UBM Electronics

EDN®

VOICE OF THE ENGINEER

NOV 17
Issue 22/2011
www.edn.com

INTELLIGENT LIGHTING

HOLDS PROMISE FOR EFFICIENT LIGHTING NETWORKS

Page 36

DESIGN FEMTOAMPERE CIRCUITS WITH LOW LEAKAGE, PART ONE

Page 30

USING MCAPI TO LIGHTEN AN MPI LOAD

Page 45

EDN contents

11.17.11

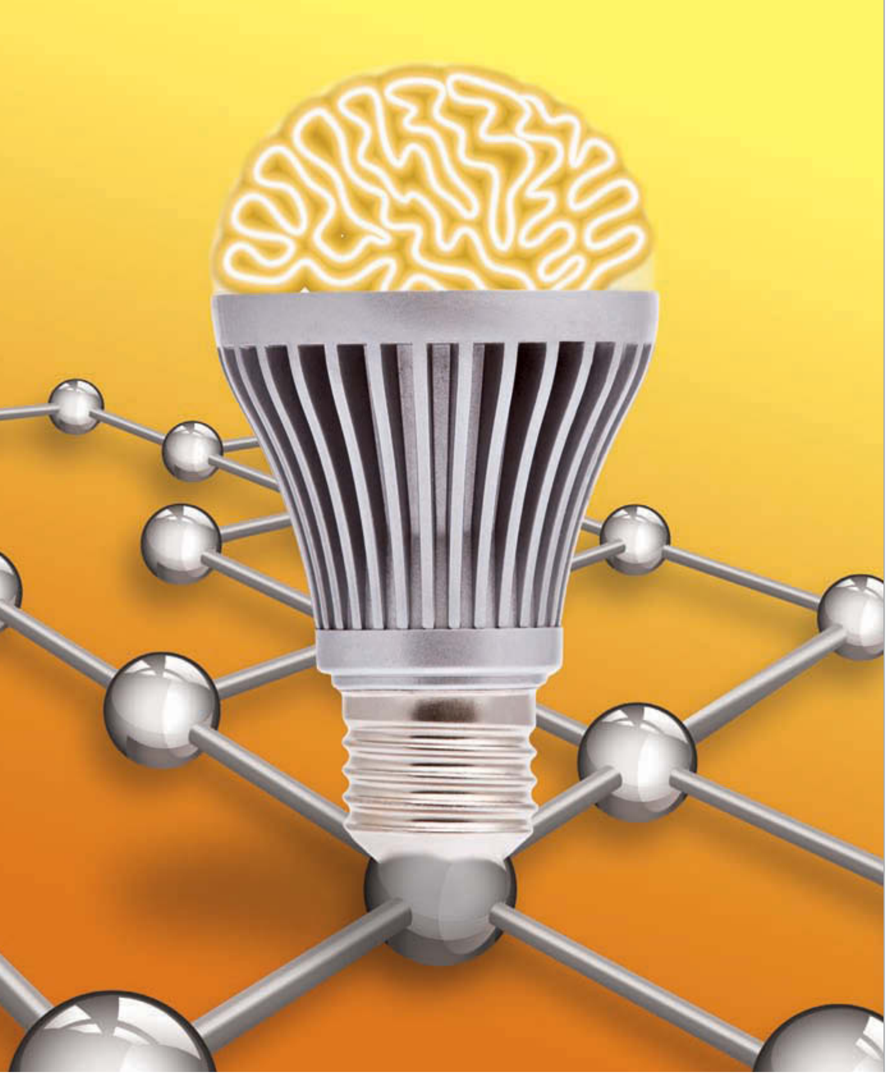

COVER IMAGE(S): THINKSTOCK

DESIGNIDEAS

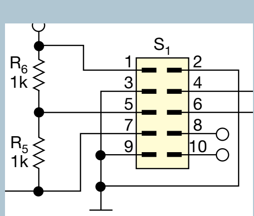

▶ Find out how to submit your own Design Idea: http://bit.ly/DesignIdeasGuide.

contents 11.17.11

pulse

DEPARTMENTS & COLUMNS

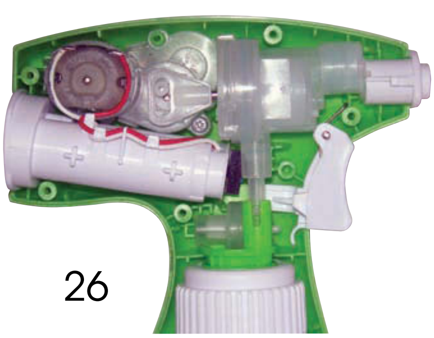

26

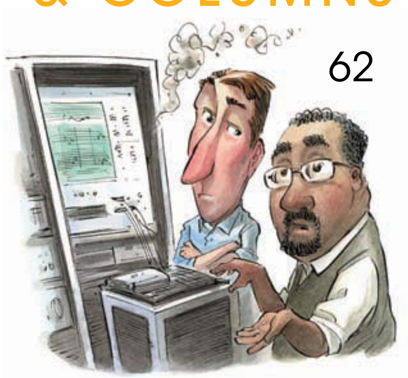

62

EDN® (ISSN# 0012-7515) is published semimonthly by UBM Electronics, 600 Community Drive, Manhasset, NY 11030-3825. Periodicals postage paid at Manhasset, NY, and at additional mailing offices. SUBSCRIPTIONS—Free to qualified subscribers as defined on the subscription card. Rates for nonqualified subscriptions, including all issues: US, $150 one year; $250 two years; $300 three years. Except for special issues where price changes are indicated, single copies are available for $10 US and $15 foreign. For telephone inquiries regarding subscriptions, call 847-559-7597. E-mail: edn@omeda.com. CHANGE OF ADDRESS—Notices should be sent promptly to EDN, PO Box 3609, Northbrook, IL 60065-3257. Please provide old mailing label as well as new address. Allow two months for change. NOTICE—Every precaution is taken to ensure accuracy of content; however, the publishers cannot accept responsibility for the correctness of the information supplied or advertised or for any opinion expressed herein. POSTMASTER—Send address changes to EDN, PO Box 47461, Plymouth, MN 55447. CANADA POST: Publications Mail Agreement 40612608. Return undeliverable Canadian addresses to BleuChip International, PO Box 25542, London, ON N6C 6B2.

Volume 56, Number 22 (Printed in USA).

EDN online

JOIN THE CONVERSATION

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

In response to "Gordon Nuttall: turning layoff lemons into start-up lemonade," by Suzanne Deffree, http://bit.ly/w5eaKj, Martine Simard-Normandin commented:

"I was laid off at 50 and did the same thing. MuAnalysis is now 10 years old. It's been a roller-coaster ride adjusting to the ever-changing market conditions of the decade. Every day brings a new challenge and a new opportunity. I am happy I took the plunge and never regretted it."

In response to "You auto know," a Tales from the Cube column submitted by Tata Motors' Vishwas Vaidya, http://bit.ly/v8I9Du, Mark Rackin commented:

"There's another reason to put those small-value bypass caps near the connector. It has nothing to do with operation in the vehicle. It has to do with ESD protection in the manufacturing process and in handling/storage before being installed and when serviced.

Having some extra capacitance on those signal pins (and power also; electrolytic caps have too much internal inductance to help for fast-rise ESD discharges) prevents low to medium ESD events from raising voltages to the level of being destructive to chips.

That's why it's a 'best-practices/lesson-learned' item in many automotive-OEM suppliers. It's not something you would be aware of in a high-humidity environment like South India, but if you start getting unexplainable field failures from car services in, say, a Chicago winter, don't be surprised!"

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

CONTENT

Can't-miss content on EDN.com

JIM WILLIAMS: CIRCUITS AS ART

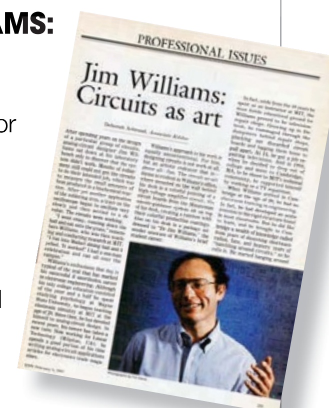

Technical Editor Paul Rako shares an article about Jim Williams' artistry, which was published in the Feb 5, 1987, issue of *EDN*.

http://bit.ly/rtRZ3B

WI-FI ALL OVER

As wireless technology moves into just about every aspect of society and industry, Wi-Fi chip sets are going into all sorts of devices—consumer, industrial, and specialized niche markets. The explosion in the market challenges even the largest chip vendors, which must keep up with trends and decide which markets to target. It also creates opportunities for small start-ups that can find success by focusing on the right niche markets.

http://bit.ly/sBumRG

ENGINEERING COMMUNITY

Opportunities to get involved and show your smarts

If there were a fire and you could save only one thing from your workbench, what would it be and why? In *EDN*'s 5 Engineers section, posted in our Voice of the Engineer blog (**http://bit.ly/VoiceOfTheEngineer**), we invite our audience to answer similar questions. Five responses are featured in the "Fun Friday" newsletter. Tune in each week to share an answer and see what your fellow EEs have to say. Subscribe to "Fun Friday" here: **https://subscribe.ednmag.com/data/edn/welcome**.

EDN

BY PATRICK MANNION, DIRECTOR OF CONTENT

Is Applied Micro's X-Gene equal to a 1960s 14-transistor radio?

This is a test; do not adjust your age: Do you recall the transistor-radio wars of the 1960s? Japan then dominated the maturing transistor-radio market, and, in an effort to differentiate themselves, Korean manufacturers started adding functionless transistors to their radios just so that they could say they had eight, 10, 14, and, eventually, 16 transistors instead of the necessary four to eight. The original four-transistor design goes back to the pioneer Regency TR-1, which Texas Instruments and Industrial Development Engineering Associates designed.

So why am I bringing this up now? Last month, at ARM TechCon, ARM announced the ARM Version 8 ISA (instruction-set architecture); hours later, Applied Micro unveiled X-Gene, the first processor employing the architecture. Oddly, parallels may exist between those 14-bit Korean radios and X-Gene—and all other multicore processors, for that matter. I'm just picking on X-Gene because it's one of the most recent.

I'll leave it up to you decide whether this technology is the Holy Grail and whether we have the tools for it.

Operating at frequencies as high as 3 GHz and touting as many as 128 quad-issue, out-of-order cores, the X-Gene connects through a coherent terabit fabric and a memory throughput of 80 Gbytes/sec. Applied Micro is touting it as the first "server on a chip" and promises it will reduce the total cost of ownership of a server farm by as much as 30% through energy savings alone.

Those promises sound good, but the hardware part bears more scrutiny, especially in the context of unstructured data. Past a certain point, the benefits of multicore do not necessarily scale linearly according to the number of cores. That we know, and many programmers and developers have spent many days and nights working on the tools and algorithms to make maximum use of the parallelism potential that multicore architectures theoretically provide. Still, the tools and scalability are lacking.

A more fundamental problem exists, though. In response to my blog about the Applied Micro X-Gene announcement, I received an e-mail from Russell Fish, co-inventor of the Sh-Boom processor and now with Venray Technology. The subject line was, "ARM X-Gene missed the Sandia Labs multicore report"—an irresistible headline. In that e-mail, Fish referred to a paper he co-authored, on behalf of Venray (**Reference 1**). It states that the memory bottleneck is a fundamental problem that no one has solved and that Patterson's Power Wall, which describes the trade-offs between memory-bus bandwidth, power, and parallelism, still applies, at least with respect to current multicore-processing architectures (**Reference 2**). Add synchronization between cores for that memory access, and the problem quickly becomes untenable.

The paper refers to an experiment in which a Sandia Labs team proved this theory. The team simulated key algorithms for deriving knowledge from large data sets. The simulations showed a significant increase in speed when you increase from two to four cores but an insignificant increase from four to eight. Moreover, for data-intensive programs, 16-core microprocessors delivered the same performance as that of two cores. The answer to the problem, according to Fish, lies at the heart of Venray Technology's TOMI Borealis design, which uses memory transistors for processing. The design has eight 32-bit microprocessors on a 1-Gbit DRAM.

I'll leave it up to you—and time—to decide whether this technology is the Holy Grail and whether we have the tools available to support it. If it's not and if we haven't solved the multicore problem, then all those hyped-up cores may lie wasted: all show, no function, like a bank of useless radio transistors tied to ground.

Comments on this topic and Fish's paper are welcome; just go to the online version of this column (**Reference 3**), which includes additional links and commentary from your peers and, possibly, Russell Fish himself. Also, stay tuned for an in-depth *EDN* feature from Fish on the problems with multicore processing and his theories on how we can overcome them. EDN

REFERENCES

1 Laughlin, Kendall, and Russell Fish III, "Economics of CPU in DRAM and TOMI technology," Venray Technology Ltd, http://bit.ly/tPG4o4.
2 Patterson, David A, and John L Hennessey, *Computer Organization and Design: the Hardware Software Interface*, Elsevier Inc, 2009, ISBN: 978-0-12-374493-7, http://bit.ly/uLi4Jk.
3 Mannion, Patrick, "Is Applied Micro's X-Gene equal to a 1960s 14-transistor radio?" *EDN*, Nov 4, 2011, http://bit.ly/v0bSxw.

Contact me at patrick.mannion@ubm.com.

Integrated Device Technology FemtoClock NG – When a Trillionth Of a Second is Just Too Long

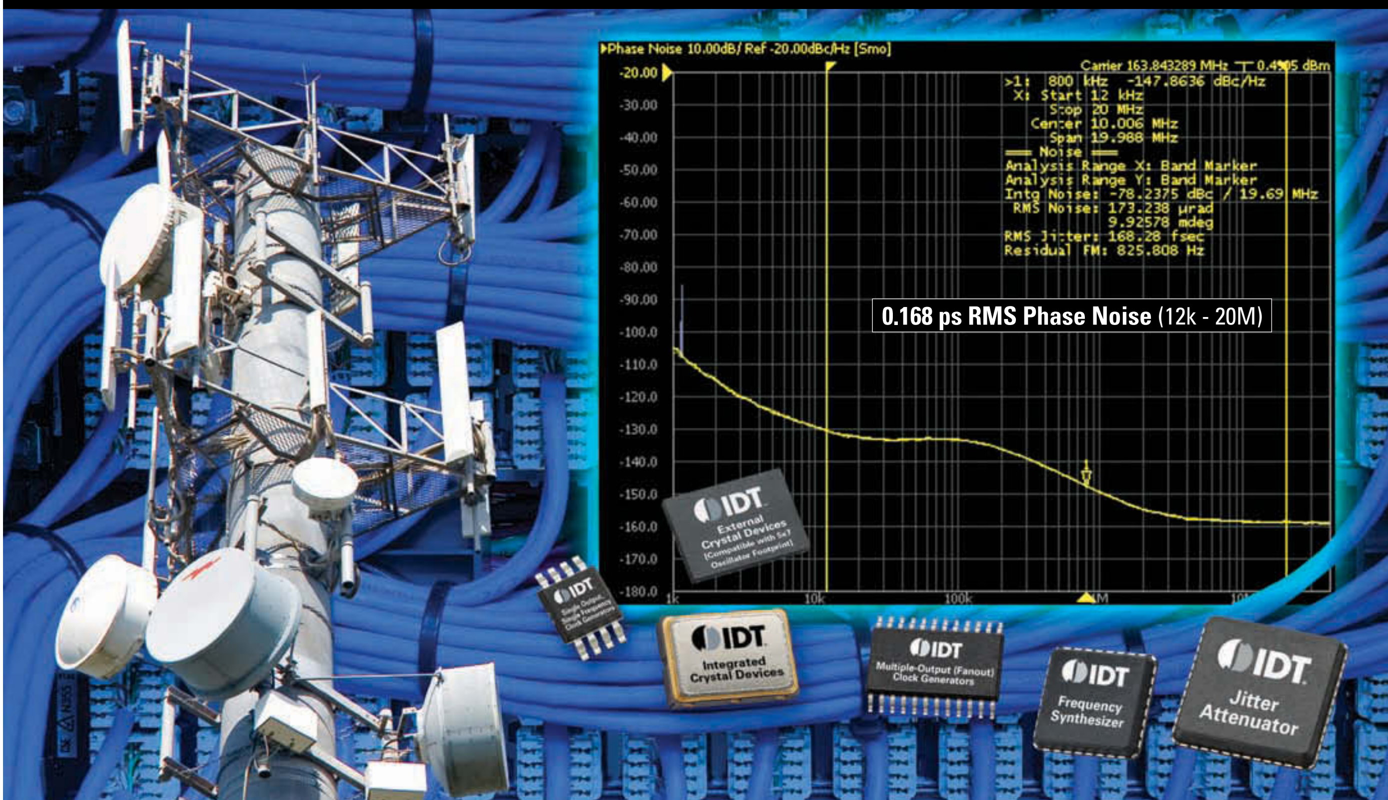

The FemtoClock Next Generation (NG) clock synthesiser family allows engineers to generate almost any output frequency from a fixed frequency crystal, and to meet the challenges of the most demanding timing applications.

Use of a single source to generate multiple clocks often leads to increased susceptibility to power supply noise and restrictions on multiplication factors. IDT's technology effectively eliminates these problems by doubling the power-supply noise rejection (PSNR) of previous generation devices while also introducing the potential for virtually limitless customization of output frequency.

An innovative fractional multiplier PLL architecture introduces the flexibility for engineers to generate any output frequency from any input frequency. And the advanced design of the FemtoClock NG family achieves greater than 80 dB of PSNR to make the devices immune to power-supply noise.

Other performance benefits of FemtoClock NG technology include low power consumption and a clocking performance of under 0.5 ps RMS phase noise jitter! The devices offer standard outputs such as differential LVPECL, LVDS and single-ended LVCMOS, providing a precise fit to any application.

With FemtoClock NG technology, IDT has eliminated the most challenging aspects of silicon-based clock design and introduced an unprecedented level of flexibility for clocking in high-performance applications.

IDT®
Integrated Device Technology

The Analog and Digital Company™

For more information visit us online at www.IDT.com

EDITED BY FRAN GRANVILLE

pulse

INNOVATIONS & INNOVATORS

ESD-protection IC targets ultra-high-speed micro-SD cards

STMicroelectronics recently announced the EMIF06-MSD03F3 micro-SD (secure-digital) ICs, which combine EMI filtering and ESD protection for applications in mobile phones, tablets, and 3G dongles using SD 3.0 UHS-I (ultra-high-speed) micro-SD cards.

According to Strategy Analytics, some 500 million mobile handsets could include a UHS-I-compliant SD-card slot. UHS-class cards will provide as much as 2 Tbytes of storage and improve users' experience for direct recording of high-definition video, playing back or sharing content, or backing up data. The exposed connections in the card slots must have protection against ESD, which users can cause simply by picking up the device, damaging the system's circuitry. The EMIF06-MSD03F3 IC protects micro-SD interfaces to the specified UHS-I speed of 104 Mbytes/sec.

The device allows either electrical or mechanical card-insertion detection, freeing designers to use the method that best suits their applications and hosts. It also integrates pull-up resistors to guarantee system behavior when no card is inserted, along with an EMI filter to block GSM interference.

The device's ±15-kV ESD protection meets the IEC 61000-4-2 standard, and the device protects all SD-card data and power lines. Other features include 7-pF-maximum load capacitance, ST's IPAD (integrated passive- and active-device) technology, and a 1.54×1.54-mm outline. The EMIF06-MSD03F3 comes in a 16-bump, 0.4-mm-pitch WLCSP and sells for 23 cents (1000).

—by Fran Granville

▷**STMicroelectronics**, www.st.com.

TALKBACK

"Remember in the 'old days' [when] TV-antenna cable was flat, clear plastic with tinned-copper wire on each side? After a hurricane my grandfather spliced his broken antenna wire with fishing line! Looked pretty much the same but didn't work at all."

—Automation engineer Bob Clarke, in *EDN*'s Talkback section, at http://bit.ly/ru6M0k. Add your comments.

The EMIF06-MSD03F3 ESD-protection IC offers features for use in mobile phones, tablets, and 3G dongles using SD 3.0 UHS-I micro-SD cards.

MachXO2.
DO-IT-ALL DEVICES DEMAND A DO-IT-ALL PLD.

With the MachXO2 PLD family, you now have everything you need to meet your low power, small form factor, high functionality and low price requirements for your next consumer or system design.

Access free reference designs and software, development kits and supporting literature at:
latticesemi.com/machxo2

Applications

- Smart Phone
- Digital Camera
- GPS
- Server
- Router
- Base Station

And more...

Benefits

- Instant-on, non-volatile
- Flexible logic from 256 LUTs to 7,000 LUTs
- Embedded memory up to 240 Kbits
- Ultra low power as low as 19µW
- On-chip user flash memory & I^2C, SPI, timer/counter
- Prices starting as low as $0.75

And more...

Lattice® Semiconductor Corporation

latticesemi.com/machxo2

Scan QR code to view MachXO2 video

Cree takes a page from Amazon's playbook with its new LED TEMPO Services

"There are two ways to build a product," says Amazon founder Jeff Bezos. "The first: A company starts with their strengths and builds to the needs of the consumer. The second: A company starts with the needs of the consumer and builds [into] the strengths of the company." According to Bezos (**Reference 1**), Amazon is an example of the first approach; it built logistical strength and educated its customers on the benefits of e-commerce. The Kindle is an example of the second; it started with consumers' need for faster delivery in a digital format and built the infrastructure of the company to satisfy that need.

