UBM Tech

EDN®

APRIL 2013
www.edn.com
Issue 4

Decoupling caps are where it's at Pg 70

"Tweet" if you're an engineer! Pg 10

Teardown: Insignia 60W-replacement LED bulb Pg 22

Design Ideas Pg 56

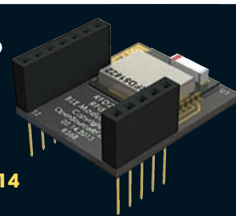

RFduino board from Open Source RF Page 14

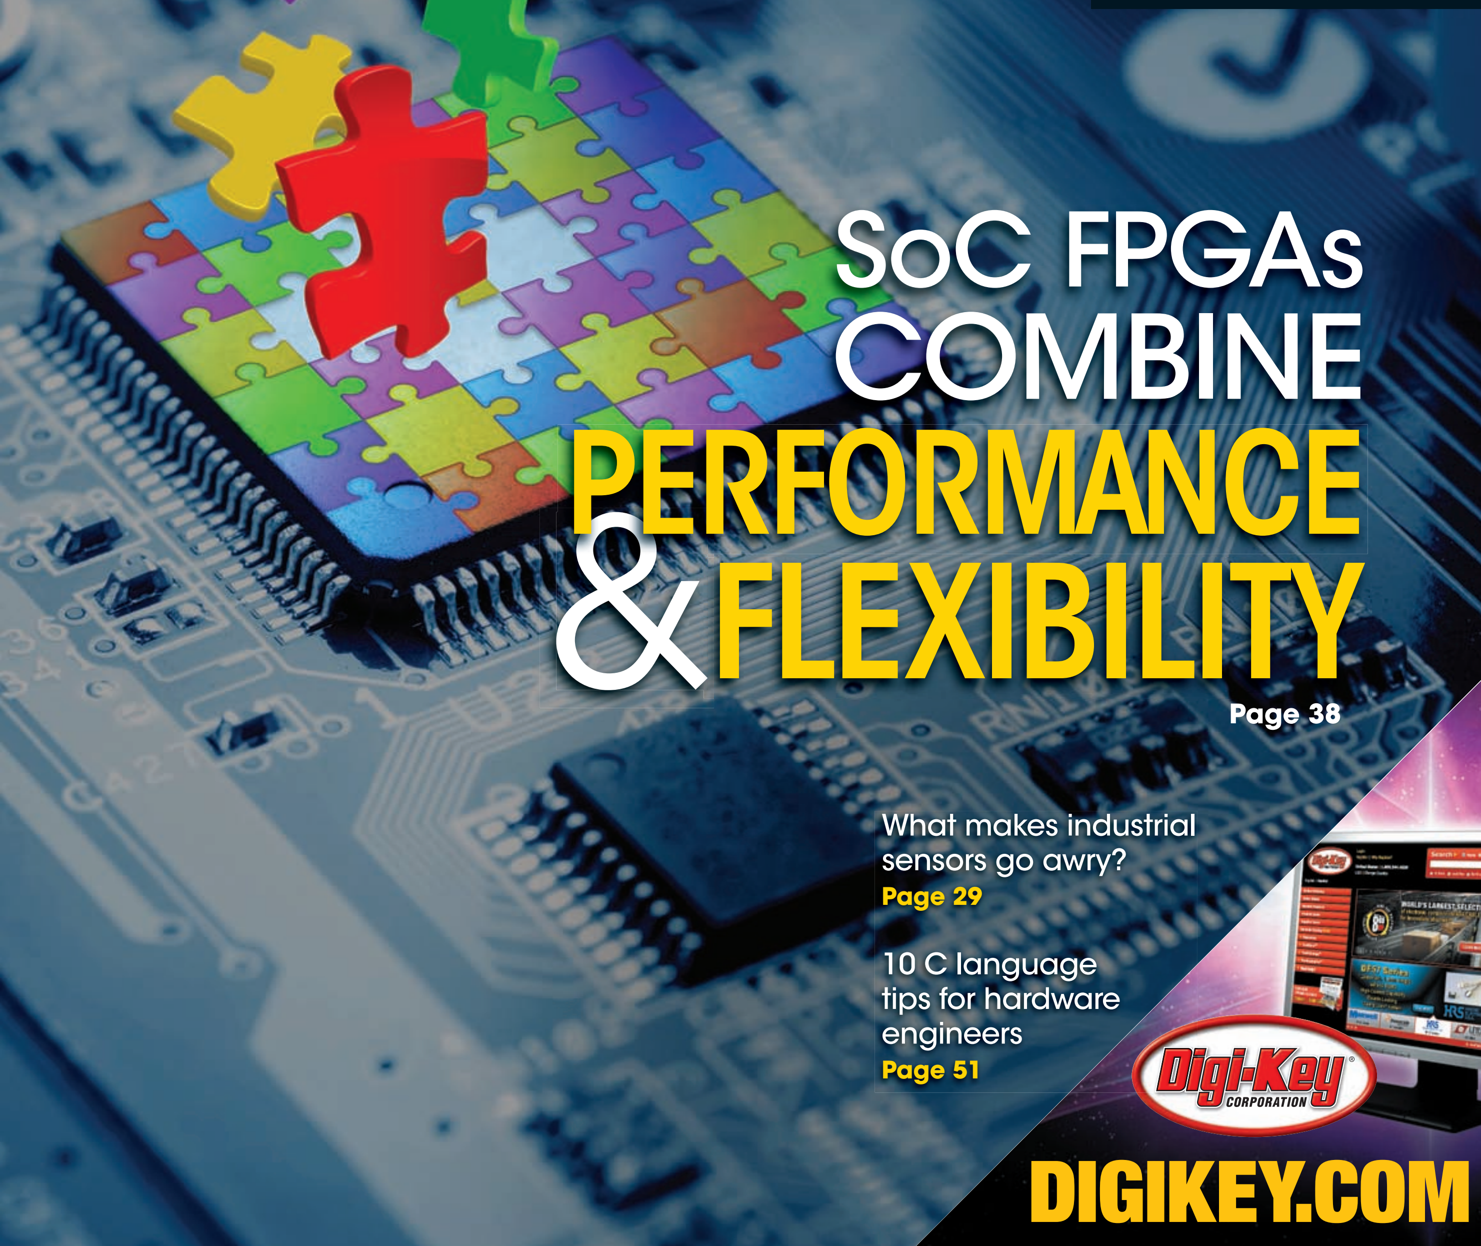

SoC FPGAs COMBINE PERFORMANCE & FLEXIBILITY

Page 38

What makes industrial sensors go awry? **Page 29**

10 C language tips for hardware engineers **Page 51**

The Trusted Leader
for Any Measurement

NI data acquisition products are the most trusted computer-based measurement devices available, offering superior accuracy and maximum performance. NI CompactDAQ systems feature more than 50 measurement-specific modules and complete breadth and depth of I/O. Coupled with the analysis and signal processing capabilities of NI LabVIEW software, this platform puts the power of graphical system design at your command so you can build any measurement system faster.

LabVIEW offers powerful analysis functions for sensor and signal data, simplifies GUI development, and helps you program the way you think—graphically.

>> Accelerate your productivity at ni.com/measurements-platform

800 453 6202

NATIONAL INSTRUMENTS

EDN contents

April 2013

SoC FPGAs combine performance and flexibility

38 Embedded-system architectures built on a combination of MCUs and FPGAs offer the kind of adaptability increasingly required to support changing demand for greater functionality across diverse applications.

by Michael Parker, Altera Corp

What makes industrial sensors go awry?

29 Sensors that measure parameters such as pressure, temperature, toxic gas, and pH are abundant and make industrial processing safer, more efficient, and less costly. Each sensor type, however, has its own set of unique characteristics, resulting in various and complex design challenges.

by Jason Seitz, Texas Instruments

IMAGE: ISTOCK

10 C language tips for hardware engineers

51 It can be common for a hardware designer to write code to test that hardware is working. These 10 tips for C—still the language of choice—may help the designer avoid basic mistakes that can lead to bugs and maintenance nightmares.

by Jacob Beningo, Beningo Engineering

COVER IMAGE: SHUTTERSTOCK; GIULIA FINI

DESIGNIDEAS

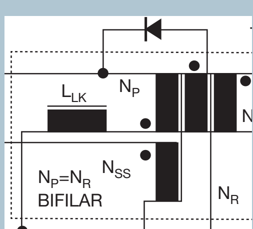

56 Recover the leakage energy of a flyback transformer

59 Double the protection of a laser driver using a 1V power supply

60 Gate-drive transformer eases multi-output, isolated dc/dc-converter designs

62 Two ICs form F/V converter

▶ Find out how to submit your own Design Idea: www.edn.com/4394666.


www.mouser.com
The Newest Products for Your Newest Designs®
GO
MORE.
mouser.com
The widest selection
of the newest products.
AUTHORIZED ecia SOURCE
Authorized distributor of semiconductors
and electronic components for design engineers.
MOUSER ELECTRONICS.
Mouser and Mouser Electronics are registered trademarks of Mouser Electronics, Inc. Other products, logos, and company names mentioned herein, may be trademarks of their respective owners.

contents April 2013

pulse

Dilbert 14

12 NXP claims sound-quality breakthrough with 9.5V boost voltage in mobile micro speakers

14 Coin-sized, Arduino-compatible computer with Bluetooth LE wirelessly connects to smartphones, tablets

16 Contactless connectivity platform targets applications that contain moving parts

16 Lattice Semi debuts miniature FPGA

16 Smartphone barcodes now readable by all POS laser scanners

17 *Star Trek*-style "tractor beam" is created on tiny scale

18 **Voices:** SparkFun's Chris Taylor: Community key to open-source hardware

DEPARTMENTS & COLUMNS

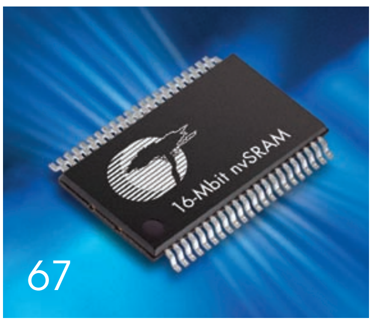

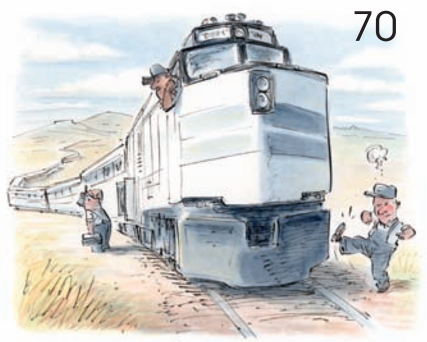

9 **EDN online:** Join the conversation; Content; Engineering community

10 **EDN.comment:** "Tweet" if you're an engineer!

20 **Signal Integrity:** Measuring nothing

22 **Teardown:** LED light shrinks size, cost with nonisolated driver

26 **Mechatronics in Design:** Excelling in the shades-of-gray real world

64 **Supply Chain:** What engineers must know about the supply chain

67 **Product Roundup:** Memory

70 **Tales from the Cube:** Decoupling caps are where it's at

EDN® (ISSN# 0012-7515) is published monthly by UBM Tech, 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. Email: 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 3609, Northbrook, IL 60065-3257. CANADA POST: Publications Mail Agreement 40612608. Return undeliverable Canadian addresses to APC, PO Box 503, RPO West BVR CRE, Rich Hill. ON L4B 4R6.

Volume 58, Number 4 (Printed in USA).

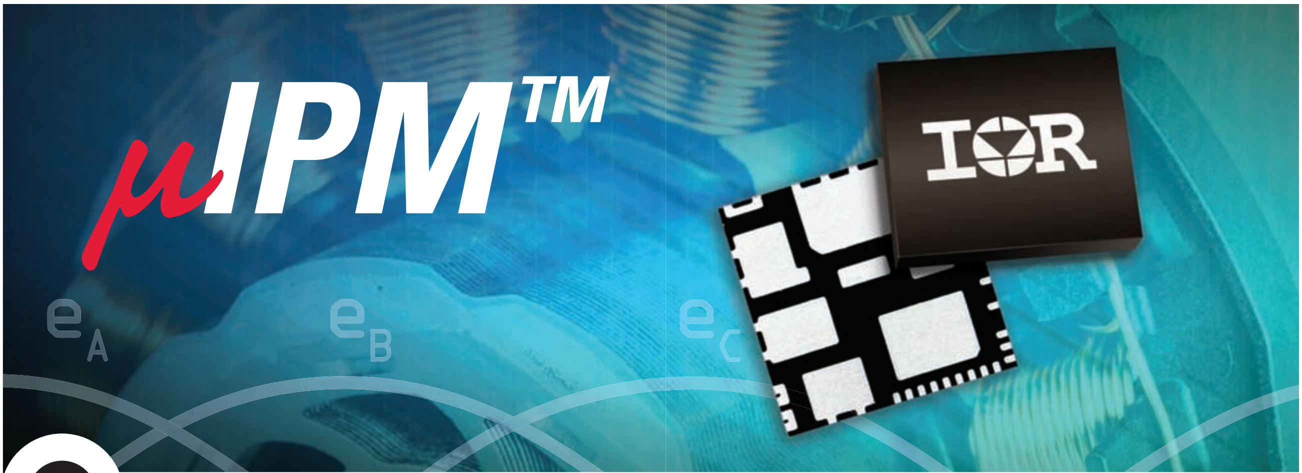

Innovative Power Module Reduces System Size

µIPM™ Power Modules Deliver up to 60% Smaller Footprint

Specifications:

Part Number	Size (mm)	Voltage	IO (DC@ 25°C)	Motor Current**		Motor Power VO=150/75VRMS	Topology
				w/o HS	w/HS		
IRSM836-024MA	12x12	250V	2A	470mA	550mA	60W/72W	3P Open Source
IRSM836-044MA	12x12	250V	4A	750mA	850mA	95W/110W	3P Open Source
IRSM836-025MA	12x12	500V	2A	360mA	440mA	93W/114W	3P Open Source
IRSM836-035MB	12x12	500V	3A	420mA	510mA	108W/135W	3P Common Source
IRSM836-035MA	12x12	500V	3A	420mA	510mA	100W/130W	3P Open Source
IRSM836-045MA	12x12	500V	4A	550mA	750mA	145W/195W	3P Open Source

Features:

- 3-phase motor control IC
- 12x12x0.9mm PQFN package offers up to 60% smaller footprint
- Eliminates the need for heat sink
- DC current ratings from 2A to 4A
- Voltage range of 250V – 500V

µIPM™ Advantages:

- Shortens design time
- Shrinks board space requirements
- Simplicity - Eliminates Heat Sink
- Replaces more than 20 discrete parts to deliver a complete motor drive stage
- Slashes assembly time and cost
- Simplifies procurement and inventory management
- Reference design kits available for quick evaluation on any 3-phase motor

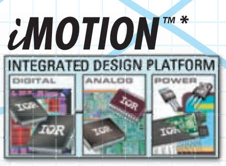

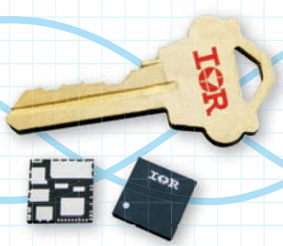

* IR's iMOTION™ (ai mo shan), representing the intelligent motion control, is a trademark of International Rectifier

** RMS, F_c=16kHz, 2-phase PWM, ΔT_{CA}=70°C, $T_A \approx$ 25°C

EDN online

JOIN THE CONVERSATION

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

In response to "My MP3 player wouldn't shut up," an entry by Bill Schweber in his Power Points blog at www.edn.com/4410592, AJ2X1 commented:

"As an old analog guy, those soft switches always made me a little suspicious. For sure, they require power be on all the time, which makes power management a bit more difficult. I did one for a TV set back in the '80s, but the actual switch was a relay. No microprocessor involved. Relay failure was about as rare as failure of the physical on/off switch (usually part of the volume control then).

Nowadays, the switch function is more likely to be a FET, but I still design the control for it to avoid the micro if I can. Pulling the plug resets everything. And though I haven't done a battery-powered design in years, I'd still prefer to hide a system-reset button somewhere to get out of the always-on condition a lost micro can get into."

In response to "Brown, lead-underpant moments," an entry from Rajan Bedi in his Out-of-this-World Design blog at www.edn.com/4410283, timbalionguy commented:

"Thanks for the great article. I'm a big fan of high-energy particle physics, and this article helped me understand some of the problems electronics can face when being bombarded with even secondary particle-beam radiation. I also very much appreciate your humor in dealing with all of this!"

***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.*

CONTENT

Can't-miss content on EDN.com

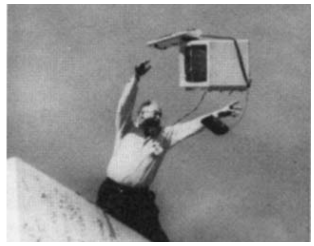

IMAGE: AL NEVES

SLIDESHOW: A JOURNEY INTO THE WORLD OF BOB PEASE

We all know "Bob Pease the engineer" from his work and published articles. Here we explore a little bit more about the man and his personality.

www.edn.com/4410514

ONE I/O LINE DRIVES SHIFT REGISTER WITH STROBE

This online-only Design Idea shows how a single microcontroller port can drive a lot of output lines through a shift register. Schematics, video, and images explain the process.

www.edn.com/4410875

ENGINEERING COMMUNITY

Opportunities to get involved and show your smarts

Serious fun at DESIGN West

You're attending DESIGN West, April 22 to 25, to learn and grow your skills, but you might as well have a little fun while you are there, too. Check out this blog post outlining some opportunities for serious fun and interacting, including an on-the-show-floor networking lounge, free beer and snacks, and a wireless-mesh-networked propeller-beanie build and prize giveaway: **www.edn.com/4410945**.

For more information and to register for DESIGN West, visit **www.ubmdesign.com**.

BY PATRICK MANNION, BRAND DIRECTOR

"Tweet" if you're an engineer!

I recently spent time completing a final review of the proceedings for the upcoming DESIGN West conference (www.ubmdesign.com) before the file went to the printer, and I came across this session title: "I <3 Android." I thought it was a typo, or some errant Android code. I instant-messaged a colleague (we were in the midst of another interminable conference call), and he clarified, with a hearty, yet subtly condescending, OMG and LOL, that it meant, "I love Android." But of course y'all knew that, right?

If I had seen it in a text from my daughter, I probably would have figured it out pretty quickly. Context is everything, however, so to come across it in a "professional" setting, within a typically "functional" conference guide, took me a bit off guard. But I should have known better.

Social media—and all its codes and shortcuts—has caught fire, even among engineers, who tend to be viewed as, well, antisocial. At least that's the perception we were working to overcome when we relaunched EDN.com last year with a view toward highlighting your thoughts and comments. Pundits and naysayers scoffed and asked, "Why? Engineers never comment. They're observers. Lurkers."

Well, eight months and more than 8000 comments later, you've certainly proven the doubters wrong. It turns out you have quite a bit to say, given the right platform.

But what about away from EDN.com? Do you "tweet" on Twitter? "Face" on Facebook? "Link" on LinkedIn? Or "+" on Google+? Our own 2013 Embedded Market Survey (bit.ly/11WzBGF) shows you do, and increasingly so. Responses to the question "In general, what sources of information do you consult to research your embedded-design decisions?" indicated usage of social-media outlets by engineers has almost doubled in the past year. Of course, stats can be misleading. By "almost doubled," I mean it's gone from 6% to 11%. But the trend is clearly upward.

Social media helps you show others what you're all about.

I have a hard time seeing how Twitter or Facebook can help you reach a design decision directly (unless you want to get fired). Indirectly, however, is another story. Social media helps you show others what you're all about, and what grabs your interest, while also helping you identify and follow like-minded people—engineers or otherwise.

How much time you want to spend doing that is up to you, of course, but eventually, as you accumulate a following and find the people whose opinions and insights you trust and respect, you can actually start to leverage those connections to get started in the right direction. You may still prefer to talk to fellow engineers for advice or input, but if you're working out of your home, or you've been "downsized" and are working independently, you don't have the option of walking down the hall to the office "guru." Online professional groups and now even social media are becoming solid options.

To help you ID some people you may find worthwhile to follow on Twitter, one of my colleagues here at UBM put together a list of 10 EEs everyone should follow (www.eetimes.com/4410448). Many individuals on the list are well known already and will be appearing live—yes, live—at our special community booth at DESIGN West. (Did I mention the show is in San Jose, April 22 to 25?) If you're coming to the conference, and you should, stop by the booth and bring your latest home-brew project (or photos) to show off. We'd love to see it, and you can bounce some ideas off the folks there who'll have their own gear to demonstrate.

Who knows? Someone from Kickstarter might stop by and offer you your first million dollars. But wait, it turns out that a solid social-media track record is a prerequisite to being an entrepreneur. Who would have thought that would be the case five, or even two, years ago? Suggestion: Find your voice in this brave new world. There's a good chance it'll pay off in the very near future. EDN

Contact me at patrick.mannion@ubm.com, or better yet, follow me on Twitter: @Patrick_Mannion! You'll find EDN *(@EDNmagazine) there, too.*

EDN

pulse

INNOVATIONS & INNOVATORS

NXP claims sound-quality breakthrough with 9.5V boost voltage in mobile micro speakers

NXP Semiconductors is launching a speaker-driver IC that enables a 9.5V boost voltage from an integrated dc/dc converter. Increasing the voltage headroom in the TFA9890 audio-driver IC prevents amplifier clipping and keeps sound quality high at maximum volume. The TFA9890 safely drives a record 4W of peak power into a standard 8Ω speaker that is typically rated at 0.5W, making a clear improvement to the sound output and quality of mobiles, tablets, TVs, and portable speakers.

Whereas traditional approaches have required cutting bass frequencies to avoid damaging the speaker, the TFA9890 builds on the advanced speaker protection introduced in the TFA9887 to enable safe operation while working at near-peak output at all times. The fully integrated protection includes adaptive excursion control, an approach that compensates for real-world changes in the acoustic environment. The IC measures current and voltage to the speaker, and uses the information to adapt the protection algorithm to account for changes such as aging, damage to the enclosure, and blocked speaker ports.

The feedback-controlled excursion-protection algorithm enables the TFA9890—a single chip that includes NXP's CoolFlux DSP, a Class-D amplifier with current sensing, and a dc/dc converter—to provide nearly twice as much power into 8Ω speakers, with sound output typically 6 to 12 dB higher than the TFA9887. In addition to increasing the speaker volume, the dc/dc converter's 9.5V boost voltage improves sound quality by increasing voltage headroom and eliminating amplifier clipping.

Other circuits and algorithms that improve sound quality include an advanced clip-avoidance algorithm, which monitors audio performance and prevents clipping even when the power supply begins to sag. Bandwidth extension increases the low-frequency response well below speaker resonance.

The intelligent dc/dc boost converter in the TFA9890 also prevents the audio system from causing battery undervoltage issues for the mobile device. The advanced, embedded algorithms require no separate licensing and allow designers to customize their audio sound quality and choose how to optimize the phone's performance.

The TFA9890 is now sampling with lead customers and was scheduled to be released in 2Q13.—**by Paul Buckley**

▷**NXP Semiconductors,** www.nxp.com

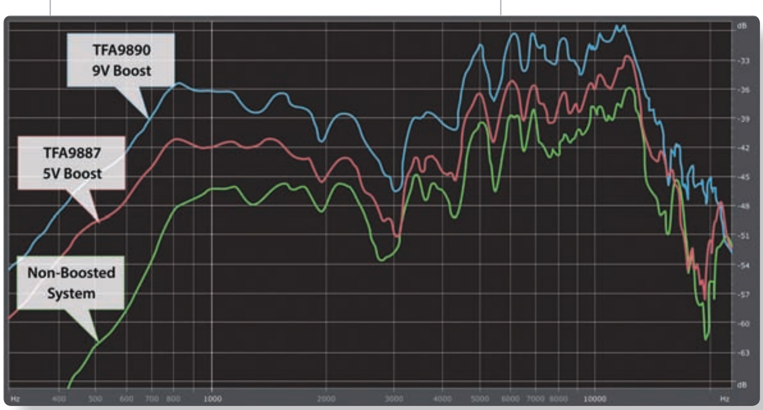

In addition to increasing the speaker volume, the NXP TFA9887 audio-driver IC's 9.5V boost voltage improves sound quality by increasing voltage headroom and eliminating amplifier clipping.

TALKBACK

"Yes, they really were the good old days. Of course, we didn't worry about GUI interfaces, and we didn't need a full-blown expensive relational database behind the scenes either."

—Author Villy Madsen, in response to reader feedback on his Tales from the Cube entry—"Where is my crowbar?"—at www.edn.com/4410773. Join the conversation and add your own comment.

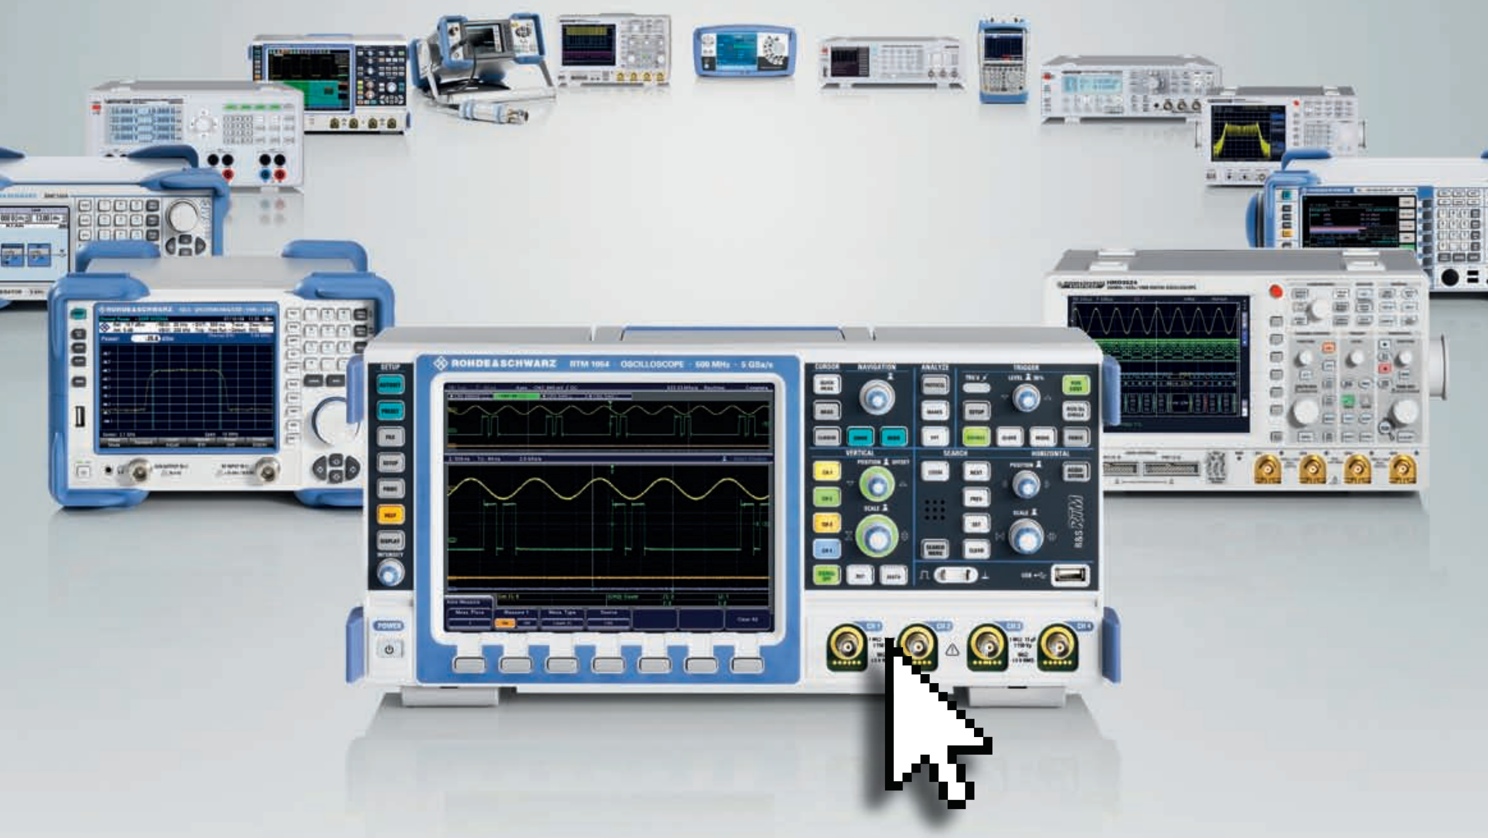

Coin-sized, Arduino-compatible computer with Bluetooth LE wirelessly connects to smartphones, tablets

Nordic Semiconductor has announced that wireless start-up Open Source RF has launched on KickStarter.com (http://kck.st/10kp55Q) the world's first Arduino-compatible open-source microcomputer that

> "Application examples include wireless multicolor RGB LED lighting and temperature sensors.

can communicate wirelessly with any Bluetooth version 4.0 (which includes Bluetooth low energy as a hallmark feature)-compatible smartphone or tablet and is based on the latest RFD51822 module. The module, in turn, is based on the Nordic nRF51822 SoC that was developed by wireless specialist RF Digital.

