UBM Electronics

EDN®

VOICE OF THE ENGINEER

MAR 3
Issue 5/2011
www.edn.com

AN EXCITING TIME

BY MICHAEL A. BRIERE

This is an exciting time in power electronics. Everything is about to change. The introduction of commercially viable GaN-based power devices in the 20 to 600 V application range will provide truly revolutionary performance/cost value to a variety of systems from lighting to motor drives and Class D-audio to ac-dc and dc-dc power supplies. This opportunity has been a long time in the making.

Ever since the discovery of the spontaneous formation of a two dimensional electron gas (2DEG) in the highly polarized AlGaN-GaN material system hetero-interface by M. Asif Khan in 1991, the potential for superior power devices has been clear. The intrinsic advantages of this material system based HEMTs of high electron mobility and saturation velocity (> 1600 cm^2/Vs, > 2.5 10^7 cm/s respectively), high carrier density (10^{13} cm^{-2}), together with the high electric field stand off capability of the AlGaN and GaN materials (> 1 MV/cm), provide the possibility of power switches with 10 to > 100 times lower on-resistance than comparable silicon based devices. In addition, the low parasitic capacitances of these intrinsically lateral devices allow for much lower switching times, an advantage leveraged in the fabrication of RF switches, as well as more efficient power devices.

The GaN-based power device program at International Rectifier, referred to as GaNpowIR®, has involved a long, costly and intense effort to turn, what has been to a large extent, an intellectual curiosity into a practical, commercially viable technology platform broadly applicable to the power electronic community. The focus has not been on achieving the ultimate performance possible within the multilayered hetero-epitaxial AlInGaN material system, or to provide devices for small niche applications, rather it has been on optimizing the performance to cost ratio within a manufacturable and reliable device technology platform.

The first issue to resolve involved the choice of substrate. A cursory review of the alternatives will show that only silicon substrates provide the necessary combination of cost, size, quality and volume of supply needed to support the broad power electronics market (currently > $ 10 B annual semiconductor content using > 10 Million six inch wafer equivalents per year). The technological challenge of growing uniform high quality AlInGaN epitaxial layers on silicon substrates, made especially complex due to significant mismatches in both the lattice constants and their thermal coefficient of expansion, has to be met. The requisite control of film thickness and compositional uniformities, epitaxial defects and macroscopic wafer warpage has been achieved on 150 mm silicon wafers through significant engineering efforts.

Next, the device design and fabrication process required additional significant attention. Here the first principle applied is that the fabrication process must be compatible with large volume silicon wafer manufacturing lines. This allows for the use of modern, highly capable equipment, as well as an overhead cost infrastructure competitive with the incumbent silicon based devices, a pre-requisite for widespread adoption. This requirement, often abbreviated as “CMOS compatibility”, requires that such III-V device stalwarts such as gold metallurgy, e-beam lithography and lift-off processes be abandoned. In addition, complex and highly costly processes involving re-growth of AlInGaN epitaxial layers for channel, drift or contact formation are likewise avoided. Standard in-process and post process control methodologies must be applied to establish a high volume production platform. The device design and construction must not compromise on performance, quality or reliability metrics established by the incumbent silicon based technologies. This requires that off-state leakage be kept below 1 µA/mm gate length, significantly less than the 1 mA/mm often employed in the GaN-based device field to date. Similarly, on to off state current ratios must be greater than one million to be comparable to state-of-the-art silicon based power devices. This requires the use of low leakage epitaxial layers, as well as insulated gates, replacing the much more common Schottky metal-semiconductor structures.

Device ruggedness in application conditions must also remain un-compromised with respect to expectations established by the incumbent silicon based technology. Large forward biased safe operating area is an important indication of such robustness. Device stability under accelerated stress conditions for extended periods of time is essential for acceptance in the power electronic community. To date, over 10,000,000 device hours of reliability data has been collected on the low voltage devices released to production by IR in early 2010, with up to 10,000 hours per device. No premature device failures have been found to date and parametric stability has been excellent. An example is shown in Figure 1, where the gate dielectric is stressed at -50 V, well above its device rating of -8.5 V for over 3,000 hours at 150ºC without deleterious effect. It is imperative that such catastrophic failure mechanisms such as the “inverse piezo-electric effect”, found

Continued inside of back wrap

UBM Electronics

EDN®

VOICE OF THE ENGINEER

MAR 3
Issue 5/2011
www.edn.com

CPUs IN FPGAs:
MANY FACES TO A TREND
Page 30

BUILD ACCURATE SPICE MODELS
FOR LOW-NOISE, LOW-POWER PRECISION AMPLIFIERS
Page 20

MONITOR PWM LOAD CURRENT
WITH A HIGH-SIDE CURRENT-SENSE AMPLIFIER
Page 39

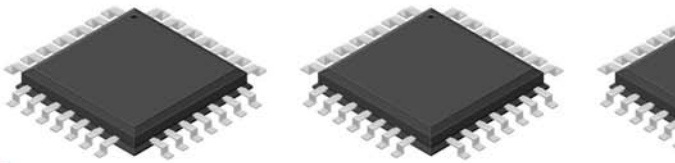

Data acquisition just got a lot easier.

More ways to control. More ways to connect.

The new Agilent 34972A Data Acquisition Switch Unit takes our best-selling Agilent 34970A to the next level. For starters, you get convenient built-in LAN and USB connectivity. Plus, you can control your data acquisition remotely via Web interface. And transfer logged data to your PC with a simple flash drive. No more expensive adapters and connectors. That's easy. That's Agilent.

NEW 34972A	34970A
USB and LAN	GPIB and RS232
Graphical web interface Benchlink data logger software SCPI programming	Benchlink data logger software SCPI programming
3-slot LXI unit with built-in 6 ½ digit DMM	
$1,845	$1,597*

*Prices are in USD and are subject to change.

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

Buy from an Authorized Distributor
www.agilent.com/find/distributors

Learn how to connect wirelessly and get more FREE measurement tips at www.agilent.com/find/Agilent34972A

Agilent Technologies

YOU NEED IT, WE HAVE IT

Digi-Key® CORPORATION

Over 515,000 Products in Stock

More Than 45,000 NEW Products Added in the Last 90 Days

440+ Suppliers

The industry's broadest product selection available for immediate delivery

www.digikey.com

1.800.344.4539

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

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

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

Applications

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

And more...

Benefits

- Instant-on, non-volatile
- Flexible logic from 256 LUTs to 7,000 LUTs
- Embedded memory up to 240 Kbits
- Ultra low power as low as 19µW
- On-chip user flash memory & I²C, SPI, timer/counter
- Prices starting as low as $0.75

And more...

Lattice® Semiconductor Corporation

latticesemi.com/machxo2

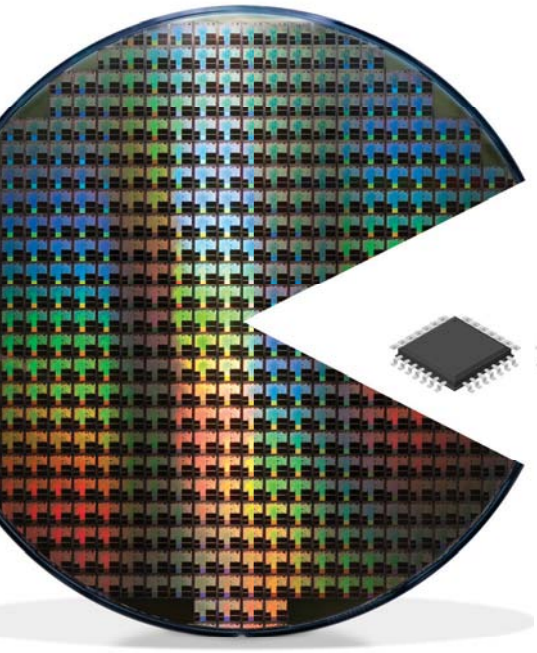

EDN contents

3.3.11

CPUs in FPGAs: many faces to a trend

by Ron Wilson, Editorial Director

Build accurate Spice models for low-noise, low-power precision amplifiers

by Don LaFontaine, Intersil

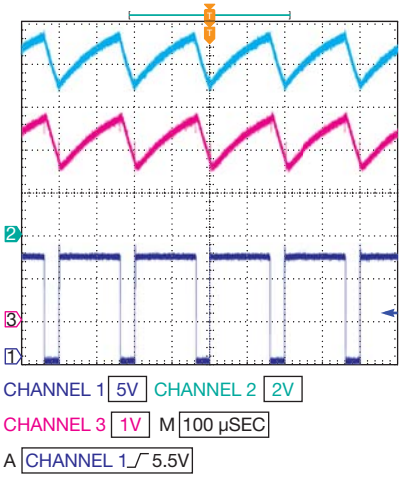

Monitor PWM load current with a high-side current-sense amplifier

by Maurizio Gavardoni and Akshay Bhat, Maxim Integrated Products

pulse

COVER: PHOTO-ILLUSTRATION BY TIM BURNS.
WAFER: PHOTOMICK/ISTOCKPHOTO.COM;
CHIPS: GRIMGRAM/ISTOCKPHOTO.COM

DESIGNIDEAS

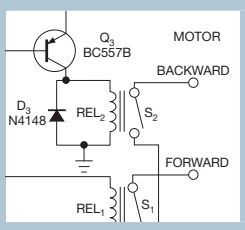

Smart Building

HVAC
(Heating, Ventilating, and Air Conditioning)

- Programmable Room Sensor
- Central Heating Unit
- Programmable Radiator Thermostat
- Burner Control Unit
- Solar Control Unit
- Air Conditioning System
- Hot Water Dispenser

Metering

- Smart Meter
- Heat Meter
- Heating Cost Meter
- Water Meter/Hot Water Meter/ Gas Meter/Pre-Paid Meter
- Power Meter (Single-Phase)
- Power Meter (Three-Phase)
- Water Meter
- Gas Meter

Network Infrastructure

ZigBee, Wireless LAN, PLC (Power Line Communication), Wi MAX, Ethernet, Optical Fiber Network, Network Server

Safety and Security

- Detector
- Smoke Detector
- Sounders and Strobes
- Access Control System
- Automatic Fire Door
- Alarm Control Panel
- Emergency Lighting Control
- Surveillance Camera
- Smart Cards
- Camera Control
- Gas Leak Alarms

Lighting

- LED Lighting
- Blind Control
- Emergency Lighting Control
- Electronic Timer
- Motion Detector
- Central Control Panel
- Electronic Ballast, Neon Lighting

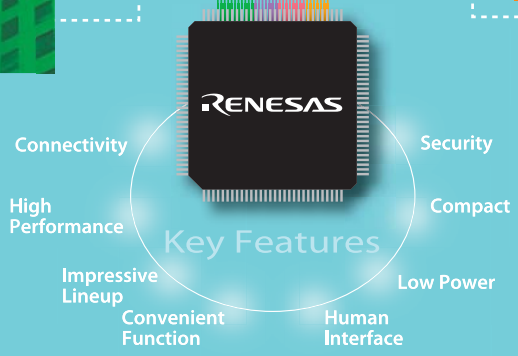

Get started now
for smart development at
www.renesas.com/smart

Renesas Electronics Corporation

Green Solutions ready for you

RENESAS

contents 3.3.11

DEPARTMENTS & COLUMNS

EDN online contents www.edn.com

ONLINE ONLY

Check out this Web-exclusive article:

Fully depleted SOI shows its stuff in CPU design

An ARM Cortex M0 paper design suggests that FDSOI (fully depleted silicon on chip) could be a strong contender at 20 nm.
→www.edn.com/110303toca

JOIN THE CONVERSATION

Have you joined *EDN*'s Electronics Design Network on LinkedIn? If not, you're missing out on some great conversation. Here's what one group member has to say about *EDN*'s LinkedIn community: "It's really nice to see smart people who are passionate about their work go at it. You guys are an inspiration. You make *EDN* one of the better groups for engineers on LinkedIn."

Join the more than 2700 engineers and executives exchanging ideas in our group.
→www.edn.com/110303tocb

EDN® (ISSN#0012-7515) is published semimonthly, 24 times per year, by UBM Electronics, 11444 W. Olympic Blvd., Los Angeles, CA 90064-1549; 310/445-4200; FAX 310/445-4299. Periodicals postage paid at Los Angeles, CA, and at additional mailing offices. SUBSCRIPTIONS—Free to qualified subscribers as defined on the subscription card. Rates for nonqualified subscriptions, including all issues: US, $150 one year; $250 two years; $300 three years. Except for special issues where price changes are indicated, single copies are available for $10 US and $15 foreign. For telephone inquiries regarding subscriptions, call 763/746-2792. E-mail: EDN@kmpsgroup.com. CHANGE OF ADDRESS—Notices should be sent promptly to PO Box 47461, Plymouth, MN 55447. Please provide old mailing label as well as new address. Allow two months for change. NOTICE—Every precaution is taken to ensure accuracy of content; however, the publisher cannot accept responsibility for the correctness of the information supplied or advertised or for any opinion expressed herein. POSTMASTER—Send address changes to EDN, PO Box 47461, Plymouth, MN 55447. Canada Post: Publications Mail Agreement 40612608. Return undeliverable Canadian addresses to Pitney Bowes Inc, PO Box 25542, London, ON N6C 6B2.

Volume 56, Number 5 (Printed in USA).

BY BRIAN DIPERT, SENIOR TECHNICAL EDITOR

Windows on ARM: for Intel, probably no cause for alarm

During Microsoft's January International CES (Consumer Electronics Show) keynote address, Steve Ballmer, the company's chief executive officer, and Michael Angiulo, corporate vice president of its planning, hardware, and PC-ecosystem team, announced that the next-generation Windows OS would run on both x86 and ARM CPU architectures. In a sense, this announcement was no surprise. I disagree, however, with the pundits who reacted to the announcement by proclaiming that Intel—and AMD, and don't forget Via—were now dead. Many of these "experts" inaccurately drew parallels between Microsoft's announcement and Apple's several-year-old transition from PowerPC to x86 for OS X and subsequently to ARM for the OS X-derived iOS. In both Apple case studies, a wholesale migration from one CPU architecture to another took place, whereas the Microsoft case involves an expansion from one architecture to two.

Why does this distinction matter? Consider consumer desktop and mobile systems, which account for most of the x86 sales volume each year. The fundamental hardware costs of ARM- and x86-based systems are comparable if not identical to each other. Assuming that cellular/broadband carrier or other service-provider subsidies don't factor into the equation, prices for the two CPUs' respective systems will be similar, as well.

As a consumer, why on earth would I choose an ARM-based variant with limited native-software support versus an x86-based alternative that continues to tap into an extensive backward-compatible code ecosystem? And, if I were a developer, why would I bother compiling and supporting an ARM-flavored version of my app or, for that matter, a "universal" variant for the limited ARM customer base? ARM's situation shows little likelihood of success in today's consumer—or enterprise—desktop- and mobile-client scenarios.

Microsoft went to great pains to point out that the code running on ARM-based hardware during the CES keynote demos was binary-compiled, not virtualized. But Ballmer and Angiulo also went to great pains to reiterate that AMD and Intel will continue as strong partners of Microsoft. Microsoft could provide dynamic binary-translation capabilities in Windows 8 for running x86 binaries on ARM hardware.

Such work-arounds are inherently imperfect, however, so why is Microsoft bothering with the ARM port at all? Servers, which are sensitive to power consumption and relatively insensitive to performance, are a possible reason. The required code support is less in this case than in the consumer market, and current draw and heat dissipation are ongoing concerns. But the enterprise is a conservative and slow-moving application area.

Looking at the suppliers' road maps, you might wonder why Microsoft is bothering to break away from its x86-exclusive embrace. I've long said that Intel won't be able to seriously compete with ARM on ARM's own turf until at least the third turn of the Atom tick-tock cadence. That prediction is coming to pass right on schedule. First-generation Atom on a 45-nm process essentially created the low-cost notebook—that is, the netbook. Second-generation Atom, also on a 45-nm process, has established a beachhead in tablets. Intel's now-mature 32-nm process is the basis for the upcoming third-generation Atom silicon spin. The next-generation 22-nm process will also inevitably yield highly integrated and power-sensitive Atom-based products. Meanwhile, AMD is translating its Fusion aspirations into product reality, and Via has migrated its dual-core Nano CPU into 40-nm-based monolithic-die form. By the time Windows 8 appears, the low-power x86 story should be in solid shape.

Perhaps, then, the more accurate question to ask is, Why is Microsoft breaking away from its ARM-exclusive embrace? With a nod to Tony King-Smith, vice president of marketing at Imagination Technologies, who gave me this idea, I believe a code-base simplification move is under way: the eventual obsolescence of Windows CE in favor of a unilateral corporate focus on modularized Windows 8 across all product segments in which Microsoft plays.

So, consider this potential future scenario: Microsoft broadens Windows 8 beyond x86 to ARM to catch the eye of mobile-electronics-equipment developers that might otherwise standardize on some other operating system and application suite. By the time Microsoft releases Windows 8, x86 CPU suppliers will be ready with silicon that's good enough from performance and power-consumption standpoints versus ARM alternatives. And developers go with x86 instead of the ARM processors they'd previously used. Wouldn't that be ironic? EDN

Read an expanded version of this column on the Brian's Brain blog at www.edn.com/110303eda.

EDN

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

EDITORIAL DIRECTOR
Ron Wilson
1-415-947-6317;
ron.wilson@ubm.com

MANAGING EDITOR
Amy Norcross
Contributed technical articles
1-781-869-7971;
amy.norcross@ubm.com

MANAGING EDITOR—NEWS
Suzanne Deffree
Electronic Business, Distribution
1-631-266-3433;
suzanne.deffree@ubm.com

SENIOR TECHNICAL EDITOR
Brian Dipert
Consumer Electronics, Multimedia, PCs, Mass Storage
1-916-548-1225;
brian.dipert@ubm.com

TECHNICAL EDITOR
Margery Conner
Power Sources, Components, Green Engineering
1-805-461-8242;
margery.conner@ubm.com

TECHNICAL EDITOR
Mike Demler
EDA, IC Design and Application
1-408-384-8336;
mike.demler@ubm.com

TECHNICAL EDITOR
Paul Rako
Analog, RF, PCB Design
1-408-745-1994;
paul.rako@ubm.com

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

SENIOR ASSOCIATE EDITOR
Frances T Granville, 1-781-869-7969;
frances.granville@ubm.com

ASSOCIATE EDITOR
Jessica MacNeil, 1-781-869-7983;
jessica.macneil@ubm.com

CONSULTING EDITOR
Jim Williams,
Staff Scientist, Linear Technology
edn.editor@ubm.com

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

COLUMNISTS
Howard Johnson, PhD, Signal Consulting
Bonnie Baker, Texas Instruments
Pallab Chatterjee, SiliconMap
Kevin C Craig, PhD, Marquette University

LEAD ART DIRECTOR
Marco Aguilera

ASSOCIATE ART DIRECTOR
Tim Burns

PRODUCTION
Jeff Tade,
Production Director
Brian Wu,
Publications Production Manager
Jeff Polman, Derric Treece,
Senior Production Artists
William Baughman, Lucia Corona,
Ricardo Esparza,
Production Artists

EDN EUROPE
Graham Prophet,
Editor, Reed Publishing
gprophet@reedbusiness.fr

EDN ASIA
Wai-Chun Chen,
Group Publisher, Asia
waichun.chen@ubm.com
Kirtimaya Varma,
Editor-in-Chief
kirti.varma@ubm.com

EDN CHINA
William Zhang,
Publisher and Editorial Director
william.zhang@ubm.com
Jeff Lu,
Executive Editor
jeff.lu@ubm.com

EDN JAPAN
Katsuya Watanabe,
Publisher
katsuya.watanabe@ubm.com
Ken Amemoto,
Editor-in-Chief
ken.amemoto@ubm.com

EXECUTIVE OFFICERS
Paul Miller,
Chief Executive Officer
David Blaza,
Vice President, UBM Electronics
Karen Field,
Senior Vice President, Content
Fred Gysi,
Chief Financial Officer
Mike Deering,
Chief Technology Officer
Stephen Corrick,
Vice President/Executive Director
Kevin O'Keefe,
Senior Vice President, Events Division
Roger Burg,
Vice President, Operations
Jason Brown,
Vice President, E-media

UBM Electronics

EDN, 33 Hayden Avenue, Lexington, MA 02421. **www.edn.com**. **Subscription inquiries:** 1-763-746-2792; EDN@kmpsgroup.com. **Address changes:** Send notice promptly to PO Box 47461, Plymouth, MN 55447. Please provide an old mailing label as well as your new address. Allow two months for the change.

UBM Electronics, 11444 W. Olympic Blvd., Los Angeles, CA 90064-1549; 1-310-445-4200; fax: 1-310-445-4299.

High-performance DSP-IP cores are ready for LTE-Advanced

Tensilica's new ConnX BBE64-128 DSP-IP (intellectual-property) cores for SOC (system-on-chip) design extend the BBE (baseband-engine) family to more than 100 billion MAC (multiply/accumulate)-operations-per-sec performance to meet requirements for the emerging LTE (long-term evolution)-Advanced standard. Tensilica manufactures the ConnX BBE64-128 DSP core for LTE-Advanced in a high-performance, 28-nm process technology to deliver more than five times the required computing power of LTE. The LTE-Advanced specification requires an increase in peak data rate from 150 Mbps to 1 Gbps, with an increase in channel bandwidth from 20 MHz to 100 MHz.

Tensilica based the BBE64-128 on a multi-slot VLIW (very-long-instruction-word) architecture that enables 128 MAC operations per cycle for maximum throughput and minimum energy consumption. It also performs common MIMO (multiple-in/multiple-out) functions, FIR (finite-impulse-response) filtering, and channel-estimation functions for LTE-Advanced. Modeless switch to Tensilica's smaller standard 16- and 24-bit instructions enables high code density for nonvector algorithms. The BBE64-128 also includes a single-cycle, 16-way complex radix-4 and radix-8 FFT (fast Fourier transform) and DFT (discrete Fourier transform) for efficiency on arbitrarily sized transformations in LTE OFDM (orthogonal-frequency-division-multiplexing) algorithms.

The company based the ConnX BBE64-UE for LTE-Advanced handset applications on a smaller feature set to minimize energy consumption and latency. It excludes some of the features of the BBE64-128, such as the option to run 128 MAC operations/cycle, to achieve higher efficiency. The BBE64-UE processor can reach approximately 300,000 GMACs/sec/W, and Tensilica manufactures the core in a low-leakage, 28-nm-process technology.

With Tensilica's automated tools, you can optimize your design for specific applications by adding custom instructions, and you can add special memory interfaces, special per-SIMD (single-instruction/multiple-data) lane look-ups, or other required functions. The customizable processor architecture allows you to tailor functions, turn them on or off, and add them during SOC design.

Tensilica's design process automatically generates the compiler for the BBE64 core to match the configuration options you choose and features full native DSP-data support, including integer, fractional, real, and complex. It automatically infers complex instructions, accelerates and vectorizes legacy code from ConnX BBE16, accelerates legacy code for industry-standard intrinsic functions, vectorizes loops with complex conditional operations, and performs ANSI C operators on vector-data types. An "analysis cockpit" provides for program analysis, including a vectorization assistant.

A complete evaluation kit for the ConnX BBE64-128 and ConnX BBE-UE cores should become available for early-access customers in the fall.

—by Mike Demler

▷**Tensilica**, www.tensilica.com.

TALKBACK

"I have a big-brother version that I have had for over 20 years. I opened it once because of a strange smell. That poor mouse."

—Reader Dick Kolman, in *EDN*'s Talkback section, at http://bit.ly/fh5HGW. Add your comments.

