulator with 25A/phase in a 6mm x 6mm TQFN

- 4.5V to 20V input voltage supply
- ±1% reference voltage accuracy over temperature
- Reference input (REFIN) for output tracking or system reference voltage
- Adjustable switching frequency from 300kHz to 1MHz
- Output sink and source current capability

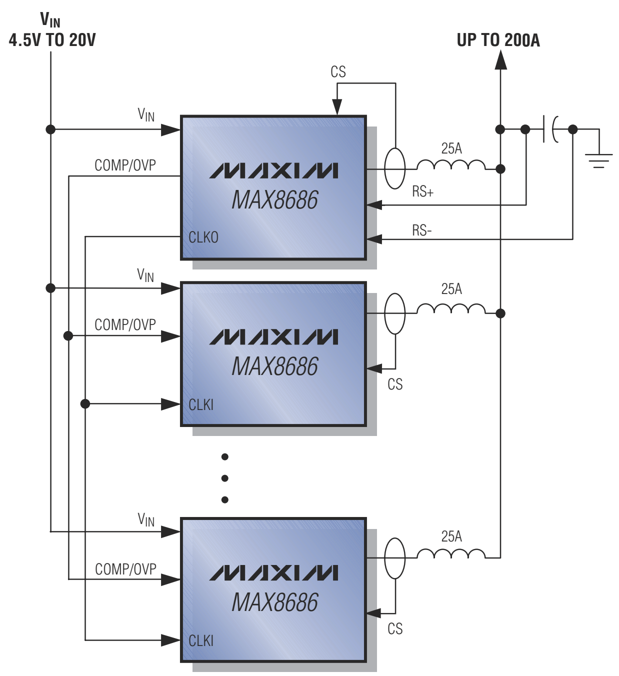

Maxim's PowerMind™ family of integrated buck regulators simplifies your toughest design challenges. These regulators provide high efficiency in compact, simple, cost-effective solutions that have the industry's lowest EMI.

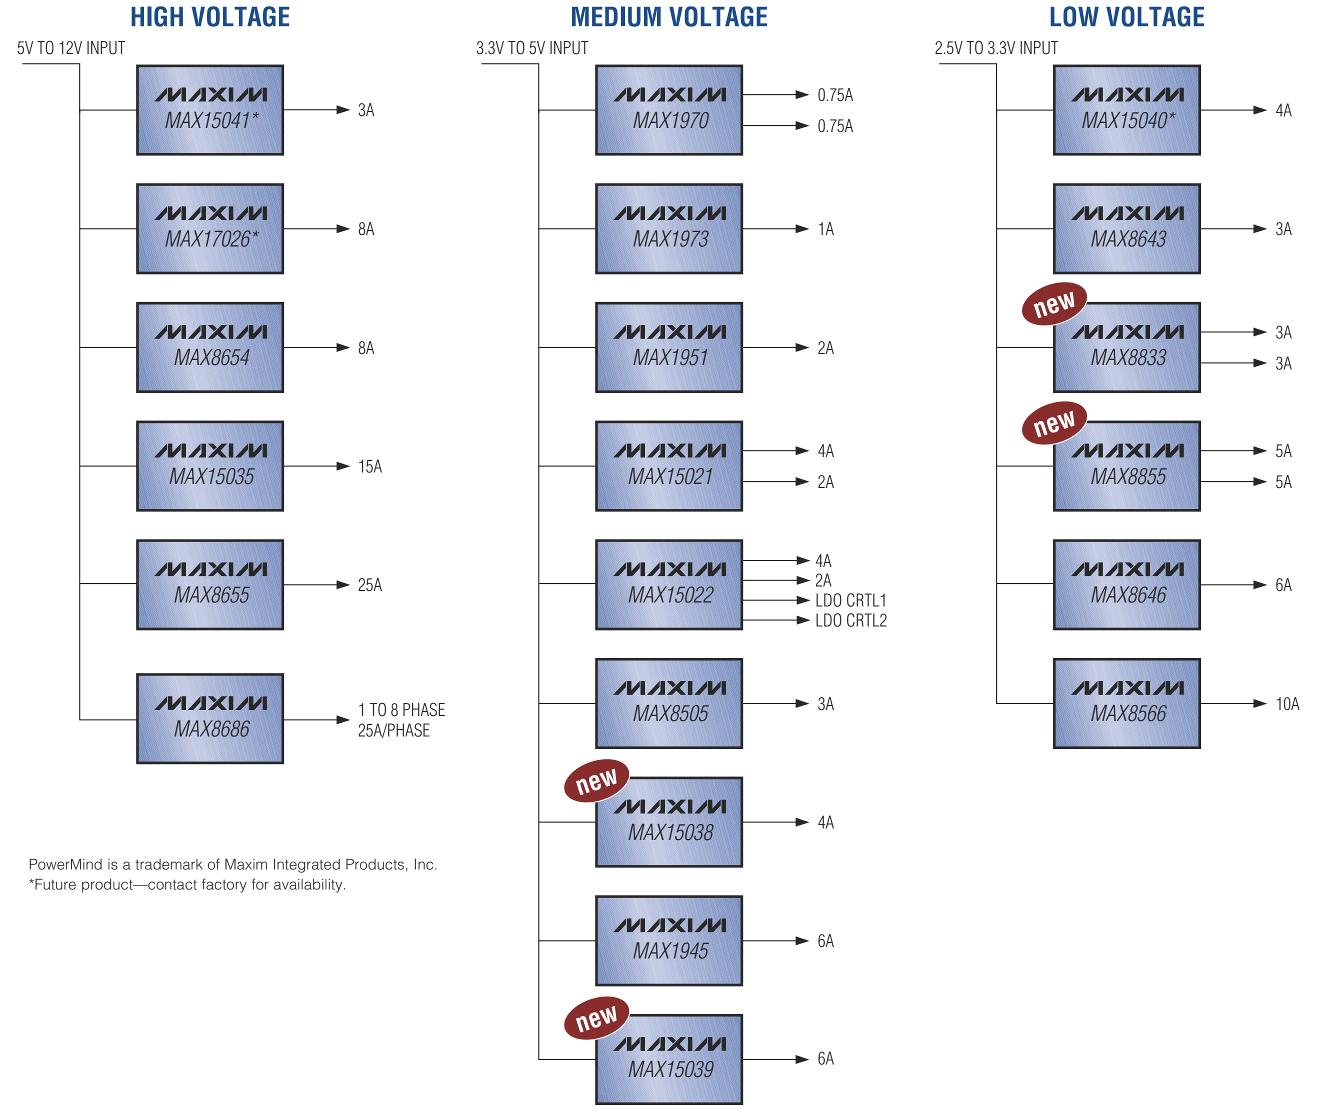

PowerMind is a trademark of Maxim Integrated Products, Inc.
*Future product—contact factory for availability.

www.maxim-ic.com/buck-regulators

Free ac/dc digital-power reference design has universal input and PFC

Microchip Technology based its new ac/dc reference design on the dsPIC33F GS series of digital-power DSCs (digital-signal controllers). The unit works with a universal input-voltage range and produces three output voltages. The ac input can range from 85 to 265V, and frequencies range from 45 to 65 Hz; continuous-power-output rating is 300W. A front-end PFC (power-factor-correction)-boost circuit converts universal ac-input voltages to a 420V-dc bus voltage. Input-power factor is greater than 0.98. An isolated buck converter uses the 420V PFC bus to create a 12V, 30A intermediate bus. The converter uses a ZVT (zero-voltage-transition) circuit to reduce losses, increase efficiency, and reduce stress on the power MOSFETs. The 12V bus then feeds a multiphase synchronous buck converter that can produce 69A at 3.3V. The 12V bus also feeds a single-phase buck converter that produces 23A at 5V.

You can download free design files and control-software code for this reference design.

The reference design has one four-layer board for digital signals and another for the other power stages. The design features soft-start capability and synchronous rectification. One dsPIC33F digital-power IC handles the PFC and primary-side ZVT-bridge control, and a second chip controls the converter stages.

You can download free documentation, including software and Gerber files, at www.microchip.com/smps. For an expanded version of this write-up, go to www.edn.com/article/CA6651496.**—by Paul Rako**

▷**Microchip Technology**, www.microchip.com.

FEEDBACK LOOP

"I love these stories; please don't stop! As tough as errors and problems can sometimes be, they offer themselves as challenging learning opportunities nonetheless."

—Engineer and *EDN* reader C Kaspereli, in *EDN*'s Feedback Loop, at www.edn.com/article/CA6648806. Add your comments.

DSP combines fixed- and floating-point operation

Texas Instruments' TMS320C6743 low-power DSP features both high-precision, wide-dynamic-range, floating-point operations and higher-performance, lower-power, fixed-point operations. The C674x core uses a superset of the fixed-point C64x+ and the floating-point C67x+ instruction sets. The operations support a wider dynamic range than do fixed-point operations.

The C6743 includes 192 kbytes of on-chip memory and consumes 490 mW in operating mode and 60 mW in standby mode. Peripheral support includes a 10/100-Mbps Ethernet MAC (media-access controller) with a management-data I/O module; a multimedia-card/secure-digital peripheral; enhanced, high-resolution PWM (pulse-width modulation); 32-bit enhanced quadrature-pulse peripherals; and two external-memory interfaces—one for slower, asynchronous memories and the other for SDRAM.

The C6743 is object-code-compatible with all TMS320C6000 devices; it is scalable and pin-to-pin compatible with C674x and OMAP (Open Multimedia Applications Platform)-L137 devices. It is available for sampling now and costs $7.85 (1000). The OMAP-L1x/C674x starter kit, TMDXOSKL137BET, which is available for $395, includes a board-specific Code Composer Studio IDE (integrated development environment) and supports DSP/BIOS kernels. For an expanded version of this write-up, go to www.edn.com/article/CA6653671.**—by Robert Cravotta**

▷**Texas Instruments**, www.ti.com/c6743dsp.

Single IC integrates both drivers and controllers for HB-LED lighting

Most solid-state lighting applications require more electronics than just a series of HB LEDs (high-brightness light-emitting diodes) and their associated power-control circuitry. Lighting features such as communication and control and dimming, among others, require additional circuitry, including high-current drive electronics.

To address these needs, Cypress Semiconductor has introduced its PowerPSoC (programmable-system-on-chip) family of integrated, embedded power controllers, which the company claims is the first single-chip device for both controlling and driving HB LEDs. The company based the controllers on the PSoC, which includes an 8-bit microcontroller, programmable analog and digital circuits, and memory. The new family also includes four constant-current regulators and four 32V MOSFETs.

The PowerPSoC includes four 2-MHz hysteretic controllers, which you can configure as buck, boost, or buck-boost, and a 32V input-voltage regulator. The underlying PSoC platform offers Cypress' CapSense programmable touch-sense control and other digital, analog, and communication options. Digital options include PWMs (pulse-width modulators), timers, and counters; analog options include ADCs and PGAs (programmable-gain amplifiers); and communication options include the DMX512 lighting-communication protocol, DALI (digital-addressable-lighting interface), SPI (serial-peripheral interface), and RS-232.

Available in a 56-pin QFN, the PowerSoC family sells for less than $10 (sample quantities). The CY3267 PowerPSoC evaluation kit sells for $175.

—by Margery Conner

▷**Cypress Semiconductor**, www.cypress.com.

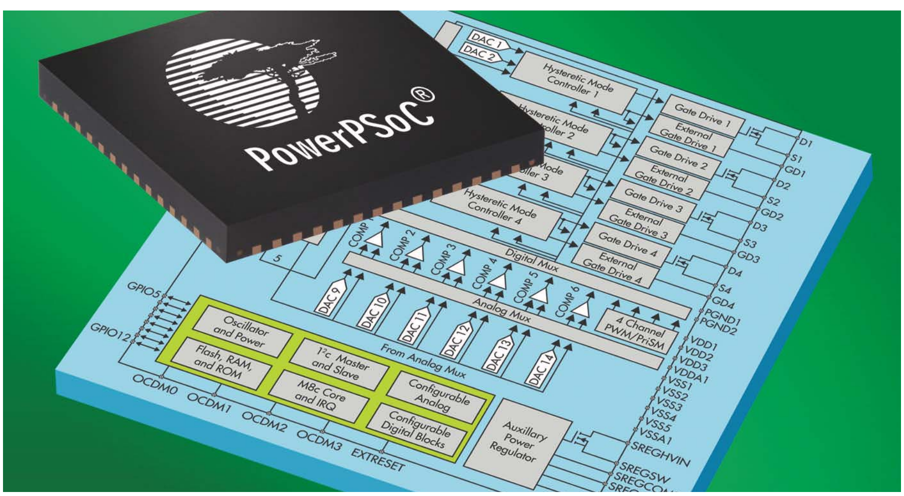

The PowerPSOC family of integrated, embedded power controllers is a single-chip device that includes an 8-bit microcontroller, programmable analog and digital circuits, memory, four constant-current regulators, and four 32V MOSFETs.

DILBERT By Scott Adams

05.28.09

AGILENT, ANITE DEMONSTRATE LTE STREAMING VIDEO

Agilent's E6620A wireless-communications test set integrates a 3GPP Release 8-compliant protocol stack.

Agilent Technologies and Anite (www.anite.com) have teamed up to demonstrate 3GPP (Third Generation Partnership Project, www.3gpp.org) LTE (long-term-evolution) interoperability at the Mobile World Congress (www.mobileworldcongress.com), which took place in February in Barcelona, Spain. The demonstration showed end-to-end streaming video with a Blue Wonder (www.bluwo.com) UE (user-equipment) chip set, Signalion (www.signalion.com) Sorbas LTE UE simulator, Agilent E6620A LTE wireless-communications test set, and Anite LTE-protocol-development SAT (stand-alone tester).

Agilent's E6620A offers real-time, system-rate network emulation for L1/L2/L3 uplink and downlink through RF or digital baseband. For an expanded version of this write-up, plus more news from the Mobile World Congress, go to www.edn.com/090528pa.

—by Rick Nelson

▶Agilent Technologies, www.agilent.com.

DSP-based instruments radically alter swept-spectrum analysis

Tektronix has announced enhancements to its RSA6000 Series of DSP-based real-time spectrum analyzers. With the aid of new hardware and software, the enhanced instruments use advanced triggering technology and real-time signal analysis to deliver unique diagnostic capabilities that speed analysis in resolving performance problems in systems for spectrum management, radar, electronic warfare, and radio communications at frequencies to 14 GHz. The enhanced units' fast measurements and short test times thus help system engineers to control project costs.

An increased need to combine digital computing with traditional RF technologies has presented spectrum-analyzer users with new challenges, such as broadband transients, DSP errors, and runt pulses.

The use of analyzers with only narrowband filters, which impose slow spectrum-update rates, makes troubleshooting of these signals increasingly complex and time-consuming. Allowed transmissions can be difficult to separate from elusive transients, with which they often overlap in time and frequency. Such interferers and hostile signals can also cause intermittent failures, which are unacceptable in high-reliability applications.

> "The analyzers capture 292,000 spectra/sec—six times as many as the previous generation.

The instruments' advanced time-, amplitude-, and DPX (digital-phosphor-technology)-triggering functions and swept-DPX capability enable faster discovery and capture of these intermittent and rapidly changing signals.

The spectrum analyzers' improved broad-sweep capability enables rapid detection of signals of interest. The DPX engine collects hundreds of thousands of spectra per second over a 110-MHz bandwidth. By enabling the instrument to sweep across its full input-frequency range of as much as 14 GHz and capture tens or hundreds of spectra in no more time than a conventional spectrum analyzer needs to capture just one spectrum, this enhancement greatly reduces the possibility of missing time-interleaved or transient signals during broadband searches. The analyzers now capture 292,000 spectra/sec—six times as many spectra per second as the manufacturer's previous generation of DPX-equipped spectrum analyzers. With this feature, you can be confident of detecting transients as brief as 10.3 μsec, which conventional spectrum analyzers miss. This feature is particularly important to those who work with such systems as software-defined radio and radar.

The DPX trigger also enables the instruments to trigger on signals within signals. You invoke this capability by pointing at the signal on the spectrum display and selecting the "trigger-on-this" feature. You can also time-qualify any trigger. The instruments' time- and amplitude-qualified triggers reduce troubleshooting time in such applications as radar and spectrum management. Suggested US prices for the RSA6000 Series instruments begin at $77,900. For those who own earlier instruments in the series, a user-installable hardware/software upgrade kit costs $12,900.

—by Dan Strassberg

▷Tektronix Inc, www.tektronix.com.

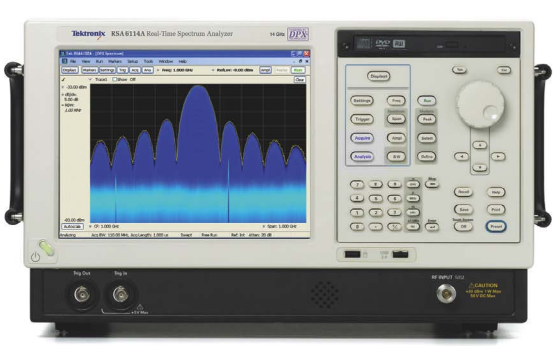

The RSA6000 Series real-time spectrum analyzers' DPX technology and extensive use of digital-signal processing achieve 100% probability of intercept and enable the newly enhanced instruments to capture and present the results of tens or hundreds of broadband sweeps in the time that conventional swept-frequency analyzers require to display one sweep.

05.28.09

VMEBUS COMPUTER COMBATS EOL ISSUES

General Micro Systems recently announced the VS275 Maritime VMEbus (Versamodule-eurocard-bus) single-board computer. The device has five expansion modules to combat EOL (end-of-life) problems by enabling seamless upgrades and new functions. Maritime uses the GMS P70 Nucleus processor module that supports either Intel's (www.intel.com) ultralow-power, 1.5-GHz Core 2 Duo or 2.16-GHz Core 2 Duo, which has 4 Mbytes of L2 cache and a 667-MHz front-side bus. PMC-X (peripheral-component-interconnect-mezzanine-card-express)- and XMC (express-mezzanine-card)-compliant sites enable plug-in configurability.

The VS275 VMEbus single-board computer has five expansion modules.

You can configure the Maritime board with a 2.5-in., 500-Gbyte SATA (serial-advanced-technology-attachment) hard-disk drive or a 256-Gbtye solid-state drive in place of the onboard PMC site. Prices for the standard version of Maritime, VS275, start at $3200 (one). For an expanded version of this article, see www.edn.com/090528pb.

—by Warren Webb

▶General Micro Systems, www.generalmicro.com.

VOICES

Is multimedia abandoning DSPs? Analog Devices disagrees

EDN recently interviewed Denis Labrecque, marketing-programs manager at Analog Devices. In the interview, Labrecque sounds off on DSPs' past, present, and future in audio, graphics, and imaging. He provides his perspectives on competition within his company's product line; against other DSP suppliers; and versus other silicon-based processing architectures, such as CPUs and FPGAs. The following is an excerpt of that interview. For the full version, go to www.edn.com/090528pc.

Can you provide a historical background on how your audio customers' needs—performance and otherwise—have evolved over time, and how Analog Devices has responded to them?

A There has been a constant demand for higher-performance floating-point audio processors. Demand started in the professional market, where programmable parts enabled new features and flexibility. Next came consumer applications driven by digital content and enabled by lower-cost parts. Analog Devices created processors with a specific mix of peripherals, memory, and ROM aimed at particular market categories.

Certain peripherals—a large number of serial ports, flexible signal-routing units, and onboard sample-rate converters—have been enthusiastically embraced by the audio community. For example, the recently announced 21469 processor provides hardware accelerators for FIR [finite-impulse-response] and IIR [infinite-impulse-response] filters as well as FFTs [fast Fourier transforms], reflecting the desire for custom room-equalization functions.

Analog Devices' work with Intel resulted in the Blackfin processor, which combines 16-bit digital-signal processing and primary-system-processor capabilities in one device. Focusing on audio, to what degree has Blackfin expanded the market that Analog Devices' products serve?

A SHARC and Blackfin fill complementary market positions. SHARC has always been and will continue to be the company's flagship high-end processor for not only the audio market but also medical and industrial. Blackfin, though a native 16-bit processor, is a very capable audio processor and is optimized to perform equally well for signal processing, control processing, and media processing, making it attractive for a broad range of multimedia applications. Its high clock speed, up to 2400 MIPS, and dual 16-bit MAC [multiply/accumulate] unit enable it to do a serious amount of audio processing. Even with the over-

head associated with 32-bit, double-precision processing, a 600-MHz Blackfin can potentially do as much processing as a 150-MHz SHARC.

