Reed Business Information (RBI™)

EDN®

VOICE OF THE ENGINEER

JULY 23

Issue 14/2009

www.edn.com

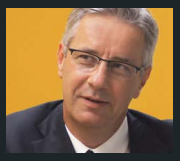

PROGRAMMABLE CHIPS

PIECING TOGETHER AN ANALOG SOLUTION

Page 28

OPTIMIZING RELIABILITY AND POWER EFFICIENCY IN EMBEDDED WIRELESS SYSTEMS

Page 23

MAKING ASIC POWER ESTIMATES BEFORE THE DESIGN

Page 36

Reach the next level

CST technology for SI/PI and EMC/EMI analysis

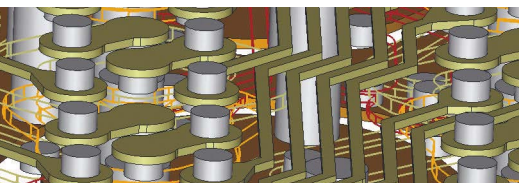

→ Get equipped for the job. CST's tools enable you to characterize, design and optimize electromagnetic devices all before going into the lab or measurement chamber. This can help save substantial costs especially for new or cutting edge products, and also reduce design risk and improve overall performance and profitability.

Solve your signal/power integrity and EMC/EMI applications by performing 2D and 3D analyses with CST's dedicated software.

For full channel characterization, use the fully integrated CST STUDIO SUITE tool set. CST MICROWAVE STUDIO® (CST MWS) is the market leading time domain tool for full-wave EM simulation. Use CST MWS to study signal integrity in full 3D for unprecedented accuracy. CST PCB STUDIO™ enables the simulation of entire board systems, including PCB, cable and connector. CST CABLE STUDIO™ lets you analyze complex cable harnesses with an unlimited number of cables. CST MICROSTRIPES™ EMC edition features advanced EMC functionality. Whatever your application is, CST has the technology to cover it.

Want to learn more? Take a look at why our customers have chosen CST technology: http://www.cst.com/Testimonials

→ Go further; draw your advantage from CST technology.

CST STUDIO SUITE – now even more choice for SI/PI and EMC/EMI simulation:

- CST MICROWAVE STUDIO®
- CST PCB STUDIO™
- CST CABLE STUDIO™
- CST MICROSTRIPES™

CST

CHANGING THE STANDARDS

MORE SEMICONDUCTORS THAN EVER*

Digi-Key® CORPORATION

The industry's broadest product selection available for immediate delivery

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

Visit digikey.com for the industry's broadest selection of semiconductors!

Authorized Distributor for more than 400 Supplier Partners.

www.digikey.com

1.800.344.4539

*Digi-Key is an authorized distributor for all supplier partners. New products added daily.

“I need the fastest DMM they make.”

They don’t make them any faster.

Nobody but Agilent makes a range of DMMs this fast, this accurate, or this reliable. Up to 1000 times more readings per second than the nearest competitor’s, and far easier to use, you’ll rip through tests in a fraction of the time. It’s what you’d expect from the leader in DMM technology.

DMM	Digits	DC Accuracy	Max Readings	Function/Range Changes	IO
34405A	5 ½	0.0250%	19 / sec	0.2 sec	USB
34401A	6 ½	0.0035%	1,000 / sec	.02 sec	GPIB, RS-232
34410A	6 ½	0.0030%	10,000 / sec	2.6 ms	GPIB, USB, LAN (LXI)
34411A/ L4411A	6 ½	0.0030%	50,000 / sec	2.6 ms	GPIB, USB, LAN (LXI)
34420A	7 ½	0.0030%	250 / sec	.02 sec	GPIB, RS-232
3458A	8 ½	0.0008%	100,000 / sec	3.0 ms	GPIB

Download the latest measurement brief and tips
www.agilent.com/find/fastestdmm

Agilent Authorized Distributors

Allied Electronics
800-433-5700

Davis Instruments
800-358-5525

Agilent Technologies

EDN 7.23.09 contents

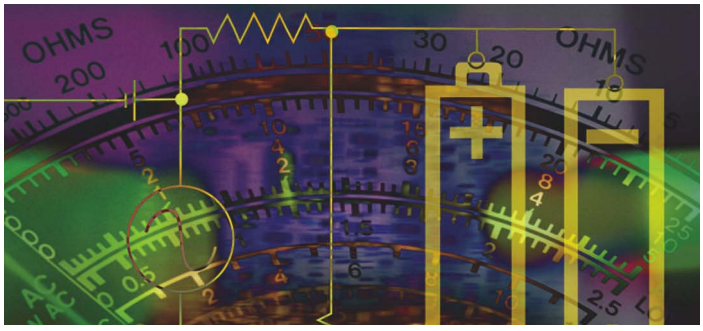

Programmable chips: piecing together an analog solution

Optimizing reliability and power efficiency in embedded wireless systems

Making ASIC power estimates before the design

pulse

DESIGNIDEAS

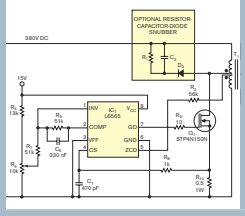

The Newest Optoelectronics

OSRAM
PANJIT TOUCH SCREENS
VISHAY
OPTEK
LEDENGIN
KINGBRIGHT
AVAGO
LITE-ON
LUMEX
LAMINA
SHARP
NORITAKE
TAOS
BIVAR
VCC
PACIFIC SILICON SENSORS
3M TOUCH SYSTEMS
NEWHAVEN DISPLAY

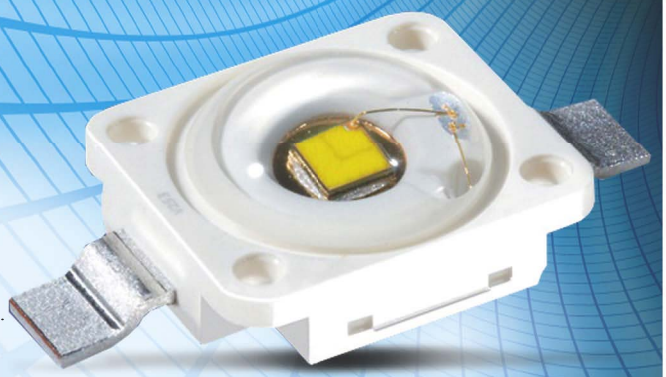

OSRAM
Opto Semiconductors

Golden DRAGON Plus LED
www.mouser.com/osram_golden_dragon_plus

www.mouser.com
Over A Million Products Online

New Products from:

Avago Technologies
Moonstone™ 3W Power LED Light Source
www.mouser.com/avagomoonstoneled

FDI Future Designs, Inc. Your Development Partner
ARM 5.7" Touch Screen LCD Kit
www.mouser.com/futuredesignsinc/a

Tyco Electronics
Authorized Distributor
Star Board LED Pixel Holder
www.mouser.com/Tyco_StarBoard

The ONLY New Catalog Every 90 Days

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

MOUSER ELECTRONICS®
a tti company

The Newest Products For Your Newest Designs

(800) 346-6873

contents 7.23.09

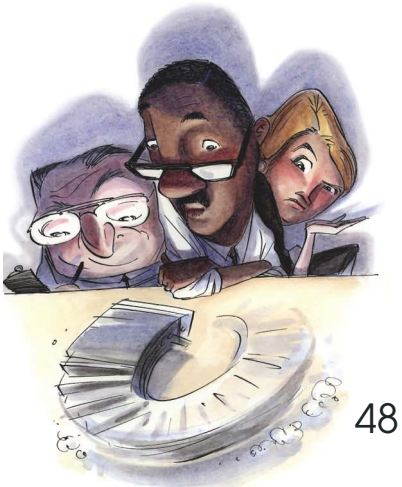

DEPARTMENTS & COLUMNS

EDN online contents www.edn.com

ONLINE ONLY

Check out these Web-exclusive articles:

Mobile-device security through virtualization
A hypervisor running in privileged mode allows a more flexible approach, providing security for mobile devices.
→www.edn.com/article/CA6668320

Applications and considerations of capacitive-proximity sensing
Capacitive-proximity sensing is a clever way to interact with user interfaces without having to physically touch the interface controls. It is based on the principles of capacitive sensing, which detect the presence or absence of a conductive object. Capacitive-proximity sensing offers several unique benefits to system and industrial designers who can effectively use it to increase device usability and aesthetics and allow increased responsiveness for devices that require configuration time at wake-up.
→www.edn.com/article/CA6666196

designideas

READERS SOLVE DESIGN PROBLEMS

The Design Ideas section (pg 40 in this issue) remains one of *EDN*'s most popular. Check out the latest Design Ideas and search our archives. You never know; you might find a solution to one of your problems.
→www.edn.com/designideas

FEED YOUR NEED

EDN's technical editors highlight notable new products, including analog and digital ICs, power components, passives, boards, and systems. To submit products for consideration, please see the information on our Editorial Opportunities page.
→www.edn.com/contribute

EDN® (ISSN#0012-7515), (GST#123397457) is published biweekly, 24 times per year, by Reed Business Information, 8878 Barrons Blvd, Highlands Ranch, CO 80129-2345. Reed Business Information, a division of Reed Elsevier Inc, is located at 360 Park Avenue South, New York, NY 10010. Tad Smith, Chief Executive Officer; Mark Finkelstein, President, Boston Division. Periodicals postage paid at Littleton, CO 80126 and additional mailing offices. Circulation records are maintained at Reed Business Information, 8878 S Barrons Blvd, Highlands Ranch, CO 80129-2345. Telephone (303) 470-4445. POSTMASTER: Send address changes to EDN®, PO Box 7500, Highlands Ranch, CO 80163-7500. EDN® copyright 2009 by Reed Elsevier Inc. Rates for nonqualified subscriptions, including all issues: US, $179.99 one year; Canada, $229.99 one year (includes 7% GST, GST#123397457); Mexico, $229.99 one year; air expedited, $399.99 one year. Except for special issues where price changes are indicated, single copies are available for $10 US and $15 foreign. Publications Agreement No. 40685520. Return undeliverable Canadian addresses to: RCS International, Box 697 STN A, Windsor Ontario N9A 6N4. E-mail: Subsmail@ReedBusiness.com. Please address all subscription mail to EDN®, 8878 S Barrons Blvd, Highlands Ranch, CO 80129-2345. EDN® is a registered trademark of Reed Elsevier Properties Inc, used under license. A Reed Business Information Publication/Volume 54, Number 14 (Printed in USA).

Buck regulators with integrated MOS

High efficiency

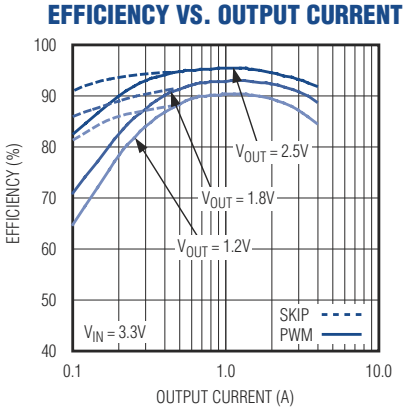

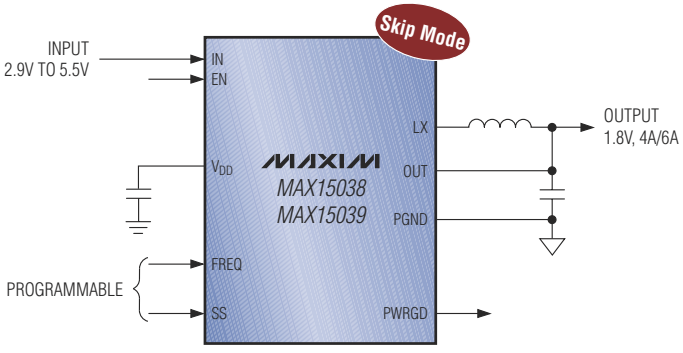

1- to 8-phase buck regulator with 25A/phase in a 6mm x 6mm TQFN

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

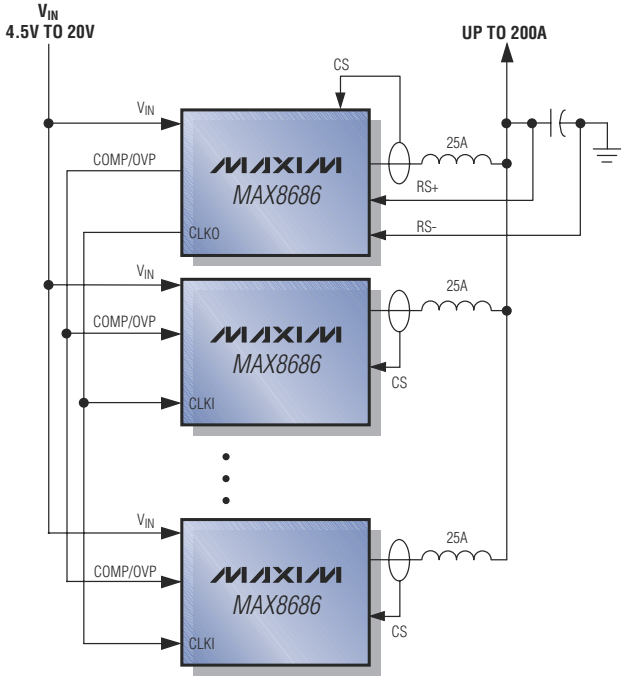

FETs simplify designs up to 25A/phase

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

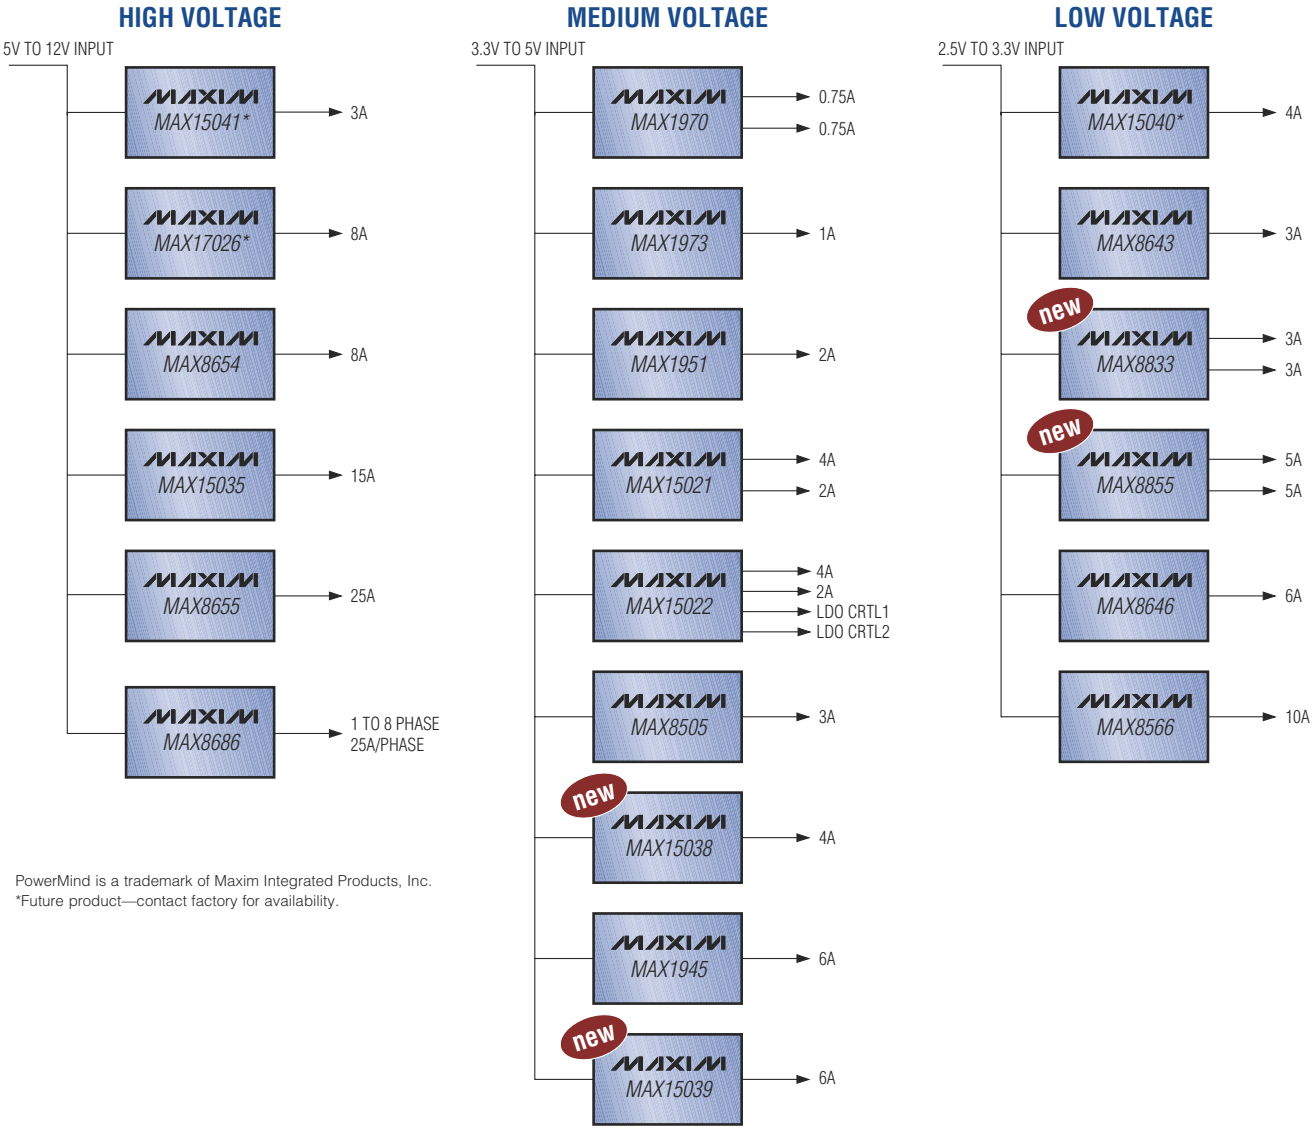

PowerMind is a trademark of Maxim Integrated Products, Inc.
*Future product—contact factory for availability.

www.maxim-ic.com/buck-regulators

MAXIM DIRECT
www.maxim-ic.com/shop

AVNET electronics marketing
www.avnet.com

MAXIM INNOVATION DELIVERED™

For free samples or technical support, visit our website.

BY PAUL RAKO, TECHNICAL EDITOR

Three killer apps and some not-so-killer

As I visited the IEEE International Microwave Symposium last month in Boston, I saw a bright future for semiconductors and electronics. For years, pundits have been bleating about the next killer application for semiconductors. Well, I see three of them coming down the pike: smartphones, LED lighting, and solid-state drives.

In smartphones, the iPhone changed everything, driving not only phone-replacement sales, but also the entire cellular infrastructure. AT&T was astonished at the amount of data that people consumed and sent with their iPhones. This phenomenon means that we need more cellular-base-station infrastructure; more data infrastructure; more touchscreens; and a massive amount of analog for cameras, flash drivers, and all the other things an advanced phone requires. In some countries, such as India, they are passing over the laptop era and jumping straight to smartphones and their cousins, netbooks.

Fear not: There will be a huge demand for analog engineers.

As for LED lighting, sure, it now pays to use it only for refrigerated coolers and 30-foot-high, difficult-to-replace architectural lighting. The semiconductor companies never cease to amaze us with the progress they make, however. Take Cree, for example. It learns about and spends on gallium nitride for its LEDs, and those resources will benefit its RF- and cellular-base-station division. Soon, LEDs will appear in general lighting, and that debut won't be big. It won't even be huge. It will be monstrous. If they can provide the energy savings of CFLs (compact fluorescent lights), use no mercury, and have lifetimes longer than one year, it will be all good for the environment.

The third killer app is solid-state-drive technology, which will replace many hard-disk drives, including those for operating systems and anything that needs fast, reliable, low-power, or shockproof data, such as the earlier-noted smartphone. This development does not help my analog pals too much, but there will be a mind-boggling amount of purchased, processed, and sold silicon for this application. Check out a video at www.youtube.com/watch?v=96dWOEa4Djs to see a computer with 24 Samsung solid-state disks. Sure, it may represent $16,000 worth of disks today, but, like LEDs and the power supplies for them, disk prices will decrease dramatically in a few years.