Like Amazon, Cree is another rare company that does both. Cree leverages its strength in high-brightness LED R&D to develop LED chips, packaged devices, and arrays, all focusing with rifle-shot accuracy on general lighting. Cree isn't interested in the backlighting market or the automotive market. It wants to make the solid-state-lighting market take off. And, in large part due to Cree's efforts, LED prices, which have been painfully high, are dropping. However, Cree also sees the quality—or lack thereof—of packaged luminaires as a major obstacle to LED acceptance. Will an LED light do what a consumer expects for as long as it should?

In general, Cree's customers are lighting manufacturers with little experience in making and designing semiconductor-based products. Although third-party labs can provide testing services, such as IES (Illuminating Engineering Society) LM-79 for LED luminaires, Cree believes that third parties miss many other aspects of end-product quality, such as chemical compatibility among materials in the luminaires and the LEDs, the effectiveness of mixing slightly different-colored LEDs for enhanced color consistency, and TM-21 LED-lifetime projections.

Cree sees a testing gap between what third-party labs can provide and what lighting manufacturers need to evaluate their design for quality, performance, and lifetime. To address this gap, the company has created TEMPO (thermal/electrical/mechanical/photometric/optical) Services, which it claims represent the accumulated advantage of Cree's extensive experience with customers' LED systems and the use of calibrated test equipment to give LED-lighting manufacturers and end users confidence in LED-product designs.

Cree offers a range of TEMPO Services to LED-luminaire makers, depending on their product-development needs. The flagship service is the TEMPO 21 Service, the most comprehensive LED-luminaire test available, which measures and analyzes a final product design before submitting it for LM-79 certification. TEMPO 21 examines all of the aspects of quality that Cree has identified as critical. In addition to a TEMPO report, the TEMPO 21 Service includes consultation time with a Cree application engineer to review the testing results and highlight possible areas for improvement in the design.

Cree is also offering two quick-turnaround testing services. The TEMPO Spot Service provides measurements of flux, efficacy, and chromaticity for luminaires and replacement lamps. The TEMPO Flash Service provides measurements of flux, chromaticity, and throw for torches and other portable-lighting designs.

TEMPO Services represent the accumulated advantage of Cree's extensive experience with customers' LED systems and the use of calibrated test equipment to give LED-lighting manufacturers and end users confidence in LED-product designs.

The company is providing TEMPO Services from its Cree Technology Centers in Research Triangle Park, NC, and Santa Barbara, CA. It plans future TEMPO Service locations for Munich, Germany; Shanghai, China; and Taiwan, China.

The flagship TEMPO 21 Service costs approximately $1200 with a two-week turnaround time; the Spot and Flash services cost less. As a cost comparison, prices for third-party LM-79 testing range from $1200 to $2500, and thermal testing costs approximately $1000; turnaround time is closer to months rather than weeks. **—by Margery Conner**

▷**Cree**, www.cree.com.

REFERENCE

1 Tunguz, Tomasz, "The two ways to build a product, according to Bezos," *Ex Post Facto*, Sept 20, 2011, http://bit.ly/nBUrnc.

DILBERT By Scott Adams

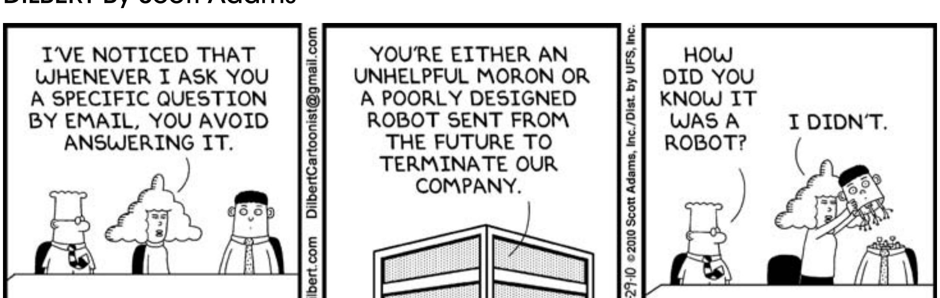

RAQ's

Rarely Asked Questions

Strange stories from the call logs of Analog Devices

Crosstalkin' Converters

Q. Should I consider crosstalk when choosing an A/D converter?

A. Certainly! Crosstalk can come about in several ways: from one signal chain on a printed circuit board (PCB) to another, from one channel within an IC to another, or through the power supply. The key to understanding crosstalk is to discover where it comes from and how it manifests itself. Is it coming from an adjacent converter, another channel of the signal chain, or from the PCB design?

The most typical type of crosstalk testing is called adjacent crosstalk. This form of crosstalk manifests itself when one channel is driven at or near full scale, while the channel or signal chain being "looked at" is open, with no signal applied. A spur that rises above the noise floor will be seen on the open channel when measuring the output frequency spectrum. This type of crosstalk defines the isolation between the open receptor channel and the driven aggressor channel.

Sometimes, open channels are robust enough to suppress cross coupling from one driven channel, but there is strength in numbers. Another crosstalk test drives all but one channel in the system with the same frequency, leaving only one channel open. In this case, the strength of all aggressors is measured through the open channel.

A third way to measure crosstalk is to drive two or more channels with different frequencies and signal strengths, testing the open channel(s) to see if the driven channels produce any cross-coupled mixing products that leak through. In this case, the mixing effect shows how the aggressor signals fall back into the band of interest.

Finally, these same three measurements can be repeated with the input signals in an overrange condition (above the full scale of the device or signal chain). This helps define how robust an open channel is when the input signal is clipped or the channel is saturated.

All of these tests should cover the full signal range and frequency range of interest for the application, as crosstalk can sometimes be caused by a poor PCB design or can manifest itself at specific operating conditions. Swapping parts out won't help. The converter or multichannel device must be thoroughly tested to make sure it is robust enough for your application

To Learn More About Testing High-Speed ADCs
http://dn.hotims.com/34948-100

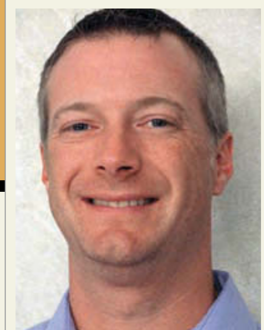

Contributing Writer Rob Reeder is a senior converter applications engineer working in Analog Devices high-speed converter group in Greensboro, NC since 1998. Rob received his MSEE and BSEE from Northern Illinois University in DeKalb, IL in 1998 and 1996 respectively. In his spare time he enjoys mixing music, art, and playing basketball with his two boys.

Have a question involving a perplexing or unusual analog problem? Submit your question to: www.analog.com/askrob

For Analog Devices' Technical Support, Call 800-AnalogD

SPONSORED BY

ANALOG DEVICES

With 64-bit ARM Version 8-based X-Gene, Applied Micro demos clean-slate approach to cloud computing

Hot on the heels of ARM's recent announcement of the 64-bit Version 8 ISA (instruction-set architecture), Applied Micro Circuits Corp demonstrated X-Gene, the first 64-bit ARM Linux running on the first ARM 64-bit hardware (see **references 1** and **2** and view video at http://bit.ly/sHNVf4). The demonstration of the core on an FPGA platform was three years in the making. Applied Micro was a strategic partner with ARM on the development of Version 8; 128 of the devices, operating at 3 GHz, will be on X-Gene when it becomes available for sampling in the second half of 2012.

Applied Micro entered the cloud-computing server market because it saw an opportunity to fundamentally change server design, recognizing a disparity between data servers' original tasks and their current tasks. That disparity is wreaking havoc with the total cost of ownership, which is based not solely on capital expenditures but also on power consumption and is rising at an incremental rate of $95 million per day.

At ARM TechCon 2011, Paramesh Gopi, president and chief executive officer of Applied Micro, demonstrated X-Gene, the first 64-bit ARM Linux running on the first ARM 64-bit hardware. The demonstration of the core's functions on an FPGA platform occurred on the same day that ARM announced its 64-bit ARM Version 8 instruction-set architecture and parallels the company's launch of the industry's first 64-bit ARM "server on a chip." The server on a chip takes a 30% chunk every year from the server farm's total cost of ownership.

At the demonstration, Andrew Feldman, founder and chief executive officer of SeaMicro, described how data and server needs have gone from internally oriented approaches in which the staff is told to wait in line, to customer-oriented cloud computing, in which wait states are not tolerable. "The work changed, but servers didn't," he says.

The wait states are a result of the "bursty" nature of Internet traffic, which can overload servers; meanwhile, downtimes mean that servers are still consuming vast amounts of power in idle mode. Feldman sees a need for small, simple CPUs to improve computation-per-unit power. His company currently uses Intel's Atom but is now shifting to the Version 8. "We will shrink the motherboard to the size of a business card and then connect them," he explains.

The processor tackles the problem from the angles of improved efficiency, hardware usage, and improved latency. From a hardware point of view, this move entails higher integration, efficient out-of-order cores, and virtualization support. The device integrates the cores with all of the networking and I/O, including PCIe and 10/40/100 GbE, all connecting through a coherent terabit fabric and an 80-Gbyte/sec memory throughput. Software support includes Lamp, MySQL, Stack, Apache Server, and Linux.

"The cloud is synonymous with Linux," says Paramesh Gopi, president and chief executive officer of Applied Micro.

—by Patrick Mannion

▷**Applied Micro Circuits**, www.apm.com.

REFERENCES

1 Barak, Sylvie, "ARM unveils 64-bit architecture," *EDN*, Oct 27, 2011, http://bit.ly/vH5RgV.

2 "AMCC demos 64-bit ARM server chip," *EE Times*, Oct 27, 2011, http://bit.ly/tsNJlr.

Freescale goes dual core for engine control

Freescale Semiconductor's new dual-core version of the 32-bit Qorivva microcontroller controls automobile power trains and aims to help car makers meet increasingly stringent fuel-economy and emission standards that the government is phasing in over the rest of the decade.

According to Freescale, power consumption is the main driver for the move to dual-core units because manufacturers can no longer crank up the clock speed of vehicles' microcontrollers to handle next-generation power-train control.

According to Richard Soja, system engineer for 32-bit automotive microcontrollers at Freescale, the Qorivva MPC5676R eliminates the need for multiple packages, reducing chip counts in cars. General Motors has already signed on to use the dual-core device for its high-end vehicles. By 2012, Freescale plans to have signed more automakers when it begins delivering samples. The device can operate in traditional diesel and gasoline engines, as well as in hybrid and all-electric vehicles, easing the transition to those new architectures as they emerge.

The device detects engine knock in real time so that engineers can tune engines for fuel efficiency. The device also provides three enhanced timing-processor units to generate precise timing signals to control fuel ignition. Freescale built the multicore, 90-nm microcontroller on the company's Power Architecture to handle direct-fuel-injection, turbocharging, drive-by-wire, and other traditional engine-control-unit functions. Each 32-bit core runs at 180 MHz, with runtime support from the Autosar real-time operating system. A unit combining a timing compiler, a debugger, and a simulator helps to develop code for the timing units.

The Qorivva MPC5676R eliminates the need for multiple packages, reducing chip counts in cars.

—by R Colin Johnson

▷**Freescale Semiconductor**, www.freescale.com.

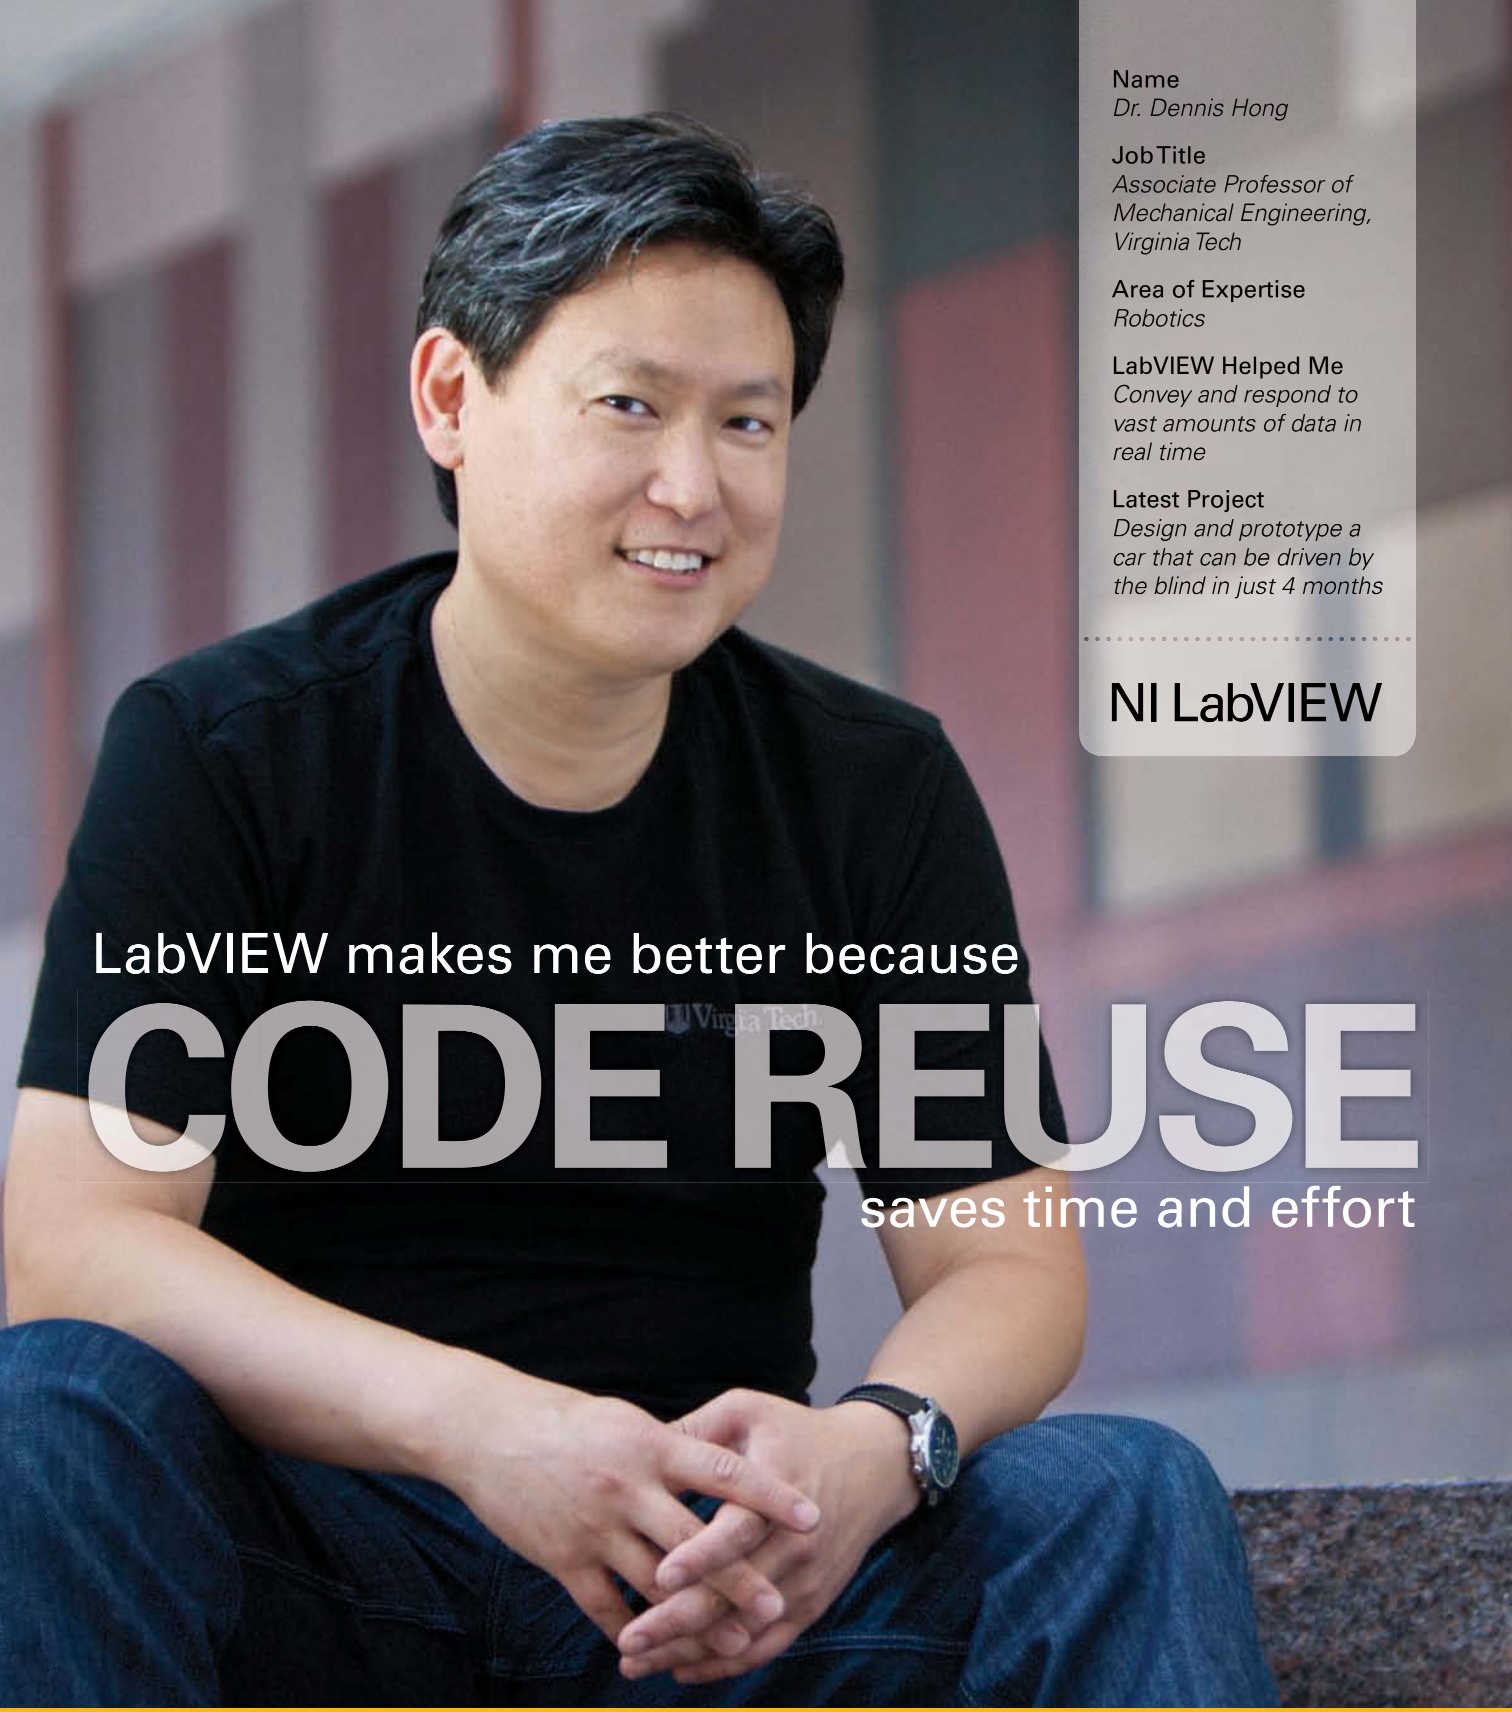
Name
Dr. Dennis Hong
Job Title
Associate Professor of Mechanical Engineering, Virginia Tech
Area of Expertise
Robotics
LabVIEW Helped Me
Convey and respond to vast amounts of data in real time
Latest Project
Design and prototype a car that can be driven by the blind in just 4 months
NI LabVIEW
LabVIEW makes me better because
CODE REUSE
saves time and effort

Modules add low-power Wi-Fi to Atmel microcontroller design in less than five minutes

With more than 4 million hot spots worldwide and the annual sale of 800 million new Wi-Fi-enabled devices, Wi-Fi continues to be a popular, readily available wireless-communication channel. It can also be a challenge to integrate into a microcontroller-based system, however, due both to its complexity, especially for those people who lack wireless-communications expertise, and to portable equipment's requirement for low power consumption. Embedded-system developers have come to expect the availability of modularized toolchain and development systems that can shave months from an embedded-system design cycle.

To fulfill these expectations, Redpine Signals has teamed with Atmel to create a Wi-Fi module using Redpine Signals' Connect-io-n and n-Link hardware, which works seamlessly with Atmel's microcontroller families and evaluation boards. The highly integrated, single-stream 802.11n Connect-io-n modules provide plug-and-play Wi-Fi connectivity to embedded devices. Wi-Fi integration can be time-consuming, according to Venkat Mattela, chief executive officer at Redpine Signals. "Whatever you need to evaluate [about Wi-Fi], you can do it within five minutes of plugging in this module," he says.

> "Wi-Fi integration can be time-consuming, but these modules let you evaluate whatever you need within five minutes."

Prices for Atmel evaluation kits, such as the AVR Xplained boards, start at $29; prices for Redpine Signals' Wi-Fi evaluation kit start at $79.

Fans of the Arduino platform, which also uses an AVR processor, may want to consider the approximately $80 RedFly-Shield kit, which also uses Redpine Signals' Wi-Fi chip set. Adafruit Industries, master of clear, illustrated tutorials for the Arduino platform, might want to consider introducing a similar Arduino Wi-Fi shield.

