Reed Business Information (RBI™)

EDN®

VOICE OF THE ENGINEER

NOV 26

Issue 22/2009

www.edn.com

FROM MAGNETIC TO SOLID STATE, SPIN-FREE:

WHAT A LONG, STRANGE STORAGE TRIP IT'S TURNING OUT TO BE

Page 24

LESSONS FROM THE LAST MILE

Page 18

EVALUATING ESD-PROTECTION COMPONENTS

Page 33

MAGNETICS IN SWITCH-MODE POWER SUPPLIES

Page 36

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.

WHERE DESIGN ENGINEERS GO FIRST!

DIGIKEY.COM

Digi-Key® CORPORATION

#1 Breadth of Product Line
#1 Availability of Product
#1 Overall Performance

Source: EETimes Distribution Study, August 2008

The industry's broadest product selection available for immediate delivery

www.digikey.com
1.800.344.4539

EDN 11.26.09 contents

From magnetic to solid state, spin-free: What a long, strange storage trip it's turning out to be

Lessons from the last mile

Evaluating ESD-protection components: Clamping voltage and dynamic resistance are crucial

What every designer should know about magnetics in switch-mode power supplies

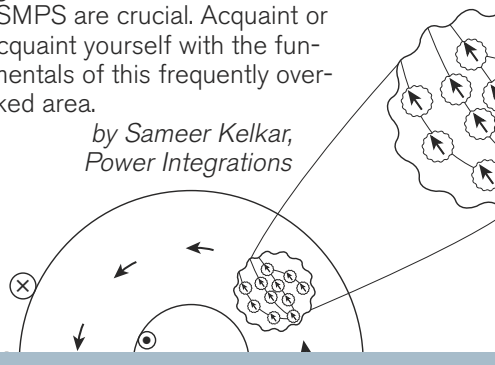

pulse

DESIGNIDEAS

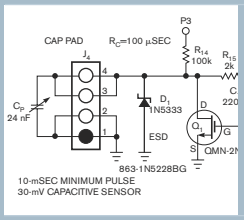

How to pick the perfect inductor for your LED driver application

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

The Newest Optoelectronics

OSRAM PANJIT TOUCH SCREENS VISHAY OPTEK LEDENGIN KINGBRIGHT AVAGO LITE-ON LUMEX Lighting Science Group NORITAKE TAOS SHARP BIVAR VCC CML

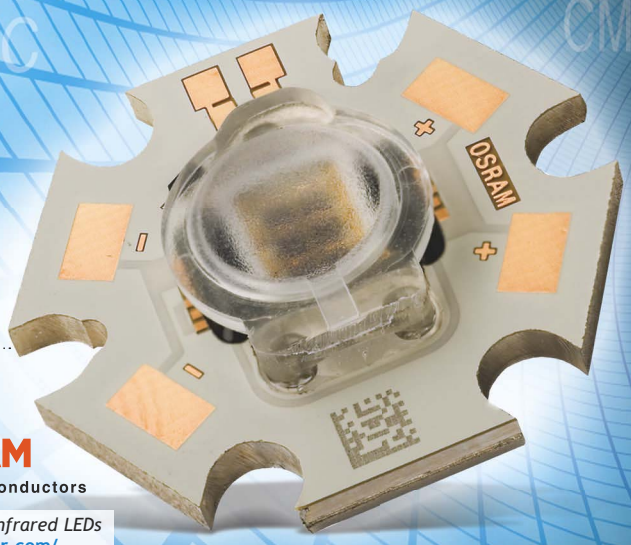

OSRAM
Opto Semiconductors

IR OSTAR® Infrared LEDs
www.mouser.com/
osramirostar

PACIFIC SILICON SENSORS 3M TOUCH SYSTEMS NEWHAVEN DISPLAY

www.mouser.com
Over A Million Products Online

New Products from:

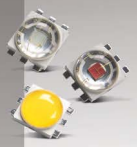

AVAGO TECHNOLOGIES
ASMT-Jx1x 1W Mini Power LED Light Source
www.mouser.com/
avago1wminiled

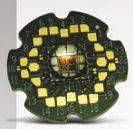

LedEngin, Inc.
40W LED Emitter
www.mouser.com/
ledengin40w

LUMEX
QuasarBrite™ Narrow Beam LEDs
www.mouser.com/lumex_
quasarbrite

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 11.26.09

DEPARTMENTS & COLUMNS

EDN online contents www.edn.com

ONLINE ONLY

Check out these Web-exclusive articles:

Specify an external reference clock to improve SERDES performance
Various serial data standards require different clock specifications.
→www.edn.com/article/CA6703242

Oscilloscope probe accessories—It's the little things that matter
Probe accessories, including tips, caps, adapters, springs, positioners, clips, and hooks, can help you conveniently make accurate measurements.
→www.edn.com/article/CA6704367

INNOVATION WANTED NOW!
Nominations for the 20th annual Innovation Awards
www.edn.com/innovation

WHAT'S UP AT EDN.COM

Visit www.edn.com and

- Sign up for our **free newsletters** or to receive your subscription **digitally**
- Check out our Hot Topics to get all the information you need about **what's HOT** right now
- Check out the **News Center**, which provides up-to-the-minute news, analysis, research, and insight on the global electronics industry
- Check out the **Design Centers**, which feature articles, new-product information, and valuable design resources in several areas

and much, much **more**!

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. John Poulin, CEO/CFO, RBI-US; Jeff DeBalko, President, Business Media; Jane Volland, Vice President, Finance. 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.

BY RICK NELSON, EDITOR-IN-CHIEF

Weak economy, anti-immigrant sentiment hit H-1B visa program

The woeful employment picture in the United States is resulting in thousands of unfilled spots in the H-1B visa program for the first time since 2003, according to a recent article in *The Wall Street Journal* (**Reference 1**). Although employers in just one day snapped up all 65,000 available visas, would-be immigrants filed only 46,700 petitions for employment as of Sept 25—about six months after employers scooped up the visas.

The article notes that, in addition to the weak economy, rising anti-immigrant sentiment in Washington and the higher costs of hiring foreign-born workers are also taking their toll on the visa program. Indian outsourcing companies such as HCL have traditionally been the largest recipients of H-1B visas, according to the article, but HCL has been hiring Americans who otherwise may have faced layoffs from companies switching work to HCL.

Would-be immigrants are also finding more opportunities at home. The article quotes Vivek Wadhwa, a scholar who has studied H-1B visas, as saying, "The best and the brightest who would normally come here are saying, 'Why do we need to go to a country where we are not welcome, … our quality of life would be less, and we would be at the bottom of the social ladder?'"

I commented on the trend for foreign nationals to stay home when I reported on a study Wadhwa conducted for the Ewing Marion Kauffman Foundation (**references 2** and **3**). The study notes that immigrant-founded US-based companies employed 450,000 workers and generated $52 billion in revenue in 2006. The *WSJ* article quotes Microsoft general counsel Brad Smith as saying that 35% of Microsoft's US patent applications last year came from new inventions by visa and green-card holders.

"While some have tried to associate the increase in foreign workers ... with the economic problems that have plagued the country, this data verifies the opposite effect," said Wadhwa when the Kauffman Foundation released its study. "If the US government and the business community could find better ways to offer good jobs in tandem with less restrictive visa policies for talented immigrants, the United States might be able to recapture many of these immigrants and their potential to help grow the US economy."

Companies such as Microsoft that benefit from the visa program contend that the current slump in the program demonstrates that the market, not Congress, should determine how many immigrants should be allowed to work in the United States. The *WSJ* article quotes Jenifer Verdery, director of work-force policy at Intel, as saying that the fact that the cap hasn't been reached this year shows that the market will temper demand.

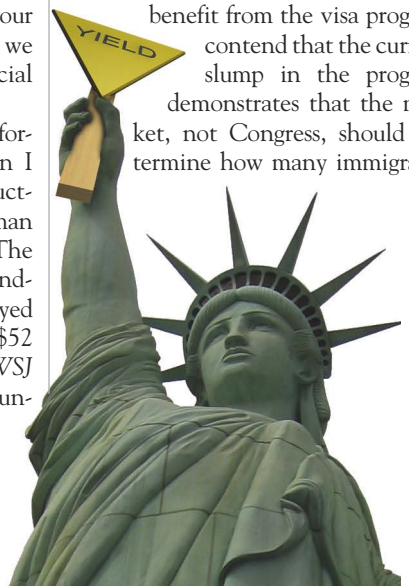

There seems to be bipartisan disagreement in Congress with that position. As the *WSJ* reports, Senator Charles Grassley, an Iowa Republican, wrote a letter to the new director of citizenship and immigration services, urging tighter controls on H-1B visas. In April, Grassley and Illinois Democrat Senator Richard Durbin introduced legislation to require companies to pass more stringent labor-market tests that would ensure they make a bigger effort to hire US workers.

The H-1B visa program is valuable, and, as the Kauffman Foundation study points out, immigrants have contributed disproportionately to the US economy's high-tech sector. If Americans are unwilling or unable to contribute their fair share, then it will be important to US economic success to attract talent from overseas. There is a role for Congress to play to provide further safeguards so that cheaper workers from abroad don't displace motivated, qualified Americans. If Congress can ensure Americans that the program works as intended, political support for expanding the program might grow. EDN

REFERENCES

1 Jordan, Miriam, "Slump Sinks Visa Program," *The Wall Street Journal*, Oct 30, 2009, http://online.wsj.com/article/SB125677268735914549.html.

2 Nelson, Rick, "Immigrant brain-drain challenges US innovation," *Test & Measurement World*, March 2, 2009, www.edn.com/091126eda.

3 Wadhwa, Vivek, et al., "America's Loss is the World's Gain," Ewing Marion Kauffman Foundation, March 2009, www.kauffman.org/uploadedfiles/americas_loss.pdf.

Contact me at rnelson@reedbusiness.com.

EDN

PRESIDENT, BUSINESS MEDIA, REED BUSINESS INFORMATION
Jeff DeBalko, jeff.debalko@reedbusiness.com
1-646-746-6573

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

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,
Web Operations Specialist
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
+44 118 935 1650;
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@rbichina.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

pulse

INNOVATIONS & INNOVATORS

WinSystems highlights Atom board, CompactFlash

WinSystems has announced its EBC-Z8510-G single-board computer, which includes an Intel (www.intel.com) Atom processor. The device measures 203×147 mm and supports the new SUMIT-ISM (stackable-unified-modular-interconnect-technology-industry-standard-module)-I/O-expansion standard plus COMIT (computer-on-module interconnect technology), which the SFF-SIG (Small Form Factor Special Interest Group, www.sff-sig.org) defines.

The $795 EBC-Z530-G includes an array of onboard peripherals and expansion options. It uses either a 1.1-GHz or a 1.6-GHz Atom and the SCH (system-controller hub)-US15W with 512 Mbytes or 1 Gbyte of DDR2 system memory. The EBC-Z510-G's I/O interface features two GbE (gigabit-Ethernet) ports, CRT and LVDS (low-voltage-differential-signaling) flat-panel video, a MiniPCIe (Peripheral Component Interconnect Express)-card interface for a wireless-networking module, four USB (Universal Serial Bus) 2.0 ports, four serial COM ports, HD (high-definition) audio, a PATA (parallel-advanced-technology-attachment) controller for both a CompactFlash and a hard disk, 48 lines of digital I/O, a parallel printer port, and a PS/2 port. Two SUMIT and legacy PC/104 connectors support additional I/O-module expansion. The ROHS (reduction-of-hazardous-substances)-compliant board operates over an industrial temperature range of −40 to +70°C for processor- and I/O-intensive applications in harsh environments.

Because the EBC-Z510-G's architecture is PC-compatible, it supports Windows XP embedded and Linux operating systems along with a software-development tool set that includes device drivers and libraries. It also supports advanced features, such as a custom splash screen, APM (advanced-power-management) and ACPI (advanced-configuration-and-power-interface) modes, and PXE (preboot execution environment).

According to WinSystems' vice president, Robert Burkle, the EBC-Z510-G is the first board to support COMIT, which targets use in SFF processor modules and baseboards. The company uses a 62×75-mm SFF-COM card, which is roughly the size of a credit card, that includes the Atom, SCH, memory, and power supplies. For more on this introduction, go to www.edn.com/article/CA6699174.

—by Rick Nelson

▷**WinSystems**, www.winsystems.com.

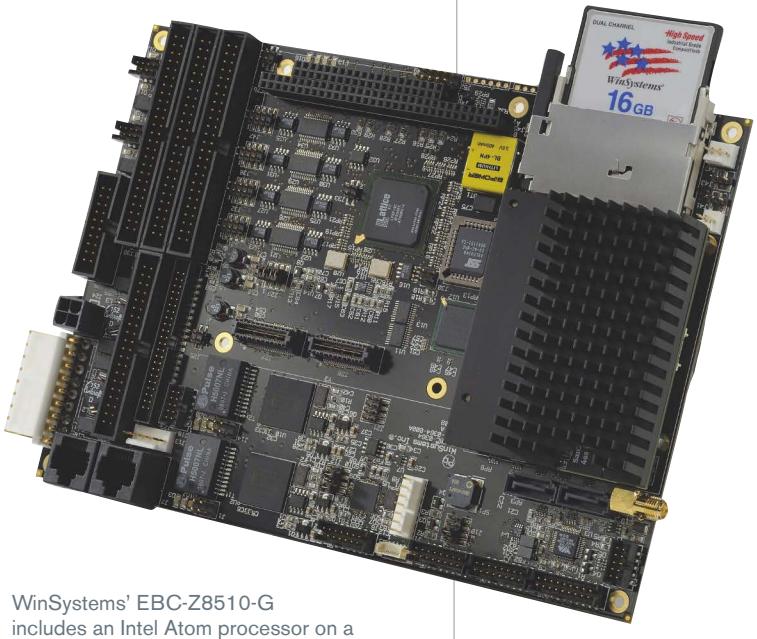

WinSystems' EBC-Z8510-G includes an Intel Atom processor on a 203×147-mm board. It supports SUMIT-ISM and COMIT, which the SFF-SIG defines.

FEEDBACK LOOP

"The debunker dog in you should get busy. ... If I'm wrong I'll eat my words (or breathe the fumes—whatever)."

—Engineer Meredith Poor, in *EDN*'s Feedback Loop, at www.edn.com/article/CA6670951. Add your comments.

High-performance MSOs feature 20-GHz analog bandwidth

Tektronix has announced the MSO70000 series of MSOs (mixed-signal oscilloscopes). The instruments can capture as many as 20 channels of data—four analog with bandwidth ranging from 4 to 20 GHz, depending on the model, and 16 digital with timing resolution of 80 psec on all models. Memory depths to 250M points are available on all channels of all models.

The MSO70000 combines a full suite of measurement capabilities that help resolve analog issues in digital systems. You can use the instruments to debug and verify in such demanding, high-speed design applications as DDR memory, high-performance ASICs, FPGAs, SOCs (systems on chips), and digital RF. The MSO70000 offers a variety of probing accessories for making minimally disruptive analog and digital connections to a DUT (device under test).

The instruments deliver advances in the discovery of problems, capture of notable events, quick searches through long records to reveal the captured events, and analysis to obtain rapid insight into the causes of anomalous DUT behavior. The devices offer as much as five times the bandwidth and timing resolution of the fastest-available integrated MSOs.

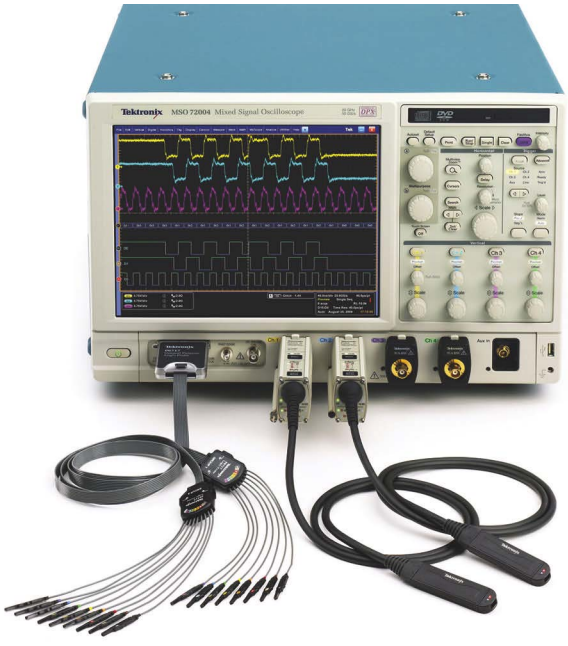

On each of their four analog channels, MSO70000 series instruments offer five times the bandwidth of other MSOs. Each of the 16 differential-input digital channels provides 80-psec timing resolution and up to 250M samples of capture memory.

Maximum sample rates are 50G samples/sec on analog channels and 12.5G samples/sec on digital channels. To minimize confusion, the analog- and digital-record durations always match; the scopes add repeated samples to the channels that are acquiring at the lower rate so that analog and digital records always contain equal numbers of samples. This feature allows you to capture long-duration events with high sample resolution and obtain time-correlated views of high-speed analog and digital data.

The MSO70000 series provides a comprehensive set of innovative solder-in probe accessories that simplify connecting to vias and fine-pitch components on tightly packed boards to acquire such signals as those on the digital-control lines of the DDR command bus. The tool set for DDR probing now also includes new BGA interposers for all variants of DDR3- and DDR2-memory components and provides access to all signals with excellent fidelity. The units work with the company's iCapture technology, which allows internal routing of selected digital signals to the analog channels for full analog evaluation, making the MSO70000 ideal for highly sensitive, fine-pitch board layouts.

The iCapture feature offers analog views of any connected digital channel, providing debugging insight across all 20 channels. The series provides serial-pattern, mixed analog and digital, logic-pattern, and bus-state triggers, which you can combine to isolate system faults that occur only during particular system states. The units provide tight timing synchronization between the analog and the digital subsystems. Timing correlation as close as 80 psec is possible, resulting in easier determination of the cause and effect of circuit behavior.

More than 30 analysis suites run on the series. You can select from the new I^2C (inter-integrated-circuit) and SPI (serial-peripheral-interface) bus-analysis tools, DPOjet (digital-phosphor oscilloscope jet) for jitter and eye-diagram analysis, DDRA (DDR analysis) for DDR-memory-bus verification, SDLA (serial-data-link analysis) for equalized-channel emulation and analysis, and SignalVu for frequency-domain display and analysis. The manufacturer's suggested US retail prices for the MSO70000 units start at $67,400.

—by Dan Strassberg

▷**Tektronix Inc**, www.tektronix.com.

DILBERT By Scott Adams

Stable quartz oscillator uses SAW technology

Targeting LANs (local-area networks) and SAN (storage-area networks), Epson Toyocom recently announced the highly stable EG-4101/4121CA SAW (surface-acoustic-wave) resonator. The part combines low jitter, low phase noise, high stability, and temperature coefficients better than those of AT-cut quartz crystals. The resonator offers a frequency tolerance of ±50 ppm and maximum phase jitter of 0.2 psec at 622 MHz over a 12-kHz to 20-MHz bandwidth. The device is available with LV-PECL (low-voltage-positive-emitter-coupled-logic), LVDS (low-voltage-differential-signaling), and HCSL (high-speed-current-steering-logic) outputs. Available frequency ranges are 100 to 700 MHz for the LV-PECL- and LVDS-output versions and 100 to 500 MHz for the HCSL version. The resonator has a supply voltage of 2.5 to 3.3V, and current consumption ranges from a maximum of 30 or 45 mA over the supply-voltage range for the LVDS version to 80 or 100 mA for the LV-PECL version.

SAW resonators and oscillators differ from SAW filters in that the resonators use a quartz rather than a ceramic element. Don't confuse SAW resonators with inexpensive silicon or ceramic resonators, which tend to have a much lower Q (quality factor) and worse initial accuracy and temperature coefficients. Because the SAW resonators operate at their fundamental-resonance mode, they lack the frequency jitter of conventional crystal oscillators that operate at a lower frequency; a PLL (phase-locked loop) then multiplies that frequency inside the chip. The operation at fundamental mode also means that the parts do not frequency-hop as quartz crystals do.

The EG-4101/4121CA device operates over a standard temperature range of −40 to +85°C or an optional range of −40 to +90°C and comes in a 7×5×1.2-mm package. It sells for $12 to $18 (1000) and is available for sampling now.

—by Paul Rako

▷**Epson Toyocom**, www.epsontoyocom.co.jp/english.

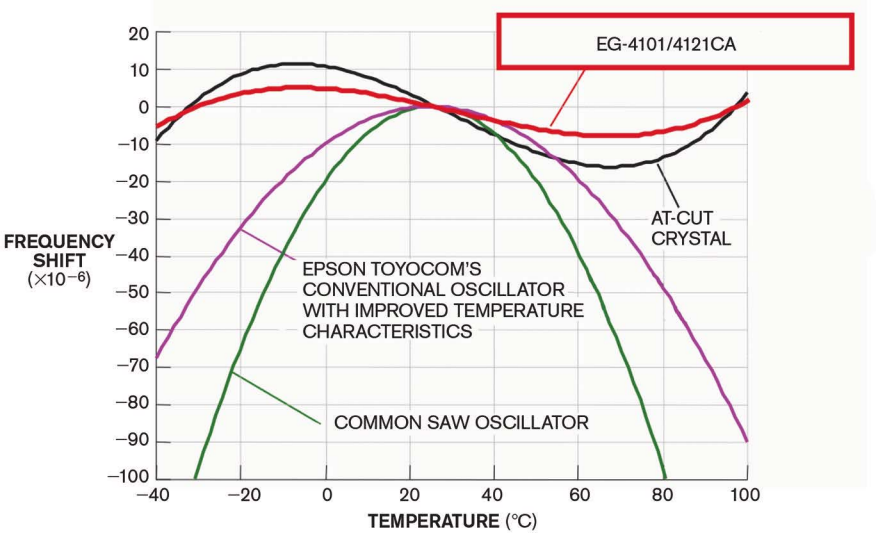

The EG-4101/4121CA series of SAW oscillators from Epson has a flat frequency variation over temperature.

CORTEX-M3 MICROCONTROLLER CUTS ENERGY CONSUMPTION

Energy Micro's new EFM32G Gecko microcontroller family sports an energy-efficient implementation of a 32-bit ARM (www.arm.com) Cortex-M3-microcontroller architecture. The family thus targets applications, such as meters, requiring extended battery life. The devices support five power modes with an operating voltage range of 1.8 to 3.8V. Active-mode current consumption is as low as 180 µA/MHz at 3V when executing code from flash memory. Standby current consumption is 900 nA at 3V with a real-time clock, a 32.768-kHz oscillator, power-on reset, brownout detection, and full RAM and CPU retention. Deep-sleep-mode current draw is 20 nA at 3V, and wake-up time from sleep mode is as fast as 2 µsec.

The 32-MHz microcontroller configurations include as much as 128 kbytes of on-chip flash and 16 kbytes of RAM. Low-power components include a 200-µA, eight-channel, 12-bit, 1M-sample/sec ADC; a 100-nA brownout detector; a 50-nA, 32-kHz real-time counter; a 100-nA-receive-mode, 9600-bps-capable UART; and a 50-nA watchdog timer with dedicated RC oscillator. The ADC supports single-ended or differential operation. The 12-bit, 500k-sample/sec DAC supports two single-ended channels or one differential channel. As many as two analog comparators are available with support for capacitive sensing with as many as eight inputs. As many as 90 GPIOs (general-purpose input/outputs) support a 20-mA drive strength. Hardware AES (Advanced Encryption Standard) with 128/256-bit encryption and decryption is available. The configurable LCD controller can drive an array of 4×40 segments.

Prices for the 32-pin devices start at $1.55 (100,000). For more details on this series, go to www.edn.com/article/CA6704374.—by Robert Cravotta

▶**Energy Micro, www.energymicro.com.**

11.26.09

China's proposed ban of rare-earth metals would affect hybrid cars, CFLs

China's Ministry of Industry and Information Technology is proposing a total ban on exports of terbium, dysprosium, yttrium, thulium, and lutetium and a restriction on neodymium, europium, cerium, and lanthanum to a total of 35,000 tons a year, which is far below global needs. Many of these metals are vital to energy-efficient technology. For example, neodymium finds use in rare-earth magnets for high-efficiency motors, and new front-loading clothes washers use rare-earth magnets in their motors.

