UBM Electronics

EDN®

VOICE OF THE ENGINEER

APRIL 7
Issue 7/2011
www.edn.com

Tales from the Cube: Band together Pg 56

EDN.comment: The 3-D IC and you Pg 8

Signal Integrity: Whang that ruler Pg 16

Design Ideas Pg 45

Make your electronics supply chain green—or else Pg 52

ARM VERSUS INTEL:

A SUCCESSFUL STRATAGEM FOR RISC OR GRIST FOR CISC'S TRICKS?

Page 24

UNDERSTANDING THE IMPACT OF DIGITIZER NOISE ON OSCILLOSCOPE MEASUREMENTS

Page 18

INNOVATIVE CIRCUIT DESIGNS TARGET PERFORMANCE IMPROVEMENT AND PRODUCT DIFFERENTIATION

Page 36

THE E-READER PARADOX

Page 42

Hello future.

Goodbye status quo.

	Agilent 2000 X-Series (MSO and DSO)	Tektronix TDS2000C Series (DSO)	Agilent 3000 X-Series (MSO and DSO)	Tektronix MSO/DPO2000 Series
Bandwidth (MHz)	70, 100, 200	50, 70, 100, 200	100, 200, 350, 500	100, 200
Max sample rate	2 GSa/s	2 GSa/s	4 GSa/s	1 GSa/s
Max memory depth	100 kpts	2.5 kpts	4 Mpts	1 Mpt
Max update rate (waveforms/sec)	50,000	200**	1,000,000	5,000
Fully upgradable	Yes	No	Yes	No
Function Generator	Yes	No	Yes	No
Notes:	**Refer to Agilent Pub 5989-7885EN for update rate measurements Data for competitive oscilloscopes from Tektronix publications 3GW-25645-0 and 3GW-22048-1 Measurements taken on same signal using Agilent MSOX2024A and Tektronix TDS2024B Screen images are actual screen captures and scopes are shown to scale			

Oscilloscopes Redefined

Starting at $1,230*

*All prices are in USD and 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

See the difference today.
www.agilent.com/find/redefined

Agilent Technologies

570,000+ PRODUCTS IN STOCK
470+ SUPPLIER PARTNERS
NEW PRODUCTS ADDED DAILY

The industry's broadest product selection available for immediate delivery

www.digikey.com
1.800.344.4539

NEDA Supplier Authorized Distributors

GET CONNECTED

Multi-lane Jitter Tolerance running too slow?

Let Centellax Save You Time

If testing jitter tolerance in multi-lane devices or high channel count SERDES with a single serial BERT gives your test technician plenty of time for more exciting things, like watching plants grow, look to Centellax to speed things up. **Our Parallel Channel BERT measures Jitter Tolerance on ALL lanes simultaneously**, at a surprisingly low cost that won't break your capital budget.

Speed Up Your Testing

Minimize cable loss with remote pattern generator and error detector heads that can be mounted close to your device under test. 4-tap de-emphasis is integrated directly in the pattern generator head, eliminating the need (and cost) of a dedicated de-emphasis processor. Our Signal Integrity Studio software ties it all together, including a parallel jitter tolerance test with unlimited channel count.

When you are ready to speed up your testing, come see us at www.centellax.com.

CENTELLAX
SIGNAL INTEGRITY MEASUREMENT

Affordability without compromise

For more information on Centellax products visit **www.centellax.com**

EDN 4.7.11 contents

ARM versus Intel: a successful stratagem for RISC or grist for CISC's tricks?

24 ARM and its licensees are striving to expand their overall market presence by tackling Intel's x86 in servers and client desktop and laptop computers. Intel has responded by attacking ARM on its own turf: handsets, tablets, and the like.

by Brian Dipert, Senior Technical Editor

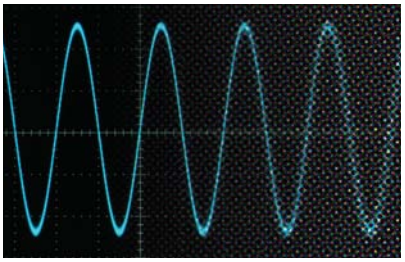

Understanding the impact of digitizer noise on oscilloscope measurements

18 How closely does the output of any ADC, waveform digitizer, or DSO follow an analog input signal? ENOB testing provides a means of establishing a figure of merit for dynamic digitizing performance.

by Jit Lim, Tektronix

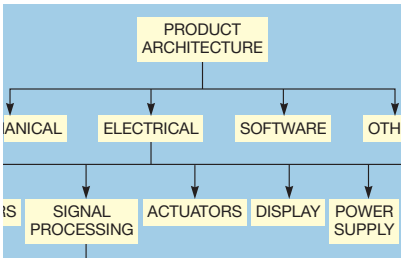

Innovative circuit designs target performance improvement and differentiation

36 Product innovation is the path to product rejuvenation. A holistic approach compels designers to consider innovation at all architectural levels. Product designers and ASIC designers can together achieve innovation breakthroughs.

by Eugene R Bukowski, Jr, Accenture

pulse

Dilbert 11

10 AWG boasts 14-bit resolution, 12G-sample/sec updates

11 Next-generation solid-state drives tout 6-Gbps SATA speeds

11 LED array enables small-form-factor MR-16 replacements

12 Researchers claim millimeter-scale computing system

12 AVX offers Spice models for tantalum, niobium-oxide capacitors

13 Eight-channel analog multiplexer meets high-reliability military specs

13 Tool for custom-IC design allows verification in real time

14 Lithography reduction yields design simplicity

The e-reader paradox

42 E-book readers require high performance, even to just display text.

by Franck Nicholls, Freescale Semiconductor

COVER: COMPOSITE IMAGE BY TIM BURNS.
GLOVE: INKTYCOON/ISTOCKPHOTO.COM;
FLAMES & BACKGROUND: SUSARO/ISTOCKPHOTO.COM

DESIGNIDEAS

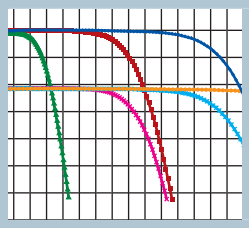

45 Low-component-count logic probe works with TTL and CMOS logic

46 Circuit implements photovoltaic-module simulator

48 Switch circuit controls lights

50 Isolated PWM suits low frequencies

Fast and Accurate
component testing
LCR meters starting at $995

- Basic Accuracy: 0.05% (SR720), 0.2% (SR715)
- Test frequencies: to 100 kHz (SR720), 10 kHz (SR715)
- Up to 20 measurements per second
- Measures R, L, C, Q and D
- 100 mV to 1.0 V test signals
- Internal bias to 2 V, external bias to 40 V
- RS-232 computer interface
- Kelvin clips, surface mount tweezers, 4-wire adaptor (opt)
- GPIB and Parts Handler Interfaces (SR720)

LCR Meters from SRS measure passive components with up to 0.05% precision. These simple to use instruments are perfect for incoming inspection, automated testing and general benchtop use. The convenient auto measure feature, binning capabilities for part sorting, and quick and easy calibration simplify test procedures. Standard RS-232, GPIB and Parts Handler interfaces make automated testing a snap. For fast, accurate component testing, the SR715 and SR720 offer unmatched value.

SR715....$995 (US list)
SR720....$1495 (US list)

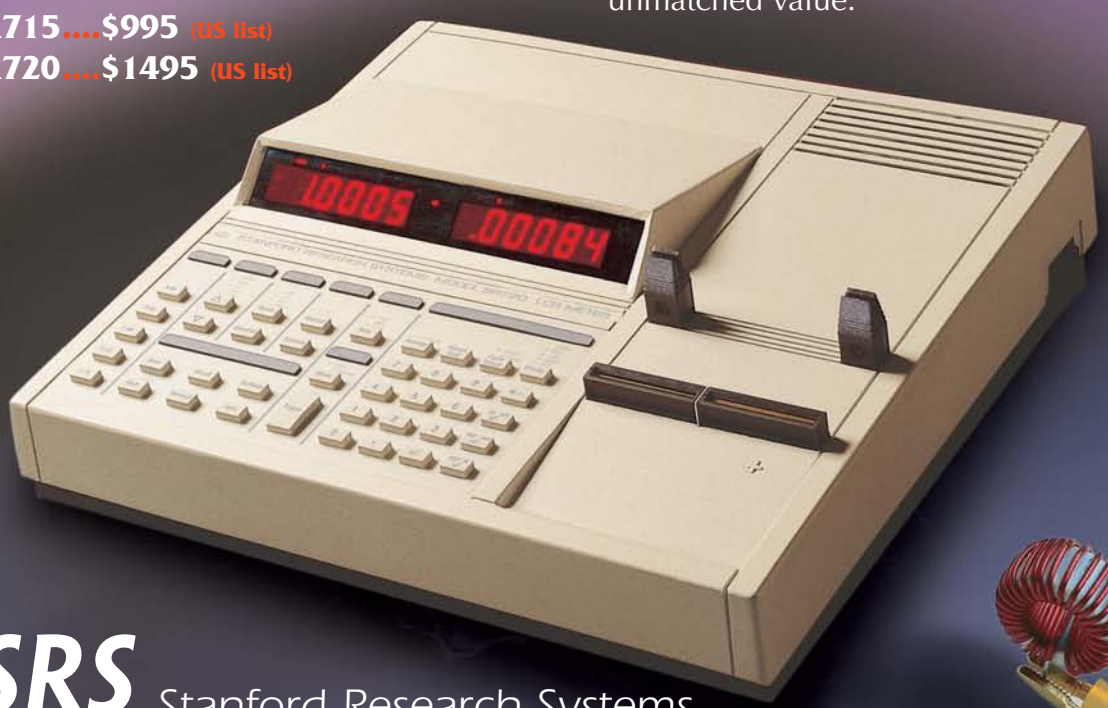

SRS Stanford Research Systems
1290-D Reamwood Ave, Sunnyvale, CA 94089 ▪ Phone (408) 744-9040
e-mail: info@thinkSRS.com ▪ **www.thinkSRS.com**

contents 4.7.11

DEPARTMENTS & COLUMNS

8 **EDN.comment:** The 3-D IC and you

16 **Signal Integrity:** Whang that ruler

52 **Supply Chain:** Make your electronics supply chain green—or else

54 **Product Roundup:** Discrete Semiconductors; Amplifiers, Oscillators, and Mixers

56 **Tales from the Cube:** Band together

EDN online contents www.edn.com

ONLINE ONLY

JOIN THE CONVERSATION

Connect with the more than 2700 engineers and executives exchanging ideas in *EDN*'s Electronics Design Network on LinkedIn.
→www.edn.com/110407toca

Readers' Choice

Here is a selection of recent articles and blog posts receiving high traffic on www.edn.com:

Baking LEDs in the oven: the Pharox 6W LED bulb
→www.edn.com/110407tocb

Counterfeiting continues to grow, but the industry fights back
→www.edn.com/110407tocc

Build accurate Spice models for low-noise, low-power precision amplifiers
→www.edn.com/110407tocd

Tablet success to infringe on PC sales
→www.edn.com/110407toce

Researchers aim at room-temperature spintronic computers
→www.edn.com/110407tocf

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 7 (Printed in USA).

BY RON WILSON, EDITORIAL DIRECTOR

The 3-D IC and you

One of the most popular topics for conference sessions lately has been the 3-D IC. Panels and papers cover a huge range of topics, but they come down to three questions: What is a 3-D IC, is it real, and what difference does it make? The question of definition is surprisingly loaded. At a recent panel, speakers divided the world of 3-D ICs into three categories. The first category covers simply stacking up independently designed dice and bonding them together, such as the stack of flash and DRAM dice on the SOC (system-on-chip) die in your cell phone. In the stacking approach, all the dice are pretested standard parts, often simply wire-bonded together using their normal I/O bonding pads, sometimes with a silicon interposer to move signals around for the best wire-bonding layout.

The second category of 3-D ICs starts in the architectural-design phase of the system. Architects and IC designers partition the system according to the best technology in which to implement each block. Logic blocks might go into a 20-nm logic process; bulk memory, into a DRAM process; and I/Os and other AMS (analog/mixed-signal) blocks, into a large-geometry, higher-voltage process. The team would design one die in each process, optimizing the interfaces among the dice for performance needs and power constraints. They would then assemble the dice into a stack, using flip-chip technology, interposers, and TSVs (through-silicon vias) to interconnect the blocks. One panelist calls this approach technology partitioning.

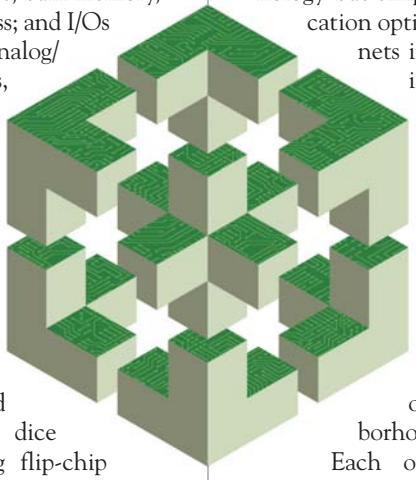

In the third category, which you could call true 3-D, designers place each cell not in a plane on one die but in the 3-D space they create by stacking many dice together. TSVs become just another element in the routing hierarchy. In this mindset, a cell may land in one corner of the third die in the stack—not because of its process technology but simply because that location optimizes timing for the nets in the area. You can imagine the complexity of a placement algorithm that must keep track of the additional delays and space required for the TSVs, the thermal- and mechanical-stress impacts of the cell on its local neighborhood, and so on.

Each of these categories has its own drawbacks and benefits. Stacking primarily delivers greater component density for space-constrained boards. Technology partitioning also offers greater density, but it further hints at cost savings and significant performance improvements from the optimized interdie interfaces. True 3-D promises the most: to take over as Moore's Law dies out, giving you a way to continue to increase transistor density.

What does this mean to designers at the board level? In an ideal world, the three approaches would become viable one after the other, providing a continuum of improving density and performance. There's a lot of work that must occur before we reach that ideal, however.

Clearly, stacking works today, if one partner—typically, a giant foundry—is tightly controlling a small supply chain. Technology partitioning is more problematic. Samsung recently demonstrated a wide-I/O DRAM die stacked on a logic die bearing TSVs, showing that this approach is not science fiction. Simon Burke, a Xilinx distinguished engineer, says that his company has internally produced an arrangement that stacks a dense FPGA-logic fabric in one technology atop a lower-density AMS die. Many issues still exist; these issues include diverse tool chains, lack of good TSV models, absence of thermal and mechanical 3-D modeling tools, and lack of standards.

Discussions of 3-D ICs come down to three questions: What is a 3-D IC, is it real, and what difference does it make?

As for true 3-D, neither design nor analysis tools exist to support it, and neither the TSVs nor the necessary die-thinning techniques are ready. KK Lin, Samsung's director of foundry design enablement, hopes that his company will have wide I/Os available to designers in 2013, but true 3-D could still lag years behind that. In short, 3-D will come—but slowly and in waves. EDN

Contact me at ron.wilson@ubm.com.

COMPOSITE IMAGE BY TIM BURNS; BLOCK PATTERN: BETACAM-SP/ISTOCKPHOTO.COM

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,
UBM Electronics
and UBM Canon (Publishing Division)
David Blaza,
Senior Vice President,
UBM Electronics
Karen Field,
Senior Vice President, Content,
UBM Electronics

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.

EDITED BY FRAN GRANVILLE

pulse

INNOVATIONS & INNOVATORS

AWG boasts 14-bit resolution, 12G-sample/sec updates

Agilent Technologies has added a high-resolution, wide-bandwidth, 8- or 12G-sample/sec modular instrument to its AWG (arbitrary-waveform-generator) portfolio. The M8190A enables radar-, satellite-, and electronic-warfare-device designers to make reliable, repeatable measurements and create highly realistic signal scenarios to test their products. Precision arbitrary-waveform generation is necessary for realistic testing of systems for detecting low-flying aircraft and for high-data-rate communications in satellite-communications systems. With resolution as high as 14 bits, the M8190A makes it easy for designers to distinguish between signals and distortion in their test scenarios and to more rigorously stress their devices. The instrument's optional 2G-sample memory lets you create longer and more realistic test scenarios.

"The M8190 allows engineers to approach reality when they create test scenarios," says Jürgen Beck, general manager of Agilent's digital- and photonic-test business. "Because the generator simultaneously offers greater fidelity, higher resolution, and wider bandwidth and produces multilevel signals with programmable ISI [intersymbol interference] and jitter at frequencies to 3 Gbps, customers can create signal scenarios that push their designs to the limit and bring new insights to their analysis."

The M8190A simultaneously offers 14 bits of resolution, as much as 5 GHz of analog bandwidth, and the ability to build realistic scenarios with 2G samples per channel of waveform memory. Compact modular AXIe (Advanced Telecommunications Computing Architecture Extensions for Instrumentation and Test Express) architecture, which targets use in high-performance instrumentation, reduces system size, weight, and footprint.

Agilent's Measurement Research Lab designed a proprietary DAC for the generator using an advanced silicon-germanium bipolar/CMOS process. The DAC operates at 8G samples/sec with 14-bit resolution and at 12G samples/sec with 12-bit resolution. At 8G samples/sec, the DAC delivers SFDR (spurious-free dynamic range) of as much as 80 dBc (decibels referenced to the carrier). This technology eliminates the trade-off between high resolution and wide bandwidth, so measurements are more reliable and repeatable and you are less likely to misinterpret glitches in the waveforms as analog output. The modular generator, whose US entry price is $79,000, works in either two- or five-slot AXIe chassis.

—by Dan Strassberg

▷**Agilent Technologies**, www.agilent.com/find/M8190.

Fitting into this two-slot enclosure, which connects to a separate laptop, the M8190A AWG provides a combination of high resolution, wide bandwidth, deep memory, and flexible waveform-segment sequencing and programming. It also fits into a five-slot AXIe enclosure, in which a plug-in computer module makes it a stand-alone unit.

TALKBACK

"I designed some of the first automatic toilet flushers in 1983 so had to deal with sifting microwatts out of a sea of garbage."

—Electronics designer Robert Capper, in *EDN*'s Talkback section, at http://bit.ly/himngN. Add your comments.

Next-generation solid-state drives tout 6-Gbps SATA speeds

Roughly 15 months after Micron (www.micron.com) unveiled its 6-Gbps, SATA (serial-advanced-technology-attachment)-supportive C300 solid-state drives and six months after the company followed them with the enterprise-targeted P300 variants, Intel has finally rolled out a 6-Gbps SATA-cognizant solid-state-drive family. Whereas Micron's latest-generation C400 drives employ NAND devices built on a 25-nm process that the company jointly developed with IM Flash Technologies (www.imftech.com), Intel has chosen the more conservative, although perhaps less cost-effective, path of sticking with 34-nm-fabricated flash-memory components.

> The transition to 6-Gbps SATA has a notable effect on sequential-access performance.

Intel's 510 series leverages the same Marvell (www.marvell.com) controller that Micron's products use. Rumor has it that Intel's internally developed controllers for earlier-generation products, such as the X25-M series, also employed SATA-system-interface IP (intellectual property) from Marvell. Intel's expertise came into play at the other end of the SATA-logic block, extending to the embedded flash-memory array.

Intel's 6-Gbps SATA-cognizant 510 series of solid-state drives touts higher sequential-read and -write speeds than its predecessor.

A notable downside of Intel's earlier-generation, 3-Gbps SATA-based solid-state drives was their comparatively slow sequential-write performance versus competitors' products. Judging from the 500-Mbyte/sec sequential-read and 315-Mbyte/sec sequential-write speed claims for the higher-capacity, 250-Gbyte 510 series variant, the company has heard and responded to the grumbling. Sequential-read and -write specifications for the 120-Gbyte product are 400 and 210 Mbytes/sec, respectively, presumably due to the fact that the controller can simultaneously access fewer storage-array components and the blocks within those components. And the transition from 3- to 6-Gbps SATA also has a notable effect on sequential-access performance: When the 250-Gbyte solid-state drive connects to a 3-Gbps SATA bus, the drive specifies sequential-read and -write speeds of 265 and 240 Mbytes/sec, respectively, whereas its 120-Gbyte sibling clocks in at 265 and 200 Mbytes/sec, respectively. Now in production, the 250- and 120-Gbyte versions of the 510 series sell for $584 and $284 (1000), respectively.

—by Brian Dipert

▷**Intel Corp**, www.intel.com.

DILBERT By Scott Adams

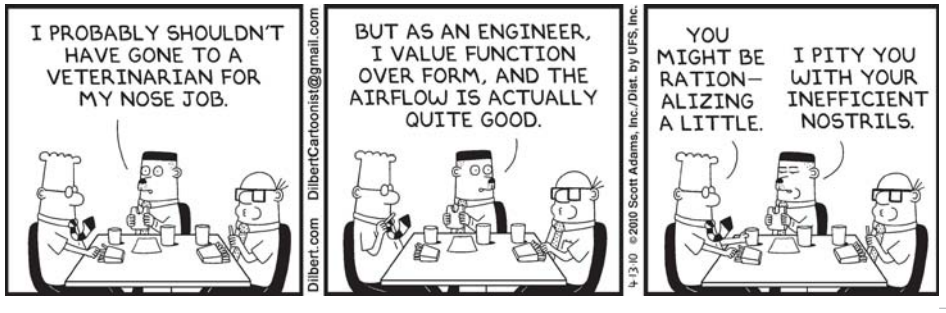

LED ARRAY ENABLES SMALL-FORM-FACTOR MR-16 REPLACEMENTS

04.07.11

In the United States, it's easy to see the world of LED lighting through Edison-bulb-illuminated glasses, but these bulbs are not the dominant global light sources. Europe, for example, has more than 1 billion 35 to 50W MR-16 halogen lights, about 2 inches long and selling for about $12 each. With the replacement market for these LEDs in mind, Cree recently introduced the MT-G LED array. It combines the company's EasyWhite color-mixing technology, a thermal resistance of 1.5°C/W, and efficacy as high as 92 lm/W (560 lm at 6W) at 85°C (3000K) in a 9.1×9.1-mm footprint.

These devices let you design an MR-16 light with one LED, packing the LED, a secondary optic, the power-management circuit, and a heat sink into the MR-16 form factor. Cree has published a reference design, including power-control circuits.

—by Margery Conner

▶**Cree**, www.cree.com.

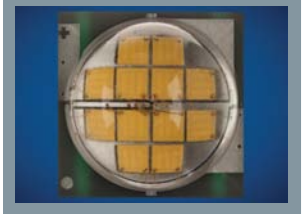

The MT-G LED array offers an efficacy as high as 92 lm/W (560 lm at 6W) at 85°C (3000K) in a 9.1×9.1-mm footprint and features Cree's EasyWhite color-mixing technology.

Researchers claim millimeter-scale computing system

Researchers from the University of Michigan have made notable progress toward mainstream millimeter-scale computing. The researchers at February's ISSCC (International Solid-State Circuits Conference) discussed a prototype implantable eye-pressure monitor for glaucoma patients that contains a complete millimeter-scale computing system. The researchers also discussed a compact radio that needs no tuning to find the right frequency and that they say could be a key enabler to organizing millimeter-scale systems into wireless sensor networks.

The researchers believe that millimeter-scale systems could enable ubiquitous computing. They point to Bell's Law, which states that a new class of smaller, cheaper computers emerges approximately once per decade. With each new class, the volume shrinks by two orders of magnitude and the number of systems per person increases. Researchers note that the law has held from 1960s' mainframes through the 1980s' PCs, the 1990s' notebooks, and the new millennium's smartphones.

"When you get smaller-than-handheld devices, you turn to these monitoring devices," says David Blaauw, a professor in the university's department of electrical engineering and computer science. "The next big challenge is to achieve millimeter-scale systems, which have a host of new applications for monitoring our bodies, our environment, and our buildings. Because they're so small, you could manufacture hundreds of thousands on one wafer. There could be tens to hundreds of them per person, and this per-

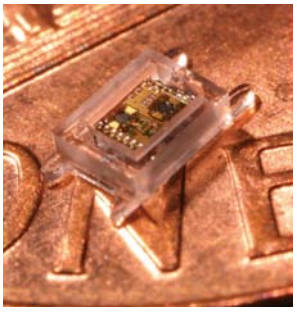

University of Michigan researchers developed a millimeter-scale computing system, which targets use in an implantable eye-pressure monitor (courtesy Cyouho Kim).

capita increase fuels the semiconductor industry's growth."

