


EDN®

VOICE OF THE ENGINEER

AUG

6

Issue 15/2009
www.edn.com



Dark side of the light
Pg 58

EDN.comment: Self-serving "synergy" Pg 5

Signal Integrity: The nature of ESD Pg 10

Design Ideas Pg 43

Supply Chain Pg 54

Product Roundup Pg 56

SPEAKING OF PORTING SOFTWARE

Page 18

AMBIENT-LIGHT SENSORS PACK IN FEATURES TO HELP APPLICATIONS GET SMARTER, GREENER

Page 13

IMS HIGHLIGHTS STUDENT INGENUITY

Page 25

TECHNIQUES FOR IMPLEMENTING HIGH-PERFORMANCE PROCESSOR CORES



Page 28

DRAM TECHNOLOGY FOR SOC DESIGNERS AND—MAYBE—THEIR CUSTOMERS

Page 34

FIRST-PASS SUCCESS IN SILICON PACKAGING

Page 38



Lighting the Way

NEC

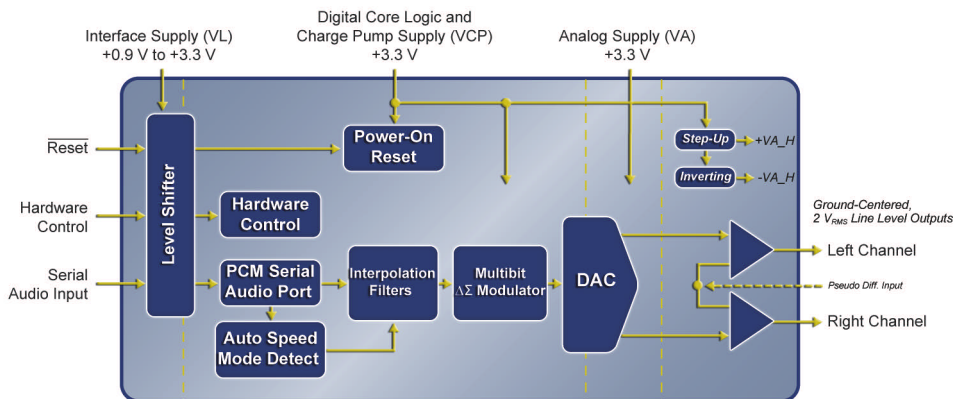
Intelligent Light Solutions from NEC Electronics America
www.am.necel.com/led

innovation



CS4353

Thermally Enhanced
4 mm x 4 mm
QFN Package
(shown at actual size)



The CS4353 is Based on a Charge-Pump Architecture and Features an On-Chip $2 V_{RMS}$ Line Driver Working from a Single 3.3 V Power Supply Without the Need for a High Voltage Regulator. Also a Ground-Centered Output Removes the Need for External DC Blocking Capacitors and Any Pop Prevention Circuit.

Finally—A No Compromises Digital-to-Analog Converter

Cut the clutter! When it comes to digital-to-analog converters, Cirrus Logic means business. The CS4353 is the only audio DAC to feature an on-chip $2.1 V_{RMS}$ line driver from a single 3.3 V power supply, allowing a $2 V_{RMS}$ voltage after the RC filter. The new architecture not only simplifies the design of analog output section circuitry but it also provides superior quality audio playback with the highest dynamic range and lowest pops and clicks of its kind. This is especially important in applications such as Blu-ray Disc® players, DTV and set-top boxes as they allow for the elimination of multiple system components and lower implementation cost.

NORTH AMERICA

+1 800-625-4084

ASIA PACIFIC

+852 2376-0801

JAPAN

+81 (3) 5226-7757

EUROPE/UK

+44 (0) 1628-891-300

CS4353 FEATURES

- 106 dB dynamic range, -93 dB THD+N
- Single-ended ground centered analog architecture
- No DC-blocking capacitors required
- Click-free stereo output
- Selectable 1 or $2 V_{RMS}$ full-scale output
- High PSRR
- Small 4 mm x 4 mm QFN package

DOWNLOAD A COPY OF THE
CIRRUS LOGIC BLU-RAY DISC®
PLAYER/RECEIVER
SOLUTIONS BROCHURE AT
WWW.CIRRUS.COM/4353EDN



LEARN MORE AT
www.cirrus.com

Cirrus Logic. We make it easier for you.

MORE SEMICONDUCTORS THAN EVER*




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

Authorized Distributor for more
than 400 Supplier Partners.

The industry's broadest product selection
available for immediate delivery

www.digikey.com
1.800.344.4539




Speaking of porting software

18 This software-porting hands-on experiment uncovers a potential audio decoder for embedded-system applications, adding audio or speech to the applications' user interfaces. *by Robert Cravotta, Technical Editor*




EDN 8.6.09 contents



Ambient-light sensors pack in features to help applications get smarter, greener

13 Ambient-light sensors help smartphones conserve display power and improve battery runtime. Newer versions combine improved features, such as spectral response and dynamic range, integrated proximity, and digital communications.

by Margery Conner, Technical Editor



IMS highlights student ingenuity

25 Competitions drive student interest in microwave theory and techniques.

by Jennifer Kempe, Contributing Editor

Techniques for implementing high-performance processor cores

28 Every step of the integration process gives the design team opportunities to move closer to its design goals.

by Gregory Schulte and Vamsi Boppana

DRAM technology for SOC designers and—maybe—their customers

34 An understanding of DRAM technology has become critical to anyone designing consumer-electronics SOCs.

by Drew E. Wingard, PhD, Sonics Inc

First-pass success in silicon packaging


38 Package issues interact with all aspects of the design flow, from chip architecture to manufacturing decisions.

by Javier DeLaCruz, eSilicon Corp

pulse Dilbert 7

- 6** 6A PSIP dc/dc converter achieves 15W/cm²
- 6** Monolithic 6A dc/dc step-down controller integrates synchronous switcher, power-output FETs
- 7** Cascade Microtech addresses power semiconductors and RFICs
- 8** Start-up offers dynamically reconfigurable logic technology
- 8** Adjustment-free units simplify precise probe positioning
- 9** **Research Update:** Iron particle in nanotube may offer archival storage; Cyborg crickets could create cooperative community networks


DESIGN IDEAS



- 43** High-speed op amp enables IR-proximity sensing
- 44** Set your lights to music
- 48** Current limiter allows large USB bypass capacitance
- 50** High-speed pulse modulator retains signal envelope



5



56

DEPARTMENTS & COLUMNS

- 5 **EDN.comment:** Self-serving "synergy"
- 10 **Signal Integrity:** The nature of ESD
- 54 **Supply Chain:** Job board targets distributors; PC market to slump 6% in 2009; Possible changes to toxic substances control act echo ROHS
- 56 **Product Roundup:** Discrete Semiconductors, Computers and Peripherals
- 58 **Tales from the Cube:** Dark side of the light

EDN online contents

www.edn.com

ONLINE ONLY

Check out these Web-exclusive articles:

Closing the ESL gap

Plenty of commentary has been written about the promise of ESL and how it remains unfulfilled. The truth is that engineers have been successfully designing with ESL tools for years.

→ www.edn.com/article/CA6665045

Tackling formal assumptions through verification planning

There are many tools for tracking assertions during verification, but what about the assumptions the design team has made about the design's environment? They need tracking, too.

→ www.edn.com/article/CA6669411



ENGINEERS' TRUE STORIES

Tales from the Cube, which appears on the last page of every issue, has become one of *EDN's* most popular sections. In *EDN's* Tales from the Cube, engineers relate their most vexing design challenges—and how they conquered them. From the cubicle to the test bench to the field, hear how your peers solve real-world problems. To catch up on any Tales you might have missed—and to learn how to submit your own—please check out our Tales archive page.

→ www.edn.com/tales

Benchmark MOSFETs

IR's new MOSFETs in TO-247 reduce $R_{DS(on)}$ by 50% over competing devices



Part Number	$B_{V_{DS}}$ (V)	$R_{DS(on)}$ (m Ω)	Q_g (nC)	I_D @ 25°C (A)
IRFP4004PBF	40	1.70	220	195*
IRFP4368PBF	75	1.85	380	195*
IRFP4468PBF	100	2.60	360	195*
IRFP4568PBF	150	5.90	151	171
IRFP4668PBF	200	9.70	161	130

* Package limited

Features

- Suited for Synchronous rectification, active ORing, industrial battery and DC to AC inverters
- Industrial grade and MSL1
- Lead free, RoHS compliant

Your **FIRST CHOICE** for Performance

For more information call
1.800.981.8699 or visit
<http://www.irf.com>

International
IR Rectifier
THE POWER MANAGEMENT LEADER

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 15 (Printed in USA).