So there you have it, three killer apps. Let the celebration begin. In a couple of years, there will be a boom in Silicon Valley the likes of which we have never before seen. Massive amounts of silicon need massive amounts of semiconductor machinery, so Applied Materials will be doing just fine. People say that solar power will also be a big application, for which Sunpower is in a better position than Nanosolar and all the other companies making cheap printable panels using CIGS (copper-indium-gallium-selenide) semiconductors. If you are going to pay $20,000 for someone to put panels on your roof, they had better be 22%-efficient Sunpower panels and not 11%-efficient thin-film panels. Some people claim 19.5% efficiency for CIGS, but that figure comes from data in a laboratory setting. I will believe the technology achieves that efficiency in other settings when *Consumer Reports* reports on it or it appears in the Digi-Key catalog.

As for the whole alternative-energy thing, I predict it will fizzle for five years until gasoline prices again increase to more than $4 and carbon taxes drive electricity prices to 25 cents per kilowatt, about twice those in California and 3.5 times those in West Virginia. Although Google claims to have a whizz-bang solar-thermal setup, the ecology freaks won't let us build power lines to the installation, so what good is it? The same goes for wind. Solar panels on our roofs are the most practical things we have going and are additional reasons for the boom in semiconductors and the machinery that makes them. Remember: LEDs need power supplies, and solar power needs inverters. The capacitor and inductor makers will have great futures, and, fear not: There will be a huge demand for analog engineers. EDN

Contact me at paul.rako@edn.com.

How to pick the perfect inductor for your LED driver application

IC reference designs are a good start. But what if you want to optimize the driver inductor for size, efficiency or price?

Or evaluate newer, high performance parts that weren't available when the reference design was created?

Then you'll want to check out the new LED Design Center on the Coilcraft web site. It's filled with interactive tools that let you compare hundreds of inductors for all LED driver topologies, including SEPIC.

Our new LED Design Center lets you:

- ✔ Search by IC to find all matching inductors
- ✔ Compare DCR, current rating, size and price
- ✔ Analyze all core and winding losses
- ✔ Request free evaluation samples

www.coilcraft.com/LED

To start your search for the perfect LED driver inductor, mouse over to **www.coilcraft.com/LED**

ORDER DIRECT
BUY.COILCRAFT.COM
800-322-2645
OVERNIGHT DELIVERY ORDER BY 5
RoHS COMPLIANT

Coilcraft®

www.coilcraft.com 800/322-2645

XILINX

RISE ABOVE

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

EDN

PRESIDENT, BOSTON DIVISION, REED BUSINESS INFORMATION
Mark Finkelstein, mark.finkelstein@reedbusiness.com
1-781-734-8431

PUBLISHER, *EDN* WORLDWIDE
Russell E Pratt, 1-781-734-8417;
rpratt@reedbusiness.com

ASSOCIATE PUBLISHER, *EDN* WORLDWIDE
Judy Hayes, 1-925-736-7617;
judy.hayes@reedbusiness.com

VICE PRESIDENT, EDITORIAL DIRECTOR
Karen Field, 1-781-734-8188;
kfield@reedbusiness.com

EDITOR-IN-CHIEF, *EDN* WORLDWIDE
Rick Nelson, 1-781-734-8418;
rnelson@reedbusiness.com

EXECUTIVE EDITOR
Ron Wilson, 1-510-744-1263;
ronald.wilson@reedbusiness.com

MANAGING EDITOR
Amy Norcross
1-781-734-8436;
fax: 1-720-356-9161;
amy.norcross@reedbusiness.com
Contact for contributed technical articles

SENIOR ART DIRECTOR
Mike O'Leary
1-781-734-8307;
fax: 1-303-265-3021;
moleary@reedbusiness.com

ANALOG
Paul Rako, Technical Editor
1-408-745-1994; paul.rako@edn.com

MASS STORAGE, MULTIMEDIA, PCs, AND PERIPHERALS
Brian Dipert, Senior Technical Editor
1-916-760-0159;
fax: 1-303-265-3187;
bdipert@edn.com

MICROPROCESSORS, DSPs, AND TOOLS
Robert Cravotta, Technical Editor
1-661-296-5096;
fax: 1-303-265-3116;
rcravotta@edn.com

NEWS
Suzanne Deffree, Managing Editor
1-631-266-3433;
sdeffree@reedbusiness.com

POWER SOURCES, ONLINE INITIATIVES
Margery Conner, Technical Editor
1-805-461-8242;
fax: 1-805-461-9640;
mconner@reedbusiness.com

DESIGN IDEAS EDITOR
Martin Rowe,
Senior Technical Editor,
Test & Measurement World
edndesignideas@reedbusiness.com

SENIOR ASSOCIATE EDITOR
Frances T Granville
1-781-734-8439;
fax: 1-303-265-3131;
f.granville@reedbusiness.com

EDITORIAL/WEB PRODUCTION
Diane Malone, Manager
1-781-734-8445; fax: 1-303-265-3024
Steve Mahoney,
Production/Editorial Coordinator
1-781-734-8442; fax: 1-303-265-3198
Melissa Annand,
Newsletter/Editorial Coordinator
1-781-734-8443; fax: 1-303-265-3279
Adam Odoardi, Prepress Manager
1-781-734-8325; fax: 1-303-265-3042

CONSULTING EDITOR
Jim Williams, Staff Scientist,
Linear Technology

CONTRIBUTING TECHNICAL EDITORS
Dan Strassberg,
strassbergedn@att.net
Nicholas Cravotta,
editor@nicholascravotta.com

COLUMNISTS
Howard Johnson, PhD, Signal Consulting
Bonnie Baker, Texas Instruments
Pallab Chatterjee, SiliconMap

PRODUCTION
Dorothy Buchholz,
Group Production Director
1-781-734-8329
Joshua S Levin-Epstein,
Production Manager
1-781-734-8333; fax: 1-781-734-8096

EDN EUROPE
Graham Prophet, Editor, Reed Publishing
The Quadrant, Sutton, Surrey SM2 5AS
+44 118 935 1650;
fax: +44 118 935 1670;
gprophet@reedbusiness.com

EDN ASIA
Luke Rattigan, Chief Executive Officer
luke.rattigan@rbi-asia.com
Kirtimaya Varma, Editor-in-Chief
kirti.varma@rbi-asia.com

EDN CHINA
William Zhang,
Publisher and Editorial Director
wmzhang@idg-rbi.com.cn
Jeff Lu, Executive Editor
jefflu@rbichina.com.cn

EDN JAPAN
Katsuya Watanabe, Publisher
k.watanabe@reedbusiness.jp
Ken Amemoto, Editor-in-Chief
amemoto@reedbusiness.jp

EDN, 225 Wyman St, Waltham, MA 02451. **www.edn.com**. Phone 1-781-734-8000. **Address changes or subscription inquiries:** phone 1-800-446-6551; fax: 1-303-470-4280; subsmail@reedbusiness.com. For a free subscription, go to **www.getfreemag.com/edn**. Reed Business Information, 8878 S Barrons Blvd, Highlands Ranch, CO 80129-2345. Include your mailing label.

Reed Electronics Group

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

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

Mentor Graphics®

THE EDA TECHNOLOGY LEADER

EDITED BY FRAN GRANVILLE

pulse

INNOVATIONS & INNOVATORS

DAC shapes sound for small speakers

Wolfson Microelectronics' new single-chip, monophonic WM9081 DAC and speaker amplifier deliver optimum sound-pressure levels and maximum intelligibility from low-cost speakers, such as those in portable navigation devices, mobile handsets, digital radios, and conference speakerphones. The IC has a high-efficiency, 2.6W speaker amplifier that can operate in either Class AB or Class D, a dynamic-range controller, and a five-band ReTune Mobile parametric equalizer. SNR (signal-to-noise ratio) is 92 dB in Class D mode, and THD+N (total harmonic distortion plus noise) is 1% at 2.4W output power into a 4Ω speaker load. The 9081 uses 30% less board area and costs 20% less in bill-of-materials expenses than a two-chip approach.

The WM9081 DAC and speaker amplifier deliver optimum sound-pressure levels and maximum intelligibility from low-cost speakers, such as those in portable navigation devices and mobile handsets.

The integrated ReTune Mobile equalizer has fully programmable coefficients, allowing the user to define speaker-frequency-response compensation and music profiles. Programmable dynamic-range control boosts low-level signals to maximize volume for improved intelligibility. The chip comes in a 28-pin COL QFN package; an evaluation system includes a board, schematics, a Linux driver, and a downloadable software GUI (graphical user interface), which allows the developer to modify an interactive version of the register map in real time and modify the device configuration while developing software settings.

—by Graham Prophet

▷**Wolfson Microelectronics**, www.wolfsonmicro.com.

FEEDBACK LOOP

"The hurt-proof playground is another example of the liberal mindset that an individual is not responsible for the result of his actions and should therefore bear no consequences from such actions. This is part of the dumbing down of society, which, unfortunately, is also a source of wealth for those in the litigation profession. It is really tragic to see our country and our world going this way."

—Engineer and *EDN* reader William Ketel, in *EDN*'s Feedback Loop, at www.edn.com/article/CA6648785. Add your comments.

Touch sensors allow for variations in parasitic capacitance

Capacitive-touch sensors can detect a finger touch by measuring capacitive changes on the touchpad. Currently, the most popular form of capacitive-touch sensors relies on a microprocessor and requires programming before use. SMSC's new CAP10xx capacitive-sensor family takes an alternative hardware-based approach that uses dedicated registers. You set these registers during a three-step tuning process that determines the system sensitivity, sets the sensor threshold to zero for the "no-touch" value, and then verifies any change in values for no-touch states. This process takes place at power-on and can adapt to system changes, such as dirt buildup on switches and manufacturing variability, including changing parasitic capacitance, in PCBs (printed-circuit boards). This tolerance allows the sensors' use with products from multiple vendors and processes.

The CAP10xx family offers an 8-kV rating, an advantage in consumer electronics. The chips also provide built-in noise filtering at the input pins, which wireless circuits or dc/dc-power regulators on the same board can cause. Options include the number of sensor inputs, slider-input capability, and LED-driver circuitry. All chips have an I^2C-bus interface; some provide an SPI (serial-peripheral interface) or a BC-Link interface. Chip packages range from a 10-pin, 3×3-mm DFN to a 32-pin, 5×5-mm QFN. Production prices range from 75 cents to $2 (10,000).**—by Margery Conner**

▷**SMSC**, www.smsc.com.

The hardware-based CAP10xx family requires no programming to set proximity-range values.

Network analyzer for active-device test adds 13.5-, 43.5-, and 50-GHz models

Agilent Technologies has added 13.5-, 43.5-, and 50-GHz models of its PNA-X network analyzer. The 43.5- and 50-GHz models support high-frequency applications, such as radar and satellite communications, whereas the 13.5-GHz model handles low-frequency devices in wireless communications. The company has also added 13.5-, 43.5-, and 50-GHz models of its PNA-X NVNA (nonlinear-vector-network analyzer). The range of these new products provides engineers who develop and manufacture active devices the flexibility to select the right instrument bandwidth for each application.

Aerospace and defense engineers who work on systems with frequencies as high as 50 GHz can now benefit from Agilent's single-connection, multiple-measurement PNA-X. This analyzer's integrated measurements, versatile hardware, and reconfigurable measurement paths address this market's key challenges: test-system costs, test complexity, throughput, accuracy, and equipment-space requirements. Typical radar, satellite, and electronic-warfare applications require complex test systems that occupy multiple instrument racks and require numerous connections to the device under test.

The PNA-X integrates the capabilities of a full rack of equipment into a single instrument, simplifying test stations, halving the number of instruments, and increasing throughput fourfold. The PNA-X also offers a single-contact approach to wafer tests. By eliminating multiple probe contacts and enabling the most accurate characterization and reliable wire-bonding, this approach significantly improves the quality of the devices produced. Moreover, engineers working on systems with frequencies greater than 50 GHz can use the PNA-X to configure a banded millimeter-wave system that operates to 0.5 THz.

The 13.5-GHz PNA-X model supports wireless communications, in which reduced test time, number of test stations, and test cost are critical. Low-noise-amplifier tests typically require separate test stations for such characteristics as small-signal gain/match, distortion, and noise figure; the PNA-X integrates these tests into one station, thereby reducing the number of stations by as much as 75% and costs by 30%.

The NVNA provides the greatest possible insight into nonlinear-device behavior. The NVNA proves especially useful for scientists who research new RF technologies and engineers who design high-performance active devices. Using NVNA, you measure X parameters and then use them to create models that you can import into Agilent's ADS (Advanced Design System) to simulate the behavior of linear and nonlinear physical components.

The configurable two- or four-port network analyzer uses a single connection for measurements of continuous-wave and pulsed scattering parameters, compression, IMD (intermodulation distortion), and noise figure. Two built-in high-performance sources offer output power of 16 dBm, harmonics of −60 dBc, and a power-sweep range of 40 dB. A wide range of measurement applications for amplifiers, converters, and modules with linear and nonlinear characteristics includes vector-noise figure, gain compression, IMD, true-differential measurements, and NVNA. Internal signal-routing switches provide increased flexibility for adding signal-conditioning hardware or additional test equipment for single-connection measurements. Internal pulse modulators and generators for fast, simplified pulse measurements permit making these measurements in 1/30 of the time necessary with analyzers that require external generators and modulators.

Prices for PNA-X start at $76,400 for a 10-MHz to 13.5-GHz unit and extend to $140,100 for a 10-MHz to 50-GHz unit. Prices for NVNA options start at $37,400. A 50-GHz comb generator, which acts as a phase reference for the NVNA, adds $16,000.

—by Dan Strassberg

▷**Agilent Technologies**, www.agilent.com/find/pna-x and www.agilent.com/find/nvna.

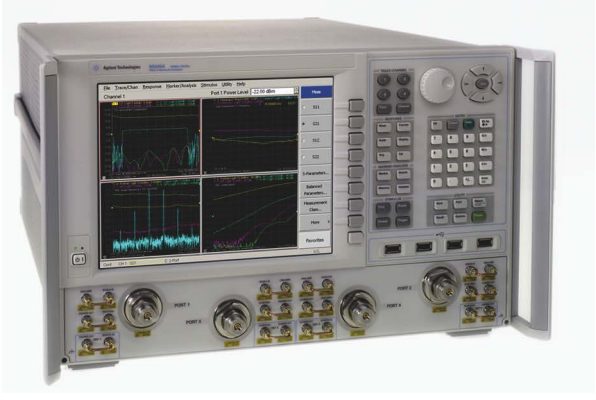

The PNA-X series of two- and four-port vector-network analyzers and nonlinear-vector-network analyzers for characterization and test of active and passive RF, microwave, and wireless-communication devices now includes 13.5-, 43.5-, and 50-GHz models.

DILBERT By Scott Adams

Isolated USB connections operate in medical and industrial equipment

Analog Devices has introduced a single-IC USB (Universal Serial Bus) isolator, which simplifies isolated USB-port implementation in system-critical medical and industrial equipment. USB ports provide a standardized, straightforward way to connect and disconnect peripheral devices to and from a computer without rebooting or turning off a system. When medical professionals need information, USB-compatible medical devices make it simple and efficient to share vital patient files or data at any time and between locations. USB-enabled devices enable health-care professionals to use a range of commercially available, cost-efficient peripherals. The demand for USB-enabled health-care devices continues to grow because equipment for patient monitoring, disease management, health and wellness, and drug delivery is now essential in maintaining the health of millions of people worldwide.

Analog Devices hopes to address these issues with the ADuM4160 USB isolator. It offers 5-kV-rms medical-grade isolation, upstream short-circuit protection, and fully isolated 1.5- and 12-Mbps data rates. IEC (International Electrotechnical Commission) 60601-1 medical-safety approvals are pending. The USB 2.0-compliant device provides fully isolated USB functions and enables medical- and industrial-design engineers to cut implementation costs by as much as 25%, reduce the size of designs by as much as 50%, and trim development time from months to weeks, according to the company.

The isolator employs the company's iCoupler technology, which combines high-speed CMOS and chip-scale-microtransformer technology. Designers can implement ADuM4160 in low- and full-speed USB-compliant systems, and it operates from the 5V USB supply or system-supplied 3.3V power using an internal regulator. The ADuM4160 also provides isolated control of the pullup resistor, allowing the peripheral to control connection timing. The device's maximum current during idle mode is 2 mA, eliminating the need for a suspend state.

"Medical designs have typically used isolated RS-232 or Ethernet connections that limit data-transfer rates, do not support plug-and-play functions, and take up excessive space," says Patrick O'Doherty, health-care-segment director at Analog Devices. "The ADuM4160 USB isolator's reinforced, medical-grade isolation makes it possible for engineers to design USB-enabled medical devices that help facilitate better patient care." Base price for the 16-lead SOIC device is $4.89 (1000).**—by Paul Rako**

▷**Analog Devices Inc**, www.analog.com.

07.23.09

LDPC TECHNOLOGY COMES TO DISK-READ CHANNELS

LSI recently launched a new era in disk-read-channel technology with the RC9500, a mixed-signal read-channel-IP (intellectual-property)-core cluster. LSI intends the block for integration with a drive vendor's IP to create a single-chip drive-electronics subsystem. The IP is LSI's–and, possibly, the industry's–first venture into 40-nm process technology, and it appears to be the first application of an LDPC (low-density-parity-check) algorithm in read-channel products. LSI intends the IP for the coming generation of 500-Gbyte/disk, 2.5-in. and 1-Tbyte/disk, 3.5-in. drives.

The 40-nm part allows the read channel to reach 4 Gbps within the power budget customers specify. At this speed, even minimal parasitics and small variations can sneak up on designers. Compounding the design problem for analog designers is the fact that the whole game in read channels is about noise. The head and media designers will push areal density until every hint of noise margin is gone, so there's nothing left over for the low-noise-amplifier designers. But low-noise design at 40 nm–especially in IP, in which you can't control the charge other circuits might be pumping into the substrate next door–is an art. It would be easy for circuit designers to lose the whole advantage of the new LDPC algorithm in the analog section.

The 40-nm process also presents issues to LSI's customers. LSI's interface to customers using this IP will be its familiar ASIC model. But everything will go better if customers bring to the table a design team and IP that are ready for the rigors of a 40-nm design. The other milestone is LDPC, a signal-recovery technique that has for years found use in such areas as satellite communications, in which SNRs (signal-to-noise ratios) are small and no one minds that the receiver fills a rack and runs slower than real time. Getting the algorithm into a manageable gate count and power consumption and getting it to keep up with a 64-Gbps bit stream are challenges. But LSI officials claim that LDPC provides at least 1 dB of improvement in SNR–a major increment in disk drives.

The algorithm required some new architectural thinking compared with previous read-channel DSP designs. LSI designers also worked on adapting the algorithm to the signal characteristics they expected from the next generation of head and media designs, which are themselves still evolving.

LSI is now offering an evaluation chip for sampling that carries the whole cluster of IP, from the analog front end through the detector-decoder block, for evaluation by key customers.—by Ron Wilson

▶**LSI, www.lsi.com.**

Thermal-management software reduces engineering time, speeds product development

Ansys Inc has announced the latest release of Ansys Icepak software, which applies fluid-dynamics technology to electronics thermal management. The 12.0 release introduces new features for PCB (printed-circuit-board) and package thermal analysis, new and enhanced technology for complex meshing geometry, and new physical-modeling capabilities. These features enable engineers who design electronic components for products such as cell phones, computers, and graphics cards to improve design performance, reduce the need for physical prototypes, and shorten time to market.