The pressure monitor comes in a slightly larger than 1-mm^3 package and fits an ultra-low-power microprocessor, a pressure sensor, memory, a thin-film battery, a solar cell, and a wireless radio with an antenna. It should become commercially available within several years.

"This is the first true millimeter-scale complete computing system," says Professor Dennis Sylvester. "Our work is unique in the sense that we're thinking about complete systems in which all the components are low-power and fit on the chip. We can collect data, store it, and transmit it. The applications for systems of this size are endless."

The processor in the eye-pressure monitor is the third generation of the researchers' Phoenix chip, which uses a power-gating architecture and an extreme-sleep mode to achieve ultra-low power consumption. The newest system wakes every 15 minutes to take measurements and consumes an average of 5.3 nW. Keeping the battery charged requires exposure to 10 hours of indoor light per day or 1.5 hours of sunlight. It can store a week's worth of information.

The researchers note that the system's radio cannot converse with other similar devices, so they have developed a consolidated radio with an on-chip antenna that needs no bulky external crystal, which two isolated devices currently use when they need to communicate. The crystal reference keeps time and selects a radio-frequency band. The process of integrating the antenna and eliminating this crystal significantly shrinks the radio system, the researchers say. They integrate the antenna through an advanced CMOS process, allowing the engineers to precisely control its shape and size and therefore how it oscillates in response to electrical signals.

"Antennas have a natural resonant frequency for electrical signals that is defined by their geometry, much like a pure audio tone on a tuning fork," says David Wentzloff, assistant professor. "By designing a circuit to monitor the signal on the antenna and measure how close it is to the antenna's natural resonance, we can lock the transmitted signal to the antenna's resonant frequency."

"This is the first integrated antenna that also serves as its own reference," Wentzloff adds. "The radio on our chip doesn't need external tuning. Once you deploy a network of these [antennas], they'll automatically align at the same frequency." The researchers are working to lower the radio's power consumption so that it is compatible with millimeter-scale batteries.

—by Suzanne Deffree

▷**University of Michigan**, www.umich.edu.

AVX offers Spice models for tantalum, niobium-oxide capacitors

AVX Corp has expanded the capabilities of its SpiTanIII software to enable designers to view all basic characteristics and parameters for tantalum and Oxicap (niobium-oxide) capacitors. The company has also introduced a new S-parameter library in Version 1.1, which includes s2p (two-port-data) files. AVX is offering models for Microsim and OrCAD PSpice as well as for OrCAD Capture schematic parts.

SpiTanIII Version 1.1 allows designers to choose a desired component by part number or by categories, including capacitance, rated voltage, case size, or series. The software displays all basic parameters for selected capacitors, such as the frequency characteristics of capacitance, ESR (equivalent series resistance), impedance, DF (dissipation factor), and ripple-current and ripple-voltage ratings. It also allows users to examine leakage-current DCL (dc-leakage)-versus-time progress and an equivalent circuit diagram and displays values of all elements. Users can export all graphs as either images or data for further evaluation.

SpiTanIII Version 1.1 includes part numbers of various SMD chip-capacitor series that debuted in 2010, and the software database features high-reliability and -performance wet tantalum-axial capacitors. You can download the software at www.avx.com/spiapps/spitan/SpiTanIII_V1.1.exe. For more news from AVX, go to http://bit.ly/hDoRoE.

—by Paul Rako

▶**AVX Corp**, www.avx.com.

Eight-channel analog multiplexer meets high-reliability military specs

Vishay Intertechnology recently expanded its family of MIL-PRF (military-performance-specification)-38535-screened analog switches and multiplexers with the release of the high-reliability, eight-channel, single-ended DG408 analog multiplexer. The device features a 44V maximum supply rating, on-resistance of 100Ω, and charge injection of 20 pC. The DG408's Enable pin allows you to use multiple devices to increase the number of input signals in an application.

The DG408 joins the previously released DG409 as a product that meets this new certification. Targeting use in high-speed data acquisition, audio-signal routing, and battery-powered and remote instrumentation, the device connects one of eight inputs to a common output, which a 3-bit binary address determines. Vishay's additional high-reliability capabilities include Group A, Group B, Group C, and Group E testing. In addition, Vishay continues to offer source-controlled drawing capabilities to allow space- and avionic-level screening.

The eight-channel, single-ended DG408 analog multiplexer features a 44V maximum supply rating, on-resistance of 100Ω, and charge injection of 20 pC.

> Source-controlled drawing capabilities allow space- and avionic-level screening.

The device comes in a choice of military-package options, including a 16-pin CerDIP, a 20-pin LCC, and a 16-pin flatpack. Samples and production quantities of the DG408AK multiplexer in the 16-pin CerDIP are available now; sample and production quantities of the DG408AZ in the 20-pin LCC and the DG408AL in the 16-pin flatpack will be available this quarter. Prices start at $168.

—by Paul Rako

▷**Vishay Intertechnology**, www.vishay.com.

Tool for custom-IC design allows verification in real time

Mentor Graphics' new Calibre RealTime platform allows designers to execute physical verification in real time during IC-layout creation. The first release provides instantaneous DRC (design-rule checking) in the SpringSoft (www.springsoft.com) Laker Version OA (Open Access) 2010.8 custom-IC-design and -layout tools, using the same Calibre decks as for a batch sign-off flow. A version of Calibre RealTime for the Mentor IC Station Version 10 custom-design environment will become available in June.

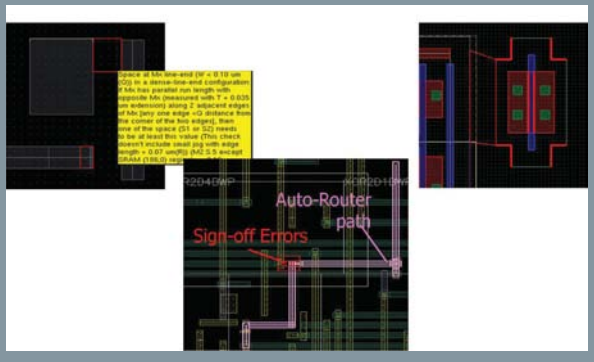

You can use Calibre RealTime to detect and define sign-off DRC errors during manual layout editing, automatic routing, or generation of automated P cells.

Calibre RealTime executes direct calls to Calibre analysis engines running foundry-qualified rule decks. User-defined custom filters allow designers to limit which checks to run, depending on design requirements and organizational processes, without modifying the foundry-qualified rule deck. When you integrate Calibre RealTime into a custom-IC-design and -layout system, the tool's engines perform incremental checking near the shapes you are editing, providing immediate feedback on design-rule violations. By running with the same rule set that you use for design sign-off, Calibre RealTime complements many layout editors' built-in checkers.

A built-in error-review tool bar in the layout-design environment enhances ease of use. In-memory checking optimizes performance. By combining Calibre RealTime with batch Calibre and Calibre RVE (results-viewing environment), layout designers can minimize the need for full-chip verification runs, shortening the production schedule. Joseph Sawicki, vice president and general manager of the design-to-silicon division of Mentor Graphics, says that the OA runtime model will enable integration with most custom design environments, including Cadence Virtuoso.

—by Mike Demler

▶**Mentor Graphics Corp**, www.mentor.com.

04.07.11

Lithography reduction yields design simplicity

Texas Instruments last month unveiled its ARM-based DaVinci-branded ICs. The video-focused TMS320DM816x is the first DaVinci-categorized product that TI built on a 40-nm lithography; its precursors use 65- and 90-nm processes. Like TI's Integra, they combine a Cortex-A8 CPU core and a C674x DSP core. The integrated peripheral mix is also similar to that of Integra, including PCIe (Peripheral Component Interconnect Express) Generation 2, SATA (serial-advanced-technology-attachment) 2.0, GbE (gigabit Ethernet), HDMI (high-definition multimedia interface), CAN (controller-area-network) transceivers, and DDR2/DDR3-memory controllers. Targeting applications such as video-surveillance systems, TI embeds video-optimized circuitry for encoding, decoding, transcoding, digital-to-analog translation, and other image-processing tasks.

> Integration yields 50% BOM-cost and one-fifth board-space reductions.

Differentiating the members of the TMS320DM816x family is the presence or absence of the PowerVR SGX530 3-D graphics engine, the clock frequency—1 GHz for the entire device or 720 MHz for the CPU and 667 MHz for other areas of the chip—and whether they have two or three HD (high-definition) video decoders. Each IC processes as many as three simultaneous 1080p, 1920×1080-pixel-resolution, progressive-scan, 60-frame/sec H.264 video streams. TI does not specify the bit rate. They also process 12 simultaneous 720p, 30-frame/sec video streams or a combination of lower-resolution streams.

A lower-end version, the TMS320DM8147, has no 3-D graphics core, whereas the TMS320DM8148 has a 3-D-augmented graphics core. They both run the ARM Cortex-A8 at 1 GHz, and the remainder of the chips operate at 750 MHz. Targeting 3W power consumption, the battery-powered devices can tackle a 1080p, 60-frame/sec video stream; three simultaneous 720p, 30-frame/sec video streams; or multiple lower-resolution streams.

TI anticipates future DaVinci family members that will be able to run the CPU and DSP at frequencies as high as 1.2 and 1 GHz, respectively. The company claims that the products' highly integrated nature reduces the system BOM (bill-of-materials) cost by 50% and decreases board space by one-fifth. Both products eliminate the need for previously required discrete components and halve the aggregate circuitry's power consumption.

DM8168 processors are now available for sampling, and prices begin at $75 (1000). The company expects the DM8148 family to become available for sampling in the third quarter, with prices starting at $51 (1000). TI is now shipping the $1995 TMDXEVM8168 evaluation module; the free EZ SDK (software-development-kit) download currently supports Linux-based operating systems. TI has scheduled Android support for the second quarter and Microsoft Windows Embedded Compact 7 support for the third quarter. **—by Brian Dipert**

▷**Texas Instruments**, www.ti.com.

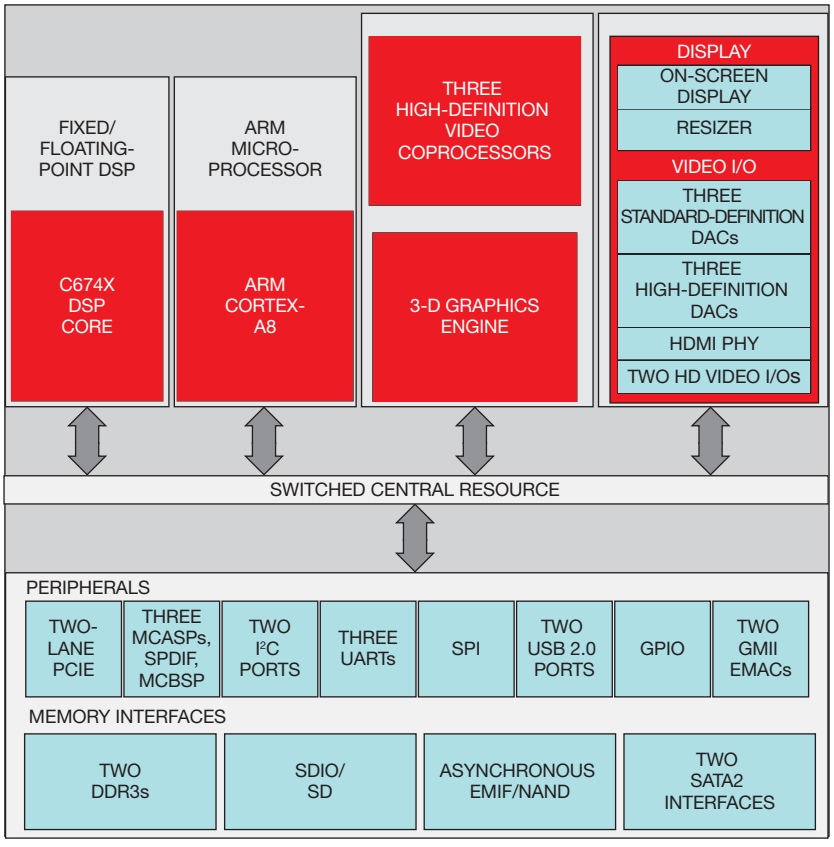

DDR: DOUBLE DATA RATE
DSP: DIGITAL-SIGNAL PROCESSOR
EMAC: ETHERNET MEDIA-ACCESS CONTROLLER
EMIF: EXTERNAL-MEMORY INTERFACE
GMII: GIGABIT MEDIA-INDEPENDENT INTERFACE
GPIO: GENERAL-PURPOSE INPUT/OUTPUT
HD: HIGH DEFINITION
HDMI: HIGH-DEFINITION MULTIMEDIA INTERFACE
I/O: INPUT/OUTPUT
I²C: INTER-INTEGRATED CIRCUIT
MCASP: MULTICHANNEL AUDIO SERIAL PORT
MCBSP: MULTICHANNEL BUFFERED SERIAL PORT
PHY: PHYSICAL LAYER
SATA: SERIAL ADVANCED TECHNOLOGY ATTACHMENT
SD: SECURE DIGITAL
SD: STANDARD DEFINITION
SDIO: SECURE DIGITAL INPUT/OUTPUT
SPI: SERIAL-PERIPHERAL INTERFACE
3-D: THREE DIMENSIONAL
UART: UNIVERSAL ASYNCHRONOUS TRANSMITTER/RECEIVER
USB: UNIVERSAL SERIAL BUS

The video-focused TMS320DM816x combines a Cortex-A8 CPU core and a C674x DSP core. Peripherals include PCIe Generation 2, SATA 2.0, GbE, HDMI, CAN transceivers, and DDR2/DDR3-memory controllers.

To get the most out of your ADC, our driver amplifiers make the difference.

A Driver for Every ADC

Low Power Differential ADC Driver
ADA4932—For true 14-bit applications to 20 MHz, or true 12-bit applications to 50 MHz. Supports single and differential inputs; only 9.6 mA quiescent current.

Ultralow Distortion, Current Feedback Differential ADC Driver
ADA4927—For driving ADCs in high gain configurations up to 20×. Supports 14-bits of resolution from dc to 100 MHz.

Level Translation, 16-Bit ADC Driver
AD8275—Integrated solution for driving precision, low voltage, single-ended, 16-bit ADCs up to 250 kSPS in ±10 V applications. Attenuates and level shifts input signals. Eliminates multiple parts and simplifies implementation.

Differential RF/IF Amplifiers
ADL5561/ADL5562—Highest linearity, lowest distortion amplifiers for driving 12-bit to 16-bit ADCs up to 500 MHz. Enable fully differential or single-ended-to-differential drive, with added flexibility of integrated on-chip user-selectable gain settings.

Amplifiers that optimize your design

You've worked hard to get the right level of performance out of your ADC, now make sure your input driver does the same. ADI's differential amp drivers deliver clean, low distortion inputs to your converter so it can operate at peak performance. Who better to turn to for these devices than the leader in signal processing, Analog Devices? In fact, we created the category of ADC drivers and offer the most options—from online design tools that help with the selection, evaluation, and troubleshooting of your circuit to the industry's broadest selection of ADC driver components. Make your ADC driving challenge easier at *www.analog.com/ADCdriver*.

MAKE ADIFFERENCE

www.analog.com/ADCdriver

ANALOG DEVICES

BY HOWARD JOHNSON, PhD

Whang that ruler

Clamp a wooden ruler to your desktop so that it overhangs the edge of the desk by about 8 inches. Now, flick the end of the ruler (**Figure 1**). It resonates, doesn't it? You can easily change the resonant frequency. Tape a few quarters to the end of the ruler and observe that the resonant frequency decreases. Shorten the length of overhang and hear it increase. If you push the resonant frequency high enough, it becomes difficult to stimulate the resonance with the soft end of your fingertip. Overcome that difficulty by depressing the end of the ruler with your finger in a way that lets the ruler slip off the hard edge of your fingernail.

At the risk of annoying your coworkers, whang the ruler over and over while you adjust the amount of overhang until you get a nice, musical resonance at about 100 Hz. Now, disconnect the clamp and hold the same length of ruler in your hand. Flick the end. Try to create the same resonant effect. I bet you can't do it.

When you clamp the ruler to the desktop, the clamp creates a low, mechanical impedance at one end of the ruler. The other end of the ruler remains free to move—the only limit being air resistance. Your stimulation creates a mechanical wave that bounces back and forth between these two endpoints, neither of which absorbs much of the mechanical energy. It therefore takes many back-and-forth cycles for the ruler to settle down. You have created a highly resonant system.

When you hold the ruler in your hand, the mechanical impedance of your hand lies close to the natural characteristic impedance of the ruler. Even if it is not a perfect match, your hand absorbs a significant portion of the energy in each cycle. The result is that the ruler cannot resonate.

The ruler supports transverse mechanical waves in just two directions: from one end of the ruler to the other and back again. Mechanical engineers call that phenomenon 1-D wave propagation. For any system like this one, an absorbing device at either end can totally damp the oscillations. Mechanical engineers use hydraulic shock absorbers, friction, air resistance, and rubber to absorb energy and create damping. Electrical circuits use resistive terminations.

Figure 1 Whanging a clamp-attached ruler causes it to oscillate for many cycles before settling down.

A more complex system, such as a child's Indian drum, supports wave motion in two dimensions. Waves on the surface of the drumhead spread and reflect in many highly varied and complex patterns. An absorbing device at just one location fails to damp the drumhead. To completely silence the drum, you must either apply absorbing material around a large fraction of the circumference or ask the child to please stop whanging. Good luck with that tactic.

PCB (printed-circuit-board) traces in a simple, linear topology behave as a 1-D-wave-propagation medium. If the trace is long enough and if it lacks any good energy-absorbing devices, it will resonate, distorting your signals.

When I say a trace has a simple, linear topology, I mean that it is a point-to-point connection or, at most, a linear-bus structure with multiple transceivers arrayed along a single trace. More complex structures, such as H distributions, star clusters, or random hairball nets, support multiple modes of oscillation and may, like a drumhead, require terminations in multiple locations.

When I say "long enough," I mean that the end-to-end trace delay is a significant fraction of the signal rise or fall time. A delay as long as one-sixth the rise or fall time is significant, especially if the trace has a particularly low-impedance driver or a large capacitive load. I simulate all such traces. Keep in mind that even short traces resonate; the resonant frequency is just so high that you may not observe its effect using logic with a particular rise and fall time. Applying logic with faster edges to the same trace might make it ring, just as flicking the ruler with your fingernail stimulates resonance at a small length.

Applying a capacitive load to a PCB trace has much the same effect as a load of quarters on the end of the ruler: It lowers the resonant frequency of the structure, making it more likely that you will notice its effects. EDN

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

UBM Canon

Your Project Acceleration Resource for Electronics Manufacturing

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
Electronics-South.com

Electronics® NEW ENGLAND

April 25–26, 2012

Boston Convention
& Exhibition Center
Boston, MA
Electronics-NE.com

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

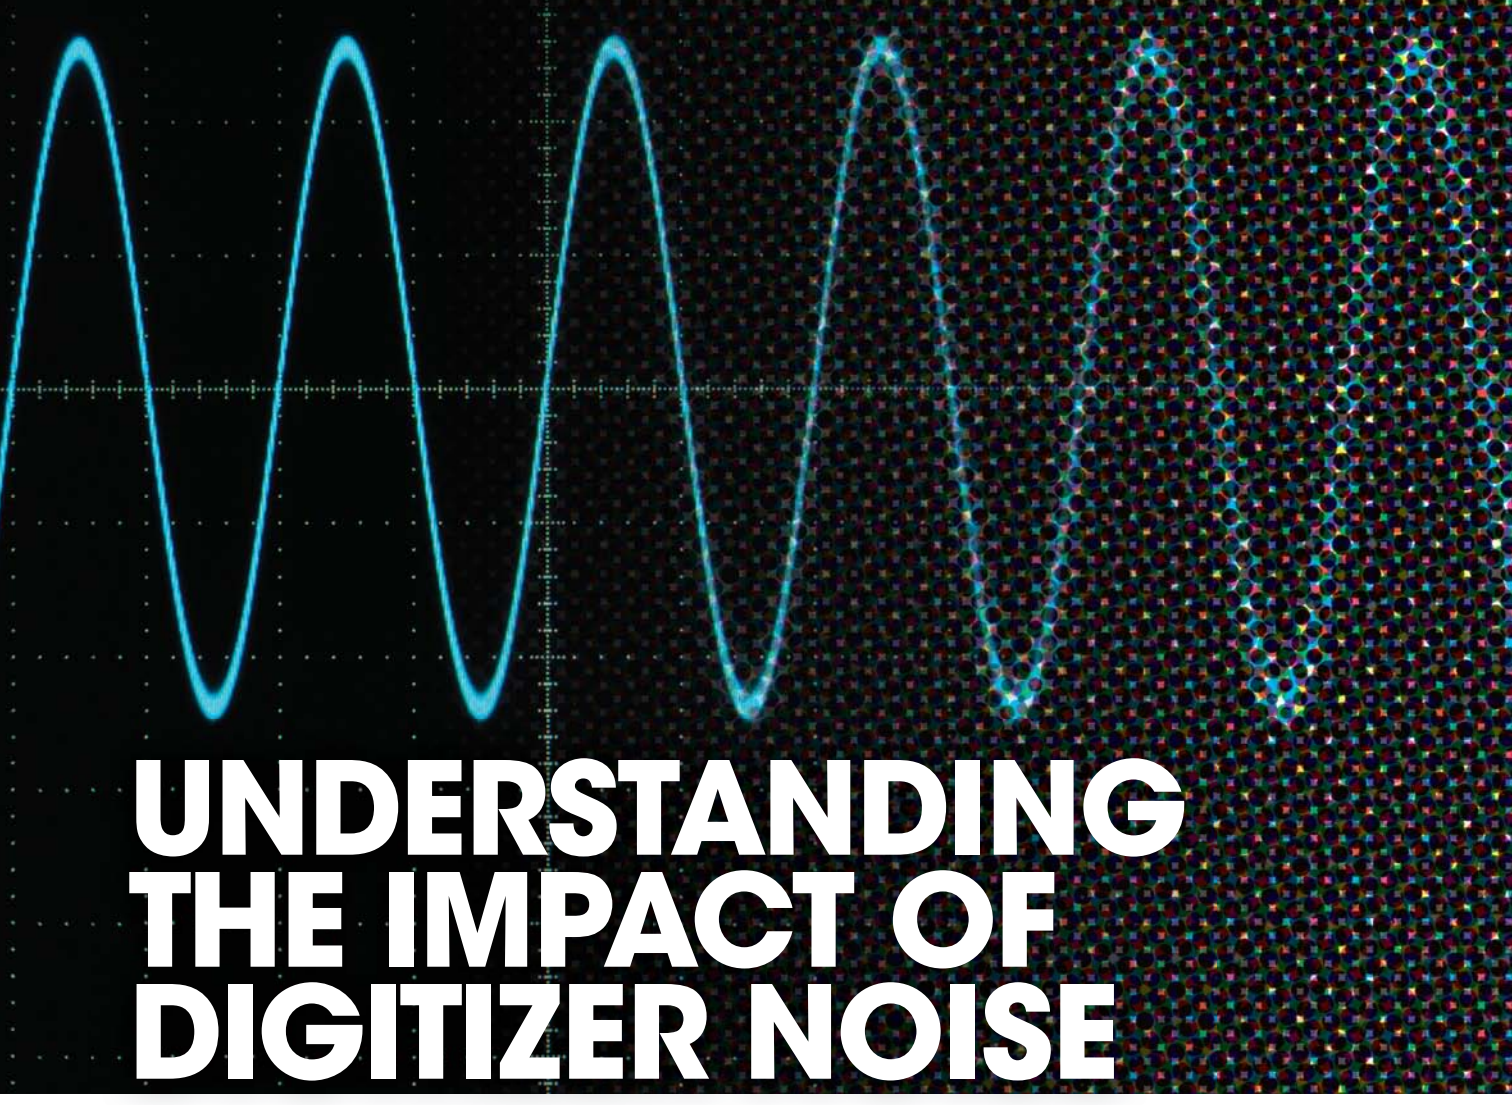

UNDERSTANDING THE IMPACT OF DIGITIZER NOISE ON OSCILLOSCOPE MEASUREMENTS

BY JIT LIM • TEKTRONIX

WHETHER YOU ARE DESIGNING OR BUYING A DIGITIZING SYSTEM, YOU NEED SOME MEANS OF DETERMINING REAL-LIFE PERFORMANCE. HOW CLOSELY DOES THE OUTPUT OF ANY ADC, WAVEFORM DIGITIZER, OR DSO FOLLOW AN ANALOG INPUT SIGNAL? ENOB TESTING PROVIDES A MEANS OF ESTABLISHING A FIGURE OF MERIT FOR DYNAMIC DIGITIZING PERFORMANCE.

