Reed Business Information (RBI™)

EDN®

VOICE OF THE ENGINEER

AUG 20

Issue 16/2009

www.edn.com

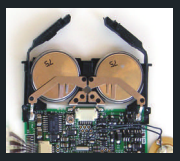

TRANSPORTING HIGH-DEF VIDEO BROADCASTS:

ARE WIRELESS NETWORKS UP TO THE TASK?

Page 24

EDN HANDS-ON PROJECT

TAKE ADVANTAGE OF OPEN-SOURCE HARDWARE

Page 20

ADDRESSING INTERLEAVED MULTI-CHANNEL MEMORY CHALLENGES

Page 32

PICO

PROVEN **Critical Component Integrity**

DC-DC Converters

2V to 10,000 VDC Output

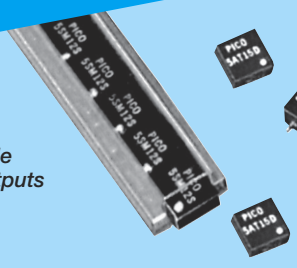

2 to 5000 VDC Outputs
Ultra Miniature
Surface Mount and Thru Hole
Single and Dual Isolated Outputs
Military Upgrades Available
AV/AV/SM/AVR Series

Military Components
-55⁰ to +85⁰C Operating Temp
Military Environmental
Screening Available
3.3 to 500 VDC Outputs
M/MV Series
Isolated - Regulated

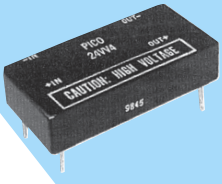

100 to 10,000 VDC Output
Proportional Control Up to 10 Watts
VV Series

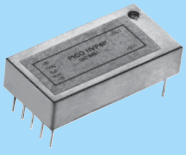

Programmable to
6000 VDC Output
HVP Series

Military Applications
-40⁰ and -55⁰ to +85⁰C
Operating Temperatures
Wide Input Range
3.3 to 350 VDC outputs
in 1/2 Brick and Ful Brick
Military Upgrades Available
LF/LM/FM Series

Wide Input Range
8 to 60 VDC Input
2 to 100 VDC Output
Isolated-Regulated
OR/IR/JR/KR Series
2-20 Watts

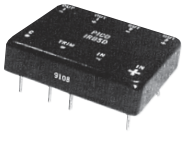

HIGH POWER-Wide Input Range
2 to 350 VDC Output
Isolated to 300 Watts
in 1/2 and Full Brick
LP/HP/XP Series

36 to 170 VDC Inputs
Terminal Strips - Thru Hole
3.3 to 48 VDC Outputs
Single and Dual Output
LV/HV Series

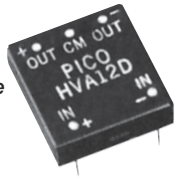

5 to 500 VDC Output
to 50 Watts
Wide Input Range
Isolated
Regulated
QP Series

Also Ac-DC single and 3 Phase.
Power Factor Corrected. 3 Watts to
2000 Watts Models.

PICO
Electronics. Inc.
143 Sparks Ave., Pelham, NY 10803

www.picoelectronics.com

Complete Listing of Entire Product Line
DC-DC Converters • AC-DC Inductors • Transformers and Inductors.

E Mail: info@picoelectronics.com

Call Toll Free:
800-431-1064

Fax:
914-738-8225

ISO 9001:2000
ANAB ACCREDITED
BSI
AS 9100 REV. B

ENTER HERE.

Digi-Key® CORPORATION

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

Over 400,000 products in stock from more than 400 supplier partners.*

An additional 1,000,000+ components can be sourced at digikey.com

The industry's broadest product selection available for immediate delivery

www.digikey.com

1.800.344.4539

*Digi-Key is an authorized distributor for all supplier partners. New products added daily. © 2009 Digi-Key Corporation, 701 Brooks Ave. South, Thief River Falls, MN 56701, USA

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

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

Mentor Graphics®

THE EDA TECHNOLOGY LEADER

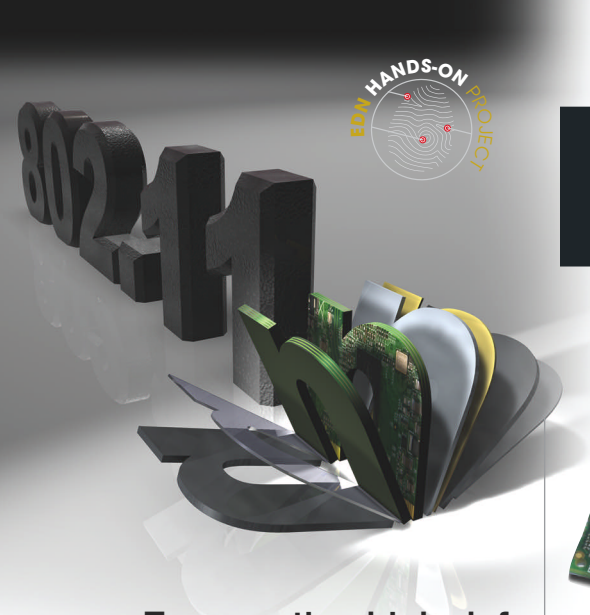

EDN HANDS-ON PROJECT

EDN 8.20.09 contents

Transporting high-def video broadcasts: Are wireless networks up to the task?

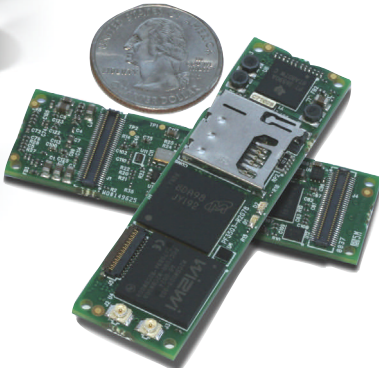

Take advantage of open-source hardware

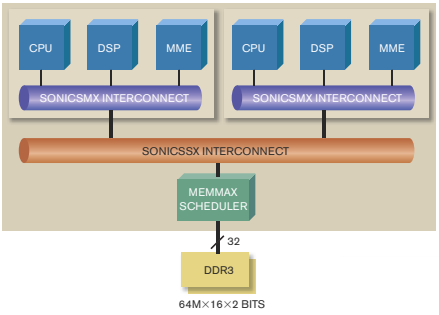

Addressing interleaved multichannel memory challenges

pulse

Dilbert 11

DESIGNIDEAS

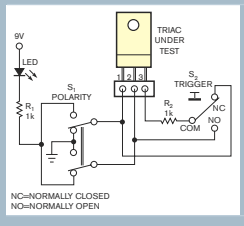

The Newest Embedded Tools

WiFi • USB • BLUETOOTH™ • WLAN • NETWORK • USB • RFID • EMBEDDED CONTROL • USB • WLAN • GPS • BLUETOOTH™ • ZIGBEE® • RFID • USB • WLAN

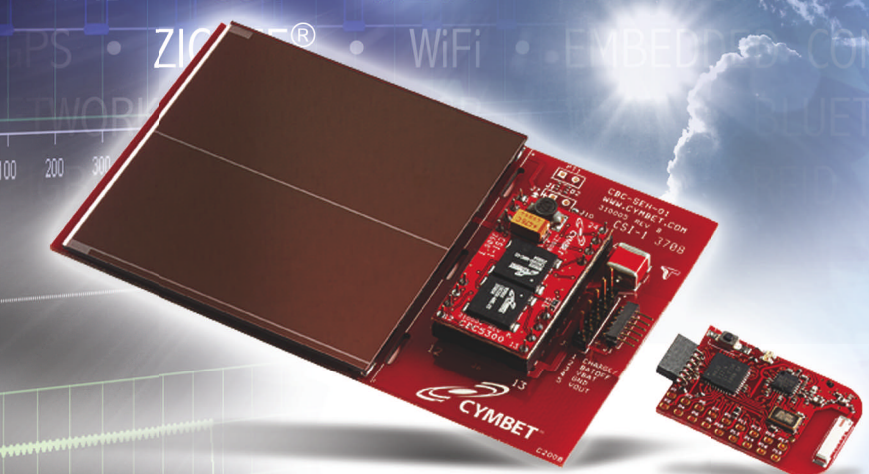

TEXAS INSTRUMENTS
Authorized Distributor

Solar Energy Harvesting Development Tool:
eZ430-RF2500-SEH
www.mouser.com/ti_ez430_rf2500_seh/

www.mouser.com
Over A Million Products Online

New Products from:

MICROCHIP
MPLAB PICkit 3 Debug Express
www.mouser.com/PICkit3debugexpress

LINX TECHNOLOGIES
Evaluation and Development Kits
www.mouser.com/linx/a

SILICON LABS
KEYBOARDEK Evaluation Kit
www.mouser.com/silabs/a

The ONLY New Catalog Every 90 Days

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

MOUSER ELECTRONICS®
a tti company

The Newest Products For Your Newest Designs

(800) 346-6873

contents 8.20.09

DEPARTMENTS & COLUMNS

EDN online contents www.edn.com

ONLINE ONLY

Check out these Web-exclusive articles:

Designing a multicore + coprocessor security system
Network and communications processors are moving to heterogeneous processor configurations for efficiency, cost, and best-of-breed flexibility.
➔www.edn.com/article/CA6662010

Choosing between an ARM7 and a Cortex-M3 processor
The Cortex-M3 offers compelling features over the ARM7, but the ARM7 currently has a larger manufacturers and tools ecosystem.
➔www.edn.com/article/CA6673293

FROM *EDN*'S BLOGS

Top 16 CEO lies
From Leibson's Law, by Steve Leibson
It seems my peanut gallery's gotten all mumbly about Guy Kawasaki's 10 lies that engineers tell. Perhaps they hit too close to home. So here's another of his lies list, 16 lies that CEOs tell and four things you wish they'd mouth instead. Maybe this'll loosen your tongues.
➔www.edn.com/090820toc1

NSF research funds could benefit EDA —or not

From Practical Chip Design, by Ron Wilson
Jeannette Wing, assistant director in charge of the National Science Foundation CISE (Directorate for Computer and Information Science and Engineering), spoke to a packed luncheon audience on the subject of the NSF and EDA.
➔www.edn.com/090820toc2

EDN® (ISSN#0012-7515), (GST#123397457) is published biweekly, 24 times per year, by Reed Business Information, 8878 Barrons Blvd, Highlands Ranch, CO 80129-2345. Reed Business Information, a division of Reed Elsevier Inc, is located at 360 Park Avenue South, New York, NY 10010. Tad Smith, Chief Executive Officer; Mark Finkelstein, President, Boston Division. Periodicals postage paid at Littleton, CO 80126 and additional mailing offices. Circulation records are maintained at Reed Business Information, 8878 S Barrons Blvd, Highlands Ranch, CO 80129-2345. Telephone (303) 470-4445. POSTMASTER: Send address changes to EDN®, PO Box 7500, Highlands Ranch, CO 80163-7500. EDN® copyright 2009 by Reed Elsevier Inc. Rates for nonqualified subscriptions, including all issues: US, $179.99 one year; Canada, $229.99 one year (includes 7% GST, GST#123397457); Mexico, $229.99 one year; air expedited, $399.99 one year. Except for special issues where price changes are indicated, single copies are available for $10 US and $15 foreign. Publications Agreement No. 40685520. Return undeliverable Canadian addresses to: RCS International, Box 697 STN A, Windsor Ontario N9A 6N4. E-mail: Subsmail@ReedBusiness.com. Please address all subscription mail to EDN®, 8878 S Barrons Blvd, Highlands Ranch, CO 80129-2345. EDN® is a registered trademark of Reed Elsevier Properties Inc, used under license. A Reed Business Information Publication/Volume 54, Number 16 (Printed in USA).

BY RON WILSON, EXECUTIVE EDITOR

SOI Industry Consortium stalks the "green thing"

In these days, only government spending seems to drive the economy, and everyone is lining up to be under one of those government spigots. This situation is tricky for the semiconductor industry because it's hard to make a case for a fabless semi company as being essential to bailing out Morgan Stanley, rebuilding our highway infrastructure, or checking the spread of swine flu. However, there is one area—energy saving—in which chips and foundries can claim some home turf.

Accordingly, just about half of new marketing programs have the word "green"—frequently capitalized—somewhere in the first paragraph of their promotional materials. Often, this approach is little more than spurious: an amplified echo of last year's key phrase, "low power." In some cases, though, the appeal to "greenness" makes sense, even without adding chlorophyll to the package epoxy.

One such situation is a new initiative—"Simply Greener"—by the SOI (silicon-on-insulator) Industry Consortium (www.soiconsortium.org). Part of the point is to hitch SOI to the green bandwagon. But there's content in there, too: One of the significant advantages of SOI is its ability to deliver a better speed-power product on a given project than a similar-geometry bulk-CMOS design. Some of the more prominent press coverage of SOI—AMD's travails and the heat problems with industry-leading game consoles, for example—may have obscured this fact. Nonetheless, it is true.

The SOI folks want to make clear the point that you can use that speed-power-product advantage to save significant power at the same speed. To underline that fact, a recent presentation gives examples of benchmark tests from ARM and IBM, which show side-by-side designs of blocks in SOI and in bulk CMOS. In ARM's case, a 45-nm-datapath design, the SOI version achieved an almost-threefold reduction in leakage and an approximately 20% reduction in dynamic power. IBM's example was more apples-to-oranges: a full-chip migration from 65-nm bulk CMOS to 45-nm SOI, resulting in an approximately one-third reduction in power and a 50% speedup. In these instances, the choice of SOI instead of the bulk process appears to be making more difference than the use of aggressive power management.

The consortium wants you to know that SOI brings built-in power savings, and it is mainstream.

The mechanism for the efficiency gain seems to be simple—probably simpler than it actually is. Because SOI builds its transistors directly over a buried insulating layer, the parasitic capacitances from the source, drain, and channel are much less than in a bulk wafer. By reducing these capacitances, a SOI transistor can operate with lower drive current and, hence, can be smaller; have a higher threshold voltage; or offer both features. Thus, both leakage and dynamic currents can be smaller at the same performance level.

A second point the SOI Consortium wants to emphasize is that SOI is available as an off-the-shelf foundry process, not just as a full-custom technology. "There's a wide range of regular users now," says Horacio Mendez, executive director of the consortium. "Almost everything IBM is building at 45 nm is in SOI, as are all of Freescale's latest networking chips. Casio is using the technology at extremely low power levels for watches, and some vendors are applying the technology in automotive applications." Foundry service is available from IBM and Chartered Semiconductor, among others.

SOI not only is a viable option for ordinary design teams but also has a road map, the organization claims. Processes are available in 65 and 45 nm, and both 32- and 22-nm processes are on the drawing boards. The main ideas the consortium wants you to know are that SOI brings built-in power savings, and it is mainstream. Those points deserve some discussion, even from teams that are tooling up to work on bulk CMOS. EDN

Contact me at ronald.wilson@reedbusiness.com.

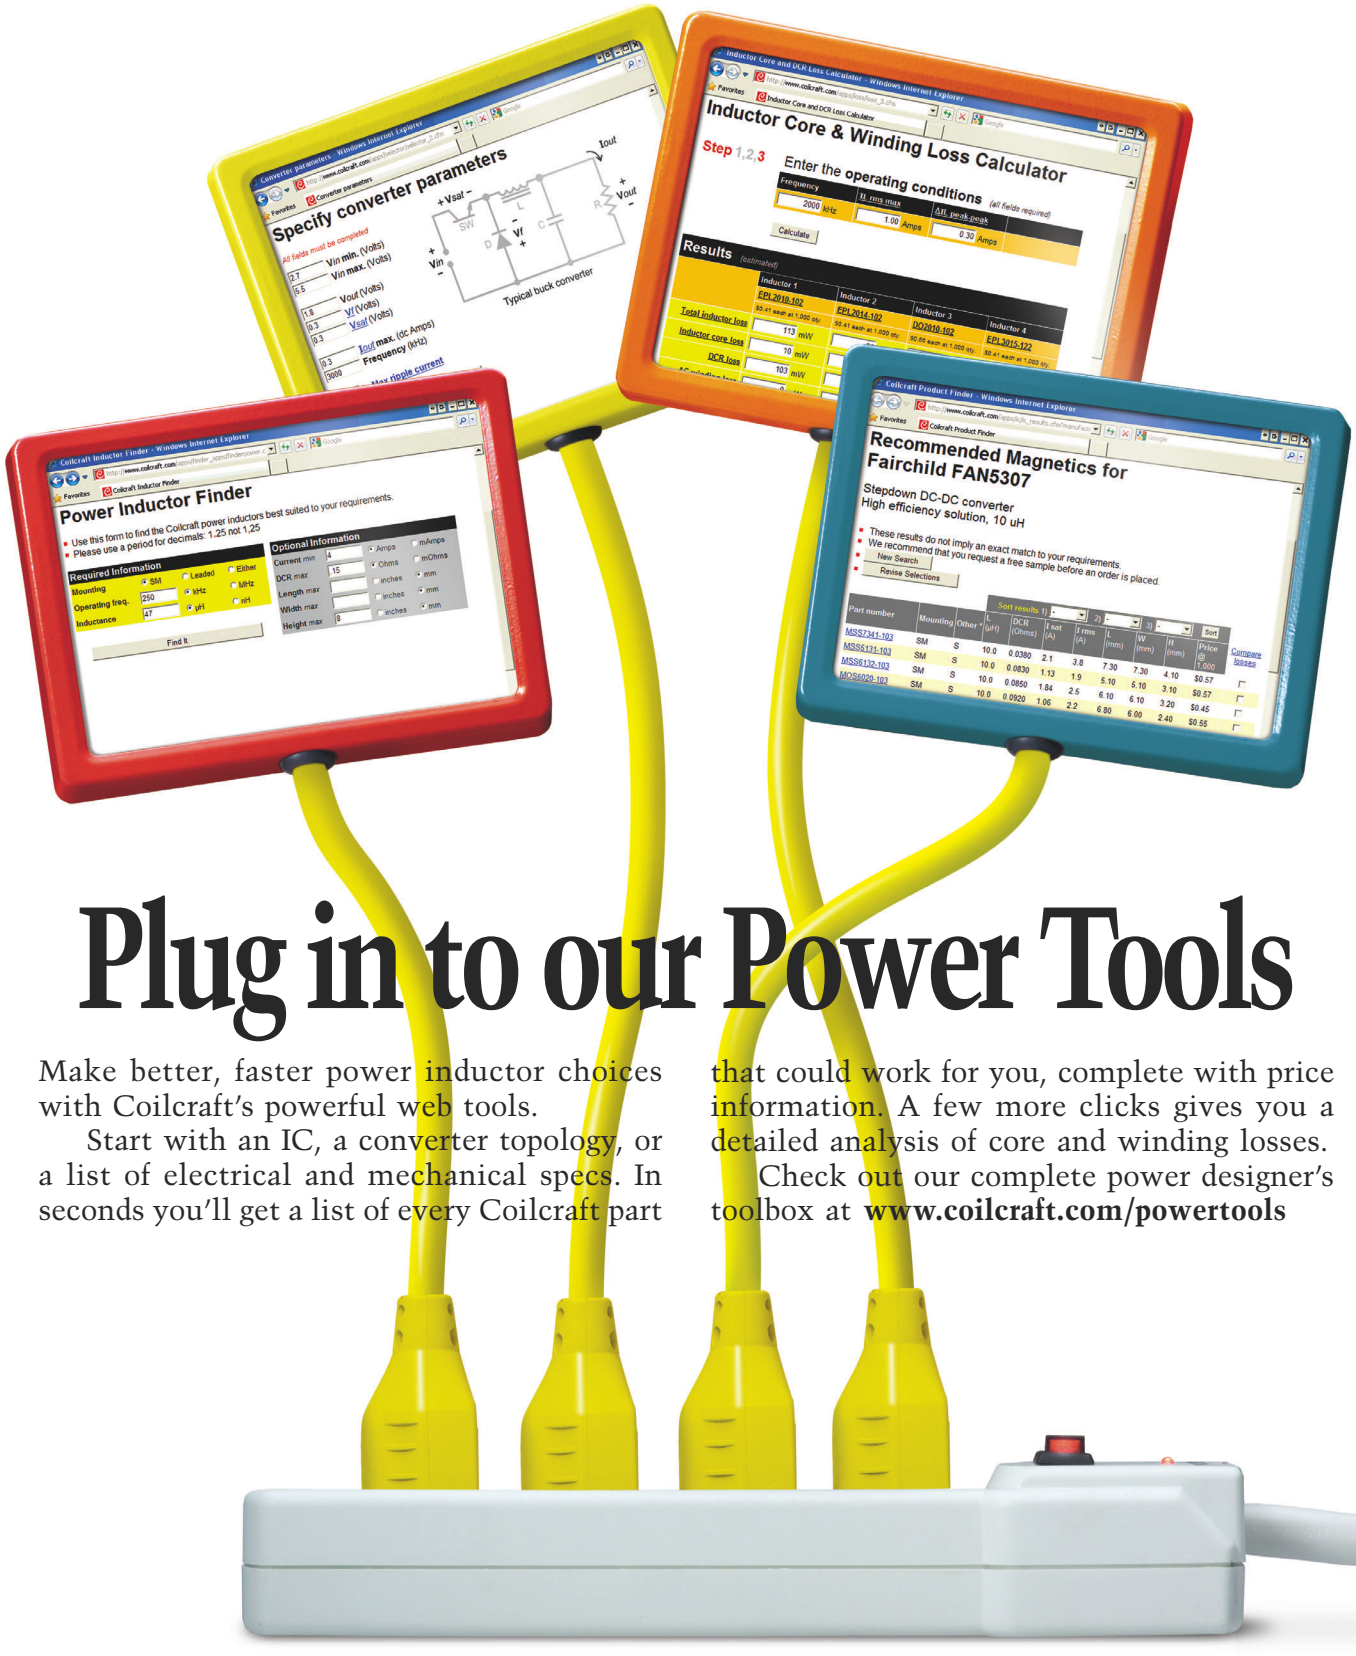

Plug in to our Power Tools

Make better, faster power inductor choices with Coilcraft's powerful web tools.

Start with an IC, a converter topology, or a list of electrical and mechanical specs. In seconds you'll get a list of every Coilcraft part that could work for you, complete with price information. A few more clicks gives you a detailed analysis of core and winding losses.

Check out our complete power designer's toolbox at **www.coilcraft.com/powertools**

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

Coilcraft®

www.coilcraft.com 800/322-2645

“I need to see the most signal detail.”