PICO

Surface Mount
(Thru-Hole Available)
Transformers and
Inductors

See Pico's full Catalog immediately
www.picoelectronics.com

Low Profile from
.19" ht.



Audio Transformers

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

Power & EMI Inductors

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

Pulse Transformers

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

Multiplex Data Bus Pulse Transformers

Plug-In units meet the requirements
of QPL-MIL-PRF 21038/27.
Surface units are electrical equivalents
of QPL-MIL-PRF 21038/27.

DC-DC Converter Transformers

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

400Hz/800Hz Power Transformers

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

Delivery—stock to one week
for sample quantities

See EEM
or send direct

for **FREE** PICO Catalog
Call toll free 800-431-1064
in NY call 914-738-1400
Fax 914-738-8225

PICO Electronics, Inc.
143 Sparks Ave, Pelham, N.Y. 10803
E Mail: info@picoelectronics.com
www.picoelectronics.com

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
Kirimaya Varma, Editor-in-Chief
kirri.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



BY BRIAN DIPERT, SENIOR TECHNICAL EDITOR

Self-serving “synergy”

The HDMI (high-definition-multimedia-interface) organization finalized its Version 1.4 specification in early June. Silicon Image, an HDMI-founder company, subsequently released its first two Version 1.4-supportive chips, the SiI9387 port processor and the SiI9334 transmitter. This timing wasn't the way things were supposed to play out. The original plan was for HDMI to publicly unveil the finalized specification, and Silicon Image to coincidentally unveil the chips, in late April.

Then, the two pushed out the embargo date for both spec and silicon to mid-May and eventually released the spec stand-alone as a two-step preview-then-final introduction, with the chips following it. Despite the multiple delays, nobody else that I know of had as of mid-July released HDMI Version 1.4-supportive chips.

Besides Silicon Image, the other HDMI-founding companies are Hitachi, Matsushita/Panasonic, Philips Consumer Electronics International, Sony, Thomson (RCA), and Toshiba. Although some of these companies have semiconductor divisions, only Silicon Image is fundamentally a semiconductor supplier. And, as a founder, it is in the unique position—unlike other HDMI-silicon vendors—of having not only access to proprietary draft versions of upcoming spec revisions, but also the ability to heavily influence the development direction of those spec revisions.

Silicon Image can start its chip-design work early, and the company garners revenue from other silicon suppliers, its supposed competitors, in the form of licensing royalties. And it has yet another notable competition-squelching arrow in its quiver. Silicon Image is an umbrella organization



Silicon Image is trying to have it both ways.

comprising three business groups, one of which is Simplay Labs, which handles industrywide compatibility testing. Each Simplay Labs validation attempt costs other companies many thousands of dollars. The validation process also takes as long as Simplay Labs deems it should take. This arrangement gives Silicon Image a convenient lock on the market through the initial and profitable portion of each HDMI specification version's lifetime.

Not surprisingly, I've heard plenty of complaints over the years from companies that try to compete with Silicon Image on the supposedly level global-standard playing field. I've also been on the receiving end of a lot of

grumbling from Silicon Image's customers—some of them even on the HDMI founders list! System suppliers would prefer to be able to source products from multiple vendors as a means of obtaining both optimum prices and assured supply. However, again and again, they find that Silicon Image is their only product option in the early stages of an HDMI-specification version. And, by the time an HDMI generation is sufficiently mature that other vendors are available, it's time for Silicon Image's reluctant customers to begin designing systems complying with the next HDMI version.

The market situation is so distorted that a collective of graphics-chip companies and display manufacturers a few years ago developed a competitive interface, DisplayPort. I've suggested many times that DisplayPort's success will be muted at best, in no small part because it took the technology's developers so long to finalize the specification. Meanwhile, HDMI was cultivating an insurmountable lead, especially in consumers' living rooms. Looking at the market today, I see no evidence that my past prognostication was off-base. However, that HDMI's detractors were motivated to spend a substantial amount of money and allocate precious resources for DisplayPort development indicates the industry's aggravation with Silicon Image's de facto HDMI dictatorship.

Silicon Image is trying to have it both ways—painting a picture of HDMI as a cozy, cooperative industry standard, when in reality HDMI is predominantly a single-silicon-supplier proprietary approach. **EDN**

Contact me at bdipert@edn.com.

⊕ For more information, go to www.edn.com/blog/40000040/post/1850046185.html.

⊕ Read a response from Silicon Image at www.edn.com/090806ed.

pulse

INNOVATIONS & INNOVATORS

6A PSIP dc/dc converter achieves 15W/cm²

Enpirion keeps raising the bar on its tiny PSIP (power-supply-in-package) dc/dc buck converters. The 6A EN5364QI and the 9A EN5394QI each have an input-voltage range of 2.375 to 6.6V and fit into an 8×11×1.85-mm package, achieving as much as 93% efficiency and a power density of 15W/cm². An external resistor divider sets the

output voltage to any user-defined value. The device has an operating frequency of 4 MHz and can synchronize with an external system clock or another EN5364/94.

Like other Enpirion PSIP products, the EN5364/94s integrate power switches, an inductor, a gate drive, a controller, and loop compensation in the 190-mm² package, requiring as few as seven external resistors and capacitors. The EN5364 and EN5394 sell for \$6.10 and \$7.50 (1000), respectively.

These figures emphasize the continuing trend toward lower prices and increased power density in the PSIP market. For example, the company about a year ago introduced the \$8.10 (1000), 9A EN5395QI PSIP converter, which fits into a 10×12×1.85-mm package and achieves approximately 11W/cm², also with as much as 93% efficiency.

—by Margery Conner

► **Enpirion**, www.enpirion.com.



Requiring as few as seven external resistors and capacitors, the 6A EN5364QI achieves 15W/cm² and efficiency as high as 93%.

FEEDBACK LOOP

“If I were asked about the future of analog, I’d say you’re going to have some very large (in terms of functionality onboard) chips, with large application departments at those companies explaining what the chips do and how their customers can use them.”

—Design engineer Chris Gammell, in *EDN's* Feedback Loop, at www.edn.com/article/CA6666235. Add your comments.

Monolithic 6A dc/dc step-down controller integrates synchronous switcher, power-output FETs

Texas Instruments' latest additions to its Swift series of power-supply ICs include point-of-load dc/dc converters in highly integrated packages. The 17V TPS54620 step-down synchronous switcher with integrated FETs comes in a 3.5×3.5-mm package and contains the control IC and power-output FETs. After adding a handful of passive components, you can pack a complete 6A power device into less than 195 mm².

Unlike some vendors' approaches to the problem, TI's package hosts a monolithic device with integrated 26-mΩ high-side and 19-mΩ low-side FETs. It downconverts supplies ranging from 1.6 to 17V, although you must provide a separate supply of at least 4.5V for control and gate-drive purposes. The inductor sets the height,

or profile, of your converter design; TI's suggested parts list's smallest suitable inductor has a 3-mm profile. The device operates at 200 kHz to 1.6 MHz, and you can lock its switching frequency to an external clock. You must provide an external compensation network of approximately five passive components, and TI's online or downloadable software package, SwitcherPro, assists with calculating those component values. The chip has a 0.8V reference and yields a 1%-accurate output across its range. Peak efficiency is 95%, and the chip maintains better-than-90% efficiency over an output range of 1 to 6A when deriving 3.3V from a 12V input. The TPS54620 costs \$1.95 (1000).—by Graham Prophet

► **Texas Instruments**, www.ti.com.

Cascade Microtech addresses power semiconductors and RFICs

Cascade Microtech is addressing the emerging energy-efficiency standards that are driving the need for accurate power-device characterization in automotive, mobile-device, transportation, and other applications. In addition, the company is targeting probing for millimeter-wave RFICs that will serve WirelessHD, automotive radar, and other applications operating in the 60- to 80-GHz range.