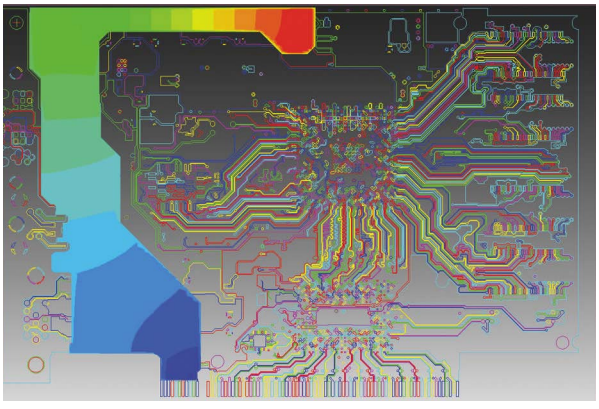

The Ansys Icepak 12.0 release introduces new features for PCB thermal analysis.

Ansys Icepak software accelerates product development by simulating the dissipation of thermal energy in electronic devices at the component, board, or system level. Based on the Ansys Fluent CFD (computational-fluid-dynamics) solver, Ansys Icepak software has a streamlined user interface that enables the rapid creation of models of complex electronic assemblies. Integration of Ansys Iceboard and Ansys Icechip capabilities into Ansys Icepak 12.0 provides engineers with an integrated platform to analyze package, PCB, and system designs. Direct importing of ECADs (electronic-computer-aided designs) allows engineers to model all components of packages and PCBs, including, traces, vias, solder balls, solder bumps, wire bonds, and dice (see "SIwave 4.0 addresses signal and power integrity," *EDN*, June 17, 2009, www.edn.com/article/CA6665970). A new PCB-trace joule-heating modeling capability, along with the ability to import dc-power-distribution profiles from Ansoft SIwave software, enhances the accurate thermal simulation of PCBs. Icepak 12.0 also offers enhanced macros; enhanced fan-modeling capabilities; parallel processing; postprocessing; and new libraries for heat sinks, thermoelectric coolers, and materials. In keeping with the Ansys multiphysics strategy, Ansys Icepak technology, in conjunction with SIwave and Ansys Mechanical products, provides a coupled approach to the electrical, thermal, and structural requirements of electronics-design engineers.

"With today's high-performance electronic devices, there is a trend to reduce device size while increasing product functionality," says Dipankar Choudhury, vice president of corporate product strategy and planning at Ansys. "This trend increases power densities in devices, which necessitates that thermal management become a design driver."

Ansys Icepak 12.0 employs two new meshing technologies—automatic multilevel meshing and Cartesian hex-dominant meshing—to enhance the software's ability to handle complex geometries and improve accuracy without sacrificing robustness. The meshing enhancements improve mesh smoothness, quality, curvature, proximity capturing, and speed. The automatic generation of accurate, conformal meshes that represent the true shape of electronic components and the solution of fluid flow and all modes of heat transfer—conduction, convection, and radiation—for both steady-state and transient thermal-flow simulations enable engineers to rapidly evaluate thermal-management issues for electronic devices, reducing development time and increasing product reliability.

Ansys Icepak 12.0 is part of the Ansys multiphysics portfolio, which also includes Ansoft electromagnetic, electromechanical, circuit, and system-behavior technologies (see "Simulation gets speed, capacity boost," *EDN*, Jan 22, 2009, pg 26, www.edn.com/article/CA6629472).

—by Rick Nelson

▷**Ansys Inc**, www.ansys.com.

ULTRA-LOW-POWER MICROCONTROLLER BRINGS ACCURACY TO SMART METERS

07.23.09

Targeting the smart-meter market, Texas Instruments has added the MSP430F471xx to its 430 family of ultra-low-power microcontrollers. The device has accuracy better than 0.1% and a dynamic range of 2400 to 1, necessary for the wide range of high and low currents that power meters encounter. Current electromechanical meters, in contrast, have only about 2% accuracy. Low-power operation is important for electronic three-phase power meters because they must run on battery-backup power during outages.

In its low-power mode, the 471xx can keep an LCD and a real-time clock alive for continued metering and reporting. Six sigma-delta ADCs monitor the three power phases, and a seventh monitors the ground return to see whether the equipment is drawing power or even if someone has tampered with the magnetic pickup. The device has multiple interfaces, including PLC (power-line communications) to communicate with the utility, and an RF interface for advanced meter infrastructure that allows in-home wireless-network communication. The device has as much as 120 kbytes of flash memory and 8 kbytes of RAM. The chip sells for $5.75 to $7.95 (1000).

—by Margery Conner

▶**Texas Instruments, www.ti.com.**

VOICES

NXP's René Penning de Vries:
Moving beyond Moore

René Penning de Vries, senior vice president and chief technology officer of NXP Semiconductors (www.nxp.com), spoke to *EDN* about how design and R&D are changing as we evolve past the traditional definitions of Moore's Law and into a new era based on value-added applications and guided more than ever before by economics. Excerpts of that interview follow.

Where do we as an industry stand with Moore's Law? Are we coming to its end?

A No, I don't see that at all. I think Moore's Law is going to have a future but only for the applications that have such a high volume that the investment associated with the next generations can be earned back. Saying that limits the application to very few that we all know about and thereby [limits Moore's Law to] very few players that we all know about. Of course, it's about the memories, it's about the big processors, and it may be about the basebands in wireless handsets. Those are the applications that can create the accumulative business that justifies the enormous investment. More and more, companies must find for themselves a place where they can make business, and, in some cases, that [idea] leads to the conclusion that that [place] is not Moore's Law or not an application based on Moore's Law.

So applications become the next industry driver.

A Yes. What you will see more and more is that the big concerns that the world is facing at this moment—the need for sustainable energy, the need for water, the need to reduce the cost of medical care—those challenges can only be realized as more and more we create intelligent solutions. Intelligent solutions in my view means optimized solutions for a particular application area. That [idea] is not necessarily based on Moore's Law technologies. On the contrary, often that [idea] will ask for technologies that can deal with high voltages, technologies that can deal with special requirements. Typically, these are technologies that are based on a relatively old CMOS generation but with a sophisticated device or a sophisticated addition on top of it that makes this specialty an option. … Technically, I do not believe Moore's Law is over. It is economics. There is a famous analogy in the industry that relates to commercial airplanes. The Concorde flew higher than the speed of light between Paris and New York. Technically, it was a very feasible proposition. Economically, it didn't work out. It was just too expensive. That's why, for commercial applications, we don't see the Concorde anymore. The development in itself did not limit evolution in the airplane industry. We saw airplanes become more efficient, less noisy, less polluting, more pleasant to travel in. There has been a continued line of innovation, not along the lines of making it faster, but along the lines of having a smaller footprint.

Considering the economy and this evolution to a "more than Moore" era, where are we in R&D and spending on R&D?

A It will change in nature. Whereas we had massive amounts of money spent in developing new technology, we now see the wealth of applications coming toward us in dedicated solutions. It will move from the foundation process technology into the area of application-specific, customer-optimized solutions. … We'll shift more toward customers, toward applications and move away from very deep process technology.

Will it be more difficult to innovate with reduced R&D investments and smaller staffs? There have been massive layoffs in the industry over the last year.

A For sure, it is a stretch. But we also know that, after the crisis, there is going to be an upturn again. To me, it's a matter of endurance. We have to sustain innovation. We have to continue to invest people, money, and brains in creating innovations that later on we can harvest.

Do you have an estimate on when that upturn will come?

A No. The only thing I know is that every day it is nearer by 24 hours.

–interview conducted and edited by Suzanne Deffree

BY BONNIE BAKER

Using an analog filter to inject noise

Sometimes things just don't make sense! For instance, your RC filter or amplifier's lowpass filter at the input of a delta-sigma ADC can produce a noisier digital output. Didn't you design the filter to reduce noise so that you'd get more instead of fewer noiseless bits from your converter? It is as easy to eliminate higher-frequency noise with an analog lowpass filter as it is to inject noise into the frequency band below the corner frequency of your filter. If your filter produces noise in the frequency band of interest, your conversion output results will be noisier than you might expect.

If you change your circuit design by reducing your filter's resistor values, you will increase the noiseless bits in the circuit. For example, the delta-sigma ADC in **Figure 1** uses a lowpass filter to reduce noise above the converter's output data rate, F_D. With this filter, use the output data rate of the delta-sigma converter to select the resistor and capacitor values in this circuit. You can use the formula $F_D = 1/(2\pi R_{FLT} \times C_{FLT})$ to calculate the values of R_{FLT} (filter resistance) and C_{FLT} (filter capacitance). This filter reduces noise by targeting the sampling frequency of the converter as the combination of $R_{FLT}/2$ and C_{FLT} goes to work (**Reference 1**).

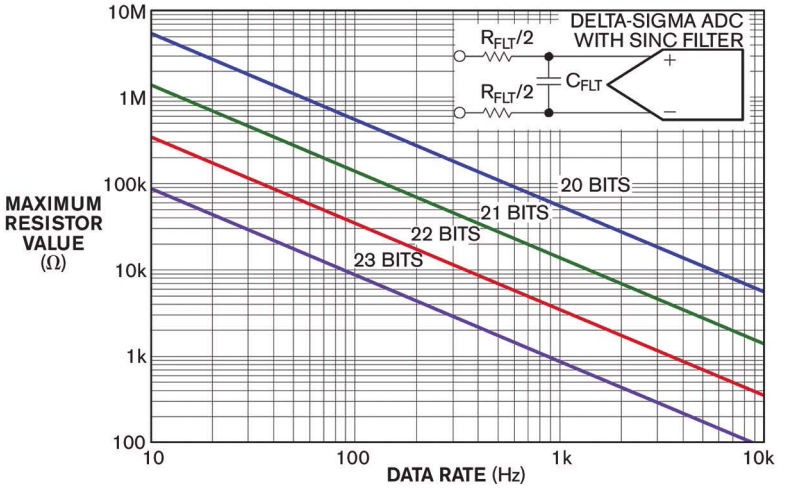

Figure 1 The combination of ($R_{FLT}/2+R_{FLT}/2$) and C_{FLT} reduces circuit noise as long as the filter's resistance noise is below the ideal converter's noise.

The missing detail in this design formula is resistor noise. There is no such thing as a noiseless resistor. The ideal resistor noise is $\sqrt{(4\times k\times T\times R\times B)}$, where k is Boltzmann's constant ($1.38\times10^{-23}\times JK^{-1}$), T is the temperature in Kelvin, R is the nominal resistance in ohms at 25°C, and B is the bandwidth of interest in hertz.

Now, let's make sense of this resistor-noise formula. The noise that you inject into your circuit up to the output data rate is equal to the resistor noise. To determine the maximum allowable resistance value in this circuit, use the following **equation**:

$$R_{FLT(MAX)} = \frac{10^{-(ER\times0.602)}}{4\times k\times T\times F_D},$$

where ER is the specified effective resolution from the ADC manufacturer's data sheet. **Figure 1** illustrates the characteristics of this formula. If you operate your 23-bit-effective-resolution delta-sigma converter at a data rate of 200 Hz, the maximum value of the filter's resistance is 4.297 kΩ or less, and $R_{FLT}/2$ is 2.148 kΩ or less. EDN

REFERENCES

1 Brooks, Michael, "13 things that do not make sense," *New Scientist*, March 19, 2005, www.newscientist.com/article/mg18524911.600-13-things-that-do-not-make-sense.html.
2 Baker, Bonnie, "Analog filter eases delta-sigma-converter design," *EDN*, June 12, 2008, pg 24, www.edn.com/article/CA6566532.
3 "Data Converters," Texas Instruments, www.ti.com/dataconverters-ca.

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

www.edn.com/bakersbest

Allied

Your R&D and MRO solutions provider

ALLIED ELECTRONICS®

On Time Delivery

Your product delivered when you need it

Technical Support

Technical support at 682.757.9204

Technical Data

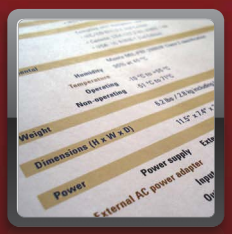

All the info you need online @ alliedelec.com

| Test & Measurement | Interconnect | Enclosures | Power | Control | Passive & Active | Assembly | Tooling |

alliedelec.com 1.800.433.5700

PRY FURTHER AT EDN.COM
Go to www.edn.com/pryingeyes for past Prying Eyes write-ups.

Simple battery technology packs in more features with complex chemistries

Lithium-ion cells are the workhorses of portable energy storage. Battery manufacturers may change the formulation of the coatings for the cathode and the anode and the characteristics of their separators, but the basic components are the same across all rechargeable lithium-ion cells. Formulations optimize battery performance for characteristics such as energy storage, rapid current draw, and cycle life.

The manufacturer welds the internal side of the end cap to the cathode tab. The black ring on the cap insulates it from the can. The cap contains a CID (circuit-interruption device). If the pressure inside the can exceeds the limit that indicates a short, an overcharge, or excessive temperature, the CID opens and permanently disconnects the cell. After inserting the end cap into the can, the manufacturer crimps the can over the cap, sealing the battery. (The second end cap is for illustrative purposes.)

A thin sheet of polypropylene acts as a separator to prevent a short circuit between the cathode and the anode but allows lithium-ion cells to flow through the separator's porous structure. The separator's porosity affects the back-and-forth flow of lithium ions as the cell charges and discharges. If the separator is very porous, lithium-ion cells can flow back and forth more quickly, but very porous separators are also less rugged during manufacturing and less durable. Another characteristic of the separator is its thickness: Thicker separators reduce the occurrence of electrical short circuits but leave less room in the can for the active material on the cathode; the cell also has lower energy density.

A shrink tube provides branding and product information and insulates the can from the outside world.

The jelly roll goes into a stainless-steel can. A welding probe goes down through the center and spot-welds the anode tab to the bottom of the can.

Washers act as plastic spacers to keep the jelly roll tightly packed and in place.

The anode is a coated, thin copper foil. A graphite coating is the most common, but some vendors use LTO (lithium-ion-titanate oxide). An LTO-coated anode gains in increased cycle life and faster charging at a heavy trade-off of lower energy density. This Imara anode uses a conventional graphite coating.

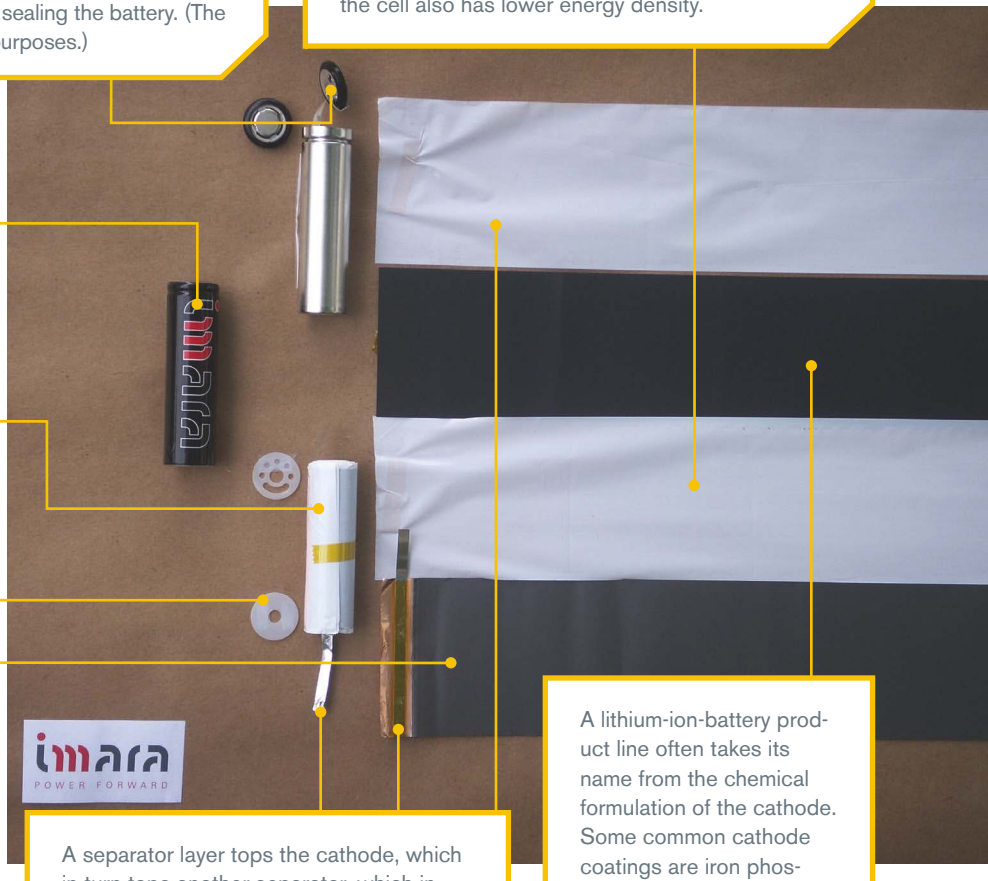

A separator layer tops the cathode, which in turn tops another separator, which in turn tops the coated anode. Manufacturers then roll up the four-layer sandwich to resemble a jelly roll. The anode current collector, or tab, hangs from the jelly roll.

A lithium-ion-battery product line often takes its name from the chemical formulation of the cathode. Some common cathode coatings are iron phosphate, cobalt oxide, and nickel-manganese-cobalt-oxide, plus the major coating ingredient, lithium.

Are you up for a challenge?

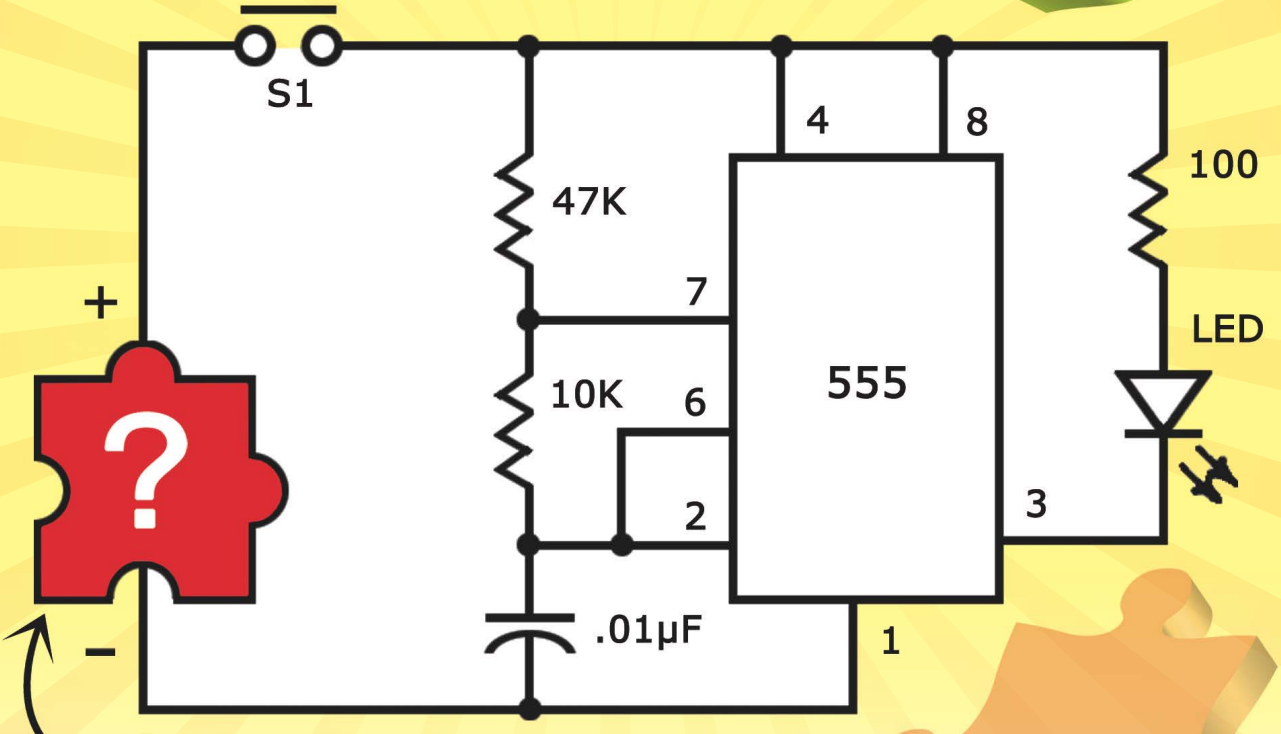

What is the missing component?

