UBM Electronics

EDN®

ELECTRONICS DESIGN NETWORK

MAY 24
Issue 10/2012
www.edn.com

POWER: A SIGNIFICANT CHALLENGE IN EDA DESIGN

Page 30

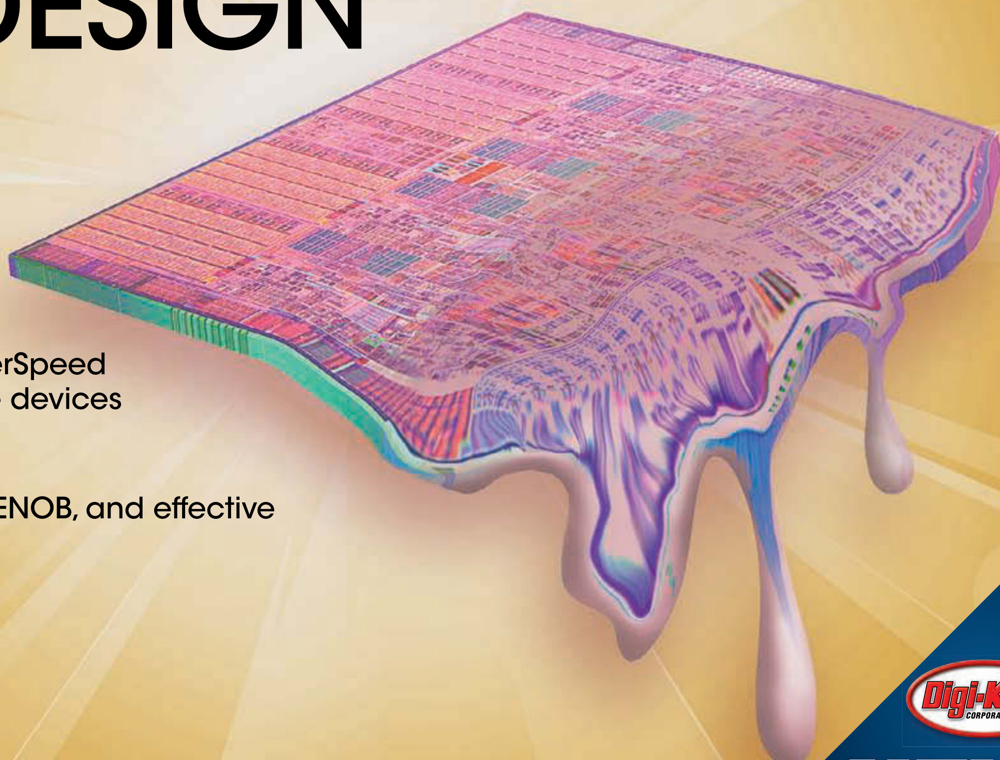

USB 3.0: bringing SuperSpeed connectivity to mobile devices

Page 25

Understanding noise, ENOB, and effective resolution in ADCs

Page 39


BEST of CLASS!
Which single distributor comes to mind as being 'best of class' for broadest overall product selection?
SOURCE: 2011 DESIGN ENGINEER AND SUPPLIER INTERFACE STUDY, HEARST BUSINESS MEDIA ELECTRONICS GROUP
2%
4%
6%
7%
8%
10%
44%
Distributor F
Distributor E
Distributor D
Distributor C
Distributor B
Distributor A
Digi-Key
SAME-DAY SHIPPING
on orders entered by 8:00 pm CENTRAL
NEXT-DAY DELIVERY
Digi-Key® CORPORATION
1.800.344.4539
DIGIKEY.COM
2 MILLION PARTS ONLINE | 500+ INDUSTRY-LEADING SUPPLIERS | NEW PRODUCTS ADDED DAILY
Digi-Key is an authorized distributor for all supplier partners. New products added daily. © 2012 Digi-Key Corporation, 701 Brooks Ave. South, Thief River Falls, MN 56701, USA
ecia
ecsn member
CEDA

YOUR MCU SHOULD WORK FOR YOU, NOT THE OTHER WAY AROUND.

IS YOUR MCU MAKING YOUR JOB HARDER THAN IT HAS TO BE?

AppBuilder Drag-and-Drop
Rapid Code Generation Utility

Introducing Precision32™

32-bit mixed-signal MCUs engineered to make your job easier.

With highly-integrated analog and peripherals and a patented dual-crossbar for maximum design flexibility, Silicon Labs' Precision32™ 32-bit microcontrollers work for you. These mixed-signal MCUs are integrated with useful features not typically found in other MCUs:

- USB oscillator for full crystal-less operation
- 5 V regulator to operate directly off USB
- 6 programmable high-drive pins that can provide up to 300 mA of current
- 16 capacitive sensing input channels for touch buttons or sliders

And the integrated analog is highly configurable—choose from 28 analog to digital channels with multiple modes (including 8, 10, 12-bit), interleaving mode (up to 2 Msps sampling), programmable auto scan to switch between channels without CPU intervention and burst mode to auto average up to 64 samples.

Easily customize your design!

The free Precision32 Eclipse-based IDE and AppBuilder drag-and-drop GUI allow you to easily set-up and adjust pin configuration, map peripherals, configure clocking trees and even generate source code—all without wading through the data sheet to configure registers.

Need MCUs that work for you? No problem.

Download white papers, data sheets, free software drivers and tools, and order hardware to start prototyping today.
www.silabs.com/32bit-MCU

SILICON LABS®
Engineering for a mixed-signal world.

EDN contents

5.24.12

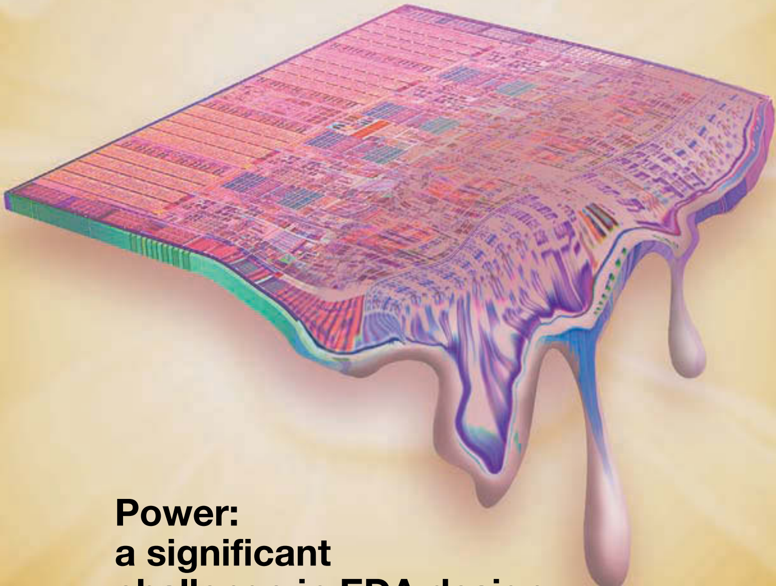

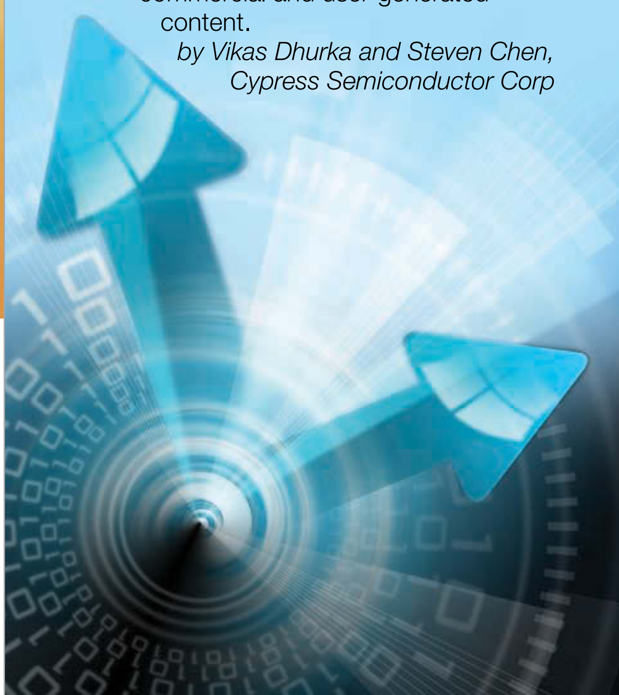

COVER IMAGE: THINKSTOCK/DAVID NICASTRO

DESIGNIDEAS

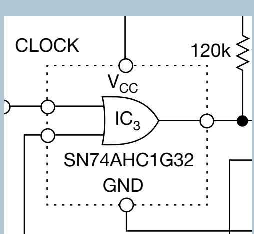

▶ Find out how to submit your own Design Idea: http://bit.ly/DesignIdeasGuide.

mouser.com
Distributing semiconductors and electronic components for design engineers.

Authorized Distributor


Texas, California, New Jersey **USA**

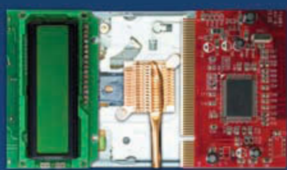
Jalisco **MEXICO**

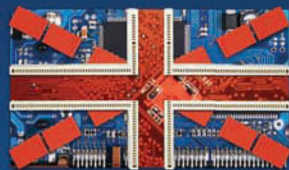
Buckinghamshire **UNITED KINGDOM**


Barcelona **SPAIN**


Brive-La-Gaillarde **FRANCE**

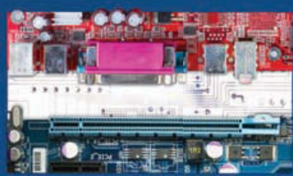
Eindhoven **THE NETHERLANDS**

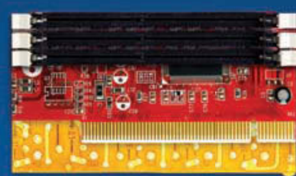
Munich **GERMANY**

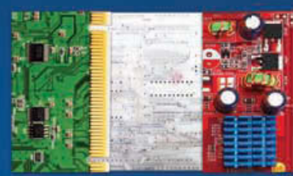
Assago-MI **ITALY**

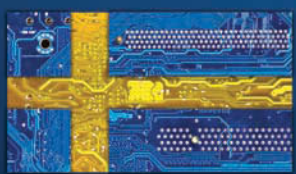
Upplands-Väsby **SWEDEN**

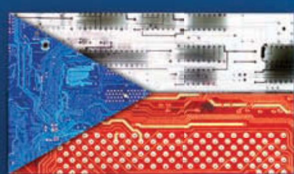
Brno **CZECH REPUBLIC**

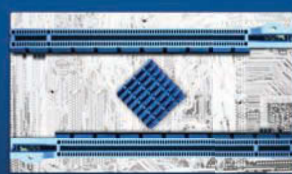
Raanana **ISRAEL**


Bengaluru **INDIA**


Hong Kong, Shanghai **CHINA**


Bangkok **THAILAND**


SINGAPORE


Taipei **TAIWAN**

With local support all over the world, **we're fluent in technology.**

Mouser delivers the components you need, on-time. And with local Technical Support and Customer Service Experts in 19 locations around the world, you'll find the newest components to launch your new design seamlessly.

mouser.com | The Newest Products for Your Newest Designs®

MOUSER ELECTRONICS®
a tti company

contents 5.24.12

pulse

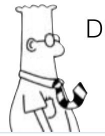

DEPARTMENTS & COLUMNS

50

54

EDN® (ISSN# 0012-7515) is published semimonthly by UBM Electronics, 600 Community Drive, Manhasset, NY 11030-3825. Periodicals postage paid at Manhasset, NY, and at additional mailing offices. SUBSCRIPTIONS—Free to qualified subscribers as defined on the subscription card. Rates for nonqualified subscriptions, including all issues: US, $150 one year; $250 two years; $300 three years. Except for special issues where price changes are indicated, single copies are available for $10 US and $15 foreign. For telephone inquiries regarding subscriptions, call 847-559-7597. E-mail: edn@omeda.com. CHANGE OF ADDRESS—Notices should be sent promptly to EDN, PO Box 3609, Northbrook, IL 60065-3257. Please provide old mailing label as well as new address. Allow two months for change. NOTICE—Every precaution is taken to ensure accuracy of content; however, the publishers cannot accept responsibility for the correctness of the information supplied or advertised or for any opinion expressed herein. POSTMASTER—Send address changes to EDN, PO Box 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 57, Number 10 (Printed in USA).

EDN online

JOIN THE CONVERSATION

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

In response to "Udacity biz model: School is accountable for students' value to employers," posted in Margery Conner's PowerSource blog at http://bit.ly/JgPgU4, Daniel van der Weide comments:

"Udacity is a tantalizing new model for teaching concepts and skills such as programming. Bricks-and-mortar schools where students interact with faculty and build relationships in that high-bandwidth medium known as human presence will be hard to replace. Even (or especially) in engineering, it's not what you know; it's who you know."

In response to "Tesla—connecting the dots," posted in Steve Taranovich's Anablog blog at http://bit.ly/IGyAtD, Pete O comments:

"I went to the Henry Ford Museum and saw a lot of things attributed to Edison but no mention of Tesla. Ford worked for Edison before he made it big in the auto world. I think Ford and Edison were good friends. Too bad that friendship distorted history, [as if] Tesla never accomplished anything."

In response to "Before exceptionalism comes can-do-ism," posted in Patrick Mannion's Design Cycle blog at http://bit.ly/HlJl4a, "Just an engineer" comments:

"I think we've let our politics drive wedges of separation between us, forgetting how much we need the spirit of working together for common goals. What used to be 'United we stand' has given way to 'my way or the highway.'"

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

CONTENT

Can't-miss content on EDN.com

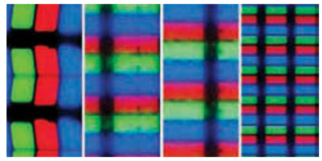

HOLD ME CLOSER, TINY DISPLAY

Why have computer displays and televisions remained at fairly conservative pixels-per-in. metrics, whereas tablets and, especially, smartphones have quickly pushed the envelope in the manner of Apple's Retina display, for example?

http://bit.ly/JtKr7J

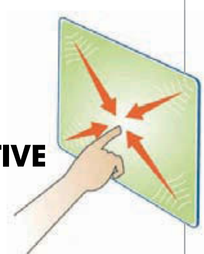

AN INTRODUCTION TO CAPACITIVE SENSING, PART 1

The popularity of this technology, especially in HIDs (human-interface devices), has grown rapidly due to its ability to reduce manufacturing cost, increase product life span by eliminating mechanical components, and enhance a product's look and feel.

http://bit.ly/IGwZ7h

ENGINEERING COMMUNITY

Opportunities to get involved and show your smarts

In *EDN*'s Tales from the Cube, engineers relate their most vexing design challenges—and how they conquered them. What's your Tale? Tell us in 700 words or so about a memorable experience you've had successfully solving an engineering problem at work or in your spare time. Visit www.edn.com/tales for samples of previously published Tales, or send your own to edn.editors@ubm.com.

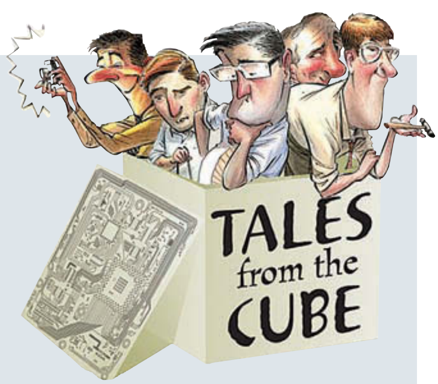

BY PATRICK MANNION, BRAND DIRECTOR

Ignore politics at your peril

When I recently heard banging on my front door, I came dashing from my office, in no mood for a door-to-door pitch. Instead, I got a chance to put my time where my mouth is with regard to engineering and politics; so far I've again failed to act—but at what price?

As engineers, we eschew politics. We're "above all that." We think that we're smarter, more creative, more logical, and more able to fix "everything." The further politicians stay out of our way, the better. On a darker slant, like most other people, we just can't stand the asinine arguments, the backstabbing, the manipulation of statistics, and the power of lobbyists and money over reason when it comes to policy. It's just too frustratingly out of—and, it seems, beyond our—control. Besides, we have better things to do, and we're too busy doing them.

Yet, we continue to complain. We complain about the lack of quality education for the next generation of engineers, about outsourcing, about government's selection of technology winners, company favoritism, the lack or overabundance of military spending, the lack of or misguided direction of energy policy, the overabundance of H1-B visas, the overabundance of government interference, and the lack of guidance. The primary complaint, however, is a lack of respect for engineers, from government through to our own chief executives—at least from those who aren't themselves engineers.

We gripe but carry on about our business. The last time I mentioned in this column anything to do with politics, I got a few angry e-mails about how *EDN* and politics should be separate. I relate to that view. *EDN* is, in many ways, the last bastion of engineering sanity. Being here is like closing the office door on all of the outside noise so that we can focus on what's real and what's doable. The sensation is real and relieving. But is it right to ignore that nether region of politics?

I was in that world doing what I do when this guy starting banging on my door. A neighbor I'd never met was going door to door asking people to sign up and get politically active in the local community. His party affiliation is irrelevant here. The point is that he was out there rallying support for his beliefs, and I decided I'd get behind him: As the saying goes, all politics are local, so I signed up. Why? Well, although we eschew politics, politics don't eschew us. They affect and control most of what we do, yet we stand by, happy in our own world, proud, right, wise, and *vulnerable*. We think that we know what's going on in Washington, and we stay away from it, but do we really know? Have you followed the numbers to see who is lobbying for our government's attention and whether it will benefit or hurt you? Neither the IEEE nor the Semiconductor Industry Association is "there for us."

I recommend that you visit www.opensecrets.org, which lists the expenditures of the top lobbyists (**Table 1**). The medical sector, from 1998 to 2012, spent more than $269 million on lobbying; pharmaceuticals spent $219 million; realtors spent $184 million; government and military contractors, including GE, at $268 million, and Northrop Grumman, at $176 million, also spent another huge amount.

It's shocking to me that there is no representation on behalf of electrical or electronics engineers. Lawyers, doctors, and others all have their say. Did you know that Google is the top Internet/computer lobbyist? In 2012, it has pumped $5 million into lobbying, with Microsoft lagging far behind at $1.8 million and other "giants" such as Intel, Cisco, and Texas Instruments falling far below $1 million.

These are good companies, but who is lobbying for engineers and engineering? I signed up on my neighbor's "get active" list, but I still haven't gone to any meetings. I'm just too busy doing things. Is that our lot? Are we just doers, at the mercy of a government beholden to everyone but the engineers driving the next generation of innovation? EDN

TABLE 1 TOP-SPENDING LOBBYISTS, 1998–2012

Lobbying client	Amount spent ($M)
US Chamber of Commerce	831.9
American Medical Association	269.5
General Electric	268.8
Pharmaceutical Research and Manufacturers of America	219.4
American Hospital Association	219.2
AARP	214.9
Blue Cross/Blue Shield	184.9
National Association of Realtors	184.4
Northrop Grumman	176.0
Exxon Mobil	173.6

Source: Center for Responsive Politics

Contact me at patrick.mannion@ubm.com, or comment directly on this column at www.edn.com/120524ed.

EDITED BY FRAN GRANVILLE

pulse

INNOVATIONS & INNOVATORS

MIMO PXI vector signal analyzer offers 780-MHz analysis bandwidth

Agilent Technologies has announced a wideband multiple-input/multiple-output PXI vector signal analyzer that supports 80- plus 80-MHz single-input/single-output and 80-MHz or 160-MHz MIMO 802.11ac measurements. The instrument also provides 780-MHz analysis bandwidth to address other wireless standards. The frequency range is specified at 10 MHz to 26.5 GHz.

With this new wideband MIMO PXI VSA, Agilent claims the industry's highest-bandwidth signal analysis. With one to four channels in a single PXIe chassis, the product provides analysis capabilities for a 4×4 MIMO configuration in a 4U-form-factor M9018A chassis.

Providing phase-coherent channel analysis, the analyzer helps R&D engineers validate their MIMO 802.11ac designs. The unit works with Agilent's 89600 VSA software.

The VSA comprises Agilent's four-channel M9362A-D01 downconverter, 12-bit M9202A digitizer, M9302A local oscillator, M9168C RF attenuator, and M9352A IF amplifier/attenuator modules for one to four measurement channels.