To help customers decide what to use, we need to know what audio application and functions they need. Is the primary function of the product audio processing for a home-theater system, automotive amplifier, mixing console, broadcast processor? Then use SHARC. However, if the product is primarily about communication, the user interface, or streaming but needs some audio features, then use Blackfin. Products falling into this category are portable media players, streaming-network nodes, or automotive head units.

A recent survey of audio-design engineers revealed robust interest in two non-DSP platforms for sound processing: x86 CPUs and FPGAs. The x86 CPU could even be the one already in a Mac or a PC, shifting expense from the audio peripheral to the more function-versatile host computer. How do you combat this attractive system-partitioning approach?

A Native processing has long been a viable solution for audio-production and -postproduction work. However, it is unlikely that the PC environment will ever achieve the low latency and real-time reliability for true professional-audio applications. Do you really want the integrity of your audio system dependent upon Windows? Do you know how threads are active on your PC? A large number of professional-digital-audio workstations in fact move native processing to internal/external audio processors.

Much has been written lately about the increasing capabilities of FPGAs. In pure processing performance, they can beat out programmable DSPs. FPGAs are good at algorithmically simple operations, such as filters and FFTs. These can be easily described in gates and implemented efficiently. We're familiar with these advantages, and [we've] embedded fixed-function hardware accelerators in the latest SHARC 21469 processor. These accelerators offload computationally intensive operations, allowing the core to focus on other tasks. The combination more than doubles the overall numeric performance of SHARC.

In many applications, we coexist nicely with FPGAs, and, in most of these cases, each product brings its special advantages to the design. Analog Devices expects to continue supplying floating- and fixed-point devices to run most algorithms. We have seen FPGAs used to implement highly specialized and customized functions that are not available in a general solution, such as handling large amounts of data movement specific to an application. The requirements are not always related to speed but instead are related to the functional needs of the system design.

—interview conducted and edited by Brian Dipert

RAQ's

Rarely Asked Questions

Strange stories from the call logs of Analog Devices

Glass Diodes May See the Light — and Hum

Q. A few months ago you discussed how to prevent the inherent high frequency (HF) noise of a switching supply from escaping into sensitive analog circuits and corrupting them. My op-amp output hums and no amount of supply filtering helps. What should I do?

A. Small children who've recently discovered riddles will tell you that bees hum because they don't know the words. Op-amps usually hum for other reasons.

The classical reason for an electronic circuit to hum was line frequency (or, more usually, double line frequency) ripple on the dc supply. This is addressed by enlarging the supply decoupling capacitors or, better, by using an electronic voltage stabilization circuit—which has the side-effect of attenuating low frequency (LF) ripple. Modern switching supplies, as we have seen, generate large amounts of HF noise, which must be imprisoned at the supply and not allowed to escape. They rarely have LF noise problems, so your supply is probably not the cause. Measure its LF noise with an oscilloscope or spectrum analyzer to confirm this.

Another cause of hum is line frequency currents in the signal line or ground (the dreaded "earth loop") due to line currents or voltages being too close to signal circuitry. Again, possible cures were discussed in the earlier "Imprison Noise" RAQ and work well at LF as well as HF. If you have successfully implemented HF noise reduction, you have probably addressed most causes of power supply related LF noise as well. Test this by operating your circuit from batteries (with simple linear IC regulators to stabilize voltages if necessary), putting a resistive load on your switching supply to simulate the op-amp's consumption. Check noise with the switching supply on and off.

If you still have hum you must look for the mechanism by which it is entering the op-amp. A surprisingly common way is by photo-current from input protection diodes. Glass packaged silicon diodes behave as photocells. If they are illuminated with modulated light, their leakage current is modulated by the light; and, if the modulation is 120 Hz (100 Hz in some countries) from line-operated fluorescent lights, the circuit incorporating them will hum. The effect is not characterized by diode manufacturers and may vary widely from one "identical" device to another. Fix the problem by using plastic diodes rather than glass ones.

---

[1] *Another reason to hum is implicit in the story of the man who went into a bar and noticed a pianist playing requests with a small monkey collecting his tips in a fez. While he listened the monkey drank his beer. He tapped the pianist's shoulder "Do you know your monkey drank my beer?" "No," replied the pianist, "but hum it - I expect I can fake it.")*

[2] *Please see the Lock Down That Noise - Don't Let It Escape RAQ at http://www.analog.com/raq/noise.*

Contributing Writer James Bryant has been a European Applications Manager with Analog Devices since 1982. He holds a degree in Physics and Philosophy from the University of Leeds. He is also C.Eng., Eur. Eng., MIEE, and an FBIS. In addition to his passion for engineering, James is a radio ham and holds the call sign G4CLF.

Have a question involving a perplexing or unusual analog problem? Submit your question to: raq@reedbusiness.com

BY BONNIE BAKER

The non-negotiable single-supply operational amplifier

Fundamental analog devices that serve applications such as high-resolution delta-sigma or SAR (successive-approximation-register) converter systems are feeling the crunch from amplifiers that have difficulty with achieving good rail-to-rail input performance. The simple rail-to-rail operational amplifier must have a transistor design that spans the power supply with minimal distortion.

The trend toward designing single-supply op amps started in the 1970s with a single differential-input stage that spanned a portion of the common-mode input range. Later, designers added a second, or complementary, differential-input stage. The two stages shared, with some distortion, the rail-to-rail input operation across the complete amplifier's rail-to-rail common-mode range (**Reference 1**). Neither of these approaches produced an amplifier adequate for the high-precision systems to span the amplifier's full common-mode input range.

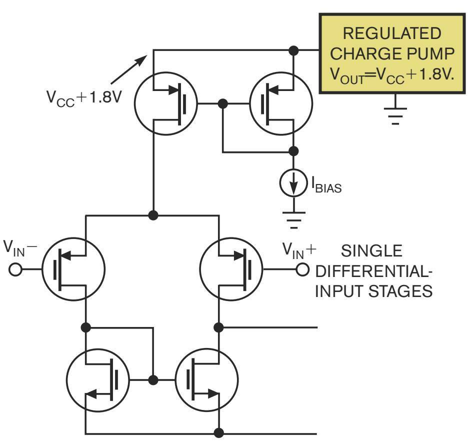

Figure 1 In this configuration, a charge pump pushes a single differential-input stage of the amplifier above the positive-power supply.

Eventually, IC designers borrowed a technology from other devices to solve this problem. They began to use the all-too-common charge pump to push a single differential-input stage of the amplifier above the positive-power supply (**Figure 1**). Amplifier designers place the switching mechanism's frequency above the amplifier's bandwidth and keep the switching noise lower than the amplifier's thermal noise floor.

The single differential-input stage with a charge pump buys you a 20- to 30-dB increase in the amplifier's CMMR (common-mode-rejection ratio). This increase has a positive effect on amplifiers in buffer configurations. You can also expect almost a tenfold decrease in the amplifier's THD (total-harmonic-distortion) performance. So, if you use an amplifier that has a charge pump in its input stage to drive high-precision SAR or delta-sigma converters, your system's performance will improve.

For example, the THD of an ADC driven by an op amp in a buffer configuration is the root-sum square of distortion contributions of the ADC and op amp. In this configuration, the system THD is:

$$THD_{SYSTEM} = 20\log\sqrt{10^{(THD_{ADC}/10)}+10^{(THD_{OPA}/10)}},$$

where $THD_{OPA}=20\log(THD_{OPA-\%}\times 100)$ and $THD_{OPA-\%}$ is the THD specification in the operational amplifier's data sheet in units of percentage.

Using these **equations**, if an operational amplifier with a complementary input stage has a THD specification of 0.004%, with an input voltage of 4V p-p, and the 16-bit SAR ADC has a THD specification of −99 dB, the system THD is −88 dB. Alternatively, if the op amp's input stage has a charge pump with a THD specification that is 0.0004%, the system THD becomes −98 dB.

Single-supply amplifiers continue to keep pace with high-resolution converters because engineers implement innovative amplifier-circuit topologies, such as an input stage with a charge pump. The charge pump is a good stopgap; however, engineers continue to demand lower system power supplies and insist on better signal integrity.EDN

REFERENCES

1. Baker, Bonnie, "Where did all the racket come from?" *EDN*, April 23, 2009, pg 18, www.edn.com/article/CA6651590.
2. "OPA365, OPA2365 2.2V, 50MHz, Low-Noise, Single-Supply Rail-to-Rail Operational Amplifiers," Texas Instruments, June 2006, www.ti.com/opa365-ca.
3. "OPA333, OPA2333 1.8V, micro-Power CMOS Operational Amplifiers, Zero-Drift Series," Texas Instruments, March 2006, www.ti.com/opa333-ca.

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

PRY FURTHER AT EDN.COM

+ Go to www.edn.com/090528pry for a more in-depth version of this write-up, including an additional image.

Trolling for gold in the BlackBerry Bold

Research In Motion's first 3G-data-capable "world" phone, the BlackBerry 9000, or Bold, gets a tech inspection courtesy of an *EDN* partnership with phoneWreck. How did RIM squeeze triple- or quad-band UMTS, triple-band Wi-Fi, a QWERTY keypad, and copious additional capabilities into a 4.5×2.6×0.5-in., 4-oz form factor?

The integrated SiRF GPS receiver works with the phone's cellular subsystems in an A-GPS mode. The phone's LED-flash-augmented, 2M-pixel camera has digital-zoom capabilities and captures 480×320-pixel video streams. Texas Instruments' TLV320AIC3106 codec handles audio capture and playback.

The processing heart of the BlackBerry Bold is Marvell's PXA930 Tavor ARM-based processor running at 624 MHz. The CPU's memory partner is a triple-technology multidie-package stack from Samsung: 1 Gbyte of MLC (multilevel-cell) moviNAND flash memory; 128 Mbytes of bootable, SLC (single-level-cell) OneNAND flash memory; and 128 Mbytes of mobile-DDR SDRAM.

The PCB (printed-circuit-board) backside (not shown) is comparatively sparse. It exposes the action end of the trackball module, along with the QWERTY keypad's membrane switches, and it reserves sufficient space for the Samsung 2.65-in.-diagonal half-VGA LCD with 64,000-color support.

The BlackBerry Bold's cellular subsystem contains Anadigics' AWT6221 and AWT6241 power-amplifier chip set, along with Infineon's WCDMA PMB5701 transceiver. The phone supports UMTS 3G-data services in the 850-, 1900-, and 2100-MHz bands, along with the 800-MHz band in a version for Japan. A Renesas power-amp-plus-transceiver combo handles the 850-, 900-, 1800-, and 1900-MHz GSM voice and data frequency-band options.

Wireless-communications options other than cellular include dual-band 802.11a/b/g Wi-Fi, which a four-chip transceiver-plus-power-amp-plus-power-management cluster from Texas Instruments implements. The BlackBerry Bold also includes Cambridge Silicon Radio's BlueCore4 Bluetooth 2.0 transceiver supporting A2DP stereo and AVRCP protocols. TI's TPS65850 handles more general systemic power-management control. Cypress Semiconductor's CYWB0124AB administers both mini-USB (Universal Serial Bus) 2.0 transfers and the microSDHC (secure-digital-high-capacity)-memory-card interface.

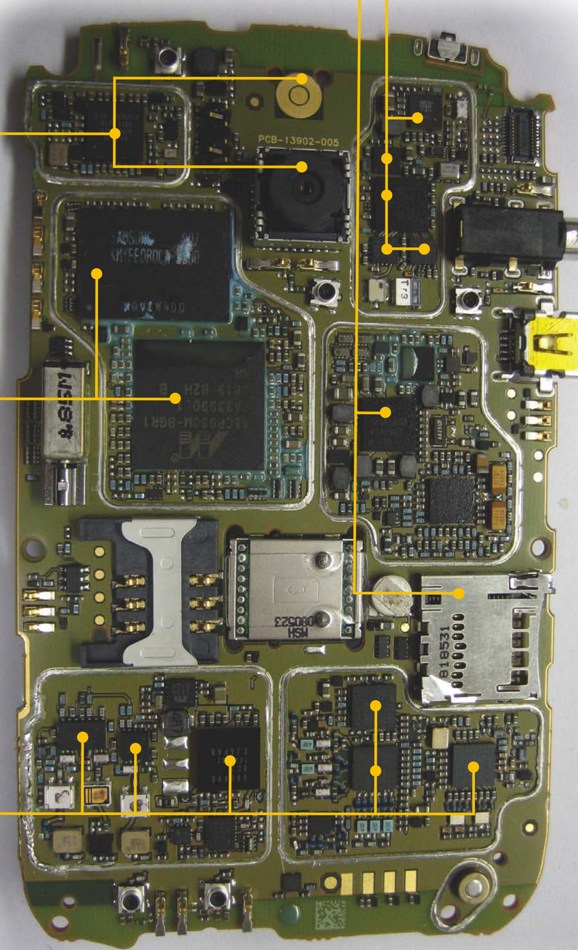

Agilent

Tektronix

LeCroy

Rohde & Schwarz

Anritsu

Keithley

Yokogawa

Tabor

Pickering

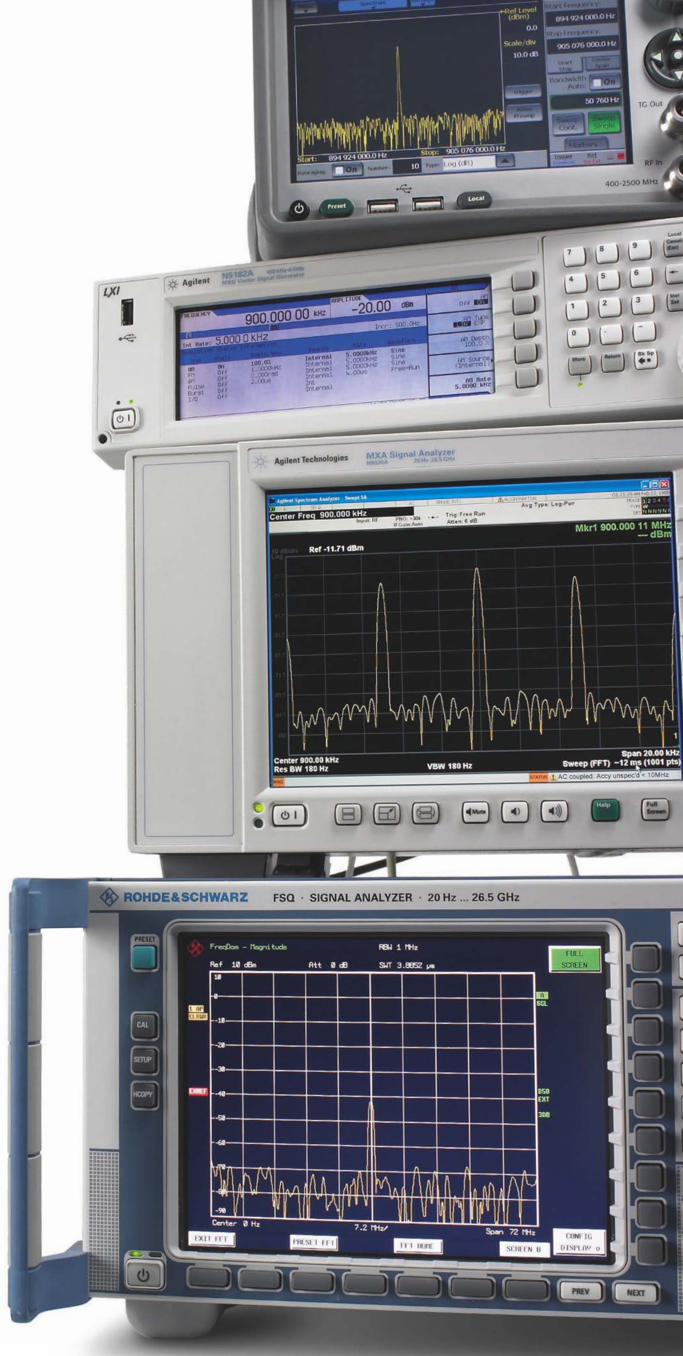

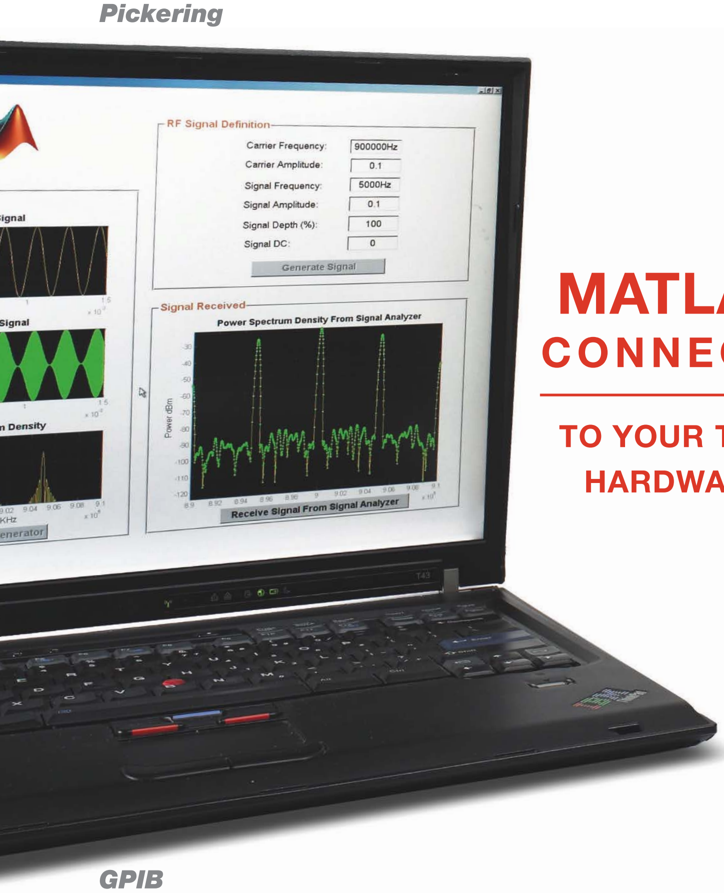

MATLAB CONNECTS

TO YOUR TEST HARDWARE

Connect to your test equipment directly from MATLAB® using standard communication protocols and hundreds of available instrument drivers.

Analyze and visualize your test results using the full numerical and graphical power of MATLAB.

For more information on supported hardware, visit www.mathworks.com/connect

GPIB

TCP/IP

VISA

USB

UDP

RS-232

The MathWorks™

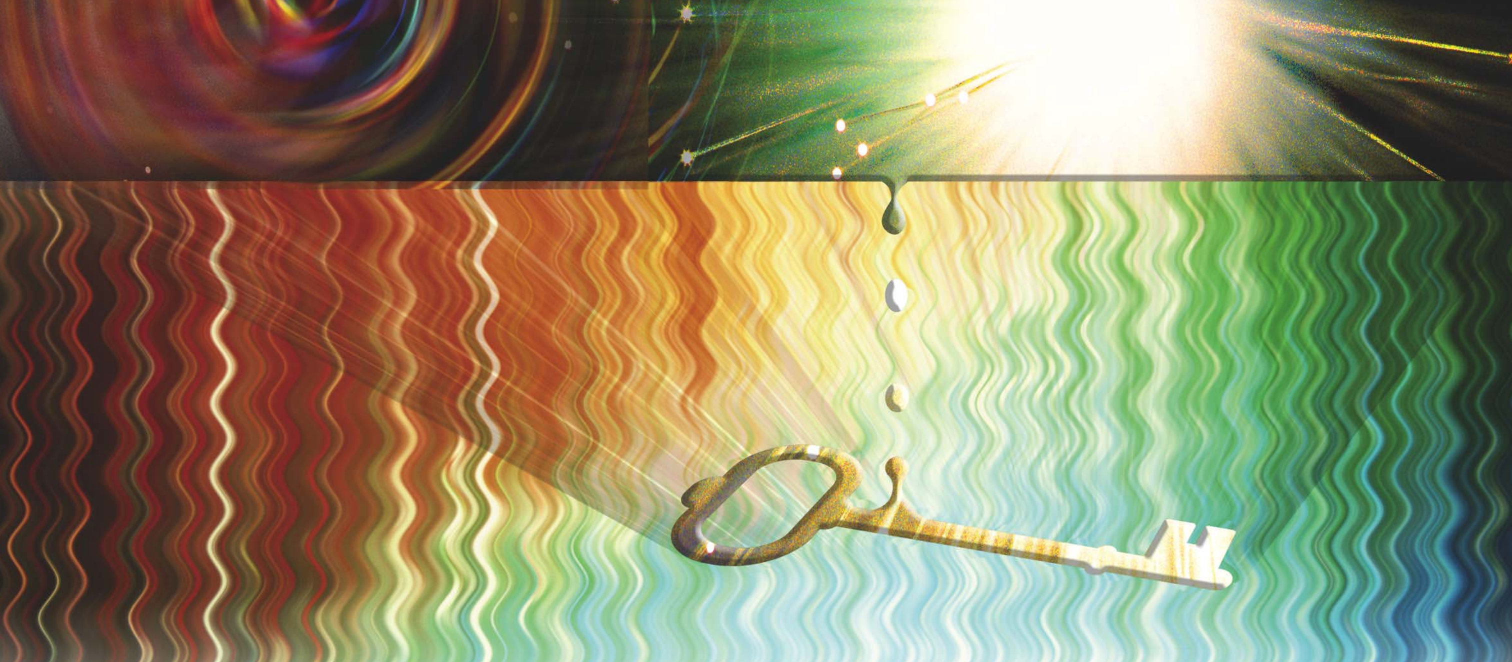

LEAKAGE CURRENT MAY BE THE GATE NO ONE THOUGHT TO LOCK THAT ALLOWS HACKERS TO CRACK HARDENED SYSTEMS.