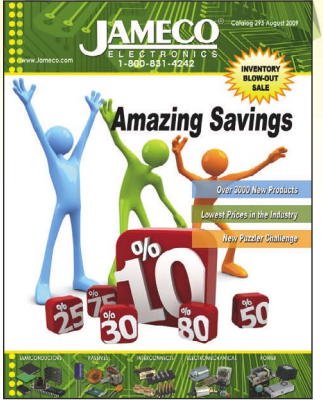

Industry guru Forrest M. Mims III has created a stumper. Can you figure out what's missing? Go to www.Jameco.com/brain to see if you are correct and while you are there, sign up for our free full color catalog. It's packed with components at prices below what you are used to paying.

JAMECO® ELECTRONICS

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

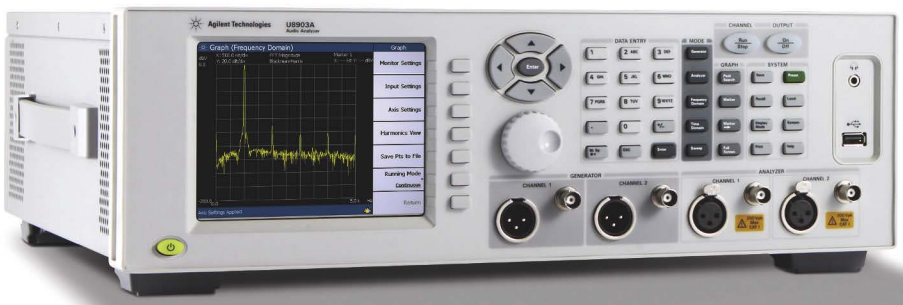

We knew you’d say that.

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

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

Trade in offer good for a limited time.

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

Agilent Technologies

OPTIMIZING RELIABILITY AND POWER EFFICIENCY IN EMBEDDED WIRELESS SYSTEMS

YOU CAN APPLY DYNAMIC-DATA-RATE AND DYNAMIC-OUTPUT-POWER TECHNIQUES TO ENSURE THAT WIRELESS MESSAGES GET THROUGH WITH MINIMAL POWER CONSUMPTION.

BY JONATHAN SUJARIT • CYPRESS SEMICONDUCTOR

The use of wireless technology in embedded-system applications, such as industrial monitoring and control, home automation, remote control, and medical equipment, continues to grow rapidly, and the number of new applications that are adopting wireless is growing even faster. These wireless systems are currently demonstrating a trend of adopting 2.4-GHz technologies because of worldwide unlicensed operation, faster data rates, and other inherent benefits over less-than-1-GHz technologies. The drawback is that these systems will all be competing with each other for airspace, as well as with other prevalent 2.4-GHz systems, such as Wi-Fi, cordless phones, and Bluetooth. It is inevitable that all of these 2.4-GHz wireless systems will eventually interfere with each other and increase the chance of communication failures. Thus, it is no longer sufficient for RF engineers to build a radio and protocol stack that can just wirelessly transmit and receive data without any defense mechanisms for interference. Designers must implement intelligent techniques so that embedded wireless systems are in fact reliable in the increasingly congested 2.4-GHz spectrum.

RF engineers face another critical challenge in power consumption. Many embedded wireless applications require that battery-powered devices last for years, not just weeks or even months. To optimize the efficiency of their systems, engineers cannot rely on using just RF components with ultra-low-current consumption. Because most low-power transceivers consume 1000 times less current in sleep mode than in transmitting/receiving mode, engineers need to turn their attention to finding ways to reduce excessive retransmitting cycles and maximize their system's sleep time. Engineers can address both reliability and power efficiency using dynamic-data-rate and dynamic-output-power techniques.

RELIABILITY

You can most easily measure the reliability of a wireless link by noting the percentage of packets that successfully

move from one device to another. In many cases, a higher percentage of successful transmissions may improve only a user's experience. However, it is a critical requirement that you cannot overlook in certain applications, such as security and medical equipment.

In typical low-power RF systems, a channel sends and receives data packets at a certain data rate. Engineers often implement a frequency-agility technique to enhance reliability by enabling a system to actively choose quieter channels when the system loses packets because the current channel is too noisy. A system with frequency agility needs a transceiver that can quickly switch channels, and it needs a protocol stack that can tell the transceiver which channel to move to. Most low-power, 2.4-GHz transceivers can quickly switch channels, but not all protocol stacks have built-in frequency agility. The latest ZigBee (www.zigbee.org) 2007 Pro specification and the Cypress (www.cypress.com) proprietary CyFi Star network protocol have frequency agility to provide a layer of defense against interference.

However, frequency agility by itself is rarely enough to ensure a worry-free link in the 2.4-GHz band. In some implementations of frequency agility, the system switches channels only when a network link fails because of excessive packet loss. This behavior is less than ideal because the network link must fail first before resuming on a new, quiet channel. Thus, frequency agility does not help to prevent failure but instead only to recover from it. This recovery capability may be sufficient for applications such as sports and leisure that can tolerate random packet loss, but some applications, such as medical equipment and industrial-process control, have low packet-loss thresholds that you cannot exceed.

Another shortcoming of frequency agility is that it assumes that a clean channel will always be available within the associated spectrum. In the 2.4-GHz band, devices such as 802.11g routers occupy 22 MHz of bandwidth, and their 802.11n counterparts can occupy as much as 40 MHz of bandwidth. With just two Wi-Fi routers occupying the entire 2.4-GHz band, other systems have little room to find clear, quiet channels, thus reducing the effectiveness of frequency agility.

AT A GLANCE

- Engineers often implement frequency agility to enhance reliability.
- In the embedded-wireless-system world, people associate lower data rates with greater robustness.
- The biggest misconception concerning power consumption is that low current means low power.
- With dynamic-data-rate technology, a wireless system can choose a rate that minimizes power in any environment.
- The output-power level of a system can dynamically change to allow for optimal power efficiency.

OPTIMIZING RELIABILITY

Although frequency agility is not sufficient on its own for providing complete reliability, the implementation of a dynamic data rate can add an extra layer of robustness to ensure a worry-free link in the 2.4-GHz band. "Dynamic data rate" refers to a system's ability to automatically switch its data rate in real time. It may seem obvious to always use the highest data rate. For example, in the realm of cell phones, the Apple (www.apple.com) iPhone seamlessly switches between EDGE (enhanced-data-rate-global-system-for-mobile-environment), 3G (third-generation), and Wi-Fi protocols to provide users with the maximum possible data rate.

In the world of embedded wireless systems, however, people consider lower data rates to be more robust than higher data rates and may prefer that feature over faster throughputs. For example, the DSSS (direct-sequence-spread-spectrum) modulation technique encodes data into longer sequences of "chips," thereby reducing the effective data rate so that you can still recover the original data amid interference. The DSSS transmitter encodes a byte of data into a 32-chip sequence that the receiver knows (**Figure 1**). Because 32 chips represent 1 byte of data, the effective data rate decreases by a factor of four. Even if interference causes the loss or corruption of some of these chips, the receiving radio can still recognize enough of the 32-chip sequence to determine the original byte of data. IEEE 802.15.4 transceivers that operate in the 2.4-GHz band use DSSS and have a fixed data rate of 250 kbps. Cypress' CyFi transceiver has a maximum data rate of 1 Mbps, but it also has DSSS data rates of 250 and 125 kbps.

DSSS proves effective in environments with random noise or short bursts of interference that cause a few chip errors. If a system does not use DSSS in this type of environment, packets may never get through because interference will continually corrupt random bits.

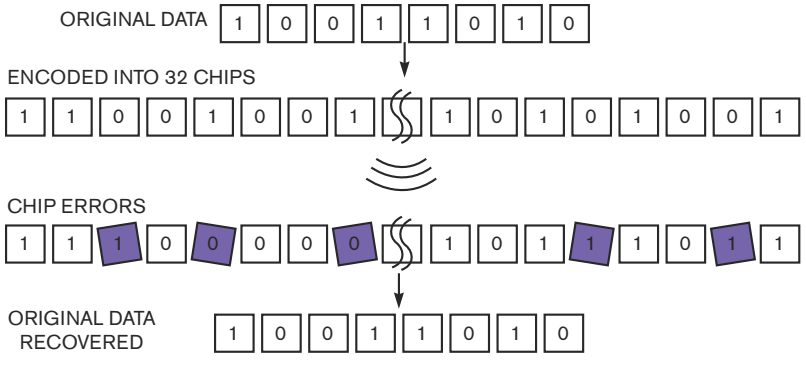

Figure 1 A DSSS transmitter encodes a byte of data into a 32-chip sequence. Because 32 chips represent a byte of data, the effective data rate drops by a factor of four. Even if interference causes the loss or corruption of some of these chips, the receiving radio can still recognize enough of the 32-chip sequence to determine the original byte of data.

DSSS is not always the most robust technique for every type of environment. Because DSSS lowers the data rate, the radio is on the air for longer periods, which increases the chances of collision with other networks. For example, if a system operates in the same channel space as a Wi-Fi router that is streaming video files, a collision could occur, causing the Wi-Fi packets to corrupt the system's packet. With Wi-Fi, it is often more effective to transmit as quickly as possible and find short time slots to fit between Wi-Fi packets.

Because different data rates are more robust than others depending on the type of interference, a reliable system can use dynamic-data-rate techniques to adapt in real time to its current environment. Both the transceiver and the protocol stack must work together to monitor the environment and continually choose the data rate that optimizes reliability. To achieve this cooperation between transceiver and protocol stack, the transceiver must support both a fast, nonencoded data rate and a slower, encoded data rate. Also, the receiving radio must determine which data rate the transmitting radio is using because the data rate is unpredictable. To alert the receiving radio to the data rate, the transmitting radio can incorporate data-rate information in the beginning of the packet overhead so that the radio can switch to the appropriate receiving mode for the payload portion of the packet.

You then need to pair a transceiver with these properties with a protocol stack that intelligently decides which data rate to use. This portion of the protocol stack is complex and is responsible for applying algorithms that always track the performance of both data rates to calculate which is better. Such integrated intelligence enables the system to achieve optimum reliability. This dynamic-data-rate technique can add an extra layer of defense against interference to other methods, such as frequency agility. In a sense, the dynamic-data-rate technique helps prevent failure, whereas frequency agility helps recover from it.

While the data rate is switching, the output-power level can also dynamically change to further improve the reliability of the wireless link. For example, if a system detects an increase in packet-error rate, you can increase the output power to overcome the interference. Higher output-power levels consume more current. Thus, a practical method of implementing dynamic output power, rather than immediately using the maximum output-power level, would be to slowly increase the output power until the packet-error rate decreases.

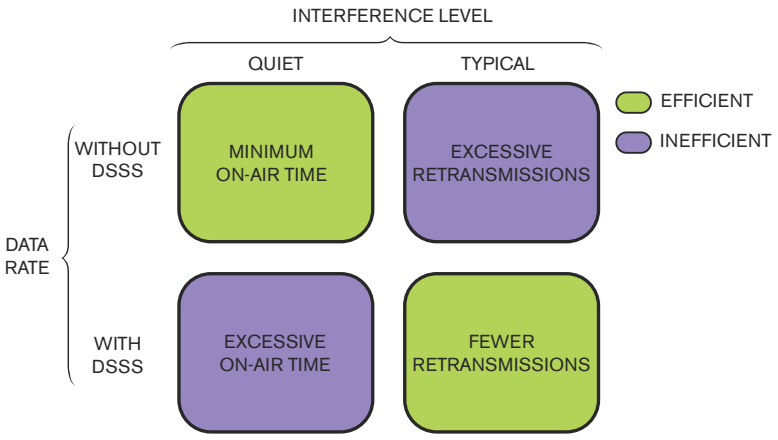

Figure 2 You need to consider environmental factors, such as interference level, when deciding on an optimum data rate and whether to employ DSSS techniques. For example, if a system is operating in a quiet channel at a slower data rate with DSSS, then it is spending excessive time on the air. Conversely, a system operating at a higher data rate without DSSS would be more susceptible to packet failures, which would result in excessive retransmissions.

POWER EFFICIENCY

Embedded-system engineers building battery-powered wireless devices are primarily interested in transceivers' current-consumption specifications. For example, they might have to choose between one transceiver, which consumes 10 mA of current during transmitting/receiving mode and 0.5 μA during sleep mode, and another, which consumes twice that amount: 20 mA and 1 μA, respectively. You might expect the engineers to choose the transceiver with half the power consumption, but there is another way to look at it.

POWER CONSUMPTION DEPENDS ON HOW WELL YOU MANAGE THE TRANSCEIVER AND NOT JUST ON CURRENT-DRAW SPECS.

Assume that, for an application, one transceiver sleeps 90% of the time; its average current would then be about 1 mA: 10 mA×10%+0.5 μA×90%. Also, assume that the other transceiver uses DSSS, so it spends less time retransmitting than the first transceiver does because of greater interference immunity. If the second transceiver sleeps 5% more than the first for the same application due to the benefits of DSSS, the second transceiver's average current would also be approximately 1 mA: 20 mA×5%+1 μA×95%. Which transceiver would you choose? If you choose the first transceiver, you will probably regret it when you later discover that the transceiver is spending all of its time retransmitting because of continual packet loss.

The biggest misconception concerning power consumption is that low current means low power. In reality, power consumption depends on how well you manage the transceiver and not just on current-draw specs. Most low-power RF transceivers consume roughly 10,000 to 20,000A more current in transmitting/receiving mode than in sleep mode. Thus, the protocol stack should try to keep the transceiver in sleep mode as much as possible.

OPTIMIZING EFFICIENCY

You can apply dynamic-data-rate techniques to maximize the percentage of time that a radio is in sleep mode, con-

sequently optimizing a system's power efficiency. Consider the case in which a wireless network is operating in a quiet channel. If the system is using a slower data rate with DSSS—for example, 250 kbps—then the system is spending excessive time on the air because the DSSS encoding is not necessary in a quiet environment (**Figure 2**). A higher data rate without DSSS—for example, 1 Mbps—would minimize the system's on-air time by transmitting as quickly as possible, resulting in more sleep time. Thus, in a channel with little or no interference, a higher data rate without encoding is a better choice for minimizing power consumption.

IT BECOMES CRITICAL FOR THE UNDERLYING TECHNOLOGY TO ACCOMMODATE THE INCREASING RF CONGESTION AND PERFORM RELIABLY.

In the case of a wireless network operating in a typical 2.4-GHz environment that is congested with interference, however, a higher data rate without DSSS would be more susceptible to packet failures, which would result in more retransmissions. If the system must continually retransmit due to packet failures, then it must spend more time in its power-hungry transmitting mode. If the system uses a lower data rate with DSSS, then the system can tolerate interference and avoid retransmissions, allowing the system to spend more time in its ultra-low-power sleep mode.

Most low-power RF technologies have a fixed data rate either with or without encoding. Therefore, they must operate inefficiently when the interference level does not favor the data rate. With dynamic-data-rate techniques, a wireless system can choose the data rate that minimizes power in any environment and always operates efficiently. If the system detects that the channel is quiet, then it will switch to the faster data rate. If the system detects that the channel is noisy, then it will choose the slower, more robust data rate.

The output-power level of a system can also dynamically change to allow for optimal power efficiency. Increasing the output power of a system can help overcome interference to reduce retransmissions. However, more output power means more current consumption. An ideal system would have a protocol stack that could calculate how much power it saves from fewer retransmissions by boosting the output power, and it would compare this power savings against how much power the boosting consumes. Another power-saving scheme is to decrease the output power to the lowest level that is still sufficient for the system to retain the same percentage of packet errors. You can accomplish this task by slowly decreasing the output-power level until the packet-error rate increases.

As the number of embedded wireless applications continues to grow and as more wireless devices find their way into applications, it becomes critical for the underlying technology to accommodate the increasing RF congestion and perform reliably. The technology should also seek to minimize power consumption so that battery life is suitable for most embedded-system applications. Dynamic-data-rate and dynamic-output-power techniques are basic yet effective for optimizing the reliability and power efficiency of wireless embedded systems. **EDN**

AUTHOR'S BIOGRAPHY

Jonathan Sujarit is a product manager at Cypress Semiconductor, where he has worked for two years. He was previously responsible for marketing embedded wireless products, including the CyFi low-power RF product. Currently, he is business-development manager, responsible for growing Cypress' North American timing-products business. Sujarit holds a bachelor's degree in electrical engineering and computer science and a bachelor's degree in business administration from the University of California—Berkeley. In his spare time, he enjoys tennis and is captain of Cypress' tennis team, which competes in the Bay Area Intramural Tennis League. You can reach him at jonathan.sujarit@cypress.com.

Go to www.edn.com/090723df and click on Feedback Loop to post a comment on this article.

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

Accelerating the **Speed of Design.**™

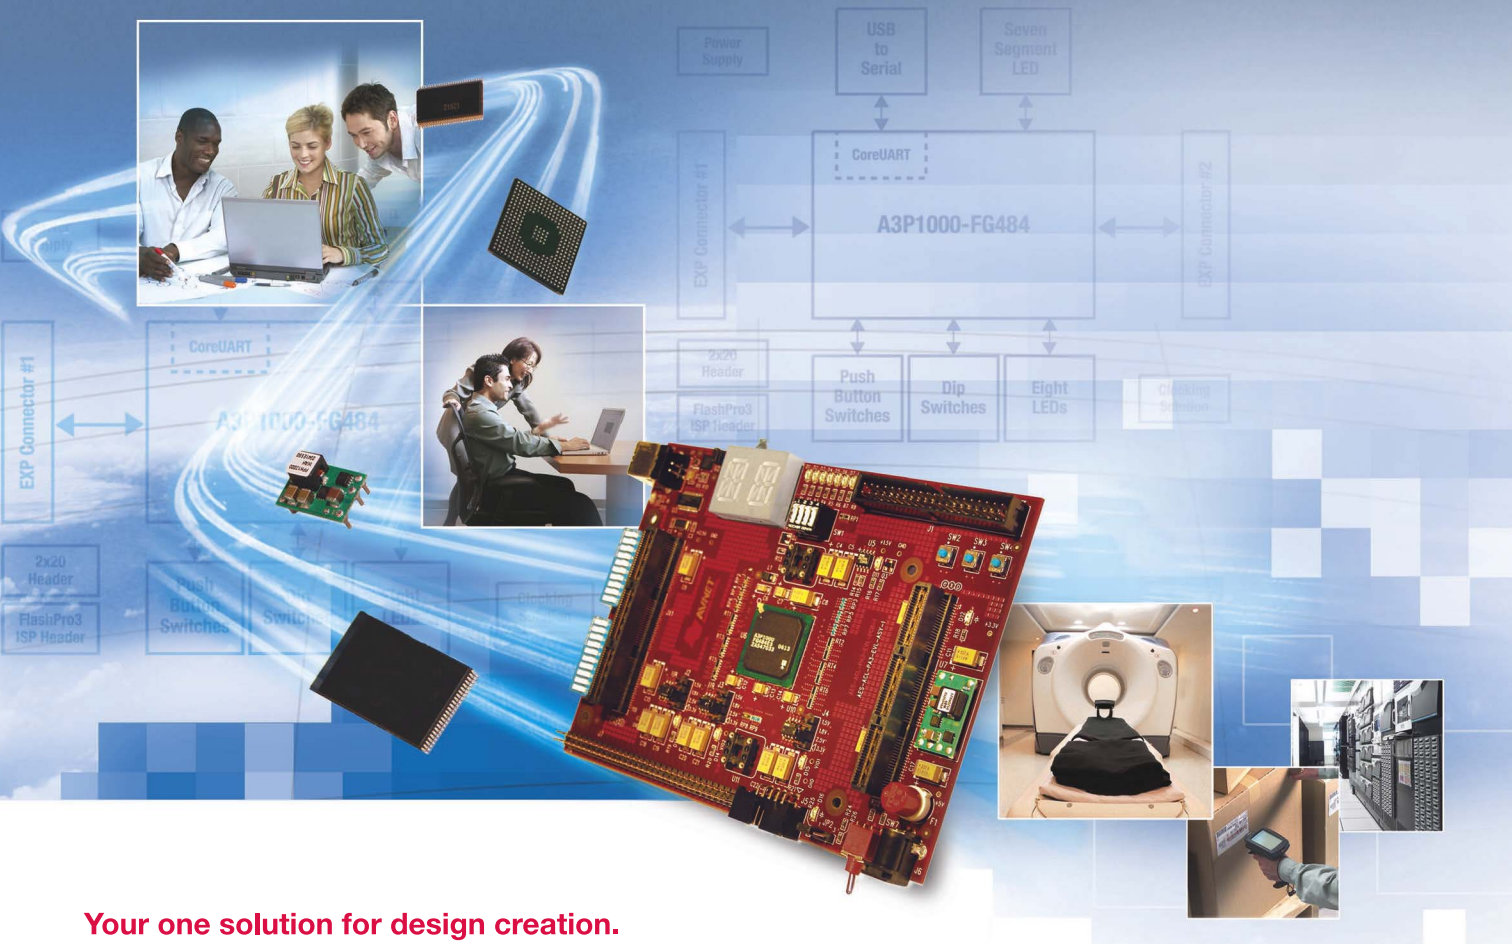

Your one solution for design creation.