According to a recent article (**Reference 1**), "No replacement has been found for neodymium that enhances the power of magnets at high heat and is crucial for hard-disk drives, wind turbines, and the electric motors of hybrid cars. Each Toyota Prius uses 25 pounds of rare-earth elements. Cerium and lanthanum are used in catalytic converters for diesel engines." Manufacturers use terbium in the phosphors of CFLs (compact fluorescent lights) to tweak their light to a more pleasant spectrum.

China is currently the only producer of some of these metals, so the country's restriction or banning of its exports will affect energy-efficient products worldwide. According to the article, China's intent is not to hold the rest of the world hostage; China needs these metals for its internal consumption.

> Neodymium is crucial for hard-disk drives, wind turbines, and the electric motors of hybrid cars.

China put many global competitors in rare-earth minerals out of business in the early 1990s by flooding the market, leading to the closure of the biggest US rare-earth mine, in Mountain Pass, CA, which Molycorp Minerals operates. The mine is one of the world's largest and richest rare-earth deposits, and the company is producing a variety of green elements there. It plans to bring the facility back into full production and re-establish domestic manufacturing capacity.

—by Margery Conner

▷**Molycorp Minerals**, www.molycorp.com.

REFERENCE

1 Evans-Pritchard, Ambrose, "World faces hi-tech crunch as China eyes ban on rare metal exports," *Telegraph*, Aug 24, 2009, www.telegraph.co.uk/finance/comment/ambroseevans_pritchard/6082464/World-faces-hi-tech-crunch-as-China-eyes-ban-on-rare-metal-exports.html.

ONLINE POWER-SUPPLY DESIGN TOOL EVALUATES 48 BILLION DESIGNS

National Semiconductor has made significant improvements to its free Webench online-design tool, which operates with most popular Web browsers. You access the tool using a one-time transfer of a database in flash format to your computer. The tool provides a speedy response as you experiment with various design configurations. It includes the new Visualizer tool to chart efficiency, footprint, and cost variables. An optimizer dial lets you establish your preference for trade-offs among footprint, efficiency, and cost.

You set a dial that causes the tool to generate 50 to 70 designs from 48 billion combinations and to select from 25 power-supply topologies, including buck, boost, buck-boost, SEPIC (single-ended-primary-inductance converter), and flyback. The tool has a database of 21,000 components from 110 manufacturers and is suitable for designs with input voltages of 1 to 100V, output voltages of 0.6 to 300V, and power as high as 300W. It can help you design for efficiency as high as 96% and switching frequencies as great as 3 MHz. The smallest-footprint design is 14×14 mm. Webench requires no registration until you run simulation or thermal analysis; at that point, you need to register for a user account to store the results.—by Paul Rako

▶**National Semiconductor, www.national.com.**

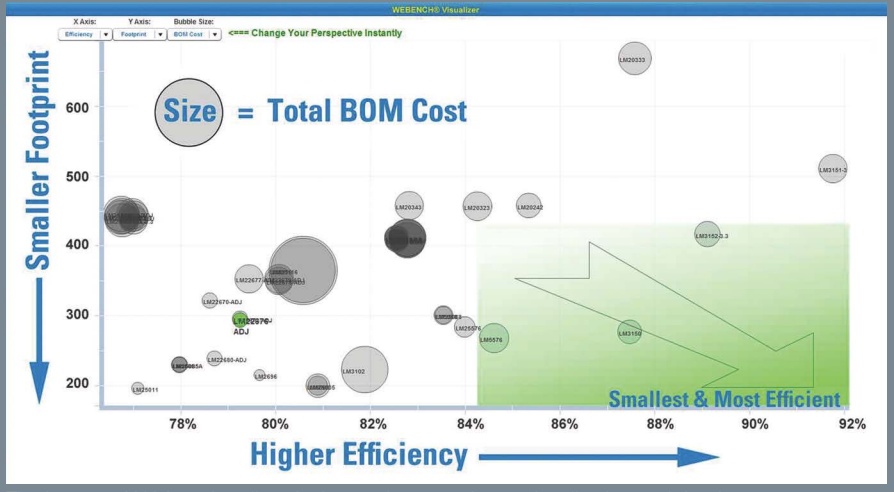

The improved Webench online-design environment has a visualization tool that lets you contrast the trade-offs of price, size, and efficiency in your power-supply designs.

11.26.09

VOICES

Tim Draper: flowing venture capital where it's needed

Tim Draper is managing director of venture capital at Draper Fisher Jurvetson (www.drapervc.com) and chairman of Bizworld (www.bizworld.org), a nonprofit organization that teaches entrepreneurship and business to children. *EDN* recently conducted an interview with him, a portion of which follows. You can read the complete interview at www.edn.com/091126pb.

What is the fundamental trade-off between venture capital and private equity?

A Both are valuable parts of the financial markets. Venture capital invests in entrepreneurs who want to build companies from nothing. Private equity invests to make existing companies more efficient. Venture capital is usually around start-ups and technology. Private equity can be for any established company in any field.

Do you have some patriotic desire to fund entrepreneurs in the United States?

A No. I fund entrepreneurs who want to change the world wherever they may be. In fact, America is driving them away. Technical immigrants on the whole create jobs for Americans. If companies in the United States become uncompetitive globally, we lose jobs.

Do you feel that philanthropic involvement is an important part of being successful?

A I mostly believe in the power of business to improve our lives. Most of my philanthropic activity has been around teaching and encouraging entrepreneurship globally.

Is regulation an essential part of a complex technological society?

A Big government was as responsible as anyone for the crash. Fannie Mae and Freddie Mac guaranteed loans they shouldn't have. Banking regulators changed the Glass-Steagall Act, which encouraged banks and investment banks to merge. In addition, the Community Redevelopment Act created a market for risky subprime loans. You can't regulate good behavior. In fact, I would argue that a freer country has fewer criminals. Our government has gone from spending 8 to 40% of our GDP [gross domestic product] over the last 100 years. Our country in effect trusts itself less than it did, and it is killing our growth.

The liquidity crisis makes it harder for you to cash out with an IPO [initial public offering], but there must be massive amounts of idle capital available for you to invest.

A That is not how it works. Without IPOs, capital

sits in places where people can't trade it, so in effect there is less of it to use for new investments. In our case, our limited partnerships invested their money with us. We then invested that money in tech start-ups. Some of those start-ups grew and created jobs and wealth, but that wealth is sitting in illiquid companies that can't seem to get public, so there is no money to return to investors. Unless investors get money back, they can't invest in more companies. Liquidity allows flexibility and creates wealth.

Instead of spreadsheets, do you look for a story, one that anticipates all the twists and turns of a creative endeavor?

A Of course. What we look for in an investment is a creative, enthusiastic chief executive officer, a motivated team, and a vision to take a unique technology to a very large global market.

Why are you a proponent of global free trade?

A If the US government forces its businesses to use any workers who are not the best for the job, it will make the United States uncompetitive globally, which will make the entire country poorer and have the effect of making the United States lose more jobs and so on until there is no business left here.

Tell us about the stock market you are attempting to create.

A Expensive regulation now costs companies on the order of $3 million a year and has made it untenable for a business that earns less than $10 million a year in profit to go public. Xchange is a new private market that allows companies to "go private" and be traded before they are big enough to go public.

Do the business plans of all your start-ups have an IPO as an exit strategy, or is a buyout perfectly acceptable?

A I am not as fond of buyouts because they limit the upside of a company's potential, and they normally lose jobs. Also, since we are always looking for companies that will define and create industries, an acquisition can keep a new industry from forming. The IPO was a great alternative that allowed a company's shareholders to trade shares without losing the company's focus or general direction. Now, however, IPOs are too expensive for most companies that would like to get liquidity for shareholders, so we started Xchange to allow companies some liquidity for shareholders without spending all the money required to comply with expensive regulations, such as Sarbanes-Oxley.

What's the most promising company you are funding?

A The next one.

—interview conducted and edited by Paul Rako

BY BONNIE BAKER

Understanding CMR and instrumentation amplifiers

The three-op-amp instrumentation amplifier in **Figure 1** is seemingly a simple configuration in that it uses a basic, decades-old operational amplifier to gain a differential input signal. The op amp's input offset-voltage error is easy to understand. The definition of an op amp's open-loop gain has not changed. The simple idea of an op amp's CMR (common-mode rejection) has been around since the beginning of op-amp time. So what is the hang-up?

Equation 1 yields the common CMR for a single op amp and instrumentation amplifier:

$$CMR = 20\log\frac{G \times \Delta V_{CM}}{\Delta V_{OUT}}, \quad (1)$$

where G is the system gain, ΔV_{CM} is the changing common-mode voltage that you apply equally to the system's input terminals with respect to ground, and ΔV_{OUT} is the change in the system's output voltage with respect to the changing V_{CM} values.

With CMR, the inner workings of the op amp are straightforward; the change of offset voltage is the only concern. Two factors influence an instrumentation amplifier's CMR. The first and most dominant factor is the balance of the resistor ratios across A_3. For instance, if R_1 equals R_3 and R_2 equals R_4, the CMR of the three-op-amp instrumentation amplifier is ideally infinite.

At a real-world level, however, the relationship of R_1, R_2, R_3, and R_4 to the instrumentation amplifier's CMR—specifically, matching the R_1-to-R_2 ratio to the R_3-to-R_4 ratio—is critical. These four resistors combine with A_3 to subtract and gain the signals from the outputs of A_1 and A_2. A mismatch between the resistor ratios creates an error at the output of A_3. **Equation 2** gives the contribution to the CMR error with respect to the relationship of these resistors:

$$CMR_{A3} = 20\log\left[\frac{100 \times (1 + R_2/R_1)}{\%ERROR}\right]. \quad (2)$$

For instance, if R_1, R_2, R_3, and R_4 are approximately the same value and the ratio of R_3 to R_4 is 1.001 of R_1/R_2, this 0.1% mismatch will cause a degradation of the instrumentation amplifier's CMR from ideal to a 66-dB level. At a gain of one, CMR_{A3} is equivalent to the CMR of the entire instrumentation amplifier.

As **Equation 1** states, the instrumentation amplifier's CMR increases as the system's gain increases—a nice feature. **Equation 1** might motivate an instrumentation-amplifier designer to ensure that there is plenty of gain available, but A_1 and A_2's open-loop gain error places a limit on this strategy. An amplifier's open-loop gain is $20\log(\Delta V_{OUT}/\Delta V_{OS})$, where V_{OS} is the offset voltage. As the gain of A_1 and A_2 increases, the offset errors from the amplifier's open-loop gain also increase. The changes in output swing of A_1 and A_2 typically span the supply rails. At higher instrumentation-amplifier gains, the open-loop gain error of the op amps dominates. These errors degrade the CMR of the instrumentation amplifier at higher gains. Consequently, the instrumentation amplifier's CMR performance values tend to reach a maximum value at higher gains.

So, from the CMR perspective, instrumentation amplifiers are systems in which various parts contribute to the CMR error at different system gains. This situation is not so mysterious when you think about the inside of this device. As you separate the parts, the picture becomes clear. EDN

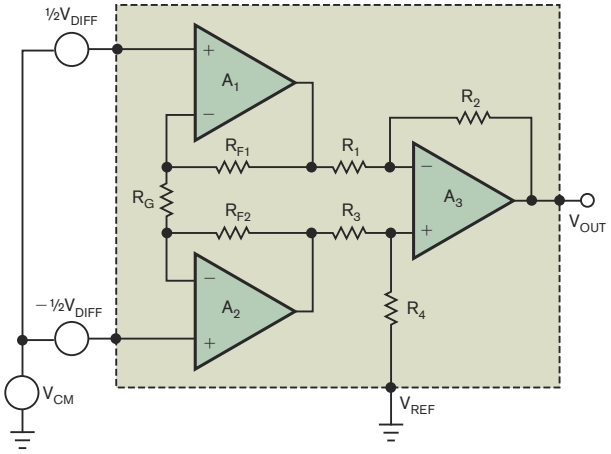

$$GAIN = \left[1 + \frac{(R_{F1} + R_{F2})}{R_G}\right]\left[\frac{R_2}{R_1}\right]$$

Figure 1 In this three-op-amp instrumentation amplifier, V_{CM} is the common-mode voltage, and V_{DIFF} is the differential input to the same instrumentation amplifier.

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.*

When it comes to creating your next embedded system, an important decision awaits you. Memory. Your design needs speed, reliability, performance and capacity to store the code and data your design demands. No problem. Numonyx has the broadest portfolio of parallel and serial NOR, NAND and phase change memory. And we offer extended temperature support with AEC-Q100 certification on many of our products and expanded design versatility with voltage support up to 5V. All designed to deliver a right-fit solution to help you shorten design cycles, reduce development costs and accelerate the roll-out of your next big idea. Find out how Numonyx memory matters for your next design.

NUMONYX® FORTÉ™ SERIAL FLASH MEMORY

Product Family	Density Range	Voltage/Solution
M25P (block erase)	512 k - 128 Mb	3V, single-I/O
M25PX (4KB block erase)	4 Mb - 64 Mb	3V, multi-I/O
M25PE/M45PE (page erase)	1 Mb - 16 Mb	3V

NUMONYX® AXCELL™ PARALLEL NOR FLASH MEMORY

Product Family	Density Range	Voltage/Solution
M29W/EW (JEDEC command set)	4 Mb - 2 Gb*	3V, page read†
P30/33 (Intel-based command set, sync burst)‡	64 Mb - 2 Gb*	1.8/3V core

Visit www.Numonyx.com/Embedded
for free access to the Numonyx Embedded Design Center.

numonyx™
innovative. memory. solutions.

*Includes stacked solutions †Available on specific densities ‡Other names and brands may be claimed as the property of others

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

Visionaire Lighting's Solar Vision Pole: shedding light on off-grid lampposts

Despite the high light efficacy of HB LEDs (high-brightness light-emitting diodes), their cost for commodity applications is still too high for them to compete head-on with older forms of lighting, such as incandescent and HID (high-intensity-discharge) lights. However, certain applications can justify paying a premium for high efficiency, long life, ruggedness, and light-color-temperature control, and these applications represent the sweet spot for HB LEDs.

The Aria-model light fixture has 48 Philips Lumiled HB LEDs. Visionaire chose these devices because they provide 100 lumens/W over a wide color-temperature range. Some HB LEDs can provide 100 lumens/W but only at a blue shade of white, typically a blue-white 6500K. Blue-white-colored lights can contribute to night-sky light pollution, which is the bane of observatories and dark-sky protectors.

One such application is solar-powered outdoor lighting for off-grid applications. Visionaire Lighting's Solar Vision Pole lamppost is especially novel because it does not use a standard rigid solar panel that requires additional bracing for wind shear and can attract the attention of scavenging thieves. Instead, a flexible solar panel encases the post and charges four gel batteries in its base. The size of the panel and the number of batteries limit the lighting to 50W, which is a weak traditional light source but makes for a strong, white-LED light. Six hours of charging is enough to run the light all night.

The amorphous-silicon flexible Solar Flex cells produce uniform power even as the sun's rays hit the round column of cells at an angle, easing the power-management task for the post. For lighting applications requiring 100 to 125W of power, the post is available in a version with a flat-mount polycrystalline panel that is both larger and more efficient than the other version.

The four gel-type battery packs in the lamp base can provide as much as 50W to the LEDs. The 12V-dc batteries each offer 30.5 Ahr. A full charge supports 40 hours of continuous illumination. The light operates at ambient temperatures as low as −76°F.

Agilent

Tektronix

LeCroy

Rohde & Schwarz

National Instruments

Anritsu

Keithley

Yokogawa

Tabor

Pickering

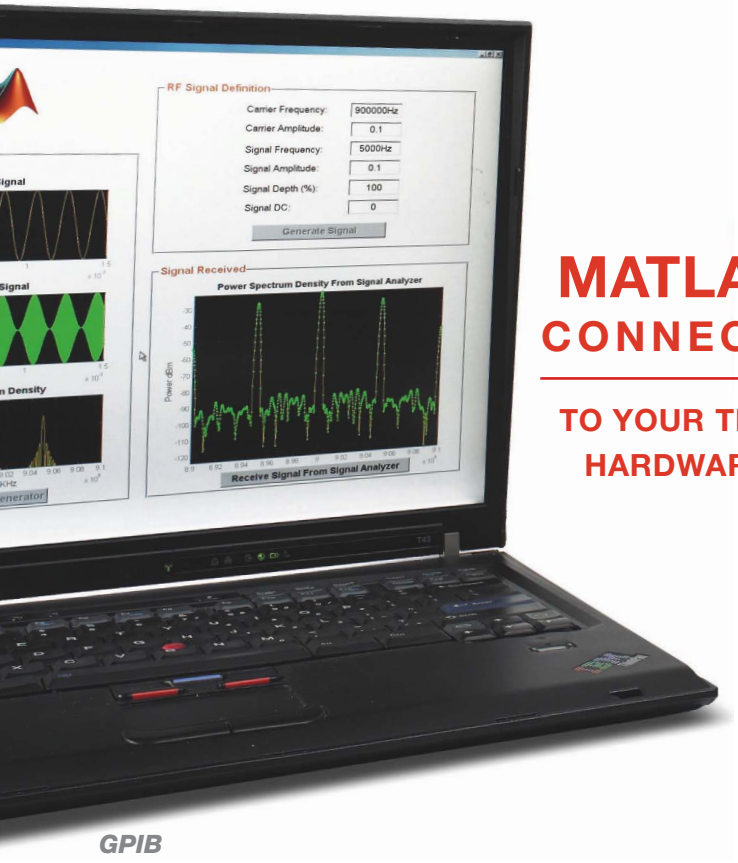

MATLAB CONNECTS

TO YOUR TEST HARDWARE

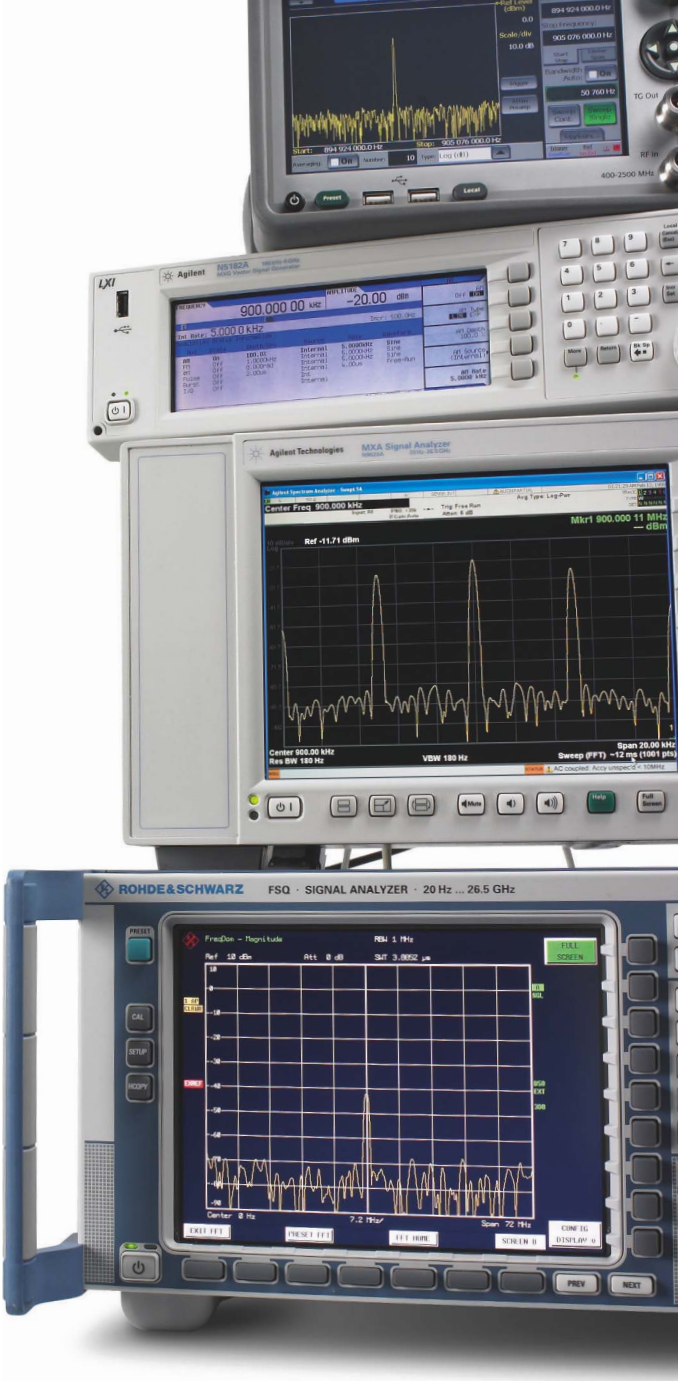

Connect to your test equipment directly from MATLAB® using standard communication protocols and hundreds of available instrument drivers.

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

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

The MathWorks™

GPIB

LXI

IVI

TCP/IP

VISA

USB

UDP

RS-232

CHIP DESIGNERS' STRUGGLES TO PROVIDE TRIPLE-PLAY HD SERVICE TO TELEPHONE, CABLE, AND WIRELESS CUSTOMERS ARE CHANGING THE NATURE OF SOC ARCHITECTURE.

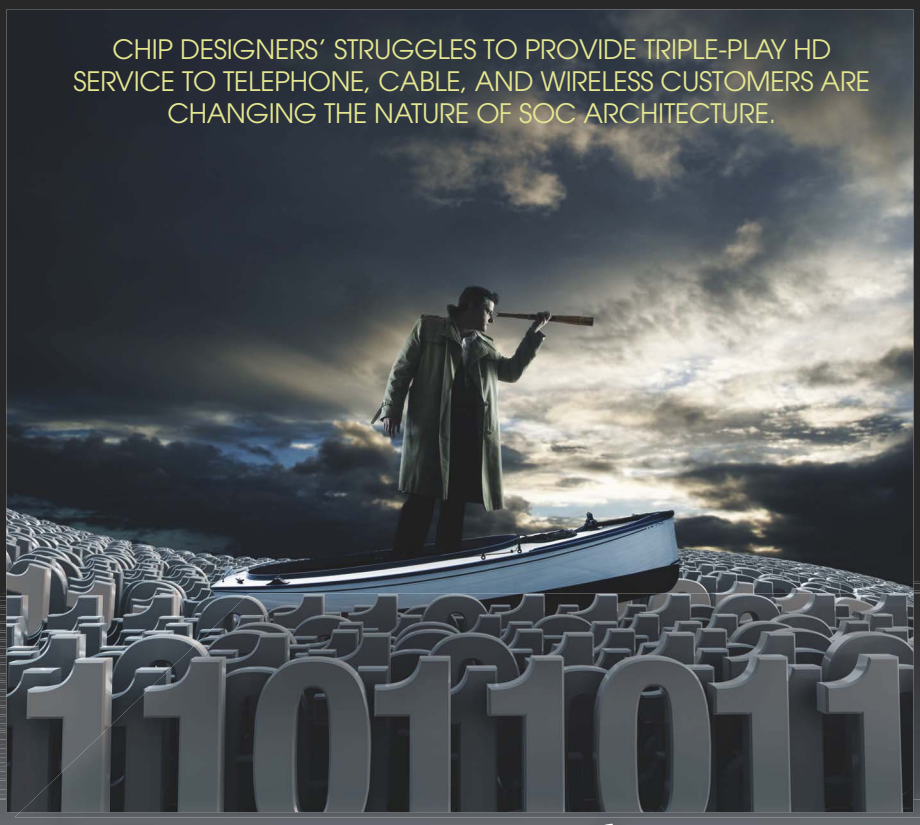

Lessons from THE LAST MILE

BY RON WILSON • EXECUTIVE EDITOR

The forces converging on the telecom and networking businesses have their roots in the changing desires of end users, and changing traffic patterns reflect those desires. For home-computer users, the mostly one-way HTML (hypertext markup language) traffic of Web browsing is gradually evolving into a rich mix of HTML, compressed HD (high-definition) video, interactive high-resolution graphics, and latency-intolerant HD audio. The heavily asymmetric traffic of Web browsing is becoming the more symmetric traffic of peer-to-peer networking.

Nowhere are these changes happening faster or with more public results than in the cellular-access networks, which are struggling to support new smartphones, such as the iPhone.

Mike Coward, chief technology officer at Continuous Computing, points out that mobile-broadband data traffic is doubling every nine months. "The iPhone does 30 times the traffic of a conventional handset," Coward says. "But that's not the bad news. Netbook users appear to create 450 times the traffic of handsets. All the operators are running up against spectrum limitations."