Leakage-power analysis ENABLES ATTACKS on cryptographic devices

BY MILENA JOVANOVIC • UNIVERSITY OF MONTENEGRO

Security requirements are increasingly stringent in applications such as e-commerce and electronic banking. Although encryption technology provides robust algorithms, the physical implementation of those algorithms usually leaks information through physical phenomena relating to the electrical operation of the devices, which an attacker could use to detect the secret key. These "side-channel" attacks use information that the hardware implementation of encryption modules leaks. This information could include correlation between data and power consumption or timing (**Reference 1**). Differential-power analysis is a well-documented, powerful side-channel attack because it allows the attacker to detect secret keys by using a measurement setup employing off-the-shelf components (**Reference 2**). The attacker relies on the fact that standard CMOS logic exhibits a dynamic-power consumption that strongly depends on the input data. For example, consider a simplified model of a CMOS inverter that receives its load from a capacitance that connects to ground. The model draws its current from a supply only for a zero-to-one output transition. In a one-to-zero transition, the energy in the output capacitance dissipates, and the circuit consumes no power for zero-to-zero and one-to-one transitions. Engineers have recently proposed many countermeasures employing both software and RTL (register-transfer-level) logic to thwart attacks through dynamic-power analysis (**Reference 3**).

Historically, the primary contributor to power dissipation in CMOS circuits has been dynamic power due to CMOS-switching activity. Dynamic power has a quadratic dependence on supply voltage and a linear dependence on clock frequency. Another important contribution to power consumption in CMOS circuits is leakage power due to parasitic currents in switched-off CMOS devices. That leakage power will most likely soon approach a level comparable with that of dynamic-power consumption (**Reference 4**). Leakage current in a CMOS design strongly depends on the input-data vector, and engineers have used this property to reduce leakage-power dissipation during circuits' standby periods (**references 5**, **6**, and **7**). They have also pro-

posed models that can estimate the input vector that produces maximum and minimum leakage current, respectively, in CMOS circuits (**Reference 8**).

Given leakage current's dependence on input values in CMOS logic, you can use leakage-current measurements to extract information about the secret data in a cryptographic core. After analyzing the dependency of leakage current on input data on a simple cryptographic core using RTL simulations, you can use statistical-analysis techniques to mount attacks. These techniques are similar to those that differential-power analysis uses and allow you to extract the secret key starting from leakage-current measurements, which are, in principle, easier to perform.

AT A GLANCE

- Differential-power analysis allows the attacker to detect secret keys by using a measurement setup employing off-the-shelf components.
- The primary contributor to power dissipation in CMOS circuits is dynamic power due to CMOS-switching activity.
- You typically use substitution boxes to obscure the relationship between the plain text and the cipher text.
- You collect side-channel information by measuring some physical quantity.

DATA DEPENDENCE

The main sources of leakage current are inverse-junction, subthreshold, and gate-tunnel current (**Reference 9**). Subthreshold current is the most important contributor in a MOS transistor biased in a weak inversion region. Designers of standard-CMOS gates base them on pullup and pulldown networks comprising parallel and series connections. Designers have developed models for the leakage current of MOS devices in series and parallel configurations.

To understand the data dependence of leakage current in standard-CMOS gates, researchers performed simulations on the standard cells of the 90-nm CMOS090 process from STMicroelectronics (www.st.com). These simulations used foundry-supplied models for the Spectre transistor-level simulator, employing five temperatures to verify the temperature dependence of leakage current (**Table 1**, which is available with the Web version of this article at www.edn.com/090528df). If you sort the leakage currents in the **table** in ascending order, the same order remains with temperature variations. For example, in a two-input NAND gate, logic input 01 can generate the maximum leakage current for all temperature values. This approach to power analysis exploits the input combination that produces the maximum or minimum leakage current of a cryptographic core. Available literature describes maximum, minimum, and other algorithms to estimate the leakage current.

The basic component of symmetric-key algorithms is the substitution box. When dealing with block ciphers, you typically use substitution boxes to obscure the relationship between the plain text and the cipher text. A substitution box provides a combinational mapping between an N-bit input word and an M-bit output word. These boxes—for example, those with four inputs and four outputs—normally use fixed tables.

Employing a truth table and a 90-nm-CMOS-process library from STMicroelectronics, Cadence (www.cadence.com) researchers synthesized a "serpent" substitution box—one having four inputs and four outputs (**Table 2**, which is available with the Web version of this article at www.edn.com/090528df). Researchers then performed extensive leakage-current simulations for all possible input combinations of the substitution box. If you sort the input combinations in the **table** in increasing order of leakage-current consumption, the results remain independent of temperature variations. Thus, you can use any thermal co-

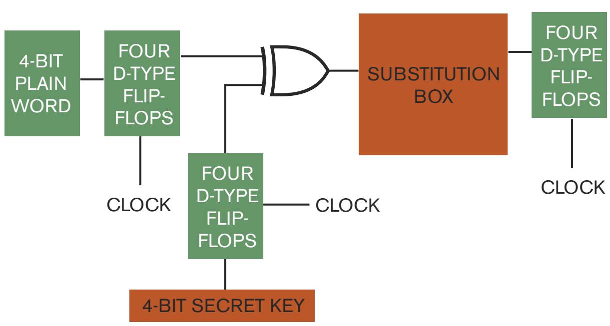

Figure 1 You can build a simple cryptographic core by adding registers to a combinational design.

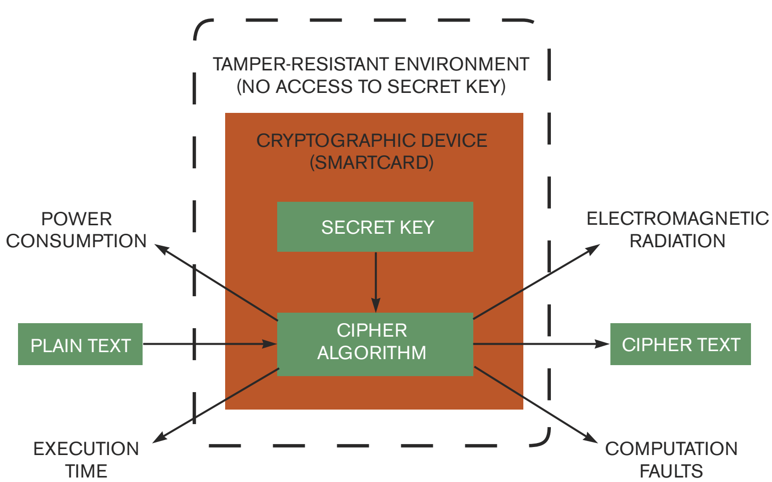

Figure 2 Passive attacks benefit from side-channel information, which you collect by measuring some physical quantity.

efficient, providing that it remains constant, in simulations or measurements. You can realize a combinational part of a simple cryptographic core by connecting XOR gates to the inputs of a substitution box. XOR gates premix any plain words and secret keys, and a substitution box ciphers the results. You can perform extensive leakage-current simulations on the combinational part of a cryptographic core for all values of keys and inputs. If you put the leakage-current values in ascending order, the order of inputs and outputs will be the same for each key, meaning that leakage current does not depend on the values of the inputs. These simulations also show that the values of leakage current differ for different keys and the same inputs. XOR gates, which are sensitive to the input changes, cause this difference. The substitution box is not the culprit because, for the same input, you can always measure the same current.

You can build a simple cryptographic core by adding registers to a combinational design (**Figure 1**). **Table 3**, which is available with the Web version of this article at www.edn.com/090528df, reports the leakage current of the registers for the input combinations and shows how the leakage current directly relates to the number of ones in a binary word. Simulations on whole cryptographic cores explore all possible combinations of plain words and keys (**Table 4**, which is available with the Web version of this article at www.edn.com/090528df). The **table** sorts leakage currents in increasing order and groups them by input keys; input columns are the inputs of the substitution box, and output columns are the outputs of the cryptographic core.

LEAKAGE-CURRENT ATTACK

Researchers typically divide attacks on cryptographic algorithms into mathematical and implementation categories, basing the implementation category on passive or active weakness. Researchers are trying to design countermeasures to thwart these powerful attacks (**Reference 10**). Passive attacks benefit from side-channel information, which you collect by measuring some physical quantity (**Figure 2**). Active attacks are more invasive because they introduce faults that result in erroneous calculations, leading to the exposure of a secret key.

The most common side channel for attacks is a device's power consumption. These types of attacks use simple power analysis, differential-power analysis, and correlation-power analysis (**Reference 11**). In a simple-power-analysis attack, an attacker uses the side-channel information from one measurement to directly determine a secret key or parts of a secret key. Differential- and correlation-power analyses are statistical at-

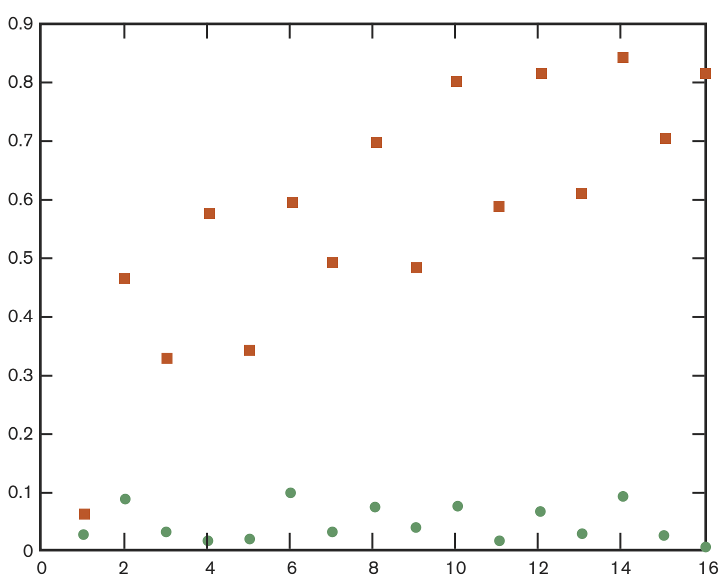

Figure 3 Results from the leakage-power-analysis attack show that all keys but the 0000 key are clearly distinguishable from the leakage-power data.

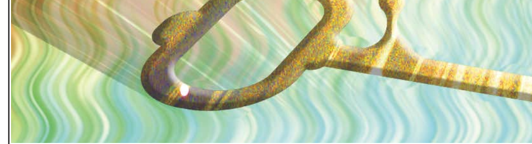

Go to www.edn.com/090528df and click on Feedback Loop to post a comment on this article.

For a list of the references cited in this article, go to www.edn.com/090528df.

tacks and need a lot of measurements and data acquisition. A correlation-power attack involves collecting data and then analyzing the collected data and uses a so-called hypothetical model of the attacked device.

You can use a statistical-analysis tool to detect the key of the cryptographic core. You can also use an attack employing the computation of the correlation coefficient between the vector of acquired leakage currents and a logic vector that uses a hypothesis of the secret key. The method uses the hamming weight, or number of ones in a binary word, of the inputs in the substitution box as a logic vector. You can obtain this weight from a key hypothesis. **Figure 3** shows the computed correlation coefficients. Orange squares denote the right-key hypothesis, and green circles denote a random-key hypothesis. The result is that all keys, except key 0000, are clearly distinguishable by using this kind of attack. Researchers are investigating why you get different results for key 0000.

This preliminary study shows the realistic possibility of using leakage currents to reveal secret keys. The result of the attack employing correlation coefficients suggests that leakage-power analysis could become a problem you should consider during cryptographic-core design, especially for cores in technologies with gates shorter than 0.1 micron that exhibit a high leakage-power consumption.EDN

AUTHOR'S BIOGRAPHY

Milena Jovanovic has been a teaching assistant in the electrical-engineering department at the University of Montenegro (Podgorica, Montenegro) since January 2007, teaching courses in electronics, computer peripherals and interfaces, and electrical-engineering materials. Jovanovic has a master's degree in electronics and is currently pursuing a doctorate in electrical engineering at the University of Montenegro. You can reach her at jmilena1983@yahoo.com.

STRESS OUT; STRAIN GAUGES IN

BY PAUL RAKO • TECHNICAL EDITOR

THE CAREFUL SELECTION AND APPLICATION OF STRAIN GAUGES CAN ENSURE GOOD RESULTS FROM THIS DIFFICULT MEASUREMENT.

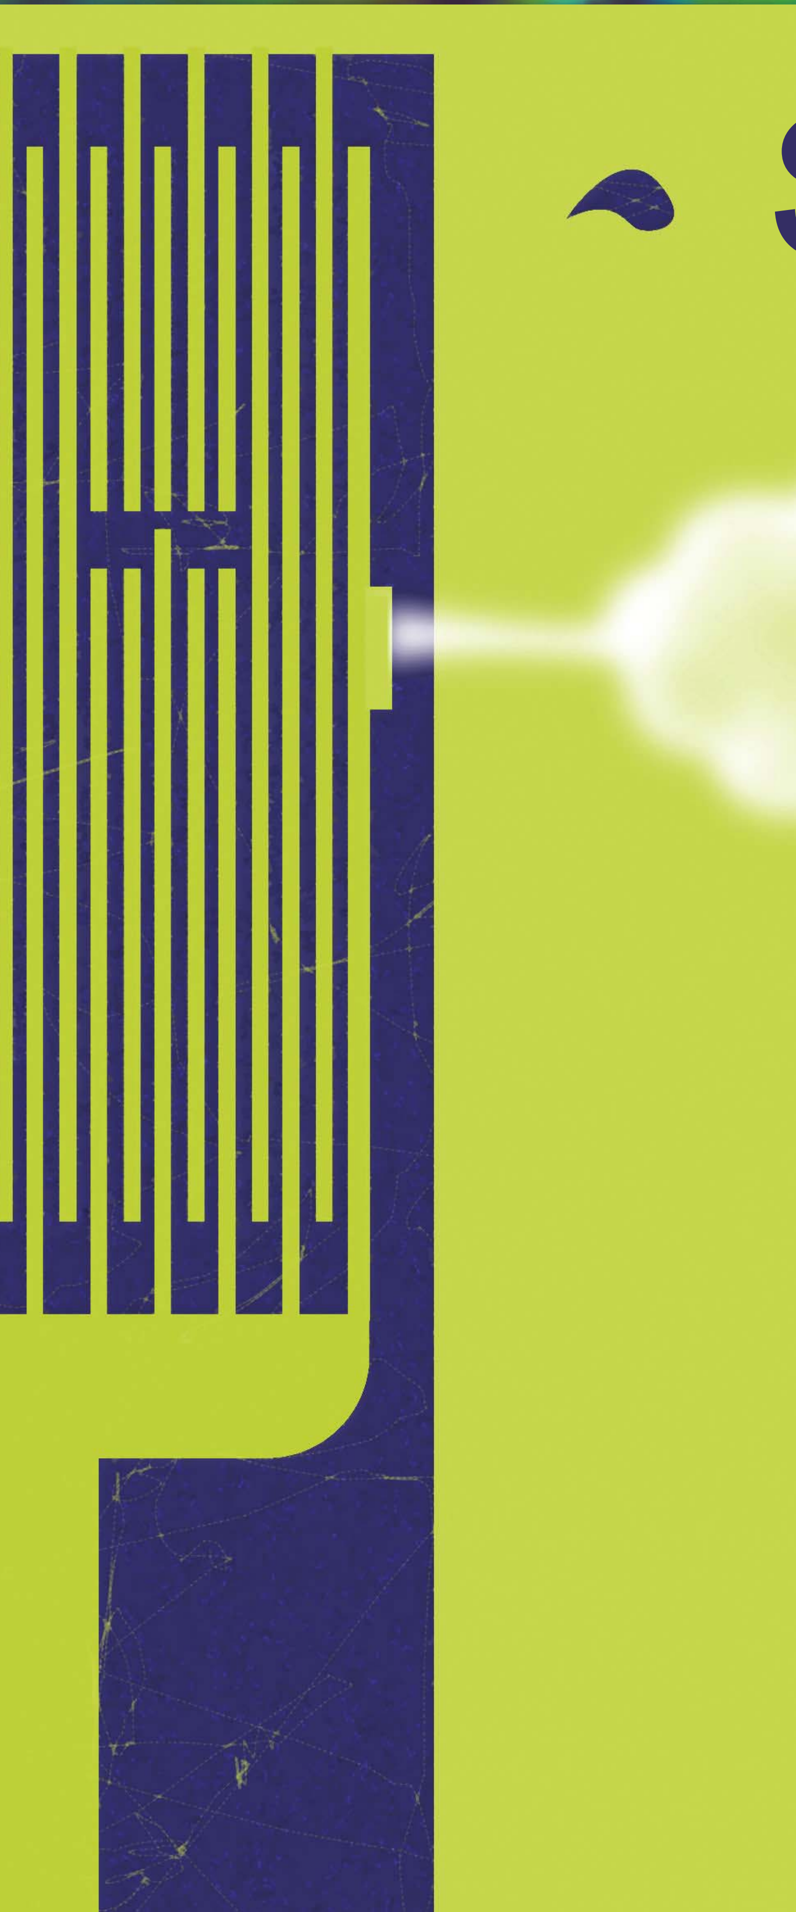

Strain gauges are the fundamental sensing elements for many types of sensors, including pressure sensors, load cells, torque sensors, and position sensors. Most strain gauges are foil types, available in a wide choice of shapes and sizes to suit a variety of applications (**Figure 1**). They consist of a pattern of resistive foil, which is mounted on a backing material. They operate on the principle that, as you subject the foil to stress, the resistance of the foil changes in a defined way. Foil gauges provide the ultimate in precision, but they also are expensive and provide small signals that are difficult to amplify. Strain gauges can also be silicon, in which metal is deposited in a thin-film semiconductor process on the silicon die. That die, often a MEMS (microelectromechanical system), forms the bendable diaphragm that responds to pressure changes. The same silicon die that carries the diaphragm structure can also have circuitry to amplify and linearize the output and compensate for any temperature effects.

By measuring strain, engineers can infer the stress of a material—an important factor because stress defines whether a part will bend or break. The stress may also represent a fluid pressure behind a diaphragm that is bending under a load. One interesting application for strain gauges involves bending: You measure the flexibility of a PCB (printed-circuit board) as a vacuum holds it in position on a test fixture. If you flex the board too far, the solder joints will break. By building a sample board with strain gauges, you can ensure that your electrical testing won't reduce the circuit's reliability, according to Swapnil Padhye, data-acquisition-product manager at National Instruments (**references 1** and **2**).

Strain gauges operate inside load cells to precisely measure force, monitor torque, or monitor pressure (**Figure 2**). They excel in measuring weight in scales, tanks, and vessels and for measuring the tension in films and strips in industrial processes. The gauges can infer the pressure in a pipe from the amount it swells, ensuring that the inside of the pipe is clean for applications such as food processes. Strain gauges in the load-bearing mounts of hoppers and bins find use in industrial processes. If you need to measure mass rather than weight, however, you must know your application's local gravity to make a precise conversion, according to Dave Cornwell, chief technology officer for Hardy Instruments. Strain gauges also find use in industrial, medical, and scientific equipment. The changes in measured strain may be slow or rapid, such as those of cyclical forces on an engine's connecting rod, which operates at tens of thousands of rotations per minute.

Mechanical engineers need strain gauges in the same way that electrical engineers need oscilloscope probes. Both groups must verify simulations, whether they are verifying finite-element mechanical models or Spice electrical models. Mechanical engineers can use strain gauges to collect real-world data on parts and structures they are designing. In addition, strain gauges are often permanent parts of designs, such as those monitoring the strain in a trestle bridge spanning a river.