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

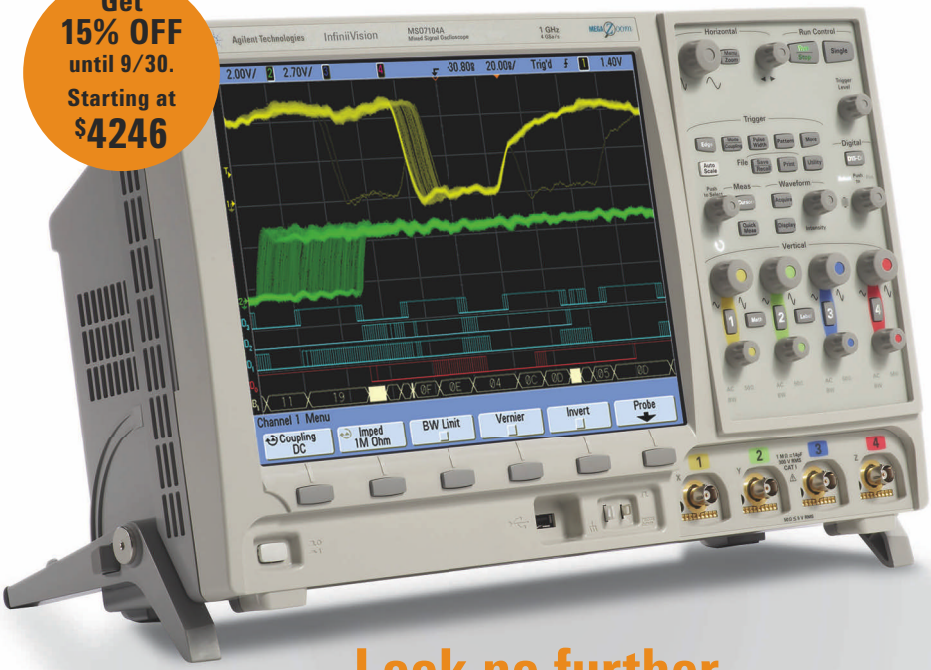

Look no further.

InfiniiVision scopes are engineered for the best signal visibility.

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

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

100 MHz to 1 GHz MSO and DSO models

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

Prices in USD and subject to change without notice.

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

Agilent Technologies

EDN

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

CONSULTING EDITOR
Jim Williams, Staff Scientist,
Linear Technology

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

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

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

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

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

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

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

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

Reed Electronics Group

pulse

INNOVATIONS & INNOVATORS

Trigger, decode, and graphing packages simplify scope-based audio-bus debugging

The serial AudioBus triggering, decoding, and graphing packages for LeCroy Corp's WaveRunner Xi and WaveSurfer Xs oscilloscopes provide a complete set of tools for precise analysis and debugging of the I²S (inter-IC-sound) bus, a newly developed serial interface for digital-audio applications. The WaveRunner AudioBus decode capability includes a unique View Audio feature that converts the digitally encoded serial-data audio signal into an analog-waveform display. This capability provides an intuitive way to understand circuit problems that can cause clipping, glitches, and other anomalies in the audio circuit, as well as a way to see the effects of digitizing the audio signal before DSP has modified it. For conventional left/right-audio or home-cinema applications, which the TDM (time-division-multiplexed) audio buses enable, View Audio can simultaneously operate on as many as four audio channels.

AudioBus uses color-coded overlays on various sections of the decoded protocol data to produce an easily understood visual display. This exclusive feature is intuitive not only to experienced audio engineers but also to users who are new to I²S's LJ (left-justified), RJ (right-justified), and TDM audio-bus data formats. In addition, decoded information condenses or expands depending on the timebase/zoom-ratio setting, simplifying both routine verification and complex troubleshooting. You can choose to decode the data into hexadecimal, binary, decimal, or decibel formats.

You can also configure the AudioBus trigger to work with I²S's LJ or RJ variants, and you can apply conditional triggering to either left- or right-channel data. Unique triggers, such as mute, clip, and glitch, help to isolate rare problems that you cannot easily detect by viewing only the decoded data. In addition, AudioBus turns the oscilloscope into a protocol analyzer with a customizable tabular display of protocol information that you can export as a Microsoft (www.microsoft.com) Excel file.

A WaveRunner Xi scope with the AudioBus TDG package provides a complete set of tools for analyzing and debugging the new serial I²S audio bus. In addition to enabling the scope to display the bus data as a color-coded analog waveform, the package transforms the scope into a protocol analyzer that can export the data as customizable spreadsheet files.

Access to the View Audio Waveform capability requires a WaveRunner scope and the $1995 AudioBus TDG (trigger/decode/graph) package. The $1345 AudioBus TD (trigger/decode) package works with both WaveRunner and WaveSurfer scopes.

—by Dan Strassberg

▷**LeCroy Corp**, www.lecroy.com.

FEEDBACK LOOP

"What would those knotheads know about innovation? Some of my best work has begun with those little, fleeting moments when something you weren't even thinking about before suddenly gels in the mind and shapes up into something neither I nor anybody else anticipated or were even looking for."

—Consulting engineer and designer Thomas Fay, in *EDN*'s Feedback Loop, at www.edn.com/article/CA6666234. Add your comments.

Intel introduces speedy, inexpensive, 34-nm solid-state drives

Although Intel and Micron's (www.micron.com) IM Flash Technologies partnership last November announced mass production of 34-nm lithography-based, 32-Gbit MLC (multilevel-cell) NAND-flash memories, Intel fabricated its first generation of MLC-derived solid-state drives, along with their SLC (single-level-cell) siblings, on 50-nm ICs. In July, Intel unveiled its second-generation 34-nm, MLC solid-state-drive products, although, at first glance, you might scratch your head at what the fuss is all about. Intel still calls the drives the X25-M family, and they come in the same 80- and 160-Gbyte capacities as their 50-nm predecessors, although they showcase a revised silver-case paint scheme.

Real-life performance improvements due to the 34-nm-fabrication reduction, especially considering recent firmware updates of the 50-nm X25-M line, are also unclear. Intel says that the X25-M offers 25% lower latency and notably faster random-write operations than its 50-nm predecessor. The reduced latency offers quicker access to data, operating at 65-µsec read latency compared with approximately 4000 µsec for a hard-disk drive. Random-write performance has increased twofold for the 80-Gbyte version and 2.5-fold for the 160-Gbyte version. According to the company, the X25-M delivers as much as 6600 4-kbyte write-I/O operations/sec—8600 for the 160-Gbyte version—and as many as 35,000 read-I/O operations/sec.

Intel also admits that it does not expect substantial gains on application-based benchmarks versus its first-generation MLC solid-state drives. Conversely, targeted synthetic benchmarks and tests will reveal more noticeable differences. Lingering delays in 6-Gbps SATA (serial-advanced-technology-attachment) system-side support are likely a notable limiting factor in fully showcasing solid-state drives' speed strengths. However, Intel is passing along 34-nm-process-lithography cost reductions to its customers. In July, the company was quoting channel prices for the 80- and 160-Gbyte X-25M of $225 and $440 (1000), respectively, a decrease from $595 and $945, respectively, at the 50-nm-based product introduction a year ago. Intel is now shipping MLC solid-state drives in the 2.5-in. hard-disk-drive form factor, with 1.8-in. counterparts becoming available by the end of this quarter. Intel currently doesn't comment on the availability of 34-nm-derived SLC solid-state drives.

—by Brian Dipert

▷**Intel**, www.intel.com.

The X25-M family comes in 80- and 160-Gbyte capacities.

DILBERT By Scott Adams

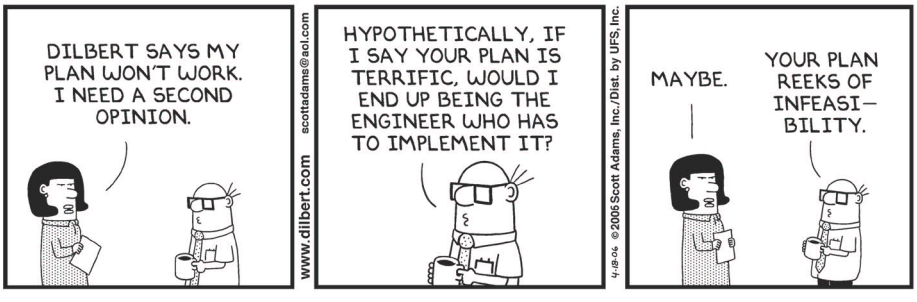

08.20.09

ROTARY-ENCODER IC MEETS ALL AUTO SPECS

The AS5163 magnetic-rotary-encoder IC from austriamicrosystems satisfies the stringent automotive-IC-protection requirements in angle-sensing applications. The device provides overvoltage protection as high as 27V, and reverse-polarity protection withstands −18V reverse polarity at the supply pins. The device also has a short-circuit-monitoring function.

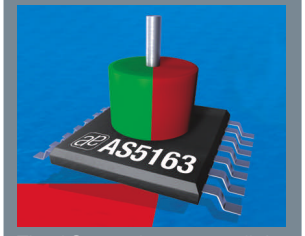

The AS5163 integrates all the functions of an analog or a digital magnetic-sensor rotary encoder and operates at automotive temperatures.

Its designers envisage that applications for the chip will include throttle-pedal sensing.

The AS5163 features a single-wire interface that you can configure as a 14-bit digital, 12-bit PWM (pulse-width-modulated), or ratiometric-analog output. You can also set it to cover any system-specific angle range by setting a start and end position of the rotational movement. The device operates at −40 to +150°C, comes in a TSSOP-14 package, and requires a 5V supply.

—by Graham Prophet

▶**austriamicrosystems, www.austriamicrosystems.com.**

Design and test combine to speed yield learning

Verigy has introduced its Yield Learning Solution, which integrates on-tester, real-time capture and analysis of electrical failures on complex SOC (system-on-chip) devices. The product combines preanalysis modules on the Verigy V93000 SOC test platform, including a design-centric analysis and visualization tool set. The Yield Learning Solution comprises the V93000 scalable test platform plus the Triage Fault Locator and YieldVision software tool sets for failure data capture and yield analysis. Triage provides on-tester fault localization and includes an on-tester sampling engine; YieldVision supports offline data analysis.

The scalable architecture of the V93000 allows for complete integration with the Yield Learning Solution. The Triage software's proprietary algorithms enable efficient data processing, and the YieldVision analysis and visualization tools reduce the time required to diagnose problems. By seamlessly linking electrical test with physical-layout data, the tools allow fast localization of the root-cause physical defects.

Diagnosing problems in nanometer-level-device design and manufacture is becoming more challenging, which makes it essential to close the loop between design, fabrication, and test, says Colin Ritchie, vice president of marketing for DFX (design-for-test/manufacture/yield) products at Verigy. Verigy's Yield Learning Solution addresses the design/fab matching that is essential for a successful business. He notes that design-for-manufacturability problems can result from lithography-unfriendly design or failure to adequately follow increasingly restrictive design rules. The inability to quickly isolate and fix such problems, he says, can lead to billions of dollars in lost annual revenue, citing VLSI Research (www.vlsiresearch.com) figures.

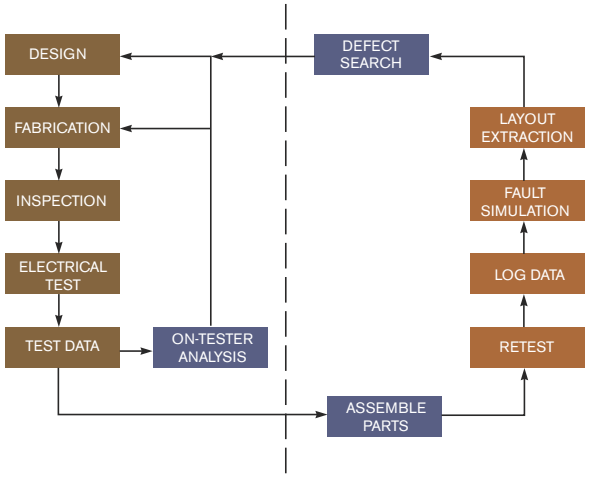

Verigy's on-tester analysis capabilities can shave weeks off the time it takes to make design revisions based on manufacturing test data, eliminating retest, fault-simulation, layout-extraction, and defect-search steps.

Ritchie notes that traditional approaches to yield diagnosis can require many days to identify design problems that lead to yield loss; such iterative approaches often require retest of failed devices plus a sequence of fault simulation and layout extraction to physically locate faults. The process not only is time-consuming but also can generate terabytes of data. In contrast, says Ritchie, Verigy's on-tester approach generates only kilobytes of data and delivers results in minutes.

The Yield Learning Solution efficiently links test back into both design and the fab, providing logic bit maps for both stuck-at and difficult-to-detect timing faults in scan chains and logic. The Yield Learning Solution provides both the accuracy necessary for the lab and the high throughput necessary for production—critical for both new-product introduction and ongoing manufacturing monitoring. Ritchie says that Triage can perform on-tester localization of blocked scan chains and hold-time faults while performing on-tester characterization. YieldVision, he says, "speaks the language of the designer and the language of the fab," providing diagnosis at the wafer, die, and component levels. Ritchie cites customer results indicating a four-week acceleration in time to market, an increase in entitlement yield of as much as 6%, and a tenfold reduction in the number of waivers required to reach entitlement yield.

—by Rick Nelson

▷**Verigy**, www.verigy.com/go/yield.

LITHIUM-BATTERY FORMULATION SUITS MEDICAL APPLICATIONS

08.20.09

EaglePicher Medical Power recently introduced a lithium-CFX (carbon-monofluoride)-battery chemistry, an enhancement of the company's proprietary CFX-battery technology. It results in a self-discharge rate of less than 1% per year at room temperature. Although lithium thionyl has a similar self-discharge rate, it increases rapidly with higher temperatures, making its use a liability in implantable medical devices, which must operate at 98.6°F.

The CFX formulation also offers energy density of 180 Whr/l—approximately 25% more than that of battery technologies that medical devices currently use. CFX lets battery-pack designers choose between a larger battery that runs for 25% longer and a 25% smaller one with an equivalent runtime.

The CFX formulation incorporates an end-of-life indicator to accurately predict battery depletion six months in advance, requiring fewer replacement surgeries for implantable medical devices. The indicator manifests itself as a sharp falloff in voltage yet still maintains adequate voltage and power to operate. The new technology will be available in a number of configurations at powers of less than 350 mAhr to 10 Ahr.

—by Margery Conner

▶**EaglePicher**, www.eaglepicher.com.

Western Digital packs 1 Tbyte into 2.5-in. disk

Western Digital recently broke its own record, becoming the first to shoehorn 1 Tbyte into a 2.5-in. hard-disk drive. This breakthrough comes on the heels of the company's becoming, in January, the first to ship a 2-Tbyte, 3.5-in. drive by squeezing four 500-Mbyte platters into the form factor. This approach was reminiscent of Hitachi's (www.hitachi.com) four-platter approach, landing that vendor in first place to cross the threshold for 1-Tbyte, 3.5-in. hard-disk-drive storage in early 2007. The company migrated to a three-platter configuration 18 months later.

Western Digital's drive, the WD Scorpio Blue, has a 3-Gbps SATA (serial-advanced-technology-attachment) interface and an 8-Mbyte RAM cache. The MSRP (manufacturer's suggested retail price) is $249.99; a 750-Gbyte version, the WD7500KEVT, sells for $189.99. External USB (Universal Serial Bus)-interface variants are also available. MSRPs for the 1-Tbyte and 750-Gbyte versions of the My Passport Essential SE are $299.99 and $199.99, respectively. All these drives feature an atypical 5200-rpm speed versus the more common 5400 rpm—whether for additional per-platter storage potential, to enhance the drives' power-consumption capabilities, or for other reasons.

Speaking of platters, the company accomplished its achievement by bumping the total per-drive platter count to three versus the more typical one- and two-platter specifications. This augmentation increases drive height to 0.49 in. (12.5 mm), thereby making the drives unusable in some ultra-thin-system designs, which rely on the more usual 0.374-in. (9.5-mm) thickness. As such, it's unclear how much if any per-platter areal-density leadership Western Digital has over competitors, such as Seagate (www.seagate.com).

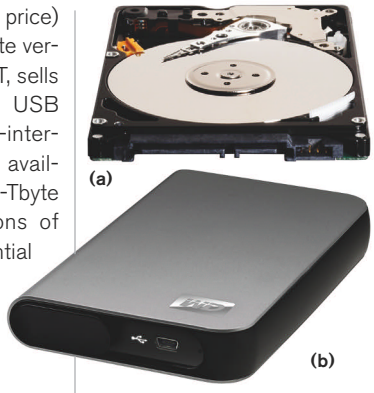

The 2.5-in., 1-Tbyte WD Scorpio Blue hard-disk drive has a 3-Gbps SATA interface and an 8-Mbyte RAM cache (a). The 1-Tbyte and 750-Gbyte versions of the My Passport Essential SE are external USB-interface versions of the Scorpio Blue (b).

Seagate is now promoting a 640-Gbyte, 2.5-in. hard-disk drive with an external USB interface. Presumably, the company based the product on as-yet-unannounced, two-platter, single-drive technology, translating to 320 Gbytes per platter. Compare this data point to the 250- and 333-Gbyte/platter specifications of Western Digital's latest offerings, and you can see how close the two companies are in this respect.

—by Brian Dipert

▷**Western Digital Corp**, www.wdc.com.

ASIC DEMULTIPLEXES TO MULTIPLE DISPLAYS FROM ONE DISPLAYPORT SIGNAL

HDMI (high-definition multimedia interface) and DVI (digital-video interface) transmit video data as continuous bit streams, whereas DisplayPort transmits the data in packets and allows for asymmetric two-way transfers. If your application is simply connecting a graphics chip to a display, packetizing creates a lot of overhead for little real benefit. If you are driving multiple displays, however, there are some advantages to all that extra work. Chip designers and mechanical engineers who have to multitask need multiple displays. Both operating-system developers and the graphics-system designers have made provisions for this requirement. The question is how to get the data to the right display. One approach has been to use an additional display-driver card per monitor in the computer chassis. Another is with an external splitter box. Both of these approaches can cost hundreds of dollars and consume 100W or so. Another alternative is to go to USB (Universal Serial Bus)-display links, but that approach quickly runs into bandwidth problems.

IDT (Integrated Device Technology) may have a better idea with the VMM (virtual-machine-monitor) 1300 PanelPort ViewXpand chip. This device, powered by the DisplayPort cable at less than 1.5W, is basically a router for DisplayPort micropackets. The chip inspects incoming micropackets, sorts them based on the pixel addresses, modifies the address to correspond to the screen coordinates on the appropriate monitor, and routes each packet to that monitor. The chip works as a hub in a dongle, in which it can route packets to any of three output ports, or in a daisy-chain topology, in which you use only two outputs, and you can string together as many chips and monitors as your graphics bandwidth allows.

The chip can also function as a protocol converter, so, when it is in hub mode, you can drive legacy display interfaces on the chip outputs. The chip routes the DisplayPort packets, so it requires no new driver software and is compatible with DisplayPort standards, including HDCP (high-bandwidth-digital-content protection) Version 1.3–a requirement for monitors that handle content-protected video.

At initialization, the chip uses identification cycles to figure out what monitors it will be driving and to map the viewing window onto the various monitors' screen coordinates, thereby setting up the routing table. The VMM 1300 is available for sampling, and production chips will become available this month.

—by Ron Wilson

▶**Integrated Device Technology, www.idt.com.**

VOICES

Asset InterTech Inc's Tim Dehne: seeking growth in embedded instrumentation

Tim Dehne, until recently a longtime executive with National Instruments Inc, has joined the board of directors of Asset InterTech Inc, a supplier of boundary-scan and embedded-instrumentation tools. Over a career stretching more than 21 years at NI, Dehne led global marketing and R&D at the company, which reported $824 million in revenues in 2008. During his tenure at NI, he held positions including vice president of strategic marketing and senior vice president of R&D. *EDN* recently spoke with Dehne. A portion of that interview follows. You can read the full interview at www.tmworld.com.

This move represents a major change for you.

A Yes, you could say that. I had been looking at a number of different options, but Asset InterTech presented a nice opportunity based on my experiences and relationship with Glenn [Woppman, Asset's chief executive officer and chairman]. It made sense to me, and I was happy to join the board.

What do you see as the similarities and differences in the directions of the two companies? National Instruments started out as a test company, but, as it grew, it's gone in lots of different directions, including control and design.

A Growth potential is one of the things that got me excited when Glenn and I first started talking after I announced my departure from National Instruments. Asset's history is in boundary scan, which basically "niched out" as a market. Asset did very well and is one of the leaders, but the space didn't grow as much as Asset and the other players in the marketplace hoped it would. But Asset is still fundamentally a hardware/software play, and that [approach] is very similar to National Instruments'.

But where the company is going is what got me excited. The semiconductor world is moving to many, many cores and even to multiple IP [intellectual-property] cores from multiple vendors. That world needs a kind of test strategy not only at the semiconductor level, where you figure out what's going on with the silicon and what may be causing some of the yield or performance issues, but also at the board level to monitor the interactions among all those different chips with all those multiple vendors' IP. Addressing these challenges involves structural test, which is different from what National Instruments does, but it seems to be an area of growth and an exciting technology development. We feel that embedded instrumentation is an area in which the company can grow to quite a good size if we do things right.

What are the prospects for mergers and acquisitions?

A I'm not saying what Asset's plan is, but, in the design space, mergers and acquisitions happen all the time. They're not as frequent in the test area, but, of course, they occur there, too. But the most important thing now with Asset is to get the company growing organically with these new initiatives with the processor-controlled test and embedded instrumentation, so that's what we are going to focus on doing.

Do you plan to form partnerships with complementary vendors?

A "Partnerships"—and I use that term loosely—typically fall into three categories. There can be a sales/distribution agreement, there can be just a marketing or message story, or there can be some in-depth joint development. The last type can be significant. Examples are the partnerships that NI had with Analog Devices [regarding NI's LabView graphical-development module for Analog Devices' Blackfin processors] and Luminary Micro [regarding a LabView graphical-development module for ARM targets on a Luminary Micro evaluation board]. These partnerships involve developers' sitting down and co-developing and creating something new in the marketplace. With joint-development efforts, generally there's money on the line; you are committing real resources, so you work harder to make those things work.

What is your role as a board member–the traditional board of directors' governance role?

A Certainly, anybody on the board has that role. But I'm a little bit different from the other board members, and I'm thankful to Glenn for recognizing the ways he could leverage my experience. When I started at National Instruments, the company employed 100 people, and I went up the marketing ranks for about a decade and then went into R&D for another decade. So given that experience base, Glenn and some of the other board members felt I could probably play a different role in addition to the corporate-governance role. I could also be a little bit more active in the marketing and R&D aspects.