Called the RFduino and making full use of the nRF51822's powerful onboard 32-bit ARM Cortex M0-based processor, this fully FCC- and CE-compliant, 2.4-GHz, wireless, coin-sized microcomputer is designed to allow electronics makers and professional developers to develop thousands of miniaturized Bluetooth low-energy applications controllable from a Bluetooth version 4.0-compatible smartphone or tablet in a very short amount of time at very low cost. Open Source RF claims that the overriding focus of the RFduino is on building new wireless applications. Makers need go no deeper into the technical design aspects than high-level application design, while design engineers have the option to use standard Nordic Semiconductor nRF51-series software-development kits to fast-track a successful prototype into production.

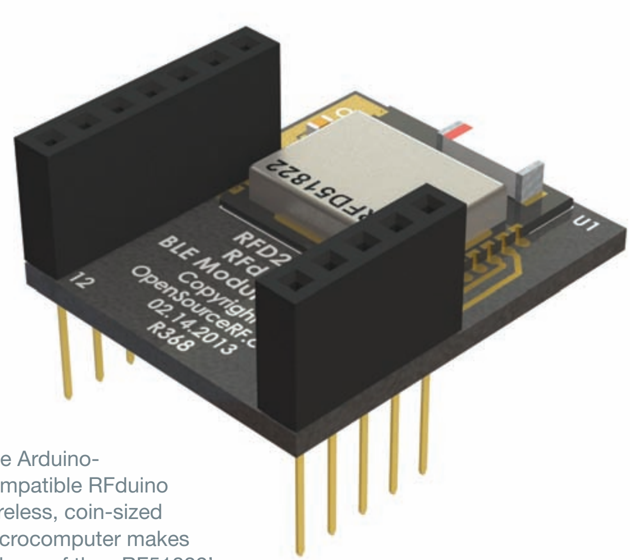

The Arduino-compatible RFduino wireless, coin-sized microcomputer makes full use of the nRF51822's powerful onboard 32-bit ARM Cortex M0-based processor. The nRF51822 is from Nordic Semiconductor.

Application examples for which Open Source RF has already developed source code include wireless multicolor RGB LED lighting; iPhone-controlled racing cars; temperature sensors; house-plant-watering sensors; proximity and motion sensors; relay switches; audio controls; robotics; theatrical props and special effects; sound, light, or button-press detectors; and various home automation and control devices.

The RFduino can be powered by anything from household outlets down to a regular CR2032 coin-cell (watch) battery. Several open-source RFduino and iPhone apps are free to use, extend, and share.

Open Source RF also offers an assortment of stackable shield accessory boards that plug directly into 0.1-in. (2.54-mm) standard spacing or solderless breadboards, or run fully standalone. These boards can also plug directly into each other to provide many combinations for quick prototyping and project building. Accessory boards included in the launch are a USB power and programming board, RGB LED and push-button board, quad servo controller board, generic prototyping board, single AAA-battery board, dual AAA-battery board, and CR2032 coin-cell-battery board.

"Over the last 12 to 18 months there has been exceptional growth in the DIY [do-it-yourself] electronics 'maker' and open-source markets supported by a growing interest in wireless generally," says Open Source RF founder Armen Kazanchian. "In addition, Bluetooth low energy has revolutionized the ULP wireless market by offering no-additional-cost smartphone and tablet connectivity as standard and the ability to make everything and anything wireless for the very first time."

—by Jean-Pierre Joosting

▷**Nordic Semiconductor**, www.nordicsemi.com

▷**Open Source RF**, www.opensourcerf.com

▷**Arduino**, www.arduino.cc

▷**RF Digital**, www.rfdigital.com

DILBERT By Scott Adams

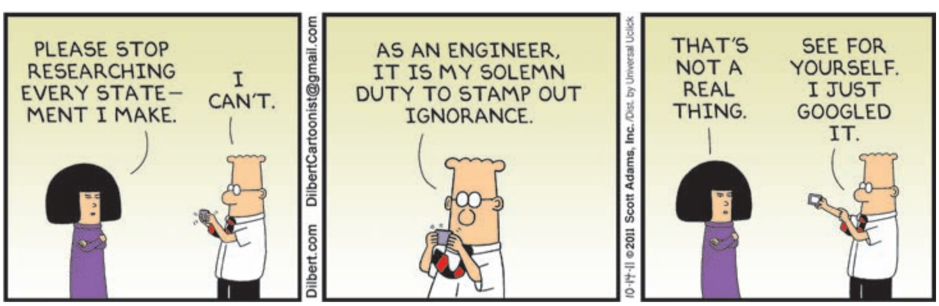

Gather, organize and control data like never before.

Agilent 34970A/72A Data Acquisition Switch Units (DAQs) have achieved market leadership for more than a decade. Perhaps it's the 34970A/72A DAQ's unmatched capability. Or the 34972A DAQ's ability to control everything remotely via mobile interface. Or the BenchLink Data Logger software that lets you scan, monitor, control and analyze your data so you can quickly act on it. Maybe it's all of the above.

34972A DAQ	34970A DAQ
3-slot LXI unit with built-in 6½ digit DMM	
USB and LAN	GPIB and RS-232
Graphical web interface	–
$1,933*	$1,673*

See a 34972A product overview. Scan or visit: http://qrs.ly/gw2wqyv

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

metrictest
PART OF THE MICROLEASE GROUP
866-436-0887
www.metrictest.com/agilent

Free $800* PC software upgrade for advanced scanning control.
www.metrictest.com/agilent/daq.jsp

* Prices are in USD and are subject to change

Anticipate __Accelerate __Achieve

Agilent Technologies

Contactless connectivity platform targets applications that contain moving parts

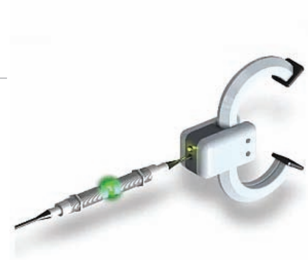

TE Connectivity's ARISO contactless connectivity enables new applications and solutions where traditional cables and connectors fail.

TE Connectivity has unveiled a new technology platform that will allow the delivery of robust power, signals, and data without any physical connectors or contact, in virtually any environment. The ARISO contactless connectivity solution seamlessly integrates wireless-power and radio-frequency (RF) technology.

The ARISO contactless connectivity approach enables new applications and solutions where traditional cables and connectors fail, reducing maintenance costs because there is no physical contact causing wear and tear on the connector. Hermetically sealed, the TE ARISO contactless connectivity couples are not susceptible to dirt, dust, or chemical fluids, making them ideal for deployment in ruggedized environments and clean environments such as food and beverage production.

Easily scalable, the platform eliminates current mechanical design barriers. It can transfer power and signals through fluids and walls. Because it is contactless, the solution minimizes risk in flammable environments. It cannot cause an arc, because it generates a magnetic rather than an electric field.

—by Julien Happich

▷**TE Connectivity**, www.te.com

Lattice Semi debuts miniature FPGA

Lattice Semiconductor Corp has announced its iCE40 LP384 FPGA featuring 384 look-up tables in packages as small as 2.5×2.5 mm. The new FPGA consumes only 21 μA (1.2V V_{CC}) static core current. Intended for portable medical monitors, mobile consumer devices, and compact embedded systems, the new device allows designers to build high-speed data-processing engines in a small footprint while maintaining low power requirements.

The new device is supported by the Lattice iCEcube2 development platform, which includes a synthesis tool with the company's placement and routing tools as well as the Aldec Active-HDL simulation software with waveform viewer and an RTL/gate-level mixed-language simulator. Additional capabilities in iCEcube2 include a project navigator, constraint editor, floorplanner, package viewer, power estimator, and static timing analyzer.

Available now in sample quantities, the iCE40 LP384 FPGA devices are offered in multiple packaging options, including 32-pin QFNs (5×5 mm), 36-ball ucBGAs (2.5×2.5 mm), and 49-pin ucBGAs (3×3 mm).

—by Stephen Evanczuk

▷**Lattice Semiconductor**, www.latticesemi.com

The iCE40 LP384 FPGA from Lattice Semiconductor allows designers to build high-speed data-processing engines in a small footprint while maintaining low power requirements.

Smartphone barcodes now readable by all POS laser scanners

Roughly 350 billion paper coupons are issued each year in the United States, with a value of more than $470 billion. Paper coupons consume 13 million trees every year, and 99% are actually never even used. The dream of paperless "mobile couponing" has remained elusive because the vast majority of in-store laser scanners cannot scan 1-D barcodes displayed on mobile phones.

In response to this issue, AMS and Mobeam have formed a strategic partnership that will accelerate the ability of smartphones to transmit barcoded content that can be read by all point-of-sale (POS) laser scanners. The partnership will provide an integrated solution consisting of AMS's light sensors and Mobeam's light-based beaming technology.

Existing red-laser scanners cannot "see" barcodes displayed on a phone screen due to random polarization of the laser beam and the screen. Mobeam's patented light-based communications (LBC) technology and software transform a barcode into a beam of light that is readable by the laser scanner.

The AMS TMD3990 utilizes the same proximity IR LED used in smartphones today for disabling touchscreen/backlight displays. Handset manufacturers will be able to integrate the technology easily and without the need to add additional components.

Full production of the TMD3990, which combines color light sensing, proximity detection, and an IR LED barcode transmit function, is scheduled for 3Q13.

—by Steve Taranovich

▶**AMS**, www.ams.com
▶**Mobeam**, www.mobeam.com

04.13

Star Trek-style "tractor beam" is created on tiny scale

A team of scientists has created a miniature real-life "tractor beam"—a beam of energy that can attract one object to another from a distance—as was depicted in the *Star Trek* television series. In this case, researchers from the University of St Andrews (Scotland) and the Institute of Scientific Instruments (Czech Republic) were able to use a light beam to draw microscopic objects toward the light source.

Normally, when matter and light interact, matter is pushed away by the radiation pressure of light, such as is observed with comet tails pointing away from the sun. The scientists, however, were able to generate a special optical field that reverses that effect and produces a "negative" force acting upon the minuscule particles that causes them to be pulled against the photon stream.

This is claimed to be the first experimental realization of the concept of such a negative optical force. Further, according to the researchers, the force is very selective to the properties, such as size and composition, of the particles that it acts on, which could allow simple and inexpensive optical sorting of micro-objects.

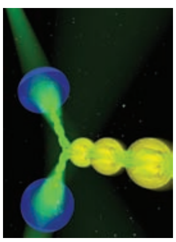

Within the experimental system, a light beam is converted into a pulling device that gathers micro-objects, similar to using a chain (courtesy University of St Andrews).

Practical areas that could benefit from this research include biomedical applications and fields involving intricate engineering. For more information, see the original article published by *Nature Photonics* ("Experimental demonstration of optical transport, sorting and self-arrangement using a 'tractor beam,'" bit.ly/Z8ir5B).—**by Rich Pell**

▷**University of St Andrews**, www.st-andrews.ac.uk

▷**Institute of Scientific Instruments**, www.isibrno.cz/index.php

VOICES

SparkFun's Chris Taylor: Community key to open-source hardware

When it comes down to it, engineer and seven-year SparkFun Electronics veteran Chris Taylor believes anything can be open source, to be shared and improved. He recently spoke with *EDN* about open-source hardware (OSH). What follows are excerpts of that discussion. Find more of the interview online at www.edn.com/4411373.

Does open-source hardware have a place in the world of professional engineering?

A That's a complicated question. There will always be a place for open source in the professional market because, in many people's opinions, it's the way to prototype. You can start from the ground up with whatever you are designing, or you can take a little bit from open-source companies that have already done the work and written the tutorials, and build up whatever prototype you are designing and go through the design phases a lot faster.

Where that gets complicated is open-source licensing. That is still nebulous. For example, the hardware could have a different license than the software or the firmware.

Open source brings up a lot of questions on "trust." Without IP, without patents, how do you know what can be trusted?

A With any open-source design, there should be a certain amount of caution. Regardless of the design, you're going to have to learn a great deal of it yourself. But in a lot of cases, we save the designer's digging through data sheets, testing, and retesting, by providing example code, boards, and layouts.

Yes, there is always going to be that fear. The open-source community tries to mitigate that [fear] as best they can by having a community around it. When a design is bad, you are going to hear about it immediately, but because it is open source, the design cycle is short and it can be rectified quickly.

If there is this community and sharing, how can engineers stay competitive with their designs?

A The most valuable element that the community lends is the improvement cycle, the feedback cycle. In our case at SparkFun, when we post a design's code online and someone takes that code, uses it, and finds an improvement or an error, we can make the improvement immediately, thanks to the community. If hardware has an improvement to be made, that is put as a comment on the product page as a forum. Because it's open

source, anyone can make an improvement. So the next version of that product is going to be better because of that community. That feedback, coupled with agility to turn out a new product, is the benefit to having the open-source community.

Why should engineers do anything open source if they're not going to get paid for it?

A In open source, a lot of that concern comes from wondering "What if someone steals my design?" Well, good; that's kind of the point. Let them steal it. If they can make it better, you can steal it right back and make yours better. That's what makes the product a quality product. Then you design the next cool thing. The information is free, but at SparkFun, the hardware is where we make the money.

Where do you see OSH in the future for engineering?

A Where you now start to see open source come into play and shaking things up is in 3-D printers and any sort of low-volume home production devices and businesses, because now we are starting to see the ability of people to create more complete objects from designs that would normally be closed source.

People are taking closed-source designs and replicating them, scanning them into 3-D files and printing them out on 3-D printers. The ability to create more and more complex, complete devices is where we are going, and it's a really exciting and adventurous new ground. This is where closed source is starting to butt up against open source.

We've always, as a company and a group of engineers, and myself personally, been interested in taking something that is totally ubiquitous but that someone wouldn't necessarily make themselves and opening that design. Anything can be open source, to be shared and improved.

Do you have any advice for established engineers who are looking to dip their toes into OSH?

A If an engineer is considering a design, by all means, open-source products are the best and fastest way—and a fun way—to do it. If you are doing a home project or planning to create an on-the-shelf product, the advice is really just go for it. You've got nothing to lose, and you have a whole community of people vetting designs for you. Not only can your good idea be profitable, but you're going to have fun in the long run.—**interview by Suzanne Deffree**

designwest

Taylor will be speaking as part of the Open Source Hardware Panel Discussion April 23 at DESIGN West, hosted by UBM Tech. Get more information and register at www.ubmdesign.com.

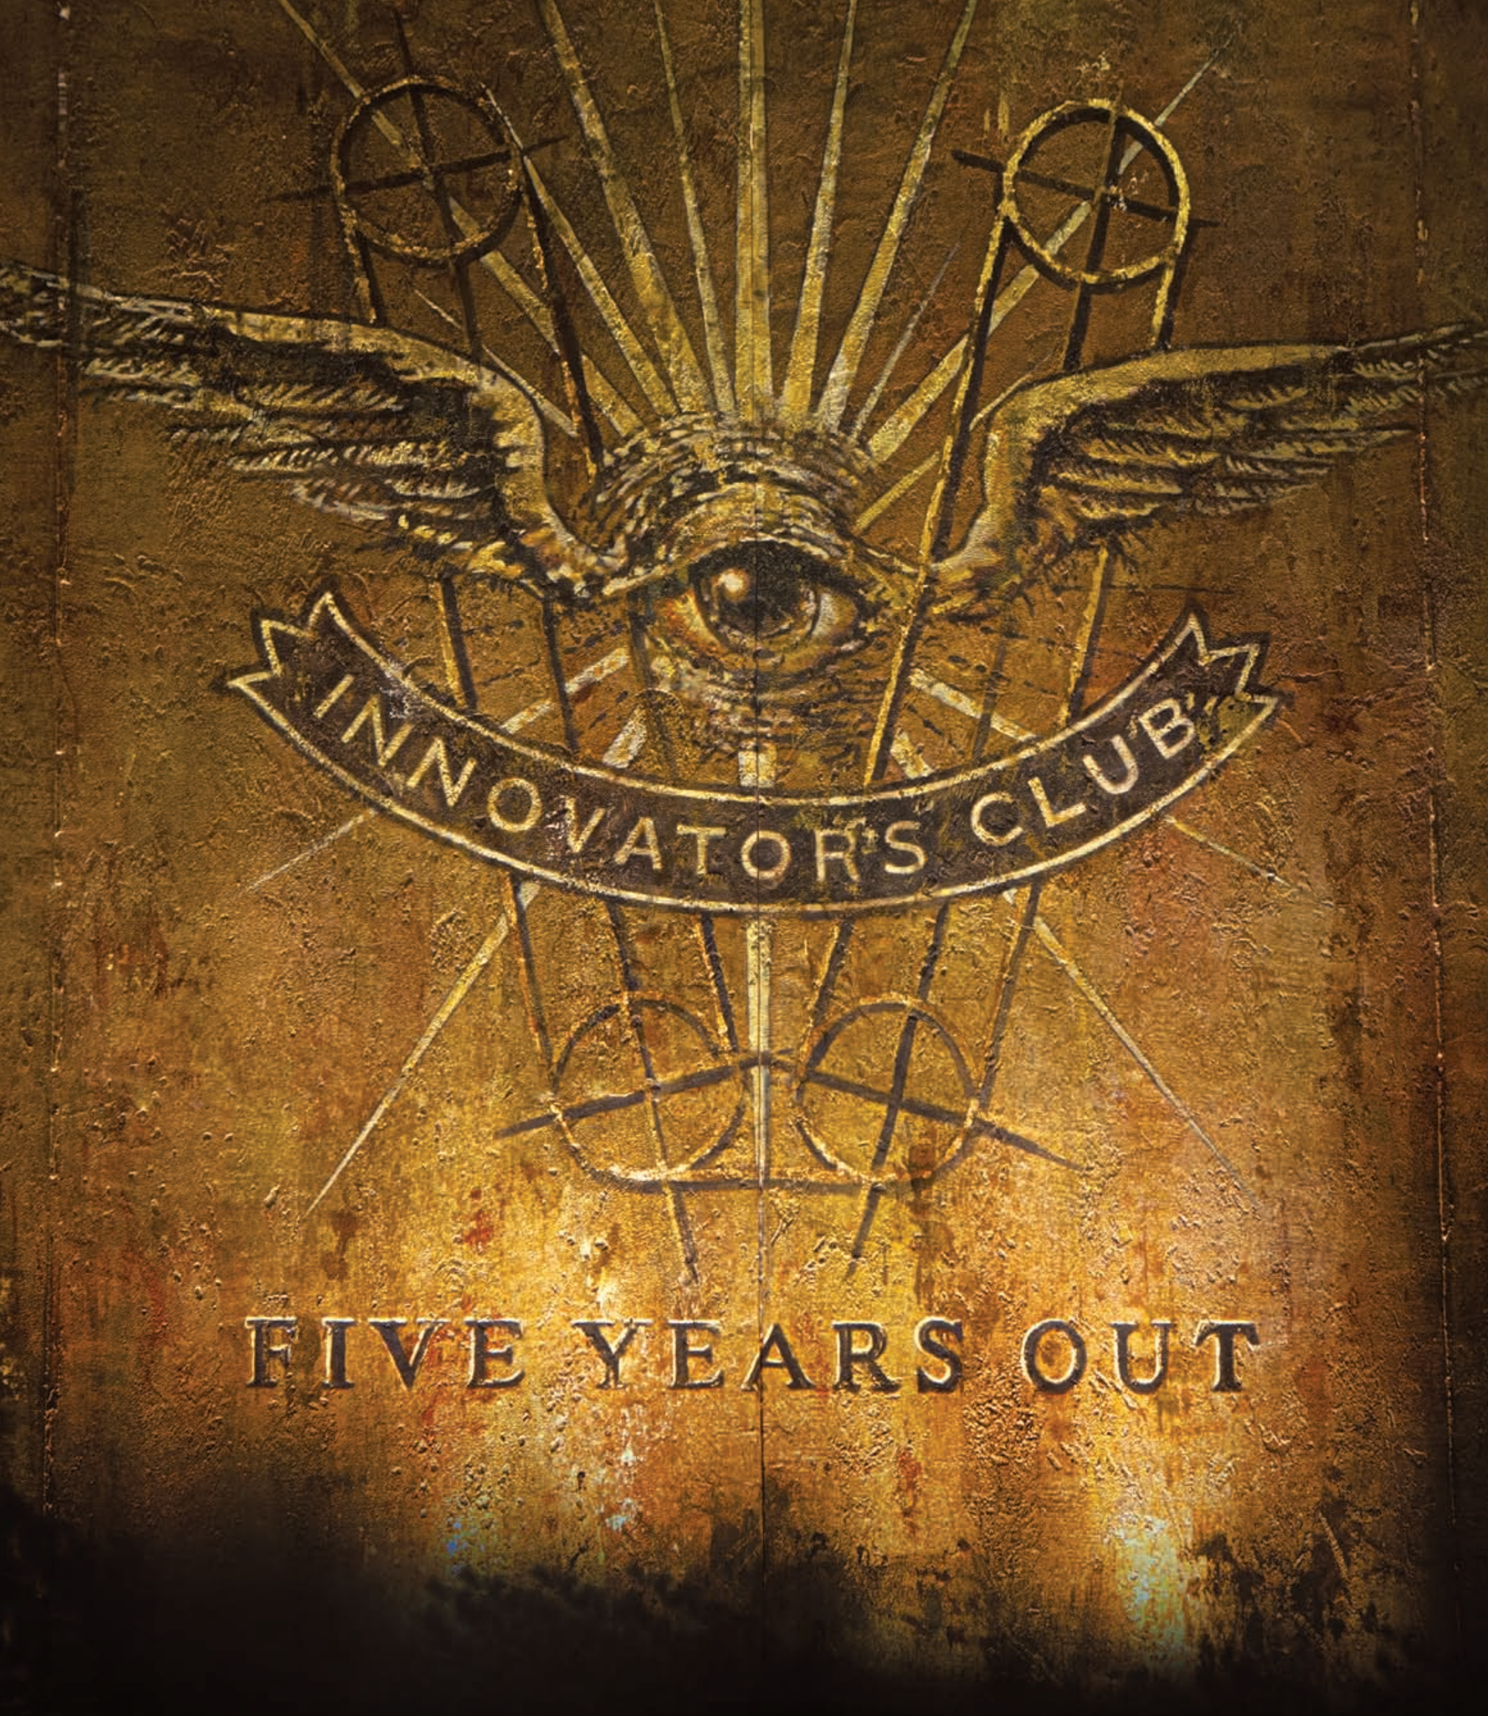
INNOVATOR'S CLUB
FIVE YEARS OUT

BY HOWARD JOHNSON, PhD

Measuring nothing

Every scope probe picks up extraneous noise. Some of that noise is self-generated, and some may be generated by the system under test. When looking at a noisy, jittery signal, how can you tell which parts of the signal are "real" and which parts derive from noise and interference? There is only one way, and that way, if you embrace it, leads to remarkable insights about noise, grounding, and the nature of digital systems.

The only way to directly observe noise and interference is to attempt to measure nothing. With your probe in place, grounded as it will be for the actual signal measurement, touch the tip of the probe to any nearby point of ground. This configuration is called a null experiment. Ideally, you should see zero, zip, nada, or, as the English call it, "naught."

What you actually observe is your own noise floor, a plethora of noise sources, a whole ecosystem of interferences all superimposed. Creative use of your trigger circuits combined with vertical averaging can often pull apart these tiny effects, deeply buried in a sea of foam, for close inspection. You can learn a great deal measuring nothing.

In theory, whatever noise the probe picks up in your null experiment will appear as noise superimposed onto your actual signal, provided the probe is held in a similar physical position. Two main things cause the noise you will see: one, currents flowing on the probe shield due to differences between the electrical potential of the digital logic ground and the scope, and two, interactions between the electromagnetic fields surrounding the device under test and the probe or probe wiring.

To determine how much noise the former source creates, keep the probe connected to its own ground, but disconnect the probe and probe ground entirely from the device under test. Keep the probe topology otherwise similar to the null experiment. This procedure eliminates the probe-shield currents, leaving only the electromagnetic pickup. If probe-shield currents are a serious problem, try a differential probe, with one leg on the signal and the other on digital logic ground. Since both inputs to a differential probe have high impedances—much higher than the impedance of a single-ended probe's ground connection—little shield current will flow during this configuration.

Every probe picks up extraneous noise.

Regarding the latter source, first determine if the noise is coming from the device under test or something else in the room. With the probe connected to its own ground, but still disconnected from the device under test, pick up the probe and wave it around. Use the probe as a magnetic-field sniffer to locate the culprit. Sometimes a fluorescent light or other circuit may induce noise in this configuration. If so, turn it off.

If electromagnetic noise seems to be coming from the device under test, check the length of the ground attachment between the scope probe and the system. The smaller you make the loop from the signal source, to the probe, and back through the probe's ground connection, the less noise your probe receives (**Figure 1**). Reduce the size of that loop, and your null-experiment results should improve.