In the company's most recent initiative to address power-characterization needs, it introduced a set of probes and accessories for its Tesla on-wafer power-device characterization system, making Tesla fully compatible with the recently released Agilent (www.agilent.com) B1505A power-device analyzer (see "Tesla power-semiconductor-characterization system debuts," *Test & Measurement World*, May 29, 2007, www.tmworld.com/article/CA6446837). The combined system offers an extended triaxial measurement range to accommodate low-noise probing of power devices with voltages as high as 2000V.

The award-winning Tesla

system meets these demands for making measurements at increasing voltage and current levels when characterizing devices fabricated using new wide-bandgap materials, such as silicon carbide and gallium nitride. Tesla offers what the company says is the industry's highest voltage and current range for on-wafer measurements: as much as 2000V triaxial or 3000V coaxial and as much as 60A pulsed or 20A continuous. With



Using a combination of Cascade Microtech's proprietary thin-film technology and coaxial-probe technology from its Infinity Probe architecture, the 110-GHz Unity-MW RFIC-engineering probe supports the precision characterization and testing of these emerging multiple-port technologies.

Agilent's B1505A power-device analyzer, the new Tesla probes take advantage of the performance of the B1505A, meeting the requirements of more advanced device-characterization applications.

"Power devices are prolific in today's semiconductor industry," says Geoff Wild, chief executive officer of Cascade Microtech. "Our customers are constantly striving to improve the efficiency of these

critical IC components. The Tesla system facilitates on-wafer CV, IV, and breakdown measurements, which in turn enables faster development cycles at an overall lower cost of test versus packaged test." Unlike other approaches, Wild adds, the system allows users to realize the full potential of their B1505As with the maximum range of voltage, current, and application compatibility.

Cascade Microtech has also announced two products that streamline engineering and production testing of high-bandwidth, short-range RFIC devices for WirelessHD, automotive radar, and other 60-GHz wireless applications,

Market researchers point to a surge in demand for devices operating in the 60- to 80-GHz range.


joining other companies at the June International Microwave Symposium addressing applications in the 60-GHz area (see "Design and test highlights at the microwave show," *Test & Measurement World*, www.tmwworld.com/blog/640000064/post/290045629.html). Using a combination of Cascade Microtech's proprietary thin-film technology and coaxial-probe technology from its Infinity Probe architecture, the 110-GHz Unity-MW (millimeter-wave) RFIC engineering probe supports the precision characterization and testing of these emerging multiple-port technologies. Using the same membrane technology, the Pyramid-MW 81-GHz production probe card supports the at-speed test of known-good die in high volume.

Cascade cites market researchers' pointing to a surge in demand for devices operating in the 60- to 80-GHz range: ABI Research (www.abiresearch.com) forecasts the installation of 1 million wireless high-definition TVs worldwide by 2012, with double-digit growth rates, and Strategy Analytics (www.strategyanalytics.com) reports that more than 2.3 million cars will have collision-avoidance radar systems, requiring more than 30 million radar sensors, by 2011.

—by Rick Nelson

▷ **Cascade Microtech**, www.cascademicrotech.com

DILBERT By Scott Adams



Start-up offers dynamically reconfigurable logic technology

Start-up Akya is offering IP (intellectual property) to allow IC designers to include reconfigurable logic on their ASSPs (application-specific standard products) or ASICs (application-specific integrated circuits). The company delivers both reconfigurable-logic fabric and IP blocks to execute commonly required functions on that fabric. Colin Dente, the company's chief executive officer, says that the reconfigurable logic can be a part of a larger chip design or form the basis of a complete IC. Akya calls its technology ART2. The ART2 architecture separates data flow and control logic to simplify design and implementation of reconfigurable structures.

Akya's technology differs from other concepts that companies have marketed as dynamically reconfigurable. "If I could think of another term to describe what we do, I would use it," says Dente, explaining that the architecture is not completely general-purpose, but that the company designed it to be just reconfigurable enough to do what is necessary in applications such as audio/video codecs and communications-signal processing. The company claims that the technology is particularly efficient at DSP functions.

Any engineer familiar with HDL (hardware-description-language)-based design can learn how to use the tools, which allow you to design a fabric that targets a specific application. This fabric comprises and draws on an IP library of processing elements. These elements are typically at the level of arithmetic functions

rather than being fine-grained, gate-level elements. To this fabric, you add an address sequencer and a programmable interconnect to yield a set of processing resources that match your design. The design tools then assist you in creating an explicit control flow for the design, and you can change that flow after the silicon becomes final. The reconfigurability materializes under the direction of the control flow. With every clock cycle, you can completely reconfigure the datapath: You specify which action each processing element is to carry out on a cycle-by-cycle basis.

Doing the control programming is somewhat similar to writing code for a processor. However, the device you work with is not a VLIW (very-long-instruction-word) processor. It does not write and read data to and from memory, and the minimal use of such transactions is one of the sources of Akya's claims of very low power usage: "In [silicon] area and power, we are right at the custom-silicon end of the [design] spectrum," Dente says, adding that portable and battery-powered products are targets for the technology. In part due to the architecture's processing efficiency, clock speed is unlikely to be a constraint for most telecommunication or audio/video-processing functions.

Thanks to the technology's reconfigurability, it reduces the risk inherent in IC design, takes half the design time of RTL (register-transfer-level) design, and lets you quickly modify designs without changing the silicon itself. With a reconfigurable block,

you might make one mask set that serves for several designs, making a custom IC viable when it otherwise would not be. The company offers a similar argument for the problem of adapting an IC to changing standards.


Akya supports ART2 with a

development kit that features one high-level language for data flow and another for control; the company also provides comprehensive training in the ART2 architecture compiler. ART2 is available now; the first commercial chip incorporating the technology is in development and will be on sale in an end product in 2010.

—by Graham Prophet

► Akya, www.akya.co.uk.

ADJUSTMENT-FREE UNITS SIMPLIFY PRECISE PROBE POSITIONING




The two-arm N2785A probe positioner provides quick and stable X-Y positioning using a simple lift-and-drop motion to set the probe in place.

Agilent Technologies' four new oscilloscope-probe positioners, the N278xA series, provide quick and stable probe positioning for measurements on PCBs (printed-circuit boards) and other devices that require hands-free probing. Unlike other probe positioners, which require multiple adjustments to lock the probe holder into place, the one-arm N2784A positioner and two-arm N2785A positioner need only a lift-and-drop motion. The probe holders' weight-stabilization technique keeps constant pressure at the probing point, holding the probe tip in position even if someone bumps the target board.

The low-cost, easy-to-use N2786A XY-axis probe holder for general-purpose probing features an easy-to-use two-legged positioner that needs no controls for adjusting its position. The 3-D N2787A positioner has a flexible, articulated arm that you can quickly set up in a variety of configurations, yet it allows fine probe-location adjustments to achieve a stable contact.

These positioners fit most Agilent passive and active probes and are compatible with many probes from other manufacturers. US prices begin at \$60.—by Dan Strassberg
► Agilent Technologies, www.agilent.com/find/probe_positioners.

08.06.09



This memory device comprises an iron nanoparticle within a carbon nanotube.

RESEARCH UPDATE

BY RON WILSON

Iron particle in nanotube may offer archival storage

Researchers continue to seek out some plausible application for carbon nanotubes. The latest effort, by a team at the University of California—Berkeley, the Lawrence Berkeley National Laboratory, and Pennsylvania State University, has discovered that a multiwall nanotube containing a particle of iron might just be an effective archival storage device.

The researchers fabricated the structure in a single step by pyrolysis of ferrocene in argon at high temperatures. The result is a cloud of nanotubes containing iron nanoparticles. The team then deposited the nanotube fragments on a convenient substrate and fabricated contacts at the ends of likely looking nanotubes, creating the experimental vehicle.

Applying a sufficient voltage

across the nanotube causes the iron particle to move because of electromigration. The team discovered that, using only the applied voltage, they could precisely control both the velocity of the particle and its position in the tube. More significant for storage applications, once you remove the high voltage, the particle stays put. If you don't intentionally move it and it is subject only to thermodynamic forces, the researchers estimate, the particle will stay in place for about 1 billion years at room temperature.

Initially, the researchers read the position of the iron particle in the nanotube with an electron microscope. This approach

would be somewhat impractical, however, in a large-scale storage system. So they investigated the electrical resistance of the nanotube with respect to the position of the iron particle and discovered that the resistance shows a close link to the position. Hence, measuring the resistance of the nanotube at a voltage low enough to not move the particle provides a nondestructive read-out of the location of the particle. This reading process is sufficiently precise to allow multibit storage within a single nanotube.