—by Margery Conner

Redpine Signals has teamed with Atmel to create a Wi-Fi module using Redpine's Connect-io-n and n-Link hardware, which works seamlessly with Atmel's microcontroller families and evaluation boards.

▷**Redpine Signals**, www.redpinesignals.com.

11.17.11

Microchip adds configurable logic to 8-bit microcontrollers

Microchip Technology Inc has added several 8-bit PIC microcontrollers that feature configurable logic and peripheral integration in 6- to 20-pin packages. The PIC10F(LF)32X and PIC1XF(LF)150X microcontrollers each feature configurable-logic cells, complementary waveform generators, and numerically controlled oscillators. The configurable-logic-cell peripherals enable software control of combinational and sequential logic, which increases the on-chip interconnection of peripherals and I/Os, thereby reducing the need for external components, saving code space and adding functions. For a video demo, go to http://bit.ly/ue9pQb.

The complementary-waveform-generator peripheral works with multiple peripherals to generate complementary waveforms with deadband control and automatic shutdown, which provides improved switching efficiencies. The numerically controlled oscillators enable linear frequency control and high resolution, which is necessary for applications such as lighting ballast, tone generation, and other resonant-control circuits.

The devices also feature power consumption of less than 30 μA/MHz in active mode and less than 20 nA in sleep mode, as well as an on-chip, 16-MHz internal oscillator, an ADC, and as many as four PWM peripherals. An integrated temperature-indicator module enables low-cost temperature measurements.

The PICdem lab-development kit now includes samples of both the PIC10F322 and the PIC16F1507, and the F1 evaluation platform is available for development with enhanced midrange-core, 8-bit PIC microcontrollers. A free configuration tool streamlines the setup process of the configurable-logic-cell module by simulating the performance of the registers and combinational logic in a GUI. The PIC10F(LF)320 and PIC10F(LF)322 are available in six-pin SOT-23 packages, eight-pin PDIPs, and 2×3-mm DFN packages. The PIC12F(LF)1501 is available in eight-pin PDIPs, SOIC packages, MSOPs, and 2×3-mm DFN packages, and the PIC16F(LF)1503 comes in 14-pin PDIPs, SOIC packages, TSSOPs, and 3×3-mm QFN packages. The PIC16F(LF)1507 and the PIC16F(LF)1508/9 come in 20-pin SSOPs, PDIPs, SOIC packages, and 4×4-mm QFN packages. Prices start at 37 cents each (10,000).—by Colin Holland

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


Name
Dr. Laurel Watts
Job Title
Principal Software Engineer
Area of Expertise
Chemical Engineering
LabVIEW Helped Me
Control multiple instruments operating in harsh conditions
Latest Project
Engineer the ultimate storm chaser
NI LabVIEW
LabVIEW makes me better because the
INTEGRATION
with hardware is so seamless

VOICES

Avnet's Harley Feldberg: embedded in the electronics supply chain

Harley Feldberg, president of Avnet Electronics Marketing, global, spoke with *EDN* following his recent ARM TechCon keynote presentation, "Staying Ahead of the Technology Curve." Feldberg discussed the changing nature of the electronics supply chain, distribution, and the semiconductor market, as well as the new Embedded Software Store (embeddedsoftwarestore.com), an information and e-commerce-based Web site focusing on the embedded-design community. Avnet and its partner, ARM, announced the initiative at the Santa Clara, CA, event. Excerpts of that interview follow.

The topic of your keynote was the changing, evolving market. And the market is constantly changing, more noticeably since the economic shifts that began in late 2008. Where is a distributor like Avnet in all of this, and how does the partnership with ARM fit into long-term strategy?

A Although distribution covers about 25% of the $300 billion global semiconductor TAM [total available market], we support about 90% of the customers. Distribution's role is the mass market—that broad amount of small and midsized customers in the tens of thousands. Simultaneously with that large market position that global distribution fills, ARM has been having tremendous success in the last couple of years in expanding its market position out from its traditional, legacy success in smartphones and tablets. We are coming together with ARM [because] its goal is [for] the architecture that originated in those couple of technologies to permeate out to a much wider, broader customer base; that is, the intersection of ARM's aims and desires and Avnet's market position.

The launch of the Embedded Software Store is really the outgrowth of our having been working more and more with ARM on different initiatives. One thing that has become very clear to both of us is that what's a little different about the embedded space is that the companies developing software to support embedded designs are generally small and dispersed all around the world. If you use my example of distribution's serving 90% of the customer base, the same phenomenon exists here. A main motivation behind this joint initiative between us is to create a more efficient connection between the tens of thousands of customers and the hundreds or maybe even thousands of small embedded-software developers.

Being that our intent is not to become a software distributor but rather to continue to be a hardware distributor, we recognize that supporting software is critical to selling hardware. This [site is] all part of ARM's desire to make it even easier to choose ARM as the design architecture by creating some efficiency on the software-development side. It fits Avnet's goals because our goal is always to sell more components, more semiconductors.

How are you feeling about the electronics supply chain overall? Many are asking the classic "Are we recovered?" as we move into the fourth quarter and see various financial and analysts' reports that call supply-chain matters into question.

A It's not an industry secret that things have clearly moderated. The $64,000 question out there is, Is this a slowdown, a pause, or just lower demand? Or are these recessionary trends we are seeing? Our opinion continues to be that this [situation] is not a 2009 redo. The things we track—like excess inventory, cancellations, schedules—have remained moderate. They are higher than anyone would like, but that is primarily because demand has come down not because inventory is piling up. My opinion is that although the technology industry has traditionally grown … at a rate higher than the GDP, it's not immune to the effects of the GDP. I think we've hit a soft patch for a while. I don't see any reasons why things are radically going to change tomorrow, but we are running our business frugally and making sure we make good business decisions. Overall, there is $300 billion of something out there, and it's not going to $200 billion, but the growth is going to be muted for a while.

Are we talking about a change in the global market, not just a slowdown in demand?

A Because of the nature of my job, I get to see all regions. If your job was only America for the last five years, [you might feel as if we] were in a recession. I have felt for a while that these [changes] are the natural impacters from what I would call a global realignment. There is still a lot of product being consumed, but it's not all in the same spot anymore. The long-term winners in distribution have to be global because you have to follow where the business goes. If you fixate on a couple of quarters [of semiconductor-industry growth], it may cause you to make short-term decisions. Clearly, it's going to be awhile before electronic and semiconductor proliferation goes down, but we are not immune to global impacters. **—interview conducted and edited by Suzanne Deffree**


Name
Peter Simonsen
Job Title
Design Engineer, Embedded Software
Area of Expertise
Renewable Energy
LabVIEW Helped Me
Perform real-world simulations with total control of the application
Latest Project
Develop a test architecture for verification of wind turbine control systems
NI LabVIEW
LabVIEW makes me better because I can
SIMULATE
real-world systems

BY BONNIE BAKER

Designing with temperature sensors, part three: RTDs

The temperature coefficient of an RTD (resistance-temperature-detector) element is positive. Most stable, linear, and repeatable RTDs are of platinum-metal construction. You can use the constant 0.00385Ω/Ω/°C to approximate the resistance change over temperature for the platinum RTD element. In contrast, the NTC (negative-temperature-coefficient) thermistor has a negative change with increasing temperature. See a comparison of resistance and temperature performance for RTD sensors and NTC thermistors in the Web version of this article at www.edn.com/111117bb.

The RTD element's resistance is much lower than that of an NTC thermistor element, which ranges to 1 MΩ at 25°C. Typical specified 0°C values for RTDs are 25Ω to 1 kΩ. Of these options, the 100Ω platinum RTD is the most stable over time and linear over temperature.

An RTD element must be excited with a stable current reference at a level that does not create an error due to self-heating. A current source that is 1 mA or less is usually adequate. Under this circumstance, the accuracy of an RTD can be ±4.3°C over its temperature range of –200 to +800°C. If higher accuracy is required, you can use the Callendar-Van Dusen equation to generate a look-up table: $R_{RTD(TA)}=R_{RTD(T0)}[1+aT_A+bT_A^2+cT_A^3(100-T_A)]$, where $R_{RTD(TA)}$ is the resistance of the RTD at ambient RTD temperature; $R_{RTD(T0)}$ is the value of the RTD at 0°C; and a, b, and c are constants, supplied by the RTD vendor.

You can implement an RTD signal-conditioning circuit in a number of ways. **Figure 1** shows an example that uses four OPA334 amplifiers, an REF5025 voltage reference, an ADS8634 ADC, and an MSP430C1101 microcontroller, all from Texas Instruments, as well as a PT100 RTD (**Reference 1**). In this **figure**, a 2.5V reference, A_1, A_2, and five resistors generate a 1-mA current source.

The signal-conditioning portion of the circuit includes A_3 and A_4. A_3 senses the voltage drop across the RTD element and cancels the RTD wire resistance errors: R_{W1}, R_{W2}, and R_{W3}. A_4 provides gain, and a lowpass filter, such as TI's FilterPro, provides the RTD's output voltage (**Reference 2**). In this circuit, the RTD element has a value of 100Ω at 0°C. If this RTD senses temperature over its entire range of –200 to +600°C, the RTD would provide a nominal 23 to 331Ω range of resistance. You can use TINA-TI to simulate the analog portion of this circuit (**Reference 3**). Within TINA-TI's examples, under the Files tab, a PT100 RTD element accurately simulates the correction of the nonlinearity of the RTD.

The circuit in **Figure 1** generates a current source that is ratiometric to the voltage reference. The ADC uses the same voltage reference to provide a ratiometric digital output. Over temperature, the ADC digitizes the changes in the RTD resistance. Although an RTD requires more circuitry in the signal-conditioning path than a thermistor or a silicon temperature sensor requires, it ultimately provides a high-precision, relatively linear result over a wider temperature range. If you use the Callendar-Van Dusen equation, this RTD circuit can achieve ±0.01°C accuracy. EDN

+ Read parts one and two of this series at http://bit.ly/rpSnOp and http://bit.ly/s6Llbu, respectively.

+ Go to www.edn.com/111117bb for the references cited in this column.

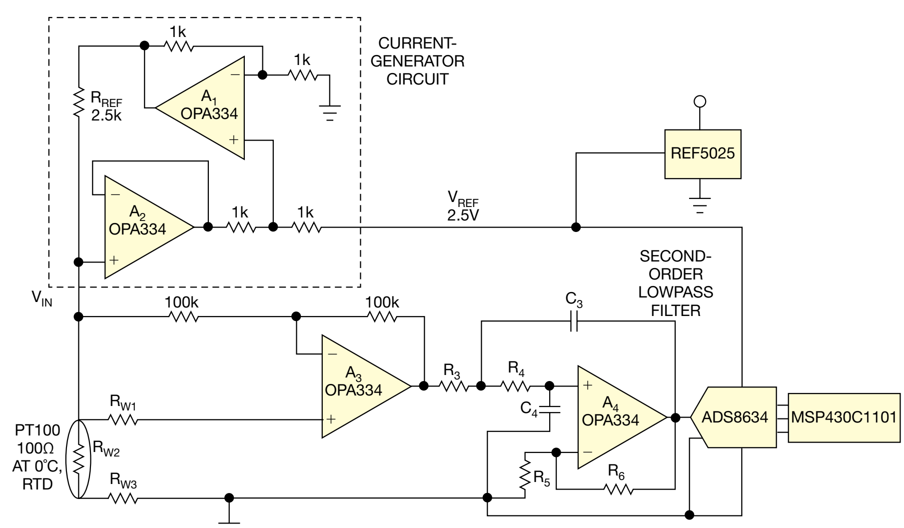

Figure 1 This implementation of an RTD circuit uses four amplifiers, a voltage reference, an ADC, a microcontroller, and a PT100 RTD.

PRY FURTHER AT EDN.COM
+ Watch a video of the inner workings of the sprayer at http://bit.ly/u6s78y.
+ Go to www.edn.com/prying eyes for more Prying Eyes write-ups.

The Scrubbing Bubbles power sprayer

Two AA alkaline batteries power this SC Johnson & Sons spray bottle. The bottle's motors and gears make a bit of noise, but the pump works fine as it dispenses the contents of the bottle. You can repeatedly refill the bottle, and plenty of battery capacity remains after pumping out a bottle of fluid. The design of the plastic housing includes adequate structural bracing and no thick sections to cause dimpling on the surface. Internal components include a microswitch and steel shafts in appropriate places; seven screws, all of the same length, hold the unit together. All of the internal components, except for a sliding trigger lock, fit into the bottom clamshell; the trigger lock mounts in the opposing clamshell. The position of the trigger lock is not critical to assembly. The clamshells fit, no matter what the position of the trigger lock.

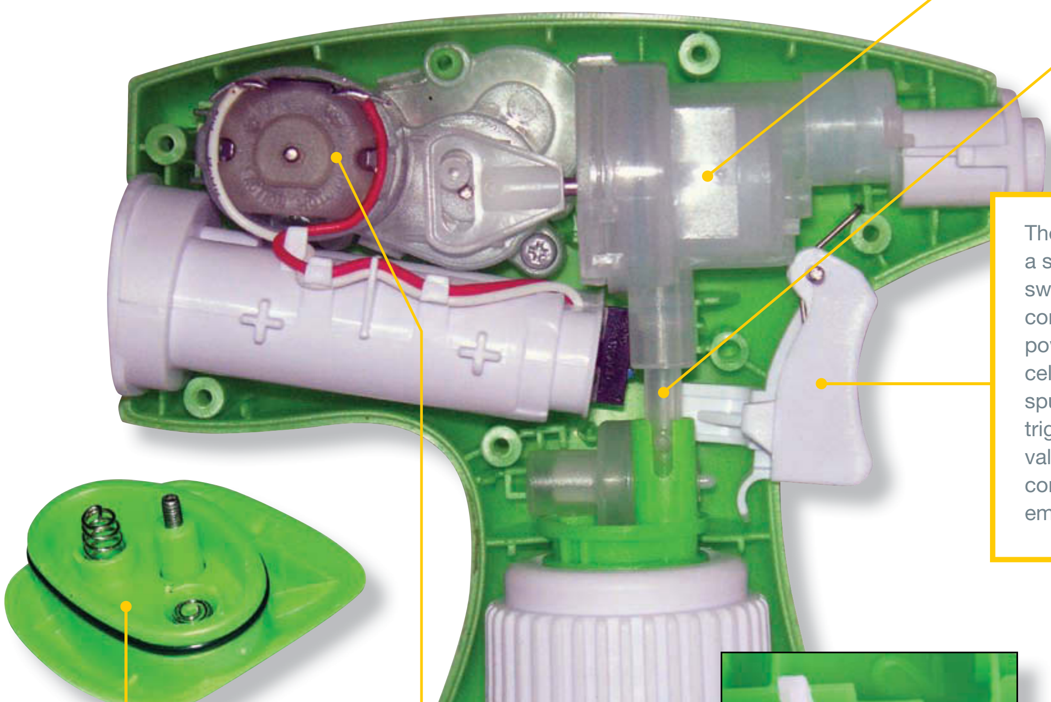

The two green cover clamshells encapsulate the pump assembly. A slotted receiver that fits over the gearbox crank drives the pump piston.

A free-floating plastic tube serves as the fluid's one-way check valve.

The trigger presses on a small, black microswitch behind the battery compartment that feeds power from the two AA cells to the motor. The spur on the bottom of the trigger operates a vent valve to allow air into the container as the pump empties the fluid.

The battery-compartment cap has an O ring to keep out moisture. The self-retaining screw has a gasket under the head. The manufacturer ultrasonically welds in a flat piece that holds the battery terminals.

One screw holds the motor and gearbox assembly onto the outer housing. Molded-in tabs in the battery compartment do a good job of wire retention.

The opposing cover has a sliding trigger lock. It works like a hammer block on a revolver. When you slide it downward, it blocks the movement of the trigger. The outside button snaps into the slide, which has a molded-in spring.


F-35
AA-1

UBM Electronics
design news
EDN
Test & Measurement World

MECHATRONICS IN DESIGN

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

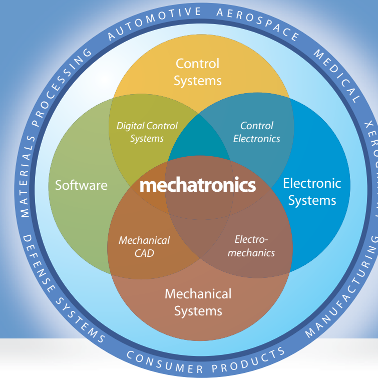

Visualizing fundamental design principles

With practice, these principles become transparent in a design.

Kevin C Craig, PhD, is the Robert C Greenheck chairman in engineering design and a professor of mechanical engineering at Marquette University's College of Engineering. For more mechatronic news, visit mechatronics zone.com.

When I was studying under Vittorio Castelli, a professor at Columbia University and a senior research fellow at Xerox, I observed, listened, and learned. His understanding and insight brought fundamental design principles to light in both what was and what was yet to be. For more than 30 years, "Rino," as we called him, guided and inspired me and others as a mentor, an educator, and an inventor with unbridled energy and passion. Mentoring is a key element in fostering innovation. Each one of you can be that mentor for a young engineer or student. What are these fundamental design principles, and how can a mentor ingrain them in an individual?

When viewing a design or creating a concept to solve a need, fundamental principles as images guide designers to achieve what they thought was impossible. As people have increasingly traded breadth of knowledge for depth of knowledge, awareness of these principles has diminished. Fundamental principles are important, however. Many of the principles that follow come from **references 1** and **2**.

First, remember the laws of nature. Predict before you build. Understanding the basic laws of nature is essential to knowing the fundamental limitations of a design, predicting how a design will perform, and knowing how to improve a machine. Second, consider simplicity versus complexity. Create designs that are explicitly simple. Keep complexity intrinsic, buried, and invisible. The less thought and less knowledge a device requires for production, testing, and use, the simpler it is.

Next, use exact constraints when designing precision structures and mechanisms; that is, apply just enough constraints—no more and no fewer—to define a position or motion. Controlled compliance can make an overconstrained design more stable, however. For example, a five-caster chair can improve load bearing, and a multiple-ball bearing can compensate for geometric errors. Also, plan load paths in parts, structures, and assemblies. Keep them short, direct, symmetric, locally closed, and easy to analyze—for example, a bicycle handbrake that the rider squeezes rather than pulls or pushes.

Remember that the forces you apply to a structure or mechanism can yield great advantages when they create new useful forces, transform or redirect themselves, balance themselves or loads, and help to distribute loads. Examples include the tubeless tire, left- and right-handed scissors, and a balanced door with an articulated hinge. Also remember to keep the functions of a design independent from one another. Everything in design is a compromise, though, and combining functions sometimes might yield benefits.

The accuracy, precision, and resolution of a machine's components and the manner in which you combine them are the most important factors affecting the quality of a machine. Always identify the directions in which accuracy and precision are most important. Before you consider performance, however, you must think about stability. Marginally stable designs work only on paper. Designs must have adequate stability margins. Beware of buckling of compression members.

Also recall Saint-Venant's Principle, which French elasticity theorist Jean Claude Barré de Saint-Venant described. It states that "the difference between the effects of two different but statically equivalent loads becomes very small at sufficiently large distances from load." In other words, several characteristic dimensions away from an effect, the effect essentially dissipates. If an effect is to dominate a system, you must apply it over several characteristic dimensions of the system.

Finally, manage friction. Friction is always present. Yet how much friction is present and the consequences of its presence are uncertain. Avoid sliding friction and use rolling element bearings whenever possible. EDN

REFERENCES

1 Slocum, Alexander H, *Precision Machine Design*, Society of Manufacturing Engineers, 1992, 0-13-690918-3.
2 Skakoon, James G, *The Elements of Mechanical Design*, ASME Press, 2008, ISBN: 978-0-7918-0267-0.

Innovation & Passion

We've been revolutionizing measurement technology since 1933. The world's first vector network analyzer. The first single-box 67 GHz spectrum analyzer. A vector network analyzer with an extraordinary frequency range of 500 GHz. The first combined phase noise and spectrum analyzer in one instrument. Our list of innovations for the aerospace and defense industries goes on.

DESIGN FEMTOAMPERE CIRCUITS WITH LOW LEAKAGE, PART ONE

CAREFULLY APPLY MATERIALS SCIENCE WHILE USING GUARDS AND SHIELDS TO REDUCE LEAKAGE.

BY PAUL GROHE • TEXAS INSTRUMENTS

Circuits that carry femtoamperes of current have many subtleties that you wouldn't normally consider in the design and layout of conventional circuitry. If you overlook these subtleties, the circuit loses low-end resolution and exhibits drift due to the components, materials, and circuit layout. Knowing the circuit's limitations and leakages and providing ways to minimize or eliminate them will lead to improved circuit performance.

The world below a picoampere is unique and plays by a different set of rules. In this world, even the mechanical parts of the circuit can become parts of the electrical circuit. Designing for operation at subpicoamp and femtoamp levels requires special techniques and compromises that normal current levels don't generally require (**Reference 1**). Unfamiliarity with or neglecting these precautions can result in endless headaches for designers. Electrical engineers will find themselves playing double roles as mechanical engineers.