One of the most common sources of errors in measurements is the presence of vertical noise, which can decrease the accuracy of signal measurement and lead to such problems as inaccurate measurements as frequencies change. You can use ENOB (effective-number-of-bits) testing to more accurately evaluate the performance of digitizing systems, including oscilloscopes. The ENOB figure summarizes the noise and frequency response of a system. Resolution typically degrades significantly as frequency increases, so ENOB versus frequency is a useful specification. Unfortunately, when an ENOB specification is provided, it is often at just one or two points rather than across all frequencies.

PHOTO-ILLUSTRATION BY TIM BURNS. SCREEN: GORDON HEELEY/ISTOCKPHOTO.COM

In test and measurement, noise can make it difficult to make measurements on a signal in the millivolt range, such as in a radar transmission or a heart-rate monitor. Noise can make it challenging to find the true voltage of a signal, and it can increase jitter, making timing measurements less accurate. It also can cause waveforms to appear "fat" in contrast to analog oscilloscopes.

THE ENOB CONCEPT

Digitizing performance is linked to resolution, but simply selecting a digitizer with the required number of bits, or quantizing level, at the desired amplitude resolution can be misleading because dynamic digitizing performance, depending on the technology, can decrease markedly as signal speeds increase. An 8-bit digitizer can decrease to 6, 4, or even fewer effective bits of performance well before reaching its specified bandwidth.

When designing or selecting an ADC, a digitizing instrument, or a test system, it is important to understand the various factors affecting digitizing performance and to have some means of evaluating overall performance. ENOB testing provides a means of establishing a figure of merit for dynamic digitizing performance. You can use it as an evaluation tool at various design stages and as a way to provide an overall system-performance specification. Because manufacturers don't always specify ENOB for individual instruments or system components, you may need to do an ENOB evaluation for comparison. Essentially, ENOB is a means of specifying the ability of a digitizing device or instrument to represent signals of various frequencies (**Figure 1**).

The **figure** illustrates that effective digitizing accuracy falls off as the frequency of the digitized signal increases. In this case, an 8-bit digitizer provides 8 effective bits of accuracy only at dc and low frequencies. As the signal you are digitizing increases in frequency or speed, performance drops to lower and lower values of effective bits.

This decline in digitizer performance manifests itself as an increasing level of noise on the digitized signal. Noise in this case refers to any random or pseudorandom error between the input signal and the digitized output. You can express this noise on a digitized signal in terms of SNR (signal-to-noise ratio): SNR=rms_{SIGNAL}/rms_{ERROR}, where rms_{SIGNAL} is the root-mean-square value of the digitized signal and rms_{ERROR} is the root-mean-square value of the noise error. The following **equation** yields the relationship to effective bits: $EB=\log_2(SNR)-\frac{1}{2}\log_2(1.5)-\log_2(A/FS)$, where EB represents the effective bits, A is the peak-to-peak input amplitude of the digitized signal, and FS is the peak-to-peak full-scale range of the digitizer's input. Other commonly used **equations** include $EB=N-\log2(rms_{ERROR}/IDEAL\_QUANTIZATION\_ERROR)$, where N is the nominal, or static, resolution of the digitizer, and, $EB=-\log_2(rms_{ERROR})\times\sqrt{12}/FS$.

These **equations** employ a noise, or error, level that the digitizing process generates. In the second **equation** above for EB, the ideal quantization error term is the rms error in the ideal, N-bit digitizing of the input signal. The IEEE Standard for Digitizing

AT A GLANCE

- ENOB (effective number of bits) is a general figure of merit for signal integrity in scopes and represents the cumulative errors across a frequency range.
- In general, the ENOB figure decreases as frequency increases.
- You should carefully evaluate ENOB performance, especially for applications involving high bit rates and fast edges.

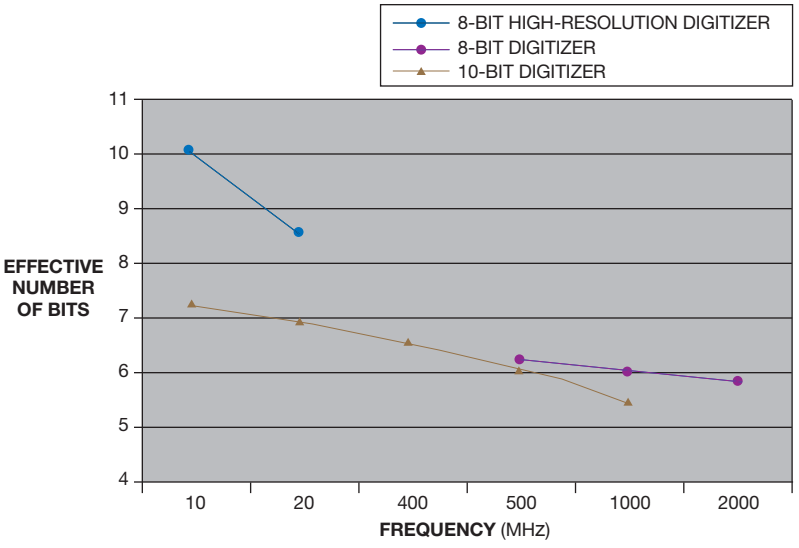

Figure 1 When comparing digitizer performance, it is important to test the full frequency range.

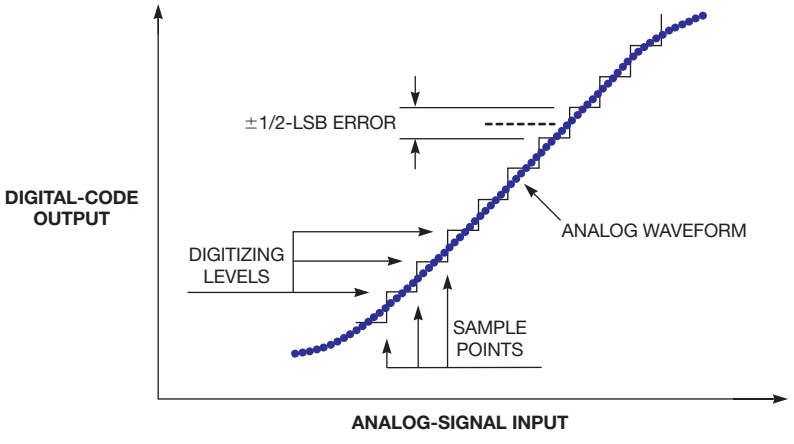

Figure 2 Quantizing errors are inherent parts of digitization.

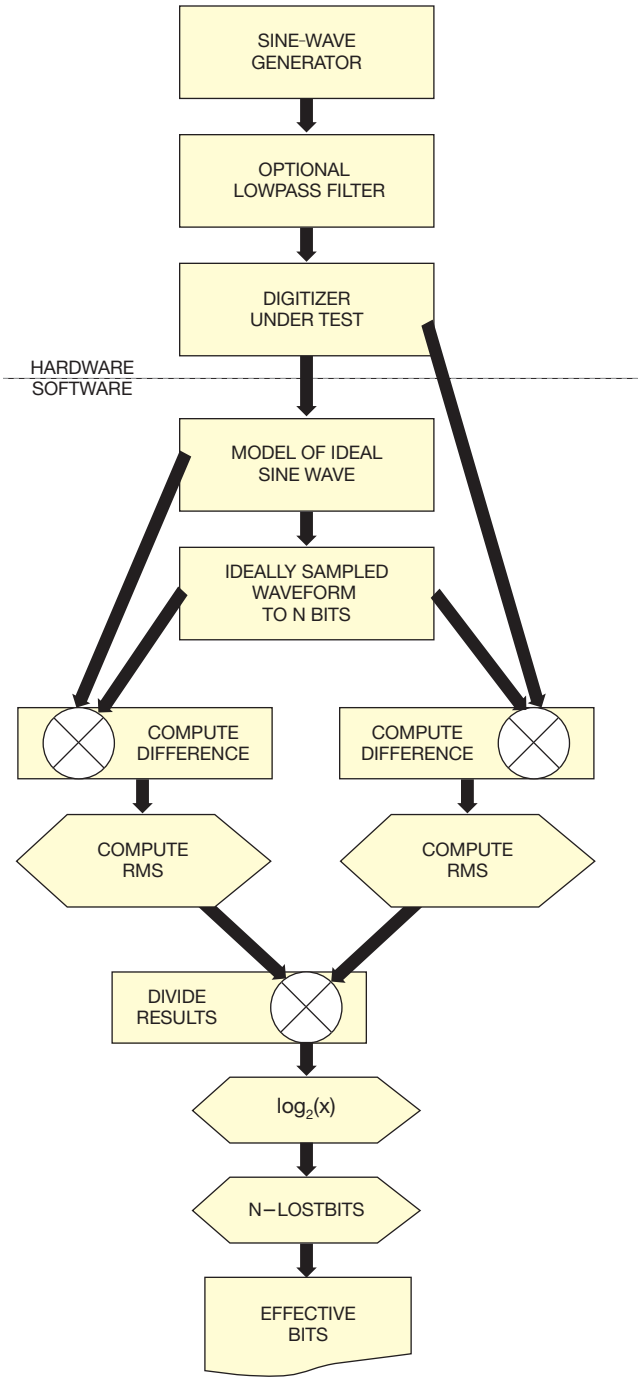

Figure 3 Measuring the ENOB figure involves applying a known, high-quality signal to the digitizer and then analyzing the digitized waveform.

Waveform Recorders (IEEE Standard 1057) defines the first two **equations** (**Reference 1**). An alternative for the third **equation** assumes that the ideal quantization error is uniformly distributed over one LSB (least-significant bit) peak to peak. This assumption allows you to replace the ideal quantization error term with $FS/(2^N\sqrt{12})$, where FS is the digitizer's full-scale input range.

These **equations** employ full-scale signals. Actual testing may use test signals at less than full-scale—50 or 90% of full-scale, for example. Improved ENOB results can improve this result, so comparisons of ENOB specifications or testing must account for both test-signal amplitude and frequency.

Noise or error relating to digitizing can come from a number of sources. Even in an ideal digitizer, quantizing causes a minimum noise or error level amounting to ±½ LSB. This error is an inherent part of digitizing (**Figure 2**). It is the resolution limit, or uncertainty, associated with ideal digitizing. A real-life digitizer adds further errors to this basic ideal error floor. These additional real-life errors can include dc offset; ac offset, or "pattern" errors, sometimes called fixed pattern distortion, associated with interleaved sampling methods; dc and ac gain error; analog nonlinearity; and digital nonmonotonicity. You must also consider phase errors; random noise; frequency-timebase inaccuracy; aperture uncertainty, or sample-time jitter; digital errors, such as data loss due to metastability, missing codes, and the like; and other error sources, such as trigger jitter.

ENOB MEASUREMENT

Beyond these error sources, still other possible sources of digitizing error exist. For example, in high-speed real-time digitizing without sample-and-hold or track-and-hold tracking, the LSBs must change at high rates to follow a quickly changing signal. This requirement increases bandwidth requirements for data lines and buffer inputs for these lesser bits. If you do not meet bandwidth requirements, quickly changing lesser bits can be dropped, lowering the ENOB.

It is often easier to measure overall performance instead of trying to distinguish and measure each error source in a digitizing system. A good place to start is by determining the digitizing system's SNR and the resulting effective bits according to the preceding **equations**. This approach provides an easily understood and universal figure of merit for comparisons.

The basic test process involves applying a known, high-quality signal to the digitizer and then analyzing the digitized waveform (**Figure 3**). The test uses a sine wave as the test signal because high-quality sine waves are relatively easy to generate and characterize. The general test requirements are that the sine wave generator's performance must significantly exceed that of the digitizer under test. Otherwise, the test will be unable to distinguish digitizing

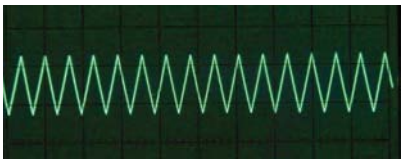

Figure 4 The oscilloscope does not display the extreme ranges of noise because they occur quickly and infrequently, and the resulting trace is thin.

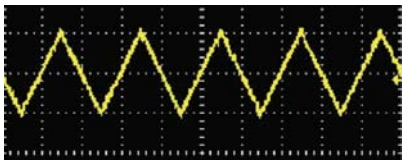

Figure 5 A DSO trace appears thicker than the analog scope's trace because every hit has the same intensity.

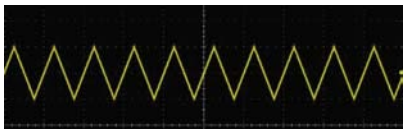

Figure 6 A DPO trace displays the signal based on the frequency of the hits.

errors from signal-source errors. It may be necessary to add filters to the source to reduce source harmonics to levels significantly below what you might expect from the digitizer under test.

To obtain an ENOB, you must compute a perfect, or idealized, sine wave and apply it to your oscilloscope, fitting it to the digitized sine wave. This approach simulates what the N-bit digitizer under test would produce if it were an ideal N-bit digitizer. You then compute the difference between the computed ideal sine wave and the perfectly sampled and digitized version. The rms value of this difference provides the ideal quantization error. You obtain the rms-error value in the ENOB **equations** by subtracting the ideal sine wave from the actual digitized sine wave and finding the rms value of the result. Alternatively, you can find the rms value of the signal and the rms error and use them to compute SNR. The final computation results in an ENOB for the digitizer. By keeping the input signal's amplitude constant for various frequencies, you can further compute ENOBs for the subject digitizer or digitizing system. You can plot these numbers against frequency to obtain a digitizer performance curve.

IT IS OFTEN EASIER TO MEASURE OVERALL PERFORMANCE INSTEAD OF TRYING TO MEASURE EACH ERROR SOURCE IN A DIGITIZING SYSTEM.

ENOB measurement combines the key digitizer system errors into a figure of merit that is easy to understand and use in comparisons. ENOB depends on the input signal's percentage of full-scale digitizer amplitude. Testing a digitizer at less than full-scale amplitude generally yields a somewhat better ENOB figure than does testing it at full-scale. Whatever test approach you use—full-scale or partial scale—the input test signal's amplitude specification should accompany the results.

SCOPE NOISE

When comparing digital oscilloscopes with analog oscilloscopes, a common misperception is that digital oscilloscopes have a higher level of vertical noise. With digital oscilloscopes, the trace may appear fatter than that of its analog-oscilloscope equivalent. A digital oscilloscope does not have higher noise levels than the equivalent analog oscilloscope, however; it just appears that way.

Analog oscilloscopes with CRT displays do not display the extreme ranges of noise because they occur quickly and infrequently (**Figure 4**), meaning that the phosphor lights quickly and infrequently, and those extremes are dim or not on the screen at all. Analog instruments do not just display voltage versus time but have a third dimension: intensity. Intensity relates to the frequency of occurrence of the signal. A DSO (digital-signal oscilloscope) shows every hit with the same intensity, no matter the frequency of pixel hits (**Figure 5**). DPOs (digital-phosphor oscilloscopes)

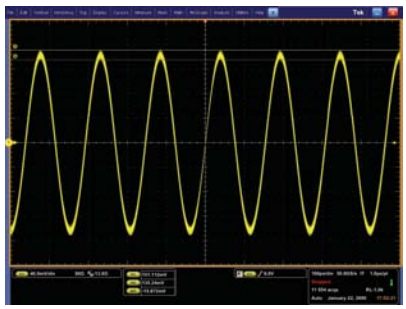

Figure 7 A test applies a 6.5-GHz sine wave to a DPO with 13-GHz bandwidth and 400-mV full-scale amplitude. The DPO also has infinite display persistence to show variations across all acquisitions.

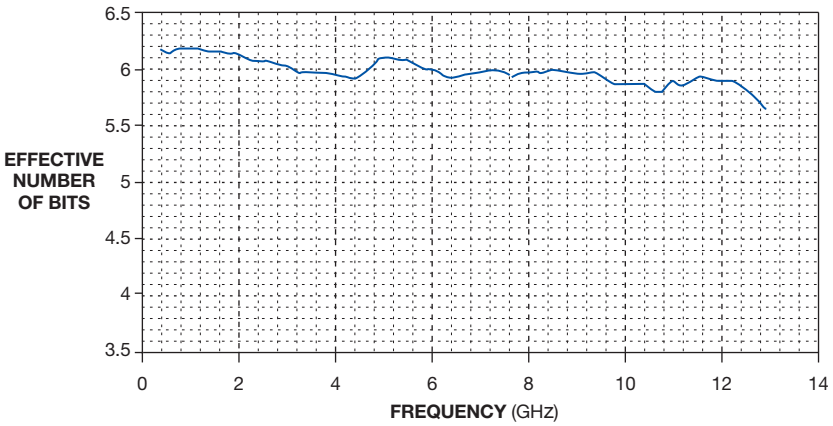

Figure 8 This result from Figure 7's test corresponds to approximately 5.9 ENOBs at 6.5 GHz.

offer a way to restore that third dimension by grading the signal employing the frequency of the hits (**Figure 6**).

REAL-WORLD SIGNAL NOISE

ENOB performance indicates noise that has an effect on both amplitude and timing measurements. To illustrate the effects of noise on amplitude, a test applied a 6.5-GHz sine wave to a Tektronix (www.tektronix.com) DPO/DSA70000B oscilloscope with a 13-GHz bandwidth and a 400-mV full-scale voltage. It also has infinite display persistence so that you can see variations across all acquisitions. With no averaging, the test run included approximately 10,000 acquisitions. The result is approximately 15.9 mV of trace thickness at the peak, representing 3% of full-scale on this oscilloscope (**Figure 7**). This result corresponds to approximately 5.9 ENOBs at 6.5 GHz (**Figure 8**). Comparative testing shows other oscilloscopes with more than 35-mV trace thickness at the peak and approximately 4.5 ENOBs using identical test setups.

ENOB EFFECTS

ENOB effects can also be seen on

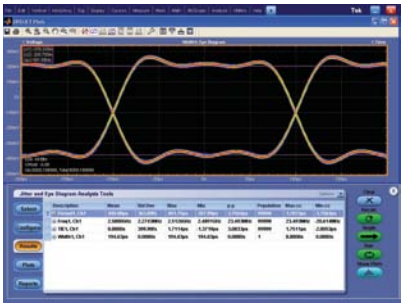

Figure 9 ENOB performance affects both amplitude and jitter on the eye diagram.

eye diagrams. ENOB performance affects both amplitude and jitter on the eye. A 5-Gbps eye diagram represents a signal bit rate that is associated with PCIe (Peripheral Component Interconnect Express) Generation 2 or USB (Universal Serial Bus) 3.0 (**Figure 9**).

The test applies this signal to the Tektronix DPO70000B oscilloscope, with the instrument set up to measure TIE (time-interval-error) jitter, affecting both jitter and amplitude noise. The measured jitter for this test was 3.08 psec p-p. In comparison testing, some oscilloscopes show more than 11 psec p-p on the same signal.

Similarly, noise also affects the eye's amplitude. In this case, a measurement of the eye height at the 50% point of the eye shows approximately 582-mV amplitude. This result compares with less than 525 mV measured on other instruments.

All digitizing systems have noise that only gets worse as speeds increase. Therefore, it is useful to have a way to evaluate the real-life noise performance of digitizing systems, including test instrumentation. ENOB is a general figure of merit for signal integrity in any analog or digital system, representing the cumulative errors across a frequency range. Generally, the ENOB figure decreases as frequency increases.

You can easily see errors relating to lower ENOB performance in real-world signals as increased noise when performing amplitude measurements and increased jitter when making jitter measurements. As the ENOB figure decreases, the measurement precision of the instrument decreases, directly equating to the margin available for the tests you are performing on the instrument. With these factors in mind, you should carefully evaluate ENOB performance, especially for applications involving high bit rates and fast edges. EDN

REFERENCE

1 "1057-1994, IEEE Standard for Digitizing Waveform Recorders," IEEE, 1994, http://bit.ly/eYo1NP.

AUTHOR'S BIOGRAPHY

Jit Lim is the senior technologist for high-speed signal analysis at Tektronix, where he has worked for 23 years. He has published numerous technical papers and holds a bachelor's degree in electrical engineering from the Massachusetts Institute of Technology (Cambridge, MA).

ARM VERSUS INTEL: A SUCCESSFUL STRATAGEM FOR RISC OR GRIST FOR CISC'S TRICKS?

ARM AND ITS LICENSEES ARE STRIVING TO EXPAND THEIR OVERALL MARKET PRESENCE BY TACKLING INTEL'S x86 IN SERVERS AND CLIENT DESKTOP AND LAPTOP COMPUTERS. INTEL HAS RESPONDED BY ATTACKING ARM ON ITS OWN TURF: HANDSETS, TABLETS, AND THE LIKE.

BY BRIAN DIPERT • SENIOR TECHNICAL EDITOR

ARM, along with its core licensees, and Intel, along with its x86 CPU competitors, have recently taken action to put to rest any remaining doubt that both camps were on a collision course—ARM touting its RISC (reduced-instruction-set-computer)-based technology and Intel backing the CISC (complex-instruction-set-computer) approach. When Intel three years ago formally introduced the first-generation Atom processor family, the company made it clear that it was aiming not just at low-end desktop and notebook PCs but also at the handheld systems in which ARM had historically dominated. In response, ARM more recently unveiled the Cortex-A15 core, whose application targets extend up to the server segment in which Intel and AMD (Advanced Micro Devices) have long reigned supreme. And at the January 2011 CES (Consumer Electronics Show), Microsoft revealed its willingness to put a nail in the coffin of the Wintel alliance by broadening upcoming Windows 8's instruction-set compatibility to encompass both ARM and x86.

Since ARM unveiled the Version 7 instruction set, the company has subdivided its product line into three segments: the highly integrated Cortex-A application processors for mobile devices, cost-sensitive Cortex-M processors for traditional microcontroller applications, and high-performance Cortex-R processors for deeply embedded real-time applications. Cellular handsets, multimedia record-and-playback devices, and other portable electronics systems incur substantial product volume shipments. That fact, along with their direct competition with Intel-architecture processors, explains why this article focuses on ARM Cortex-A CPUs (see **sidebar** "Intel's potential multiphase response"). For more, see *EDN*'s coverage of Cortex-M, Cortex-R, and alternative-market Cortex-A products, such as Ambarella's iOne (**Reference 1**).

ARMING FOR BATTLE

Any discussion of ARM's potential success in currently x86-dominated ecosystems must begin with an understanding of ARM's business model and current product offerings. From a fiscal viewpoint, as an IP (intellectual-property) developer, ARM depends largely on the market success of its licensees, which fall into conventional- and architecture-license camps. Conventional licensees implement predesigned cores in their SOC (system-on-chip) designs, a more straightforward path to bringing products to market, which conversely limits each licensee's ability to differentiate its products from those of competitors. Architecture, or instruction-set, licensees, on the other hand, have more design flexibility but also incur incremental corresponding design challenges. Although they must, as their name implies, retain full ARM instruction-set backward compatibility, they can also build on that suite with proprietary instructions, as well as make other more fundamental circuit alterations and enhancements. ARM currently has few architecture licensees, including Intel, Marvell, Microsoft, Nvidia, and Qualcomm, for example.

As for the Cortex-A proliferations currently available to conventional licensees, the Cortex-A8 builds on the Cortex-A5 foundation, which many licensees bypassed in the transition from the ARM11. Cortex-A8 offers dual-issue superscalar support and deepens the per-core pipeline from eight to 13 stages as a means of boosting clock rates at the potential expense of IPC (instruction-per-clock) efficiency. Cortex-A8 requires an upgraded FPU (floating-point unit), which was optional on ARM11. Cortex-A8 also requires the 64-bit, SIMD (single-instruction/multiple-data) Neon computing engine. Moving from that point to the ascendant Cortex-A9 involved several enhancement steps. Implementing Neon, an FPU, or both became a design-trade-off decision versus an implementation requirement. ARM shrunk the per-core pipeline to nine stages but retained more-than-1-GHz performance, thanks to lithography reductions, and out-of-order support further improved the average delivered IPC.