► **University of California—Berkeley**, <http://berkeley.edu>.


► **Lawrence Berkeley National Laboratory**, www.lbl.gov.

► **Pennsylvania State University**, www.psu.edu.

CYBORG CRICKETS COULD CREATE COOPERATIVE COMMUNITY NETWORKS

Crickets, cicadas, katydids, and some other insects make noise by beating their wings. More significantly, they can change the pitch of their wing beats and, hence, the frequency of the noise in response to changes in pitch of other nearby insects. Swarms of these insects form peer-to-peer wireless networks that can propagate information at relatively high speed across significant areas—perhaps tens of kilometers.

The insects use this ability for their own inscrutable purposes. At least they did until they encountered Benjamin Epstein, PhD, vice president of special projects at OpCoast, a provider of



Electronically modified crickets would be the nodes in a local-area sensor network.

custom products for networking, modeling and analysis, software and database systems, and security and lawful interception. Epstein had the idea that by wiring electronics into the wing muscles on cicadas he could in principle artificially modulate the wing beats, in effect loading information into the insects' network and watching it propagate.

Epstein landed a research contract from the US Department of Defense to explore the feasibility of the idea for creating a living network of battlefield toxic-gas detectors. Other applications include mine safety and survivor detection in natural disasters.

The idea is to create a package of electronics, including a gas sensor, perhaps an audio sensor in case the insect proves uninterested in rebroadcasting a strange signal by itself, a wing-muscle stimulator, a tiny bit of intelligence, and some sort of power supply. This package has to be small enough to be implantable in the insect and to not interfere with its short-term

survival. The size issue, Epstein says, was not so much a concern with cicadas, which can be rather large. It does become an issue with more common and smaller insects, such as crickets, which are more desirable hosts because they can thrive in a wider range of climates.

In operation, the electronics payload would be passive until the detector found whatever it was looking for. The package would then stimulate the insect's wing muscles, causing the bug to emit a particular pitch. Presumably, neighboring insects would pick up the new pitch and repeat it.

► **OpCoast**, www.opcoast.com.

60.90.80



BY HOWARD JOHNSON, PhD

The nature of ESD

Alex Ching from BJ Pipeline Inspection Services in Canada designs large, complex digital systems. His product incorporates analog, digital-logic, battery, and chassis grounds, which connect at various “single-point connections” using wires, 100-k Ω resistors, and TVSDs (transient-voltage-suppression diodes). When he strikes the product with an ESD (electrostatic discharge), errors appear on sensitive internal LVDS (low-voltage-differential-signaling) data lines. Ching

has sprinkled ESD-protection chips throughout the product in a fruitless attempt to control that problem. He says that, when an ESD event activates these protection chips, they are supposed to dump ESD currents into the digital-logic ground. From the digital-logic ground, he assumes that the currents flow through a single-point connection to the battery-ground network and, from there, through the parallel combination of a 100-k Ω , high-voltage resistor and an SMBJ5.0CA TVSD to the chassis. “Since the ESD-protection chips are so far from the chassis, ESD currents must travel a long distance on the digital-logic ground before passing through the battery ground to the chassis,” Ching says. “Is this long ESD current path causing noise on my LVDS logic?”

Ching faces a common problem: In response to various noise issues that have emerged over the life of his product family, engineers have modified the grounding structure, probably several times, resulting in a “confused” grounding architecture. I cannot understand his whole system from this brief description, but I can perhaps discuss the general nature of an ESD transient.

ESD transient currents spread far and wide.

ESD is a fast event. As the ESD currents surge through your system, they follow, at each instant, what seems the best path available. If, for example, an ESD transient current reaches a junction of two wires, it splits evenly, half going down each wire, with no regard to where the wires lead. One path may lead to a good ground, whereas the other leads to an open circuit. The current cannot discern the better path to ground without first propagating to the ends of the two pathways. Only then, after considerable time and many reflections pass, do the lower-frequency elements of the ESD transient begin to recognize what you may think of as the best path to ground.

As a result of this behavior, the single-point ground connections in Ching’s product do almost nothing to guide ESD transient currents. Once inside the product, ESD transient currents spread far and wide regardless of any ground jumpers, 100-k Ω resistors,

and TVSDs that may exist. For example, ESD currents can easily jump directly from board to board through no connection other than the parasitic capacitance of one system element to another. In other cases, current flowing on one path can, through the action of mutual inductance, easily induce destructive voltages in an unrelated pathway.

To control ESD transient currents, you must design the system so that ESD events have only one clear path to ground at every point. For example, terminating a shielded cable onto a metal-shell connector with 360° shield contact around the connector and between the connector and the metal system chassis provides an adequate means of shunting ESD transients from a cable shield to the chassis. The transients have nowhere else to go.

Although TVSDs do a good job limiting the voltage across their terminals, they do not prevent current transients from entering your system. Once current passes through a TVSD, it must still return to ground. Generally speaking, if you mount the TVSD anywhere on your PCB (printed-circuit board), then you allow ESD current to flow on the PCB, and you may lose the war for that reason alone. The best designs shunt ESD currents back to earth before entering sensitive areas of the product.

Ching needs to properly diagnose his system before making any more changes. An engineer familiar with EMI (electromagnetic interference) can determine precisely where the ESD currents are flowing within the product and then design adequate means of protecting the product. Doug Smith of D.C. Smith Consultants provides some great tips on his Web site, www.emcesd.com, about diagnosing the flow of current from large, fast transient events. **EDN**

www.edn.com/signalintegrity

"I need a function generator that generates efficiency, too."



Signal quality is just the beginning.

You're facing faster design cycles and more complicated testing protocols. PC connectivity is more critical than ever, too. Agilent function/arbitrary waveform generators help shorten your test times with legendary signal quality and purity — along with state-of-the-art I/O connectivity. Plus they're bench-friendly, rack-able, reliable and easy to use. That's Agilent. That's smart.

Function/Arbitrary Waveform Generator Series

33210A—10MHz	\$1,195*
33220A—20MHz	\$1,860*
33250A—80MHz	\$4,571*



Add a FREE External Reference worth up to \$495**
www.agilent.com/find/smartarb

© 2009 Agilent Technologies, Inc.
*Pricing in USD and subject to change
**Limited time offer

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



Agilent Technologies

160 Megawatts of Green Engineering

MEASURE IT – FIX IT



Since 2001, engineers at Siliken Renewable Energy have used the National Instruments graphical system design platform to produce solar panels generating 160 MW of renewable solar energy.

MEASURE IT

Acquire

Acquire and measure data from any sensor or signal

Analyze

Analyze and extract information with signal processing

Present

Present data with HMI, Web interfaces, and reports

FIX IT

Design

Design optimized control algorithms and systems

Prototype

Prototype designs on ready-to-run hardware

Deploy

Deploy to the hardware platform you choose



Siliken Renewable Energy, one of the world's fastest-growing manufacturers of solar panels, trusts NI LabVIEW software for applications ranging from research and development to automated test. Like Siliken, companies around the world implement the NI graphical system design platform to create environmentally friendly products, technologies, and processes. Using modular hardware and flexible software, they are not only testing and measuring existing systems but also creating innovative ways to fix the problems they find.

>> Download the solar power resource kits at ni.com/160

888 293 3091



AMBIENT-LIGHT SENSORS PACK IN FEATURES TO HELP APPLICATIONS GET **smarter,**

greener

AMBIENT-LIGHT SENSORS HELP SMARTPHONES CONSERVE DISPLAY POWER AND IMPROVE BATTERY RUNTIME. NEWER VERSIONS COMBINE IMPROVED FEATURES, SUCH AS SPECTRAL RESPONSE AND DYNAMIC RANGE, INTEGRATED PROXIMITY, AND DIGITAL COMMUNICATIONS. THESE FEATURES ARE READYING THE SENSORS FOR USE IN LARGE-SCREEN AND AUTOMOTIVE DISPLAYS AND SOLID-STATE LIGHTING.

ALSs (ambient-light sensors) have been around for years, but they are now seeing increased use due to the success of smartphones, such as Apple's iPhone. The iPhone uses an ALS to reduce power demands and eke out battery life by adjusting the display lighting for ambient-light conditions. The iPhone also uses a proximity detector, a close cousin of the ALS, to reconfigure itself in response to user actions. Today, smartphones, with their trademark large LCDs, are the main users of ALSs, but future applications can realize energy savings and increased ease of use. For example, a large-screen LCD TV must

adjust its LED backlighting to the ambient light. Similarly, a room with SSL (solid-state-lighting) illumination can change its lighting based on natural lighting or to suit its occupant's mood. Automobile lighting can accommodate day or night driving or reflect the brightness of streetlights, saving power and providing a better user experience.