IMAGE: MASTERFILE

Quality & Reliability

Expect only the best. Products of exceptional quality and services that focus on your needs. We help protect your investment with reliable, long-term customer support and calibration tools for self-maintenance and emulation of legacy products. Comprehensive solutions of the highest standards: that’s our philosophy.

This three-part article guides you through the tricky and unconventional design techniques you need to create successful low-current circuits. This first part defines and describes the designs that carry these low currents. It explains the problems that arise when you design these circuits and examines the application of shielding and guarding methods. Part two will examine how your component selection affects the performance of your low-leakage circuits and discuss how noise creeps into low-leakage designs. Part three will provide detailed PCB-design techniques and show a real-world example of a low-leakage design. It will also describe how to verify the performance of your low-leakage-design techniques.

AT A GLANCE

- Low-current circuits must have low leakage.
- Leakage, noise, and stray capacitance can affect your design.
- You must understand and apply insulator properties in your circuit.
- Use guard rings, shields, and enclosures to protect sensitive circuit nodes.

LOW-CURRENT APPLICATIONS

To put things into perspective, 1A equals 6,241,500,000,000,000,000, or 6.24^{18} electrons/sec; 1 pA, or 1^{-12}A, equals 6.24 million electrons/sec; and 1 fA equals 1^{-15}A, or 6240 electrons/sec. In the subpicoamp world, there are three common enemies: current leakages, noise sources, and stray capacitance. A good low-current design must minimize the effects of these common enemies and strike a balance between optimal performance and product manufacturability. You will need special techniques and materials that may be incompatible with conventional production flows.

These high-impedance circuits often go directly into an amplifier input with no parallel-resistive connections. Examples of these circuits include pH probes, gas-sensor amplifiers, medical sensors, sample-and-hold circuits, and three-amplifier instrumentation amplifiers. The circuits can have input impedances into the teraohm range. A transimpedance amplifier, or current-to-voltage converter, is often used at these low current levels. You see this circuit configuration in noninverting amplifiers, photodetector amplifiers, current-to-voltage converters, and photomultiplier circuits. The amplifier's inverting input node and its feedback elements are critical nodes. The current leakage in this node determines the ultimate accuracy of the device. Higher-current circuits, such as low-frequency filters and logarithmic amplifiers, also benefit from low-leakage-design techniques. They will have extended dynamic range, with improved low-end accuracy and lower drift than nonoptimized designs.

CAUSES OF DISTURBANCE

Dirty PCB traces can cause leakage at low currents. The dirt between the traces or across insulating materials—not the trace or wire itself—causes the leakage, serving as a conductive medium between two conductors. Dry dirt in itself may not cause a problem. A combination of dirt with moisture, salts, and oil, however, becomes conductive. The concept here is simple: Keep things clean.

Moisture is the instigator of most leakage problems. When moisture combines with environmental salts and other contaminants, its conductivity increases. Insulation, PCBs, and other hydroscopic materials absorb the moisture, decreasing the electrical resistance of the materials and leading to increased leakage between the conductors.

Contamination between conductors can also create a galvanic reaction in the presence of the right combination of materials and moisture. Moist and salty dirt between a copper trace and a zinc-plated screw or an aluminum case will generate a current between the materials. This current is detrimental to your measurement and causes corrosion of the materials. Because the moisture level varies over the course of the day, season, and geographical location, it creates a moving baseline leakage that is difficult to remove. These leakages change hourly, weekly, or yearly, depending on the environment and the season.

The particles and moisture in air as it moves over a conductor generate a small charge, so you should protect the input circuit from moving air currents. Make sure that fan-cooling airflow does not blow directly over sensitive nodes. Airflow can also cause dust and moisture to accumulate on the conductors and components.

You must take into account the properties of insulating materials in your design. These materials come into direct contact with low-level signals, usually through the connectors, the supports, or the PCB. In the electronics industry, the most common insulators are fiberglass, glass, ceramic, PVC (polyvinyl chloride), epoxy, and Teflon. Each material has its own weaknesses and strengths. Dry air is a good insulator. Keeping conductors in the air can provide the lowest-leakage results. Air does have a low breakdown voltage, however, which limits this technique in high-voltage applications. PTFE (polytetrafluoroethylene) and FEP (fluorinated-ethylene propylene), more commonly known as Teflon, have the best leakage and high-voltage characteristics of common insulating materials, but they are expensive, soft, and difficult to machine. Teflon PCBs are expensive because of the material and the extra steps the fabrication process requires.

Ceramic, although a good insulator, tends to be piezoelectric. Ceramic self-generates charge when it is subjected to stress or impact. It also readily absorbs moisture if it is not sealed or glazed. Although glass is a good insulator, it displays some of the piezoelectric properties of ceramics. IC packages use a molded glass-epoxy compound that allows for currents lower than 1 fA. Epoxy is an excellent, low-cost insulator; however, it is hydroscopic and can absorb moisture over time. Many components, connectors, and wire insulation use PVC, which can generate charge if flexed or rubbed against another conductor, just as combing your hair can generate current. For this reason, PVC insulation in and around the input circuit should be avoided.

It might seem logical to build the ultimate low-leakage layout entirely on a slab of Teflon. This can be a bad idea, however. Because Teflon is a good insulator, any charge deposited on its surface will slowly dissipate. If a sensitive node is nearby, the accumulated charge will lead to slow settling or drift. A better approach is to cover a large

surface area with a guarded conductive plane. Although this approach seems counterintuitive to the desire for low leakage, you should minimize the use of insulators. Insulation must provide isolation, but using too much of it provides a surface to accumulate extra charge.

For low-voltage circuits, an aluminum standoff topped with a small piece of Teflon insulation works better and is less expensive than using an entire standoff made of Teflon. If the circuit will be handling high voltages, you need a Teflon standoff for its better insulating properties. For ac circuits, the narrow insulator has higher stray capacitance that may cause other problems. As in all analog design, you must consider many trade-offs. PCBs have a large influence on low-leakage design because the PCB material is in intimate contact with all of the circuit nodes. The performance of your circuit is only as good as the performance of the PCB material. As with RF circuits that operate at gigahertz speeds, you should consider the PCB as an active component. Most PCBs' material characteristics and development focus on high-frequency RF designs. Manufacturers gear PCB specifications toward circuits that operate at these speeds. They give a nod to low-current requirements by specifying a volume resistivity. The manufacturer's specifications are for the fresh laminate material before processing—not the finished product. That product comprises a sandwich of laminates, bonding glues, fillers, solder masks, and silk-screen that make up a PCB.

The most common PCB material is FR4 (flame-resistant 4), which comprises epoxy-impregnated fiberglass cloth. Manufacturers compress this epoxy under high pressure to form a solid board. FR4 has good electrical properties, but it is not the most desirable material for low-current circuits. You can improve the performance of FR4 using special layout and circuit techniques.

When performance is more important than cost, you can use exotic Teflon or ceramic hybrids, such as Rogers Corp's Duroid hybrid substrate materials, targeting use in microwave and ultra-high-speed digital circuits. The materials' excellent controlled dielectric properties can result in two- to three-times-lower stray capacitance and leakage than those of FR4, but at a cost two- to five-times higher.

The boards also require special PCB-fabrication processes and etching, which some PCB-fabrication houses may be unable to accommodate. The Rogers soft, bendable 3003 material, which is ceramic-reinforced PTFE, requires backing for mechanical stability. Rogers 5880, a glass-reinforced PTFE, gives the best low current and stray capacitance, but it is brittle and cracks easily. It is possible to create a hybrid board, with advanced materials for the critical layers and FR4 for noncritical layers and mechanical stability. This approach is expensive and requires using an advanced board house, however.

Use caution with solder-mask placement. Although solder masks generally help reduce moisture infiltration into the PCB material, surface-charge problems might arise with large areas of Teflon. A better approach uses a bare-copper guard-plane area around sensitive nodes. To prevent oxidation, either solder-level or plate the bare-copper guard area with gold or tin.

SHIELD, GUARD, ENCLOSE

You wire a metallic shield, case, or enclosure to a ground or a common potential. At high impedances, however, these shields create problems with stray capacitance and leakages. Examine, for example, a circuit with a 2.5V input voltage and with 2.5V across the stray-capacitance and leakage paths (**Figure 1**). The 2.5V across the leakage resistance creates a leakage current, and the 2.5V source voltage charges or discharges the stray capacitance, which takes some time to get through the high source impedance and affects the measurement's settling time.

Guards are important in subpicoamp designs because they can cancel the input-leakage currents and most of the fixture capacitance. You drive the guard to a potential equal to the input-signal level. You apply a buffered output derived from the measurement amplifier. This guard acts as a subshield, surrounding and protecting the input-signal lines. External leakages now flow into the low-impedance guard instead of the input traces (**Figure 2**). This approach yields only a few millivolts of potential difference instead of 2.5V across and smaller current flows through leakage resistor R_{LEAK} and stray-capacitance capacitor C_{STRAY}. As a bonus, guards also reduce the input capacitance through a bootstrapping effect. Performed correctly, this approach can cancel out fixture and cable capacitance. Unfortunately, you cannot cancel out the amplifier's input-stage capacitance.

Locate the input traces and all of the sensitive feedback components on your PCB within the perimeter of the thick copper-guard traces (**Figure 3**). Then, remove the solder mask from this area to reduce surface charges. Buffer amplifier A_2 drives the guard

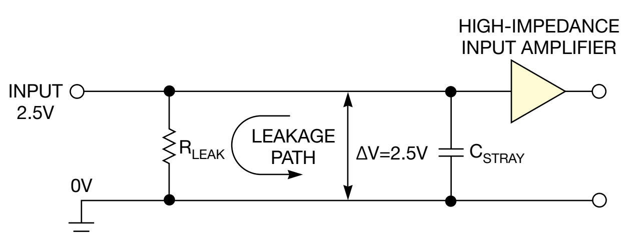

Figure 1 Current leakage and stray capacitance can cause problems in your low-current circuits.

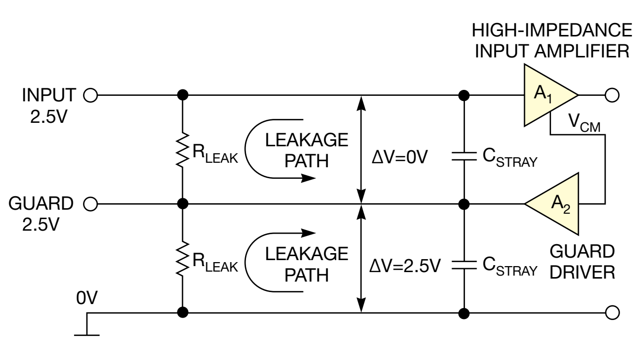

Figure 2 Adding a guard ring between the input node and ground will reduce leakage and capacitive-loading effects.

FOR MORE INFORMATION

Agilent Technologies
www.agilent.com

Keithley Instruments
www.keithley.com

Rogers Corp
www.rogerscorp.com

Trompeter
emersonconnectivity.com

ring. In the inverting and transimpedance designs, you drive the guard to the same potential as the noninverting input's node and feed the potential on the noninverting pin to the guard buffer. The noninverting node is low-impedance, and the buffer does not affect the circuit's operation. The guard should cover the entire input section, the inverting node, and the feedback resistor. Extend it as far into the sensor circuit as possible without affecting sensor operation.

When designing in the noninverting mode, drive the guard to the same potential as the inverting input node through a buffer. This node follows the input signal through the feedback action of the amplifier. Take care that the capacitance of the buffer's input does not cause peaking due to capacitive loading of the inverting node. The guard-driver amplifier should be unity-gain-capable and protected from short circuits and external overvoltages. The bandwidth of the buffer should be slightly wider than the main circuit's bandwidth to reduce phase-lag errors. Avoid a peaking response in the guard buffer to prevent system instability. A grounded shield protects the circuit from external noise and EMI by shunting the noise to ground. Because the grounded shield generally does not follow the input voltage, it does not cancel the capacitance caused by the guard.

In the previous examples, you buffer the guard line from a circuit node using a

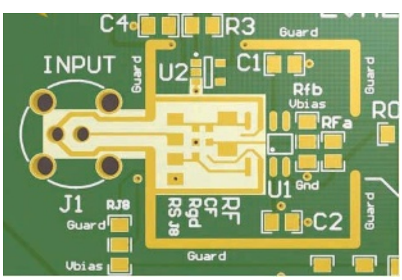

Figure 3 The thick copper trace on this PCB acts as a guard ring. The gold-plated traces prevent corrosion. Remove the solder mask in the guarded area to reduce leakage.

separate amplifier, providing a low impedance to drive the shields and coaxial-cable guards. If you need to guard a small location, you can derive a local guard from the opposite input terminal. Keep in mind that the local guard also adds capacitance to the node to which it connects. This capacitance can lead to peaking in noninverting-amplifier configurations. If the opposite node is high-impedance, the guard can introduce external noise into the summing node unless you shield the guard itself. Do not use an unbuffered guard to drive external circuitry. Use it only for the immediate area surrounding the device.

Keep in mind that the guard is not ground, and ground is generally not a guard. The guard lines should not carry any currents other than the leakages, and you should treat them as signal lines. For effective designs, use guards and grounds together. The guard surrounds the input trace, and the grounded shields protect the guards from external interference. When you lay out a PCB, place a guard plane or guard traces below the sensitive traces. Be careful not to break up the power or ground layer too much. Surround the input circuit in a guarded cocoon using metallic shields on visible component sides and guard traces on layers below the sensitive nodes.

You should enclose your low-current circuits in a sealed environment. If possible, include a desiccant pack to absorb any traces of moisture. The wiring and control shaft entry and exit points should be airtight. You can use triaxial cables and connectors for low-current measurements. The cable contains both an outer grounding shield and an inner guard shield around the center conductor, extending the guard out to the measurement point. Commercial test equipment often uses Trompeter 70-series triaxial BNCs. Agilent prefers the three-lug style, while Keithley prefers the two-lug style. **EDN**

REFERENCE

1. *Low Level Measurements Handbook*, Sixth Edition, Keithley Instruments, 2004.

AUTHOR'S BIOGRAPHY

Paul Grohe is an applications engineer for Texas Instruments' Precision Systems group. He attended the College of San Mateo (San Mateo, CA).

INTELLIGENT

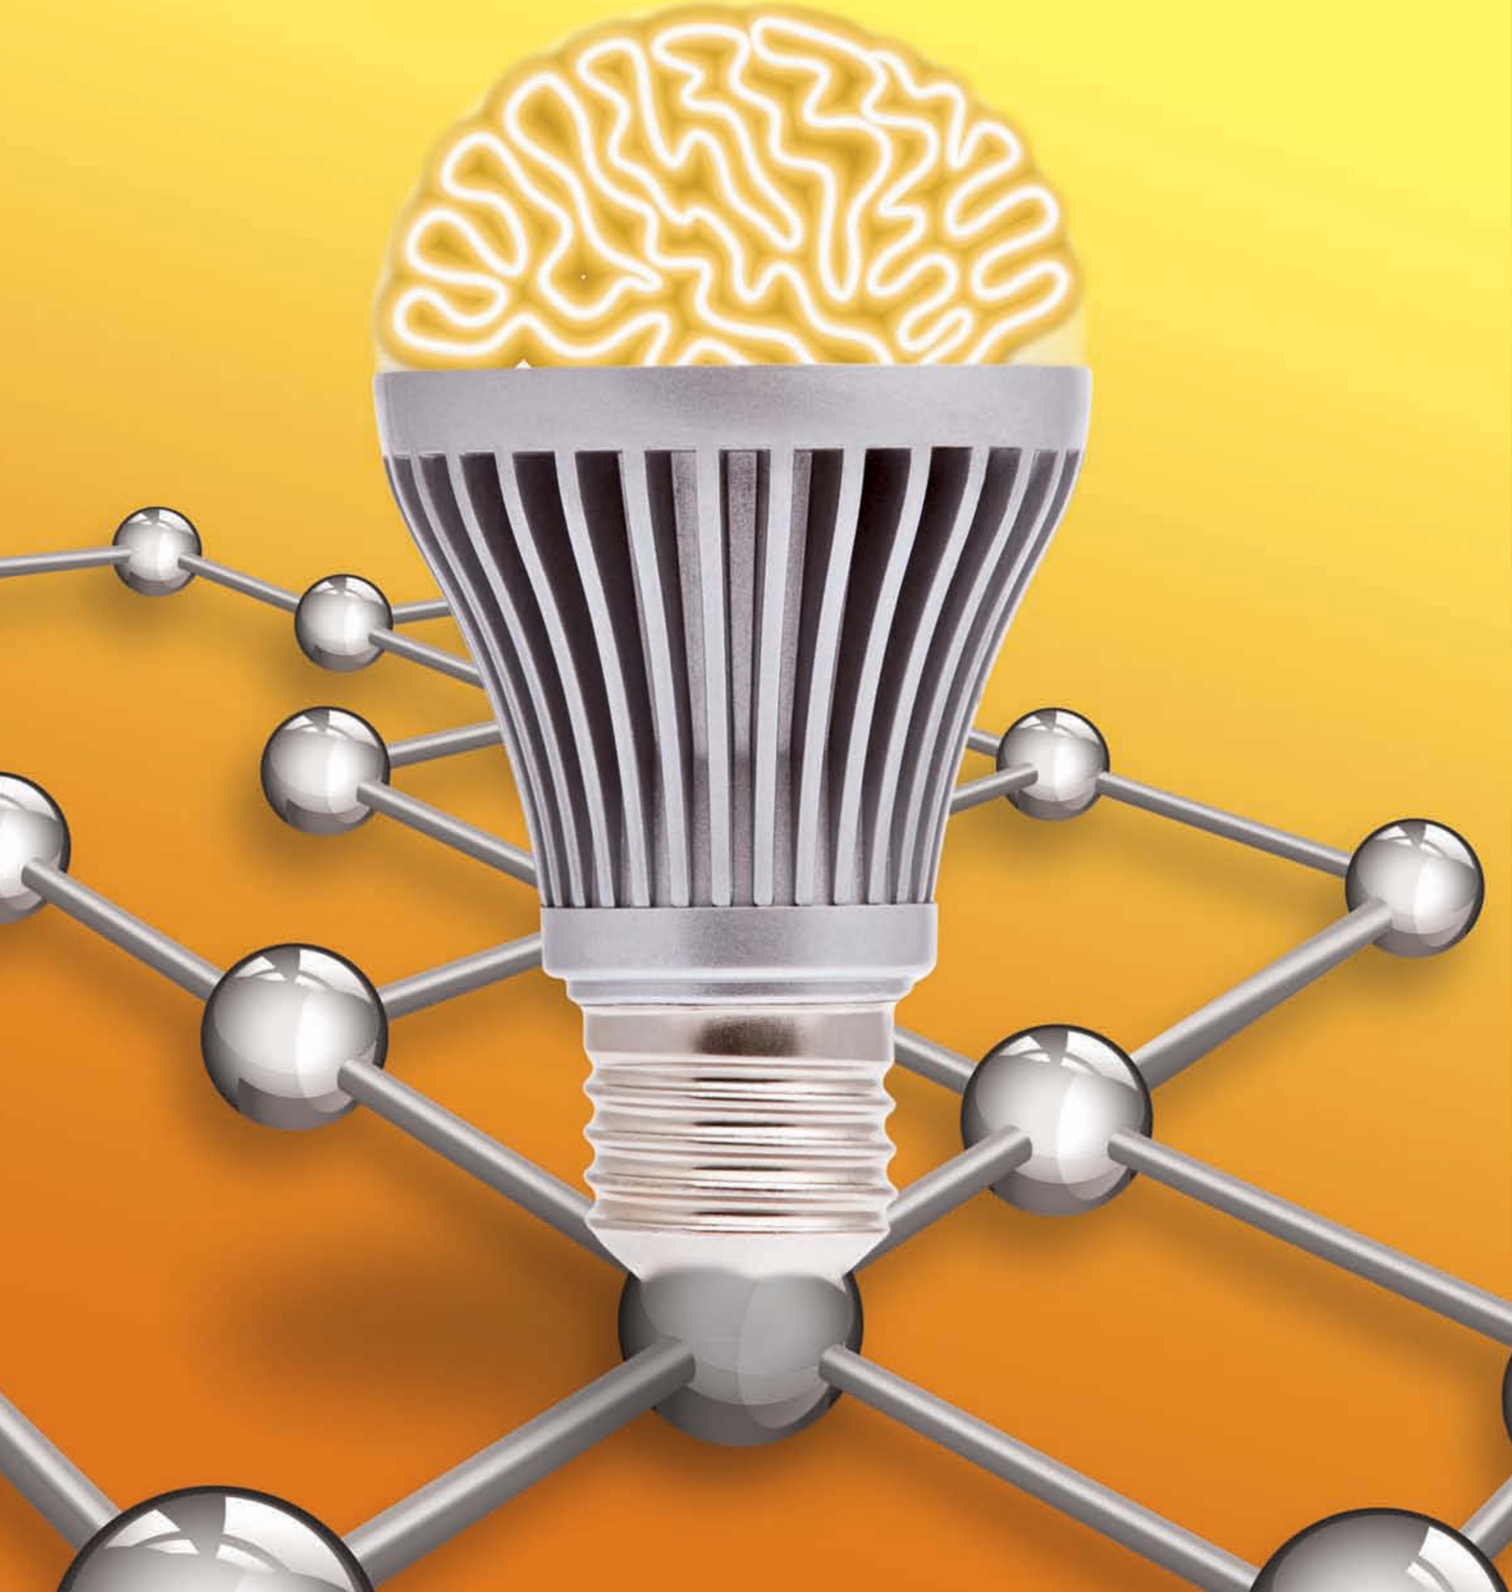

LIGHTING HOLDS PROMISE FOR EFFICIENT LIGHTING NETWORKS

AT THE RECENT DESIGNING WITH LEDs WORKSHOP, A KEYNOTE PANEL OF EXPERTS ADDRESSED THE EMERGING TOPIC OF INTELLIGENT LIGHTING. HERE ARE SOME HIGHLIGHTS OF THE DISCUSSION.