Conventional licensees' core-implementation constraints don't rule out notable innovations. Take, for example, the Fast14 dynamic-logic and signal-encoding techniques that Samsung licensed from Intrinsity for use in Samsung's Cortex-A8-based Hummingbird SOC. After acquiring Intrinsity a year ago, Apple subsequently implemented those techniques in its Cortex-A8-based A4 CPU. Fast14 enabled both Apple and Samsung, which also acted as the A4 SOC's foundry source, to obtain notably higher clock speeds at a given process node than other Cortex-A8 licensees could accomplish. Conversely, architectural licensee Qualcomm developed the ARM Version 7-compliant Scorpion architecture, which, initially on 65-nm lithography, is currently in 45 nm.

Scorpion achieved dual-core status in mid-2010 with the 1.2-GHz MSM8260 and MSM8660. Functionally, it is an intermediary step between Cortex-A8 and Cortex-A9, supporting some but not all of Cortex-A9's out-of-order instruction-execution capabilities. Scorpion-based Snapdragon SOCs also implement Cortex-A8- and A9-compliant floating-point and Neon SIMD engines. Scorpion implements the floating-point engine in a pipelined fashion, and the SIMD engine, at 128 bits, is twice as wide as the one in Cortex-A9.

Two generic terms, "application processor" and "baseband processor," bear mentioning, although in practice they are becoming increasingly irrelevant, thanks to single-die integration trends. An application processor, as a seminal 2004 presentation from BDTI (Berkeley Design Technology Inc) describes, runs user applications, supports complex operating systems, emphasizes multimedia processing, supports Java and other virtual machines, and implements security features (**Reference 2**). The companion baseband processor, on the other hand, tackles wireless communications, including various cellular voice and data protocols.

Various ARM licensees have divergent perspectives on baseband- versus application-processor segmentation. Texas Instruments, for example, announced in October 2008 that it would phase out its participation in the baseband business, which it viewed as highly commoditized and therefore insufficiently profitable. Qualcomm, on the other hand, has embedded baseband capabilities within most of its Snapdragon CPUs, befitting the company's MSM (Mobile Station Modem, formerly QSD for Qualcomm Semiconductor) product-name prefix, although some basebandless APQ (Application Processor Qualcomm) devices also exist. Nvidia, conversely, has been adamant that it makes no sense to burden the silicon area or constrain the technology development of an application processor with baseband capabilities. It may also make little sense to put a cellular-cognizant application processor into, for example, an optionally or by-default Wi-Fi-only tablet-computer design, therefore explaining, for example, the dual-core 1.2-GHz Qualcomm APQ8060 in Hewlett-Packard's first-generation TouchPad tablet.

AT A GLANCE

- ARM's conventional and architecture licensees have varying means of developing products employing the company's diverse instruction sets and cores.
- Nvidia has clearly broadened its corporate focus beyond or, depending on your perspective, away from PC graphics and toward the burgeoning ARM-SOC (system-on-chip) market.
- Texas Instruments, late to the Cortex-A9 era, is determined to quickly and solidly regain its long-standing ARM momentum.
- Apple and Samsung direct their semiconductor groups' SOC outputs at their system divisions. Samsung does so with an inconsistent sourcing strategy, however.
- Qualcomm and Marvell are leveraging their architecture licenses to create differentiated ARM-based products.

A FOCUS ON MULTIMEDIA

Nvidia obtained notable aspects of its ARM-SOC program through the acquisitions of MediaQ, which it announced in August 2003, and of PortalPlayer, which it announced in November 2006. The company's initial ARM11-based Tegra SOCs achieved limited market success, aside from being the processing nexus of Microsoft's Zune HD portable multimedia player

Figure 1 Nvidia's Tegra 2 SOC is seemingly the reference CPU for Google's Android 3.0 Honeycomb operating system in Motorola's Xoom tablet (a), but the integrated processor also finds use in modern high-end smartphones, such as Motorola's Atrix 4G (b), whose mating dock, display, and keyboard enable it to approximate a full PC experience (c). LG's Optimus 2X might use Tegra 2, but the newer Optimus 3D went with TI's OMAP 4 instead, due to its higher-resolution video-encoding support (d). RIM's BlackBerry PlayBook tablet is another OMAP 4 design win (e), whereas Samsung's 10.1-in. Galaxy Tab gave Tegra 2 the nod. The company's Galaxy S II smartphone bafflingly splits processor loyalties between Nvidia and Samsung's own semiconductor group (f). HP's TouchPad is a Qualcomm dual-core Snapdragon advocate (g).

and short-lived Kin cellular handset. The company bypassed both the Cortex-A5 and the Cortex-A8 generations to come up with its next-generation Tegra 2 products, which it formally unveiled at the January 2010 CES. A year ago, Nvidia had a somewhat-tarnished industry reputation because the raft of Tegra 2-based tablet computers and other devices that company Chief Executive Officer Jen-Hsun Huang claimed at CES would shortly materialize hadn't done so.

What a difference a year makes. Nvidia had achieved first-to-market status in the dual-core Cortex-A9 generation by several quarters versus competitors such as Texas Instruments. In doing so, it gained Google's nod in the reference design for tablet-targeted Android Version 3 Honeycomb (**Figure 1**). Several high-end smartphones that debuted at the 2011 iterations of CES and MWC (Mobile World Congress) in mid-February also harness its capabilities, although its reported inability to decode high-profile 1080p H.264-encoded video led to its 11th-hour replacement in the Boxee set-top-box design by Intel's Atom-based CE4100 CPU. Similarly, LG chose a Texas Instruments OMAP (Open Multimedia Applications Platform) 4-based CPU for its Optimus 3-D handset, although Tegra 2 had received the nod in the Optimus 2X phone LG had introduced less than two months earlier. LG chose the OMAP 4 because TI's SOC could encode 3-D content in real time in 1080p resolution, whereas Nvidia's CPU had only 720p capabilities (**references 3** and **4**).

Tegra 2 comes in T20 and AP20 variants, encompassing the Tegra 230 and Tegra 250, respectively. Both CPUs run at 1 GHz, but the GPU (graphics-processing unit) in T20 uses a 333-MHz clock and DDR2 system memory, befitting the larger screens and batteries in

tablets. Its handset-focused AP20 peer, in contrast, operates the GPU at 300 MHz and interfaces to the more-power-stingy LP (low-power) DDR2 SDRAM. Strictly speaking, it's a triple-ARM configuration because it also contains an ARM7 core for overall SOC management. You can suspend both Cortex-A9 cores, albeit not individually, when, for example, only the audio, imaging, graphics, and other dedicated processing resources are in use (**Figure 2**).

Nvidia leaked documentation in late January suggesting that it planned to announce AP25 and T25, or Tegra 2 3-D, during the first quarter of this year. Both variants upped the ARM core clock speed to 1.2 GHz, with GPU speeds up to 400 MHz. It's now unclear, however, whether they'll ever see the light of day because the quad-core follow-on AP30 and T30 SOCs, or Tegra 3, code-named Kal-El, have also appeared, courtesy of an aggressive development cycle. Nvidia obtained first Kal-El samples from its foundry partner just 12 days before MWC, during which it was showing robust graphics and video demos running at 2560×1600-pixel resolutions, twice the per-frame pixel count of conventional 1080p video.

The target per-core Cortex-A9 clock speed for Kal-El is 1.5 GHz. Unlike Tegra 2, Kal-El embeds a Neon SIMD vector FPU for each CPU core. Like Tegra 2, however, Kal-El continues to rely on a unified 1-Mbyte pool of L2 cache memory, which twice as many CPU cores as before share, as well as a 32-bit system-memory interface. The Nvidia-designed GPU is not only faster in Kal-El than in Tegra 2 but also ups the core count from eight to 12. Both Tegra 2 and Kal-El employ a 40-nm manufacturing process, thereby explaining the die-size boost from 49 mm^2 on Tegra 2 to roughly 80 mm^2 on Kal-El. Nvidia's marketers believe that Kal-El will deliver roughly five times the aggregate speed of Tegra 2—twice the CPU performance and three times the graphics performance—and consume no more—and, in some cases, notably less—power on a workload-

INTEL'S POTENTIAL MULTIPHASE RESPONSE

You might think, faced with ARM's formidable development road map and equally formidable laundry list of licensees, coupled with the loss of the x86 instruction-set lock in upcoming Windows 8, that Intel would be reeling and ready to retreat in its legacy stronghold computing markets. First, though, recall that the company has formidable semiconductor-manufacturing resources that it can apply to problems such as these, in the form of new process lithographies, which the foundries that most ARM licensees use find challenging to pace, and in the form of a multifab network. Last September, company officials explicitly stated their plans for a future move of Atom-based SOCs (systems on chips) down to at least the 15-nm process node, if not further.

Intel originally manufactured Atom on a 45-nm lithography in CPU form; the companion chip set used a 90-nm process. It has now progressed down to the 32-nm process level in the form of the highly integrated Medfield chip set for handsets and Oak Trail chip set for tablets. Oak Trail offers full PCI (Peripheral Component Interconnect)-bus support for Windows 7 cognizance (Reference A). Intel's partnership with Nokia appears to be on the ropes due to Nokia's recent embrace of the Windows Phone 7 operating system to the exclusion of both legacy Symbian and Intel-co-developed—and, presumably, x86-favoring—MeeGo. Nevertheless, Intel executives remain confident that Medfield-based handsets will enter production this year. Numerous system manufacturers are demonstrating Oak Trail-based tablets running both Microsoft-developed and alternative operating systems.

Instruction-set compatibility is seemingly decreasing in importance over time with consumers, due to the consolidating number of file formats that require support, the increasing number of OS-agnostic applications that comprehend those formats, and the growing trend of storing both applications and their data in the "cloud" rather than on client systems. However, instruction-set compatibility—specifically for x86—remains critically important for software developers, who prefer to reuse known functional code snippets and associated development tools whenever possible.

Keep in mind, as well, that AMD and Via Technologies are also aggressively developing low-power x86 products, such as AMD's Bobcat core, which first-generation Fusion graphics-inclusive CPUs employ, and Via Technologies' now-monolithic dual-core Nano X2 CPU. As such, if Microsoft's broadening of Windows 8 beyond x86 to ARM represents an aspiration to consolidate on one operating-system code base, thereby obsoleting Windows CE-based products, ARM suppliers could find themselves facing formidable x86 competition in the near future in their traditional handheld-system strongholds (Reference B).

Even if ARM can tip the overall semiconductor-shipment balance in its favor over time, don't count out Intel and Via. Intel has again obtained an ARM-architecture license, after having turned its back on ARM in mid-2006 through the Xscale sale to Marvell, by virtue of its acquisition of Infineon's wireless group. Via also develops ARM-based SOCs for netbooks and other mobile-electronics devices. Although these chips seemingly contradict the company's x86-based products, you can make a credible argument that it's better to compete against yourself than against other companies.

REFERENCES

A Dipert, Brian, "x86 processors: Continued innovation is a welcome contradiction," *EDN*, Feb 4, 2010, pg 19, http://bit.ly/ik9EMK.

B Dipert, Brian, "Windows on ARM: for Intel, probably no cause for alarm," *EDN*, March 3, 2011, pg 10, http://bit.ly/eL4XSO.

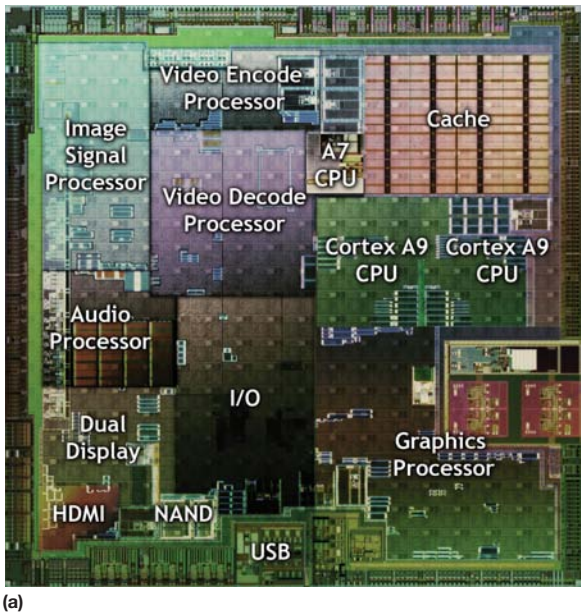

(a)

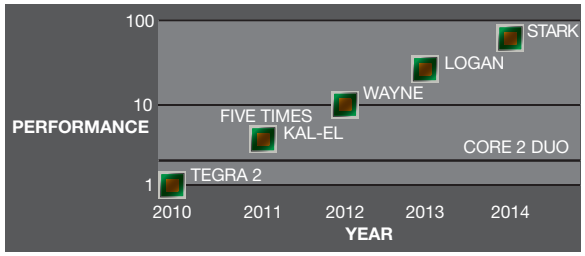

(b)

Figure 2 Nvidia's Tegra 2 combines two Cortex-A9 cores with shared L2 cache and a host of dedicated-function silicon resources (a). The company's aggressive product road map shows new releases on a yearly cadence, with dramatic performance boosts in each iteration (b).

dependent, per-CPU core basis.

Nvidia executives optimistically claim that Kal-El will be in production-grade tablets as early as August and that smartphone-inclusive products will be available in time for the holiday shopping season. Would-be customers and partners should use appropriate caution, however, considering overly enthusiastic corporate prognostications made in the past. Nvidia plans to introduce Kal-El successors, also with superhero-reminiscent names, on a yearly cadence, and the company claims that these products will lead to a roughly 100-times overall performance boost over Tegra 2 by 2014.

Mobile systems don't encompass all of Nvidia's ARM aspirations. The company recently boosted its SOC-license status to the architectural level, which it plans to harness on the supercomputer-targeted Project Denver. Company executives teased media attendees of CES about the existence of Denver, but industry observers don't currently know much about it.

TI MOVES

Long-standing ARM licensee Texas Instruments has to date followed a predictable SOC-development path, combining ARM CPU cores with Imagination Technologies PowerVR graphics cores and, in some cases, augmenting the cores' capabilities with internally developed dedicated-function logic blocks and general-purpose DSP resources. A traditional early adopter of each new ARM-core iteration, the company has admitted that it was late to the Cortex-A9—that is, OMAP 4—era, thereby opening the door for competitors such as Nvidia. Judging from the flurry of recent product announcements and customer adoptions, TI seems disinclined to repeat the same mistake.

Take OMAP 4, for example. TI announced at MWC that it was shipping the initial product-family member, the OMAP 4430, in production volumes, for use in systems such as the LG Optimus 3D and the first tablet from RIM (Research in Motion), the BlackBerry PlayBook (**Figure 3**). Like Nvidia's Tegra 2, the OMAP 4430 shares a common 1-Mbyte pool of L2 cache between two Cortex-A9 cores running at 1 GHz. However, the OMAP 4430 is also unique in several respects. For example, it employs a PowerVR SGX 540 GPU running at approximately 300 MHz—roughly 50% faster than that of competitors' SGX 540-based SOCs. It also interfaces to system memory over two 32-bit LPDDR2 ports operating at

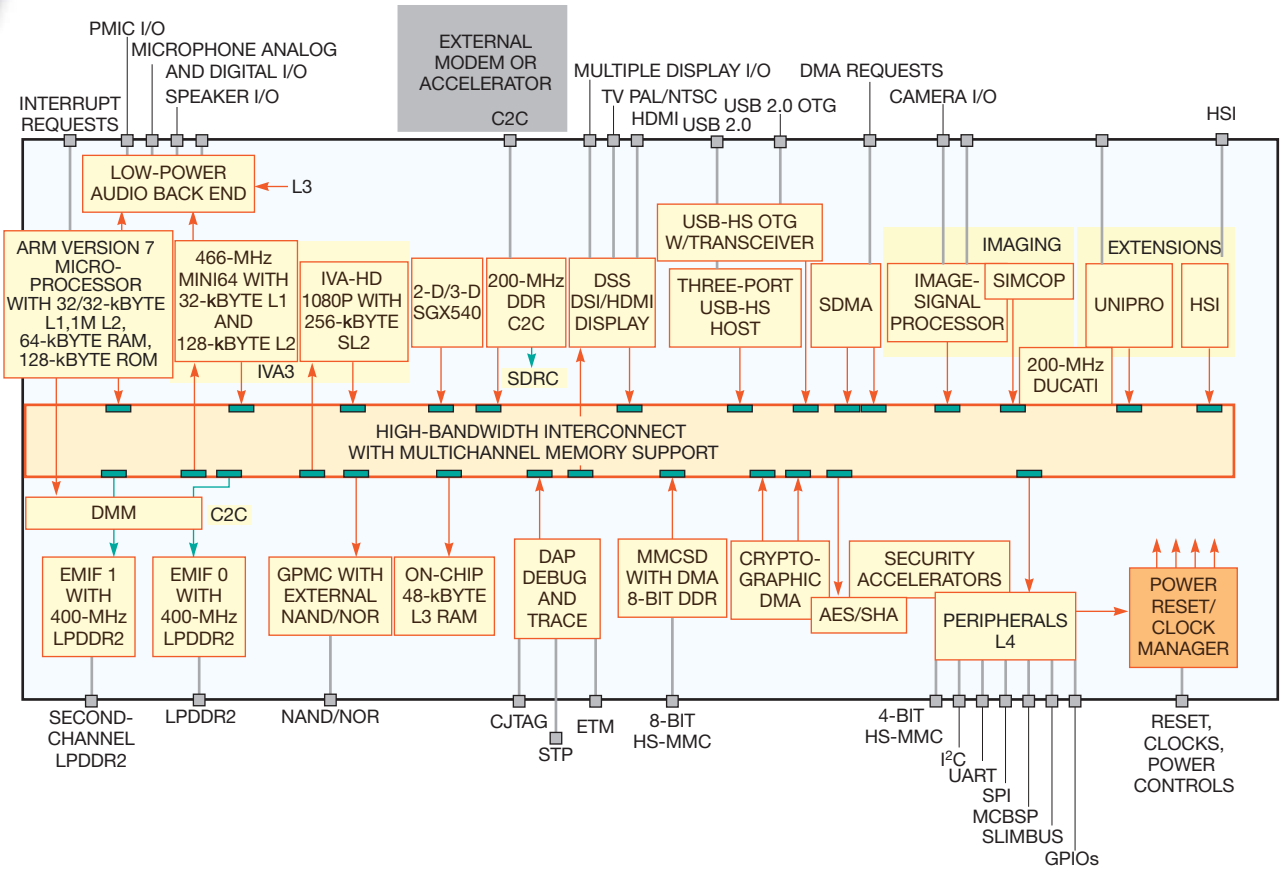

AES: ADVANCED ENCRYPTION STANDARD
C2C: CPU TO CPU
CJTAG: COMPACT JOINT TEST ACTION GROUP
DAP: DEBUG ACCESS PORT
DDR: DOUBLE DATA RATE
DMM: DYNAMIC MEMORY MANAGER
DRAM: DYNAMIC-RANDOM-ACCESS MEMORY
DSI: DIGITAL SERIAL INTERFACE
DSS: DISPLAY SUBSYSTEM
EMIF: EXTERNAL-MEMORY INTERFACE
ETM: EMBEDDED TRACE MACROCELL
GPIO: GENERAL-PURPOSE INPUT/OUTPUT
GPMC: GENERAL-PURPOSE MEMORY CONTROLLER
HD: HIGH DEFINITION
HDMI: HIGH-DEFINITION MULTIMEDIA INTERFACE
HS: HIGH SPEED
HSI: HIGH-SPEED INTERFACE
HSMMC: HIGH-SPEED MULTIMEDIA CARD
I²C: INTER-INTEGRATED CIRCUIT
IVA: IMAGING/VIDEO/AUDIO ACCELERATOR
L1: LEVEL ONE CACHE
L2: LEVEL TWO CACHE
L3: LEVEL THREE CACHE
L4: LEVEL FOUR CACHE
LPDDR: LOW-POWER DOUBLE DATA RATE
MCBSP: MULTICHANNEL BUFFERED SERIAL PORT
MMCSD: MULTIMEDIA CARD/SECURE DIGITAL CARD
NTSC: NATIONAL TELEVISION SYSTEM COMMITTEE
OTG: ON-THE-GO
PAL: PHASE-ALTERNATING LINE
RAM: RANDOM-ACCESS MEMORY
ROM: READ-ONLY MEMORY
SDMA: SYNCHRONOUS DIRECT-MEMORY ACCESS
SDRC: SYNCHRONOUS DRAM CONTROLLER
SHA: SECURE HASH ALGORITHM
SIMCOP: STILL-IMAGE COPROCESSOR
SL2: SECONDARY LEVEL TWO CACHE
3-D: THREE-DIMENSIONAL
2-D: TWO-DIMENSIONAL
UART: UNIVERSAL ASYNCHRONOUS RECEIVER/TRANSMITTER
UNIPRO: UNIFIED PROTOCOL

Figure 3 Comparatively late to the Cortex-A9 party, Texas Instruments is determined to play catch-up with OMAP 4. The OMAP 4430 employs a PowerVR SGX 540 GPU running at approximately 300 MHz.

speeds as high as 400 MHz, and each CPU core includes Neon SIMD vector floating-point capabilities in the form of ARM's MPE (media-processing engine).

In December, TI also unveiled the more advanced OMAP 4440, which the company fabricates, like its 4430 sibling, on a 45-nm process. Available for sampling now, with production slated for the second half of this year, it runs the Cortex-A9 cores at 1.5 GHz and delivers a 25% increase in overall graphics performance, according to TI. Other enhancements include 60-frame/sec, 1080p video-decoding capabilities; support for as many as two 12M-pixel camera sensors; an HDMI (high-definition-multimedia-interface) 1.4 video output; and automatic stereoscopic 3-D displays. At MWC, TI highlighted these capabilities, along with gesture-based user interfaces and integrated DLP (digital-light-projector) features, as Me-D experiences that rationalized OMAP 4's performance potential. The company also showcased Epos technology, which allows users to take notes and draw using a pen or a stylus. The writing tool comes with an Epos-patented ultrasonic transmitter, and the OMAP processor selects the transmitted signals using three microphones to accurately determine the location of the pen

GRAPHICS, THE OTHER KEY DISTINCTION FOR ARM-BASED SOCs

Stepping back and redirecting your corporate focus is sometimes a smart move. Such is the to-date corporate summary for graphics provider Imagination Technologies (formerly, VideoLogic), which, with partners such as NEC and STMicroelectronics, originally spent several largely fruitless years attempting to compete against ATI Technologies, now a division of AMD; Nvidia; and others in the PC-graphics market (references A and B). The company's tile-based-rendering approach is amenable to memory-constrained systems, such as mobile-electronics devices, in which absolute graphics-API (application-processor-interface) compliance is also not strictly necessary. So the company regrouped, reloaded, and aimed its attention at ARM. The resultant partnership between the two IP (intellectual-property) providers has for many years been mutually beneficial, as PowerVR-graphics technology found its way into most of ARM's SOC (system-on-chip)-based system designs.

Success, however, inevitably attracts competitors. Ironically, at least some of Imagination Technologies' competition nowadays comes from ARM itself, which in mid-2006 acquired another tile-based-rendering IP provider, Falanx Microsystems. After rebranding Falanx's products as Mali, the ARM-sourced graphics technology has to date found its most visible backer in Samsung, which employs it in the Exynos dual-core Cortex-A9 SOC.

Another PowerVR competitor, Qualcomm, not only does its own ARM CPU designs by virtue of its architecture license but also in 2009 purchased the former ATI Technologies' Xilleon handheld-graphics group, which it now rebrands as Adreno. Fellow ARM-architecture licensee Marvell isn't currently in Imagination Technologies' camp, either; instead, it's leveraging the graphics expertise of Vivante. And it's probably no surprise to learn that Nvidia is also using internally developed graphics cores in its Tegra series of ARM-based SOCs.