These gauges are not the only way to measure strain. For example, you can instead make an epoxy-plastic model of a part, heat it, apply loads, let it cool, and

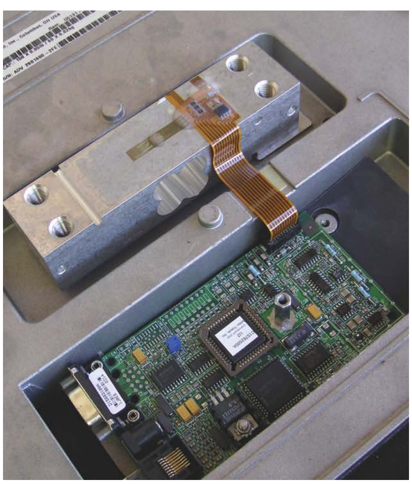

Figure 2 This Mettler Toledo scale has a precision load cell with a memory chip for linearization and calibration.

then illuminate it with polarized light. The light produces colored fringes that correspond to the strain on the plastic. Princeton Professor Robert Mark used this method to model the flying buttresses of Gothic cathedrals. This work shows why they have survived for centuries: The buttresses are in a state of pure compression—that is, compression everywhere and for all wind and snow loads. If there were any tension on these buttresses, they would fall apart because they are just stacked stones (**Reference 3** and **Figure 3**). Another approach to measuring strain involves the use of StressCoat, a brittle lacquer, which engineer Greer Ellis invented in 1942 while at strain-gauge manufacturer Magnaflux Corp (**Reference 4**). In this approach, you paint the part with the lacquer, apply design loads to the part, and observe the cracking in the coating. StressKote markets a similar product. Some engineers may dismiss this method, instead relying on computer simulations and FEA (finite-element-analysis) approaches. Real-world loads on real-world parts are far more reassuring, however, than pretty pictures on a computer screen.

Vishay developed another novel method, PhotoStress, which combines the intuitive visualization of StressCoat and the flexibility of polarized-light viewing. The method uses polarized light to illuminate a proprietary optical film. You contour the film to your part, apply the design loads, and illuminate the part with polarized light, letting you see the strain patterns in the part. Optical transducers on a polariscope also give quantitative measurements of strain, and companies such as Vishay provide liquid-photosensitive coatings for casting contourable sheets.

One of the greatest problems engineers have in applying strain gauges is the existence of so many uncontrolled variables. With voltmeters and light sensors, the manufacturers control most of the precision; you need only to connect the voltage probes to the circuit or to shine light on the sensor. With strain-gauge measurements, however, you must first select a gauge from hundreds or thousands of types, select a location for the device, prepare the surface, and bond the gauge to the part you are measuring. You make connections between the strain gauge and the measuring amplifier. Figuring out these processes is not the end of your troubles, however. You also need to ensure that you stay within the temperature range of the gauge, that you linearize the gauge output, and that you fully understand the relationship between stress and strain in the part you are measuring—a fundamental requirement.

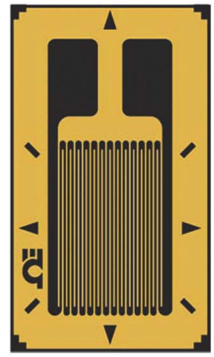

Figure 1 The most common type of strain gauge is made from a metal foil (courtesy Omega Engineering).

Another problem is the fact that some materials, such as fiberglass and carbon fiber, are anisotropic—having properties that differ according to the direction of the measurement. In these cases, the fibers are often oriented in a certain direction, and the relationship between stress and strain depends on the applied direction as well as the interaction of the directions inside the material. You can see that effect when you open or close window blinds. The blinds don't break during this process because there is little pressure on

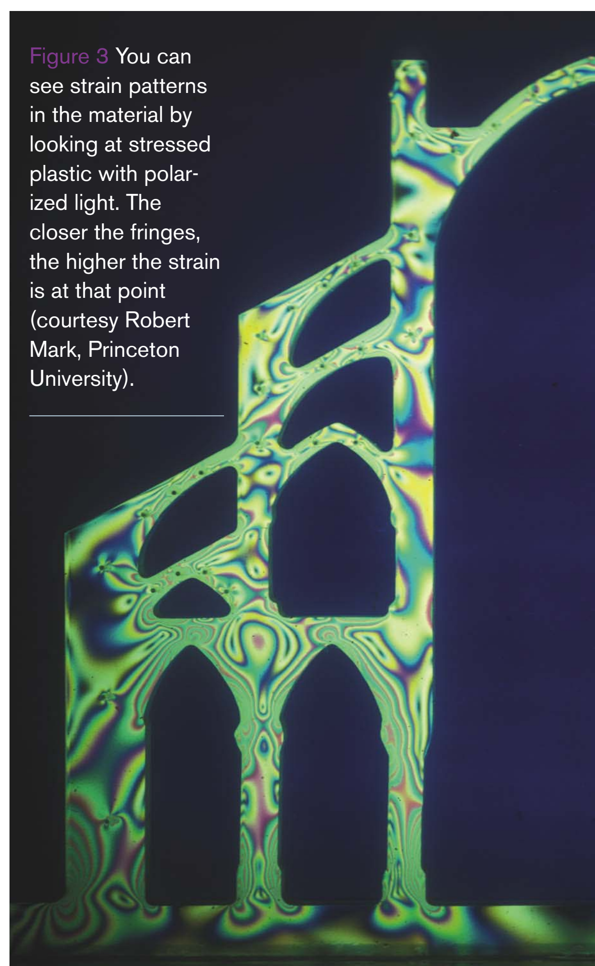

Figure 3 You can see strain patterns in the material by looking at stressed plastic with polarized light. The closer the fringes, the higher the strain is at that point (courtesy Robert Mark, Princeton University).

AT A GLANCE

- Strain gauges can be made of metal foil, silicon, or piezoresistive materials.
- Strain gauges provide a small change in resistance, so they find use in bridge configurations.
- The last step in manufacturing a gauge is when you glue it to your part.
- Watch out for residual stresses in parts. These stresses can cause catastrophic failure under light loads.
- Strain-gauge amplifiers are expensive because they are difficult to make. You should have a good reason for designing your own bridge-interface circuits.
- Taking good measurements may require weeks instead of hours.

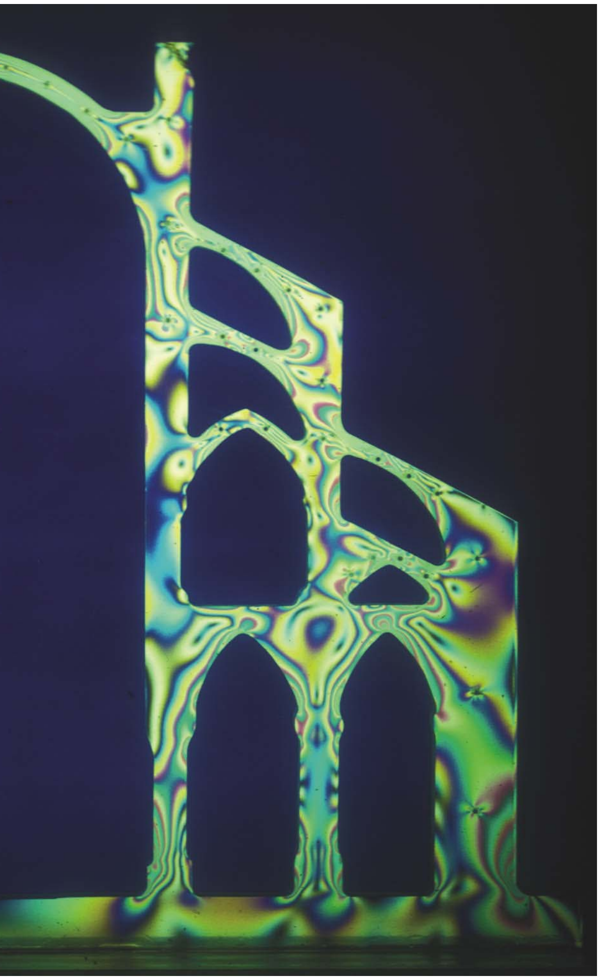

them. Bending those blinds over your knee presents a stiffer structure, which fails catastrophically if you subject it to any significant amount of strain. You can evaluate strain in anisotropic materials by using a "rosette" strain gauge, which allows you to simultaneously measure strain in two or three directions (**Figure 4**).

Another tricky, uncontrolled variable is the captive stress that exists in the part you are measuring. You may remember from your statics and dynamics courses that bridges do not connect at both ends. The mathematical calculations cannot solve for an overconstrained system. You face the same problem in taking strain measurements. Bolting the part into place and pulling it into alignment with a wrench causes a significant amount of stress and strain on the part. If you then glue a strain gauge to that part, the output will display zero strain, even though the part may be on the verge of breaking or may have already bent when you muscled it into position. "For making accurate measurements, the gauge is important, but the steel or aluminum it sits on is more important," says Hardy's Cornwell. To avoid linearity and hysteresis problems, Cornwell suggests using special alloys and heat-treating the part after machining to relieve local stresses from the machining operations. A strain gauge also averages the strain over its area. A hole near the gauge causes a stress concentration with large strains, but the gauge averages that concentration with the strain along the rest of its length and indicates a lower strain. "You need to select a gauge that is the appropriate length for the strain field or stress concentration you are looking at," says Tom Rummage, a senior application engineer at Vishay.

More subtle problems can also occur. An outside layer of a casting, for example, may harden first. Then, as the inside of the part solidifies, the part's cooling generates residual stresses. You have to realize that any casting, weldment, or machined part with surface stresses has static internal stresses far beyond what you can trivialize as second- or third-order effects. As always, you should experiment and collect data. "Stress cannot exist at a free boundary," says Rummage. "Put a strain gauge down that has three elements a certain distance away from where you are going to drill a hole. As you drill the hole, you create that free boundary. If it collapses in, it was under compression. If it pulls away, it was under tension. By knowing those three gauges and their angular relationship to one another, you can calculate the residual stress the part was experiencing." By verifying internal residual stresses in a part, you can then design-in a way to accommodate them. That approach might be heat-treating, using a different casting alloy, or taking a set of measurements to prove that no strain that approaches the elastic-limit or fatigue-failure guidelines exists anywhere in the part. Make sure that there are not multiple molds, different processes, or new vendors that could create variances in the internal stress. Even with these variances, however, the design may be perfectly acceptable because prestressed concrete depends on pre-existing internal stresses to meet design loads.

You must know the static and dynamic loads of what you are measuring. Select a gauge that works with the expected strain, but also consider shock loads and the effects of momentum and point-loading on the material. You also must make sure that static discharges will not damage the sensor electronics. "Everything normally is designed for CE [Conformité Européenne], which is the human-body model," says Hardy's Cornwell. "In a factory, you have the fork-lift model. When someone runs around in a fork lift, he gets a lot more voltage than CE stipulates." Cornwell explains that, once the operator lowers a pallet onto a platform scale, a giant arc jumps from the forks on the fork lift. If the operator doesn't use a ground strap, the only ground-return path is through the load cells and the strain-gauge wires. In addition, the strain gauge may be subject to fatigue failure if you strain it too many times over too large a range. Also remember that the modulus of elasticity may differ under compression and tension in your material. It is not a common problem, but it highlights the fact that a good strain-gauge engineer must know mechanics, materials, electronics, physics, and the theory of experiments.

Because your procedures and design can have a large effect on the validity of strain-gauge measurements, it is always a good idea to include the strain-gauge vendor's applications engineers in your plans. You may have concerns about using the vendor's strain gauge in your application. Some vendors, such as Omegadyne, can address those concerns: They apply the gauge for you, using all the expertise their engineers have accumulated over the years. Omegadyne can custom-design a gauge for you in as little as two weeks, according to William Hamilton, a design-and-manufacturing engineer at the company.

Don't underestimate the importance of strain-gauge measurements: Although you can slap a gauge on a part and have an answer in an hour, the answer will be wrong. In your rush, you might glue down the gauge with an epoxy that hardens in five minutes. This fast-drying epoxy not only releases heat but also heats or shrinks as it hardens. This condition places strain in the gauge, which then yields erroneous readings. Similarly, you can't just slap a foil gauge on a thick blob of epoxy because the distance between the foil gauge and the part's surface provides a substantial error. You must measure the strain of the part, not the strain on ⅛ in. of epoxy between the part and your gauge.

Rather than rushing a measurement, conduct a series of experiments that prove the validity of strain-gauge selection and mounting. "The biggest problem our customers have is making the right choice of strain gauge," says Rob Carney, OEM-sales manager at Omegadyne. "Accuracy, stability, temperature range, elongation, and test duration are all important factors." Once you mount the gauges, make sure that they always return to zero when you remove the strain, that there is no hysteresis, and that they provide good repeatability. You should correlate those measurements to a NIST (National Institute of Standards and Technology) standard and then take the measurement. The gauges themselves are often the smallest cost you face. The larger costs are in the mounting and characterization of the gauges, as well as the circuitry and test equipment you need to record the results. For this reason, buying a strain-gauge-conditioning system may be better than designing your own circuitry.

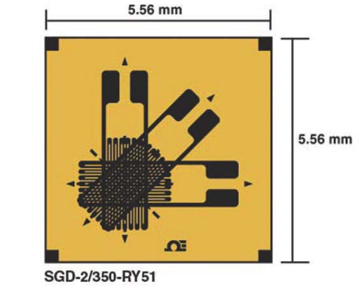

Figure 4 This rosette-type strain gauge has three gauges to measure strain in three directions at once (courtesy Omega Engineering).

STRAIN FUNDAMENTALS

To understand why strain-gauge instruments are sophisticated and costly, you must understand their fundamentals. A microstrain is a change of one-millionth in the resistance of a strain gauge, meaning that a bridge factor of two and an excitation of 1V yield 0.5 μV per microstrain. "Take a rope that is 15.8 miles long," says Vishay's Rummage. "If you pull that rope to a uniform strain of one microstrain, [the result] is 1 in. If you are not careful with your surface preparation, that one part in 1 million may not be achievable." Some inexperienced engineers and academic researchers try to use an ohmmeter for this measurement, but they soon realize that the data it yields is unusable. To make the change in strain become a large percentage change in the sensed measurement, experienced engineers incorporate the strain gauge into a Wheatstone bridge, a four-resistor device that nulls out errors and makes the change in strain a large percentage change in sensor output. Better yet, if you use four strain gauges in proper orientation, you get four times the signal amplitude and sensitivity. To infer the resistance change, you provide the bridge with an ac or a dc excitation voltage. The ac approach has certain advantages, such as nulling out the thermocouple effects of the lead-wire material you are soldering to the gauge's foil material. These thermocouple potentials do not change potential when the bridge excitation changes, so you can null out the dc error when you use synchronous demodulation to extract a dc value from the ac signal.

Due to improvements in operational amplifiers, a four-gauge full bridge is not always necessary to get a measurement. You can instead use a quarter-bridge configuration in which only one active strain gauge and three passive resistors complete the bridge. Alternatively, you can use a half-bridge with two gauges in one leg. This approach cancels out the temperature coefficient of the strain gauges. You then complete the bridge with two passive resistors that also share the same temperature coefficient. The bridge configuration ensures that the temperature coefficients of the gauges and passive resistors are ratiometric and cancel out. With some amplifiers, you need not even mount the second gauge to measure strain. Instead, you can just use a "dummy" gauge as a passive resistor, as long as it is at the same temperature as the active strain gauge.

Although canceling out the temperature coefficient of the strain-gauge ma-

Figure 5 This strain-gauge amplifier costs $1149 and provides 24-bit measurement accuracy, excitation, bridge completion, and a host of other features (courtesy National Instruments).

terial is important, you must also deal with a more basic temperature-compensation problem: The material that you are measuring has a thermal coefficient of expansion, meaning that, as you heat the material, it expands, and the strain gauge you glued to the specimen will expand along with the material, providing an erroneous output. Thus, your gauge produces an output even though no strain exists in the material. To avoid this problem, carefully select a gauge material that has the same temperature coefficient of expansion as the material that you are trying to measure. Gauge manufacturers make a large variety of both positive- and negative-expansion coefficients that allow you to null out this temperature problem. The thermal coefficient of expansion might differ in each direction in a carbon fiber or other nonisotropic material.

"A strain gauge gives a very small change in resistance, so the resulting voltages are very small," says David Potter, market-development manager at National Instruments. "Depending on where you are using them, the environment, and how long the wires are, the SNR [signal-to-noise ratio] can be pretty low." It might be a good idea to take data from the gauges without excitation. This approach gives you a measurement of the noise because the gauge

⊞ Go to www.edn.com/090528cs and click on Feedback Loop to post a comment on this article.

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

has no valid output without excitation. You can now appreciate why a strain-gauge amplifier is an expensive and specialized instrument (**Figure 5**). It has to complete, excite, and offset-null the bridge; measure small signals; reject noise; provide antialiasing filters for the analog-to-digital conversion; and buffer the output to reduce noise entering the signal path. Top-grade instruments also provide shunt-range calibration and remote sensing, in which a separate pair of leads ensures the precise control of the voltage at the bridge. If the instrument can excite an ac bridge, the instrument must also demodulate the excitation to give you a dc signal. A lab technician using a cell phone nearby can cause some unexplained sensor signals. "You can ground the wiring shield at the amplifier and leave it open at the part," says Vishay's Rummage. "This scheme is the most prevalent."

Rummage also advises providing grounding at the part and leaving the shield open at the instrument if your design has noise problems. You may also want to try grounding the cable shield at both ends. As with mounting and selecting the gauge, the measurement must also try to control all the variables. Once you select a gauge and connect it to the amplifier, you must control the measuring environment or at least conduct experiments that yield good data. Keep both the strain gauges and the amplifier within their specified temperature ranges.

Errors and nonlinearities can arise in myriad places and can degrade the accuracy of your measurement. Take frequent, controlled calibration runs. If you can verify the strain-gauge readings with an expensive load cell and a NIST-certified amplifier, you can take measurements over temperature, humidity, and any other conditions that might affect the material you are measuring. Once you have this data, you can correct the raw strain-gauge readings in software such as The MathWorks' Matlab or National Instruments' LabView. Do a calibration run both before and after you take the measurement to confirm that you have not damaged the gauge, glue, or part you are measuring. A calibration run cancels out most second- and third-order effects, such as the transverse sensitivity of the foil gauge. Transverse sensitivity arises because a foil gauge does

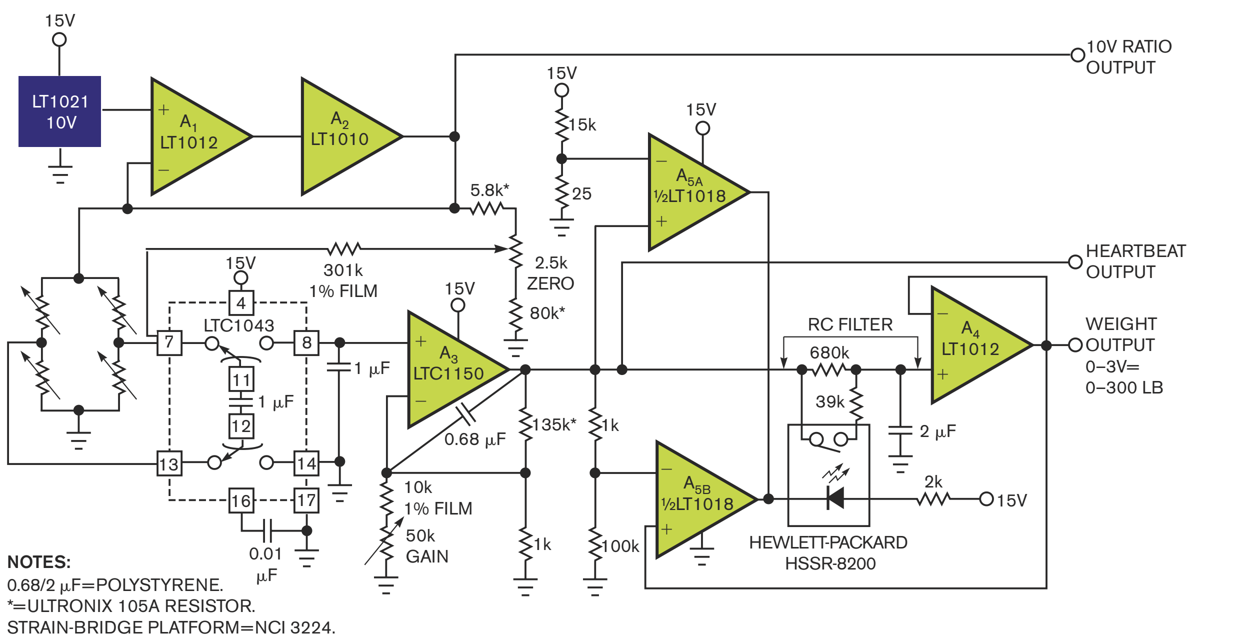

Figure 6 This strain-gauge circuit can sense the flow of blood in your body and weigh you to within 0.01-lb accuracy (courtesy Linear Technology).

not reject all strain in the orthogonal direction.