BY MARGERY CONNER • TECHNICAL EDITOR

During the September 27 Designing with LEDs Workshop in Boston, moderator Carol Lenk, chief technology officer of ReliaBulb, and other industry experts addressed the emerging topic of intelligent lighting. Panelists were Eric Holland, vice president of engineering at Lighting Science Group, which last winter partnered with Google for the light bulbs used in the Android@Home demonstration; David Ewing, chief technology officer at Synapse Wireless, which provided the low-cost, low-power wireless network used in 2010's *Tron: Legacy* movie; and Paul Wilson, regional marketing director, Americas, for NXP.

IMAGE(S): THINKSTOCK

ERIC HOLLAND, LIGHTING SCIENCE GROUP

Holland spoke about the value proposition of intelligent lighting and how that idea relates to ROI:

"Wikipedia's entry on intelligent lighting is a bit antiquated, referring to it solely in the context of theatrical lighting. But the entry is still relevant because what you're starting to see now is the migration of more advanced features typical of the theatrical market [such as defining and programming a 'scene' that defines the light color, intensity, and application] over into consumer and commercial infrastructure. Obviously, the value proposition is a lot different, but the objective is the same. Lighting Science serves the consumer, commercial, and infrastructure markets, and we're finding that the needs for each one of those markets vary greatly.

"For example, when utilities talk about intelligence, it's more about monitoring and control. This [idea] was

UTILITIES ARE ASKING FOR 'INTELLIGENT FIXTURES,' AND THEY CAN'T REALLY ARTICULATE WHAT THEY MEAN BY THAT TERM.

—Eric Holland

AT A GLANCE

- Intelligent lighting can shift the focus from the cost of an individual lighting fixture to the cost and feature benefits of a lighting network.
- Make sure that the network is easily software-updatable because features will change.
- Internet Protocol-addressable lights will likely be important in lighting networks.
- Standby power must be low to minimize power costs.

a change for us because we're used to talking about the efficiency of LEDs, but efficiency is still something of a hard sell to the utilities for streetlighting because streetlights are on during off-peak rate times. When utilities look at LEDs, it's still a matter of cost and maintenance: What is it going to take to maintain that lighting?

"A few years back, I was involved in a project with another company where we were doing wide-area monitoring of high-pressure sodium lighting. It was a pretty good value proposition around intelligence for those lights because the failure modes for high-pressure sodium are somewhat tricky. It's hard to predict when they start to fail from cycling; also, you go out there in the daytime to repair it, and all the lights are off anyway. With LEDs, the value proposition for monitoring is a little more sketchy [because LED failure modes are more predictable].

"We're seeing a lot of requirements from utilities for 'intelligent fixtures,' and, when we go back and ask the utilities what they mean by that [term], they can't really articulate a clear answer. More often than not, they just want a 0 to 10V interface so that they can dim it at some point in the future. Occasionally, they can actually specify a system that they want to use with their fixture or exactly what it is they want to do with it, but, in general, they just know that intelligent lighting is coming but not the details.

"In commercial, we're seeing a big value proposition for intelligence, particularly on the savings side. There are a lot of rebates available to building owners and operators to install LED fixtures, and a lot of that [motivation] is driven by decreasing their demand at peak hours. LEDs have the ability to dim in a very efficient manner and also to do occupancy sensing. For parking garages, this [idea] is a very big deal. If you can dim the lights to 50% power, but the perceived light doesn't drop by 50%, you can lower the energy requirements in peak demand for that garage.

"The consumer market is really about coolness and convenience. [Panel moderator] Carol Lenk touched on this [idea] in her introduction. There are the geeks who are willing to put in the new switches, the special devices, [and] the new wiring, but, when we talk about the average consumer and what it's going to take to get massive adoption, it has to be as simple as replacing a light bulb and offering alternatives [in which consumers] don't have to change out their switches. But even more important is, What's the surrounding ecosystem? What does your app look like? Everybody knows home automation hasn't happened. To make intelligent lighting in the home work, it has to have apps. A year from now, we need to be able to look in the Android marketplace and the iOS marketplace and see that they're full of apps for solid-state lighting, with everything organized so that it really doesn't matter whose system is underneath."

DAVID EWING, SYNAPSE WIRELESS

Ewing captured the likely "gotchas" you can expect to encounter in designing an inexpensive, ubiquitous, completely autonomous lighting network:

"Synapse has designed wireless capability into outdoor lighting, and now we're seeing interest in indoor lighting. We're starting to see more intelligence being embedded in both wired and wireless lighting systems, and we're starting to see all the same challenges we've had in our wireless mesh networks.

"There are a lot of choices for control networks for lighting. There are the traditional wired ways of doing it; in Europe, you see a lot of DMX [digital multiplex] and DALI [Digital Addressable Lighting Interface], which is in both Europe and the United States, plus you see traditional industrial protocols, [such as] Modbus.

"Sometimes, our wireless-control efforts are basically cable replacements.

WILL THE SENSORS FUNCTION AS SIMPLE VACANCY DETECTORS, OR WILL THE NETWORK NEED TO KNOW INSTANTLY WHEN SOMEONE WALKS INTO A ROOM?

—David Ewing

This [approach] is really important in retrofit applications, but it can also make a lot of sense for new designs because of the expense of laying the control cable. Wireless is not the only option for intelligent-lighting control, and power-line carrier is one option. However, [power-line carrier has] many challenges: You will need phase couplers, and there are issues with bandwidth and with interference.

"In the wireless space, ZigBee is a familiar approach; it's been around for 10 years. For large networks, there just haven't been many success stories. If you were to try to outfit this convention center with ZigBee-networked lights, there would be problems, [including] latency issues and the commissioning of individual lights and controllers. It doesn't scale well.

"Most of us are familiar with Wi-Fi: That's what most of us see in our everyday dealings with wireless control. But Wi-Fi is expensive, both in its implementation and in its power usage. What's new is 6LoWPAN [Internet Protocol Version 6 over low-power wireless personal-area networks], the 'new hotness.' It's an IP-centric world. Everyone's carrying smartphones. Ultimately, individual light bulbs are going to connect into an IP universe; 6LoWPAN is a new protocol that takes IP Version 6, the next-generation Internet protocol, all the way down to individual lights.

"There are a lot of choices for control mechanisms. DMX has been around for a long time in both theatrical and commercial applications. At the lowest level is 0 to 10V control for simple dimming applications. But generally for intelligent lighting, we need more than a simple 0 to 10V range of commands.

"Within the range of different network protocols, there are several different issues. [One is] battery life: The lights themselves generally have power available, but the sensors may be separate from the lights and away from ac power, which means there are power trade-offs. For example, will the sensors function as simple vacancy detectors and yield a five-year battery life, or will the network need to know instantly when someone walks into the room, which requires more current? People don't necessarily know what they want, so we get to try to make this system easy enough and adaptable enough and flexible enough without making it so complex that they won't know how to use it.

"[Another issue is] commissioning: This [term] refers to associating the address of a particular device with the actual physical device. If you have installers go through and install lights in a building, you can't count on those guys to do any kind of commissioning; you're lucky if they simply get things wired correctly. If you place the onus of commissioning on installers, you're almost invariably doomed to failure. Nobody wants to do commissioning. So a lot of the effort we've been involved in is to make it commissioning-free so the installer with almost no training can test it out, and then we come back through with intelligent location-based technology that allows you to assign lights to locations and define scenes.

"Remote software upgrades [are also important]. All these lights are running software on platforms such as 8051 microcontrollers. You're going to want to upgrade it, especially because this is a new area; we don't know ultimately what the software challenges will be. So it's important to be able to upload new code.

"In addition to these issues, there are also customer expectations in the areas of commercial, industrial, and outdoor lighting. [As part of] cloud-based lighting management, the customer will want to be able to monitor and control a lighting system from a smartphone, tablet, or laptop, including setting schedules, reconfiguring zones, setting scenes, and monitoring power and any sensors included in the hardware. Security is important. Systems will need to use AES 128 [128-bit Advanced Encryption Standard]. You want to feel [as though] you're connecting to your bank; that's how secure it has to be.

"[Systems must also be] scalable: Streetlighting systems will routinely have 5000-plus lights. Things can't break down. [And systems must be] fast: When a customer pushes a light switch, they expect it to come on instantly. So the control signal can't make a trip all the way back to the Internet—a round trip to do control."

PAUL WILSON, NXP SEMICONDUCTORS

Wilson discussed some of the hardware components that intelligent lights require and the importance of power control and efficiency in a successful intelligent-lighting network:

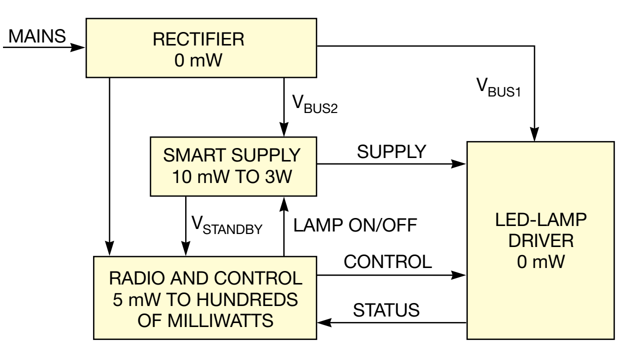

Figure 1 An intelligent LED lamp requires LED-lamp-driver, rectifier, smart-supply, and radio and control components, each drawing different amounts of power.

TABLE 1 WI-FI AND LOW-POWER-RF COMPARISON

Protocol	Wi-Fi	Low-power RF
Number of lamps	200	200
Standby-lamp-radio-power draw (mW)	500	48
Active-lamp-radio-power draw (mW)	1000	48
Standby-radio-adapter/gateway draw (mW)	NA	48
Active-radio-adapter/gateway draw (mW)	NA	48
Total annual standby power (kWhr)	876	84.5
Total annual power at 30% active, 70% standby (kWhr)	1138.8	84.5
Annual standby-power cost at 10 cents/kWhr	$87.60	$8.45
Annual cost at 30% active, 70% standby at 10 cents/kWhr	$113.88	$8.45

"Lighting represents one of the world's greatest opportunities for significant energy savings, as up to 25% of home-energy usage is from lighting. We all know the energy-savings potential of LED lighting in replacing the existing base of incandescent lights. We estimate that another 30% of savings can be gathered by applying [and combining] these LED lights with intelligent lighting: Use light when you really need light. With smart controls, you can do dimming, scenes, profiles, adjustments, monitoring, [and] preventive maintenance. [In addition] intelligent lighting enables participation in utilities' demand-response programs, resulting in reduced tariffs.

"Some of the key elements of an intelligent-lighting network are the switches, sensors, controllers, and the wirelessly enabled smart lamps themselves. If you want to have access to them though smartphones, tablets, and other Internet-connected devices, then you need to have some sort of gateway. I'll talk a little about what's going on inside that wirelessly enabled smart lamp because there is a price to pay in terms of power when you're putting all this smart technology inside a light bulb.

"**Figure 1** shows what's inside a smart LED bulb; ac power goes into a rectifier and feeds into both the LED driver/power supply and a 'smart' power supply for the wireless transceiver. These two parts—the smart power supply and the radio transceiver—are necessary to make the LED bulb both intelligent and wireless.

"One of the gotchas in smart lighting is now that you have the smart power supply, it must be constantly on and drawing power. Even when the switch is in the off position, the light must be listening for controller signals. Fortunately, intelligent light bulbs, which are often centrally positioned in every room, can make ideal network routers, but they do draw continuous amounts of power even when off, so standby power can impact system power efficiency.

"Standby power can vary dramatically based on power-supply topology. For example, a low-cost linear supply can consume as much as 3W; in a 13W LED you'll negate the power savings of using an LED lamp, so that's not a really smart way of doing it. … A better choice is a buck topology with about 10 mW of standby power. State of the art

FOR MORE INFORMATION

Advanced Encryption Standard
http://csrc.nist.gov/publications/fips/fips197/fips-197.pdf

Digital Addressable Lighting Interface
www.dali-ag.org

Digital Multiplex
www.dmx512.com

Google
www.google.com

Lighting Science Group
www.lsgc.com

Modbus Organization
www.modbus.org

NXP Semiconductors
www.nxp.com

ReliaBulb
www.reliabulb.com

6LoWPAN
http://datatracker.ietf.org/wg/6lowpan/charter

Synapse Wireless
www.synapse-wireless.com

Wi-Fi Alliance
www.wi-fi.org

ZigBee Alliance
www.zigbee.org

for a radio transceiver for the wireless portion is about 17 mA of current. The networking stack you choose has an impact, as well. David [Ewing] mentioned 6LoWPAN. I think the key is to get the code size small and then keep the RF-transceiver- and microcontroller-power needs as low as possible. Using a 'mostly asleep' duty-cycle design where the transceiver and microcontroller don't have to be constantly on also helps cut down on standby power.

"As an example of how real this issue is, a 10W bulb on for four hours a day uses 40 Whr a day. If the standby power for the device is 1W ... it will consume about 20 Whr a day; 33% of the power is consumed in standby. At 100 mW, it falls to 2 Whr, or just 5% of the electricity bill at 100 mW standby power; NXP demonstrated at LightFair last May its 6LoWPAN chip set operating with a duty cycle of about 10% on and listening for control signals, resulting in a standby power of about 30 mW and reducing standby power to negligible amounts.

IN AN INTELLIGENT-LIGHTING SYSTEM, USERS CONTROL DIMMING FROM A SMARTPHONE.

—Paul Wilson

"Because virtually everyone has Wi-Fi in their homes, why not just use Wi-Fi inside the bulbs? Wi-Fi is a fine system and supports a high data rate, but it's also very-high-power. When you compare Wi-Fi to a low-power, 30- to 70-duty-cycled RF technique, Wi-Fi costs about $113 to operate; low-power RF is $8.45 [**Table 1**]. There's also a problem with [the fact that] Wi-Fi is more of a star network than a mesh, and getting across the network can be a problem.

"Dimming control becomes much simpler when designing for smart lighting networks. Currently, LED lamps must be able to dim with existing TRIAC [triode-alternating-current] phase-cut dimmers, and including this circuitry in an LED lamp costs about 50 cents. In an intelligent-lighting system, users control dimming from a tablet, a smartphone, or a wireless device. Eliminating the TRIAC-dimmable circuitry will help offset the cost of including wireless circuitry and intelligence in LED lamps.

"Once a lamp is 'smart,' many other functions can be implemented with relative ease, such as color-mixing and color-temperature control, brightness, power monitoring, end-of-life prediction, and occupancy response. The sell of LED lighting is no longer the cost of the bulb; intelligent lighting changes it to a system sell and system benefits."EDN

You can reach Technical Editor Margery Conner at 1-805-461-8242 and margery.conner@ubm.com.

BY SVEN BREHMER, POLYCORE SOFTWARE INC • MARKUS LEVY, THE MULTICORE ASSOCIATION • AND BRYON MOYER, INDEPENDENT CONSULTANT

Using MCAPI to lighten an MPI load

USE MCAPI TO LESS EXPENSIVELY DELIVER MPI PERFORMANCE IN A SYSTEM WITH BOTH LIMITED RESOURCES AND LIMITED REQUIREMENTS.

HPC (high-performance computing) relies on large numbers of computers to perform tough jobs. One computer often acts as a master, parceling out data to processes that may be located anywhere in the world. The MPI (message-passing interface) provides a way to move the data from one place to the next. Normally, MPI would be implemented once in each server to handle the messaging traffic. With servers using many cores, however, it can be expensive to use a complete MPI implementation because MPI would have to run on each core in the computer in an asymmetric-multiprocessing configuration. On the other hand, the MCAPI (Multicore Communications API)—a protocol designed with embedded systems in mind—more efficiently moves MPI messages around the computer.

HEAVYWEIGHT CHAMPION

The well-established MPI HPC protocol is robust enough to handle the problems that might be encountered in a dynamic network of computers. For example, such networks are rarely static. MPI must be able to handle a variable number of nodes in the network—due to updates, maintenance, the purchase of additional machines, or even a user's inadvertent unplugging of a physical network cable. Even with a constant number of servers, those servers run processes that may start or stop at any time. MPI thus includes the ability to discover who is on the network.

At the programming level, MPI reflects nothing about computers or cores. It knows only about processes. Processes start at initialization, and this discovery mechanism builds a picture of the arrangement of the processes. MPI allows for flexibility in the creation of a topology. When everything is up and running, however, a map of processes can be used to exchange data. A given program can exchange messages with one process inside or outside a group or with every process in a group. The program itself has no idea whether it is talking to the computer next to it or to one on another continent. A program doesn't care whether a computer running a process with which it's communicating is single-core or multicore, homogeneous or heterogeneous, or symmetric or asymmetric. It knows only that it wants to send an instant message to a process. The MPI implementation on the computer must ensure that the messages reach the targeted processes.

Due to the architectural homogeneity of symmetric multicore implementations, achieving this goal is simple. An OS instance runs over a group of cores and manages them as a set of identical resources, naturally spreading a process over the cores. A multithreaded process can take advantage of the cores to improve computing performance; nothing else must be done.

However, symmetric multiprocessing starts to bog down with more cores because adding cores also bogs down bus and memory access. For computers designed to help solve big problems as quickly as possible, it stands to reason that more cores in a box is better, but only if the computer can effectively use them. To avoid the limitations of symmetric multiprocessing, you can instead use asymmetric multiprocessing for systems with multiple cores.

With asymmetric multiprocessing, each core or subgroup of cores runs its own independent OS instance, and some might even have no OS at all, running on "bare metal." Because a process cannot span more than one OS instance, each OS instance and, potentially, each core runs its own processes. So, whereas a symmetric-multiprocessing configu-

TABLE 1 MPI/MCAPI COMPARISON

	MPI	MCAPI
Topology	Dynamic	Static
Coupling	Loose	Tight
Locality	Not local	Local
Timing	Loose	Tight
Resources	Ample	Limited

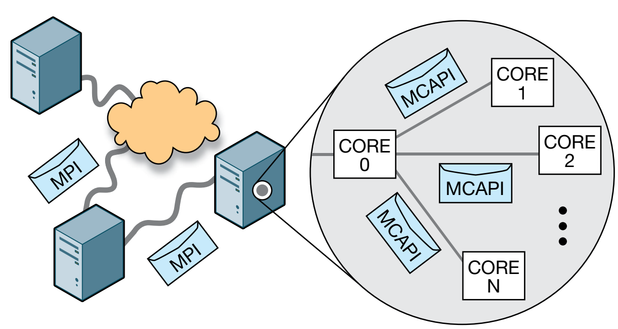

Figure 1 The accelerator cores run MCAPI instead of MPI, meaning that MPI messages run between the servers but MCAPI messages run between the cores in the server.

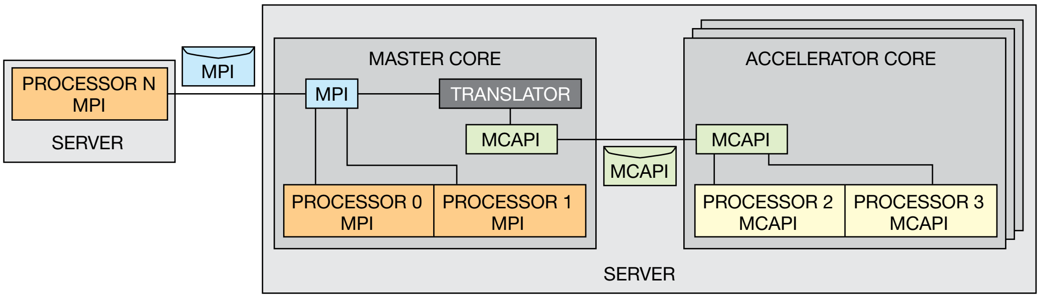

Figure 2 A translator converts any messages moving between the MPI and the MCAPI domains.

ration can still look like one process, asymmetric multiprocessing looks like many processes, even if they are multiple instances of the same process.

In this configuration, each OS must run its own instance of MPI to ensure that the network represents its processes and feeds it any messages coming its way. The environment connecting the cores within a closed box—or even on one chip—is smaller than the network within which MPI must operate. It also typically has fewer resources than a network does. MPI thus has too many features for communication within a server.

DIFFERENT ROLES

Although they may look similar in spirit, MPI and MCAPI play different roles. MPI comes from the HPC world; MCAPI, from the embedded-system world. They thus have different characteristics, including topology, coupling and locality, resources, and timing, which are complementary to each other (**Table 1**).

The network over which MPI runs may change configuration at any moment either physically or by starting and stopping processes. In contrast, an embedded system is static. For the most part, it is physically impossible to disconnect the components in an embedded system. Even when you use something like a PCI card to add computing power, it's not a plug-and-play configuration: The PCI slot makes it possible to add a board, but, once added, it's generally expected that the board will remain there. Thus, MCAPI doesn't need the performance to deal with topology changes.