—interview conducted and edited by Rick Nelson

RAQ's

Rarely Asked Questions

Strange stories from the call logs of Analog Devices

Some Chips Have Moving Parts!

Q. Is it true that over the years electrical devices have evolved to where they have no moving parts at all?

A. While it is true that any machinery becomes more reliable as the number of moving parts where friction can cause wear is reduced, there are actually integrated circuits (ICs) which only work because of moving parts on the surface of the chip.

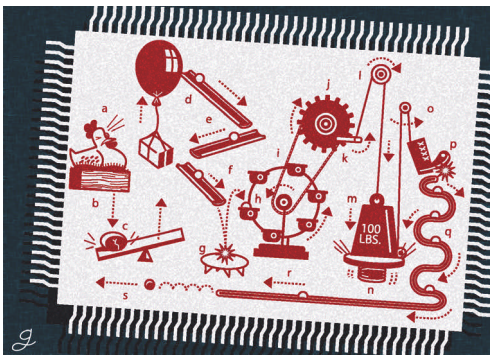

These are known as Microelectromechanical Systems, or MEMS. They use standard IC process technology to make structures in metal, silicon and silica on the surface of a chip. Such structures may be designed to move and thus perform many useful functions.

Moving parts of such chips flex, but do not usually bear on other surfaces, so friction is not often a problem. They are usually made of silicon, which has very low mechanical hysteresis with deformation, and therefore great resistance to fatigue. Silicon does not change its properties or suffer damage even when flexed many trillions of times.

The first commercial MEMS devices with visibly moving parts were accelerometers. Electronic accelerometers once cost hundreds or thousands of dollars, today the least expensive cost only a dollar or two making it economical to use them in inexpensive gadgetry, air bag deployment (their first major application), joysticks for computer games, shock protection for disk drives and athletes ankles, keystone correction in projectors, orientation detection in hand-held monitors, and a thousand other uses.

Traditional gyroscopes rotate, but it's quite possible to make a gyroscope that vibrates rather than rotating; and this can easily be done in a MEMS structure. MEMS gyroscopes are used wherever rotation is measured; their low cost enables hitherto unaffordable applications such as optical image stabilization, safety controls in motor vehicles and short-range inertial navigation for GPS receivers when no satellite is visible.

MEMS structures allow the manufacture of high quality microphones on an IC chip. These are smaller, more cost-effective and reliable than any other microphone technology and are starting to replace the electret microphone in many applications.

The linked articles describe all these chips with moving parts, and their many uses, in much more detail.

To Learn More About MEMS Technologies
http://designnews.hotims.com/23114-101

Contributing Writer James Bryant has been a European Applications Manager with Analog Devices since 1982. He holds a degree in Physics and Philosophy from the University of Leeds. He is also C.Eng., Eur. Eng., MIEE, and an FBIS. In addition to his passion for engineering, James is a radio ham and holds the call sign G4CLF.

Have a question involving a perplexing or unusual analog problem? Submit your question to: raq@reedbusiness.com

For Analog Devices' Technical Support, Call 800-AnalogD

BY BONNIE BAKER

One circuit provides system resolution and 12-bit accuracy

Handheld meters, data loggers, and automotive and monitoring systems typically require a multiplexed system with the low-cost combination of high accuracy and high system resolution. A system that can handle this diversity requires a multiplexer, a gain cell, and an ADC. A feasible approach has a 10-channel PGA (programmable-gain amplifier) teaming up with a medium-speed, 12-bit SAR (successive-approximation-register) ADC (**Figure 1**). The single-supply, 10-channel PGA has a rail-to-rail I/O with a gain adjustment of 1 to 200V/V. The PGA's low-noise performance of 12 nV/$\sqrt{Hz}$ at 10 kHz is appropriate for a 12-bit system.

The analog interface between these two devices includes an operational amplifier in a buffer configuration and an RC circuit. The 12-bit, capacitor-based SAR ADC has an inherent sample/hold function and requires the RC circuit, which facilitates the charging action of the ADC's input structure. The calculated value of the PGA noise, referred to the output, is equal to the PGA's noise density at 10 kHz (12 nV/$\sqrt{Hz}$) times the square root of the PGA's closed-loop bandwidth times the square root of $\pi/2$. The multiple of $\sqrt{(\pi/2)}$ accounts for the noise in the frequency region beyond the PGA's bandwidth. You then multiply this number by the gain of the PGA. The following **equation** uses a PGA gain of 16V/V: $PGA_{RMS\text{-}NOISE}$=12 nV/$\sqrt{Hz}\times\sqrt{(1.6\ MHz\times\pi/2)}\times$16V/V=304 µV rms. The ADC noise of 431 µV rms from this converter is well below 1 LSB or 1.22 mV in this 5V system. The noise from the buffer amplifier, which is 39 µV rms, contributes little or no noise to this system.

The combined noise of the PGA, op amp, and ADC is 529 µV rms, which is still less than 1 LSB of the 12-bit converter. You calculate this value using a root-sum-square equation or the following **equation**: Noise referred to output=$\sqrt{(PGA_{RMS\text{-}NOISE}{}^2 + \text{op amp}_{RMS\text{-}NOISE}{}^2 + ADC_{RMS\text{-}NOISE}{}^2)}$. The equivalent 12-bit accuracy of this system when the PGA is in a gain of 16V/V is 0.432 LSB: Noise referred to output$\times 2^N$/FSR (full-scale range), where N is 12 and FSR is 5V/V. If you look at this system across the PGA's gain range of 1 to 200V/V, you find that the PGA dominates the noise-contribution portion in this circuit. Once the PGA gain exceeds approximately 125V/V, this system no longer matches the 12-bit-accuracy criterion. However, the system's referred-to-input LSB voltage size becomes smaller (**Figure 2**). The trade-off for a smaller LSB is a decrease in the system's effective number of bits.

The system in **Figure 1** provides an adequate gain range for the PGA when 12-bit accuracy is required and an equally adequate gain range when good system resolution is required. EDN

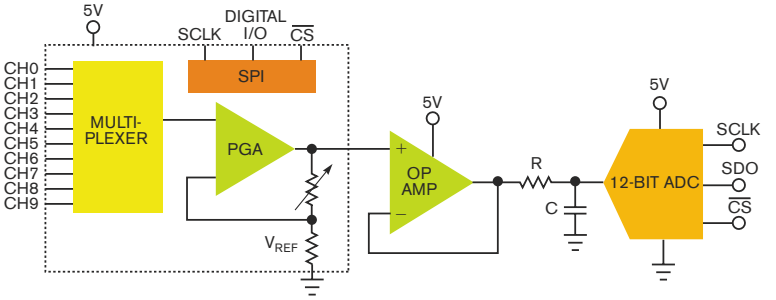

Figure 1 This system uses a multiplexed PGA and an operational amplifier that drives a 12-bit converter.

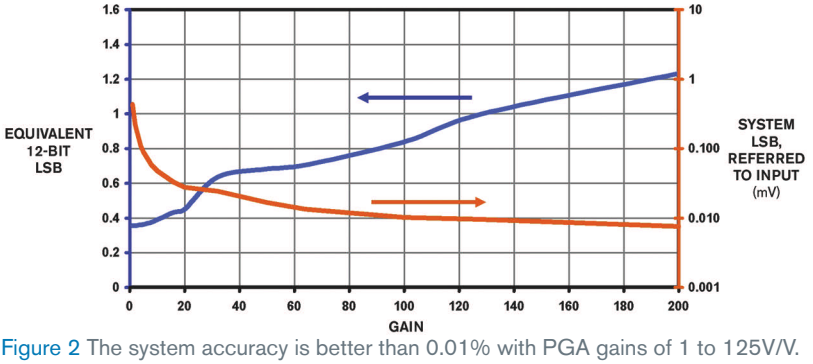

Figure 2 The system accuracy is better than 0.01% with PGA gains of 1 to 125V/V. With PGA gains of 125 to 200V/V, the system accuracy is better than 0.02%.

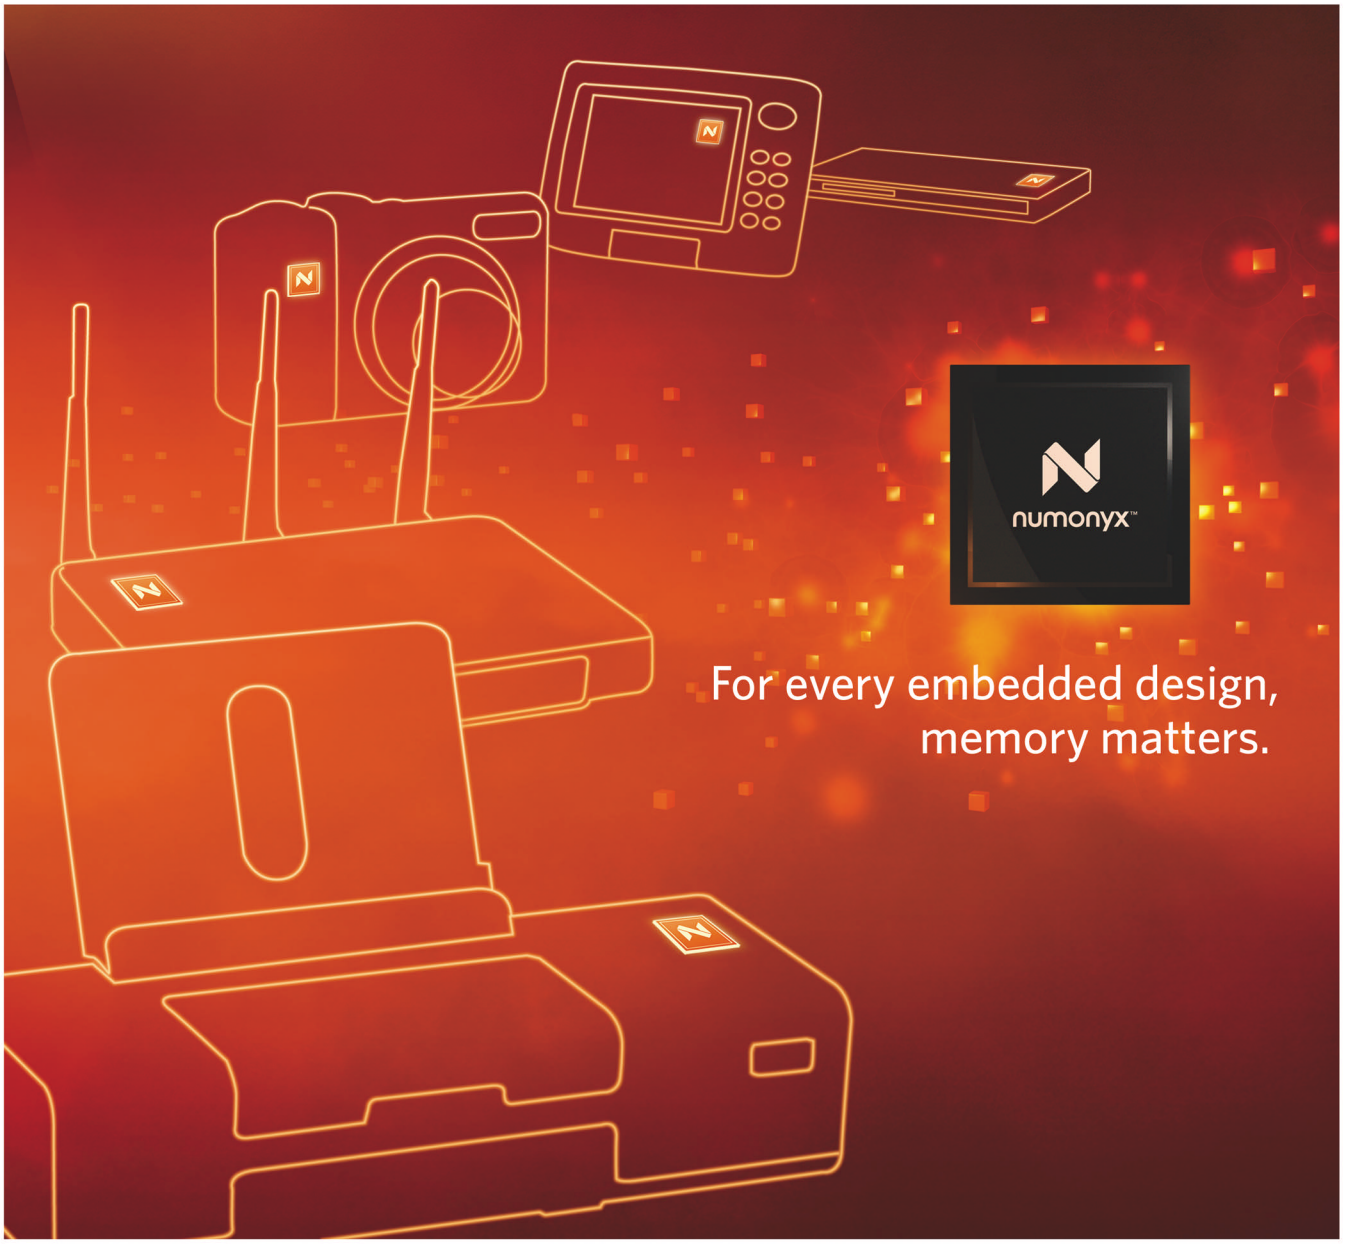

When it comes to creating your next embedded system, an important decision awaits you. Memory. Your design needs speed, reliability, performance and capacity to store the code and data your design demands. No problem. Numonyx has the broadest portfolio of parallel and serial NOR, NAND and phase change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

NUMONYX® FORTÉ™ SERIAL FLASH MEMORY

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

NUMONYX® AXCELL™ PARALLEL NOR FLASH MEMORY

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

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

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

numonyx™
innovative. memory. solutions.

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

Rex 6000 PDA

Franklin Electronic Publishers introduced the Rex electronic organizer at the 1997 Comdex show. The Rex was a Type II-compatible, 3.4×2.1×0.13-in., 1.4-oz PCMCIA, or PC-Card. The company later introduced a version with 512 kbytes of memory. Xircom bought the Rex line in 1999 for $13.25 million. Intel acquired the Rex as part of its acquisition of Xircom, which was completed in March 2001. A month earlier, Xircom had introduced the Rex 6000, which Citizen Watch Company of Japan produced and marketed as the DataSlim-2. The organizer held phone numbers, a to-do list, an appointment calendar, a memo list, and several utilities, such as a calculator and some games. The Rex 6000 had 2 Mbytes of flash memory and 32 kbytes of RAM. It initially sold for $150. Intel canceled the project shortly after its acquisition of Xircom.

Two 3V CR-2016 lithium batteries with 90-mAhr capacity power the PDA. The operating system shows a low-battery icon when power decreases to 2.86V and emits a warning at 2.82V.

The five hardware buttons connect through a six-circuit ribbon cable to a connector on the motherboard. The ribbon goes under the battery clip and snakes through a gap in front of the connector.

A piezoelectric actuator lies on top of the chips to provide alarm and beeping sounds. Plastic tape over the back of the case ensures that the actuator and the battery do not short out.

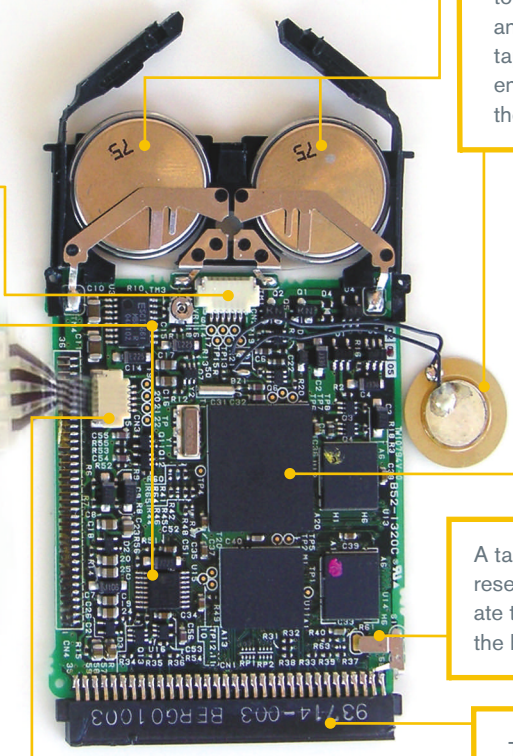

The 4.3-MHz Rex 6000 CPU is a Toshiba microprocessor that is compatible with the Zilog Z80. A separate 32.768-kHz watch crystal provides the real-time clock, two Fujitsu 29DL164BD-90 chips provide flash memory, and the fourth digital chip is an LCD controller.

Analog chips include an Epson SCI7661 charge-pump dc/dc converter to make a negative bias for the LCD. A Texas Instruments quad, 10-bit, 85k-sample/sec TLV1544 ADC senses the touchscreen. Other analog chips include several linear regulators.

A tab of metal forms a crude reset switch that users actuate through a small hole in the back of the case.

The gray-scale, 240×120-pixel, 2.6-in.-diagonal LCD has no backlight. The display connects to a ribbon cable soldered to the motherboard. The row- and column-driver chips mount on the LCD glass.

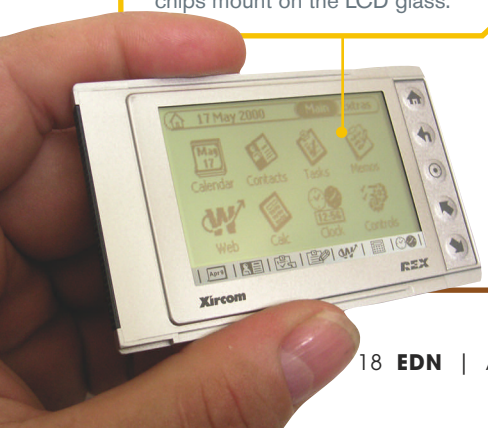

The resistive touchscreen connects to the motherboard by a small, five-pin ribbon connector with only four used circuits. Icons on the bottom of the touchscreen provide seven soft buttons for the four primary PDA functions as well as a network download, a calculator, and a clock function.

The PCMCIA form factor of the Rex 6000 allows you to synchronize it by sliding it into the PC-Card slot of a laptop computer. Alternative offerings are RS-232 and USB docks. The Rex does not use the PCMCIA bus to communicate; it uses the UART channel that is part of the PCMCIA standard. The USB cradle also emulates a serial UART.

Agilent

Tektronix

LeCroy

Rohde & Schwarz

National Instruments

Anritsu

Keithley

Yokogawa

Tabor

Pickering

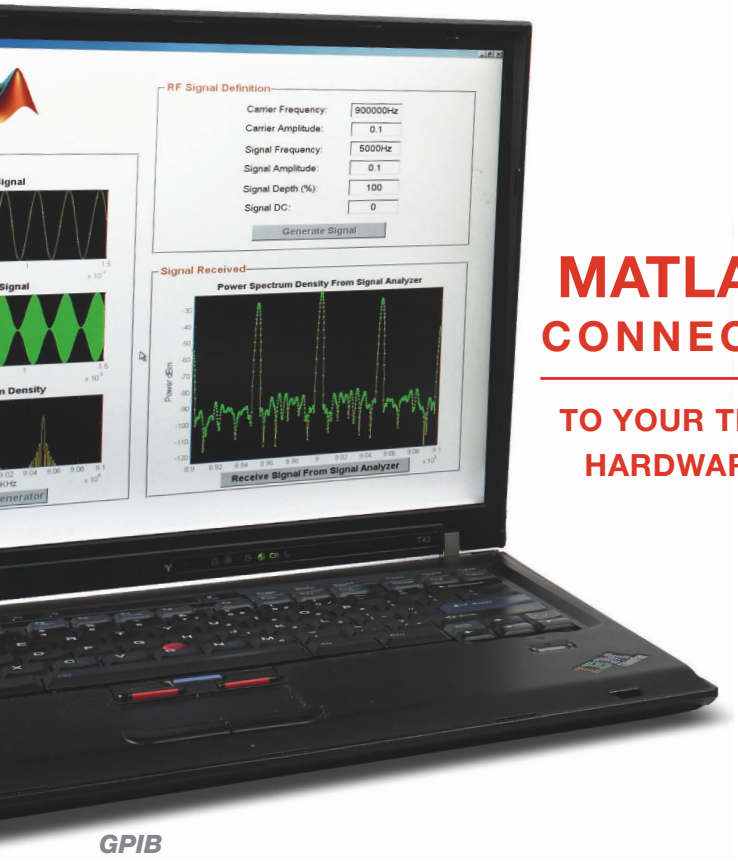

MATLAB CONNECTS

TO YOUR TEST HARDWARE

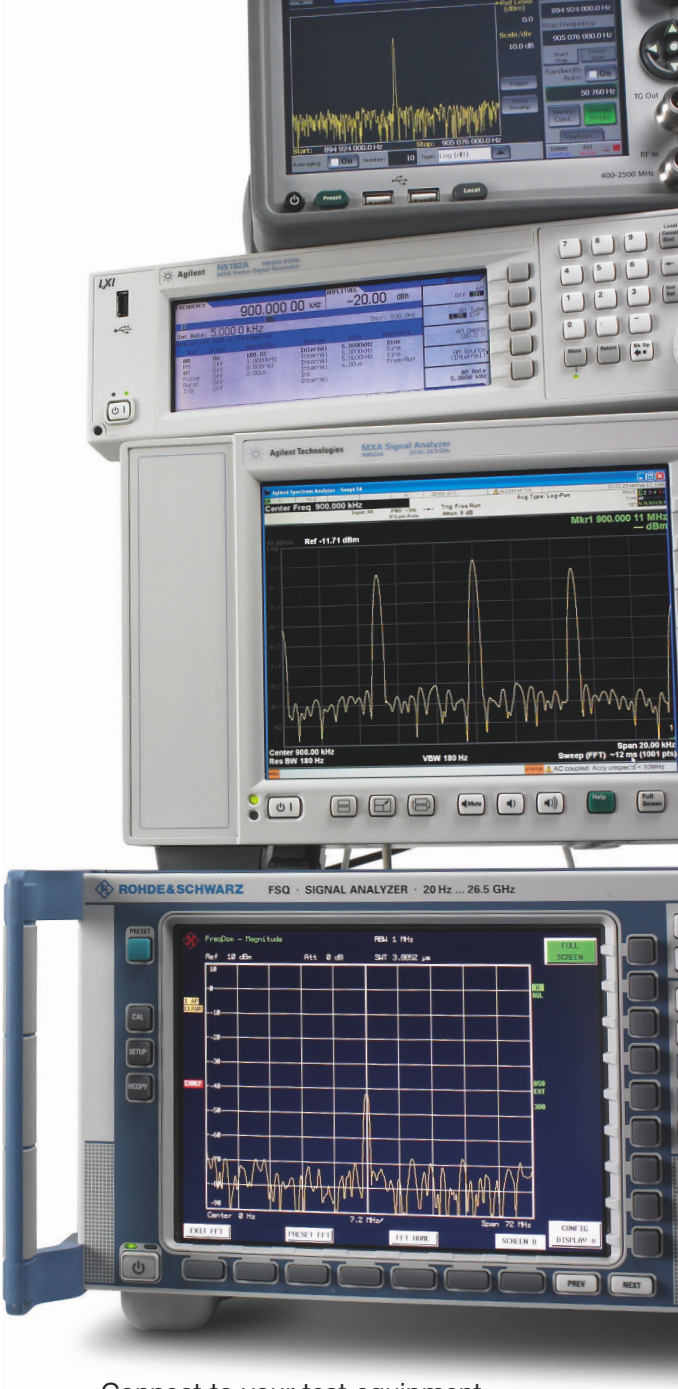

Connect to your test equipment directly from MATLAB® using standard communication protocols and hundreds of available instrument drivers.