The iPhone is not the end of the story, either. Handset designers are pressing ahead with plans for mobile devices that can display and capture HD video. "Even with LTE [long-term evolution], there's not enough air bandwidth to give everyone HD video in their palm," Coward says. And a movie viewer in every palm is not the worst-case scenario. "Peer-to-peer traffic from netbooks and video sharing can be network breakers," he warns.

Mobile services must live within the physics of their air interfaces and thus face the most acute problem. Even cable- and telephone-service providers are under pressure, however. "While US-based broadband customers game, e-mail, and social-network over 384-kbps or 3-Mbps links, our counterparts in Korea, China, or Japan are real-time gaming and sharing video on 40- to 100-Mbps links," says Bruce Tolley, vice president of corporate marketing at Solarflare Communications. "A common deployment in Japan and China is IEEE 802.3ah

PON [passive-optical-network] fiber to the building, with 100-Mbps VDSL [very-high-speed-digital-subscriber-line] tails into the houses. This [bandwidth] is more than many of us have available in our corporate networks here in the United States."

So US cable and telephone operators are scrambling to upgrade, driving optical fiber as close as possible to the customer premises and then bridging the so-called last mile with cable or twisted pairs. "Cable operators will be strong in triple-play [voice, data, and video] in the United States," says Greg Fisher, vice president and general manager of Broadcom's carrier-access business. "With new VDSL technology, the telephone-company operators should be able to provide 50 to 100 Mbps on their copper for short distances." That ability is a big deal for the carriers. Verizon believes it can charge more than $100 per user per month for that kind of service. So nearly everyone is in the same boat. Sooner or later, bandwidth limitations will keep customers from using the network as they wish.

AT A GLANCE

- Triple-play use models are threatening today's networks.
- Carriers are rushing to increase speed and to shape traffic.
- Traffic shaping and security require fast deep packet inspection.
- A new generation of silicon architectures is rising to the challenge.

Security is yet another issue lurking beneath the surface of the shift in network use. Network-application providers, ISPs (Internet-service providers), and carriers all must protect themselves from denial-of-service attacks and intrusion. And carriers must protect their subscribers. "People are not talking enough about mobile security," Coward warns. "The first time there is a big intrusion into smartphones, users are going to blame their carriers." The same argument could apply just as well to fixed-service providers.

A NOSIER, SMARTER NETWORK

According to networking experts, the solutions to both of these problems—

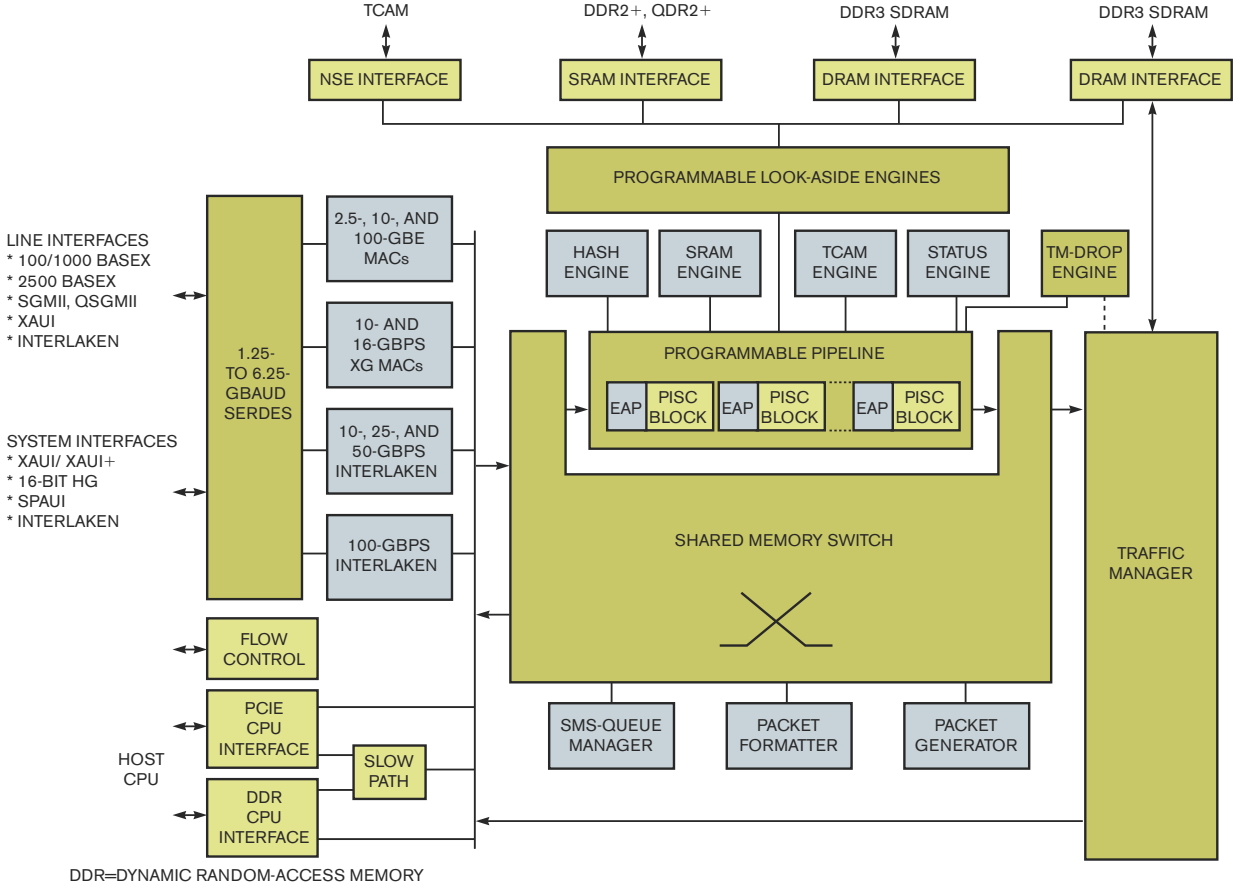

DDR=DYNAMIC RANDOM-ACCESS MEMORY
EAP=ENGINE-ACCESS POINT
HG=HOME GATEWAY
MAC=MEDIA-ACCESS CONTROL
NSE=NETWORK-SEARCH ENGINE
PCIE=PERIPHERAL COMPONENT INTERCONNECT EXPRESS
PISC=PACKET-INSTRUCTION-SET COMPUTER
QSGMII=QUAD SERIAL GIGABIT MEDIA-INDEPENDENT INTERFACE
SERDES= SERIALIZER/DESERIALIZER
SGMII=SERIAL GIGABIT MEDIA-INDEPENDENT INTERFACE
SMS=SHARED MEMORY SWITCH
SPAUI=SERIAL-PERIPHERAL-ATTACHMENT-UNIT INTERFACE
SRAM=STATIC RANDOM-ACCESS MEMORY
TCAM=TERNARY CONTENT-ADDRESSABLE MEMORY
TM=TRAFFIC MANAGEMENT
XAUI=10-GBPS ATTACHMENT-UNIT INTERFACE
XG=X GIGABIT

Figure 1 Xelerated's HX330 is an evolution of the programmable-pipeline strain of architectural thinking.

running out of bandwidth and security—begin in the same place: with knowing what is in the packets traversing the network. To make the most of what bandwidth they have, carriers must shape the traffic that passes through their domains. And to protect themselves and their customers, carriers and service providers must identify and destroy pernicious packets. Both of these processes require inspection of the packets as they pass through switches, routers, and even, some argue, line cards. But where to perform this inspection, how deeply to look into the packet, and what to do with the resulting information are all debated issues, the resolutions to which are greatly influencing silicon design.

"Bandwidth gets very expensive in access networks," says Kent Fisher, chief systems engineer at Freescale Semiconductor. "So there is a lot of incentive for carriers to parse the packet stream, identify the applications that are using the packets, and apply protocols and traffic shaping to get the most out of their bandwidth."

DPI (deep packet inspection)—looking deep enough into a packet to identify its payload—has many attractions. DPI allows a switch or router to prioritize and schedule individual packets—for example, giving latency-intolerant audio packets an immediate departure, making sure that the packet stream for an HD-video player gets its required minimum bandwidth, and scheduling HTML packets for a browser before data packets for a file swap. And detecting virus-bearing or denial-of-service traffic can also require looking at the payload. More controversial is the revenue aspect of the question. DPI allows a carrier to identify and charge extra for packets associated with premium services or to impede packets from rival services.

Other issues surround DPI, as well. If packets are encrypted and you can get the key, then inspection requires decrypting and re-encrypting each packet. "About half the time, there is no way to read encrypted traffic, so you have to rely on statistical techniques to guess what the packets are," says Continuous Computing's Coward.

And DPI is hard work. Instead of just breaking apart the header on each packet, you have to read the whole thing and, in the worst case, run it through a regular-expression processing algorithm to detect embedded patterns that can indicate data types or the presence of a virus. Particularly in software, that task takes a lot of cycles and a lot of energy. "With all of their requirements, mobile operators are asking us for 20 times more processing work per packet than in yesterday's systems," Coward says.

WHO DOES THE WORK?

Who will do all this work is another difficult issue. "Classification and QOS [quality-of-service] processing have to happen from end to end of the network, even in the metro networks," says Freescale's Fisher.

"You don't want to end up doing deep classification at really high bit rates," however, says Syed Shah, a systems architect at the company. "It's much more

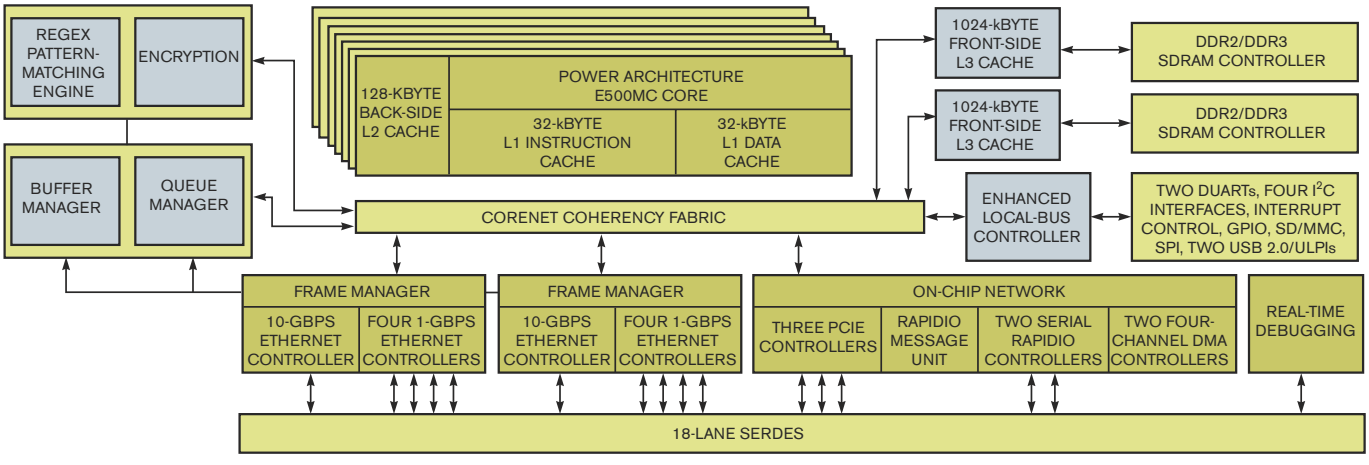

DDR=DOUBLE DATA RATE
DMA=DIRECT-MEMORY ACCESS
DUART-DUAL UNIVERSAL ASYNCHRONOUS RECEIVER/TRANSMITTER
GPIO=GENERAL-PURPOSE INPUT/OUTPUT
L1=LEVEL 1
L2=LEVEL 2
L3=LEVEL 3
MMC=MULTIMEDIA CARD
PCIE=PERIPHERAL COMPONENT INTERCONNECT EXPRESS
SD=SECURE DIGITAL
SDRAM=SYNCHRONOUS DYNAMIC RANDOM-ACCESS MEMORY
SERDES=SERIALIZER/DESERIALIZER
SPI=SERIAL-PERIPHERAL INTERFACE
ULPI=USB 2.0 TRANSCEIVER-MACROCELL INTERFACE/LOW-PIN INTERFACE
USB=UNIVERSAL SERIAL BUS

Figure 2 Freescale's 4080 family processors bear a family resemblance to other heterogeneous multicore-processor architectures.

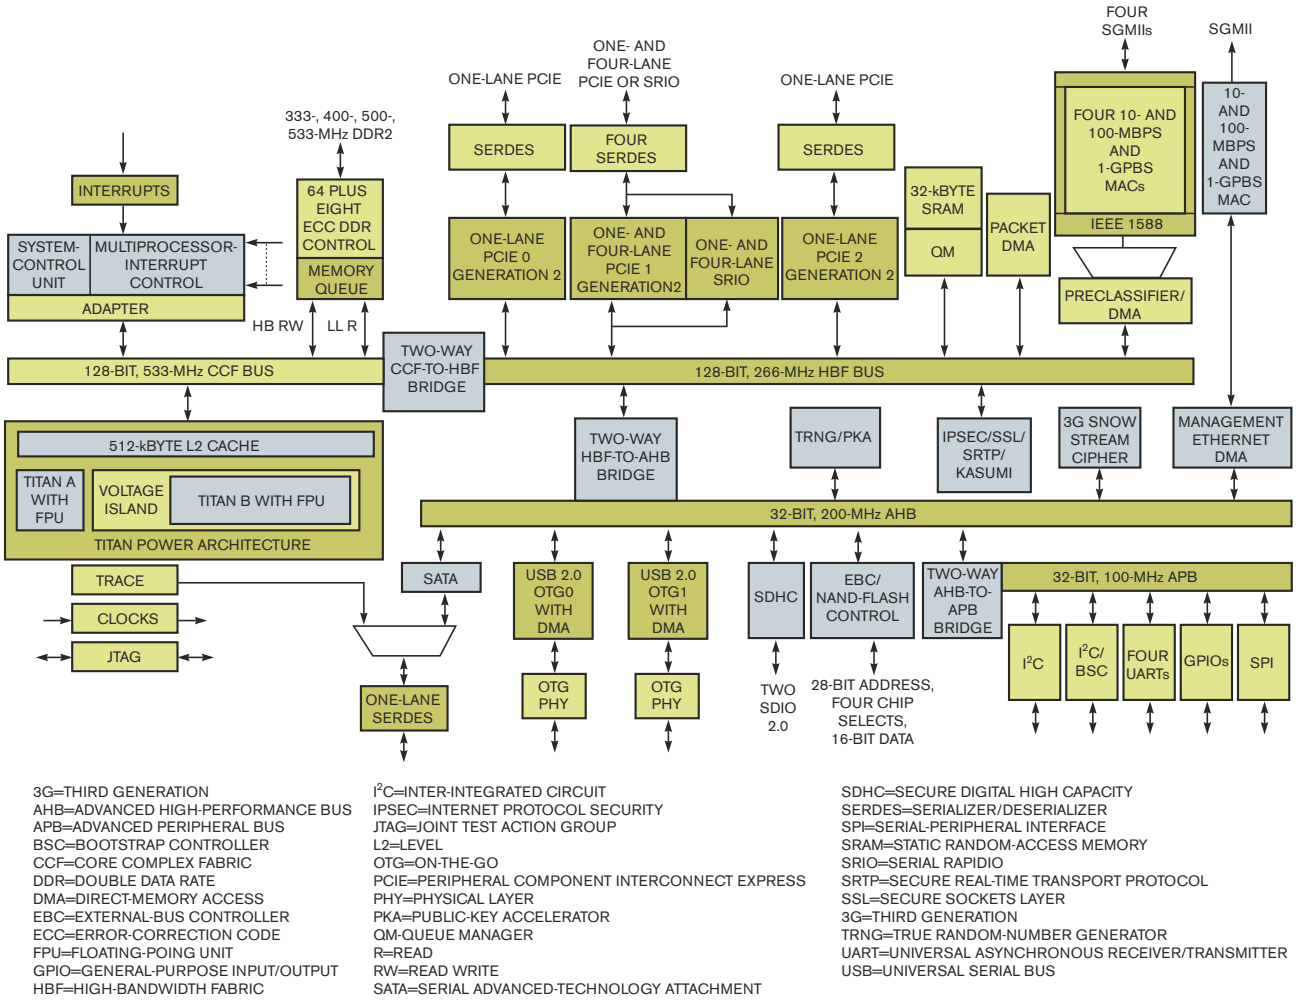

3G=THIRD GENERATION
AHB=ADVANCED HIGH-PERFORMANCE BUS
APB=ADVANCED PERIPHERAL BUS
BSC=BOOTSTRAP CONTROLLER
CCF=CORE COMPLEX FABRIC
DDR=DOUBLE DATA RATE
DMA=DIRECT-MEMORY ACCESS
EBC=EXTERNAL-BUS CONTROLLER
ECC=ERROR-CORRECTION CODE
FPU=FLOATING-POING UNIT
GPIO=GENERAL-PURPOSE INPUT/OUTPUT
HBF=HIGH-BANDWIDTH FABRIC
I²C=INTER-INTEGRATED CIRCUIT
IPSEC=INTERNET PROTOCOL SECURITY
JTAG=JOINT TEST ACTION GROUP
L2=LEVEL
OTG=ON-THE-GO
PCIE=PERIPHERAL COMPONENT INTERCONNECT EXPRESS
PHY=PHYSICAL LAYER
PKA=PUBLIC-KEY ACCELERATOR
QM-QUEUE MANAGER
R=READ
RW=READ WRITE
SATA=SERIAL ADVANCED-TECHNOLOGY ATTACHMENT
SDHC=SECURE DIGITAL HIGH CAPACITY
SERDES=SERIALIZER/DESERIALIZER
SPI=SERIAL-PERIPHERAL INTERFACE
SRAM=STATIC RANDOM-ACCESS MEMORY
SRIO=SERIAL RAPIDIO
SRTP=SECURE REAL-TIME TRANSPORT PROTOCOL
SSL=SECURE SOCKETS LAYER
3G=THIRD GENERATION
TRNG=TRUE RANDOM-NUMBER GENERATOR
UART=UNIVERSAL ASYNCHRONOUS RECEIVER/TRANSMITTER
USB=UNIVERSAL SERIAL BUS

Figure 3 Applied Micro's chip architecture resembles nothing so much as the architecture of the networks in which it will find use.

feasible to inspect the packets in the access network." Network architects recognized this situation years ago and came up with ideas such as MPLS (multiprotocol label switching) and the QOS bits in the IPv4 (Internet Protocol Version 4) header. In these schemes, a classification engine inspects each packet at or near its source and leaves a marker at Layer 2 or 3, indicating the priority the packet requires. Switches and routers deeper in the network then need not perform deep inspection.

"We see businesses trying to aggregate data and voice traffic from multiple ISPs and to route the packets using QOS bits or VLAN [virtual-local-area-network] tags," says Michael Durant, vice president of engineering at Arcturus Networks. This approach can in principle keep most of the switching decisions at Layer 2. "But browser vendors compete with each other on things like audio quality," Durrant continues. "So they routinely set the QOS bits very high. That practice creates artificially stringent QOS demands."

Even with all applications playing fairly, legitimate differences can exist in objectives between an application trying to impress a user, a base station trying to manage overloaded channels, a backhaul aggregator, and the metro network, for instance. So boxes deeper in the network may want to take a peek at suspect packets. "[Still,] I don't think you really need DPI either in the line card or in the metro network," says Thomas Eklund, vice president of marketing and business development at packet-processing chip vendor Xelerated. "Depending on regulations, inspection probably makes the most sense integrated into the access-network fabric. There, you must classify each packet through Layer 4. Beyond that [layer], I would argue it isn't really necessary."

You must also consider the government regulations that Eklund calls Layer 8. The Federal Communications Commission's sudden interest in network neutrality—the idea that the net, including carriers, should treat every packet the same—is of particular concern to equipment and silicon providers. Just what this doctrine means and how it might turn into regulation are areas of anxious debate. "Net neutrality appears to forbid DPI," Eklund observes. "But security, user demands for QOS, and the carriers' need to generate revenue may all require DPI." Such conflicts typically lead to politically driven instability in

regulations and, hence, create a need for great flexibility in switches and routers throughout the network.

ADDRESSING THE SILICON

From this statement of the problem, you can generalize about the kind of silicon that the next generation of access multiplexers, base stations, and carrier-Ethernet switches and routers will require. First, these chips will have to be fast. Wire speed for a VDSL2 twisted pair may be 100 Mbps. Deeper into the network, all transmission is optical, and a speed requirement of 10s of gigabits per second is not unusual. Switch and router boxes can't run below wire speed and depend on big buffers to make up the difference if carriers are succeeding in getting high channel usage because there would never be enough dead time in which to work through the buffer. And some new media types, notably audio conferencing and videoconferencing, are highly intolerant of the latencies big buffers would create.

The chip or chips must also be able to perform packet inspection. Just how deep that inspection must go is a matter of great uncertainty. As a generalization, however, the closer to the edge of the network a chip will sit, the more likely it is to have to do DPI. After inspection, the hardware will have to classify the packet and place it in the right queue for export. Further, the system will have to support a growing array of administrative, bookkeeping, supervisory, and error-recovery protocols.

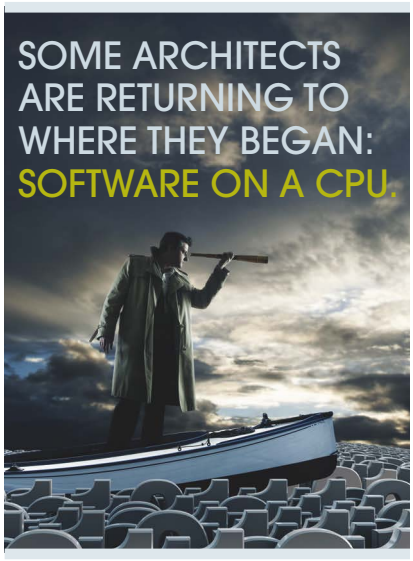

What does all this mean for the silicon? In simpler times, the hardware was just a fast CPU with a lot of memory—sometimes, just an embedded PC. All the functions were in the software. As speeds and functions both grew, however, their product outran Moore's Law. At that point, the hardware architecture split into two planes. Sequential, control-oriented tasks stayed in a CPU in the control plane. The much faster but readily parallelizable packet processing moved to more specialized hardware in the data plane.

Under growing pressure, the data plane evolved further. As data rates grew too high for CPUs to keep up, some architects developed network processors—essentially, microcontrollers with tightly coupled hardware accelerators to handle the bottleneck tasks. Other design teams went in a different direction: a hardware pipeline. Fixed-function hardware engines could keep up with very high wire speeds; if the sequence of

tasks in packet processing remained the same, simple data flows between pipeline stages eliminated many of the loads and stores inherent in a CPU-centric architecture, saving time and power.

But as protocols grew more diverse and complex, the fixed functions and fixed topologies broke down. Pipeline stages began to look like programmable accelerators. "Life is too risky now for fixed-function pipelines," says Xelerated's Eklund. "Programmability is not necessary only in the access fabric. It has to go much deeper into the network." Pipelines also sprouted thickets of conditional bypass and feedback paths and, eventually, accelerators of their own until the pipeline became just the central engine in a network of processing elements (**Figure 1**).

THE EVOLVING ENGINE

This growing complexity is erasing the distinction between the control and the data planes. At the same time, process migration is yielding less increase in circuit speed. As a result, some architects are returning to where they began: software on a CPU. This time, though, the CPU is a multicore cluster with both general-purpose processors and specific accelerators. Toby Foster, senior product marketing manager at Freescale, describes such a device (**Figure 2**). "The QorIQ chip family employs multiple e500 Power Architecture cores to cover applications from line cards to base stations and infrastructure," he says. "As the control and data planes merge, we see multicore chips with datapath accelerators—a queue manager, a crypto engine, a regular-expression matcher—encroaching on the traditional ASIC approach."

With all these cores, the traditional bus-based interconnect structure is failing, as well. To get the bandwidth the chip needs, architects provide each processing site, including the accelerators, with local caches, and they may tie everything together through a nonblocking switch fabric. If architects then provide hardware coherency across the caches and fabric, the programming model for the chip can approximate coding for a single CPU.

Even with good cache design, however, scheduling data movement under software control in such a chip involves a lot of work. "Traffic management in a multicore chip creates access issues," warns Satish Sathi, senior principal engineer at Applied Micro. "And these issues involve fairness, QOS, and conflicts for resources. You can resolve them in software, but that [approach] creates overhead."