The two-channel wideband MIMO PXI VSA sells for $114,000; multichannel-analysis capability for the M9362A-D01 is available through a free download of the M9202A software. **—by Janine Love**

▷**Agilent Technologies**, www.agilent.com.

TALKBACK

"'Who are *you*, my lesser-life-form electronics technician, to tell me, *the engineer*, what the problem is?' [I've] been there more than once, so this [sentiment] didn't surprise me."

—Jim Sanders, former electronics/computer technician, in *EDN*'s Talkback section, at http://bit.ly/JeY7qo. Add your comments.

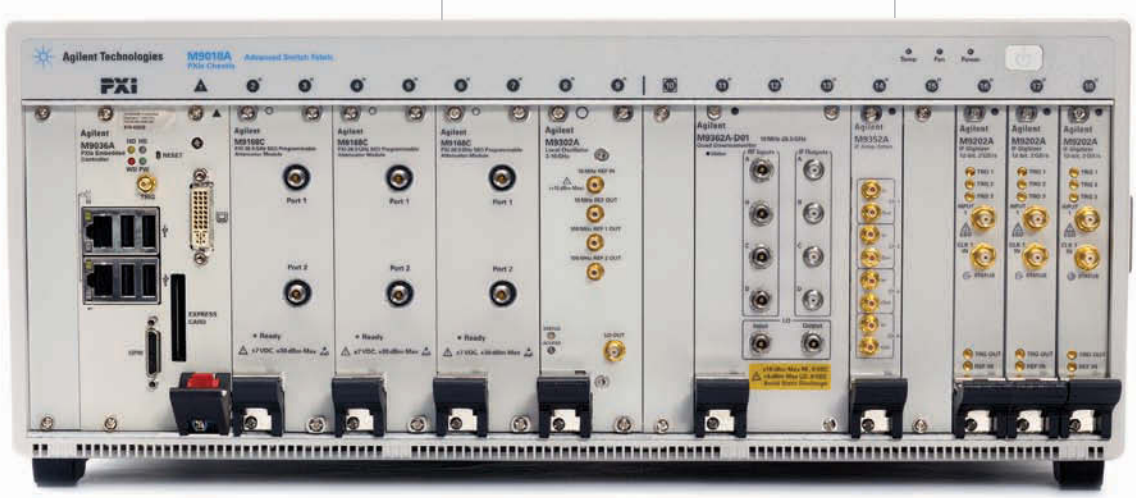

With one to four channels in a single PXIe chassis, the MIMO PXI VSA provides analysis capabilities for a 4×4 MIMO configuration in a 4U-form-factor M9018A chassis.

RAQ's

Rarely Asked Questions

Strange stories from the call logs of Analog Devices

Switching Bits

Q: With no signal at the input to my ADC, why are the output data bits switching?

A: People unfamiliar with high-speed ADCs might expect a converter's digital output to be constant given a static analog input, but this is similar to expecting to see a simple dc offset error at the output of an op amp with no input signal. If you remove the input signal to an amplifier circuit and measure the output voltage with a DVM, the reading will tell you the amplifier's offset. The DVM is averaging the displayed results (using an ADC!), however, so it will not reveal the noise at the amplifier's output. Measuring the noise requires a scope or a spectrum analyzer.

Like any other component in the signal chain, ADCs have their own thermal noise contribution. Thus, if you want to verify that the ADC is behaving as expected with no input signal, you need to capture a block of data and average it, just as the DVM does with the amplifier circuit. High-speed ADCs typically float to their midscale code, plus or minus any inherent offset, so the resulting average output code should be within the ADC's offset specification. While analyzing the block of data you've captured, you can easily verify the noise performance of the ADC. The data sheet specification is *input referred noise*, which is specified in LSB rms. The measurement is known as a *grounded input histogram* test, although the nomenclature stems from early converters that had a bipolar input range centered on ground, so shorting the input to ground was equivalent to having no input signal. Modern high-speed converters typically operate with a single supply, so their input common-mode voltage is the midpoint of the front-end analog power supply rather than ground. Luckily for you, the histogram test is performed by capturing a block of data with no input signal, which you've already done. Instead of averaging the captured output data, however, make a histogram instead. A typical high-speed ADC might specify 1-LSB rms input noise, so you might see a Gaussian distribution with the offset ±3 codes. The input referred noise is calculated as the standard deviation of the captured data.

So, getting back to the original question, the wideband noise contributed by the ADC will cause the outputs to switch, even with no input signal. Good luck with the rest of your debug.

To Learn More About ADC Noise
http://dn.hotims.com/41001-100

David Buchanan received a BSEE from the University of Virginia in 1987. Employed in marketing and applications engineering roles by Analog Devices, Adaptec, and STMicroelectronics, he has experience with a variety of high-performance analog semiconductor products. He is currently a senior applications engineer with ADI's High Speed Converters product line in Greensboro, North Carolina.

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

For Analog Devices' Technical Support, Call 800-AnalogD

Innovative IC lets you know when lightning will strike

Like most other people, engineers know that lightning can damage systems, circuits, and humans. Knowing when lightning is approaching and how far away it is would be helpful, even if you've already used good design practice, carefully placed suppression devices, and common sense. To date, such knowledge has required relatively large, costly, boxlike instruments. Austriamicrosystems plans to radically change that situation with its AS3935 Franklin lightning-sensor IC.

The company claims that this tiny, low-power IC is the first to perform this function. The device requires a user interface, a few small passives, and a microcontroller for management. It enables lightning-detector key fobs; personal and portable warning instruments, such as global-positioning systems; "fish finders" and golf accessories; and larger installations in networks, smart grids, and power systems.

Unlike other detectors, which sometimes give false alarms, the device uses a magnetic-field sensing coil as part of its RF receiver to detect the electrical emissions from lightning that is 1 to 40 km away. It assesses using signal strength and by analyzing the waveform using proprietary, integrated algorithms employing standard meteorological-survey patterns and data. As a result, local disturbances, such as microwave ovens, nearby motors, or similar transient sources, cannot fool the device, according to Johnsy Varghese, marketing manager at the company. Such false alarms frequently cause people to turn off similar detectors.

The AS3935 Franklin lightning-sensor IC uses an analog front end and embedded algorithms to sense the approach and proximity of lightning.

The 4×4-mm, 2.4 to 5.5V AS3935 is a power miser, too, with drain of just 60 µA in listening mode, in which it spends most of its time. It needs seven passive components and has I²C and SPI ports for interfacing to the support microcontroller. The AS3935 embeds all signal-analysis algorithms and requires no processing effort on the microcontroller's part. Further, an automatic antenna-tuning mode reduces the tolerance constraints on the front-end components and eases manufacturing concerns.

The vendor has done field trials in Finland and at the Florida Institute of Technology, a leading center for lightning research, and tests are now under way in South Africa. The company compared field data with data from the National Lightning Detection Network and has designed and tested the devices for both cloud-to-ground and cloud-to-cloud lightning. Many other instruments detect only cloud-to-ground lightning.

The AS3935 comes in a 16-lead MLPQ package and sells for $3.55 (1000). An evaluation kit comprising an evaluation module, a lightning-emulation unit, a thumb drive with a graphical user interface, and hardware- and software-design files, is also available.

—by Bill Schweber

▷**Austriamicrosystems**, www.austriamicrosystems.com/Lightning-Sensor/AS3935.

Medical wall-mount power supplies provide as much as 15W

Power Sources Unlimited Inc has announced the TR15RAM series of ac/dc medical-grade wall-mount switching power supplies. The series provides as much as 15W of continuous output power and meets EN60601-1/UL60601 third edition and EN55011, FCC Class B emissions limits. The TR15RAM series also meets Energy Star 2.0 Efficiency Level V requirements with efficiency ratings as high as 82%, as well as the requirements of global energy-efficiency standards, such as EISA2007, CEC, and the European ERP Directive.

Specifications include a universal ac input of 90 to 264V ac at 47 to 63 Hz and output voltages of 5 to 24V dc. Other features include interchangeable ac plugs, continuous short-circuit protection, overvoltage protection, and a variety of dc-output plugs, cable lengths, and types. Prices start at less than $19 (OEM quantities).

—by Steve Taranovich

▷**Power Sources Unlimited**, www.psui.com.

DILBERT By Scott Adams

Compact LCOS pico projector cuts costs

Microdisplays, including pico projectors, could be the next big things in consumer electronics. A couple of examples of possible incarnations for tiny, LED-based handheld projectors include Micron's $99 Pop Video for the iPhone, iPod, and iTouch. Approximately the same size as the iPhone, the device allows users to share photos and videos by turning any flat, plain surface into a large display and giving the technology a social dimension, according to proponents.

Commercial ventures such as stores and restaurants also benefit from turning any flat, plain surface into a display, which technologies such as the Light Blue Optics Light Touch projector make possible.

Several years ago, seeing the potential for pico projectors, Micron purchased Displaytech, a developer of LCOS (liquid-crystal-on-silicon)-display technology. Micron combines the technology with high-brightness LEDs to make its qHD optical engine. Intersil, meanwhile, combines the qHD engine with its new fully integrated ISL97901 buck/boost RGGB (red/green/green/blue) LED driver for pico-projector-display applications. The optical-engine/driver-IC Pico-qHD fits on a 1.7×2 in. PCB. Intersil claims that the $2.50 (1000) Pico-qHD is the least-expensive and the smallest production-ready LED-LCOS pico projector.

Intersil's Pico-qHD reference design includes an LED driver, a video processor, a power-management IC, a buck-boost converter, and a battery charger.

Intersil is offering an LED-LCOS-pico-projector reference design that includes the ISL97901 buck-boost LED driver, the TW8835 LCD video processor, the ISL9307 power-management IC, the ISL9110 buck-boost converter, and the ISL9230 battery charger. The optional ISL58333 optoelectrical-IC sensor for automatic, real-time white-balance control maintains image quality and reduces optical-calibration cost.

Prices for the reference design start at $13,995, including technical support and a full hardware kit with design files and firmware. According to Intersil officials, Intersil's tech support and code more than justify the price.

—by Margery Conner

▷**Intersil**, www.intersil.com.

Stereo-vision-camera reference design uses TI's OMAP

Embedded-design-services company E-con Systems has announced a stereo-vision-camera reference design that it based on Texas Instruments' OMAP (Open Multimedia Applications Platform)/DM37x family of processors and Gumstix's Overo computer-on-module series. Code-named Capella, the reference design targets companies that want to integrate stereo vision into their products.

The product features 3-D analytical applications, such as machine vision; mobile robotics, including obstacle detection; and automotive applications, including computer-aided navigation and driver guidance. It also targets 3-D object-reconstruction applications, such as details of sculptures, monuments, and faces; 3-D video recording and playback; and 2-D synchronous-dual-camera applications requiring two simultaneously captured images of the same or different objects.

The reference design includes a DM3730-based Gumstix Overo WaterStorm computer on module, a Gumstix Tobi expansion board and E-con's camera daughterboard. It also comes with a software-development kit with full source code for application development on V4L2 or OpenCV. E-con also provides sample applications with source code for stereo-image capture and depth measurement.


The Capella reference design targets companies that want to integrate stereo vision into their products. It features 3-D analytical applications, such as machine vision.

Capella runs Linux 2.6.35 and has Ethernet, an HDMI video port, a USB OTG port, and a serial interface for a Linux console. The stereo-engine driver exposes three V4L2 interfaces to capture left, right, and combined left/right video streams. The software-development kit contains sample applications using both the V4L2 interface and the OpenCV stack. The Overo module comes with all of the necessary software with V4L2 stereo-camera drivers and the OpenCV library.

A pair of 1/3-in. global shutter monochrome Aptina CMOS image-sensor MT9V024s power the unit's daughterboard. The sensors have an active area of 752×480 pixels and can deliver synchronous parallel 736×480-pixel monochrome video of 8-bit/pixel gray-scale images at 30 frames/sec. The cameras come with prealigned and precalibrated M12 lenses on an S-mount lens holder, and customers can select and calibrate the lens for their requirements. The two cameras are separated by 100 mm. This baseline allows accurate measurement to a depth of 3m. E-con will customize the product for customers who want to increase the baseline or change the lenses.

—by Steve Taranovich

▷**E-con Systems**, www.e-consystems.com.

Anticipate ___Accelerate ___Achieve

Infiniium 90000 Q-Series

Sound-expanding IC yields audio analog of 3-D for smartphones, tablets

Small, portable sound sources, such as smartphones and tablets, provide a constrained in-depth audio sensation. The LM48903 from Texas Instruments strives to alleviate that constraint by providing a more immersive audio sensation, even from the headphones or closely spaced speakers that users often employ with these devices. The IC manipulates the audio tracks to provide a spatially expanded, 3-D-like sound impression, with user-programmable parameters and configuration.

The IC also integrates a pair of 18-bit stereo ADCs and matching Class D amplifiers, which can continuously drive as much as 2W per channel into 4Ω loads with total harmonic distortion plus noise of less than 1%, along with I^2S and I^2C digital interfaces. A demo of the LM48903 shows an impressive selection of movies on a tablet computer; equally impressive audio comes from small external speakers less than a few inches apart. The demo uses external speakers, but TI intends that customers will package this IC within a tablet or a smartphone as part of an internal audio-output path.

The LM48903 spatial processor creates a more immersive experience from flat, constrained audio tracks, using the available headphone or speakers.

The vendor provides software tools so that OEMs designing the IC into their products can adapt it without developing or tuning algorithms. A Web-based speaker-array-designer-generator tool creates the spatial-audio coefficients for desired effects. An Android driver and evaluation board are also available. You can see a video explaining how spatial audio works at http://bit.ly/HU7ycm, or visit http://bit.ly/HU7Ccd for more information. The LM48903 is available in a 30-bump, 2.7×3.2-mm micro-SMD package and sells for at $1.75 (1000).

—by Bill Schweber

▷**Texas Instruments**, www.ti.com.

Single-chip 10-GbE processor uses massive parallelism

Broadcom Corp recently announced its fourth-generation Ethernet processor, which it claims is the first chip to use massive parallelism by virtue of its 64 packet-processing cores running at 1 GHz. Providing full-duplex, 100-Gbps performance, it can also provide a dozen 10-Gbps channels. "By 2015, there will be twice as many devices connected to the Internet as there are people in the world," says Dan Harding, senior director of marketing, infrastructure, and networking at Broadcom. Many of those people will be streaming video. "As a result of this increasing demand for bandwidth, the core of the network is going to need upgrading to 100-GbE over the next four years," he adds.

According to Broadcom, by the end of 2012, the number of Internet-connected devices will exceed 7 billion. Over the next four years, most of the content that users access from mobile devices will be high-bandwidth streaming video. Further, application downloads will balloon to 47 billion per year. To meet this demand, Internet service providers are quickly adopting 100-GbE, which should grow at a rate of 170% over the next five years, according to Infonetics Research Inc.

Broadcom claims that it has addressed the need for more bandwidth with higher levels of integration, which enable it to reduce by 80% the power and area of its fourth-generation network processor. With its multithreading support and specialized accelerators, the new network-processor technology can offload many tasks that previously required external FPGAs. This device's seven cores handle tasks such as algorithmic look-up and packet generation, according to Nicolas Tausanovitch, senior product-line manager of infrastructure and networking at the company. By taking on tasks that external FPGAs and expensive SRAM previously had to perform, the chip decreases the complexity and bill-of-materials cost for line cards.

Using 40-nm design rules for its array of 64 packet processors, the BCM88030 also includes seven on-chip accelerators for common functions, including a programmable algorithmic look-up engine for massive Internet Protocol Version 6 tables using low-cost DDR-3 DRAM, an algorithmic-access-control list using Broadcom's proprietary knowledge-based processor, and a high-speed packet parser and classifier.

—by R Colin Johnson

▶**Broadcom**, www.broadcom.com.

VOICES

Nathan Seidle: open source, open mind, open for business

As an undergraduate electrical-engineering student in 2003, Nathan Seidle founded SparkFun Inc, an online retail store that describes itself as selling "the bits and pieces to make electronics projects possible." Since then, the company has expanded to more than 450 original products, which its growing and varied customer base of EEs, hobbyists, artists, students, do-it-yourselfers, and others is snatching up by the shopping cartful. Seidle recently spoke to *EDN* about SparkFun's approach to open-source hardware; the "maker" movement; and ways to spread engineering education, including the company's new learn.sparkfun.com site. Excerpts of that conversation follow. For the full interview, go to www.edn.com/120524pa.

What's your definition of a "maker"?

A It's the whole do-it-yourself movement—the person who ceases to talk about something and the person who actually does something [or] anyone actually using their hands. Whenever I go to MakerFaire, the folks that have no real formal training and let their creativity run are the people doing some of the most interesting things.

Are we seeing a resurgence of the maker movement?

A Yes. We make very sure at SparkFun, as it pertains to the maker movement, that we separate ourselves from the trend. We try to be the tool maker for all the makers out there. Let us provide you with the picks and the shovels so you can go build your great art piece or your master plan.

How does open-source hardware fit into all of this?

A SparkFun ended up playing in the open-source-hardware world in a roundabout [way]. We were sharing our schematics, data sheets, and how we did stuff from the beginning because we figured it would enable our customers to be more successful at their projects. Over time, we realized that, if we shared even more files—for instance, the PCB-layout files or the firmware or whatever with a given product—customers could take our products and sort of mash them up and come up with their own product or project. That enables all of our customers to be more successful. It can also be a bad thing because our competitors can see our firmware, and they can duplicate our PCBs. But they could do that anyway. We think it's important to use open-source hardware because it enables all of our customers rather than temporarily slowing down a few of our competitors.

How does open source fit into the company's overall business strategy?

A It's something that I believe is an alternative to intellectual property. At the same time, SparkFun uses it as an interesting business lever or push. By open-sourcing our products, we are enabling our competition to copy us. Right now, we come out with a product, and, within 12 to 14 weeks, our competition has copied it. That [fact] means that we have 12 to 14 weeks to come up with the next revision or the next innovation or the next whatever it is to stay ahead of the competition. We use open source as a business motivator, as a driver. We use it to keep sharp.

Less than 10 years after launching SparkFun from your bedroom, the company now has more than 130 employees, a 50,000-sq-foot facility, and more than 450 original products. What's next?

A We are either a big fish in a small pond or a tiny fish in a gigantic pond. Compared with the Digi-Keys of the world, we are a speck. Compared with *Make* magazine or Adafruit, we're good competitors, all of similar size.

The next big thing for SparkFun? I would love to show folks that they may enjoy electronics without knowing it's electronics. We have been pretty successful in lowering the barrier to entry to tinkering, hobby-electronics making. We have made it easier for folks to play with stuff. I had to wait until my sophomore or junior year of college to ever play with a microcontroller. What would happen if we allowed ninth and 10th graders to play with microcontrollers? How much more successful are they going to be in whatever pursuit they have later in their academic career?

You started the company in your early 20s. Has your age ever been an obstacle in building a business and career?

A Are there challenges at a young age? Definitely. At the same time I tell a lot of folks who are thinking about starting a business, if you wait until you have a family, a mortgage, and a paycheck, that's a whole lot of responsibility and a lot of risk to go off and start your own business. I was very fortunate starting it when I did because I had no responsibilities.

Any advice for young engineers just starting out?

A We are in an interesting time. If you think back to the 1990s and how a young person would go from zero to 60, from pretty much obscurity to the front page of something, really the only route was to be a professional athlete or a musician making it overnight. In the past couple of years, we have seen so many companies and so many young people break huge because they had the right friends, the right support network, and the right idea. ... It's a pretty amazing time.

—interview conducted and edited by Suzanne Deffree

BY BONNIE BAKER

EMI problems? Part four: It could be conduction

Radiated electromagnetic-interference noise can originate from an unintended source and an unintentional antenna (see **parts 1**, **2**, and **3** of this series, available at www.edn.com/120524bb). Conductive EMI noise also can originate from a radiating EMI-noise source or from board components. Once your board receives conductive noise, it resides in your application circuit's PCB traces. Common radiated EMI-noise sources include the elements that the previous articles discuss as well as onboard SMPSs (switch-mode power supplies), connection wires, and switching or clocked networks.