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

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

> Proactive engineering support
> Factory certified FAEs

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

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

AVNET® Memec

Accelerating Your Success™

1 800 332 8638
www.em.avnet.com

Avnet Green Initiative

NEDA Supplier Authorized Distributor

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

AS CONFIGURABLE AND PROGRAMMABLE ANALOG CHIPS GROW MORE POPULAR, THEY ARE RAISING SERIOUS QUESTIONS ABOUT THE TRADITIONAL WAY TO DO BOARD-LEVEL DESIGN.

BY RON WILSON • EXECUTIVE EDITOR

PROGRAMMABLE CHIPS PIECING TOGETHER AN ANALOG SOLUTION

Configurable analog ICs have been available in various forms for years, seemingly making less of an impression on the market than their digital cousins, FPGAs (field-programmable gate arrays). The current economic situation suddenly makes the concept of programmable analog standard products more attractive, however. The need to stay on the leading edge of design practice in a product area but without a lot of design cost and without the end-user demand that could justify taking any one design into production fits perfectly with the benefits of these devices. Their day may finally be at hand.

If this is programmable analog's day in the sun, though, a big question is on the table: What design method should you use with these chips? Does the traditional, intuition-based analog flow—depending as it does mainly on breadboarding for verification—make sense? Or will users find flows such as those that designers use for FPGAs—system-level languages on the front end and simulation for verification—more appropriate? The answer is both complex and instructive.

A VARIETY OF ARCHITECTURES

One reason for this complexity is the huge range of architectures living under the collective description of "programmable analog." At one extreme are function-specific chips that have an important degree of user configurability but without changing their basic function. At the other extreme, some chips are the analog

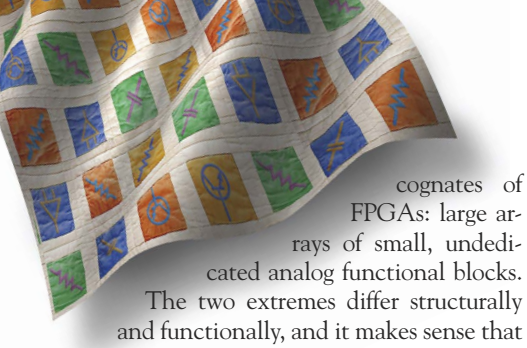

cognates of FPGAs: large arrays of small, undedicated analog functional blocks.

The two extremes differ structurally and functionally, and it makes sense that they would have different requirements.

Consider two examples. At one end is Lattice Semiconductor's family of programmable power controllers (**Figure 1**). These multifunction chips control both the sequencing and the trimming of the several power supplies on a modern PCB (printed-circuit board). For the most part, they are simple PLDs (programmable-logic devices) that Lattice based on its programmable-logic technology. The chips also include precision ADCs, programmable threshold monitors, and DACs to sense output-supply voltages and provide fine-trimming voltages to the supplies' feedback loops.

At the other extreme, consider the FPAAs (field-programmable analog arrays) from Anadigm. These devices are in essence uncommitted arrays of the components necessary for assembling switched-capacitor analog-signal-processing circuits. The purpose of programming in the FPAAs is not merely to set circuit parameters, but also to create the circuit topology. Even the term "programmability" has two meanings in these products. The methods that designers use to deal with them will certainly also differ.

AT A GLANCE

- The economy is stimulating interest in programmable analog components.
- No one design flow exists for using programmable analog chips in a system.
- The design flow depends on the metaphor the vendor offers for understanding the chip.
- In the future, these design flows may be mixed-signal versions of today's FPGA (field-programmable-gate-array) flows.

DIFFERING METAPHORS

Chip developers tend to conceal the structure of programmable chips from their users and substitute in its place a metaphor for visualizing the architecture. They hope that this visualization will be both more familiar to users and more relevant to the problem at hand. Metaphors vary with the structure of the underlying silicon and with the vendor's view of the customer and the customer's problem. For instance, in the early days of digital PLDs, vendors represented the innards of a PLD as sets of NAND gates driving the inputs of big NOR gates, following the then-standard way of expressing logic functions as minterms.

When FPGAs arrived, their vendors initially described them as large arrays of logic cells, with each cell containing a few gates and a flip-flop. Both of these metaphors closely resemble the actual circuitry of the chip. FPGAs have now become so large that the metaphors have become more abstract. Today, the typical way to think of an FPGA is as a blank slate onto which one writes RTL (register-transfer-level) logic, with conveniently scattered clusters of bulk memory, DSP (digital-signal-processing) blocks, and high-speed I/O cells. No one any longer attempts to make the metaphor reflect the circuitry.

You can see a similar distinction in the metaphors vendors choose for their programmable analog devices. Lattice, for example, describes its power-controller ICs in terms close to the actual components on the die: digital and analog inputs, limit comparators, an ADC, a programmable-logic array, and a bank of DACs. In contrast, Anadigm does not discuss its chips to most users in terms of capacitors, configurable amplifiers, ladder networks, or programmable analog switches. The company instead describes analog functional blocks—op amps, filters, and so forth. According to Simon Dickinson, Anadigm's chief operating officer, when the device will become a component of a larger design, the company will sometimes encourage the user to think of the chip as a fixed-function or multifunction black box without reference to its contents.

In general, the more specific the function of the programmable analog chip, the more literal the metaphor is likely to be. If a programmable architecture is highly adaptable, however, the vendor can take one of two approaches: a metaphor as application-independent as the

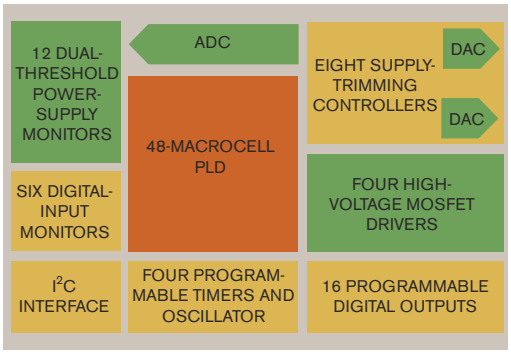

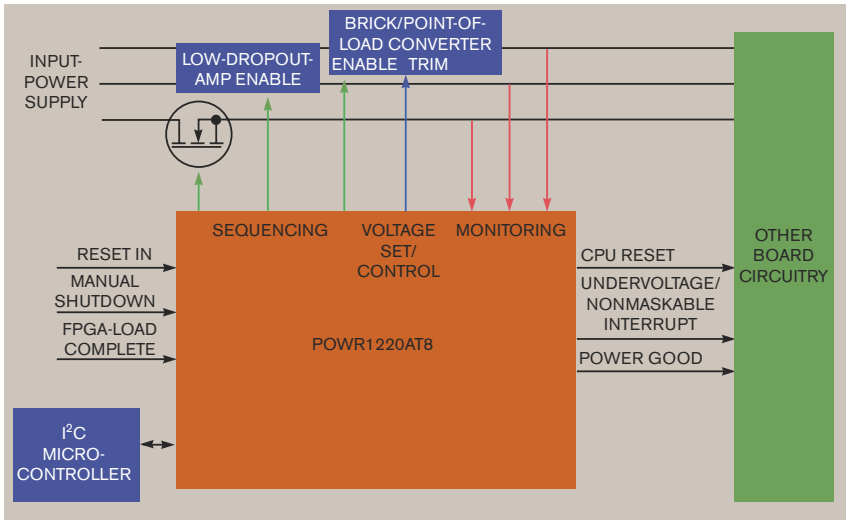

Figure 1 Lattice's 1220AT8 provides programmable supervision for a set of power supplies on a board.

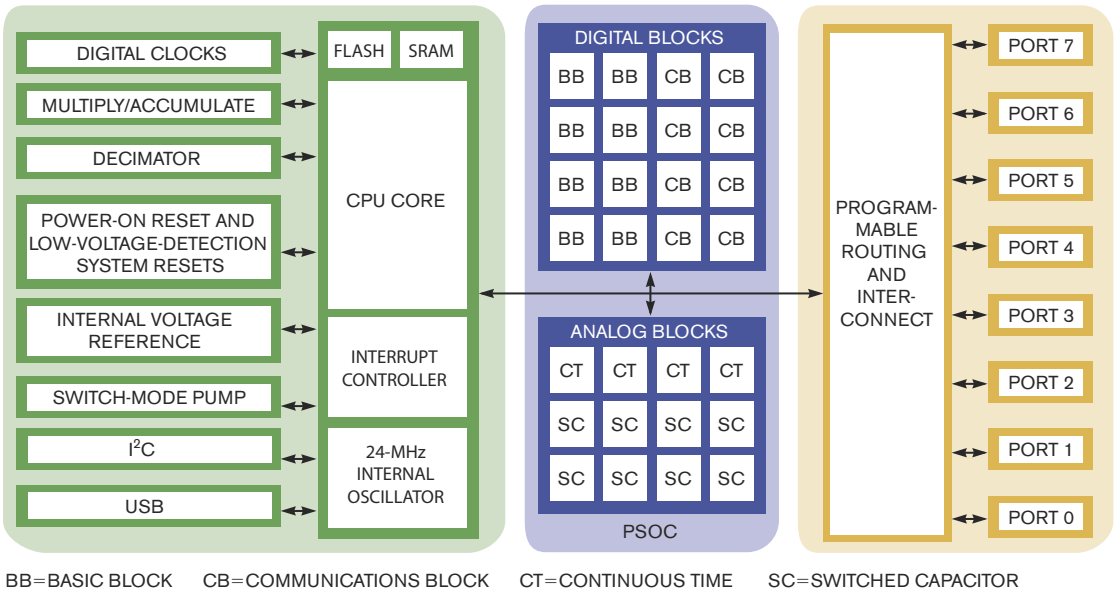

Figure 2 The Cypress pSOC combines a microcontroller with a programmable array of both analog and digital blocks.

underlying architecture, such as Verilog-A or Spice netlists, or a metaphor that conceals the programmable architecture, representing it as merely a parameterized fixed-function chip. The importance of the metaphor is that it—rather than the underlying silicon architecture—determines the design method users will adopt. Case studies illustrate this point.

SOME SAMPLE METHODS

"People are still using discrete chips for controlling reset signals, watchdog timing, and the like," says Shyam Chandra, Lattice's marketing manager, referring to power-controller ICs. "I counted … 400 chips just for reset generation [in On Semiconductor's line], and there must be 100 hot-swap-controller ICs on the market." So Lattice chose a programming metaphor that resembles the

way designers would use a handful of smaller, fixed-function chips.

To develop design specifications, Lattice-product users comb through the data sheets for the processors, memories, FPGAs, and other devices on their boards and determine the necessary power-up sequence for each chip. Compiling these sequences generates the complete state diagram for the controller's logic. But it is a manual task, Chandra warns. "About 75% of the time, the first pass isn't correct," he says. "People make mistakes, and there are often gray areas in the data-sheet specifications, as well." Hence, Lattice has provided a simple programming-language and waveform-simulation tool so that users can code the sequence and watch it in action without destroying anything. When the sequence is right, the tool configures the state engine in the Lattice chip so that it can drive the reset signals and MOSFET gates that sequence power on the board. Similarly, Lattice provides a tool to assist in setting up the trimming and margin-sensing circuits of the chip. Using a library of known dc/dc-converter parts, the software takes in the user's margining and voltage requirements and sets up the on-chip comparators and DACs and the necessary resistor values to both monitor the converter outputs and drive the trim inputs.

The notion of a standard-product metaphor goes beyond the world of power control. In a different market, Actel offers configurable-AFE (analog-front-end) blocks on the Fusion family of FPGAs. The company is similarly supportive of configuring its programmable analog sections but similarly skeptical about users' interest in board-level simulation. "I would say less than half of our customers do any kind of full-board simulation," says Mark Nagel, a field-application engineer at the company. "At the chip level, we do provide a tool that can generate waveforms for analog stimulus, which you can then feed through an ADC module to provide digital output for use in a [Mentor Graphics'] ModelSim simulation of the FPGA logic. But our AFE topology is pretty fixed. For the most part, … our users tend to think through their needs for resolution, sample rate, and so on ahead of time; configure the AFE; and try it. You can query the analog nodes in the AFE with the part on the board, and you can look at the digital outputs using a Synplicity embedded logic analyzer in the FPGA portion of the chip." So users typically don't explore the AFE with simulation tools; they simulate just the digital logic.

Another view comes from Cypress Semiconductor, whose pSOCs (programmable systems on chips) offer a relatively rich configurable array of analog components that couple tightly with an array of digital building blocks and a microcontroller core (**Figure 2**). The general-purpose nature of Cypress' silicon architecture could justify a high-level-language metaphor and a synthesis-simulation design flow. The company has chosen a different direction, however. "Our metaphor is a parts catalog, not a programmable analog array," explains Jason Baumbach, a field-application engineer at the company. "Presenting the users with thousands of configuration registers is not helping them. Instead, we offer a catalog of 'user modules,'" configured portions of the analog array on the chip. They appear to the user as off-the-shelf analog components, however. Users draw up schematics of the analog portion of their designs employing the user modules, and then they generally go directly from the schematic to a bread-

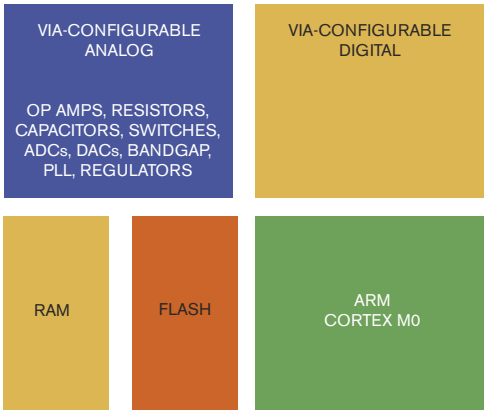

Figure 3 Triad's Mocha offers via-programmed, rather than field-programmable, technology in the programmable analog domain.

THESE COMPANIES MAKE NO ATTEMPT TO PROVIDE A BOARD-LEVEL SIMULATION TOOL THAT WOULD SHOW HOW THEIR CHIPS WORK WITH THE OTHER CIRCUITS ON THE BOARD.

board using a Cypress pSOC development board. "We don't see most people using Spice or even [The MathWorks'] Matlab," Baumbach says. "For the most part, the things they are trying to do in the analog circuitry are pretty simple."

OTHER VIEWS

For the most part, these companies make no attempt to provide a board-level simulation tool that would show how their chips work with the other circuits on the board or even a detailed simulation of how the analog and digital portions of the chip work with each other. The users' experience and the information users extract from data sheets serve this purpose. "In the power-management world, simulation means reading the schematic and thinking about it," observes Lattice's Chandra. But Cypress sees evolution taking the company in another direction. The analog sections of the pSOC chips continue to become more substantial. And customers are increasingly exploiting the fact that the integral microcontroller can reconfigure the analog array on the fly. This capability is useful, but it undermines the parts-catalog metaphor. "The devices keep getting more powerful," says Mark Saunders, a product-marketing manager at the company. "So we need to keep pushing on our abstractions."

In contrast, austriamicrosystems AG is experimenting with a broader approach to simulation for relatively simple chips. "We see two kinds of engineers," explains Bruce Ulrich, director of marketing for linear and wireless products at the company. "There are purists who just have a cultural issue with simulation. They see it as somehow sleazy. But there are also designers using off-the-shelf dc/dc converters who don't think of themselves as experts and find it useful to have a tool that helps them experiment with switching frequencies and loads to see what's going to happen."

Toward this end, the company has linked to its home page a version of Transim Technology's WebSim linear-simulation engine. The page has models of austriamicrosystems' power ICs, a netlist-capture tool, the simulation engine, and a tool for generating bills of materials. Users can breadboard a power subsystem on the site, explore parameters and performance, and get a parts list. "This is not Spice-accurate; it's a linear approximation," Ulrich emphasizes. "But it's close enough to give you an idea of how the circuit will behave and warn you when you are about to have switching issues, noise problems, or instability." The company recently added the capability to its site but is considering an extension of the model libraries so that users can explore other parts from the company's product lines.

Triad Semiconductor is in a different situation. Like Cypress, the company makes chips that combine a programmable analog array with a microcontroller (**Figure 3**). In this case, however, the programmability is in the form of factory-programmed vias rather than field-programmable flash cells. Most of the time, according to Reid Wender, Triad's vice president of marketing and technical sales, the company's users start out not in simulations but by breadboarding their designs with discrete, off-the-shelf analog components for the analog-signal paths and an FPGA for the digital logic and ARM Cortex M0 core. "Most of these designers have preferred off-the-shelf components, and they ask our ASIC to match those specifications," Wender says.

The customer shares this breadboard with Triad, which uses the board's behavior to define the transfer functions for the analog-signal paths in the ASIC. The Triad team then uses a chip-level simulation tool to configure the analog array to match the transfer functions on the customer's breadboard. Triad shares its simulation results with the customer,

For links to the Web sites of companies this article mentions, go to www.edn.com/090723cs.

settles on a via mask, and configures banked wafers for a total turnaround time of about four weeks. The customer can then drop the Triad silicon into its breadboard and proceed with verification. The hard part of this flow, according to Wender, is understanding the integration of analog-signal-processing paths with software running on the Cortex. "We have been working with Keil to simulate mixed-signal peripherals. But, eventually, we are going to need a full analog/mixed-signal simulator on the desktop—something that can pull together Verilog, Spice, and software simulations on the desktop for a low price," he says. "We are still searching."

THE FUTURE

A view of the future comes from a continuing research project at the Georgia Institute of Technology, where Paul Hasler, a professor of electrical engineering and computer science, has for a decade been working on FPAAs. The current project involves constructing large arrays with approximately 1000 analog components and thousands of switch-level devices organized into 100 computational analog blocks. "These chips are maybe 10 times the capacity of existing commercial programmable analog arrays," Hasler says. "In analog-signal-processing performance, we can put the equivalent of 1 teraMAC [trillion multiply/accumulate instructions] of signal processing into one chip for a few hundred milliwatts." Hasler and his team have created a full design flow employing the metaphor of analog-signal processing rather than the details of the chip design. "One of our big chips has about 100,000 programmable parameters," Hasler says. "You can't deal with that level of complexity manually, so we approach programming with a block-level signal-processing metaphor."

Even so, the sheer complexity of the FPAAs demands an ASIC-like design flow. Attempting to troubleshoot a 1000-component analog design on the breadboard would be hopeless. So the FPAA flow employs two levels of simulation. The flow starts with [The MathWorks'] Simulink and a library of computational elements for which Hasler's team has created Spice netlists. The user can perform system simulations in Simulink and then by substitution create a Spice netlist, which then goes to a chip compiler that produces the equivalent of an FPGA-programming file. "We can compile most valid Spice netlists, but not all of them will produce efficient designs," says Hasler. "At the Spice level, a user must learn to work with the tools to make best use of the silicon. At the Simulink level, that work is mostly done in the library." The team is now developing tools that can extract a Spice netlist with accurate parasitics from the switch-level programming file and perform, in effect, a layout-versus-source comparison. "Taking back-annotation all the way back to the Simulink level will be a little trickier," Hasler adds.