Applied Micro's approach is hardware-based virtualization. In effect, Sathi explains, the control software sets up a route through the engines on the chip for each category of packets. A network of queues and a hardware-arbitration engine then steer the packets through the maze of engines, buses, and bridges (**Figure 3**). "The arbitration engine does dynamic arbitration based on actual end-to-end congestion on the chip," Sathi says. "Each packet gets inspected at the end of each task and routed to its next stop."

It's not a coincidence that this scenario sounds remarkably like a network—with nodes, routers, heterogeneous interconnect, and virtual channels. Increasingly, networking chip architectures are leaving behind the idea of a CPU core with accelerators on a bus and the concept of a CPU controlling a data-plane pipeline. Instead, the chips are becoming miniature models of the networks they will serve: heterogeneous collections of processing and routing sites, heterogeneous interconnect, virtual connections, and hardware-supported explicit routing of packet streams. The ideal we are approaching is the ability to define a virtual data-flow machine for each packet flow on an underlying fabric of programmable engines. Therein may lie the future not only of networking ICs but also of the SOC (system on chip) itself. EDN

FOR MORE INFORMATION

Applied Micro
www.appliedmicro.com

Arcturus Networks
www.arcturusnetworks.com

Broadcom
www.broadcom.com

Continuous Computing
www.ccpu.com

Freescale Semiconductor
www.freescale.com

Solarflare Communications
www.solarflare.com

Verizon
www.verizon.com

Xelerated
www.xelerated.com

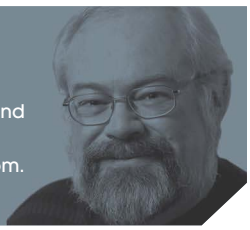

You can reach Executive Editor Ron Wilson at 1-510-744-1263 and ronald.wilson@reedbusiness.com.

FROM MAGNETIC TO SOLID STATE, SPIN-FREE:

What a long, strange storage trip it's turning out to be

BY BRIAN DIPERT • SENIOR TECHNICAL EDITOR

TO SERIOUSLY COMPETE WITH HARD-DRIVE MAKERS, SEMICONDUCTOR VENDORS MUST AMASS A ROBUST, SUSTAINED SUPPLY OF SILICON FOR SOLID-STATE DRIVES. THEY ALSO MUST ADDRESS PLENTY OF MISCONCEPTIONS ABOUT THE NEWER TECHNOLOGY'S CAPABILITIES AND LIMITATIONS.

Flash-memory-based solid-state drives have recently stirred up the staid storage industry, and their initial success stories foretell a potentially stellar future. Consider, for example, how rapidly they've taken over the formerly robust market for 1.8-in. hard-disk drives. Also consider their significant influence on smaller-form-factor hard-disk drives' lackluster initial unveilings. A notable percentage of netbook, tablet, and other alternative mobile computers, especially those running Linux operating-system variants, contain solid-state drives instead of hard-disk drives. Thin and light conventional notebook PCs running Windows and OS X are also well along the conversion path.

Enterprise-computing applications might at first glance seem to be poor candidates for solid-state technology, given its substantially higher cost than the magnetic alternative at the high capacities that this market segment requires. Yet, by virtue of its low energy consumption, increased reliability, and ultrafast read rates, the technology is making notable progress in conquering the corporate world. Consider that, to maximize hard drives' performance, IT departments have long formatted only the platters' fastest-access portion, wasting the rest of the drive and thereby blunting the argument that hard drives cost less than their nonmagnetic counterparts. Consider, too, that a number of applications, including smartphones, PDAs (personal digital assistants), digital still cameras, and videocameras, that might have formerly gone with—and, in some cases, in initial product generations, did go with—hard drives have migrated en masse to solid-state storage.

Proponents of both approaches dispute the extent of solid-state technology's potential to obsolete its predecessor, and the industry has yet to determine a winner. These debates illustrate a number of fundamental misrepresentations of solid-state drives' strengths and shortcomings. *EDN* readers' feedback to past editorial coverage reveals similar misunderstanding (**Reference 1**). This article attempts to clear up at least some of that confusion.

Cost, power consumption, and any other all-important comparative factors aside, the differences between hard-disk and solid-state drives boil down to a few fundamental points. First, solid-state drives, especially those in applications with random-access patterns, read data substantially faster than hard drives can, assuming the absence of any storage-to-system-interface bottlenecks. In contrast, solid-state drives, especially those in applications with random-access patterns, write data much slower than hard drives do. Also, once the solid-state drive has depleted its inventory of spare capacity, block-erase delays become a larger percentage of the total write latency. Further, unlike with a hard-disk drive, flash memory is not fully bit-alterable. Although flash memory can change ones to zeros on a bit-by-bit basis, converting even a single zero back into a one requires erasing the entire block containing the bit.

Another difference is that flash memory eventually "wears out" after extended erase cycles. However, it tends to do so on a block-by-block basis and in a predictable manner that the media controller can easily detect far in advance and compensate for in a variety of ways. Many hard-drive failures, in contrast, are abrupt and systemic. Further, because solid-state drives are semiconductor-based, they are notably more immune to the effects of abrupt shock, sustained vibration,

and other types of jarring. They're also comparatively impervious to environmental interference, such as from magnetic fields.

Most of today's flash memory uses a conceptually common floating-gate cell structure (**Figure 1**). Fast-random-read-access NOR and slower—albeit less expensive on a per-bit basis—NAND technologies differ predominantly in their cell-to-cell interconnect schemes. In its default erased state, the transistor turns on—that is, outputs a one—when the memory device's integrated address-decoding circuitry activates the necessary array row and column lines to select it. With the transistor programmed by means of additional electrons stored on its floating gate, it cannot turn on when address inputs select it and therefore outputs a zero to read attempts.

Altering the stored value of a flash-memory transistor involves the application of higher-than-normal voltages to various transistor junctions, thereby creating the necessary electric fields to affect electron flow onto or off the floating gate. Initial flash-memory generations required off-chip generation of these voltages; nowadays, most devices employ on-die high-voltage pumps for this function. Theoretically, you could alter a transistor's value—both for one-to-zero programming and for zero-to-one erasure—on a bit-by-bit basis, as is the case with EEPROMs (electrically erasable programmable read-only memories), FRAMs (ferroelectric random-access memories), MRAMs (magnetic RAMs), and battery-backed SRAMs (static RAMs). The necessary signal-routing, isolation, and other circuitry, however, would use too much die area and would therefore be too expensive at the IC capacities that bulk-storage applications require.

FROM AN EFFICIENCY STANDPOINT, SOLID-STATE DRIVES AND THE NAND CHIPS WITHIN THEM PREFER TO WRITE DATA IN APPROXIMATELY 4-KBYTE CHUNKS.

With modern NAND flash memory, on the other hand, you can erase blocks in approximately 512-kbyte increments. Bit-by-bit programming is possible; from an efficiency standpoint, however, solid-state drives and the NAND chips within them prefer to write data in approximately 4-kbyte chunks. This preference reflects the cost-versus-performance sizes of the RAM buffers on the flash-memory die. The bulk-alteration requirement differentiates flash memory not only from other nonvolatile semiconductor-

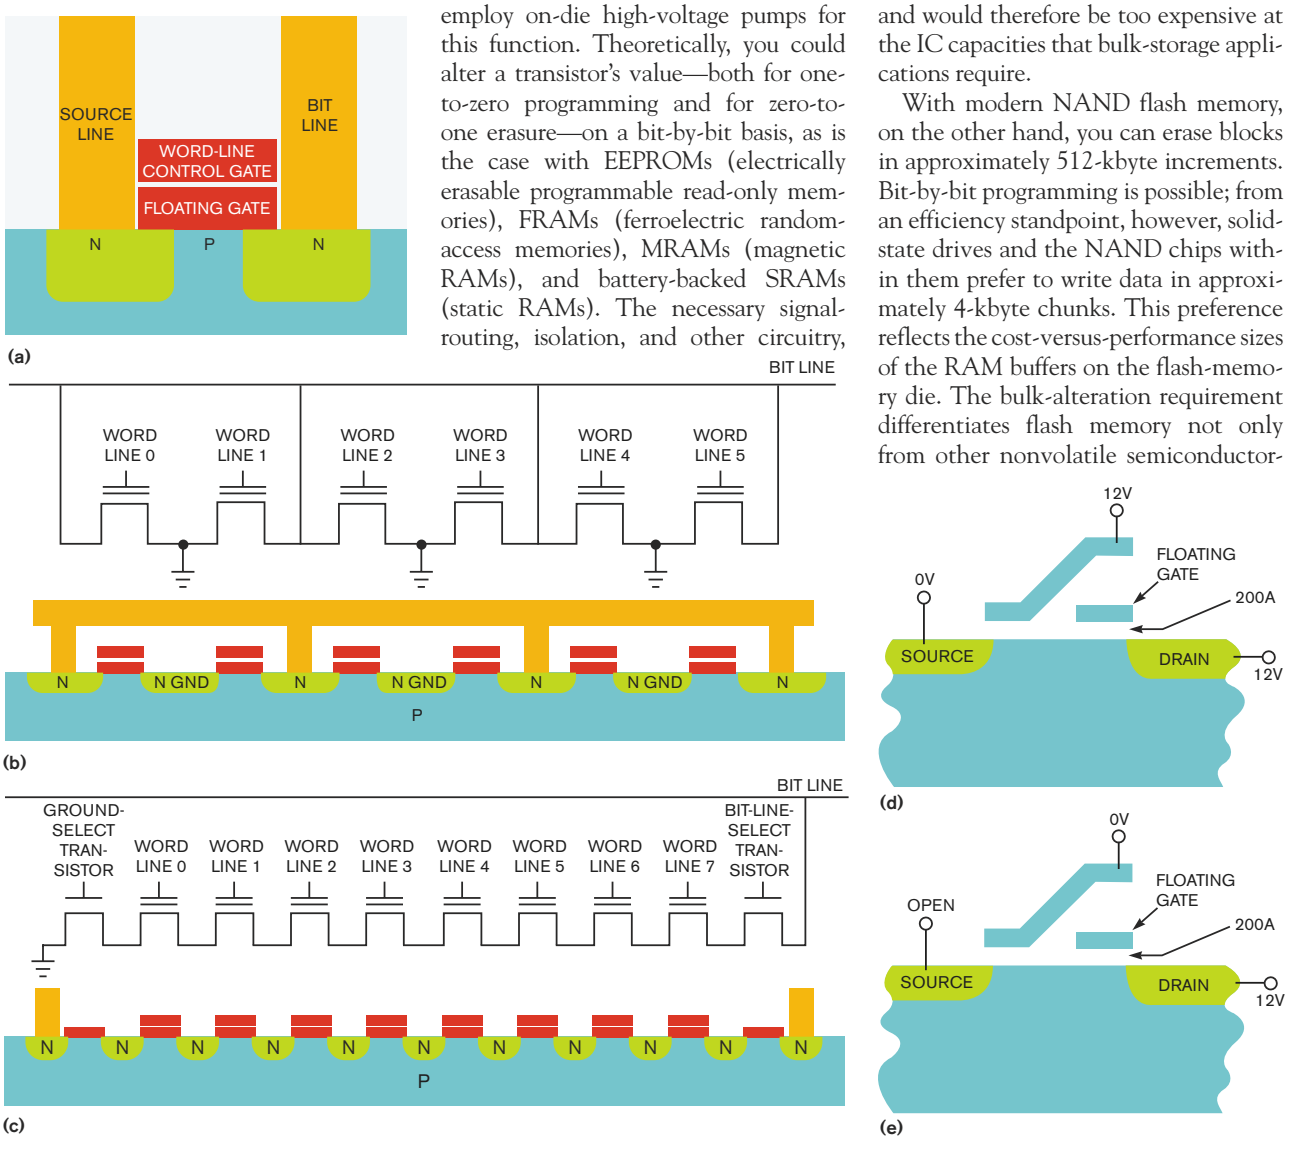

Figure 1 Most single-transistor flash-memory cells operate in a conceptually similar fashion (a) regardless of whether they interconnect in a NOR (b) or a NAND (c) scheme. Bit-by-bit or more efficient page-by-page programming places incremental charge on the floating gate (d), thereby counteracting an applied turn-on voltage during subsequent reads. Block-by-block erasure (e) removes this charge surplus (courtesy the Wikipedia Foundation).

storage technologies but also from hard-disk drives. The repeated electron flow across the thin silicon layer between the flash-memory transistor's substrate and floating gate and through incremental program and erase cycles stresses the oxide. At first, electrons inadvertently become trapped in the oxide lattice, impeding the flow of other electrons and slowing subsequent program and erase operations. Eventually, the oxide breaks down by rupturing, for example, leading to fundamental transistor failure. This demise tends to disrupt the function of the entire erase block that contains the affected transistor.

Modern flash memories come in both SLC (single-level-cell) and MLC (multilevel-cell) variants; today's MLC variants are primarily 2-bit-per-cell devices (**Figure 2**). With SLC flash memories, the voltage-sensing circuitry that connects to the array transistors' outputs can be relatively simple because it needs to discern only one voltage threshold and because the transistor's one and zero output voltages have substantial margin to this threshold. However, with a 2-bit-per-cell MLC flash memory, three voltage thresholds—that is, four levels—require discernment during reads. The programming operation for placing the necessary amount of electron charge onto the transistor's floating gate is similarly precise, and the effects of supply-voltage and operating-temperature variation and cycling further complicate this operation. Spansion claims that its MirrorBit MLC technology approach somewhat reduces the need for precise electron placement. Nonetheless, the concept remains largely relevant.

It's probably no surprise that MLC reads and writes—that is, programs—are substantially slower than their SLC counterparts and that the maximum block-cycling specifications for MLC memories are on average an order of magnitude less than those of SLC chips. These fundamental trade-offs are necessary for obtaining a lower per-bit cost for MLC storage devices (**Table 1**).

Now, consider Intel and Micron Technology's new 3-bit-per-cell memories and that Sandisk recently began producing 4-bit-per-cell X4 devices (**Reference 2**). The difference between any two sequential voltage levels and, hence, decoded-bit combinations with these new devices is on the order of 100 or so electrons or fewer in some cases. This situation represents a profound challenge for semiconductor-process and -product engineers. By potentially hampering both performance and data dependability, it calls into question the chips' suitability for applications requiring highly reliable storage. Then again, folks not too long ago were saying the same thing about 2-bit-per-cell MLC flash memory.

AT A GLANCE

- Solid-state storage offers many benefits over the currently dominant rotating magnetic alternative, but key differences require consideration and accommodation.
- Operating systems and their underlying file systems make assumptions about mass storage's long-latency random accesses and full bit alterability. Neither assumption is valid in the flash-memory era.
- Both operating systems' and system firmware's full support for a new ATA (advanced-technology-attachment) command promises to substantially simplify the flash-media controller's task, thereby boosting write performance.
- Operating systems can make key customizations to their functions once they know that they're talking to a solid-state drive instead of a hard-disk drive.
- Hardware evolutions, by discarding vestigial interfaces and extraneous functions, can notably optimize systems' solid-state drive implementations.

CONTROLLER CHOICES

The media controller may be a hardware-centric device, a software-fueled CPU, or any combination thereof. The hardware-versus-software choice of a controller involves trading off cost, performance, and power consumption versus flexibility and the ability to upgrade. Whatever its composition, the media controller acts as a bridge between the flash memories and the conventional hardware and software interfaces that the CPU, core-logic chip set, and other subsystems expect. The controller also manages the data stored in the single-component or multicomponent merged flash-memory array to avoid "hot-spot" overcycling of any erase blocks in the array, ideally as a background function that is invisible to the host both in access time and in any other regard. The controller leverages flash memory's strengths and mitigates its read- and write-speed weaknesses. The result is, with any luck, at least on par with—and, ideally, much faster than—the hard-drive alternative.

One perhaps obvious way of boosting effective solid-state-drive performance at the expense of incurring higher power consumption is to access multiple components in parallel using several address-, data-, and control-line channels

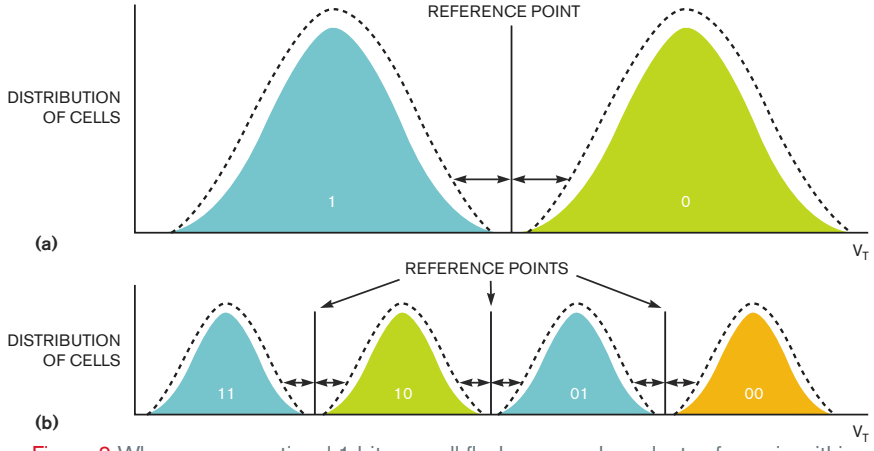

Figure 2 Whereas conventional 1-bit-per-cell flash memory has plenty of margin within the threshold-voltage envelope between a sensed zero and a sensed one (a), 2-, 3-, and 4-bit-per-cell technologies are more challenging to reliably implement across supply-voltage, temperature, and erase-cycling ranges (b) (courtesy Micron Technology).

TABLE 1 REPRESENTATIVE 1- AND 2-BIT-PER-CELL NAND-FLASH-MEMORY SPECIFICATIONS

Technology	Erase specification (cycles)	Program-page size (kbytes)	Erase-block size (kbytes)	Random-read latency (µsec)	Per-page program time (kbytes)	Per-block erase time (msec)
SLC (1-bit-per-cell) flash memory	100,000	4	512	25	250	2
MLC (2-bit-per-cell) flash memory	10,000	4	512	50	900	2

between the controller and the flash memories. You can then not only simultaneously read, program, or erase multiple array elements, but also juggle multiple operations with different ICs. For example, you could read from one while writing or erasing another if the system's access profiles justify this added level of controller complexity. In choosing a multichannel scheme, however, you also multiply the granularity of the solid-state drive's capacity and the effective sizes of program pages and erase blocks. This situation might warrant the choice of a flexible controller design that can run in either single-channel or multichannel mode.

Modern flash memories exhibit significant disparities between program-page and erase-block sizes and between program and erase times. The controller should, therefore, manage the media in such a way that background-erase operations for wear-leveling purposes—which manufacturers also commonly call housekeeping, garbage collection, and merging—on a component or a block within that component don't collide with system-write-request-initiated foreground programming operations on that same component or block or, for that matter, foreground system-read requests. Embedding a large RAM cache on the solid-state drive, much like the buffers on modern hard drives, can also be an effective collaborator to system-side buffering in mitigating any perceived decrease in performance that these housekeeping tasks incur. The trade-off of this approach, however, is that it requires more parts.

Reads and writes were traditionally the only required storage functions because, unlike with flash memory, you could fully overwrite hard-drive media on a bit-by-bit basis. A file-deletion request causes the file system to update its internal tables accordingly, but it historically didn't pass that information to the drive. Hence, the solid-state-drive controller is unaware that it can do background cleanup to free up the relevant pages and blocks containing them for future writes. The necessary erase and program operations occur only after the file system requests an explicit overwrite of the LBAs (logical-block addresses) associated with the drive's now-invalid PBAs (physical-block addresses). These operations are then unfortunately in the foreground where they adversely affect perceived read and write speed.

Manufacturers typically ship solid-state drives from the factory with spare "fresh" capacity, which is invisible to the operating system. The controller uses this capacity to delay the inevitable onset of the noted performance-strapping scenarios. However, good news is on the way in the form of the "trim" command, which the T13 Technical Committee of the INCITS (International Committee for Information Technology Standards) is now standardizing as part of the ATA (advanced-technology-attachment) command set. At press time, the T10 Technical Committee had not yet revealed its plans for the SCSI (small-computer-system-interface) command set. Before a system uses the trim command, it interrogates the drive to determine rotation speed. If it encounters a 0-rpm response, the system assumes that it is dealing with a solid-state disk and does further queries to determine whether trim support exists along with other relevant parameters. The trim command informs the drive that pages stored within the array are no longer valid and are therefore candidates for housekeeping. Deletion of a file within a trim-cognizant operating system results in the sending of relevant information for the corresponding LBAs to the drive's controller.

Although the trim command can dramatically improve sustained solid-state-drive performance in applications requiring many file deletions, it's ineffective in cases in which file updates

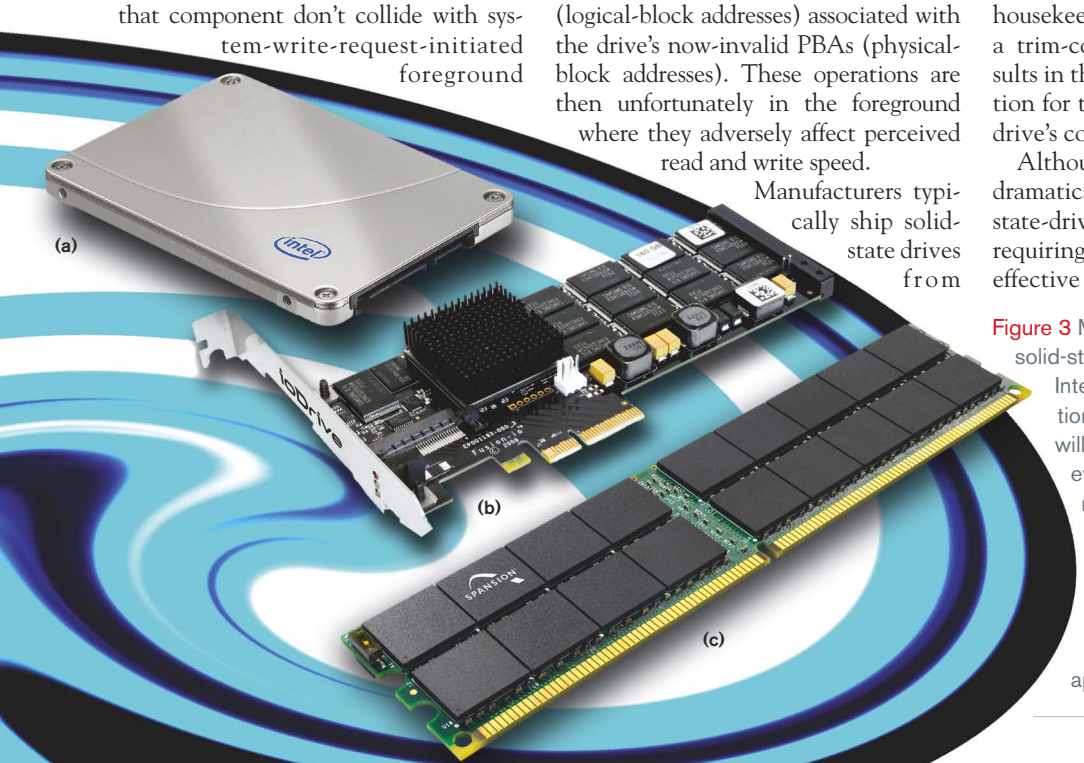

Figure 3 Most of today's solid-state drives, such as Intel's X25 units, use conventional storage interfaces and will benefit from those interfaces' evolutionary performance improvements (a). More revolutionary approaches migrate to alternative system interfaces with closer proximity to the CPU, such as Fusion-io's approach with PCIe (b) and Spansion's approach with DRAM (c).

occur, such as when you open a document for editing and then save the updated version or with Microsoft Outlook's PST database format. More generally, it exposes the weakness inherent in the strong linkage between LBAs and PBAs in FFSs (flash file systems). Sandisk's venerable FFS, along with the FTL (flash-translation-layer) technology the company obtained when it acquired M-Systems in mid-2006, strives to provide PBA independence for frequently updated files, such as the Windows Registry and the FAT (file-allocation table).

The company unveiled its ExtremeFFS at January's CES (Consumer Electronics Show) and both implements it in its products and makes it available for licensing. ExtremeFFS further severs explicit LBA-to-PBA linkage, thereby claiming to boost random write speeds by a factor as great as 100 times. ExtremeFFS makes less efficient use of the flash-memory media to accomplish this objective; Sandisk declines to provide specifics. Given the burgeoning capacities available with lithographies such as Intel and Micron's latest 34-nm process, however, the incremental ExtremeFFS overhead will over time become less of a practical issue.