At its most basic, an ALS consists of a photodiode or a phototransistor. However, a simple light-sensitive semicon-

ductor is insufficient because the ALS must be "photopic," meaning sensitive to the same frequency spectrum as the human eye (Figure 1). Incandescent and HID (high-intensity-discharge) lights emit 50 to 60% of their radiation in the nonvisible IR (infrared) range as heat. According to Oleg Steciw, product-marketing manager for ALS products at Intersil, you should use the HID with the best spectral response you can find. Otherwise, he says, "You'll be in a

room, and, suddenly, the backlight will go haywire because there's some external light source that you can't even see, wreaking havoc."

Werner Mashig, application engineer on Arrow Electronics' lighting team, explains, "[Some] manufacturers put IR-filter [compounds] into the epoxy to filter out the IR light so that the sensor will respond like the human eye."

Another approach is to use multiple photodiodes in the ALS. "One photodi-



The proximity sensor in a 3G (third-generation) iPhone is within the red circle, and the ambient-light sensor is the green part to its left. The iPhone's speaker is the gray, mesh-covered oblong (courtesy iFixit). Light and proximity sensors are often next to a handheld device's speaker because both the sensors and the speaker require access to the outside world. A speaker is a better choice than a microphone because of the humidity that's likely next to a microphone.

ode is a broadband one that sees everything from 300 to 1100 nm,” says Carlo Strippoli, vice president of marketing and sales for TAOS (Texas Advanced Optoelectronic Solutions). “The second diode is a dedicated IR photodiode and serves to monitor the IR reaching the sensor and then subtracting it from the light received at the broadband photodiode.”

Fluorescent-light sources, which are more efficient than incandescent or HID lights, emit almost none of their radiation in the IR range, but they may exhibit a 60-Hz flicker that can cause an ALS to trigger when it’s not supposed to. The newer digital ALSs integrate ADCs that convert the photocurrent to a digital signal to interface to a digital-communication bus. The ADC can serve double duty by filtering out optical noise, such as 60-Hz flicker, through high-resolution sampling. Rohm’s BH17xx series integrates a 16-bit ADC that produces 1-lux resolution over a range of 0 to 65,000 lux. Two measurement-resolution levels allow selection between sampling time and performance. In the high-resolution sampling mode, the ADC filters out optical noise. The lower-resolution mode with its shorter sampling time suits applications such as GPSs (global-positioning systems), in which the light-level changes are dynamic: A GPS system will probably operate in an automobile’s interior or in natural light. The ideal ALS exhibits uniform light sensitivity regardless of the light source.

“Digital is the direction ambient-light

AT A GLANCE

Although smartphones are the major users of ALSs (ambient-light sensors), look for their ability to bring intelligence and power savings to such applications as laptops, TV screens, and automotive and room lighting.

The most important feature for an ALS is its ability to see light as the human eye does. This task requires filtering out the IR (infrared) spectrum.

Vendors are integrating more features, including digital interfaces and proximity sensors, into ALSs.

More complex devices perform RGB (red/green/blue) sensing.

sensing is going,” says TAOS’ Strippoli. “It allows you to put multiple sensors on a single two-wire bus,” such as the I²C (inter-integrated circuit). This feature is especially important for flip phones. A digital bus minimizes the number of wires at the hinged interface where the cell phone flips up.

An analog interface requires at least two wires for every sensor. Analog ALSs are still good fits for some designs, such as those in which the voltage or current output of the ALS directly drives the lighting subsystem, those lacking a microcontroller or an available ADC input, and those low-end designs in which price is the dominating feature (Figure 2).

In the past, ALSs could vary from

part to part in the amount of current a given amount of light produces. Such variability makes it difficult to design for a tight sensitivity range. “The manufacturers are [now] doing a great job of binning the components to give more consistency across the design so there’s not as much variation of the photocurrent,” says Arrow’s Mashig. He suggests looking at the specification for photocurrent versus brightness to check the tightness of manufacturers’ binning.

A low-power lighting system is especially important for battery-powered devices, and this requirement includes the ALS itself. In general, both analog and digital versions of ALSs have a shutdown or sleep mode, during which the sensor operates at approximately 1 μ A. Because of the relative simplicity of analog ALSs, they require less power than their digital counterparts. For example, a representative digital ALS draws 190 μ A in active mode and 1 μ A in power-down mode due to the integration of the ADC; an analog equivalent of the part draws 97 and 0.4 μ A, respectively. However, the total power consumption is comparable to or a little less than that of an analog ALS with a separate ADC.

In addition to an ALS, smartphones often use proximity detectors. Apple’s integration of a proximity detector in the iPhone prompted a move toward making handheld consumer devices more intelligent when interacting with their users (see sidebar “iPhone puts proximity detection on the map and in your face”).

IPHONE PUTS PROXIMITY DETECTION ON THE MAP AND IN YOUR FACE

The Apple iPhone packs several sensors into its slim profile: an ALS (ambient-light sensor), an accelerometer, and a proximity sensor. In addition, the display itself is a giant touch sensor, and that fact could pose a problem when the phone is in use next to a user’s face. Apple solved the problem of inadvertent activation of the screen by including a proximity sensor that detects proximity and turns off the touchscreen when the phone is 3 to 5 cm from a user’s face (Reference A).

REFERENCE

A “iPhone 3G Teardown,” www.ifixit.com/teardown/iphone-3g-s/817/1.




Figure 1 The human eye is sensitive to wavelengths of approximately 380 to 780 nm, peaking at approximately 555 nm. The wider red curve shows the sensitivity of a representative ambient-light sensor.

Because of the close links in both technology and usage between ALSs and proximity detectors, ALS vendors are starting to add proximity detection to the list of integrated features in ALSs. “The ISL29011 drives an external IR LED so that it synchronizes the transmittivity of the LED and then captures the reflection off the object in front of it,” says Intersil’s Steciw. “You want the sensing range to be within about 3 to 5 cm.”

The placement of the IR LED outside the sensor package gives designers more flexibility in where they place the device or what they’ll place it behind, says TAOS’ Strippoli. “The iPhone puts [the proximity detector] behind a glass that blocks about 95% of visible light,” he adds. “So if you use a device that gives you just a [fixed] single output, you get a very low signal.”

For discrete proximity sensors, it’s still common to keep the IR-radiating




Figure 2 Although digital ALSs are now available, analog sensors are still popular for many applications. Microsemi’s Best Eye processing provides a nearly perfect photopic light-wavelength-response curve. The sensor output feeds into a wide-dynamic-range compression amplifier that provides accurate resolution over five decades of ambient light.

LED inside the sensor package. Avago recently introduced the APDS-9120 proximity sensor, which combines a built-in signal-conditioning IC, an emitter, and a detector into a package that offers both analog- and digital-output options. Like Steciw, Strippoli views the power-saving requirements of portable devices driving the trend in packaging proximity detectors along with the ALS but sees it as part of the move toward greener products. He believes that Asian countries in particular are likely to mandate the ability to tell when a viewer is using a large screen or monitor by monitoring proximity.

ALSs in smartphones detect light intensity but provide no information about the color spectrum. A recent development in ALSs is the ability to perform RGB (red/green/blue) sensing, a necessary feature for large-screen LCDs. For the best viewing experience, these

TAKING ADVANTAGE OF LIGHT SENSORS WITH MICROCONTROLLERS RUNNING DALI

By Bobby Wong, NEC Electronics America Inc

In our energy-conscious world, one simple way to reduce energy consumption is by adjusting office lights to take advantage of the available natural light. Light sensors can operate in multiple locations to detect the amount of naturally occurring ambient light. With the appropriate lighting system, users could accordingly adjust office lights to produce the desired amount of total lighting necessary for each area. Sensors have proved that they can dramatically enhance lighting systems—from improving energy efficiency by sensing ambient light to improving color by detecting light output. Although sensors provide the data, the lighting system still needs an intelligent microcontroller to receive and process the data and adjust the lights accordingly.