Understanding why this, and a hundred other tricks, works for noise abatement is all part of the art of digital design. If you want to master that art, follow my advice: "Measure not the thing you know; measure naught." EDN

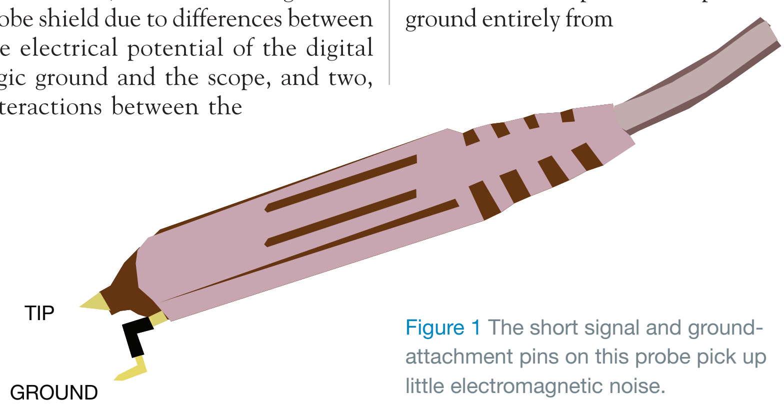

Figure 1 The short signal and ground-attachment pins on this probe pick up little electromagnetic noise.

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

SNAP DRAGON
locking CR2032 holder
BATTERYHOLDERS.COM

WANT MORE?

- Read more about the Insignia Teardown, and see additional photos, at www.designingwithLEDs.com.
- Go to www.edn.com/learning/teardowns for more Teardowns.

LED light shrinks size, cost with nonisolated driver

LED-bulb prices are dropping. A year ago you could expect to pay $50 for a Philips dimmable 60W-replacement LED bulb, while today you can go to Best Buy and purchase its house-brand 8W, 800 lumens Insignia 60W-replacement bulb for just $17. What has changed in LED-bulb design to allow this price drop? Tearing apart the bulb gives us a look into some design trends in LED lighting, such as how the LEDs are placed within the bulb and what driver architecture is used.

The Insignia bulb has a shape similar to the familiar incandescent light, with the addition of three metal heat-sink fins and a plastic bulb instead of glass.

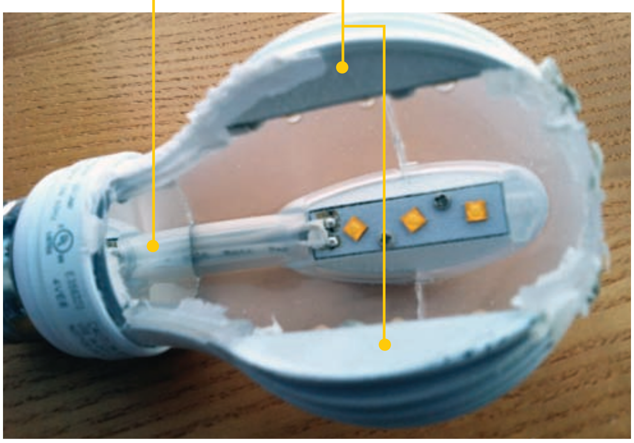

The plastic bulb cover was removed with a Dremel tool, exposing the six Cree white LEDs that illuminate the bulb's light-mixing chamber, which allows an even glow with no pixelation. The metal fins on which the LEDs are mounted serve both to elevate the LEDs and as heat sinks. At the bottom of the mixing chamber is a paper-thin aluminum reflector that helps reflect the light up and out of the bulb. All of the electronics for this bulb lie beneath the mirror in the base of the bulb, in a separate and encapsulated compartment.

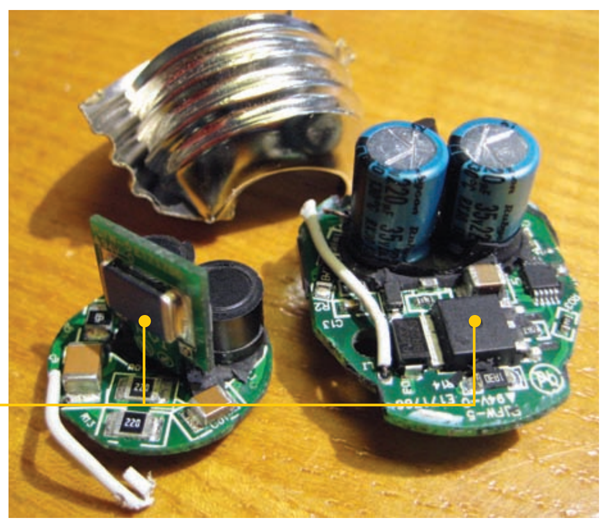

Removing the rubbery potting compound shows that the electronics are mounted on two PCBs that nestle together. Here, the PCBs are separated and next to the bulb's base.


INNOVATORS SWEAT THE HUGE STUFF.
Bringing breakthrough products to market faster for less. Innovators don't tackle small problems. That's why they start with Wind River. They know they'll get the industry's most comprehensive device development portfolio—an unrivaled set of solutions, service and support that's made us a world leader in embedded software for more than three decades. But beyond that, they'll get the safety, security and reliability they need to confidently pursue their vision, and transform their dreams into reality. No sweat.
WIND RIVER
INNOVATORS START HERE.
www.windriver.com/innovators

For comparison purposes, the driver electronics are shown for an LED-light teardown from about a year ago. Not only is the packaging quite different, but also there are a lot more electronics. For example, the older design has three electrolytic capacitors and a very large transformer.

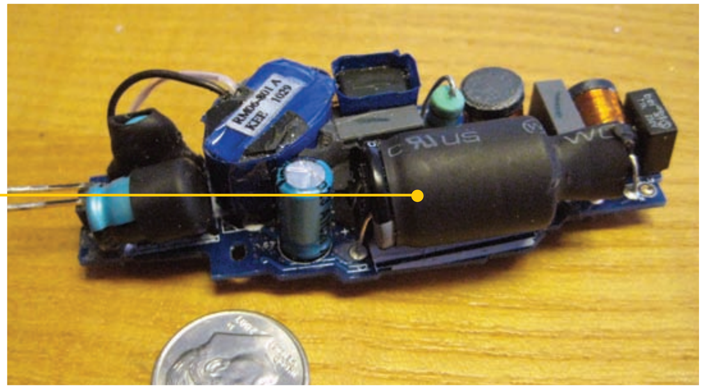

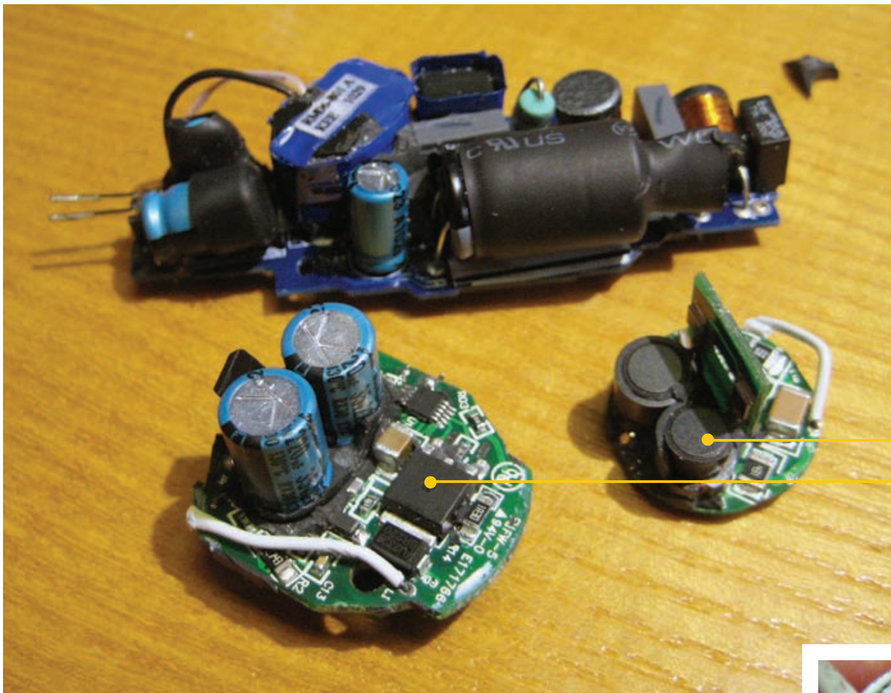

The two different generations of drivers are shown side by side: The Insignia driver has just two relatively small e-caps. This setup raises the question: What LED-driver IC does the bulb use, and how does it allow such a tiny driver?

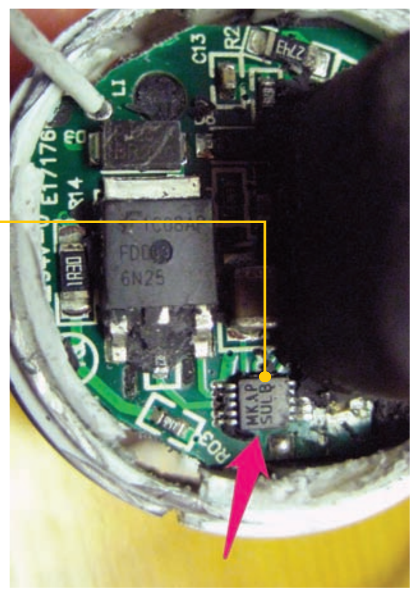

Question answered: The "SUL B" on the tiny IC is the marking for the Texas Instruments LM3445. There is no transformer, indicating that the LED-driver design is nonisolated. The incandescent-bulb design itself is nonisolated. If you broke the glass in an incandescent bulb while it was plugged in, you would have direct access to the ac-line power. Clearly, nonisolated designs can be made fully compliant with UL specifications. Note: While a nonisolated ac/dc LED-driver design can both be safe and meet UL specifications, developing and testing a nonisolated offline LED driver in the lab requires stringent lab-safety procedures.

Dimming is an important bulb characteristic for the US market. I used a Lutron Maestro dimming switch, with a programmable dimming control, and did a side-by-side comparison with an incandescent bulb. The Insignia dimmed consistently and smoothly, with a dimming profile similar to the incandescent bulb. Watch a video of the dimming test at http://bit.ly/145ROBK.


THiNK
QUICK-FIT TERMINALS
THiNK Keystone UL Recognized Quick-Fit Terminals
• UL Recognized Components • Easy, safe and quick way of connecting and disconnecting wires • Cost effective PCB connectivity • Sturdi-Mount features for increased PCB mounting strength • Durable Tin Plating for excellent solderability • Wide variety of configurations to suit most applications • Male tab sizes for .110"; .187"; .205"; .250" • Female receptacle sizes for .110"; .187"; .205"; .250"; .312" • Custom terminals can be manufactured to your specifications • Request Catalog M60
KEYSTONE
ELECTRONICS CORP.
It's what's on the inside that counts.
Tel: (718) 956-8900 • (800) 221-5510
e-mail: kec@keyelco.com • Web: www.keyelco.com
There's a keystone in every great invention ... Follow, THiNK & Learn at:

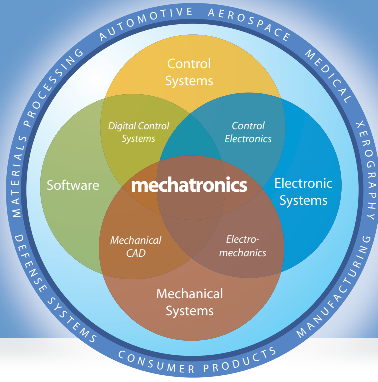

Excelling in the shades-of-gray real world

Gray-box modeling combines the physical and empirical to solve problems.

By Kevin C Craig, PhD

In a world where problems are often ignored and allowed to fester for months or years, engineers do not have that option, as engineering problems ignored may lead to financial collapse or, worse, loss of life. Engineers solve problems to help people, and they do that with a sense of urgency. In many situations, a combination of human-centered design and state-of-the-art technology yields feasible and sustainable solutions. In more complex situations, physical insight may be incomplete, and engineers perform experiments to validate what they do understand and inform what they don't. This approach to problem solving is called gray-box modeling (**Figure 1**).

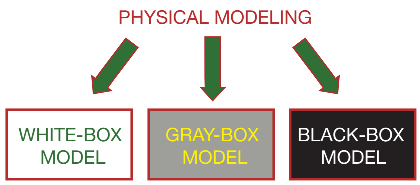

Figure 1 Approaches to problem solving include white-, gray-, and black-box modeling. The approach used depends on what is known about a system.

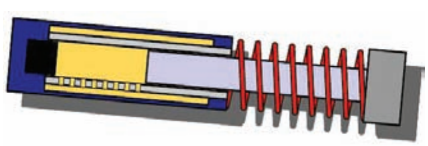

Figure 2 A shock absorber comprises a cylinder surrounding a movable piston.

Let's use as an example a shock absorber (**Figure 2**). It consists of a cylinder surrounding a movable piston. Moved by a shock from the outside, the piston compresses oil inside the cylinder through holes in the wall, thus dampening oscillations. A spring pushes the piston back to its original position, and a rubber stopper prevents the piston from impacting the walls of the cylinder when shocks are too strong. The shock absorber comprises the interaction of the mechanical movements of rigid bodies, the viscoelastic dynamics of fluids, the elastic behavior of springs, and the deformations of elastic-plastic materials. Looking from the outside, we are aware of only the phenomenological properties. We observe aspects such as nonlinear stiffness, nonlinear viscous damping at high frequencies, and hysteretic effects at low frequencies, but we are not able to assign these phenomena to the individual parts of the shock absorber.

The shock absorber is integrated into a suspension system that must be designed and controlled. Mathematical equations are needed to predict the behavior of both the shock absorber and the integrated suspension system. A mathematical model of this physical system therefore must be created, and this model is based on simplifying assumptions. Depending on the nature of the simplifying assumptions, models of varying complexity and fidelity result.

Information about the real system comes from the inside and the outside. Looking from the inside, we apply the laws of nature, together with the constitutive equations of the components, to the physical model to generate the mathematical equations of motion. These are solved by numerical simulation to predict the behavior of the physical model, which must be experimentally verified. Because we use our insight and understanding of the way the system works to create the model, we call this model—an approximate image of the physical system—a white-box model.

Looking from the outside, measurements alone give no insight into the real system, and thus no understanding of how the real system works is brought into the construction of the model. A mathematical model is chosen that fits optimally the measured data. This type of model is called a black-box model.

In reality, modeling is always something in between these two views, resulting in a gray-box model. White-box models are approximations of reality and always need experiments to identify parameters in the model, validate model predictions, and show where the model is deficient. The set of possible black-box models should always be guided by some knowledge of the inner workings of the real system. Since the focus of the engineer is solving the problem, the gray-box modeling approach is intuitive and obvious. Engineers contribute to society not only technological solutions to problems but also a solution process that transcends boundaries. EDN


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


IF WE ™
help people with
a damaged skeleton,
could they walk again?
A robotic skeleton – a dream
our software could bring to life.
Innovative thinkers everywhere use the 3DEXPERIENCE software platform from Dassault Systèmes to explore the true impact of their ideas. Insights from the virtual world help companies combine electronic, mechanical and software technologies to design innovative solutions that overcome human limitations. How long before bionic humans are real life, not science fiction?
3D
V+R
3DEXPERIENCE
It takes a special kind of compass to understand the present and navigate the future.
Discover our Industry Solution Experience:
HT body
3DS.COM/HIGHTECH
DASSAULT SYSTEMES
IF WE ask the right questions
we can change the world.

WHAT MAKES INDUSTRIAL SENSORS GO AWRY?

SENSORS THAT MEASURE PARAMETERS SUCH AS **PRESSURE, TEMPERATURE, TOXIC GAS, AND pH** ARE ABUNDANT AND MAKE INDUSTRIAL PROCESSING SAFER, MORE EFFICIENT, AND LESS COSTLY. EACH SENSOR TYPE, HOWEVER, HAS ITS OWN SET OF UNIQUE CHARACTERISTICS, RESULTING IN VARIOUS AND COMPLEX DESIGN CHALLENGES.

BY JASON SEITZ • TEXAS INSTRUMENTS

One of the most popular measurements in industrial processes is temperature. Temperature can be measured by a variety of sensor types, including thermocouples, resistance-temperature detectors (RTDs), and thermistors.

To measure the largest temperature range, a system designer often uses a thermocouple. For example, a Type C thermocouple has a measurement temperature range between 0 and 2320°C. The thermocouple's principle of operation is based on the Seebeck effect, where, if two dissimilar metals are placed together, a voltage is produced that is proportional to the temperature at the junction. Thermocouples are bipolar devices that produce a positive or negative voltage, depending on the sensing, or "hot" junction, temperature relative to the reference, or "cold" junction, temperature. First, you need a bias to the thermocouple so it won't rail against ground in a single-supply system. Next, measure the cold-junction temperature to obtain the temperature being measured. One drawback to thermocouples when compared with other temperature sensors is limited accuracy, which is typically worse than ±1°C.

IMAGE: ISTOCK

If the system requires greater precision over a reduced temperature range—say less than 660°C—a designer can implement the measurement with an RTD, which can be as accurate to below ±1°C. RTDs are resistive elements whose resistance depends on the ambient temperature in which they are placed. They come in two-, three-, and four-wire configurations. Increasing the number of wires increases accuracy. RTDs require an excitation in the form of a current source. Current-source values often are 100 μA to 1 mA to handle PT100 (100Ω at 0°C) and PT1000 RTDs (1000Ω at 0°C).

For accuracy up to ±0.1°C, with the trade-off of an even smaller temperature range (less than 100°C), you can use thermistors. Similar to RTDs, thermistors' resistance also changes with temperature. Thermistors typically are connected in a resistor-divider configuration where the other resistor in the divider is the same as the nominal value (value at room temperature, 25°C) of the thermistor. One end of the thermistor is connected to the supply voltage; the other end is connected to the other resistor, which in turn is connected to ground (**Figure 1**). To determine temperature, measure the voltage at the divider's center point. You would expect +V/2 at 25°C. For any deviations from this, you can calculate the thermistor's resistance and use a look-up table to determine the ambient temperature being measured.

In summary, temperature sensors need bias (voltage or current). When it comes to thermocouples, cold-junction compensation is needed. The Texas Instruments LMP90100 is a 24-bit sensor analog front end (AFE) with four differential or seven single-ended inputs, two matched programmable current sources, and continuous background calibration (**Figure 2**). This integrated configurable chip is suitable for addressing the various design challenges associated with temperature sensors.

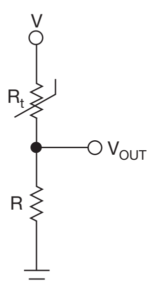

Figure 1 A thermistor circuit is shown with one end connected to the supply voltage and the other to the other resistor and the ground.

AT A GLANCE

- Temperature, an important measurement in industrial processes, can be determined using thermocouples, resistance-temperature detectors, and thermistors.
- Strain gauges and load cells that utilize the Wheatstone-bridge circuit are popular for measuring pressure, force, and weight.
- Electrochemical cells typically are used to measure various gases. An alternative is nondispersive infrared technology.
- Using an appropriate analog front end improves measurement accuracy while reducing design complexity.

Strain gauges and load cells that utilize the Wheatstone-bridge circuit are popular implementations for measuring pressure, force, and weight. Any strain or stress on the gauge causes a resistance change and subsequent voltage differential change on the sensor output (**Figure 3**). The voltages that come out of these sensors are low, typically in the millivolt range. To make the most accurate measurements, amplify this small voltage up to the data converter's full dynamic range. A programmable-gain-amplifier (PGA) stage allows for interfacing with multiple sensors as well as optimum flexibility. This stage should have low noise, low offset, and low-offset drift to ensure the best system performance.

These types of sensors also require excitation in the form of a bias voltage. A common failure mode for pressure sensors is incorrect measurements due to opens or shorts in the bridge. Even harder to detect are out-of-range signals caused by sensor damage or degradation over time. A way to catch all of these failure modes is to incorporate a diagnostic circuit. This circuit injects a small current, sometimes referred to as "burnout" current, into the resistor ladder of the Wheatstone bridge and then measures the resulting voltages.

If, for example, the bridge outputs are at the same potential (+V/2), is it because there is no pressure on the gauge or because you have a failure in the system that is shorting the outputs? Injecting current into one of the differential outputs and measuring the differential voltage between the outputs can answer this question. In normal operation, the differential voltage will be the voltage drop across the resistors in the bridge. However, if there is indeed a short, there will be little or no voltage drop.

In short, Wheatstone-bridge sensors require an excitation voltage, a low-noise/offset PGA, and diagnostics. The LMP90100 also mates well with these types of sensors. Its continuous background sensor diagnostics allow detection of open-circuit, short-circuit, and out-of-range signals. By injecting burnout currents into a channel after it has undergone a conversion, it avoids burnout-current injection from interfering with the channel's conversion result. Diagnostics provide continuous noninvasive failure detection, aid in root-cause analysis, and minimize system downtime.

Electrochemical cells typically are used to measure a wide variety of toxic and nontoxic gases, such as carbon monoxide, oxygen, and hydrogen. They are based on the principles of chemical oxidation and reduction, and produce a current in proportion to the measured gas. Most cells comprise three electrodes: working (WE), counter (CE), and reference (RE). The WE oxidizes or reduces the target gas and produces a current proportional to the gas concentration. The CE balances the generated current, and the RE maintains the working electrode potential to ensure the proper region of operation. Electrochemical cells are intended to interface with a potentiostat circuit. This potentiostat circuit provides current (and biasing, if required) to the CE. It maintains the WE at the same potential as the RE and converts the output current from the WE into a voltage using a transimpedance amplifier (TIA).

Electrochemical sensors, like many sensors, have a dependence on tempera-

JOIN THE CONVERSATION

To comment on this article, go to **www.edn.com/4407697**.

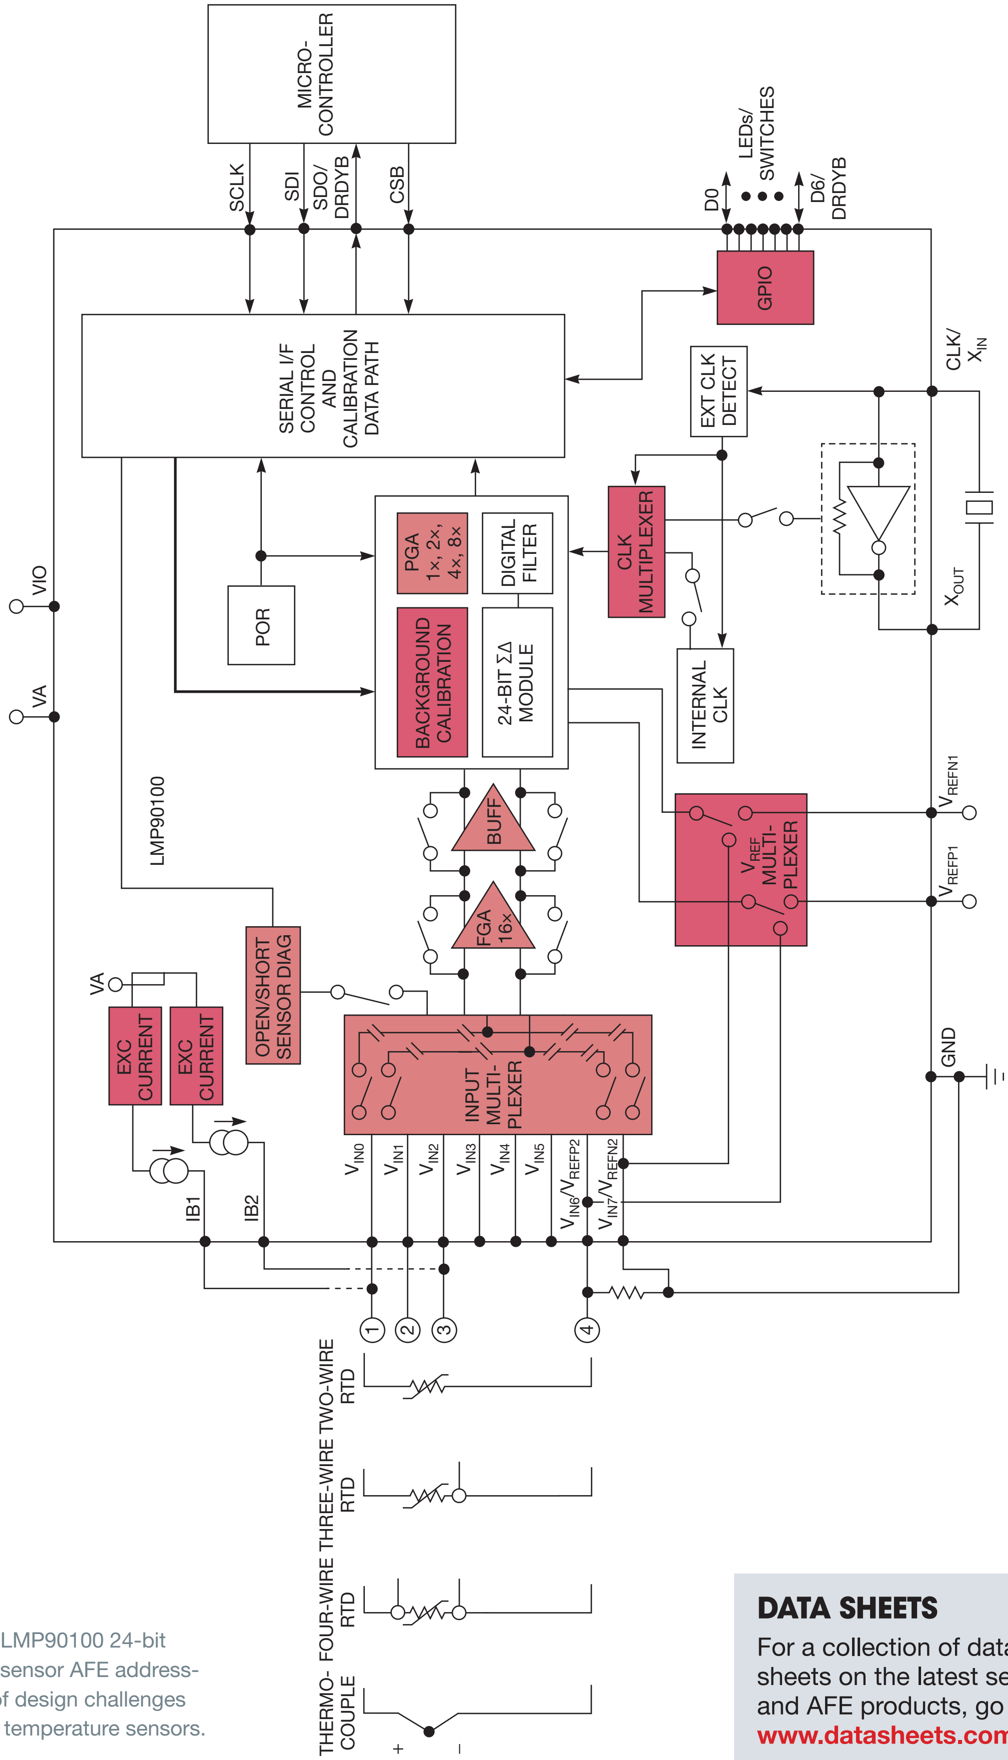

Figure 2 The LMP90100 24-bit configurable sensor AFE addresses a variety of design challenges inherent with temperature sensors.