A CAREFULLY CRAFTED SOFTWARE/HARDWARE IMPLEMENTATION CAN OPTIMALLY BENEFIT FROM FLASH MEMORY'S UNIQUE CAPABILITIES.

SYSTEM OPTIMIZATIONS

Solid-state units are currently shoehorning themselves into legacy hard-drive designs to leverage that technology's huge market, thereby jump-starting the solid-state-storage ramp-up (**Figure 3**). But as solid-state drives become more prevalent and presumably, therefore, a from-the-start implementation choice, a carefully crafted software/hardware implementation can optimally benefit from flash memory's unique capabilities. Consider, for example, common file-system operations, such as periodic automatic disk defragmentation, prefetching, file-location optimization, and system-side caching. These features all aim to compensate for hard drives' head-relocation and platter-rotation latencies, which cause slow random read accesses. Neither of these latencies is a factor with solid-state drives. Eliminating such workarounds can consequently improve system cost, power consumption, and other key variables and can reduce flash-media cycling.

Microsoft's latest Windows 7 operating system makes such adjustments when it detects a solid-state drive's presence in a system using the same scheme it uses for assessing trim support (**Reference 2**). Trim cognizance extends beyond simple file-deletion operations to encompass the full range of related functions, such as partition formats and system snapshots. Avoiding random-location writes whenever possible to boost performance can benefit both hard-disk- and solid-state-drive technologies. Integrated file-compression support for flash-memory-housed data can reduce

the per-bit cost gap between the technologies. Windows 7 and its peers also are more scrupulous about, for example, ensuring that partition- and file-location endpoints align with—rather than overlap—flash memory's write-page boundaries. And solid-state drives may provide an opportunity to reduce the system DRAM's budget requirement to less than it was in the hard-drive-centric past, thanks to solid-state's fast-access—at least for reads—virtual-memory-paging scheme.

System-hardware interfaces provide another opportunity for optimization. Except perhaps with extremely high-rotations-per-minute, enterprise-tailored units, hard drives tap the bandwidth capability of modern storage interfaces, such as 3-Gbps SATA (serial ATA) and SAS (serial attached SCSI), only when doing transfers to and from the drive's RAM buffer. Solid-state drives conversely can make more meaningful use of the performance potential of SATA and SAS, and the two technologies' performance gap will only increase in the upcoming 6-Gbps serial-storage-interface generation (**references 3** and **4**). Similar disparities are likely with the upcoming 4.8-Gbps USB (Universal Serial Bus) Version 3 and with Intel's embryonic Light Peak optical-interface technology.

But why restrict yourself to a legacy storage interface at all? Companies such as Fusion-io have figured out that PCIe (Peripheral Component Interconnect Express)-based add-in cards can boost performance by moving the solid-state drive closer to the CPU with which it's interacting. Giving the flash memory, either on a module or directly attached to the system board, a dedicated interface to the chip set affords an even closer linkage. Intel uses this approach with its Turbo memory cache, for example. The approach incurs a trade-off, however, in that it makes it more difficult for the end user to later alter the system-memory allocation. Alternatively, you can use the DRAM bus, as Intel's 28F016XD flash memory attempted to do in the mid-1990s and as Spansion's EcoRAM does today. In such a configuration, you might even be able to dispense with a dedicated flash-memory-controller chip, instead employing software running on the host CPU or circuitry within the chip set's logic. EDN

REFERENCES

1 Dipert, Brian, "Solid-state drives challenge hard disks," *EDN*, Nov 13, 2008, pg 25, www.edn.com/article/CA6611643.

2 "Engineering Windows 7," Microsoft Corp, 2009, http://blogs.msdn.com/e7/archive/2009/05/05/support-and-q-a-for-solid-state-drives-and.aspx.

3 Dipert, Brian, "Speedy simplicity: serial-storage interfaces," *EDN*, Jan 22, 2004, pg 33, www.edn.com/article/CA373882.

4 Dipert, Brian: "Interface overkill? Is eSATA necessary for your next system design?" *EDN*, May 10, 2007, pg 48, www.edn.com/article/CA6437950.

You can reach
Senior Technical Editor
Brian Dipert
at 1-916-760-0159,
bdipert@edn.com,
and www.bdipert.com.

Over 5000 Standard PICO Transformers and Inductors

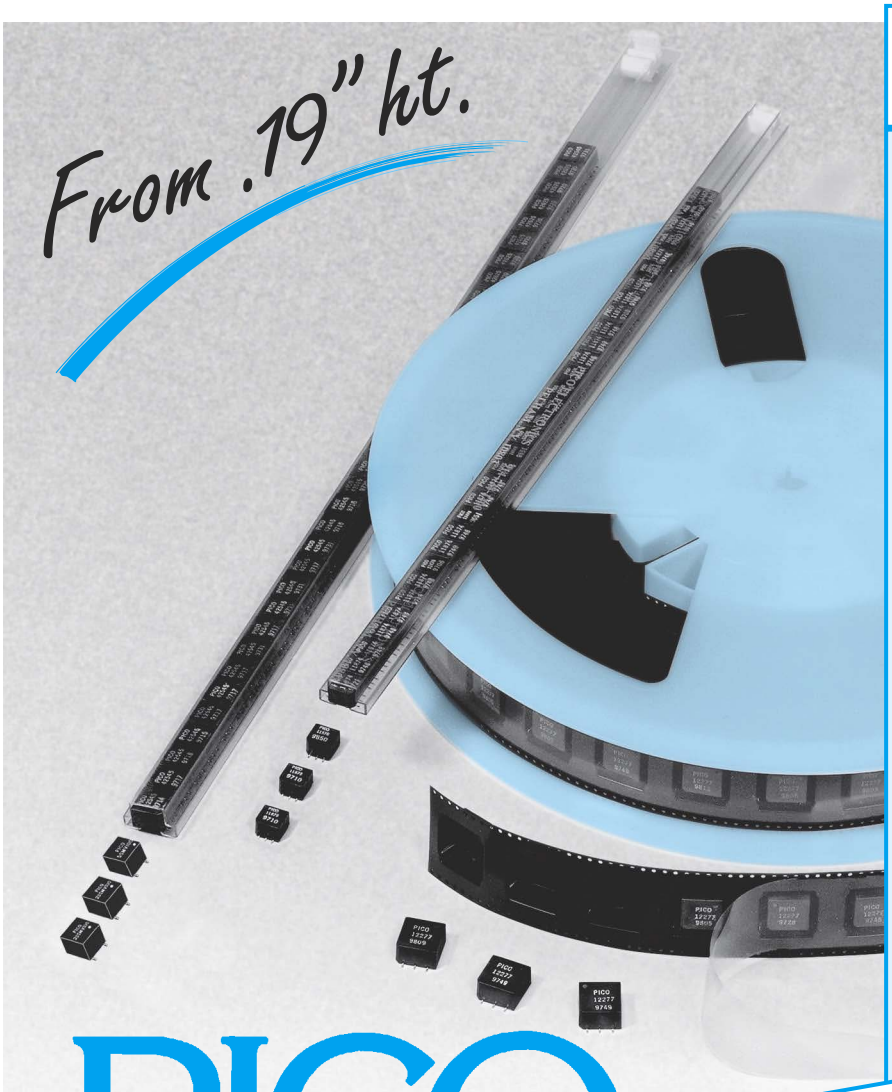

Surface Mount
(and Plug-In)

Audio Transformers
Impedance Levels 10 ohms to 250k ohms, Power Levels to 3 Watts, Frequency Response ±3db 20Hz to 250kHz. All units manufactured and tested to MIL-PRF-27. QPL Units available.

Power & EMI Inductors
Ideal for noise, spike and Power Filtering Applications in Power Supplies, DC-DC Converters and Switching Regulators

Pulse Transformers
10 Nanoseconds to 100 Microseconds. ET Rating to 150 Volt Microsecond, Manufactured and tested to MIL-PRF-21038.

Multiplex Data Bus Pulse Transformers
Plug-In units meet the requirements of QPL-MIL-PRF 21038/27
Surface units are electrical equivilents of QPL-MIL-PRF 21038/27.

DC-DC Converter Transformers
Input voltages of 5V, 12V, 24V And 48V. Standard Output Voltages to 300V (Special voltages can be supplied). Can be used as self saturating or linear switching applications. All units manufactured and tested to MIL-PRF-27.

400Hz/800Hz PowerTransformers
0.4 Watts to 150 Watts. Secondary Voltages 5V to 300V. Units manufactured to MIL-PRF-27 Grade 5, Class S (Class V, 155°C available).

DELIVERY-stock to one week for sample quantities

PICO Electronics, Inc.

143 Sparks Ave. Pelham, NY 10803 • ***Call Toll Free*** 800-431-1064 • fax: 914-738-8225

Send for free 180 pg PICO Catalog

See PICO's full Catalog Immediately on the internet

www.picoelectronics.com

E-mail: info@picoelectronics.com

INTRODUCING THE MOST POWERFUL WEBSITE IN THE WORLD

– literally.

At www.ttiinc.com/power, engineers will find power management trends, information and products – all in one place.

All the power management tools you spend time searching for online are now in one toolbox. Including trend reports by Dennis Zogbi (passives) and Ron Bishop (connectors), webcasts from top suppliers and more. Bookmark this website and discover a world of power.

tti®

NEVER SHORT ON SOLUTIONS

1-800-CALL-TTI

www.ttiinc.com/power

POWERED BY:

VISHAY®

KOA

muRata

Tyco Electronics

FCI

molex®

Littelfuse®

BY CHI T HONG • CALIFORNIA MICRO DEVICES

Evaluating ESD-protection components: Clamping voltage and dynamic resistance are crucial

A CHANGING PRODUCT LANDSCAPE AND NEW DESIGNS CALL FOR IMPROVED PROTECTION AGAINST ESD STRIKES ON COMPONENTS. A LOW-VOLTAGE DEVICE DOESN'T NECESSARILY HAVE GREATER PROTECTION. PROTECTION COMES FROM LOW CLAMPING VOLTAGE AND LOW DYNAMIC RESISTANCE.

Lighter, smaller consumer devices, such as laptops, cell phones, and iPods, use ASICs with geometries as small as 90 nm. At those geometries, even small levels of ESD (electrostatic-discharge)-induced voltage and current can cause catastrophic failure. Other potential sources of ESD strikes are end users who touch I/O connector pins while hot-plugging peripherals using USB (Universal Serial Bus) and HDMI (high-definition-multimedia-interface) connectors. The ESD you generate by walking across a carpet—potentially, a 1-nsec, 30A pulse—is enough to destroy an ASIC. Chip makers are also reducing the standard level of on-chip ESD protection. For these reasons, ESD-protection devices are critical to a design's success. How do you go about selecting the best ESD-protection device?

Conventional wisdom relies on using IEC (International Electrotechnical Commission) 61000-4-2, which the organization accepted as a standard in 1995. The IEC developed the IEC 61000-4-2 rating test to measure how well a chip could sustain an ESD attack in a finished product-application environment. The standard refers to a contact ESD of 8 kV or an air ESD of 15 kV. The ESD rating on a chip tells you only the amount of voltage that the protection device can survive, however. It does not guarantee that the DUP (device under protection) can survive because you must protect the DUP from residual current when the ESD device cannot shunt most of the peak pulse current to ground. The standard also can't tell you how much residual current reaches the component. You need to know the ESD device's clamping voltage and dynamic resistance. Therefore, although it would be simple to specify an ESD-protection device on the basis of the IEC 61000-4-2 rating alone, it's more difficult to ensure that you've protected your design.

During an ESD strike, the ESD-protection device should shunt most of the peak pulse current to ground. However, residual current flows through the DUP, damaging or destroying it (**Figure 1**). The power across the DUP results from high clamping voltage times high residual-current joule heating due to high dynamic resistance. This combination poses the greatest danger of ESD damage (**Figure 2**). The peak pulse current equals the shunted current through the ESD device plus the residual current. So the larger the shunt current, the smaller the residual current for a device's clamping voltage. Also, a device's shunt resistance equals the clamping voltage divided by the dynamic resistance. In other words, a device's shunt current and dynamic resistance are inversely proportional to each other. Therefore, an ESD device with a higher dynamic

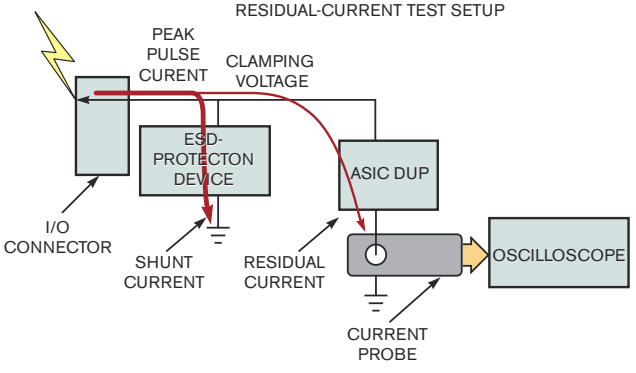

Figure 1 Residual current flows to the ASIC device under protection.

TABLE 1 DYNAMIC RESISTANCE FOR DETERMINING PROTECTION

Device 1	Device 2
8.8V clamping voltage at 1A	14V clamping voltage at 1A
0.7Ω dynamic resistance	3Ω dynamic resistance
Clamping voltage for 30A strike: 8.8V+29A×0.7Ω=29.1V	Clamping voltage for 30A strike: 14V+29A×3Ω=101V

resistance allows more residual current to flow to the DUP.

Dynamic resistance and clamping voltage determine how well an ESD-protection device protects against residual current. During routine operations, the protection device must maintain a high impedance. However, when an ESD strike hits, the protection device rapidly shunts ESD energy to ground. Data sheets may list dynamic resistance. If the data sheet you are looking at doesn't list this spec, however, you can calculate it from a graph of the device's clamping voltage versus peak pulse current for IEC 61000-4-2 Level 4, in which the peak pulse current is 30A. Clamping voltage matters more than operating voltage when comparing protection devices. During an ESD strike, an ESD device's lower clamping voltage minimizes ESD damage due to joule heating. So both a low clamping voltage and low dynamic resistance provide a more accurate metric of how an ESD device will work during a strike (**Figure 3**).

Because dynamic resistance is the yardstick by which you should measure protection devices, the question is how to calculate this parameter if the data sheet does not list it. You can easily calculate dynamic resistance as a slope on a graph of clamping voltage versus peak pulse current, with a peak pulse current of 0, 1, 2, and 3A and so on. After 1A, the slope of dynamic resistance is close to linear. You can also calculate this parameter from clamp-voltage numbers with corresponding peak pulse current because a device's dynamic resistance is the slope of the graph equal to one clamping voltage minus another clamping voltage divided by the peak pulse current minus another peak pulse current. The formula for determining clamping voltage at a peak pulse current of 30A for IEC 61000-4-2 is to add the breakdown voltage to the peak pulse current and multiply that result by the dynamic resistance.

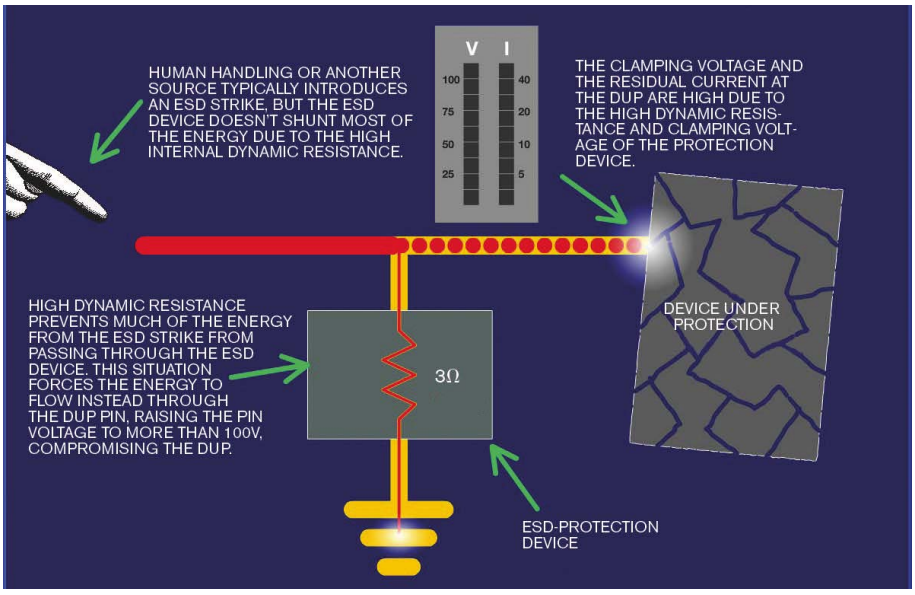

Figure 2 Catastrophic failure of the DUP may occur when the ESD device's dynamic resistance is high.

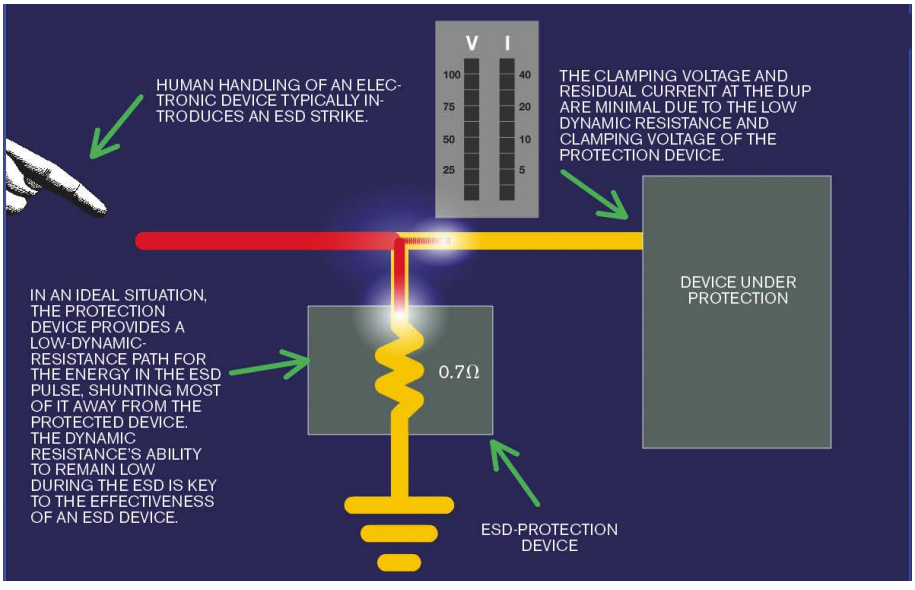

Figure 3 High dynamic resistance increases residual current to the DUP. Low dynamic resistance decreases residual current to the device under protection. A protection device offering low dynamic resistance makes a path for an ESD strike and shunts current away from the device under protection, thereby minimizing residual current to the DUP.

TABLE 2 ESD-DIODE-PROTECTION DEVICE VERSUS SUPPRESSOR/VARISTOR PERFORMANCE

Voltage (kV)	ESD diode	Device under protection	Peak clamp voltage (V)	Residual current (A)	Residual peak maximum/total (VA)
4	Pass	Pass	47.97	4.68	152.02
6	Pass	Pass	68.23	6.83	351.66
12	Pass	Pass	137.33	14.43	1787.4
Voltage (kV)	**Suppressor/varistor**	**Device under protection**	**Peak clamp voltage (V)**	**Residual current (A)**	**Residual peak maximum/total (VA)**
4	Pass	Pass	162.33	14.78	1405.65
6	Pass	Fail	181.77	16.24	2253.56

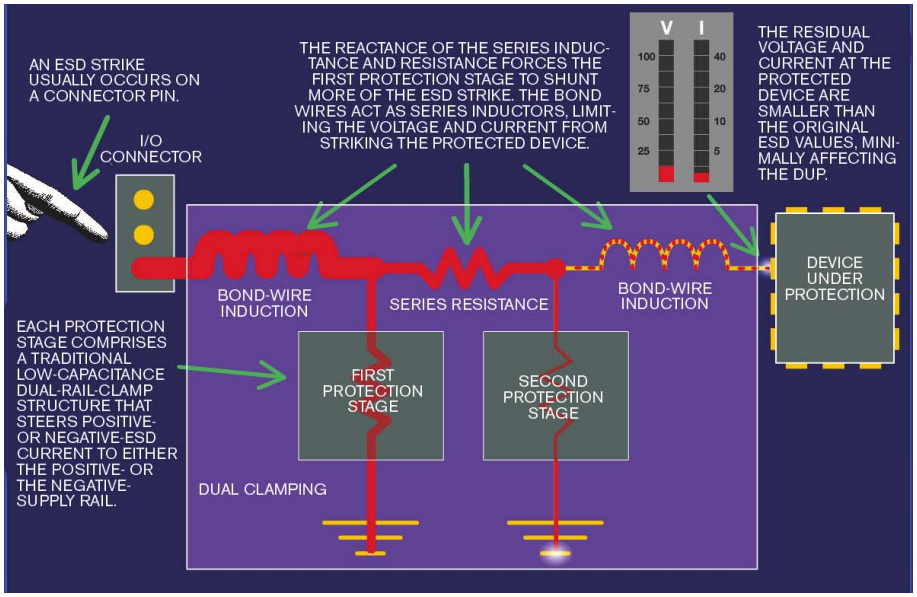

Figure 4 A two-stage ESD-protection architecture provides dual-stage clamping to limit both clamping voltage and residual current from reaching the DUP.

The residual current flowing through a "protected" chip is proportional to the dynamic resistance of the protection device versus the resistance in the rest of the circuit. Dynamic resistance is the most important factor in determining protection from ESD. All other things being equal, a 5V ESD diode is only marginally better than a 3.3V diode (**Table 1**). When comparing devices, bear in mind the dynamic resistance, which affects the residual current, rather than the breakdown voltage.

ESD diodes and suppressor/varistors have different performance. **Table 2** shows the specs of a system that survived a 12-kV strike using a diode-protection device but failed at 6 kV using a high-dynamic-resistance suppressor. Note that the varistor survived, but the system failed. Note also that both the clamping voltage and the residual current are much higher with varistors.

The availability of two-stage ESD-protection architectures means engineers need not choose between signal integrity and ESD protection. A two-stage ESD architecture offers more protection for a DUP when a single-stage ESD architecture is insufficient for providing lower dynamic resistance and therefore higher residual current. The first stage acts as a traditional ESD device, lowering clamping voltage and dynamic resistance, and the second stage further reduces clamping voltage and residual current (**Figure 4**).

The fact that a device has low clamping voltage doesn't mean that it offers higher protection. When creating your design, therefore, the most important parameters for comparison between ESD-protection devices are clamping voltage, dynamic range, and the total number of protection stages in the device. These parameters let you know before you specify an ESD device whether that device is right for your application. EDN

Go to www.edn.com/ms4335 and click on Feedback Loop to post a comment on this article.

AUTHOR'S BIOGRAPHY

Chi T Hong is a senior application engineer at California Micro Devices (Milpitas, CA). He received bachelor's and master's degrees from the University of Southern California (Los Angeles). For more information on two-stage ESD devices, go to the XtremeESD products page at www.cmd.com.

ACKNOWLEDGMENT

Thanks go to Allen Tung of California Micro Devices for his help with this article.

BY SAMEER KELKAR • POWER INTEGRATIONS

What every designer should know about magnetics in switch-mode power supplies

POWER IS OFTEN AN AFTERTHOUGHT IN SYSTEM DESIGN, BUT THE CHOICE AND DESIGN OF THE MAGNETIC ELEMENTS AT THE HEART OF AN SMPS ARE CRUCIAL. ACQUAINT OR REACQUAINT YOURSELF WITH THE FUNDAMENTALS OF THIS FREQUENTLY OVERLOOKED AREA.

The application of electromagnetics has been in practice for more than a century: In 1831, English chemist and physicist Michael Faraday invented the transformer, although he called it an induction coil. Unfortunately, engineering schools rarely provide instruction in practical magnetics relevant to SMPS (switch-mode-power-supply) applications. Part of the problem is that the classic design equations for magnetics target sinusoidal waveforms, but SMPSs operate with rectangular waveforms.