The Tensilica ConnX BBE64 DSP-IP cores support LTE-Advanced infrastructure and handset applications.

Generator produces pulses to 330 MHz; functions, noise to 500 MHz

Agilent Technologies has expanded its pulse/function/arbitrary-noise-waveform generator family to help design and test engineers improve the efficiency and precision of testing higher-speed, higher-bandwidth, analog-, digital-, and mixed-signal devices. Engineers must improve quality yet must get to market faster with products whose features differentiate them from those of competitors. These conflicting requirements necessitate expanded test capabilities during development. The Agilent 81160A provides innovative functions and streamlined setup to help you more quickly complete expanded test protocols.

The generator eliminates the need for cumbersome multi-instrument setups for stress-testing devices. The 81160A provides versatile waveforms as well as signals with an intrinsic jitter of 7 psec rms. The instrument targets use in general-purpose bench tests and advanced serial-data stress tests. It generates 330-MHz pulses, optionally generates pulse patterns, and generates 500-MHz functions and arbitrary-noise waveforms with a 2.5G-sample/sec data rate and 14-bit vertical resolution. Selectable crest factors for white gaussian noise allow you to determine how much distortion to apply to a device during stress testing to meet various serial-bus standards. Glitch-free timing-parameter adjustments let you change the frequency without dropouts or glitches and enable continuous operation without rebooting or resetting the device under test. Arbitrary bit patterns allow simple pattern settings to emulate capacitive loading of the channels and eliminate the need for complex measurement setups to test designs to their limits. Noise waveforms' maximum duration—that is, the period during which patterns do not repeat—is 20 days, yet you can exactly and repeatedly reproduce brief portions of such waveforms to quickly investigate unusual behavior of the system under test.

The 14-bit-resolution, 2.5G-sample/sec-data-rate 81160A extends the frequency range of the manufacturer's pulse/function/arbitrary-noise-waveform generator family to 330 MHz for pulses and 500 MHz for functions and noise.

The suggested US entry price for a single-channel version of the 81160A 330-MHz (pulse)/500-MHz (function/arbitrary-noise) generator is $16,900. A two-channel version costs $20,000.

—by Dan Strassberg

▷**Agilent Technologies**, www.agilent.com/find/81160.

Analog, NI team up on simulation tool

Analog Devices Inc and National Instruments recently announced their collaboration on a new release of NI's Multisim component-evaluation tool with added features to provide engineers with an easy-to-use environment for the simulation of linear circuits using Analog Devices' components. The free component-evaluation tool is available on Analog Devices' Web site.

This edition of the tool allows engineers to design larger, more complex circuits and easily import their own models into the tool. By matching more than 1000 of Analog Devices' amplifiers, switches, and voltage references to more than 550 of National Instruments' simulation models, designers have free access to the simulation environment, allowing them to easily experiment with circuit designs and reduce system-development time and cost.

"Finding the right solution for a linear circuit often means sorting through many products and specifications," says Steve Sockolov, director of the Precision Signal Conditioning Group at Analog Devices. "We offer innovative ways to help customers quickly find what they need."

The NI Multisim Component Evaluator—Analog Devices Edition provides engineers the ability to build simulatable circuits to evaluate a library of Analog Devices' operational amplifiers, switches, and voltage references. It also provides Spice-parser improvements, updated BSIMs (Berkeley Short-Channel IGFET Models), support for advanced parameters and enhanced digital-simulation accuracy, improved design communication with on-page connectors, and a new WYSIWYG net-naming system.

—by Paul Rako

▷**Analog Devices**, www.analog.com.

▷**National Instruments**, www.ni.com.

DILBERT By Scott Adams

Lighting platform brings building-block approach to LED lighting

Some companies are chasing after the retrofit market of LEDs that fit into incandescent lighting systems, but Tyco Electronics instead targets the next generation of LED-lighting systems, designing them from the ground up to exploit the advantages of solid-state lighting. With its new Nevalo lighting system, Tyco is betting that a modular approach will prove to be the winning one.

By concentrating on new lighting wins, the Nevalo system need not fit into current form factors of light bulbs, luminaires, or tubes and instead introduces a building-block approach to lighting design. The Nevalo platform includes light modules, drivers, optics, wiring systems, and software tools to evaluate thermal management and price and performance budgeting.

The Nevalo system includes more than 60 LED-light modules ranging from 300 to 3400 lumens in form factors for common lighting applications; optics in total internal reflection and reflector styles; drivers with constant-current output, dimming-control capabilities, and temperature monitoring; and a new ribbon-based, four-wire wiring-configuration system keyed and color-coded for ease of manufacturing. The system includes circuit-protection devices and heat sinks that match the LED-light modules to help provide appropriate thermal management and life expectancy.

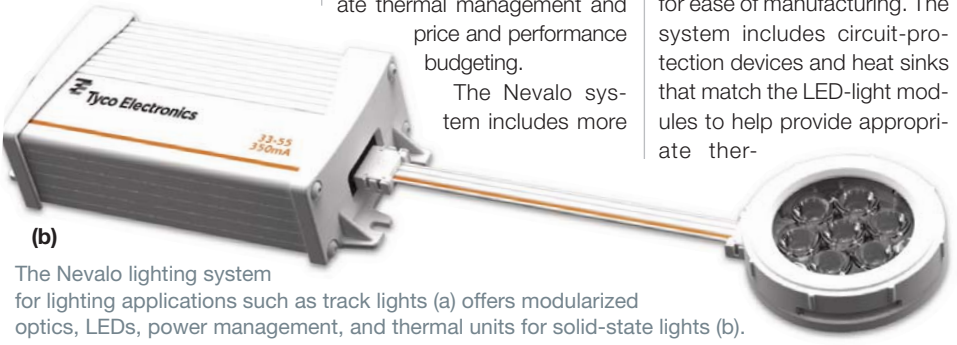

The Nevalo lighting system for lighting applications such as track lights (a) offers modularized optics, LEDs, power management, and thermal units for solid-state lights (b).

The LED-light-module subsystem comprises a PCB (printed-circuit board), including LEDs, thermal-interface material, a socket for circular applications, and matching optics. By mixing and matching the components, designers can create down lights, sconces, and track-lighting equivalents. Prices for system-development kits range from $249 to $399.

—by Margery Conner

▷**Nevalo**, www.nevalo.com.

Molybdenite provides advantages over silicon, graphene

Researchers at Switzerland's EPFL (École Polytechnique Fédérale de Lausanne) are touting molybdenite, or MoS_2, for use in electronics (Reference 1). Using the material would enable manufacturers to develop smaller and more energy-efficient electronic chips. The mineral is abundant in nature, but no one had previously studied it for use in electronics.

"It's a 2-D, very thin material and easy to use in nanotechnology," says Andras Kis, a professor at EPFL. "It has real potential in the fabrication of very small transistors, LEDs, and solar cells." Molybdenite is also less voluminous than silicon, a 3-D material. "In a 0.65-nm-thick sheet of molybdenite, the electrons can move around as easily as in a 2-nm-thick sheet of silicon," Kis says. "But it's not currently possible to fabricate a sheet of silicon as thin as a monolayer sheet of molybdenite."

Researchers also claim that molybdenite can yield transistors that consume 100,000 times less energy in a standby state than traditional silicon transistors. Manufacturers must use a semiconductor with a gap to turn a transistor on and off, and molybdenite's 1.8-eV (electron-volt) gap is ideal for this purpose. This gap also gives it an advantage over graphene, which lacks this gap.—by Suzanne Deffree

▶École Polytechnique Fédérale de Lausanne, www.epfl.ch.

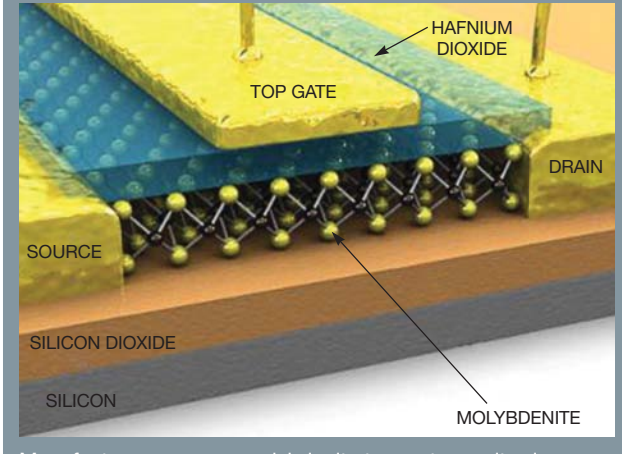

Manufacturers can use molybdenite to create an ultra-low-power FET by acting as its channel on a silicon-on-insulator substrate.

REFERENCE

1 **Radisavljevic, B, et al, "Single-layer MoS_2 transistors," *Nature Nanotechnology*, January 2011, http://bit.ly/ih8lrR.**

Processor touts 3G HSDPA, Android

Broadcom's new BCM2157 HSDPA (high-speed-downlink-packet-access) baseband processor integrates simultaneous 3G (third-generation) HSDPA connectivity at downstream rates as fast as 7.2 Mbps with advanced smartphone features for mass-market mobile handsets. The BCM2157 modem also supports WCDMA (wideband-code-division-multiple-access) and EDGE (enhanced data rates for global-system-for-mobile-communications-evolution) networks.

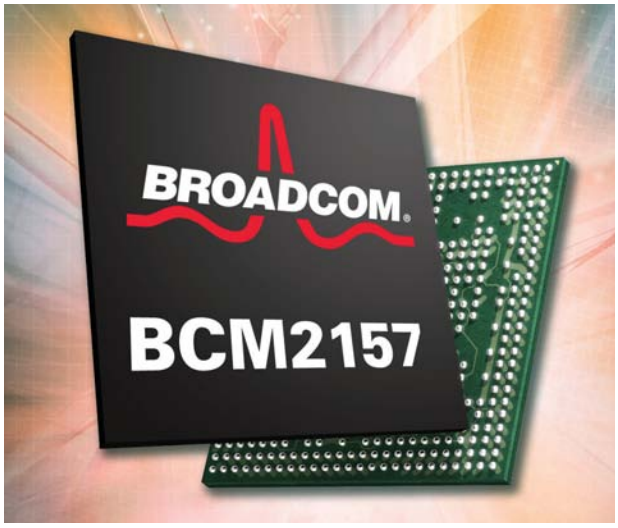

The BCM2157 integrates a baseband processor for 3G HSDPA cellular communications with an application processor for low-cost Android smartphones.

> Support for mobile-hot-spot functions enables handsets to share a 3G connection with eight devices.

The dual-core device employs a dedicated, 500-MHz ARM11 processor to support multitouch, advanced multimedia technology and video applications with resolution as fine as 432×320-pixel HVGA (half-size video-graphics array), a 5M-pixel camera, and the ability to encode and decode H.264 video at 30 frames/sec.

Broadcom manufactures the BCM2157 in a low-power, 65-nm CMOS process. Support for mobile-hot-spot functions enables handsets to share a 3G connection with as many as eight simultaneous devices or users through Wi-Fi. Two SDIO (secure-digital-input/output) ports provide for connections to off-chip wireless-LAN transceivers and SD (secure-digital) MMC (multimedia cards).

The device also provides a USB (Universal Serial Bus) 2.0 port, along with audio amplifiers for driving 100-mW headsets, dual digital microphones, and a built-in five-band audio equalizer. The audio DSP supports Broadcom's M-Stream technology with the GSM (global-system-for-mobile)-communications SAIC (single-antenna-interference-canceling) standard. An on-chip PWM (pulse-width modulator) drives a vibrator to notify users of received calls.

The BCM2157 modem supports 3G dual SIM (subscriber-identification-module)/dual-standby technology and AMR-WB (adaptive multirate wideband) for HD (high-definition) voice calls over 3G networks. Dual-SIM is not in the Android standard, but it is a growing requirement in parts of the world that lack number portability. With two SIM cards, subscribers can migrate between carriers and receive calls on different numbers.

The BCM2157 supports Android version 2.2 Froyo but not the Adobe Flash-player 10.1 feature of that version of Google's operating system, which requires more advanced processors. The device also integrates Broadcom's proprietary InConcert technology, which supports interoperability of devices in the 2.4-GHz frequency range, coordinating transmissions to optimize throughput in Bluetooth, Wi-Fi, GPS (global-positioning-system), and NFC (near-field-communications) applications. Digital-audio ports include I^2S (inter-integrated-circuit sound) and PCM (pulse-code modulation), with a UART (universal asynchronous receiver/transmitter) to interface to transceivers for Bluetooth earpieces or FM-band radios. A BSC (Broadcom serial-control) bus connects to an off-chip PMU (power-management unit).

The BCM2157 platform employs technology from the earlier dual-core BCM2153 architecture, which targets Symbian (www.symbian.com) and Windows Mobile operating systems and finds use in Samsung's (www.samsung.com) bada handsets.

The device is now available for sampling to early-access customers, and the company plans its first commercial launch for this quarter.

—by Mike Demler

▷**Broadcom**, www.broadcom.com.

MICREL INTRODUCES POL BUCK REGULATORS

03.03.11

Micrel recently introduced the MIC26400, MIC26600, and MIC26950 POL (point-of-load) synchronous buck regulators, all of which feature digitally modified adaptive on-time-control architectures and operate over an input-supply range of 4.5 to 26V and an adjustable output voltage as low as 0.8V. The devices employ Micrel's Hyper Speed Control architecture, which allows for fast transient response, reduces output capacitance, and enables high-input-voltage and low-output-voltage operation. They also include undervoltage lockout to ensure proper operation under power-sag conditions, internal soft-start to reduce inrush current, foldback-current limit, "hiccup"-mode short-circuit protection, and thermal shutdown. Typical accuracy is ±1%, and the devices operate at a switching frequency of 300 kHz. Prices for the MIC26400, MIC26600, and MIC26950 start at $1.82, $2.20, and $2.77 (1000), respectively.—by Paul Rako

▶**Micrel, www.micrel.com.**

BY HOWARD JOHNSON, PhD

Field cancellation

Figures 1, **2**, and **3** illustrate the pattern of magnetic fields surrounding one microstrip trace on a typical PCB (printed-circuit board). The trace appears in cross-sectional view as a small black rectangle. The line below the trace represents a solid plane layer within the board, which may contain other signal layers not shown. Current on the trace flows into the page in all three cases.

Figure 1 shows the pattern of magnetic lines of force resulting from the current on the trace, without considering current on the reference plane. This pattern of lines looks like a nested set of concentric circles. As you move to a distance *r* away from the trace, the field intensity falls off as 1/*r*. An infinite straight wire suspended in free space generates this same familiar pattern.

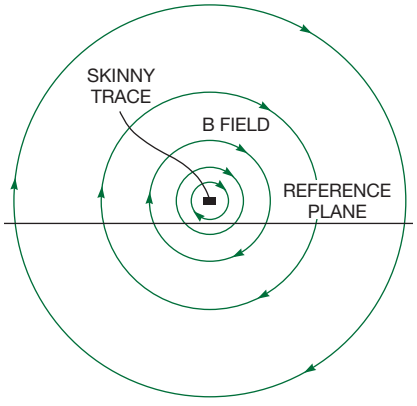

Figure 1 One current by itself creates magnetic fields in simple concentric patterns.

In a singled-ended system, outgoing current on the trace stimulates an equal and opposite current returning to its source through the power and ground system. **Figure 2** depicts the magnetic field resulting from that current. At frequencies higher than approximately 1 MHz on a typical PCB, current returning to its source flows along the top surface of the solid reference plane, parallel to the outbound trace. The current bunches strongly right underneath the signal trace, with a distribution that falls off rapidly as you move away from the trace on either side. The pattern of magnetic fields from this distribution makes football-like shapes. Where each line of force intercepts the solid plane, it "bends" in proportion to the profile of current flowing in the sheet at that location. These lines of force circulate in a direction opposite those in **Figure 1**.

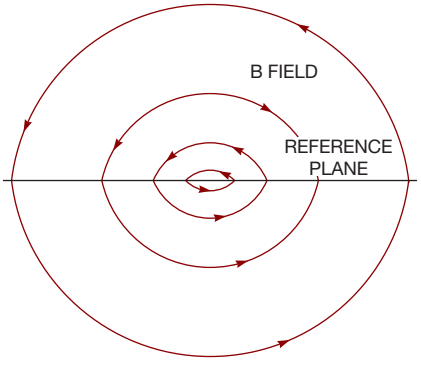

Figure 2 Current on the reference plane creates its own pattern of magnetic fields.

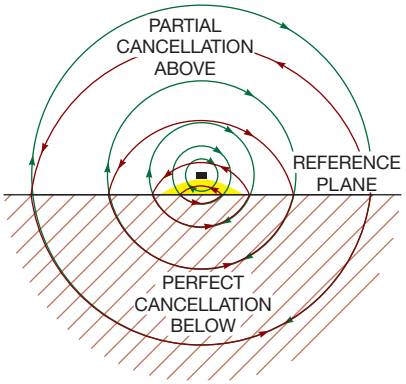

Figure 3 The two field patterns cancel perfectly in the region below the plane.

When you superimpose the two patterns, the fields in the bottom half of the diagram cancel perfectly, assuming a reference plane of infinite extent and perfect conductivity (**Figure 3**). The field intensity above the trace does not cancel perfectly, but it substantially cancels, vastly reducing both crosstalk and EMI (electromagnetic interference) on that side. The field intensity increases only in the small yellow region directly under the trace, at which point it approximately doubles. The net benefit is fantastic. There is no simpler or more effective means of combatting crosstalk and EMI than a solid reference plane.

I love the field-cancellation benefit, but who knew that fields do what they do so well? James Clerk Maxwell, in his seminal text, *A Treatise on Electricity and Magnetism*, published in 1891 by Clarendon Press, explains in volume two, article 654, that current flows in the solid reference plane in reaction to the current flowing in the trace. The reference-plane currents, called eddy currents, form in a pattern that neutralizes the component of the magnetic field normal to the reference plane. The solution in **Figure 3** attains that goal, perfectly canceling the vertical component of the magnetic field at the surface of the solid reference plane. Were some portion of the field to behave differently, it would set up eddy currents in the plane that would counteract the miscreant field, forcing it back into the shape shown in **Figure 3**.EDN

Howard Johnson, PhD, frequently conducts technical workshops for digital engineers. Visit his Web site at www.sigcon.com.

Twice the Power
Half the Size

Visit us at APEC Booth #108

New family of high-density power components enables fast, adaptable designs

V·I CHIP FACTORIZED POWER

Example Application

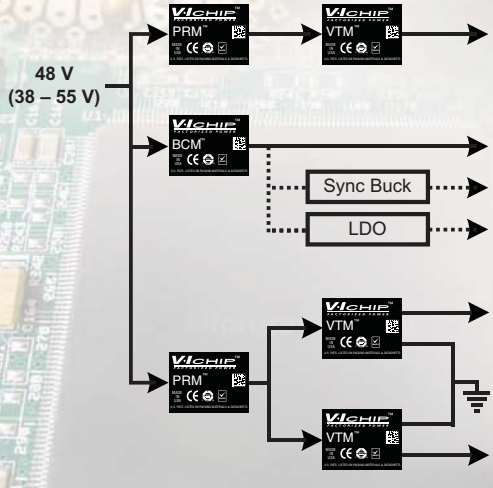

Half-Chip	V_{IN} (V)	V_{OUT} (V)	I_{OUT} (A)	P_{OUT} (W)
BCM	48	12	10	120
PRM	48	5-55	4	200
VTM	48	12	10	135
		6	20	100
		4	25	115
		2	40	88.5
		1.5	50	80

- 1/32 brick footprint
- Up to 97% efficiency, up to 1300 W/in³
- 200W PRM™ Regulator
- 50A VTM™ Current Multiplier
- 120W BCM™ Bus Converter

Half-size V•I Chip modules increase design flexibility for configurable, high-density, low-profile power system solutions and effectively address lower power applications with better efficiency and state-of-the-art performance.

Visit
www.vicorpower.com/rp/half-chip_edn
for more information
and to order samples online.

All parts available from stock NOW!

800-735-6200 **vicorpower.com/rp/half-chip_edn** VICOR

BY PALLAB CHATTERJEE, CONTRIBUTING TECHNICAL EDITOR

New devices for nanoelectronics

Nanoelectronics holds the promise of advanced capabilities for computing, RF and communications, and memories. One of the challenges in exploiting these capabilities is being able to manufacture, model, and test these devices. Manufacturers are creating these new devices not only in smaller versions of traditional devices, as in the standard semiconductor industry for CMOS, but also in new device construction and device materials.

Many of these devices use graphene as a key material in gates, epitaxial layers, and interconnects. The challenge with using graphene as a switch or a transistor material is that it lacks the bandgap that normally occurs in CMOS processes. As a result, manufacturers have created new structures, such as doped "dual"-gate devices. These devices provide the high-performance transporting and conduction aspects of the graphene and still allow for use in common-architecture circuit design. Manufacturers construct the dual-gate devices both vertically with a top gate and a back gate beneath the graphene layer or laterally with adjacent gates. The manufacturers can fabricate these devices with modifications of standard CMOS processes (**Figure 1**). These devices offer performance well beyond that of standard CMOS devices; experimental data from IBM (www.ibm.com) shows a typical device frequency as high as 170 GHz.

Carbon nanotubes are also making headway in circuits, finding use both in high-performance conductance channels as enhancements to standard CMOS and in interconnect. Vias and contacts use the nanotubes for interconnect to provide low-resistance, high-current pathways. Vendors form the contacts and vias using an alloy of titanium/lead/gold/carbon nanotubes.

For communications, the high speed available in nanotechnology is driving advancement in optical interconnect and devices. Recent advances in zinc-oxide-silicon radial heterojunctions have produced nanorods and nanopillars that exhibit dramatic improvements over standard photodetectors. The increased vertical surface area for the vertical nanorod demonstrates a high gain in photoconductivity due to the enhanced surface electron-hole separation. This gain leads to an approximately 2.5-times-higher responsivity to ultraviolet light over a planar structure of the same thickness.

On the memory side, a number of new technologies are in progress. These nanotechnology memories include PCM (phase-change memory), RRAM (resistive random-access memory), MTJ (magnetic-tunnel junction), MRAM (magnetoresistive random-access memory), STT (spin-transfer-torque), MQCA (magnetic-quantum cellular automa), and traditional flash devices. All of these nonvolatile memory technologies are working on retention and low-power read and write access with error-correction schemes. These characteristics, including the ability to read, program, and rewrite the state of cell at high density, are driving the new memory investigations. These technologies all focus on manipulation of magnetic or material properties of the storage cell. With nanomanufacturing, these cells are now small enough to manipulate and large enough to let you know whether programming has occurred.

These technologies are trying to surpass the density and power of the NAND-flash cell. The cell size for NAND flash, now down to 0.0028 square microns for a 25-nm process, is having scaling issues because the cell stores so few electrons as a delta between states. As a result, there will soon be no measurable difference in the states. The scaling has progressed to the point at which, on a 25-nm process, a 5% variation in cell delta, due to line edge roughness, is approximately three silicon lattices. The current state of NAND-flash technology is 64-Gbit MLC (multilevel-cell) design. EDN

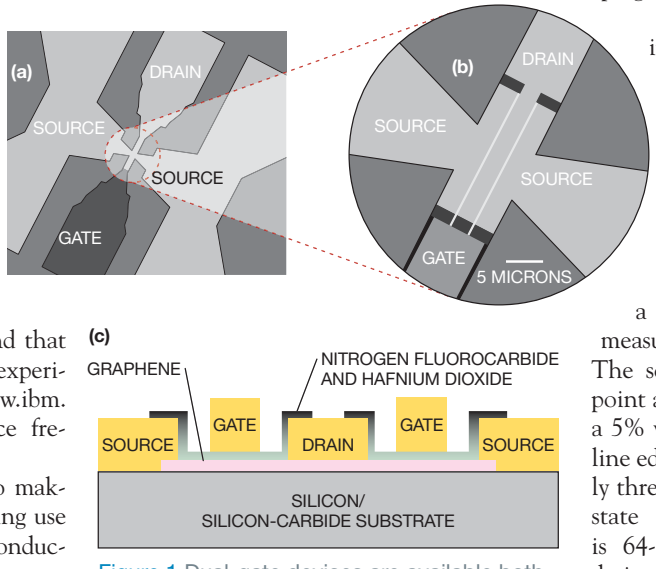

Figure 1 Dual-gate devices are available both vertically and laterally (a and b). A top-gated FET (c) employs graphene (courtesy IBM).