Both Linear Technology and Analog Devices have contributed literature for those engineers brave enough to design their own bridge-excitation and amplification circuits (**references 5** and **6** and **Figure 6**). If you are designing a low-cost product that requires a strain gauge, you may have to design your own circuitry. Engineers who must measure strain as part of a product-development cycle should rely on measuring experts, such as Omega, Vishay, and National Instruments. Remember: Test equipment and circuitry cannot make up for a botched gauge selection or installation. Select the proper gauge material and type. Then, decide whether to use a full-, half-, or quarter-bridge configuration. Select the right mounting place and epoxy. Watch out for those captive stresses that render your measurements meaningless. Then, make sure you get the strain-gauge signal to your amplifier. If possible, solder the instrument's lead wires to the gauge. Any connectors that are not gold-plated cause gross errors in the measurement. Use first-class test equipment and understand the design of a bridge-measurement circuit. Frequently calibrate your system and make sure that FEA matches your real-world-strain measurements.

Both electrical and mechanic engineers put an absurd amount of faith into computer simulations because they believe a computer cannot make mistakes. FEA engineers often adamantly believe that simulations are accurate—only to find that captive stress in the part or a mesh-selection error in simulation gives erroneous results. "A lot of people completely and abjectly trust a finite-element model to tell them where the direction and magnitudes of the strains are," states Vishay's Rummage. "Those are assumptions they have made, and they have to be validated." If you carefully take your measurements and understand all the aspects of strain measurement, you should stand your ground against simulations. EDN

FOR MORE INFORMATION

Analog Devices
www.analog.com

Hardy Instruments
www.hardyinstruments.com

IPC
www.ipc.org

JEDEC
www.jedec.org

Linear Technology
www.linear.com

Magnaflux
www.magnaflux.com

The MathWorks
www.mathworks.com

Mettler Toledo
www.mt.com

National Institute of Standards and Technology
www.nist.gov

National Instruments
www.ni.com

Omegadyne
www.omegadyne.com

Omega Engineering
www.omega.com/pressure/psc.html

Princeton University
www.princeton.edu/cee

StressKote
www.stresskote.com

Vishay
www.vishay.com/company/brands/micromeasurements

REFERENCES

1 IPC Association Connecting Electronics Industries/JEDEC (Joint Electron Device Council)-9704, "Printed Wiring Board Strain Gage Test Guideline," June 2005, www.ipc.org/TOC/IPC-9704.pdf.

2 Padhye, Swapnil, "Using IPC/JEDEC-9704 & 9702 Standards for Strain Gage Testing of Your Printed Wiring Boards," National Instruments, ftp://ftp.ni.com/evaluation/daq/ekit/pcb_strain_gage_testing.pdf.

3 Mark, Robert, *Experiments in Gothic Structure*, MIT Press, April 1982, ISBN-10: 0-262-13170-6, ISBN-13: 978-0-262-13170-4, http://mitpress.mit.edu/catalog/item/default.asp?tid=9531&ttype=2.

4 Ellis, Greer, "Method of Determining Strain Concentration in Rigid Articles," US Patent Office, Sept 8, 1942, www.google.com/patents?id=-MtgAAAAEBAJ&dq=Greer+Ellis.

5 Williams, Jim, "Bridge Circuits: Marrying Gain and Balance," Linear Technology Corp, Application Note 43, June 1990, www.linear.com/pc/downloadDocument.do?navId=H0,C1,C1154,C1009,C1026,P1213,D4134.

6 Kester, Walt, "Practical design techniques for sensor signal conditioning," Analog Devices Corp, 1999, www.analog.com/static/imported-files/seminars_webcasts/49470200sscsect2.pdf.

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

More measurements. More data. Less cost.

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

Agilent 34970A switch/measure unit

Agilent 34980A switch/measure unit

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

www.agilent.com/find/daqswitch

Agilent Authorized Distributor

METRICTEST™
6,000 instruments. One source.

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

Agilent Technologies

BY JOHN NIETO • INPHI CORP

Evolving to DDR3 technology

THE MOVE FROM DDR2 TO DDR3 REQUIRES NEW DESIGN TECHNIQUES TO IMPROVE SIGNAL INTEGRITY AND GAIN THE MAXIMUM BENEFIT FROM THE LATEST MEMORY-INTERFACE TECHNOLOGY.

Applications demanding higher system bandwidth and lower power, such as converged notebooks, desktop PCs, and servers, continue to drive the evolution of industry standards, including DDR3 (double-data-rate 3), as the JEDEC (Joint Electron Device Engineering Council) Solid State Technology Association defines it. The latest DDR3-memory standard, JEDEC JESD79-3A, supports these needs and the requirements of emerging dual-core- and multicore-processor systems. DDR3, the latest DDR-memory-interface technology, differs from the well-established DDR2 standard in several areas, including data rate, operating voltage, and logic (**Table 1**).

DDR2-memory-interface technology addresses and supports current system requirements. As memory requirements and technology continue to evolve, however, the standards must adapt to enable approaches that deliver even more functions. DDR3 offers significant advantages over previous DDR generations. It supports data rates as high as 1600 Mbps per pin with an operating voltage of 1.5V, a 17% reduction from the previous generation of DDR2, which operates at 1.8V. DDR3's built-in power-conservation features, such as partial refresh, can be important in mobile-system applications in which battery power is not necessary just to refresh a portion of the DRAM not in active use. DDR3 also has a specification for an optional thermal sensor that could allow mobile-system engineers to save further power by providing minimum refresh cycles when the system is not in high-performance mode.

DDR3 uses eight internal banks compared with DDR2's four to further speed systems by allowing advance prefetch, which reduces access latency. This speed should become more apparent as the size of the DRAM increases in the future. The I/Os for DDR3 use the JEDEC standard SSTL (stub-series-terminated logic) 15, which employs 1.5V logic, whereas DDR2 uses

TABLE 1 DDR3 VERSUS DDR2

	DDR2	DDR3/DDR3L
Rated speed (Mbps)	400 to 800	800 to 1600
Drain-to-drain voltage/supply voltage (V)	1.8±0.1	1.5±0.075 (DDR3)/ 1.35+0.1/−0.7 (DDR3L)
Internal banks	Four	Eight
Termination	Limited	All DQ signals
Topology	Conventional T	Fly-by
Driver	OCD calibration	Self-calibration with ZQ
Thermal	No	Optional

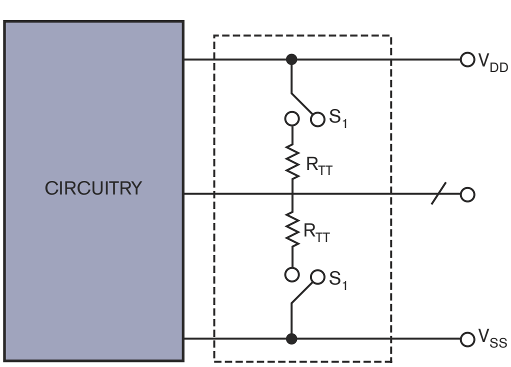

Figure 1 You can switch the on-die termination and select the effective value through a pullup/pulldown-resistor network.

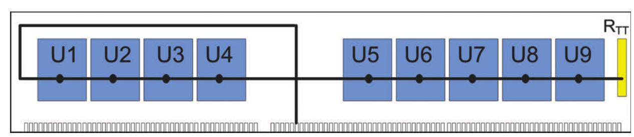

Figure 2 The DDR3 fly-by topology finds use in command/address and clock signals to improve signal integrity.

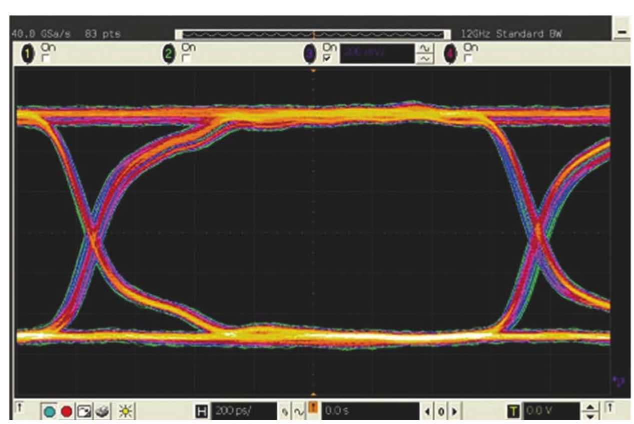

Figure 3 With the optimal termination value, driver impedance, and output inversion enabled, the postregister DDR3 data-eye plot exhibits a wider opening with minimal ring-back.

JEDEC standard SSTL18, which uses 1.8V logic. The DDR3 architecture fully uses ODT (on-die termination), ZQ (zero-quotient) calibration, and a fly-by topology for improved signal integrity. With such a high demand for lower-power memory and cost-saving technologies, JEDEC is now defining a DDR3L (DDR3-low-voltage) node.

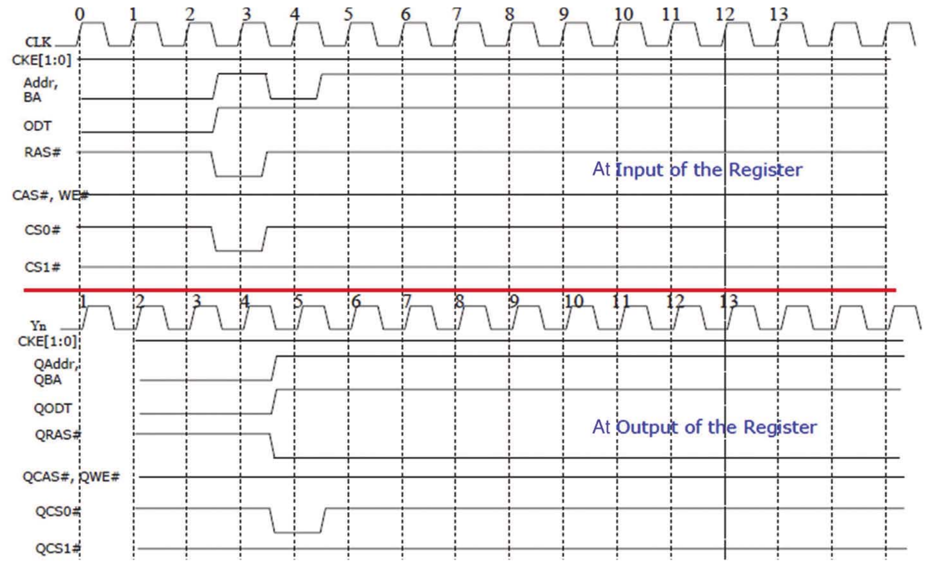

Figure 4 Normal operation inverts half of the address bus on the postregister to prevent simultaneous-switching-output noise.

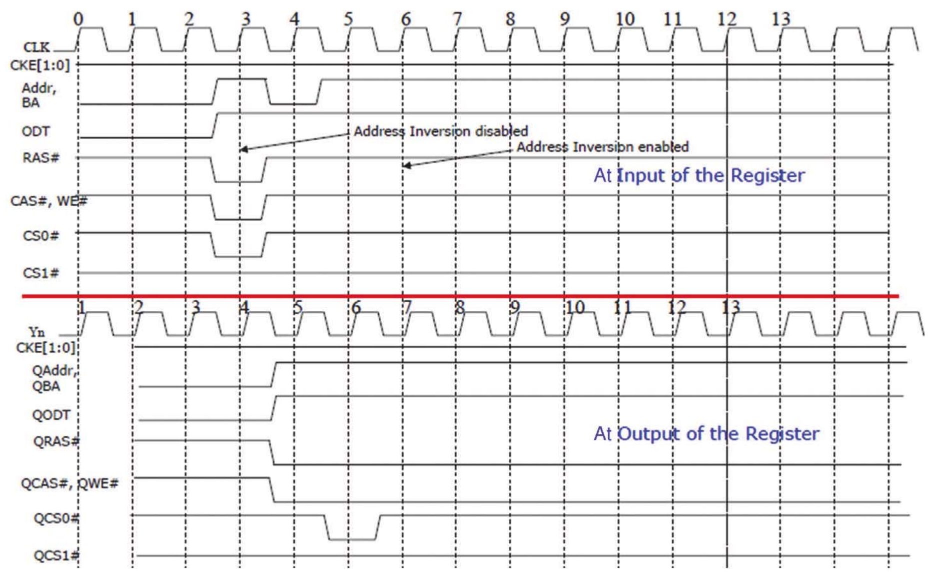

Figure 5 During an MRS command, the timing through the register automatically switches to driving the data for three clock cycles to account for the simultaneous-switching-output effect.

IMPROVING SI

Because DDR3 runs at higher memory speeds, the signal integrity of signals traveling through the memory module becomes more important. DDR3 uses a fly-by topology instead of the T branches that DDR2-module designs use. Thus, the address and control lines are in a single path chaining from one DRAM to another, whereas DDR2's T topology branches on the modules. Fly-by topology eliminates the mechanical line-balancing requirement and uses an automatic signal-time delay that the controller generates during memory-system training. Each DDR3 DRAM chip has an automatic leveling circuit for calibration and to memorize the calibration data.

DDR3 implements several impedance-calibration sequences to improve signal integrity. It uses long-ZQ calibration after power-up and periodically uses short-ZQ calibration during normal operation to compensate for voltage and temperature drift. These calibration sequences vastly improve the connectivity between the output driver of the DRAM and the PCB (printed-circuit-board) trace. A ZQ pin on the DRAM connects to an external precision resistor that adjusts the output-driver impedance as well as the ODT to match the trace impedance, thus reducing impedance discontinuity and minimizing reflection on the signals. The use of external precision resistors reduces the effect of variations due to process, voltage, and temperature and maintains a tight tolerance for better-controlled impedance values. DDR2, on the other hand, employs on-chip resistors, which can exhibit larger variations. The system and DRAMs also use dynamic ODT for improved signaling, especially for higher speeds.

You can switch the ODT on and off at the DRAM and select the effective value through a network of pullup and pulldown resistors (**Figure 1**). The dynamic nature of the circuit provides optimal command-, address-, control-, data-, and strobe-bus termination for better signal control and improved margins than those that DDR2 achieves using static termination on the motherboard. To further improve signal integrity, DDR3 also uses its fly-by topology for command/address and clock signals (**Figure 2**). It routes the signals to the DRAMs linearly and to the edge of the card at its bus termination. For registered DIMMs (dual-inline-memory modules), an IC component buffers the command, address, control, and clocks. This approach helps reduce the number of stubs and stub lengths that normally would be in DDR2's T topology; however, this approach also introduces flight-time skew between the clocks and data strobes at the DRAMs. You can compensate for this flight-time skew from the controller side on the motherboard by performing a leveling technique for deskew, which puts the DRAMs through a training sequence for tuning the DRAM clock.

The internal core speed of the DRAM basically remains unchanged in the transition from DDR2 to DDR3. DDR2 currently has a maximum bandwidth of 800 Mbps per pin but can extend to 1066 Mbps. To meet bandwidths as high as 1600 Mbps, DDR3 uses an 8-bit prefetch rather than the 4-bit prefetch that DDR2 uses. As a result, for every read or write operation, the technology accesses 8 bits in the DRAM core.

Registered DIMMs from previous generations, DDR and DDR2, exhibit excellent performance for systems that require higher bandwidths and better throughput efficiency. DDR2 employs at least one register and a PLL instead of two separate components; DDR3 employs a single monolithic-IC chip, which integrates the register and PLL. This integrated part features programmable-drive strength, input-bus termination, and output inversion, effectively reducing SSO (simultaneous-switching-output) noise. For power savings, the device includes output inversion, the ability to float the outputs, and the ability to power down the chips using input logic states. Another power-saving technique comes from programming the output drivers' impedance. Enhanced termination techniques and impedance matching vastly improve the signal integrity of the data eye using an integrated device, which in turn increases the margin in the timing budget.

Figure 6 The very-low-profile DIMMs are essentially the same as the standard DIMMs but at 60% of the height.

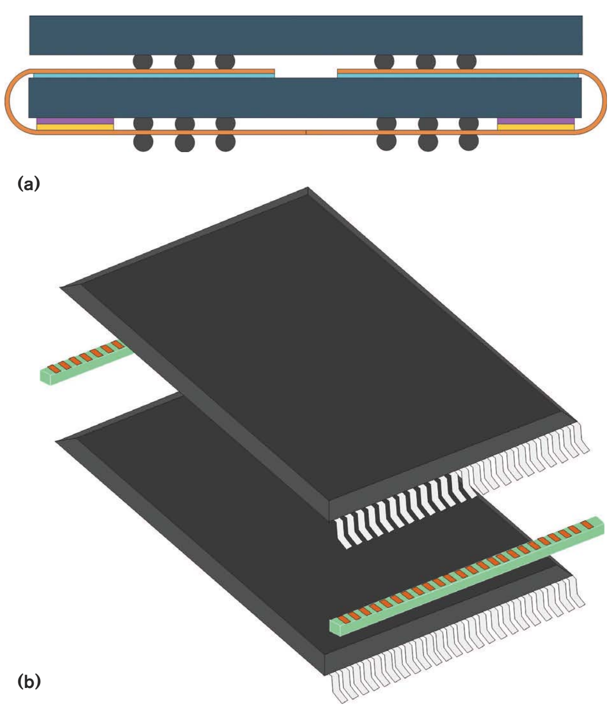

Figure 7 You can use a flex circuit and a PCB interposer to stack BGA (a) or TSOP (b) DRAM.

Using the optimal termination value, driver impedance, and output inversion, the postregister data eye exhibits a wider opening with minimal ring-back (**Figure 3**). The preregister also exhibits similar data-eye quality if it receives sufficient setup-and-hold time from the controller to prevent any bit errors. The register-PLL component uses on-die termination to maximize throughput efficiency and achieve higher speeds without redesigning or doing any extensive modifications to the system.

PROGRAMMING CAPABILITY

The memory controller sends commands to the memory that force it into a mode of operation. The states of CS (chip-select), RAS (row-address select), CAS (column-address select), and WE (write enable) define the commands. The MRS (mode-register-set) command is one of the first in system initialization. It allows programming of the configuration registers in the DRAM for various functions, features, and modes, making the memory module flexible for applications. The configuration registers are at their default values at power-up and programmed accordingly for proper operation.

The MRS command is present in both DDR2- and DDR3-DIMM applications. In DDR3-registered DIMMs, however, the register-IC chip, or register, which you can also program, can decode the command and switch to a different timing mode by driving the received data for three clock cycles instead of the conventional one. In normal operation, the system intentionally inverts half of the address bus on the postregister to prevent SSO noise (**Figure 4**).

An MRS command disables address inversion to allow correct access to the DRAMs, consequently increasing the propagation delay due to SSO noise. To account for SSO effects, the timing through the register automatically switches to driving the data for three clock cycles (**Figure 5**). This mode was not present in DDR2-registered-DIMM applications because the registers for DDR2 are strictly buffers without any programmability and decoding logic. You can program the register and the DRAMs for various functions, features, and modes. The JEDEC DDR3 register-specification document provides register mapping.

INCREASING MEMORY CAPACITY

With data centers needing more capacity, space is at a premium. Few options are available for maximizing space. One is the usage of more blade servers, which would require internal pe-

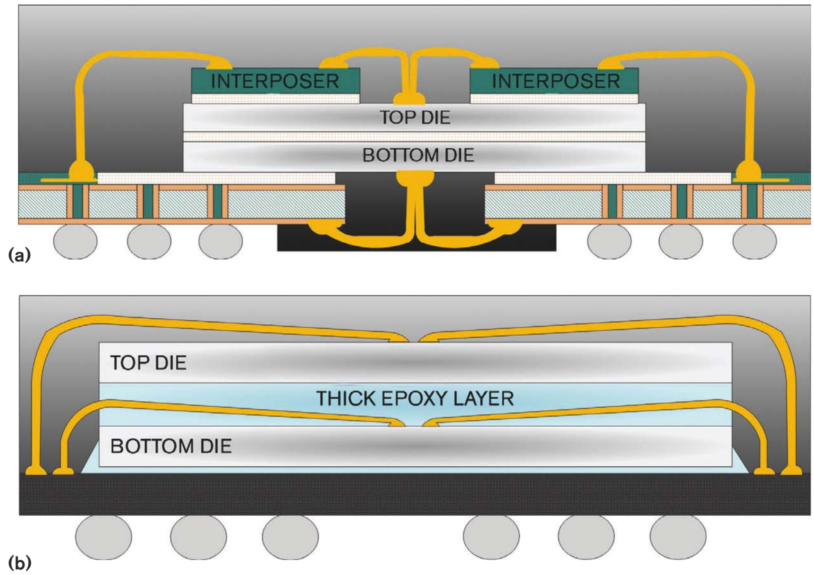

Figure 8 One version of die-stack-packaging technology has an interposer (a); another version does not (b).

ripherals to change to perhaps a smaller form factor, such as hard disks, memory, or I/O cards. JEDEC in 2005 established a VLP (very-low-profile) standard for smaller-form-factor DIMMs. As blade servers become more common, so will the VLP DIMMs, and the technology promises to become more of a presence as users begin to adopt DDR3-memory technology. The VLP DIMMs are essentially the same as the standard DIMMs but at 60% of the height (**Figure 6**). The height reduction improves airflow in the system, thus improving cooling. Other applications needing VLP DIMMs are ATCA (Advanced Telecommunications Computing Architecture), Micro ATCA for the telecommunications market, and embedded single-board computers. Other small form factors for DIMMs are SO-DIMM (small-outline DIMM) and mini-DIMM.