The starting point for understanding magnetics is to look at the relationships between current flow and electric and magnetic fields. **Figure 1** shows a simple air-cored winding. A current-carrying conductor creates its own magnetic field (B), which produces flux lines around the conductor. In this example, 10 turns of wire carry a dc current, and each turn creates its own magnetic field. The fields combine to create a concentrated and fairly linear field within the winding; the field diverges and weakens outside the winding. The magnetic field inside the winding is the primary storage area for energy, but the external field can also store a significant amount.

If you place an object comprising a magnetic material, such as iron, within the winding, the magnetic field exerts an EMF (electromotive force) on the object. If you then place a second winding within the field and the primary winding is carrying ac current so the field is changing with time, the magnetic field will induce a current to flow within the second winding. Lenz's Law, which Russian chemist and physicist Heinrich Lenz postulated in 1834, states that an induced current always flows in a direction opposing the motion or change causing it.

Thus, you can describe the properties of a magnetic field in terms of its intensity or its density. The magnetic-field intensity defines the field's ability—in ampere turns per meter—to exert forces on magnetic poles. The magnetic-flux density (B) is the ability of the magnetic field, in teslas, to induce an electric field when it changes. This property introduces the dimension of time.

Two laws—Ampere's and Faraday's—jointly govern the relationship between magnetic components and their characteristics you see from the terminals. Ampere's Law, which French physicist and mathematician André-Marie Ampère postulated in 1826, relates the integrated magnetic field around a closed loop to the electric current passing through the loop. Faraday's Law, which Faraday postulated in 1834, states that the induced EMF or EMF in any closed circuit equals the time rate of change of the magnetic flux through the circuit.

You may wonder why magnetic circuits require cores. Answering this question requires consideration of another characteristic, permeability—a measure of the amount of flux a magnetic field can push through a unit volume of material. You would not expect the winding in **Figure 1** to perform well as an electromagnet because it has no core. However, if you insert an iron core in the center of the windings, it can make a powerful electromagnet because the permeability of iron is about 10,000 times that of free space, enabling the concentration of a relatively large amount of magnetic flux between the windings. Permeability is roughly analogous to conductivity in the electrical realm. **Table 1** shows the equivalence between the magnetic and the electrical domains. Just as a conductor is a conduit for energy to flow in the form of an electrical current, a high-permeability magnetic material acts as a conduit for energy to flow as magnetic flux.

It is important to account for leakage in magnetic circuits. Many parallels exist between the electrical and the magnetic realms. However, compared with free space, the conductivity of common conductors, such as copper, at approximately 10^{20}, is much higher than the permeability of magnetic materials, at approximately 10^4. Thus, you can easily ignore leakage currents,

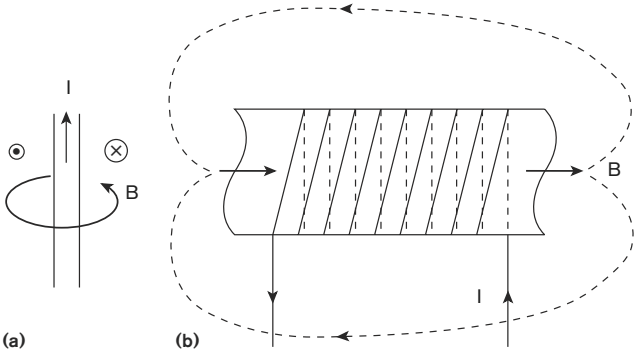

Figure 1 In a simple air-cored winding, a current-carrying conductor creates its own magnetic field, which produces flux lines around the conductor (a); 10 turns of wire carry a dc current, and each turn creates its own magnetic field (b).

TABLE 1 EQUIVALENCE BETWEEN MAGNETIC AND ELECTRICAL DOMAINS

Magnetic		Electrical	
Symbol	**Parameter**	**Symbol**	**Parameter**
ϕ	Flux	I	Current
H	Magnetic-field intensity	E	Electric-field intensity
B	Magnetic-flux density	J	Current density
MMF	Magnetomotive force	EMF	Electromotive force
μ	Magnetic permeability	σ	Electrical conductivity
R	Reluctance	R	Resistance
L	Inductance	C	Capacitance
LG	Air gap	D	Dielectric

but not leakage flux, in low-frequency systems. Although permeability is analogous to conductivity, it is not a linear characteristic for many materials and takes a different value depending on what previously occurred (**Figure 2**).

Figure 2a shows the relationship between magnetic-field intensity, H, and magnetic-flux density, B, in a ferromagnetic material. The slope of this curve at any given time is the instantaneous permeability (μ_R) of the material. For low values of intensity, permeability is constant and relatively high. However, for larger values of intensity, permeability starts decreasing to the point that the material starts resembling free space (μ_R=1). Thus, you need a larger and larger magnetic-field intensity to produce a small increase in the density field. At this point, the material reaches saturation.

The area between the rising BH curve and the vertical axis represents energy stored in the material. If you then reduce the field intensity from the saturation point, you can recover energy from the material; however, the energy you recover is less than that stored, so the BH curve follows a different path. The result for a complete cycle is that the BH curve forms a closed S shape. The area the S encloses represents the hysteresis curve, or the total lost energy in the cycle. The area of the curve is a function of the frequency; thus, at higher frequencies, the area of the hysteresis curve increases, and so do the losses.

The total flux, or flux linkage, relates to the electrical current through the inductance constant. Thus,

$$L_M = \frac{N\phi}{I},$$

where L_M is the inductance constant, $N\phi$ is the flux linkage, and I is the electrical current. Further,

$$\phi = A_E \times B;\ B = \mu_0 \times \mu_R \times H;\ H = \frac{N \times I}{l_M},$$

where A_E is the cross-sectional area of the core, B is the magnetic-flux density, μ_0 is the permeability of free space, μ_R is the relative permeability of the core material, H is the magnetic-field intensity, N is the number of turns, I is the current, and l_M is the magnetic path length. Therefore,

$$L_M = \frac{N^2}{\frac{l_M}{\mu_0 \times \mu_R \times A_E}} = \frac{N^2}{R}.$$

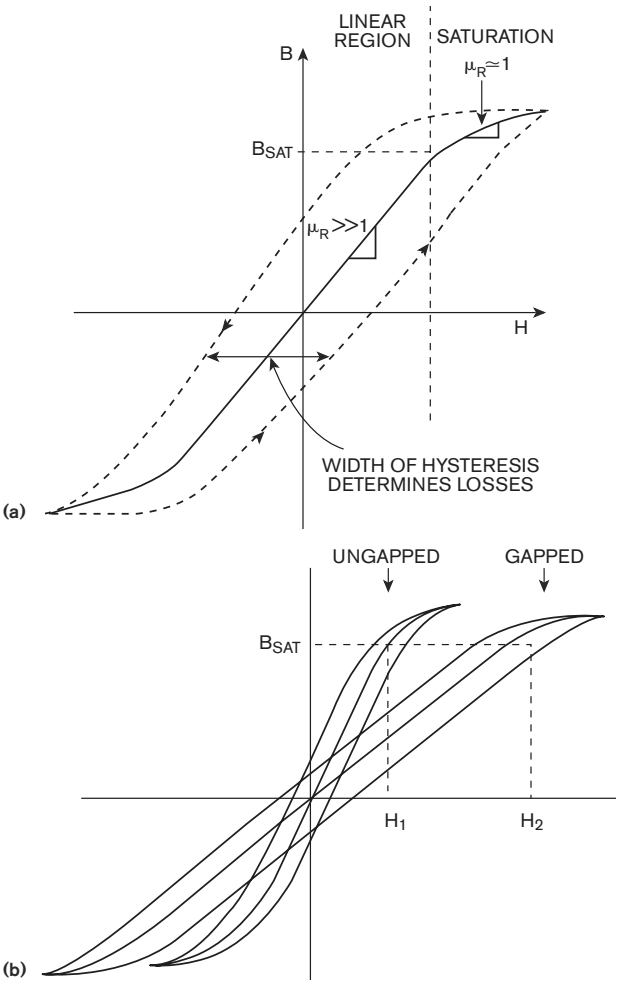

Figure 2 A BH curve shows the relationship between magnetic-field intensity, H, and magnetic-flux density, B, in a ferromagnetic material (a). The slope of this curve at any moment is the instantaneous permeability of the material. For low values of H, permeability is constant and relatively high. The introduction of an air gap tilts the BH curve to the right (b). The ungapped core would saturate at a field intensity of H_1, whereas a gapped core can be useful at a field intensity as high as H_2, where H_2 is greater than H_1. Because current, I, is the prime driver of magnetic-field intensity, you can push more current through the core without saturating it.

The parameter

$$\frac{l_M}{\mu_0 \times \mu_R \times A_E}$$

is called reluctance, R, and is purely material- and geometry-dependent. It is analogous to resistance in the electrical domain.

TRANSFORMERS, INDUCTORS, AND SMPSs

You normally construct transformers using an iron core because iron is highly permeable, enabling it to be efficient at transferring energy from the primary to the secondary winding. The purpose of the transformer core is not to store energy but to

act as an instantaneous conduit. Practically speaking, a transformer does store energy in its magnetizing and leakage inductances. These inductances degrade performance, and the goal of transformer design is normally to minimize them. Mutual inductance is a measure of the coupling between the primary and the secondary winding. Leakage inductance occurs when the magnetic flux does not fully couple to the secondary winding.

An inductor stores and releases energy to smooth the current through it. A flyback transformer is an inductor with multiple windings; it stores energy it takes from the input during one portion of the switching period—that is, the on-time—and then delivers energy to the output during a subsequent interval—that is, the off-time. For a core to act as an efficient conduit for magnetic flux, it must contain highly permeable material. Such materials are inherently incapable of storing significant energy, however. When an application requires energy storage, you accomplish the task by creating one or more nonmagnetic gaps in series with the core and store the energy across the gap. The following **equation** defines stored energy in a magnetic circuit per unit volume:

$$W = \frac{1}{2} \times \frac{B^2}{\mu},$$

where W is the energy stored and μ is the permeability of the material. Because magnetic materials are highly permeable, they can store little energy. The addition of an air gap reduces the effective permeability and allows for energy storage.

Figure 2b shows the effect of the introduction of an air gap on the BH characteristic of the magnetic circuit. The air gap tilts the BH curve to the right. The ungapped core would saturate at a field intensity of H_1, whereas a gapped core can be useful up to a field intensity of H_2, where H_2 is greater than H_1. Because current is the prime driver of magnetic-field intensity, you can effectively push more current through the core without saturating it.

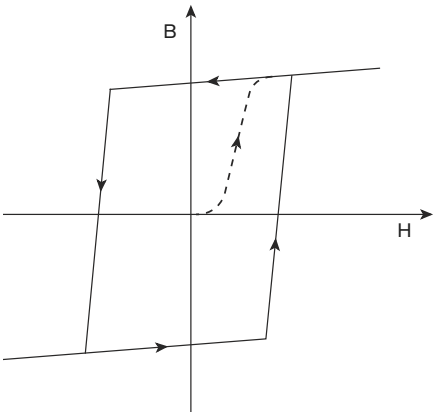

Figure 3 Ideal magnetic materials have a square-loop characteristic with high permeability and insignificant energy storage until they finally reach saturation.

This concept is analogous to storing energy across the insulation layer between the two plates of a capacitor. You create the air gaps by constructing the core with either physical discrete air gaps or from a granular composite material with distributed air gaps. The magnetic particles are isolated from one another in a solid, nonmagnetic binder. In this way, the gap is effectively distributed throughout the core. The core's function in a flyback transformer remains the same whether you construct it with distributed air gaps or composite material: It provides the flux-linkage path between the primary winding and the gap and between the gap and the secondary windings.

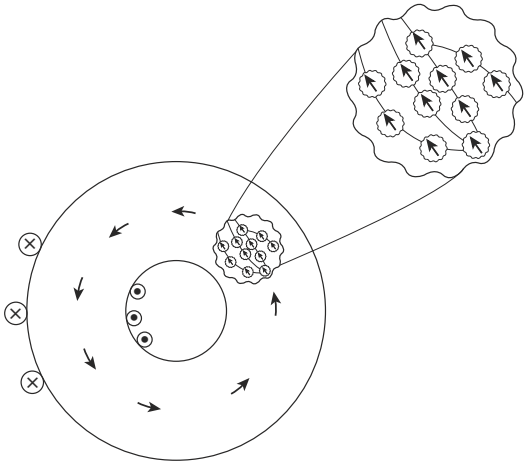

Figure 4 At low flux densities, flux tends to concentrate in the easiest—that is, lowest-resistance—paths, in which the particles are in close proximity.

MAGNETIC-CORE MATERIALS

Because of its low cost and high-saturation flux density, laminated silicon steel is the most popular material for low-frequency applications; however, it exhibits high core losses at higher frequencies. At higher frequencies, cores require more exotic materials, such as low-loss amorphous metal alloys and composite powdered-metal cores, including powdered iron, Kool Mu, and Permalloy powder and ferrites.

Ferrites are the most popular core materials for SMPS applications because they exhibit low losses and are inexpensive. Ferrites are ceramic materials manufacturers develop by sintering a mixture of iron oxide with oxides or carbonates of either manganese and zinc or nickel and zinc. The main disadvantage of ferrite is that, being a ceramic, the core is mechanically less robust than other materials and may be unacceptable in high-shock environments.

Ideal magnetic materials have a square-loop characteristic with high permeability and insignificant energy storage until you finally drive them into saturation—that is, when they develop a sharp saturation characteristic (**Figure 3**). With metal-alloy cores, the characteristic approaches the square-loop characteristic of that in **Figure 3**. Composite metal-powder and ferrite cores exhibit a rounded, or soft, characteristic due to the particulate structure of the core material. In composite metal-powder cores, nonmagnetic gaps exist between the discrete magnetic particles. Similar nonmagnetic occlusions occur among the sintered particles in ferrite cores. These tiny gaps cause the distribution of flux and flux changes across the entire core, rather than at a discrete flux-change boundary.

At low flux densities, flux tends to concentrate in the paths with the lowest resistance, in which the particles are in close proximity (**Figure 4**). As the flux density increases, these paths are the first to saturate. Any incremental flux must then move to adjacent paths in which the magnetic material has not saturated but the gap is somewhat wider. This process continues, effectively widening the incremental distributed

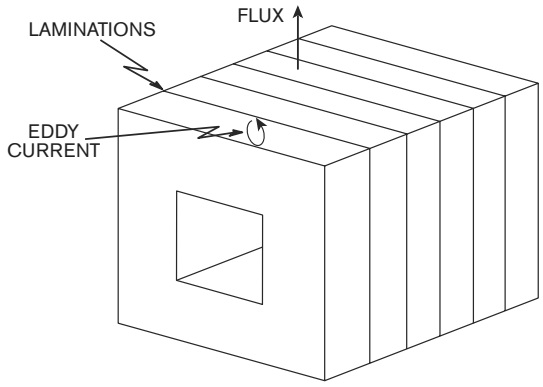

Figure 5 To minimize eddy current, you can divide transformer cores into electrically insulated laminations. The flux divides between the laminations, and the laminations present a much higher resistance than does the bulk core. The result is that the laminated core presents an effective eddy-current resistance many times greater than that of an equivalent solid core.

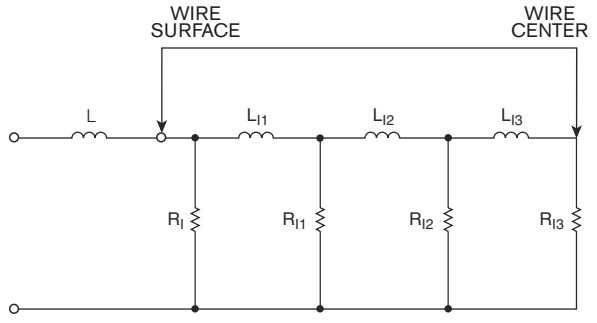

Figure 6 At dc or low frequency, the effect of inductance is trivial, and the current distributes itself evenly from the surface to the center, minimizing losses. However, at high frequency, the inductive reactive of L_I becomes larger, so L_I effectively blocks the current from flowing in the center of the wire and concentrates at the surface.

gap as the flux increases. The incremental permeability and inductance thus progressively decrease, creating the rounded shape of the BH characteristic.

When you add a discrete gap to a ferrite core for energy storage in a filter or flyback application, the rounding of the ferrite characteristic disappears because the high reluctance of the gap swamps it, and the inductance characteristic becomes linear until it reaches saturation—another advantage of using ferrite cores in SMPS applications.

EDDY-CURRENT LOSS

Eddy current is an induced current that flows around the core. Eddy-current loss can occur in both transformer cores and windings at high frequencies. The parasitic eddy-current characteristic is a function of the volts you apply per turn and duty cycle. At high SMPS frequencies, eddy currents can cause serious problems.

You can think of the core itself as a single-turn secondary winding that links to all the windings. It induces a voltage, equal to the volts per turn you apply to the windings, around the core's periphery. This voltage induces a current flow around the core, resulting in energy losses that manifest themselves as heat due to the internal resistance of the battery. If the core comprises high-resistivity material, such as ferrite, the eddy current is low, and eddy-current loss is insignificant in SMPS applications. For metal-alloy cores, resistivity is low; in solid-metal cores, eddy current would cause a shorted turn. One approach to this problem is to break the core into electrically insulated laminations (**Figure 5**). The flux divides between the laminations, and they present a higher resistance than that of the bulk core. Thus, the laminated core presents an effective eddy-current resistance many times greater than that of an equivalent solid core. For this reason, all iron-alloy transformer cores have laminations.

WINDINGS AND HIGH FREQUENCIES

The behavior of electric currents in high-frequency windings can differ greatly from those in low-frequency applications, resulting in high-frequency skin effects and proximity effects. To understand the skin effect, consider that electricity is like water: Both always take the easiest path available—that is, they follow paths requiring the lowest expenditure of energy. At low frequencies, electricity accomplishes this goal by minimizing heat-caused losses. At high frequency, current flows in the paths that minimize inductive energy. Conservation of energy causes high-frequency current to flow near the surface of a thick conductor, even though this flow may result in higher losses (**Figure 6**).

In the **figure**, you can see the surface and the center of the wire. L is the inductance of the wire. L_I is the inductance within the wire from the surface to the center, and R_I is the distributed resistance. At dc or low frequency, the effect of L_I is trivial, and the current distributes itself evenly from the surface to the center, minimizing loss. However, at high frequency, the inductive reaction of L_I becomes larger, so L_I effectively blocks the current from flowing in the center of the wire, and it concentrates at the surface. The penetration, or "skin," depth is the distance from the conductor surface to a point at which the current density is one electric-field current times the surface current density. At 100 kHz in copper, penetration depth is 0.24 mm, meaning that, at a 100-kHz frequency, the largest diameter of wire you can use is 0.48 mm—that is, AWG #25. If your design requires thicker wire, you should instead consider paralleling two thinner wires.

Another important high-frequency effect that has an impact on transformer windings is the proximity effect. When two conductors thicker than the penetration density are in proximity and carry current in the same direction, the magnetic-flux lines are denser near the wire junction (**Figure 7**). Consequently, current tends to flow along the halves of the wire that are not in close proximity to each other. You must consider this effect when deciding the strategy for placement of transformer windings.

Proximity losses increase exponentially as the number of layers increases. To combat proximity losses, the cores for high-frequency SMPS applications often have a window shape with the winding much wider than it is high, thus minimizing the number of layers. This approach is not a panacea, however, because stretching the windings increases capacitance between

them. An interleaving winding strategy, on the other hand, alternately places the primary and secondary windings on top of each other, avoiding the need for an elongated window. Contrary to popular belief, in high-frequency SMPSs, it is often better to leave the window area unused rather than pack it with copper to reduce dc resistance. The increased layers can increase ac losses by as much as 10 times as compared with dc losses.

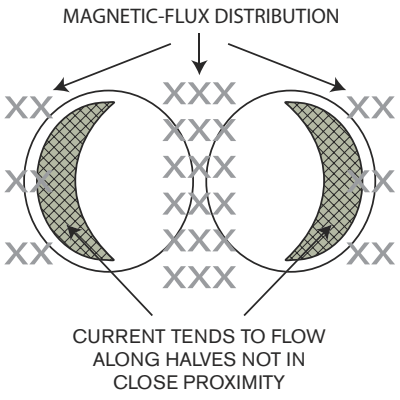

Figure 7 When two conductors thicker than the penetration depth are in proximity and carry current in the same direction, the magnetic-flux lines are denser near the wire junction. Consequently, current tends to flow along the halves of the wire that are not in close proximity to each other.

OTHER CONSIDERATIONS

Minimizing losses is vital to SMPS design, but minimizing EMI (electromagnetic interference) is also critical. One cause of EMI is interwinding capacitance that couples common-mode noise from the primary to the secondary windings. You can reduce this capacitance by using a Faraday shield—a thin foil or metallized film in the intervening space between the primary and secondary windings. The shield should be thinner than the penetration density to avoid eddy-current losses and should connect to the low-impedance—that is, nonswitching—node of the transformer's primary. Special winding configurations, such as those that Power Integrations (www.powerint.com) developed, can also reduce EMI (**Reference 1**). The unterminated windings on the transformer act as electrostatic shields within the transformer. The transformer shields cancel some of the switching signals, significantly attenuating the composite signal across the parasitic capacitance and significantly reducing the transformer's parasitic (capacitive) coupling paths (**Figure 8**).

Core hysteresis, eddy-current, and winding losses, which generate temperature rise, occur in all transformers. In buck-derived applications, under fixed-frequency operation, volt seconds and flux swing are constant. Hysteresis loss is therefore constant, regardless of changes in input voltage or load current. Core eddy-current loss is proportional to the input voltage; thus, worst-case loss occurs at high input voltages. Winding losses are proportional to the duty cycle, which is greatest at low input voltage.

SMPS designers must select cores that will remain within a safe working temperature under any of these worst-case conditions. They should also consider start-up and transients when a sudden increase in load current occurs. The control loop

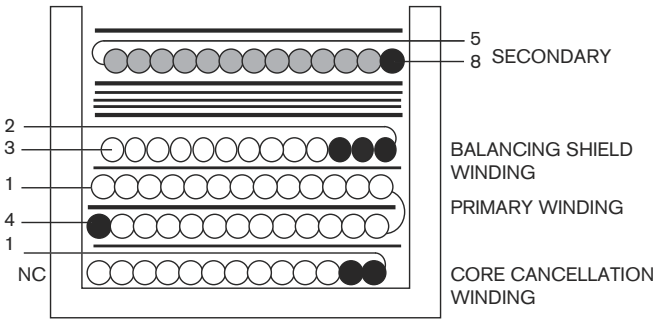

Figure 8 The unterminated windings on the transformer act as electrostatic shields within the transformer, canceling some of the switching signals, attenuating the composite signal across the parasitic capacitance, and reducing the transformer's parasitic (capacitive) coupling paths.

calls for full current, pushing the duty cycle to its absolute-maximum limit. If the input voltage is at the maximum, the volt seconds you apply to the transformer windings could be several times larger than normal, driving the core into saturation. Fortunately, modern control ICs include soft-start and sophisticated current- or volt-second-limiting circuitry to protect the system from both operational and fault conditions.

To design a flyback transformer for an SMPS, first define the circuit parameters and select the core material and geometry. Next, determine the peak current in the circuit and the maximum flux density and flux swing. After making those determinations, tentatively select the core's shape and size and determine the loss limit for flux density or core losses. You then calculate the number of turns, the gap length, the conductor size, and the winding resistance. Last, calculate flux density, winding loss, total loss, and temperature rise and then adjust these values to the size of the core. Available software, such as PI Expert from Power Integrations, can help with this process (**Reference 2**). The available packages can automate power-supply design, including component selection and transformer design, enabling even less-experienced design engineers to successfully complete an SMPS design that complies with international standards for efficiency, safety, and EMI. Detailed knowledge of magnetics is not necessary, but it helps to have a basic understanding of what the software is doing. **EDN**

REFERENCES

1 "Power Supply Design" and "PI Expert Support" technical forums, www.powerint.com//forum/ask-pi-engineer.
2 "PI Expert Design Software," www.powerint.com/pi-expert.

AUTHOR'S BIOGRAPHY

Sameer Kelkar is a senior applications engineer at Power Integrations, where he designs and develops tools, design ideas, and reference-design kits for the company's products. He also provides technical support to internal divisions to ensure product quality. Kelkar has a master's degree in electrical engineering from the University of Minnesota (Minneapolis, MN).