Pallab Chatterjee is on the IEEE Nanotechnology Council.

Power Density AND Performance — Get the Best of Both Worlds with SuperSwitcher II™

Family of DC-DC Buck Regulators Saves Design Time and Cost

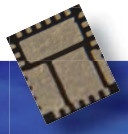

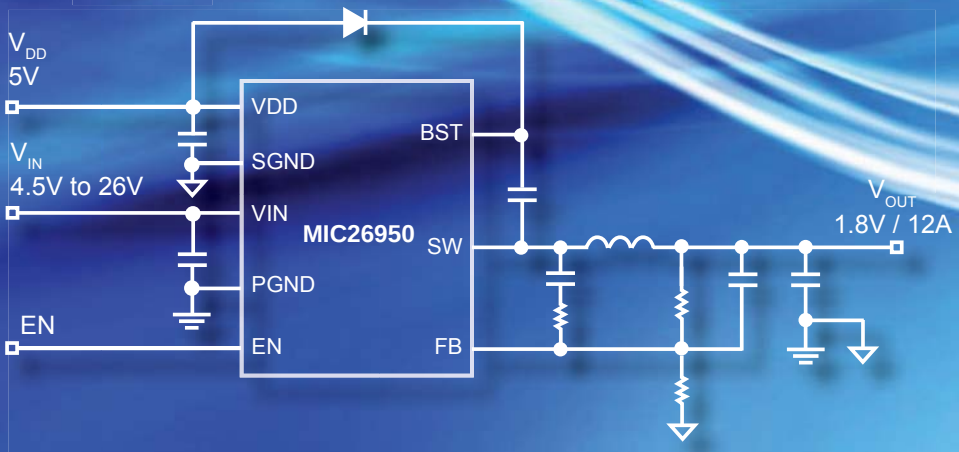

Part Number	Input Voltage	Output Voltage	Output Current	Package
MIC26400	4.5V to 26V	0.8V to 5.5V	5A	MLF®-28L
MIC26600	4.5V to 26V	0.8V to 5.5V	7A	MLF®-28L
MIC26950	4.5V to 26V	0.8V to 5.5V	12A	MLF®-28L
MIC27600	4.5V to 36V	0.8V to 5.5V	7A	MLF®-28L

Applications

In today's power-hungry applications, there is an insatiable need to deliver high levels of current while maximizing efficiency and minimizing board space. Micrel's SuperSwitcher II™ high power density buck regulators are an ideal fit.

The SuperSwitcher II™ integrated FET DC-DC buck regulators are optimized to deliver world class power density coupled with simplicity of use. Ultra-fast transient response, high efficiency, extremely low duty cycle operation and Any Capacitor™ stable are all merged together to make the SuperSwitcher II™ family easy to use with a reduced BOM cost.

For more information, contact your local Micrel sales representative or visit Micrel at: www.micrel.com/ad/superswitcherII.

Key Features

- Any Capacitor™ Stable – No Stability Issues with Distributed Capacitance, ESR Changes
- High Delta V Operation Can Eliminate Intermediate Supplies and Improve System Efficiency ($36V_{IN}$, $0.8V_{OUT}$)
- Fast Transient Response Allows for Smaller, Less Expensive Output Capacitors
- Constant Frequency Operation Simplifies EMI Design

MICREL®
Innovation Through Technology®

www.micrel.com

BUILD ACCURATE SPICE MODELS FOR LOW-NOISE, LOW-POWER PRECISION AMPLIFIERS

UNDERSTAND SPICE-MODEL CONSTRUCTION TO APPRECIATE AC AND DC PERFORMANCE.

BY DON LAFONTAINE • INTERSIL

System engineers require accurate models for all types of ICs, and they require Spice models to run comprehensive circuit simulations. Early Spice models had few nonlinear elements, minimizing simulation time at the cost of accuracy, but new methods let you increase the number of nonlinear elements and improve accuracy. You can create a multistage model for low-noise, low-power operational amplifiers. The model employs work from Analog Devices (**Reference 1**) and requires several architectural changes to model a low-noise, low-power precision amplifier. The model architecture processes the input signal through eight stages. You can easily calculate the parameters for the eight stages with a handheld calculator. To understand the model creation, you must have experience using Spice.

Although higher-speed amplifiers have multiple poles and zeros, this model is for a single-pole, 10-MHz amplifier. It lets you simulate the amplifier's key ac and dc parameters. The model includes ac parameters for flicker and flatband noise, slew rate, CMRR (common-mode rejection ratio), gain, and phase. The model's dc parameters are V_{OS} (input offset voltage), I_{OS} (input offset current), quiescent supply current, and output-voltage swing. The model uses the 25°C typical parameters (**Reference 2**).

The closer you model the input stage to the actual amplifier, the more accurate your results will be. You can achieve an accurate ac representation of the am-

PHOTO-ILLUSTRATION BY TIM BURNS. SPICE RACK: PALEKA19/ISTOCKPHOTO.COM

plifier's performance using a few of the process parameters of the input-stage transistors or MOSFETs. This model's architecture lets you model amplifiers with split supplies. There is no ground reference in any of the signal-processing blocks. After the differential-to-single-ended conversion, all internally generated node voltages are referenced to the midpoint of the power supplies, much like the actual operation of an amplifier.

EIGHT CASCADED STAGES

You build the model's circuit schematic with eight functional blocks (**Figure 1**). The only circuitry resembling an amplifier is the input stage (see **Listing 1**, which is available in the Web version of this article at www.edn.com/110303df). All other stages process the input signal with voltage-controlled current sources or voltage-controlled voltage sources. The stages might also include diodes, dc supplies, resistors, capacitors, and inductors.

The voltage-noise stage generates the flicker and flatband noise. To generate a flatband noise with only 4 nV/√Hz, set all diodes' and transistors' model parameters, kf and af, to zero and one, respectively. To lower the model's noise floor to single-digit nanovolts, it may be necessary to reduce the model's Johnson noise (**Reference 3**) by reducing the resistance values where possible.

Before reducing the resistor values, you calculate the standard resistor values and complete all the simulation tweaks. You then tweak the voltage-noise stage by reducing the resistor values to 1Ω and recalculating the transconductance and time constants of the stages to maintain the same transfer function. You can usually set resistors R_5, R_6, R_9, R_{10}, R_{11}, and R_{12} to 1Ω. You need not reduce the Johnson noise if the amplifier you are modeling has hundreds of nanovolts of input noise. The initial noise simulations will tell you whether this step is necessary. Once you set the model's flatband noise below the amplifier's noise floor, you can change the flicker and flatband noise by adjusting D_N, R_{17}, and V_5.

The input stage of the selected op amp comprises five bipolar transistors that model the actual IC configuration. For op amps with NMOS or PMOS stages, see **sidebar** "Typical input stages" in the Web version of this article at www.edn.com/110303df. The input stage includes a current supply to model I_{OS}, a voltage supply to model V_{OS}, and a voltage-controlled voltage source. R_1 and R_2 account for the CMRR of the device.

The first gain stage sets the combined

AT A GLANCE

- Model your op amps with detailed transistor-level input stages.
- Set model parameters to replicate the low noise of your amplifier.
- Properly characterize the supply-current usage to get an accurate model.
- A procedure helps you set values in the eight stages of the model.
- Correlate your simulations with real-world results.

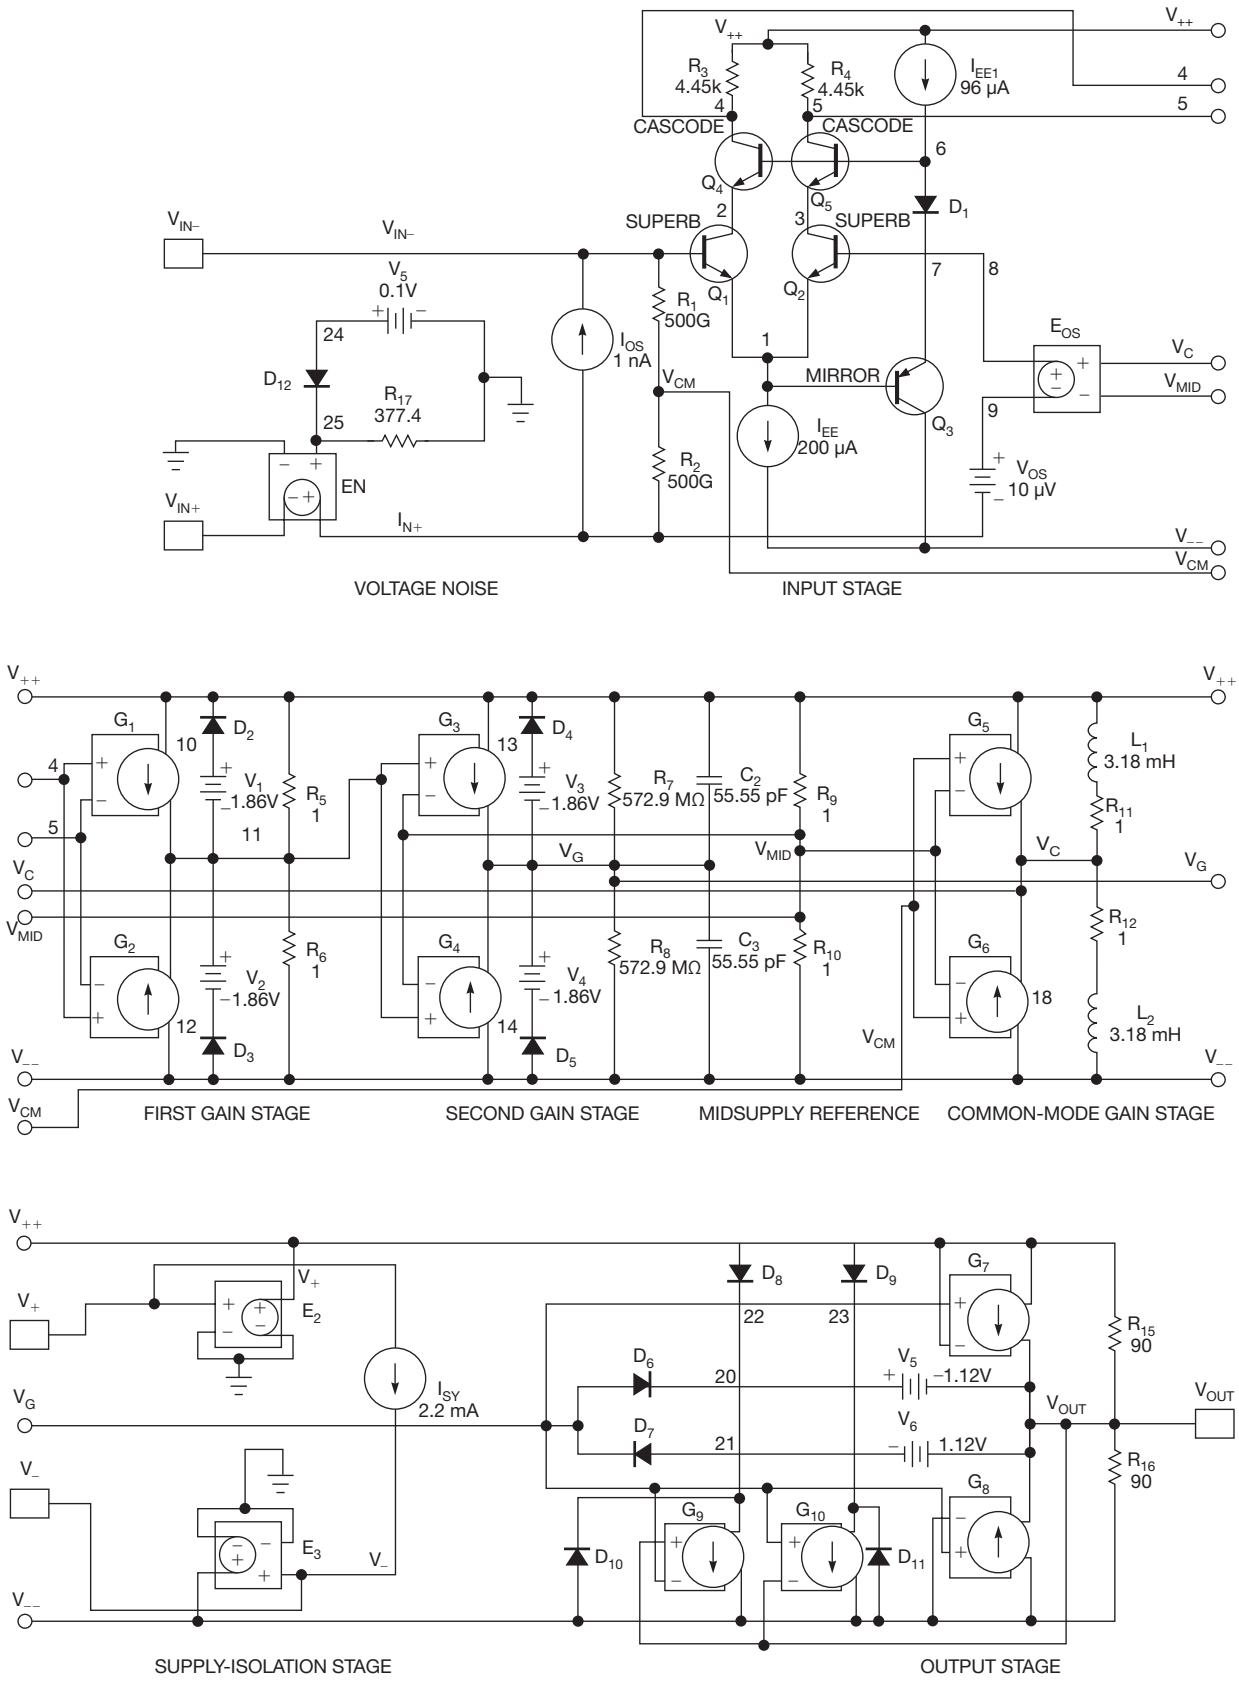

Figure 1 This Spice model has eight stages.

gain of the input stage and the first gain stage to one, simplifying your calculation to determine the slew-rate-limiting components in the second gain stage. Diodes D_2 through D_5 and V_1 through V_4 clamp the output-voltage swing. Increasing the values of V_1 through V_4 will result in decreasing the maximum output-voltage swing. The maximum output-voltage swing, V_{OH}, is controlled by adjusting V_1 and V_3, and the minimum output-voltage swing, V_{OL}, is controlled by V_2 and V_4. Voltage limiting must take place in the open-loop gain stage; otherwise, succeeding nodes could attempt to simulate the generation of huge (hundreds of kilovolts) signals.

In the second gain stage, you set the open-loop voltage gain, bandwidth, and slew rate of the amplifier by adjusting transconductance blocks G_3 and G_4 and components R_7, R_8, C_2, and C_3.

The midsupply reference stage comprises two equal-value resistors, R_9 and R_{10}. You use these resistors to generate a midsupply reference voltage. You set the resistor values to 1Ω to reduce the Johnson voltage noise of the model. The high current that flows through these resistors does not appear in the simulation results because of the supply-isolation stage.

The common-mode gain stage comprises two voltage-controlled current sources that drive two equal-value resistors in series with an inductor that connects to the supply rails (see **sidebar** "How the voltage-controlled-current-source stage works" in the Web version of this article at www.edn.com/110303df). The inductors simulate the fall-off of CMRR at higher input frequencies. You control the current sources with the input-common-mode voltage relative to the midsupply voltage. You generate the common-mode voltage with resistors R_1 and R_2 in the input stage. You set the transconductance of each source based on the reciprocal of the associated resistor value divided by the CMRR of the amplifier at dc. The inductors add a Z-plane zero to the common-mode gain, which is equivalent to adding a pole to the CMRR. After you scale and frequency-shape the common-mode voltage, you add it back into the input stage with a voltage-controlled voltage source, E_{OS}.

The supply-isolation stage comprises two voltage-controlled voltage sources and a current source. This stage enables you to program the total supply current of the amplifier with one entry in the node list. It also isolates the internal supply currents from the external supply current that the simulator sees, enabling the model to provide the correct quiescent supply current for low-power amplifiers with low voltage noise. The design trade-off is disabling the output-supply-current circuitry. By disabling this circuitry, the model does not account for the load current.

The output-stage operation is not obvious. After receiving appropriate frequency shaping, the amplifier's output signal appears as a voltage referenced to the midsupply at the inputs to G_7 and G_8. Transconductance blocks G_7 and G_8 drive two equal-value resistors that connect to the supply rails. They act as active current generators. Both G_7 and G_8 generate just enough current to provide the desired voltage drop

TABLE 1 OP-AMP SPECS

Parameter	Value	Comments
Quiescent supply current	0.0022A	
V_{CC}	15V	
V_{EE}	–15V	
I_{EE}	200 μA	Differential input-current source, or tail current
Slew rate	3.6×10^{6}V/sec	
F_{P1}	5 Hz	Dominant pole from Figure 3
A_{VOL}	264,000V/V	128.43 dB
V_{OS}	1×10^{-5}V	
I_{OS}	1×10^{-9}A	
Temperature	25°C	
V_T	0.0257V	
CMRR	3.16×10^{7}V/V	150 dB
F_{CM}	50 Hz	Common-mode pole
R_{OUT}	45Ω	
I_{SC}	45 mA	
V_{OH}	13.7V	Maximum output voltage
V_{OL}	–13.7V	Minimum output voltage

across its parallel resistor (**Figure 2**).

By simulating the right output resistance, you reduce the dc open-loop gain at the amplifier load. You model output short-circuit protection with diodes D_6 and D_7, along with dc supplies V_5 and V_6. Under fault conditions, the output voltage clamps to the previous frequency-shaping stage. You set the output short-circuit-current limit by adjusting the values of V_5 and V_6.

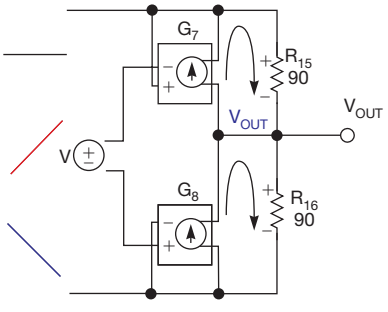

Figure 2 A voltage-controlled voltage source's output stage has three modes of operation. When the input voltage is constant, the voltage drop across the resistors causes them to equally oppose each other. When the input voltage rises, the current in R_{15} decreases, the current in R_{16} increases, and the midpoint voltage increases. When the input voltage decreases, the current in R_{15} increases, the current in R_{16} decreases, and the midpoint voltage decreases.

CALCULATE PARAMETERS

To calculate the Spice model's parameters, you need various specifications from the op amp's data sheet and IC-design information (**Table 1**). Putting the following **equations** into an Excel spreadsheet will enable you to change critical specs and quickly see the effect on the op amp's performance. You enter E_{OS} and I_{OS} from the data sheet directly into the model. You can often get the input-differential-capacitance spec, if your model requires it, from the data sheet. You determine the values of C_2 and C_3 by relating the input differential-pair tail current, I_{EE}, to the amplifier's slew rate. This step establishes the maximum frequency for the single-pole RC network and sets the unity-gain bandwidth. Amplifier data sheets rarely state I_{EE}, so you must get this value from the IC designer. For the amplifier of **Reference 2**, the **equation** works out to 55.55 pF:

$$C_2=C_3=\frac{I_{EE}}{\text{SLEW RATE}};$$

$$C_2=C_3=\frac{200\times10^{-6}}{3.6\times10^{6}}=55.55\text{ pF}.$$

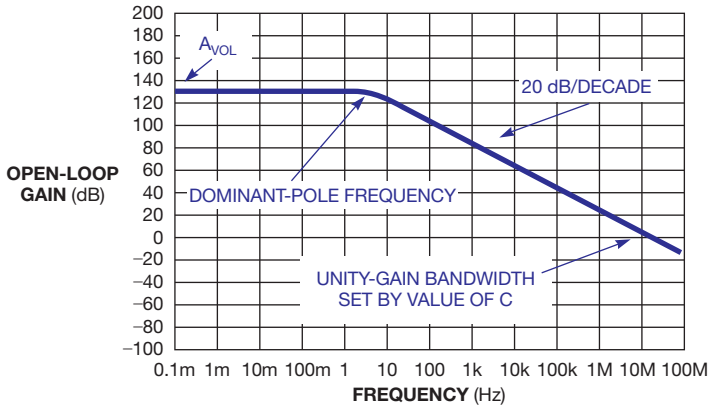

Figure 3 You use the A_{VOL}-versus-frequency chart in the IC's data sheet to determine the dominant-pole frequency.