Standard and VLP DIMMs have various densities and ranks. DIMMs are typically 64- or 72-bit-wide words. A *rank* is an identical arrangement of memory banks on a module. The following **equation** shows how to determine the DIMM's total memory density: DRAM density/8×bus width/DRAM width× no. of ranks=total density. The configuration and density of the component de-

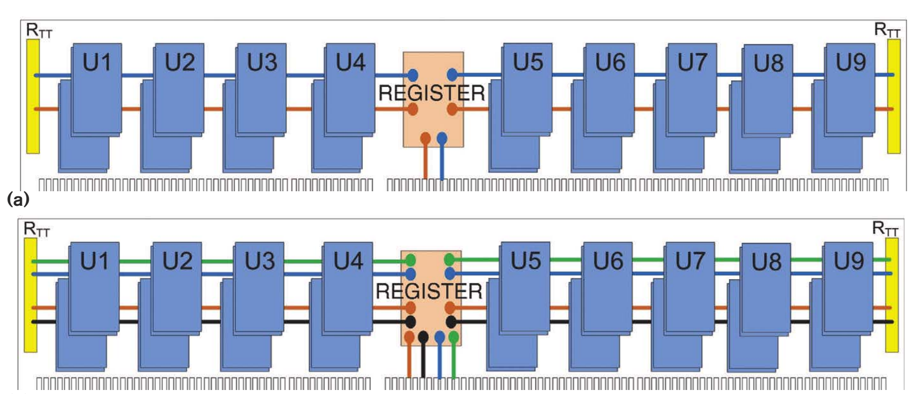

Figure 9 The dual-rank module requires only two chip-select signals (a), whereas quad-rank-registered DIMMs with stacked dual-die DRAM require four (b).

termine the number of ranks on a module. For example, two ranks of 512-Mbit DRAMs with 4-bit-wide data on a module with a 64-bit-wide bus would be 512 Mbits/8×64/4×2=2 Gbytes.

The most common ranks are single and double, but the total number of possible ranks is four. The quad-rank module is relatively new to DDR DIMMS, and until recently, memory controllers did not support it. Due to the ever-increasing demand for higher-density DIMMs, however, today's CPUs support quad-rank modules. In a relatively mature DDR2-memory market, JEDEC has defined and standardized only two quad-rank DIMMs, and two more are in development. In contrast, in an emerging DDR3-memory market, six quad-rank modules are in development.

A JEDEC-standard memory module in a planarity configuration has a maximum of 36 DRAM chips. DRAM and module vendors have developed several stacking technologies to further increase memory capacity on any memory module, however. One stacking technology is the use of a PCB or flex-circuit interposer to stack BGA (ball-grid-array) or TSOP (thin small-outline package) DRAM (**Figure 7**). For improved manufacturing reliability, die-stack packages are other alternatives to stacking DRAMs. This technology allows you to place two or four die in one chip package. A couple of die-stack-packaging technologies use a window chip-scaled package with an interposer (**Figure 8**). In this approach, both die face in opposite directions, or you wire-bond both die in the same direction. In theory, a quad-die chip would maximize the number of DRAM die from 36 to 144 on a DIMM. However, due to electrical limitations, the maximum number achievable is now 72.

Two registered DIMMs implement dual- and quad-rank modules with dual-die stacked DRAMs (**Figure 9**). The dual-rank module requires only two chip-select signals, whereas the quad-rank module requires four, one for each rank. The implementation of quad-die-stacked DRAMs on modules will

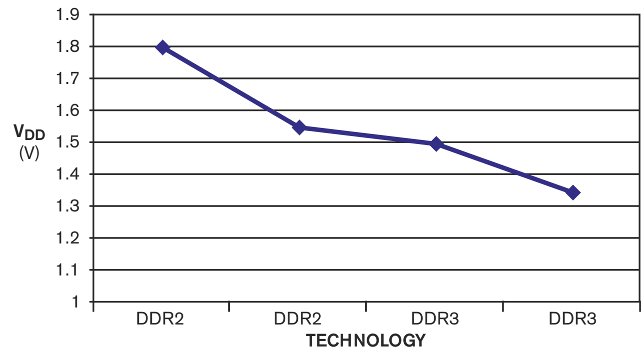

Figure 10 DDR3 uses 25% less power per DIMM than does DDR2.

be commonplace as higher-density modules and smaller-form-factor modules, such as VLP-registered DIMMs, SO-DIMMs, and mini-DIMMs, become more mainstream.

LOW-POWER DDR TECHNOLOGY

As data centers and server farms increase their capacity to meet higher bandwidth demand and increasing DDR-memory-interface data rates—and as companies grapple with rising energy costs—overall power consumption becomes more of a concern. The importance of energy saving is beginning to take precedence. The DDR market is currently addressing the need for low-power, cost-saving approaches. The DDR2 technology operates at 1.8V, which is approximately a 17% increase from DDR3, but companies are pushing for reducing that node to 1.55V. A similar trend is beginning to develop in DDR3 in which the original standard operating voltage is 1.5V but is decreasing to the 1.35V operating-voltage node, which is another 10% reduction.

For a list of references, go to www.edn.com/090528ms4317.

Go to www.edn.com/ms4317 and click on Feedback Loop to post a comment on this article.

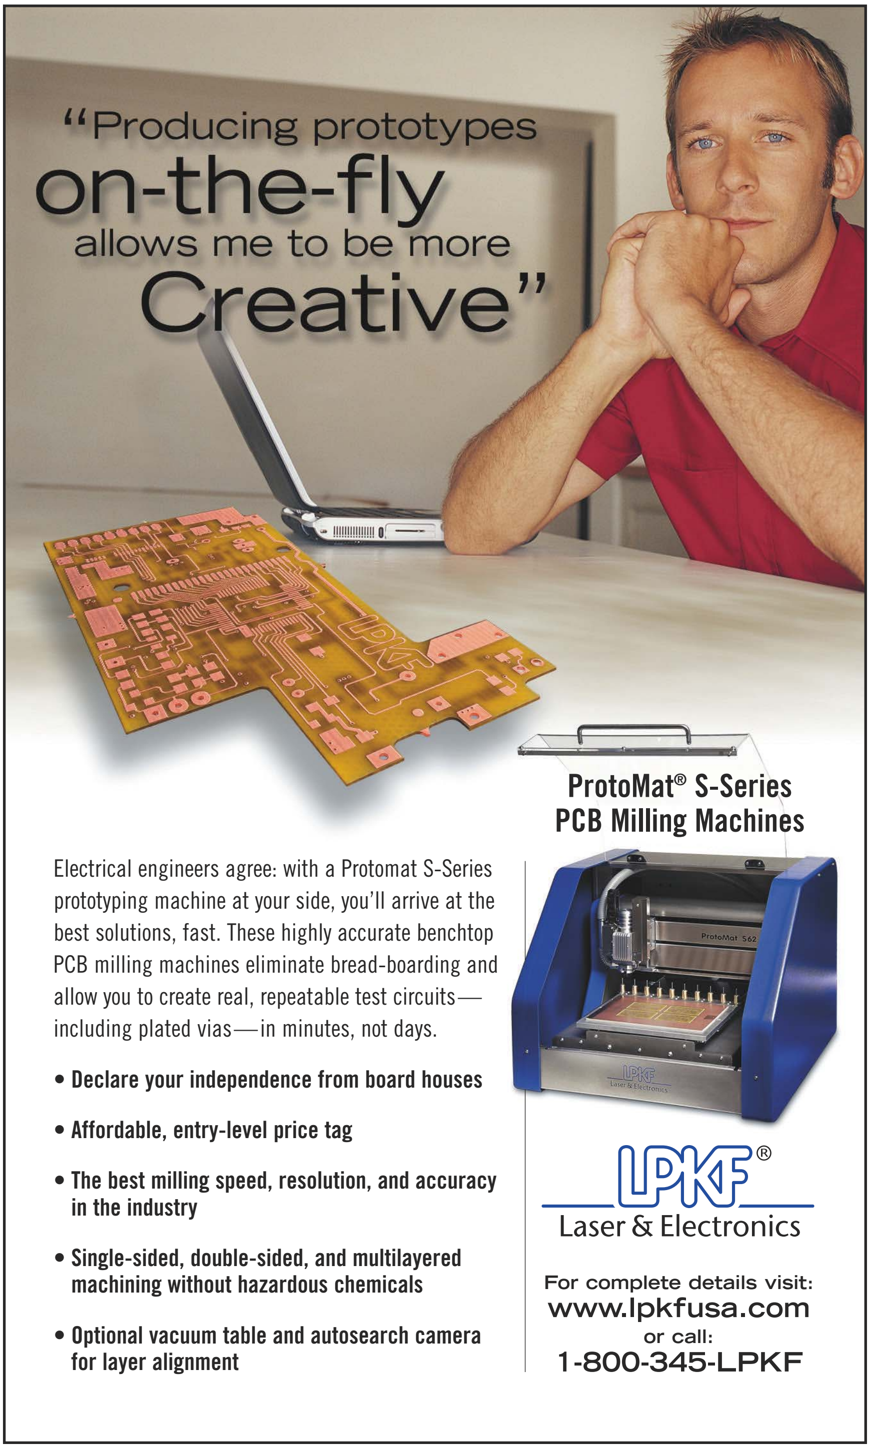

Figure 10 shows a linear downward trend of the operating voltage for the DDR technology. Going from 1.8 to 1.35V yields a 25% reduction in power per DIMM alone, with some servers supporting as many as 18 DIMMs. This reduction is significant not only in the core of the memory and logic chips of the DIMMs, but also in the I/Os of both, which consume most of the power. As the operating voltage scales down, so will the memory-interface technology. Manufacturers are studying other areas of memory-interface technology for power reduction. These areas include implementing special power-down features and modes when the system is idling.

From mobile-system applications to notebooks to enterprise servers, a constant demand exists for higher bandwidth, lower power, and improved throughput efficiency. The industry is examining all aspects of the DDR-memory-interface technology to meet the needs and demands. The onus is on the DDR-chip vendors to adapt and provide IC chips for DDR DRAMs, registers, and PLLs with improved signal performance. JEDEC has defined innovative methods to improve signal integrity from one generation of DDR to the next to facilitate memory-driven-application requirements. It is clear that DDR3 memory is addressing the need for higher-density modules with lower power and more efficient throughput.EDN

AUTHOR'S BIOGRAPHY

John Nieto is a senior applications engineer at Inphi Corp, where he is responsible for DDR2 and DDR3 registers and PLL support for DIMM-application simulations. He has a bachelor's degree in electrical engineering from the University of California—San Diego.

Breaking Through Performance Barriers

Speedster™ is the World's Fastest FPGA

New Speedster 1.5 GHz FPGAs

Breaking through the speed barriers of traditional FPGAs, fully reprogrammable Speedster devices from Achronix Semiconductor are capable of operating at 1.5 GHz peak performance.

Speedster's unique 40 lanes of 10.3 Gbps embedded SerDes and four independent 1066 Mbps DDR2/DDR3 controllers enable extremely high I/O throughput to match the device's outstanding internal performance. At 3X faster, Speedster opens up new worlds of application design previously unavailable to engineers using traditional FPGAs.

Familiar Silicon. Familiar Tools. Fast Time-to-Market. Speedster uses familiar LUT-based fabric and standard synthesis and simulation tools, so designers can use standard RTL.

High-Performance Applications

- Networking
- Telecommunications
- Encryption
- High-performance computing
- Digital signal processing
- Test and measurement
- Military and aerospace

See what's possible:
www.achronix.com/speedster

achronix
SEMICONDUCTOR CORPORATION

Current-sense monitor and MOSFET boost output current

Gyula Diószegi and János Nagy, Divelex Ltd, Budapest, Hungary

A previous Design Idea describes a programmable current source that used a three-terminal National Semiconductor (www.national.com) LM317 adjustable regulator (**Reference 1**). Although that circuit lets you program the output current, the load current flowed through the BCD (binary-coded-decimal) switches. However, you may find it difficult to purchase BCD switches that can handle more than 25 mA, limiting the circuit's output current. By applying the simple, four-pin Zetex (www.zetex.com) ZXCT1010 current-sense-monitor chip, you can boost current because it doesn't flow through BCD switches (**Figure 1**). The load current results in a voltage on the sense resistor R_{SENSE}. The voltage on R_1, the 100Ω resistor, is the same as that on R_{SENSE}, generating an output current on R_1: $I_{OUT} \times 100 = I_{LOAD} \times R_{SENSE}$, and $V_{OUT} = I_{OUT} \times R_{OUT}$, where I_{OUT} is the output current, I_{LOAD} is the load current, and V_{OUT} is the output voltage. You can apply the output voltage as a control voltage to regulate the load current.

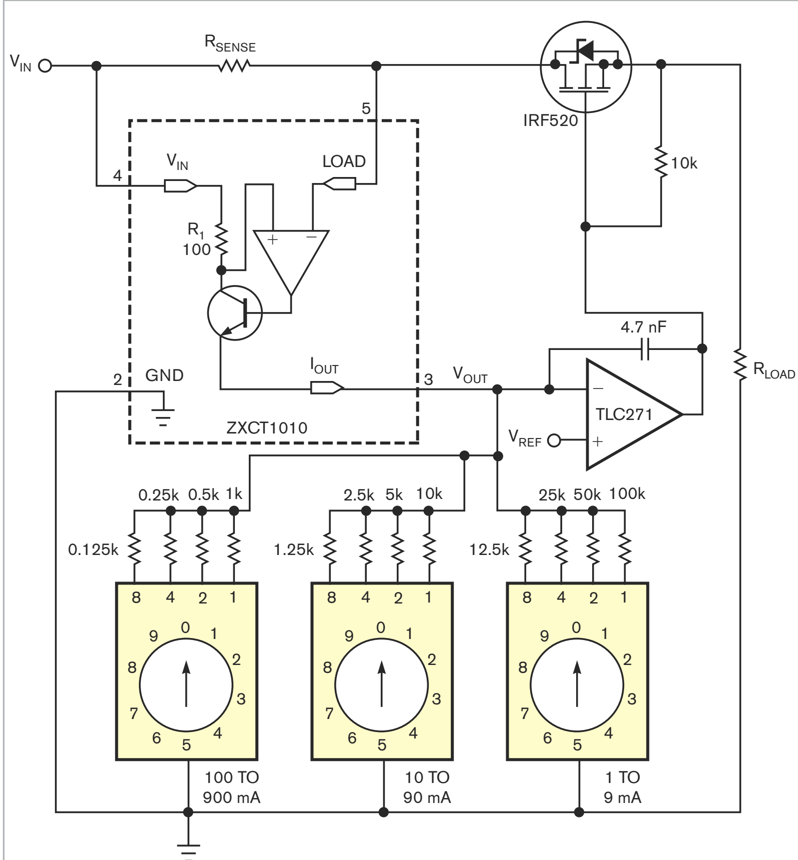

Figure 1 Passing current through a MOSFET and regulating it with a current-sense monitor bypasses the BCD switches, letting you increase load current.

One application for this circuit would be to refill accumulators in portable devices. In this case, the circuit works at 18V. The Fairchild Semiconductor (www.fairchildsemi.com) IRF520 is an N-channel, power-MOSFET chip in an aluminum heat sink with as much as 9.2A current and 0.27Ω drain-to-source resistance to connect the load current. An op amp controls the IRF520 in the feedback of the load current. In this application, the maximum output current is 1A, and the value of the sense resistor is 0.1Ω. The PCB (printed-circuit board) can also have this small resistance value, which you calculate using

DESIGN NOTES

Ultralow Power Boost Converters Require Only 8.5µA of Standby Quiescent Current

Design Note 465

Xiaohua Su

Introduction

Industrial remote monitoring systems and keep-alive circuits spend most of their time in standby mode. Many of these systems also depend on battery power, so power supply efficiency in standby state is very important to maximize battery life. The LT®8410/-1 high efficiency boost converter is ideal for these systems, requiring only 8.5µA of quiescent current in standby mode. The device integrates high value (12.4M/0.4M) output feedback resistors, significantly reducing input current when the output is in regulation with no load. Other features include an integrated 40V switch and Schottky diode, output disconnect with current limit, built in soft-start, overvoltage protection and a wide input range, all in a tiny 8-pin 2mm × 2mm DFN package.

Application Example

Figure 1 details the LT8410 boost converter generating a 16V output from a 2.5V-to-16V input source. The LT8410/-1 controls power delivery by varying both the peak inductor current and switch off time. This control scheme results in low output voltage ripple as well as high efficiency over a wide load range. Figures 2 and 3 show efficiency and output peak-to-peak ripple for Figure 1's circuit. Output ripple voltage is less than 10mV despite the circuit's small (0.1µF) output capacitor.

The soft-start feature is implemented by connecting an external capacitor to the V_{REF} pin. If soft-start is not needed, the capacitor can be removed. Output voltage is set by a resistor divider from the V_{REF} pin to ground with the center tap connected to the FBP pin, as shown in Figure 1. The FBP pin can also be biased directly by an external reference.

The $\overline{SHDN}$ pin of the LT8410/-1 can serve as an on/off switch or as an undervoltage lockout via a simple resistor divider from V_{CC} to ground.

Ultralow Quiescent Current Boost Converter with Output Disconnect

Low quiescent current in standby mode and high value integrated feedback resistors allow the LT8410/-1 to regulate a 16V output at no load from a 3.6V input with about 30µA of average input current. Figures 4, 5 and 6 show typical quiescent and input currents in regulation with no load.

The device also integrates an output disconnect PMOS, which blocks the output load from the input during shutdown. The maximum current through the PMOS is limited by circuitry inside the chip, allowing it to survive output shorts.

Compatible with High Impedance Batteries

A power source with high internal impedance, such as a coin cell battery, may show normal output on a voltmeter, but its voltage can collapse under heavy current demands. This makes it incompatible with high current DC/DC converters. With very low switch current limits (25mA for the LT8410 and 8mA for the LT8410-1), the LT8410/-1 can operate very efficiently from high impedance sources without causing inrush current problems. This feature also helps preserve battery life.

Conclusion

The LT8410/-1 is a smart choice for applications which require low standby quiescent current and/or require low input current, and is especially suited for power supplies with high impedance sources. The ultralow quiescent current and high value integrated feedback resistors keep average input current very low, significantly extending battery operating time. The LT8410/-1 is packed with features without compromising performance or ease of use and is available in a tiny 8-pin 2mm × 2mm package.