DATA SHEETS

For a collection of data sheets on the latest sensor and AFE products, go to **www.datasheets.com**.

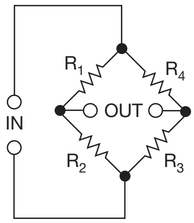

Figure 3 A Wheatstone-bridge circuit is shown. Any strain or stress on the gauge causes a resistance change and subsequent voltage differential change on the sensor output.

ture. To enable the best performance, measure the cell's temperature. Make appropriate temperature corrections based on that cell's performance versus temperature plots, which can be found in the data sheet.

The sensor, gas type, and gas-concentration level dictate how much current will be output at the sensor's working electrode. To handle this variability, use a TIA with adjustable gain. Currents in the level between one to hundreds of microamperes are possible, so having a TIA gain in the one to hundreds of kilohms range is sufficient.

Different sensors require different biasing; some require a zero bias. Be aware of these requirements so that the current produced by the sensor meets specifications. Whether the cell goes through an oxidation (CO) or reduction (NO_2) reaction to the measured gas determines if the cell produces a current into or out of the WE, respectively. The voltage at the TIA's noninverting pin should be level-shifted appropriately to ensure maximum gain without saturating the amplifier's output in single-supply systems. For example, the TIA produces an output voltage governed by the equation $V_{OUT}=-I_{IN}\times R_{FEEDBACK}$, where I_{IN} is current going toward the TIA across the feedback resistor. If the current into the TIA is positive (reduction reaction), V_{OUT} will be negative in reference to the noninverting pin voltage. That voltage should be raised to avoid railing the output to the negative supply.

Basically, it is essential that electrochemical cells have temperature correction and a potentiostat circuit that provides current sinking/sourcing, voltage biasing, current-to-voltage conversion, and level shifting. The TI LMP91000, a configurable AFE potentiostat, satisfies these functions. It contains a complete potentiostat circuit with sink and source capability along with programmable TIA gain, electrochemical-cell bias, and internal zero voltage. Moreover, this sensor AFE contains an integrated temperature sensor and comes in a small 14-pin, 4-mm^2 package that allows positioning the device directly under the electrochemical cell for accurate temperature compensation and improved noise performance.

Not all gases can be accurately measured with an electrochemical cell. An alternative is to use nondispersive

LIKE THIS?

For more articles like this one, go to **www.edn.com/design/sensors**.

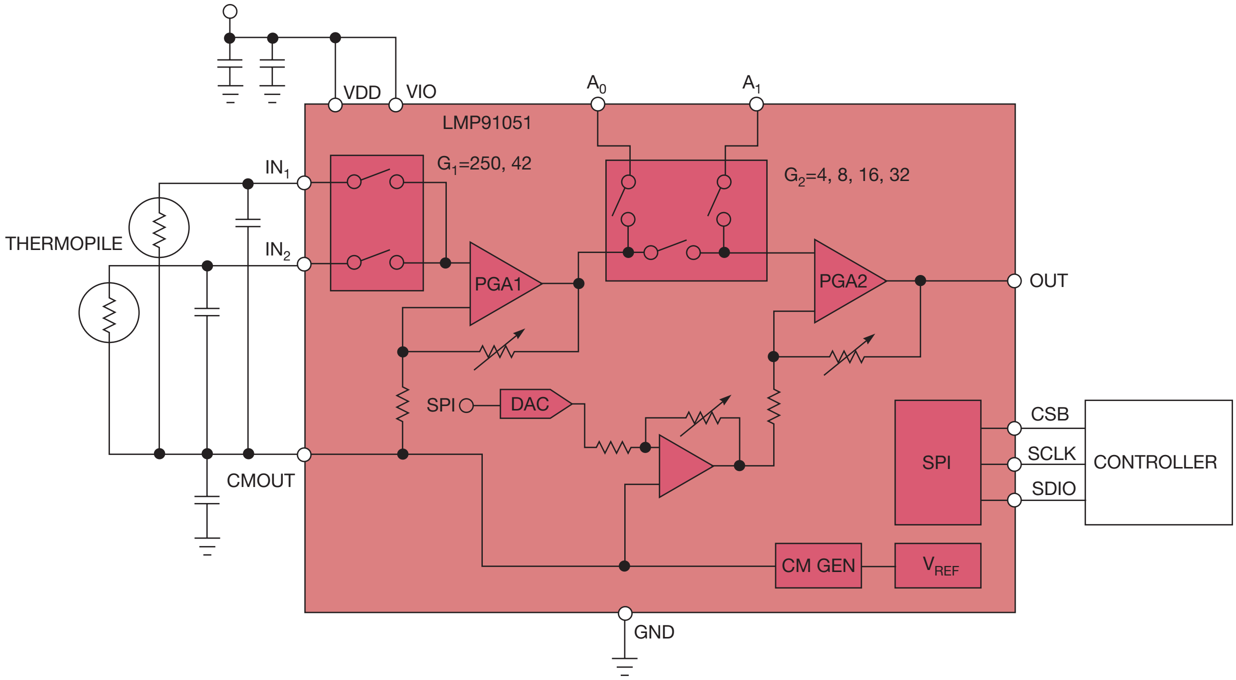

Figure 4 The level of integration provided by the LMP91051 configurable sensor AFE for NDIR sensing reduces design time, board space, power, and cost.

infrared (NDIR) technology, a type of IR spectroscopy. IR spectroscopy is based on the principle that most gas molecules absorb IR light. (Absorption occurs at a specific wavelength.) The amount of light absorbed is proportional to the gas concentration. Specifically, NDIR passes all IR light through the gas sample and uses an optical filter to isolate the wavelength of interest.

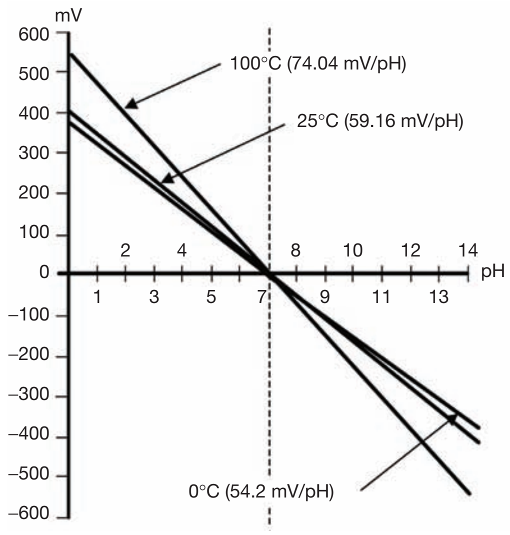

Figure 5 A pH electrode produces a voltage output that is linearly dependent upon the solution's pH being measured.

Typically, a thermopile with a built-in filter is used to detect the amount of a specific gas. For instance, because CO_2 has a strong absorbance at a wavelength of 4.26 μm, a bandpass filter is used to remove all light outside of this wavelength. Along with CO_2 and alcohol detection, NDIR gas sensors also can be used to detect greenhouse gases and refrigerants such as Freon.

A major obstacle with NDIR systems is accurately determining, over time, whether changes in the light transmitted to the detector are actually due to absorption from gas rather than deterioration of the light source or chamber contamination. Calibration is possible at the beginning of the NDIR system operation. To combat light-source deterioration and chamber contamination over time, however, ongoing calibration is required. This calibration can be expensive and time-consuming and is just not feasible in long-term field operation.

One way to solve this problem is to employ a reference channel in your system. This reference channel contains a detector that measures the light source in a range where no absorption occurs. Gas concentration now is determined by the ratio of the two transmitted light quantities. Any errors due to the light source's deviating are now cancelled out. This deviation results in long-term drift, which occurs over long periods of time. Hence, there is no need to simultaneously sample both the reference and active channel. You can use an input multiplexer to switch between the two channels, reducing system cost and complexity while maintaining accuracy.

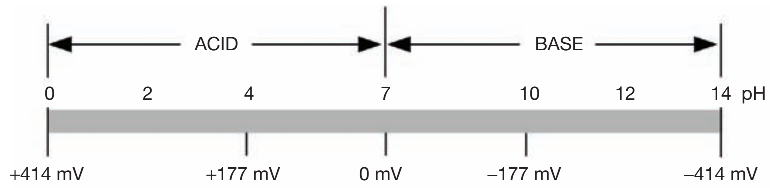

Figure 6 As the pH increases, the voltage produced by the pH-measuring electrode decreases.

Thermopiles used as IR detectors in NDIR systems produce a voltage based on the amount of incident light they receive in watts. The measured gas type, its absorption coefficient, and the gas-concentration range all affect the amount of incident light on the thermopile detector. The result is thermopile output voltages, typically in the range of tens of microvolts. Therefore, you need to design supporting electronics with the ability to amplify the thermopile output voltage with different gains. This can be handled by an AFE with a built-in PGA. Gain settings in the range of hundreds to thousands of V/V are required to amplify the small thermopile signal to the system's full-scale ADC, and achieve maximum system accuracy.

Another factor in NDIR system design is knowing how to handle the significant offset voltages associated with thermopile sensors. The thermopile is expected to have an offset component larger (up to 1 mV) than the actual signal, which limits the system's dynamic range. A way to minimize this problem is to integrate offset compensation into the system's electronics. One solution is to use a DAC to compensate for the measured offset. The system microcontroller can capture the level of offset and remove the offset by programming the DAC to shift the output toward the negative rail, zero scale. This solution utilizes the ADC's complete dynamic range, minimizing ADC resolution requirements.

Also, due to the thermopile's offset voltage, you need to bias the thermopile above ground. You can do this with a common-mode generator, which applies a common-mode voltage to the sensor. This approach level-shifts the thermopile sensor signal away from the negative rail, allowing for accurate sensing in the presence of sensor offset voltages.

Again, NDIR systems need a reference channel, adjustable amplification, offset compensation, and biasing. The LMP91051 configurable sensor AFE can handle this task for NDIR sensing applications (**Figure 4**). It has a

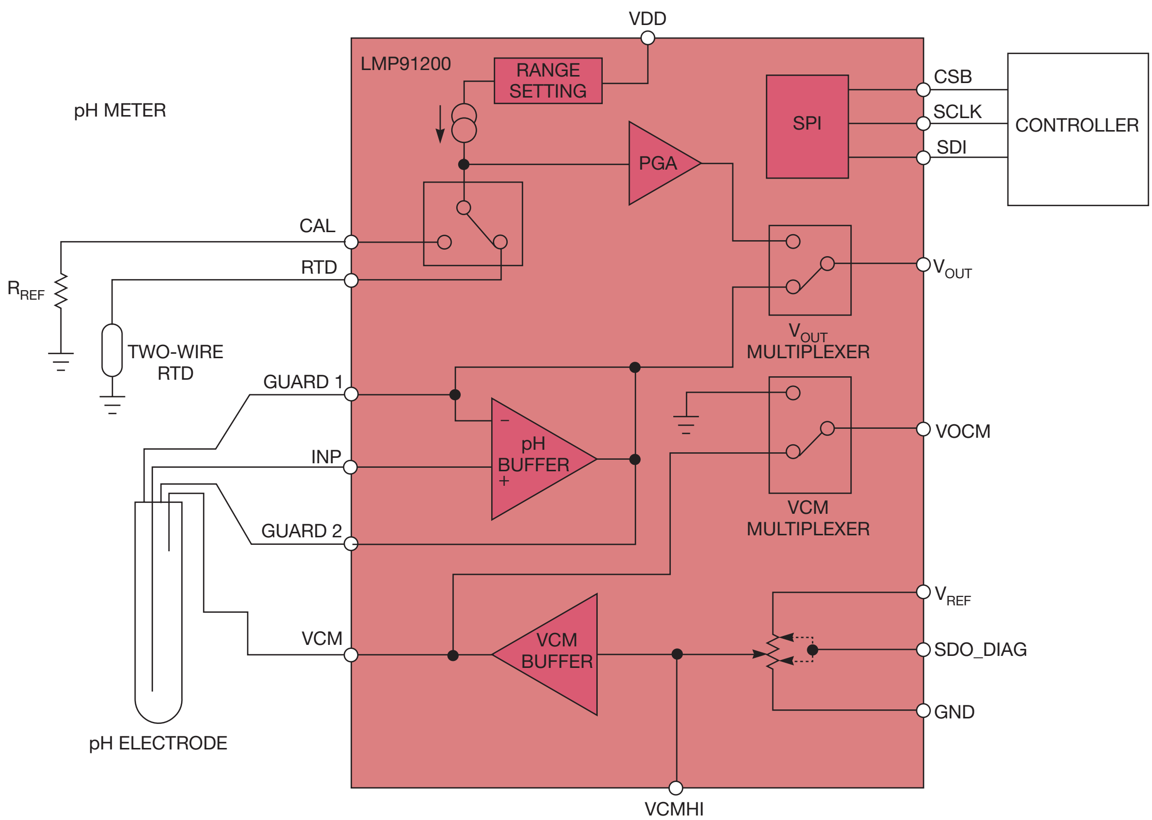

Figure 7 The LMP91200 configurable sensor AFE for chemical sensing addresses the need for high impedance, a low-input bias current interface, common-mode voltage, and temperature compensation.

dual-channel input to support an active and reference channel, PGA, adjustable offset cancellation DAC, and common-mode generator. The LMP91051's ability to integrate these important NDIR system blocks reduces design time, board space, power, and cost.

A pH electrode measures hydrogen-ion (H+) activity and produces an electrical potential or voltage. The pH electrode operates on the principle that an electric potential develops when two liquids of different pH come into contact at opposite sides of a thin glass membrane. These pH electrodes use the same principle to measure pH in a variety of applications, including water treatment, chemical processing, medical instrumentation, and environmental test systems.

The pH electrode is a passive sensor, which means no excitation source (voltage or current) is required. However, it's a bipolar sensor whose output can swing above and below the reference point. Therefore, in a single-supply system, the sensor needs to be referenced to a common-mode voltage (often half-supply) to prevent it from railing to ground.

The source impedance of a pH electrode is very high because the thin glass bulb has a large resistance, typically in the range of 10 to 1000 MΩ. This means that only a high-impedance measurement circuit can monitor the electrode. Furthermore, the circuit should have low-input bias current since even the smallest current injected into the high-impedance electrode creates a significant offset voltage and introduces measurement error into the system. Also, it is possible that current drawn from the pH electrode, while the system is shut down, can degrade the sensor over time. It is important, therefore, to keep the input bias current low, even when power is not supplied to the measurement circuit.

The pH electrode produces a voltage output that is linearly dependent upon the solution's pH being measured. The transfer function and pH scale in **figures 5** and **6** show that as the solution's pH increases, the voltage produced by the pH-measuring electrode decreases. Note that a pH electrode's sensitivity varies over temperature. Looking at the pH electrode transfer function shows that the sensitivity linearly increases with temperature. Due to this behavior, it is critical to know the solution's temperature being measured and compensate the measurement accordingly.

In the end, pH sensors require a high-impedance, low-input bias current interface, common-mode voltage, and temperature-compensation ability. The LMP91200 sensor AFE for chemical sensing addresses these functions (**Figure 7**). You can interface easily with RTDs with its programmable current source. Temperature-measurement accuracy is further enhanced with the multistep temperature-measurement feature, which removes error in the temperature signal path. The device's input bias current is in the range of only tens of fA at 25°C, minimizing error when connecting to a high-impedance pH electrode. Finally, the bias current is just hundreds of fA when the device is powered off, minimizing electrode degradation due to current draw over time.

Some of the design challenges associated with industrial sensors include the need for excitation, gain, temperature compensation, offset cancellation, current-to-voltage conversion, high-impedance interface, and diagnostics. Using an appropriate AFE improves measurement accuracy while reducing design complexity. EDN

REFERENCES

1 Seitz, Jason, "Designing CO_2 and alcohol-sensing applications," *ECN*, April 2012, http://bit.ly/ZdGV9u.

2 Seitz, Jason, "Designing with pH electrodes," Application Note AN-1852, October 2008, http://bit.ly/147ZPX2.

3 Seitz, Jason, "Designing with thermocouples," Application Note AN-1952, March 2009, http://bit.ly/Yrwz6P.

4 Stout, Mike, "Detect faults between a sensor and amplifier," *Electronic Design*, December 2010, http://bit.ly/YrwBvs.

AUTHOR'S BIOGRAPHY

Jason Seitz is a staff applications engineer in TI's Sensor Signal Path group, where he works on precision, low-power, and low-voltage analog systems. Seitz received his bachelor's degree in electrical engineering from the University of California at Davis, and his master's degree in electrical engineering from Santa Clara University, California.

SoC FPGAs COMBINE PERFORMANCE & FLEXIBILITY

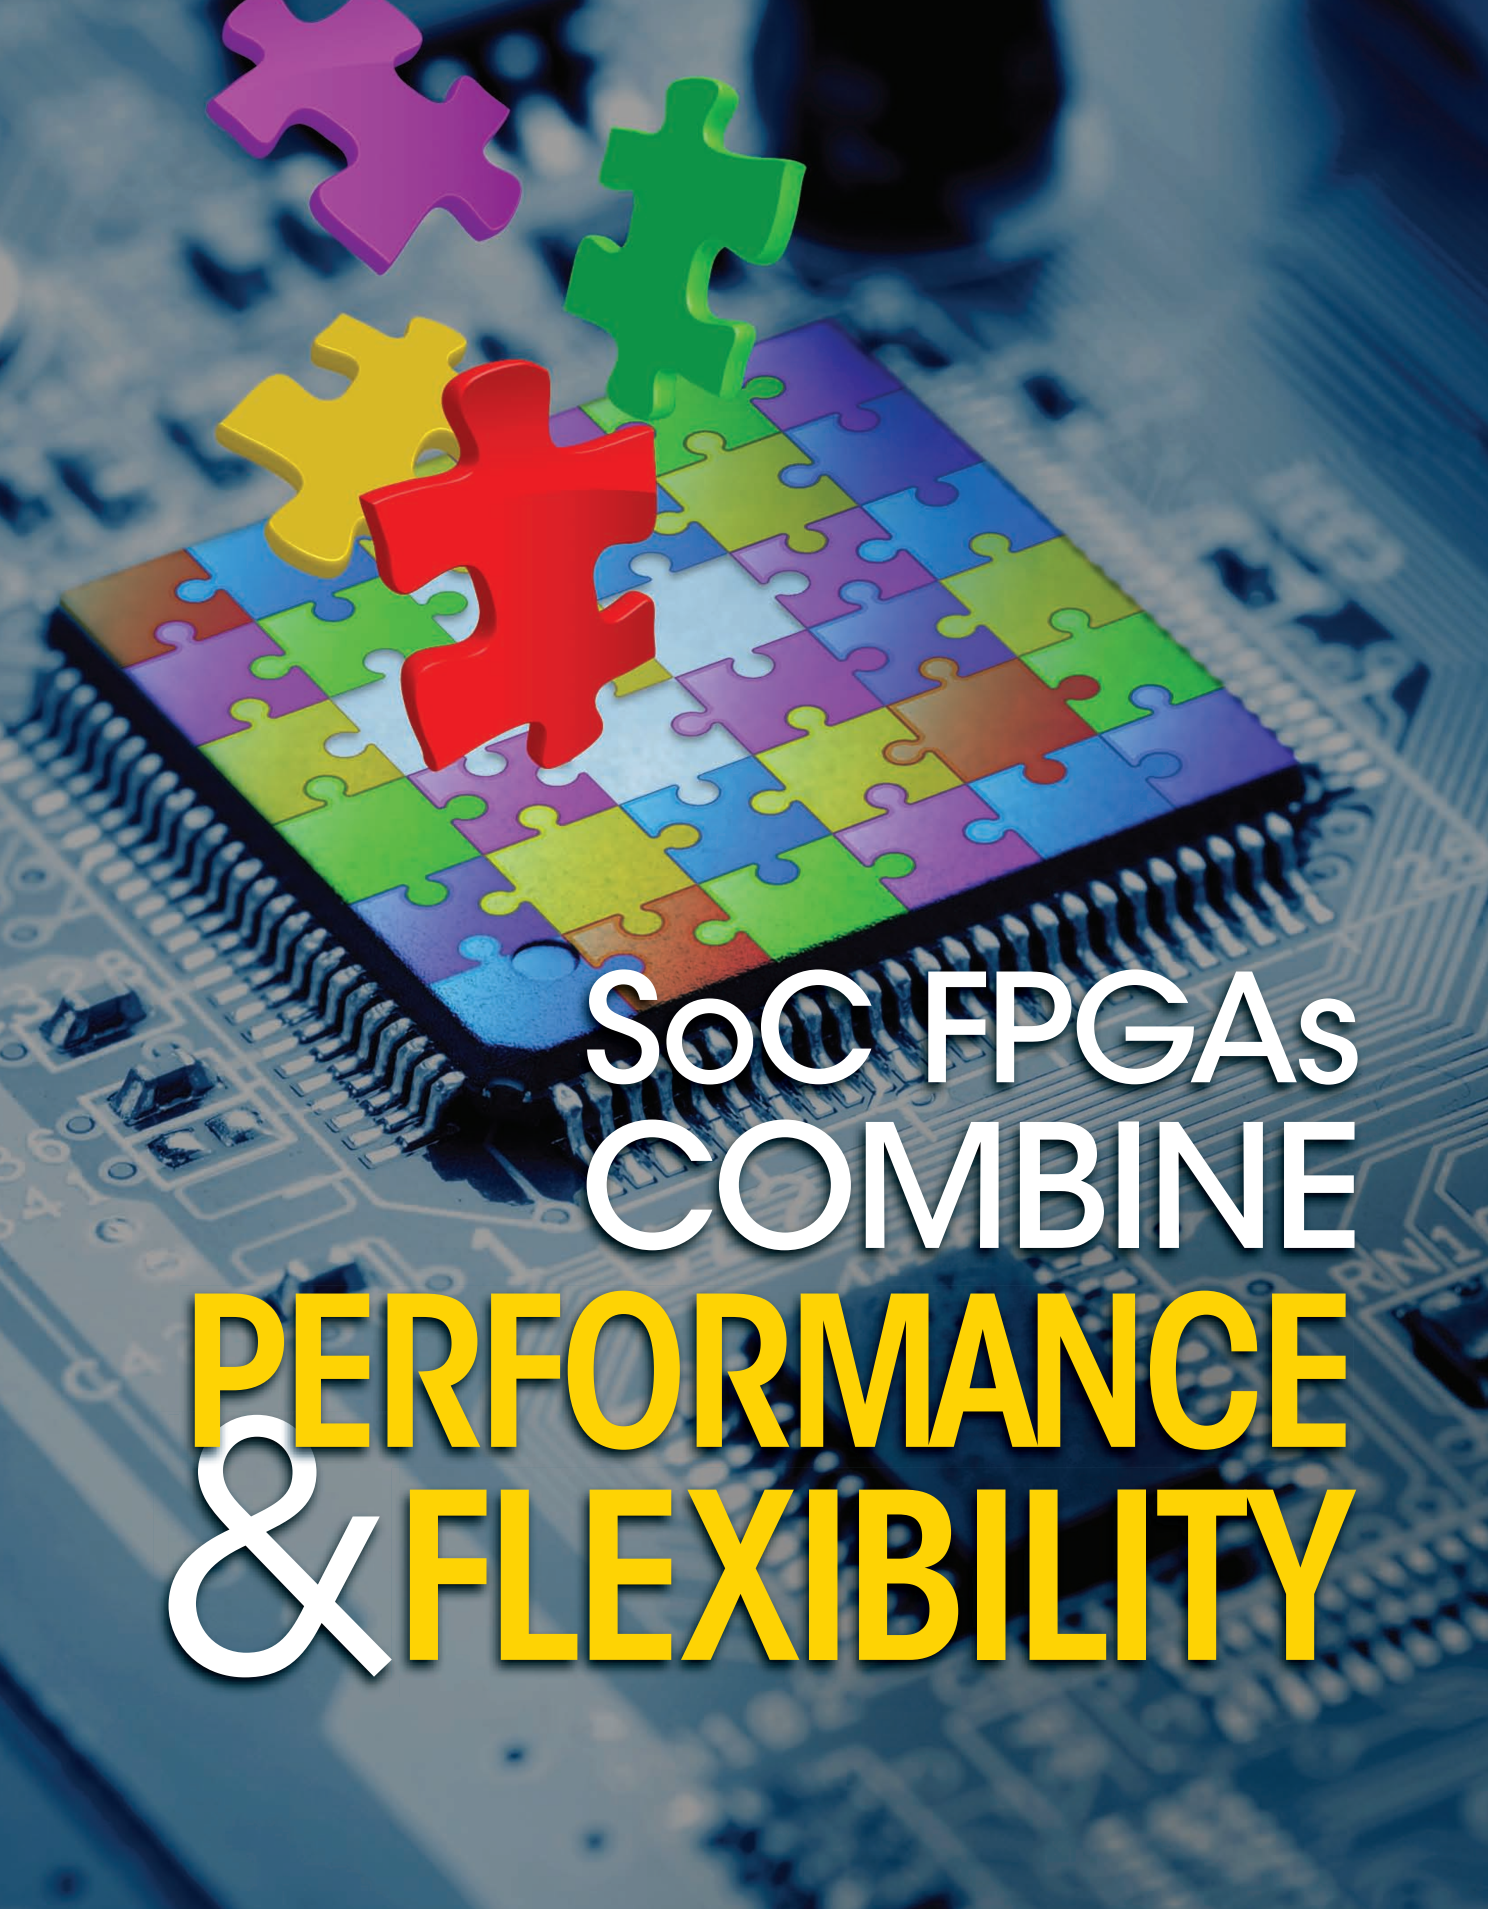

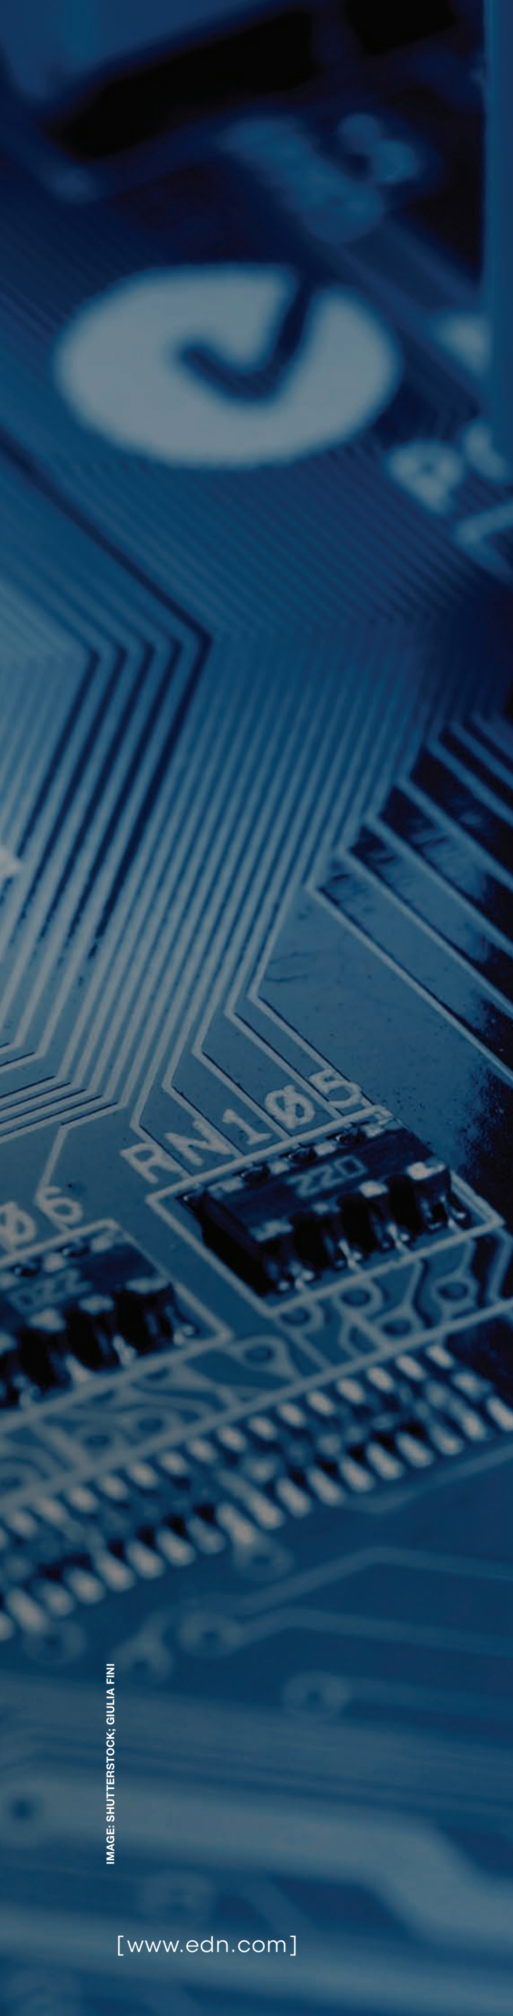

IMAGE: SHUTTERSTOCK; GIULIA FINI

EMBEDDED-SYSTEM ARCHITECTURES BUILT ON A COMBINATION OF MCUs AND FPGAs OFFER THE KIND OF ADAPTABILITY INCREASINGLY REQUIRED TO SUPPORT CHANGING DEMAND FOR GREATER FUNCTIONALITY ACROSS DIVERSE APPLICATIONS.