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

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

GPIB

LXI

IVI

TCP/IP

VISA

USB

UDP

RS-232

The MathWorks™

Many designers are familiar with open-source software, such as Linux, in which the source code is available to all. However, fewer are familiar with organizations offering open-source hardware. These organizations release free information, including schematics, BOM (bill-of-materials) information, and PCB (printed-circuit-board)-layout data, covering the overall hardware design. Designers with this information can build or add to a freely available design. In many cases, open-source software supports the original design, providing additional advantages. Some aspects of open-source hardware go beyond the sharing of the design itself. These aspects can save time and money for not only hardware developers but also PCB designers and fabricators, contract manufacturers, and even software developers.

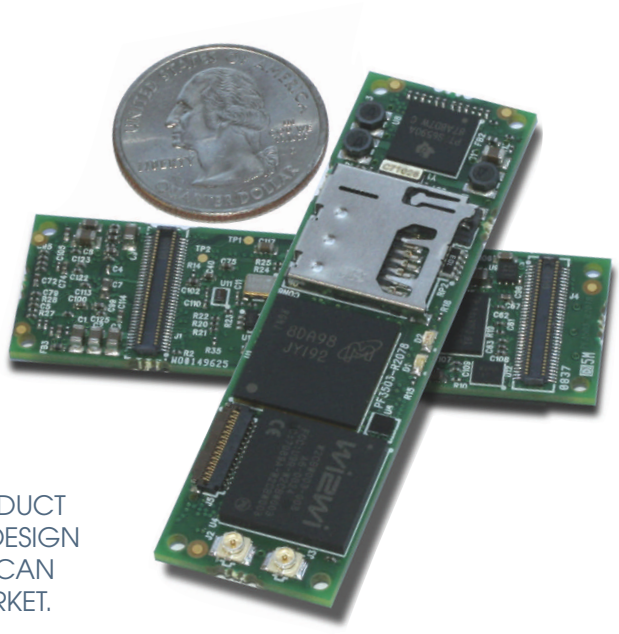

BASING YOUR PRODUCT ON A REFERENCE DESIGN OR DEMO BOARD CAN SPEED TIME TO MARKET.

TAKE ADVANTAGE of OPEN-SOURCE HARDWARE

BY GERALD COLEY • TEXAS INSTRUMENTS

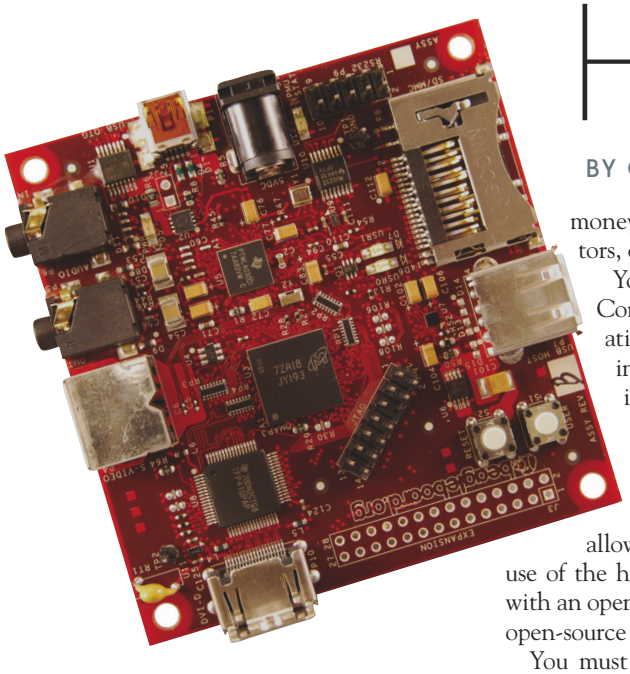

Figure 1 The BeagleBoard (bottom left) and the Overo expansion boards from Gumstix (top right) are examples of available open-source hardware.

You can license open-source projects from organizations such as Creative Commons, which offers the Attribution-ShareAlike licensing program. Creative Commons stipulates that a user must attribute the open-source work in the manner that the original designer specifies but not in a way that indicates that the original designer endorses the user's work. Likewise, if users provide that work as open-source hardware, releasing it back to the community for access by others, then they must provide that work under the same Attribution-ShareAlike licensing (**Reference 1**). Other licenses, such as the modified BSD (Berkeley Software Distribution), allow for the assignment of copyrights and provide certain restrictions to the use of the hardware design (**Reference 2**). Be sure to read the license that comes with an open-source design before using it. If users are considering creating their own open-source design, they need to figure out which license works best for them.

You must consider several factors, including power, cost, and documentation, when selecting an open-source-hardware platform. Make sure that the license provides comprehensive, high-quality documentation, including schematics, BOM, and PCB data. The documentation must have the support of a large community of users, and it must align with your product's needs.

Some popular open-source-hardware platforms include Gumstix, Arduino, and

the BeagleBoard (**Figure 1**). Gumstix uses an open-source-hardware model in posting the schematics and layouts of all the company's Overo-series expansion boards. Arduino employs a microcontroller as its hardware and has its own community of designers and hobbyists. Arduino's schematics come in Eagle and PDF (portable-document format), and the PCB information is in Eagle. The Creative Commons license covers licensing for the platform, and extensive libraries and software support are available. The BeagleBoard uses a Texas Instruments OMAP3530 (**Reference 3**). A large, Linux-based open-source-software community supports the BeagleBoard, and schematics are available in PDF and Cadence's OrCAD. The BOM is in Microsoft Excel, and PCB information is in Cadence's Allegro and Gerber files.

Due to its high technology level, the BeagleBoard presents some interesting challenges for users of the OMAP3530 device, but you can overcome the challenges by taking advantage of open-source hardware. The 515-pin OMAP3530 device uses 0.4-mm-pitch balls and supports POP (package-on-package) technology, which mounts the memory devices on top of the processor. This technology can create challenges, including schematic design and PCB layout, fabrication, and assembly.

AT A GLANCE

- Open-source hardware offers an advanced start on your design.
- Open-source software complements open-source hardware.
- Open-source hardware prepares your PCB (printed-circuit-board)-fabrication and -assembly houses for high-volume production.
- You may want to share your improvements by making them open-source additions, as well.

STREAMLINE THE DESIGN

You can reduce risk by basing designs on open-source hardware and taking advantage of a proven design that has operated successfully in the past. Thus, you can work from a known starting point and easily see what's there, what's missing, and what is unnecessary because you have access to both the finished hardware and the complete design. It also saves development time. Having access to the complete BOM with part numbers enables you to quickly adjust it. You can look for places to substitute your favorite capacitors and resistors and get access to the information on any unfamiliar parts. As long as the cost and specification of the part in use are on target, you should have no problem with using it. If you are unfamiliar with a part, however, you can improve the design by selecting a part better suited for your design's needs.

For instance, the OrCAD schematic tool lets you quickly and easily add devices to the schematic. It takes advantage of the unused pins on the processor and even replaces devices on the open-source-hardware design. Creating components can take a lot of effort, so using those on the schematic saves time and reduces the risk of errors. Alternatively, you can use a PDF version of the hardware schematic to create schematics with your favorite tool. This approach can be an advantage because re-creating the design gives you more in-depth knowledge of how the design works and where problems might arise. The BeagleBoard's documentation includes a fairly detailed reference manual that can answer questions about the design. Users also can access the support community for additional help. Don't be afraid to ask for help; open-source hardware is all about learning from other people's mistakes.

Using the Allegro file format is the shortest path to completing a layout. You can import footprints from this database to your library, saving a significant amount of time. Because you already have a working board, you can use these footprints with confidence. If your board or prototype is similar to the BeagleBoard design and you need to add several components, you can simply do an ECO (engineering-change order), meaning that you use the current layout and add only key components so that you don't disturb the basic layout.

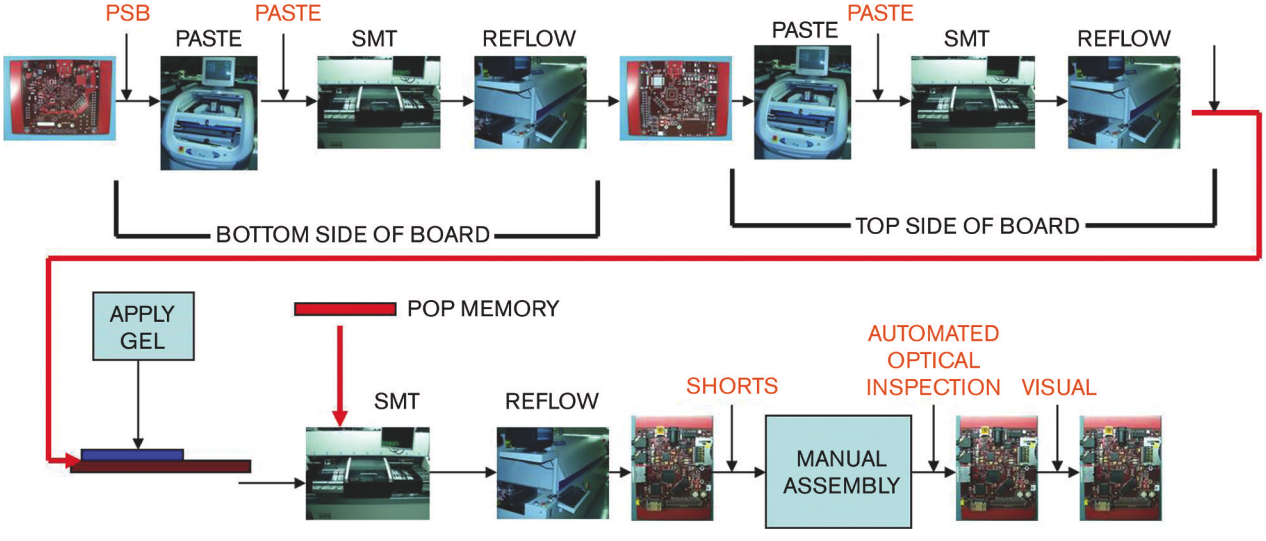

Figure 2 A board-assembly flow begins with a PCB and ends with a visual inspection of the work.

Allegro has a useful free viewer for reading design files. This viewer exposes the layer stacks and all of the PCB information, which provides a guide when designing the board on your own tool. In this way, you can see board routing, trace widths, and trace spacing. Another option is to convert the database to other tools using available and third-party tools. You must be careful not to lose design information when using this method, however. You can use Gerber files to double-check the information and as learning tools. These files are not user-friendly but can be helpful in creating your layout. The OMAP3530 device also includes a design-guideline application note to assist in layout.

GETTING UP TO SPEED

Many new devices use more advanced technology, including blind vias, stacked vias, narrow traces, and via-in-pad technology, than the technology from your PCB supplier. Although the suppliers may be able to handle these advances, they may not have previously needed to provide it. If a user is comfortable with a PCB vendor because it's a known entity, though, it may be worth getting the vendor to the point at which it can handle the new needs. However, depending on your requirements and confidence level, you may be better off going to a vendor that has experience with the type of board you are working on. Using the available CAD data, the PCB vendor can get up to speed and prototype an initial board so that you are confident that the company can build it. If the vendor needs to adjust certain aspects of the design, it can modify the Gerber files for its normal processes.

Expected yield determines the cost of building boards. Building the board as it is gives vendors insight into yield. They can then prevent any problems that arise when building the new board. For example, if they have problems with via-in-pad technology, you should avoid using it in your design. The assembly house also may encounter problems in assembling your board. You can get the assembly house up to speed by using open-source hardware. Alternatively, you can use your own assembly house, if applicable, to build the entire board or a subset of the board.

In many cases, it can take several passes to work out the wrinkles in assembly. By using the open-source-hardware design and assembling as many boards as you need, you can solve any problems. It is better to work on these problems with a proven and tested board instead of your first prototype, on which it may be more difficult to find issues. The BeagleBoard also has a POP assembly guideline (**Reference 4**). As in the case of the PCB-layout guidelines, it also uses the BeagleBoard design (**Figure 2**).

There is always pressure to complete the hardware so that you can start the software verification. Basing your design on open-source hardware allows an early start on software development by using the available open-source-hardware boards to begin the development effort. Because open-source hardware typically also has a software component, the software team can jump-start its efforts by using the available code. This approach allows the hardware team to make sure it gets the design right on the first pass.

In addition, the software team may discover that it needs to change the hardware to improve performance or to add a feature. Making a lot of changes to the basic design quickly diminishes this advantage, however. Make sure that the basic design components, including the memory, power management, and key debugging peripherals, remain the same. Additionally, pay careful attention to the available software to see whether it offers the necessary applications or functions. Although software developers can write the software themselves, they typically would rather use the already-available software.

In summary, using the OrCAD schematic design tool or implementing your own version using open-source hardware reduces risk and saves development time. You then need to focus only on what to add or remove from the basic design to complete the final design. You can use the Allegro CAD files or Gerber files to provide a map with which

FOR MORE INFORMATION

Arduino
www.arduino.cc

BeagleBoard
www.beagleboard.org

Berkeley Software Distribution
www.bsd.org

Cadence
www.cadence.com

Creative Commons
http://creativecommons.org

Gumstix
www.gumstix.com

Realtime Technology
www.realtimedsp.com.cn

Texas Instruments
www.ti.com

you can reuse whatever works correctly. Next, select a PCB vendor to accurately build the technology and save time and cost for the prototype run. This approach increases the chance of first-pass success and reduces the number of costly problems that would require re-spins and debugging during the prototype phase. Then, focus on getting the assembly house to build the open-source-hardware board. Now that the hardware prototypes are ready to run software and you have tested the software from the open-source-hardware board, there is a much higher chance of success.

GIVING BACK

Open-source hardware is about sharing work with others for everyone's benefit. It is acceptable if you never meant for the product to be open. You need not make your changes available to the community. In the spirit of open-source hardware, however, it's beneficial for all parties to provide upgrades and additions to the community whenever possible so that the next user can add other enhancements. When you add a function to hardware, it affects the software, which adds a reason to enhance and improve the overall performance of the software to take advantage of the new feature. As developers design products based on this design, another community member has perhaps added the function with the already-completed software work to help make it a better product.

In the future, more companies will offer varying levels of open-source hardware to their customers and the community at large, creating an environment in which developers spend most of their efforts on improving rather than re-creating the design. The community can benefit from this common goal, so keep your eyes open for the next entrant in the world of open-source hardware. EDN

Go to www.edn.com/090820df and click on Feedback Loop to post a comment on this article.

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

REFERENCES

1 "License Your Work," Creative Commons, http://creativecommons.org/choose.

2 Nelson, Russell, "The BSD License," Berkeley Software Distribution, Oct 31, 2006, www.opensource.org/licenses/bsd-license.php.

3 "OMAP3530 Applications Processor," Texas Instruments, www.ti.com/omap3530-oshw.

4 Gutierrez, Keith, and Gerald Coley "PCB Assembly Guidelines for 0.4mm Package-On-Package (PoP) Packages, Part II," Texas Instruments, April 2008, www.ti.com/litv/pdf/spraav2.

AUTHOR'S BIOGRAPHY

Gerald Coley is a systems engineer at Texas Instruments, where he has worked for 13 years. Coley focuses on the OMAP processors and has designed more than 20 development boards, including the BeagleBoard.