Conducted EMI noise is a consequence of the normal operation of switching circuits in combination with parasitic capacitances and inductances. **Figure 1** illustrates the origin of EMI-noise sources that can couple into PCB traces. Vemi1 originates with switching networks, such as clocking signals or digital-signal traces. The mode of coupling for these noise sources is through the parasitic capacitance between traces. These signals transmit current spikes into a neighboring PCB trace.

Vemi2 again originates with switching networks or from an antenna on the PCB. The coupling mode for these noise sources is through the parasitic inductance between traces. This signal transmits voltage disturbances into the neighboring PCB trace.

The third EMI-noise source is a consequence of neighboring wires in a cable. Signals traveling down these wires can produce a crosstalk effect.

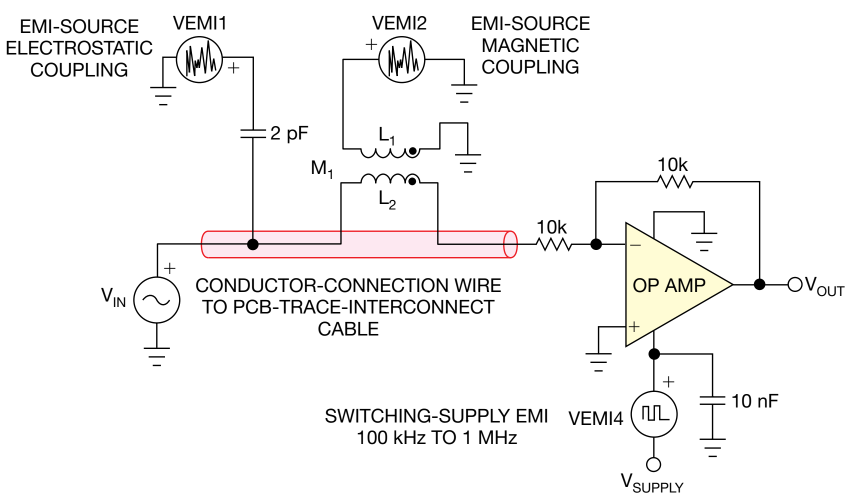

Figure 1 Vemi1 noise originates with switching networks, such as clocking signals or digital-signal traces. Vemi2 noise originates with switching networks or from an antenna on the PCB. The third EMI-noise source is a consequence of neighboring wires in a cable. Switch-mode power supplies generate Vemi4 noise.

SMPSs generate Vemi4. The noise from the SMPS rides on the power-supply traces and appears as Vemi4 signals. During normal operation, SMPS circuits create conducted EMI opportunities. The on and off switches in these power supplies generate large, discontinuous currents. These discontinuities are present at the input of buck converters, the output of boost converters, and at both the input and the output ports of flyback and buck-boost topologies. The discontinuous currents from the switching action generate a voltage ripple, which travels through the PCB traces to other parts of the system. Input-voltage ripple, output-voltage ripple, or both from an SMPS can compromise the operation of the load circuit.

Figure 2, available online at www.edn.com/120524bb, shows an example of the frequency content of a dc/dc buck SMPS input running at 2 MHz. The fundamental frequency content of SMPS-conducted noise is 90 to 100 MHz.

Conducted emissions are either differential-mode or common-mode interference. A differential-mode-interference signal appears between the circuit's input terminals, such as signal and ground. The current flows through both inputs with the same phase. However, the direction of one current input is equal to and opposite from the direction of the other current input. The load at the end of these two inputs creates a voltage that changes with the current magnitude. This change in voltage between line one and the differential reference creates noise or communication errors in the system.

Common-mode interference occurs as you add a ground loop or an undesirable current path to the circuit. If an interfering source exists, common-mode currents and a common-mode voltage develop on both lines, with the ground loop acting as the common-mode reference. Both differential-mode and common-mode interference require filters to combat the detrimental effects of EMI.

In next month's column, we will address circuit approaches to solving EMI-noise problems. EDN


FIVE
YEARS
OUT.
AVIATION: More comfort, more safety, more innovation and of course, more electronics. Are you Five Years Out? Then you're probably working with us.
ARROW
Five Years Out

PRY FURTHER

Go to www.edn.com/pryingeyes for more Teardowns.

Inside the Schick Hydro microcontroller-powered wet razor

Microcontrollers are everywhere, including the oddest of places, such as a razor—the Schick Hydro Power Select customizable power razor, to be precise. I'm not sure why you'd need microcontroller-controlled customizable power or the packaging, waterproofing, and power issues that go along with it, but I gave it a shot. After all, Schick promises it "allows men to interact with their razor in a new way." I never realized that I longed to interact with my razor in new ways, but Schick has created a demand—a craving, if you will—and I just had to try it. Alas! It wasn't all that much better than interacting with my regular arm-powered, hand-controlled razor, so I instead took it apart. It's a bit more interesting to go inside and see how the designers implemented a wet razor, powered from one AAA battery and featuring three vibration settings, easy-to-read indicators, and a haptic sensor. Why they would do it is another question.

Removing the cap exposes the AAA battery. These batteries have a peak voltage of 1.8V but drift down to 0.9V over time.

The battery holder attaches firmly to the PCB using a solder joint and clips.

The internals, with plastic and rubber sealing removed, expose the unbalanced motor that vibrates the razor. Unlike a toothbrush motor with an actuator that rotates the heads, this motor doesn't connect to the razor blades. It simply causes the razor to vibrate. If anyone has found this vibrating feature on a razor to be useful, please let me know. I may be missing something.

A Microchip MCP1624 low-voltage input regulator boosts the battery's voltage to the 2V that the Microchip PIC10F222 microcontroller requires. The regulator maintains that 2V over the life of the battery, even as it drifts over time from 1.8 to 0.9V.

A vibration-stage indicator uses LEDs.

At the heart of the Hydro Power is the PIC10F222 microcontroller, again from Microchip. This six-pin, 8-bit device has 768 bytes of flash memory that sells for 40 to 79 cents (www.datasheets.com). It controls almost everything, including the motor, through an on-chip PWM controller.

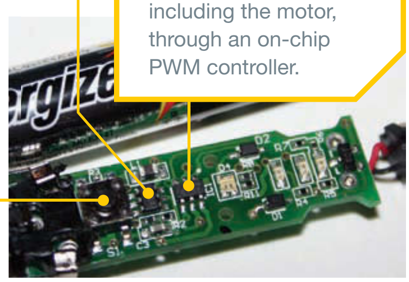

Batteries are inserted here.

A three-stage power button sets the vibration. The razor rotates through all three settings and then goes into off mode. It required too much pressure to activate due to its sealing cap.

The razor has a three-stage power switch.

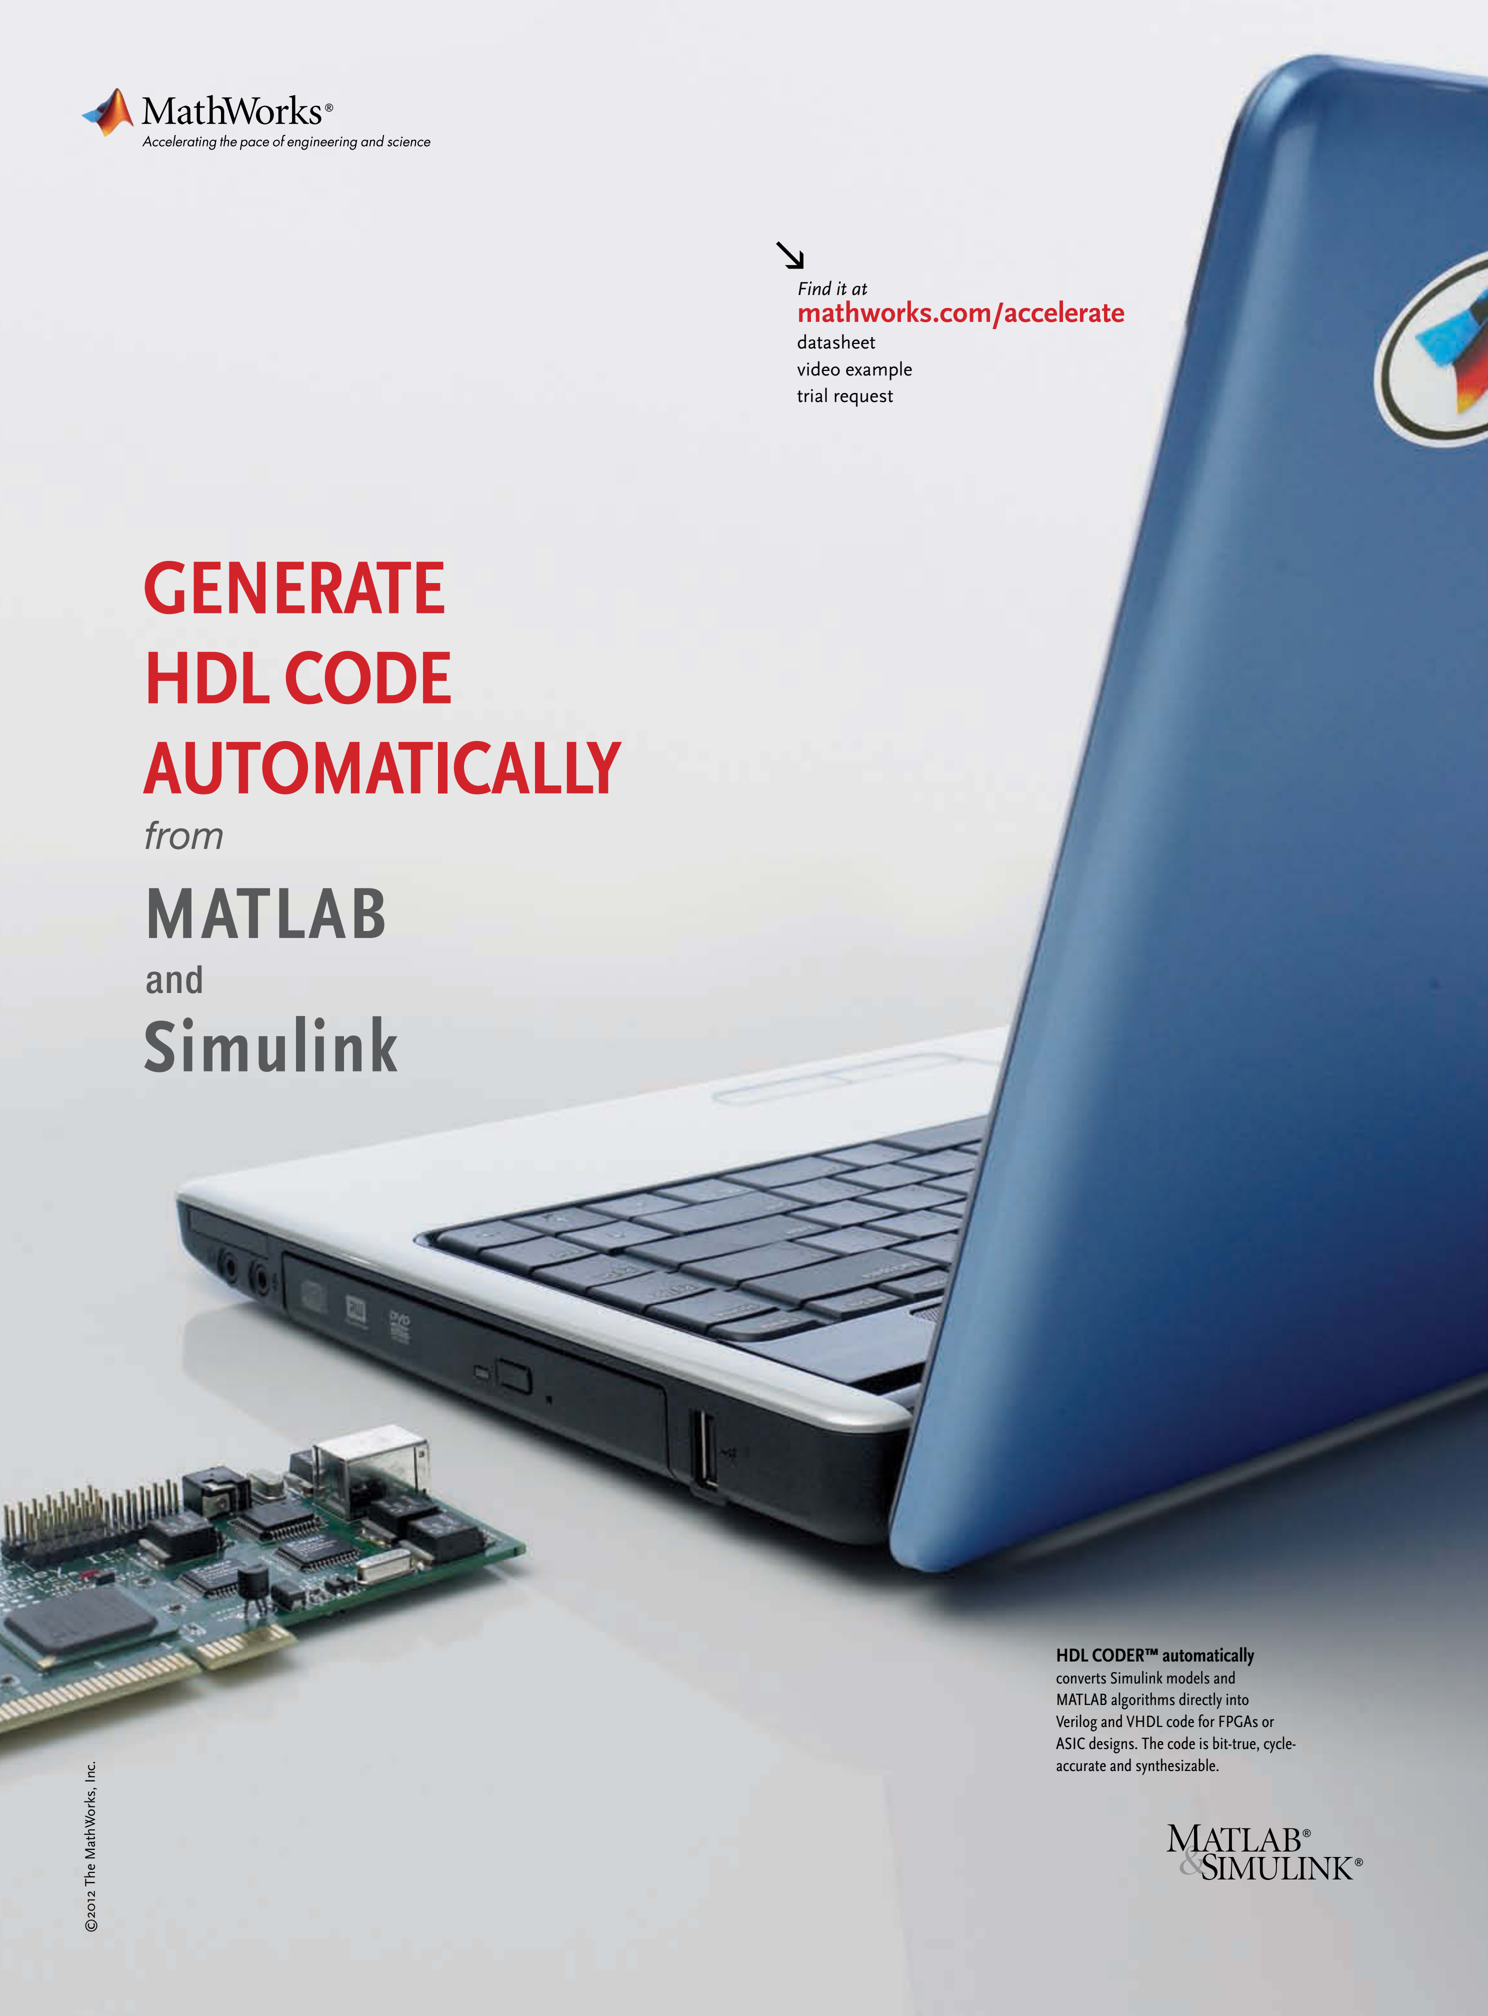
MathWorks®
Accelerating the pace of engineering and science
Find it at
mathworks.com/accelerate
datasheet
video example
trial request
GENERATE
HDL CODE
AUTOMATICALLY
from
MATLAB
and
Simulink
HDL CODER™ automatically
converts Simulink models and
MATLAB algorithms directly into
Verilog and VHDL code for FPGAs or
ASIC designs. The code is bit-true, cycle-
accurate and synthesizable.
MATLAB® & SIMULINK®
©2012 The MathWorks, Inc.

UBM Electronics

DesignNews EDN

MECHATRONICS IN DESIGN

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

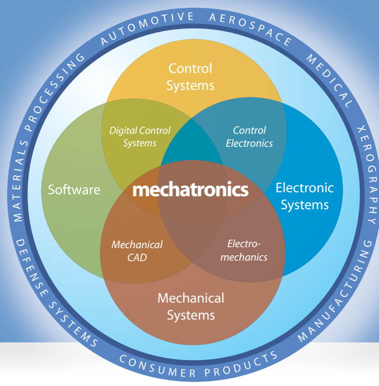

Handling time delays

As engineering systems become more complex, understanding time delays is essential.

We experience it every morning as we struggle to find the right water temperature in the shower. Time delays are everywhere! They arise in engineering, biology, physics, economics, and the environment. As engineering systems become more complex, multiple sensors, actuators, and controllers introduce multiple delays, particularly in interconnected and distributed systems. In a dynamic system, changes do not happen instantaneously, so applying an otherwise-correct control decision at the wrong time could result in catastrophe. From your first exposure to feedback-control systems, you learn to always conserve phase. What is the relationship between time delay and phase, and how do they relate to the stability of a feedback-control system? Does time delay always degrade the performance of a feedback-control system?

Time delays arise in control systems from delays in the process itself, represented by a time constant, τ, and a natural frequency, ω_N, in transfer functions; from delays in the processing of sensed signals; and from delays in the implementation of a digital-control system as a result of sample-and-hold, calculation, and velocity estimation, in which the total time delay can be one to two times the sample period.

The input-output time delay shifts the signal in time (**Figure 1**). The frequency response of a time delay is exact, with a magnitude of one and a phase angle that decreases linearly with frequency. A greater time delay corresponds to a more rapid increase of phase lag with frequency.

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

Why does time delay, and correspondingly phase lag, almost always cause a system to become unstable? Physically, an imbalance between the strength of the corrective action and the system's dynamic lags results in the application of corrective action in the wrong direction. Mathematically, when the denominator of the closed-loop transfer function equals zero, the system becomes unstable. Because this denominator is one plus the open-loop transfer function, when the open-loop transfer function equals −1—that is, the magnitude is one and the phase angle is −180°—the closed-loop system is marginally stable.

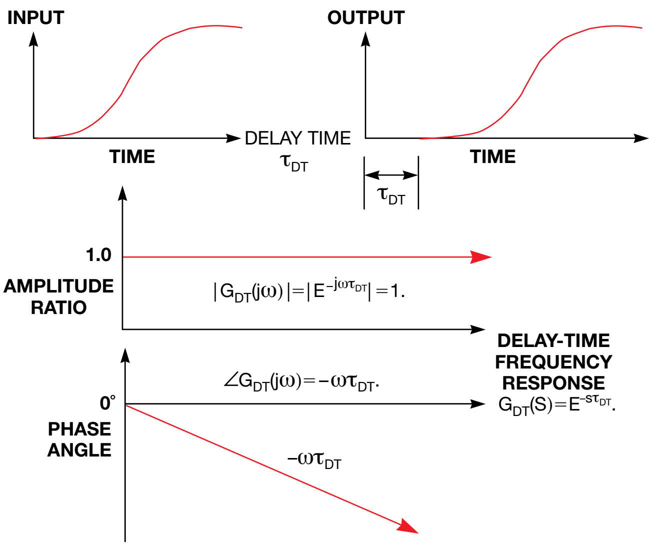

Figure 1 The input-output time delay shifts the signal in time.

How close is a stable system to becoming unstable? Because of model uncertainties, it is not merely sufficient for a system to be stable; it must have adequate stability margins. Stable systems with low stability margins work only on paper. Classic control quantifies uncertainty in a way that assumes that either gain changes or phase changes occur. The tolerances of gain or phase uncertainty are the gain margin and phase margin.