You then determine the values of R_7 and R_8 by applying the dominant-pole frequency **equation**. This **equation** sets the breakpoint for the RC network. You observe f_{P1}, the dominant-pole frequency, by inspecting the open-loop gain chart in the amplifier's data sheet (**Figure 3**). The amplifier of **Reference 2** yields a value of 573 MΩ:

$$R_7=R_8=\frac{1}{2\pi f_{P1} C_{2/3}};$$

$$R_7=R_8=\frac{1}{2\pi (5)(55.55\ pF)}=572.958\times10^6.$$

You use the A_{VOL} from the data sheet and the resistance value $R_{7/8}$ to determine the transconductance of the voltage-controlled current source:

$$G_3=G_4=\frac{A_{VOL}}{R_{7/8}};$$

$$G_3=G_4=\frac{2.64\times10^6}{572.958\times10^6}=4.6\times10^{-3}.$$

For the amplifier of **Reference 2**, the transconductance blocks solve to 0.0046. The model uses the first gain stage to set the combined gain of the input stage and the first gain stage to one. You can now calculate the voltage at the input of G_3 and G_4 to cause 200×10^{-6}, the tail current, to flow through R_7 and R_8:

$$g_m=\frac{I}{V}\longrightarrow V_{G3/G4}=\frac{I}{g_m}=$$

$$\frac{200\times10^{-6}}{4.6\times10^{-3}}=43.4\times10^{-3}.$$

During slew-rate limit, the 200×10^{-6} current sink clamps the current through resistor R_3 or R_4. With positive input voltage, R_4 carries the current. Negative input means that the current will pass through R_3. This current flows through the 4.45-kΩ resistor, resulting in a voltage drop of 0.89V. This voltage drop appears at the input to G_1 and G_2. To set the combined gain of the input stage and the first stage to one, you must calculate the transconductance of G_1 and G_2 so that their output voltage equals 0.0434V when 0.89V is at their inputs. If you set the resistor value in parallel with the outputs of G_1 and G_2 to 1Ω, the voltage will equal the current, and you can solve for the transconductance of G_1 and G_2:

$$G_1=G_2\longrightarrow g_m=\frac{I}{V}=$$

$$\frac{43.4\times10^{-3}}{890\times10^{-3}}=48.77\times10^{-3}.$$

If the design-review document is unavailable, set R_3 and R_4 to 1Ω for the calculation of the voltage appearing at the inputs to G_1 and G_2. If you can get the collector currents from the IC designer, you can enter them directly into the model:

$$R_3=R_4=4.45\ k\Omega.$$

You set $V_{1/3}$ and $V_{2/4}$ voltages for the maximum output-voltage swing:

$$V_{1/3}=V_{CC}-(V_{OUTMAX})+V_T\ln\left(\frac{2I_{EE}}{I_S}\right);$$

$$V_{2/4}=(-V_{OUTMAX})-V_{EE}+V_T\ln\left(\frac{2I_{EE}}{I_S}\right),$$

where V_T=0.02585V at T=25°C and I_S=1×10^{-12}A for both diodes. The model clamps the output voltage at a value equal to $V_{CC}+V_{D2/4}-V_{1/3}$ for positive input-voltage swings and $-V_{EE}-V_{2/4}+V_{D3/5}$ for negative input-voltage swings. For noise considerations, select a 1Ω value for R_{11} and R_{12}:

$$R_{11}=R_{12}=1\Omega.$$

Calculate the values for G_7 and G_8:

$$G_7=G_8=\frac{1}{R_{11/12}\times CMRR}.$$

By inspecting the CMRR-over-frequency chart in the data sheet, you can estimate the common-mode pole frequency, f_{PCM}. This frequency lets you solve for the inductor values:

$$L_1=L_2=\frac{R_{11/12}}{2\pi f_{PCM}}.$$

You set the output-stage transconductance equal to the reciprocal of $2R_{OUT}$:

$$G_7=G_8=G_9=G_{10}=\frac{1}{2R_{OUT}}.$$

This step results in unity gain through G_7, G_8, G_9, and G_{10}. You establish the output resistors as $2R_{OUT}$ to account for the fact that the output current must appear to be coming from one supply rail:

$$R_{15}=R_{16}=2\times R_{OUT}.$$

You can solve now for V_5:

$$V_5=I_{SC}R_{OUT}-V_{D6}.$$

You create the absolute value of V_6 by taking the absolute value of the first term:

$$V_6=|I_{SC}R_{OUT}|-V_{D7}.$$

SIMULATION RESULTS

Once you calculate all the model parameters, you can run the simulation. The simulations correlate well to the device's specified performance from the data sheet (**figures 4** through **9**). The IC's data sheet provides more comparisons of simulations to real-world results. Modern computing pow-

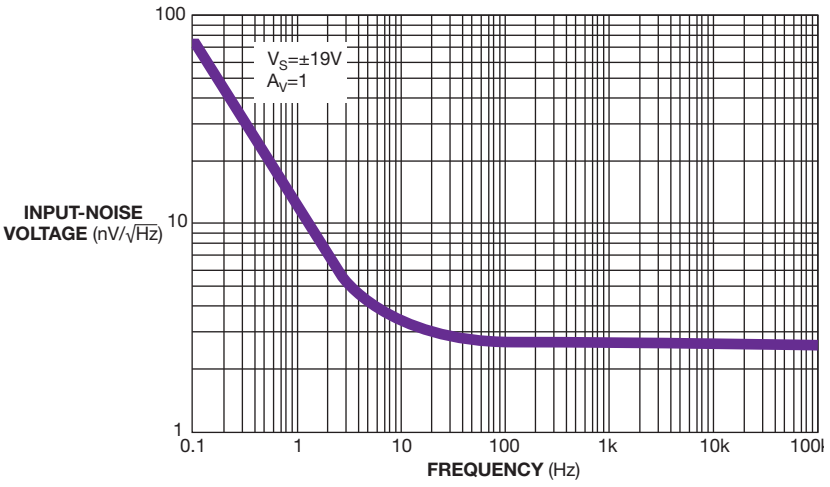

Figure 4 The measured input voltage noise spectral density has a flicker-noise corner and a 2.44 nV/√Hz flatband noise.

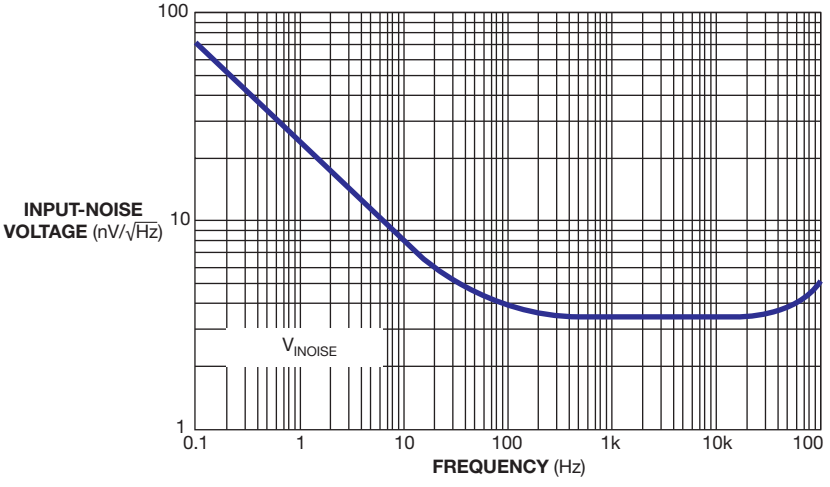

Figure 5 The simulated input voltage noise closely matches the flicker noise. The flatband noise is higher at 4 nV/√Hz.

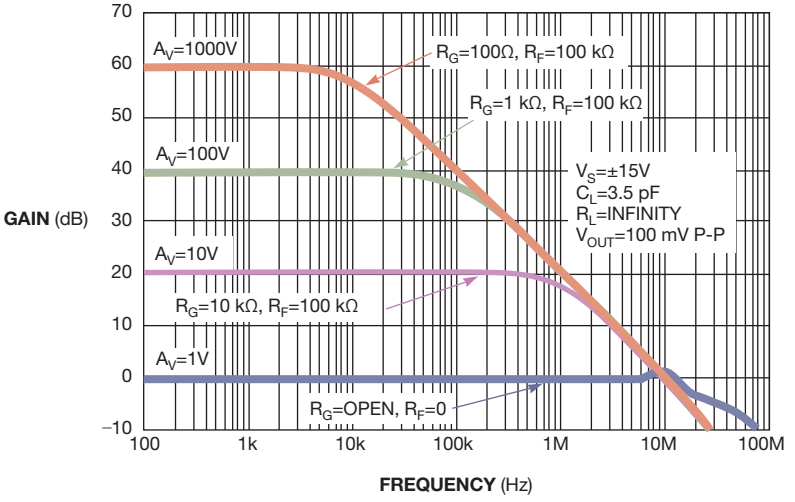

Figure 6 The measured closed-loop gain conforms to a single dominant-pole response except when R_F is 0 Hz.

UBM Canon

Your Project Acceleration Resource for Electronics Manufacturing

Electronics® NEW ENGLAND

April 6–7, 2011

Boston Convention
& Exhibition Center
Boston, MA

Electronics-NE.com

Electronics® WEST

February 14–16, 2012

Anaheim Convention Center
Anaheim, CA

ElectronicsWestShow.com

NEW

Electronics® SOUTH

May 2–3, 2012

Charlotte Convention Center
Charlotte, NC

Find the suppliers, tools, and services you need to make your product, process, and business more efficient, cost-effective, and profitable.

For information on attending or exhibiting, please call 310/445-4200 or visit CanonTradeShows.com

Produced and Managed by UBM Canon, 11444 W. Olympic Blvd., Los Angeles, CA, 90064-1549 Tel: 310/445-4200 UBMCanon.com

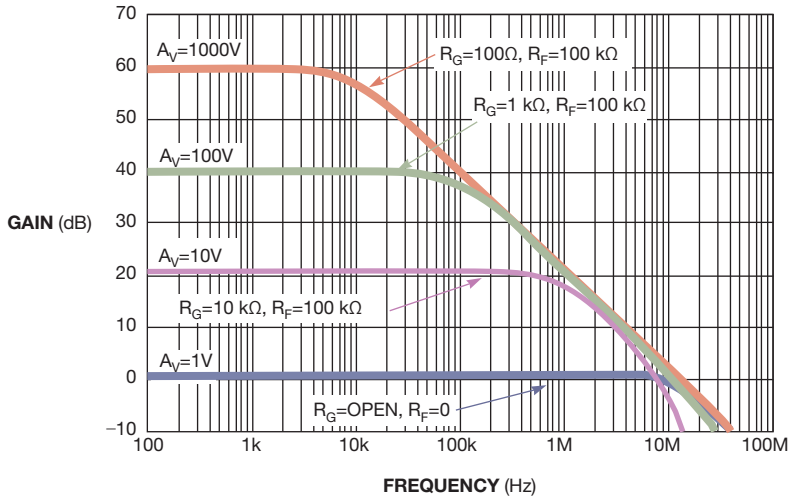

Figure 7 The simulation results closely conform except when A_V is 10V.

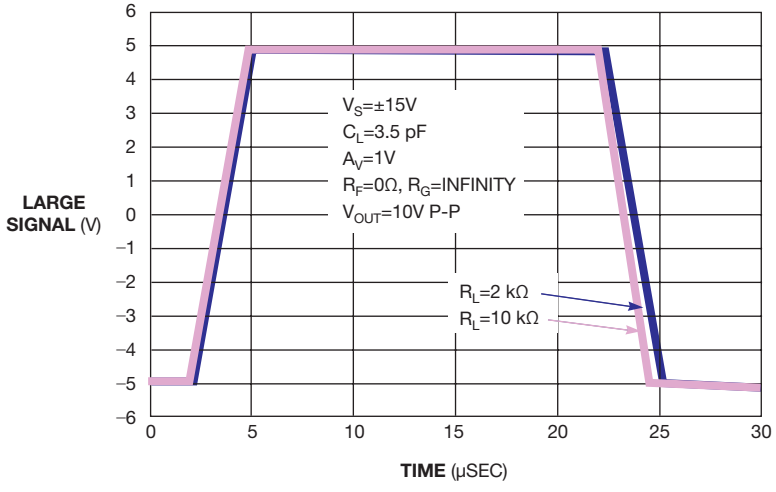

Figure 8 This measured large-signal response comes from the IC's data sheet.

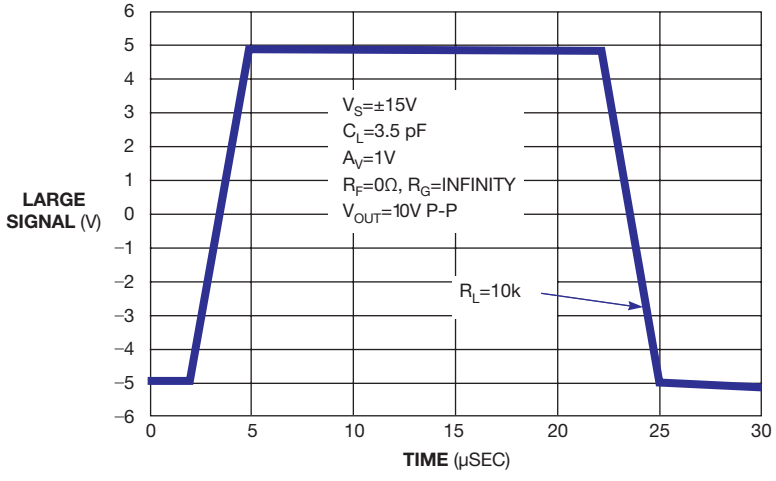

Figure 9 The simulated large-signal response is nearly identical to the measured results.

er quickly solves a model with five bipolar transistors in the input section. These transistors each have specific model parameters. Using this approach yields simulation results close to those of measurements of the part. In addition, this procedure allows you to model single-digit nanovolt noise parameters. Along with the ability to model low supply currents, this procedure is ideal for modern low-noise, micropowered amplifiers. EDN

ALONG WITH THE ABILITY TO MODEL LOW SUPPLY CURRENTS, THIS PROCEDURE IS IDEAL FOR MODERN LOW-NOISE, MICROPOWERED AMPLIFIERS.

REFERENCES

1 Alexander, Mark, and Derek F Bowers, "SPICE-Compatible Op Amp Macro-Models," Application Note AN-138, Analog Devices, February 1990, http://bit.ly/f1itjZ.

2 "Precision Single and Dual Low Noise Operational Amplifiers, ISL28127, ISL28227," Intersil Data Sheet, Dec 16, 2010, http://bit.ly/i0tcDD.

3 Bowers, Derek F, "Minimizing Noise in Analog Bipolar Circuit Design," *IEEE Proceedings of the 1989 Bipolar Circuits and Technology Meeting*, Sept 18, 1989, pg 107, http://bit.ly/ftdxMk.

AUTHOR'S BIOGRAPHY

Don LaFontaine is senior principal application engineer and senior engineering manager for Intersil's analog/mixed-signal product line (Palm Bay, FL), where he has worked for 30 years, focusing on precision analog products. He has a bachelor's degree in electrical engineering from the University of South Florida (Tampa, FL). His personal interests include family, surfing, and dirt-biking. You can reach him at dlafonta@intersil.com.

UBM Canon

New England's Project Acceleration Resource for Electronics Manufacturing

Electronics NEW ENGLAND

April 6-7, 2011

Boston Convention & Exhibition Center
Boston, MA

Your face-to-face meeting place for the latest technologies in electronics manufacturing, including:

- Electronic Manufacturing Services
- Testing & Inspection Equipment
- Automated Assembly Machines & Systems
- Board Level Components
- Computers/Peripherals/Software
- PCB Productions Equipment
- Embedded Devices

. . . and much more

Your Project Acceleration RESOURCE

Supported by:

EDN

CIRCUITS ASSEMBLY
The Journal for Surface Mount and Electronics Assembly

Register Today at:
Electronics-NE.com
Please use Promo Code AA

Produced and Managed by UBM Canon, 11444 W. Olympic Blvd., Los Angeles, CA, 90064-1549 Tel: 310/445-4200 ubmcanon.com

CPUs IN FPGAs: MANY FACES TO A TREND

WHETHER AS SYNTHESIZABLE SOFT CORES OR HARD CORES ON THE DIE, CPUs ARE SHOWING UP IN MORE FPGA DESIGNS, BRINGING WITH THEM IMPORTANT CHALLENGES FOR DESIGNERS.

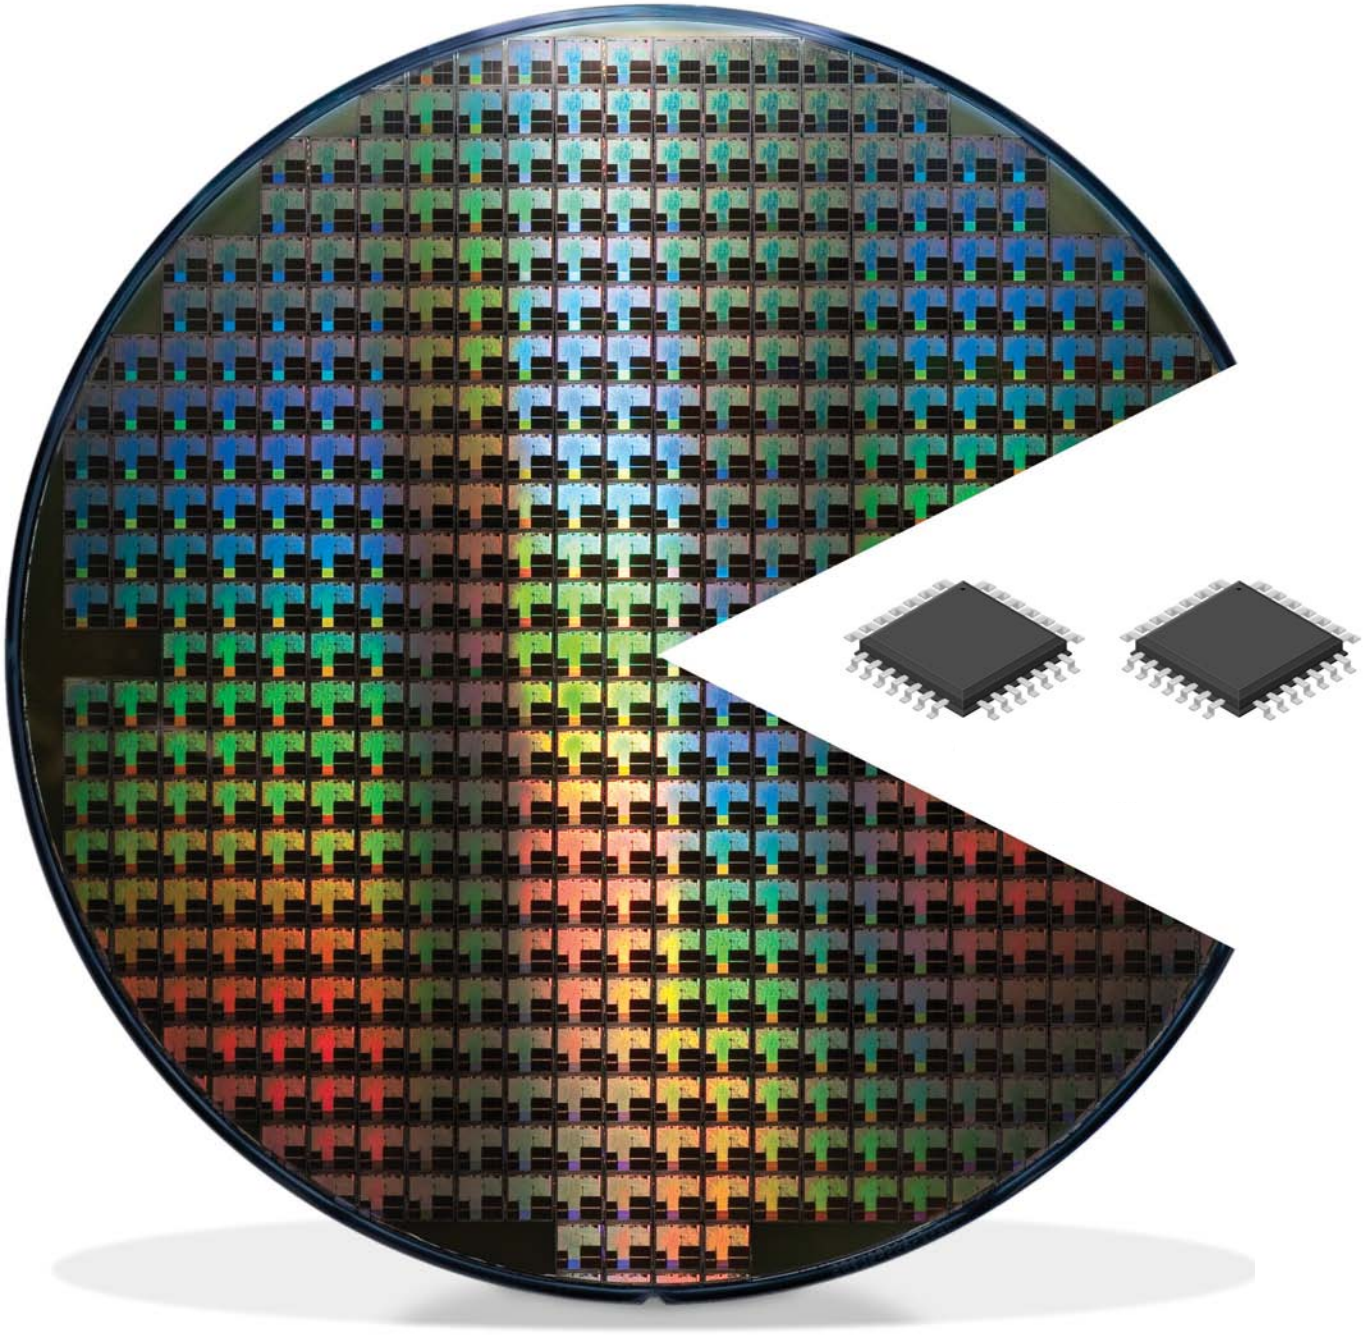

BY RON WILSON • EDITORIAL DIRECTOR

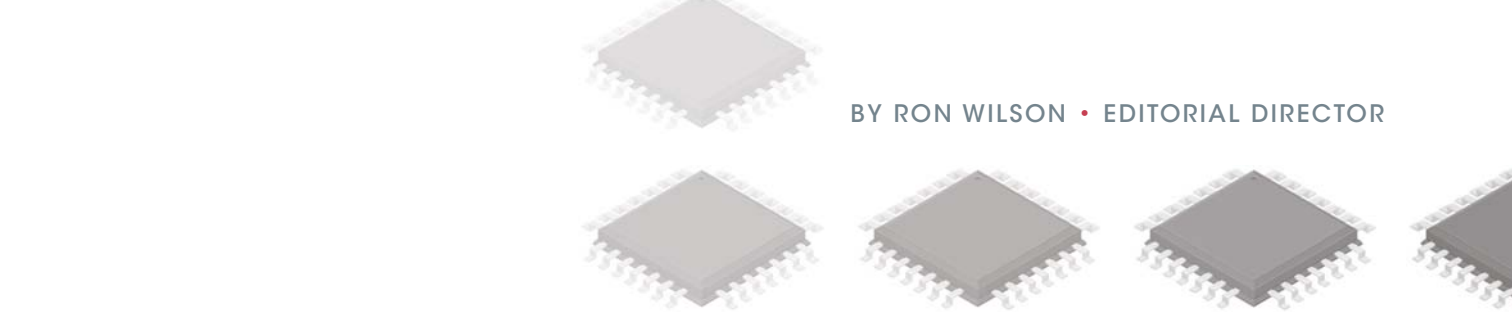

CPU cores in FPGAs have a history reaching back to the early years of the FPGA's existence and a future extending far into the realms of microcontrollers and ASSPs (application-specific standard products). We are now at an inflection point in that trajectory, facing manifold options. CPU cores may be soft—synthesizable cores that go into the FPGA's programmable logic—or hard—cell-based blocks that the FPGA vendor builds directly onto the die. The CPU architecture may be industry-standard, proprietary to the FPGA vendor, or unique. Processing capability spans tiny 8-bit microcontroller cores and 32-bit CPU clusters with DSP extensions. All this diversity conceals profound differences in implementation flow, in system performance, and in debugging access, all of which demand exploration.

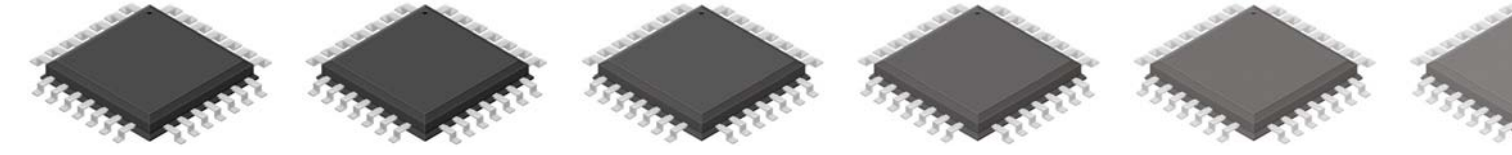

EXPLAINING DIFFERENCES

To understand the current nature of CPUs in FPGAs, it helps to have a little perspective. Early FPGAs, confined by their limited capacity, lived beside a microprocessor or a microcontroller, implementing glue logic or, more grandly, a bus bridge. With Moore's Law and human nature on the march, however, FPGAs grew. Users began to dream of implementing their own CPU designs—every digital designers' Holy Grail—in FPGAs.

Early attempts had modest goals: 4-bit microcontrollers, for example. Even in these low-gate-count undertakings, however, designers began to discover issues about implementing a CPU core in programmable fabric.