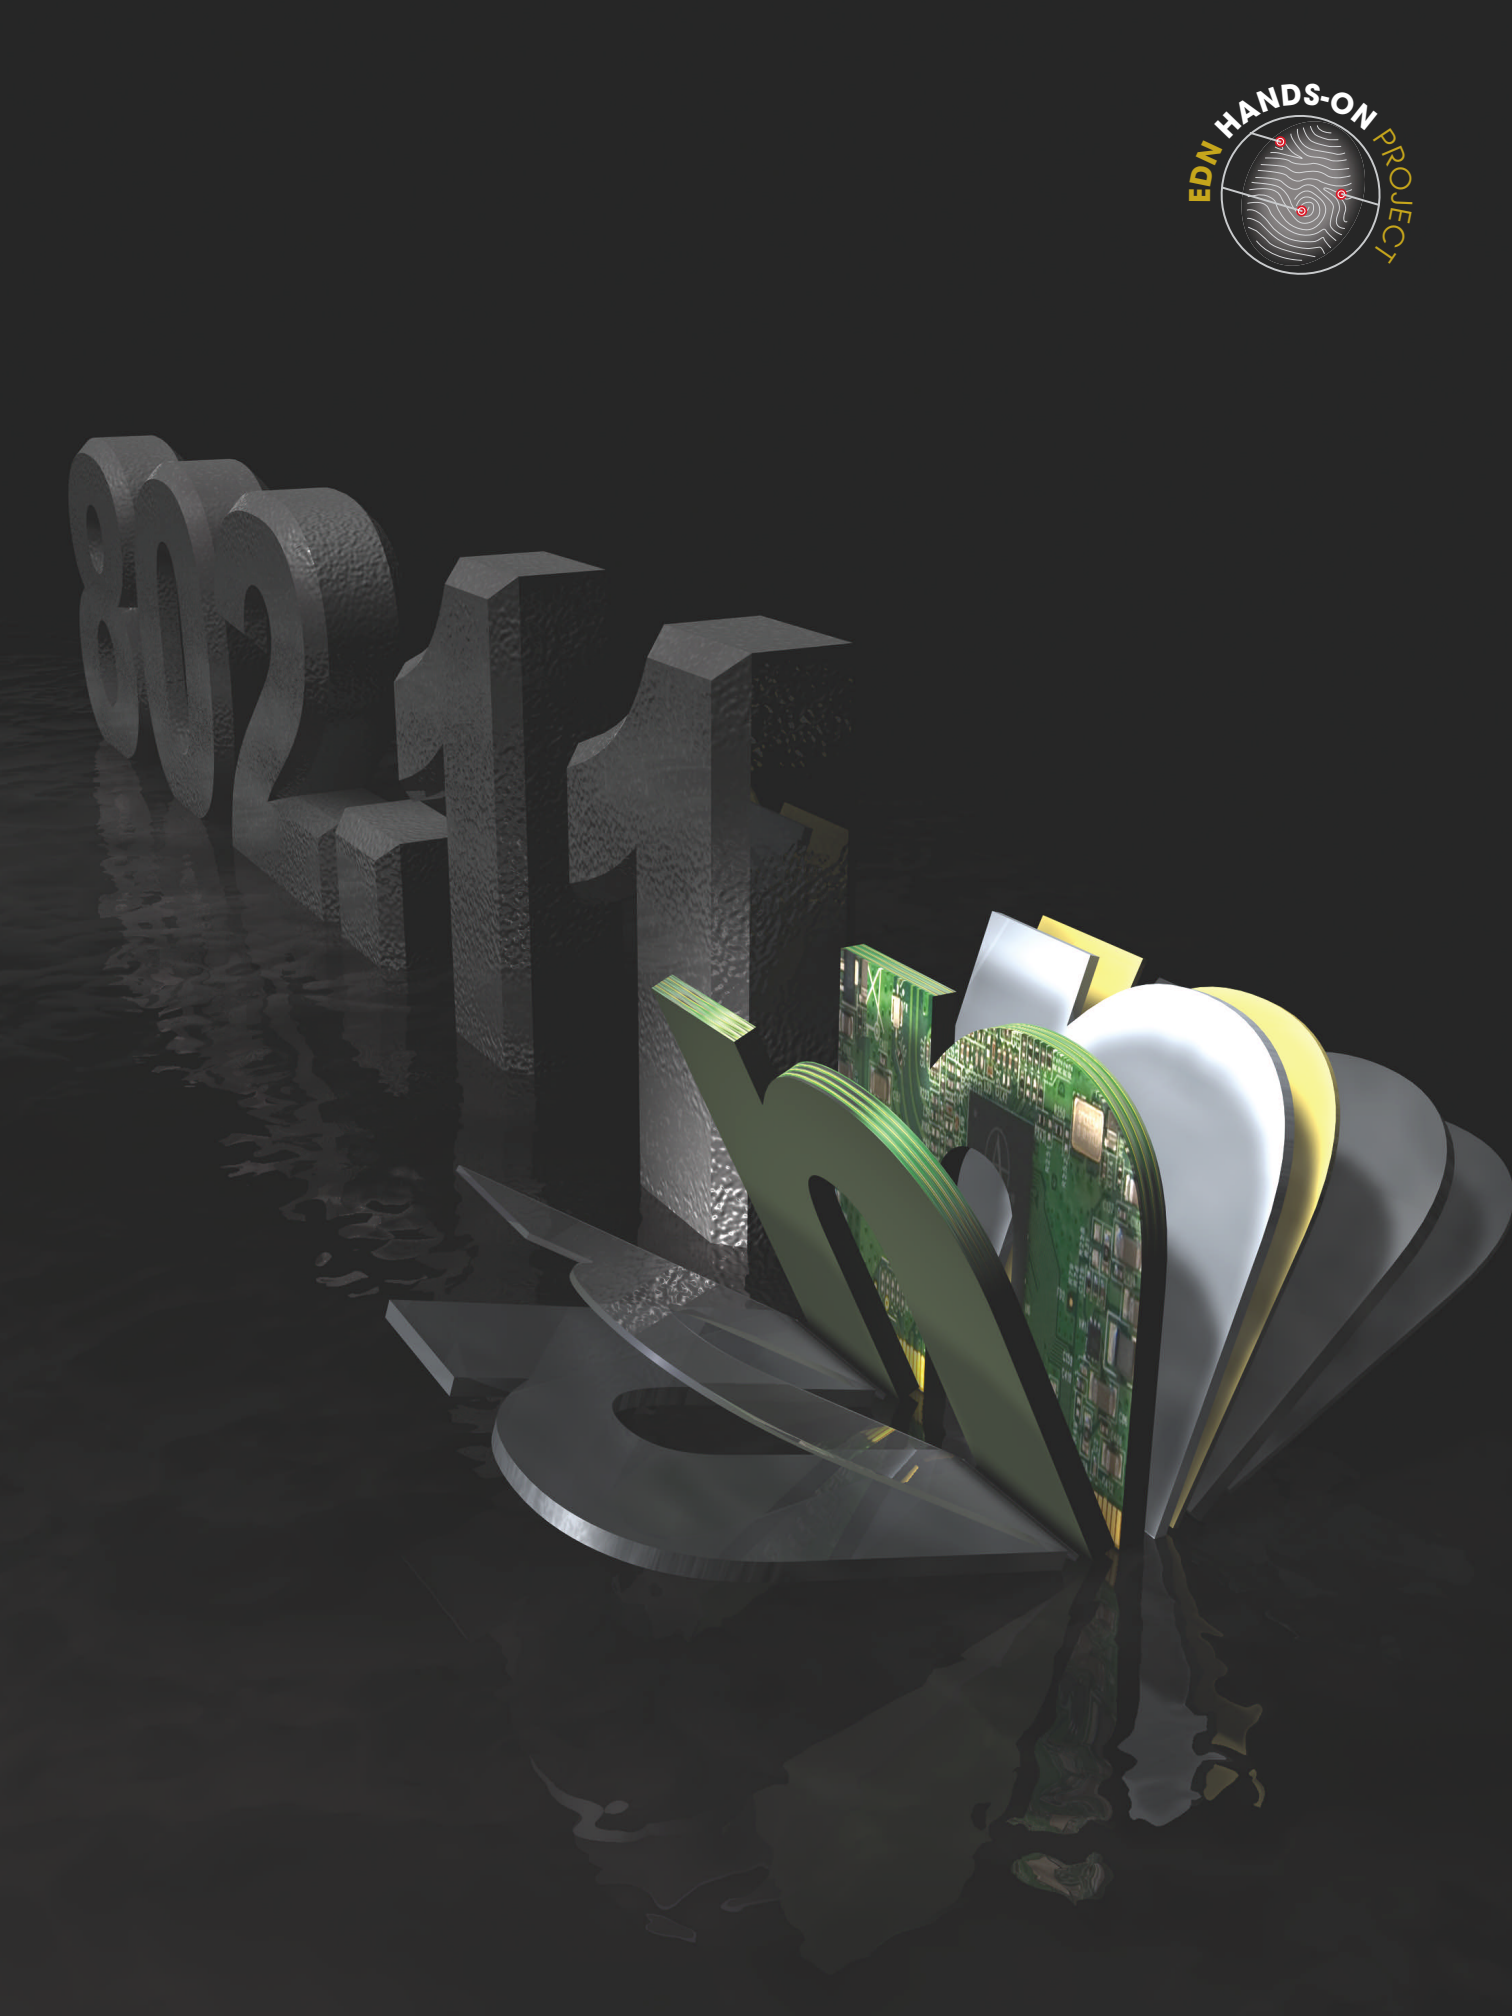
EDN HANDS-ON PROJECT
802.11

HIGH-RESOLUTION VIDEO-STREAMING SUPPORT IS SUPPOSEDLY A KEY JUSTIFICATION FOR 802.11N VERSUS ITS 802.11A AND 802.11B/G PREDECESSORS. SO WHY DOESN'T IT DELIVER ON ITS PROMISES?

TRANSPORTING HIGH-DEF VIDEO BROADCASTS: ARE WIRELESS NETWORKS UP TO THE TASK?

BY BRIAN DIPERT • SENIOR TECHNICAL EDITOR

It all seemed so simple, at least at first glance. I have a Microsoft Windows Vista Ultimate-based notebook PC, a Dell XPS M1330. I have two Xbox 360s, one in the living room and the other in the bedroom (**Figure 1**). I have reasonably solid over-the-air television reception at my home office (**Reference 1**). I'm not using the XPS M1330 as a work PC because I've migrated to mostly Apple systems—in some cases running Windows XP virtualized. So I decided to convert the Dell PC to a PVR (personal video recorder), leveraging its built-in Media Center capabilities and streaming both live television and recordings to the game consoles acting as Media Center Extenders.

I don't have Category 5 Ethernet cable running to any of the three LAN nodes, so I at first tried connecting them to the router and each other using Netgear's XAV101 HomePlug AV power-line-networking adapters (**Figure 2**). This arrangement worked fairly well, especially once I stuck ac-noise filters on the refrigerator and furnace fans' power connections (see **sidebar** "Revisiting power line"). Erratic bandwidth still resulted in more frequent glitches in the playback than I preferred, sometimes but not always when ceiling or window fans or other potential power-grid noise sources were operating. With annoying regularity, one or both Extenders would also refuse to connect to the laptop until I power-cycled one, some, or all of the HomePlug AV adapters in use.

Finally, one day I threw up my hands in frustration, determined to find some alternative way—other than crawling under the house and punching holes in floors to string Category 5 cable—of interconnecting these nodes. Ideally, I hoped to completely dispense with power-line networking in my LAN. (I was also using HomePlug AV to tether an Insteon home-automation controller to the router, thereby making the controller LAN- and WAN-accessible.) The XPS M1330 embeds a Broadcom 802.11n transceiver, my Apple router is 802.11n-capable, and 802.11n-based bridges and switches claim to allow legacy devices to leverage the IEEE's latest and greatest wireless-networking technology. Could 802.11n be my mentor to Media Center nirvana?

THE DEVIL'S IN THE DETAILS

Before diving into the sordid step-by-step story, here is some important background information. First, the Broadcom Wi-Fi IC in the XPS M1330 is a BCM4328, which Dell refers to as the Wireless 1505 Module, and, in conjunction with its mated MIMO (multiple-input/multiple-output) antenna array, it is dual-stream- and dual-band-capable (**Reference 2**). Conversely, the XPS M1330's wired-Ethernet transceiver, a Broadcom 59XX-series IC, is not GbE (gigabit-Ethernet)-cognizant; it supports only 10- and

100-Mbps Ethernet. The XPS M1330, which is in the living room and less than 25 feet away from the same-room router, sources video streams whose parameters are also critical to this project's outcome. MPEG-2-based Media Center serves files that are on average larger and, therefore, have higher playback bit rates than those that more modern video codecs, such as MPEG-4 and VC-1, create. Microsoft offers four quality-versus-bit-rate settings, but they apply only to analog recordings.

With ATSC (Advanced Television Systems Committee) sources, Media Center does no re-encoding and otherwise does not alter the incoming MPEG-2 video and Dolby Digital audio data; Microsoft simply embeds it as is within the proprietary DVR-MS "wrapper" format. You should assume, therefore, that each live-TV or recording audio-plus-video stream you want to route around your network is worst-case roughly 20 Mbps, accounting for DVR-MS overhead beyond the 19.2-Mbps ATSC bit rate. Media Center employs UDP (User Datagram Protocol) as the transport protocol, along with RTP (Real-Time Transport Protocol) for multimedia streaming and RTSP (Real-Time Streaming Protocol) for control functions.

In addition to Xbox 360s, the network nodes in my living room—12 feet away from the router—and in my bedroom also both include a Sony PlayStation 3. An Apple TV is in the living room, and the bedroom contains both a Roku Netflix Player and SoundBridge. I had been using a 10/100-Mbps Ethernet switch at each node to share the HomePlug AV connection among multiple pieces of gear because the HomePlug AV adapters aren't GbE-capable. (I have rarely had more than one piece of gear simultaneously active at each node.)

AT A GLANCE

- Power-line networking's long-standing shortcomings encouraged me to assess the validity of 802.11n's performance promises.
- Windows Media Center's reliance on MPEG-2 results in substantially higher bit rates than necessary with more modern video codecs.
- A single-channel, 5-GHz 802.11n topology couldn't deliver sufficient sustained bandwidth, due in part to my intra-LAN streaming scheme.
- A dual-channel approach improved the wireless network's peak speeds but also resulted in mysterious packet-dropping glitches.
- Gigle Semiconductor's Mediaxtream power-line-networking technology doesn't yet deliver on its GbE (gigabit-Ethernet) claims, at least in my setup.

Although all of these client devices support at least one IEEE 802.11 flavor, I wanted to simplify and optimize the performance of my migration from HomePlug AV to 802.11n. I therefore planned to swap out each of the 10/100-Mbit wired-Ethernet switches for D-Link's DAP-1522, which combines a four-port GbE-capable switch and a dual-band, dual-stream 802.11n subsystem. The Insteon controller's single-client network node would require the use of only Linksys' simpler WGA600N or WET610N bridge devices.

BAND, ENCRYPTION CHOICES

A bit of upfront research fortunately saved me some later hassles. An online review of the D-Link DAP-1522 revealed that the unit had subpar performance in 802.11n's 2.4-GHz band versus the 5-GHz alternative and that it performed worse with WEP (Wired Equivalent Privacy) and WPA (Wireless Protected Access)-plus-TKIP (Temporary Key Integrity Protocol) encryption than with the more modern WPA/AES (Advanced Encryption Standard) combination (**Reference 3**). My Apple 802.11n-cognizant router also supports bonding together two wireless channels to boost the resultant bandwidth capability only in the 5-GHz ISM (industrial/scientific/medical) band. I was motivated to

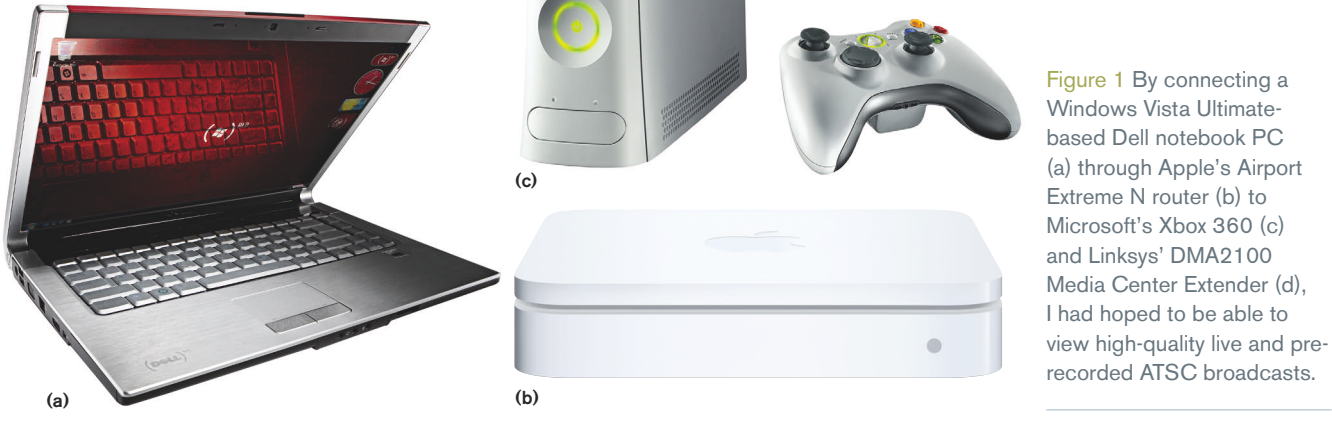

Figure 1 By connecting a Windows Vista Ultimate-based Dell notebook PC (a) through Apple's Airport Extreme N router (b) to Microsoft's Xbox 360 (c) and Linksys' DMA2100 Media Center Extender (d), I had hoped to be able to view high-quality live and prerecorded ATSC broadcasts.

lift the high-performance section of my Wi-Fi LAN above the already-cluttered 2.4-GHz spectrum that microwave ovens, wireless-surround-sound-speaker-transmitter/receiver combos, neighbors' access-point signals, and other broadcasters populate. Also, in my diminutive open-air geodesic dome, the 2.4- versus 5-GHz-range discrepancy was not a practical concern.

The 5-GHz ISM band is comparatively crystal-clear in my rural locale. I therefore bound the Dell XPS M1330-to-router-to-D-Link DAP-1522 chain by means of a WPA-plus-AES-encrypt-

REVISITING POWER LINE

I have so far been unable to dispense with HomePlug AV in my LAN, so I've spent some time determining whether I could improve the technology's robustness. In a sense, the power-line approach has an inherent advantage: The adapters can directly transfer data between them over the power grid with minimal router interaction. However, surge protectors and noise filters are equal parts curses and blessings for power-line networking. You can't plug an Ethernet-to-power-line adapter into them because the filter circuitry siphons off the networking data stream that's multiplexed on the ac-waveform carrier signal. Their omission from power-grid noise sources is equally debilitating to the power-line network, however.

Some signal attenuators are fairly obvious, notably motor-based products, such as stand-alone fans, heaters, air conditioners, refrigerators, vacuum cleaners, hair dryers, and the like. Other more obscure noise sources include the switching power supplies in ac/dc converters and battery chargers. Companies such as Cal-Lab sell specialized hardware that combines an unfiltered outlet for the power-line-networking adapter and a filtered connection, which also protects against lightning and other power surges, for noise-generating gear (Figure A). Similarly, Intellon recently sent me two PowerNet 200 HomePlug AV adapters that Monster Cable sells; the adapters integrate two filtered and protected power outlets.

Gigle Semiconductor recently partnered with Belkin to unveil Belkin's F5D4076 gigabit power-line-networking adapters, which Gigle based on its GGL541 IC. The GGL541 supports both HomePlug AV, which operates in the 2- to 28-MHz band, and Gigle's proprietary Mediaxtream technology, which uses the 50- to 300-MHz band. Like 5-GHz Wi-Fi versus 2.4-GHz 802.11, Mediaxtream's higher frequency delivers potentially higher performance. Indicative of this promise, the F5D4076 includes a 1-GbE (gigabit-Ethernet) transceiver, whereas consumer HomePlug AV adapters belie their "200-Mbps" marketing claims by embedding only 10/100-Mbps PHY (physical-layer) interfaces.

However, again as with 5-GHz versus 2.4-GHz wireless, Mediaxtream has notably shorter usable range than does HomePlug AV. The initial production firmware in the Belkin adapters selects either HomePlug AV or Mediaxtream mode, depending on the power-grid characteristics that the ICs' embedded DSPs determine at power-up. As you can see from the Network Performance Tuner plot, the adapters have selected HomePlug AV mode in my setup. In fact, they run slightly slower than my Netgear HomePlug AV-dedicated hardware, even at a two-node deficit. Gigle is working on firmware improvements, both to increase the number of supported nodes and to bond the HomePlug AV and Mediaxtream channels together rather than using a more elementary either-not-both approach.

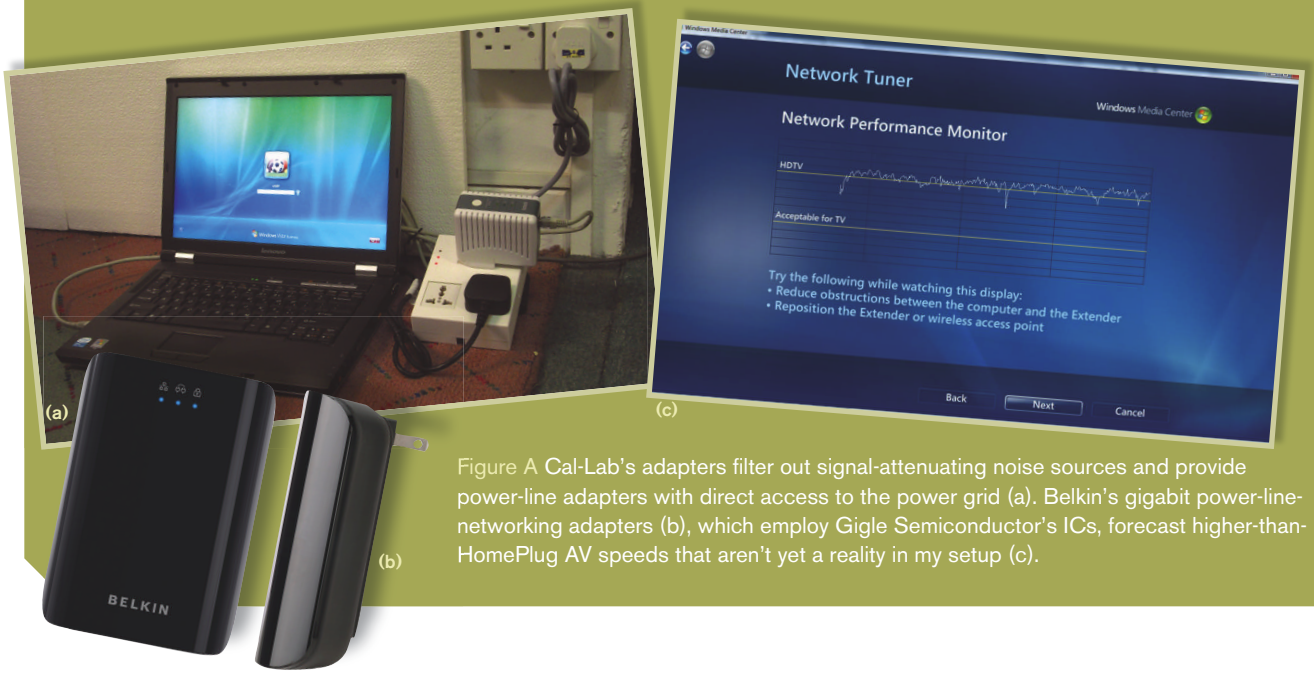

Figure A Cal-Lab's adapters filter out signal-attenuating noise sources and provide power-line adapters with direct access to the power grid (a). Belkin's gigabit power-line-networking adapters (b), which employ Gigle Semiconductor's ICs, forecast higher-than-HomePlug AV speeds that aren't yet a reality in my setup (c).

Figure 2 Netgear's HomePlug AV adapters (a) were the initial means by which I connected various network nodes. Power-line networking's unreliability encouraged me to evaluate the 802.11n wireless alternative, enabled by D-Link's integrated switch plus access point (b), and two generations' worth of Linksys single-client bridges (c and d).

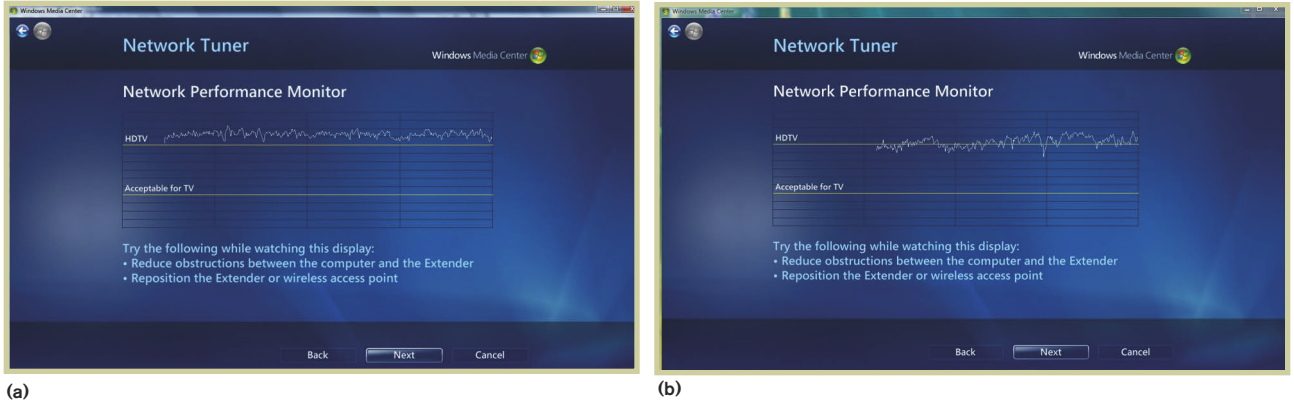

Figure 3 A HomePlug AV-only topology might be somewhat flaky (a), but it delivers faster streaming video than a single-channel, bonded 802.11n alternative (b).

ed 802.11n wireless spur running on Channel 149—that is, 5.745 GHz. I also tried other 5-GHz-band channels during later debugging. Initial streaming-attempt results were horrible: Windows Media Center's Network Performance Tuner utility measured best-case speeds of less than 8 Mbps. Nothing I tried improved the situation until I noticed that the laptop's Broadcom-sourced Windows Vista-driver suite carried a publication date of December 2006. Dell was still shipping it in the system that I had purchased in July 2008 (**Reference 4**)!

A visit to Dell's support Web site unearthed a slightly newer driver dated October 2007 available for downloading. At press time, Dell had made available no newer version. Installing the 2007-dated driver notably improved the average speed of the wireless link but still not to a level at which it would reliably sustain streaming of a high-definition recording from the laptop to the game console (**Figure 3**). The Network Performance Tuner generates plots that define 22 Mbps as the requisite HDTV (high-definition-television)-bandwidth threshold and 8 Mbps as the acceptable-for-TV bandwidth. My 802.11n network's bandwidth capability also woefully undershot the 150-Mbps, single-stream and 300-Mbps, dual-stream claims of the technology's backers.

Pondering the problem uncovered a possible partial explanation, which several Wi-Fi-silicon-vendor representatives later confirmed. If my DSL (digital-subscriber-line) connection were capable of 20-Mbps sustained speeds (it isn't), and if I were watching a 20-Mbps video from the Internet (I can't), the incoming data would enter the LAN through the router's wired-Ethernet WAN port, and it would then stream to the Xbox 360 over Wi-Fi. However, my video-streaming setup was intra-LAN in nature, and I was therefore using one Wi-Fi channel for two simultaneously transmitting, 20-Mbps data streams: one from the laptop to the router and another from the router to the game console. Even though that 802.11n channel had a 40-MHz-wide bonded-spectrum footprint, it was still insufficient for shouldering the entire bit load. Temporarily disabling the XPS M1330's 802.11n transceiver and instead connecting its Ethernet port to a Linksys WGA600N bridge yielded no improvement and

Figure 4 The Apple router's lack of simultaneous dual-band support necessitates the inclusion of an 802.11g access point for legacy LAN clients (a). Netgear's WNHDE111 created a second 5-GHz bonded 802.11n channel (b), and I also used Apple's Airport Express N access point during debugging (c). Regardless of their frequency bands, the simultaneous operation of two 802.11n signals resulted in quality-degrading packet-loss glitches (d), which did not occur when part of the source-to-destination span relied on HomePlug AV (e) or Category 5e wired Ethernet (f). Diagnostics reports point to the router's switch as the culprit (g).

proved that the laptop's wireless subsystem wasn't the weak link.