Coupling refers to the strength with which two systems interconnect. Networks are loosely coupled, so breaking the network shouldn't affect a computer's ability to function, except to the extent that, if it needs something across the network, it can't get it, and you must hope that the programmer created a graceful way to handle this situation. At the other end of the scale, an embedded system is typically restricted to one box. If the system has multiple cores, the cores of the processor connect tightly because they share a hard-wired bus, perhaps some memory, and the same silicon crystal.

Coupling closely ties to the concept of "locality": A network may connect you to a computer halfway around the world; two cores are typically separated by microns. Whereas MPI must handle loosely coupled nonlocal nodes, MCAPI can assume tight coupling and close proximity.

The resources available to handle message passing also scale as you go from the network level down to the processor. It's a straightforward matter to add storage to a network; it's impossible to add on-chip RAM to a processor. The storage you can add to a network is huge; the fixed storage on a processor is limited. Thus, the resources available for managing MPI tend to be greater; MCAPI must operate on a budget.

Response time is also a consideration. Moving a message around the world takes time, and that time is not deterministic. Send the same message multiples times, and it will take different routes that have different delays. By contrast, many embedded systems have stringent real-time requirements that must be met. Milliseconds matter. MCAPI can therefore be quick and responsive in a way that MPI can't be.

A FEATHERWEIGHT STEPS IN

Unlike with MPI, The Multicore Association designed the MCAPI specification to be lightweight so that it can handle interprocess communication in embedded systems, which usually have considerably more limited resources. Although MCAPI works differently from MPI, it still provides a basic, simple means of getting a message from one core to another. You can thus use MCAPI to less expensively deliver MPI performance in a system with both limited resources and limited requirements.

To bring MCAPI into an MPI design, consider a program using MPI, which uses few MPI constructs that just send and receive simple messages. The idea is to designate one master core within the server to run a full MPI service plus a translator for all other accelerator cores in the box. The accelerator cores will run MCAPI instead of MPI, meaning that MPI messages will run between the servers but MCAPI messages will run between the cores in the server (**Figure 1**).

For those program instances running on the accelerator cores, you then replace the MPI calls with the equivalent MCAPI calls. For that reason, this approach works only for simpler uses of MPI; many MPI constructs have no MCAPI

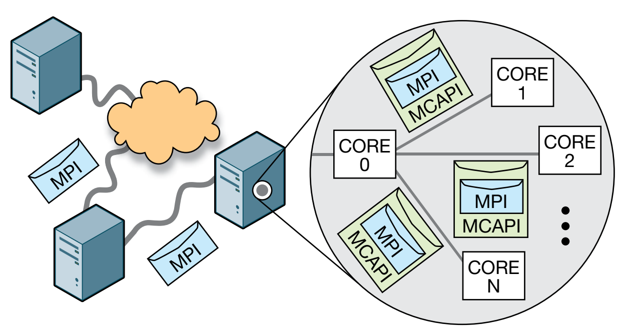

Figure 3 One approach is to keep all of the original MPI calls in the program for the master and the accelerator cores and then wrap the MPI messages in the MCAPI messages to get them to the cores.

DesignCon® 2012

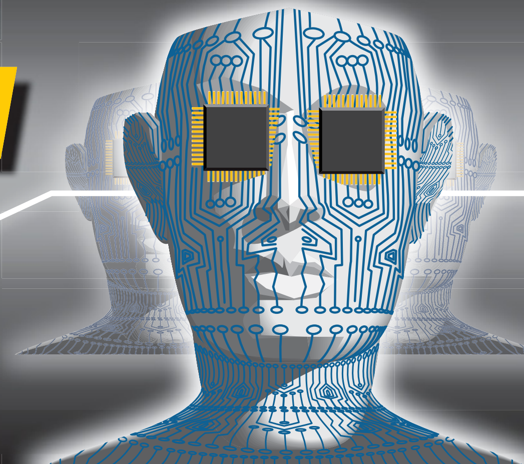

WHERE CHIPHEADS CONNECT

Conference: January 30 - February 2
Exhibition: January 31 - February 1
Santa Clara Convention Center | www.designcon.com

Don't miss the definitive event for chip, board, and systems designers.

Join thousands of engineering professionals who make the decision to start the year off right with DesignCon!

INDUSTRY TRACKS:

- Chip-Level Design for Signal/Power Integrity
- Analog and Mixed-Signal Design and Verification
- FPGA Design and Debug
- System Co-Design: Chip/Package/Board
- PCB Materials, Processing and Characterization
- PCB Design Tools and Methodologies
- Memory and Parallel Interface Design
- High-Speed Serial Design
- High-Speed Timing, Jitter and Noise Analysis
- High-Speed Signal Processing, Equalization and Coding
- Power Integrity and Power Distribution Network Design
- Electromagnetic Compatibility and Interference
- Test and Measurement Methodology
- RF/Microwave Techniques for Signal Integrity

KEYNOTE SPEAKERS:

Tuesday, January 31

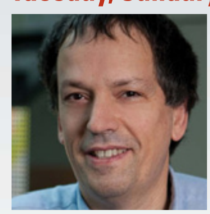

Ilan Spillanger
VP Hardware and Technology, Interactive Entertainment Business Unit, **Microsoft**

Wednesday, February 1

Prith Banerjee
SVP Research, **Hewlett Packard** *and Director of* **HP Labs**

Register today at www.designcon.com.
Use promo code PRINT to save 15% off any conference package. Expo registration is FREE. The first 50 people to use promo code PRINT will get last year's conference proceedings for free.

OFFICIAL HOST SPONSOR:

Agilent Technologies

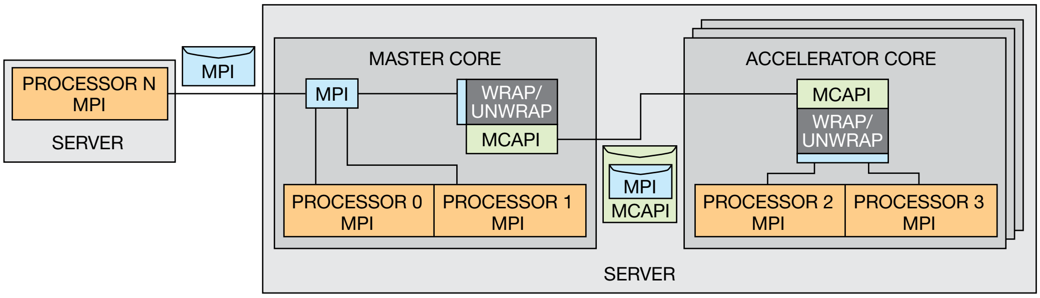

Figure 4 The approach in Figure 3 requires MPI on the master core, MCAPI on the master and accelerator cores, and a small wrapping and unwrapping facility to tunnel the MPI messages over MCAPI.

equivalents. A translator converts any messages moving between the MPI and the MCAPI domains (**Figure 2**). The cost of this arrangement lies in the fact that the program must be edited and recompiled to use MCAPI instead of MPI for the accelerator cores. This approach also complicates program maintenance due to the existence of two versions of the program: one using MPI and one using MCAPI.

Alternatively, you can keep all of the original MPI calls in the program for the master and the accelerator cores and then wrap the MPI messages in the MCAPI messages to get them to the different cores (**Figure 3**). To make this approach work requires MPI on the master core, MCAPI on the master and accelerator cores, and a small wrapping and unwrapping facility to tunnel the MPI messages over MCAPI (**Figure 4**). The trick is that this wrapping service must replicate the MPI API, even though it's simply stubbing out the actual MPI functions. The wrapper then drops onto the accelerator cores as a library that masquerades as the MPI library so that the processes running in the accelerator cores feel as if the system is meeting their MPI needs. In the master core, the wrapper must additionally represent the processes running on the accelerator cores so that the MPI messages route properly when going to the accelerator cores.

A message can follow several paths. For example, Processor N sends a message to Processor 0. In this case, Processor 0 is running on the master core with full MPI service, so this scenario can be handled as a standard MPI message. In another case, Processor N sends a message to Processor 3. Here, Processor 3 is running on an accelerator core, so the master, which recognizes that Processor 3 is on another core, receives the message. The master wraps the MPI message in an MCAPI message and sends it to the other core. The accelerator core accepts the MCAPI message and unwraps the MPI message, which is now available to Processor 3.

In a third scenario, Processor 1 sends a message to Processor 2. In this case, both processes are in the same server but on different cores. The master core wraps the MPI message in an MCAPI message, which the accelerator core unwraps for consumption by Processor 2. Another case has Processor 2 sending a message to Processor 3. Here, both processes are on the same core, but the core doesn't have MPI running. So the message goes to the full MPI implementation on the master core over MCAPI, which routes the message right back. This approach may sound inefficient, but it's faster than having the messages cross the Internet between servers.

In yet another alternative, Processor 2 sends a message to Processor N. The accelerator core wraps the message as MCAPI and sends it to the master core. Once the master core unwraps the message, the MPI service can route the message to the other server. Again, only one core needs to run MPI; the other cores run MCAPI. Although MCAPI assists in moving MPI messages, the processes exchanging MPI messages have no idea that anything but MPI is running. Using any of these approaches, MCAPI lets you effectively use the extra cores in modern servers in an HPC network. EDN

ACKNOWLEDGMENT

This article originally appeared on EDN's *sister site,* MCU Designline, *http://bit.ly/oYNwWO.*

AUTHORS' BIOGRAPHIES

Sven Brehmer is president of Polycore Software Inc.

Markus Levy is president of The Multicore Association.

Bryon Moyer is an independent consultant.

Factory Hardened

Our signal-chain solutions withstand the toughest conditions

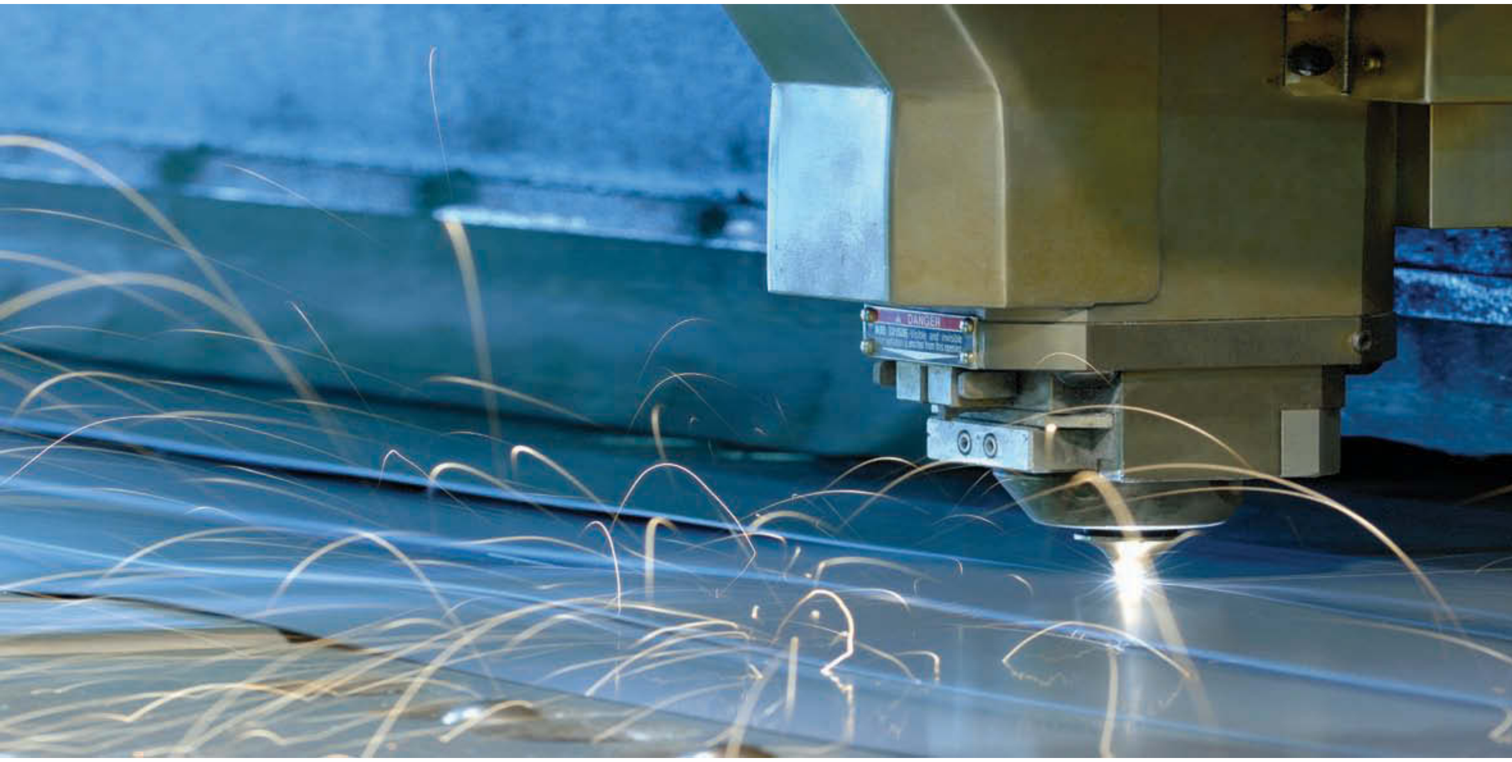

To withstand anything, you can't just use anything. Maxim's industrial signal-chain solutions are built from the ground up to withstand extreme temperatures and destructive electrical events. Industry-leading ESD protection, latchup immunity, and fault tolerance are built in. So you get a smaller, simpler solution that is tough enough for the factory floor.

Get the Toughest Signal-Chain Solutions

- Data interfaces offer up to ±35kV ESD protection to prevent damage to industrial fieldbuses
- High-voltage op amps withstand ±8kV ESD and increase latchup immunity
- 72V analog multiplexers eliminate external protection circuitry
- Multichannel, simultaneous-sampling delta-sigma ADCs provide redundant paths for increased fault tolerance

See for yourself. Order free samples of these products today.

www.maxim-ic.com/factory-hardened

MAXIM
INNOVATION DELIVERED®

MAXIM DIRECT www.maxim-ic.com/shop
AVNET electronics marketing www.em.avnet.com/maxim
Digi-Key www.digikey.com/maxim
MOUSER ELECTRONICS www.mouser.com/maximic

Add extra output to a boost converter

Vladimir Oleynik, Moscow, Russia

Designers use step-up-converter ICs in battery-powered portable equipment. These chips usually provide one output with a fixed or an adjustable voltage. Some chips contain an LBI/LBO (low-battery-in/low-battery-out) function. The chip manufacturer intends for these pins to be used for monitoring a low-battery condition and to warn gadget owners when a battery goes flat. You can instead use this function to provide an extra voltage output.

The Maxim MAX756 boost converter provides a fixed output of 3.3 or 5V at 300 and 200 mA, respectively (**Figure 1**). The input voltage can range from 0.7 to 5.5V. For low-battery detection, the part has on-chip circuitry comprising a comparator, a reference, and an open-drain MOSFET. When the voltage at the LBI input is lower than its threshold level of 1.25V, the MOSFET at the LBO output sinks current to ground.

DIs Inside

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

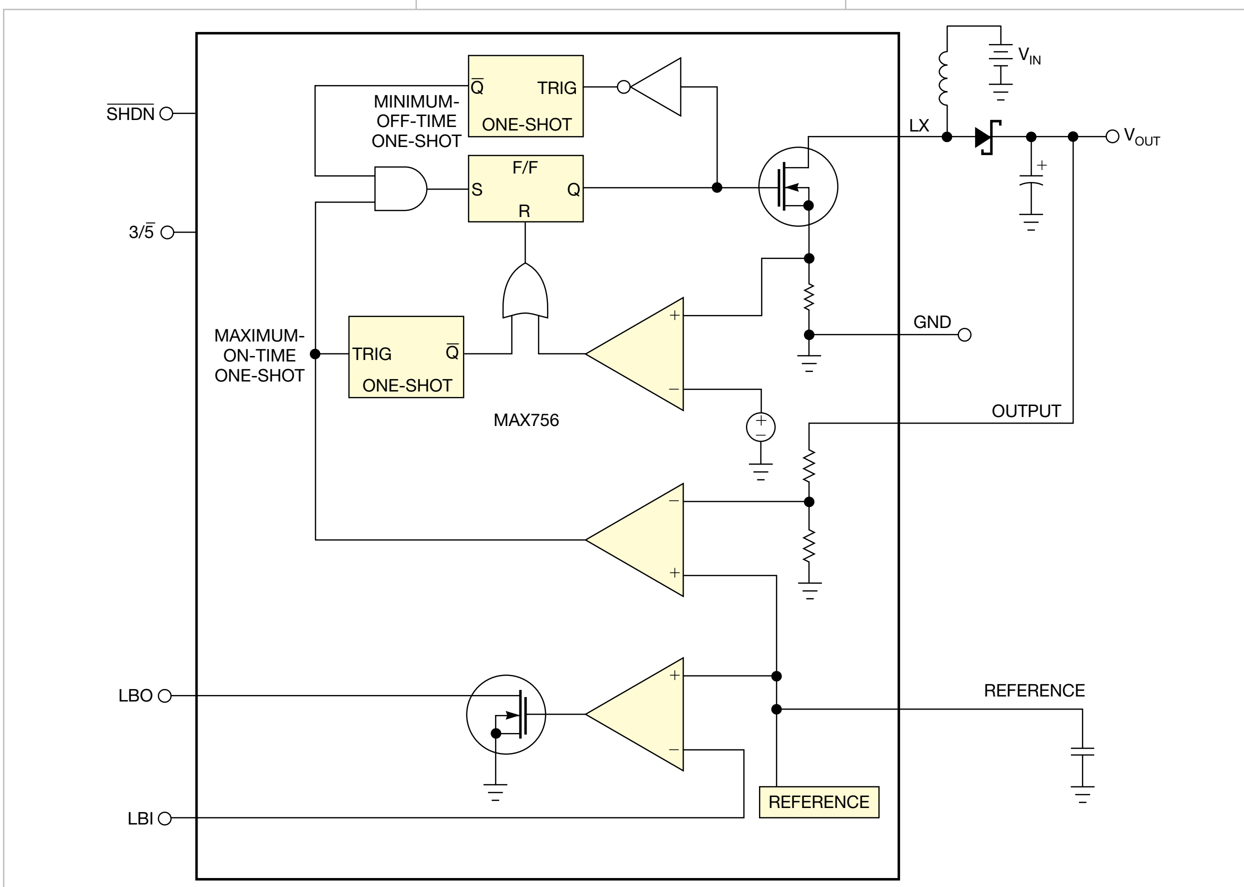

Figure 1 You can use the battery-monitoring circuit in this fixed-output boost converter to make a secondary output voltage (courtesy Maxim).

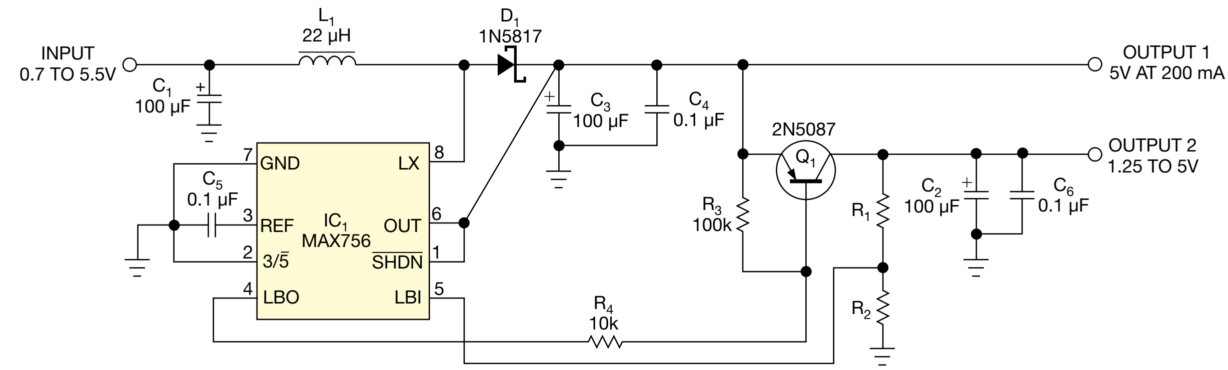

Figure 2 You can use the LBI and LBO pins of this boost-converter IC to operate a low-dropout secondary output.

You can use these components to make a second output with a regulated voltage (**Figure 2**). R_1 and R_2 determine the secondary output voltage according to the following **equation**: Output 2 $=V_{REF}(R_1+R_2)/R_2$, where V_{REF} is the reference voltage, which is 1.25V for this chip.

You can set Output 2 from 1.25 to 5V as long as it is less than Output 1. Because Output 2 is derived from Output 1, the total output current for both outputs should not exceed 200 and 300 mA for Output 1 and Output 2 voltages of 5 and 3.3V, respectively.

You can also use the LBI/LBO function to make a second boost converter (**Figure 3**). The CD4093 quad Schmitt-triggered NAND gates, inductor L_2, R_2 through R_4, Q_1, D_1, C_1, and C_2 compose this boost converter. Add C_1 and R_2 to IC_{1B} to make a free-running oscillator that IC_{1A} gates on. For the values of R_2 and C_1 in the **figure**, the oscillator frequency is approximately 17 kHz. R_1 pulls up the open-drain LBO output.