Do these developments have Imagination Technologies worried? Perhaps, but you wouldn't be able to tell from the calm demeanor of Marketing Vice President Tony King-Smith, who points out, for example, the steep performance, functionality, software compatibility, and other learning curves that Imagination Technologies has scaled and that other companies will also need to surmount to be credible alternative sources. King-Smith cites, for example, the BlackBerry Playbook as a successful PowerVR case study without explicitly admitting that other graphics suitors also had pursued RIM (Research in Motion).

To keep competitors in its rearview mirror, Imagination Technologies will need to keep developing unique, compelling new products. Addressing that concern, the company rolled out its Series 6 technology, code-named Rogue, at February 2011's MWC (Mobile World Congress). Rogue promises a 20- to 100-fold performance boost over today's quad-core graphics engines and counts ST-Ericsson among its initial licensees.

From a longer-term perspective, Imagination Technologies in mid-December announced its intention to acquire Caustic Graphics, a ray-tracing technology pioneer. Even if Imagination Technologies' ARM market share diminishes over time, King-Smith is confident that alternative-architecture SOC designs will more than pick up the slack. The company is a key MIPS partner, for example, and supplies the graphics-accelerator cores that Intel's Atom chip sets use.

REFERENCES

A Dipert, Brian, "Getting glitzy with graphics for embedded systems," *EDN*, April 1, 1999, pg 79, http://bit.ly/h6OrkH.

B Dipert, Brian, "Balancing in three dimensions," *EDN*, April 27, 2000, pg 54, http://bit.ly/eyxQbb.

or the stylus. This approach allows users to take notes either on the screen or off screen next to the mobile device and with or without paper or ink, according to TI (**Reference 5**).

Texas Instruments also announced plans last August to support the Cortex-A15, or Eagle, core architecture. It's probably no surprise, then, that the company was first to unveil product plans a week before MWC, even though the industry still knows little about Cortex-A15, aside from its support for extended memory addressing and hardware virtualization. TI implemented the core on a 28-nm process lithography and consolidated it under the OMAP 5 marketing umbrella. Initial Cortex-A15-based products will encompass two CPU cores running at frequencies as high as 2 GHz. Imagination Technologies' PowerVR SGX 544 GPU handles graphics-processing tasks, although core counts and clock speeds are not yet public information.

OMAP 5 integrates dedicated function blocks for video, imaging and vision, digital-signal-processing, display, and security tasks. The SOC also dedicates two ARM Cortex-M4 processors to real-time processing. Like OMAP 4, its system-memory interface comprises two 32-bit channels in at least two technology variants. The OMAP 5430 comes in a 14×14-mm POP (package-on-package) module with LPDDR2 memory support, and the OMAP 5432 comes in a 17×17-mm BGA package and supports faster DDR3/DDR3L SDRAM. TI most likely is targeting the OMAP 5430 at smartphones and the 5432 at tablets. Regardless, Texas Instruments is clearly confident in ARM's claims that Cortex-A15 applications can extend from servers to mobile systems. OMAP 5's core-shared L2 cache size doubles from that of OMAP 4 to 2 Mbytes.

INTERNAL SOC MUSCLE

To some, the word "Apple" and the phrase "semiconductor supplier" may not automatically go together, but the company has a long history of designing chip sets for its 68K- and PowerPC-based Macs. Commensurate with the transition to Intel-architecture microprocessors, the company also turned over core-logic development to its new CPU partner. Apple has now redirected its IC focus on its ARM-based products, such as the iPad, iPhone, iPod touch, and second-generation Apple TV. Presumably, Apple decided that its high product-shipment volumes rationalized internal IC-design projects. Eliminating the middleman semiconductor supplier's pricing markup would more than counterbalance these projects' R&D costs. Two key corporate acquisitions—of PA Semi three years ago and of Intrinsity—fueled the development of Apple's A4 SOC, which is currently in the iPhone 4, first-generation iPad, fourth-generation iPod touch, and second-generation Apple TV. Employing a single-core Cortex-A8 processor architecture, the A4 also embeds a PowerVR SGX 535 GPU and 640 kbytes of L2 cache memory. Its CPU clock frequency is system-dependent: 800 MHz on the iPod touch, for example, versus 1 GHz on the iPad.

The A4 CPU displaced Samsung-designed ARM SOCs in earlier iPhone and iPod touch systems. Ironically,

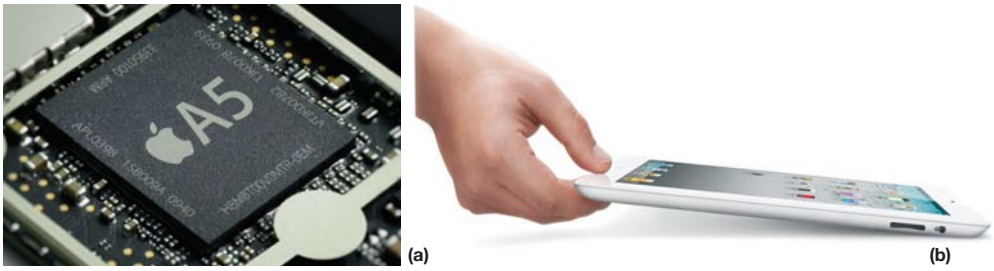

Figure 4 Apple's recently introduced dual-core 1-GHz A5 CPU (a) forms the processing nexus of the company's iPad 2 (b), although SOC specifics are currently scant.

however, Samsung continues to act as Apple's foundry for the A4 and supplies the DRAM in the SOC's POP module. Equally ironic, Samsung's S5PC110A01 Hummingbird SOC is similar to Apple's A4, with the exception of Hummingbird's L2-cache size of 512 kbytes and its inclusion of a PowerVR SGX 540 GPU versus the SGX 535 in the Apple A4. Die plots of the two ICs confirm that they're unique designs; nonetheless, the common lineage is indisputable.

Like Apple, Samsung's semiconductor group exclusively devotes its Hummingbird allocation to its handset division. At MWC, the company unveiled its latest-generation Galaxy S II high-end smartphone, which it based on the next-generation Exynos SOC. The dual-core, 1-GHz, Cortex-A9-based Exynos SOC, formerly code-named Orion, migrates from the PowerVR graphics in Hummingbird to ARM's own Mali 400MP GPU core (see **sidebar** "Graphics, the other key distinction for ARM-based SOCs").

As it turns out, however, at least two versions of the Galaxy S II will exist. The GT-I9100 model uses Exynos. Reflecting the "may-not-be-applicable-in-some-regions" qualifier on the spec sheet, however, an equivalent GT-I9103 variant employs the SOC engine of a semiconductor competitor, Nvidia's Tegra 2 (**Reference 6**). Samsung also based its 10.1-in. variant of the Galaxy Tab tablet series on Tegra 2, although the 7-in. model employs Hummingbird. The reason for the two-tier Galaxy S II differentiation is unclear, and Samsung isn't forthcoming with an explanation. It might, for example, be a function of low initial Exynos product yields, or it may reflect limited Samsung fab capacity due to the amount of foundry supply that partner and competitor Apple consumes.

Apple recently announced its next-generation A5 CPU and unveiled the iPad 2 (**Figure 4**). Industry participants currently know little of the SOC, aside from its 1-GHz clock speed, its dual-core CPU, and its graphics subsystem's ability to deliver as much as nine times the graphics performance of the PowerVR SGX 535 in the first-generation iPad, according to the company. As with the A4, the A5 may closely relate to the Cortex-A9-based Samsung Oxynox, with variance in L2 cache size, graphics-core version, and perhaps GPU supplier. As third-party developers obtain development systems and their hardware analyses leak into the public domain, you'll undoubtedly find out more about the A5's specifications and lineage.

FLEXIBLE DESIGN

Marvell and Qualcomm have more design flexibility than their competitors by virtue of their ARM-architecture licenses. Qualcomm's Scorpion microarchitecture has served the company well through several process-lithography generations, with many design wins over several years, across multiple customers, and across diverse mobile operating systems. However, its pseudo-Cortex-A9 performance capabilities have become dated in the era of true Cortex-A9 designs, especially the emerging multicore variants.

In response, Qualcomm chose February's MWC to unveil Krait, its next-generation proprietary core, which re-

tains compatibility with the ARM Version 7 instruction set. Initially targeting the 28-nm process node and running at per-CPU core speeds as high as 2.5 GHz, Krait will initially come in three product variants, each interfacing to system memory over a common dual-channel bus. It also serves as the launch platform for Qualcomm's next-generation Adreno GPU cores. First to become available for sampling this quarter is the MSM8960, whose dual CPU cores are asynchronously and independently controllable. The MSM8960 also contains the Adreno 225, which Qualcomm claims delivers eight times the performance of the original Adreno core.

Following the MSM8960 to the sampling pedestal are two other Krait proliferations, both of which Qualcomm scheduled for initial availability in early 2012. The single-core MSM8930 mates with the Adreno 305, which should deliver more than six times the performance of the original Adreno. At the product family's high end, the quad-CPU-core APQ8064 pairs with the integrated quad-core Adreno 320 GPU, with as much as 15 times the original Adreno's capabilities. All product-family members will integrate Wi-Fi, GPS (global-positioning-system), Bluetooth, and FM connectivity transceivers and will support NFC (near-field communication) and autostereoscopic 3-D still-image and video capture and playback. MSM products will additionally include LTE (long-term-evolution)-only or 3G (third-generation)/LTE combo cellular modems.

Marvell obtained its ARM-architecture license by virtue of its 2003 purchase of Asica, and most of its design team through the mid-2006 acquisition of Intel's ARM-based Xscale product line. Having expanded beyond its Intel-sourced PXA Series origins, the company's products also subdivide into the 100, 300, 500, 600, and 1000 Armada-family tiers, with a variety of feature differentiations. They use either

FOR MORE INFORMATION

Ambarella
www.ambarella.com

AMD
www.amd.com

Apple
www.apple.com

ARM
www.arm.com

BDTI
www.bdti.com

Boxee
www.boxee.tv

Caustic Graphics
www.caustic.com

Hewlett-Packard
www.hp.com

Imagination Technologies
www.imgtec.com

Intel
www.intel.com

LG Electronics
www.lg.com

Marvell
www.marvell.com

Microsoft
www.microsoft.com

MIPS
www.mips.com

Motorola
www.motorola.com

NEC
www.nec.com

Nokia
www.nokia.com

Nvidia
www.nvidia.com

Qualcomm
www.qualcomm.com

Research In Motion
www.rim.com

Samsung
www.samsung.com

ST-Ericsson
www.stericsson.com

STMicroelectronics
www.st.com

Texas Instruments
www.ti.com

Via Technologies
www.via.com.tw

Vivante
www.vivantecorp.com

the ARM Version 5-compliant Sheeva PJ1 or the ARM versions 6- and 7-compliant Sheeva PJ4 CPU-core technology. Other differentiating features include the number of per-SOC cores, along with their clock speeds; cache sizes and the number of cache levels; system-memory-bus width and speed; the degree of on-chip graphics- and video-processing horsepower; and the variety and number of other integrated peripherals.

To date, Marvell has had limited success in the computing and communications applications that this article discusses, with the exception of being a key supplier to RIM. The company keeps trying, however. Last September, for example, Marvell unveiled the Armada 628, containing three Sheeva PJ4 Cortex-A9-class CPU cores in a nod to their two-issue, limited instruction-reordering capabilities. Two of the cores run at 1.5 GHz, the third clocks in at 624 MHz, and the on-chip graphics can process 200 million triangles/sec. Two months later, Marvell unveiled a quad-core, 1.6-GHz Sheeva PJ4-based SOC containing as much as 2 Mbytes of L2 cache and targeting servers. Curiously, though, at the January 2010 CES, the company had claimed that its first quad-core ARM device would serve the mass consumer market and high-volume gaming applications (**Reference 7**).

Speaking of the consumer market, Marvell at MWC announced the 40-nm-based PXA978 World Phone communications processor, running at 1.2 GHz and including a cellular-modem supporting 3G UMTS (Universal Mobile Telecommunications System) and China's TD-SCDMA (time-division-synchronous-code-division/multiple access), with HSPA (high-speed-packet-access) support. EDN

REFERENCES

1 Dipert, Brian, "Ambarella's iOne: augmenting an image-processing foundation with ever-increasing integration," *EDN*, Dec 22, 2010, http://bit.ly/hDf8sM.

2 "Selecting Application Processors for Mobile Multimedia," Berkeley Design Technology, 2004, http://bit.ly/gJEzs8.

3 Davies, Chris, "Optimus 3D vs Optimus Tab: Not all HD 3D video is created equal," *SlashGear*, Feb 14, 2011, http://bit.ly/gEPmti.

4 Dipert, Brian, "Coming soon: 3-D TV," *EDN*, April 8, 2010, pg 24, http://bit.ly/dbBAnv.

5 "TI's OMAP platform spurs Me-D experiences: Welcome to the next mobile dimension," Texas Instruments, Feb 14, 2011, http://bit.ly/eNv3vg.

6 Galaxy S2 Specification, Samsung, http://bit.ly/ek8IKb.

7 "Marvell Announces Another Breakthrough in Chip Technology: World's First Quadruple Core Processor for ARM Instruction Set," Marvell, Jan 6, 2010, http://bit.ly/eZTxNk.

BY EUGENE R BUKOWSKI, JR • ACCENTURE

Innovative circuit designs target performance improvement and differentiation

PRODUCT INNOVATION IS THE PATH TO PRODUCT REJUVENATION. A HOLISTIC APPROACH COMPELS DESIGNERS TO CONSIDER INNOVATION AT ALL ARCHITECTURAL LEVELS. PRODUCT DESIGNERS AND ASIC DESIGNERS CAN TOGETHER ACHIEVE INNOVATION BREAKTHROUGHS.

Product innovation is the path to product rejuvenation. It is a deliberate expression of creativity that follows a scientific approach and a strategy that adds value and closes performance gaps. Relying on conventional techniques to reduce costs may fail to close performance gaps with a competitor because the competition may have discovered a new or an improved way of doing something.

Consumers are in an endless pursuit to do more with less and to get more for less. Products that can differentiate themselves from others can create competitive advantages. Innovative products differentiate themselves from the competition by providing better ways to deliver more value to the customer. Mature organizations know how to apply innovation to all segments of product architectures—from small components and packaging to the supply-chain process. A holistic approach to product innovation provides the best opportunity to close performance gaps, differentiate products, and reduce costs.

Electronic design is no stranger to innovation. The need to apply innovation to IC design has swelled with the deployment of ASICs. When economic advantage exists, many electronic designs now use ASICs. This trend means that a larger proportion of product innovation involving electronic design must occur at the semiconductor level. ASIC designs are one more level away from the true voice of the customer. This fact presents challenges to product designers because they are not typically directly involved in ASIC design.

The trend toward more ASIC usage also presents an opportunity for product innovation. Developers can exploit ASIC design through the selection and refinement of subcircuits and with their adaptations to new situations. Innovative circuit designs can close performance gaps and differentiate products. Product designers and ASIC designers must work together to optimize their success with product innovation in alignment with customer priorities.

PRODUCT INNOVATION

Product innovation occurs when someone uses an invention or a new idea to change how a current product works. Innovation can involve gradual, fundamental, or step-function product changes. Innovation is not about inventing something new, but it should be a creative process. Innovation involves finding new and improved ways of doing things for commercialization. Following an established method for innovation is significantly more reliable than simply employing serendipity.

Many versions of innovation exist, but most include a common set of steps (**Figure 1**). The organization must first validate the product need and the business case with a well-defined, investigated, documented, and communicated concept. The organization then uses creative procedures to generate a plethora of ideas, which it evaluates, ranks, and selects for further development. Selected ideas are developed into concepts and presented for further evaluation. The organization also evaluates concepts, ranking and selecting among them for further development. Selected concepts are the basis of prototypes for additional evaluation. Finally, the organization determines which one of the prototypes satisfies the original need and therefore justifies commercialization.

According to the Henderson-Clark model, innovation is separable into two dimensions (**Reference 1**). Modular innovation requires new knowledge involving one or more components of a design. Architectural innovation involves new methods in deploying the linkage among all the components of a design. Combining extents of the modular and architectural dimensions can result in various types of innovation. Breakthrough innovation occurs only when designers revolutionize both the modular and the architectural dimensions. For example, the advent of television remote control was a radical innovation because both the modular innovation—the infrared devices—and the architectural innovation—the remote-control electronics—were revolutionary.

Innovation may occur at multiple levels of product archi-

tecture. It is possible to segment the architecture of any product design into subarchitectures and further break down each segment as appropriate into multiple layers (**Figure 2**). Because each design component can have its own subarchitecture, breakthrough innovation may need to occur at multiple levels of the overall product architecture. Thus, breakthrough product innovation often requires radical innovation for electronic design.

Figure 1 Product innovation subdivides into several key steps.

OPPORTUNITIES FOR ASIC DESIGN

Many electronic designs now use ASICs rather than discrete components when economic incentive to do so exists. As a greater proportion of electronic systems' value-added content transfers to ASICs, radical innovation and differentiation become more difficult for product designers. Much of the opportunity for modular innovation continues to shift to ASIC designers. Therefore, product-design organizations that have cultivated close working relationships with their ASIC suppliers have developed a competitive innovation advantage over those without such relationships.

Many applications can no longer afford to partition subsystems into separate bipolar analog and MOS (metal-oxide-silicon) digital ICs. Economic incentive, application-size requirements, or both can combine analog and digital functions onto a single IC. This trend is especially true in the medical-device industry, which typically has a higher proportion of digital circuits than analog circuits; hence, fabrication of mixed-signal devices often occurs using MOS or mixed bipolar-MOS technologies.

Circuit designers once used innovation to develop a better way to build on-chip filters using MOS technologies. Signal-processing systems require precision amplification and filtering. These filters often require large values of resistance; however, standard MOS technologies did not offer this feature without significant economic penalty. Designing circuits with off-chip discrete components is undesirable because doing so consumes limited chip I/O and limited space on PCBs (printed-circuit boards).

Using switched-capacitor techniques for filters and signal processing is an example of breakthrough innovation that involves circuit design. Engineers applied this development to multiple cascading levels of electronic design for both the modular and the architectural dimensions of innovation (**Table 1**). Standard MOS capacitors switched at high frequencies using conventional digital clock signals can imitate large resistance values.

This fundamental change at the filter level represents a revolutionary form of modular innovation. The circuit designs required to configure these new filters for signal processing are revolutionary architectural innovations at various levels. Switched-capacitor technology at the ASIC level was a breakthrough innovation because it enabled the production of SOCs (systems on chips).

AUTOMATIC COMPENSATION

Process variation adversely affects every form of production, including that of ASICs. ICs exhibit relatively small process variation at any position across a die, meaning that you can match operating features with a high level of certainty across multiple devices within a die. For example, the threshold voltage for the same type of MOSFET (MOS field-effect transistor) is virtually identical at any position within a die.

ASIC-process parameters may concurrently present significantly large process variation among silicon-wafer lots. For example, MOSFETs' threshold voltage can vary ±20% between wafer lots. This potential process variation and other forms of variations can complicate generating designs that produce the same level of precise outputs for all acceptable outcomes of wafer production. In addition to process variation, analog-IC designers must also consider variation from the application environment. For example, variation in both power supply and temperature may affect the ability of analog subsystems to deliver consistently high performance.

You can apply techniques such as Lean Six Sigma for incremental continuous improvement to reduce the process variation of wafer fabrication. Lean Six Sigma combines lean manufacturing processes with Six Sigma, which seeks to improve the quality of process outputs by identifying and removing the causes of defects and minimizing variability in manufacturing and business processes. However, circuit designs also require uniquely competitive ways to add value. Analog-circuit designs must perform with precision even when MOS technologies and application environments present challenging variations. Electronic systems must be able to adjust their own configurations and compensate for situations. These needs require radical innovation.

Circuit-design techniques that automatically compensate for application and process variation might be more economical than alternatives. Organizations may also consider

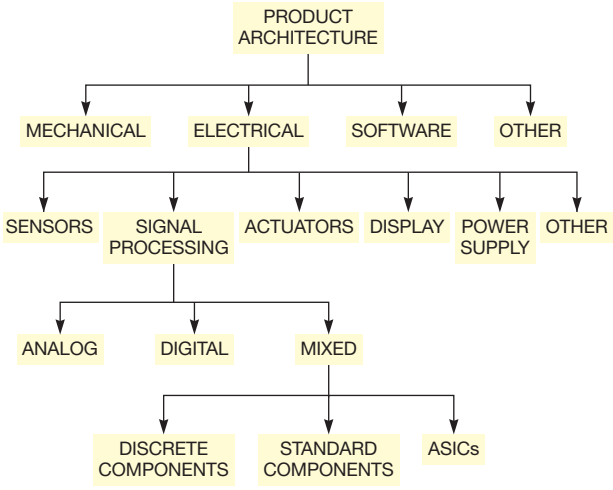

Figure 2 Product-architecture segmentation breaks down into multiple layers.

TABLE 1 CASCADING EFFECT OF THE DIMENSIONS OF INNOVATION

Level of product architecture	**Modular** (component)	**Architectural** (linkage)	Examples of breakthrough innovation
Electronic system	ASIC	Electronic design	Standard application chip sets, self-adjusting electronic designs
ASIC	MOS subsystem	Chip configuration	System on chip, automatic on-chip compensation
MOS subsystem	MOS subcircuit	Signal processing	Switched-capacitor signal processing, process-parameter signal processing
MOS subcircuit	Semiconductor device	Device configuration	Switched-capacitor filters, process-parameter sensors

changing application requirements or modifying MOS processes. However, these conventional approaches are sometimes impractical. Automatic on-chip compensation techniques for process variation could be the next breakthrough innovation involving electronic design. Although alternatives to this approach exist, product innovation compels designers to explore all such ideas.

To achieve breakthrough innovation for automatic compensation, both the modular and the architectural dimensions of innovation must be revolutionary. First, you must develop a library of sensors, involving modular innovation that can accurately measure a variety of application and on-chip process parameters. These sensors represent new "components." You must then use architectural innovation to develop new circuit-design techniques to use this new information for the other value-added subsystems. These approaches represent a new way of configuring subsystems to be self-adjusting and "smart" to compensate for unfavorable operating conditions.

ON-CHIP MOSFET THRESHOLD VOLTAGE

A subcircuit for the on-chip sensing of process variation is one example of modular innovation (**Figure 3**). This subcircuit captures the exact value of the on-chip P-channel MOSFET threshold voltage with zero body effects, V_{T0}. This technique presents the threshold voltage as an exact voltage quantity, a bias current that a resistor value establishes, or an exact rational-number multiple of threshold voltage. Although this version applies to N-well CMOS technologies, you can also apply it to P-well, dual-tub, and bipolar-CMOS technologies.

The threshold-voltage sensor contains uniquely biased P-channel transistors and a feedback loop with high gain for control. You must isolate each P-channel transistor in its own N well. This approach became possible at no extra cost only with the advent of CMOS process technologies. The new way of physically laying out a MOS transistor and eliminating the threshold voltage's body effect represents a level of modular innovation at the MOS-subcircuit level. The P-channel transistor represents a component of the sensor.

The feedback loop is a subarchitecture that uses the qualities of the P-channel devices. The op amp controls the steady-state bias of the P-channel transistors to capture and present the exact on-chip threshold voltage. This technique represents the architectural innovation of a linkage for the sensor. The subcircuit becomes a component that enables other MOS subsystems to self-adjust for process variation.

To reveal the exact on-chip threshold voltage of the P-channel transistor, this innovation first requires a means of preventing the body effect from occurring. The body effect occurs when the bulk-to-source voltage of a P-channel transistor is greater than 0V. The body effect causes the threshold voltage to vary as a function of the bulk-to-source voltage. The source channel of a P-channel transistor in its own N well, or body, can connect directly to its N-well terminal, thereby preventing the body effect. As a result, the threshold voltage of two matched P-channel transistors at the same position of a die remains equal regardless of how they are biased in the circuit.