BY MICHAEL PARKER • ALTERA CORP

With the integration of dual ARM A9 CPU cores, a complete set of ARM peripherals, the ability to implement in hardware either fixed or floating-point signal processing, and unmatched I/O flexibility, the latest FPGA system-on-chip devices can perform what used to require a complete circuit card containing dozens of chips. A perfect example is next-generation motor control.

Next-generation motor-control systems are designed to deliver maximum motor efficiency using very fast control loops that exceed the capabilities of processor-only solutions. The inner control loops, implementing what is known as field-oriented control (FOC), require transforms best performed in floating point. FPGA SoCs that combine dual processor cores and FPGA fabric are ideal for this application, combining in a single device the capabilities of general-purpose processors and high-performance logic for specialized algorithms (see sidebar "How FPGAs and multicore CPUs are changing embedded design").

MOTOR COMMUTATION

Electric motors rely on basic electromagnetic principles. Forces due to magnetic attraction and repulsion are used to generate torque, as well as the resulting rotary motion of the motor.

Magnetic fields are generated in both the rotor and the stator, through the use of either permanent magnets or electromagnets. Proper alignment of the magnetic fields as the rotor turns is required.

The magnetic fields must continuously change to maintain an alignment that produces torque throughout the rotor's full 360 degrees of rotation. This process is called commutation. The most basic method used is the dc brush motor. This motor design has permanent magnets in the stator and one or more electromagnet pairs in the rotor. In order to eliminate the sparking and wear of brushes for mechanical commutation, the motor design can be turned inside out by placing the permanent magnets on the rotor, and the electromagnets in the stator. Replacing mechanical with electronic commutation can achieve the same effect, and is known as a brushless dc motor.

FIELD-ORIENTED CONTROL

In a dc brush motor, the control loops can simply drive the motor current variable, and the brushes perform the mechanical commutation function, although in a suboptimal fashion. With FOC, sinusoidal commutation will be performed electronically in an optimal fashion, using integrated control loops to maximize the magnetic field components producing useful torque and minimize magnetic field components that do not yield torque (perhaps, for example, merely exerting force on the motor bearings). The object of FOC is to ensure the magnetic fields are precisely orientated at all times to produce maximum torque, eliminate torque ripple, and thereby increase motor efficiency and lower the cost of ownership of the system.

The FOC function shown in **Figure 1** uses Park and Clarke transforms as well as PI control loops for torque (useful oriented magnetic field direction) and flux (magnetic field direction producing no torque). PI controllers are used rather than the PID controllers typically seen in control systems. (The derivative term is not used.)

Motors normally use three independent phases. These phases can be replicated around the stator, giving the number of poles or windings. In a three-phase system, the sum of the currents is defined as zero by Kirchhoff's law. This means that a three-phase current vector (a, b, and c components) can be expressed as two orthogonal phases (α, β components), using the Clarke transform. The Clarke transform is valid at a given rotation angle of the rotor.

Using the Park transform, these current vectors are mapped onto the rotating plane of the spinning motor. This results in the α, β components being mapped into q (quadrature) and d (direct) components. The transform requires the rotor angle, which is the input, often determined by a quadrature encoder attached to the rotor shaft. The

AT A GLANCE

- Given the rapid growth of new standards and protocols as well as increasing pressure to speed time to market, embedded-system design is due for a disruptive paradigm change.
- The latest FPGA system-on-chip devices can perform what used to require a complete circuit card containing dozens of chips.
- SoC FPGAs for real-time applications such as motor control provide both integration benefits and the ability to scale performance.

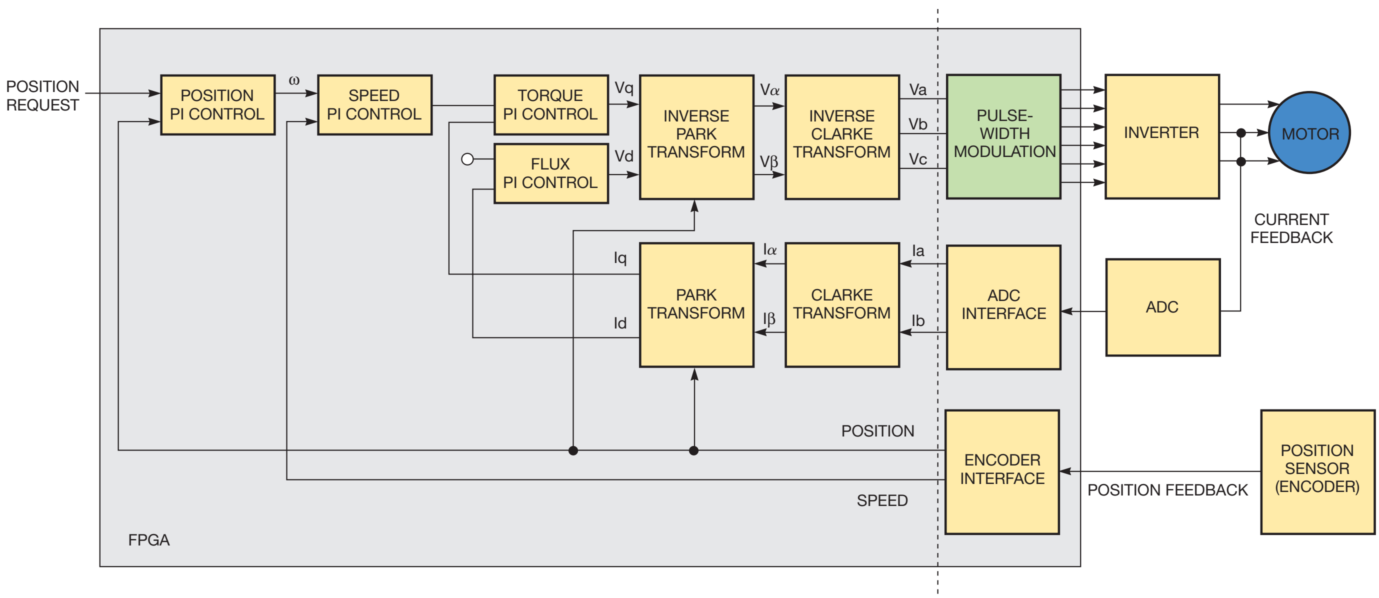

Figure 1 This FOC-based motor-control system uses Park and Clarke transforms as well as PI control loops for torque and flux.


The goal of EDS 2013 is to strengthen your business and the electronics industry as a whole by focusing on the future. We're looking at where we're going and how new technologies can help us evolve.
It's time to look at the future of the electronics industry
This year we're giving you even more ways to connect with the world's top manufacturers, distributors and sales representatives. In addition to the important meetings with your channel partners, you can relax, recharge and reconnect at the all-new Connections Café. And be sure to catch up with your industry colleagues at the always-exciting Mission: Possible networking reception. Join us for the "can't miss" electronics distribution event of the year.
Register today at EDSConnects.com or call 312-648-1140.
EDS
mission:vision
MAY 6-9, 2013
The Cosmopolitan Las Vegas
edsconnects.com
An industry service of
ecia
Electronic Components Industry Association

Clarke and Park transforms, therefore, need to be continuously calculated as the motor rotates.

REQUIREMENTS

While motor-control applications are mechanical, the rate at which the drive circuits must be updated and the current, position, and speed sensors read can be quite high. A reasonable scenario might be a motor operating at a maximum of 12,000 RPM, or 200 revolutions per second (RPS). If we use a rule of thumb that a minimum of 80 samples is needed to generate a well-shaped sinusoidal current waveform, then the required sample rate will be the equal to motor speed (RPS) × 80 × number of motor pole pairs.

The number of motor pole pairs is the number of electromagnetic windings in the stator. For a motor at 12,000 RPM with eight windings or pole pairs, this works out to a sampling rate of 128,000 per second, and a processing latency of 7.8 μsec.

Each application will be different, but often the system must be designed such that it can acknowledge and process all interrupts and update the motor drivers within 5 μsec or less. The process is to sample the feedback motor currents, position, and speed, and use FOC to calculate updated motor current driver values at a rate of 200,000 per second with a 5-μsec processing latency. For a processor-based system, this implies an interrupt rate of 200 kHz, which can be a significant challenge for general-purpose processors with caches, operating systems, and non-vectored interrupt controllers.

PROCESSING SYSTEM

Use of an SoC FPGA system can allow a more optimal and flexible implementation. The SoC contains two 800-MHz ARM Cortex A9 microprocessor systems.

The ARM A9 processor is a more general-purpose, high-performance processor, but it is not optimized for demanding real-time applications with guaranteed response times. Real-time performance limitations can be mitigated, however, by taking advantage of the integrated programmable logic of the SoC FPGAs. The programmable logic can be used for PWM drive circuits, versatile interfaces to any ADCs and DACs, position and speed sensor interfaces, safety cut-off circuits, a proprietary network or MAC hardware interface, and more. It can also implement the FOC and control loops with less latency and much faster response time than are possible using processor-only systems. Typical system response time is less than 2 μsec, which is about an order of magnitude faster than most processor-only systems can support.

FOC HARDWARE

A design tool such as DSP Builder from Altera can be used to take a Simulink representation of a design and implement it directly into FPGA logic. DSP Builder also generates floating-point logic, just like the simulation, to provide greater dynamic range and numerical stability than fixed-point implementations allow. The resources used to implement in the smallest SoC FPGA are shown in **Table 1**.

All interfaces to the FOC blocks are memory mapped by Altera's DSP Builder and Qsys tools into the ARM

TABLE 1 HARDWARE RESOURCES TO IMPLEMENT IN SoC FPGA (5CSEA2)

FPGA hardware resources and performance	Logic elements	Variable precision DSP blocks	M10K memory blocks	F_{MAX}/processing latency
Floating-point FOC, including Park, Clarke, and PI controllers	7.7K	2	1	120 MHz 1.74 µSEC
Available resources	25K	36	140	
5CSEA2 SoC FPGA (%)	31	5.5	0.7	

processor system. This allows for software control of the desired torque output, for the gains of the PI loops controlling torque and flux, for monitoring operations, and for other functions.

CONTROL-LOOP HARDWARE

The position and speed control loops may also be implemented in logic, similar to the PI control used in the FOC function. The gain stages of the different PID circuits can be updated using memory-mapped registers to the ARM processor. If needed, this can provide extremely low latency and fast response at the same rate as the commutation function. Sub-5-µsec total response of the control and FOC is easily achievable, which can provide more stability for very-fast-reacting systems. Just as in FOC, these control loops can be implemented in hardware from Simulink models, using DSP Builder to provide floating-point hardware in the FPGA logic.

Another consideration is that multiple axes of motor control can usually be added to the FOC hardware with a small increase in logic size and processing latency. A four-axis motor controller may have 20% higher logic use and minimal latency increase over a single-axis implementation. In a software-based implementation, the latency will scale in a roughly linear fashion.

CONTROL-LOOP SOFTWARE

Motor control loops are most commonly implemented in software using C code. Due to the effect of the FOC's isolating the control loops from the motor commutation, the interrupt rates and latencies of the motor control can be determined by the response of the system, which is normally much slower than the PWM commutation circuits. Simulation and analysis are required to determine the minimum update rates for expected performance and stability; however, an update rate of 10 kHz is normally adequate. In this case, the torque settings to the FOC would be updated at this lower rate (every 100 µsec), based upon feedback and calculation of the current position and speed information.

This SoC approach allows for the hardware to be used for the FOC "inner loop" and the ARM A9 for the motor control "outer loop." The inner loop can guarantee latency under 2 µsec in the FPGA hardware. The ARM can achieve very reasonable processing latency under most conditions and is ideal for outer-loop processing where interrupt latency requirements are much more relaxed, and the consequences of missing an interrupt are not catastrophic.

ARM A9 INTERRUPT LATENCY

The following is based on an analysis by the Altera SoC engineering team of real-time system performance of the ARM processors contained in low-cost SoC FPGAs, specifically interrupt latency. Interrupt response time can vary with system configuration, operating systems, cache configuration, and processing-task load.

The team considered the following scenarios:

- **Stage 1:** Highly constrained system on a single core to establish a benchmark for the lowest latency that can be achieved (To measure the maximum achievable performance, this stage did not use an OS, and the tests are implemented using bare-metal drivers.)
- **Stage 2:** Similar to stage 1, but the

Make an IMPACT

A whole new way of doing business is emerging – one where businesses of all sizes can make a difference and have an impact. Business4Better (B4B) Conference and Nonprofit Expo will bring together midsized business, large corporations, and nonprofits to truly define mutually beneficial partnerships and the impact they can have on societal causes.

Explore the benefits of business-integrated social responsibility programs and powerful partnerships. Join leading companies, innovative nonprofits, and individuals inspired to take action and transform business impact on social change.

Check out the lineup of speakers:

Neil Bush
Chairman
Points of Light

Premal Shah
President
Kiva

Chris Jarvis
Co-Founder & Partner
Realized Worth

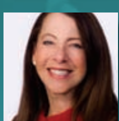

Carol Cone
Global Practice Chair
Edelman Business +
Social Purpose

Jeff Hoffman
President
Jeff Hoffman & Associates

Chad Bolick
Associate Director of
Partnerships
FSG

Dev Patnaik
CEO
Jump Associates & Author
"Wired to Care"

Bea Boccalandro
President
VeraWorks

Be a part of B4B. Be a Business4Better.

Register now at Business4Better.org
for Early Bird rates.

All registration proceeds are donated to B4B's nonprofit partners

Business4Better
The Community Partnership Movement

May 1-2, 2013 | Anaheim

PRODUCING PARTNERS
City of Anaheim
Freeman
Veraworks
Knowledge Marketing

PREMIER UNDERWRITER
Jive Software

NONPROFIT PARTNERS
B Lab
Points of Light Institute
OneOC
Stanford Social Innovation Review

tests are implemented using an RTOS (μC/OS-II) and the measurements are repeated over a number of iterations (1024) to get a statistical model of the interrupt latency

- **Stage 3**: Similar to stage 2, with multiple background tasks running on the RTOS alongside interrupt latency measurements (This background task writes data onto the UART port in an infinite loop, and in addition, another background task is performing memcpy in the external memory in an infinite loop, which generates AXI read/writes that may require a long time to complete.)
- **Stage 4:** Emulation of a larger system where not all code can fit within the L1 cache (In this case, the critical code is placed in the L2 cache. Since the cache is likely to be large enough to hold the whole of the background processing code, after each interrupt, L1 cache flush is triggered. The objective here is to measure the interrupt latency achievable when using the SoC FPGA with a real-time operating system in a real-world environment where cache flushes play a significant role in system performance.)

Cache (L1 & L2) has the most influence on the interrupt latency (**Table 2**). In a stage 4 situation, with the instruction code locked in L2, the latency is significantly lower than the same system running from external memory. With L1 enabled, it is lower still. The latency will be much worse in the stage 3 and 4 cases, however, due to the background tasks that cause the ISR code to be replaced with other background tasks.

The interrupt response latency needs to be added to the processing time of the FOC (**Table 3**). A basic FOC algorithm was benchmarked, including trigonometric functions, Clarke/Park transforms, PID controller, and inverse Park/Clarke transforms.

Due to the processor performance of the ARM A9, a software-based FOC

TABLE 2 MEASURED INTERRUPT RESPONSE LATENCY (μSEC)

Cache configuration	Stage 1	Stage 2	Stage 3	Stage 4
Cache disabled	2.53	2.65	2.87	2.89
L2 enabled	0.57	0.54	3.53	1.5
L1 enabled	0.16	0.19	1.29	1.29
L1 and L2 enabled	0.16	0.19	1.28	1.3

TABLE 3 MEASURED FOC LATENCY (μSEC)

Cache configuration	Interrupt response (stage 4)	FOC processing	Total latency
L2 enabled	1.5	3.27	4.77
L1 enabled	1.29	0.83	2.12
L1 and L2 enabled	1.3	0.73	2.03

HOW FPGAs AND MULTICORE CPUs ARE CHANGING EMBEDDED DESIGN

By Sanjay Challa, National Instruments

With the explosion of embedded devices in the past few decades, many improvements have been made in both the hardware components and software tools. Despite this innovation and growth, however, traditional embedded-system design approaches have evolved little if at all and are increasingly proving to be a hurdle. Given the increasingly rapid growth of new standards and protocols as well as increasing pressure on design teams to deliver to market more quickly, embedded-system design is due for a disruptive paradigm change.

With the accelerating growth of advances in hardware technologies and software tools, the challenge posed by integration is set to rise. This challenge, if unaddressed, will result in more expensive end products and can prevent experimentation, growth, and delivery of more innovative designs to the marketplace.

STANDARD EMBEDDED ARCHITECTURE

In the general computing marketplace, standardization has resulted in more robust operating systems, more refined end applications, and advances in the underlying hardware components. The lesson learned is that time saved in avoiding the integration effort of custom hardware architectures and associated software components results in better end solutions, which are delivered to market faster.

For the embedded space, a corresponding standard architecture needs to be flexible enough to adapt to diverse use cases while providing an avenue for updates. Given these constraints, the most robust architecture for standardization in the embedded design space is a microprocessor and an FPGA working alongside each other as a single unit (Figure A). Together, these two elements enable substantial flexibility in designs.

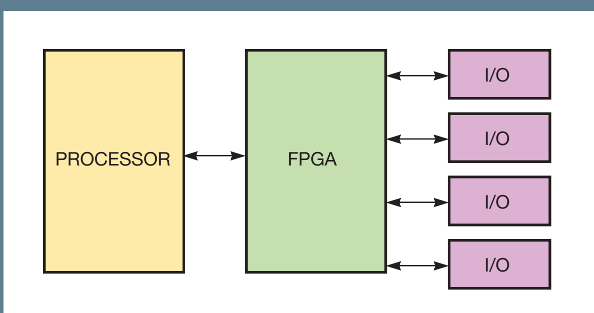

Figure A In this standard hardware architecture, the combination of a processor and an FPGA enables flexibility while making it possible for standardization that can utilize higher-level tools to make substantial gains in the design workflow. The processor makes it possible to reuse existing code libraries, while the FPGA allows for the flexible implementation of custom algorithms.

FPGAs offer the benefits of hardware determinism and reliability without the up-front cost and rigidity of ASIC design. Additionally, the ability to load new logic and redefine the connections in the FPGA fabric makes it possible for designers to future-proof designs and benefit from more robust updates without requiring any substantial modifications to hardware.

The combination of processors and FPGAs in embedded-system design is growing in many industries. Embedded-systems developers are using designs based on several processors and FPGAs. The FPGAs are used to take accurate, high-speed measurements or run time-critical algorithms. Meanwhile, the processors run a real-time operating system to handle lower-frequency control loops or provide Ethernet communication to other distributed nodes and facilitate remote data access, system management, and diagnostics.

HIGHER-LEVEL TOOLS

A key benefit of a standard architecture is that more capable and optimized high-level tools can be developed and used for design. Higher-level tools make it possible for domain experts to be more closely involved in embedded-system design with smaller and more efficient design teams. As a result, more complex products can be pushed to market sooner with smaller design teams.

General-purpose computing provides evidence for the efficiencies that can be gained in application development with higher-level design tools and languages. Unsurprisingly, the embedded marketplace has started to witness the growth of higher-level design tools, including the Xilinx AutoESL C-to-Gates high-level synthesis tool, Mentor Graphics Catapult C Synthesis tool, and NI LabVIEW ultimate system-design software.

AUTHOR'S BIOGRAPHY

Sanjay Challa is a product manager for embedded software at National Instruments, with a focus on real-time operating systems and FPGA-based embedded systems. He joined the company in 2010. Challa received his bachelor's degree in biomedical engineering from the Georgia Institute of Technology (Atlanta).

JOIN THE CONVERSATION

To comment on this article, go to **www.edn.com/4411150**.

controller has been shown to be able to meet the 5-µsec requirement. There is little margin, however, and interrupt response times are often the most nondeterministic part of the system. While the results shown are the longest latencies seen under these test conditions, much longer times may occur on occasion, due to the statistical nature of the processing-task load. However, in the case where interrupts are occurring at rates on the order of 10 to 20 kHz, or 50 to 100 µsec, it appears that the ARM Cortex A9 can meet the real-time requirements with extremely high margins, and very low probability of not servicing interrupts in a timely manner. By isolating the ARM interrupts from the much higher-speed FOC commutation requirements, real-time performance is virtually assured under a wide variety of operating conditions and processing loads.

The use of SoC FPGAs for real-time applications such as motor control provides not only integration benefits but also the ability to scale performance as needed. This approach allows high-rate, deterministic functions (inner loop) to be implemented in hardware, while lower-rate, more dynamic and complex processing (outer loop) can take place in software, providing the best of both worlds to the system designer. **EDN**

ACKNOWLEDGMENT

A version of this article originally appeared on EDN sister site Embedded.com at www.embedded.com/4402474 and www.embedded.com/4395443.

AUTHOR'S BIOGRAPHY

Michael Parker is principal architect for DSP product planning at Altera Corp. He joined the company in January 2007 and has more than 20 years of DSP wireless engineering design experience with Alvarion, Soma Networks, TCSI, Stanford Telecom, and numerous start-up companies. Parker holds a master's degree in electrical engineering from Santa Clara University (Santa Clara, CA) and a bachelor's degree in electrical engineering from Rensselaer Polytechnic Institute (Troy, NY).


Economy oscilloscopes redefined.
InfiniiVision 2000 X-Series
70 to 200 MHz
New memory depth up to 1 Mpts
New serial decode options*
New 5-year warranty
Starting at $1,258**
InfiniiVision 4000 X-Series
200 MHz to 1.5 GHz
12.1-inch touch screen
InfiniiScan Zone touch triggering
Update rate to 1 million wfms/s
Memory depth up to 4 Mpts
Starting at $5,611**
InfiniiVision 3000 X-Series
100 MHz to 1 GHz
8.5-inch display
Update rate to 1 million wfms/s
Memory depth up to 4 Mpts
Starting at $2,929**
Oscilloscopes Redefined. With its brilliant 8.5-inch screen, fast update rate, full upgradability and 5-in-1 instrument integration, the Agilent 2000 X-Series oscilloscope is already first in class. Now with more memory depth, new serial decode options and 5-year warranty, it's accelerating ahead. Experience the complete InfiniiVision line for oscilloscope features and performance you need at prices that will surprise you.
See a video demo.
Scan or visit
http://qrs.ly/992zypm
Agilent and our Distributor Network
Right Instrument.
Right Expertise.
Delivered Right Now.
TESTEQUITY
800-732-3457
www.testequity.com/agilent
See the InfiniiVision X-Series oscilloscopes
www.testequity.com/Agilent_Scope
© Agilent Technologies, Inc. 2013
*I2C, SPI, CAN, LIN, RS232, UART **Prices are in USD and are subject to change
Anticipate Accelerate Achieve
Agilent Technologies

BY JACOB BENINGO • BENINGO ENGINEERING

10 C language tips for hardware engineers

IT CAN BE COMMON FOR A HARDWARE DESIGNER TO WRITE CODE TO TEST THAT HARDWARE IS WORKING. THESE 10 TIPS FOR C—STILL THE LANGUAGE OF CHOICE—MAY HELP THE DESIGNER AVOID BASIC MISTAKES THAT CAN LEAD TO BUGS AND MAINTENANCE NIGHTMARES.

On its own, the software-development process has numerous hazards and obstacles that require navigation in order to successfully launch a product. The last thing that any engineer wants is challenges resulting from the language or tool that is being used. It often makes sense or is necessary for the hardware designer to write code to test that the hardware is working or, in resource-constrained cases, develop both hardware and embedded software. The language of choice is still C, and despite the advances in tools and structured programming, time and again basic mistakes occur that lead to bugs and maintenance nightmares. In an attempt to avoid these C-programming pitfalls, here are 10 C language tips for hardware engineers.

TIP #1: DON'T USE "GOTO" STATEMENTS

A couple decades or so ago, when computer programming was in its infancy, a program's flow was controlled by "goto" statements. These statements allowed a programmer to break the current line of code and literally go to a different section of code. **Listing 1** shows a simple example.

Programming languages eventually began to incorporate the idea of a function, which allows the program to break off to a section of code. Rather than requiring another goto statement, however, when completed the function returns to the next instruction after the function call. **Listing 2** shows an example. The result was improved program structure and readability, and ever since, this has been considered the appropriate way to write a program. The very sight or thought of goto statements can cause software engineers to cringe and shudder in distaste. One of the main reasons is that a program with goto's sprinkled throughout is very difficult to follow, understand, and maintain.