However, saving energy should not reduce productivity. A smart microcontroller-based lighting system would allow users to override the automatic light-level sensors when necessary and “remember” programmed user settings to enhance the users’ experience.

Although multiarea lighting control, sensor input/processing, and scene setting may sound complicated, the DALI (digital-addressable-lighting-interface) protocol for white-light control in offices and factories already implements many of these features. Companies space these sensors and lights throughout their facilities, and the devices therefore require a network. The DALI network

can control as many as 64 lights with 64 generic controls, such as slider dimmers and sensors. Each area light can store as many as 16 scenes, and each scene stores a digital-dimming level of 0 to 255. When a sensor provides ambient-light input to the microcontroller, the microcontroller can send a DALI command through the network to any of the 64 lights and control them to dim to a specific scene setting. The DALI protocol is also extensible, allowing a supplier to include vendor-specific features for added value. Some microcontrollers have specialized hardware for driving lights from fluorescent tubes to LEDs, and they simplify the support for a DALI network. Unlike discrete light drivers, these microcontrollers can process sensor inputs and intelligently control lights in a wide area using the DALI protocol to produce the optimized amount of light and save energy along the way (Reference A).

REFERENCE

▲ “Lighting-control solutions,” NEC Electronics America, www.am.necel.com/applications/lighting.

AUTHOR’S BIOGRAPHY

Bobby Wong is technical-marketing engineer for the multipurpose-microcontroller strategic-business unit at NEC Electronics America Inc.

displays must match their backlighting to the color temperature of the ambient lighting (Reference 1). The LCD controller uses the RGB ALS output to tune the RGB HB (high-brightness) LEDs to match the ambient lighting: Backlighting for a fluorescent-lit room has a different color temperature from that of a natural- or incandescent-lit room. In addition, as RGB LEDs age, their color changes slightly, calling for an additional RGB ALS in the backlight itself to sense and give feedback to drive the compensation for the LEDs' color change. Intersil, TAOS, and Rohm all offer RGB sensors.

SSL is an emerging application for RGB ALSs. In this application, color sensors provide feedback to a room's lighting-control system to adjust the light intensity, color, and color-temperature output of the HB LED-based luminaires. Lighting-control information is more complex than the simple on/off-light-switch information that room lighting currently uses, and lighting designers must be familiar with communication

protocols. The DALI (digital-addressable-lighting-interface) protocol, which theatrical lighting has used for years, is one possible approach (see sidebar "Taking advantage of light sensors with microcontrollers running DALI").

Automotive lighting also needs ALSs. Night-driving applications have for years used simple photosensors to turn lights on and off, but more complex ALSs optimize cabin lighting for safe driving and for aesthetics, such as colored lighting

and light-intensity variation. Like most other automotive components, ALS specifications must include operation over the wider temperature and vibration range. **EDN**

REFERENCE

Conner, Margery, "The direction of light: Electronic and thermal improvements bring advances to lighting technologies," *EDN*, Feb 5, 2009, pg 26, www.edn.com/article/CA6632368.

FOR MORE INFORMATION

Apple
www.apple.com

Arrow Electronics
www.arrow.com

Avago Technologies
www.avagotech.com

Intersil
www.intersil.com

Microsemi
www.microsemi.com

NEC Electronics
www.necel.com

Optek
www.optek.com

Osram
www.osram.com

Panasonic
pewa.panasonic.com

Rohm
www.rohm-semiconductor.com

Sharp
www.sharpmeg.com

TAOS
www.taosinc.com

Toshiba
www.toshiba.com

Vishay
www.vishay.com

Go to www.edn.com/090806df1 and click on Feedback Loop to post a comment on this article.

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

You can reach Technical Editor Margery Conner at 1-805-461-8242 and mconner@connerbase.com.



Noises removed by Samsung capacitor

Feedthrough Type VLC -Vertically Laminated Capacitor

(0402 X5R 2.2uF 6.3V, 0402 X6S 0.47uF 4V)

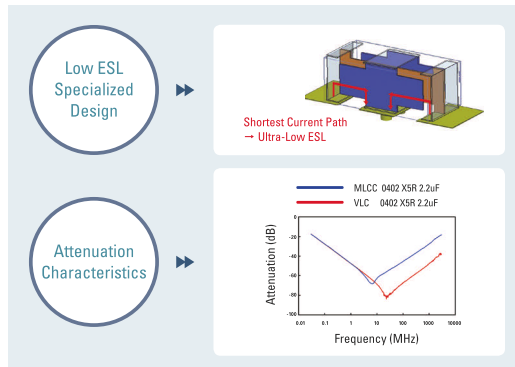
VLC (Vertically Laminated Capacitor) has vertically oriented internal electrodes that effectively reduce the current loop area in the device. Therefore, it provides ultra-low equivalent series inductance (ESL) despite its relatively simple terminal structure. This innovative capacitor is recommended for various package applications (CPU, GPU), mobile phone, and laptop, etc.

FEATURES

- Ultra Low ESL
- Rated Voltage 6.3 or 4Vdc
- X5R or X6S TCC

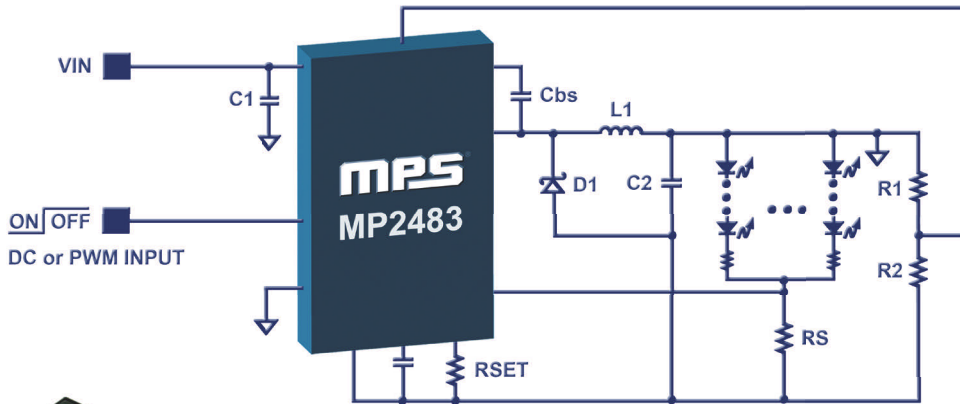
ADVANTAGES

- Excellent Bypassing Characteristic
- Saving Space by Replacing 4 MLCC's
- No Tomb-stoning



Drive Up To Eight 700mA LEDs

2.5A, 50V Programmable Frequency White LED Driver




Get The Most From a Single IC!


MP2483 Features

- 2.5A Maximum Output Current
- Wide 4.5V to 50V Operating Range
- 0.15Ω Internal Power MOSFET Switch
- Adjustable Switching Frequency up to 1.4MHz
- Analog and PWM Dimming
- 0.2V Reference Voltage
- No minimum LED required
- Open Strings Protection
- Output Short Circuit Protection

Competition



MPS WLED Driver



DC to DC Converters CCFL / LED Drivers Class D Audio Amplifiers Linear ICs





SPEAKING OF PORTING SOFTWARE

BY ROBERT CRAVOTTA • TECHNICAL EDITOR


The goal of this hands-on project was to port a common set of software across a variety of processor architectures: two ARM Cortex-M3 ports using Atmel and Texas Instruments processors, a port to a Microchip PIC32 device, and a port to an Atmel AVR32 processor. Each port hit some snags, but they all successfully completed the goal of a working port of the target software. During the project, I realized that the software for the porting effort, Vorbis Tremor, might also make a good candidate for embedded-system developers to consider when exploring how to add audio to their embedded-system design.

I adopted two lessons I learned from a previous project about accelerating software with hardware that spans multiple vendors (**Reference 1**). This project uses a common set of software in each port so that everyone can benefit from identifying the differences and similarities in each effort. The second lesson I learned was to ensure that each vendor provided engineering support so that its development kit can be part of the project. This requirement serves multiple

purposes, most notably not overwhelming me with more than 70 development kits.