The threshold-voltage sensor performs by means of the high-gain op amp, which uses negative feedback through a source follower to ensure that the voltage to the negative input of the op amp remains exactly equal to one-half of the power supply ($V_{DD}/2$) in steady state. The two matched resistors at the positive input of the op amp establish the reference voltage. Thus, the P-channel transistors, in the absence of the body effect, have an operating point that must satisfy the following **equation**: $(1\times W/L)(V_{SG1}-V_{T0})^2=(4\times W/L)(V_{SG2}-V_{T0})^2=(4\times W/L)(V_{SG3}-V_{T0})^2$, where W/L is the transistor's width-to-length, or aspect, ratio and V_{T0} represents the on-chip P-chan-

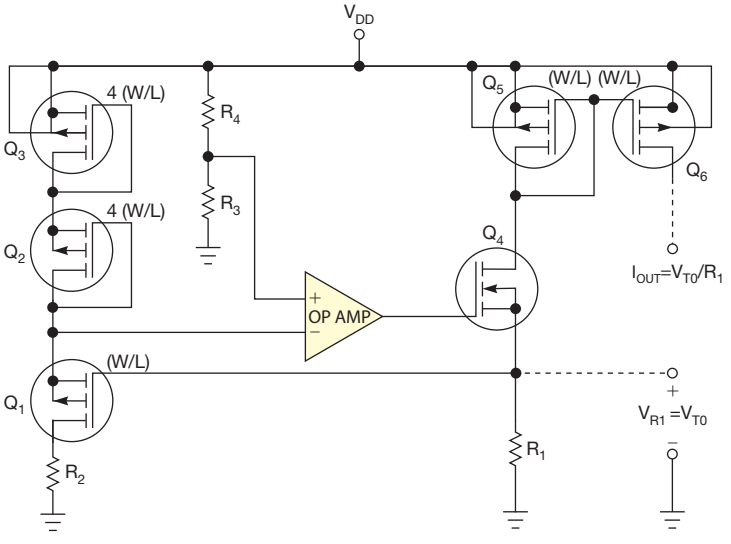

Figure 3 A threshold-voltage subcircuit for the on-chip sensing of process variation is one example of modular innovation.

DESIGN NOTES

High Efficiency, High Density 3-Phase Supply Delivers 60A with Power Saving Stage Shedding, Active Voltage Positioning and Nonlinear Control for Superior Load Step Response

Design Note 489

Jian Li and Kerry Holliday

Introduction

The LTC®3829 is a feature-rich single output 3-phase synchronous buck controller that meets the power density demands of modern high speed, high capacity data processing systems, telecom systems, industrial equipment and DC power distribution systems. The LTC3829's features include:

- 4.5V to 38V input range and 0.6V to 5V output range
- 3-phase operation for low input current ripple and output voltage ripple with Stage Shedding™ mode to yield high light load efficiency
- On-chip drivers in a 38-pin 5mm × 7mm QFN (or 38-pin FE) package to satisfy demanding space requirements

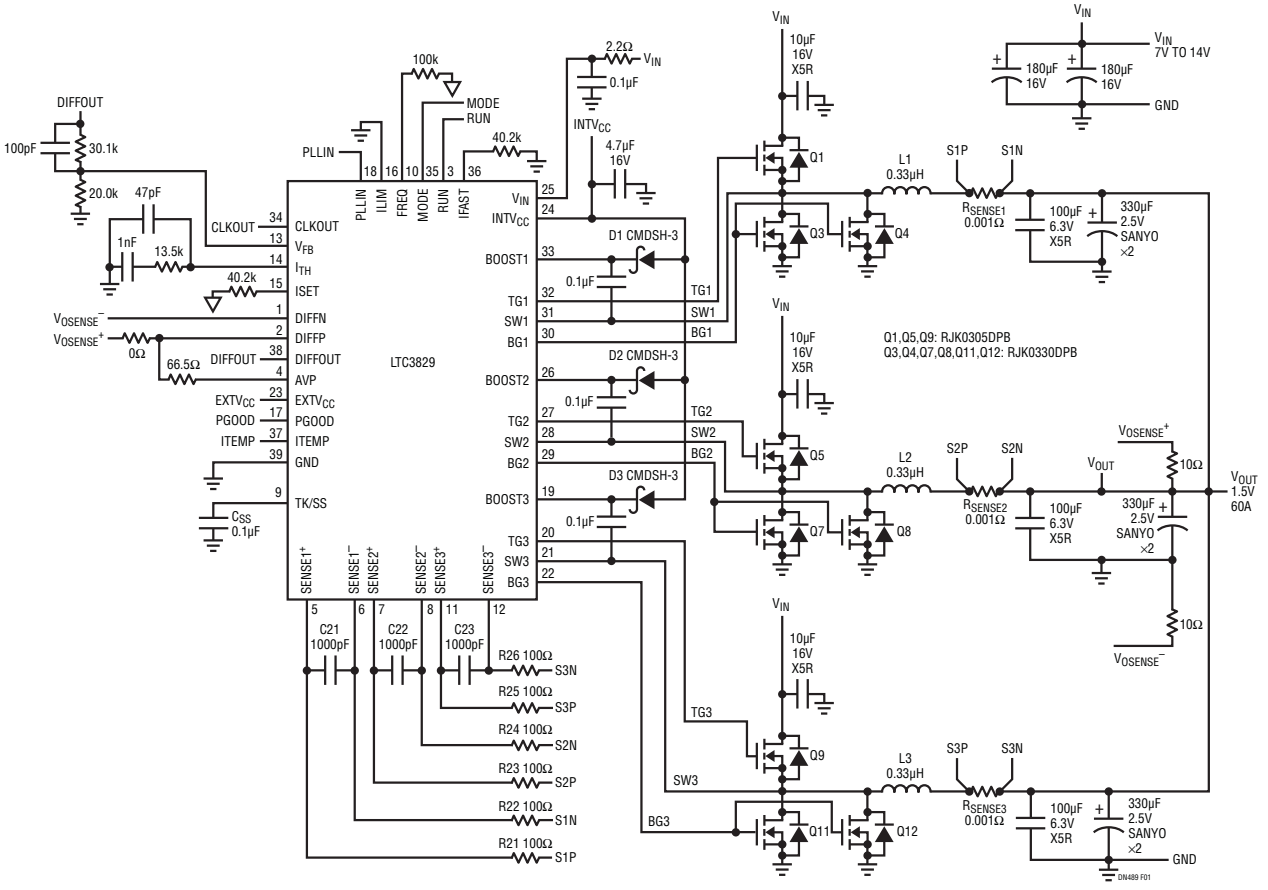

Figure 1. A 1.5V/60A 3-Phase Converter Featuring the LTC3829

- Remote output voltage sensing and inductor DCR temperature compensation for accurate regulation
- Active voltage positioning (AVP) and nonlinear control ensure impressive load transient performance

1.5V/60A, 3-Phase Power Supply

Figure 1 shows a 7V to 14V input, 1.5V/60A output application. The LTC3829's three channels run 120° out-of-phase, which reduces input RMS current ripple and output voltage ripple compared to single-channel solutions. Each phase uses one top MOSFET and two bottom MOSFETs to provide up to 20A of output current.

The LTC3829 includes unique features that maximize efficiency, including strong gate drivers, short dead times and a programmable Stage Shedding mode, where two of the three phases shut down at light load. Onset of Stage Shedding mode can be programmed from no load to 30% load. Figure 2 shows the efficiency of this regulator at over 86.5% with a 12V input and a 1.5V/60A output with Stage Shedding mode, dramatically increasing light load efficiency.

The current mode control architecture of the LTC3829 ensures that DC load current is evenly distributed among the three channels, as shown in Figure 3. Dynamic, cycle-by-cycle current sharing performance is similarly tight in the face of load transients.

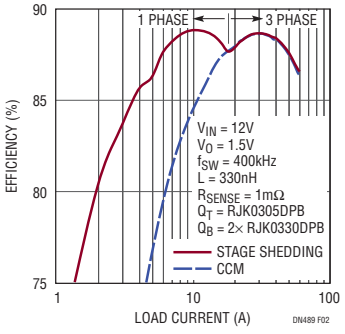

Figure 2. Efficiency Comparison of Stage Shedding vs CCM

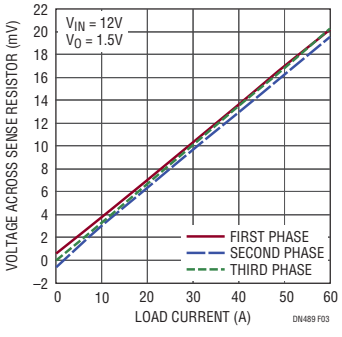

Figure 3. Current Sharing Performance Between Phases

Data Sheet Download

www.linear.com

A fast and controlled transient response is another important requirement for modern power supplies. The LTC3829 includes two features that reduce the peak-to-peak output voltage excursion during a load step: programmable nonlinear control or programmable active voltage positioning (AVP). Figure 4 shows the transient response without these features enabled. Figure 5 shows that nonlinear control improves peak-to-peak response by 17%. Figure 6 shows that AVP can achieve a 50% reduction in the amplitude of voltage spikes.

Conclusion

The LTC3829's tiny 5mm × 7mm 38-pin QFN package belies its expansive feature set. It produces high efficiency with a combination of strong integrated drivers and Stage Shedding/Burst Mode® operation. It supports temperature compensated DCR sensing for high reliability. AVP and nonlinear control improve transient response with minimum output capacitance. Voltage tracking, multichip operation and external sync capability fill out its menu of features. The LTC3829 is ideal for high current applications such as telecom and datacom systems, industrial and computer systems.

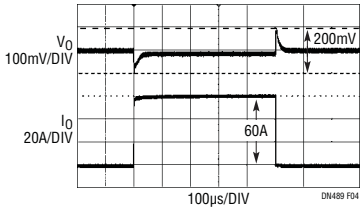

Figure 4. Transient Performance without AVP and Nonlinear Control

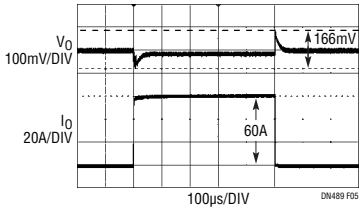

Figure 5. Transient Performance with Nonlinear Control

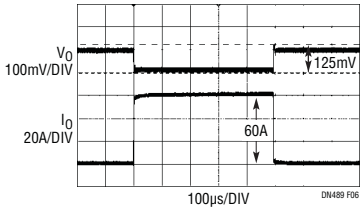

Figure 6. Transient Performance with AVP

For applications help,
call (408) 432-1900, Ext. 3967

nel MOSFET threshold voltage with no body effect—that is, when the bulk-to-source voltage is 0V. The process parameters in the **equation** are equal and cancel each other out.

Note that P-channel transistors Q_2 and Q_3 match. Each resides in its own N well, and their source terminals connect to their respective bulk terminals. In addition, their gate terminals connect directly to their drain terminals. Their function is somewhat similar to that of two matched resistors. Therefore, their gate-to-source voltages, which in this configuration also equal their drain-to-source voltages, are exactly equal to one-fourth of the power-supply voltage: $V_{SG2}=V_{SG3}=V_{DD}/4$.

Also note that the feedback transistor Q_1 also matches transistors Q_2 and Q_3, except that its aspect ratio is only one-fourth of their values. It also resides in its own N well, and its source terminal connects directly to its bulk terminal. Due to the squared relationship between the gate-to-source voltage and the drain current of the MOSFET, this orientation means that the gate-to-source, or overdrive, voltage of Q_1 beyond V_{T0} must be exactly twice that of either transistor Q_2 or transistor Q_3. Thus, the circuit has an operating point that must satisfy the following **equation**: $[(V_{DD}/2)-V_{SG1}]=V_{T0}$. Because the source potential of feedback transistor Q_1 must remain equal to one-half of the power supply in steady state, the voltage quantity, which the source follower forces at the gate terminal of Q_1, must remain exactly at V_{T0} in steady state. This sensor is an example of revolutionary modular innovation.

MANY ELECTRONIC DESIGNS NOW USE ASICs WHEN ECONOMIC INCENTIVE TO DO SO EXISTS.

To eliminate the effect of channel-length modulation, you can replace resistor R_2 with another P-channel transistor that matches the aspect ratios of transistors Q_2 and Q_3. This technique ensures that the drain-to-source voltage of transistor Q_1 remains exactly one-fourth of the power supply in steady state, just like that of transistors Q_2 and Q_3. Furthermore, you can replace the two matched resistors at the positive input of the op amp with two matched P-channel transistors.

NEXT STEPS

To achieve breakthrough innovation that automatically compensates for wafer-lot process variation, both the modular and the architectural dimensions of innovation must be revolutionary. Therefore, building a complete component library that can sense adverse variation and unfavorable operating conditions requires many more subcircuits like the threshold-voltage sensor. In addition, revolutionary circuit configurations are necessary to effectively use this information to improve on-chip system performance. This form of robust self-adjusting electronics could be the next radical innovation in electronic design.

Perhaps this trend marks the beginning of a new conversation between product designers and ASIC suppliers. Recognizing the need for innovation begins with the voice of the customer, and the product-design organization should know the customer better than anyone else. To achieve breakthrough innovation for product applications, you must drive the innovation all the way down into the ASIC level. Using a holistic approach to innovation, product designers and ASIC designers should work together for success and consider customer priorities. EDN

REFERENCE

1 Henderson, Rebecca M, and Kim B Clark, "Architectural Innovation: The Reconfiguration of Existing Product Technologies and the Failure of Established Firms," *Administrative Science Quarterly*, Vol 35, pg 9, March 1990, http://bit.ly/ewuTJT.

AUTHOR'S BIOGRAPHY

Eugene Bukowski is a manager and a management consultant at Accenture for the company's Process and Innovation Performance service line. Prior to joining Accenture, he commercially designed analog MOS ICs for both IBM and General Motors. He also has published 10 inventions, six of which are patented. Bukowski has electrical engineering degrees from Purdue University (West Lafayette, IN) and Duke University (Durham, NC), along with a master's degree in management from Kettering University (Flint, MI). In addition, he is a Lean Six Sigma Master Black Belt, including Design for Lean Six Sigma, a Shainin Red X Master problem solver, and a senior member of the American Society for Quality.

FRANCK NICHOLLS • FREESCALE SEMICONDUCTOR

The e-reader paradox

E-BOOK READERS REQUIRE HIGH PERFORMANCE, EVEN TO JUST DISPLAY TEXT.

Many people assume that e-book readers require only minimum processing power to render a basic text page, but this assumption is wrong. Displaying text may be easy, but displaying it quickly—at the speed of paper—requires high performance. To ensure the wide acceptance of e-readers, they must provide a user experience closer to that of reading a good-old paper book and to the experience that traditional consumer-electronics devices offer. Today's e-readers face these challenges: to be as quick as paper and as quick as LCD-based devices. High-performance microprocessors, such as those operating at 800 MHz on the ARM Cortex-A8 architecture, can enable better usability and new use scenarios for e-readers.

AS QUICK AS PAPER

Go into your favorite electronics store and select two e-readers. Press the next-page button and note how long it takes for a page to turn. You may see a noticeable difference between the two readers. Page-turn time is the biggest challenge for e-readers. Although you may tolerate the fact that your PC takes a full second to render a page full of pictures, it becomes painful when you need to wait in the middle of a sentence, especially when this happens 300 times in a row in a single book.

It may seem paradoxical, but users expect a higher performance from a light e-reader than from a big notebook because they are comparing the e-reader to a paper book, not to other electronic devices. A printed book is an instant-on device; as soon as you turn the page, the next one is before your eyes. Therefore, e-readers require a high-performance processor, even for text only.

Table 1 compares the classes of processors that will find use in e-readers this year. Because of the need for low-power consumption and compactness, all use the ARM architecture. The **table** shows that the recent ARM Cortex-A8 core enables four times the speed of the low-end ARM9. Processors confirm this increase. **Table 2** shows the time these processor take to open a PDF (portable-document-format) file. The benchmarks use a standard LCD screen.

One way to secure a maximum bandwidth on the CPU is to offload it from all display-control tasks. E-reader screens require a lot of preprocessing. The current generation uses an external controller that sits between the processor and the display, but this controller is expensive.

A next-generation chip, such as Freescale Semiconductor's (www.freescale.com) i.MX508 e-reader processor, integrates the controller directly in the silicon, in hardware rather than in software, to help ensure both high performance and low power consumption. CPU performance has a direct impact on the time it takes to render a new page, whether as text or as an image.

AS QUICK AS AN LCD

Most e-readers use EPDs (electrophoretic, or electronic-paper, displays), and the dominant player is E Ink (www.eink.com). EPDs have many benefits over LCDs for reading e-books. For example, EPDs are bistable, which means that they can hold an image without the need for updating and therefore consume no power except when turning a page. They are also reflective, meaning that they require no backlight, making it more comfortable for a reader's eyes and enabling readability in sunlight. In short, they are like paper (**Figure 1**).

This technology has downsides, however. First, EPD technology has a slow display-frame rate. Changing the color

TABLE 1 ARM-BASED PROCESSOR COMPARISON

Type of processor	Typical speed (MHz)	DMIPS/MHz	Dhrystone performance (DMIPS)
ARM9 (ARM926EJ-S)	400	1.06	424
ARM11 (ARM1136J-S)	532	1.18	628
ARM Cortex-A8	800	2.07	1656

TABLE 2 BENCHMARK TO OPEN AND RENDER A PDF

Type of processor	65 pages of text (243-kbyte PDF)	Four pages of color ads (731-kbyte PDF)	Benchmark
ARM9 (Pocketbook, Fnacbook, Oyo, Bookeen, and others)	Less than 2 sec	Approximately 10 sec	400-MHz ARM9 133-MHz DDR2
ARM11 (Kindle, Nook, Sony Reader)	Approximately 1.5 sec	Approximately 6 sec	532-MHz ARM11 133-MHz DDR
ARM Cortex-A8 (next-generation 2011; for example, Freescale i.MX508 processor)	Less than 1 sec	Approximately 2 sec	800-MHz Cortex-A8 200-MHz DDR

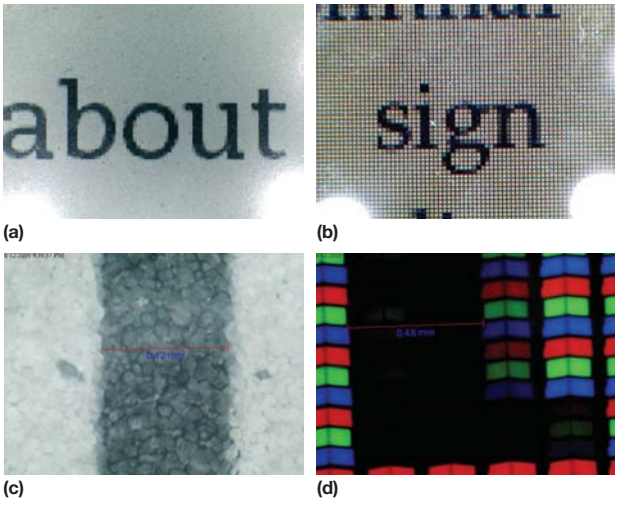

Figure 1 EPDs have an advantage over LCDs for reading e-books. Compare an EPD at 60-times magnification (a) to an LCD at the same magnification (b) and an EPD at 400-times magnification (c) to an LCD at the same magnification (d).

of a microcapsule from black to white requires a waveform operation. EPDs have several waveform modes for gray-scale depth, and they also have various update times; longer waveforms produce better gray-scale accuracy. For instance, a black-and-white update of a microcapsule takes 260 msec, but a 16-level gray-scale update requires 600 msec on an 85-Hz E Ink display.

EPDs also generate "ghosting"—the persistence of the previous microcapsule's state. To minimize ghosting, developers use a 16-gray-scale update of the entire display to drive all microcapsules, not just those that need to change to a known state—typically, black—before migrating them to the final image. For this reason a "flash" takes place on each page turn (**Figure 2**).

POWERFUL PROCESSORS CAN NOW DECODE A PAGE FASTER THAN EPDs CAN RENDER IT. ELECTRONICS ARE NO LONGER THE LIMITATION.

Developers can avoid these drawbacks and make an EPD look as if it's performing like an LCD, but these approaches require preprocessing and subsequently higher performance. One approach is the quick page flip. The black flash and the E Ink waveform's update time limit the page-turn speed. As a result, users cannot quickly flip through an e-book. To enable page flipping just as with a paper book, developers use a partial update to update only those pixels that need to change. It enables 4 frames/sec; a black-and-white update lasts 260 msec. Because this approach may generate ghosting, the EPD must perform a full update with flash when the user stops flipping to clear the screen. The challenge is then to have enough upfront processing power to quickly send new pages to the EPD.

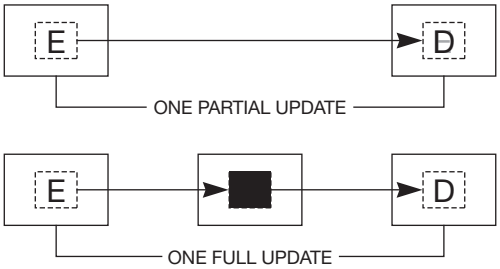

Figure 2 The black flash minimizes the ghosting effect.

Another approach to making an EPD mimic the performance of an LCD is with animation. Although an E Ink screen at 85 Hz can process approximately 4 frames/sec with a 260-msec update time, it is still far from the LCD user experience. However, applications can use regional updates to improve the effective frame rate of the display. A regional update refreshes only a portion of the screen, down to one pixel.

For instance, the i.MX508 e-reader processor embeds an E Ink controller that enables 16 concurrent updates, meaning that 16 regions of the screen can change at once. The theoretical maximum frame rate is then 64 updates/sec—16 parallel black-and-white updates every 260 msec. Although this technology cannot enable standard video playback, it can drive animations for enhanced user interfaces. The prerequisite is to have enough performance upfront for the EPD controller to deal with the parallel updates and the subsequent collisions; hence, a power processor is necessary. New e-readers will support new LCD-like user interfaces and still offer all the benefits of electronic paper.

A few e-readers, including a Sony device, now use regional updates. The handwriting application requires no flashing to draw lines; the screen concurrently receives small regional updates the size of a few dots. With the generation of processors using the ARM Cortex-A8 architecture, readers will be able to draw at the same speed on EPDs as on LCDs.

PHYSICS VERSUS ELECTRONICS

New e-readers will require high-end processors because speed—especially page-turn speed—will be higher. Powerful processors can now decode a page faster than EPDs can render it. Electronics are no longer the limitation. Physics will remain the bottleneck. Higher processing performance paves the way to new user interfaces and new applications, thanks to clever use of EPDs. With the creation of application stores for e-readers, developers will find in microprocessors a field of innovation. A new generation of color EPDs will soon hit the market, enabling higher frame rates; e-reader processors must keep headroom for the near future. EDN

AUTHOR'S BIOGRAPHY

Franck Nicholls is a product marketer at Freescale Semiconductor, where he is responsible for marketing the company's application processors to consumer products in Europe. His two areas of interest are tablets and e-readers. Nicholls graduated from Telecom ParisTech (Paris, France).

88dB SFDR @ 100MHz

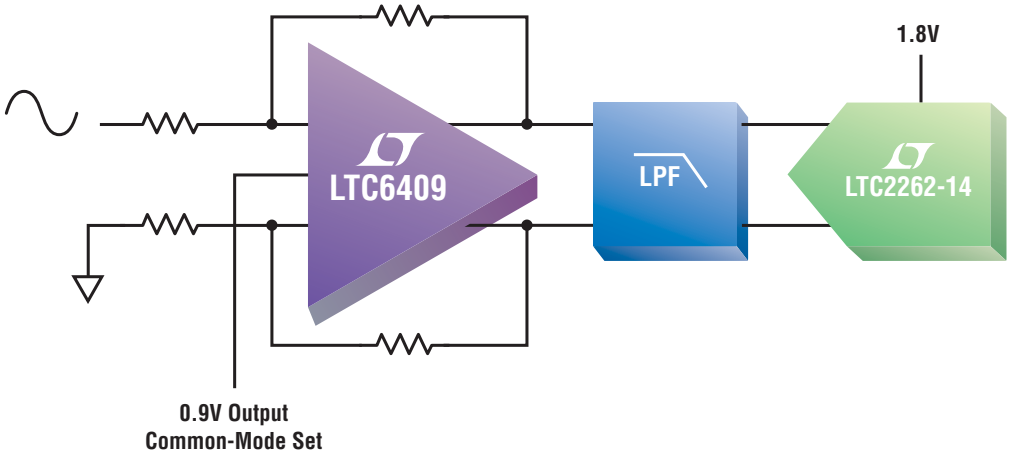

Unleash Your High Speed ADC

With 1.1nV/√Hz input noise density and 88dB SFDR performance at 100MHz, the LTC®6409 enables your high speed ADC to achieve outstanding performance. Its input common mode range includes ground and its output common mode can be set as low as 0.5V, making the LTC6409 the perfect choice for driving AC- or DC-coupled signals into the latest 1.8V data converters. Fully specified over the –40°C to 125°C temperature range, and available in a tiny 3mm x 2mm QFN package, the LTC6409 combines excellent AC performance with flexibility, robustness and a minimal footprint.