DUAL CHANNELS=GLITCHES

Apple's latest routers and router-plus-hard-disk-drive products, Time Capsules, can single-handedly support simultaneous 2.4- and 5-GHz wireless networks, but my second-generation Airport Extreme N router is single-band, operating at either 2.4 or 5 GHz but not both at once. I had therefore already attached a WEP-encrypted Belkin F5D7130 access point to it for use with legacy 802.11g devices (**Figure 4**). Given the initial subpar results for single-channel 802.11n, I further expanded my access-point topology with Netgear's 5-GHz-only WNHDE111 device, thereby creating an additional 802.11n bonded beacon. This second signal, on 5.18-GHz Channel 36, does not overlap and is spectrally as far away as possible from the Apple router's 5.745-GHz Channel 149 signal.

My initial approach involved streaming from the Dell laptop to the WNHDE111, from there to the router over Category 5e cable, and from the router to the DAP-1522 over the router's built-in 802.11n facilities. Although the average bandwidth of the dual-channel, 5-GHz approach was notably higher than with its single-channel predecessor, as-yet-unseen glitches randomly emerged. From the Network Performance Tuner plot, you can see how significantly they impeded bandwidth; they were also as much as 5 seconds wide, and they therefore created egregious degradations in image and sound quality.

I ruled out ambient interference as their cause by shutting off every other potential wireless beacon regardless of its transmitting frequency. I followed with an intensive sweep of the ISM spectrum to confirm an absence of noise from neighbors' electronics. I tried out a variety of 5-GHz bonded-channel combinations to confirm noninteraction between them, and I also reversed the datapath through the network, disabled SSID (service-set-identifier) broadcasts, attempted streaming between 5- and 2.4-GHz 802.11n variants, and even shut off the router's Wi-Fi system and instead relied on external access points.

As before, I tried disabling the laptop's Wi-Fi transceiver, instead using the WGA600N bridge adapter. To rule out the DAP-1522, I also streamed to a Linksys DMA2100 Media Center Extender, which contains a built-in 802.11n subsystem. Nothing helped. Time and again, I encountered glitches only when I was streaming data through the Apple router's internal switch between two internal or external 802.11 access points. Conversely, the glitches disappeared when the laptop-to-console span consisted of an 802.11n-plus-Category 5 or 802.11n-plus-HomePlug AV hybrid combination. This hybrid topology represents my current workaround.

Neither Apple's Airport Extreme N router nor its Airport Express N access point, which I also used in debugging, provided the statistical reporting necessary to determine which part of the LAN chain was generating the glitches. Fortunately, I had access both to diagnostics utilities for the laptop and to status screens on the WNHDE111 access point, which revealed that the wireless portions of the topology were robust. Conversely, they bolstered my belief that the packet drops were occurring within the router's switch subsystem. Unfortunately, my attempts to contact Apple bore no fruit, and a recently released router firmware update did not improve performance. Key semiconductor suppliers to the Airport Extreme N design, Atheros and Broadcom, have also not commented on a possible root cause, perhaps so that they won't anger their customer.

See the "Transporting high-def video broadcasts" posts at www.edn.com/briansbrain for supplemental information on this article's topics.

Go to www.edn.com/090820cs and click on Feedback Loop to post a comment on this article.

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

THE SAGA CONTINUES

As time and personal-network bandwidth allow, I'll be replacing the Airport Extreme N router with other potential LAN-controller candidates: Apple's second-generation Time Capsule, which the company based on a third-generation router design; D-Link's DIR-825; Linksys' WRT600N; and Netgear's WNDR3300, WNR3500, and WNDR3700. All of these routers, except for the WNR3500, can operate simultaneously at 2.4 and 5 GHz, enabling me to retire the Belkin F5D7130 802.11g access point.

Ordinarily, the Apple Time Capsule would be the only feasible alternative router candidate because I rely on Apple's Mac OS 10.5 Time Machine feature for system backup—historically, to an external hard-disk drive that I tethered to the Airport Extreme N router over USB (Universal Serial Bus). However, I've recently moved my backups to a Netgear/Infrant ReadyNAS NV+ network-storage device, which supports Time Machine protocols through a recent firmware update (**Reference 5**). I hope that at least one of these alternative routers' switches exhibits no baffling dropped-packet glitches between two 802.11n transceivers' channels. And, if single-channel 802.11n capabilities end up being sufficient, perhaps I can retire the WNHDE111 802.11n access point, too. EDN

REFERENCES

1 Dipert, Brian, "Thin air: ATSC reception isn't always easy," *EDN*, May 14, 2009, pg 20, www.edn.com/article/CA6656302.

2 Dipert, Brian, "802.11n: a complicated spec to be is about to become even more messy," *EDN*, April 15, 2009, www.edn.com/blog/400000040/post/1970043397.html.

3 Higgins, Tim, "D-Link DAP-1522 Review: Dual-band Draft 11n for the Masses?" SmallNetBuilder, June 2, 2008, www.smallnetbuilder.com/content/view/30455/96.

4 Dipert, Brian, "CES: IEEE 802.11n: the vendor-neutering of a once-promising standard," *EDN*, Jan 8, 2007, www.edn.com/article/CA6405195.

5 Dipert, Brian, "Accelerating consumers' NAS adoptions: assessing your product options," *EDN*, June 25, 2009, pg 30, www.edn.com/article/CA6666226.

ACKNOWLEDGMENT

I'm particularly grateful to Netgear for supporting this project by providing a preproduction version of the WNDR3700 for testing purposes.

FOR MORE INFORMATION

Apple
www.apple.com

Atheros Communications
www.atheros.com

Belkin International
www.belkin.com

Broadcom
www.broadcom.com

Cal-Lab
www.cal-lab.com

Dell
www.dell.com

D-Link
www.dlink.com

Gigle Semiconductor
www.gigle.biz

Intellon
www.intellon.com

Linksys
www.linksysbycisco.com

Microsoft
www.microsoft.com

Monster Cable Products
www.monstercable.com

Netgear
www.netgear.com

Roku
www.roku.com

Sony
www.sony.com

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

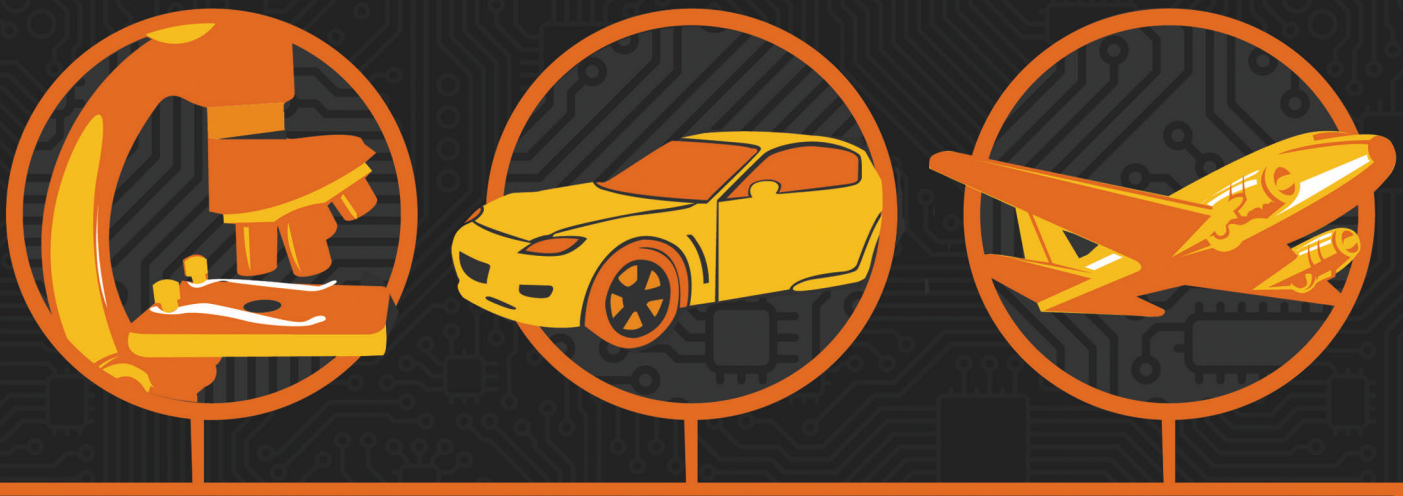

NO MATTER WHAT THE IDEA

YOUR PCB PROTOTYPES SHOULD BE THE EASY PART

QUOTE & ORDER PCBS ONLINE AT **WWW.SUNSTONE.COM** OR CALL **1-800-228-8198**

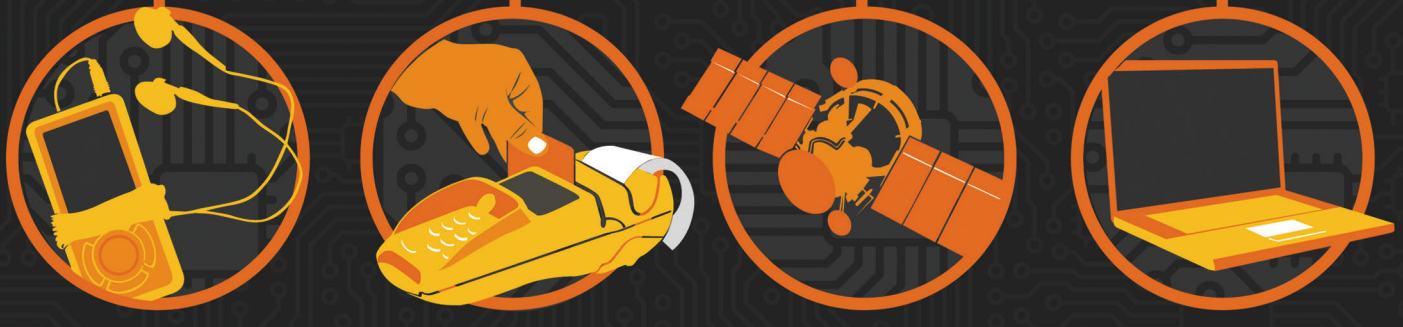

SUNSTONE CIRCUITS

THE EASIEST PCB COMPANY TO DO BUSINESS WITH

ValueProto™

PCBexpress®

Full Feature

Sunstone Circuits® pioneered the online ordering of printed circuit boards and is the leading PCB solutions provider with more than 35 years of experience in delivering quality prototypes and engineering software. With this knowledge and experience, Sunstone is dedicated to improving the PCB prototyping process from quote to delivery (Q2D®).

Did You Know? Sunstone Offers:

- Controlled impedance testing
- Fine lines and spacing [.003]
- RoHS compliant finishes
- Free 25-point design review
- Free shipping & no NRE's
- Flex / Rigid Flex Boards
- Online Quote & Order
- PCB123® design software
- RF / Exotic Materials
- Over 99% on-time or early delivery
- Best PCBs in the industry
- Live customer support 24/7/365

BY DREW E WINGARD, PhD • SONICS INC

Addressing interleaved multichannel memory challenges

INTERLEAVING ADDRESSES IN MULTIPLE DRAM CHANNELS CAN GREATLY IMPROVE MEMORY BANDWIDTH, BUT IT IS NOT A TRIVIAL TASK.

Achieving the total external-memory-bandwidth requirements of consumer SOCs (systems on chips) at acceptable costs gets more difficult with each generation of electronic systems and the DRAM technologies that serve them. SOC designers must optimize for total DRAM efficiency to enable the system to reach the highest performance and to minimize DRAM-subsystem costs. As SOCs move to more advanced DRAM technologies, such as DDR3, that have larger minimum burst lengths, challenges arise in delivering optimal memory throughput. Using multichannel DRAM subsystems helps maximize efficiency by ensuring that the DRAM bursts are smaller than the processor- and I/O-interface accesses.

WHY MULTICHANNEL?

Discussions of DRAMs introduce terms such as "banks," "ranks," and "channels" (**Reference 1**). Each term describes a mechanism for arranging multiple arrays of DRAM cells, either within the same chip or across multiple chips, to increase memory density, memory performance, or both. The multibank architecture of current SDRAM devices enables SOC DRAM controllers to improve performance by exploiting parallelism across the banks inside a DRAM chip to hide the page-closing and -opening penalties of the internal DRAM cell arrays. You could employ the same principle with multiple chips, but few consumer SOCs support multiple ranks of DRAM devices, in which multiple DRAM chips connect to the same data signals to increase the total supported memory size, because few consumer-electronics systems need that much memory.

The substantial increases in performance in consumer SOCs require both more total on-chip processing and substantial increases in DRAM bandwidth. Although the data-pin bandwidth of DDR SDRAM devices has improved over time, these improvements have not kept pace with the requirements of several key consumer-SOC markets, such as HDTV (high-definition television). As a result, the total number of data pins necessary to satisfy DRAM-bandwidth needs is growing for such SOCs.

The combination of increasing minimum burst length associated with DDR3 and the increasing data-pin count leads to substantial increases in the minimum burst size of single-channel DRAM systems. When these bursts become larger than the data objects that the initiators on the SOC are accessing, effective throughput declines as the DRAM transfers become less efficient.

One approach to this challenge is to further increase the available DRAM bandwidth to compensate for the loss of efficiency. However, the bandwidth increase must come without further increase to the DRAM burst size, or you lose even more efficiency. A better approach is to introduce multiple channels.

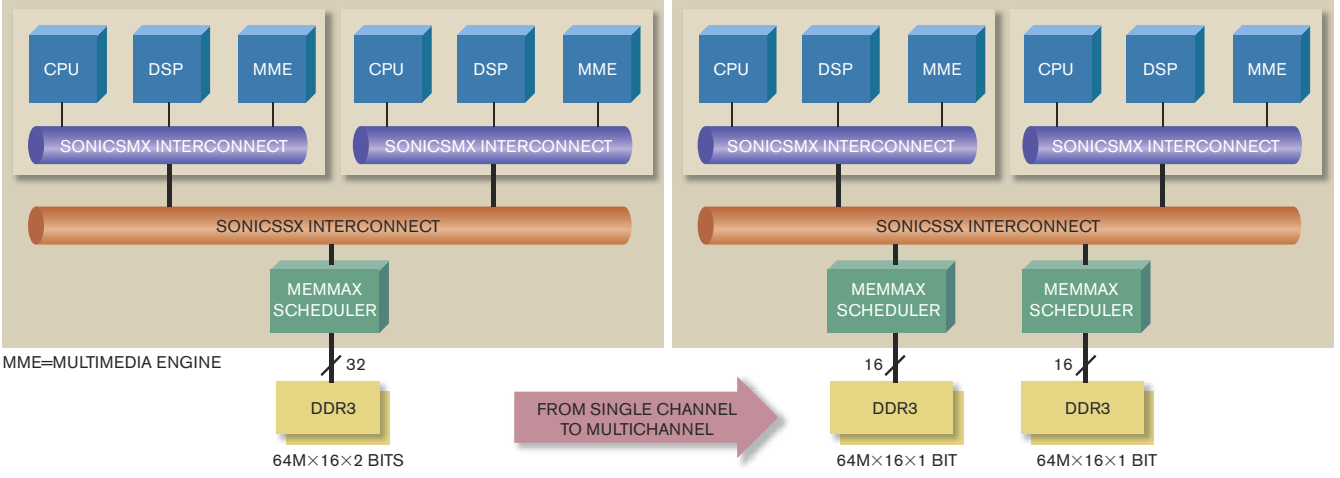

Figure 1 Moving from a single DRAM channel to two channels can increase peak bandwidth.

The key benefit of a multichannel DRAM system is an improvement in access efficiency due to shorter bursts that more closely match the size of the data types transferring to memory. Note that the DRAM bursts are smaller but not shorter because only the word is smaller. The prefetch degree of the SDRAM type in use still largely determines the burst length.

THE KEY BENEFIT OF A MULTICHANNEL DRAM SYSTEM IS AN IMPROVEMENT IN ACCESS EFFICIENCY.

A second benefit of the multichannel system is a likely reduction in efficiency loss due to page-closing and -opening delays because an N-way multichannel system has N times as many banks and can therefore manage N times more open pages. However, the fact that the pages are only 1/Nth as large mitigates this benefit, so some transactions that might have found their page to be open in a single-channel system—due to a previous transaction's accessing memory locations at a similar address, for example—can find a closed page instead in a multichannel system. In a consumer SOC with many initiators, enough data flows are often waiting for access to DRAM that the benefit of more open pages outweighs the smaller page.

To quantify the efficiency benefits available from a multichannel memory system, consider the migration of a production HDTV SOC from a DDR2 SDRAM baseline to a next-generation design based on DDR3 devices. A static analysis of the memory traffic from the original design—currently in high-volume production—assigns the memory efficiency to be 100% as a point of reference. For the next generation of this SOC, compare two memory-system architectures: a single-channel DDR3 system with a word size of 32 bits and a two-channel system in which each channel has a word size of 16 bits (**Figure 1**). Note that both configurations offer the same peak DRAM bandwidth and use the same number and type of DRAMs.

TABLE 1 MEMORY-SYSTEM EFFICIENCY

	DDR2	DDR3	DDR3
Channels	One	One	Two
Data word width (bits)	32	32	16
Effective bandwidth (%)	100	84	100

You might expect the single-channel DDR3 design to have lower memory efficiency because DDR3 devices use a prefetch-8 architecture and thus have a minimum efficient burst length of 8 words, whereas the reference system has the 4-word burst characteristics of DDR2. When you consider the efficiency loss due to both minimum DRAM bursts that are larger than the SOC traffic patterns and address-alignment problems that also result from larger bursts, you can see that the single-channel DDR3 system loses a substantial amount of efficiency. In contrast, the dual-channel system has longer bursts that are half as wide, so the basic burst size and alignment do not change. Therefore, the dual-channel DDR3 system is as efficient as the single-channel DDR2 system of the same size. **Table 1** shows the efficiency difference, which directly translates into usable memory-system bandwidth.

Note that this static analysis ignores any of the bandwidth increases that are available due to the higher operating frequencies available with DDR3 SDRAMs. It compares only the relative efficiency of the memory usage. It also ignores the likelihood that the DDR3 design will have higher overall performance requirements and, thus, higher total external-memory-bandwidth requirements, which are likely to outstrip the frequency benefits of the DDR3 transition. More important, the analysis assumes that you can schedule the memory-system traffic for maximum efficiency in all cases and balance it evenly across the dual-channel configuration.

Most scenarios require you to balance traffic across the channels to deliver multichannel benefits. For example, applications that require high memory efficiency normally have a fair amount of queuing to cover the latency between memory requests and responses, including the arbitration delays on the SOC among the initiators. If a lightly loaded channel runs out of requests to service while the other channels are busy, then that channel's throughput and the efficiency benefits of the multichannel approach decrease. If the channels act as independent regions in the memory map, the SOC architect, accelerator designer, firmware developer, or application developer should carefully allocate data structures in memory so that the initiators' access to those structures is well-balanced over periods as short as a few microseconds. It is difficult to manage this task with a dual-channel memory system and nearly impossible with the four-channel systems of the future.

Another challenge with this static load-balancing approach is that transaction-ordering requirements often prevent a single initiator from sharing memory bandwidth from multiple channels at once. This problem arises from the nature of the initiators' communication protocols, which specify that request and response order should match, and from the fact that DRAM channels have significant latency variations. So the response to a first request to a first channel may likely be unready for delivery until after the response from a second channel is ready. Using flow control to hold off the second channel is typically unacceptable because it would cause the DRAM to stop servicing requests while waiting, which reduces DRAM throughput. The result is that static load balancing generally requires that most initiators communicate only with a single channel, which makes sharing and load balancing substantially more difficult.

WHY IS INTERLEAVING THE ANSWER?

The ideal approach for load balancing a multichannel DRAM system would be one that achieves excellent balancing of traffic, is largely independent of the number of channels, and requires no extra work in the design of either the initiators or the software that controls them. Rather than treating the channels as independent memory regions with the resulting load-balancing challenges, interleaving the channels in the address space enables them to appear as a single, logical memory region and offers the promise of achieving all of these goals.

To understand why channel interleaving can achieve automatic load balancing, it is important to understand the memory-access behavior of the initiators in the SOC. The initiators access data structures in DRAM, so the type of stored data or the processing algorithm in use with the data determines the expected access patterns to those data structures. Streaming accelerators, communication processors, cameras, displays, and

CHOOSING THE INTERLEAVING BOUNDARIES IS A KEY TASK IN ACHIEVING GOOD LOAD BALANCE.

I/O interfaces all normally manage large data buffers in DRAM. They typically access those data structures using midsized to long burst transactions, with each burst transaction starting at the sequential address following that of the previous transaction, until they reach the end of their buffer. In contrast, a CPU accesses a variety of data structures in memory, from small control structures through large buffers. Because most CPUs in consumer SOCs have internal caches, the dominant accesses to DRAM are cache-line-sized, and the principles of spatial and temporal locality teach that many of the accesses in a time window have similar addresses.

For many consumer SOCs, the video decoder offers a particularly challenging case. The video decoder accesses several frames of uncompressed video images, each of which may comprise 2 million pixels of storage, and an incoming stream of compressed data that tells the decoder how to decode the next frame. The incoming stream frequently tells the decoder to fetch an arbitrary macroblock from the frame storage, with each macroblock comprising a 2-D set of pixels. Because the required macroblock can start at any pixel address, it is unlikely that the macroblock will fit nicely into a few DRAM bursts. Furthermore, the total required macroblock bandwidth is high enough that mapping the macroblocks into efficient DRAM transfers is essential.

Earlier interleaving often focused on improving memory bandwidth by accessing multiple physical DRAMs in an interleaved manner to improve pin bandwidth. These systems relied on spreading a burst transaction across several DRAMs. This approach differs greatly from one in which you intentionally create smaller channels of memory to reduce, rather than increase, the burst size. Some of the latest PCs, including servers and desktops, also use multiple interleaved channels. These channels are normally interleaved at CPU cache-line boundaries in an attempt to best exploit the spatial locality of computing applications. These boundaries are finer than optimal for consumer SOCs, in which the access patterns tend to have more regularity.

The choice of interleaving boundaries can greatly affect the load balancing. Some of the accelerators access data structures in a fairly predictable pattern, in which the address of the next memory request is spaced a fixed distance away from the previous one. These strided accesses reduce the channel balance when the stride value is a multiple of the interleaving size because consecutive accesses may map to the same channel. Because there can be a close correlation between memory data structures and strided access patterns, the designer can optimize overall throughput by selecting different interleaving boundaries for different regions of memory based on the data structures they each store.