The presence of delays may be either beneficial or detrimental to the operation of a dynamic system. Judicious introduction of a delay may stabilize an otherwise-unstable system, using, for example, a wait-and-act control strategy, or may improve steady-state tracking error. The impact of delays continues to grow in many fields, including the control of distributed systems, such as energy and computing grids. EDN

USB 3.0: BRINGING SUPERSPEED CONNECTIVITY TO MOBILE DEVICES

WITH THE RISE OF RICH MULTIMEDIA CONTENT ON EVERYTHING MOBILE, USERS CAN LEVERAGE USB 3.0'S ULTRAHIGH SPEED FOR MOVING LARGE AMOUNTS OF COMMERCIAL AND USER-GENERATED CONTENT.

BY VIKAS DHURKA AND STEVEN CHEN • CYPRESS SEMICONDUCTOR CORP

As mobile handheld devices continue to become essential parts of our lives, developers are inventing new applications every day to support our on-the-go lifestyles. Smartphones and tablets, for example, are increasingly finding use in media playback. Improving the consumption and transport of content, however, has proved to be one of the most challenging tasks. As a result, improved storage capacity and data-transfer rates are both crucial in mobile-system designs. Although the storage capacity of mobile devices doubles every year to keep up with increasing content sizes, data-transfer rates remain frustratingly low.

IMAGE: SHUTTERSTOCK

THE EMERGENCE OF USB 3.0

According to In-Stat, manufacturers in 2015 will ship more than 280 million mobile phones with USB 3.0 capability. USB 3.0 brings a tenfold bandwidth improvement over USB 2.0, achieving a raw throughput of 5 Gbps, or 600 Mbytes/sec—nearly 480 Mbytes/sec of actual data throughput after accounting for the 20% protocol overhead. Because USB 3.0 is data-agnostic, it can stream any type of content, including high-definition video. USB 3.0 has also raised the bar in power provisioning and management. It allows a device to draw as much as 900 mA from a PC during operation and still maintain plug-and-play capabilities. With several power-saving modes, such as idle, sleep, and suspend, USB 3.0 optimizes the operating life of battery-powered mobile devices.

With more than 6 billion devices enabled and nearly 100% penetration in mobile handheld devices, USB is now the most established interface. Because USB 3.0 is backward-compatible with USB 2.0, consumers should have no trouble understanding and using it.

USB 3.0 is equally familiar to both consumers and tablet designers. Two design philosophies are spearheading the shift toward mobile computing in the tablet form factor. Conventional-laptop manufacturers, including Asus, Acer, and Toshiba, opt for top-down miniaturization, whereas mobile-handset makers, including Apple, Samsung, and HTC, are taking a bottom-up approach. As USB 2.0 has held a firm presence in both fronts, tablet designers with a background in either laptops or handsets will find moving to USB 3.0 a straightforward process.

CONNECTOR CONCERNS

Although the standard A connector for USB 3.0 has the same dimensions as for USB 2.0, there are concerns with the size of the USB 3.0 microconnector for phones, tablets, and other handheld devices. The USB 3.0 microconnector has a unique configuration, where the USB 3.0 lines run alongside the USB 2.0 connector (**Figure 1**). This connector has a larger profile than that of the USB 2.0 connector, which causes a concern for mobile-device manufacturers striving for an ultracompact form factor. Manufacturers are wary of adding yet another bulky connector to their systems. The USB 3.0 microconnector has a total width of 12.85 mm (**Figure 2**).

AT A GLANCE

- The increase in multimedia-content consumption on mobile devices calls for higher storage capacity and faster connectivity.
- USB 3.0's 5-Gbps raw signal rate can handle high-capacity multimedia-content transfers and high-definition video streaming.
- The USB 3.0 microconnector is smaller than the USB 2.0 and HDMI microconnectors together, saving real estate on mobile-system designs.
- Although Wi-Fi offers wireless convenience, USB 3.0 provides battery-charging capabilities and ultrahigh data rates, making it the most versatile connectivity choice available.

In addition to a USB 2.0 microconnector, most high-definition handhelds today have an HDMI (high-definition-multimedia-interface) microconnector for streaming high-definition video. Manufacturers can, however, omit micro HDMI and instead stream high-definition video over an MHL (mobile high-definition link), a new proposed mobile video/audio interface for directly connecting mobile devices to high-definition TVs and external displays. Because MHL is connection-agnostic, USB 3.0 can instead carry its audiovisual signals, avoiding the need for another HDMI microconnector.

The USB 3.0 microconnector is smaller than USB and HDMI microconnectors together. The HDMI microconnector measures 6.4 mm long, and the USB 2.0 microconnector measures 7.8 mm long. Taking the separation between these two connectors into account, the result is a total connector length of about 15 to 16 mm, or about 25% larger than a USB 3.0 microconnector. Thus, mobile-device manufacturers can produce elegant designs using one multipurpose USB 3.0 microconnector that enables faster content data transfer, rapid battery charging, and high-definition video streaming over MHL.

EXTENDED THROUGHPUT

Storage in most mobile devices employs either SD (secure-digital) card or eMMC (embedded-multimedia-card) flash technology. The SD 2.0 standard supports a maximum of 25 Mbytes/sec, whereas the eMMC 4.3 standard supports a maximum of 52 Mbytes/sec. Neither can effectively use the full USB 3.0 data pipeline. This underusage, however, changes with the new SD 3.0 standard that raises SD performance to 104 Mbytes/sec. Similarly, the latest eMMC 4.41 specifications will achieve a 104-Mbyte/sec upgrade. SD 3.0 cards are available, with micro-SD versions coming soon; eMMC 4.41-compliant flash devices are in full production.

USB 2.0 is barely supporting the current storage standards (**Figure 3**). With the UHS (ultra-high-speed) II- and UFS (universal-flash-storage)-based devices poised to hit the market by the end of 2012, USB 2.0 will soon become the rate-limiting link in the overall data-transfer pipeline. Mobile-system manufacturers must adopt USB 3.0 into their new designs to take full advantage of these storage improvements.

The amount of digital content on mobile devices will continue to increase with the proliferation of high-definition video-recording features. A 10-minute

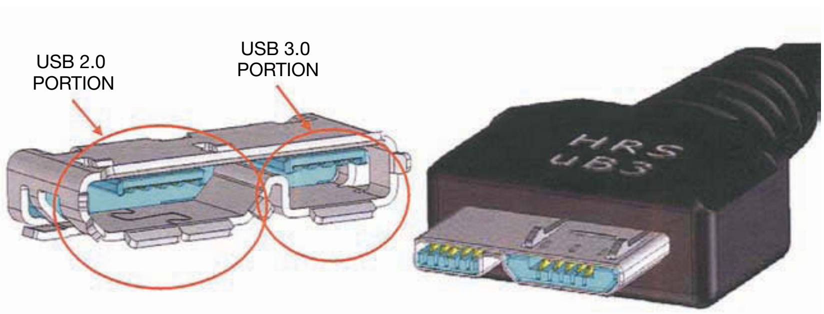

Figure 1 In the USB 3.0 microconnector, the USB 3.0 lines run alongside the USB 2.0 connector. This connector has a larger profile than that of the USB 2.0 connector, which causes a concern for mobile-device manufacturers.

home movie can easily take up multiple gigabytes of storage. A common usage scenario today is for users to transfer these movies to a PC for editing, sharing, and playback. Although this scenario would be a considerable task for USB 2.0, USB 3.0 SuperSpeed offers instant gratification with less waiting.

With so much data to store, handset and tablet vendors are looking into employing RAID (redundant-array-of-independent-disks) technology to increase data throughput. Using a RAID-0 configuration, large data are striped and stored simultaneously in two SD or eMMC devices. RAID-0 works with USB 3.0 by effectively doubling the maximum throughput of storage devices and using more of USB 3.0's available bandwidth. Today, the transfer of an 8-Gbyte movie over the fastest USB 2.0-enabled phone requires more than seven minutes at 18 Mbytes/sec. Upgrading to USB 3.0 broadens the data-transfer pipeline and allows data transfers to occur at peak current-storage limitations, meaning that an internal bridge supporting a USB 3.0 port and two storage devices in a RAID-0 configuration can operate at the maximum performance of the storage device. This rate is now approximately 150 to 200 Mbytes/sec. At this rate, the transfer time of an 8-Gbyte movie decreases to 41 sec (**figures 4 and 5**).

USB 3.0 can also help to reduce the manufacturing cost of devices. Many handheld-device vendors use USB during factory manufacture to preload apps, music, movies, and other content onto devices. Smartphone operating systems and preloaded content larger than 2 Gbytes allow manufacturers to preload phones with USB 3.0 at 10 times USB 2.0 speed, translating to a more efficient process and greater cost savings.

SYNC AND GO

In the era of cloud computing, an abundance of cloud services, such as iCloud, Microsoft SkyDrive, and Dropbox, has emerged to allow users to wirelessly synchronize their data without going through a PC. The need for wired technology in mobile devices is up for debate. Wi-Fi is the most popular wireless standard today, with a 93% attachment rate on smartphones and tablets, according to IMS Research. Its untethered convenience has brought countless benefits to our daily life, including surfing the Web at a local Starbucks, sending e-mail just before catching a flight, or taking a picture on one device and having it appear on another through the cloud.

In its latest revision, Wi-Fi 802.11n supports MIMO (multiple-input/multiple-output), in which multiple antennas coherently resolve more data than a single antenna can. Wi-Fi 802.11n also provides a 40-MHz-wide channel for increased throughput and can operate in the less-congested 5-GHz ISM (industrial/scientific/medical) band. Although 802.11n's theoretical bandwidth of 600 Mbps is still no match for the 5-Gbps bandwidth of USB 3.0, it is good enough for moving content between mobile devices and PCs.

Wi-Fi 802.11n's benefits are diminished in mobile designs, however. The maximum 600-Mbps bandwidth typically occurs only under ideal operating conditions in which the 40-MHz-wide channel uses four spatial streams. Because each spatial stream requires an antenna and an ADC, it is not cost-effective for mobile-system vendors to implement all four sets. To be effective, each antenna also needs distant placement; the compact mobile form factor severely constrains this requirement. Furthermore, using a 40-MHz-wide channel is often impractical in the congested 2.4-GHz ISM band, which is already full of Bluetooth devices, microwave ovens, and other common RF devices. As a result, the practical Wi-Fi bandwidth on any mobile device decreases significantly.

Wi-Fi also lacks the battery-charging benefit of USB 3.0 devices. Digital content transfers wirelessly between mobile devices and PCs, concurrently draining the device's battery unless you connect it to a wall charger or a PC using USB (**Figure 6**). Thus, mobile devices can never achieve full "cordless" synchronization using Wi-Fi unless the amount of transferred digital content is negligible. Wi-Fi is ideal for surfing the Web, sending e-mails, and listening to radios over the air; however, it is not the choice for large, fast content transfers. With USB 3.0's versatile connectivity, high bandwidth, and battery-charging capabilities, it is likely to coexist with Wi-Fi on mobile devices for the foreseeable future.

Although USB 3.0 presents a promising technology for mobile devices, it must undoubtedly overcome hurdles to reach the same level of success as USB 2.0. Because USB 3.0 operates at a higher frequency, new challenges, such as signal integrity and cable length, arise. Designers can overcome these challenges, however, by following best-design practices.

Currently, only a few USB 3.0-enabled PCs exist; thus, mobile devices cannot connect to many PCs with USB 3.0, regardless of whether the mobile devices implement it. Now that Intel's next-generation Ivy Bridge CPU with native USB 3.0 host support is out in the market, more PCs will come with standard USB 3.0 ports. In addition, early USB 3.0 adopters must develop their own USB 3.0 driver, which adds another layer of overhead and design complexity. Microsoft has announced robust USB 3.0-driver support in Windows 8 operating systems, and ubiquitous, seamless USB 3.0 support will be available throughout the market.

USB is still the most versatile connectivity option for everything mobile. USB 3.0 offers clear benefits, such as high bandwidth and battery charging, making it ideal for large data transfers and high-definition video streaming. USB 3.0 will employ the strong user awareness from USB 2.0 and develop a broad ecosystem of support from hardware to software. The time is right for mobile-system manufacturers to introduce SuperSpeed USB 3.0 in next-generation mobile devices. EDN

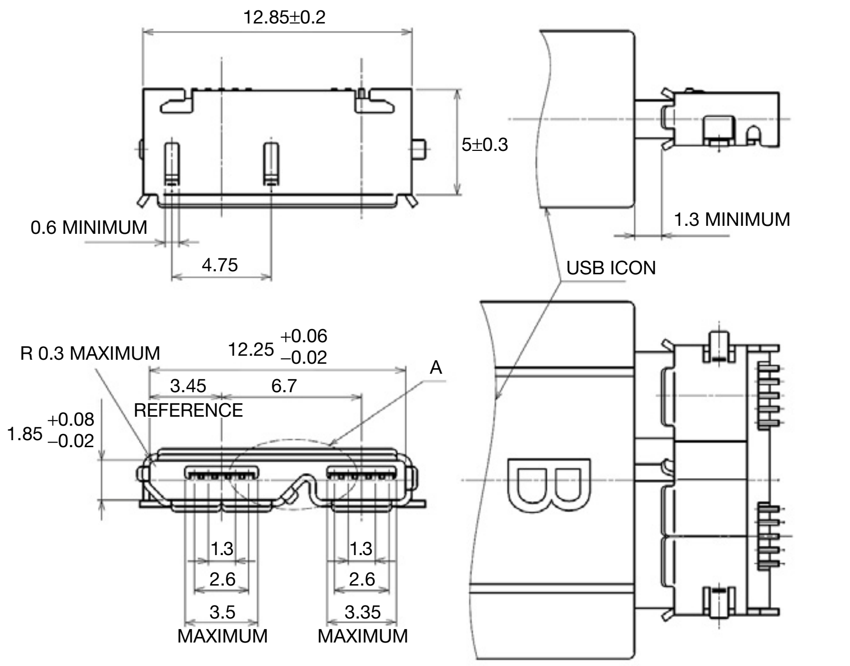

Figure 2 With a total width of 12.85 mm, this connector has a larger profile than the USB 2.0 connector (courtesy USB 3.0 specification).

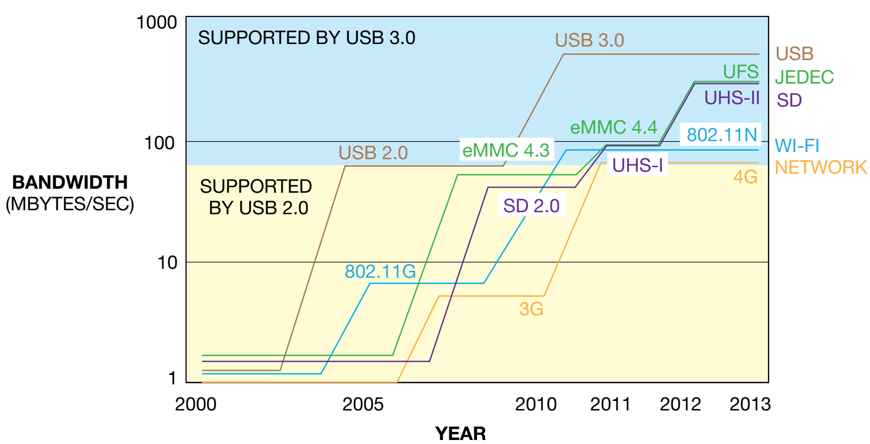

Figure 3 As part of the interface evolution, USB 2.0 is barely supporting today's storage standards.

AUTHORS' BIOGRAPHIES

Vikas Dhurka is a senior product-marketing manager at Cypress Semiconductor. Since joining Cypress in 2002, he has worked on various products and end markets and now manages the new-product portfolio for the TrueTouch family of products. Dhurka has a master's degree in electrical engineering from the University of Southern California

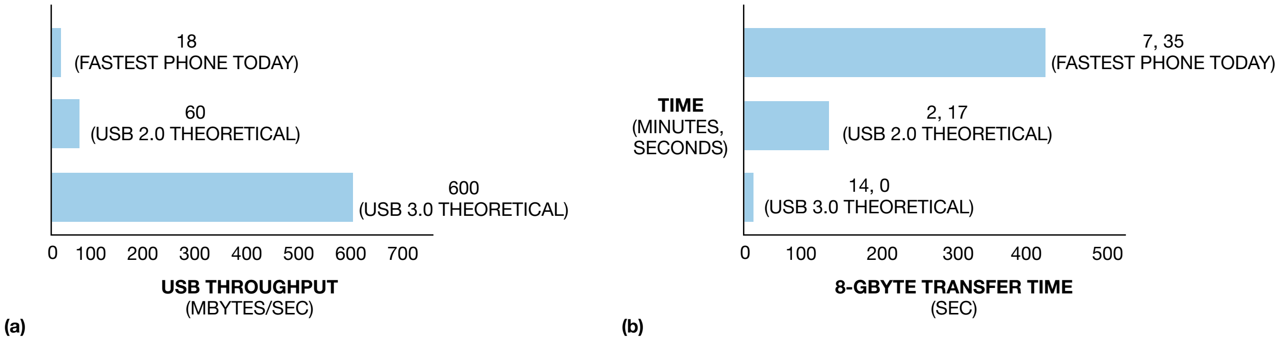

Figure 4 USB 3.0 can theoretically provide data-transfer rates as high as 600 Mbytes/sec (a), so an 8-Gbyte file theoretically takes less than 14 seconds (b). In practice, however, the speed of storage memory limits the data-transfer rate. Using internal bridge and RAID-0, the storage memory speed is 150 to 200 Mbytes/sec, so an 8-Gbyte file takes 41 sec.

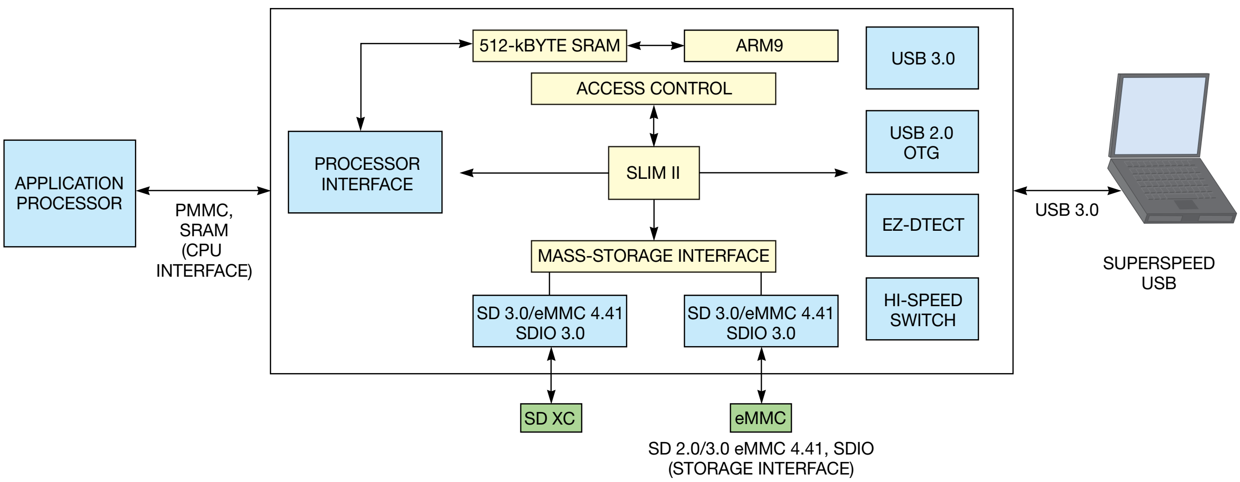

Figure 5 An internal bridge, such as Cypress' Benicia USB 3.0 storage controller, enables efficient data flow through portable devices and provides a RAID-0 configuration.

Figure 6 Wirelessly transferring digital content between mobile devices and PCs concurrently drains the device's battery unless it connects to a wall charger or a PC through USB (a). USB 3.0 has versatile connectivity, high bandwidth, and battery-charging capabilities that make it likely to coexist with Wi-Fi on mobile devices for the foreseeable future (b).

(Los Angeles) and a bachelor's degree in electronics and communication from Gujarat University (Ahmedabad, India).

Steven Chen is a product manager working in the data-communications division of Cypress Semiconductor. He holds a bachelor's degree in nanotechnology engineering with a management-sciences option from the University of Waterloo (Waterloo, ON, Canada).