Some issues are structural. Implementers of discrete CPU cores have available a mixture of standard-cell and full-custom techniques. These riches allow architects to use structures that give them the best performance for the power and area, knowing that the implementation team has the tools to do its part of the job. Those structures may not make any sense in an FPGA fabric, however. For instance, FPGAs and shared tristate buses do not get along well together. Another example would be the fast—often single-cycle—multiport register files at the heart of CPU design. In advanced architectures, the general-register files may have eight or more ports, allowing different pipeline stages to initiate address calculations, do scalar arithmetic, and save results, all in the same cycle. The only options for implementing general registers in an FPGA, however, are the inherently single-port LUTs (look-up tables) in the logic cells or the block RAMs. Either option requires a wrapper of logic cells to emulate large-scale multiport access. The result, once you synthesize it into an FPGA, may be functionally correct, but either it significantly lowers the CPU's maximum frequency or the designer must significantly reorganize the execution pipeline.

Critical paths in general present problems for the FPGA adaptation of a CPU. A designer of silicon cores can simply minimize the number of negative-slack paths and then throw at them a whole battery of tricks: low-threshold cells, high-drive cells, custom cells, and hand placement and routing. None of these techniques is readily available to the typical FPGA user.

As FPGA users began to lust after more substantial cores, such as the elegant ColdFire or the ubiquitous ARM7, implementation realities limited the attractiveness of the whole idea. Many designers began to believe that CPU cores in FPGAs were big, slow, power-hungry, and generally not worth the trouble.

Engineers don't stay frustrated for long, though. About 10 years ago, de-

signers at Altera concluded that they could design a CPU core from a clean sheet of paper for an FPGA implementation—omitting features or structures that didn't fit into the FPGA's way of doing things—and come up with reasonable specifications for most applications. The result, confirming their optimism, was the Altera Nios core. Other vendors, including Xilinx with the MicroBlaze, joined. Both FPGA-vendor applications teams and third-party designers added their hands to the task, adapting industry-standard cores from ARM, MIPS, and other sources.

WHAT IS THE POINT?

At this stage, you might wonder whether there is a practical need behind all this activity. If you want a simple, slow CPU, you can buy an off-the-shelf microcontroller chip for less than $1. So why bother putting a core into an expensive FPGA? The answers to that question are revealing.

"Nios cores go into about 30% of the devices we ship," says Chris Balough, Altera's senior director of software, embedded, and DSP marketing. "The customers who use Nios have pulled the architecture in many directions." Those words form a bit of an understatement.

"Most often, customers use Nios as a system controller," says Nirmal Kari, embedded-product marketing manager at Altera. "After that use, the most common applications are as a coprocessor, in multiprocessor clusters, or as a DSP core. Then, some users will implement state machines using a core." The wide variety of Nios applications has led to an accumulation of both supporting hardware blocks and core implementations (**Figure 1**). Economy-level and high-speed variants of Nios are now available, as are a custom instruction generator, a memory manager, and a Linux port.

AT A GLANCE

- Both hard and soft CPU cores are available to FPGA users.
- CPU cores can serve a range of purposes in an FPGA design.
- The vendor's view of the purpose has influenced tools and support for individual cores.
- Verification and software debugging for designs with embedded cores are not trivial issues.

Users' motives are equally diverse. The justification is sometimes a complex analysis of total cost of ownership, showing that avoiding an external, single-sourced microprocessor or microcontroller over the life of the product will save money. The justification is sometimes technical, as when the pinout to interface the FPGA to an external processor is unavailable. At other times, the decision is obvious, as when a task changes from time to time or is slow enough to do in software and the necessary space in the FPGA is going unused.

Xilinx's MicroBlaze has evolved along a similar path but perhaps with a different emphasis. "Much of the use of MicroBlaze is to implement state machines," says Vidya Rajagopalan, Xilinx's vice president for processing products. "Ironically, the bigger the FPGA, the smaller the CPU cores people seem to put into it."

This trend has not kept MicroBlaze from feature expansion. MicroBlaze has a memory-management unit, memory-protection hardware, and single-precision floating-point hardware, as well as the obligatory Linux port, according to Navanee Sundaramoorthy, Xilinx's manager for embedded-platform marketing. The two seemingly contradictory points of view may reflect Xilinx's increasing focus on its EPP (Extensible Processing Platform), an FPGA with a hard ARM Cortex A9 microcomputer cluster. The company's positioning anoints EPP as the correct approach to application processing in FPGAs and focuses MicroBlaze on more deeply embedded uses, such as sequencing and control of other blocks within the fabric. Altera, in contrast, has not yet fully disclosed its coming hard-core products and still recognizes Nios as spanning the range from Linux-based application processing to state machines.

Microsemi, now owner of Actel, has a different orientation: providing microcontroller alternatives, according to senior product-development manager Tim Morin. Consequently, Morin focuses on the company's SmartFusion chips with their hard Cortex M3 CPUs. The prod-

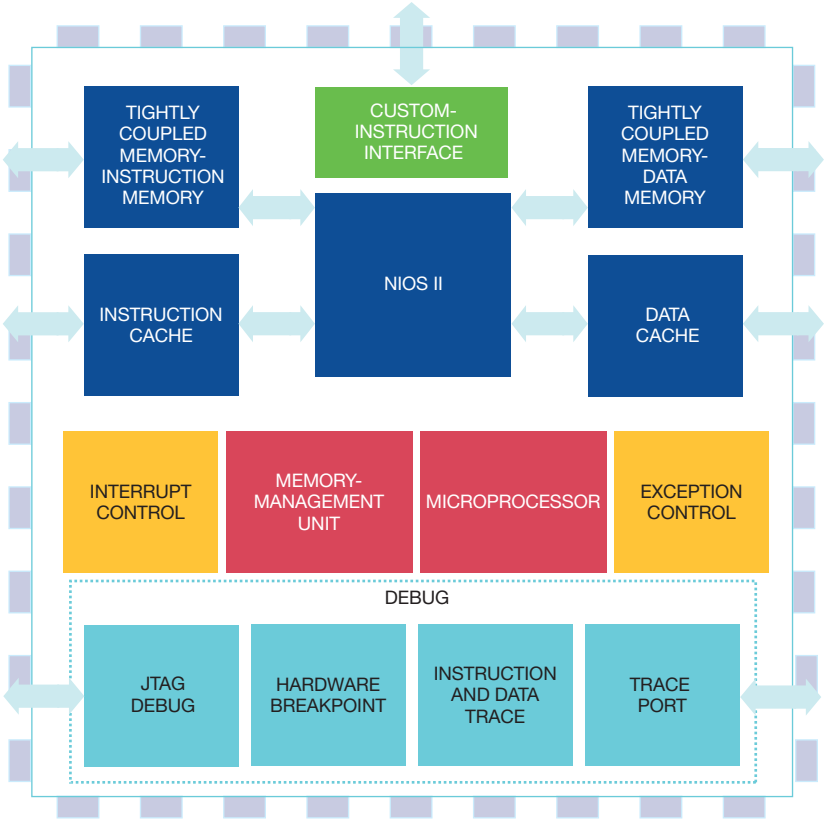

Figure 1 Nios has evolved from a simple controller to a capable application-level CPU, leading to an accumulation of supporting hardware blocks and core implementations.

PHOTO-ILLUSTRATION BY TIM BURNS. WAFER: PHOTOMICK/ISTOCKPHOTO.COM; CHIPS: GRIMGRAM/ISTOCKPHOTO.COM (PREVIOUS PAGE)

uct line also includes soft cores, however, such as an 8051, a Cortex M1, and a military-oriented Leon, all for use in the non-hard-core A3P and Fusion FPGA families.

Lattice Semiconductor has taken yet another approach with its 32-bit architecture. LatticeMico32 is, unlike the other soft cores, an open-source IP (intellectual-property) block, according to Mike Kendrick, software-product planning manager at the company. The CPU is structurally similar to many compact modern 32-bit cores, with a five-stage pipeline, provision for both caches and fast local RAM, and substantial interconnect and peripheral IP.

The decision to go open-source has interesting implications for the modest-sized designs Lattice targets. Kendrick says that most LatticeMico users fit into one of two profiles. Some use the core to implement complex but low-performance functions, such as connectivity to an external processor or user-interface hosting. Others use the core for system-management purposes: bringing up the system at power-up; power, performance, and safety monitoring; and debugging. None of these applications requires earth-shaking speed, use of the latest real-time operating system, or great numerical power. LatticeMico focuses on these applications.

In this context, the choice to go open-source becomes important, rather than merely convenient. Without the need for absolute performance, LatticeMico can remain a simple design that users with access to all the RTL (register-transfer logic) can readily understand. The RTL access offers the implicit ability to simply go into the RTL source, see what's going on, and adapt it.

INTERCONNECT CHOICES

Whatever their markets address, all the CPU cores must deal with buses and memories, but history and application focus influence how each core approaches these issues. Nios, which preceded most modern bus standards, has its own interconnect system: Avalon. Avoiding the problems of shared physical buses, Avalon generally uses a switch matrix. Bowing to the momentum of ARM, however, Altera recently announced that Avalon could accommodate an ARM AXI (advanced extensible interconnect) structure by putting gaskets around it. Xilinx also recently announced adoption of AXI as its core interconnect IP.

The physical layer is not the whole solution to the user's problem. "Customers don't want a core; they want an ecosystem," says Vin Ratford, senior vice president of marketing at Xilinx. So the vendors provide not only the CPU core but also a set of soft memory blocks, soft peripherals, and a system-building graphical user interface that does the dirty work of assembling the blocks, gluing them together, and generating synthesis and mapping scripts. In the case of Xilinx's MicroBlaze, the ecosystem can create sophisticated micro-

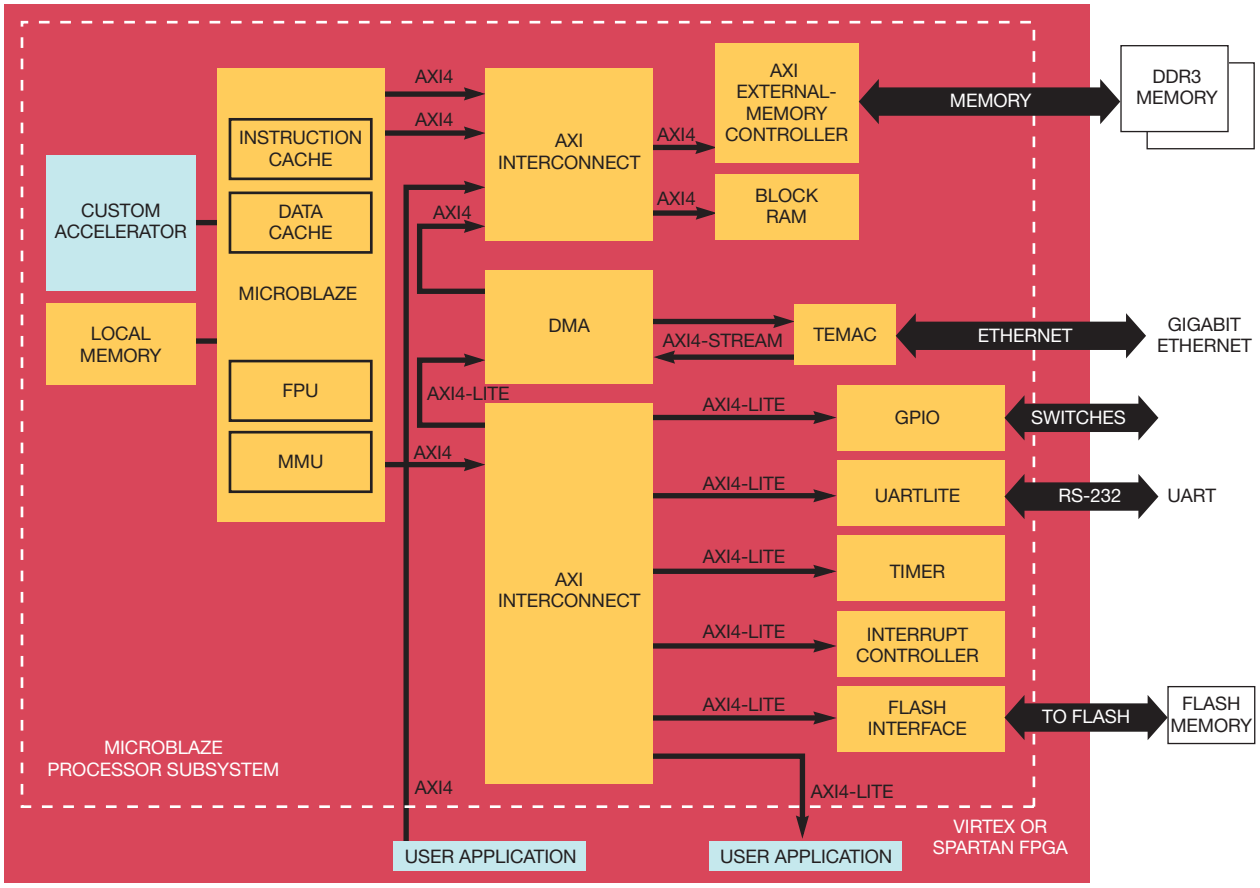

AXI: ADVANCED EXTENSIBLE INTERFACE
FPU: FLOATING-POINT UNIT
MMU: MEMORY-MANAGEMENT UNIT
TEMAC: TRIMODE ETHERNET MEDIA-ACCESS CONTROLLER
UART: UNIVERSAL ASYNCHRONOUS RECEIVER/TRANSMITTER

Figure 2 A fully developed MicroBlaze system relies on the AXI interconnect scheme for nearly all its connectivity.

computers almost entirely from Xilinx-provided IP blocks and automatically generated interconnect (**Figure 2**).

All the major vendors have taken this approach. "Our tool provides the IP to stitch together the CPU, memory, peripherals, and AHB [advanced high-performance-bus] interconnect," says Microsemi's Morin. "The customer just drags and drops blocks in a graphics environment."

Lattice maintains its open-source stance with the interconnect IP as well as the CPU core. LatticeMico uses the Wishbone open-source interconnect approach, which Kendrick says superficially resembles Altera's Avalon but in detail differs greatly, so that interfaces for the two standards are not interchangeable. Like Altera, Lattice offers a graphics-system-building utility that generates an RTL design from block-diagram-style input.

THE MEMORY PROBLEM

Fast RAM is another problem. Application processors running Linux typically have large set-associative caches. Processors with relatively static task mixes often use faster, simpler direct-mapped caches or more deterministic RAMs, such as ARM's TCM (tightly coupled memory). The RAM itself isn't a big problem for modern FPGAs, which are rich in fast block RAM. Linking the CPU core to the RAM and implementing a controller can be more fraught with difficulties, however.

Microsemi, with its microcontroller focus, simply interfaces block RAM to the CPU core. "ARM designed the Cortex M3 to be cacheless," Morin says, "so we provide two dedicated 32-kbyte hard SRAMs on switch-matrix ports on SmartFusion. For the soft Cortex M1 core, we don't offer caches. Instead, the user can specify TCM through the system-building tool, and we configure a wrapper to attach the TCM to the CPU core."

Xilinx is similarly conservative with MicroBlaze. "We can implement caches either in block RAM or in LUTs," Rajagopalan says. "The caches are usually directly mapped. You can do a set-associative controller, but the added complexity would cut into maximum frequency."

Altera, with its wider range of target applications, keeps its options open. "We implement L1 cache for Nios as a soft controller working with on-chip memory blocks," Kari says. "We can do set-associative caches with little impact on maximum frequency, but the result on system performance depends on the software. It can go either way."

ACCELERATORS

One of the nominal attractions of

cores in FPGAs is the ability to implement application-specific hardware accelerations in the programmable fabric. Just how you go about accomplishing that task, however, varies from chip to chip. The Nios tools offer users the ability to define a custom instruction. The tool then generates the changes in decoding and execution logic necessary to integrate the new instruction into the Nios pipeline. This approach can work well if a single operation is the bottleneck on a critical loop. Alternatively, you can create a custom accelerator, often by directly synthesizing RTL from the C code you are trying to accelerate. You would then compile an Avalon wrapper for the new block and drop the result into your design.

Xilinx takes a similar two-level approach with MicroBlaze. "There are two ways to do accelerators," says Sundaramoorthy. "Large accelerators should attach to AXI. Smaller ones—say, a FIR [finite-impulse-response] filter or a saturating multiplier—couple directly to the MicroBlaze pipeline." Nominally, because cores are in RTL, you can do anything you want to extend the operation of a third-party core. The knowledge of the core design you would require, however, and the impact on the resulting system's correctness and performance might be your responsibility rather than the core vendor's.

HOW TO DEBUG

Designers face several issues in verifying an FPGA design with an embedded CPU: verifying the CPU core, verifying the microcomputer as a subsystem, and debugging the software. The difficulty of each issue depends on the core in question. FPGA vendors claim that they have thoroughly verified their own hard or soft CPU cores and peripheral IP. They further claim that the subsystems you build with their system-building tools are correct by construction. Little dispute of these claims exists; beyond them, though, things get more complicated.

Vendors may make different—or no—claims about third-party CPU cores. Also, if you plan to modify a core with a custom acceleration, even using the vendor's tools, make sure that you and the vendor agree on a verification plan. A careful verification manager might want to reverify the entire CPU. Do you have the tools, expertise, and test bench necessary to do so?

Then there is system verification. "It's true that vendor tools can generate correct interface RTL," says Adam Sherer, verification director of product marketing at Cadence. "But the vendor can't preverify system interactions. We have seen a lot of FPGA-design teams who have been successful with midtier designs get into trouble with CPU cores. They add a CPU, a bus, and some peripherals, and those additions take the design beyond what their program-and-try verification style can handle. They need high-speed models and an ASIC-level verification method to deal with the complexity they have created." Such a verification methodology may entail not only simulation models of all

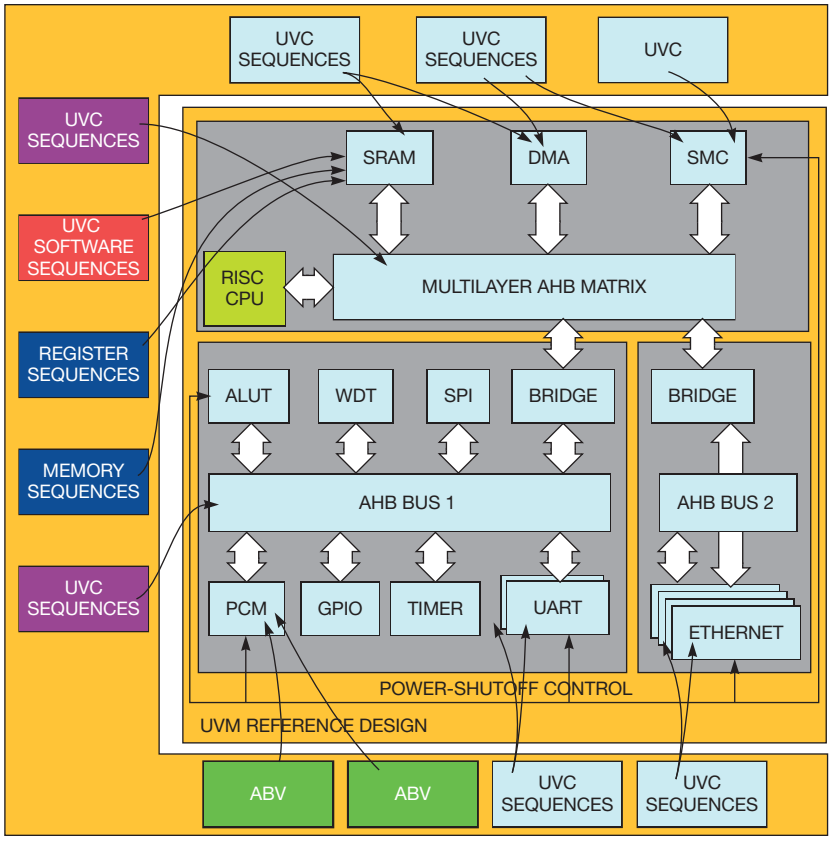

ABV: ASSERTION-BASED VERIFICATION
AHB: ADVANCED HIGH-PERFORMANCE BUS
ALUT: ADDRESS LOOK-UP TABLE
APB: ADVANCED PERIPHERAL BUS
DMA: DIRECT-MEMORY ACCESS
GPIO: GENERAL-PURPOSE INPUT/OUTPUT
PCM: PULSE-CODE MODULATION
SMC: SMART MEMORY CONTROLLER
SPI: SERIAL-PERIPHERAL INTERFACE
UVC: UVM VERIFICATION COMPONENTS
UVM: UNIVERSAL VERIFICATION METHODOLOGY
WDT: WATCHDOG TIMER

Figure 3 Assembling the tools and verification IP necessary to fully test a microcomputer subsystem in an FPGA can be a significant undertaking.

the functional blocks in the microcomputer subsystem but also third-party verification IP for these blocks, plus a test bench that can manage and track the progress of the verification effort (**Figure 3**). This process is more complex and skill-intensive than the processes many FPGA teams have used before.

Lattice's Kendrick agrees that system verification can be a complex issue for many design teams. To limit the complexity, Kendrick says, LatticeMico users often start with a verified open-source reference design and make changes to adapt the design to their requirements. This approach at least attempts to limit the scope of the verification problem, so that the design is likely to be at least partially functional on first pass. "Some customers use RTL simulation for the initial stages of bringing up the design, but many do not," Kendrick says. "Many just make careful modifications to the working reference design and then go directly to the board to try it out."

As for the software-debugging question, the SOC (system-on-chip) world increasingly relies on internal debugging kernels, such as ARM's CoreSight. These kernels may be in the hard cores planned for EPP and similar products, but the vendors don't always say. CoreSight does appear to be in the SmartFusion M3 core, for example. Such support on vendor-proprietary cores is another matter, however. MicroBlaze has JTAG (Joint Test Action Group) access that allows single-step and breakpoints, according to Sundaramoorthy. Similar capabilities exist on Nios, but your favorite debugging tool may not support the facility on the core you select. LatticeMico's in-core monitor, for example, interfaces to a Lattice debugger that, although specific to the company, complies with Eclipse specifications.

For deeper inspection of system operation, you can create a custom test—an assertion, essentially—and compile it into the FPGA. But with compilation times for big FPGAs approaching a day, this option may be less than attractive. You can also use the built-in debugging capabilities in the logic fabric—Chipscope, in Xilinx's case—to examine the operation of the design down to the individual FPGA register. Without detailed knowledge of the implementation, however, linking FPGA-logic cells to source-code execution would be a monumental job. Cadence's suggestion of starting with a fast system-level model before moving to hardware seems more appropriate.

As a whole, the hard and soft CPU cores now available to FPGA users represent a range of capabilities, sizes, and costs, and they find use across purposes ranging from hosting Linux and a complete system-control task load to accelerating arithmetic operations, performing system-management functions, and implementing a state machine in an easily modifiable way. Vendors offer different approaches to creating systems, integrating custom hardware, and debugging, but the cores have in common the ability to greatly complicate system verification and software debugging. It is essential at the outset of a design in which you will use a CPU core to understand these implications and to have a thorough, agreed-upon verification and software-development plan. EDN

FOR MORE INFORMATION

Actel/Microsemi
www.actel.com

Altera Corp
www.altera.com

ARM
www.arm.com

Cadence
www.cadence.com

Eclipse
www.eclipse.org

Lattice
www.latticesemi.com

MIPS
www.mips.com

Wishbone
cdn.opencores.org/downloads/wbspec_b3.pdf

Xilinx
www.xilinx.com


You can reach Editorial Director Ron Wilson at 1-415-947-6317 and ron.wilson@ubm.com.