LISTING 1 USE OF A GOTO STATEMENT

```
void main(void)
{
   int Count = 0;

   // Do Something

   // Wait for a while
   WAIT:
   if (Count < 5000)
   {
      goto WAIT;
   }

   // Do more stuff
}
```

TIP #2: USE FOR(;;) OR WHILE(1)

If goto's are out, some hardware engineers may wonder how an infinite loop can be created for the program. After all, this may have been done before by creating a goto statement that returns to the top of main. The answer is to take advantage of the looping statements that are already built into the C language: for and while (**listings 3** and **4**).

The loop conditionals in the **listings** are relatively straightforward. The for loop is nothing more than the for conditional with no conditions. The while loop, on the other hand, will execute as long as the statement is true, which is the same as having any nonzero value for the condition.

TIP #3: USE THE APPROPRIATE CONDITIONAL STATEMENT FOR THE JOB

Program execution time can be highly dependent on the type of conditional structure that is selected for making a decision, in addition to the readability of the code. Many hardware engineers are familiar with the use of the simple if statement. Sometimes, however, the engineer doesn't realize that if the first condition isn't correct, an else or else if statement can be

LISTING 2 USE OF A FUNCTION TO CONTROL FLOW

```
void main(void)
{
    // Do Something

   // Wait for while
   Delay (50);

   // Do more stuff
}
```

LISTING 3 USE OF AN INFINITE FOR LOOP

```
// Using a for loop to create the
   infinite program loop
void main(void)
{
   // Initialize the system
   Sys_Init();

   // Run the program forever
   for(;;)
   {
      // Run the cool application
   }
}
```

LISTING 4 USE OF AN INFINITE WHILE LOOP

```
// Using a while loop to create the
   infinite program loop
void main(void)
{
   // Initialize the system
   Sys_Init();

   // Run the program forever
   while(1)
   {
      // Run the cool application
   }
}
```

used. This can save the processor time by not having to evaluate another conditional statement. In the before code in the example shown in **Listing 5**, if Var is equal to one, the code will still check to see if the Var is equal to zero. However, in the after code that uses the else, only the first statement is evaluated, and then the code moves on, thereby saving clock cycles and making the code more clear.

The if/else if/else statements still may not always be appropriate. If there are a number of possible conditions that need to be checked, a switch statement may be more appropriate. This allows the processor to evaluate the statement and then select from a list of answers what it should do next rather than continually evaluate a bunch of conditions. **Listing 6** shows an example that corresponds to the same type of example shown in **Listing 5**.

The moral of the story is simply to keep alternative conditional-statement options open and select the most appropriate for the job. This approach will make it easier to understand the flow of the program by making the structure straightforward and could squeeze extra clock cycles out of the processor.

LISTING 5 USE OF IF/ELSE INSTEAD OF JUST IF

BEFORE

```
if(Var == 1U)
{
   // Do something neat
}

if(Var == 0U)
{
   // Do something cooler
}
```

AFTER

```
if(Var == 1U)
{
   // Do something neat
}
else
{
   // Do something cooler
}
```

TIP #4: AVOID USING ASSEMBLY LANGUAGE

The natural language for a microprocessor is assembly language instructions. Writing a program in the low-level machine language can result in more efficient code for the processor. Humans, however, don't naturally speak this language, and as experience has shown, writing assembly language results in misunderstanding. Misunderstanding then leads to improper maintenance, or worse, and the result is a system overridden with bugs. A general tip is to avoid the use of assembly language. The detailed truth of the matter is that most compilers now compile very efficient code. Developing in the higher languages, such as C/C++, results in a more organized structure, which is easier to understand and maintain, and produces overall better code. **Listing 7** shows an example that compares the use of assembly and C code to increment a 32-bit variable.

With that said, there are still occasions when it is appropriate to use assembly language, but those times are scarce. The first recommended time is when developing a boot-loader. In this instance, during start-up it may be necessary to optimize how quickly a decision is made to boot the application or the boot-loader. In this case, assembly code for the branch decision may make sense. Another case is when developing a control loop that has tight timing requirements on a DSP. In order to squeeze every clock cycle out of the device, it may make sense to code the control loop in assembly. If the task at hand is appropriate for using assembly, make sure that it is well documented so that future developers (or future versions of yourself) can understand what the code is doing.

TIP #5: TAKE ADVANTAGE OF MODULARITY

One of the most common experiences the author has when taking on a new project that was started by hardware engineers is the atrocious organization of the code. It isn't uncommon to find that code consists of a single main module with in excess of 25,000 lines. In these applications, everything is global, functions are sparse, and goto statements rule the organization of the code. This was the norm 15 years ago, but not anymore! Programming in C/ has given engineers the ability

LISTING 6 USE OF SWITCH STATEMENTS

```
switch(Var)
{
   case 1:

      // Do something neat

   break;

   case 0:

      // Do something cooler

   break;

   default:

     // Should never get here
     // Error or something

   break;
}
```

to break up their code into separate functional modules. This eases navigation of the code but also allows an engineer to use object-oriented techniques such as encapsulation. Where it makes sense, organize code into logical modules. It'll take a little bit more time up front (a few minutes), but in the long run it will save many long nights and debugging headaches!

TIP #6: WRITE LASAGNA NOT SPAGHETTI CODE

Beningo is an Italian name, and as with any good Italian, my love of pasta is a given. When comparing pasta with software, spaghetti and lasagna come to my mind. Spaghetti is chaotic; noodles intertwine, going this way and that way and resulting in a complete lack of any type of structure. Writing unstructured code is exactly like spaghetti: With each bite, you have no clue what you are going to get!

On the other hand, there is lasagna! The noodles are layered, giving the meal structure. Code developed using layers not only is easier to understand, it has the potential to have one layer removed and a new layer added, basically allowing for reuse and ease in maintainability. **Figure 1** shows an example of a simple software model that uses the lasagna model.

TIP #7: USE DESCRIPTIVE VARIABLE NAMES

One of the barriers to writing great software that is understandable and easy to maintain is the naming convention of variables. In an effort to keep variable names short, it is common for developers

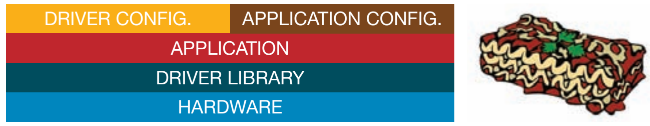

Figure 1 Code developed using layers is easier to understand, and it has the potential to have one layer removed and a new layer added, allowing for reuse and ease in maintainability.

to create shortened, cryptic mnemonics that only they can understand once in a blue moon. Modern languages allow for hundreds of characters to be included in a variable name. To keep things clear, "call a spade a spade," as the phrase goes, rather than something else. This will make the variable name obvious not only to the developer but also to future maintenance teams. **Listing 8** shows an example.

TIP #8: USE #PRAGMA STATEMENTS SPARINGLY

In the C language, there is a special type of statement known as #pragma. These statements often handle nonstandard syntax and features. They should be avoided as much as possible because they are nonstandard and will not port from one processor to the next. Some compilers may require them for tasks such as defining an interrupt service routine. In these instances, there may be no way around using a #pragma. If possible, keep all of the #pragma statements together in one module or in a couple of modules. This will help to ensure that when the code is ported, there are only a few places to update the code rather than areas sprinkled throughout the code base. This will also help prevent a nightmare when the ported code is compiled for the first time.

TIP #9: ERRORS AREN'T ALWAYS AS THEY SEEM

One of the gotchas to look out for when debugging a C program is compiler errors. Depending on the sophistication of the compiler, when an error is detected, more often than not the error lies somewhere in the program other than where the compiler is indicating. The reason for this has to do with the steps that the compiler takes to generate the program. The types of errors are generally consistent, so there are a few errors an engineer can look for that nine times out of 10 are the culprit:

LISTING 7 INCREMENTING A VARIABLE IN ASSEMBLY VERSUS C

```
INC  [_Sys_Tick+3]      ;Increment Byte LSB          Sys_Tick++;
JNZ  .GoReturn          ;Jump is no overflow

INC  [_Sys_Tick+2]      ;Increment Byte 1
JNZ  .GoReturn          ;Jump is no overflow

INC  [_Sys_Tick+1]      ;Increment Byte 2
JNZ  .GoReturn          ;Jump is no overflow

INC  [_Sys_Tick]        ;Increment Byte MSB
```

ASSEMBLY | C CODE

DESIGN NOTES

High Voltage CMOS Amplifier Enables High Impedance Sensing with a Single IC – Design Note 513

Jon Munson

Introduction

Accurately measuring voltages requires minimizing the impact of the instrument connection to the tested circuit. Typical digital voltmeters (DVMs) use 10M resistor networks to keep loading effects to an inconspicuous level, but even this can introduce significant error, particularly in higher voltage circuits that involve high resistances.

The solution is to use high impedance amplifiers in an electrometer configuration, so only miniscule amplifier input current comes from the test node. To make the input current the lowest possible value, field-effect transistors (FETs) are traditionally employed at the inputs of these circuits. FETs are generally low voltage devices, and introduce voltage offset uncertainty that is difficult to eliminate. Monolithic amplifiers exist that incorporate FET inputs, but these are often very low voltage devices, especially those using typical CMOS fabrication methods, so their utility is limited in high voltage applications. Enter the LTC®6090, a CMOS amplifier that can handle over $140V_{P-P}$ signal swings with sub-mV precision, ideally suited to tackle the problem.

The LTC6090 Easily Solves High Voltage Sensing Problems

The LTC6090 combines a unique set of characteristics in a single device. Its CMOS design characteristics provide the ultimate in high input impedance and "rail-to-rail" output swing, but unlike typical CMOS parts that might be powered by 5V, the LTC6090 can operate with supplies up to ±70V. The device can hold its own in the small-signal regime as well, featuring typical V_{OS} under 500µV and voltage-noise density of $11nV/\sqrt{Hz}$, yielding a spectacular dynamic range. With the high voltage operation comes the possibility of significant power dissipation, so the LTC6090 is available in thermally-enhanced SOIC or TSSOP packages. It includes an overtemperature output flag and an output disable control that provide flexible protection measures without additional circuitry.

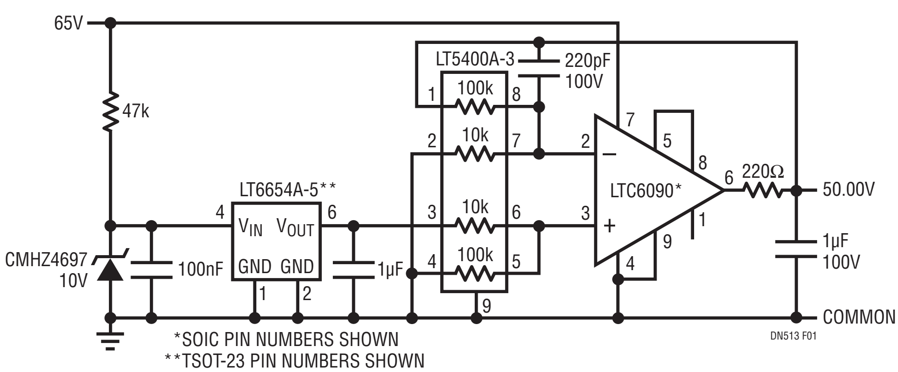

Figure 1. High Voltage Precision Reference

Accurate 50.00V Reference

The LTC6090 is capable of 140V output levels in single-supply operation, so amplifying a quality 5V reference is a simple matter of using accurate resistor networks to maintain precision. The LT®5400 precision resistor array handles voltages up to 80V, so utilizing the 10:1 ratio version for a gain of 10 is an easy way to produce an accurate 50V calibration source without any adjustments required. Figure 1 shows a circuit amplifying the LT6654A 5.000V reference to 50.00V with better than 0.1% accuracy. The circuit may be powered from 55V to 140V, with 65V being a useful supply voltage furnished from the optional portable supply shown as part of Figure 2.

The LTC6090 is set up with a 1µF output capacitance to provide excellent load-step response. The capacitance is isolated from the op amp with a resistance that forms an effective noise-reduction filter for frequencies above 700Hz. The precision LT5400A-3 resistor network provides 0.01% matched 10k/100k resistances that, along with the absence of loading by the high impedance CMOS op amp inputs, forms a highly accurate amplification factor. Input offset voltage of the LTC6090 contributes <0.03% error, while the LT6654A contributes <0.05%.

The entire circuit of Figure 1 draws about 4mA of quiescent current and can drive 10mA loads.

Simple Large-Signal Buffer

The LTC6090 behaves as an ordinary unity-gain-stable operational amplifier, so constructing an electrometer-grade buffer stage is simply a matter of providing 100% feedback with the classic unity-gain circuit. No discrete FETs or floating biasing supplies are needed.

As shown in Figure 2, the LTC6090 can easily be powered with a split supply, such as a small flyback converting battery source. This basic circuit can provide precision measurement of voltages in high impedance circuitry and accurately handle signal swings to within 3V of either supply rail (±62V in this case). With input leakage current typically below 5pA, circuit loading is essentially inconsequential (<V_{OS}) for source impedances approaching a gigaohm. The useful full-swing frequency response is over 20kHz.

Conclusion

The LTC6090 is a unique and versatile high voltage CMOS amplifier that enables simplified high impedance and/or large signal swing, very wide dynamic range amplification solutions.

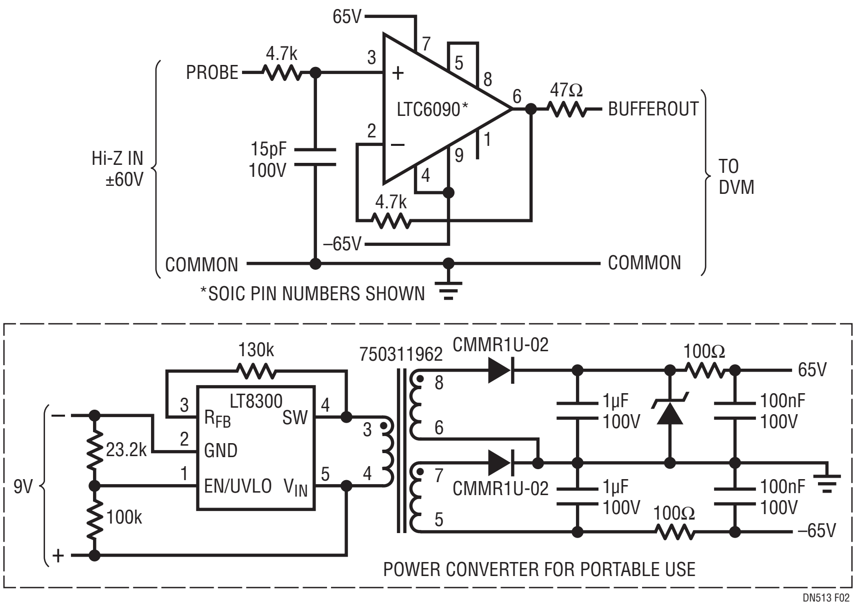

Figure 2. Buffered Probe for Digital Voltmeter

Data Sheet Download

www.linear.com/6090

For applications help,
call (408) 432-1900, Ext. 3481

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

dn513f LT/AP 0413 196K • PRINTED IN THE USA

LISTING 8 VARIABLE NAMING

```
int Frq;          int Frequency;
int Btn;          int Button;
int MtrState;     int MotorState;
int Spd;          int Speed;
```

• Watch for missing #include files. This can result in the developer looking at a perfectly good line of code, but because the necessary header files aren't included, the compiler flags it as an error, indicating that something is not defined.

• Watch for missing semicolons. One of the most common errors when writing C code is to forget the semicolon at the end of a statement.

• Watch for missing brackets. Brackets are another common mistake and are either left out by accident or because mistyping produces a different character.

• Watch for missing commas. In complex definitions, it's easy to forget the comma!

In general, when a strange compiler error pops up, look around at what might have been compiled immediately before that line. Odds are that is where the mistake is! It may be one line up, half a page away, or in a completely different file.

```
\
int
_,l;\
char*I,
*O[]={"",
"gjstu","t"
"fdpoe","uij"
"se","gpvsui",\
"gjgui","t"
"jyui","tfwf"
"oui","fjhiui",
"ojoui","ufoui",\
"fmfwfoui","uxfmgu"
"i","b!qbsusjehf!jo!"
"b!qfbs!usff/\xb\xb",""
"uxp!uvsumf!epwf"
"t-\xb","uisff!gsf"
"odi!ifot-!","gpvs!d"
"bmmjoh!cjset-!","gjwf"
"!hpme!sjoht<\xb","tjy!h"
"fftf!b.mbzjoh-!","tfwfo!t"
"xbot!b.txjnnjoh-\xb","fjhiu"
"!nbjet!b.njmljoh-!","ojof!mbe"
"jft!ebodjoh-!","ufo!m"
"pset!b.mfbqjoh-\xb","fm"
"fwfo!qjqfst!qjqjoh-!","ux"
"fmwf!esvnnfst!esvnnjoh-!",""
"Po!uif!","!ebz!pg!Disjtunbt!n"
"z!usvf!mpwf!hbwf!up!nf\xb","boe"
"!"};int putchar(int);int main(void
){while(l<(sizeof O/sizeof*O-2)/2-1){
I=O[_=!_?sizeof O/sizeof*O-
3:_<(sizeof(O)/sizeof*O-2)/2?
sizeof O/sizeof*O-2:_==(sizeof(
O)/sizeof*O-2)/2?++l,0:_<(sizeof(
O)/sizeof(*O))-3?(_-1)==(sizeof(O)/
sizeof*O-2)/2?sizeof O/sizeof*O-1:_-1
:_<sizeof(O)/sizeof*O-2?l+1:_<sizeof(O)
/sizeof*O-1?l+(sizeof O/sizeof(*O)-2)/2:(
sizeof(O)/sizeof*O-2)/2];while(*I){putchar(
*I++-1);}}
return 0;}
```

Figure 2 A good programmer writes clean code that is easy to understand and maintain, not the fewest lines of code!

Readers weigh in

▶ **"Funny ... in 40 years I've only needed to code GOTO a dozen or so times, and each time the nasty mess of added state, loops, and breaks to get rid of them was just not worth the coding complexity. This is kinda like the function pointer discussion ... don't use them unless you really need to, and even then, code a couple alternative ways without them to see if there are any other constructs that both work and are clear and easy to maintain."—John.Bass**

▶ **"You could add to #9 to beware of EXTRA semicolons. I once handed off a project to another engineer, and it took him two days to find my:**

```
if(whatever = TRUE);
{
stuff;
}
```

I've watched out for that ever since then."
—DavidPK

JOIN THE CONVERSATION

To share your own comments on this article, go to **www.edn.com/4408338**.

Don't give up! With some experience, finding the difficult ones eventually becomes second nature.

TIP #10: GOOD PROGRAMMERS DON'T NECESSARILY WRITE FEWER LINES OF CODE

It is a common misconception that a good programmer can write fewer lines of code than an average programmer to do something. Don't get sucked into this idea! A good programmer has a well-thought-out and -structured code base. Variables are nicely named and encapsulated, with few global variables existing in the system. Functions should be short and concise. If the code looks confusing and it would be more clear to write more lines of code, then do so! Check out the online awards for writing the most confusing C code. A good programmer writes clean code that is easy to understand and maintain, not the fewest lines of code (**Figure 2**)! **EDN**

AUTHOR'S BIOGRAPHY

Jacob Beningo is a Certified Software Development Professional (CSDP) who specializes in the development and design of quality, robust embedded systems. He has written technical papers on embedded design methods and taught courses on programmable devices, boot-loaders, and software methods. Beningo holds bachelor's degrees in engineering and physics from Central Michigan University (Mount Pleasant, MI) and a master's degree in space-systems engineering from the University of Michigan (Ann Arbor, MI).

designideas

READERS SOLVE DESIGN PROBLEMS

Recover the leakage energy of a flyback transformer

Todor Arsenov, Spellman High Voltage Electronics Corp, Hauppauge, NY

The classical technique for demagnetizing the transformer in any forward converter is to implement a second winding bifilar with the primary winding to ensure continuous flow of the magnetizing current when the power switch (typically a power FET) turns off. Such a circuit generally limits or clamps the FET's drain-to-source voltage to two times the dc supply-rail voltage. The same technique—using this recuperating winding—can be successfully implemented in a flyback topology to deal with the leakage-inductance problem.

Note that in any flyback converter, the flyback transformer (multiwinding inductor) is far from perfect; leakage inductance (primary to secondary) is as much as 5% of the magnetizing primary inductance[1]. The leakage inductance (L_{LK}) is effectively in series with the power FET (drain connection). Complicating matters, the parasitic output capacitance of the FET (C_{OSS}) forms a series-resonant circuit with L_{LK}. When the FET turns off, very high overvoltage and ringing can occur. The higher the Q of the circuit, the higher the ringing voltage. This situation will likely cause significant EM interference and, due to the elevated FET drain voltage, lower the FET reliability.

Figure 1a shows a flyback converter with such a recuperating winding added on a modified demonstration board

DIs Inside

59 Double the protection of a laser driver using a 1V power supply

60 Gate-drive transformer eases multi-output, isolated dc/dc-converter designs

62 Two ICs form F/V converter

▸To see and comment on all of *EDN*'s Design Ideas, visit www.edn.com/designideas.

(STMicroelectronics' Viper17L[2]). Some important considerations: Resistors R_{S1} and R_{S2} are sense resistors used for monitoring the currents; scope measurements for currents are measured directly across these resistors. The transformer ratios are the same as in the original transformer. The recuperating winding, N_R,

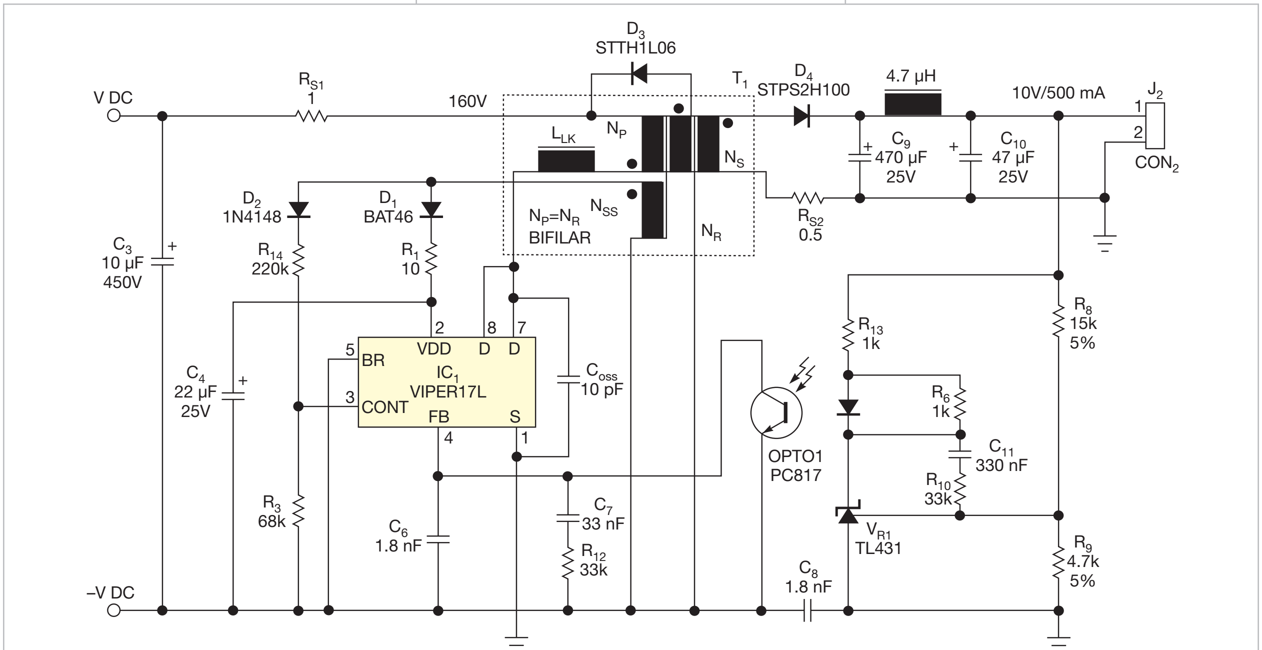

Figure 1a A passive snubber is realized by adding a bifilar winding, N_R, and a diode, D_3, to a flyback transformer.

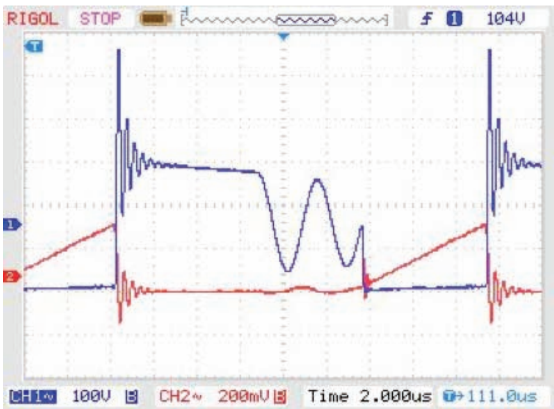

Figure 1b Without any snubbers, the ringing on voltage (Channel 1, blue) and current (Channel 2, red) waveforms can be quite big.

is magnetically coupled tightly to only the primary winding, N_P, by making these two windings bifilar. Bifilar windings are made by simultaneously winding two wires side by side around the magnetic core, or bobbin; this approach maximizes coupling and tightly matches the parasitic capacitance and inductance. The coupling between the primary and the other windings is not as important.