POWER:

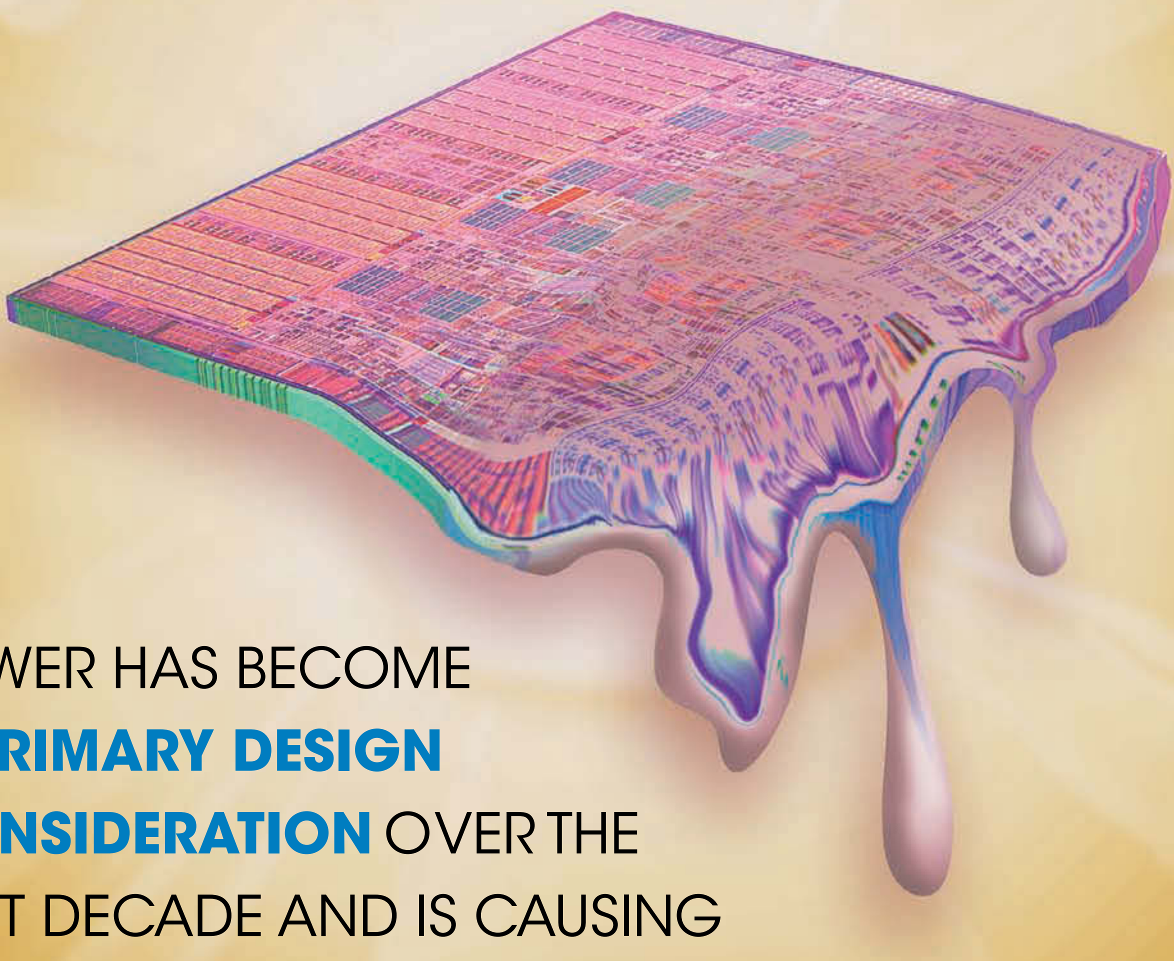

POWER HAS BECOME A **PRIMARY DESIGN CONSIDERATION** OVER THE PAST DECADE AND IS CAUSING SOME BIG CHANGES IN THE WAY THAT ENGINEERS DESIGN AND VERIFY SYSTEMS. **PHYSICS NO LONGER PROVIDES A FREE RIDE.**

A SIGNIFICANT CHALLENGE IN EDA DESIGN

BY BRIAN BAILEY • CONTRIBUTING TECHNICAL EDITOR

Power is the rate at which energy is consumed—not a hot topic 10 years ago but a primary design consideration today. A system's consumption of energy creates heat, drains batteries, strains power-delivery networks, and increases costs. The rise in mobile computing initially drove the desire to reduce energy consumption, but the effects of energy consumption are now far-reaching and may cause some of the largest structural changes in the industry. This issue is important for server farms, the cloud, automobiles, chips, and ubiquitous sensor networks relying on harvested energy.

The reason for the sudden change is that physics was helping with process technologies down to 90 nm. With each increasingly smaller node, however, voltages decreased, creating a corresponding drop in power. In general, power budgets remained fixed even as developers integrated additional capabilities. With smaller geometries, voltage scaling is more difficult and is failing to keep up. As voltages approach the threshold voltage, switching times increase. To compensate, designers lowered threshold voltages, but doing so caused a significant increase in leakage and switching currents.

IMAGE: INTEL; SHUTTERSTOCK

Every stage in the design flow—from software architecture to device physics—affects power consumption. Although each team can locally optimize power consumption, no single group can create a low-power design. Conversely, any one group can destroy it. This situation is creating a new need for cooperation and cross-discipline tooling. Power issues do not stop on chip. They spread to interconnect topologies, board and system design, power controllers, and so on. Current EDA tools do not build in the concept of power, meaning that designers are adopting retrofit approaches rather than rebuilding from the ground up.

AT A GLANCE

- Although each team can locally optimize power consumption, no single group can create a low-power design. Conversely, any one group can destroy it.
- Power estimation is an accurate science. However, this idea is true only when you have a complete design and the correct set of vectors.
- Processors are generally the least energy-efficient approach to an arbitrary problem, but you can often implement them in a small area, given their multitude of functions.
- The power-delivery network should be able to sustain the load without compromising voltage integrity.

THE ROLE OF PHYSICS

The power a chip consumes is the sum of switching, or dynamic, power and passive, or leakage, power. The dynamic component of power is due to the capacitive load of a design. This component charges through a PMOS transistor whenever a net makes a transition from zero to one. The energy drawn from the power supply is equal to the capacitive load multiplied by the square of the voltage. The system stores half of this energy in the capacitor; the other half is dissipated in the transistor. For the one-to-zero transition, no additional energy is drawn from the power supply, but the charge dissipates in the NMOS transistor. Assuming that the node changes at frequency F, then dynamic power is $FC_LV_{DD}^2$, where C_L is the capacitive load and V_{DD} is the voltage. Although other forms of dynamic-power consumption exist, they are much smaller.

Reducing the voltage has a considerable effect due to the voltage-squared term. Unfortunately, performance also relates to voltage because the increased voltage causes an increase in the gate drive, $V_{GS}-V_T$, where V_{GS} is the gate-to-source voltage and V_T is the threshold voltage. In older technologies, the leakage power was insignificant. As device sizes have decreased, leakage has become more significant in a number of areas, including gate-oxide tunneling, subthreshold voltages, reverse-bias junctions, gate-induced drain, and gate current due to hot-carrier injection.

Silicon dioxide is the typical material for insulation. At low thickness levels, electrons can tunnel across it. This relationship is exponential, meaning that halving the thickness increases the leakage by a factor of four—not an issue until transistor geometries decreased to less than 130 nm. Using high-k dielectrics instead of silicon dioxide provides similar device performance, allowing a thicker gate insulator and thus reducing this current.

Transistors have a gate-to-source threshold voltage below which the subthreshold current through the device decreases exponentially. As supply voltages decrease to reduce dynamic-power consumption, the threshold voltage also decreases, resulting in less gate-voltage swing below the threshold to turn off the device. Subthreshold conduction varies exponentially with gate voltage.

The formation of a reverse bias between diffusion regions and wells or between wells and the substrate causes small reverse-bias-junction leakages. A high-electric-field effect in the drain junction of MOS transistors causes gate-induced drain leakage, which is typically handled in the fabrication technology. Gate-current leakage is due to the drift of the threshold voltage in short-channel devices and relates to high electric fields within the device. The fabrication technology also primarily controls this effect.

Designers have made a trade-off between dynamic- and static-power consumption. The reduction in voltage has reduced the dynamic power but increased the static power. Consider a typical chip in a cell phone. When the device is operating, leakage accounts for about 10% of the consumed power; the other 90% is dynamic power. When a cell phone is in standby mode, however, which could be 90% of the time, little dynamic activity occurs in the chip. Minimizing both types of power is thus equally important.

Devices continue to show improvement in power consumption. For example, the Samsung 28-nm low-power process delivers 35% less active and standby power at the same frequency than the company's 45-nm low-power process, and the 28-nm process offers a 60%-of-active-power reduction at the same frequency compared with 45-nm low-power system-on-chip designs. Taiwan Semiconductor Manufacturing Co's 28-nm high-performance, low-power process consumes more than 40% less standby power than the company's 40-nm low-power process. GlobalFoundries, meanwhile, offers three power levels for its 28-nm node (**Figure 1**).

Moore's Law continues unabated, and chips are packing more functions into each device. According to Colin Baldwin, director of marketing for Open-Silicon, customers can design the next-generation device with similar unit cost and twice the amount of performance

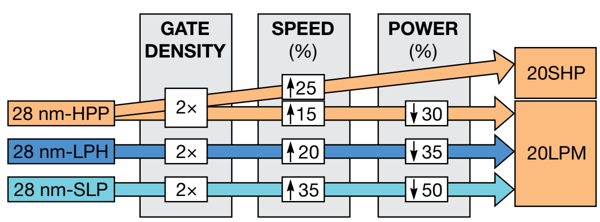

Figure 1 Taiwan Semiconductor Manufacturing Co's 28-HPL process has more than 40% less standby power than the company's 40-LP process. Global Foundries, meanwhile, offers three power levels for its 28-nm node.

but with an overall increase in power, even though power per device has decreased. Clock frequencies, another variable, are slowly drifting upward, but less quickly than the process entitlement in many markets. Open-Silicon finds that most customers are trying to integrate added function with only slightly increased overall power. Thus, to maintain the same overall power, it becomes necessary to look at the energy savings that can be made during other parts of the design flow.

OPTIMIZE AND COMPARE

Design involves estimation and optimization. Estimation allows you to make comparisons between possible implementation options. In addition, you can make optimizations automatically, or you can do so with tool assistance at various levels of abstraction. According to Arvind Shanmugavel, director of application engineering for Apache/Ansys, power estimation is an accurate science only when you have a complete design and a correct set of vectors. Until you have completed the design, everything is, by definition, an estimate of what will happen in the design. You should be looking for large and relative changes rather than absolute numbers in the power budgets during the early phases of the design. You can expect a 20% deviation between the RTL (register-transfer level) to silicon and a 10% deviation between gates to silicon, according to Venki Venkatesh, director of engineering at Atrenta.

If a tool states that one possible approach would consume less total energy than an alternative approach, this statement must be correct; otherwise, the tool may cause the selection of an inferior approach. Unlike area and performance, power is vector-related, and you may thus need to run several simulations to get a representative sample of the design's activity. For example, consider the choice of applying random data into an audio processor versus more typical speech data. **Figure 2** shows the transition activity for a few registers in a finite-impulse-response filter (**Reference 1**). For an architecture that does not destroy the data correlation, the speech data switches 80% less capacitance than does the random input. The sequencing of operations can result in large variations of the switching activity due to these temporal correlations.

Some companies, however, believe that you can get close enough using statistical methods employing expected activity from counters and other recognizable pieces of logic. You can now optimize energy consumption in many ways—most at RTL or below. According to Shanmugavel, clock gating is among the common techniques for minimizing active-power consumption. Shutting off the clock for a circuit prevents any toggle activity of the clocks or registers in a design. Another technique is to employ voltage islands, which lowers the operating voltage of a design and quadratically reduces the switching component of active power. Designers

POWER IS VECTOR-RELATED; YOU MAY NEED TO RUN SEVERAL SIMULATIONS TO GET A REPRESENTATIVE SAMPLE.

use voltage islands for areas of a chip in which performance and speed are noncritical but can save power.

DVFS (dynamic-voltage/frequency scaling) is by far the most complex active-state power-management technique. This approach changes the active operating voltage and frequency depending on the demand of the load. During high-load conditions, the voltage and frequency are at nominal conditions, and the chip or unit functions to its fullest extent. During low-loading conditions, the voltage or the frequency scales down to perform at a lower speed but provides lower active-power consumption. Designers realize this technique through a combined hardware/software approach.

On-die voltage regulators meet the demands of various active- and static-power requirements. ICs usually have off-chip voltage-regulator modules that can supply the voltage and current requirements for active states. However, designers are increasingly using on-die regulators as the number of voltage domains and the need for these voltage domains to respond faster to the demand current increase.

Stacking ICs that communicate with one another to minimize the signal interconnect is an emerging trend in low-power design. According to Apache's Shanmugavel, manufacturers often stack processors and memory

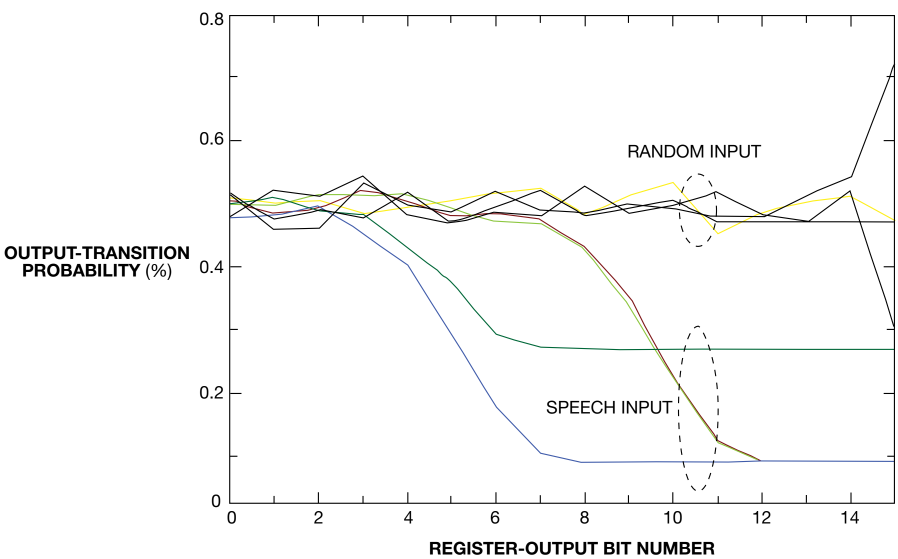

Figure 2 For an architecture that does not destroy the data correlation, the speech data switches 80% less capacitance than does the random input. The sequencing of operations can result in large variations of the switching activity due to these temporal correlations.

Figure 3 To get a feel for the adoption rates of these techniques, Synopsys collects data from customers using its Global User Survey. The top, left-most column, for example, means that 10% of survey respondents have data center and networking as a primary application using back-biasing or well-biasing. Note that percentages equal more than 100 because the survey accepted multiple answers (courtesy Synopsys Inc, Global User Survey, 2011).

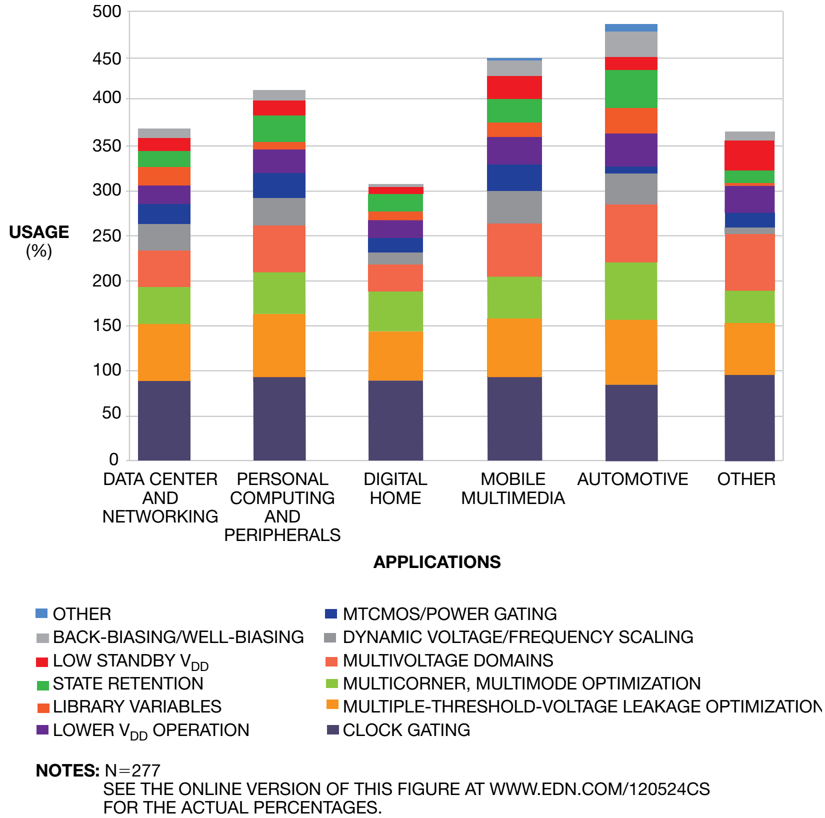

over a silicon interposer that makes connections using TSVs (through-silicon vias). These interposers provide a low-capacitance signal interconnect between the die, thus reducing the I/O's active-power consumption. As the cost of 3-D ICs begins to decrease and as designers get a better understanding of thermal impact, a migration to 3-D ICs will occur across the industry.

To minimize static-power consumption, designers can use power gating, which provides the maximum amount of leakage-power savings for a device in standby. Shutting off the clocks for the functional units reduces active power, but the unit still consumes leakage power. Designers must understand several trade-offs in power gating before implementing the design.

One of the oldest techniques for reducing leakage power is to swap nominal-threshold-voltage gates with high-threshold-voltage gates. In CMOS, the subthreshold leakage is inversely proportional to the threshold voltage. Higher-threshold-voltage devices have a lower leakage envelope than do lower-threshold-voltage devices but come at a cost of larger delays. Designers must perform a careful trade-off analysis to achieve optimal leakage savings using this technique.

Another approach for reducing static power, active back-biasing, increases the bias voltage of the substrate nodes in CMOS gates to reduce the leakage current. This biasing technique essentially increases the threshold voltage of a unit or the entire chip during standby modes, thus decreasing the leakage power. To get a feel for the adoption rates of these techniques, Synopsys collects data from customers using its Global User Survey (**Figure 3**).

Besides performing RTL optimization, designers are also developing tools that allow the estimation and exploration of architectures at the system level. Power is a system-level issue, and some designers find that they cannot think of power in the bottom-up manner used today for chip assembly and verification. Designers traditionally designed chips for maximum flexibility; with the current costs of designing chips, that flexibility remains a big consideration. As with everything else, however, flexibility brings a price. Processors are generally the least energy-efficient approach to an arbitrary problem but can often be implemented in the smallest area, given their multitude of functions.

VERIFICATION

Power adds another layer of complexity that designers must verify. It requires additional tool support that manufacturers are cobbling into those now on the market. Power adds several new devices to the design, such as isolation logic, power switches, level shifters, and retention cells.

However, according to Krishna Balachandran, director of low-power-verification marketing at Synopsys, power optimizations may also involve sequential RTL transformations that must be verified against the original RTL. The absence of such verification can lead to nonfunctioning systems on chips or higher-than-desired leakage. Simulation approaches may be too slow, not cost-effective, and not exhaustive, leading to incomplete verification coverage of power optimizations. Traditional formal-equivalence tools typically target verification of combinational transformations and are inadequate for the kind of changes typically necessary for power optimizations. Most commercially available formal-equivalence tools also suffer from capacity and performance limitations that must be overcome to handle low-power designs with complex power architectures and hundreds of power domains. A new class of formal-equivalence tools with high capacity and performance targeting verification of sequential transformations must evolve to meet these new requirements.

According to Lauro Rizzatti, general manager of Eve-USA, power optimization also presents challenges for EDA vendors. Many low-power techniques are generally incongruent with RTL simulation or emulation, which abstracts out any notion of voltage. Designers must adapt these digital tools to support power intent and low-power-optimization techniques for implementation.