For these projects, it is imperative to choose a scale of work that is neither too trivial nor too ambitious. An original candidate for the software to port was benchmark code from the EEMBC (Embedded Microprocessor Benchmark Consortium) because many companies use the benchmark software on their processors, meaning that more compa-

THIS SOFTWARE-PORTING HANDS-ON EXPERIMENT UNCOVERS A POTENTIAL AUDIO DECODER FOR EMBEDDED-SYSTEM APPLICATIONS, ADDING AUDIO OR SPEECH TO THE APPLICATIONS' USER INTERFACES.



```
using namespace std;
string title(wxw/"codehive - Revision: ");
title << _v(SvnRevision);
vent.ship();

void Svn::OnChange(wxw) {
    // pass the event to the editor
    title << "codehive - Revision: " << SvnRevision;
    // ...
}

// ...

title << wxw/" - ";
SubTitle(title);
<< editor->GetTitle();
}
```




Figure 1 The SAM3U-EK evaluation kit supported the port to Atmel's ARM Cortex-M3 processor.




Figure 2 The ATEVK1105 evaluation kit supported the port to Atmel's AVR32 processor.

nies would be able to participate in the project. However, the focus of the project was on the porting effort rather than the optimization of the code and processor architecture, and these benchmarks focus on the processing performance of the core. For these reasons, I eventually abandoned the idea of using benchmark code in favor of using an audio codec because it combines the requirements for real-time performance with a real-world interface for storage, retrieval, and playback of the audio. I also needed no expensive equipment to tell whether the code was running in real time because I had access to sensitive and free signal processors—my ears—to tell whether the audio was playing too fast or too slow or that the processor was missing data.

While exploring the idea of using an audio codec, including MP3, for the porting target, I discovered the open-source Ogg Vorbis audio-compression format and its Tremor library, a fixed-point implementation of the Vorbis decoder. Using a fixed-point decoder would allow more processor architectures to participate in this project.

Borrowing from another lesson I learned from my earlier life as an engineer, I framed the project description to avoid as much bias as possible in how each team approached the porting effort (Reference 2). This type of uncertainty when specifying a project often yields unexpected benefits, and this project was no exception. The project required the porting effort only to provide

a mechanism to store and access audio files and be able to play them on some output. Each vendor could choose any processor and any development tools it wished to complete the porting effort. Each team made different choices in solving the porting challenge, and one of those design decisions highlighted why the Tremor player might be a good candidate for embedded-system developers to consider.

FOUR PORTS

For this project, each engineer was free to select development boards and

tools. The porting effort took an average of a day and a half from choosing the target, understanding the open-source software, and making the necessary changes to the software to complete the port. After that, I duplicated the porting effort over the phone with each team. This approach saved me a lot of time, and it gave me access to the thought process of each person. It also meant that each project was unique rather than a refinement of my own effort with each development kit. In each porting effort, some things worked smoothly, but there was always something that did not proceed the way we would have liked. I will share the hiccups but without specifying with which team it occurred.

Atmel participated in two ports, an AVR32 and an ARM Cortex-M3 (Figure 1). A different team member performed each port, and each took a different approach. The AVR32 port used the ATEVK1105 evaluation kit (Figure 2). Atmel released this new board this year at the ESC (Embedded Systems Conference) in San Jose, CA. We used the AVR32 Studio development-tool set. The audio output went through an adapted DAC for wave playback using software from a previous project that used this peripheral. We performed the port in two stages. The first stage linked the .ogg audio file into the executable file. The second stage accessed the .ogg audio file using code from a FAT (file-allocation-table)-library example through a data-flash device. This two-stage ap-

AT A GLANCE

▣ Contemporary compilers are competent at creating executables of open-source software for embedded-processor targets.

▣ Compilers cannot automate the porting of real-world interfaces and managing dynamic-memory structures. This task remains a vendor's job or a hands-on job.

▣ Embedded development tools must support a complex ecosystem of host and target architectures; this arrangement provides many opportunities for unexpected behaviors to manifest themselves in the tools.

▣ The Tremor audio decoder is a candidate worth considering for embedded-system designers exploring whether to add rich audio data to their interfaces.




Figure 3 The Explorer 16 development board supported the port to Microchip's PIC32 processor.



Figure 4 The DK-LM3S9B96 development kit supported the port to Texas Instruments' ARM Cortex-M3 processor.

proach helps isolate delay sources.

The audio codec uses dynamic allocation, which can be a significant source of delay if you are not careful about external-memory accesses and garbage-collection events in the heap. In the case of the AVR device, the multiply function proved to be an area for optimization in part because it handles big- and little-endian representation and it does not take advantage of the extra hardware resources available on the AVR processor to improve multiplication performance.

In addition to isolating sources of delay, this two-stage approach made me realize that I didn't necessarily need to store the audio stream on an external storage device because an embedded system often does not allow the user to access the data, and rarely will it even change the audio stream during the life of the application.

The Atmel ARM Cortex-M3 port used the new SAM3U-EK evaluation kit. This board is so new that a complete set of driver code and samples were unavailable for all of the peripherals, including a DAC driver. The Cortex-M3 is a new generation of the ARM architecture and does not directly benefit from legacy code from earlier architectures, such as the ARM7. However, the support library for this architecture will grow, especially as more M3 devices from a growing list of vendors become available. In a sense, this project was an early adopter of this board and processor (Reference 3). The project used the

Yagarto (yet-another-GNU-ARM tool chain). The engineer considered other tool chains, such as IAR, but, due to time constraints with a learning curve involving the allocation library, the engineer used Yagarto. The engineer stored and retrieved the audio file from the SD (secure-digital) card using sample open-source code to manage the file system. Onboard NAND flash could also have stored the audio files.

The Atmel M3 port uses a ring-buffer implementation to feed the DAC/DMA engine transfer; this approach differs from the ping-pong buffers the other ports used. The ring buffer allows the buffer tuning to adjust not only the buffer size but also the number of buffers to optimize performance by tracking how many of the buffers were full over time. For example, with a 22-kbps sample, a 2-kbyte buffer resulted in 60 to 70% full buffers, whereas a 4-kbyte buffer resulted in less than half the buffers being full. The porting effort progressed starting from 8-kbps samples. This approach exposed efficiency issues in the allocation of memory. The CPU's usage with the 8 kbps was 10%, but usage shot up to 60% at 22 kbps. The higher bit rates caused more access to external memory, which introduced significant delays. Possible optimizations include changing the dynamic-memory allocation code as well as some manual managing of the heap to straddle external and internal memory.

The Microchip PIC32 port used the Explorer 16 development board (Fig-

ure 3) with a customized board and the MPLab Real ICE (in-circuit emulator). We used the MPLab Academic version for the software-development tool chain. The PIC32 device is pin-compatible with earlier 16-bit devices and uses the same peripheral blocks as the PIC24. The porting used legacy code by recompiling the PIC24 code. We stored and retrieved the audio files with an SD card and played through the PWM (pulse-width modulator) using a ping-pong-buffer implementation.

The Texas Instruments ARM Cortex-M3 port used the DK-LM3S9B96 development kit with the Keil μ Vision3 software-development tool chain (Figure 4). This effort required rewriting the allocation routines to avoid dynamic allocation during playback; this task included explicitly defining a stack and a heap space. We stored and retrieved the audio files with the SD card using sample code. The audio output used I²S (inter-IC-sound) demonstration code, which included volume control and a touch-screen scroll interface. Other optimization options would address the multiplication macros.

PROBLEMS

Each porting effort ran into problems. Some of these problems were early-adopter problems, such as when you are using newly released resources. For example, one of the boards had an earlier version of the firmware that had a problem that the manufacturer fixed in a

later version of the firmware. From this situation, I learned that there should be a straightforward way to update the firmware or version information on the board to avoid sending out a board with a problem that the manufacturer has already fixed.

To make this project more interesting, I used a 64-bit Vista desktop. None of the original port efforts used this machine, so, although this approach did not stop the projects, it did cause some stalls. In one case, we learned that we had to explicitly install the software-development tools as administrators by right-clicking the setup.exe file and specifying “run” as administrators. In two projects, we had trouble with getting my desktop to properly recognize the board through the USB (Universal Serial Bus). In one case, it required finding the 64-bit version of the .inf file in a different directory from the directory in the 32-bit version. In another case, it required adding the missing 64-bit information to the 32-bit version. Apparently, 64-bit Vista has not been a big issue, but I expect that more developers will in the near future be using 64-bit Vista hosts. These types of problems help to illustrate the challenges facing development-tool support teams as they work to support not only several host operating systems, but also different versions of these operating systems.