EDITED BY MARTIN ROWE
AND FRAN GRANVILLE

designideas

READERS SOLVE DESIGN PROBLEMS

Inspect solar cells without a microscope

Chun-Fu Lin and Tai-Shan Liao,
National Applied Research Laboratories, Hsinchu, Taiwan

Solar cells convert light energy into electricity, making them a renewable energy source. Solar-cell manufacturers often use SEMs (scanning electron microscopes) to detect defects in solar cells while they're still in wafer form. Although SEMs can see down to a solar cell's grain structure, they can be slow because their scan area is small. A SEM must scan a wafer many times to cover it.

Instead of using a SEM, you can use an SWIR (shortwave-infrared) camera system to detect defective cells. You can take advantage of a solar cell's electroluminescence signature to find defects on a solar cell. A cell's light has a wavelength of about 1.1 micron, which results when you apply a forward bias voltage and forward operating current of at least 7A to the cell. An SWIR sensor can provide an image of an entire wafer, eliminating the need to scan the wafer. The sensor identifies defects by detecting a wafer's electroluminescence.

Figure 1 shows the system, which uses an SWIR sensor that converts an image into an analog voltage. A preamplifer boosts the signal to a level sufficient for an ADC in a digital-processing module to digitize the analog signal at 10M samples/sec.

The ADC's digital output travels through an LVDS (low-voltage-differential-signaling) data interface to a Dalsa (www.dalsa.com) frame-grabber card in a computer. Custom image-processing software, written in C++, processes the data, producing an image of the entire wafer on the computer's screen.

The board containing the sensor, preamplifier, and ADC also has a microcontroller, which generates a clock signal for the timing of the sensor and the ADC. An RS-232 communica-

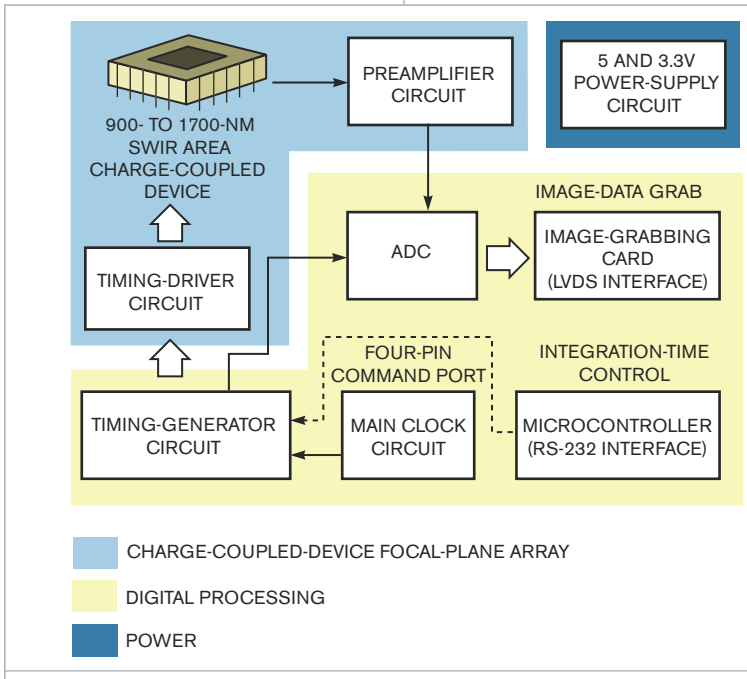

Figure 1 An ADC digitizes an analog signal from an SWIR sensor and sends the signal to a frame grabber for processing.

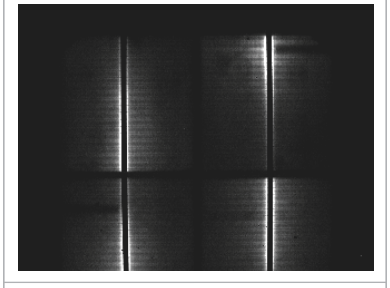

Figure 2 An electroluminescence image of solar cells shows dark areas that indicate failed cells.

DIs Inside

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

tions port on the Atmel (www.atmel.com) microcontroller allows it to communicate with a PC to get commands from the user who set parameters such as the SWIR sensor's operating mode. A timing-driver circuit sends the clock signal to the SWIR sensor.

Figure 2 shows the image from the SWIR camera circuit. This image shows the intensity distribution of the cell's light output. A homogenous intensity-distribution image is essential for a high-quality solar cell, but solar cells always show some inconsistencies. All defects resulting in a local reduction of the carrier concentration are visible on the electroluminescence image as dark bars. EDN

Solar-powered sensor controls traffic

Larry K Baxter, Capsense, Lexington, MA

Have you ever sat in your car waiting for the light to turn green when nobody's using the cross street? This wait is due to the fact that the sensors controlling these traffic signals—in one large-suitcase-sized box per intersection—are classically dumb, with relays, cams, and switches, although they now may include software that accepts data from local sensors, automobile-sized inductive loops buried in the asphalt. Modern controllers have gained some intelligence. For example, they may share

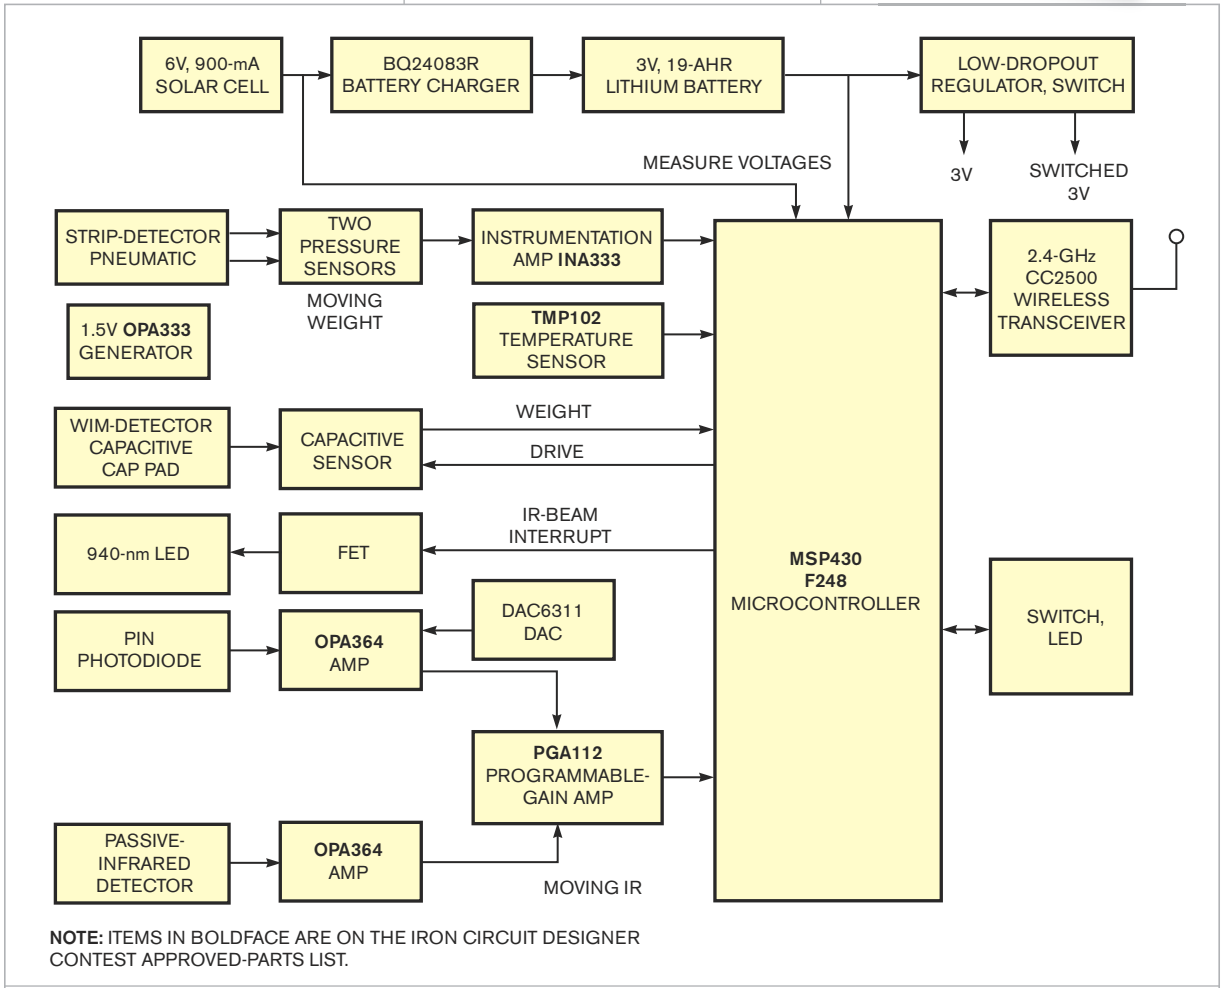

Figure 1 Most of the circuit amplifies outputs from four sensors, digitizes them with the MSP430's 12-bit ADC, does some preprocessing, and messages the controller.

data with nearby intersections, respond to radio requests from emergency vehicles, and sometimes take commands from a traffic-control center. This Design Idea describes the TSP (traffic-sensor post), a more accurate, effective, inexpensive, and easy-to-install approach to monitoring traffic flow. These sensors measure vehicle location and speed in four or more streets at an intersection or at a distance from the intersection for early warning. A second application of this technology, the WIM (weight-in-motion) sensor, weighs moving trucks.

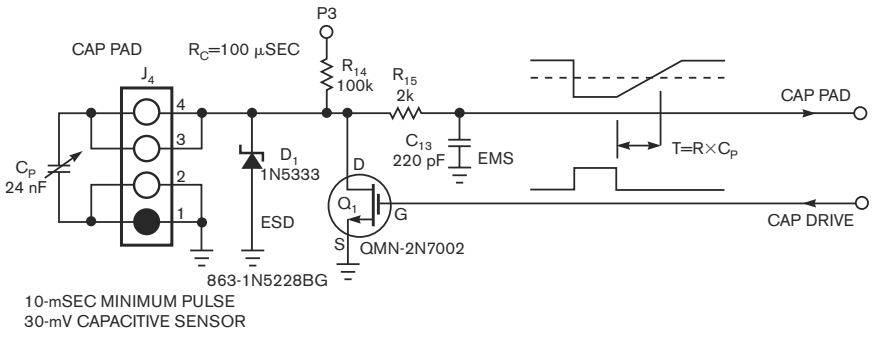

Figure 2 The Cap Pad sensor has a nominal capacitance of about 24 nF at rest, with a change of about 7% full-scale when a truck passes.

The circuit comprises a wireless, solar-powered sensor array that handles all the data collection at an intersection (**Figure 1**). Cities can install these sensors at each of the four corners of an intersection for full coverage. The sensors send data to the single controller box over IEEE 802.15.4 in a star network. The approach combines four sensors in an inexpensive, low-maintenance, 6-in.-diameter, 6-foot-tall post. You can build the circuit into the post that holds the traffic lights, or you can use it stand-alone. Not all TSPs require all four sensors; you can select those that your application needs based on usage. The TSP is the first wireless approach to this problem, and one of the sensors, the Cap Pad, provides a huge advantage over current expensive and inaccurate WIM sensors (**Figure 2**).

The TSP uses a PIR (passive-infrared) sensor that looks 10 microns into the deep-IR band for moving IR sources. This technology finds use in inexpensive motion-detecting lamp controls and senses vehicles from 30 feet away. The detection range is good, the parts are cheap, and the beam can see through a layer of dirt. It can't measure speed, distance, or direction.

The TSP also uses conventional pneumatic tubes. Rubber tubes are stapled to the asphalt and feed two pressure sensors. This approach accurately measures speed, but permanent installations cannot use it because it gets damaged easily. Municipalities often deploy pneumatic tubes to measure traffic volume in road construction.

The Cap Pad comprises a 10-in.×12-foot sandwich of three 0.05-in.-thick stainless-steel sheets separated by two 0.05-in.-diameter closed-cell urethane-foam layers (**Figure 3**). You capacitively measure the 0.025-in. deflection of the pad under a truck's tire to weigh the axle. One Cap Pad can handle the WIM requirements, and using two can add speed and direction information. You use multiple pads to handle multilane roads. The Cap Pad can be fastened to the asphalt with adhesive or pavement tape or buried under as much as an inch of asphalt for protection. Its materials cost is only a couple hundred dollars, a huge saving over the piezoelectric WIM sensors currently in use.

The TSP also uses a near-IR transmitter/receiver using a pulsed LED for transmission and a PIN (positive-intrinsic-negative) photodiode for reception. Both need cylindrical lenses to focus the beam to a 2°-wide, 5°-high ellipse that covers a remote retroreflective screen, as in highway signs, or to the IR sensors on another TSP. A multilayer optical bandpass filter that removes visible light further improves the range.

Precision capacitive sensors can measure an air gap between adjacent metal plates to subnanometer accuracy. Unfortunately, accuracy in the WIM application requires flat and parallel surfaces, and the Cap Pad has neither. Capacitive sensors can also accurately measure a force on adjacent flat plates with a restoring spring, but flatness and parallelism are still requirements. Maintaining parallelism over a 10-in. pad would be difficult, and roads are seldom flat.

If compression of the air pockets in closed-cell foam provides the restoring force, however, the resulting spring constant changes from the conventional F=K×x of springs or cantilevered beams to F=P0×H/(H−x), where F is force, P0 is atmospheric pressure, H is the starting gap, and x is the dis-

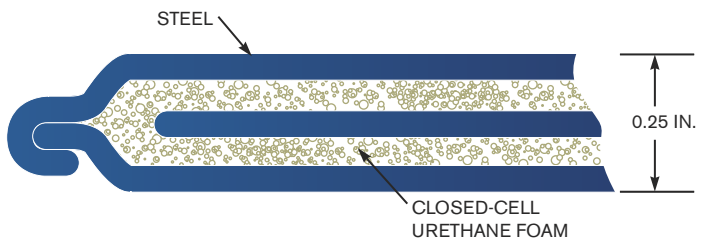

Figure 3 The Cap Pad sensor is a 10-in.×12-foot sandwich of three 0.05-in.-thick stainless-steel sheets separated by two 0.05-in.-diameter closed-cell urethane-foam layers.

Industry's smallest power-management IC for portable media players

Integrates Li+ charger, Smart Power Selector™, three step-down converters, and a WLED step-up converter in a tiny 4mm x 4mm package

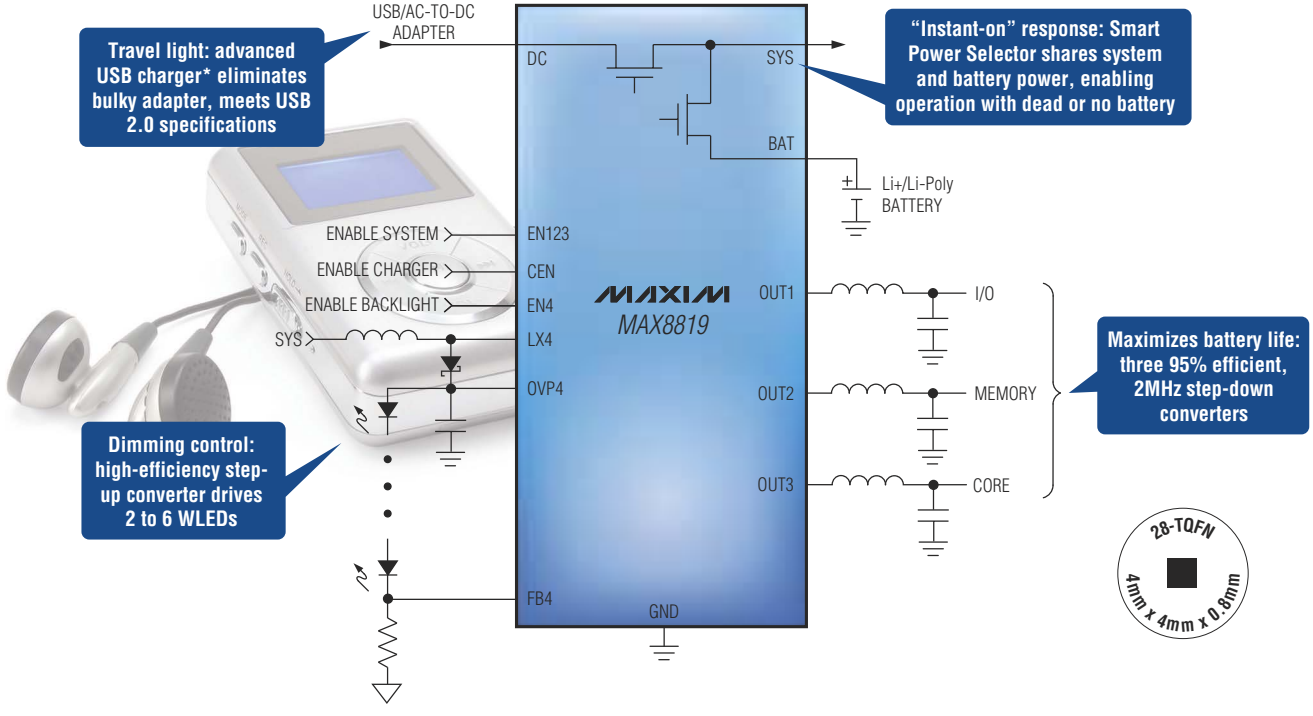

~88mm² solution area, 1mm solution height

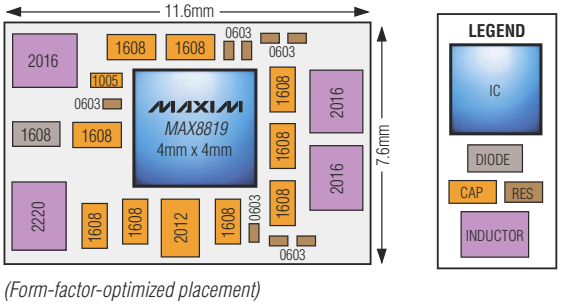

(Form-factor-optimized placement)

Evaluation kit

Smart Power Selector is a trademark of Maxim Integrated Products, Inc.
*U.S. Patent #6,507,172.

www.maxim-ic.com/MAX8819-info

MAXIM DIRECT.
www.maxim-ic.com/shop

AVNET electronics marketing
www.em.avnet.com/maxim

MAXIM
INNOVATION DELIVERED™

For free samples or technical support, visit our website.

placement. The result of this **equation** is that the capacitance of the pad varies linearly with applied force, and the surfaces of the Cap Pad no longer need to be parallel or flat. It accurately measures a force regardless of its size.

Most of the circuit amplifies outputs from the four sensors, digitizes them with the MSP430's 12 bit-ADC, does some preprocessing, and messages the controller. The 6V solar panel, 40 IXYS (www.ixys.com) solar cells in series, charges a 19-Ahr, 3V, lithium-polymer battery through IC_1. Low-dropout regulator/switch IC_2 regulates battery output at 3V. The battery generates more than 4V at full charge and 3.2V at the end of charge, and the low-dropout regulator at 42 mA generates only 50 mV. IC_2 also switches active-mode 3V power.

The road-strip sensor senses the 0.1- to 1-psi pulse when a car drives over the pneumatic tubes. A 400Ω silicon bridge sensor differentially outputs approximately 50 mV. Instrumentation amplifiers IC_3 and IC_4 boost the output to a few volts. The pressure sensor, as well as the Cap Pad and the PIN sensor, has a quiescent level with no traffic. A timer detects the no-traffic state and stores this level in RAM, updating every second to follow slow offset drifts from environmental factors, so sensor offset accuracy is not critical. The pressure sensor's scale accuracy—at approximately 30%—is relatively uncritical, but the Cap Pad's scale accuracy should be a few percentage points or less. All sensors must have good resolution.

IC_5 handles accurate temperature measurements, which are necessary for the Cap Pad, whose temperature dependence results from the elastic modulus change of polyurethane. The Cap Pad has a nominal capacitance of about 24 nF at rest, with a change of approximately 7% full-scale when a truck passes. The Cap Drive pulse discharges this capacitance at a 700-Hz rate, and a 100-kΩ resistor charges it to 3V with a 240-µsec time constant. A timer times the number of pulses it takes to cross the internal $V_{DD}/2$ reference using the internal comparator, and, because you can clock the timer at 12 MHz, the resolution is 1%. You can get increased resolution by timing out the nominal quiescent pulse width and capturing the pulse's level at that point with the 12-bit ADC.

The Cap Pad's sandwich construction shields the active element from electromagnetic interference, but a 3W zener diode cleans up any remnant lightning strokes. The IR LED drive is a 20-to-1 current mirror to handle LED voltage variation. A DAC handles the PIN photodetector's offset because the extreme night-to-day dynamic range would overrange the 12-bit DAC. The PIR sensor turns moving deep-IR targets into bipolar millivolt voltage pulses with its special segmented lens and dual-element pyroelectric detector. A PGA (programmable-gain amplifier) selects and variably amplifies the PIR sensor's signal and the PIN signal. The timer uses standard connections.

For a power budget, more schematics, and more details of this circuit, see the Web version of this article at www.edn.com/091126dia.EDN

Self-oscillating H bridge lights white LED from one cell

Luca Bruno, ITIS Hensemberger
Monza, Lissone, Italy

You can build a self-oscillating H bridge by replacing the pull-up collector resistors of a classical BJT (bipolar-junction-transistor) astable multivibrator with PNP BJTs (**Figure 1**). Because this circuit oscillates at supply voltages as low as 0.6V, you can use it in general low-voltage, low-power push-pull applications. You can, for example, drive a diode-capacitor charge pump to generate negative supply voltage in battery-powered systems. This Design Idea shows how to use it to light a white LED from one cell without an inductor.

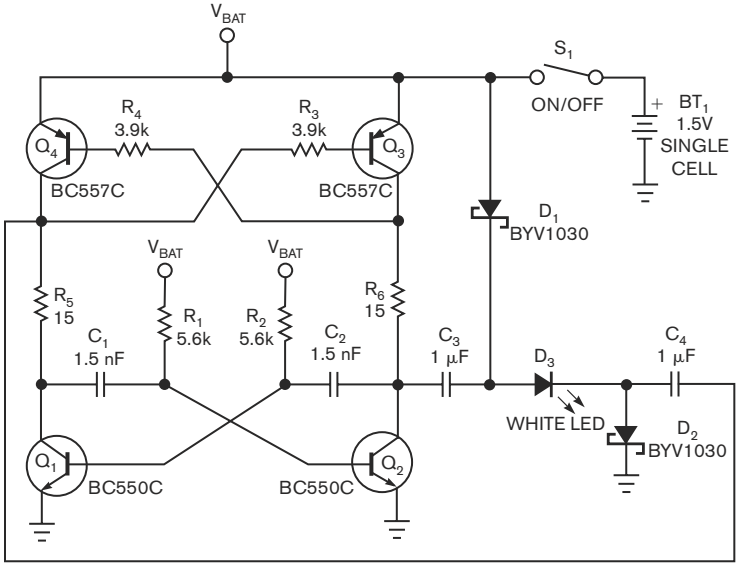

Figure 1 Resistors R_1 and R_2 and capacitors C_1 and C_2 set the oscillation frequency.

Smallest low-frequency RF products have highest performance and efficiency with lowest current draw

300MHz to 450MHz Transmitters

Part	Features	Maxim Package Area (mm^2)	Closest Competition Package Area (mm^2)	Maxim Current Consumption (mA)	Closest Competition Current (mA)	Maxim Tx Power (dBm, max)	Closet Competition Tx Power (dBm, max)
MAX1472	ASK	9	32	9.1	14	10	10
MAX7044	ASK, clock output	9	32	13.8	14	13	10
MAX7057/ MAX7058	Frac-N programmable frequency ASK/FSK, dual-frequency ASK	60	50	8.5mA ASK, 13mA FSK	17	10	10

300MHz to 450MHz Transceivers

Part	Features	Maxim Package Area (mm^2)	Closest Competition Package Area (mm^2)	Maxim Current Consumption (mA)	Closest Competition Current (mA)	Maxim Tx Power (dBm, max)	Closest Competition Tx Power (dBm, max)
MAX7030/MAX7031/ MAX7032	315/345/433.92MHz ASK, 308/315/433.92MHz FSK	25	49	8.5 to 12.5	16	10	10

300MHz to 450MHz Receivers

Part	Features	Maxim Sensitivity (dBm)	Closest Competitor Sensitivity (dBm)	Maxim Current Consumption (mA)	Closest Competitor Current Consumption (mA)
MAX1473/MAX7033	ASK, 3.3V/5V, AGC, AGC hold (MAX7033)	-114/-114	-113	5.5 at 3V	5.0 at 5V
MAX1471	ASK/FSK, polling timer	-114 ASK, -108 FSK	-113 ASK, -105 FSK	7.0 at 2.4V	5.7 at 5V
MAX7042	FSK, 3.3V/5V	-110 FSK	-97 FSK	7.0 at 2.4V	9.0 at 2.7V

www.maxim-ic.com/LFRF-info

MAXIM DIRECT
www.maxim-ic.com/shop

AVNET electronics marketing
www.em.avnet.com/maxim

MAXIM INNOVATION DELIVERED™

For free samples or technical support, visit our website.