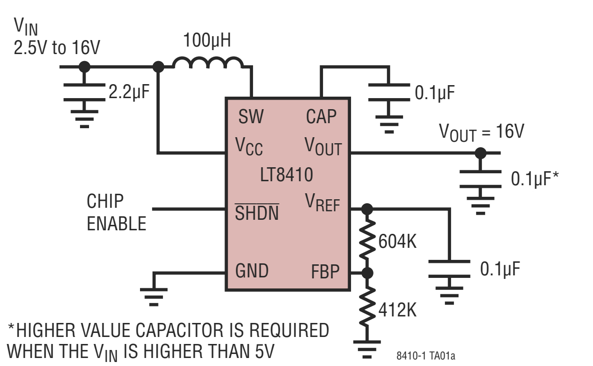

Figure 1. 2.5V-16V To 16V Boost Converter

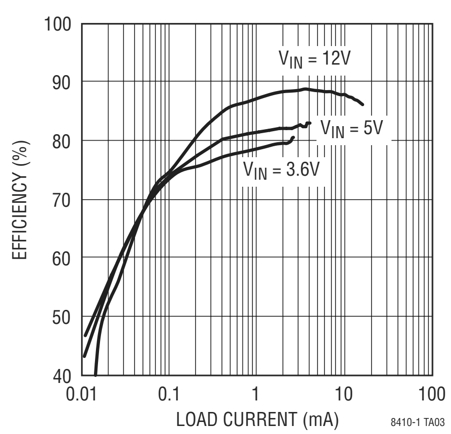

Figure 2. Efficiency vs Load Current For Figure 1 Converter

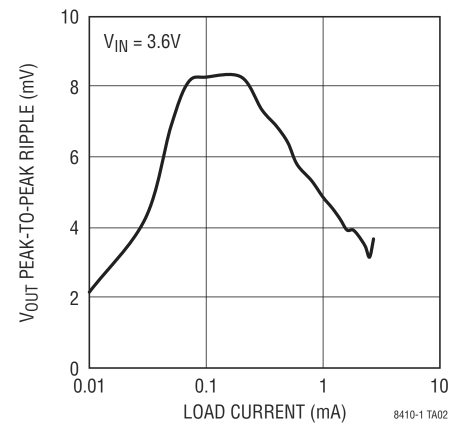

Figure 3. Output Peak-to-Peak Ripple vs Load Current for Figure 1 Converter at 3.6V

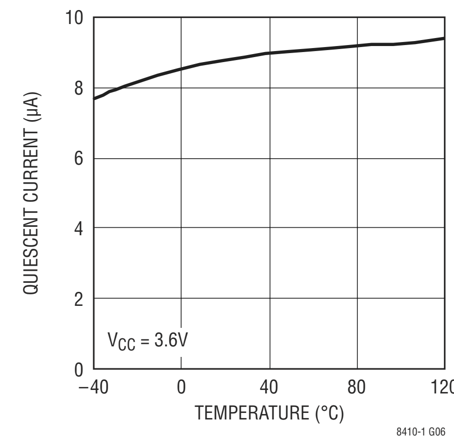

Figure 4. Quiescent Current vs Temperature (Not Switching)

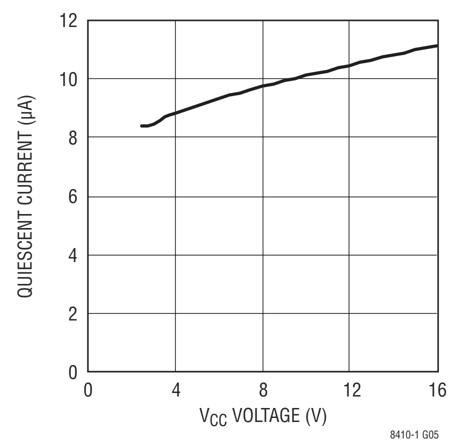

Figure 5. Quiescent Current vs V_{CC} Voltage (Not Switching)

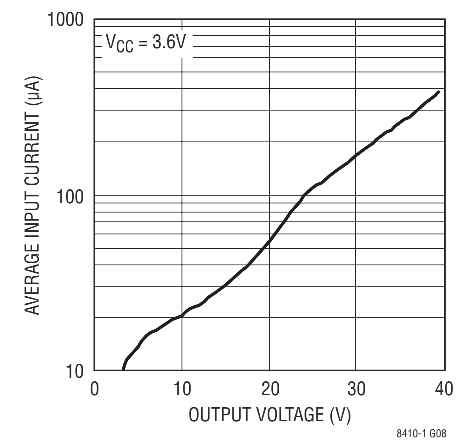

Figure 6. Average Input Current in Regulation with No Load

Data Sheet Download

www.linear.com

For applications help, call (719) 593-1579, Ext. 245

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

dn465 LT/TP 0509 155K • PRINTED IN THE USA

the cuprum material's 35-micron-thick layer. The BCD switches are in parallel and connect from 125Ω to 100 kΩ to adjust the output voltage on the op amp's negative input. The equations to calculate resistor values are: $V_{SENSE}=R_{SENSE}\times I_{LOAD}$, $I_{OUT}=R_{SENSE}\times I_{LOAD}/100$, and $R_0=V_{REF}\times 100/(R_{SENSE}\times I_{LOAD})$. If you choose a value of 0.1Ω for the sense resistor and a value of 0.1V for the reference voltage, the **equation** becomes $R_0=100/I_{LOAD}$.

Applying this **equation**, you can calculate the four weighting resistors of the three BCD switches, which you can determine when the current flows through only that resistor. For currents of 800, 400, 200, 100, 80, 40, 20, 10, 8, 4, 2, and 1 mA, the corresponding resistances would be 0.125, 0.25, 0.5, 1, 1.25, 2.5, 5, 10, 12.5, 25, 50, and 100 kΩ. If the load current is 1A, then the output current is only 1 mA, and, if the load current is 1 mA, then the output current is only 1 μA. Note that the IRF520's surface is on the drain potential. EDN

REFERENCE

1. Guy, John, "Programmable current source requires no power supply," *EDN*, June 12, 2008, pg 70, www.edn.com/article/CA6566536.

Multiplexed, programmable-gain, track-and-hold amplifier has instrumentation inputs

W Stephen Woodward, Chapel Hill, NC

ADCs need adequate signal-acquisition analog interfaces to perform at their best. The classic general-purpose ADC front end includes multiple channels of differential input, digitally programmable gain, and track-and-hold capability. This Design Idea presents a new, complete, high-performance, low-parts-count ADC front end that implements the standard ensemble of functions (**Figure 1**). However, it also incorporates the concepts of the flying-capacitor differential input and the divergent-exponential negative-time constant that an earlier Design Idea describes (**Reference 1**). This Design Idea adds to that circuit multiplexed inputs and a versatile track-and-hold function.

The multiplexer address and the state of the hold-mode bit control signal acquisition and conditioning. With a hold state of zero and the multiplexer's address equal to the selected input channel, the flying capacitor, C_1, connects to the positive and negative differential-input terminals, which acquire the input voltage. Moving the hold state to one isolates C_1 from the input. Then, the multiplexer's address becomes zero, and the hold state returns to zero, initiating regenerative negative-time-constant exponential amplification of the input voltage. From that point until the point when hold reasserts and a connected ADC samples and converts the output voltage, the input voltage and the output voltage are divergent exponential functions of time, with a gain equal to $2^{(1+t/10\ \mu sec)}$.

Building on the assets of that earlier design, this new circuit has the desirable features of multiple instrumentation-style differential inputs. Also, neither resistor matching nor the CMR (common-mode rejection) of the op amp limits the circuit's CMR. Stray-capacitance issues do have an effect on CMR, but you can minimize this capacitance by careful circuit layout. The circuit also has rail-to-rail inputs and virtually unlimited programmable gain. Further,

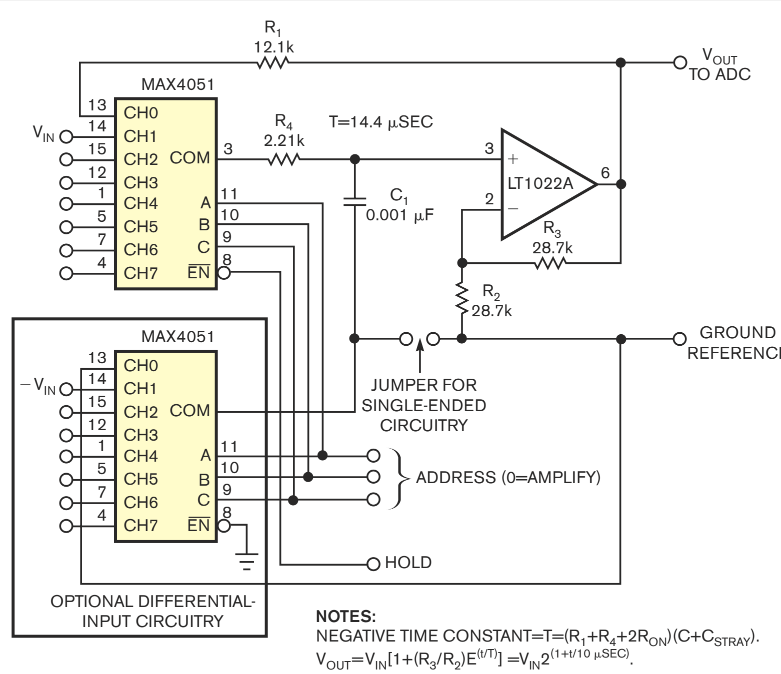

Figure 1 This high-performance, low-parts-count ADC front end implements the standard ensemble of functions.

only the resolution of the amplify interval's timing limits gain-set resolution (**figures 2** and **3**). This circuit also has ±10V output-amplitude capability—two to four times greater than that of monolithic digitally programmable-gain instrumentation amplifiers.

The inherent noise and dc accuracy of the chosen op amp, the accuracy and repeatability of the timing of exponential generation, ADC sampling resolution, and RC-time-constant stability are the main limits on signal-processing performance and the amplifier's precision—for example, its gain-programming accuracy, dc error, noise, and jitter. In the circuit, 1 nsec of the amplify-interval timing error or jitter equates to 0.007% of gain-programming error.EDN

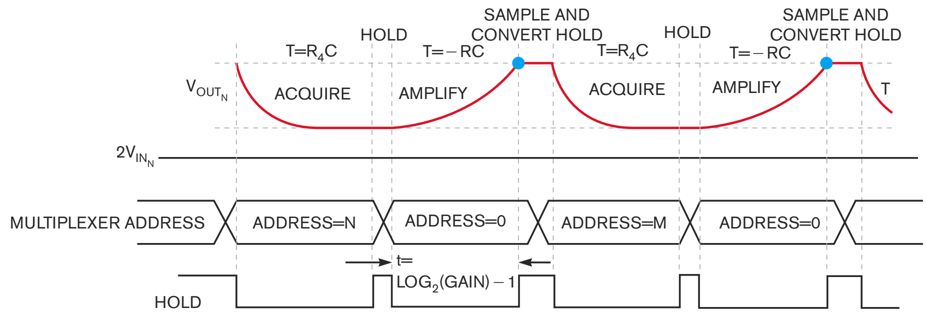

Figure 2 Only the resolution of the amplify interval's timing limits gain-set resolution.

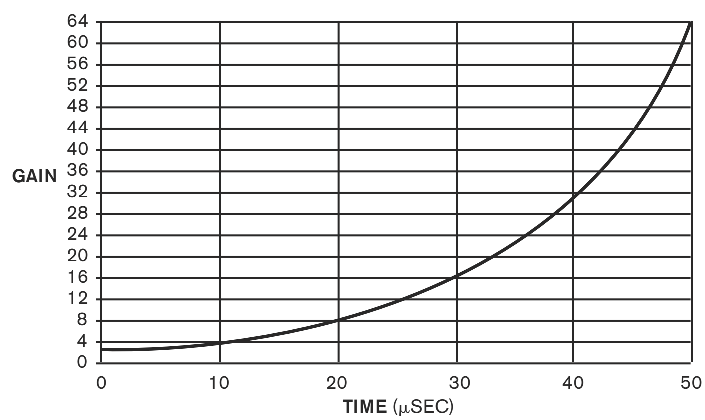

Figure 3 This graph of input- and output-voltage gain shows the time elapsed since the track/amplify-logic transition.

REFERENCE

1 Woodward, W Stephen, "Flying capacitor and negative time constant make digitally programmable-gain instrumentation amplifier," *EDN*, Feb 5, 2009, pg 48, www.edn.com/article/CA6632372.

Simple circuit smoothly drives stepper motors

Uwe Schüler, Institute of Physiology, Tübingen, Germany

The circuit in this Design Idea drives low-power, unipolar stepper motors using only a shift register, a few resistors, and low-power transistors. Adding an inexpensive 4053 analog switch allows bidirectional switching. Compared with other simple stepper-motor-drive circuits, it has better-than-half-step characteristics (**Figure 1**).

After power-up, all shift-register outputs are in a zero state. Pin QP3 feeds back to the serial input through an inverter—transistor Q_5 in **Figure 2** and analog-switch IC_2 in **Figure 3**. The circuit generates a sequence of four ones and then four zeros. You can use this pattern to drive, for example, NPN transistors with emitters that tie to ground and collectors that tie to the stepper-motor coils. However,

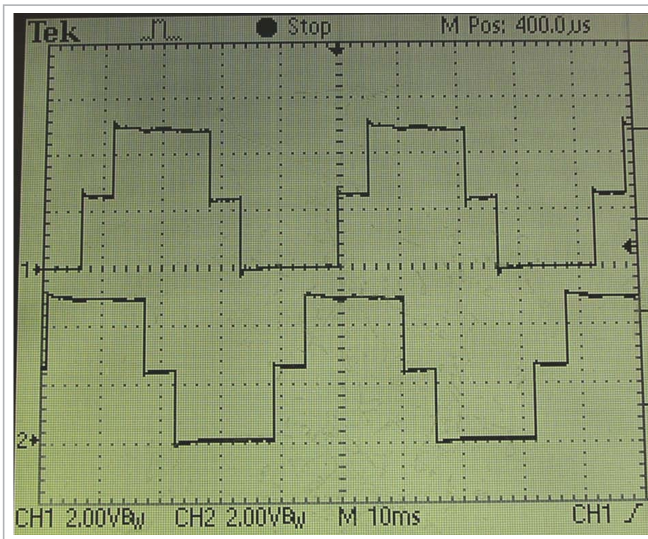

Figure 1 An oscilloscope snapshot shows the base voltages of Q_1 and Q_2 in figures 2 and 3.

to achieve smoother drive characteristics, the shift-register outputs drive four simple DACs, each comprising two identical resistors.

These DACs can generate output voltages of 0, 2.5, and 5V to drive four emitter followers. A snapshot from an oscilloscope shows the base voltages of Q_1 and Q_2 (**Figure 1**). They come close to a quarter-step drive pattern. The circuit can use almost any 8-bit shift register.EDN

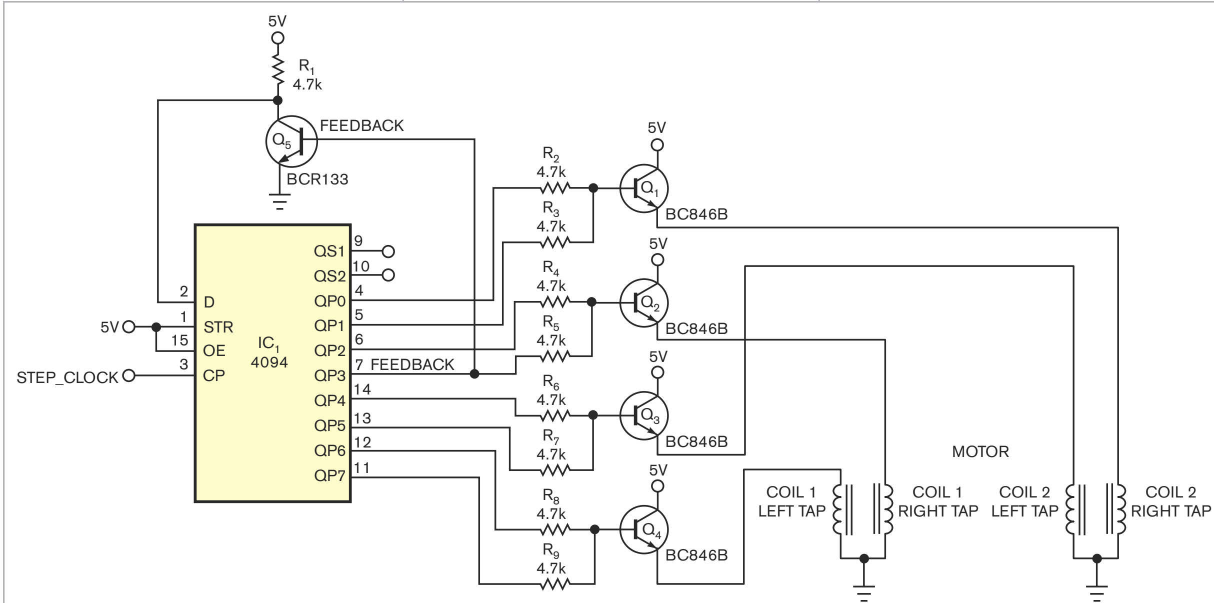

Figure 2 This circuit drives low-power, unipolar stepper motors using only shift-register IC_1 and a few resistors and transistors.

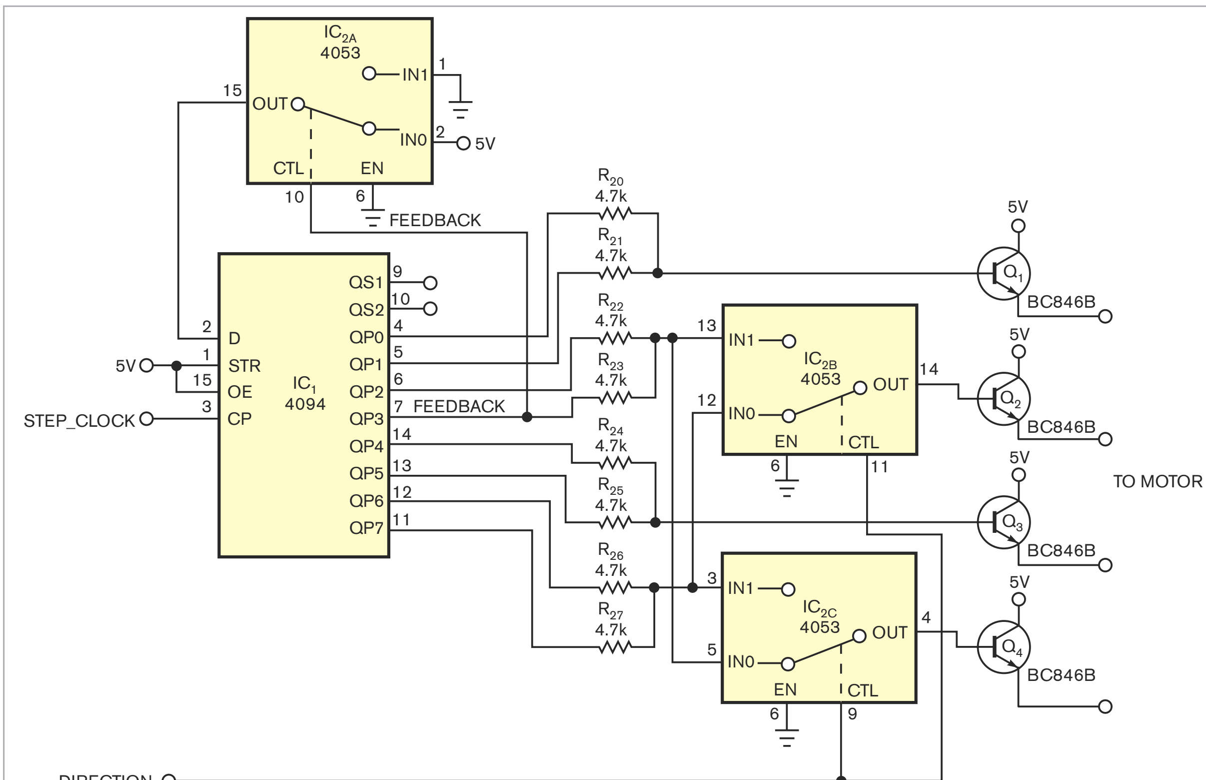

Figure 3 This circuit enhances the one in Figure 2 by adding an inexpensive 4053 analog switch, allowing bidirectional switching.


XILINX
COMPLEX DESIGN MADE LOGICAL
Cut development time as much as 50% while adding breakthrough performance, power and cost benefits. The ISE Design Suite 11 enables new domain-specific design for Targeted Design Platforms, allowing you to focus on building design differentiation.
Download a free evaluation at www.xilinx.com/ise.
IP
C/C++
Verilog
VHDL
FFT
MicroBlaze
PCIe
Ethernet
GMACS
DRC
I/O
MGT
MHZ
ISE
DESIGN SUITE
© Copyright 2009 Xilinx, Inc. XILINX, the Xilinx logo, and ISE Design Suite are trademarks of Xilinx in the United States and other countries. All other trademarks are the property of their respective owners.

Excel spreadsheet yields RLC best-fit calculator

Alexander Bell, PhD, Infosoft International Inc, New York, NY

Commercial off-the-shelf software such as Microsoft (www.microsoft.com) Excel lets you automate engineering functions (**references 1 through 3**). This Design Idea explains how you can use Excel to calculate the values of two passive components—resistors, inductors, or capacitors—from the standard E-Series, which comprises E6, E12, E24, E48, E96, and E192, that you can use in circuits such as filters. The application's results depend on whether you select a parallel- or a series-connected topology.