Such flows may be the future, even for far-simpler components, says John Pierce, director of mixed-signal-simulation marketing at Cadence. "The traditional approach just isn't making it any more, even for fixed-function devices. You have to look at what will happen when you integrate your programmable component into the system—not just at how to program it."

There are problems still to solve, Pierce adds. Intuitively, the right place to start a board-level simulation is with Matlab or a similar tool. Getting from a transfer-function view to a switch matrix is nontrivial, however. Issues arise even at the circuit-simulation level. "Verilog-A or SystemVerilog [and similar languages] don't want you to change configuration-register settings during a run," he says. If you try to model the configuration registers and analog switches as part of the device netlist, however, the simulation can just explode, especially if you are using switched-capacitor techniques. "The techniques for simulating programmable analog in a systems environment exist," Pierce says. "The challenge is bringing them into our Verilog-AMS [analog/mixed-signal] world." EDN

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

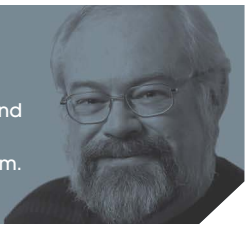

No need to duplicate yourself

BuyerZone saves you time and money on COPIERS AND MORE.

Get free advice, pricing articles, ratings and quotes on Business Purchases from digital copiers, to phone systems, to payroll services and more. Visit BuyerZoneBusiness.com today to make fast and efficient business purchase decisions. At BuyerZone, you're in control. We're just here to make your work easier.

Join the millions who've already saved time and money on many of their Office Equipment purchases by taking advantage of BuyerZone's FREE, no obligation services like:

- Quotes from multiple suppliers
- Pricing articles
- Buyer's guides
- Supplier ratings
- Supplier comparisons

REQUEST FREE QUOTES NOW!
Call (866) 623-5564 or visit
BuyerZoneBusiness.com

BuyerZone
Where Smart Businesses Buy and Sell

A division of
Reed Business Information.

BY BOB EISENSTADT • ALCHIP TECHNOLOGIES

Making ASIC power estimates before the design

A STRUCTURED APPROACH CAN GIVE YOU MEANINGFUL POWER ESTIMATES WHEN THERE'S STILL TIME TO INFLUENCE PROCESS AND ARCHITECTURE DECISIONS.

ASIC design teams face a growing problem. ASIC power-consumption estimates that take place before the design phase lack both scope and credibility. These shortcomings affect process, design, and IP (intellectual-property) selection. As the design complexity increases and power-consumption requirements tighten, this problem becomes more pervasive, often causing design rework, schedule slippage, higher NRE (nonrecurring-engineering) costs, and other challenges.

To address these issues, engineers should use a structured approach and a set of heuristic power-estimation rules for making effective predesign ASIC power estimates. Given that IC design flows and design requirements differ greatly, you may need engineering judgment and creativity to extend this approach across a range of design situations. In any event, these concepts should provide a good starting point. Companies that outsource portions of their ASIC design flows should engage early with their ASIC design partners to improve the quality of their predesign power estimates.

From an analytical viewpoint, generating power-consumption estimates should be straightforward. However, predesign power estimation is often given a low priority and may rely on the wrong mix of people to support the effort. In general, the team responsible for process, IP, and design selection should also own the predesign power estimates. This team will need augmentation if it is too distant from the design-implementation effort or if it has little hands-on power-estimation experience. For example, a software-centric architect plus a continually distracted manager will probably generate inadequate predesign ASIC power estimates, causing suboptimal process, design, and IP selection. That team needs designers with hands-on experience, and it may also need to include low-power-design experts from the ASIC design partner. The structured approach lends itself to distributing tasks and recombining results so that team members can share the workload.

GETTING STARTED

Power estimation involves dynamic power and static, or leakage, power. This article focuses on making approximate estimates; thus, it ignores various second-order effects, such as crossover current. To simplify the power-estimation process, perform the analyses of static and dynamic power estimations separately. If possible, prioritize the power-minimization requirements by determining the relative importance of static versus dynamic power consumption. For example, assume that an ASIC design targets a battery-powered application in which it will be dynamically active 1% of the time and leaks current the other 99% of the time. In this case, minimizing the static power should take priority over minimizing the dynamic power.

IP selection and process selection are both critical but often seem like recursive problems. That is, IP selection drives process selection, and process selection drives IP selection. Fortunately, most applications should align with a requirement for low-power, general-purpose, or high-performance features. The application usually determines the choice of libraries and IP. The libraries and IP in turn help determine a target process. For example, most libraries and IP for portable devices should be available in low-power processes, and most libraries and IP for high-bandwidth networking should be available in high-performance processes. A check of library and IP availability will quickly help eliminate process choices and reduce the scope of the power-estimation effort before you start your design.

Once you have settled on a process node and type, you have bounded your operating voltage (**Table 1**). Note that the core voltage for low-power processes is 0.2V higher than the core voltage for general-purpose processes. So a low-power process has lower static power and higher dynamic power than the corresponding general-purpose process. This situation provides strong motivation for determining the relative importance of dynamic versus static power for your design.

USE SPREADSHEETS

Engineers prefer to use spreadsheets for early-phase power estimation because they are easy to use, help organize data, and support what-if design scenarios. Some companies may attempt to apply more elaborate approaches from the EDA industry. EDA tools generally provide their most effective results when you use them on well-defined designs. Before design start, the process, libraries, IP, and so forth remain undefined, which makes it difficult to estimate power with EDA tools. Using spreadsheets to estimate power consumption represents a rare design-community preference for software from Micro-

TABLE 1 NOMINAL ASIC-PROCESS CORE VOLTAGES

Process (nm)	Low-power process (V)	General-purpose process (V)
90	1.2	1
65	1.2	1
40	1.1	0.9

soft (www.microsoft.com), for example, rather than from the EDA industry.

MEMORY POWER ESTIMATION

Memory provides an excellent starting point for making initial power estimates. Complex designs may contain hundreds of types of memory, and the total memory power is often a substantial—if not the dominant—power component. You need neither a mature netlist nor powerful EDA tools to estimate memory power and study memory-power trade-offs. You need only the supplier's memory specifications or access to memory compilers to generate the memory specifications. Start with a sample of memory configurations for quick estimates and comparisons. Focus on memory configurations that consume the most power: those with high instance counts, that run at high frequencies, and that have large bit widths. As the design architecture matures, gradually include memories in the memory power estimates until you've covered all the memories in the design.

IF A HIGH-THRESHOLD-VOLTAGE TRANSISTOR-BASED MEMORY IS TOO SLOW, YOU COULD USE A SIMILAR COMPILER TO GENERATE A SIMILAR MEMORY.

Contemporary memory compilers have various compilation options. Try the initial compilation runs without using power-saving features to obtain baseline data for supplier-to-supplier comparisons. You can then go back and explore power-saving features to help optimize the target implementations as the design matures. Track each memory type's width, depth, and power consumption. Also track the memory compiler's column-multiplexer option, which affects memory aspect ratios, performance, and power consumption. Use of inconsistent column-multiplexing options would invalidate data from comparative memory-compilation runs.

Vendor memory-power specifications can vary greatly, so designers need to work through the specifications to understand the conditions and make valid comparisons across memory suppliers. Memory specifications should break out dynamic power consumption employing read and write activity rates. However, the specification's footnotes may bury activity-rate information. If the activity rates remain unknown, ask the memory vendor about them. Also note that memory output ports usually drive relatively small loads, so dynamic power due to memory output loading should be far lower than the memory's net dynamic power consumption.

Data sheets don't provide exhaustive information that covers all cases of interest, so use scaling to generate design-specific data. When possible, you should also verify memory-power scaling against the actual memory-specification data to help validate the scaled results.

SCALING RULES

If you ignore second-order effects, you can assume that switching capacitance is constant across process and temperature. Thus, you can also assume that dynamic memory power—related to switching capacitance—is constant over process and temperature for rough power-estimation purposes. For example, if a high-threshold-voltage transistor-based memory is too slow, you could use a similar compiler to generate a similar memory based on faster, higher-leakage, low-threshold-voltage transistors. The high- and low-threshold-voltage versions of the memory should have equivalent layouts and dynamic power, but their leakage power differs greatly.

Memory dynamic power should scale as a function of voltage squared: $P=CV^2f$, where P is the power dissipated, C is the effective power-dissipation capacitance, V is the operating voltage, and f is the effective transition frequency. Because switching capacitance is fairly constant over process and temperature, for estimation purposes dynamic power becomes a function of just the voltage squared and the transition frequency. Fortunately, these variables are independent of each other, so you can focus on them separately. For example, if the memory vendor has provided best-case dynamic power at 1.32V, reflecting a 10% supply variation, but the design operates at or below 1.26V, reflecting a 5% supply variation, the dynamic-power-consumption scaling factor would be $(1.26)^2/(1.32)^2=0.91$.

To verify that the memory vendor's dynamic power scales as the square of the supply voltage, check the memory vendor's dynamic-power-spec data at typical and best-case conditions—say, 1.2 and 1.32V, respectively. Then, calculate the ra-

TABLE 2 SIMPLIFIED MEMORY-POWER SPREADSHEET

Memory	Word count	Bit width	Memory-instance count	Column multiplexer	Read activity (%)	Write activity (%)	Frequency (Hz)	Per-memory dynamic power (mW)	Per-memory static power (mW)	Dynamic power (mW)	Static power (mW)
One read/one write, high threshold voltage	7744	12	15	Eight	0.25	0.25	80	0.775	0.0034	11.64	0.051
One read/one write, low threshold voltage	4384	72	Two	Eight	0.25	0.25	270	9	0.025	18	0.050
One read/one write, high threshold voltage	5760	64	One	Eight	0.25	0.25	135	4.25	0.009	4.25	0.009
One read/one write, high threshold voltage	14,528	12	Two	16	0.25	0.25	80	0.925	0.005	1.85	0.010
Total										35.74	0.12

tio of typical dynamic power to best-case dynamic power to see whether it yields $(1.20)^2/(1.32)^2$, or 0.826.

Memory dynamic power should scale linearly with frequency and activity rates: $P=CV^2fA$, where A is the activity rate, which can range from 0 to 100%. Many memory specifications provide simple power activity constants that you should multiply by frequency and activity rates. So you should include these power activity constants and memory frequencies in the memory-power spreadsheet to help automate the calculations. For example, assume that memory-vendor specs provide dynamic-power consumption at fixed activity rates, such as 20%, and the design runs at a maximum activity rate of 10%. You should scale down the 20% activity rate to derive power consumption with a 10% activity rate, using a scaling factor of 10%/20%, or 0.5.

Memory-read and memory-write activity rates should use different power activity factors. However, if read and write activity rates remain consistent, you can combine and scale them together. For example, assume a memory's read and write activity rates are both 15%, and you need to calculate what would happen with a 25% activity rate. You could combine these 15% read and write activity-rate results and scale them with the factor of 25%/15%, or 1.67. If the read and write activity rates differ, with a read activity rate of 25% and a write activity rate of 5%, for example, then don't combine the dynamic read and write power-consumption calculations. Instead, calculate the read-activity power separately from the write-activity power and then combine them.

Memory static power should remain fairly constant across activity rate and frequency. Thus, static-memory-power calculations should be independent of dynamic-memory-power calculations, and static-power calculations should use separate spreadsheet columns. Memory static power instead varies over process, voltage, and temperature. The highest leakage occurs with the best-case process and highest voltage corner at the highest temperature. These conditions differ from best-case timing conditions, which use the best-case process and voltage corner at the lowest temperature. It follows that you must carefully read the vendor's memory specs to ensure that the leakage data the vendor used in the power estimates reflects the highest leakage-power conditions, not the best-case timing.

Memory static power varies exponentially with temperature, so it doubles approximately every 15°C for 90-nm and larger processes. For newer, smaller processes, leakage tends to be less sensitive to temperature. You can study the effects of temperaure on leakage by checking the memory-spec leakage data at, say, 25 and 85°C and calculating the temperature change for leakage to double. For example, 85−25=60°C and 4×15=60°C, so leakage should increase by approximately 2^4, or 16 times, from 25 to 85°C. As another example, if the leakage specifications are at, say, 85°C but the design's maximum junction temperature is 100°C, the leakage-power estimate should be double the leakage power the vendor specifies at 85°C.

Table 2 shows a simplified memory-power spreadsheet. A more complete spreadsheet should include memory area, bit-count totals, read and write dynamic power per megahertz, and other variables. Also note that the memory in the second row with low threshold voltage has a much higher performance requirement than the other memories. To support higher performance, implement this memory with higher-leakage, low-threshold-voltage transistors; this instance accounts for 40% of the total static, or leakage, power. Efforts to minimize memory power should focus on memory configurations consuming the most power, and your spreadsheet should make those configurations evident.

MAKE INITIAL STANDARD-CELL POWER ESTIMATES WITHOUT POWER-SAVING TECHNIQUES, SO YOU WILL HAVE BASELINE DATA.

IP AND STANDARD-CELL POWER ESTIMATION

IP power estimation is similar to but less complex than memory power estimation. IP instance counts are usually much lower than memory instance counts, and there are usually fewer specialty-IP choices with fewer power-saving features, which simplifies power estimation. Other factors such as IP design risk, cost, and area may outweigh IP power consumption when you compare specialty IP from different suppliers. You should base your initial IP power estimation on the vendor's specifications. Enter the IP-power data into the power spreadsheets and separate the static- and dynamic-power components.

Make initial standard-cell power estimates without power-saving techniques, so you will have baseline data. Again, separate the static- and dynamic-power components. Predesign-phase estimates are difficult to make because you don't yet know the standard cell's instance count, multiple-threshold-voltage mix, routing capacitances, and clock-tree capacitances. In short, you don't know the capacitance in the CV^2f equation to make dynamic power estimates.

To work around not knowing the capacitance, you can scale the power results from a similar previous design, assuming that power data that includes routing capacitance is available. If you are targeting the same process node, scale for instance count and for frequency. If targeting a smaller process node, scale for instance count, frequency, voltage squared, and capacitance. Note that capacitance increases as spacing shrinks. For example, if gate-area scaling is 50%, use 60 to 90% for capacitance scaling.

An alternative approach to making dynamic power estimates involves scaling clock-tree power from a previous similar design. To do so, estimate the clock power per flip-flop that the previous design consumed. For example, assume a representative clock tree has 100,000 flip-flops and consumes 5 mW when including clock-tree routing capacitance. This scenario yields an average estimated power consumption of 0.05 μW per flip-flop, including clock-tree routing capacitance. Next, estimate clock-tree load counts in the new design. If necessary, scale for load capacitance due to library and process changes. To estimate flip-flop-load scaling factors, compare the clock-pin input capacitances of representative flip-flops in the previous design with similar flip-flops in the new design. Next, scale for frequency and the load count and then estimate the power per clock in the new design. Clock trees consume approximately one-half of the standard-cell dynamic

power in many ASICs, so double the clock-tree power to estimate total standard-cell power.

Estimating standard-cell static power is also problematic. One approach relies on scaling leakage-power results from a similar previous design, assuming that such data is available. If targeting the same process node and library, then scale the instance count to make the leakage-power estimate. You may also need to adjust for the relative mix of multiple-threshold-voltage cells. If targeting the same process node but with a different library, scale for instance count and for relative library leakage. You can base library-leakage scaling on the relative leakage of a small representative sample of standard cells. If targeting a different process node, scale for both instance count and relative library leakage.

The second approach for standard-cell leakage-power estimation relies on standard-cell library data. Select a representative sample of standard cells. You may need multiple samples of standard cells to account for different libraries and multiple-threshold-voltage-cell variations. Estimate the leakage for each sample of standard cells. Next, estimate the total instance count reflecting each sample of standard cells. Finally, scale the standard-cell samples to represent corresponding estimated design-instance counts and calculate the total static-power estimate.

For quick power estimates, start with a representative sample of key components. As the design matures, narrow your process and IP choices and introduce more component data. Start with baseline data without power-saving features. Then, evaluate and work in appropriate power-saving features as the design progresses. To simplify the effort, try to separate the analysis of static and dynamic power. Also, focus on memory, IP, and standard-cell power separately. Recombine memory, IP, and standard-cell power estimates when necessary. IC design projects differ greatly, so you will need flexibility and engineering judgment to handle design-specific power-estimation issues. **EDN**

> **Go to www.edn.com/ms4324 and click on Feedback Loop to post a comment on this article.**
>
> **For more technical articles, go to www.edn.com/features.**

AUTHOR'S BIOGRAPHY

Bob Eisenstadt is currently principal engineer at Alchip Technologies (Santa Clara, CA), a rapidly growing fabless-ASIC company. For the past 22 years, he has worked in a variety of ASIC design and consulting positions in Silicon Valley. He has extensive hands-on ASIC design experience in architecture, RTL (register-transfer-level) design, logic verification, timing analysis, physical design, and physical verification. In 2003, Eisenstadt co-founded Silicon Mosaic, where he developed and patented a generic power-gating product. He holds a bachelor's degree in electrical engineering from Cornell University (Ithaca, NY) and master's degrees in both electrical engineering and business administration from Santa Clara University (Santa Clara, CA).

EDITED BY MARTIN ROWE
AND FRAN GRANVILLE

designideas

READERS SOLVE DESIGN PROBLEMS

Quasiresonant flyback converter easily charges energy-storage capacitors

Todor Arsenov, Toronto, ON, Canada

Designers often use chargers with flyback topologies to quickly charge energy-storage capacitors (**references 1** and **2**). In a flyback topology, the energy transfer takes place only when the charger's power MOSFET is off, which effectively isolates the power switch from the load, comprising high-energy storage-capacitor banks. Thus, the voltage levels on the circuit transformer's secondary can vary from zero to a predetermined value and corresponding energy level without any significant stress on the components on the primary side of the transformer.

The classical flyback capacitor charger operates in CCM (continuous-conduction mode). Flat-topped, short-duration current pulses on the transformer's secondary charge the storage capacitors (**Reference 3**). Unfortunately, this charging strategy requires complex control circuitry to limit both the secondary current and the capacitor voltage. Most circuits use a specialized PWM (pulse-width-modulation)-controller IC, which increases the overall cost of the charger. Another disadvantage of the CCM is the small portion of energy that accumulates during the on-time of MOSFET conduction:

DIs Inside

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

$$\Delta W = \frac{1}{2} \times (L_P \times I_{P_{PK}}^2 - L_P \times I_{P_{OFFSET}}^2),$$

where $I_{P_{OFFSET}}^2$ signifies the initial nonzero primary current at the beginning

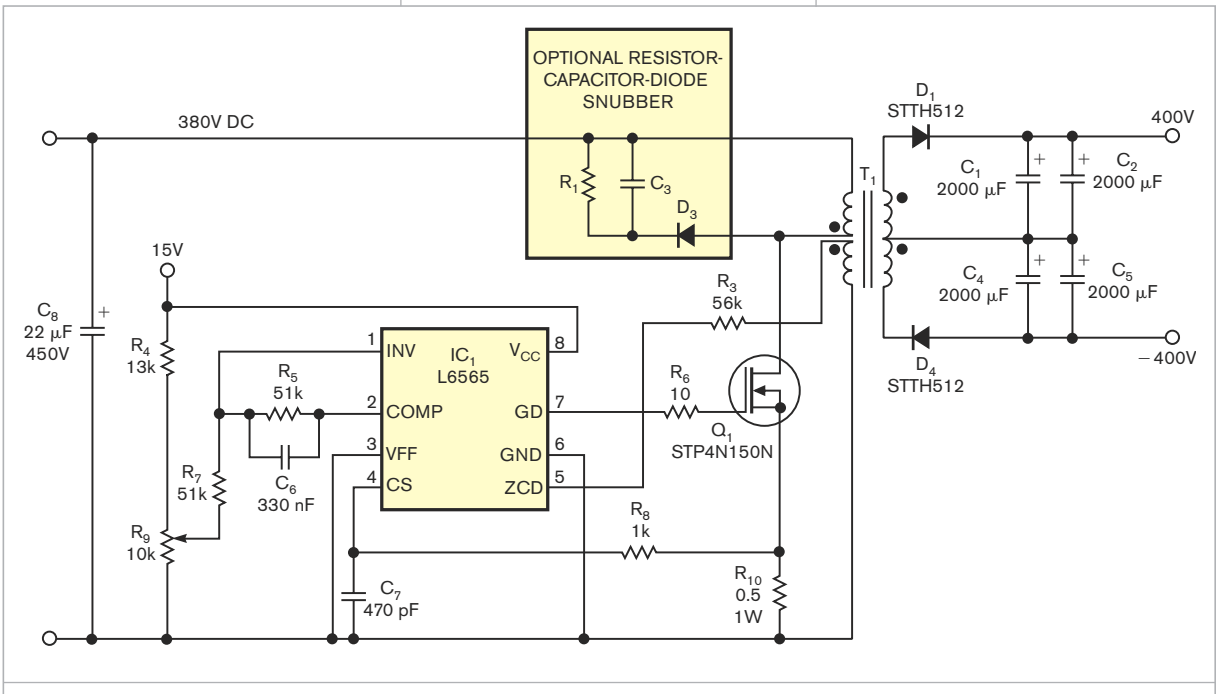

Figure 1 A flyback capacitor charger uses BCM operation.

of the on-time interval.