Features

- Unity Gain Stable
- 1.1nV/√Hz Input Noise Density
- 10GHz GBW @ 100MHz
- DC- or AC-Coupled Inputs
- 0.5V to 3.5V V_{OCM}
- –40°C to 125°C Fully Specified
- Tiny 3mm x 2mm QFN Package

Differential ADC Drivers

Part Number	80dBc HD2/HD3 (MHz)	Input Referred Noise (nV/√Hz)	Voltage Gain
LTC6409	110	1.1	R-set
LTC6406	30	1.6	R-set
LTC6404-1	15	1.5	R-set
LTC6400-20	120	1.9	20dB
LTC6416	90	1.8	0dB

Info & Free Samples

www.linear.com/6409

1-800-4-LINEAR

www.linear.com/ampsflyer

LINEAR TECHNOLOGY

EDITED BY MARTIN ROWE
AND FRAN GRANVILLE

designideas

READERS SOLVE DESIGN PROBLEMS

Low-component-count logic probe works with TTL and CMOS logic

Aruna Rubasinghe, University of Moratuwa, Sri Lanka

The circuit in **Figure 1** uses the LM358 dual op amp running as a comparator, plus a few other inexpensive components, to make a TTL (transistor-transistor-logic)/CMOS-logic probe. The circuit gets its power from the circuit under test, which lets it work with TTL or CMOS logic. The IC_{1A} and IC_{1B} op amps come in an LM358 package. Switch S_1 selects the TTL or the CMOS mode of operation. The green LED shows logic low, and the red LED shows logic high.

The noninverting input of IC_{1A} and the inverting input of IC_{1B} connect to the test probe. The circuit uses 90% of the power-supply voltage as CMOS-logic high and 2.7V as TTL high. It uses 0.7V as logic low for both TTL and CMOS because their logic-low levels approach 0.7V. Voltage divider R_3/R_4 divides voltage from 2.7V zener diode D_1, providing 1.35V at IC_{1A}'s noninverting input and IC_{1B}'s inverting input. Diode D_2 is forward-biased, and the 0.7V voltage across D_2 becomes the lower limit, which represents logic low. You set this voltage on the noninverting input of IC_{1B}.

In TTL mode, the voltage at the inverting input of IC_{1A} is 2V. In CMOS mode, the voltage at the inverting input of IC_{1A} is nearly 90% of the input voltage through voltage divider R_6/R_7. When the probe is in its high-impedance state in either CMOS or the TTL mode, IC_{1A}'s inverting input voltage is greater than its 1.3V noninverting input voltage. IC_{1A}'s output is low. IC_{1B}'s 1.3V inverting input voltage is greater than its 0.7V noninverting input voltage. The output of IC_{1B} is also low, and both LEDs are off.

In TTL mode, when measuring logic high, IC_{1A}'s 2.7V inverting input voltage is less than its noninverting input voltage, which is the probe voltage. IC_{1A}'s output is high. IC_{1B}'s inverting input voltage, which is the probe voltage, is greater than its 0.7V noninverting input voltage. IC_{1B}'s output is, therefore, low. The red LED turns on, indicating logic high.

When measuring logic low, IC_{1A}'s 2.7V inverting input voltage is greater than the voltage at its noninverting input, which is the probe voltage. Thus, IC_{1A}'s output is low. IC_{1B}'s inverting input voltage, which is the probe voltage, is greater than its 0.7V noninverting input voltage. Thus, IC_{1B}'s output is high. The green LED turns on, indicating a logic low.

In CMOS mode, when measuring logic high, IC_{1A}'s inverting input voltage, which is 90% of the supply voltage, is greater than the voltage at its noninverting input. The output is thus high. IC_{1B}'s inverting input voltage, which is the probe voltage, exceeds that of its 0.7V noninverting input voltage, and IC_{1B}'s output is low. The red LED turns on, indicating logic high.

When measuring logic low, IC_{1A}'s inverting input voltage, which is 90% of the supply voltage, exceeds the voltage at its noninverting input. IC_{1A}'s output is low, and IC_{1B}'s output is high because its inverting input voltage is higher than 0.7V at its noninverting input voltage. The green LED turns on, indicating logic low. When the probe's pin is pulsing, both LEDs alternately turn on and off at the pulse frequency. EDN

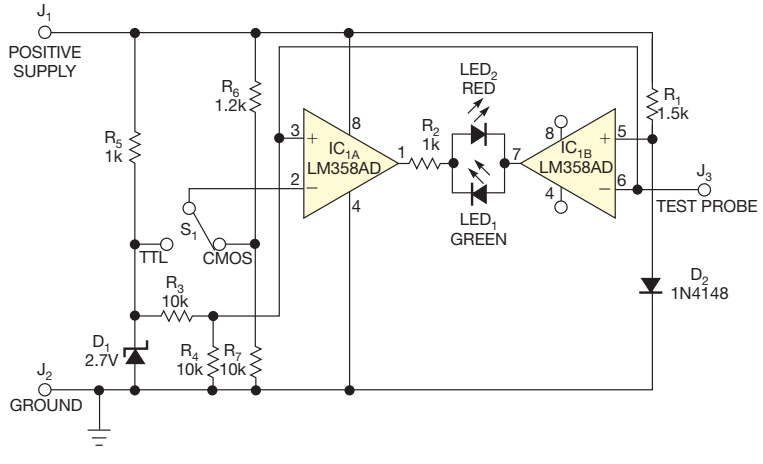

Figure 1 Two comparators and some voltage dividers determine the status of a TTL or a CMOS-logic signal.

DIs Inside

46 Circuit implements photovoltaic-module simulator

48 Switch circuit controls lights

50 Isolated PWM suits low frequencies

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

Circuit implements photovoltaic-module simulator

José M Blanes and Ausiàs Garrigós, University Miguel Hernández, Elche, Spain

Electronics engineers often use photovoltaic-module simulators to test dc/dc-power converters, inverters, or MPPT (maximum-power-point-tracking)-control techniques. The use of these simulators lets you work in the laboratory with predefined photovoltaic conditions, thus avoiding the drawbacks of real photovoltaic modules. Various commercial simulators are available, but they are often expensive.

This Design Idea presents a simple circuit that works as a photovoltaic-module simulator using a dc-voltage source as its input. The circuit employs the simplest equivalent circuit of a photovoltaic module: a current source in parallel with a diode. The output of the current source is directly proportional to the irradiance, and the characteristics of the parallel diode change with temperature.

The current source in **Figure 1** operates at a short-circuit current as high as 2A and employs two MJ15023 power bipolar transistors, Q_2 and Q_3, working in the linear region. The value of the current is proportional to the transistor's base current. You can control the short-circuit current with potentiometer R_2, whose change in value forces a variation in Q_1's collector current and, thus, the base current of Q_2 and Q_3.

The circuit has 50 MUR1520 diodes that connect in series and are in parallel with the current source. You can short-circuit these diodes in groups of 10, so you can choose the number of series diodes: 10, 20, 30, 40, or 50. If you need more precision, you can increase the number of diode groups in the simulator, providing more output curves.

Figures 2 and **3** show the current- and power-voltage characteristics of

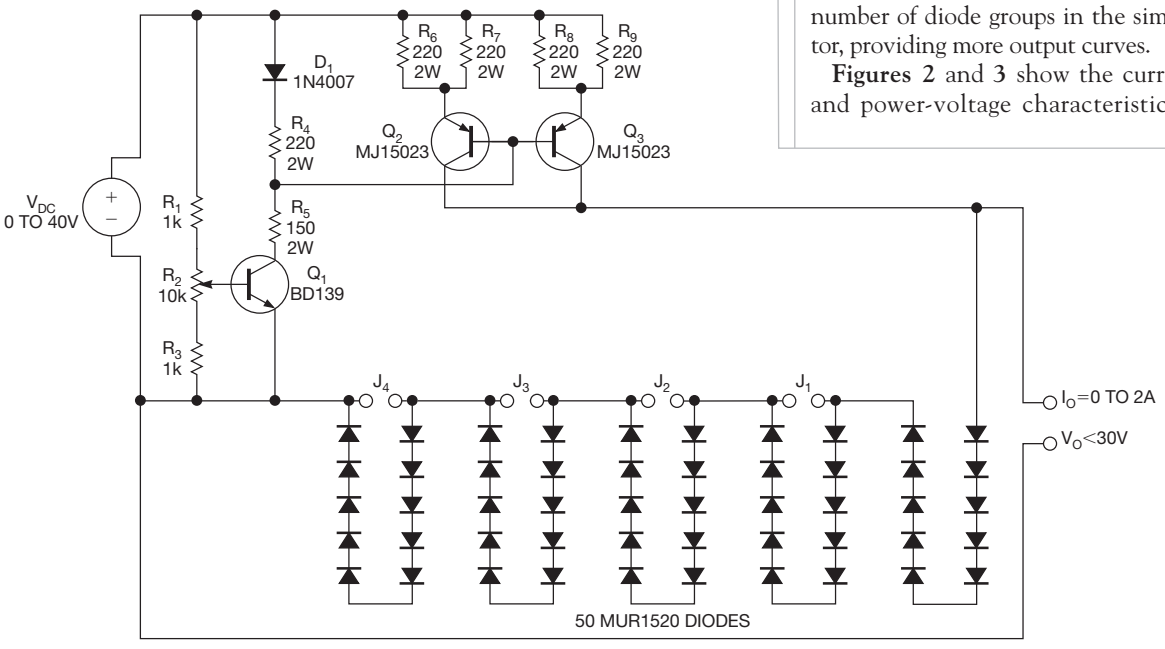

Figure 1 You can change the photovoltaic characteristics of this simulator by modifying the current-source output value or connecting 10 to 50 diodes in series. You can short-circuit these diodes in groups of 10.

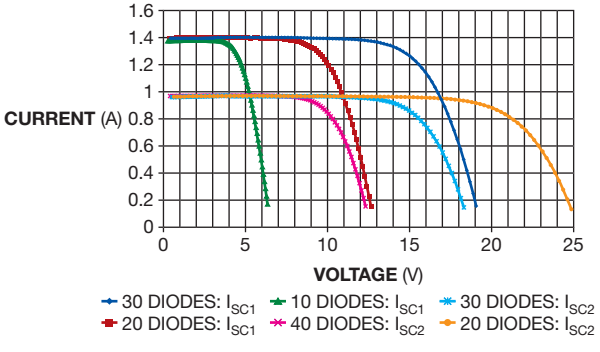

Figure 2 Modifying the current-source output value also modifies the short-circuit current.

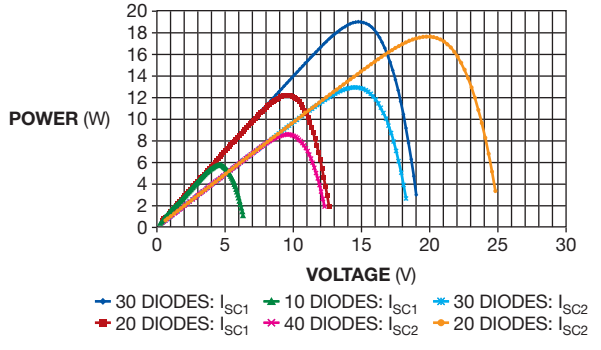

Figure 3 For each combination of diodes, the maximum power-point voltage changes.

Get an edge: smart grid solutions

From metering to grid monitoring and the communications in between

System designers use Maxim's field-proven technologies to reduce risk while cutting design and deployment cycles for smart grid applications. Building on more than two decades of expertise, Maxim's turnkey solutions are easily modified to meet new standards and usage requirements, helping you protect your infrastructure investments.

Our IC solutions give the grid a higher IQ

- Exceed measurement accuracy requirements so your design doesn't just meet today's spec but tomorrow's as well
- Extend communications across transformers to reach more endpoints and reduce deployment costs
- Improve power quality and delivery while protecting assets and lowering O&M costs

Smart Grid Solutions Guide

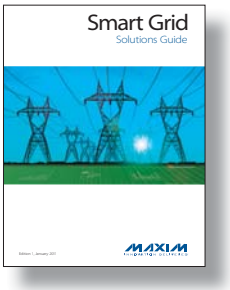

Smart grids start here:
www.maxim-ic.com/grid

MAXIM
INNOVATION DELIVERED®

MAXIM DIRECT
www.maxim-ic.com/shop

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

For a complete list of Maxim's sales offices and franchised distributors, visit **www.maxim-ic.com/sales**

the photovoltaic-module simulator for two current-source values and a different number of diodes. Extracting these curves required the use of an electronic load (**Reference 1**). The short-circuit current of the simulated module changes, modifying the current-source output value, and, when you connect or disconnect diodes, it varies the open-circuit voltage of the module. You can use digital circuits to control the simulator to create photovoltaic patterns. EDN

REFERENCE

1 Ausiàs Garrigós and José M Blanes, "Power MOSFET is core of regulated-dc electronic load," *EDN*, March 17, 2005, pg 92, http://bit.ly/fz8geK.

Switch circuit controls lights

C Castro-Miguens, University of Vigo, Spain,
and JB Castro-Miguens, Cesinel, Madrid, Spain

Cities and towns worldwide are considering and installing LED streetlights to help save electric energy, reduce costs, protect the environment, and improve lighting for their citizens. Despite this trend, the lamps' turn-on/turn-off time control is receiving little attention.

A suitable control can achieve an important energy saving because lights can operate too late, too early, or both, wasting energy or providing insufficient light. Using a twilight switch can significantly reduce energy consumption in all types of lamps (**Figure 1**). It offers a cost-effective, compact, and reliable way of providing lighting time control.

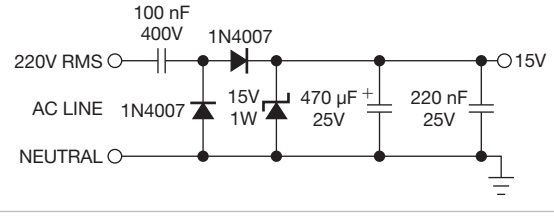

Figure 1 This charge pump feeds the twilight switch from the ac line with high efficiency.

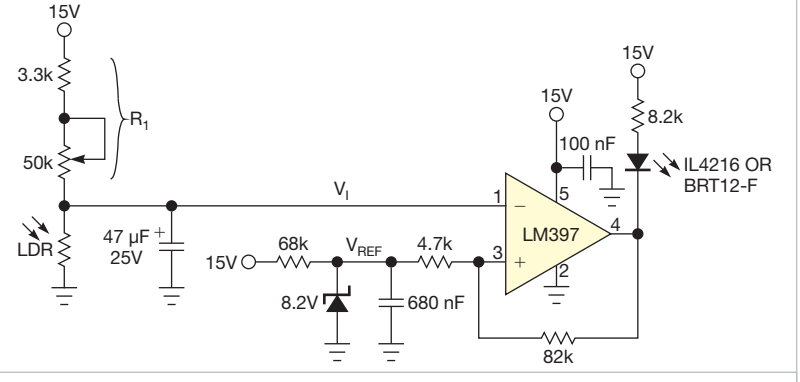

Figure 2 The circuit uses a light-dependent resistor to measure the ambient-light level.

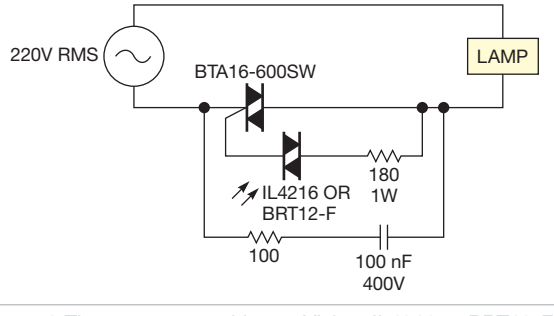

Figure 3 The comparator drives a Vishay IL4216 or BRT12-F optocoupler with a TRIAC output.

The circuit does not use a relay. Therefore, it has no moving parts, and it is not prone to contact oxidation. It uses a TRIAC (triode alternating current) that can switch hundreds of watts.

The circuit requires little power. It uses a charge pump to feed the circuit from the ac line, drawing less than 37 mW for a 220V-rms ac line. It uses just a few low-cost components.

You can adjust the circuit's darkness and illuminance level that switches the light on and off using only onboard potentiometer R_1. The circuit automatically turns on the lamps at nightfall and turns them off at daybreak. You can use it with incandescent lights, fluorescent lights, or LEDs.

The circuit uses an LDR (light-dependent resistor) to measure the ambient-light level (**Figure 2**). Be sure that the LDR you use has a spectral response similar to that of the human eye to achieve good performance. It uses a hysteresis comparator because a basic comparator configuration oscillates or produces a noisy output when the illumination level is close to the edge between natural light and darkness. Hysteresis creates two switching thresholds in the circuit: The upper threshold voltage is 8.47V for the rising input-voltage change from natural light to darkness, and the lower threshold voltage is 7.75V for the falling input-voltage change from darkness to natural light. The relationship between the 82-kΩ and the 4.7-kΩ resistors controls the 0.72V hysteresis. This value is adequate to avoid the false triggering that light noise can cause.

THE CIRCUIT DOES NOT USE A RELAY, AND SO IT HAS NO MOVING PARTS.

When the ambient light falls below the level that R_1 sets, the input volt-

SPONSORED BY AVNET Embedded

age, V_1, rises above the upper threshold voltage and the output of the comparator decreases, switching on the TRIAC. When the ambient light rises above the level that R_1 sets, the input voltage decreases below the lower threshold voltage and the output of the comparator increases, switching off the TRIAC.

You must provide a mechanical isolation between the LDR and the lamp light to prevent the formation of a feedback path to the LDR. Otherwise, the lamp light will cause an oscillation at the comparator's output and then in the lamp's state. The BTA16-600SW, which is available from many sources, is suitable for switch lamps operating at more than 2000W.

The comparator drives a Vishay (www.vishay.com) IL4216 or BRT12-F optocoupler with a TRIAC output (**Figure 3**). The optocoupler, in turn, drives the BTA16-600SW TRIAC that controls the lamp. EDN

Isolated PWM suits low frequencies

Tim Regan, Linear Technology, San Jose, CA

Many industrial- and medical-system circuits require isolation from the mains-ac power. You can often send a signal across the isolation barrier using a small transformer; transformers do not pass low-frequency signals well. The circuit in this Design Idea converts a low-frequency PWM (pulse-width-modulated) signal to a higher frequency, which you pass across the transformer, and retains the duty cycle. Once the frequency is on the other side, you can then convert the PWM signal back to an analog voltage.

The circuit converts a 1-kHz PWM signal to a 100-kHz signal with the same duty cycle (**Figure 1**). This 100-kHz signal easily couples across an isolation transformer. You then filter it to provide a dc control voltage on the isolated side. IC_2, an LTC6992-2, is a voltage-controlled PWM IC. A voltage ranging from 0 to 1V on the MOD input pin varies the duty cycle of the output from 5 to 95%. The duty cycle does not reach 0 or 100%, which would not pass through the transformer. Resistor R_{SET} fixes an internal master-oscillator frequency of 100 kHz. The voltage on the DIV pin sets a divider ratio. An internal 4-bit ADC translates the analog voltage to a digital-divider value. With the input shored to ground, the divider ratio is one, and the circuit outputs the oscillator frequency.

R_2 and C_2 make amplifier IC_1 a voltage integrator. It servo-controls the voltage at the MOD pin of IC_2. R_1 and C_1 filter the input signal to an average dc value. The integrator compares this value with the average dc value of the IC_2 output that you filter with R_3 and C_3. This step forces the 100-kHz output signal to the same duty cycle as the 1-kHz input signal. The time constants of these filters should be much longer than the clock period to minimize duty-cycle jitter. You use a 500-msec time-constant network for the 1-kHz PWM signal and a 1-msec time-constant network with the 100-kHz PWM signal. The amplitude of each signal must be the same for an accurate match of the two duty cycles. For this reason, IC_2's power-supply voltage is the same supply the PWM controller uses. Any supply-voltage variation affects each signal in the same way, providing insensitivity to power-supply variation.

IC_2 has 20 mA of output current to drive the primary of the isolation transformer, and comparator IC_3 squares up the 100-kHz PWM signal on the isolated side of T_1. You can use this output directly as a digital-control signal if necessary. In this circuit, you filter the signal with a 1-msec RC lowpass network and then buffer the signal with voltage-follower amplifier IC_4, yielding a dc analog-control voltage.

The circuit accurately replicates a stepped increase in the input PWM duty cycle (**Figure 2**). The circuit operates from a 5V supply. The average voltage changes from 1 to 4V when you change the input duty cycle from 20 to 80%. The slow change is due to the 500-msec time-constant filter, an acceptable scenario in cases in which a gradual change in the isolated control signal is acceptable. You can circumvent the slow response at the expense of some overshoot if you need a faster response. For this task, you use an anticipator circuit, as a previous Design

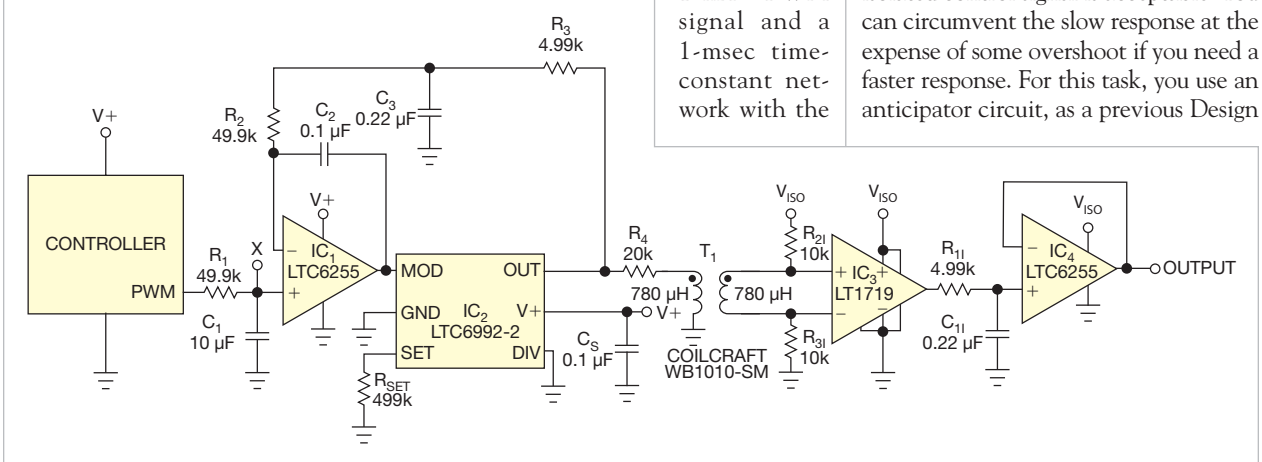

Figure 1 This circuit creates a 100-kHz PWM signal from a 1-kHz PWM signal to send it across an isolation transformer. The output is an integrated dc control voltage.

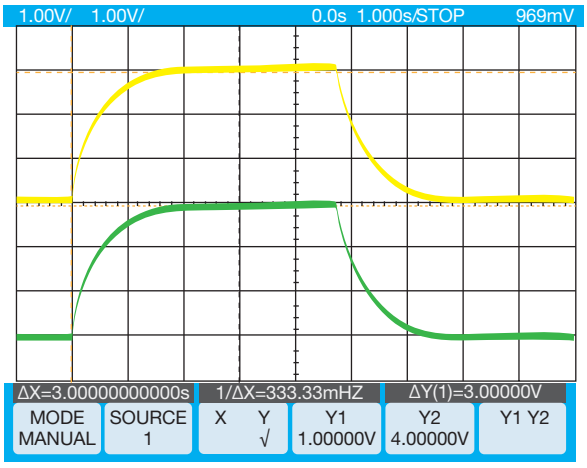

Figure 2 The circuit gives a gradual exponential response to a 20 to 80% step change in duty cycle. A 500-msec time constant filters the input PWM (yellow). The filtered, 100-kHz isolated signal closely matches the input (green).