The calculations appear in an Excel spreadsheet that you can download from the online version of this Design Idea at www.edn.com/090528dia. The VBA (Visual Basic for Applications) source code for this project resides in a single code module (**Listing 1**, which is also available with the online version of this article). It contains three main public functions, FitR(), FitL(), and FitC(), and several private auxiliary functions. The key algorithm loops through the range of values, trying to find the best fit for the target. There is an inner loop for the first value of RLC and an outer loop for the second one.

Figure 1 shows the user interface. You can enter the user-defined functions FitR 1234, P, or E192 into any cell of the Excel worksheet. The cells accept four arguments and return a text string containing the best-fit values, R_1 and R_2 in this case, and the relative error of approximation. **Table 1** shows the functions' parameter list. For better readability, the spreadsheet returns the values of R_1 and R_2 in commonly used electrical-engineering format by applying a scientific-to-engineering format-conversion function, E2BOM().

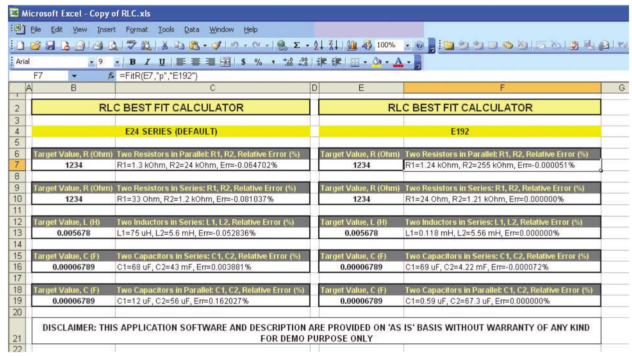

Figure 1 In the user interface, you enter the user-defined functions FitR 1234, P, or E192 into any cell of the Excel worksheet.

The computation engine for electrical resistance and inductance components uses the same formulas: a simple sum of the resistance for the series connection and a sum of conductance for parallel topology, whereas, in the case of the capacitors, the formula is vice versa. You can also fine-tune the functions by changing the constant values corresponding to the upper and lower search limits (**Listing 1**). Thus, you can extend the search range and increase the accuracy, although this process requires more computation time. If you use Microsoft Office 2007, you must contend with an increased security level and set the proper permission level to run the VBA content of the Excel workbook.

This approach is essentially a desktop application, extending the functions of the popular Excel application. You can install the application on either a computer or a network. To further extend its accessibility and bring it to the global level, you should consider an online Web application. The modern RIA (rich-Internet-application) concept and corresponding development tools, available on the market, let you build Web applications with the level of interactivity and responsiveness close to those of the desktop application. A Web-based application provides for easy implementation and maintenance. The user needs only a Web browser. Web applications are essentially platform-independent and globally accessible. Web-based applications of the RLC calculator don't require the user's machine to have MS Office. You can also place the application in password-protected directories from which you can control access to them. A demo version of an online RLC best-fit calculator incorporates the latest set of Microsoft technologies, such as ASP.NET, C#, and Ajax, providing a rich user experience with high interactivity and responsiveness (**Reference 4**).EDN

REFERENCES

1 Bell, Alexander, "Add CAD functions to Microsoft Office," *EDN*, March 21, 2002, pg 94, www.edn.com/article/CA200384.

2 Bell, Alexander, "Simplify computer-aided engineering with scientific-to-engineering conversion," *EDN*, Sept 30, 2004, pg 94, www.edn.com/article/CA454644.

3 Bell, Alexander, "Voice feedback enhances engineering calculator," *EDN*, July 11, 2002, pg 108, www.edn.com/article/CA231578.

4 Bell, Alexander, "Best Fit RLC Calculator," www.alexanderbell.us/RLC/RLC.aspx.

TABLE 1 FUNCTIONS FITR(), FITL(), AND FITC() PARAMETER LIST

No.	Parameter	Description	Required
1	R	Target value	Yes
2	ParSer	Topology: parallel or serial connection	Yes
3	ESeries	Standard series: E6, E12, E24, E48, E96, or E192	No: Default value is E24
4	ExtSearch	Flag to use preferred search limit or extended	No: Default is preferred search range

Automatically turn secondary lamp on or off

Vladimir Oleynik, Moscow, Russia

You may sometimes need to turn on a secondary device, such as a lamp or an alarm, when a device that is normally on loses power. You can build a simple circuit using just a transformer and a relay for this purpose. In the circuit, a primary load is in series with an ac-mains transformer (**Figure 1**). The transformer connects in an unusual way. Its usual secondary low-voltage winding is Winding 1, and its primary ac-mains winding is Winding 2. Under these conditions, the main lamp's voltage is slightly less than during its ordinary operation—the ac-mains voltage minus the voltage drop over Winding 1. That situation is acceptable in most cases because the lower voltage doesn't greatly affect the operation of the load—that is, the luminosity of the main lamp. Select Winding 1 to match the main load's current needs. In this circuit, a 220V, 50-Hz ac voltage appears at Winding 2.

Connect a relay to Winding 2 so that the secondary loss connects to the relay's NC (normally closed) terminal. Use a relay with a winding that can operate at 220V, 50 Hz for your ac-mains voltage. For example, you can use a TR91-220VAC-SC-C relay from Tai-Shing Electronics Components Corp (www.tai-shing.com.tw). This relay's coil operates at a 220V, 50-Hz, SPDT (single-pole/double-throw) commutation of 240V ac under a 40A load.

Using an SPDT relay adds flexibility in controlling the spare load. It lets you switch a load on or off with no need for additional electronic components. In the **figure**, a spare lamp turns on when the main lamp burns out because the secondary load connects to the relay's NC contact.

Select a transformer whose secondary winding (Winding 1 in the **figure**) has a low-rated voltage that provides sufficient current for the main load—the lamp. Match the relay's rated coil voltage to the ac-mains voltage and frequency specifications. EDN

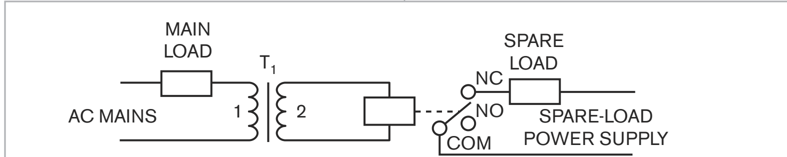

Figure 1 A transformer and a relay are all you need to control a secondary load should the main load fail.

productroundup

SENSORS AND TRANSDUCERS

MEMS motion sensor suits mobile phones

The LIS352AX MEMS sensor operates with any supply voltage over a 2.16 to 3.6V range and remains stable over time and temperature. The sensor uses a three-axis accelerometer that provides acceleration values in absolute analog outputs. It uses an internally regulated voltage, making the measurements insensitive to typical variations in power-supply voltage in battery-operated devices. The sensor has a direct connection to the battery without the need for a separate voltage regulator, making the device suitable for space- and cost-constrained battery-operated devices. The device has high stability over a wide temperature range for zero-gravity offset and sensitivity and provides a ±0.3-mg°/C offset-drift temperature variation. Available in a 3×5×1-mm³ plastic package, the LIS352AX MEMS sensor costs $1.30 (10,000)/year.

STMicroelectronics, **www.st.com**

Differential Hall-effect sensor uses peak-detecting algorithm

The three-wire ATS617LSG differential analog Hall-effect sensor switches in response to changing magnetic fields. The dynamic, peak-detecting sensor uses an analog peak-detecting algorithm, automatic gain control, and an integrated tracking capacitor, providing accurate edge detection corresponding to a rotating ferrous target. The peak-detecting algorithm suits use in camshaft systems that can sense a range of target geometries, including fine-pitch, encoded targets. The ATS617LSG differential Hall-effect sensor costs $2.55.

Allegro MicroSystems, **www.allegromicro.com**

High-resolution MEMS motion sensor measures tilt, shock, and acceleration

Measuring tilt, shock, and acceleration and operating at a 1.8V primary supply voltage, the ADXL346 digital three-axis iMEMS smart-motion sensor provides 4-mg/LSB resolution, enabling you to measure inclination changes as small as 0.25°. The three-axis smart-motion-measurement system includes selectable measurement ranges and measures dynamic acceleration from motion or shock and static acceleration, such as gravity, allowing use as a tilt sensor. The device automatically modulates power consumption in proportion to the output-data rate and saves additional power by offering a standby mode or by switching to sleep mode during periods of inactivity. A FIFO-memory block stores 32 sample sets of X-, Y-, and Z-axis data and offloads the FIFO function from the host processor. This function allows the host processor and other peripherals to go into sleep mode until needed. Available in a 3×3×0.95-mm LGA-16 lead package, the ADXL346 costs $3.04 (1000).

Analog Devices, www.analog.com

TEST AND MEASUREMENT

Trace probe targets ARM Cortex-M3 cores

Aiming at ARM Cortex-M3 cores, the J-Trace high-speed hardware-trace probe provides plug-and-play compatibility with the vendor's Embedded Workbench integrated-development environment. The device accelerates debugging, allowing designers to perform development on the target hardware. The trace probe connects to the host PC running Windows over a USB port, which also provides power. The device also works as a standard JTAG debugging probe through a separate JTAG port. The J-Trace high-speed hardware trace probe costs $1195.

IAR Systems, www.iar.com

3U load board confirms VITA 46/48 specifications

The 3U VPX load board helps confirm that the chassis meets the VITA 46/48 power specifications for VPX and assists in locating hot spots in the enclosure. Load-board functions test a system's cooling capabilities by applying the load to the power supply for verification and creating necessary heat to confirm chassis cooling. Locating hot spots enables the user to redirect airflow and prevent overheating. The card also features a microcontroller-based 100W maximum stepped load control. The 3U VPX load board costs $1000.

Elma Bustronic, www.elmabustronic.com

EDN ADVERTISER INDEX

Company	Page
Achronix Semiconductor	43
Agilent Technologies	8, 36
Analog Devices	17
austriamicrosystems AG	54
BuyerZone	55
Coilcraft	7
Digi-Key Corp	1
Electro Technik	25
Express PCB	41
Jameco Electronics	27
Keil Software	40
Kepco	41
Lattice Semiconductor	53
Linear Technology Corp	45
	46, 49
LPKF Laser & Electronics	42
LS Research	56
The MathWorks	21
Maxim Integrated Products	10, 11

Company	Page
Mentor Graphics	15
Micro Crystal	26
Microsoft	35
Mill Max Manufacturing	9
Mouser Electronics	4
National Instruments	2
Numonyx	C-3
Pico Electronics	C-2
	5, 34
Sealevel Systems	26
Sensirion AG	57
Tadiran Electronic Industries	19
Trilogy Design	57
Vicor Corp	22
	C-4
Xilinx Inc	51

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

EDN product mart

This advertising is for new and current products.

 BY WALTER LINDENBACH

Full speed ahead!

A gas plant was having a speed-control problem with 1800-hp Solar Saturn jet-turbine-driven gas compressors—the same engine that US Air Force jet fighters used in Vietnam. For some reason, the over-speed sensor was failing—a bad situation because only the load of the compressor limited the turbine speed, and if the compressor was "running dry," there was no load. It could then continue to accelerate to as much as 36,000 rpm, at which point it would begin to "throw blades." In other words, the little steel blades on the turbine rotor would begin to come off at a speed such that they would penetrate a half-inch of steel casing and could then keep on going through a number of people if they were in the way. Not good!

The engine's electromechanical speed-limit switch failed regularly because of vibration. A tachometer generator produced an alternating output voltage whose frequency was proportional to the rate of shaft rotation, and we were to use this signal to produce an ignition-off output signal if the turbine rose to a speed of 108% of normal.

The circuit consisted of a reference oscillator and a counter to compare it with the input frequency. The counter was made with the ancient HTTL (high-threshold transistor logic) operating at 15V. The tachometer-input signal was the clock input, and the reset signal to the counter was supplied by the reference oscillator operating at a fraction of the tachometer-input signal frequency produced when the jet turbine was operating at 108% of maximum speed. Then, if the jet-turbine speed exceeded the upper setpoint, the counter would overflow, producing the ignition-off output signal.

We scheduled installation for a −20°F December day. I put the precious box and my tools into the car and drove to the gas-plant site in Carstairs, AB, Canada. The technician took me to the compressor building. We had ear protectors, but the noise and vibration rattled me. The compressor building was about 0°F, and I couldn't wear gloves during the installation.

The technician shut down one of the compressors, though the noise did not noticeably diminish, and motioned to me to go ahead. The first thing I found was 24V on a terminal strip. That discovery rattled me some more, but, when he came back, he said, "Oh, sorry. I forgot the standby battery."

When the installation was complete, I checked and rechecked the wiring. It just had to be right: Turbine blades turning me into mincemeat was a powerful motivation. The technician turned on the control panel and began the start-up sequence, and the tachometer needles began to turn. Advancing the throttle, he brought the turbine up to 40% of full speed, then sharply spun the throttle knob up to full speed!

My heart nearly stopped as the tachometer needles spun around. Then, there was a blood-curdling scream and roar as of a great beast in agony, but—luckily—my limit sensor operated, and the turbine stopped. The tachometer needles began to turn toward zero.

I began to breathe again and asked the technician, "Why did you do that?"

"Oh," he said, "I knew it would work." EDN

Walter Lindenbach started and operated Calgary Controls Ltd (Calgary, AB, Canada) from 1970 to 1990, at which point he discovered an allergy to work and retired. You can reach him at lindenbachw@shaw.ca.

+ www.edn.com/tales

DANIEL VASCONCELLOS

Memory for the mind's eye.

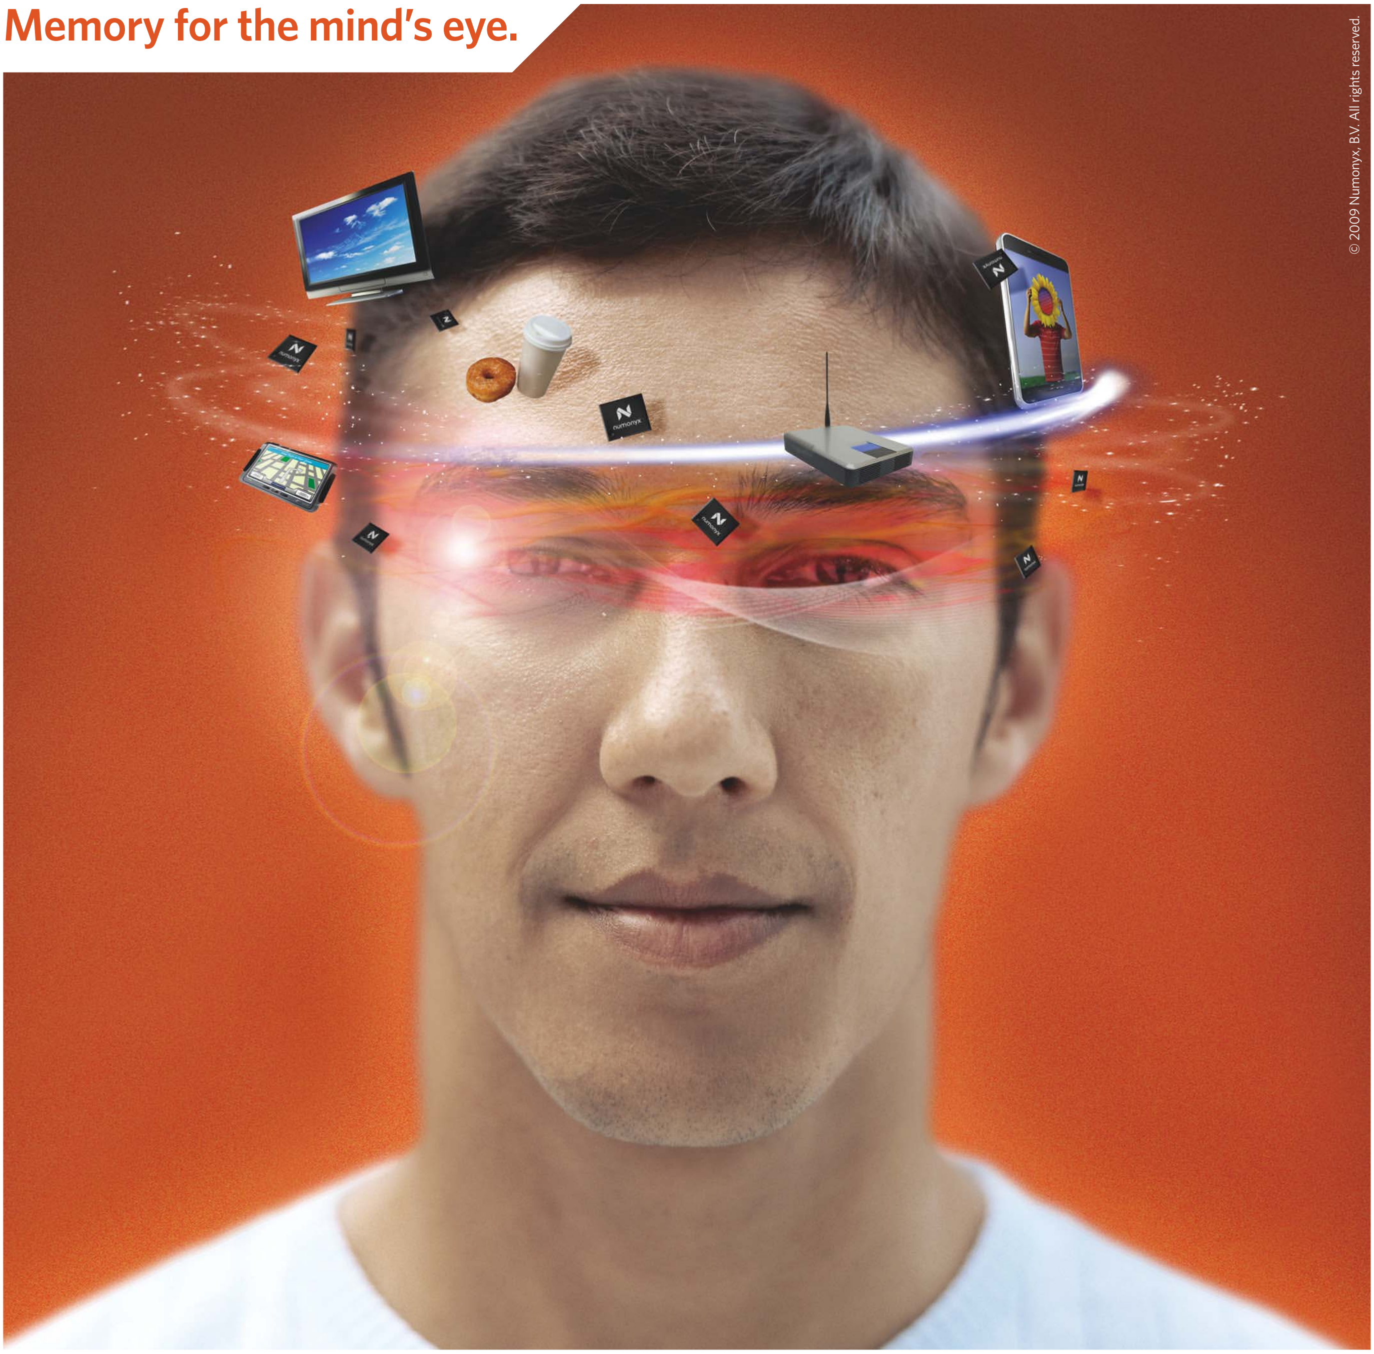

Numonyx® Axcell™ Flash Memory

More headroom for your imagination.

Numonyx® Axcell™ is the world's leading parallel NOR flash for embedded applications. Its easy-to-use command sets, robust security features, leading performance and broad design support free your mind for imagination. See why Numonyx Axcell delivers everything you want to accelerate your next design—and above all, excel. Request your free solutions toolkit today at **www.numonyx.com/go/axcell**.

AXCELL™ PARALLEL NOR FLASH MEMORY

Product	Voltage	Performance	Density
P30	1.7V - 2.0V, 1.7V - 3.6V	X16 52 MHz Burst	64 Mb - 2 Gb*
P33	2.3V - 3.6V	X16 52 MHz Burst	64 Mb - 2 Gb*
M29EW	2.7V - 3.6V, 1.65V - 3.6V	X8, X16 Page	32 Mb - 2 Gb*
M29W	2.7V - 3.6V	X8, X16 Page**	4 Mb - 256 Mb
M29DW	2.7V - 3.6V	X8, X16 Page	32 Mb, 128 Mb - 256 Mb

* Stacked Solution

** Page mode supported on specific densities.

numonyx™

innovative. memory. solutions.