LINEAR TECHNOLOGY

DESIGN NOTES

Step-Down µModule Regulator Produces 15A Output from Inputs Down to 1.5V—No Bias Supply Required

1.5V to 5.5V Input, 0.8V to 5V Output from a 15mm × 15mm × 4.32mm LGA Package

Design Note 488

Alan Chern and Jason Sekanina

15A High Efficiency Output from a Low Input Voltage

The LTM®4611 is a switch mode, step-down DC/DC µModule® regulator in a compact 15mm × 15mm × 4.32mm LGA surface mount package. The switching controller, MOSFETs, inductor and supporting components are housed in the package. With a built-in differential remote sense amplifier, the LTM4611 can tightly regulate its output voltage from 0.8V to within 300mV of V_{IN} and deliver 15A output efficiently from 1.5V to 5.5V input.

Only a handful of components are needed to create a complete point-of-load (POL) solution with the LTM4611 (see Figure 1). The C_{SS} capacitor provides smooth start-up on the output and limits the input surge current during power-up. C_{FF} and C_P set the loop-compensation for fast transient response and good stability. The output voltage, 1.5V, is set by a single resistor, R_{SET}.

Efficiency is exceptional, even down to the lowest input voltages, as shown in Figure 2.

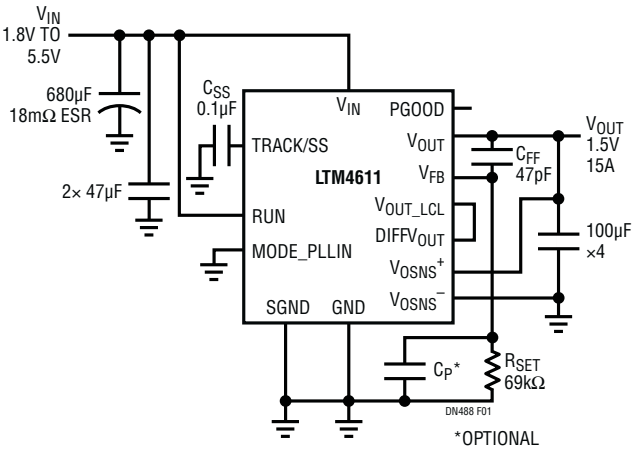

Figure 1. 1.8V_{IN} to 5.5V_{IN} to 1.5V_{OUT} with 15A Output Load Current

Input and Output Ripple

Output capacitors should have low ESR to meet output voltage ripple and transient requirements. A mixture of low ESR polymer and/or ceramic capacitors is sufficient for producing low output ripple with minimal noise and spiking. Output capacitors are chosen to optimize transient load response and loop stability to meet the application load-step requirements by using the Excel-based LTpowerCAD™ design tool. (Table 5 of the LTM4611 data sheet provides guidance for applications with 7.5A

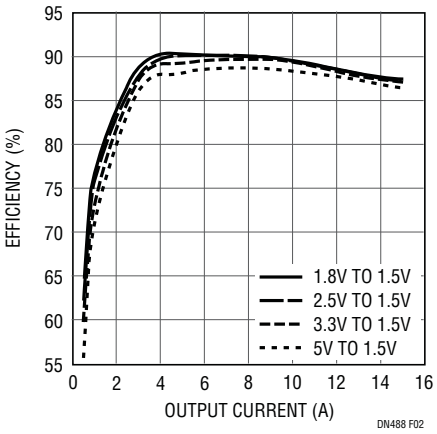

Figure 2. Efficiency of Figure 1 Circuit

QR code
m.linear.com/4611

Scan this code with your smart phone or click to view informative videos.

Download a QR code reader application to scan and read QR codes from your phone.

video.linear.com/56

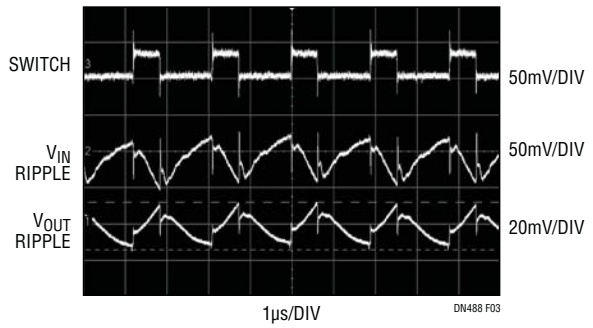

Figure 3. $5V_{IN}$ to $1.5V_{OUT}$ at 15A Output Load

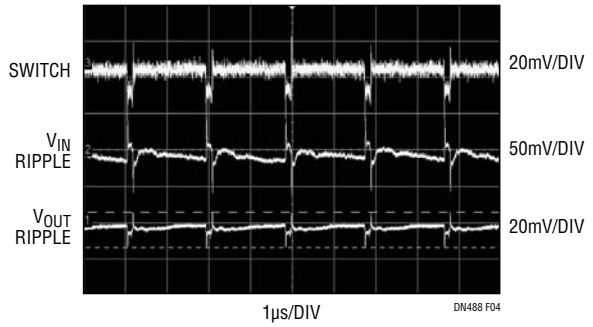

Figure 4. $1.8V_{IN}$ to $1.5V_{OUT}$ at 15A Output Load

video.linear.com/57

load-steps and 1µs transition times.) For this design example, four 100µF ceramic capacitors are used. Figures 3 and 4 show input and output ripple at 15A load with 20MHz bandwidth-limit. View the associated videos to see the test methodology, as well as ripple waveforms without bandwidth limiting.

For this design, the choice of input capacitors is critical due to the low input voltage range. Long input traces can cause voltage drops, which could nuisance-trip the µModule regulator's undervoltage lockout (UVLO) detection circuitry. Input ripple, typically a non-issue with higher input voltages, may fall a significant percentage below nominal—close to UVLO—at lower input voltages. In this case, input voltage ripple should be addressed since input filter oscillations can occur due to poor damping under heavy load current. This design uses a large 680µF POSCAP and two 47µF ceramic capacitors to compensate for meter-long input cables used during bench testing.

Thermally Enhanced Packaging

The device's LGA packaging allows heat sinking from both the top and bottom, facilitating the use of a metal chassis or a BGA heat sink. This form factor promotes excellent thermal dissipation with or without air flow. Figure 5 shows the top view thermal imaging of the LTM4611 at a power loss of 3.5W with no air flow, when converting 5V to 1.5V.

Internal self-heating of the LTM4611 remains quite low even at a low 1.8V input voltage due to its micropower bias generator that enables strong gate drive for its power MOSFETs. Figure 6 shows a power loss of 3.2W with hot spots slightly changed from their positions with a 5V input—the nominal surface temperature is 60°C. Watch the associated videos to see the test set-up and watch 200 LFM of air flow cool the unit by 10°C.

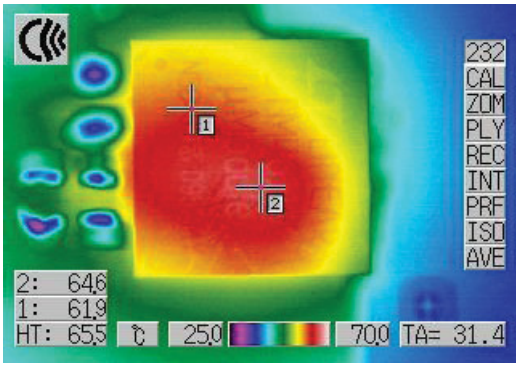

Figure 5. $5V_{IN}$ to $1.5V_{OUT}$ at 15A Output Load. 3.5W Power Loss with 0LFM and 65°C Surface Temperature Hot Spot

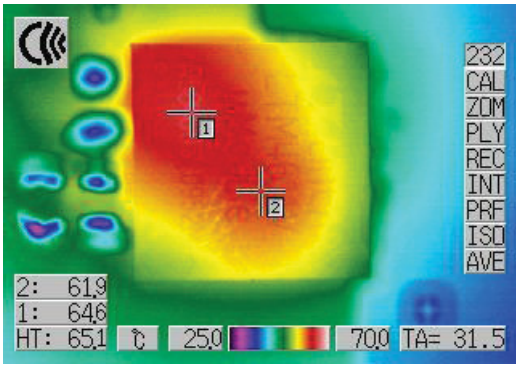

Figure 6. $1.8V_{IN}$ to $1.5V_{OUT}$ at 15A Output Load. 3.2W Power Loss with 0LFM and 65°C Surface Temperature Hot Spot

video.linear.com/55

Conclusion

The LTM4611 is a step-down µModule regulator that easily fits into POL applications needing high output current from low voltage inputs—from 1.5V to 5.5V. Efficiency and thermal performance remain high across the entire input voltage range, simplifying electrical, mechanical and system design in data storage, RAID, ATCA, and many other applications.

Data Sheet Download

www.linear.com/4611

For applications help,
call (408) 432-1900, Ext. 2593

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

dn488f LT/AP 0311 226K • PRINTED IN THE USA

BY MAURIZIO GAVARDONI AND AKSHAY BHAT • MAXIM INTEGRATED PRODUCTS

Monitor PWM load current with a high-side current-sense amplifier

AN INNOVATIVE CIRCUIT ALLOWS YOU TO MONITOR DIODES' RECIRCULATING AND POWER-SUPPLY CURRENT.

Accurate high-side current sensing is essential for automotive-control systems, such as those for EPS (electric power steering), automatic gear shifting, transmission control, engine fuel-injection control, braking-valve control, and active suspensions. All of these applications require precise regulation of the current through a motor or solenoid. The current sets the torque of the motor or the drive of the solenoid. You can design a precision, high-side current-sense amplifier to monitor inductive load currents over a wide range of input common-mode voltages. The circuit is suitable for applications in which the input common-mode voltage becomes negative due to inductive kickback, reverse-battery conditions, or transient events.

SENSE CURRENT IN EPS

Unlike conventional power-steering systems, an EPS system has no hydraulic pump or fluid. Instead, it features an electric motor that attaches to the steering rack through a gear mechanism. When the driver turns the wheel, a steering sensor detects the wheel's position and rate of rotation. The system then feeds this information, along with input from a steering torque sensor in the steering shaft, to the power-steering-control module. To determine the required steering assistance, the control module also takes inputs from the vehicle's speed-sensor, traction-control, and stability-control systems (**Figure 1**).

An interface with the power module then allows the control module to regulate the amount of current in the motor. Increasing the motor current increases the power assistance, and vice versa. You control motor current by feeding a PWM (pulse-width-modulated) voltage across the motor with an H bridge (**Figure 2**). An associated truth table summarizes the various modes of operation for a full H-bridge circuit (**Table 1**). The motor presents an inductive load. You determine torque by averaging the resulting ripple current. This current represents the resulting power assistance the circuit provides to the driver.

You use a current-measuring device to monitor motor current and provide real-time feedback to the control module. The module adjusts the PWM duty cycle until the current reaches its target value. You can insert a low-value sense resistor in series with the current path to produce a small reduction in voltage. A current-sense amplifier on that differential voltage indicates the current magnitude.

Current-sensing approaches include low side, high side, and on the motor. For low-side current sensing, you place the sense resistor between the H bridge

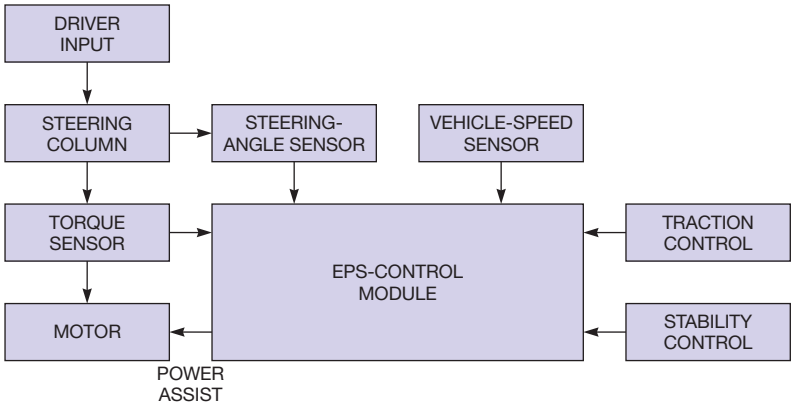

Figure 1 To determine the required steering assistance, the control module also takes inputs from the vehicle's speed-sensor, traction-control, and stability-control systems.

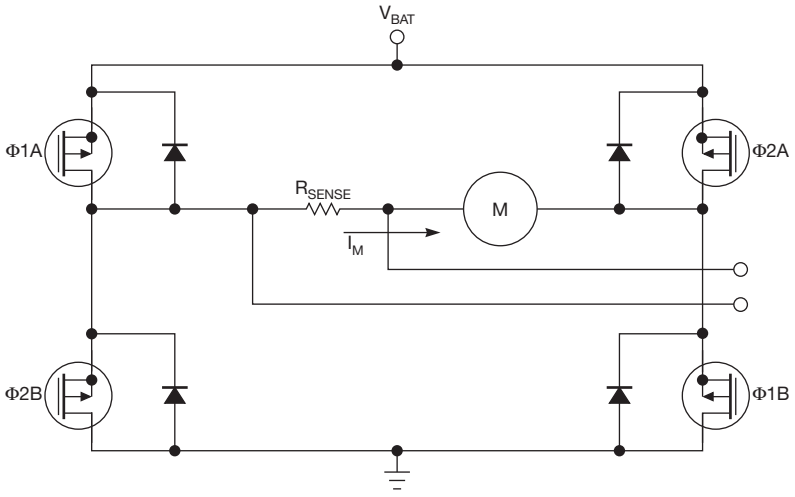

Figure 2 The phase of the signals to the four gates of this H-bridge circuit determines motor direction and speed.

and ground on the bottom of the dc bus. For high-side current sensing, you put the sense resistor between the positive battery terminal and the H bridge on the high side of the dc bus. You can perform output PWM current sensing on the motor itself.

These alternatives require trade-offs. The low-side approach is convenient but adds undesirable resistance in the ground path. It also lacks the diagnostic ability to detect a short-to-ground fault. Neither the high-side nor the low-side approach lets you continuously monitor current in the recirculation diode. PWM current sensing does allow sensing of the diode current and adds no resistance to the ground path.

A PWM current-measuring circuit entails performance constraints that are far from trivial. The circuit must contend with common-mode voltages that swing all the way from ground to the battery voltage. Thus, to reject common-mode excursions, the circuit must have not only a high range of input voltage corresponding to this swing but also an excellent CMRR (common-mode-rejection ratio) at the switching frequency and at relevant edge-rate-induced frequencies.

Common-mode transients and a minimum duty cycle for the PWM signal also impose a stringent requirement for settling time in the current-sense amplifier. For an accurate and linear response, the current-measurement circuitry must have high gain, high accuracy, and low offset voltage. Because human steering is part of the control loop, linearity and accuracy are especially critical. Any nonlinear-

TABLE 1 H-BRIDGE TRUTH TABLE

Φ1A	Φ1B	Φ2A	Φ2B	Condition
On	On	Off	Off	The motor is powered between the battery voltage and ground; its current increases and flows as the arrow indicates
Off	On	Off	On	The current flows as the arrow shows but decreases and circulates through Φ1B, Φ2B, and the sense resistor
Off	Off	On	On	The motor is powered between the battery voltage and ground; its current increases and flows in a direction opposite to the arrow
Off	On	Off	On	The current flows in a direction opposite to the arrow and decreases and circulates through Φ2B, Φ1B, and the sense resistor

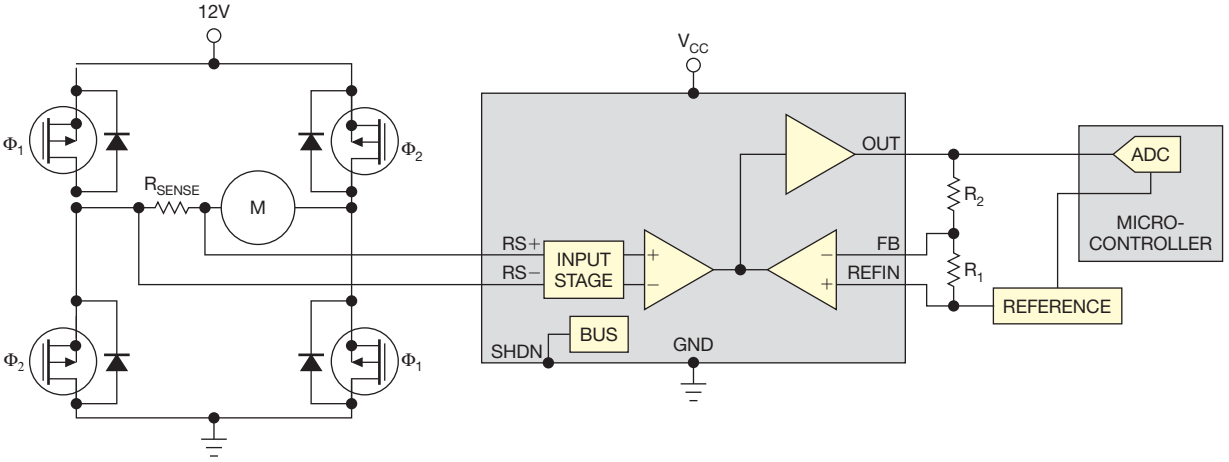

Figure 3 You can construct a PWM-compatible, H-bridge current-sensing circuit.

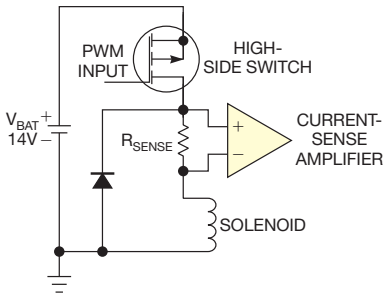

Figure 4 This typical solenoid drive circuit employs a high-side switch.

ity in the circuitry can impair the driving experience by causing oscillation or vibration when you oversteer the vehicle.

You connect the motor in a full H-bridge configuration (**Figure 3**). This approach lets you reverse the motor-applied voltage polarity, which enables it to rotate in either direction. The IC withstands –20 to +75V common-mode voltages, making it immune to inductive flyback, load-dump transients, and reverse-battery faults. The device also integrates an instrumentation amplifier. Its indirect current-feedback architecture provides precision current sensing with a maximum input offset voltage of 400 µV and a gain error of 0.6%. The external reference voltage supports the bidirectional current sensing that a full H bridge requires and senses unidirectional current when operating with a half H-bridge circuit. In a bidirectional system, the output voltage equals the reference voltage when the sensed current is 0A. Both adjustable- and fixed-gain flavors enable this part to provide maximum flexibility over a range of applications.

SOLENOID-DRIVE SENSING

Solenoids find wide use as electromechanical switches in vehicles. A starter solenoid, for example, delivers a large electric current to the starter motor, which in turn sets the engine in motion. However, several automotive-control systems employ a solenoid drive for precision control. For example, a diesel-engine system for railroads relies on solenoids as sophisticated electronic-control valves. They inject the appropriate quantities of fuel directly into the engine cylinders at high pressure. You accurately control the timing of these valves with the engine-control unit to ensure synchronization with the diesel engine. The result is a relatively "green" engine that makes less noise, produces fewer emissions, and is more fuel-efficient. Other applications for solenoid control include automatic gear shifting, transmission control, braking control, and active suspensions.

THE CHALLENGES OF SENSING SOLENOID CURRENT FOR REGULATING PWM FREQUENCY ARE SIMILAR TO THOSE IN THE H-BRIDGE APPLICATION.

The high-side switch is typically a FET whose gate you control with a PWM signal (**Figure 4**). When the FET is on, it connects the solenoid to the 14V battery voltage, producing a current that charges the solenoid coil. When the FET turns off, solenoid current discharges through the clamp diode and shunt resistor. By regulating the PWM frequency and duty cycle, you determine the resulting average ripple current in the solenoid. This current in turn controls the force the system applies to the actuator.

The challenges of sensing solenoid current for regulating PWM frequency and duty cycle are similar to those in the H-bridge application. Common-mode voltages at the input of the current-sense amplifier range from the battery voltage to the slightly negative level of the clamp diode's drop voltage. Typical solenoids require a few amperes of current, so a clamp diode that withstands such current may develop a forward voltage higher than 1V.

Again, the current-sense amplifier's wide input common-mode range and fast settling in response to common-mode variations suit this application. The main difference between this application and that of the H bridge is that the solenoid current always flows in the same direction; the current-sense amplifier, therefore, needs only to be unidirectional. The IC becomes a unidirectional current-sense amplifier when you connect its reference input to ground.

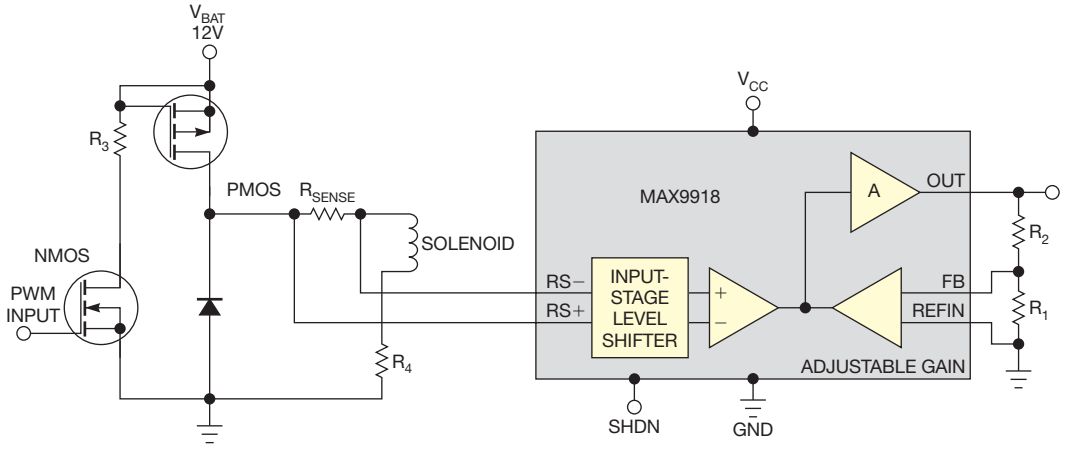

Figure 5 Engineers prototyped this solenoid-drive circuit in the lab.

Introducing EDN's New Microsite on LEDs

Your complete source for up-to-the-minute information on LEDs and their applications.