Transistors Q_1, Q_2, Q_3, and Q_4 form the H bridge, which acts as a simple charge-pump converter and requires only two small, inexpensive ceramic capacitors, C_3 and C_4, to perform its function. When Q_2 and Q_4 are on, capacitors C_3 and C_4 charge to the battery voltage through forward-biased Schottky diodes D_1 and D_2. When Q_1 and Q_3 are on, they discharge the capacitors through resistors R_5 and R_6 and the LED. Because this process repeats at a high rate of speed, the LED appears always on.

The circuit oscillates with a frequency based on time constants R_1C_1 and R_2C_2. During discharge, the voltage that develops across resistors R_5 and R_6 and the LED remains approximately constant because of the high switching frequency. The measured value, for a nominal 1.5V battery voltage, is 3.8V—enough to drive a white LED with a forward voltage of 3 to 3.5V. Resistors R_5 and R_6 set the LED's peak current and limit the possible current spikes that a push-pull output stage can produce.

Choosing the astable oscillator's frequency involves a trade-off between the time necessary to charge capacitors C_3 and C_4 and the need to reduce their discharge. For a given capacitance value of C_3 and C_4, you must experiment to find the optimum frequency. With the component values in **Figure 1**, the frequency and the duty cycle are about 66 kHz and 50%, respectively, and the LED's drive current is a square-wave signal with 20-mA peak value and 10-mA average value. The LED dims gradually as the battery voltage decreases, and the LED is off when the battery voltage falls below 0.9V. For high efficiency, use small-signal transistors with high dc current gain and low collector-to-emitter saturation voltage. Note that the circuit can drive any type of LED; in this case, you should increase current-limiting resistors R_6 and R_5 to achieve the LED-drive current your application requires. EDN

Low-cost LCD-bias generator uses main microcontroller as control IC

Tom Hughes, Dannemora, Auckland, New Zealand

LCD circuits often require a −10V voltage at 2 to 15 mA to bias a graphics-LCD-driver IC. You can usually accomplish this task with an external charge-pump IC, such as Maxim's (www.maxim-ic.com) ICL7660, but that approach adds cost to the design. Instead, you can control a buck-boost switch-mode regulator using the same microcontroller that sends data to the LCD. In addition, you can sequence the power rails under software control, as some types of LCD controllers require.

The circuit includes IC_1, an Atmel (www.atmel.com) Attiny15 microcontroller (**Figure 1**), which provides regulation with 200-mV-p-p ripple at a 30-mA load current when supplying −10V. **Listing 1**, which is available in the online version of this Design Idea at www.edn.com/091126dib, lets you download the source code, which uses only 4.8% of the total CPU time to achieve the stated regulation, even with a relatively low-speed clock frequency of 1.6 MHz.

To minimize CPU time, the software uses the 8-bit on-chip PWM (pulse-width modulator) to drive Q_1. With the on-chip ADC in free-running mode, the microcontroller generates a hardware interrupt with a period of 7.69 kHz. The interrupts have one drawback: If they stop, the circuit can go out of regulation. Thus, you must take care when using interrupts with long processing times. The Attiny15 uses an on-chip, 16× PLL (phase-locked loop) to drive the PWM timer. You can achieve a PWM carrier frequency of 100 kHz, which allows the use of a relatively low-capacitance filter capacitor, C_1.

Two constants in the source code let you alter the bias voltage of the circuit's output voltage. These constants employ basic buck-boost-converter theory (**Reference 1**). The following **equation** defines the maximum 8-bit constant, or threshold, that the ADC reads on the chip: $51.2\times\{V_{CC}-[(V_{CC}-V_{MAX})/(R_4+R_5)]R_5\}$, where V_{MAX} is the maximum desired output voltage and V_{CC} is the supply voltage. To achieve optimum operation, increase the PWM signal's duty cycle when you need higher voltages. Use the following **equation** to determine the 8-bit PWM's value: $255-V_{OUT}/(V_{OUT}-V_{IN})\times255$, where V_{OUT} and V_{IN} are the output and input voltages, respectively. In practice, however, if you keep the current at less than 2 mA, this requirement is less important.

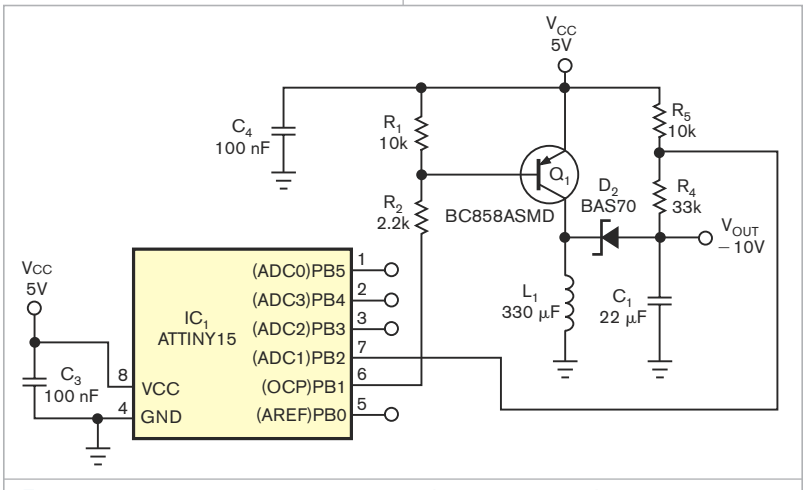

Figure 1 An Attiny15 microcontroller provides regulation with 200-mV-p-p ripple at a 30-mA load current when supplying −10V.

The circuit can deliver currents that Q_1's collector current predominantly delivers. This current is the peak output current that the circuit can safely deliver. The following **equation** calculates the current: $I_{OUTMAX}=(V_{IN}\times0.08)/V_{OUT}$, where I_{OUTMAX} is the maximum output current. If your design needs higher current, then substitute a BC327 for Q_1. Additionally, the inductor should have a maximum rms (root-mean-square) current value of at least twice the peak output current and preferably be a low-ESR (equivalent-series-resistance) type to maximize circuit efficiency. **EDN**

REFERENCE

1 Hart, Daniel W, *Introduction to Power Electronics*, First Edition, pg 202, Prentice Hall, Oct 25, 1996, ISBN-10: 0023511826, ISBN-13: 978-0023511820.

productroundup

OPTOELECTRONICS/DISPLAYS

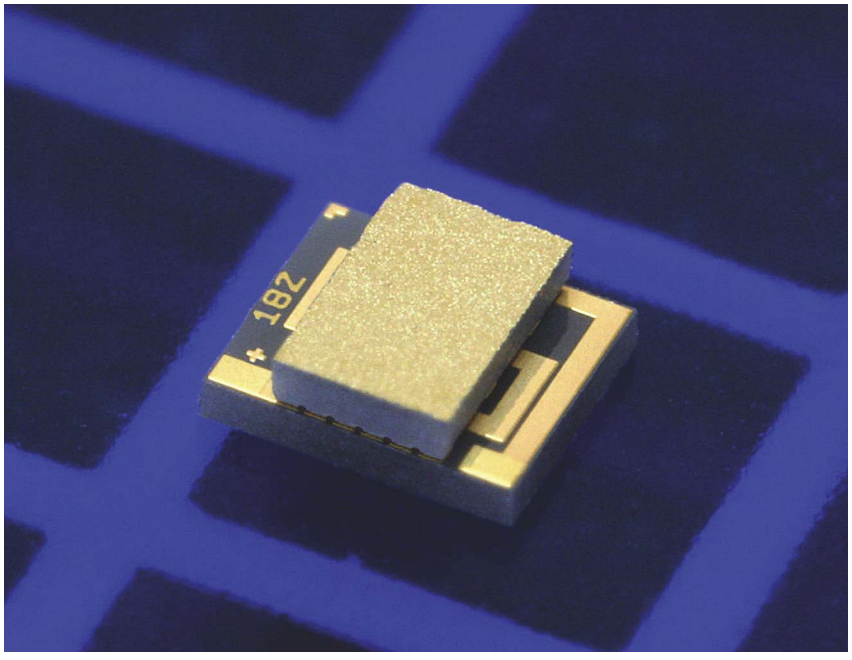

Thermoelectric cooler meets 3000g shock-testing standard

Passing the 3000g shock test, the OptoCooler HV14 thermoelectric cooler meets requirements for mechanical shock for military standard 883E Method 2002. At 85°C, the high-voltage, heat-pumping thermoelectric module operates at a 2.7V maximum voltage and pumps 1.5W of heat. The device creates a 60°C maximum temperature differential between the hot and the cold sides with no heat load, suiting cooling and temperature control of optoelectronic devices, including laser diodes, avalanche photodiodes, and high-brightness LEDs. Available in a 3×3-mm footprint, the OptoCooler HV14 costs $18 (1000).

Nextreme Thermal Solutions, www.nextreme.com

LED driver features temperature-management control

The LM3424 high-brightness LED driver with temperature-management control supports the use of an online design environment. The vendor's Webench LED Designer tool supports the device, which enables lighting designers to build thermally reliable systems. The driver works in indoor, outdoor, and automotive applications and enables programming of temperature and slope breakpoints, allowing safe operation of the LEDs. During an overtemperature condition, the product's thermal-foldback circuitry reduces the regulated current through the LEDs. A reduced current dims the LEDs to a programmed range, and the LEDs remain within that range until returning to a safe operating temperature. The device can drive as many as 18 high-brightness LEDs in series with a 2A output current and allows regulation of currents based on buck,

boost, SEPIC, flyback, and buck-boost topologies. Operating over a 4.5 to 75V input range, the PWM controller operates at an oscillator frequency as high as 2 MHz. Available in a TSSOP-20 package, the LM3424 high-brightness LED driver costs $1.75 (1000).
National Semiconductor,
www.national.com

Adhesive sheets target touchscreen and flat-panel displays

The self-wetting adhesive protector sheets reduce surface damage to touchscreens and flat-panel displays in use or during shipping, preserving the optical quality and extending the life of the touch device. A self-wetting adhesive applies itself by spontaneously wetting out on smooth surfaces with no or low finger pressure, adhering to a display or touchscreen. The adhesive provides high shear resistance and prevents sheet displacement. The repositionable sheets come in antiglare, clear hard-coat, and antireflective coated versions of the clear protector sheets. Optical properties of the clear protector sheet include 91% transmissivity with 1% haze, and the antiglare protectors provide 90% transmissivity with 8% haze. Available in a maximum 0.007-in. thickness, the protector sheets have a −30 to +70°C operating temperature. The protector sheets cost $1 each.
Fujitsu Components America,
www.fujitsu.com

Miniature, 0.5W LED package measures 3.5 mm²

The miniature, 0.5W OVS5Mx-BCR4 LED package series has a 120° viewing angle and a water-clear lens. The package has a 150-mA power dissipation of 0.48W for white, warm-white, and blue LEDs; 0.51W for green LEDs; and 0.33W for red, amber, and yellow LEDs. The devices provide 4000-, 4500-, and 5000-mcd luminous intensity for the red, amber, and yellow packages, respectively. Luminous flux is 25 lm for the white, warm-white, and green LEDs and 6 lm for the blue LEDs. The devices operate over a −40 to +100°C temperature range. Measuring 3.5×3.5×1.2 mm in a surface-mount housing, the OVS5MxBCR4 series costs 25 cents (1000) in amber, red, and yellow; 45 cents for white and warm white; 52 cents for blue; and 72 cents for green.
Optek Technology,
www.optekinc.com

EDN ADVERTISER INDEX

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

BENOIT LÉVEILLÉ • RF ELECTRONICS TECHNICIAN

Hawk eyes, analog equipment trump expensive digital test set

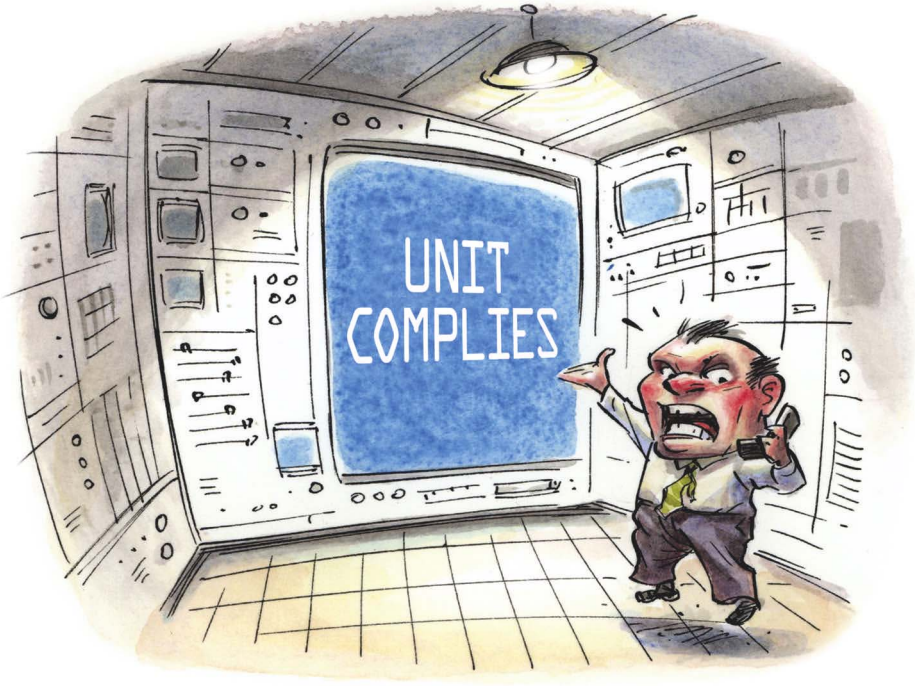

DANIEL VASCONCELLOS

A few years ago, I was a repair technician in RF electronics at a service center for cable-TV equipment. This equipment included everything from commercial satellite receivers and cable amplifiers to head-end TV modulators. My co-workers and I worked with old analog spectrum analyzers, and we interpreted most good and bad signals through visual analysis with instruments at the appropriate setting and scale.

I had been working as a bench technician for about eight years in the company, so my eyes were experienced in seeing problems in any kind of signal.

One day, one of our best customers, complaining of a "noisy picture," sent us a well-known company's TV modulator. Because we were not the authorized repair center, we advised the customer that we would send it to the nearest authorized service shop. Nevertheless, when we received the modulator, we hooked it up to a TV set; the noise on the display was evident at first glance. I took some readings with my old Hewlett-Packard spectrum analyzer and found that the inband noise floor was 40 dB below that of the video carrier, whereas it should have been 60 dB below—and even lower with equipment from this recognized manufacturer. The out-of-band spectrum, on the other hand, was excellent for all signals. We sent the unit to the authorized service shop with our description of the problem.

We received the unit back three weeks later and were astonished to see that it was still defective even after the service shop claimed it had no problem; it was exactly as it was when we sent it. Our engineer called the service shop's technicians and told them the whole story, and they asked us to send the unit back to them. Because we had the unit in hand, however, our engineer asked me to troubleshoot—but not repair—the unit. I deduced that the noise was inband, meaning that something was wrong with the modulation circuit. I traced the problem to the video circuit and found that a faulty transistor was attenuating the video baseband signal. As a result, the AGC (automatic gain control) was boosting the tiny signal and, thus, the noise, explaining why the out-of-band response was so clean. We noted the problem and sent the unit back to the service shop.

A few days later, our engineer received a call from the shop. The manufacturer's technician told him that his shop had found no problem with the unit. Incredulous, our engineer started to argue about the high inband noise level, the noisy picture, and so on. After a few minutes of arguing, the technician at the shop ended the discussion by declaring, "Don't tell me I am wrong! I checked this modulator on a $50,000 test set!" The technician's big, expensive digital unit had probably displayed a "unit complies" message with some mystic numbers and readings. Our technician asked theirs to ship it back, and we decided to repair the unit at no charge to our customer.

Now that most instruments are digital, I am suspicious of what they tell me. Do I see a difference between the scope's sampling rate and the signal rate I measured? I double- and triple-check each response at different settings. From this experience, I learned that nothing replaces the knowledge and skill you learn over the years. A keen eye, some common sense, a good deal of logic, and a simple and appropriate instrument are worth more than the most expensive digital apparatus you work with as a "pushbutton" operator. EDN

Benoit Léveillé is an RF electronics technician in Saint-Eustache, PQ, Canada.

⊞ www.edn.com/tales

BY PANCH CHANDRASEKARAN

Get Up to Speed with Multi-Gigabit Serial Transceivers

Are you being asked to make your next-generation product design connect to a high-bandwidth network with an unfamiliar or a yet-to-be-defined protocol? Are you making the transition from parallel to serial I/O chip-to-chip communications? Or do you just need the highest-serial bandwidth, most reliable multi-gigabit transceivers the industry has to offer?

You are not alone. Serial connectivity is no longer just the design domain of communications engineers. Today, a growing number of designers in the consumer, automotive, industrial control, broadcast equipment, aerospace and defense markets are being tasked with developing products that employ multi-gigabit serial transceiver technology to communicate with next-generation, high-demanding, high-speed networks.

Today's quest for more bandwidth and better efficiency means that many designers like you must quickly get up to speed with the analog nuances of multi-gigabit serial transceivers. Luckily, Xilinx® has engineered its Targeted Design Platforms to help you.

Over the last two decades, the world's top telecommunications companies have relied upon Xilinx FPGAs' mix of logic functionality, high-speed memory, parallel connectivity, and serial I/O capabilities to create every generation of modern communications equipment.

With many years of experience serving the communication markets, Xilinx is able to help designers in a broad set of markets to quickly master gigabit serial transceiver technology and leverage it to create innovative products.

		MAINSTREAM	HIGH-END	ULTRA HIGH-END
	LINE RATES	▪ Up to 3.2Gpbs	▪ Up to 6.5Gpbs	▪ Beyond 11Gpbs
	FOCUS	▪ Simplicity, ease of design, and results	▪ Increased capabilities and performance	▪ Achieving breakthrough bandwdith
MARKET SEGMENTS	Wired	▪ Low-cost, efficient protocol bridging	▪ Advanced protocol mapping, performance optimized backplanes	▪ Cutting edge protocol support, advance processing capabilities
	Wireless	▪ Low-cost, Femto/ Picocell deployment	▪ Cost/Power efficient, mainstream deployment	
	Video Broadcast	▪ Efficient processing and routing capabilities	▪ Accelerated encoding and processing	▪ High-data aggregation/ routing, advanced processing capabilities
	Consumer	▪ Cost-effective integration, simplified serial interfaces		
	Automotive	▪ Low-power, cost-optimized, and flexible serial IO connectivity		
		SPARTAN LXT FPGAs	VIRTEX LXT FPGAs	VIRTEX HXT FPGAs

In 2009, Xilinx introduced both of its next-generation FPGA families — the high performance Virtex®-6 family and low-cost Spartan®-6 family providing Xilinx customers access to a full line of serial transceiver-rich FPGAs that can easily maintain line rates from up to 3.2Gbps in the Spartan-6 family, to up to 6.5Gbps in its Virtex-6 LXT devices, all the way beyond 11Gbps in the Virtex-6 HXT devices.

What's more, Xilinx announced this full range of serial connectivity-enabled FPGAs as the foundation of its new Targeted Design Platform strategy—combining the full line of transceiver-rich FPGAs with world-class tools, validated IP, reference designs, training, and support all delivered in domain and market specific kits. The Targeted Design Platforms allow customers to get their products to market faster than ever before.

With the ISE® Design Suite, customers can now get started on their designs targeting Virtex-6 HXT devices, as well as enhanced support for Spartan-6 LXT and Virtex-6 LXT and Virtex-6 SXT devices. Over the next several months, Xilinx will deliver a number of connectivity enabled design kits targeting wired, wireless broadcast video, packet processing, and traffic management application using either Virtex-6 devices for high performance, or Spartan-6 devices for low cost.

To learn more, visit the Connectivity Page at www.xilinx.com/connectivity where you'll find documentation, videos, links to software and IP downloads, and much more for helping you get up to speed no matter which fast lane you're on.

About the Author: Panch Chandrasekaran is the Sr. Product Marketing Manager at Xilinx Inc. (San Jose, Calif.). Contact him at more_info@xilinx.com

Isolated RS485

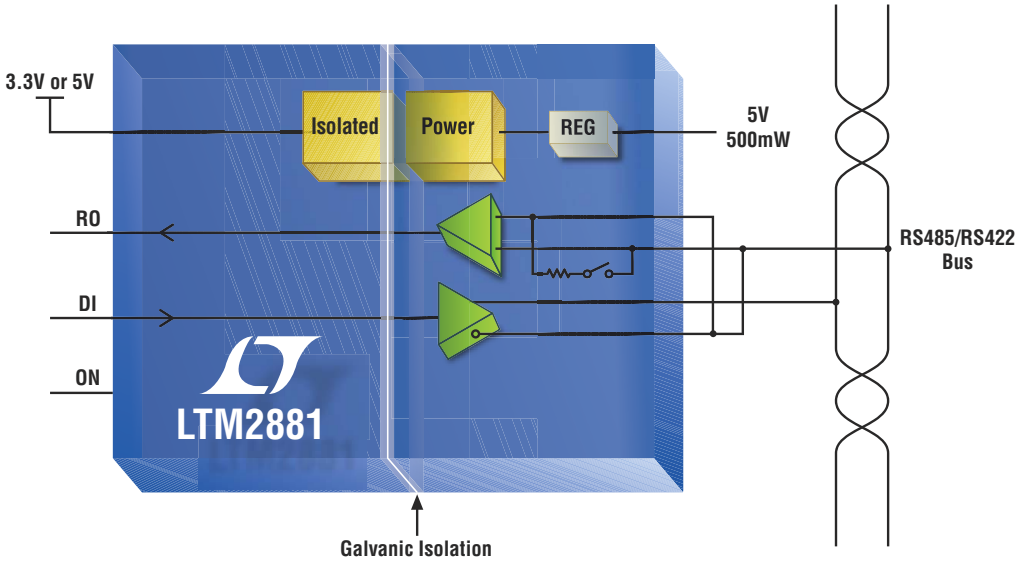

Complete 20Mbps µModule® Transceiver Includes 2500V_{RMS} Isolated Power—No External Components Required

The LTM®2881 is an isolated RS485 transceiver that guards against large ground-to-ground differentials. The LTM2881's internal inductive isolation barrier breaks ground loops by isolating the logic level interface and line transceiver. An onboard DC/DC converter provides power to the transceiver with an isolated 5V supply output for powering additional system circuitry. With 2500V_{RMS} galvanic isolation, onboard secondary power and a fully compliant RS485 transmitter and receiver, the LTM2881 requires no external components and provides a tiny, complete µModule solution for isolated serial data communications.

Features

- Isolated RS485/RS422 Transceiver: 2500V_{RMS}
- Integrated Isolated, Low EMI DC/DC Converter
  - Up to 500mW Surplus Power
- 3.3V or 5V Input Supply Voltage (LTM2881-3/LTM2881-5)
- 20Mbps or Low EMI 250kbps Data Rate
- High ESD: ±15kV HBM
- Common Mode Transient Immunity: 30kV/µs
- Integrated Selectable 120Ω Termination
- Small Footprint, Low Profile (11.25mm x 15mm x 2.8mm) in Surface Mount LGA & BGA Packages

LTM2881 Demo Board

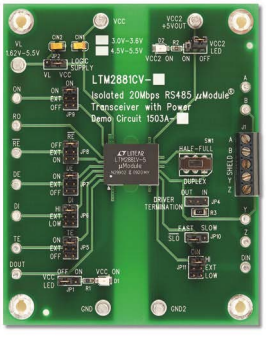

Info & Free Samples

www.linear.com/2881

1-800-4-LINEAR

Free Telecom, Datacom and Industrial Power Products Brochure

www.linear.com/48vsolutions

LINEAR TECHNOLOGY