Only this limited portion of energy transfers from the primary to the secondary sides and enters the storage capacitor. Therefore, you can considerably increase the amount of energy transferable to the capacitive load if the converter can operate in BCM (boundary-conduction mode). The secondary current becomes zero, the power MOSFET turns on, and the primary current builds from zero. Thus, a bigger portion of energy accumulates during every consecutive on-time interval:

$$\Delta W = \frac{1}{2} \times L_P \times I_{P_{PK}}^2 .$$

With all other conditions equal, BCM operation ensures faster accumulation of a predetermined amount of energy because of the bigger stored portions of energy during the on-time intervals. Many converter circuits that operate using BCM incorporate PWM controllers that implement BCM operation for capacitor charging. These circuits often use Maxim (www.maxim-ic.com) MAX8622 or Linear Technology (www.linear.com) LT3468 ICs. These ICs are specialized devices to accommodate BCM operation.

You can, however, implement flyback BCM operation without these specialized parts. Manufacturers implement BCM in the variable-frequency versions of flyback converters, which are quasiresonant, ZVS (zero-voltage-switching) converters that commonly find use in TV SMPS (switched-mode-power supplies). For example, you can use the STMicroelectronics (www.st.com) quasiresonant-SMPS-controller L6565 to build a flyback capacitor charger working in BCM (**Reference 4**). Doing so eliminates the need for using a specialized chip for capacitor chargers.

Figure 1 shows the power stage of a charger using the ST L6565. It achieves the BCM using a second primary winding on T_1 that feeds the transformer-sensing input at the ZCD pin of the L6565. The voltage of this winding is a scaled-down replica of the drain-to-source voltage of power MOSFET Q_1. When the circuit interrupts the secondary current—indicating full demagnetization of T_1—it detects the minimum of the first valley of ringing, and the L6565 turns on the MOSFET. This action eliminates the idling and zero-phase-time intervals, thus establishing BCM. The elimination of the zero-phase-time intervals greatly reduces the charging time of the storage capacitors.

At the beginning of the charging sequence, the output voltage is low because of the large capacitance values. The secondary current decreases slowly. The reflected voltage on the primary side is too low to trigger the ZCD pin of the L6565. Thus, the L6565's

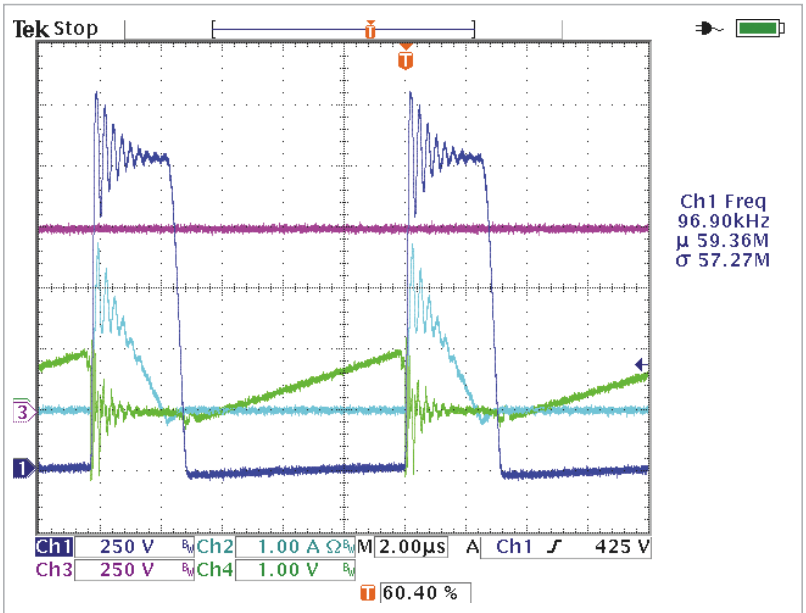

Figure 2 When the secondary current reaches 0A, the MOSFET turns on, and the primary current increases from 0A.

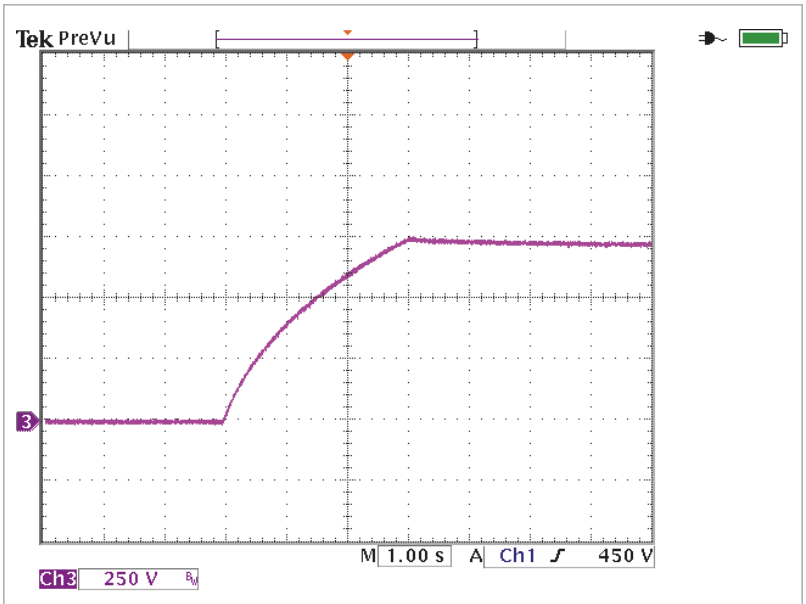

Figure 3 The output-capacitor voltage reaches its full level in about 3 seconds.

initial starting timer sets the switching frequency to 2.5 kHz at the start of charging. The output voltage across the storage capacitors increases to a point at which the switching frequency becomes variable because of the demagnetization of the transformer core. **Figure 2** shows that, as soon as the secondary current (Channel 2) becomes 0A, the power MOSFET turns on, and the drain-to-source voltage decreases (Channel 1). At that time, the primary current again increases (Channel 4). At the output voltage close to full charge, the switching frequency is approximately 100 kHz. **Figure 3** shows the total voltage of 750V across C_1, C_2, C_4, and C_5 within a 3-second charging time.

The waveforms in **figures 2** and **3** are evaluation measurements of a prototype capacitive charger using the L6565 and power MOSFET STP4N150. A low-power PFC (power-factor-correction) stage, using transition-mode-PFC controller L6562, delivers the input-bus voltage of 380V dc. This configuration ensures not only the dc-voltage bus for the power stage of the charger, but also a high power factor during the charging phase. EDN

REFERENCES

1. Lan, Rayleigh, and Hunter Chen, "Flyback Charge Xenon Flash Capacitors," *Power Electronics Technology*, March 2007, http://powerelectronics.com/mag/703PET24.pdf.
2. Creel, Kirby, "Expedite Transformer Calculations for Flybacks," *Power Electronics Technology*, January 2008, http://powerelectronics.com/passive_components_packaging_interconnects/magnetics/transformer-calculations-defibrillator-capacitors-0108.
3. Sokal, NO, and R Redl, "Control algorithms and circuit designs for optimally flyback-charging of an energy storage capacitor (e.g., for flash lamp or defibrillator), *IEEE Transactions on Power Electronics*, Volume 12, Issue 5, September 1997, http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?tp=&arnumber=623007&isnumber=13547.
4. "L6565 Quasi-Resonant SMPS Controller," STMicroelectronics, http://eu.st.com/stonline/products/literature/ds/7587/l6565.pdf.

First-event detector has automatic-reset function

Vasil Borodai, Zaporozhje, Ukraine

The circuit in **Figure 1** lets you indicate which game player presses a button first. Each button has a corresponding LED that indicates the pressing of the button. All other LEDs remain locked out until someone presses a reset button. When a player presses a pushbutton, the corresponding optoisolator turns on, which illuminates the appropriate indicator LED. The LED remains on after the player releases the pushbutton. The voltage at Point A pulls down to nearly 3.7V, which you determine by adding the forward voltage of the optoisolator's internal LED, the phototransistor's voltage, and the LED's voltage: 1.3+0.6+1.8V=3.7V. The green LED then turns off.

Beginning at time T_1 (**Figure 2**), no other player can change the situation by pressing a pushbutton because switching on any other optoisolator requires a voltage exceeding 3.9V. Resistor R_1 depends on V_{PS} such that $R_1=(V_{PS}-V_{D1})/I_{OPTOLED}$, where V_{PS} is the power-supply voltage, V_{D1} is the voltage of diode D_1, and $I_{OPTOLED}$ is the current of the optoisolator LED. Thus, for a 9V power supply, R_1 has a value of 1.5 kΩ. When a player presses the reset button, the player LEDs turn off, and the green LED illuminates. The voltage at Point A returns to 9.2V (time T_2 in **Figure 2**).

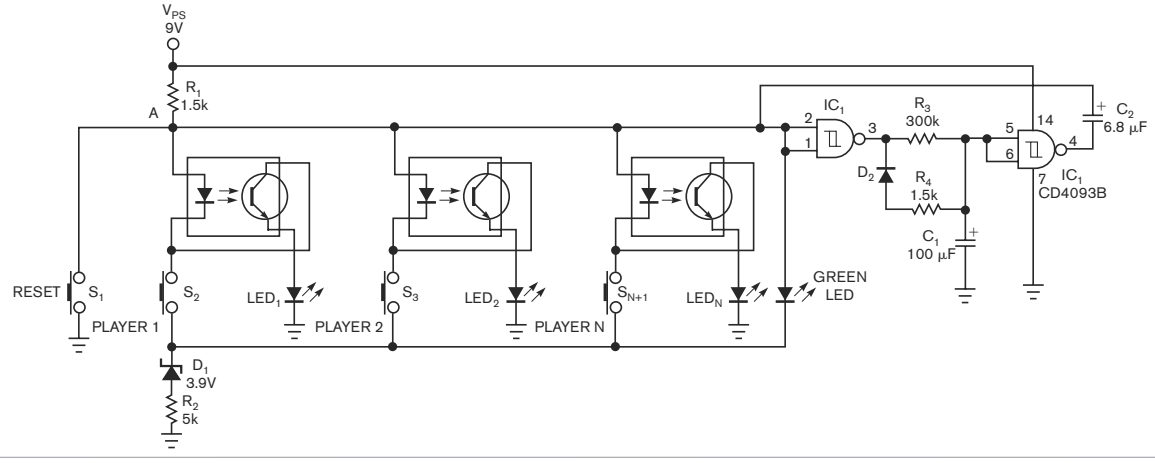

Figure 1 This circuit lets you indicate which game player presses a button first.

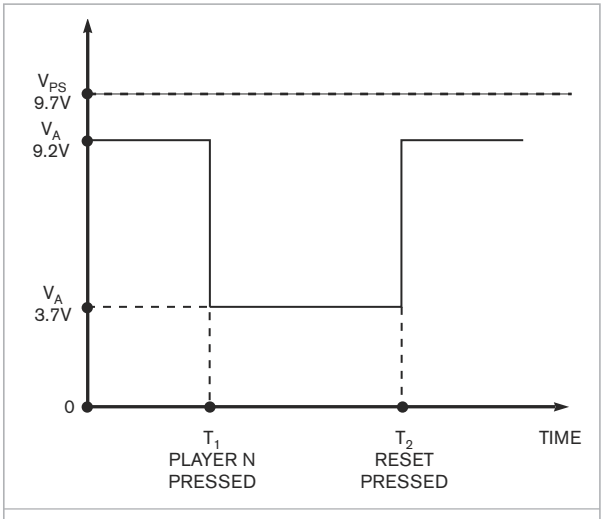

Figure 2 When a player presses the reset button, the player LEDs turn off, and the green LED illuminates. The voltage at Point A returns to 9.2V.

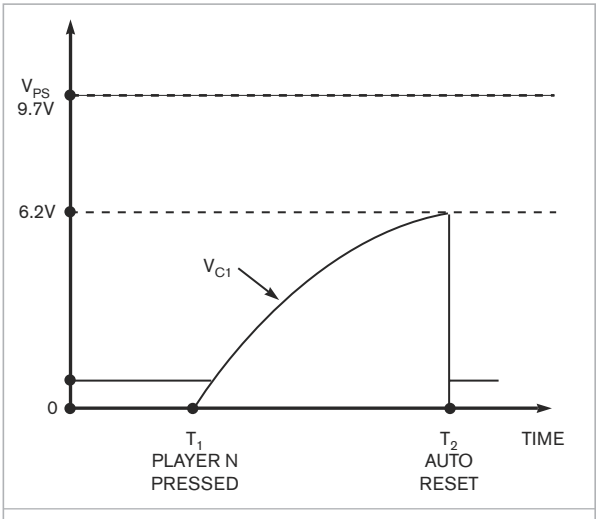

Figure 3 When Point A drops to 3.7V, the inputs at IC_1, pins 1 and 2, go low, and the output at Pin 3 goes high, charging C_1.

You can also add an auto-reset feature to the circuit. When Point A drops to 3.7V (time T_1 in **Figure 3**), the inputs at IC_1, pins 1 and 2, go low, and the output at Pin 3 goes high, charging C_1. After about 30 seconds (time T_2 in **Figure 3**), C_1 has enough voltage to force IC_1's Pin 4 low. R_3 and C_1 determine the charging time. A pulse of current flows through C_2, which forces the voltage at Point A to nearly 2V. That action momentarily interrupts the current in any optoisolator LED. As a result, the circuit automatically resets, and the green LED lights. IC_1's Pin 3 goes low, which discharges C_2 through R_2, resetting the circuit to its original state. EDN

Signal-powered linear optocoupler provides isolated control signal

Mitja Rihtarsic, Škofja Loka, Slovenia

The circuit in **Figure 1** provides an isolated control voltage, such as 0 to 10V. In the low part of the range, 0V to approximately 2V, the controlled device is off. Therefore, the upper part of the range must be as linear as possible. You can meet this requirement using a linear optocoupler, such as Vishay's (www.vishay.com) IL300 or Avago Technologies' (www.avagotech.com) HCNR200 or HCNR201.

These optocouplers each comprise an LED and a photodiode on the transmitting side and an identical photodiode on the receiving side. Because of this construction, the emitted light from the LED should cause the same current to flow in both photodiodes. The current through the photodiode on the receiving side, feedforward current I_{FF}, is the output current, and you must set this current in proportion to the transmitted signal voltage, V_1. This current equals the feedback current, I_{FB}, through the transmitter-side photodiode. A feedback loop around the emitting side of the optocoupler keeps the feedback current in proportion to the transmitted signal. When the feedforward current and the feedback current are equal, the output current is proportional to the transmitted signal.

The hidden cost, however, is a power supply. You need some power on both

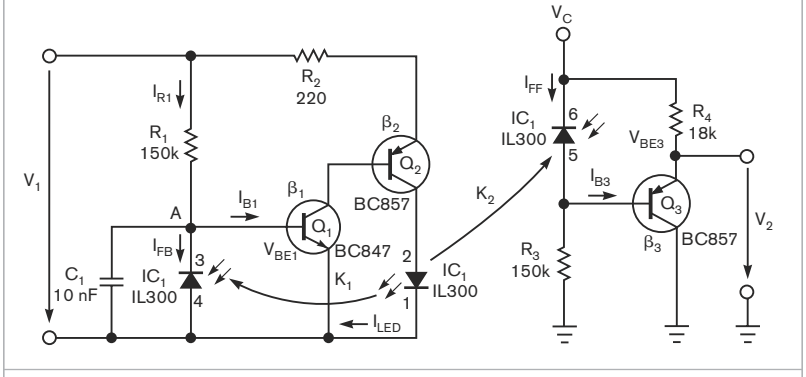

Figure 1 Optocoupler IC_1 isolates the control circuit's input and output.

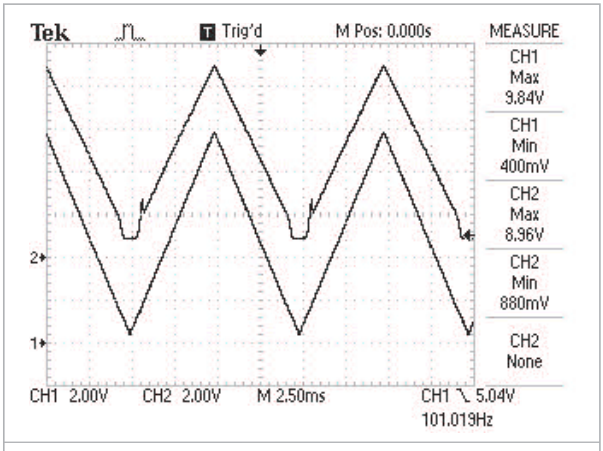

Figure 2 The output voltage (upper trace) turns off when the input voltage (lower trace) gets too low.