POWER-DELIVERY NETWORKS

According to Dermott Lynch, vice president of marketing at Silicon Frontline Technology, power devices typically operate at 70 to 90% efficiency, which results in a loss of 10 to 30% of overall system power. Ely Tsern, vice president and chief technology officer for Rambus' semiconductor-business group, adds that more aggressive power-mode transitions, with finer-grained power domains, will result in faster transition of local supply currents, which in turn can induce greater di/dt supply noise for sensitive local circuits, especially analog circuits.

Shanmugavel cautions, however, that, under all conditions, the power-delivery network should be able to sustain the load without compromising the voltage integrity. For example, when a global clock transitions and a functional unit turns on to perform a task, a transient-current demand occurs. This transient current can be three to five times that of the nominal current, depending on the functional block, which places an enormous load on the power-delivery network. You must validate the transient voltage noise on the network under these circumstances. EDN

REFERENCE

1 Chandrakasan, Anantha P; and Robert W Brodersen, "Minimizing Power Consumption in CMOS Circuits," University of California—Berkeley, http://bit.ly/HIetoD.

You can reach Contributing Technical Editor Brian Bailey at brian_bailey@acm.org.

BY STEVE LOGAN • MAXIM INTEGRATED PRODUCTS INC

Understanding noise, ENOB, and effective resolution in ADCs

NOISE, ENOB, AND EFFECTIVE RESOLUTION ARE CRITICAL PARAMETERS WHEN CHOOSING WHAT YOU NEED.

One of the major trends for ADCs is the move toward higher resolution. The trend affects a range of applications, including factory automation, temperature sensing, and data acquisition. The need for higher resolution is leading designers from traditional 12-bit SAR (successive-approximation-register) ADCs to delta-sigma ADCs with resolutions that reach 24 bits. All ADCs have a certain amount of noise, including both input-referred noise, which is inherent to the ADC, and quantization noise, which is the noise that occurs while the ADC is converting. Specifications such as noise, ENOB (effective number of bits), effective resolution, and noise-free resolution in large part define how accurate an ADC is.

Consequently, understanding the performance metrics relating to noise is among the more difficult aspects of moving from a SAR to a delta-sigma ADC. With the current demand for higher resolution, designers must develop a better understanding of ADC noise, ENOB, effective resolution, and signal-to-noise ratio.

HIGHER RESOLUTION

In the past, a 12-bit SAR ADC was often good enough to measure a variety of signals and voltage inputs. If an application needed finer measurement, a designer could add a gain stage or a PGA (programmable-gain amplifier) in front of the ADC. For 16-bit designs, a designer's choice is still primarily SAR ADCs but also includes some delta-sigma ADCs. For designs that need more than 16 bits, however, delta-sigma ADCs are becoming more prevalent. SAR ADCs currently are at an 18-bit limit, whereas delta-sigma ADCs are expanding their presence at 18, 20, and 24 bits. ADCs' prices have decreased considerably in the last 10 years, and the units have become simpler to use and more widely understood.

EFFECTIVE RESOLUTION

The following **equations** define effective resolution in bits: Effective resolution=$\log_2$[full-scale-input-voltage range/ADC-rms noise], or, more simply, effective resolution=$\log_2[V_{IN}/_{VRMSNOISE}]$. Do not confuse effective resolution with ENOB. The most common method for measuring ENOB uses a fast-Fourier-transform analysis of a sine-wave input to the ADC. The IEEE standard 1057 defines ENOB as equal to $\log_2$[full-scale-input-voltage range/(ADC-rms noise×$\sqrt{12}$)]. SINAD, defined as the SNR-and-distortion ratio, and ENOB are used to measure the ADC's dynamic performance. Therefore, SINAD=(rms input voltage/rms noise voltage), where rms noise is

$$1/M\left[\sum_{m=0}^{M-1} E_{AVMFM}\right].$$

Here, E_{AVM} is the residual of X_{AVM} and X_{AVMFM} is the averaged magnitude spectral component at a given discrete frequency after a discrete Fourier transform.

Effective resolution and noise-free resolution measure the ADC's noise performance at essentially dc, which does not factor in spectral distortion, including total harmonic distortion and spurious-free dynamic range. Once you know the ADC's noise and input range, calculating effective resolution and noise-free resolution becomes simple.

The ADC's input-voltage range is based on the reference voltage. If the ADC integrates a PGA, you must also factor the PGA into the voltage range. Some delta-sigma ADCs include PGAs to increase the gain of small signals. The newest ADCs with PGAs often specify the noise as less than 100 nV rms. Although this noise figure looks impressive compared with that of older ADCs, the figures often employ a small input range because the small range ultimately amplifies to fit a wider portion of the ADC's active range, based on the reference voltage. So, the effective resolution and noise-free resolution of these ADCs may be worse than those of ADCs without PGAs.

Consider, for example, a 24-bit ADC with a PGA setting of 128. It offers 70-nV-rms noise with a reference voltage of 2.5V and an input range of $\pm V_{REF}$/PGA(±2.5V/128=39.1 mV). The effective resolution is, therefore, $\log_2[V_{IN}/V_{RMSNOISE}]=\log_2$[39.1 mV/70 nV]=19.1 bits. Using the same ADC with a PGA setting of one, the noise rises to 1.53 μV rms. With an input range of 5V(±2.5V/1), the effective resolution becomes 21.6 bits. The best practice is to check the ADC's data sheet for the input range that you need.

NOISE-FREE RESOLUTION

Noise-free resolution uses the peak-to-peak-voltage noise rather than the rms noise. The following **equation** defines the noise-free resolution in bits: Noise-free resolution=$\log_2$ [full-scale-input-voltage range/ADC peak-to-peak noise].

TABLE 1 MAX11200 SAMPLE RATE VERSUS NOISE

Data rate (samples/sec)						
2.4576-MHz internal oscillator for 60-Hz rejection	2.048-MHz internal oscillator for 50-Hz rejection	ADC noise (µV rms)	Bipolar noise-free resolution (bits)	Bipolar effective resolution (bits)	Unipolar noise-free resolution (bits)	Unipolar effective resolution (bits)
1	0.83	0.21	22.3	24	21.3	24
2.5	2.08	0.27	22	24	21	23.7
5	4.17	0.39	21.4	24	20.4	23.1
10	8.33	0.57	20.9	23.6	19.9	22.6
15	12.5	0.74	20.5	23.2	19.5	22.2
30	25	1.03	20	22.7	19	21.7
60	50	1.45	19.5	22.2	18.5	21.2
120	100	2.21	19	21.7	18	20.7

Noise-free resolution=$\log_2[V_{IN}/\text{V-p-p noise}]$. You can think of noise-free, or flicker-free, resolution in terms of a 5½- or 6½-digit multimeter in the lab. If the last digit on the display is stable and not flickering, the data-output word is better than the noise level of the system. Using a crest factor of 6.6 as an example, the peak-to-peak noise is 6.6 times the rms noise. As a result, the effective resolution is 2.7 bits higher than the noise-free resolution. Using the same noise and reference values, the noise-free resolution is 18.9 bits.

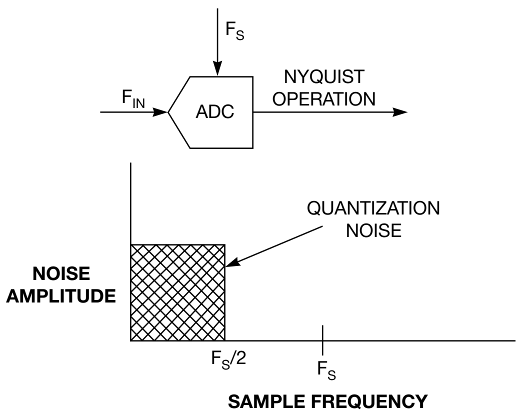

Figure 1 The noise performance of a standard ADC is worse than that of a delta-sigma device.

NOISE-FREE COUNTS

Noise-free counts are another metric that precision systems use to evaluate ADC performance, especially in applications such as weigh scales, which may require 50,000 noise-free counts. You can calculate this value by converting the noise-free resolution into counts by a factor of 2^N. For example, using the formula 2^{10}, an ideal 10-bit ADC has 1024 noise-free counts. An ideal 12-bit ADC has 4096 noise-free counts. Again, using the same noise-free resolution values, that example would yield $2^{18.9}$, or 489,178 noise-free counts.

OVERSAMPLING WITH DELTA-SIGMA ADCs

Delta-sigma ADCs use an oversampling architecture, which means that the ADC's internal oscillator/clock runs at a higher frequency than the output-data, or throughput, rate. Some delta-sigma ADCs can vary the output-data rate, allowing designers to optimize sampling for higher speeds with worse noise performance or for lower speeds with more filter-

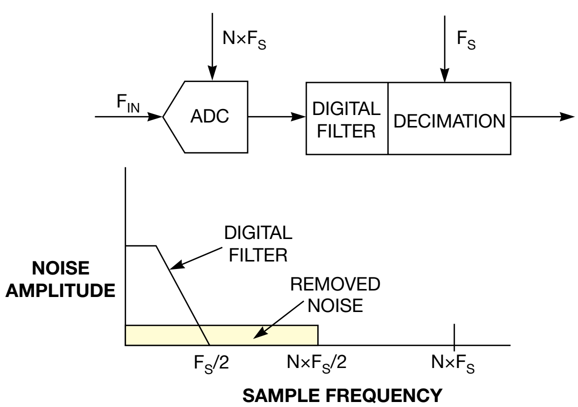

Figure 2 Noise performance improves in an ADC with oversampling by a factor of N, a digital filter, and decimation.

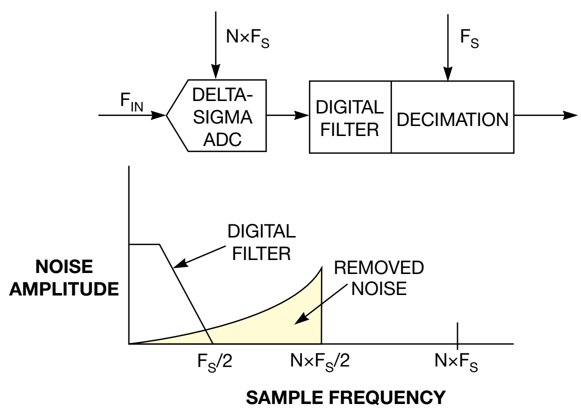

Figure 3 In an ADC with oversampling by factor of N, noise shaping, a digital filter, decimation, and noise (yellow) in the ADC's input frequency band of interest decrease significantly.

ing, noise shaping—that is, pushing the noise into the frequency band outside the area of measurement interest—and better noise performance. Many new delta-sigma ADCs offer the effective-resolution and noise-free-resolution results in table form, making it easy to compare the trade-offs.

Table 1 shows an example ADC's data rate, noise, noise-free resolution, and effective resolution in both bipolar modes and unipolar input modes. The 24-bit MAX11200 ADC can measure either bipolar or unipolar inputs. It operates from a 2.7 to 3.6V supply, and the reference can be biased up to the supply. The bipolar values are based on the maximum input range of ±3.6V; the unipolar measurements are based on a 0 to 3.6V input range.

Designers can program the MAX-11200's internal oscillator through the software for 2.4576 MHz for 60-Hz rejection at the lower data-rate settings or for 2.048 MHz for 50-Hz rejection at lower data rates. At either data rate, the ADC noise is the same. Therefore, the resulting noise-free resolution and effective-resolution values are consistent. You can apply an external oscillator for a 55-Hz notch filter that yields good rejection at both 50 and 60 Hz.

Effective bipolar resolution has a 24-bit maximum because the output data word is 24 bits long. At the three slowest data-rate settings, the ADC's noise level is low enough that the effective resolution is better than 24 bits if the ADC were to output more than 24 bits of data on the serial interface. Unless the data-output word limits it, effective resolution is always 2.7 bits better than the noise-free resolution. Noise shaping allows delta-sigma ADCs to achieve low noise and high precision.

NOISE SHAPING, FILTERING

Figure 1 shows a standard ADC's quantization noise; **Figure 2** details an ADC that includes oversampling, a digital filter, and decimation. Most ADC cores that use oversampling are delta-sigma units. Oversampling by a factor of N spreads the noise over a wider frequency band, whereas the digital filter removes a large portion of the noise. **Figure 3** details a delta-sigma modulator that adds noise shaping to the same blocks as those in **Figure 2**. Pushing the noise to disproportionately higher frequencies makes the noise in the frequency band of interest ultralow. Techniques such as this one allow delta-sigma-ADC manufacturers to achieve noise figures of less than 1 μV rms.

With their oversampling capability and inherently low noise, delta-sigma ADCs represent excellent design choices for systems that require higher resolution. As designers must resolve even smaller signals, a firm understanding of ADC noise, effective resolution, ENOB, and noise-free resolution becomes integral to choosing the right ADC.EDN

ACKNOWLEDGMENT

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

AUTHOR'S BIOGRAPHY

Steve Logan received a bachelor's degree in electrical engineering from San Jose State University (San Jose, CA).

Simple anticipator circuit improves on earlier idea

Nicholas Lockerbie, Scottish Universities Physics Alliance, Glasgow, Scotland

A self-biasing thermistor circuit had an overall dominant thermal/electrical time constant as small as 25 msec but had a significantly rounded response following a square (thermal) input pulse of 100 msec; the source was a pulsed blue LED. A single-op-amp-based "anticipator-like" circuit generated a more faithful picture of the actual input excitation, and improved on the circuit described in a previous Design Idea (**Reference 1**). The lowpass filter in **Figure 1**, with an RC time constant of approximately 25 msec, models the circuit.

The input voltage models the stimulus to the transducer—a pulsed thermal flux, say—and the voltage response, V_X, at the output of the transducer forms the input to the corrective anticipator circuit, whose output is the voltage signal, V_{OUT}. In the complex frequency, s, you can write the output of the lowpass filter as a function of its input, V_{IN}, in the form of the following well-known **equation** for a single-pole response:

$$\frac{V_X}{V_{IN}}=\frac{1/sC}{R+(1/sC)}=\frac{1}{1+sCR}.$$

The single pole of this lowpass response is at frequency $s=-(1/CR)$.

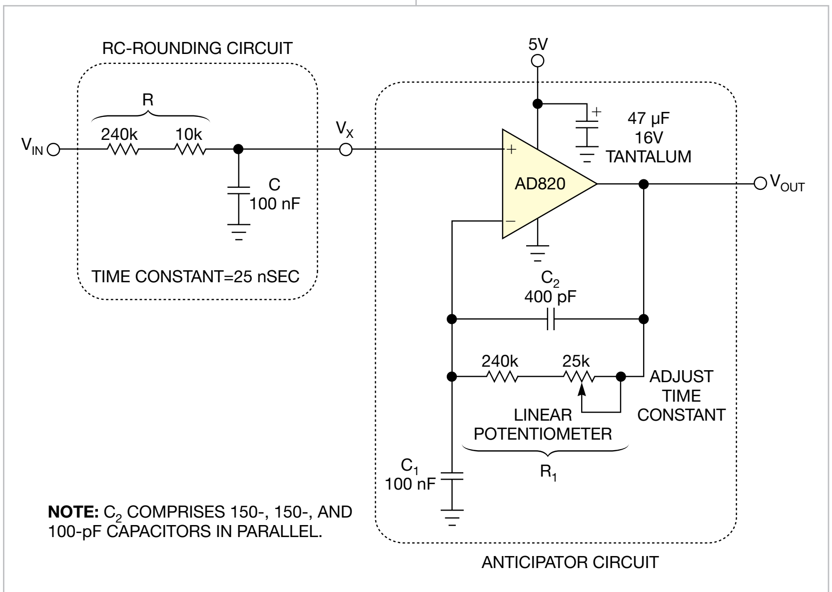

Figure 1 This RC circuit models a single dominant time-constant transducer, producing a rounded, intermediate voltage, V_X, in response to an input signal, V_{IN}. With V_X as its input, the anticipator circuit reconstructs the original signal from V_X, such that $V_{OUT}=V_{IN}$.

DIs Inside

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

Ignoring for now the 400-pF capacitor in the feedback loop of the circuit of **Figure 1**, the voltage gain arising from its noninverting op-amp configuration is

$$\frac{V_{OUT}}{V_X}=1+\frac{R_1}{(1/sC_1)}=1+sC_1R_1.$$

Consequently, this response has a single zero at $s=-(1/C_1R_1)$.

If you cascade the responses of the RC-rounding transducer and anticipator circuits, then the overall input-to-output system function becomes

$$\frac{V_{OUT}}{V_{IN}}=\frac{1+sC_1R_1}{1+sCR}.$$

Therefore, by choosing the time constant of the anticipator to be equal to that of the transducer—that is, by choosing here $C_1R_1=CR$, pole-zero cancellation occurs in the system function, such that $V_{OUT}=V_{IN}$.

In other words, the anticipator circuit's output now becomes identical to the unmodified input stimulus to the transducer. Furthermore, this situation must be true no matter what temporal form the input signal takes.

Figures 2, **3**, and **4** show an input waveform to the transducer; the corresponding transducer's output signal;

Securing Energy

Proven security solutions protect the grid's endpoints

Attacks to the smart grid can occur from any point on the network. From stealing electricity and customer data to threatening critical assets, cyber criminals can deploy threats to the grid miles from the control room.

Maxim's Newport smart meter development platform protects the endpoints of the grid with advanced cryptography, physical-attack-detection technology, and sensitive data protection so the action doesn't stop.

Proven Security

- Accelerators support asymmetric schemes including RSA, DSA, and elliptic curve
- True hardware random number generation generates stronger challenges and keys
- Tamper inputs detect physical attacks
- Built-in IP protection for supply chain and manufacturing process

Get Started with Built-In Security

Newport Development Platform

Free whitepaper: "Securing Smart Meters."

www.maxim-ic.com/newport

MAXIM
INNOVATION DELIVERED®

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

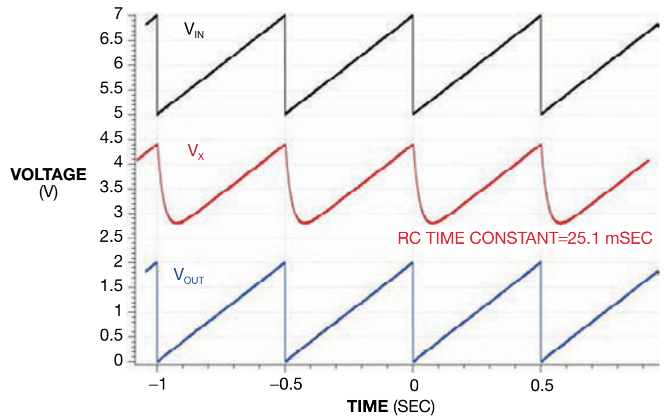

Figure 2 The reconstruction of the input sawtooth waveform shows that the circuit can operate properly with complex, possibly repetitive waveforms.

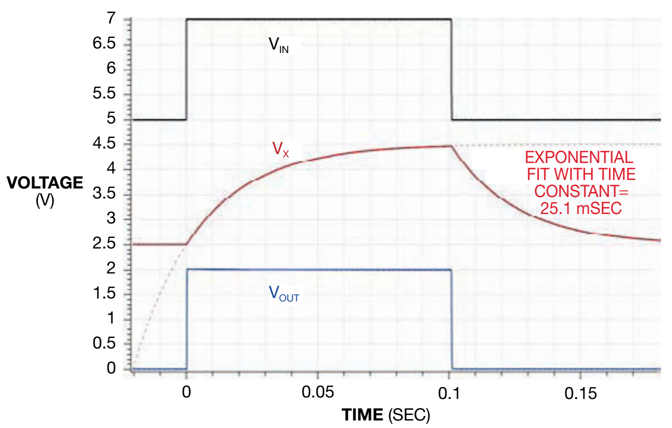

Figure 3 The output of the anticipator jumps immediately to the asymptotic limit of a rising or falling exponential transducer response.