The challenge of balancing the traffic loading increases with the number of channels. The interleaving approach can balance the traffic as long as most initiators regularly access each of the channels—in other words, as long as the number of channels times the size of the interleaving boundary is smaller than the range of addresses the initiators are accessing. Thus, optimizing the interleaving boundaries is related to the number of channels. Choosing the interleaving boundaries is a key task in achieving good load balance, and you may achieve better balance when you divide the DRAM address space into several subregions with different interleaving boundaries. Additionally, because the best boundary choice can depend on the data structure and access patterns and because these parameters can both change based on the operating mode of the SOC, it is sometimes valuable to change these boundaries when the mode changes.

A major benefit of interleaving the channels is in isolating the initiators from the channel configuration, which enables simpler design and much greater reuse of the processors, accelerators, and software for the SOC. However, this isolation ensures that the initiators are not channel-aware, thus increasing the importance of maintaining high throughput for initiators that are accessing several channels at once. Computer systems maintain throughput by building reordering buffers near the initiators,

so channels that service requests sooner than an initiator is ready to receive them must have a place to store their responses. However, the large number of initiators in a consumer SOC and the growing depth of the DRAM-access pipeline mean that the total amount of required storage for this reorder buffering would be too large and expensive in the markets these SOCs serve.

It is equally important to ensure that the SOC memory controllers have a great degree of flexibility for scheduling traffic to the DRAM channels to ensure the highest efficiency and throughput. The interleaving system should therefore limit neither the number of transactions that can be outstanding to channels nor the controller's ability to schedule the transactions that it has received.

One approach that could address many of these challenges would be to manage the interleaving in the memory controller itself. This approach has the advantages that it localizes the information about the number of channels and interleaving boundaries into one location on the SOC, minimizing the disruption to other architectures and hardware, and allows the use of shared buffering in the controller to manage ordering and allow many transactions to be outstanding across the channels. However, such an architecture requires most of the system communications to pass through a single point on the SOC, which is likely to create a performance bottleneck in both wire routing and memory efficiency. This memory-efficiency loss can result from the same access-granularity problem that happens in a single-channel DRAM system: that the internal interface carrying the memory traffic may become so wide—to carry the total DRAM bandwidth for the SOC—that a DRAM burst for a single channel may not pack efficiently into the internal-interface word.

INTERLEAVED MULTICHANNEL TECHNOLOGY MANAGES FLEXIBLE INTERLEAVING BOUNDARIES AMONG MULTIPLE DRAM CHANNELS IN THE INTERCONNECT.

The SonicsSX interconnect from Sonics (www.sonicsinc.com) uses another approach, which employs IMT (Interleaved Multichannel Technology). IMT manages flexible interleaving boundaries among multiple DRAM channels in the interconnect, rather than in the RAM controller, providing the benefits of automatic load balancing and high throughput without creating performance bottlenecks or requiring reordering buffers. You measure the DRAM efficiency as the fraction of the DRAM clock cycles during which a useful data word is transferring to or from DRAM. Although achieving DRAM efficiency of 60% is relatively straightforward for most designs, targeting efficiencies of 75 to 90% is more challenging and normally requires substantial analysis and optimization during the SOC-design phases.

Choosing the right multichannel architecture involves selecting the proper number of DRAM channels, allocating the DRAM address space across one or more multichannel memory regions, and selecting the interleaving characteristics for each region. The SOC designer normally chooses the minimum number of channels that provide the required memory-system efficiency and throughput. This configuration generally minimizes system costs because it results in the minimum DRAM costs and reduces the number of DRAM-related control and address pins on the SOC.

Single-channel SOCs normally treat DRAMs as single pools of address space that all of the initiators share. The software that executes on the host CPU during booting allocates some of this DRAM to specific uses and initiators, and the operating system dynamically allocates the rest of the DRAM space. In an interleaved multichannel system, the strided access patterns of some initiators can cause channel imbalances with certain interleaving boundaries. When several such initiators share the memory system, the designer may wish to allocate multiple subregions in the logical DRAM address space with different interleaving boundaries that better match the access characteristics of the initiators that share each subregion.

When multiple subregions are in use, the operating system normally allocates one, and designers have optimized this

subregion for more general-purpose traffic, typically with a small interleaving boundary—perhaps at the cache-line level, as in a PC. The booting process or the device drivers normally assign the other subregions for collections of initiators with more predictable accesses. Such subregions are likely to have larger interleaving boundaries. Designers can support multiple operating modes with multiple address subregions by allocating more total DRAM address space than the DRAMs contain and allowing the subregions to alias their contents onto each other in memory. This approach relies on careful allocation of data structures in the subregions to ensure that you do not simultaneously allocate a given area in physical DRAM to multiple data structures in different subregions.

Designers normally map each subregion onto all of the physical DRAM channels. But aggressive power-management schemes, in which some of the channels may power down in some operating modes, provide an example in which a designer may populate and power some subregions with only enough channels to deliver the required memory throughput for those modes.

Once designers know the number of channels in a subregion and the initiator-access characteristics, they can choose the interleaving boundary for the subregion. The choice of an interleaving boundary is critical for achieving good load balance among the traffic that targets that subregion.

ESTIMATING RESULTS

SOC designers do performance estimation and analysis using spreadsheets, cycle-accurate simulation in SystemC, and RTL (register-transfer-level) simulation. Although each technique has its advantages and disadvantages, many designers apply several of the techniques to a design. It is therefore important to have a consistent vocabulary for performance-oriented characteristics of the design with instrumentation to measure or calculate those characteristics. Key measurement parameters include throughputs, latencies, and DRAM channel efficiencies, which you measure in the timing domain of the initiators that generated the traffic. It is useful to have access to these performance results both as aggregated summary data and at the granularity of individual transactions.

SOC designers must compare these performance results with the quality-of-service requirements of the system. Knowing the granularity of individual-transaction results helps SOC designers debug simulation results to understand why the performance may be different from what they expect. Tooling that helps designers track transactions as they propagate from the initiator, across the interconnect, through the memory scheduler, and into the DRAMs increases visibility into challenging performance-debugging situations. Once designers gain insight into the reasons for performance results, they can use it to modify the SOC and memory-system configuration to further optimize performance.

The substantial increases in DRAM-bandwidth requirements of consumer SOCs and the prefetch architecture of DDR SDRAMs have caused single-channel DRAM systems to lose substantial efficiency as bursts become larger than SOC transactions. This scenario leads designers to select multichannel DRAM systems. However, only interleaved multichannel systems that support multiple subregions with different interleaving boundaries deliver the automatic load balancing and hardware and software transparency necessary for consumer SOCs. EDN

> ⊞ Go to www.edn.com/ms4326 and click on Feedback Loop to post a comment on this article.
>
> ⊞ For more technical articles, go to www.edn.com/features.

REFERENCE

1 Wingard, Drew, "DRAM technology for SOC designers and–maybe–their customers," *EDN*, Aug 6, 2009, pg 34, www.edn.com/article/CA6674038.

AUTHOR'S BIOGRAPHY

Drew E Wingard, PhD, co-founded Sonics in September 1996 and is currently chief technical officer and secretary. He received a bachelor's degree in electrical engineering from the University of Texas—Austin and master's and doctorate degrees in electrical engineering from Stanford University (Stanford, CA).

Allied

Your R&D and production solutions provider

ALLIED ELECTRONICS®

Small Volume Order

Fulfilling small orders for prototyping is what we do best

Breadth of Product

Over 1.2 million products from a wide range of technologies

Special Sourcing

Contact your Allied Sales Rep for hard-to-find electrocomponent products

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

alliedelec.com 1.800.433.5700

XILINX

RISE ABOVE

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

EDITED BY MARTIN ROWE
AND FRAN GRANVILLE

designideas

READERS SOLVE DESIGN PROBLEMS

Triac tester allows for manual or automatic operation

Abel Raynus, Armatron International, Malden, MA

Triacs are bidirectional ac switches that can control loads with currents as high as 25A rms at voltages as high as 600V. They find wide use in motor-speed, heater, and incandescent-lamp controls. Logic triacs are especially attractive for microcontroller-driven devices. You can activate a triac directly from microcontroller-output ports because of the triac's trigger current of only 3 to 10 mA. As with any electronic device, triacs can have some internal problems that you can detect before using them in a design.

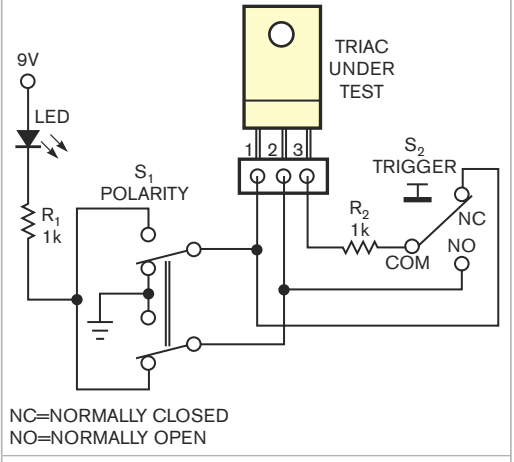

Figure 1 A triac tester uses a switch to reverse the polarity of the test signal.

Figure 1 shows a simple and inexpensive test fixture that tests the L2004F31, L2004F61, L2004L1, and L4004V6TP triacs from Littelfuse (www.littelfuse.com), but you can use it to test any other leaded triac because all the standard packages, including TO-220AB, TO-202AB, TO-251, and IPak, have the same pin layout. An IC socket provides easy insertion of a triac under test. You can also apply this idea to SMDs (surface-mount devices), provided that you can find or create an appropriate test socket. Polarity switch S_1, a DPDT (double-pole/double-throw) device, lets you check conductivity in both directions. Trigger switch S_2, a momentary SPST (single-pole/single-throw) pushbutton device, activates the triac under test by connecting the gate (Pin 3) with MT_2 (Pin 2) through resistor R_2 (**Figure 1**).

The test takes less than 5 seconds and comprises four steps (**Table 1**). An LED indicates the result of each step to the test operator. A triac is good if

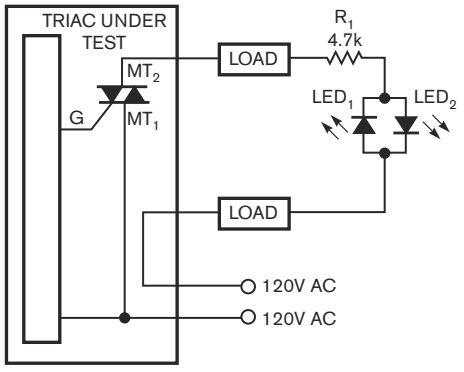

Figure 2 With a resistive load, the tester uses two LEDs to indicate pass and fail in both directions.

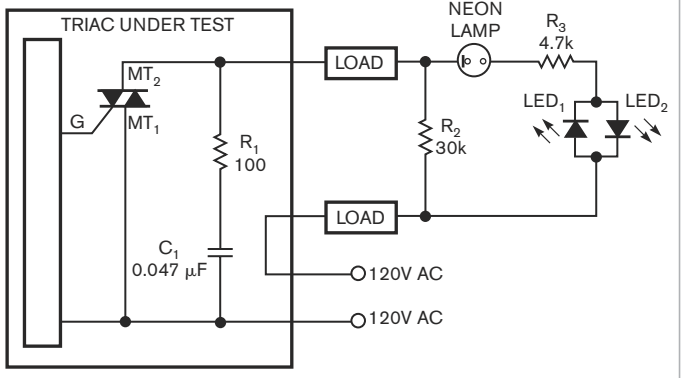

Figure 3 For an inductive load, add a neon lamp to minimize leakage current.

DIs Inside

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

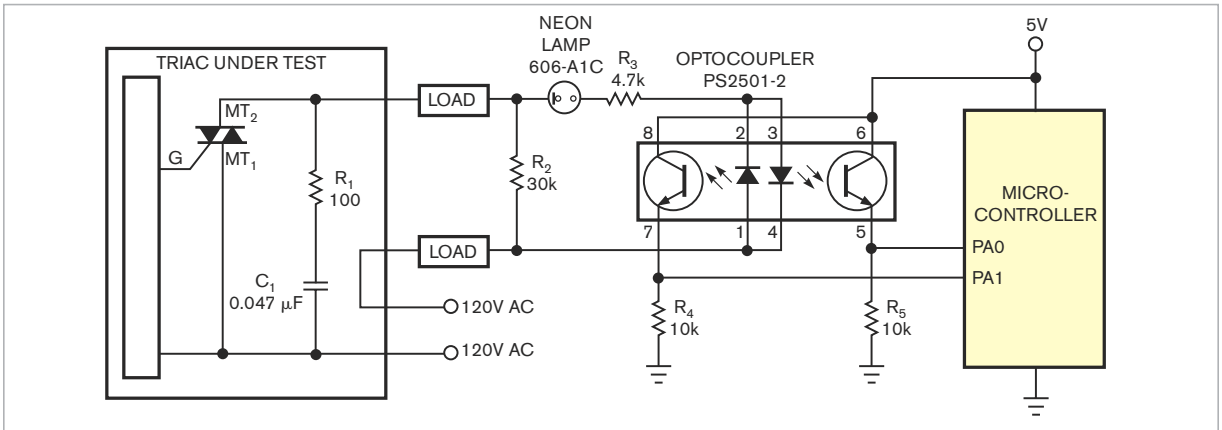

Figure 4 An optocoupler isolates the triac from ground.

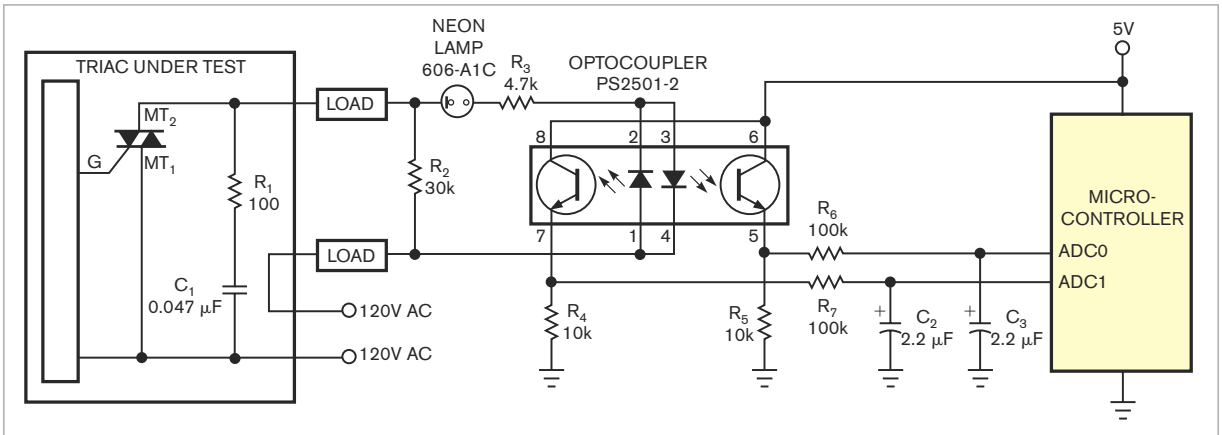

Figure 5 RC filters let you use PWM signals.

it passes all four tests. You should perform another triac test during manufacturing to ensure that there is no problem with the subassembly board and that the triac works properly. This test saves time and labor in case you detect a problem after assembling the entire product. You perform this test with the triac soldered into place on the board. You use the nominal power-supply voltage of 120/220V ac. The test should have minimal influence on the DUT and should use minimal time and labor. This test uses the triac tester in place of a load. The connection from the tester to the DUT can vary, and be sure to take some safety measures when connecting 120/220V ac.

You use a different test fixture for triacs that drive a resistive load, such as an incandescent lamp or a heater (**Figure 2**). Each LED checks conductivity in one direction. When the triac is closed, both LEDs should be off. When it is open, both LEDs should be on. In the case of an inductive load, such as a motor, use an RC snubber circuit comprising C_1 and R_1 in parallel with the triac (**Figure 3**). Unfortunately, the snubber circuit introduces a small current leakage into the test circuit even when the triac is closed. The circuit in **Figure 3** shows you how

TABLE 1 TEST FOR TRIACS

Step no.	Operations	LED status	Result
1	Insert triac under test into the socket; turn on power	Off	OK
		On	Shortage inside triac
2	Push and release trigger switch S_2	Off	Break inside triac
		Stays on	OK
		On but goes off after you release S_2	Bad "hold" function in triac
3	Move polarity switch S_1 into another position	Off	OK
		On	Shortage inside triac
4	Push and release trigger switch S_2	Off	Break inside triac
		On	OK
		On but goes off after you release S_2	Bad "hold" function in triac

to avoid this problem using resistor R_2 and a neon lamp with an ac breakdown voltage of 95V.

The indicators of the test result in **figures 1**, **2**, and **3** are LEDs. Sometimes, the triac test is part of a multitasking test system that checks other components or parameters of the whole device, which includes the triac. This test involves a sequence of measurements, and a system operator gets only one of two possible signals: pass or fail. These tests use a microcontroller-based system. Thus, all the interface signals should be in digital format: high or low.

You can also use analog signals by activating the microcontroller's ADCs. This approach is less preferable, however, because of the limited number of ADCs in low-end microcontrollers and more complicated software. Interfacing the triac under test with the microcontroller creates no problem if the triac's MT_1 pin is grounded. In most cases, MT_1 and MT_2 are isolated from the ground. When this scenario occurs, you can use an optocoupler, such as the PS2501-2 from California Eastern Laboratories (www.cel.com, **Figure 4**). It comprises two optically coupled isolators containing LEDs and NPN phototransistors with a maximum voltage of 80V.

If the triac output comprises a sequence of pulses, such as a PWM (pulse-width-modulated) signal for motor-speed or lamp-brightness control, then use a lowpass RC filter before the microcontroller's ADC inputs (**Figure 5**). The time constant of this filter, $\tau = R_6 \times C_2$, depends on the PWM signal period and duty cycle. The measurement in the chain of tests should start no earlier than $3-5\tau$. Using the microcontroller's ADC requires additional firmware. To avoid this requirement, you can compare the voltage after the filter with a reference voltage with a comparator, such as the LM393 from National Semiconductor (www.national.com), to produce a logic-high level for the microcontroller's input. **Reference 1** describes an alternative approach with minimal external components for the expense of the firmware complication. EDN

REFERENCE

1 Raynus, Abel, "Microcontroller detects pulses," *EDN*, July 24, 2008, pg 58, www.edn.com/article/CA6578137.

Handheld DMM copes with logic nanosecond-pulse-width waveforms

Marián Štofka, Slovak University of Technology, Bratislava, Slovakia

When testing sequential-logic circuits, you may find that, although the repetition frequency of a logic signal is within the range of your DMM (digital multimeter), you can't measure it. The displayed frequency value is either dubious or chaotically changing in time. The DMM may also behave as if there were no signal. Any of these undesired states might appear when the duty cycle of the measured waveform is either close to zero or is approaching one—in other words, when the width of a pulse—high or low—is much narrower than the repetition period of these pulses. This problem occurs because you can't expect a DMM with an upper frequency limit of perhaps 200 kHz to measure 100-nsec-wide pulses, even if the repetition rate of these pulses is well below the upper limit of the DMM's frequency range—perhaps just 5 kHz. For a rough estimation of bandwidth for measuring at a pulse width of 100 nsec, consider this pulse to be a half-period of a square-wave signal. Use the following **equation** to calculate the required bandwidth:

$$B \approx \frac{1}{2T_W} = \frac{1}{2 \times 10^{-7}} = 5\ \text{MHz}.$$

This frequency is well beyond the bandwidth of most DMMs. The second cause of failing to measure the repetition rate of logic waveforms with too-low or too-high duty cycles lies in the internal ac coupling of the DMMs during frequency measuring. Due to this coupling, the decision threshold of an internal comparator, which you derive from the mean value of the measured waveform, is close to either the low or the high level of this waveform. In the case of narrow pulses, the operation of the internal comparator becomes ambiguous, and any noise in the measured waveform or that the comparator itself generates may cause an error.

You can address the problem by placing a binary divider between the source of a logic signal and the DMM. The binary divider comprises IC_1, a positive-edge-triggered, D-type flip-flop (**Figure 1**). The supply pin of IC_1 connects to the supply terminal of the tested logic

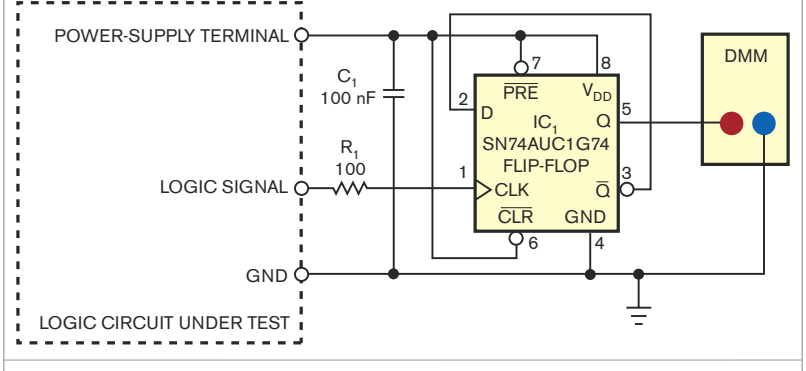

Figure 1 A binary divider turns low- or high-duty-cycle waveforms into square wave so that you can measure their frequencies.

circuit. Therefore, you can run the logic at any industry-standard supply voltage of 1.2, 1.5, 1.8, or 2.5V. In testing 3.3V logic, use an external 2.5V source to supply IC_1. The internal protective diodes at Pin 1 of IC_1, along with resistor R_1, reduce the voltage swing at Pin 1 to an acceptable level in such a case.

A square-wave signal is at the output of the binary divider (**Figure 2**). The DMM no longer sees nanosecond pulses at its measuring terminal. You have only to multiply the displayed frequency value by two to obtain the correct frequency. Due to relatively low values of R_1 and of the input capacitance, approximately 2.5 pF, at the clock input of the flip-flop, you need not worry about frequency compensation. The time constant of $R_1 \times C_{IN}$ is merely 0.25 nsec. The width of pulses—either low or high—at the input of the circuit can decrease to 1 nsec. EDN

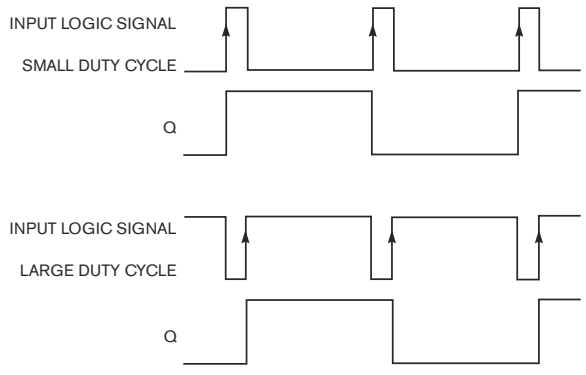

Figure 2 The flip-flop output, Q, produces a signal with a 50% duty cycle.