When the voltage at the LBI pin is lower than 1.25V, the LBO pin is low, thus allowing operation of the IC_{1B} oscillator. IC_{1C} and IC_{1D} drive

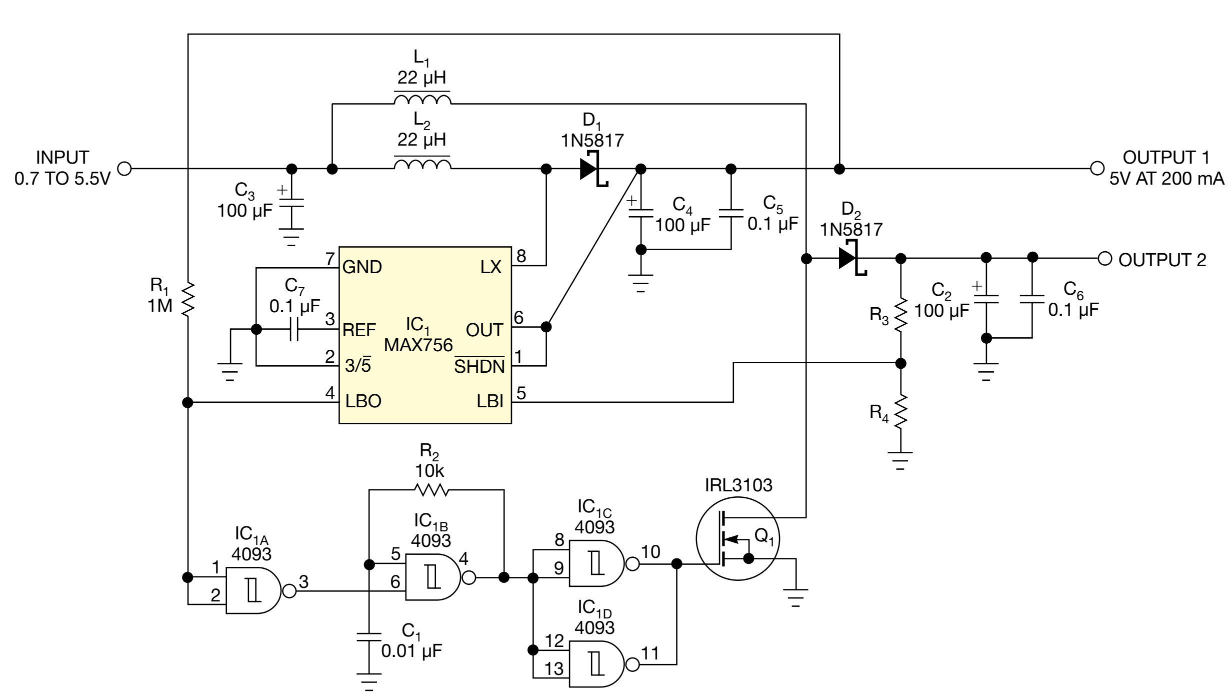

Figure 3 If you need a boosted secondary output, you can use the LBI and LBO pins to make another switching boost converter.

power MOSFET transistor Q_1. When Q_1 is on, it pulls current from inductor L_1. When Q_1 is off, this energy charges capacitor C_2 through flyback diode D_1. You apply feedback with resistor divider R_3 and R_4 to determine the Output 2 voltage, according to the following **equation**: Output $2=1.25V\times(R_3+R_4)/R_4$. IC_1 gets power from Output 1.

The voltage at Output 2 is a function of the output current and the input voltage (**Figure 4**). If you have adequate input voltage, the output graph shows a flat section where the IC's regulation is effective (**figures 5** and **6**).EDN

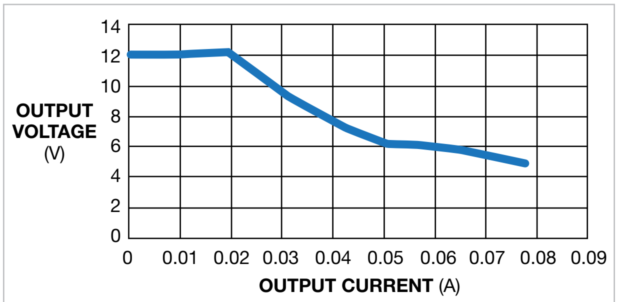

Figure 5 With a 2V input, the circuit maintains regulation to an output current as high as 20 mA.

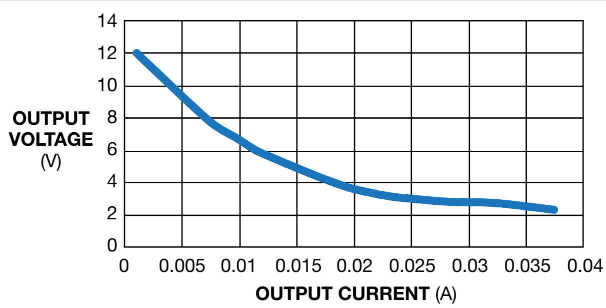

Figure 4 With only 1V input, the circuit cannot hold regulation, and the output voltage drops directly with output current.

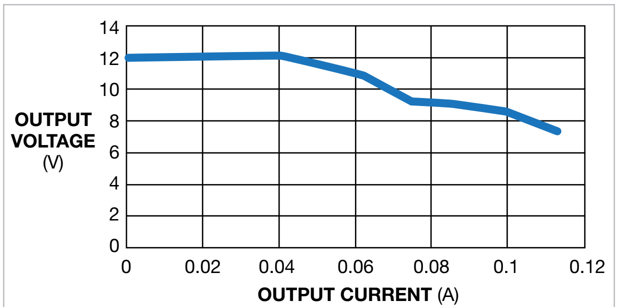

Figure 6 With a 3V input, the circuit holds regulation to an output current as high as 40 mA.

Fabricate a high-resolution sensor-to-USB interface

Zoltan Gingl, University of Szeged, Szeged, Hungary

The circuit in this Design Idea combines a mixed-signal microcontroller, a USB UART (universal asynchronous receiver/transmitter), and a novel adaptable analog sensor-input circuit. It allows you to connect many types of sensors to the design's two analog-input channels, control the device, and read measurement data on a USB host. The USB connection powers the circuit. You can control the device from your computer with simple commands; even terminal software can make the measurements. The 8051 core allows for easy programming with freely available tools, such as IDEs (integrated development environments), debuggers, and C compilers.

The design is based on an $8 microcontroller that features an 8051 architecture, as well as a PGA (programmable-gain amplifier) and a 24-bit sigma-delta ADC (**figures 1, 2,** and **3**). Microcontroller IC_1 has an input multiplexer allowing differential or single-ended mode. It also has two DAC outputs and can provide five unassigned digital-I/O pins (**Figure 1**). One output pin drives D_1 under program control. The remaining digital pins are used to configure the two analog-input ports. You also send the microcontroller's reference output to one of the analog-input ports. Four remaining digital pins interface with the USB's UART chip (**Reference 1**).

A 3.3V linear regulator, IC_2, powers the microcontroller (**Figure 2**). You power USB chip IC_1 directly from the USB port through a ferrite bead and a filter network. This popular and reliable USB UART chip lets you communicate with a computer using any operating system. Op amp IC_4 buffers the microcontroller's reference output (**Figure 3**).

Two configurable analog ports allow you to connect many sensor types using two three-input connectors, each of which has a ground pin (**Figure 4**). One ground pin provides 3.3V power, and the other outputs the buffered reference voltage—nominally, 2.5V. Wire

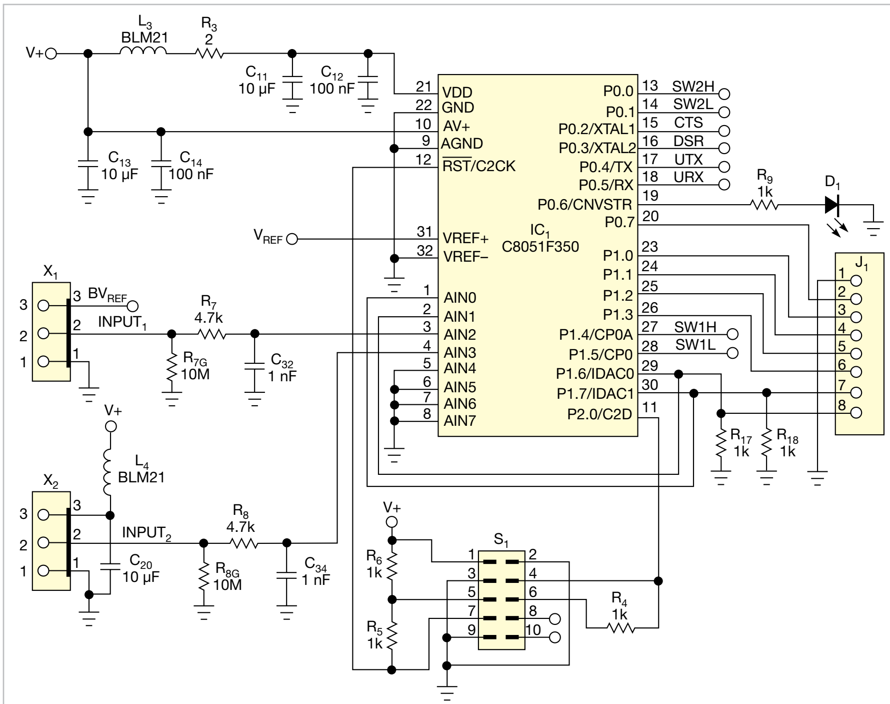

Figure 1 Microcontroller IC_1 has an input multiplexer allowing differential or single-ended mode and two DAC outputs, and it can provide five unassigned digital-I/O pins.

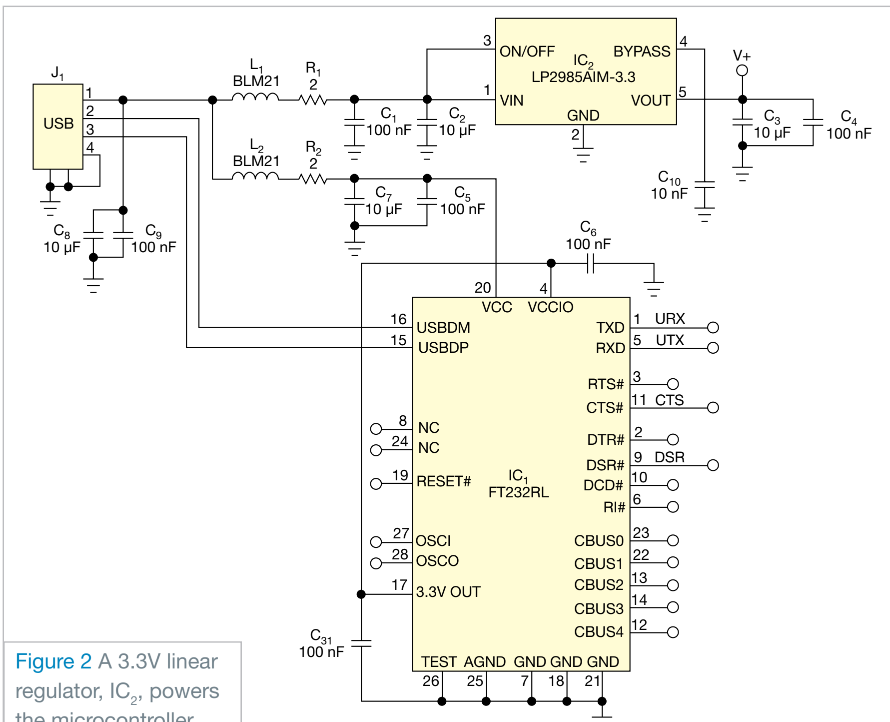

Figure 2 A 3.3V linear regulator, IC_2, powers the microcontroller.

the central pins of the two connectors to the microcontroller's analog-input multiplexer. In this way, you can either measure two single-ended voltages or use these two connectors as differential inputs. Both inputs have individually switched pullup and pulldown resistors, R_{10}, R_{11}, R_{14}, and R_{15}.

CURRENT-OUTPUT SENSORS CAN BE CONNECTED AS YOU WOULD CONNECT PHOTODIODES—BETWEEN THE GROUND AND THE INPUT PINS.

The analog-input architecture allows you to directly connect many kinds of sensors. For example, you can connect a thermistor or a photoresistor between the ground and the input pins and switch on the pullup resistor to form a voltage divider; the on-chip ADC can directly digitize this voltage divider's output (**Figure 5**). This approach also features ratiometric operation, meaning that the ADC uses the same reference as the driving voltage of the voltage divider. Current-output sensors can also be connected as you would connect photodiodes—directly between the ground and the input pins. Switch the pulldown resistor so that the photocurrent develops a voltage.

The high-resolution ADC and PGA allow direct connection of thermocouples (**Figure 6**). You achieve the required bias point by switching on both the pullup and the pulldown resistors on one channel. You can use directly connected bridge-type sensors, such as load cells and pressure sensors, by switching off all of the internal resistors. In these cases, you should operate the ADC in differential mode. Leaving all of the switches open also suits use in potentiometer inputs or IC sensors, such as the SS49E Hall-effect magnetic-field sensor.

When using directly connected sensors, you should consider source impedance, signal range, filtering, and noise

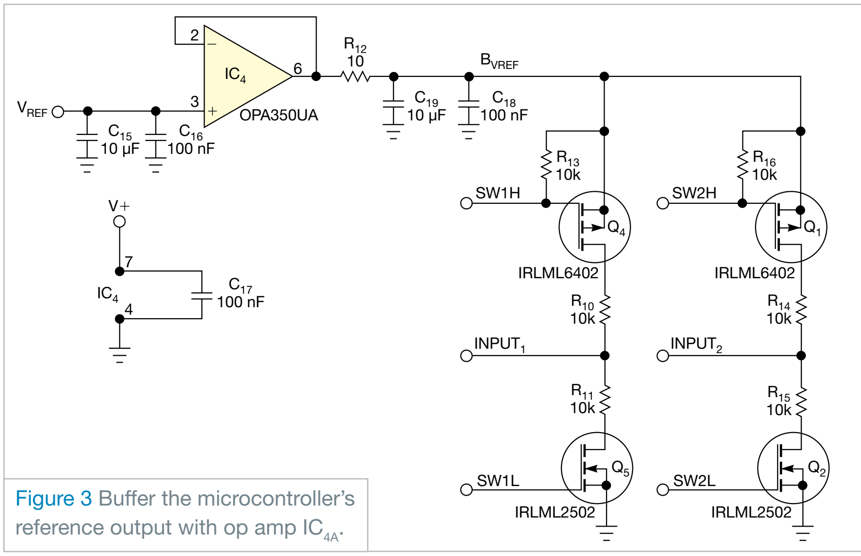

Figure 3 Buffer the microcontroller's reference output with op amp IC_{4A}.

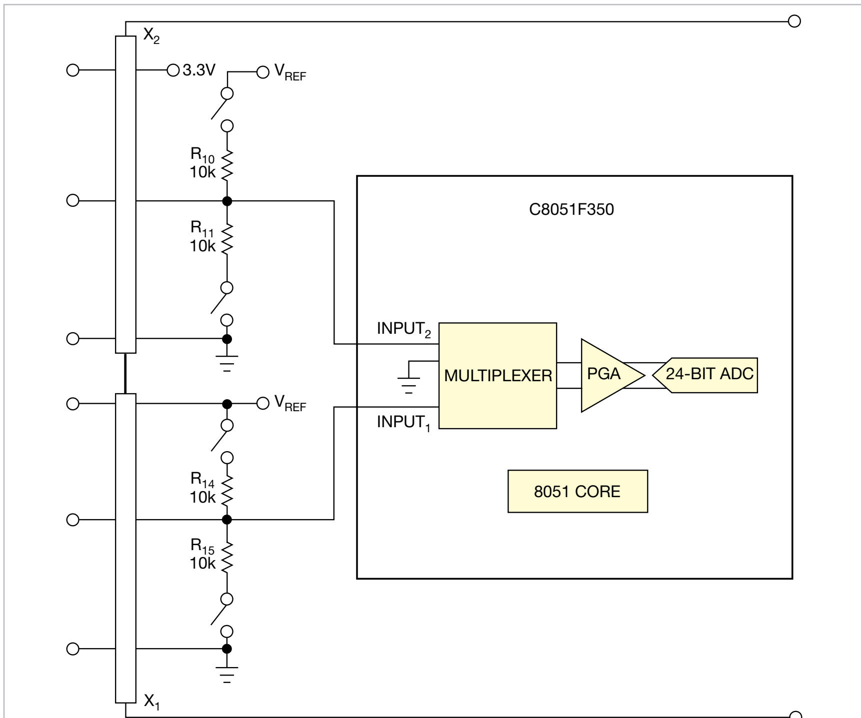

Figure 4 You create two configurable analog ports that allow you to connect many sensor types using two three-input connectors, each of which has a ground pin.

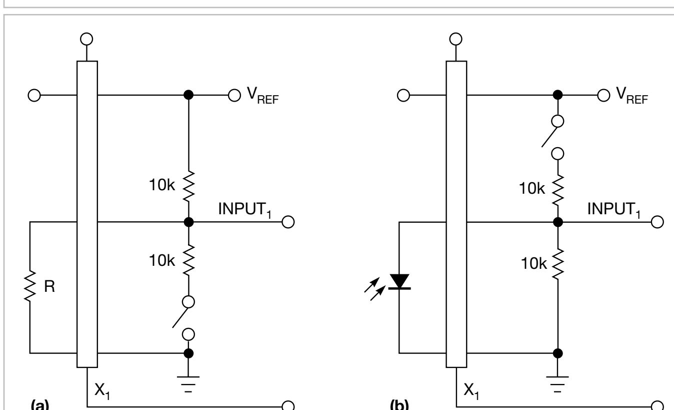

Figure 5 You could configure analog-input port X_1 for a resistive sensor (a) and for a current-output sensor (b), such as a thermistor or a photodiode.

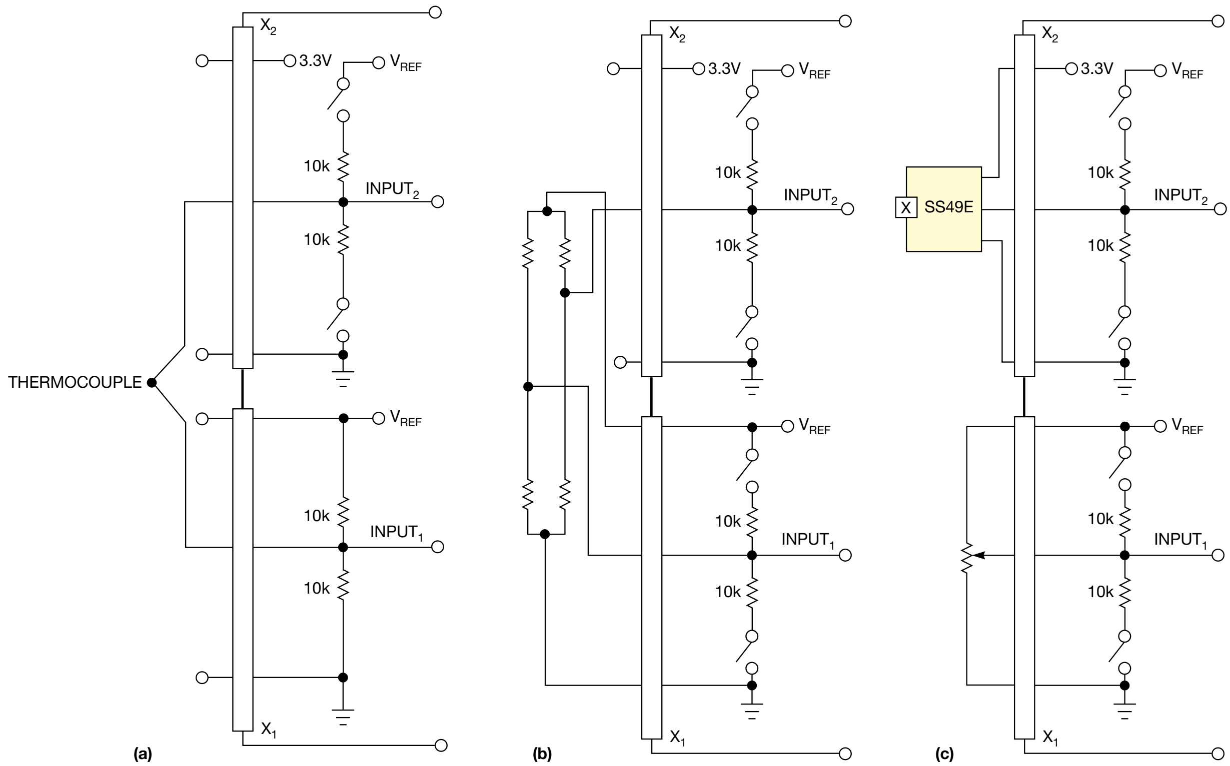

Figure 6 You can directly connect small-output voltage sensors (a), resistor bridges (b), or IC sensors and potentiometric sensors (c).