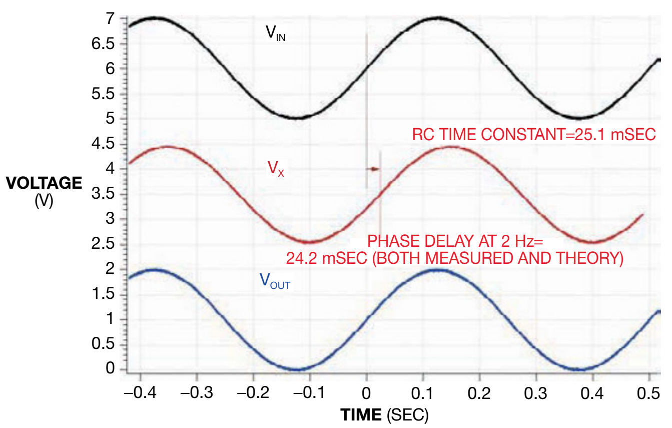

Figure 4 The transducer-induced phase shift of a sine-wave signal having a single frequency component is properly back-corrected in time, and the slightly attenuated transducer output has been restored at the output.

and the final anticipator circuit's output, which is derived from the available signal, V_X. The **figures**, respectively, show these waveforms for three forms of 2-Hz frequency input: a sawtooth, a pulse input with 20% duty cycle, and a sine-wave input. In each **figure**, the waveforms for V_{IN} and V_X have been offset, vertically, by 5 and 2.5V, respectively. The sawtooth and sine-wave input signals have 1V offsets to keep the inputs positive for the single-supply op amp. In practice, a 400-pF capacitor is necessary for optimal stability because it prevents the op amp's closed-loop gain from tending toward infinity, with rising frequency. Also, although the AD820 is operating in single-supply mode, it can also operate in split-supply mode to handle bipolar signals.

FOR COMPLEX, POSSIBLY REPETITIVE WAVEFORMS, THE RECONSTRUCTION OF THE INPUT SAWTOOTH WAVEFORM PROVIDES COMPELLING EVIDENCE OF THE CIRCUIT'S PROPER OPERATION.

The anticipator circuit performs as you would expect, producing closely equal input and output voltages for all of the tested input waveforms. For more complex, possibly repetitive waveforms, the reconstruction of the input sawtooth waveform provides compelling evidence of the circuit's proper operation in such a case (**Figure 2**). Like the original circuit, the output of this anticipator jumps immediately to the asymptotic limit of a rising or falling exponential transducer response (**Figure 3**).

In **Figure 4**, on the other hand, the transducer-induced phase shift of a sine-wave signal having a single frequency component has been properly back-corrected in time, by −24.2 msec, and the slightly attenuated transducer output—in this case, to 95% of the input amplitude—has been restored at the output to a level of 99.7%. The dc component of the input voltage also has communicated properly and in all cases to the output.

An important general caveat is that the intermediate signal, V_X, must remain sufficiently large, following the rounding action of the transducer, for the anticipator circuit to accurately reconstruct the input stimulus. EDN

REFERENCE

1 Regan, Tim, "Anticipator circuit speeds signal settling to a final value," *EDN*, March 17, 2011, pg 58, http://bit.ly/h7qZPo.

Perform the XOR/XNOR function with a diode bridge and a transistor

Raju Baddi, Tata Institute of Fundamental Research, Pune, India

When designing logic for higher-than-usual supply voltages, such as 24V, you can use a voltage regulator with a standard logic family and interface it through level shifters. Alternatively, if the logic is not too complex and the speed is not extensively high, you can build gates from discrete components and operate them directly from the available voltage. Discrete-component AND, OR, and NOT functions are relatively straightforward, but XOR and XNOR functions usually require combining several of the basic AND, OR, and NOT functions.

This Design Idea presents an unusual method of performing the exclusive functions with two resistors, four diodes, and one transistor. The NPN configuration results in the XNOR operation, and the PNP configuration gives the XOR operation.

Consider the XNOR circuit in **Figure 1a**. When either of the gate inputs, A or B, is at an opposite logic state, a voltage drop of the high voltage minus the low voltage minus 1.2V is available to forward-bias the base-emitter junction. The transistor turns on, and the logic-zero voltage at the collector is approximately $0.6+V_L+V_{CE}$, where V_L is the low voltage and V_{CE} is the collector-to-emitter voltage. When inputs A and B are at the same logic state, you cannot forward-bias the transistor's base-to-emitter junction, so output Y is at the supply voltage.

The choice of 6.8 kΩ at the collector is based on driving the A and the B inputs with standard transistor-transistor or CMOS logic, and you can change it to suit your application. The CMOS 4000 family can reliably source or sink 1 mA with a 5V supply. Low-speed TTL can source 0.4 mA and sink 8 mA. A 0.4-mA logic-one drive is sufficient for the base current, but the logic zero at either A or B forms the emitter current and is of more concern with regard to the limited 1-mA sink of the CMOS. With a net 1-mA current and approximately 250 μA remaining for the output load, you must choose a 6.8-kΩ resistor (0.75 mA×6.8 kΩ) to drop approximately 5V.

Next, consider the XOR configuration in which logic zero at either A or B is relevant for the base and logic one is relevant for the emitter. The logic-one voltage available at Y is V_H–0.6V–V_{CE}, and logic zero is approximately 0V but current-limited through the collector resistor.

The concern here is with the TTL logic-one output current of approximately 0.4 mA, which is the transistor's emitter current. With the choice of 10 kΩ for the collector resistor, the voltage drop across it can reach nearly 4V. This level would be tolerable for driving CMOS loads but not for TTL, which demands a current of at least 0.4 mA for logic-zero inputs when Y is at logic zero. The 10-kΩ resistor cannot provide this current. However, using the previous XNOR configuration and another inverting transistor after Y, you can obtain the XOR function (**Figure 1b**). The XOR seems to be suitable only for the CMOS/TTL input at A and B and capable of driving only CMOS at output Y. EDN

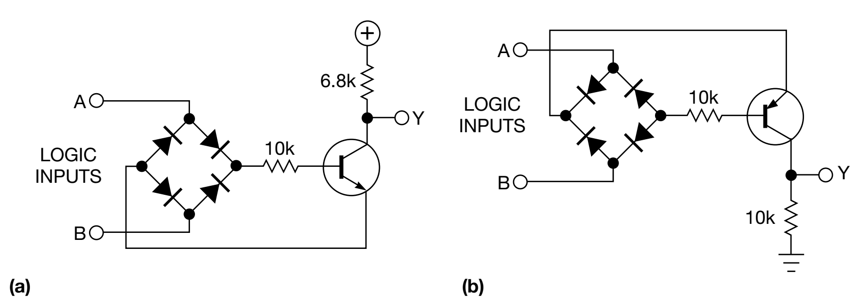

Figure 1 Discrete implementations of XNOR (a) and XOR operations (b) allow you to run logic at higher supply voltages than standard logic families.

Use an op amp as a set/reset flip-flop

Aruna Prabath Rubasinghe, University of Moratuwa, Moratuwa, Sri Lanka

You can make a set/reset flip-flop from two NAND or NOR gates or use readily available set/reset flip-flop ICs, such as the 74HC279 quad-set/reset latch. The drawback of these methods is that they require a large amount of space to form the flip-flop. Even if you need only one flip-flop, you must use a large IC package with this approach. However, you can use a spare rail-to-rail op amp to perform the needed latch function (**Figure 1**).

This approach has a low space requirement because the rail-to-rail-input/output MCP6022 is a dual-op-amp package, meaning that you can build two set/reset-latch circuits in a small amount of space. By choosing the appropriate op amp, you can also operate this circuit at nonstandard supply voltages if necessary.

You set the op amp's inverting input to 2.5V using the R_1/R_2 voltage divider. You cannot pull it lower than 2.5V due to D_4's being reverse-biased. You can, however, drive the input to a logic high. Series diodes D_1 and D_2 apply positive feedback to the op amp's noninverting input, which R_3 pulls low if both the output and the set input are low.

With the output low, applying a 5V

pulse to the set input forward-biases D_3; D_1 and D_2 remain reverse-biased. The resulting 4.4V at the op amp's noninverting input drives the output high, forward-biasing D_1 and D_2 and latching the noninverting input at 3.8V—much higher than the inverting input voltage, even after the set input returns to low. If you then drive the reset input high, the inverting input, at 4.4V, becomes higher than the noninverting input's 3.8V, driving the output low. When the reset input returns to low, D_4 reverse-biases and the 2.5V at the inverting input holds the output low. Note that the input pulses must be of sufficient duration to allow for the op amp's delay and slew rate.

Unlike a true set/reset latch, both inputs' being simultaneously active high yields unknown results. You can introduce a slight voltage offset by using a series resistor in either the set or the reset input to ensure a desired logic state under this condition. CMOS logic drives rail to rail and can interface directly with this circuit. However, driving with transistor-transistor logic requires pullup resistors at the TTL outputs to ensure sufficient input voltage. You may need to increase the value of the 10-kΩ resistors to avoid loading the pullup resistors.

The selected op amp must have rail-to-rail-output capability. You can use op amps without this ability by increasing the supply voltage as necessary to compensate the output voltage for the required logic level. You can change the trigger-voltage level by adjusting R_1 and R_2 and by using low-voltage drop diodes, such as germanium or Schottky types.

This circuit has relatively low input impedance, but it is usually not a problem because most logic-output signals can drive this load. You can change the input impedance by adjusting the resistor values as necessary. EDN

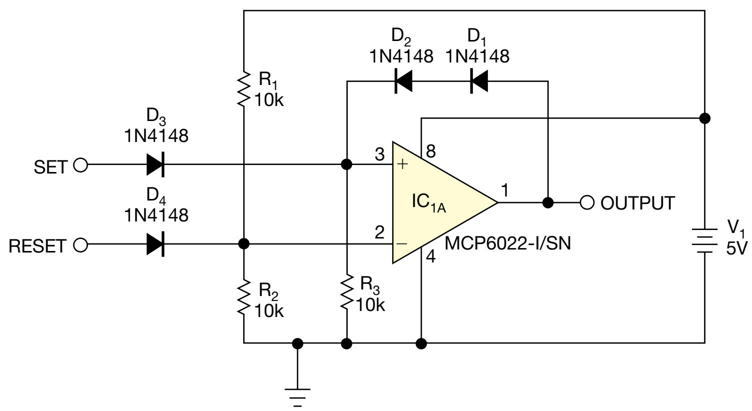

Figure 1 Positive feedback and diode steering latch the op amp's output high or low when you apply a positive-going pulse to the set or the reset input.

Generate a pulse width that is inversely proportional to the square root of an analog voltage

Marián Štofka, Slovak University of Technology, Bratislava, Slovakia

The circuit in **Figure 1** is an edge-triggered monostable circuit, based on a previous design of an edge-triggered parabolic-pulse generator (**Reference 1**). This circuit involves a simple but significant modification of the earlier generator by disconnecting the input comprising IC_3 and S_2 (in the original design) of the first of the integrators of a cascade from the source of the reference voltage, V_{REF}, and connecting it to the input-voltage terminal in **Figure 1**.

The width of the output pulse at output Q in this circuit is

$$T_Q(V_{IN})=T_{QMIN}\times\sqrt{\frac{V_{REF}}{V_{IN}}},$$

where

$$T_{QMIN}=\frac{5\tau_{IL}}{3}=2\tau_{IQ},$$

and τ_{IL} and τ_{IQ} are time constants of the first and the second integrators in the cascade comprising **Reference 1**'s IC_{2D} and IC_{2C}, respectively.

Although the monostable would be functional with just this modification, the logic circuitry of IC_1, IC_2, and IC_3 of **Figure 1** adds another feature. The added logic ensures that the generator ignores the next trigger pulse that comes within the monostable's busy state.

In this way, the generator's integrator capacitors can discharge close to 0V with an error not exceeding 0.4%, even at relatively high trigger frequencies, exceeding the value of $1/[T_Q(V_{IN})]$. Consequently, the output pulses at a given input voltage are of constant width, even when the triggering period closely approaches or is smaller than the width of the output pulse.

The subcircuit comprising IC_1 and IC_2 generates an RST (reset) signal, the trailing edge of which determines the end of one cycle of operation of the monostable. The RST signal in this circuit inhibits retriggering of the monostable within the interval in the Q output's low-to-high transition and the RST signal's high-to-low transition. For this purpose, the triggering signal's clock is ORed in IC_3 with the RST signal (**Figure 2**).

The next effective triggering is thus enabled just after the trailing edge of the RST pulse. The leading edge of the RST pulse occurs roughly when

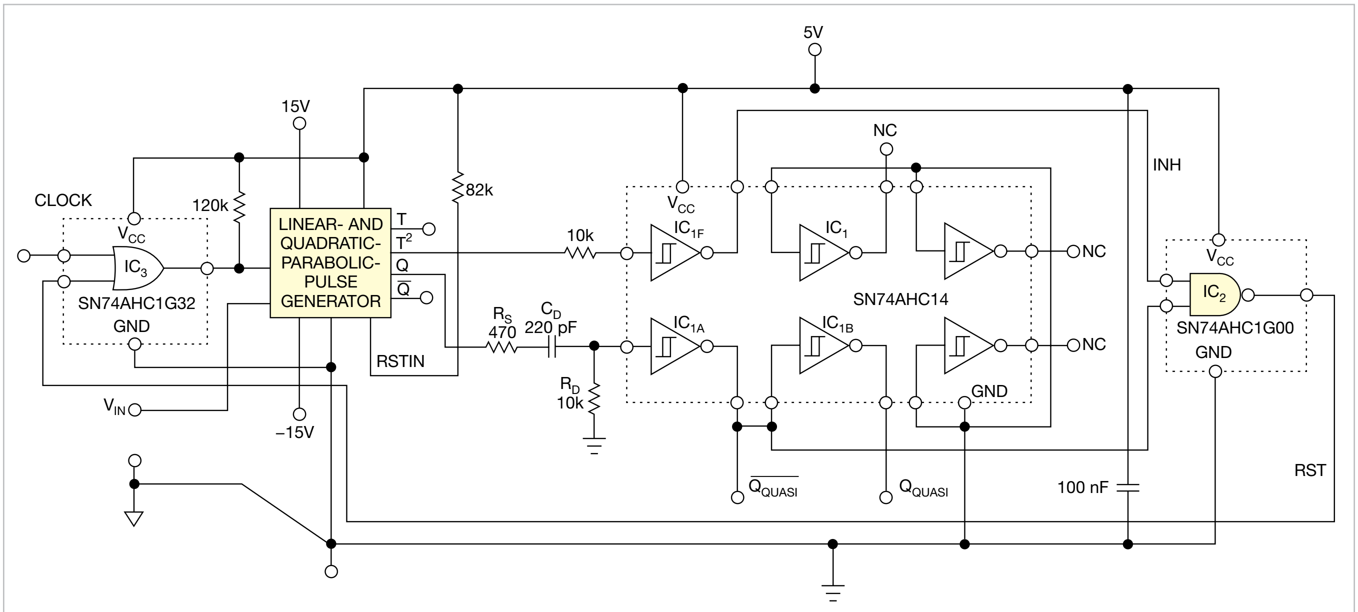

Figure 1 A low-to-high transition at the clock input triggers the monostable. The width of the generated pulse at the complementary Q and Q̄ outputs is a mathematically defined nonlinear function of the analog input voltage of 0 to 3V.

the quadratic-parabolic voltage, V_{OQ}, reaches half its peak voltage, V_{PEAK}. The trailing edge of the RST pulse is delayed with respect to the instant of V_{OQ}'s dropping below $V_{PEAK}/2$. The auxiliary time constant, $(R_D+R_S)C_D$, of the $R_S/C_D/R_D$ network at the input to IC_{1A} defines this delay.

Experimental evaluation shows that the relative error of the output pulse width,

$$\delta_{TQ}=\frac{T_{QMEAS}(V_{IN})}{T_{QMEAS}(V_{REF})}\times\sqrt{\frac{V_{REF}}{V_{IN}}}-1,$$

is negative, not exceeding -8×10^{-4} for an input voltage of approximately 200 to 3000 mV, with a reference voltage of 3000 mV, which **Reference 1**'s IC_1 sets.

The error then rises in magnitude, reaching a maximum of $\delta_{TQ}=-2.337\times10^{-3}$ at an input voltage of 99.925 mV. By further lowering the input voltage, the negative error decreases in magnitude and is $\delta_{TQ}=-1.113\times10^{-3}$ at an input voltage of 9.915 mV. At an input voltage of 3.08 mV, the relative error is positive, $\delta_{TQ}\approx2.9\times10^{-3}$. Further decreasing the input voltage causes the positive error to rise rapidly, reaching 3% at an input voltage of 1.065 mV. Note, however,

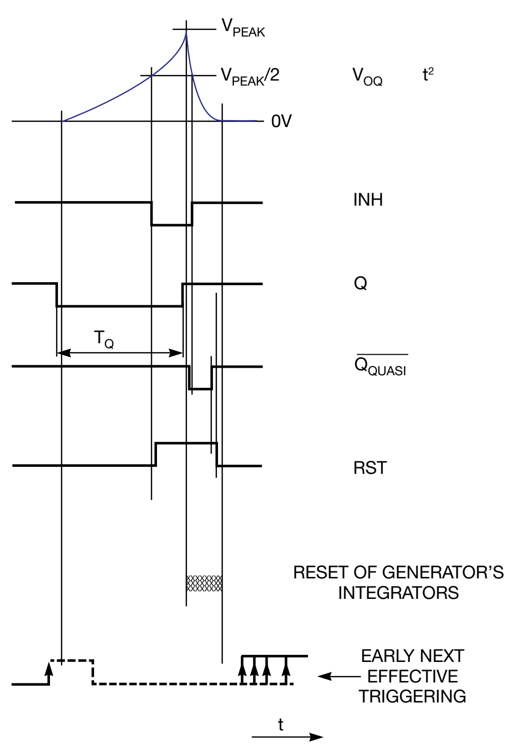

Figure 2 The high level of the generated RST-logic signal prevents any eventual low-to-high transitions at the clock input from triggering the monostable until the integrators of the generator reset in a defined way.

that the input-voltage span is almost 3000-to-1. The trigger frequency is 2 or 200 Hz.

You can obtain almost the same pulse widths at trigger frequencies of 2 kHz, 200 kHz, and 2 MHz. The relative change of pulse width due to trigger-frequency variation is comparable to δ_{TQ} values or lower. A full-scale input, with the input voltage equal to the reference voltage, achieves a measured pulse width of 445.44 µsec.

With the V_{OQ} output, you can also use the circuit as a precision quadratic-parabolic-timebase generator; the input voltage controls the generator's speed. EDN

REFERENCE

1 Štofka, Marián, "Positive edges trigger parabolic timebase generator," *EDN*, July 28, 2011, pg 51, http://bit.ly/yYvu3F.

How the aftermarket and counterfeiting are linked

At first glance, it may seem that back-end services, such as collection, recycling, and disposal, have little to do with front-end concerns, such as counterfeit components, but they have more to do with each other than you might think. Manufacturers with strong partnerships in the aftermarket are more likely to be able to reduce counterfeiting problems with their parts. To understand why, consider that the primary sources of counterfeit components are factory-made parts that are rejected for sale. They are often cheap substitutes or salvaged waste components that fail to meet quality requirements, leading to potential failures, according to officials at IHS iSuppli.

Because these components contain recyclable materials, such as silicon, or hazardous elements, such as cobalt and lead, component makers outsource the collection and disposal of the parts to third parties that specialize in recycling and reclamation. These companies deal with all kinds of industries, not exclusively electronics. Therefore, they may understand the risk associated with a dysfunctional semiconductor. In some cases, they sell parts destined for scrap to another party that remarks or manipulates a failed part to look new.

The second source of counterfeit parts is PCBs that are also destined for the scrap heap. The boards are diverted and the components are picked off of them, often by hand. The third parties then remark, refurbish, and sell these parts as new. The electronics industry spends millions of dollars on detection and inspection equipment, trying to spot these parts as they enter the supply chain.

> Misrepresenting brands, such as Intel or IBM, by selling old components or systems as new represents professional suicide for a distributor.

Every sold counterfeit part represents a loss of revenue to distributors or their suppliers. As such, many distributors inspect parts before they enter the warehouse, as they leave the warehouse, and when they return to the warehouse. Suppliers and most OEM customers also conduct incoming and outgoing inspection. The redundancy of these efforts costs the supply channel time and money, yet counterfeit components still get through. Supply-chain participants in 2011 reported 1363 verified counterfeit-part incidents worldwide, a fourfold increase from 324 in 2009, according to IHS iSuppli. These numbers mark the first time the reported number of incidents in one year exceeded 1000, a total that could encompass millions of purchased parts.

Distributors are increasingly getting involved in the aftermarket. In addition to providing an extension of their core front-end services, including procurement, logistics, and fulfillment, they are moving toward the back end, as well. They take back used goods, repair them if they can, and dispose of them if they can't. These back-end services require many of the same things front-end services do: warehouses, logistics, and an expertise in electronics.