Build a simple complementary-bracket-pulse generator

Horst Koelzow, Global Thermoelectric, Calgary, AB, Canada

When building push-pull switching power converters or motor controllers, you often need alternating pulses with a small amount of dead time between them to minimize simultaneous conduction in output-switching devices. Switching controller ICs have this feature, but they usually operate within closed loops to minimize IC pin count. When optimizing switching output stages, you may need open-loop control. **Figure 1** shows how you can build such a generator with just two common ICs. As a bonus, both the overlapping, P-channel drive and the nonoverlapping, N-channel drive are available simultaneously.

The circuit's input, Pin 10 of IC_1, comes from clock generator IC_{2F}. A slightly delayed and inverted version occurs at IC_1's Pin 9 from IC_{2A}. IC_1 then decodes the original and delayed inputs to form the desired outputs (**Table 1**). Because IC_1 is an analog demultiplexer, you can set its outputs either active high or active low with pull-up or pull-down resistors. You determine the high or low inactive state by tying the X or Y pins to either the power-supply voltage or ground. Depending on the state of IC_{1A}'s A and B inputs, internal switches in IC_1 close between X and X0 to X and X3, as well as from Y and Y0 to Y and Y3. Buffers IC_{2B} through IC_{2E} buffer and invert the resulting outputs. You can use the remaining gate as a variable-frequency or variable-duty-cycle generator. You determine the dead time,

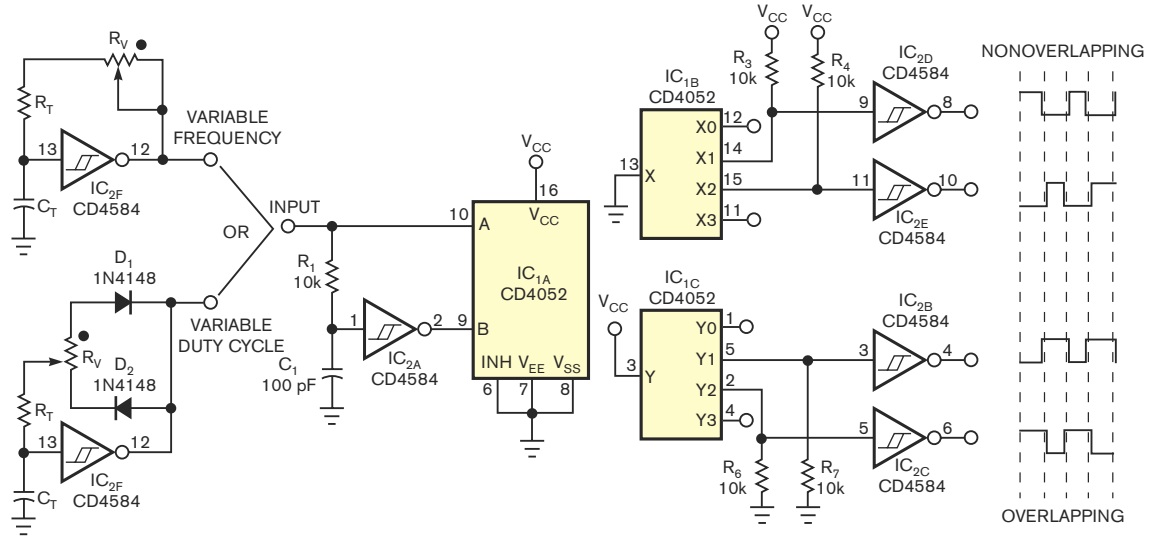

Figure 1 You can build a simple pulse generator with just two commonly available ICs.

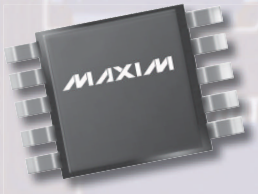

HB LED drivers for the widest range of lighting applications

Flexible and robust with wide input-voltage and output-power ranges

Maxim's HB LED drivers reduce the size, cost, and complexity of lighting designs. Our portfolio includes 29 ICs covering the full range of linear and switch-mode topologies. All devices operate over -40°C to +125°C, support wide input-voltage ranges, and integrate advanced protection features.

Architectural and street lighting

MAX16820/MAX16822/MAX16832

- Low-cost buck drivers
- Up to 65V input

MAX16826

- 4-string switching driver with I²C

MAX16824/MAX16825

- 3-string linear drivers with SPI™

M400

MR16 lights

MAX16834

- Compatible with electronic transformers and triac dimmers
- Enables flicker-free, dimmable MR16s
- Up to 12W of LED power

LuxDot™

Offline lighting

MAX16834

- > 0.90 input power-factor correction (patent pending)
- Compatible with triac dimmers
- Low component count

PAR20

Solar-powered lighting

MAX16821

- Synchronous rectification for driving high-power, common-anode LEDs
- > 94% efficiency, delivers up to 150W

MAX16834

- Boost and buck-boost topologies

EV kits available

SPI is a trademark of Motorola, Inc.
M400 and PAR20 product photos courtesy of LEDtronics, Inc.
LuxDot is a trademark of LedEngin, Inc. Photo courtesy of LedEngin, Inc.
Remaining photo: iStockphoto.com

www.maxim-ic.com/LED-Light

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

AVNET electronics marketing
www.avnet.com

MAXIM
INNOVATION DELIVERED™

For free samples or technical support, visit our website.

which is independent of frequency or duty cycle, using the time constant of R_1 and C_1. Depending on output-device characteristics and switching frequency, output buffers may require an additional stage, or you can replace them with MOSFET-gate-driver devices. Supply voltage is not critical but should be high enough to guarantee that output devices fully turn on. In general, a higher supply voltage allows for higher-speed operation. The MC14xxx series of ICs is the same as the CD4xxx series. If you need higher-frequency operation at lower supply voltages, then use the 74HC4xxx-series devices. All of these ICs are available from a number of manufacturers, including Texas Instruments (www.ti.com, **Reference 1**) and On Semiconductor (www.onsemi.com, **Reference 2**). EDN

TABLE 1 ORIGINAL AND DELAYED INPUTS

Pin 9 (Input B)	Pin 10 (Input A)	Sequence
0	1	Phase A
1	1	Dead time
1	0	Phase B
0	0	Dead time

REFERENCES

1. "CD4051B, CD4052B, CD4053B," Texas Instruments, October 2003, http://focus.ti.com/lit/ds/symlink/cd4052b.pdf.
2. "Semiconductor and Integrated Circuit Devices," On Semiconductor, www.onsemi.com/pub_link/collateral/mc14584b-d.pdf.

Power-miserly voltage reference needs just one pin

Peter T Miller, Applied Inspirations LLC, Bethlehem, CT

The supply rail normally powers a microcontroller's voltage-reference source. In power-critical battery-operated applications, the constant drain, even of a few 10s of microamps, can be prohibitive. This situation requires adding a pin to turn the reference voltage on and off. By adding a 0.1-µF capacitor in parallel with the voltage reference and a simple bit of software that you can download from the online version of this Design Idea at www.edn.com/090820dia, you'll need just one pin to both power and read the reference voltage.

When you connect the voltage reference as in **Figure 1**, the software configures the Microchip (www.microchip.com) PIC chip's V_{REF} (reference-voltage) pin as a switched-on output. After approximately 300 µsec, the voltage across the capacitor stabilizes at 1.225V.

There is an initial overshoot when the ZXRE4041 powers up. The pin is then reconfigured as an analog input for the ADC's reference-voltage source. The reference voltage quickly drops by 20 mV in the next 50 µsec as the ZXRE4041 shuts down. With a 0.1-µF capacitor, the voltage then slowly drops 60 mV over 2 msec because of leakage. Although this delay is exponential, the rate is so slow that, for practical purposes, you can consider it linear for this short time window.

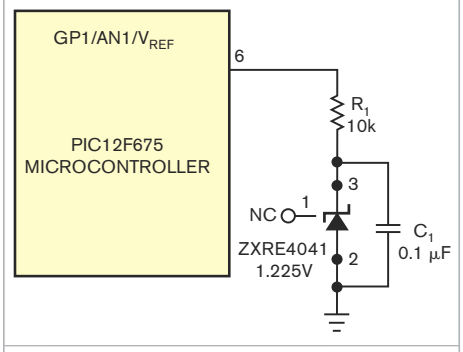

Figure 1 A voltage reference and a capacitor provide a reference voltage for a microcontroller.

You must also consider that the ADC also draws current through the 10-kΩ resistor during conversion, causing voltage drop. Although Microchip doesn't characterize this voltage drop in its documentation, tests consistently measured a drop of 80 mV for several devices, giving a calculated current of 6.67 µA. Using a conservative internal 4-MHz clock and allowing an ADC clock of frequency oscillation divided by 16 for operation at the minimum operating voltage, one conversion takes 45 µsec. This action slightly drains the capacitor, but this drainage appears to be only 2 or 3 mV. Calculations of initial watt-seconds minus watt-seconds used yield even lower values. Subtracting these fixed, repeatable losses from the initial steady-state 1.225V yields a new reference voltage of $1.225V_{REF}-0.020V$ shutdown drop-0.080 IR drop$=1.145V$.

Allowing 75 µsec to do the analog-to-digital conversion, store the value, and set up for the next conversion on another channel, 11 conversions will result in the last one's reference voltage being lower by 22.5 mV—that is, 10 conversions×75 µsec×(60 mV/2000 µsec). This error is only 1.9% compared with the first conversion's results.

If you just need an approximate voltage for a consumer product, for example, to warn of low battery voltage, you can use an LED instead of the ZXRE4041. Just change the value of R_1 to 300Ω to provide sufficient current to turn on the LED. Although LEDs lack the temperature stability of dedicated voltage-reference chips, the variation may be acceptable for the application because most consumer products find use within the comfort range of humans. If an LED is already part of the system, then the voltage-reference cost is only that of the software. Using this technique, an LED can now provide status-indicator, photodetector, and voltage-reference functions and enter a zero-power state using only software to reconfigure the changes. EDN

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

300MHz to 450MHz Transmitters

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

300MHz to 450MHz Transceivers

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

300MHz to 450MHz Receivers

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

www.maxim-ic.com/LFRF-info

MAXIM DIRECT
www.maxim-ic.com/shop

AVNET electronics marketing
www.avnet.com

MAXIM
INNOVATION DELIVERED™

For free samples or technical support, visit our website.

productroundup

COOLING AND ENCLOSURES

Sequential connecting box provides a secondary lock for added safety

Linking generators to low-voltage networks, the PowerLock Box sequential connecting box uses 2U of rack space in a 19-in. rack. The devices come in standard packages that seal only when you mate the connectors and in a sealed package with a lid covering the connector ports. A secondary lock activates a microswitch joining a two-pin connector on the back of the unit to a circuit breaker or an alarm circuit. Meeting IP65 standards, the environmentally sealed connector box comes in 400 and 600A versions. The devices have a 1000V-ac/1500V-dc voltage rating with a −30 to +85°C operating-temperature range and 500-connection life cycles. Prices for the PowerLock Box sequential connecting box range from $500 to $600, depending on quantity.

ITT Interconnect Solutions, www.ittcannon.com

240V NEMA-rated enclosures have European-style power plugs

The 240V, NEMA-rated equipment-enclosure series with built-in European-style power plugs requires no plug adapters or converters. Applications include heaters, cooling fans, and vented enclosures, and options include ABS plastic or fiberglass construction, allowing designers to protect both wired and wireless communications equipment. Prices for the 12 enclosures range from $49.99 to $469.99.

L-com Inc, www.l-com.com

Waterproof computer uses heat-pipe cooling technology

The waterproof, rugged, small-footprint, fanless WPC-500F computer features power, video, serial, and USB connectors that couple through watertight, locking, bayonet-style connectors used in military-designed hardware. Meeting IP67.NEMA 6 environmental specifications, the watertight construction survives liquids, chemicals, dust, and dirt. Using no cooling fans, the device cools the internal CPU using advanced heat-pipe technology. An aluminum chassis acts as a heat sink, dissipating internal heat, and provides noise-free operation. The computer uses Intel's Atom processor and has a 10 to 30V-dc input-power range. Measuring 191×298×75.5 mm, the basic configuration of the WPC-500F costs $1995.

Stealth.com, www.stealth.com

DC blowers suit hybrid-electric vehicles

The ODB600PT and the ODB-9733 dc-blower series come in 12, 24, and 48V models. The blowers feature brushless-dc motors, locked rotor and polarity protection, and auto-restart functions. Available in 1800-, 2500-, and 3000-rpm speeds, the blowers have a 25- to 35-cfm airflow range. The ODB600PT and the ODB9733 measure 120×120×32 and 97×94×33 mm, respectively, and suit use in cooling fuel cells in hybrid-electric vehicles, telecom equipment, and other battery-powered products. The ODB600PT dc-blower series costs $23.50 for the 12 and 24V models and $27.50 for the 48V model. The ODB9733 series costs $18.50 for the 12 and 24V models and $20.50 for the 48V model.

Orion Fans, www.orionfans.com

Condenser uses airflow in the chassis to cool CPU

Suiting use in noise-sensitive workstations, the SilentFlux passive pro cooler high-efficiency condenser uses airflow within the chassis to cool the CPU. The condenser uses bubble-

pump technology, a nonmechanical, self-contained system that uses CPU-generated heat to create hot liquid and gas bubbles, which move through hermetically sealed, closed-loop tubing and a condenser/radiator to cool the system. The bubbles continue circulating, condensing, and returning, creating the bubble pump. The aluminum design reduces the need for high-power, high-noise airflow. The SilentFlux passive pro cooler costs $30 (1000).

Noise Limit, www.noiselimit.com

INTEGRATED CIRCUITS

ADC enables daisy-chaining of as many as eight devices

The four-channel, simultaneously sampling MAX11040 ADC allows daisy-chaining of as many as eight devices, providing a 32-channel sampling capacity. The 24-bit converter provides cascade SPI, QSPI, and Microwire interfaces. Featuring a minimum of 90 dB of SINAD and 91 dB of SFDR, the device aims at industrial power-grid-protection equipment, medical EKG/EEG equipment, and applications requiring accurate conversion of simultaneously sampled channels between 0.25k and 64k samples/sec. The converter provides overvoltage protection against ±6V voltages. Available in a TSSOP-38, the ADC costs $13.45 (1000).

Maxim Integrated Products, www.maxim-ic.com

Multiplexers provide rail-to-rail analog capabilities

The eight-channel DG508B and dual four-channel DG509B CMOS analog multiplexers enable ±15V high-precision, low-noise signal switching, suiting use in data-acquisition, medical-system, and test-and-measurement equipment. Features include rail-to-rail analog capability, 3-pA channel off-leakage, 3-pF parasitic capacitance, and 2-pC charge injection. Operating over a −40 to +125°C temperature range, the multiplexers consume 10 µA at 25°C with a 600-µA maximum consumption over temperature. The devices have a 200-MHz, −3-dB bandwidth with crosstalk and off-isolation of −40 dB at 100 MHz. Available in SOIC-16 and TSSOP-16 packages, the DG508B and DG509B CMOS analog multiplexers cost $1.15 (1000).

Vishay Intertechnology, www.vishay.com

EDN ADVERTISER INDEX

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

ERNEST TANNER • LATTICE SEMICONDUCTOR

Finger on the trigger

DANIEL VASCONCELLOS

I've been working in the electronics field for more years than I like to admit. The equipment and techniques we have available now are much more powerful than I could have imagined—a good thing, given that the problems we have to solve have evolved just as much.

My company is a semiconductor manufacturer that makes CPLDs and FPGAs. Because most of my career has been in the lab, I spend most of my time testing and debugging. Sometimes, a customer's design is not working properly, and I have to figure out what the problem is. If it turns out to be a problem with the design and not with our part, we like to have solid proof before blaming the mistake on the customer.

One customer's board had pulses—definitely coming from our part—that shouldn't have been there. Most of the prototypes had the unwanted pulses, so it wasn't just one broken chip. Ideally, that signal was normally high with low-going pulses of varying widths—17 µsec being the longest. The problem pulses were about 100 nsec wide and appeared fairly regularly.

Just for the record, this design looked like a well-done job. We could find no obvious problems.

PROGRAMMABLE LOGIC ALLOWS YOU TO EASILY MODIFY THE DESIGN.

While probing the input signals, we noticed that the problem disappeared when the probe was on a particular input. Aha! We'd all seen this one before. There was a small glitch; it had to be very small for the 0.7-pF capacitance of the probe—a FET probe, in this case—to eliminate it.

Fortunately, the customer had submitted good information about the design, and we were able to inspect the schematic for the board and the source code for the programmable part. The design used that signal all over the board as a clock. It wasn't a very fast clock, but the edge rates were approximately 1 nsec. The signal was an old-fashioned TTL (transistor-transistor-logic)-level signal, with an input low voltage of 0.8V, an input high voltage of 2V, and an output high voltage of approximately 3V. I had enough information to convince me, but I wanted more. Just to test the theory, we soldered a 10-pF capacitor to the trace, and the problem went away. No engineer I know would use this as a solution, but it was good confirmation.

We had previously not had a good trigger, so it was hard to get a good look at one part of the signal. We took a lot of single-shot pictures and determined that the longest pulse on the output was about 17 µsec. Because this was 2006, not 1976, we had some nice options for triggering. Setting the trigger to capture a negative-going pulse greater than 15 µsec gave us a nice, stable look at the output. The glitches were occurring at consistent places, which made it possible to zoom in and take a good look at the suspicious clock. The falling edge had a little bump starting at about 0.6V and peaking at 0.9V before starting back down.

Still, we wanted to get even more evidence. One good feature of programmable logic is that it allows you to easily modify the design to bring out an internal signal. We created a buffer that brought the clock signal back out to an unused pin. If the bump was not enough to cause the glitch, then the buffered output should be clean. Otherwise, we'd get a significant pulse. We got a nice, clean 1-nsec pulse at a voltage as high as 2.3V. And we were able to give the customer a detailed report that settled the issue for good. EDN

Ernest Tanner is an application engineer at Lattice Semiconductor.

Go to www.edn.com/090820tales and click on Feedback Loop to post a comment on this article.

www.edn.com/tales

No need to duplicate yourself

BuyerZone saves you time and money on COPIERS AND MORE.

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

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

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

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

BuyerZone
Where Smart Businesses Buy and Sell

A division of
Reed Business Information.

PoE+ Compliant PDs for the Power Hungry

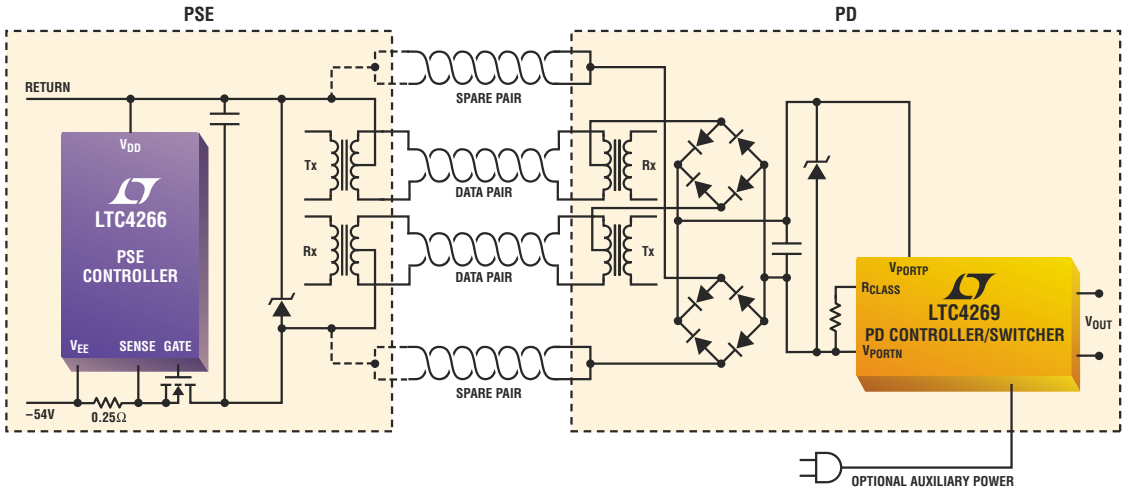

High Efficiency, Low EMI PD Controllers with Flexible Auxiliary Input

Linear Technology now offers a family of Power over Ethernet interface controllers for emerging PoE+ applications, including high performance wireless access points, PTZ security cameras and RFID readers. The LTC®4269 is a complete PoE+ compliant solution that integrates both an IEEE 802.3at PD controller and a high efficiency, low EMI switching regulator with the option of either a synchronous NoOpto flyback or a synchronous forward converter. Both options provide an optional auxiliary input and support the new PoE+ standard with 2-Event Classification signaling.

PoE Solutions from LTC

Part Number	PSE/PD	Channels	Description	Package (mm)
LTC4264	PD	1	Pre-Standard High Power PD PoE Interface Controller (35W)	4 x 3 DFN-12
LTC4265	PD	1	IEEE 802.3at High Power PD PoE+ Interface Controller (25.5W)	4 x 3 DFN-12
LTC4268-1	PD	1	Pre-Standard High Power PD PoE Interface Controller with Integrated Synchronous Switching Regulator (35W)	7 x 4 DFN-32
LTC4269-1	PD	1	IEEE 802.3at High Power PD PoE+ Interface Controller with Integrated Synchronous No-Opto Flyback Switching Regulator (25.5W)	7 x 4 DFN-32
LTC4269-2	PD	1	IEEE 802.3at High Power PD PoE+ Interface Controller with Integrated Synchronous Forward Converter (25.5W)	7 x 4 DFN-32
LTC4266	PSE	4	Quad IEEE 802.3at PoE+ Controller with DC Disconnect	7 x 7 QFN-36, SSOP-36
LTC4263-1	PSE	1	IEEE 802.3af PSE Controller with Internal Switch	4 x 3 DFN-14, SO-14

Info & Free Samples

www.linear.com/PoE

1-800-4-LINEAR

Free Telecom, Datacom and Industrial Brochure

www.linear.com/48vsolutions

LINEAR TECHNOLOGY