In **Figure 1b**, it can be seen that without any clamping (D_3 disconnected), the voltage at the FET's drain (IC_1, pins 7 and 8) due to the ringing reaches 560V peak. The primary current is shown magnified in **Figure 2a**. At the moment the FET turns off, the primary current (magnetizing current) remains constant, charging the capacitance, C_{OSS}. This is indicated by the step waveform. The magnetizing current remains constant, as diode D_4 on the secondary side is still not conducting; this can be seen from the secondary current waveform in **Figure 2b**. The short period of time after the turn-off (when D_4 is not yet conducting) is the time when the series-resonance circuit's C_{OSS} is charged. Corresponding to the time when the FET's drain voltage, V_{DS}, becomes high

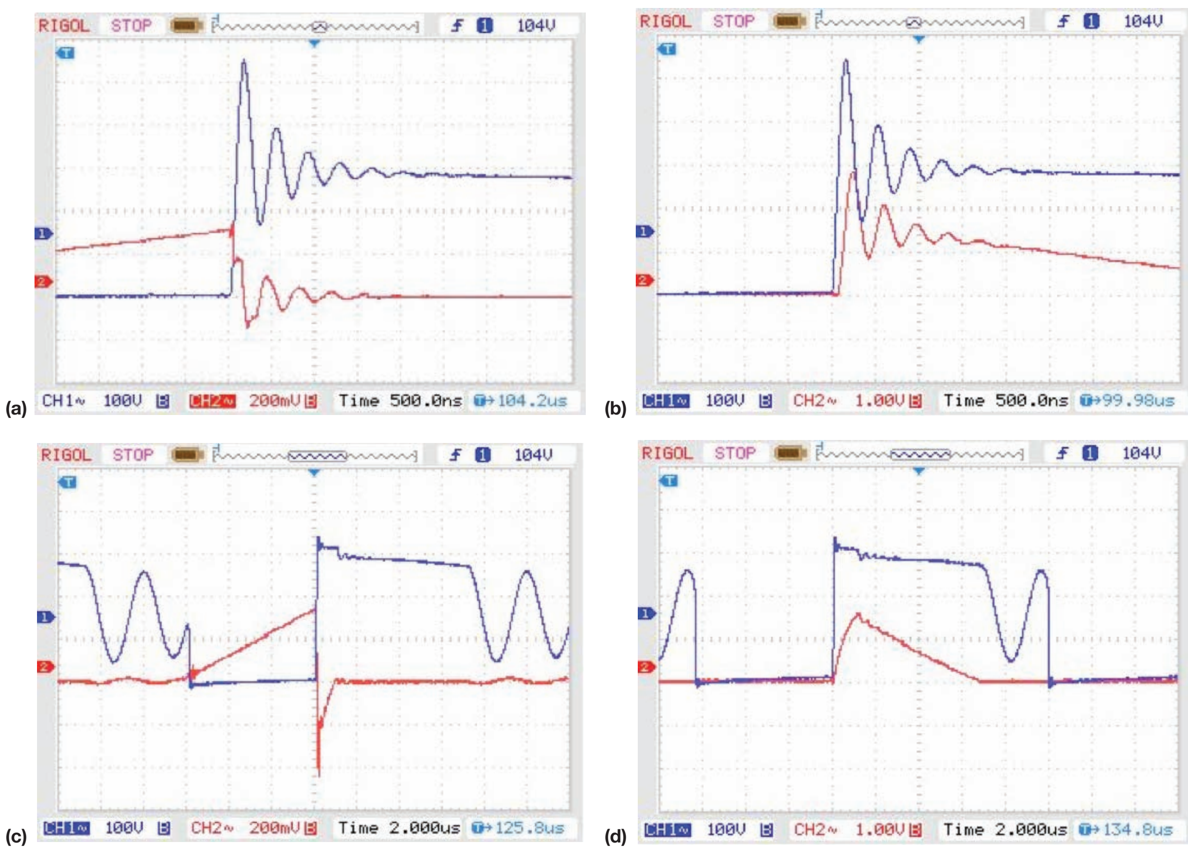

Figure 2 In these results from the series tank circuit, there is ringing at the FET's drain (V_{DS}, Channel 1, blue) after the primary current (Channel 2, red) charges C_{OSS} following turn-off (a). V_{DS} is again shown as Channel 1; the charging of C_{OSS} delays the secondary current through D_4 (Channel 2) by just under 100 nsec (b). The bifilar winding, N_R, steers the primary current (Channel 2) back to the power rail and clamps the switch voltage (Channel 1) (c). The leakage flux fights the current transfer; the secondary current (Channel 2) rises to an equilibrium peak value until the leakage energy is fully recuperated (d).

enough, D_4 becomes forward biased and the energy stored in the series-resonant tank circuit is released. The energy stored is a function of the resonant circuit's Q factor and is surprisingly high.

When the recuperating winding, N_R, and diode D_3 are connected to the power-supply rail, a completely different process is observed. The recuperating winding simply bypasses the parasitic C_{OSS}, steering the accumulated leakage energy back to the supply rail. In **Figure 2c**, it should be noted that the negative surge of the primary current (Channel 2) is actually the current flowing from the recuperating winding. The secondary diode, D_4, is forward biased immediately (**Figure 2d**); as the secondary current (Channel 2) rises to the steady-state peak value, the primary current diminishes to zero. As $N_P=N_R$, this assures limiting of the V_{DSS} to two times V_{DC}.

Eliminating the excessive ringing by leakage-energy recuperation is clearly an advantage as all other passive RCD snubber techniques dissipate this energy and thus lower the efficiency of the converter. Limiting the maximum V_{DS} to two times V_{DC} is acceptable bearing in mind that most monolithic embedded converters incorporate high-voltage power FETs. (STMicroelectronics' Viper17, for example, has an impressive 800V rugged power section[3].) Bifilar windings are readily available from most of the transformer vendors, or, for in-house production, a Multifilar magnet wire can be used to make these windings. **EDN**

REFERENCES

1 Dixon, Lloyd H, *Magnetic design handbook*, Texas Instruments, 2001.

2 "EVALVIPER17L-7W demonstration board," STMicroelectronics, June 2008, http://bit.ly/145jGHF.

3 "VIPER17 data sheet," STMicroelectronics, June 2010, http://bit.ly/Z8xOqM.

Double the protection of a laser driver using a 1V power supply

Tai-Shan Liao, National Applied Research Laboratories, Instrument Technology Research Center, Hsinchu, Taiwan

An excessive level of light from a laser pointer, even if only for a short duration, can be harmful if it enters the human eye either directly or through reflection from a shiny object. Most countries, therefore, have laser safety requirements that limit the maximum emission level. This Design Idea describes a laser driver that works even with a single 1.5V cell discharged to 1V, and uses dual current-control transistors to improve reliability against shorting and allowing excessive laser current and light emission.

In **Figure 1**, the transistors Q_1, Q_2, and Q_3 compose a negative impedance, which can be described approximately as $Z \approx -\beta(V_{DD}-V_{BE})/R_F$. Assume that all of the transistors have the same current gain (β), and V_{BE} is the base-to-emitter voltage of all transistors. Feedback is provided through R_F, and R_1 bias controls the collector current of Q_1. Inductor L_1 and parasitic capacitance form a resonant circuit that oscillates due to the negative impedance, resulting in about 3.5V pk-pk at Q_1's collector, with the battery at 1V. Schottky diode D_1 and C_1 form a half-wave rectifier that provides about −3V for the laser cathode; with V_{DD} at 1V, this provides a 4V working range to overcome the laser threshold.

Q_5 and Q_6 control the laser current. The photodiode built into the laser assembly monitors the light intensity and sends negative feedback through Q_4 to bias Q_5 and Q_6 to the proper collector current for the constant desired laser intensity. The Q_5 and Q_6 pair is series connected so that if one should fail shorted, the other will still maintain the laser current at a safe level. The probability of failure of two transistors at the same time is far lower than the probability of failure of a single transistor.

Editor's note: Due to variations in laser and photodiode efficiency, R_7 might need to be adjusted to ensure the laser output is within safety-regulation limits. **EDN**

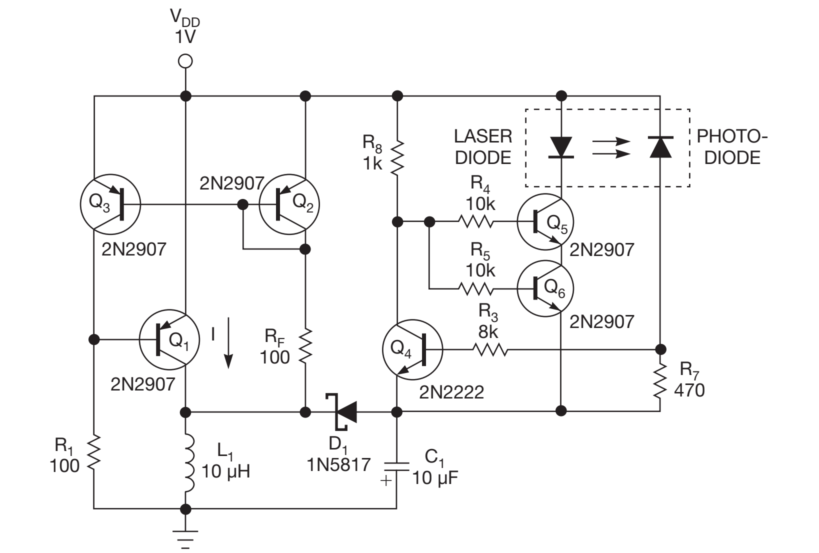

Figure 1 You can use this dc-dc step-up circuit and dual current-control transistors to safely power this laser from an almost-discharged battery.

Gate-drive transformer eases multi-output, isolated dc/dc-converter designs

Robert N Buono, Crestron Electronics Inc, Rockleigh, NJ

Oftentimes the biggest obstacle in designing an isolated dc/dc converter is the transformer design, a prospect that sometimes discourages designers from undertaking an otherwise straightforward design task. You can take advantage of the characteristics of an off-the-shelf gate-drive transformer and produce four separate isolated dc outputs. Gate-drive transformers are actually ideal for low-power dc/dc power transfer, because they have already been optimized for a high product of voltage and time (ET, or volt-microsecond, product) as well as for low leakage inductance.

A core with high permeability and low core loss at high switching frequency (F_{SWX}) supports the typical 10 to 15V applied primary voltage and the typical 500-nsec to 5-μsec on-time of switching frequencies between 100 and 500 kHz. This range of voltage and time is in the range needed for this dc/dc design. Also, a core geometry and winding configuration has already been chosen for low leakage inductance in order to produce fast rise and fall times, as well as low ringing. Lastly, the wire gauges used are sufficient for dc/dc-converter applications handling winding currents in the tens-of-milliamps range without excessive copper losses.

The Pulse Electronics P0585 gate-drive transformer has five windings, each with an identical number of turns[1]. One winding is wound with triple-insulated wire (TIW); the other four windings are standard magnet wire. You drive the TIW winding as the primary to provide a 3-kV_{RMS} primary-to-secondary breakdown rating. The breakdown voltage rating between the four secondary windings is not specified, but this type of wire insulation is typically used in offline supplies where up to 400V can be seen between windings.

Isolated outputs offer great flexibility. They provide a convenient way to break ground loops, power remote circuits at different ground potentials, and allow for simple negative or positive output voltage polarity selection. **Figure 1** shows the four secondaries of this transformer, creating four separate, equal-voltage outputs. You can, however, wire these secondaries in various series/parallel combinations to produce a myriad of output voltage/current combinations.

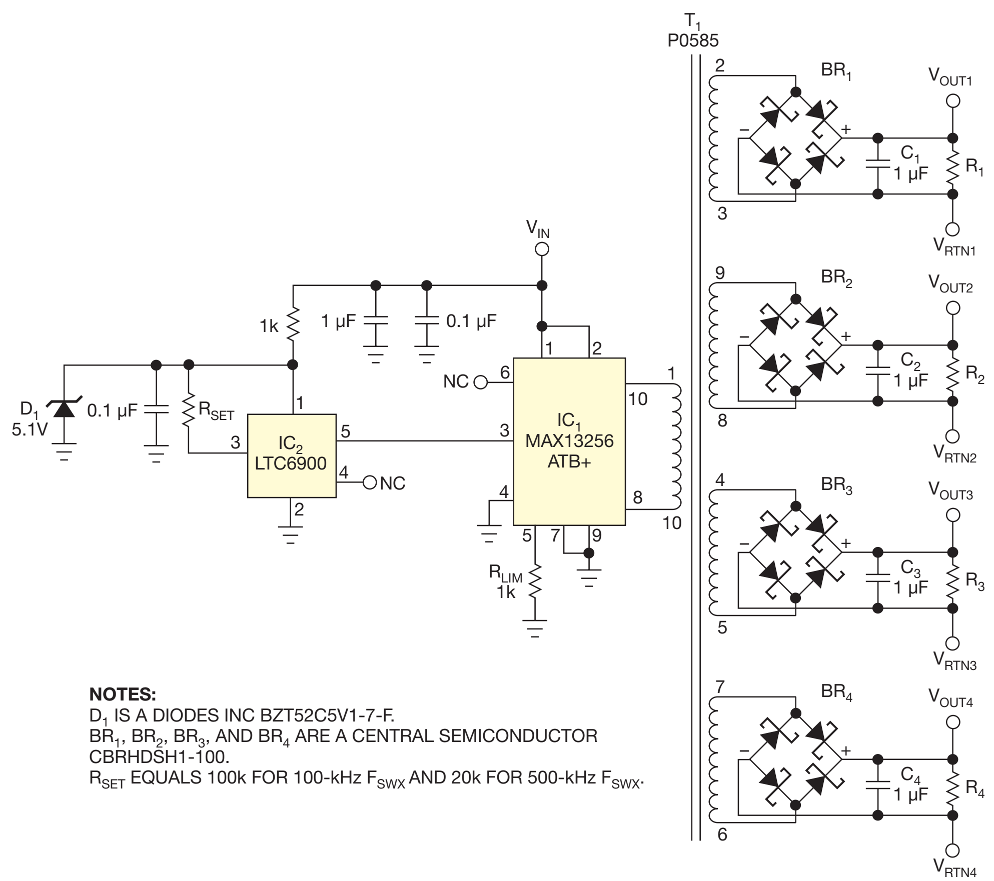

Figure 1 IC_1 drives the transformer primary with a 50%-duty-cycle square wave to avoid core saturation.

The Maxim MAX13256 H-bridge transformer driver (IC_1) is an ideal part for this application. It incorporates all of the functions needed for a standalone, transformer-isolated dc/dc converter. Its internal FETs withstand 36V and are configured as two separate push-pull outputs, which drive a transformer primary with a precise 50% duty cycle to avoid core saturation. It also incorporates adjustable and robust internal current limiting, so the outputs are protected against short circuits and recover nicely upon fault removal. It also incorporates undervoltage lockout (UVLO) to prevent switching activity when the input voltage is too low.

The Linear Technology LTC6900 clock source (IC_2) was added to allow precision adjustability of the switching frequency. The MAX13256 does have an internal clock, but most users would probably prefer to set the switching frequency themselves for overall system compatibility or EMI reasons. The MAX13256 accepts an external TTL-level clock, and its UVLO feature ensures that IC_2 is up and running before the ramping V_{IN} of IC_1 reaches its turn-on threshold. The value of R_{SET} determines the output frequency of IC_2, which is set to twice the desired switching frequency of IC_1.

Table 1 shows the measured results for 10, 12, and 15V input voltages at switching frequencies of 100 and 500 kHz. Due to the high switching frequency, low leakage inductance, and use of Schottky bridge rectifiers, the outputs produce very low voltage ripple of less than 20 mV pk-pk, even with low-value (1 μF) surface-mounted ceramic output capacitors. The **table** also shows the efficiency as well as the output voltage versus load current variation due to the outputs being unregulated. The outputs could be followed by linear regulators if lower noise or more tightly regulated dc outputs are desired.

For these measurements, full load is the loading that produces a measured primary current of 500 mA peak. This is the minimum current-limit threshold for the MAX13256 when R_{LIM} is 1 kΩ. Some designers may want to operate at less than these empirically determined full-load current levels for more margin against spurious overcurrent tripping. The light-load output voltage rise observed with the higher switching frequency is due to the fact that this is a snubber-less design for simplicity and high efficiency. As the switching frequency rises, more leakage inductance energy is generated, which transfers to the secondary windings and increases the measured output voltages.

Following is a brief tutorial on checking the transformer operating parameters against the data-sheet specifications. The P0585 transformer has a maximum ET product of 95 Vμsec. This calculation is the product of the maximum voltage impressed across the primary winding and the maximum time that voltage is present (on-time). Since the MAX13256 drives the transformer primary with a precise 50% duty cycle, the maximum ET product will occur with a 15V input voltage. At the lowest switching frequency of 100 kHz in this case, the maximum on-time at 100 kHz is 5 μsec. The maximum ET product is therefore 75 Vμsec, which meets the specification.

The peak flux-density spec is 2100 Gauss. To calculate peak flux density, equations 2A and 2B provided on the data sheet are based on V_{IN} and the switching frequency. Again, peak flux density is produced under operating conditions of V_{IN}=15V and a 100-kHz switching frequency. Note that in equation 2A, "D_{ON}" is the duty cycle of 50%, or 0.5, not time in microseconds. Under these conditions, the calculated peak flux density is 1512 Gauss, which meets the data-sheet specification.

Core loss is calculated using the formulas provided on the transformer data sheet. The results are 0.468W at 100 kHz and 0.117W at 500 kHz, which is lower than at 100 kHz due to the lower ET product.

The copper loss of 93.75 mW was calculated using the formula provided on the transformer data sheet. This simplified formula calculates copper loss based on I^2R losses in the windings and does not consider skin or proximity effects in the windings. Therefore, there is no frequency dependence in these simplified results, which are based on ±500-mA peak current in the primary winding and ±125-mA peak currents in each of the four secondary windings.

Using the temperature-rise formula from the transformer data sheet and the total losses calculated above (561.75 mW at 100 kHz), the predicted temperature rise of the transformer is 37.2°C.

This Design Idea uses the P0585 gate-drive transformer, but other (smaller) off-the-shelf gate-drive transformers can be used, especially if fewer outputs are needed, and at less current. Just be sure to check the transformer's maximum volt-μsec specification as in the example described here. EDN

REFERENCE

1 "P0585 transformer data sheet," Pulse Electronics, http://bit.ly/ZVsWYb.

TABLE 1 MEASURED RESULTS AT INPUT VOLTAGES OF 10, 12, AND 15V DC

Switching frequency	Measurements	V_{IN}=10V DC	V_{IN}=12V DC	V_{IN}=15V DC
100 kHz	V_{OUT} at I_{OUT} (light load)	9.82V at 1.31 mA	11.82V at 1.58 mA	14.77V at 1.97 mA
	V_{OUT} at I_{OUT} (full load)	8.19V at 110.29 mA	10.26V at 103.97 mA	13.17V at 111.51 mA
	P_{OUT} at full load (each output, total)	0.9W, 3.61W	1.07W, 4.27W	1.47W, 5.87W
	I_{IN}, P_{IN} at full load	450.4 mA, 4.5W	428.28 mA, 5.14W	463.33 mA, 6.95W
	Efficiency (P_{OUT}/P_{IN}) at full load	80.22%	83.07%	84.46%
500 kHz	V_{OUT} at I_{OUT} (light load)	12.92V at 1.72 mA	15.3V at 2.04 mA	18.74V at 2.5 mA
	V_{OUT} at I_{OUT} full load)	8.09V at 108.94 mA	10.2V at 103.36 mA	13.12V at 111.08 mA
	P_{OUT} at full load (each output, total)	881 mW, 3.53W	1.05W, 4.2W	1.46W, 5.84W
	I_{IN}, P_{IN} at full load	445.45 mA, 4.45W	426.26 mA, 5.12W	461.72 mA, 6.93W
	Efficiency (P_{OUT}/P_{IN}) at full load	79.33%	82.42%	84.13%

Originally published in the May 26, 1983, issue of *EDN*

Two ICs form F/V converter

Peter Winship, University of California, Berkeley

Using only six components, you can configure a circuit (**Figure 1**) whose output voltage is proportional to its input frequency. Moreover, only three of the components—capacitor C_0, resistor R, and the OP-07 op amp—must exhibit low drift for stable operation over temperature. The circuit provides linear operation well into the megahertz region.

The average current (I_{AVG}) from the 40106 Schmitt-trigger inverter's ground pin (pin 8) is linearly dependent on the frequency at which C_0 is discharged into the op amp's summing junction. The op amp forces this current to flow through the 13.33-kΩ feedback resistor, producing a corresponding voltage drop. The output voltage is

$$V_0 = -V_{CC}RC_0f,$$

where f is the input frequency. Adjust the 10-kΩ potentiometer to calibrate the converter.

The 1- and 0.1-µF capacitors smooth the transients that result from the rapid switching. For the **figure**'s values, the output ranges from 0 to −10V for inputs of 0 to 10 kHz. If you need higher-frequency operation, you must consider the effects of rapid switching on the CMOS inverter's supply current. Because a CMOS IC's power dissipation is proportional to frequency, you can simply add its supply current to the capacitor's discharge current in the calculations.

You can make a frequency summer by exploiting the fact that there are six Schmitt triggers per package: Attach a capacitor to each inverter's output, apply a different frequency to each input, and obtain V_0 proportional to the sum of the input frequencies:

$$V_0 = -V_{CC}R(C_1f_1 + C_2f_2 + \ldots + C_6f_6).$$

Moreover, you can extend the technique by paralleling additional ICs.

With the **figure**'s component values, the F/V converter yields ±0.4% max nonlinearity for inputs of 0 to 10 kHz. EDN

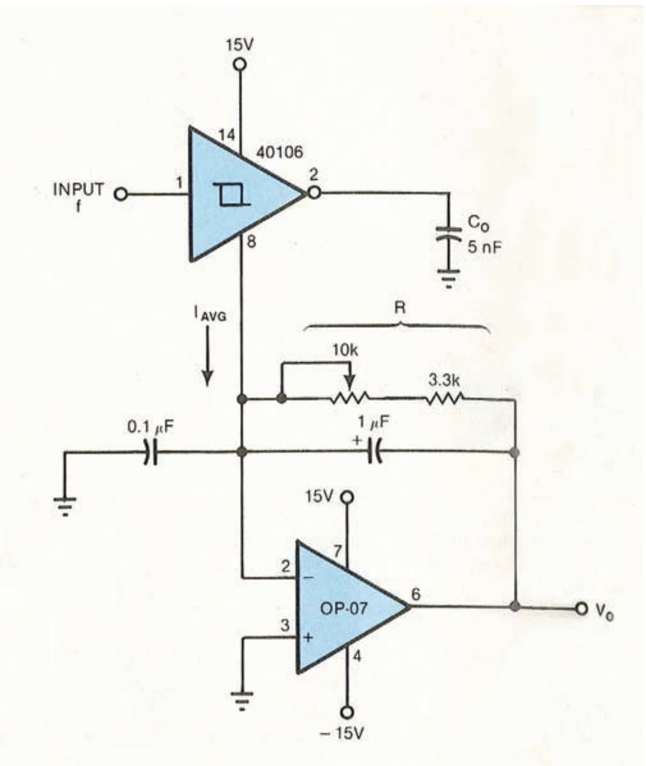

Figure 1 Using a capacitively loaded IC's supply-current/frequency dependence, this F/V converter yields 0 to −10V output with 0- to 10-kHz input frequencies. By paralleling more Schmitt triggers, you can use the circuit as a frequency summer.

What engineers must know about the supply chain

Accenture and others have reported that company stock prices can drop (bit.ly/WRwdJj) 7.5 to 12% following severe supply-chain disruptions (bit.ly/WRwbBb). PricewaterhouseCoopers estimates this drop at 9% and says that, for two-thirds of those companies, stock performance will still lag behind that of competitors a year after the disruption (pwc.to/14q9atK).

Electronics designers have a strong influence on supply-chain design. Their architectural- and technology-platform decisions certainly affect the security of supplies, and they play a major role in setting life-cycle costs. For years, companies have tried to bring operations and technology together—with varying degrees of success—by placing supply-chain people on development teams. But the fact remains that, the majority of the time, designers select the manufacturers and vendors that make it into their supply chain.

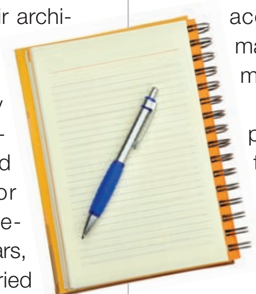

Proper supply-chain design can affect time to market, stock price, and the optimization of cost through benchmarked negotiations—10% or more, according to Freebenchmarking.com's Component Cost Estimator (registration required) (bit.ly/X8bMKk). That should be enough to capture a designer's imagination. Careers can be made on such gains.

What do top designers need to know about supply-chain design?

LEVERAGE

Leverage is a critical factor in negotiations. Without it, you take what is being offered; with it, you obtain amazing deals. Designers have tremendous leverage when selecting vendors for new products. Designers should not award a design win to any vendor until production-level pricing has been found acceptable through market-price benchmarking.

In a major design project, leverage the opportunity for price concessions by your production components. This step is critical with single-sourced devices with high IP content, such as FPGAs, chip sets, and processors. Benchmarking component prices before signing a contract is also critical on short-life-cycle ODM products.

REGISTRATION

Component registrations can happen without the OEM's knowing it. Registration locks a company into a supply channel and essentially eliminates any hope of ongoing cost reduction on a component. Even minimal or no engineering support from a distributor or representative can lead to a registration.

Demand to know whether you are being registered, and ask for justification as to why.

SOURCING

Supply-chain disruptions, supplier failures, and quality problems can happen at any time. Make sure you have specified at least two manufacturers for all components in your design (except single-source ones). Don't cheat and think that specifying a 5% version of a component as a second source for a 10% version counts; it still creates a single-point-of-failure risk. Alternate manufacturing sources also create leverage for better pricing.

Focus on a few component manufacturers to concentrate spending. Make sure these manufacturers deserve your business by being service-oriented, cost-effective, and reliable sources of acceptable, quality devices.

Specify components to concentrate the purchasing to the fewest component values possible. For example, don't specify many decoupling capacitor values on a circuit if one value will do.

SIMULATION

Simulate your design to ensure it works under all operating conditions of voltage, temperature, and manufacturing tolerance. For each component, make notes on the design file as to the level of testing that would be required to qualify a new component for substitution. Notes about factors such as a component engineering paper study, emissions testing, or re-simulating would greatly assist when components become obsolete or future cost-reduction opportunities are identified.

PORTABILITY

Design decisions on testing and assembly can lock a company into a specific contract manufacturer, because the cost of moving is prohibitive. Products that require multimillion-dollar testers are an example, particularly if a low-cost test method could have been used instead.

There may not be a reason to move from a particular contract manufacturer, but the prohibitive cost of moving still eliminates a degree of freedom to reduce leverage. Where possible, design so that products can be moved from factory to factory, manufacturer to manufacturer, or country to country with relative ease.

REGULATIONS

Companies have growing responsibilities to shareholders and government entities to report on risk as well as environmental and social performance. These requirements have become more demanding since the 2008 financial collapse and require more due diligence.

Some suppliers cooperate in these risk assessments, but others are less willing to do so. Make sure compliant, socially responsible companies that are willing to share information receive your business.

—by Ken Bradley, president, Lytica

This story was originally posted in two parts by EBN: *bit.ly/16TTqi4 and bit.ly/WRxS1v.*

MEMORY

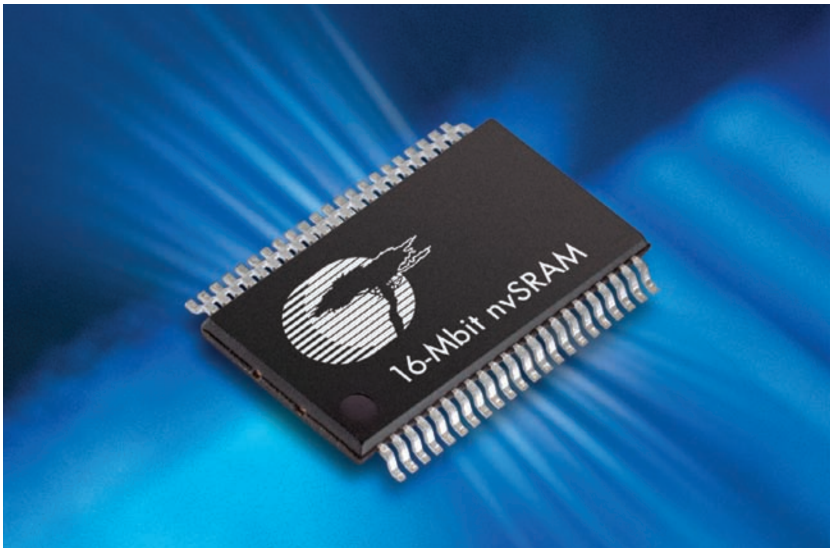

Cypress adds 16-Mbit nvSRAMs and synchronous NAND interfaces

A 16-Mbit nonvolatile SRAM family can interface directly with Open NAND Flash Interface (ONFI) and Toggle NAND bus controllers. The family addresses solid-state drives (SSDs) for enterprise systems, high-end programmable-logic controllers (PLCs), high-speed data/error loggers in storage, and networking equipment. The 16-Mbit nvSRAMs offer access times as low as 25 nsec; infinite read, write, and recall cycles with 20-year data retention; and an optional integrated RTC, which enables time-stamping of critical data to be logged. The ONFI NAND nvSRAM supports the ONFI 3.0 NV-DDR (100 MHz) as well as the ONFI 3.0 NV-DDR2 (200 MHz). The Toggle NAND interface nvSRAM is compatible with Toggle 2.0 NAND controllers for DDR operation at 200 MHz. Both the ONFI and the Toggle versions support single-channel operations in ×8 and ×16 data bus widths, and dual- and quad-channel operations in ×8-bit data bus width, which allows operations of up to 400 million transactions per second.