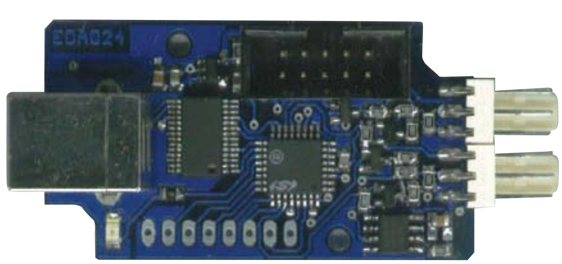

Figure 7 The design fits into a 2.36×1.38-in. enclosure.

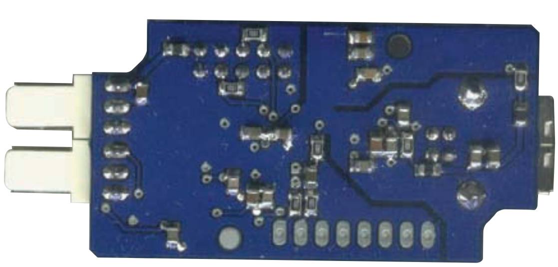

Figure 8 The underside of the PCB holds several passive components.

pickup (**references 2** and **3**). You might need to add external buffer amplifiers or a more precise voltage reference. The availability of a reference voltage and 3.3V power on the analog ports makes this setup possible. You can also use the DAC outputs in connector J_1 to trim a value or to provide an arbitrary voltage to the sensors. Note that J_1 also has five digital-I/O pins (**Figure 1**).

The design fits into a 2.36×1.38-in. enclosure (**Figure 7**). The underside of the PCB holds several passive components (**Figure 8**). **Reference 4** provides details and enables you to download the entire design, as well as CAD/CAM files, bills of materials, and software. EDN

REFERENCES

1 Kopasz, Katalin; Peter Makra; and Zoltan Gingl, "Edaq530: a transparent, open-end and open-source measurement solution in natural science education," *European Journal of Physics*, Volume 32, February 2011, pg 491, http://bit.ly/nGXz0o.

2 *C8051F35x Delta-Sigma ADC User's Guide*, Silicon Laboratories, 2005, http://bit.ly/qg4jgl.

3 Kester, Walt; James Bryant; and Joe Buxton, "High resolution signal conditioning ADCs," Analog Devices, http://bit.ly/ncGvNb.

4 Gingl, Zoltan, "EDAQ24 24-bit microcontroller sensor-to-USB interface and data logger," 2011, http://bit.ly/pHCk47.

Originally published in the January 23, 1986, issue of *EDN*

Converters yield droop-free S/H circuit

TG Barnett, The London Hospital Medical College, London, UK

In low-frequency applications, many monolithic sample/hold circuits suffer a droop rate that can cause an unacceptably large output error. The S/H circuit in **Fig 1** eliminates droop error by operating two 8-bit multifunction converters back to back. The circuit requires a 5V supply and accepts analog inputs between 0 and 2.5V (although you can scale and offset any input signal to fall within this range).

The analog input is applied to the inverting input of an LM324 op amp (IC_1), which is wired as a comparator. The op amp and the IC_2 multifunction converter form a ramp-and-compare A/D converter. (Because the Ferranti ZN435 multifunction converter includes a voltage-output D/A converter, an 8-bit up/down counter, a 2.5V bandgap reference, and a clock generator, you can configure the device either as an A/D or as a D/A converter.) The converter's internal counter counts from 0, producing a positive-going ramp at the analog output.

When the ramp voltage exceeds the analog input, the comparator output goes high and sets IC_5's Q_1 output high, thus inhibiting IC_2's clock and stopping the counter. IC_2's digital outputs are connected to the digital inputs of IC_3, which is wired as a D/A converter. The D/A converter provides the S/H circuit's analog output.

The output will remain in a hold state until you reset the monostable multivibrator (IC_4), whose outputs apply simultaneous reset pulses to IC_2 and IC_5. The circuit then resamples and holds a new value of analog input. The S/H circuit provides 8-bit hold accuracy for analog input frequencies as high as 1 kHz; you can use a faster op amp for IC_1 for higher-frequency operation. EDN

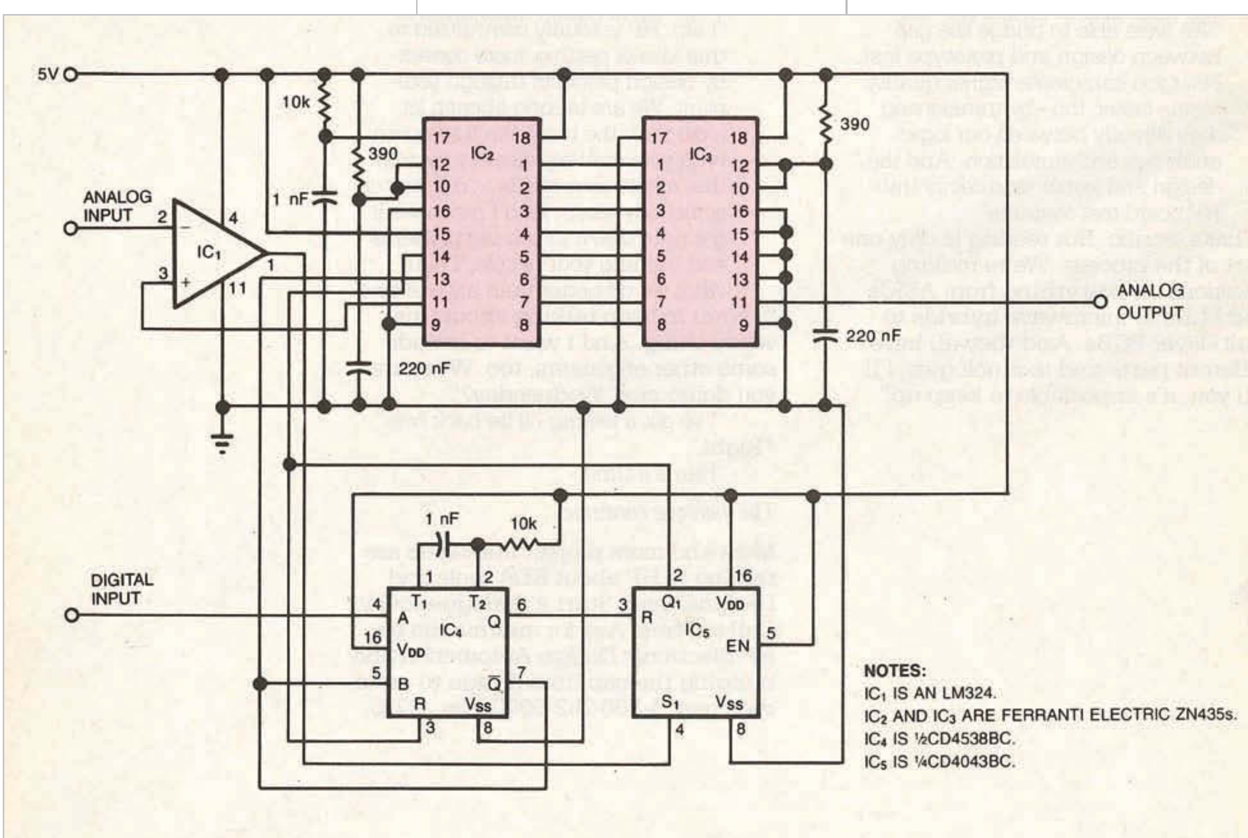

Figure 1 This S/H circuit provides unlimited hold time without droop by digitizing the analog input (IC_2) and converting back to analog using a D/A converter (IC_3).

productroundup

OPTOELECTRONICS AND DISPLAYS

Designers in the industrial, automotive, and commercial sectors looking for optoelectronic components want reliable isolation. An optocoupler is just the device to provide optimum signal isolation because it is extremely resistant to EMI and voltage fluctuations. Optocouplers combine in one package an LED on the input side, which converts an electric signal into light, and a light-receiving element on the output side, which converts light into an electric signal. High-end-system manufacturers are increasingly adopting the devices to provide input-output isolation for glitch-free circuitry in standard electrical equipment. Vendors are heeding the call for strict clearance and creepage specifications from the IEC, designing parts with stretched small-outline packages.

"Industrial" seems to be the buzzword in the display sector, as well. TFT-LCD manufacturers are making displays with specifications that meet the rugged demands of this market. These features include excellent viewability, resistance to mechanical shock and vibration, long operating life, and performance in extreme temperatures, including reaching full brightness on power-up, regardless of ambient-temperature conditions.—**by Ismini Scouras**

Sharp's LED-backlit LCDs target industrial applications

Targeting use in factory automation, gaming, medical devices, point-of-sale terminals, transportation, and test and measurement, these large-format, LED-backlit LCDs feature a built-in LED driver. The 10.4-in., VGA-landscape LQ104V1LG81 and LQ104V1DG81 feature an LVDS interface and a digital interface, respectively. Both feature 450-nit brightness, an 800-to-1 contrast ratio, and a –30 to +80°C operating-temperature range. The 12-in., SVGA-landscape LQ121S1LG84 features a 50,000-hour backlight life, a –30 to +80°C operating-temperature range, 450-nit brightness, and an 800-to-1 contrast ratio. The 15-in., XGA-landscape LQ150X1LG91 features 350-nit brightness, an 800-to-1 contrast ratio, 110°-vertical and 160°-horizontal viewing angles, and a –20 to +70°C operating-temperature range. The LQ104V1LG81 and LQ104V1DG81 sell for $325 each; the LQ121S1LG84 and LQ150X1LG91 sell for $350 and $375, respectively.

Sharp Microelectronics of the Americas, www.sharpsma.com

Avago's ACPL-K4xT optocouplers aim at automotive applications

Targeting use in onboard chargers and other high-voltage systems in hybrid and electric vehicles, the ACPL-K4xT optocouplers offer working voltage as high as 1140V. They feature an operating-temperature range of –40 to +125°C. Current consumption can reach as low as 1.5 mA/channel. Additional features include a supply-voltage range as high as 20V, zero off-state current, and 30-kV/µsec common-mode rejection at a common-mode voltage of 1500V. The one-channel K43T, two-channel K44T, and one-channel K49T sell for $2.23, $2.97, and $1.55 (1000), respectively.

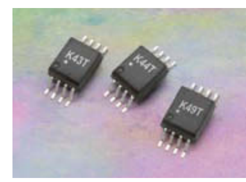

Avago Technologies,
www.avagotech.com

Vishay offers CNY64ST, CNY65ST optocouplers

Targeting solar and wind-turbine installations, the VDE-certified CNY64ST and CNY65ST optocouplers operate in Category IV installations with transient-overvoltage protection of 12,000V and recurring peak voltage of 1450V. The devices offer creepage distances of 9.5 amd 14 mm, respectively, with working voltages up to 600V. Insulation distance is greater than 3 mm, current-transfer ratios are 50 to 300% at 5 mA, and operating-temperature range is –55 to +85°C. Prices for the 64ST and 65ST are $1.70 and $1.90 (1000), respectively.

Vishay Intertechnology,
www.vishay.com

Optrex launches TVL-55682x TFT-LCD line

The TVL-55682x family of wide-format 10.1-, 11.6-, and 14-in.-diagonal TFT LCDs suits applications in high-volume consumer products. The 10.1-in., 5.2-mm-thick device has 1024×600-pixel resolution, 2.8W power consumption, 200 cd/m² brightness, a 500-to-1 contrast ratio, and a standard LVDS interface. The 11.6-in., 3.7-mm-thick unit has 1366×768-pixel resolution, 4.2W power consumption, 220-cd/m² brightness, and a 500-to-1 contrast ratio. The 14-in., 5.2-mm-thick LCD has 1366×768-pixel resolution, 4.8W power consumption, 200-cd/m² brightness, and a 500-to-1 contrast ratio. Prices for the 10.1-in. unit start at less than $108 (high volumes). The 11.6- and 14-in. units sell for $133 and $165, respectively.

Optrex America, www.optrex.com

EDN ADVERTISER INDEX

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

EDN productmart

This advertising is for new and current products.

Where there's smoke

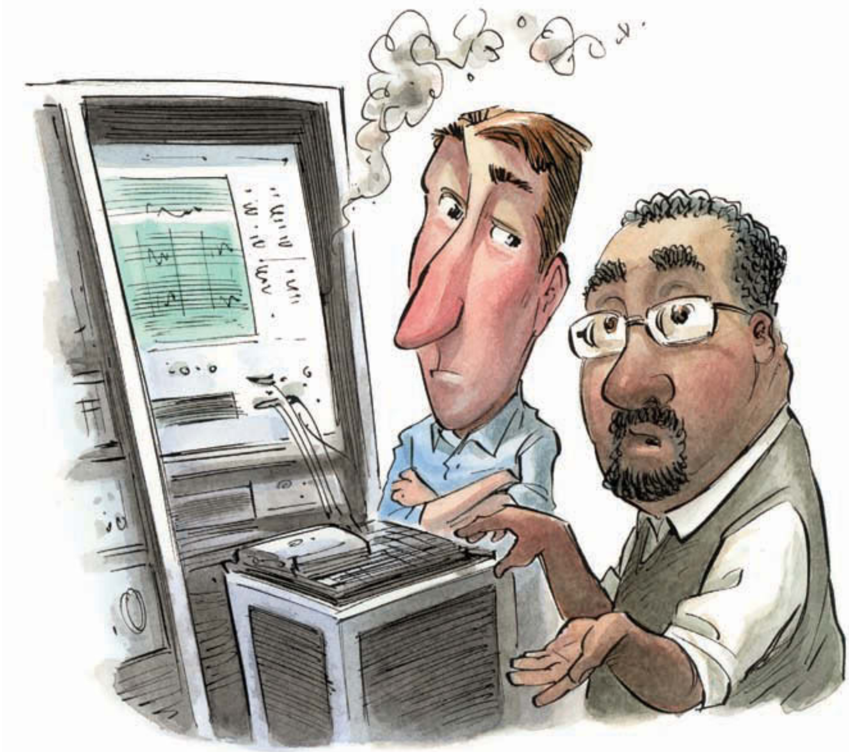

I received an e-mail one day, requesting that I investigate an automated test station that I had brought online several months prior. The support engineer explained that the system had inexplicably shut down during some UUT (unit-under-test) troubleshooting and that the engineer running the test smelled smoke in the area. I visited the test area and asked the engineer to walk me through his exact procedure for troubleshooting.

All of my assumptions about what had happened involved operator error. If the wrong relays had been commanded to close, they could have shorted out the power supplies and triggered their overcurrent-protection mode. The engineer showed me five times exactly what he had done. I didn't believe him the first four times, but I soon realized that it was not operator error!

I decided I'd try finding the source of the smoke. When I saw no obvious signs of burned resistors or scorched wires in the test adapter, I got out my tools and starting pulling apart the test station one card at a time. The cross-point-relay matrix had a burned adapter card in front of it. The damage was so extensive that it left scorch marks on a nearby card.

A relay problem had obviously shorted the power, but how could that be the case when the software never commanded the wrong relays to close? How could the damage have been so extensive when the power supplies are current-limited?

The only way to figure out what had happened was to examine the burned card and trace back the shorted traces to the cause of the problem. The burned lines connected a 26V-ac output of the UUT to the digitizer in the test station. With a 1-MΩ impedance on the digitizer, the voltage and current did not even come close to the maximum limits of the relays. To make it even more perplexing, the engineer never ran those tests during his troubleshooting.

Now, with more questions than I had before my trip to the test area, I packed up my stuff and headed back to my desk to pore over all of the information, hoping to spark some theories. First, I drew the entire schematic, including the relays from the cross-point-matrix card, the test adapter, and the portion of the UUT involved. I checked the current rating of the adapter card that had burned. I then stepped through each line of code in the software and started reading data sheets, line by line, for hours. I could find no reason that the card had burned.

When reading through the cross-point-matrix data sheet, I noticed a warning below the voltage and current ratings of the relays. It stated that operators must protect the relays from voltage transients. As soon as I read that warning, it hit me: The UUT component that the engineer was measuring was the output of a transformer. An inductive device, such as a transformer, generates a voltage transient whenever you remove power because of the collapsing magnetic field. This transient could sometimes reach thousands of volts, far exceeding, for short periods, the voltage limit of the relays.

I concluded that, over several months of properly running the UUT, the cross-point matrix was receiving damaging voltage spikes with each power cycle until a relay finally failed and shorted the circuit. To make matters even worse, the shorted transformer could supply more than 10A, explaining why the adapter card was so badly damaged. I learned a valuable lesson about designing with reed relays that day.

First, I assured the engineer who had performed the troubleshooting that he was not at fault. I then added an MOV (metal-oxide varistor) to suppress the voltage transients and a fuse to prevent the MOV from shorting out at the end of its life. I then included a voltage-divider circuit to further isolate the relays from the transformer in case another short occurred. EDN

Chris Fazekas is a systems test engineer at Lockheed Martin Missiles and Fire Control (Orlando, FL).

DANIEL VASCONCELLOS

New Concept in BIG Battery Charging

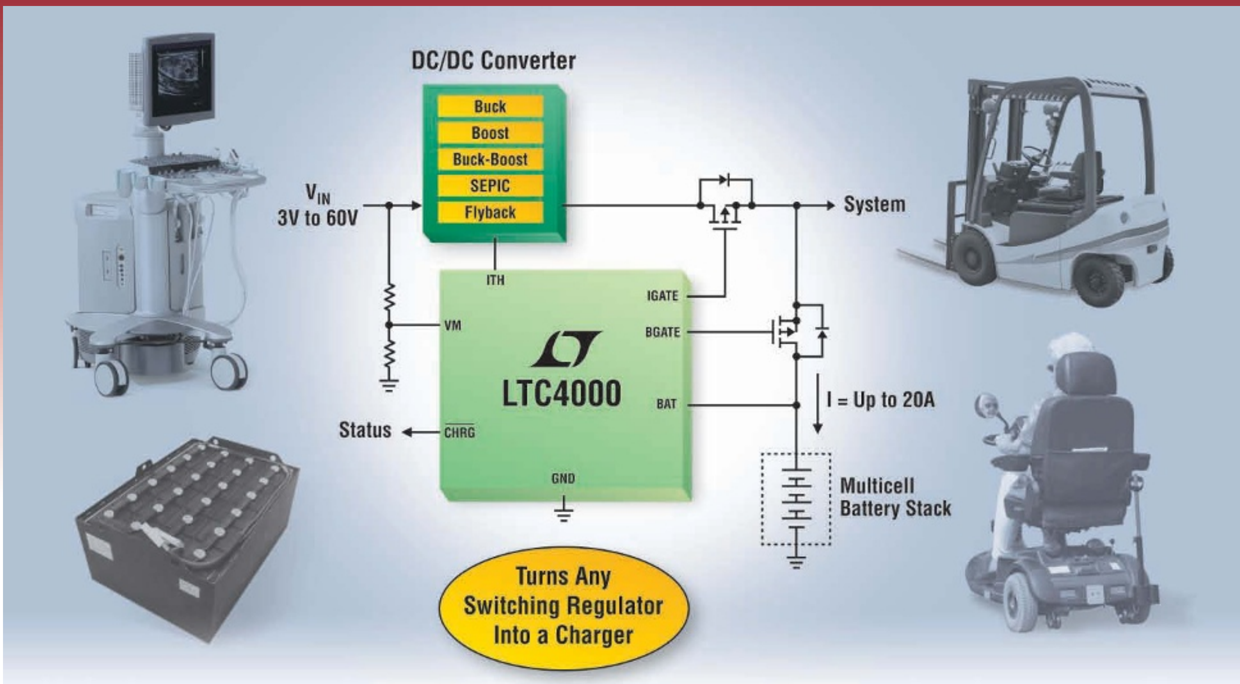

Complete Solution: PowerPath™ Control & Termination with No Software

The LTC®4000 is a high voltage controller and power manager which, when paired with a DC/DC converter, becomes a full-featured battery charger solution. The LTC4000 is capable of driving virtually any topology, including buck, boost, buck-boost, SEPIC and flyback converters. Its intelligent PowerPath manager provides power to the system load when input power is available, enabling instant-on operation even with a deeply discharged battery. A full-featured controller, the LTC4000 can charge a variety of battery types including lithium, nickel and lead acid-based chemistries. Highly accurate charge current and float voltage, as well as on-board termination, ensure safe and accurate charging.

Features

- Input/Output Voltage: 3V to 60V
- Charge Currents up to 20A
- Input Ideal Diode for Low Loss Reverse Blocking and Load Sharing
- Programmable Input and Charge Current: ±1% Accuracy
- ±0.2% Accurate Programmable Float Voltage
- Programmable C/X or Timer-Based Charge Termination
- NTC Input for Temperature Qualified Charging

Charge Current Profile vs V_{BAT}

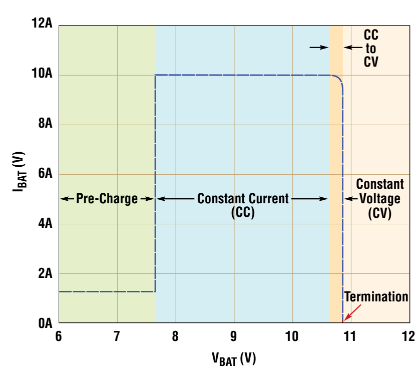

V_{FLOAT} = 10.8V

Info & Free Samples

www.linear.com/product/LTC4000

1-800-4-LINEAR

www.linear.com/batsolutions

LINEAR TECHNOLOGY