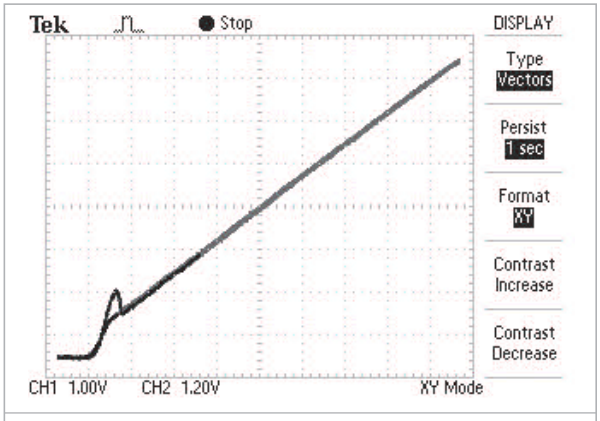

Figure 3 An XY plot of the circuit's input and output voltages shows linearity once the input voltage is high enough to power the circuit.

sides of the signal path. The circuit in this Design Idea uses power from signal voltage V_1 to supply a feedback loop in the transmitting side similar to the way some circuits in a 4- to 20-mA loop get power from the loop current. Both photodiodes operate in reverse-biased, photoconductive mode. The currents through them are proportional to incident-light flux, which feedback gain K_1 and forward gain K_2 describe.

$$K_1 = \frac{I_{FB}}{I_{LED}}, \quad (1)$$

where I_{LED} is the LED's current, and

$$K_2 = \frac{I_{FF}}{I_{LED}}. \quad (2)$$

A description of the circuit begins with a sum of the dc currents at Node A.

$$\frac{V_1 - V_{BE1}}{R_1} = I_{FB} + I_{B1}. \quad (3)$$

The gains of both transistors amplify current I_{B1} into the base of Q_1. The amplified current then flows through the LED.

$$I_{LED} = \beta_1\beta_2 \times I_{B1}. \quad (4)$$

Equations 1 through **4** yield the output feedforward current:

$$I_{FF} = (V_1 - V_{BE1}) \times \frac{1}{R_1} \times \frac{K_2\beta_1\beta_2}{K_1\beta_1\beta_2 + 1}. \quad (5)$$

When the product of feedback gain K_1 and transistor gains β_1 and β_2 is much greater than one, you can cancel out the transistors' gains, yielding a characteristic that is linear:

$$I_{FF} = (V_1 - V_{BE1}) \times \frac{1}{R_1} \times \frac{K_2}{K_1}; K_2\beta_1\beta_2 >> 1. \quad (6)$$

The ratio of feedback gain K_1 and forward gain K_2 is transfer gain K_3. Because K_1 and K_2 are similar, K_3 is approximately one. In reality, K_3 may deviate, but it changes less than K_1 or K_2 alone:

$$K_3 = \frac{K_2}{K_1}. \quad (7)$$

Equation 6 subtracts the base-to-emitter voltage from the input voltage. Although the base-to-emitter voltage is not constant, it is desirable to remove it. You accomplish this task using the emitter follower in the receiving circuit. The output voltage, V_2, is a sum of voltage across R_3 and the base-to-emitter voltage of Q_3.

$$V_2 = (I_{FF} + I_{B3}) \times R_3 + V_{BE3}. \quad (8)$$

You can use a different **equation** to yield the feedforward output current:

$$I_{FF} = \left((V_1 - V_{BE1}) \times \frac{1}{R_1} - I_{B1}\right) \times K_3. \quad (9)$$

You can rearrange **equations 8** and **9** as:

$$V_2 = V_1\frac{R_3}{R_1}K_3 + \left(V_{BE3} - V_{BE1}\frac{R_3}{R_1}K_3\right) + R_3(I_{B3} - I_{B1}K_3). \quad (10)$$

In the first term in **Equation 10**, the ratio of resistors R_3 and R_1 is approximately 1-to-1. You must be careful with the transfer gain, K_3, which is the reason that K_3 remains in **Equation 11**.

$$V_2 \approx K_3 \times V_1. \quad (11)$$

When K_3 is one, voltages V_{BE1} and V_{BE3} cancel each other to some degree. Therefore, **Equation 11** omits the second term in **Equation 10**. Base current I_{B3} depends on resistor R_4 and the output load. When you can set both base currents to be equal, the last term would cancel out, too. The values of resistor R_2 and capacitor C_1 must be small enough so that transistors Q_1 and Q_2 don't saturate. C_1 enhances stability.

Figure 2 shows the necessary voltage for the circuit to begin operation. The output voltage (upper trace) has flatness at its lowest voltages as opposed to the input voltage (lower trace). **Figure 3** shows the two signals' linearity. Dividing the measured maximum of voltages V_1 and V_2 yields 0.91V. A test circuit uses an IL300, which has a gain of 0.851 to 0.955. The measurement meets the requirements of **Equation 11** despite the **equation**'s simplifications. EDN

Dark-activated switch needs only three components

Abel Raynus, Armitron International, Malden, MA

Assume that you have a device that receives its power from the main 120 or 220V-ac line and you need to add a switch between the ac line and the device so that the device works only when it is dark. Although you may think this task would be trivial, it is difficult to find a workable approach because most of the published schematics need 6 to 12V-dc power supplies and relays. Several off-the-shelf dark-activated switches, such as devices from Suns International (www.suns-usa.com), are available, but they're expensive for a consumer product. After looking at products from dozens of Web sites, you may decide to make your own. The solution is simple and inexpensive.

The circuit in **Figure 1** employs an internally triggered triac, which Teccor Electronics (www.teccor.com) originally developed. The primary purpose of any triac is bidirectional-ac switching. The Quadrac triac has a built-in triggering device with the threshold-voltage level of approximately 40V. To achieve this level, the circuit uses a voltage divider comprising a photocell and resistor R_1. When you light the photocell, its voltage drop is lower than the triggering level of the threshold voltage, and Q_1 is locked, so the load disconnects from the ac line. When it becomes dark, the peak voltage amplitude on the photocell increases to 40V, opening Q_1 and making the load connect to the power line.

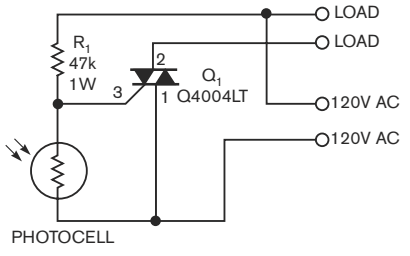

Figure 1 This dark-activated switch needs only a photocell, a resistor, and a triac to switch between the ac line and the device.

The choice of Q_1 depends on the load current and ac-line voltage. This circuit uses the Q4004LT from Littelfuse (www.littelfuse.com) with a maximum current of 4A rms and a voltage of 400V. You can use any photocell, but this circuit uses an off-the-shelf model and accordingly uses a value of 47 kΩ for R_1 to achieve reliable switching. For an inductive load, add a 100Ω resistor in series with a 0.1-μF capacitor between pins 1 and 2 of Q_1. EDN

productroundup

PASSIVES

Inductors shield EMI-susceptible applications

Using an integral ferrite shield, the 1200RS and the 1200LRS inductor series target power-line filtering in low- to medium-current applications, such as switching power-supply circuits, that are susceptible to EMI. The 1200LRS series features 6.8-µH to 22-mH inductances and a 4.7A-dc current rating; the 1200RS series provides 10-µH to 1-mH inductances and a 4.6A-dc maximum current rating. Each inductor series has a 10.5-mm diameter×8.5-mm through-hole, radial-leaded package. The devices operate over a −40 to +85°C temperature range, and prices range from 41 to 50 cents (high volumes).

Murata Power Solutions, www.murata-ps.com

Multitapped silicon-chip resistors target use in hybrid electronics

The 110R series of silicon-chip resistors uses high-stability, self-passivating, and moisture-resistant tantalum-nitride resistor elements, enabling the development of quick-turnaround prototypes in hybrid-microelectronics labs. Comprising a string of 20 resistors, each chip allows wire-bonding for a range of resistances. The strings include a set of 10 2.5-kΩ resistors and a set of 10 25-kΩ resistors, allowing for a 275-kΩ total resistance. Comprising 21 wire-bond pads, the devices provide values of 2.5 to 275 kΩ. Operating over a −55 to +125°C temperature range, the resistors provide a 250-ppm/°C TCR with 10-ppm/°C-maximum-TCR tracking between resistors. Each 110R resistor measures 34×34 mils, and prices range from $3.90 to $7.25, based on tolerance and volume.

OnChip Devices, www.onchip.com

Resistors allow failsafe operation under continuous overloads

Providing the surge-handling characteristics typical of wire-wound devices, the WWF (wire-wound-fusible) series of conformal-coated devices also features a repeatable-fusing characteristic. This feature works well in applications subject to constant or sustained overload, which results in excessive heat buildup. The devices allow adjustments for quick fusing. Thus, they generate minimal heat and slow fusing for applications in which fusing resistors can open too quickly. The devices can withstand high voltage surges and allow inherent increases or decreases in voltage. One device in the series survives 6-kV inrush pulses without failing; however, it fails at 120 or 240V-ac steady-state overloads without flame or excessive heat. Prices for the units range from 20 cents to $1.75 (250 to 2500).

Stackpole Electronics, www.seielect.com

Tantalum capacitors come in surface-mount, molded packages

The M34 and the M35 wet-tantalum capacitors operate over a −55 to +85°C temperature range and at temperatures as high as 125°C with voltage derating with a ±20% standard capacitance tolerance. The M34 provides a capacitance range of 10 µF at 125V to 120 µF at 25V; the M35 has a capacitance range of 1.7 µF at 125V to 220 µF at 6V and features a 3V reverse-voltage capability at 85°C. Available in a surface-mount, molded package, the M34 and the M35 cost $41.15 and $29.10, respectively.

Vishay Intertechnology, www.vishay.com

Power resistor meets high-temperature soldering requirements

Using a low-inductive-resistive film, the PWR263S-20 power resistor meets high-temperature soldering

requirements, enabling improved automated board assembly. Features include 20W power consumption at 25°C, 6.5°C/W thermal resistance, and a 0.5Ω to 130-kΩ resistance range. A D²Pak surface-mount package suits space-constrained designs requiring a lead-free soldering temperature. Prices for the PWR263S-20 power resistors range from $1.75 to $2.55 (5000).

Bourns, www.bourns.com

INTEGRATED CIRCUITS

Clock buffers have 75- and 100-fsec jitter

The 12-channel ADCLK954 LVPECL and the ADCLK854 LVDS/CMOS clock-fanout buffers provide 75- and 100-fsec jitter, respectively. The clock buffers feature a 9-psec skew, suiting use in high-speed ADCs and DACs, wireless-infrastructure equipment, medical-imaging applications, and industrial applications. The 4.8-GHz ADCLK954 provides two selectable differential inputs using the input-select control pin. Both inputs have 100Ω on-chip termination resistors. The clock buffers operate over a −40 to +85°C temperature range. The ADCLK954 comes in a lead LFCSP-48 package with 12 or 24 channels and costs $5.95 (1000); the ADCLK854 comes in a lead LFCSP-40 package with 12 channels and costs $6.95 (1000).

Analog Devices, www.analog.com

12- and 16-bit precision DACs are pin- and software-compatible

The pin- and software-compatible 16-bit MAX5138 and 12-bit MAX5139 precision DACs accept three-wire SPI-, QSPI-, Microwire-, and DSP-compatible serial interfaces. Suiting industrial-control and -automation applications, the devices include a power-down feature that reduces power consumption by 200 nA and an integrated voltage reference with a 100-ppm/°C temperature coefficient. Measuring 3×3 mm, the MAX5138 and MAX5139 cost $3.39 and $1.60 (1000), respectively.

Maxim Integrated Products, www.maxim-ic.com

EDN ADVERTISER INDEX

Company	Page
Agilent Technologies	2
	22
	C-3
Allied Electronics	19
austriamicrosystems AG	45
Avnet Electronics Marketing	27
BuyerZone	35
Coilcraft	9
CST GmbH	C-2
Digi-Key Corp	1
Electro Technik	26
EMA Design Automation	31
Express PCB	39
Ironwood Electronics	47
Jameco Electronics	21

Company	Page
Keil Software	32
Linear Technology Corp	C-4
LS Research	39
Maxim Integrated Products	6,7
Mentor Graphics	12
Mill Max Manufacturing Corp	11
Mouser Electronics	4
National Instruments	33
Pico Electronics	5
	34
Trilogy Design	47
Xilinx Inc	10

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

EDN product mart

This advertising is for new and current products.

"Dog" PLL chases its own tail

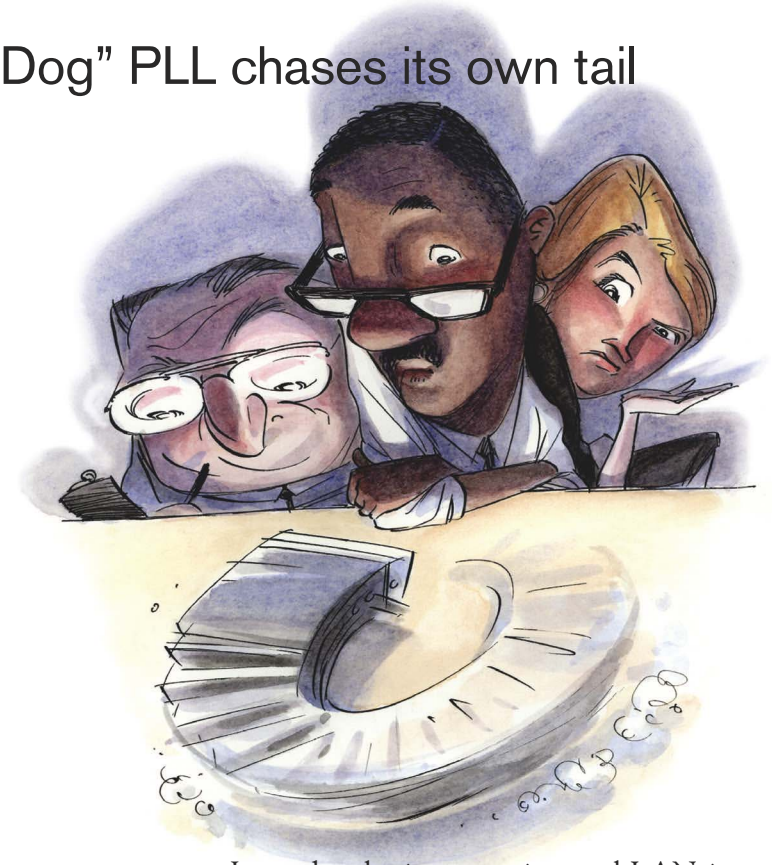

Long ago, I was developing experimental LAN-interface hardware to transmit data on a synchronous-RF carrier from a master unit to remote devices that simultaneously returned data on a synchronous-RF carrier two octaves lower. I amplified the received carriers to TTL (transistor-transistor-logic) levels for digital-PLL (phase-lock-loop) CDR (clock and data recovery).

One of the biggest problems was getting the master unit's receiver PLL to recover the returned carrier without jittering all over the place when the loop-filter bandwidth was small. No matter what I tried, the LC (inductor/capacitor) oscillator on the wire-wrapped breadboard simply would not run in a stable fashion at a narrow loop bandwidth. It insisted on wild phase gyrations around the phase-lock point.

I suspected that synchronous digital noise from both the local transmitting oscillator and returned carrier was pulling the PLL oscillator. I verified this suspicion by disconnecting the PLL-control voltage from the oscillator varactor, leaving only the assumed stray noise coupling. Sure enough, when I manually tuned the oscillator close to the operating frequency, the device locked strongly and stably to the digital noise without jitter.

Rebuilding the LC oscillator in a metal box to fully shield it from digital electromagnetic fields and filtering the power-supply- and control-voltage inputs resulted in an amazing improvement. I learned that you must protect PLL oscillators from digital-synchronous-noise influences from supply rails and stray electromagnetic coupling.

About a year later, I was developing hardware that exchanged data between a remote slave and a master unit using two optical fibers. I had designed my receiver's clock recovery for both the slave and the master to rely on well-behaved LC tanks rather than ornery PLLs. A colleague, on the other hand, had designed the master unit's optical-link clock generator to use a PLL digital IC with an RC (resistor/capacitor) oscillator. The master clock rate was 16.384 MHz, but the optical links required 19.456 MHz for pattern-synchronization overhead, and the budget did not allow for a more stable VCXO (voltage-controlled-crystal oscillator).

The optical link's transmitter PLL insisted on doing those wild phase gyrations. "Impossible," I thought. All noise transients from digital transitions occur just after the master oscillator's switching transition, so, in theory, it should not be susceptible to its own noise, but it was. Disabling the master's remote-unit receiver allowed the PLL to run stably and verified the cause of the problem: Its own stray digital noise was returning to haunt it through the long path to the remote unit and back again. This round-trip delay was unpredictable; every 10m of cable length was equivalent to a 360° shift in returned noise relative to the clock period. Because the oscillator was divided by 19, phase detection occurred only once every 19 clock cycles. In the intervening cycles, the oscillator was free to become a happy wanderer. With the remote unit 1 km away, the oscillator was tracking the influence of its phase from 10 µsec and 200 clock periods earlier. This PLL was chasing its own tail.

We tamed the problem by building the VCO on a small PCB (printed-circuit board) with an unbroken bottom ground plane supported on a standoff above the main PCB. The oscillator's PCB ground plane helped to shield the topside circuit from the evil digital influences below. EDN

REFERENCE

1 Rabinovich, Rick, "Power-rail filtering improves PLL performance," *EDN*, March 19, 2009, www.edn.com/article/CA6645280.

Contact design consultant Glen Chenier at glen@teetertottertreestuff.com.

DANIEL VASCONCELLOS

“I need to see the most signal detail.”

Get 15% OFF until 9/30. Starting at $4246

Look no further.

InfiniiVision scopes are engineered for the best signal visibility.

It’s a bold claim — let us prove it to you.
Agilent engineered the custom MegaZoom III technology in our Agilent InfiniiVision scopes so you get the best possible signal visibility on every measurement you make.

1. **Fastest uncompromised update rate** shows you critical signal detail and gives you statistically valid measurement results in seconds.
2. **Industry’s only hardware accelerated mixed signal oscilloscope** gives you analog, logic, and protocol views.
3. **Customize your scope** with software applications including mask testing, RS-232/UART, CAN/LIN and more.

100 MHz to 1 GHz MSO and DSO models

Take the 5-minute Scope Challenge online and register to win a free Agilent MSO7104 scope.
www.agilent.com/find/bestscopes

Prices in USD and subject to change without notice.

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

Agilent Technologies

V_{IN} = 1.8V to USB

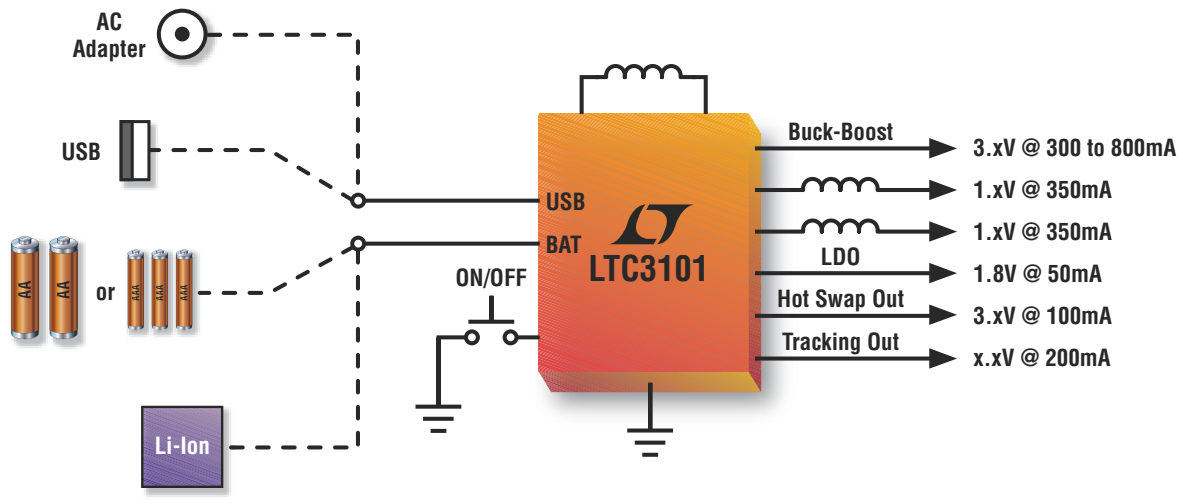

95% Efficient & 38µA Total I_Q

Our LTC®3101 brings a new level of performance to devices that can be powered from multiple sources: 2 or 3 AA (or AAA) cells, a Lithium-Ion/Polymer battery, an AC adapter or a USB port. Its switching PowerPath™ manager seamlessly transitions between the different input sources while maintaining two "always alive" rails, a 50mA LDO and a 200mA MAX output that tracks the higher voltage input. It also provides three high efficiency switching DC/DC converters, a protected 100mA Hot Swap™ output and pushbutton On/Off control in a compact 4mm x 4mm QFN package.

Features

- Seamless & Automatic Transition Between Multiple Input Sources
- V_{IN} Range: 1.8V to 5.5V
- Buck-Boost Converter with V_{OUT} Range: 1.5V to 5.25V
- Dual 350mA Sync Buck Regulators with V_{OUT} Down to 0.6V
- Pushbutton On/Off Control
- 4mm x 4mm x 0.75mm QFN-24

Efficiency Curve

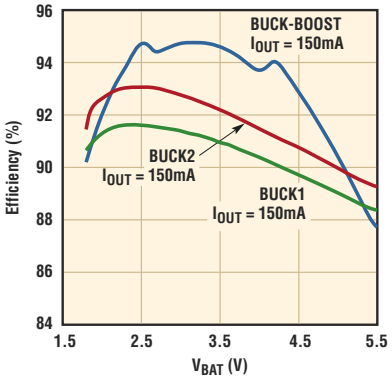

Info & Free Samples

www.linear.com/3101

1-800-4-LINEAR

Free Portable Solutions Brochure

www.linear.com/portsolutions

LINEAR TECHNOLOGY