Cypress Semiconductor, www.cypress.com

LSI FSPs enhance performance of Ultrabooks

SandForce SF-2200/2100 Client Flash Storage Processors (FSPs) are specifically designed to satisfy the power-consumption requirements of Ultrabooks while also providing the performance benefits of flash-based technology. LSI FSPs are the intelligence that manages flash memory in solid-state drives (SSDs). The new LSI SandForce power-management optimizations enable SSD manufacturers to develop drives that extend battery life and reduce system-resume times, critical requirements in today's growing class of thin-and-light Ultrabook platforms.

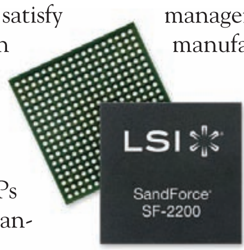

LSI, www.lsi.com

Micron offers NAND-based SSD for data centers

Micron Technology is leveraging its NAND-based flash-memory devices in a new solid-state storage module for SATA caching in data centers. The P400m conforms to standard 2.5-inch, 7-mm form factors and is offered in densities of 100, 200, and 400 Gbytes. The drives include software and firmware features such as XPERT (Extended Performance and Enhanced Reliability Technology) and can perform 10 drive fills per day for five years, roughly the equivalent of writing every photo posted to Facebook for 311 days straight, or 78 billion photos. The drives combine low latency and low power consumption with multiple features for onboard data-loss prevention.

Micron Technology, www.micron.com

Toshiba develops zero-wait-state flash memory for ARM MCUs

The TMPM440F10XBG high-speed flash-memory device provides zero-wait-state performance at 100 MHz for its ARM-based embedded microcontrollers. The Nano Flash-100 product merges high-speed programming based on NAND-flash-memory-cell device technology with NOR-flash-memory circuit technology. The device allows for 100-MHz performance and the higher densities required to match the performance of ARM-based microcontrollers.

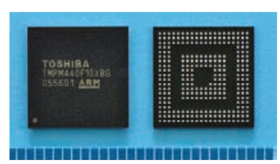

Toshiba, www.toshiba.co.jp

productroundup

Tek scope package tackles DDR4, DDR3L, and LPDDR3

DDRA analysis software for the MSO/DPO/DSA70000, DPO7000, and MSO/DPO5000 series oscilloscopes now has full electrical verification and conformance test support for JEDEC DDR4, DDR3L, and LPDDR3 memory standards. The software now supports all variants of the DDR standard. By automating test setup and execution, DDRA ensures designs are in full conformance with memory standards. If a memory system fails conformance testing, engineers can switch to debug mode and use such tools as Visual Trigger to isolate events of interest for deeper root-cause analysis with the DPOJET jitter and eye measurement suite. The company also offers new interposers (DIMM, package-on-package, and embedded) for DDR4 and LPDDR3. The controlled impedance and trace-length matching of these interposers in combination with de-embedded filters for the oscilloscope provide an accurate representation of the signal. Prices start at $5910. DDRA support for DDR4, DDR3L, and LPDDR3 will be a free upgrade for existing users.

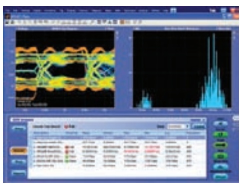

Tektronix, www.tek.com

4-Mbyte flash from TI guarantees operation in harsh environments

The SM28VLT32-HT is a high-temperature, nonvolatile 4-Mbyte flash-memory device guaranteed for at least 1000 hours of operating life in harsh-environment applications. The flash device is designed to eliminate the need for costly screening and qualification testing of industrial-grade components for temperature ranges outside data-sheet specifications. The device features a temperature range

from −55°C to +210°C, an SPI interface that simplifies design and packaging and reduces pin count, and small package integration into multichip modules for systems with limited board space. The SM28VLT32-HT is sampling now in an 8×25-mm ceramic flat pack, with Known Good Die (KGD) package options and volume quantities expected to be available in 1Q13.

Texas Instruments, www.ti.com

SPI-flash devices from Microchip offer enhanced memory for battery-operated designs

The SST25PF020B, SST25-PF040B, and SST25PF080B, offering 2-, 4-, and 8-Mbit memory, respectively, feature fast program and erase times, an operating voltage range of 2.3 to 3.6V, an active read current of 10 mA (typical) at 80 MHz, standby current of 10 µA (typical), and low-power operation in a small-footprint package. The devices are intended for battery-operated accessories, sensors, and equipment applications in consumer, medical, and industrial markets. Partitioned into uniform 4-kbyte sectors with 32- and 64-kbyte blocks, the devices enable flexible partitioning of program and data code in the same memory block while offering fast erase and program performance, with sector and block erase times as fast as 18 msec, and erasing the entire device in 35 msec. Word-programming time is 7 µsec using AAI (Auto Address Increment), and

reliability for the devices is rated at 100,000 endurance cycles (typical) and greater than 100 years of data retention. The SST25PF020B starts at 53 cents each, in 8-lead 150-mil SOIC, 8-contact USON (3×2-mm), or 8-contact WSON (6×5-mm) packages (10,000). The SST25PF040B starts at 66 cents each, in 8-lead 150-mil SOIC, 8-lead 200-mil SOIC, or 8-contact WSON (6×5-mm) packages (10,000). The SST25PF080B starts at 81 cents each, in 8-lead 150-mil SOIC, 8-lead 200-mil SOIC, or 8-contact WSON (6×5-mm) packages (10,000).

Microchip Technology, www.microchip.com

32-bit RX MCU platform from Renesas extends memory options

The 32-bit RX MCU platform has expanded with 120 new devices in the RX631 and RX63N series. The MCUs feature up to 256 kbytes on-chip RAM. The devices offer an optional DEU (Data Encryption Unit), supporting AES encryption/decryption at 128-, 192-, and 256-bit key lengths, and an RTC feature with antitamper detection and time-stamp functionality designed to reduce power consumption with 2V RTC operations. The RX63N/RX631 devices will be available in mass production by December 2013, with pricing ranging from $6.35 for 1-Mbyte flash/192-kbyte SRAM in an LQFP 100-pin package without CAN (10,000) to $9.50 for 2-Mbyte flash/256-kbyte SRAM in an LQFP 176-pin package with CAN (10,000). The RX210 devices were also added to the low-power, low-voltage RX200 series. They feature up to 1-Mbyte on-chip flash, while achieving active-mode power consumption of 7.2 mA at 50-MHz operation and a standby mode of 0.4 µA. The MCUs feature an extended supply range, operating between 1.62 and 5.5V with reprogramming supported across the full sup-

ply voltage range, and a range of on-chip peripherals and timers, including watchdog timer, RTC, and motor-control timers. The RX210 devices were expected to be available in mass production by the end of March 2013. The RX210 device with 1-Mbyte flash, 96K RAM, and 100 LQFP is priced at $5.13 per unit (10,000), while devices with 1-Mbyte flash, 96K RAM, and 144 LQFP are priced at $5.98 per unit (10,000).
Renesas Electronics, am.renesas.com

Toshiba microSDHC memory card supports HD-content viewing

New microSDHC memory cards, based on SeeQVault technology, provide a platform for storing and distributing high-definition content, and maximum write and read speeds of 20 and 40 MBps, respectively. SeeQVault applies bidirectional authentication with a unique identifier and public key infrastructure to ensure a high-level secure environment for storing and distributing HD content. The technology provides access to the latest digital content and can be used for playback of content in a smartphone, tablet, or TV. The initial product line-up will cover 16- and 32-Gbyte capacities compatible with the Ultra High Speed Bus I and the UHS Speed Class 1.

Toshiba Electronics Europe, www.toshiba-components.com

2001 Electronic Components offers 64-Mbit SPI Flash

The A25LQ64 supports the standard SPI and features a software protocol allowing operation on a simple three-wire bus while it is in single input/output mode. With high-performance normal, dual, and quad read outputs, the SPI can clock frequencies up to 66, 84, and 104 MHz, respectively. The module supports Quad Peripheral Interface (QPI) read mode, enabling the user to

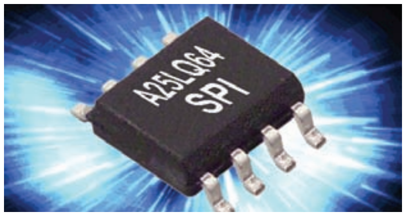

take advantage of the quad I/O serial flash and fast read data. The A25LQ64 produces a maximum low active read current of 25 mA at 104 MHz, 20 mA low active read, and low active erase/programming currents at 84 MHz, and draws a maximum standby current of 10 μA in standby mode. The high-speed performance module allows for a maximum program time of 3 msec, a sector erase time of 150 msec, a block erase time of 500 msec, and a chip erase time of 25 sec. The SPI flash contains auto-erase and auto-program algorithms, which verify data at a selected sector, block, or page. The A25LQ64 SPI flash comes in an 8- and 16-pin SOP, 8-pin WSON, and 24-ball BGA, and is able to retain its programmed state for 10 years. Pricing starts at $1.47 (1000) in a BGA 24 package.
2001 Electronic Components, www.2k1.co.uk

VIA Labs' USB 3.0 to NAND-flash-controller chips improve performance

The VL752 and VL753 SuperSpeed USB 3.0 to NAND-flash controllers deliver enhanced performance using the latest NAND-flash technologies including 19-, 20-, and 21-nm geometries and have been certified by the USB Implementers Forum (USB-IF) for SuperSpeed operation, ensuring interoperability and backwards compatibility. The two-channel VL752 and single-channel VL753 offer data-transfer speeds up to 280 MBps, flexible bad-block management, support for the latest 200-MHz DDR NAND flash, and a robust ECC engine for greater data integrity. The package sizes have been

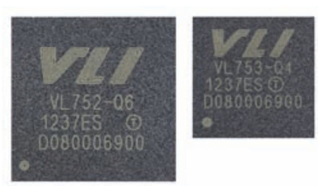

reduced to a QFN68 8×8-mm package for the VL752 and a QFN48 6×6-mm package for the VL753 to enable the smallest-form-factor designs. The standard prices are $1.50 for the VL752 and $1 for the VL753.
VIA Labs, www.via-labs.com

Micron's 128-Gbit NAND-flash device measures 146 mm^2

Using 20-nm process technology to move from a multilevel memory cell to a triple level yielded a 128-Gbit NAND-flash device that measures only 146 mm^2. The memory is optimized for removable-storage options such as USB drives and flash cards, and stores 3 bits of information per cell. It is sampling to select customers now, and is expected be in volume production during the second quarter of 2013.
Micron Technology, www.micron.com

Decoupling caps are where it's at

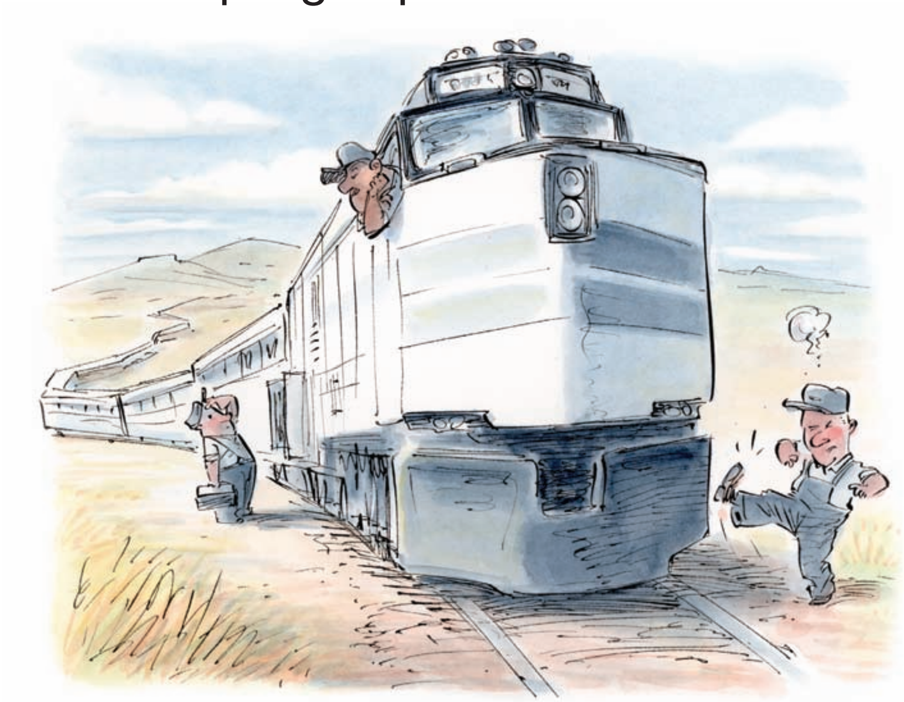

DANIEL VASCONCELLOS

The company I had transferred to in the Midwest back in the 1980s made locomotives and was plagued by a long-term field problem. An overcurrent-protection circuit would sometimes trip when the locomotive controller was cold. This situation posed a serious problem for operation of the locomotive, because if the engine lost power, trains would not be able to operate and pull reliably, and the locomotive engine became a million-dollar paperweight.

A unit exhibited the problem at the factory in the rail yard, and the assembly crew gladly brought the controller back to the engineering lab for evaluation. The box was ice-cold and displayed the apparent overcurrent condition. Test engineering fired up the automated test procedure and quickly trapped the fault.

The design was set up so that the microprocessor would select an analog-to-digital channel to read back, and then read back the value for scrutiny. The problem here was the value coming back was way too high, and was equivalent to reading back many, many more amps on the traction motors than desired.

Closer inspection showed that the selected channel was being written satisfactorily by the microprocessor but wasn't getting held reliably by the latch. Inspecting the power-supply (V_{CC}) lines to the chip showed a dip in the supply coincident with the write of the register. So the analog-to-digital channel we thought we were reading was not the intended one.

The board layout showed a long V_{CC} trace from the decoupling capacitor to the IC before being stitched down to the V_{CC} plane. The dip in the V_{CC} caused the latch/register IC—some multi-D-latch variety of the 74ALS374, if you remember the pre-VLSI days, called MSI (medium-scale integration) days—to lose "consciousness." When V_{CC} returned, the latched address value was scrambled and was pointing to the wrong analog-to-digital channel to be read.

A product from Rogers Corp, which I came across while reading some trade literature, provided the solution to this problem. Rogers made a nifty device that fit right below the DIP and inserted nicely under the existing IC package. We quickly obtained samples, tacked them piggyback-style onto the tops of the offending latch, and then watched the V_{CC} droop diminish well within limits. The latch could remain functional during the microprocessor write, and the system could read back the correct channel. Multiple quick fixes later, we had our solution without having to re-lay out the board! Because we had found the root cause, the test results and field repairs went smoothly and quickly.

A PRODUCT FROM ROGERS CORP PROVIDED THE SOLUTION. THE NIFTY DEVICE FIT RIGHT BELOW THE DIP AND INSERTED NICELY UNDER THE EXISTING IC PACKAGE.

Packaging technologies have changed, but decoupling capacitors are the reason every electronic design can actually function. If you don't believe me, try to build a board without any decoupling caps and send me the scope photos of the V_{CC} lines near ICs!

So the lesson learned was that the V_{CC} and grounds must be fed and cared for judiciously, or all hell will break loose. In addition, I swore to myself that going forward I would never design a system where you can write but not read back a register to verify the write. In some ways it reminded me of the write-only memory, but with much worse consequences. (Throughout my career, I have accumulated a few of these principles and collectively refer to them as, well, Maxim's maxims.) EDN

Maxim "Jay" Skender has been an electronic design engineer, for both analog and digital (FPGA and embedded microprocessor) applications, for most of his 28 years in the industry. He received a bachelor's degree in electrical engineering from California Polytechnic State University (San Luis Obispo, CA) in 1985.

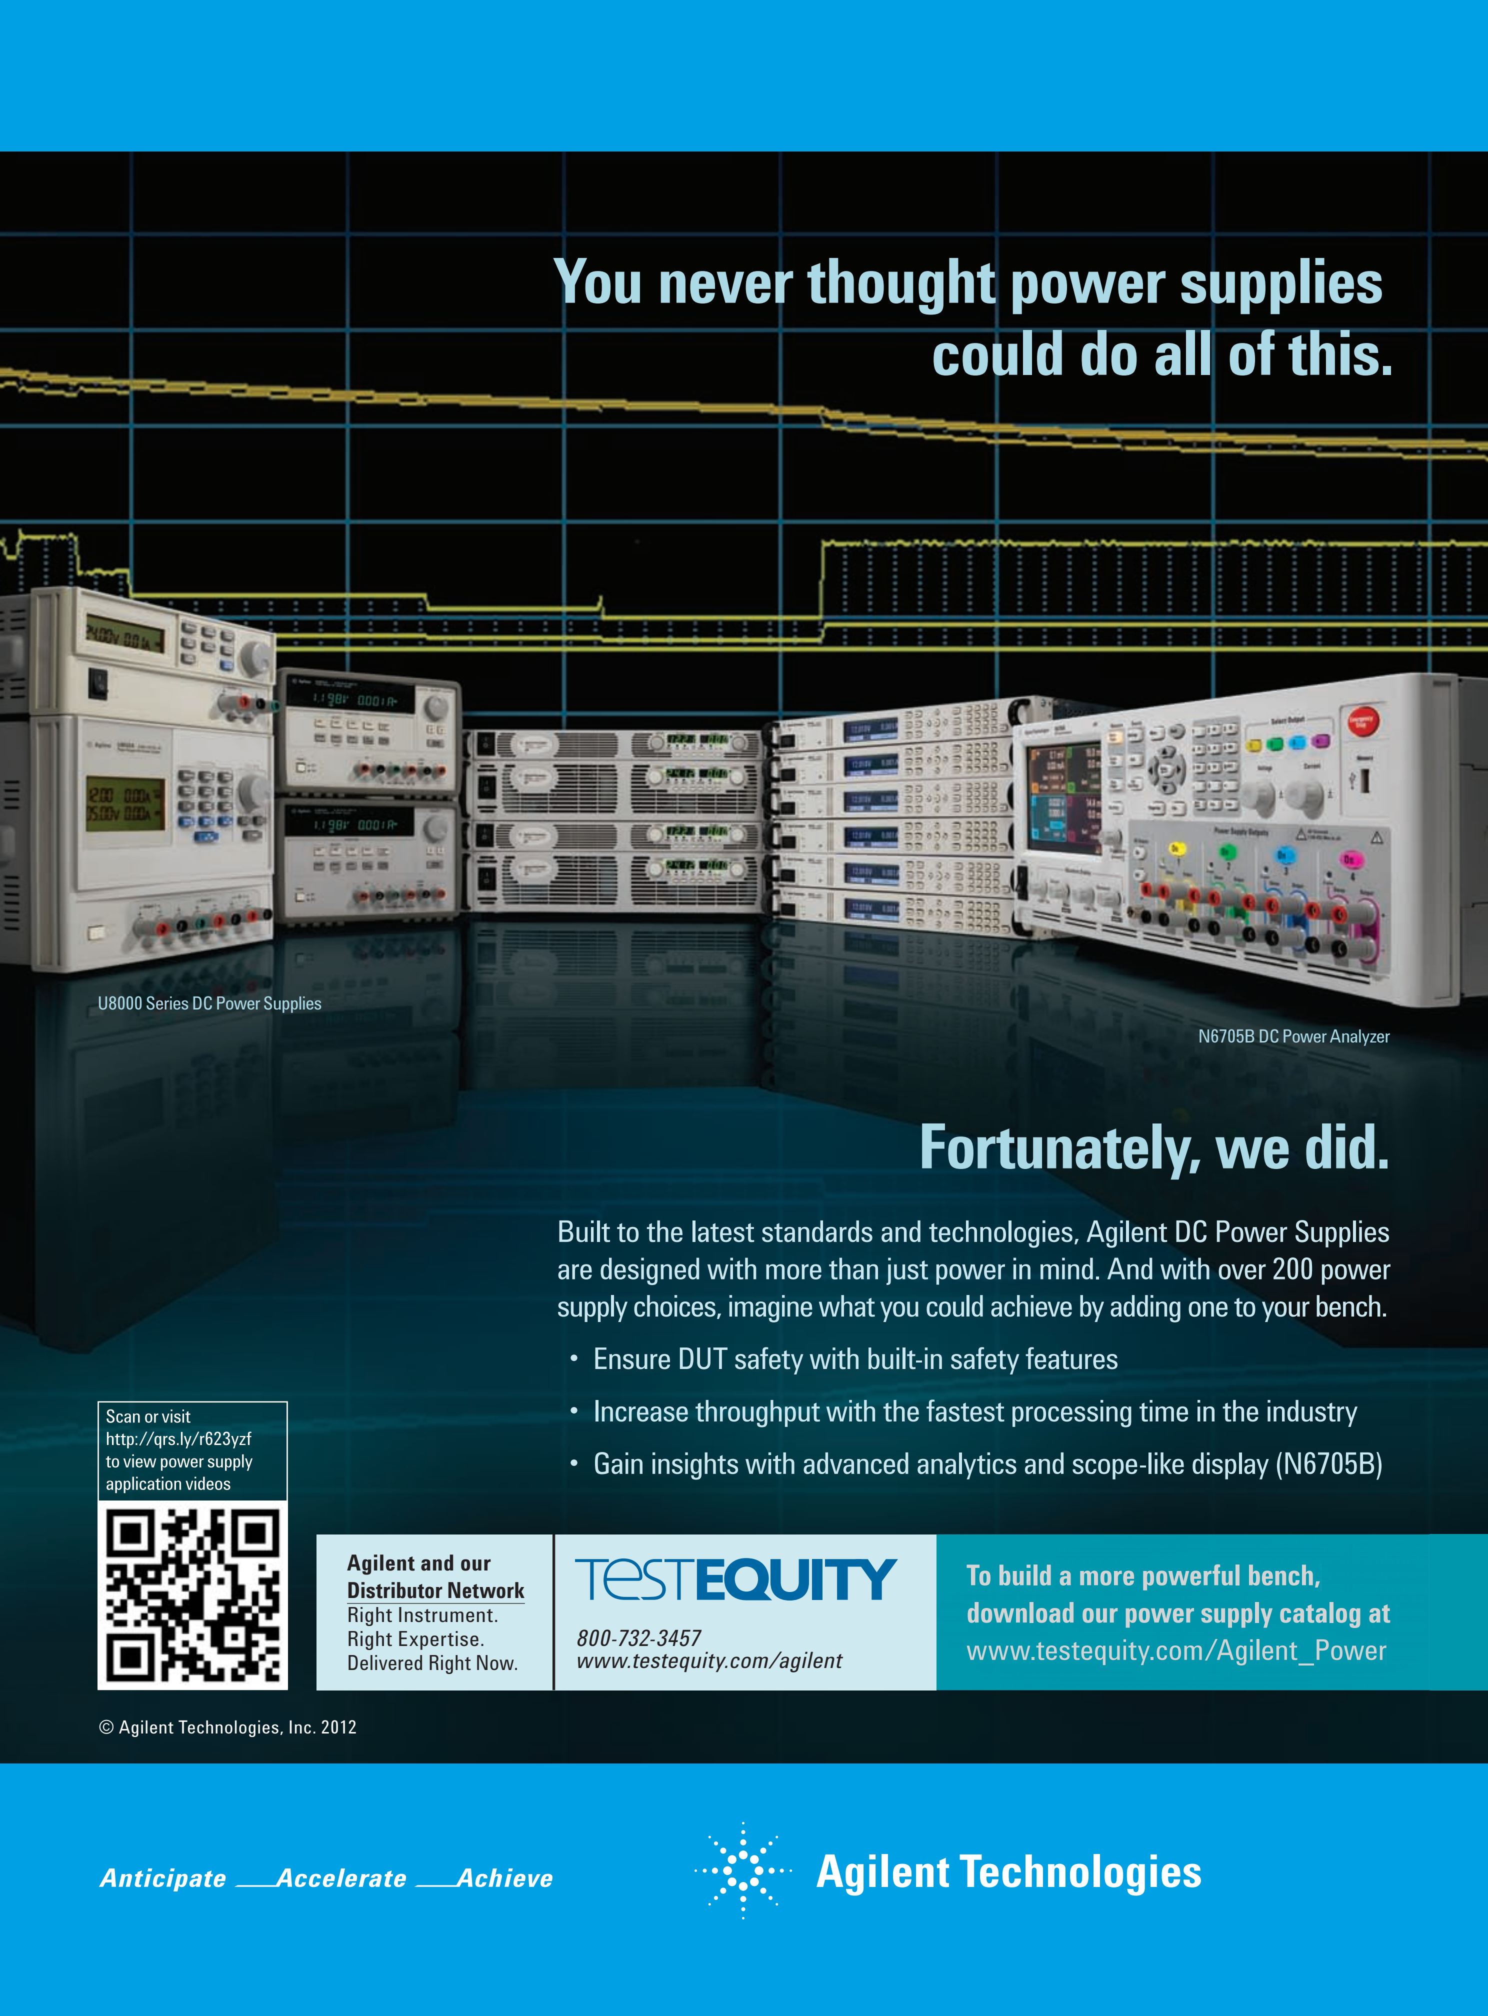
You never thought power supplies could do all of this.
U8000 Series DC Power Supplies
N6705B DC Power Analyzer
Fortunately, we did.
Built to the latest standards and technologies, Agilent DC Power Supplies are designed with more than just power in mind. And with over 200 power supply choices, imagine what you could achieve by adding one to your bench.
• Ensure DUT safety with built-in safety features
• Increase throughput with the fastest processing time in the industry
• Gain insights with advanced analytics and scope-like display (N6705B)
Scan or visit http://qrs.ly/r623yzf to view power supply application videos
Agilent and our Distributor Network
Right Instrument.
Right Expertise.
Delivered Right Now.
TESTEQUITY
800-732-3457
www.testequity.com/agilent
To build a more powerful bench, download our power supply catalog at www.testequity.com/Agilent_Power
© Agilent Technologies, Inc. 2012
Anticipate Accelerate Achieve
Agilent Technologies