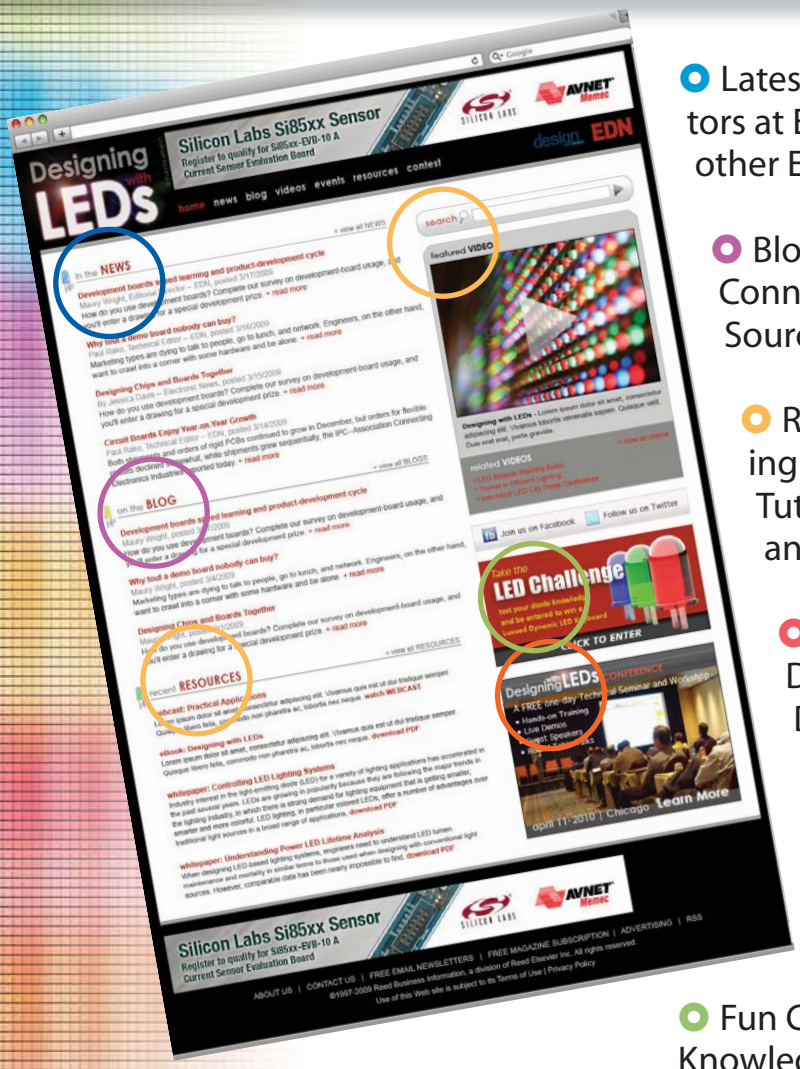

- Latest News from the Editors at EDN, Design News and other Engineering Magazines
- Blogs, including Margery Conner's Popular "Power-Source" Blog
- Resource Center, Including White Papers, eBooks, Tutorials, Webcasts and More
- Videos, Including Design Tips, Application Demos, and More
- Updates on Future In-Person Workshops and Seminars, like this one!
- Fun Challenges to Test Your Knowledge of LEDs

Check It Out Today at www.edn.com/leds

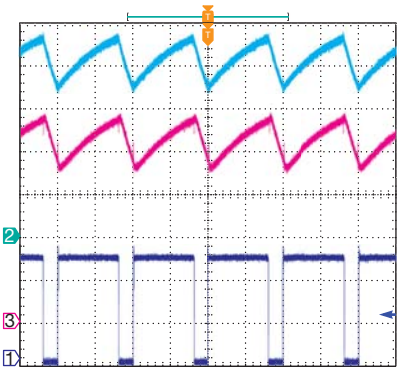

Figure 6 Waveforms at a PWM frequency of 5 kHz illustrate the operation of the circuit in Figure 5, with a duty cycle of 80%. The top trace is the voltage across R_4, the middle is the output of the current-sense amplifier, and the bottom is the PWM signal at the bottom of the PFET.

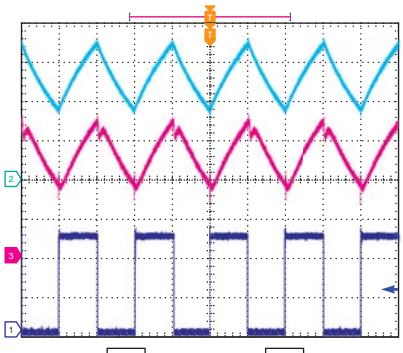

Figure 7 Waveforms at a PWM frequency of 5 kHz illustrate the operation of the circuit in Figure 5, with a duty cycle of 50%. The top trace is the voltage across R_4, the middle is the output of the current-sense amplifier, and the bottom is the PWM signal at the bottom of the PFET.

LAB RESULTS

You can prototype a typical application circuit for solenoids in the lab (**Figure 5**). You emulate the solenoid with a 2-mH inductor having a low ESR (equivalent series resistance) of 1.6Ω. The sense resistor is 100 mΩ, and a value for R_4, not present in the actual solenoid circuit, of 15Ω limits the maximum solenoid current, as the following **equation** shows:

$$I_{MAX}=V_{BAT}/(R_{SENSE}+ESR+R_4)= 12V/(0.1+1.6+15)\Omega=0.72A.$$

This maximum current value is the theoretical limit the circuit reaches when the inductor is fully charged. The resistor and inductor values set the circuit's time constant to approximately 0.12 msec, which is equivalent to 8.3 kHz. You set a gain of 80 with the external resistors R_1 and R_2 having values of 1 and 79 kΩ, respectively.

Waveforms at a PWM frequency of 5 kHz illustrate the operation of the circuit in **Figure 5**, with duty cycles of 80 and 50% (**figures 6** and **7**, respectively). The top waveform is the voltage across R_4, which is proportional to the current flowing in the inductor. The middle waveform is the output of the current-sense amplifier, and the bottom waveform shows the PWM signal at the drain of the PFET. Higher duty cycles yield higher current.

Thus, a precision, high-voltage, high-side current-sense amplifier allows accurate measurement with smaller sense resistors. It handles the bidirectional motor currents that you derive from H bridges, such as those in EPS systems, as well as the unidirectional solenoid currents in automatic gear shifting, transmission control, braking control, and active suspensions. EDN

AUTHORS' BIOGRAPHIES

Maurizio Gavardoni is a product definer for amplifiers and comparators in the multimedia-business unit of Maxim Integrated Products. Before joining Maxim, he worked as a mixed-signal-system designer for several automated-test-equipment companies. He holds bachelor's and master's degrees in electrical engineering from the University of Milan (Milan, Italy).

Akshay Bhat is a strategic-applications engineer with Maxim Integrated Products. He holds a master's degree in electrical engineering from San Jose State University (San Jose, CA) and a bachelor's degree in electrical engineering from the University of Mumbai (Mumbai, India). Previously, Bhat worked at Larsen and Toubro Ltd, where he assisted the Indian Space Research Organization in the testing and development of communication satellites.

Energy Harvesting Now

Free Power from Thermal, Kinetic & Solar Energy

Our new analog IC solutions enable the commercial deployment of energy harvesting from a variety of "free" energy sources. An appropriate transducer placed on the energy source delivers an electrical signal that our products convert and condition into usable power. These revolutionary ICs consume only nanoamps of current to provide high efficiency power conversion with minimal external components.

Energy Harvesting IC Family

Part Number	Description	Energy Source
LTC®3105	400mA boost converter with MPP control and 250mV start-up	
LTC3108	Ultralow voltage boost converter and system manager	
LTC3109	Auto-polarity version of LTC3108	
LTC3588	Piezoelectric energy harvesting power supply	
LT®3652/HV	Power tracking 2A solar battery charger	
LTC4070	Nanoamp operating current shunt Li-Ion battery charger	

Info & Free Samples

www.linear.com/energyharvesting

1-800-4-LINEAR

LINEAR TECHNOLOGY

EDITED BY MARTIN ROWE
AND FRAN GRANVILLE

designideas

READERS SOLVE DESIGN PROBLEMS

Design provides single-port-to-dual-port SDRAM converter

Yu-Chieh Chen, Instrument Technology Research Center, National Applied Research Laboratories, Hsinchu, Taiwan

In multimedia and video applications, DRAM buffers large frames in the image coding and decoding between an image-sensing device and an image-display device (**references 1** and **2**). It also plays a key role in image-frame-rate converters by matching the displaying format of the monitor (**Reference 3**). Accessing dual-port SDRAM is a complex task, especially when two data read or write operations are occurring simultaneously into a single-port SDRAM during high-speed operation.

SDRAM supports asynchronous control, performs reading and writing operations, and allows multiple CPUs to read or write in parallel. To reduce production costs and avoid complexity, engineers typically prefer conventional one-port SDRAM. You can, however, use an FPGA to convert a single-port SDRAM to function as a dual-port SDRAM.

The single-port-to-dual-port SDRAM converter comprises two read-FIFO (first-in/first-out) and write-FIFO memory blocks. A pair of read and write controllers each control a read-and-write data bus. **Figure 1** shows the architecture of the proposed converter, and **Figure 2** shows the operating flow of the read-and-write controller. The write controller first stores the writing data in write-FIFO and executes the SDRAM write command when the WKO (write-keep-out) signal falls to a low level, which means that the read operation is complete. Meanwhile, the RKO (read-keep-out) signal goes high until the write operation is complete. The procedure prevents the read signal from occurring during the writing operation.

The read-data operation is similar to the write-data operation. The read controller first waits for the RKO signal to fall low if it is high when the read operation starts. The read controller then executes the read command and stores the data from SDRAM to read-FIFO memory. The read controller simultaneously pulls the WKO signal high until the read operation is complete. **Figure 3** shows the hardware logic of the smart controlling strategy, which the following **equation** also describes:

$$F_{COM} = [MUX_{READCOM,WRITECOM} \times \overline{RKO} \times \overline{WKO}],$$

where MUX is the multiplexer. For the read example, the SEL (select) signal switches to the read mode and then waits for the RKO signal to fall low if the RKO signal is at a high level when the read command executes (**Reference 1**).

This strategy doubles the high-level period of both the read operation, or RKO signal, and write operation, or WKO signal. These signals rise to logic high before executing each read or write operation. This step ensures that the read operation does not interrupt the write operation and that the write operation does not interrupt the read operation.

Figure 4, available with the Web version of this Design Idea at www.edn.com/110303dia, shows the simulation waveform of the converter, which is a part of the image-sensing and -displaying sequence. The iSDRAM_Read signal is high, indicating that the SDRAM is performing a read operation. At the same time, the iSDRAM_Write signal is high during the SDRAM-writing op-

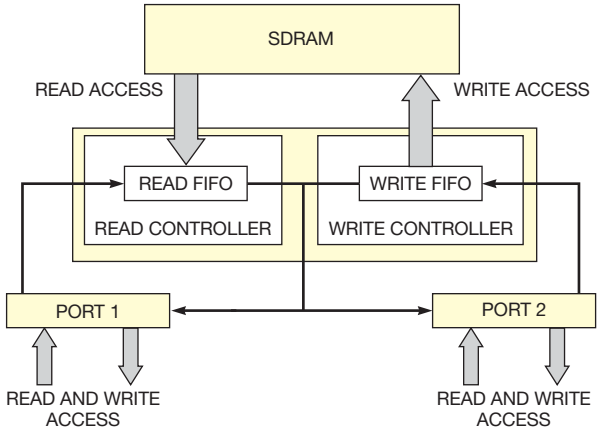

Figure 1 The converter architecture uses two ports for separate read and write operations.

DIs Inside

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

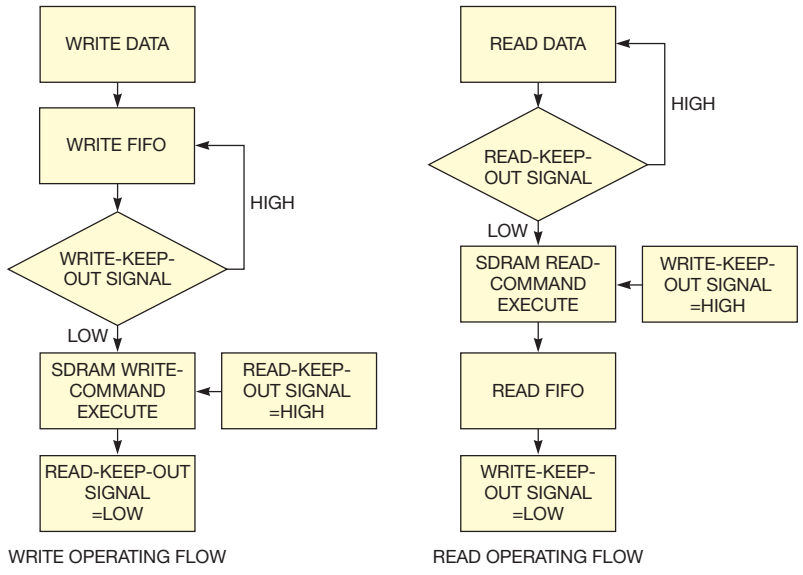

Figure 2 Read and write operations occur when the RKO and WKO signals are at logic-low states.

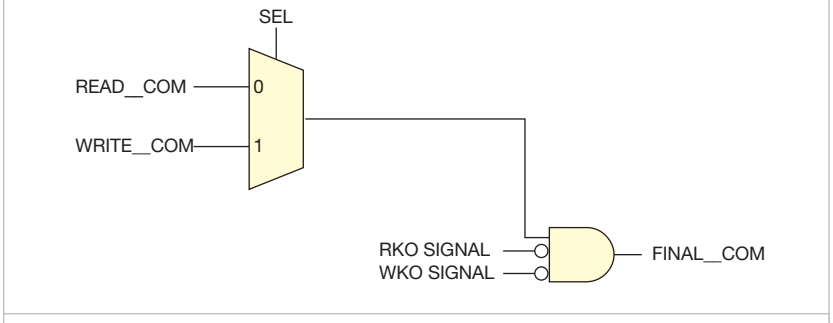

Figure 3 The hardware logic uses a multiplexer to select write or read signals.

eration. The iSDRAM_ReadKeepOut signal goes high before each iSDRAM_Write signal and lasts two times longer than each iSDRAM_Write signal. This setup ensures that the writing command does not execute during the read operation. For the same reason, the read command doesn't execute during the writing operation.

The write-clock operation is initially ahead of the read-clock operation. After several time delays, however, the write-clock operation catches up with the read-clock operation. At this point, the controller holds the data and waits until the iSDRAM_ReadKeepOut command falls to a low level before executing the read command. Using this read-and-write strategy, the controller resolves the conflict of executing read and write commands during write or read operations. EDN

REFERENCES

1 Zhu, Jiayi; Peilin Liu; and Dajiang Zhou, "An SDRAM controller optimized for high definition video coding application," *Proceedings of the IEEE 2008 International Symposium on Circuits and Systems*, May 2008, pg 3518, http://bit.ly/h6xaz0.
2 Jack, Keith, *Video Demystified*, Fifth Edition, Newnes Publishers, 2007, ISBN: 0750683953.
3 Yim, Myung-Sik, "Image decoding apparatus having frame rate transformation function," US Patent No. 5917950, June 29, 1999, http://bit.ly/i4GR98.

Zener diode protects FPGA inputs

Rick Collins, Arius, Frederick, MD

Although 5V-powered logic is still common in many applications, most FPGAs support interface levels of 3.3V and lower. When you connect an FPGA to higher voltage levels, the FPGA's application notes commonly suggest that you use the PCI (Peripheral Component Interconnect)-bus clamp diodes in the FPGA's I/O blocks with an external series-limiting resistor to prevent damage to the FPGA (**Figure 1**). The PCI clamp diode limits the voltage to a level that doesn't harm the input, and the resistance value limits the current to a safe level that doesn't harm the PCI clamp diode. This approach works well in designs with low-speed signals.

However, when you use this approach with a higher-rate signal, the effects of the parasitic RC filter distort the signal

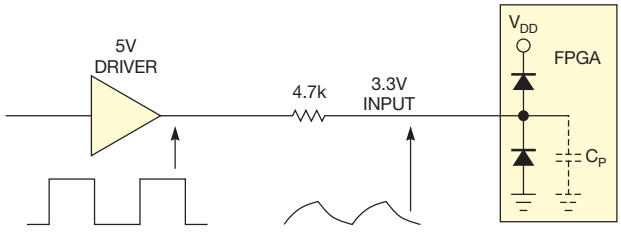

Figure 1 Engineers often use a series resistor to interface 5V logic levels to FPGAs.

Breakthrough image quality *and* low power in a compact solution!

8-channel front-end plus ADC receiver solution provides the highest level of performance at low power for all ultrasound imaging modalities

MAX2078

65mW/ch, octal front-end with CWD beamformer

+

MAX1437B

96mW/ch, octal, 12-bit, 50Msps ADC with serial outputs

- Up to 4dB better SNR than closest competitor
- Enhanced 2D image quality
- Improved sensitivity in pulsed Doppler and color flow imaging

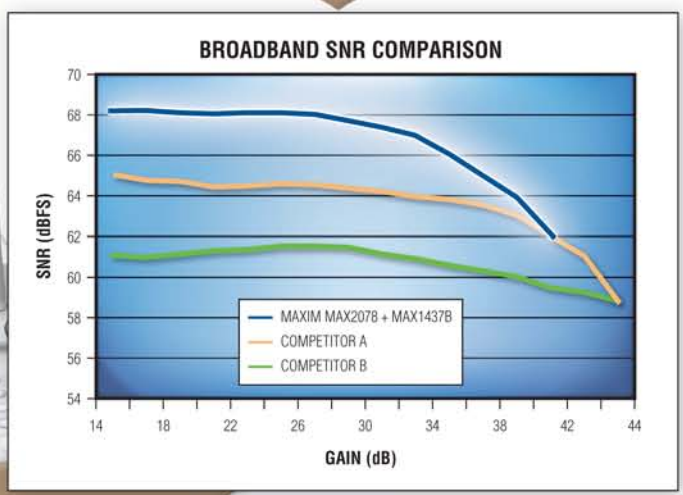

Available with or without an integrated continuous-wave Doppler (CWD) mixer-based beamformer

Part	Function	Channels
MAX2077 + MAX1437B	Ultrasound front-end + 12-bit, 50Msps ADC with serial LVDS outputs	8
MAX2078 + MAX1437B	Ultrasound front-end with CWD mixer-based beamformer + 12-bit, 50Msps ADC with serial LVDS outputs	8

www.maxim-ic.com/Ultrasound

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.

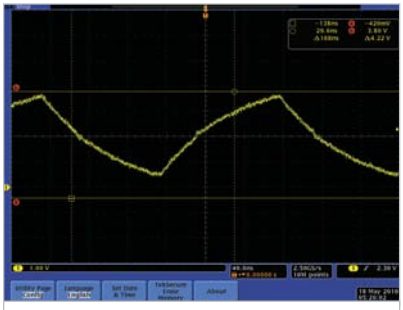

Figure 2 A parasitic RC filter slows the signal's edges.

(**Figure 2**). The circuit from the FPGA's application notes requires a change, which you can accomplish without redesigning the PCB (printed-circuit board). In this case, substituting a zener diode for the resistor shifts the signal level without causing excessive distortion (**Figure 3**). The zener diode works with the PCI clamp diode and the internal pulldown resistor to set the voltage level at the input pin.

To set the static level at the input pin, you must enable the FPGA's internal pulldown resistor to prevent the PCI clamp diodes from being driven too hard when the input is continuously high. The current from the pulldown resistor is smaller than the rating current of the zener diode. Low-voltage zener diodes also have round "knees" in the avalanche IV (current-to-voltage) curve. This curve results in a zener voltage that's lower than the rated value, so you need to use a higher-voltage zener diode. The diode should also have a low capacitance. The CZRU52C3, a 3V zener diode from Comchip (www.comchip.tw.com), works well, reducing the circuit's voltage by 2V (**Figure 4**).

Some parasitic effects in the zener diode will create other distortion to the waveform. The parasitic capacitance of the diode causes the diode to initially look like a short to the signal edges from the 5V driver. The FPGA pin will see a high-voltage overshoot for approximately 10 nsec, quickly decaying to the rated level of the input pin. The RC time constant of the pin capacitance and the pulldown resistance result in a slower drop to the final value, which the zener diode and the pulldown resistance determine. **Figure 5** shows a detailed view of the leading edge. EDN

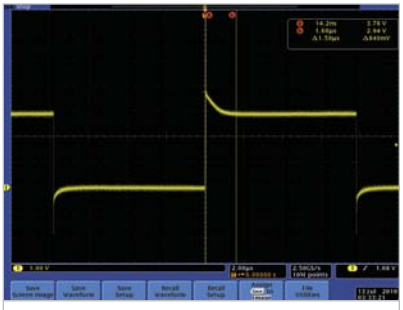

Figure 4 An oscilloscope trace shows the signal with overshoot.

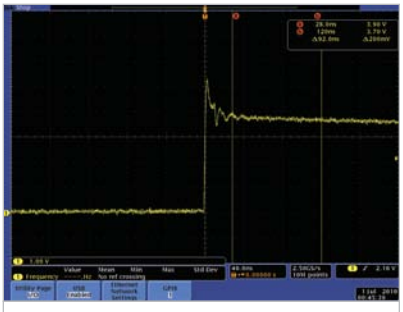

Figure 5 The brief overshoot of the leading edge of the pulse does not damage the FPGA inputs.

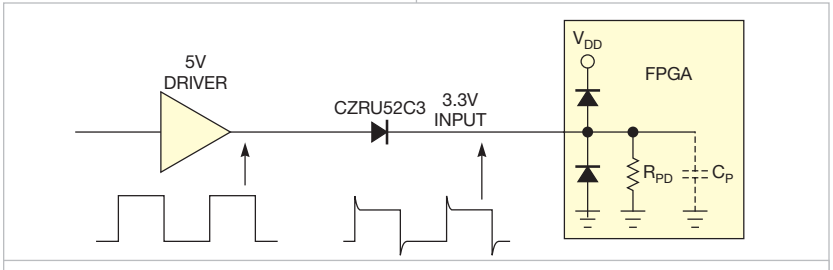

Figure 3 Replacing the resistor with a zener diode shifts the voltage level with little distortion.

Bicolor LED indicates 10 states

Abel Raynus, Armatron International, Malden, MA

Most microcontroller-based devices have several states of operation. Pushbutton switches often control these states. To minimize costs, many designs use only one switch; the number of presses sets the microcontroller's state. In the circuit in this Design Idea, lighting an LED indicates the chosen state.

A large selection of bicolor LEDs is on the market in various shapes and colors from Kingbright (www.kingbrightusa.com), Optosupply (www.optosupply.com), and other suppliers. These LEDs provide new opportunities to use just one LED to indicate 10 states. You set the states by pressing a pushbutton switch that connects to a microcontroller's external interrupt pin (**Figure 1**). After each push, the interrupt subroutine sets the device to the next state (**Figure 2**).

State 0 is standby mode with the LED off. The device waits for you to press a switch. The LED's color indicates the next three states. For example, a Kingbright WP5A9EGW12.3SF LED yields red, green, and orange. Add

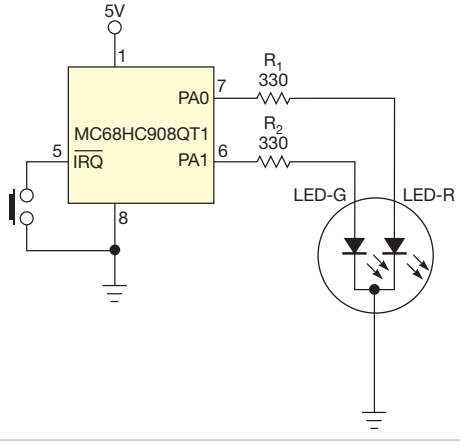

Figure 1 A bicolor LED and a pushbutton switch are all you need to set and indicate a microcontroller's state.

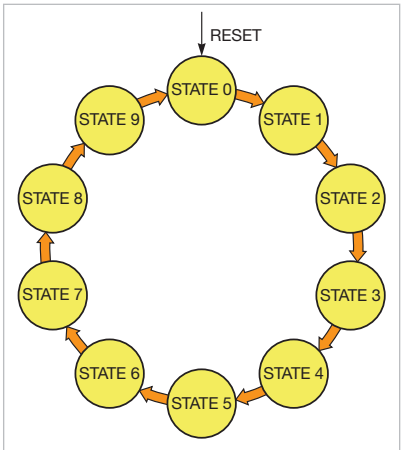

Figure 2 Pressing the switch advances the microcontroller's state.