Distributors, such as Arrow and Avnet, also sell computer equipment and systems, so they deal with everything from microprocessors to plastic and steel enclosures. They also "work" for their suppliers, including Intel, AVX, and IBM. Misrepresenting those brands by selling old components or systems as new represents professional suicide for a distributor.

As disposal and recycling become more than just good ideas in the electronics industry, selecting the right aftermarket partner is critical.

—by Barbara Jorgensen, *EBN* Community Editor

This story was originally posted by EBN: *http://bit.ly/IRtyn4.*

OUTLOOK

APPLE TO REMAIN MEDIA-TABLET KING IN 2012

Despite the growing number of Kindle Fire adopters and the arrival of Microsoft-based media tablets, Apple Inc should hold its majority stake of sales to the market in 2012. According to Gartner Inc, worldwide media-tablet sales to end users should reach 118.9 million units in 2012, a 98% increase from 2011 sales of 60 million units. Apple's iOS continues to be the dominant media-tablet operating system. The research company estimates that the OS will account for 61.4% of worldwide media-tablet sales to end users in 2012.

"Despite PC vendors' and phone manufacturers' wanting a piece of the pie ... we have seen limited success [besides] Apple with its iPad," says Carolina Milanesi, research vice president at Gartner, which projects that Microsoft tablets will this year make up a meager 4.1% of media-tablet sales and grow to 11.8% of sales by the end of 2016. Windows 8 is Microsoft's official entry in the media tablet market.

Android tablets, meanwhile, should account for 31.9% of media-tablet sales in 2012 at 37.9 million devices and should grow to 137.7 million devices in 2016, according to the company.

—by Suzanne Deffree

productroundup

SENSORS/TRANSDUCERS

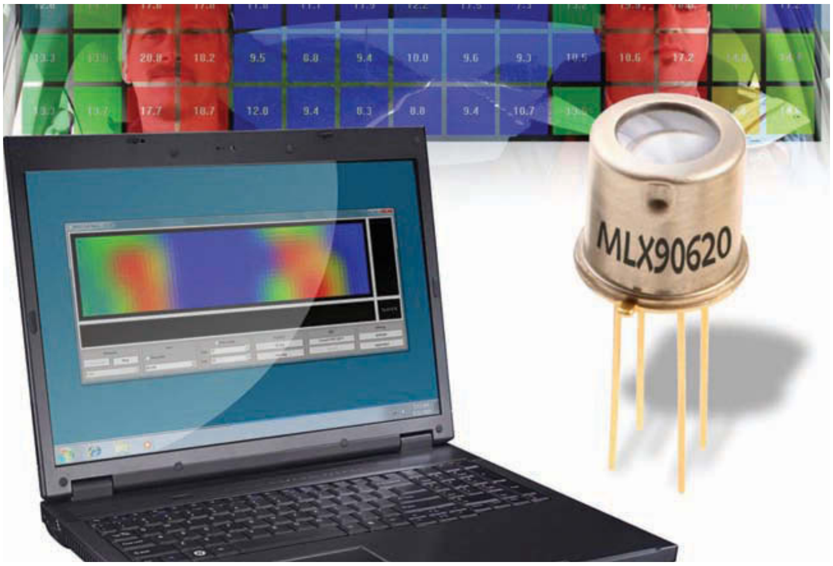

Melexis MLX90620 FIRray sensing device delivers low-cost thermography

The 16×4-element, FIR (far-infrared) MLX90620 thermopile sensor array covers a −20 to +300°C temperature range. The device produces a map of heat values for the target area in real time, avoiding the need to scan the area with a point sensor or use an expensive microbolometer. The device immediately captures 64-pixel images in 2-D and integrates an amplifier and an ADC in every pixel. It offers an adjustable frame rate of 0.5 to 64 Hz and maintains accuracy of ±1.5°C when operating at 0 to 50°C. Field-of-view options include 60×15 and 40×10°. An I^2C-compatible digital interface and a triggered mode for synchronization with a control unit allow designers to use the MLX90620 individually or combined in multiple devices to form an array with a larger imaging resolution. Packaged in an industry-standard TO-39 can, the device sells for $53.77 (100).

Melexis, www.melexis.com/MLX90620

IDT TS3000GB0A0 sensor supports 1.8V

The TS3000GB0A0 temperature sensor supports a 1.8V power supply and meets the JEDEC specification for Grade B temperature sensors over the temperature range of −20 to +125°C. An ADC enables programmable resolution as high as 12 bits with a conversion time of less than 100 nsec, irrespective of the resolution. The TS3000GB0A0 supports both SMBus and I^2C interfaces and integrates on-die power management to minimize power consumption during critical modes. The $1.12 (10,000) sensor also features input-glitch filtering and power-up voltage hysteresis to enhance fault tolerance.

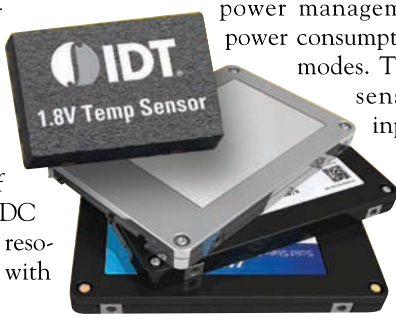

IDT, www.idt.com

Allegro A1357 sensor IC has wide ambient temperature range

The customer-programmable A1357 Hall-effect linear IC has an ambient temperature range of −40 to +150°C and features a PWM, current-sourced, two-wire output with a frequency of 1 kHz, which is proportional to an applied magnetic field. The device converts an analog signal from its internal Hall-sensor element to a digitally encoded PWM output signal. A BiCMOS, monolithic circuit in the A1357 integrates a Hall element; precision temperature-compensating circuitry to reduce the intrinsic sensitivity and offset drift of the Hall element; a small-signal, high-gain amplifier; and proprietary dynamic offset-cancellation circuits. The sensor IC targets the automotive and industrial markets and comes in a lead-free, three-pin SIP with 100%-matte-tin lead-frame plating. It sells for $1.50 (1000).

Allegro MicroSystems, www.allegromicro.com

Tektronix PSM series power sensors/meters speed measurements

The PSM RF power-meter series comes fully calibrated over its entire operating temperature range, eliminating the need for sensor zeroing and meter-reference calibration. The compact PSM3000, PSM4000, and PSM5000 target use in continuous-wave and pulse-modulation measurements. They include Microsoft Windows-based power-meter-application software for controlling the meter, displaying readings, and recording data.

productroundup

The sensors feature 2000 readings/sec to reduce test times. Products in the PSM series have a dynamic range of –60 to +20 dBm and frequencies of 10 MHz to 26.5 GHz. The PSM3000 provides true average-power measurements, independent of signal modulation and bandwidth. The PSM4000 delivers average-power measurements and adds pulse- and peak-power measurements for gathering basic data on pulsed RF and microwave signals. The PSM5000 provides the same measurements as the PSM4000 and adds pulse profiling for signal viewing and characterization in pulsed RF and microwave systems. Prices start at $2590.

Tektronix, www.tek.com

Libelium's dual RFID-ZigBee sensors enable NFC applications

This RFID/NFC module for the Waspmote sensor platform allows the use of sensor data in location-based services, such as asset tracking, supply-chain monitoring, intelligent shopping, or access management. The device uses RFID/NFC passive sensors and ZigBee active sensors, making asset tracking more accurate. Product-management software provided with the module allows users to, in real time, access information about remaining stock; storage and transportation conditions; expiration dates; and consumer profiles, including the time a customer spends at a shelf. This technology also covers security applications, including access control, to ensure maximum privacy and authentication requirements. The RFID/NFC interface reads the information from passive tags, such as cards, key rings, and stickers, and transmits it using the ZigBee radio to an Internet gateway, which uploads it to a cloud server. It can also send the identification data to the cloud using the Wi-Fi radio, which performs secure connections with Web servers. This RFID/NFC radio completes Waspmote connectivity, which supports ZigBee, Wi-Fi, Bluetooth, and 3G/GPRS, making it compliant with any wireless technology. The module is available in both 125-kHz and 13.56-MHz bands.

Libelium, www.libelium.com

Microchip MCP9808 sensor guarantees 0.5° accuracy

The MCP9808 silicon temperature sensor guarantees an accuracy of 0.5° from −20 to +100°C and temperature resolution of 12 bits (0.0625°C/LSB). The MCP9808 comes in 2×3-mm MSOPs and eight-pin DFN packages. Additional features include shutdown, an undervoltage/overtemperature monitor, and a critical-temperature alert. The device also has user-selectable measurement resolution of 0.5, 0.25, 0.125, and 0.0625°C and user-programmable temperature limits. Operating voltage range is 2.7 to 5.5V, typical operating current is 200 μA, and typical shutdown current is 0.1 μA. The device has an I²C/SMBus-compatible, two-wire interface, and prices start at 84 cents each (5000).

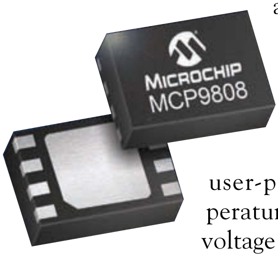

Microchip Technology,
www.microchip.com

Sensortechnics digital pressure sensors feature 24-bit resolution

These custom piezoresistive pressure sensors deliver 24-bit resolution and a total error band of 0.1% typical over pressure ranges starting at 2.5 mbar full-scale. A low-noise amplifier and a 24-bit ADC provide high-resolution digital signals with high SNRs. The devices also achieve response times of 250 μsec. The onboard microcontroller enables programmable correction algorithms. The vendor allows modifications of the digital sensor interface and can program custom communication protocols. Other optional functions include status requests and diagnostic messages.

Sensortechnics,
www.sensortechnics.com

Freescale MMA65xxKW sensors target air-bag systems

The Xtrinsic MMA65xxKW family of accelerometers employs intelligent sensing technologies for front- and side-crash detection. Safing sensors reside in a vehicle's main electronic-control unit if the air-bag system uses satellite sensors. Satellite inertial sensors sit around the perimeter of the car to detect front and side impacts. Working together, these sensors assist the vehicle in determining optimal air-bag deployment. The MMA65xxKW sensors work at 105g and a 12-bit data output with 18.2 LSB/g, enabling improved sensor resolution with no programming necessary by the manufacturer and wider dynamic-range measurement for small vehicles. Compatible with a standard SPI protocol, the sensors feature an arming-pin function that reduces the risk of data corruption in the main crash sensor of the air bag. The MMA65xxKW devices incorporate a single- or dual-axis overdamped lateral inertial sensor in a QFN package. Prices start at $3.69 (10,000) for the MMA6519KW, MMA6525KW, and MMA6527KW and at $2.99 for the MMA6555KW and MMA6556KW. The Tower System development platform helps designers combine modules and add code.

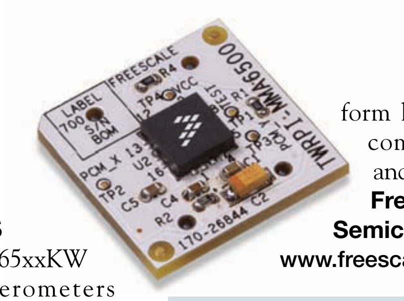

Freescale Semiconductor,
www.freescale.com

EDN ADVERTISER INDEX

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

Lightning strikes

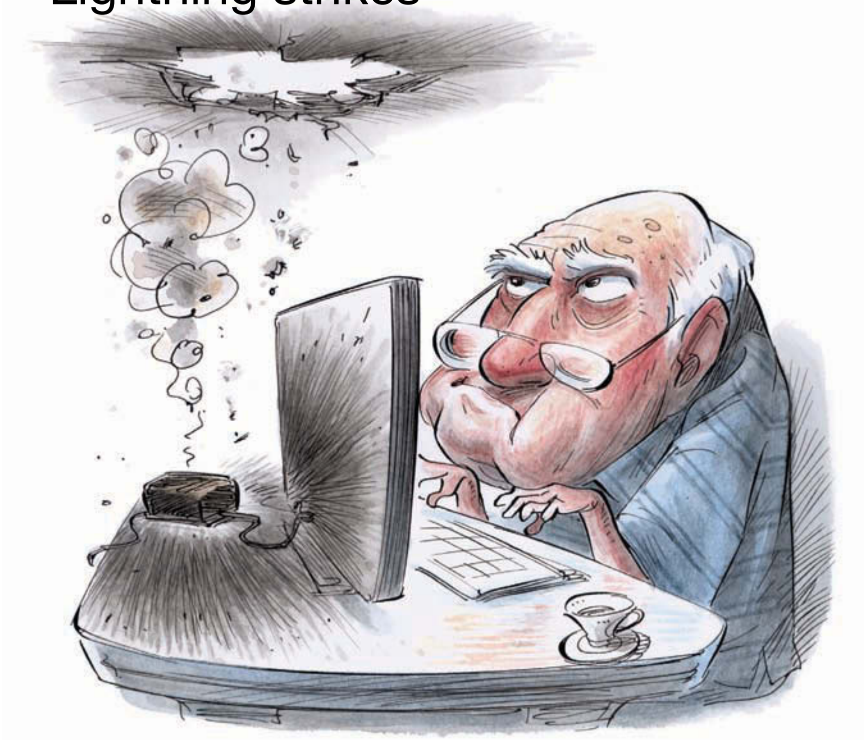

When I was working for a modem supplier in the 1980s, customers in the southeastern United States, particularly Florida, were returning the units after they failed catastrophically. The modems all came back with burned circuitry in an onboard SMPS (switch-mode power supply). There was not much left to analyze. Florida gets a lot of lightning storms, which can cause big problems for electronic devices. With modems, lightning normally causes damage in the telephone-line-interface circuitry, but most of our returns were due to failures of the switch-mode ac-power supplies. The problem usually involved a complete meltdown of the power supply and its components, making analysis of the root problem almost impossible. We performed lots of testing using simulated lightning strikes to the ac-power line, but the testing exposed no apparent problem.

It was time to change the testing strategy and look at lessons learned from previous experience. When issues occur, users are often instructed to turn a device off for 10 seconds or more before powering back on because powering back on too quickly sometimes causes circuits to misbehave. Taking enough time to fully reset to a known state is necessary to ensure a predictable power-on sequence. Lightning can induce surges on the ac-power line. If these surges are large enough, they can destroy electronic devices. However, it is more common for short interruptions of ac power to occur when protective devices in substations of the power company disconnect momentarily, resulting in a corresponding interruption of ac power to customers.

I set up a test by interrupting ac power for short durations to find out whether our power supply had such a problem. Everything was fine for interruptions lasting more than a second. As I continued to test with interruptions of less than a second, I found that, at approximately 750 msec of power-line interruption, our power supply would fail; heat was destroying many components around the main MOSFET's switch-mode transistor. It was difficult to analyze what had happened after the PCBs and their components had burned or melted. When I set the interruption time to less than 300 msec, however, the problem disappeared. It seemed that our product was susceptible to failures with ac-power-line interruptions of 300 to 750 msec.

With the help of a digital oscilloscope that could collect and record data until a trigger event stopped data collection and one that also could define complex triggering algorithms, we uncovered the root problem. Using multiple voltage and current probes on the SMPS MOSFET, we defined a trigger event for the time during which power dissipation in the transistor exceeded its rating. This task required creating a trigger event based on multiplying the instantaneous source-to-drain voltage by the drain current and calculating average power. By inducing ac-power interruptions shorter than 300 msec and gradually increasing the time of interruptions to 300 to 750 msec until failure occurred, the oscilloscope would trigger and stop data capture, which allowed analysis of conditions that occurred before the failure.

The oscilloscope showed that a circuit associated with ensuring an orderly shutdown after loss of ac power was operating in unexpected ways. The ac-power-line interruptions of the critical time allowed the main switch-mode transistor to enter a linear operating region, which caused enough dissipation to destroy it. The ensuing rush of current caused a general meltdown of surrounding circuitry, including the PCB.

Luckily the fix was simple. We soldered a discrete silicon diode to the back of the PCB at the factory, eliminating any future problems. EDN

Jim Sylivant is a professional engineer in Apex, NC.

DANIEL VASCONCELLOS

Isolated 12-Port PoE

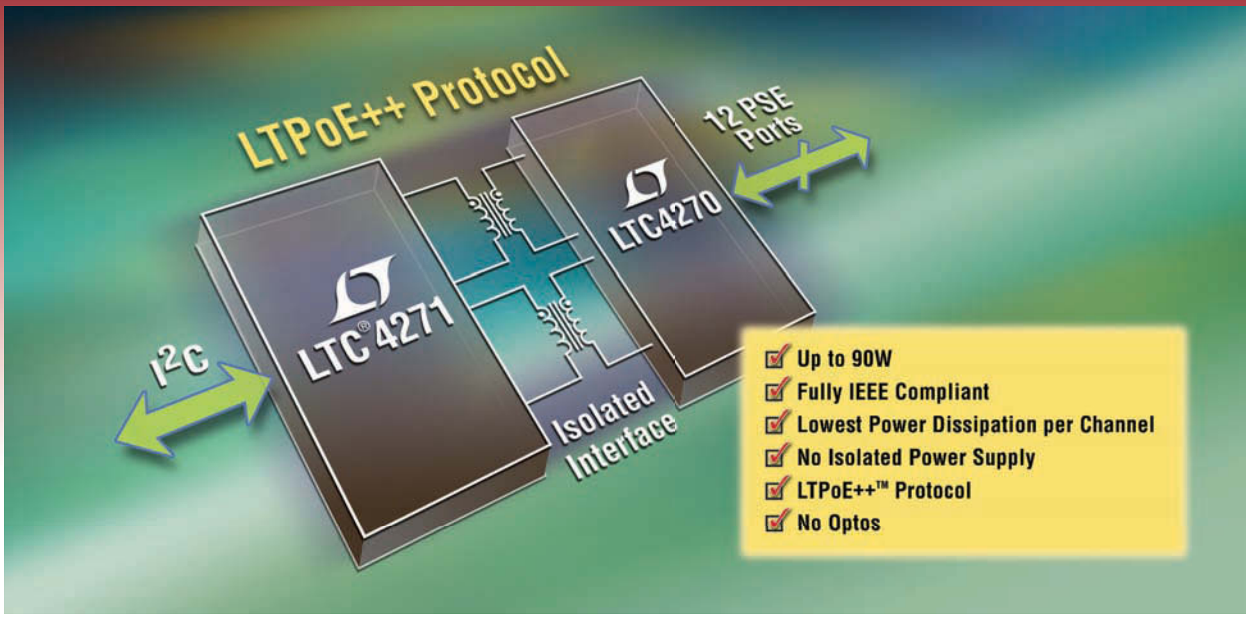

Highly Integrated Solution Dramatically Reduces BOM Cost

Linear's new 12-port PSE controller chipset for PoE, PoE+ and LTPoE++ systems significantly reduces Power over Ethernet system costs. The LTC®4270/4271 chipset's transformer-isolated communication protocol simplifies designs by replacing the expensive opto-couplers and complex isolated 3.3V supply used in traditional designs. LTPoE++ PSEs and PDs use Linear's proprietary signaling scheme to mutually identify and provide PD power up to 90W, while ensuring backward compatibility and interoperability with equipment based on IEEE PoE standards.

Features

- 12 PSE Channels Compliant with IEEE 802.3at Type 1 & 2
- Chipset Provides Electrical Isolation
  - Eliminates up to 6 High Speed Opto-Couplers
  - Eliminates Isolated 3.3V Power Supply
- Low Power Dissipation
  - 0.25Ω Sense Resistance per Channel
- Very High Reliability 4-Point PD Detection
- Available in Three Power Grades
  - A Grade: LTPoE++ 35W to 90W
  - B Grade: PoE+ 25.5W
  - C Grade: PoE 13W

Linear Technology 12-Port PSE Controllers

Part Number	Isolation	LTPoE++	PoE+	PoE	PD Power
LTC4271/70A	Transformer	•	•	•	90W
LTC4271/70B	Transformer		•	•	25.5W
LTC4271/70C	Transformer			•	13W
LTC4272A	External	•	•	•	90W
LTC4272B	External		•	•	25.5W
LTC4272C	External			•	13W

www.linear.com/PoE
1-800-4-LINEAR

LINEAR TECHNOLOGY