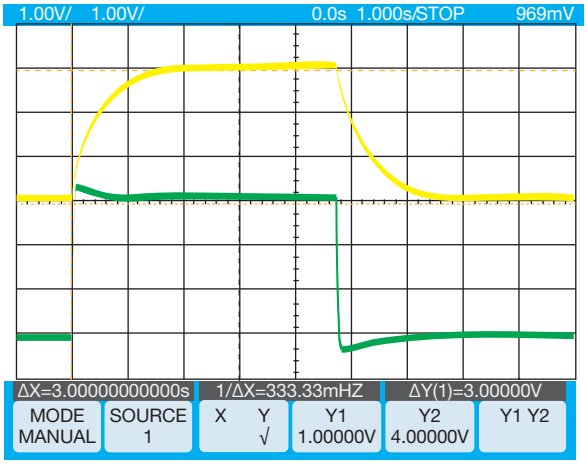

Figure 3 An anticipator circuit speeds the output response (green).

Idea describes (**Reference 1**). Adding the anticipator circuit at the X node of **Figure 1** results in a faster response to a final value (**Figure 3**).

You add another voltage-controlled PWM IC to re-create the 1-kHz PWM signal on the isolated side (**Figure 4**). You use amplifier-integrator circuit IC_6 to servo-control the duty cycle of IC_7. Resistor R_{SETI} programs IC_7 for 1-kHz, 5- to 95%-duty-cycle operation. The circuit forces the 1-kHz output duty cycle to equal the 100-kHz input-signal duty cycle. Again, the supply voltage for comparator IC_5 and the PWM device must be the same. If you want minimum duty-cycle ripple, set filter R_{2I} and C_{2I} to have a 500-msec time constant. Unfortunately, this approach would make the response time of the reconstructed output slow. You cannot use the anticipator circuit when you re-create the slow PWM signal because the slow signal is now the dependent variable in the circuit, and fast jumps in the feedback voltage would result in the loop's continuously hunting and never settling to a final value. Use instead a 10-msec time-constant filter on the 1-kHz output to obtain a reasonable response time and then minimize duty-cycle ripple with an additional lowpass network comprising R_{4I} and C_{4I}.

Note that the feedback signal for integrator IC_1 in **Figure 1** is on the negative pin and that the feedback signal to re-creation integrator IC_6 is on the positive pin. If you connect the fast-responding input signal to the positive input of IC_6, it would cause a large overshoot in the duty cycle of the output signal and a long recovery time. You compensate for the polarity reversal by biasing the DIV pin of IC_2.

Setting the DIV pin above $V_{SUPPLY}/2$ programs the output-control polarity to change from 95 to 5% duty cycle with an increasing voltage applied to the MOD pin. An increase in the 100-kHz signal's duty cycle now ramps down the MOD pin and increases the 1-kHz output signal's duty cycle to match it. EDN

REFERENCE

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

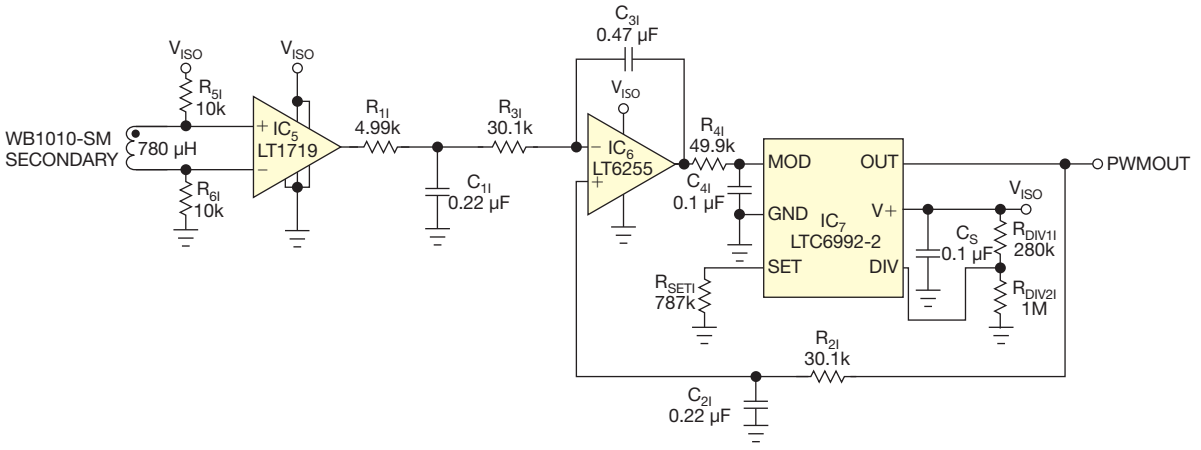

Figure 4 This circuit re-creates the 1-kHz PWM control signal on the isolated side. It cannot use the anticipator circuit, but R_{4I} and C_{4I} provide extra lowpass filtering, helping to reduce output jitter.

Make your electronics supply chain green—or else

Regulations and the need to save energy are compelling supply chains in the worldwide electronics industry to take a turn toward sustainable approaches. From the moment the design process begins, design engineers need to specify "green" parts, design for green manufacturing, design light packaging, and ensure that their products are recyclable or returnable at the end of life.

The demand to make the electronics supply chain green is coming from multiple directions. "There's pressure coming from the regulators, corporate, customers, investors, and competitors," says Pam Gordon (**photo**, left), president of Technology Forecasters Inc (www.techforecasters.com). "Sometimes executives do care about the environment, but there are enough external pressures that the executive does need to be an environmentalist."

Just a few years ago, the EU (European Union) imposed a green focus on the electronics supply chain with its ROHS (restriction-of-hazardous-substances) directive, which called for the removal of six materials deemed hazardous from electronics design. The EU's REACH (registration, evaluation, authorization, and restriction of chemicals) ordered companies to detail a range of chemicals in their products. In recent years, products' end of life has become an issue, whether it involves design for recycling or the patchwork of take-back legislation emerging from individual states in the United States.

Ultimately, the power to force change comes from regulations that countries and states have passed. "The order of priority in the green supply chain is legislation, legislation, and corporate image," says Kenneth Stanvick, senior vice president at DCA (Design Chain Associates, www.designchainassociates.com). Yet some companies shrug off the regulations. "Many companies are having a hard time recognizing that—from a legal perspective—they're breaking the law if they don't comply," he adds.

The electronic-waste laws in the United States have brought every corner of the electronics industry into environmental compliance. "The green supply chain is getting some traction because of the e-waste legislation; 24 states have legislation," says Patrick Penfield, director of supply-chain executive programs at Whitman School of Management at Syracuse University (http://whitman.syr.edu). "Those 24 represent 65% of the population in the United States. All of the laws, except California's, use the producer-responsibility approach," meaning that the maker must take back the product at the end of its life.

Many small to midsized companies are turning to distributors to help facilitate take-back programs. "Distributors such as Avnet and Arrow are helping customers with end-of-life and reverse-logistics issues," says Gerry Fay, senior vice president of global-supply-chain and strategic accounts at Avnet Inc (www.avnet.com).

The silver lining in the pressures to be environmentally sound is cost savings that come from reduced energy consumption. A small revolution is occurring in reduced energy use in manufacturing and in the logistics of getting parts to one location and sending them to another. "Green behavior can be very beneficial from a cost standpoint," says Colin Campbell (**photo**, center), vice president of supply chain at Newark (www.newark.com). "It takes planning and coordination with suppliers, but manufacturers can consolidate freight and ship orders every week instead of every day."

Energy has also become one more consideration in design. Both corporate and consumer customers expect companies to offer low-power-consumption products. "The foremost issue in energy is lower-power design. It's not a new concept, but it has become a huge theme," says Andrew Femrite (**photo**, right), engineering-solutions manager at Arrow Electronics Inc (www.arrow.com). "People are looking for low power in all their applications. People are looking for energy harvesting, where the product no longer runs on batteries."

The companies that fare best in the transition to a greener electronics industry are those that move quickly and aggressively. "Being reactive to the host of increasing regulation from around the world is where the tax comes in," says Technology Forecasters' Gordon. "Companies that get ahead of it can save money and reduce costs. The companies that are the boldest, the most consistent, and the most serious with executive commitment will be the best off."

—by Rob Spiegel

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

productroundup

DISCRETE SEMICONDUCTORS

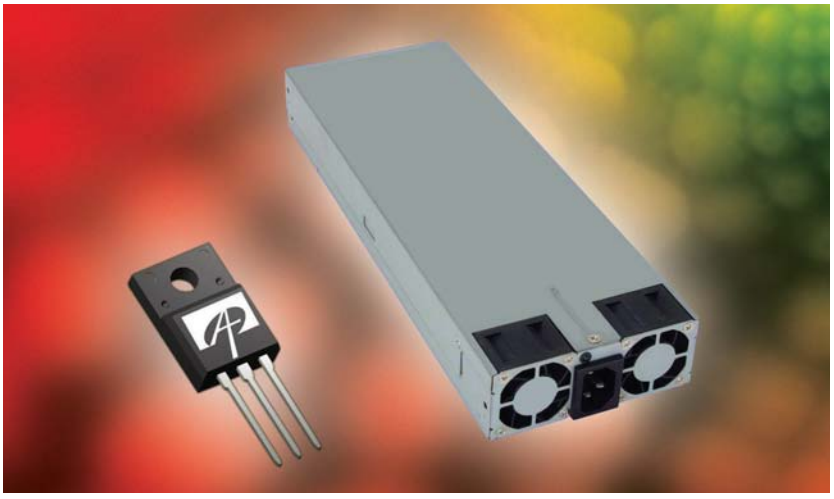

MOSFET targets high-voltage primary-side applications

Targeting use in high-voltage primary-side applications, the 600V AOTF27S60 MOSFET features a maximum drain-to-source resistance of 0.16Ω. Drain-to-source resistance times the gate charge is 3.64Ω-nC. The device sells for $3.60 (1000).

Alpha and Omega Semiconductor, www.aosmd.com

30V MOSFET comes in Power-SO-8 package

The 30V, PSMN1R0-30YLC MOSFET features an on-resistance of 1.4 mΩ at 4.5V and comes in an LFPak Power-SO-8 package. Applications include high-performance

dc/dc conversion, such as in synchronous-buck regulators, synchronous rectifiers in isolated power supplies, and power OR-ing. The device sells for $1.07 (10,000).

NXP Semiconductors, www.nxp.com

MOSFET family aims at dc/dc-switching applications

The IRF6811 and IRF6894 DirectFET power MOSFETs target use in 12V-input, synchronous-buck applications, including next-generation servers, desktops, and notebooks. The IRF6811 features a typical drain-to-source on-resistance of 2.8 and 4.1 mΩ at 10 and 4.5V, respectively; a gate-to-source voltage of ±16V; and a typical gate charge of 11 nC. The IRF6894 features a typical drain-to-source on-resistance of 0.9 and 1.3 mΩ at 10 and 4.5V, respectively; a gate-to-source voltage of ±16V; and a typical gate charge of 29 nC. Prices begin at 75 cents and $1.75 (10,000), respectively.

International Rectifier, www.irf.com

Medium-voltage MOSFETs come in CanPak package

The 60 to 150V OptiMOS MOSFETs come in a CanPak package. Thermal resistance on the top side of the package is 1.5K/W versus 55K/W for a traditional DPak. The MOSFETs' on-resistance ranges from 2.8 to 28 mΩ. Prices begin at 40 cents (10,000).

Infineon Technologies, www.infineon.com

TMBS rectifiers have 10 to 60A current ratings

This family of 45V TMBS (trench-MOS-barrier-Schottky) technology rectifiers is available in four

power-package types and features current ratings of 10 to 60A. With typical forward-voltage drops as low as 0.28V at 5A, the rectifiers reduce power losses and improve efficiency in high-frequency dc/dc converters, switch-mode power supplies, freewheeling and OR-ing diodes, and reverse-battery protection in desktop PCs, servers, and LCD TVs. The devices have a maximum junction temperature of 150°C. The 10, 20, 30, 40, and 60A devices sell for 40 cents, 60 cents, 80 cents, $1, and $1.20, respectively.

Vishay Intertechnology, www.vishay.com

AMPLIFIERS, OSCILLATORS, AND MIXERS

Differential amplifier drives 100-MHz signals

The LTC6409 fully differential amplifier has 1.1 nV/√Hz noise and –88-dBc distortion. Using a SiGe BiCMOS process, the device features a 10-GHz gain-bandwidth product, 100-dB SFDR to 40 MHz, and 1% settling

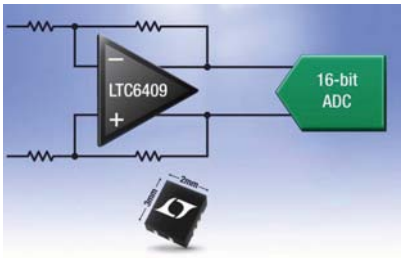

time of 1.9 nsec. The device drives high-speed ADCs and has an input common-mode range of 0 to 3.5V and an output common-mode range of 0.5 to 3.5V when using a 5V supply. The unity-gain-stable amplifier has an output current as high as 95 mA and operates from a 3 or a 5V supply. Maximum power consumption is 56 mA, and shutdown mode reduces this current to 500 µA. Turn-on time is typically 160 nsec. The device targets use in pulsed-signal applications, such as radar-signal-processing, imaging, high-speed-test-and-measurement, and communications applications. It comes in a 2×3-mm QFN package and operates over the 0 to 70, –40 to +85, and –40 to +125°C temperature ranges. Prices start at $4.50 (1000).

Linear Technology,
www.linear.com/6409

36V precision op amps suit use in industrial applications

The OPAx209 precision-op-amp series targets use in fast, high-precision data-acquisition applications, such as automated test equipment, medical instrumentation, and professional audio preamplifiers. The devices feature 2.2-nV/√Hz noise density; 0.1- to 10-Hz noise of 130 nV p-p; 18-MHz gain bandwidth; and settling time to 16-bit accuracy, or 0.0015%. Offset voltage is 150 µV, and the devices operate over a single- or dual-supply range of 4.5 to 36V. The single-input OPA209 comes in SO-8, MSOP-8, and SOT23-5 packages and sells for 95 cents (1000). The dual-input OPA2209 comes in SO-8 and MSOP-8 packages and sells for $1.65, and the quad-input OPA4209 comes in a TSSOP-14 package and sells for $2.90.

Texas Instruments, www.ti.com

Stereo line driver targets digital TVs, DVD recorders

The 2V-rms WM3100 stereo line driver operates from a 3.3V supply to drive digital TVs, DVD recorders, Blu-ray players, gaming consoles, and set-top boxes. It delivers a 108-dB SNR and –95-dB THD+N and has pins to control mute and output modes. It supports differential and single-ended inputs in both inverting and noninverting configurations. In a 14-pin SOIC package, the WM3100 sells for 44 cents (1000). Evaluation boards are also available.

Wolfson Microelectronics,
www.wolfsonmicro.com

Current-sense amps support I²C output

The MAX9611 and MAX9612 current-sense amplifiers integrate a 12-bit ADC and a gain block that you can configure as either an op amp or a comparator. The devices integrate an I²C-controlled, 12-bit, 500-sample/sec ADC. The I²C bus is compatible with 1.8 and 3.3V logic. Input-common-mode voltage range is 0 to 60V, and programmable full-scale sense voltages are 440, 110, and 55 mV. Operating over –40 to +125°C, the devices are available in 10-pin µMAX packages. Prices start at $1.50 (1000).

Maxim Integrated Products,
www.maxim-ic.com

Band together

In the 1980s, I started a job with a company that produced high-end analog stereophonic-audio systems. One of the products was a system comprising a 200W amplifier, a tape deck, a turntable, and a 10-band graphics equalizer. The equalizer split the audio-frequency range of 20 Hz to 20 kHz into 10 bands, allowing users to individually adjust the level of each band, using vertically sliding controls, from –12 to +12 dB. In its stereo version, it had 10 bands for the left channel and 10 bands for the right channel. The position of the sliders represented the frequency response. Most graphics equalizers then used one coil or one inductor per band, but this one used operational amplifiers instead of inductors.

Testing the equalizers involved feeding the output of an audio-frequency generator to the equalizer and reading its output on a center-zero voltmeter calibrated in decibels. For covering one band of the equalizer, the audio-frequency generator had to be set to the correct frequency range and then its dial had to be rotated to the center frequency of the band under test. I adjusted the generator's output to show 0 dB at the output on the voltmeter if the slider of the band was in midposition. Moving the slider up caused the reading on the output needle's voltmeter to show 12V; moving it down resulted in a reading of –12V. Because the stereo version of the equalizer had 20 bands, I had to repeat this procedure 20 times.

I first attacked the time-consuming generation of the correct audio frequency for the band under test without selecting range switches and turning dials. If I could rig up 10 audio-frequency generators and then tune each one to its center frequency, I could then reduce the task of testing to selecting the required audio oscillator for the band under test and then adjusting the band slider to see the response on the output voltmeter.

I selected a Wien-bridge oscillator for this task because these devices are simple, using familiar op amps. Their frequency of interest depends on the values of two resistors and two capacitors. I would need 10 oscillators, each with a different frequency, leaving the question of stability, which I handled by using two diodes in the op amp's feedback path.

Using LM324 op amps, I built 10 identical Wien-bridge oscillators and adjusted their resistance and capacitance to generate the audio frequency corresponding to the center frequency of each of the equalizer bands. Because all 10 oscillators would always be active, I used a CD4066 analog switch to each oscillator's output as a means of selecting the required output. I tested crosstalk by switching in one oscillator and feeding its output to the equalizer. I attenuated the equalizer's output by sliding down the matching frequency band's slider control, sliding up all the other controls, and checking the output on the output voltmeter. The output stayed put at –12 dB, so the crosstalk was better than 24 dB, considering that each band was capable of a total gain of 24 dB.

To sequentially select the analog switches, I used a CD4017 decade counter with 10 outputs. An astable multivibrator comprising a 555 timer generated the clock signal for the counter. I used the CD4017's Enable input as a stop/run control for the clock, and I added a potentiometer to the 555 to vary the speed of the clock signal as a testing-speed control.

I bundled all of these electronics inside a box with the controls and connectors outside. I timed myself to completely test the equalizer. At the fastest, it took 1.5 seconds per band, working out to 30 seconds for the 20 bands, excluding about 10 seconds to change from the left to the right channel. I added 10 LEDs to provide a direct visual indication of which slider to move next. Within 24 hours, there was no backlog, and production jumped from 10 per day to 200. EDN

Kunal Ghosh is a project manager in Hyderabad, India.

DANIEL VASCONCELLOS

Enabling the Most Energy Efficient Home Networks

Reduce Power and BOM Costs with Legendary Reliability and Robustness

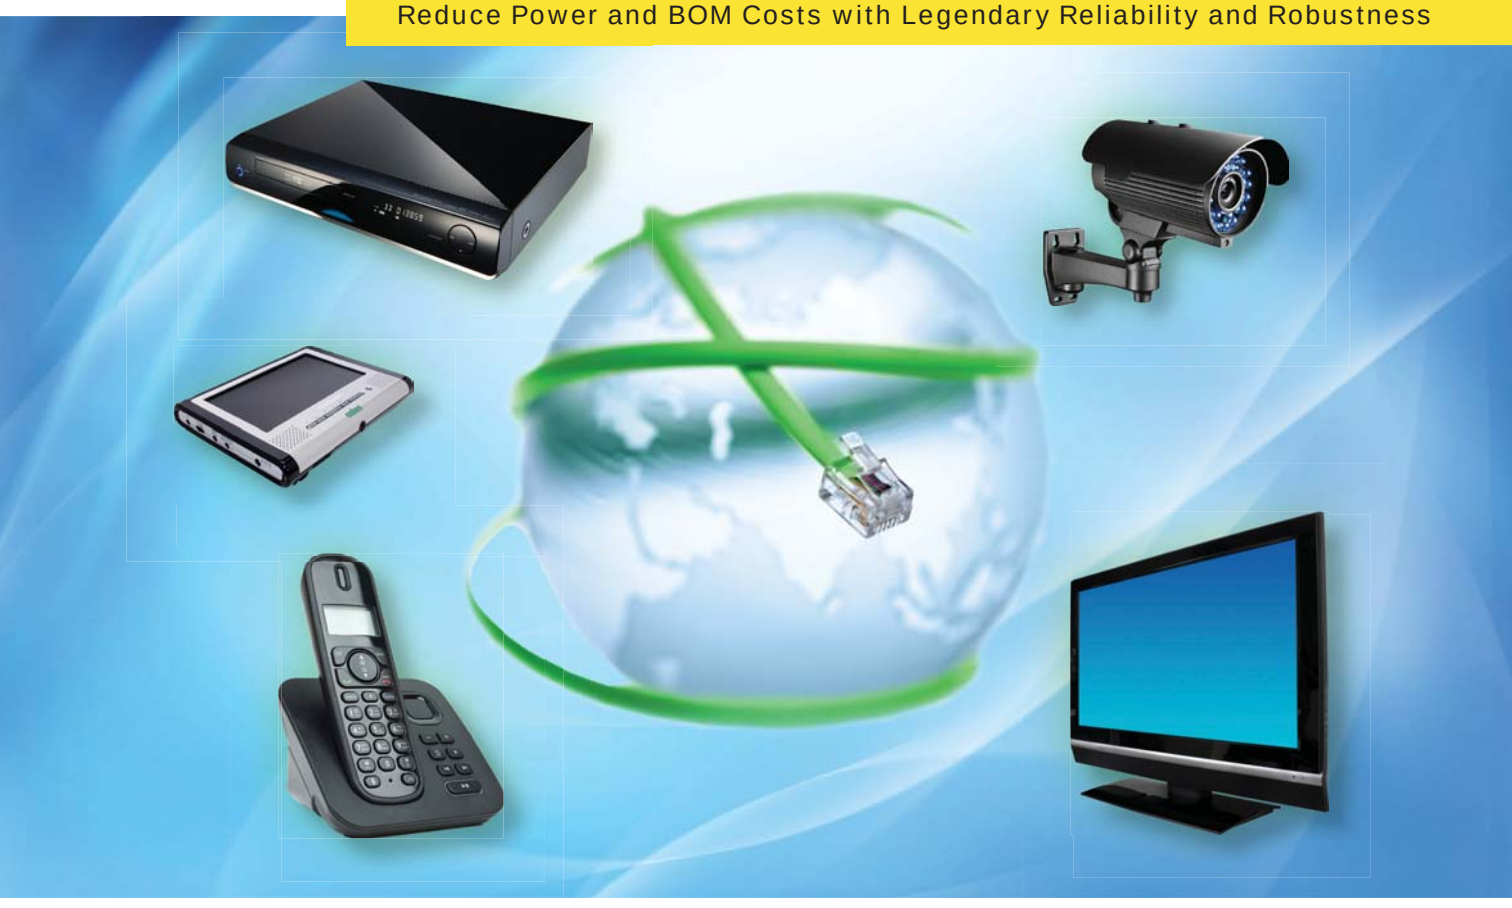

As a world leader in Ethernet for the Digital Home, Micrel is a preferred choice for attaching consumer electronics to home networks today. Featuring energy efficient designs, with advanced physical layer and switch technology, Micrel offers the industry-leading low power line of Ethernet devices:

- PHY Transceivers
- Single-port Controllers
- Switch on a chip ICs

Micrel's family of Ethernet devices has been adopted and proven by major consumer electronic OEMs around the globe, with hundreds of millions of ports deployed.

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

Enables a Wide Range of Applications

- Digital TVs
- Media Players
- Mobile Internet Devices (MID)
- Set Top Boxes
- Residential Gateways

Micrel Advantage

- Green and highly integrated designs, reducing power consumption and board real-estate
- Advanced power management
- Low EMI emissions and high-EMI immunity

MICREL®

Innovation Through Technology®

www.micrel.com

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.

TestEQUITY
800-732-3457
www.testequity.com/agilent

Learn how to connect wirelessly and get more FREE measurement tips at www.testequity.com/Agilent_DAQ

Agilent Technologies