TABLE 1 COLOR AND BLINKING PATTERNS FOR 10 STATES

State	Red LED	Green LED	Indication
0	Off	Off	No light (standby)
1	On	Off	Red
2	Off	On	Green
3	On	On	Orange
4	Blinks	Off	Red blinks
5	Off	Blinks	Green blinks
6	Blinks	Blinks	Orange blinks
7	On	Blinks	Red and orange blink
8	Blinks	On	Green and orange blink
9	Blinks	Blinks	Red and green blink

blinking to indicate the remaining states. You can independently control the red and the green LEDs, yielding six more combinations of indication (**Table 1**).

Using the microcontroller's internal oscillator as a clock source simplifies the circuit because the oscillator needs no external components and its 5% tolerance is sufficient for this application. Firmware sets the built-in pullup resistor for external interrupt input.

The clock speed and the values of the prescaler and time-counter-modulo registers determine the rate of LED blinking. The internal oscillator generates a 12.8-MHz clock, resulting in a 3.2-MHz bus frequency, with one cycle equal to 0.3125 µsec. By choosing a prescaler value of 64, a time-counter-modulo register can calculate a cycle of 50 times the timer period in milliseconds. For example, for 1 second of blinking, you should set the time-counter modulo to 50,000, or $C350 in hexadecimal.

Using tables rather than polling in programming allows you to significantly reduce the size of firmware (**Reference 1**).You can download firmware assembly code from the online version of this Design Idea at www.edn.com/110303dib. This method is applicable to any microcontroller. You can even increase the number of indicated states by using different blinking rates. EDN

REFERENCE

1 Raynus, Abel, "Tables ease microcontroller programming," *EDN*, April 22, 2010, pg 76, http://bit.ly/brXaq7.

Relay driver switches two relays with one pin

Gunther Kraut, PhD, Egmating, Germany

Independently switching two loads, such as relays, usually requires two microcontroller-I/O ports. Two ports let you control the loads so that you can switch them both off, switch them both on, or switch one on and the other off. If you don't need both relays on at once, you can control the remaining three states using only one port. You can also use a port's high-impedance state to control the relays.

You can, for example, use this application to control a motor. The motor's rotational direction depends on which of its two phase lines you select. This approach works with either classic electromechanical relays or SSRs (solid-state relays). If both relays are open, the motor is off. For electromechanical relays, you can use the driver circuit in **Figure 1**. Setting the microcontroller pin to logic one causes Q_1 to turn on relay REL_1, and the motor runs forward. Setting the pin to a logic zero turns on Q_2, which turns on Q_3. This action closes the switch of REL_2 and causes the motor to run in the opposite direction. If the microcontroller port is in a high-impedance state, Q_1, Q_2, and Q_3 cannot deliver current because the constant voltage of 1V at the base of Q_2 is too low to reach the threshold of the base-emitter junctions of Q_1 and Q_2 plus the diode drop of D_1. Both relays are off, and the motor is off. The supply voltage through a voltage divider or an emitter-follower circuit can provide the constant voltage of 1V. Free-wheeling diodes D_2 and D_3 prevent the collectors of Q_1 and Q_2 from overvoltage when the relay's coils are off. You can use almost any small-signal NPN and PNP transistors in this circuit. The selection of diode D_1 is also not critical.

The circuit is simpler for SSRs because virtually all microcontroller ports directly drive the LEDs (**Figure 2**). A logic one supplies the LED in S_1, whereas a logic zero supplies the LED in S_2, making the corresponding TRIAC (triode-alternating-current) switches conductive. If the port is in the high-impedance state, no current can flow through the LEDs because the constant dc voltage of 1.2V is below the sum of the threshold voltages of the LEDs. Voltage-dependent resistors R_3 and R_5 and snubber circuits C_1, R_4, C_2, and R_6 protect the SSRs. These protection circuits' values depend on the respective load. EDN

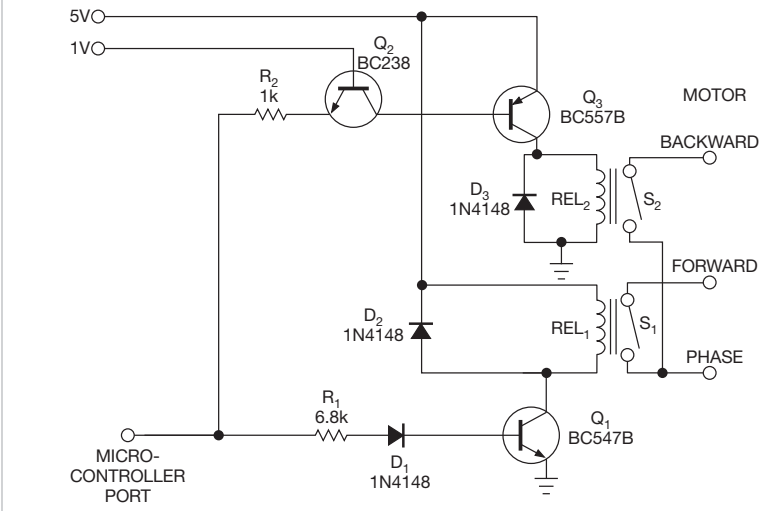

Figure 1 Relays respond to logic one, logic zero, and the high-impedance states of the microcontroller port.

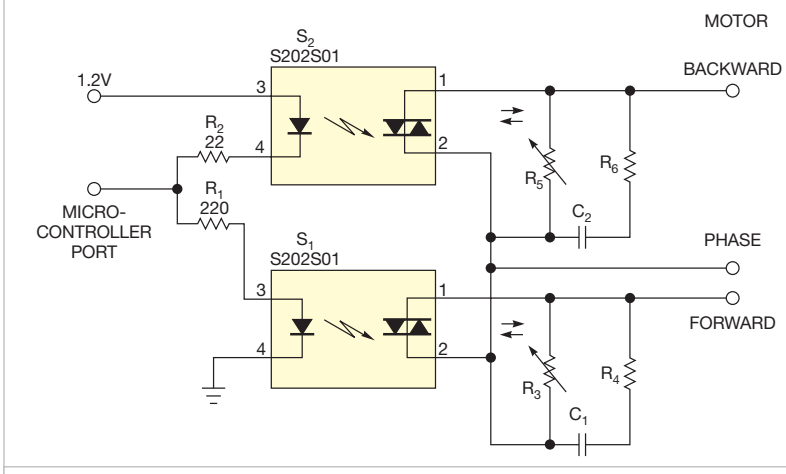

Figure 2 The circuit is simpler for SSRs because virtually all microcontroller ports directly drive the LEDs.

Counterfeiting continues to grow, but the industry fights back

The problem of counterfeit components in the electronics supply chain continues despite increased law enforcement. The problem increased last year because of the lengthening leadtimes on many parts. To ensure continued production, many OEMs turned to the open market, in which counterfeit parts are plentiful. Yet, even as the flow of counterfeit parts increases, players in the industry have adopted mitigation strategies.

"When the market gets tight, counterfeiting goes up," says Ed Smith (**photo**, left), president of Avnet Electronics Marketing Americas (www.avnet.com). "Customers miss the forecast, and they start buying from unauthorized outlets. When the leadtimes come into line and the market stabilizes, counterfeiting comes down." In the past year, though, counterfeiting has become more pronounced in both the ups and the downs of the market. "Counterfeiting is higher than it was during the last stable period," says Smith. "The counterfeiting is worse than the last good times, and it's also worse than the last bad times."

Counterfeiters clean and remark scrap parts. "Some of the counterfeiters remove used parts from electronic scrap," says P Steve Gregory, president and chief executive officer of PC Components (www.pccomponents.com). "They clean

up and resurface the parts so they look like new; they then remark the parts with identical date codes and lot numbers so they can pass them off as new."

Some counterfeiters are manufacturing their own parts. The manufacturing of parts sometimes occurs during downtime in the Asian plants where the original contracts occurred. "These counterfeiting manufacturers are getting better," says Gregory. "It used to be just generic capacitors, but now they're manufacturing more complex parts."

China is still the culprit in much of the electronics counterfeiting, but that situation could change if the Chinese government cracks down on counterfeiters. "Most of the counterfeit parts are still coming from China," says Mark Snyder, president of ERAI (www.erai.com), which offers risk-management services. "We expected to see an insurgence from India, but it hasn't happened yet."

Cracking down on China won't necessarily curb the problem, he adds. "As we work with the government to get a handle on the counterfeit parts coming out of China, it will just go to another part of the world."

One of the fears in the electronics industry is that counterfeit parts could find their way into the DOD (Department of Defense) or NASA (National Aeronautics and Space Administration) and compromise weapons or space systems, with catastrophic results. DOD and NASA are driving a lot of the law enforcement, according to Robin B Gray, Jr (**photo**, right), president and chief executive officer of ECIA (Electronic Components Industry Association, www.eciaonline.org). "They have concerns about counterfeit products entering their supply chain," he says. "They seem to be on the verge of either regulation or maybe even legislation that would mandate that, whenever possible, manufacturers must buy components from an authorized source."

The recent practice in the military of buying outside authorized channels has heightened concern about counterfeit parts in the military-supply chain. "The military should be nervous," says Avnet's Smith. "Your risk is negligible if you buy everything from an authorized distributor. The military used to buy only from authorized distributors. Now, they do some contracting at the lowest cost. Sometimes they don't have a good forecast, so sometimes they're buying outside the authorized market."

Law enforcement in this area has improved, and federal law-enforcement agents are arresting counterfeiters. "[The Customs Department] is beginning to develop policies, and those policies are working," says Smith. "They are making us sign agreements that we buy from the manufacturer and not the third parties, and they're asking us to test the products to make sure they come from the OEMs."

Companies are reluctant to reveal that they've run into counterfeit parts because it may be bad for business. "Law enforcement is doing a reasonable job, but they have to deal with companies that don't want to expose the fact that they bought counterfeit parts," says Smith.

Although counterfeiting is still a major problem in the electronics industry, vigilance and improved law enforcement are helping to curb the damage. "Counterfeiting is not going away anytime soon," says ERAI's Snyder. "Everyone in the supply chain is doing a better job. When we find a real counterfeiter, we send it to the federal authorities, and they chase it down."

—by Rob Spiegel

productroundup

POWER SOURCES

500W power supplies target medical, industrial applications

The single-output, 360 to 500W CSS500 series of power supplies accepts a wide ac-input range; has a low profile of less than 1U; and targets use in medical, dental, and industrial-automation applications. The power-factor-corrected supplies feature a universal 90 to 264V-ac, 47- to 63-Hz input, allowing worldwide use. They also feature a leakage current of less than 300 µA at 264V ac, 63 Hz, which is important for many medical and dental applications. The CSS500 withstands 4 kV ac, input to output. Output voltages are 12, 24, 30, 36, 48, 54, or 57V dc, and power-saving efficiency is as high as 92%. With convection cooling, requiring no fans, these supplies provide as much as 360W of output power. For applications requiring higher power, with forced-air cooling of 30 cfm or 3.1m/sec, output power increases to 500W. The convection-cooled version has a U-channel chassis that measures 4.7×8×1.5 in. The devices sell for $175 (100) each.

TDK-Lambda Americas, www.us.tdk-lambda.com/lp/products/css-series.htm

70W LED-lighting supplies support more-than-70,000-hour lifetime

The 70W LDS70 LED-lighting standard power supplies support lifetime requirements greater than 70,000 hours. The devices provide constant-current and constant-voltage modes. Constant-current mode features programmable current levels; constant-voltage mode has no-load to full-load operation. Both 12 and 58V models have no minimum-load requirement and have comprehensive short-circuit, overvoltage, and overtemperature protection. The devices' input-voltage range is 100 to 240V ac for U models and 120 to 277V ac for H models, and active power-factor correction is 0.9. Prices for the LDS70 series start at $79 (50).

Emerson Network Power, www.emersonnetworkpower.com

Encapsulated ac/dc converters have universal input of 90 to 264V ac

The BP series of encapsulated ac/dc converters operates from –40 to +70°C at half load and –40 to +60°C at full load with convection cooling. The devices are available in 5, 10, 15, and 30W versions with universal input of 90 to 264V ac and regulated output of 3.3, 5, 12, 15, 24, or 48V. The devices come in PCB-, chassis-, and DIN-rail-mountable packages as small as 1.75×2.25×0.75 in. Prices range from $59 to $139 (single units).

Bear Power Supplies, www.bearpwr.com

Six-output dc/dc µModule regulator operates from –40 to +150°C

The LTM8008 dc/dc µModule regulator guarantees operation from –40 to +150°C. It includes a 3 to 72V-input onboard SEPIC with buck- and boost-conversion capability and six linear outputs in a 15×15×2.82-mm LGA package. The device regulates a 3 to 72V input supply to 5.6V output, which powers the regulators to generate five 5V rails and one 3.3V rail with output currents of 150 to 500 mA. Other features include a programmable operating frequency of 100 kHz to 1 MHz and programmable soft-start. The LTM8008 sells for $19 (1000) each.

Linear Technology Corp, www.linear.com/8008

EDN ADVERTISER INDEX

Company	Page
Agilent Technologies	4
Anaren Microwave	25
austriamicrosystems AG	49
Digi-Key Corp	C-1, 5
Interconnect Systems Inc	23
International Rectifier Corp	C-1, C-2, C-3, C-4, 9
Ironwood Electronics	53
Lattice Semiconductor	6
Linear Technology	37, 38, 44
Maxim Integrated Products	47
Micrel Semiconductor	19
MicroPower Direct	35
Mill-Max Manufacturing Corp	11
Pico Electronics Inc	24, 43
Renesas Technology Corp	8
Signal Consulting Inc	21
Stanford Research Systems Inc	33
UBM Canon Trade Events	27
Vicor Corp	17

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

EDN productmart

This advertising is for new and current products.

Cheaper to keep 'er

I was working at a semiconductor manufacturer, and we had just designed an evaluation platform that showed the excellent performance of our new state-of-the-art product. This platform included the board with a power switch and LED indicators in a case, a power brick, and the software to run the board. We had done thorough worst-case design analyses, reliability studies, power-cycling studies, and all the rest of what you normally do for new products. We found no problems, the design got the go-ahead, and we put it into production.

A few weeks later, someone on our team discovered a less expensive capacitor with the same ratings as the original, and we made a BOM (bill-of-materials) change after completing the additional evaluation process on the new part. Within a couple of months of the change, though, a customer returned one of these platforms with a charred capacitor. You guessed it: It was the replacement part. We assumed that this was just a statistical failure and continued to use the replacement capacitors with no further thought.

Two or three days later, though, we received another return with the same problem. At this point, we began to suspect that the problem was more than one of statistical failure. The 47-µF capacitor was rated at 16V and had 12V impressed upon it. There were four other identical capacitors on the board that seemed to be fine. While we were puzzling over the problem, the same problem occurred on one of the boards we kept in the lab.

We began to wonder what could cause the capacitor in this position to fail while identical capacitors in other positions did not. After a few days, we found that another capacitor in the same position on another laboratory board failed as we turned on the power switch on the board. The technician said that the other lab failure also happened at the power switch's turn-on. We now wondered what could cause a problem only upon application of power. By this time, we had another return and a crisis on our hands.

Thinking about the power arrangement, we realized that we had shipped a power brick with the boards. The power brick was generally on when the power switch on the board was on, meaning that the filter capacitor of the power brick was fully charged to 12V. Could it be that the surge current, upon turning on the switch on the board, could destroy the capacitor?

We looked at the specifications of the original capacitor and the current capacitor for handling surge currents and found that neither one had a surge rating. But why didn't the power-cycling testing show any problems?

After much thought and discussion, we realized that the power FETs we used to turn power on and off in the power-cycling testing have a finite on-resistance that would effectively limit the surge current when we applied power. This could account for the lack of problems we found in the power-cycling tests. Each capacitor type, having no surge-current rating, was from a different manufacturer, meaning that the capacitors could reasonably have different capabilities for handling surge currents. Furthermore, we had no guarantee that the original capacitor type, without a surge-current rating, would never have a problem. The surge-current tolerance could well depend upon individual capacitors' manufacturing lots.

We added a 1Ω wire-wound resistor in series with the board input to limit the inrush current that occurred upon turning on the power switch on the board. Afterwards, we never again encountered the problem. When building the platform, we had considered using a capacitor that had a reasonable surge rating, but such a capacitor would have cost more than a non-surge-rated capacitor plus the resistor. So we resolved a vexing problem and learned the value of anticipating a product's end use and the differences between automated testing and actual use. EDN

Nicholas Gray is an engineer in San Jose, CA.

DANIEL VASCONCELLOS

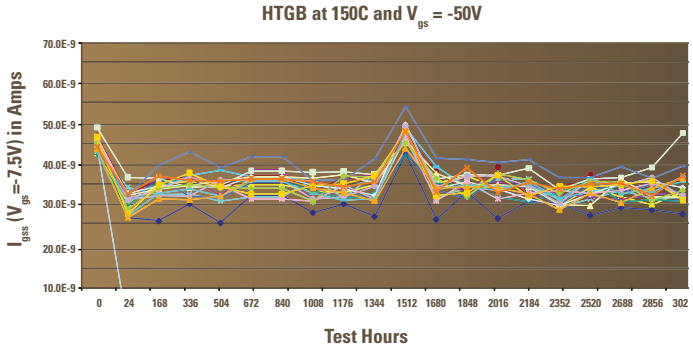

Figure 1: Long term stability of gate dielectric for first generation LV GaNpowIR HEMT, gate leakage measurement in nA, under accelerated stress conditions, Vgs=-50 V and 150ºC over 3,000 hours. Lg= 0.3 µm, Wg = 960 mm.

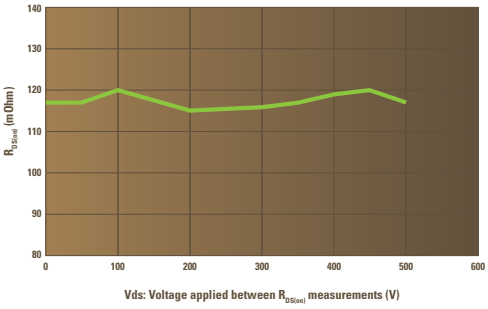

Figure 2 : Stability of on state resistance for first generation 600 V GaNpowIR HEMT prototypes after the application of reverse bias, Vds, measured within 1 µs of switching to on state. Measured at 25ºC.

in metal-semiconductor gated GaN-based HEMTs be eliminated. Commonly reported trapping related instability phenomenon such as current collapse or dynamic Rdson must likewise be minimized beyond concern. Figure 2 shows the Rdson measured within 1 µs of applying varying reverse bias conditions for early 600 V GaNpowIR prototypes. As can be seen, the commonly reported trapping effects have been effectively minimized in this platform.

It has been the resolution of the above issues that has occupied the energies of the team at International Rectifier over the past 7 years. It is with great satisfaction that these have been addressed sufficiently to support the targeted release of the first commercially available 600 V rated GaN-based power devices by the end of 2011, in agreement with the schedule announced in September 2008. The team of scientists and engineers at International Rectifier are to be congratulated on this remarkable achievement, which adds to the long heritage of innovation of the company.

The availability of cost effective, high quality, robust, high performance GaN-based power devices will enable truly innovative improvements in power electronic density, efficiencies and costs in the coming years. From more efficient solar panel based inverters to higher density efficient permanent magnet based motion control systems to lighter weight and denser inverters for electrified vehicles, as well as next generation integrated dc-dc power supplies for electronics, GaN-based power devices will revolutionize the industry. This technology platform strongly supports the objective to enable lower system costs to promote the adoption of efficient systems that significantly reduce worldwide power consumption. Figure 3 shows the significant advantages of GaN-based power stages in high frequency, dense dc-dc converters designed to support electronic loads such as microprocessors. Figure 4 shows a comparison between state of the art 600 V rated silicon trench IGBTs and prototype first generation 600 V rated GaNpowIR devices in terms of conduction*switching loss figure of merit. As can be seen, the GaN-based devices perform remarkably better. A potential order of magnitude further improvement in performance for GaN-based power devices is possible over the coming decade.

It is tempting to relate the introduction of these devices to that of power MOSFETs some 30 years ago, which revolutionized power electronics by enabling compact cost effective switching regulator based power supplies. However, it appears likely that this paradigm shift will have aspects of the greater character of the revolution in data processing which occurred through the development of large scale and very large scale integrated circuits. This is due to the intrinsically integratable nature of the lateral GaN-based device platform, allowing integration without performance compromise, a feat not possible in vertical silicon based power device technology. This will provide an entirely new level of performance, density, robustness and cost effectiveness.

It is indeed an exciting time in power electronics.

Dr, Michael A. Briere is an Executive Scietific Consultant at ACOO Enterprises LLC under contract to International Rectifier, where he has led the GaN effort since late 2004.

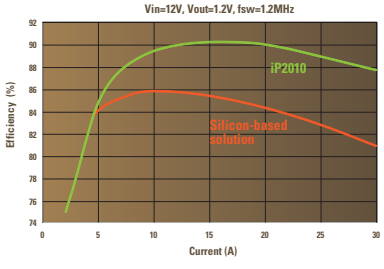

Figure 3: Measured conversion efficiency for 12 V to 1.2 V buck converters, including losses in driver and output inductor, for first generation low voltage GaNpowIR HEMT based power stage, iP2010 compared to results for state of the art silicon based solutions at a switching frequency of 1.2 MHz.

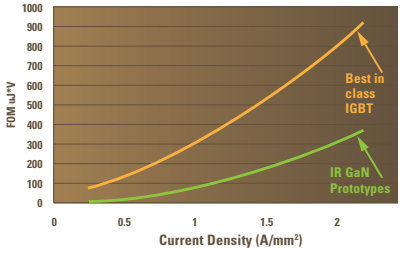

Figure 4: Measured performance figure of merit, Vds(on) * (Eon+Eoff), at 25ºC in uJ*V, for first generation 600V GaNpowIR prototype HEMTs compared to best in class, latest generation trench IGBTs.

For more information on GaNpowIR® visit www.irf.com

1947: IR opens for business, launching selenium rectifiers

1954: Commercializes germanium rectifiers

1959: A pioneer in silicon controlled rectifier (SCR)

1962: Introduces new epitaxial process for producing industry's most stable high voltage SCRs

1979: Patents first commercially viable Power MOSFET (HEXFET®)

1980: Introduces high reliability power semiconductor devices for space programs

1983: Launches first commercially viable high voltage power ICs

1996: Introduces MOSFET and Schottky in a single package

1950 **1960** **1970** **1980** **1990**

2000

Introduces wafer level package (FlipFET™)

2001: Invents integrated building blocks for DC-DC converters (*iPOWIR*™)

2002: Invents first dual-sided surface mount package (DirectFET®)

2003: Introduces *iMOTION*™ Integrated Design Platform

Invents scalable multi-phase architecture (xPHASE®)

2006: Introduces AC-DC Smart Rectifier™ IC

2007: Introduces Sup*IR*Buck™ Integrated Voltage Regulators

2010

Introduces first commercial GaN based power stage device

2008: Introduces Automotive dedicated devices

Introduces GaN-based Power Device Technology Platform

See IR's Latest Innovations at APEC 2011

Six Decades of Technology Leadership

EXPECT MORE

INNOVATIVE power semi technology

INNOVATIVE application-specific chip sets

INNOVATIVE packaging designs

INNOVATIVE power systems architectures

A rare blend of benchmark silicon, advanced circuit design, control architecture, packaging and applications know-how.

When you think IR, think innovation.

For over 60 years, engineers have relied on International Rectifier's benchmark technology to transform crude electricity into clean, efficient power for home appliances, satellites, automobiles, portable electronics, computing, telecommunications, and a host of other systems.

for more information call 1.800.981.8699 or visit us at www.irf.com

International IOR Rectifier

THE POWER MANAGEMENT LEADER

Visit us at Booth 409

APEC® Fort Worth Convention Center 07-09 March 2011