In one porting effort, the development-tool installation DVD was a blank disk. Fortunately, quickly downloading and online access of all files eliminated that problem, but the problem marred an otherwise-excellent experience with these kits. In another case, the manufacturer had to separately ship a power cord because not all kits included a power cord. The reason for this omission was to help keep the cost of the kits down and to avoid filling your drawers with too many redundant power cords.

A snag occurred when I tried to plug in a serial port between my desktop and the development board. Imagine my surprise when I realized that my computer had a dozen USB ports but no serial port. Two other computers that I recently purchased also had no serial port. You might need a serial port, but do not assume that manufacturers still include them.

Despite all of these problems, the

✚ For related blog posts about embedded processing, go to www.edn.com/blog/1890000189.html.

bring-up on the boards was usually smooth and straightforward. We would set up and power up the board and then verify that the preloaded software was operating properly. After that, we would select some code, compile it in the tool set, load it onto the board, and then verify that it was operating properly. Starting from a known condition and adding one more step into the tool-chain flow in this way helped us identify where a problem might originate and how to address it. Likewise, with the porting effort, adding peripheral ports one at a time or increasing the bit rate in steps helped isolate where logic and performance problems were originating.

NO SILVER BULLET

In this porting, we had to rewrite the input and output of the audio data. The Vorbis implementation uses stdio for handling the input and output. However, the authors of the software recognize that this mechanism is not appropriate for embedded-system applications, so the code includes a callback structure, or callbacks, that allows a developer to provide custom functions for these important I/O functions, including decoding a Vorbis stream from a memory buffer.

A big reason for doing this experiment was to demonstrate that there are no silver bullets for porting code—especially embedded code. None of these ports used an operating system on the target system. As a result, accessing the peripherals required an explicit effort by the developer. This effort might include pulling code from a library or from sample code or, in an early-adopter phase, writing the code yourself.

Additionally, the software may exhibit issues depending on how the memory architecture has changed. Unfortunately, compilers are weak in this area and do not provide as much automation or assistance as developers could use. However, Tremor comes in three main versions. A general-purpose implementation targets processors with access to large off-chip memory, a low-memory version trades memory space for more instructions during execution, and a third version con-

tains low-memory code for processors without byte addressing.

The Ogg Vorbis specification is in the public domain and is free for commercial or noncommercial use, making the format an interesting candidate for embedded-system applications. Developers can independently write Ogg Vorbis software that is compatible with the specification for no charge and without restrictions of any kind. As embedded-system applications expand and the richness of the user interface expands beyond blinking lights and simple buzzers, such as those on coffee pots and washing machines, developers may want to consider the Tremor implementation. Embedded systems need not support file sharing that a rich user-multimedia environment might have to do, and they can take advantage of the static nature of the audio messages they might include to provide a new and cost-effective differentiating feature. As for choosing a processor and the development-support tools, this exercise demonstrated that vendors that are serious about supporting this capability may want to perform an optimization of the Tremor code and offer it as a reference-design implementation. **EDN**

REFERENCES

- 1 Cravotta, Robert, “Accelerate your performance,” *EDN*, Nov 11, 2004, pg 50, www.edn.com/article/CA476908.
- 2 Cravotta, Robert, “Valuing uncertainty,” *EDN*, Jan 5, 2006, pg 38, www.edn.com/article/CA6294179.
- 3 Cravotta, Robert, “Welcome to the jungle,” *EDN*, Oct 30, 2003, pg 39, www.edn.com/article/CA330073.

FOR MORE INFORMATION

ARM
www.arm.com

Atmel
www.atmel.com

EEMBC
www.eembc.org

IAR
www.iar.com

Microchip
www.microchip.com

Texas Instruments
www.ti.com

Xiph.Org Foundation
www.xiph.org

You can reach
Technical Editor
Robert Cravotta
at 1-661-296-5096
and rcravotta@edn.com.






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.



The path of least resistance



...for high density interfaces.

SEARAY™

- 500 I/O open pin field array for routing flexibility
- 7mm to 17.5mm stacks, 10.5 GHz in differential configuration
- Low insertion and extraction forces
- Standard and custom jumper lengths and pin counts

SEARAY™ Right Angle

- Reduced skew and impedance mismatch for micro backplanes

HD Mezz™

- Elevated stack heights up to 35mm
- Rated at 9.5 GHz single ended
- Up to 299 I/Os

HD Mezz is a trademark of Molex Incorporated

DP Array™

- Easier differential routing
- Up to 168 pairs, up to 17mm stacks

samtec
TRANSMISSION LINE SOLUTIONS

www.samtec.com/tls

Students from across the globe gathered at the 2009 IMS (International Microwave Symposium) in Boston June 7 through 12 to demonstrate their innovative approaches to the microwave industry. This year, the IEEE MTT-S (Microwave Theory and Techniques Society), which sponsors the IMS, challenged students to participate in three design competitions, as well as the symposium's traditional student-paper competition. Each of the events pushed students to work beyond their fundamental training with the technology to create designs that function successfully at a professional level.


David Yu-Ting Wu, a student at the University of Waterloo (Waterloo, ON, Canada, uwaterloo.ca), presents his winning prototype for the high-efficiency-power-amplifier-design competition to a group of MTT-S judges (courtesy MTT-S).

BY JENNIFER KEMPE • CONTRIBUTING EDITOR

IMS HIGHLIGHTS

STUDENT INGENUITY

COMPETITIONS DRIVE STUDENT
INTEREST IN MICROWAVE THEORY
AND TECHNIQUES.



MTT-S organizers say that the student-paper competition is one of the most active segments of the IMS technical program. This year, students—most of them at the graduate level—submitted 235 papers, accounting for 25% of all the submissions to the symposium. The technical-program committee reviewed each paper and judged the submissions based on quantitative content, presentation quality, and MTT-S interest in the subject matter. From all the submissions, the judges selected 24 papers as finalists for presentation at the symposium. Student papers this year addressed topics including reconfigurable microwave filters, millimeter-wave antenna arrays, outphasing transmitters, dielectric-resonator filters, substrate-integrated waveguides, UWB (ultrawideband)-over-fiber systems, power amplifiers, cognitive radio, RF MEMS (radio-frequency microelectromechanical) switches, and implantable medical devices for transcutaneous transmission.

James Komiak of BAE Systems (www.baesystems.com), who volunteers to help organize the student program as a member of the MTT-S technical committee, considers the student-paper competition to be an asset to both the students and the professionals who work with them. “The quality of the student papers has been just amazing in what they’re achiev-

AT A GLANCE

▣ The student-paper competition at the MTT-S (Microwave Theory and Techniques Society)-sponsored IMS (International Microwave Symposium) benefits both the students and the professionals who work with them.

▣ The first student-design competition aimed to engage students with practical applications of design rather than theory.

▣ This year, the symposium offered competitions for designing a packaged UWB (ultrawideband) filter and a low-noise amplifier.

▣ Most of the students employed a trial-and-error approach to adjusting their designs to meet the official specifications.

ing on a professional level,” he says. “The program allows them to network and make contact with the judges, so it’s an important step in their academic and future professional career.” The opportunity for professional achievement associated with the paper competition has led to the success of the program, as participation in recent years has grown.

INTEREST IN MICROWAVE

Although the symposium has offered the student-paper competition

since 1984, the student-design competition is only five years old. Members of the MTT-S technical committees observed a lot of student interest in computers, whereas they were neglecting microwave technology, explains Komiak. Looking to revitalize student attendance at the convention, the technical committees launched the first student-design competition to engage students with practical applications of design rather than mere theory. Competitors designed, constructed, and measured a highly efficient power amplifier at a frequency of 1 to 20 GHz and with an output-power level of 5W or more.


“The students enjoy [the challenge] because it’s something concrete that proves their work and their achievements,” says Komiak. After winning honorable mention in this year’s competition, Paul Saad, a doctoral student at Chalmers University of Technology (Göteborg Sweden, www.chalmers.se/en), is proud of his achievement at the symposium. “It’s good for my career, and it’s also good to help my university be known in this field,” he says.

FILTER AND AMP CONTESTS

After the success of the first year, the technical committees have widened the program, incorporating new topics for the design competitions. The committees gear these topics toward the practical applications that working engineers face every day and represent an opportunity for students to demonstrate their preparation for a professional career. This year, the symposium offered two new design events—a competition for a packaged UWB filter and a competition for a low-noise amp.

The contestants have only five months to plan, build, and test their designs, so many teams have trouble meeting all the specifications of the competition. The leader of the winning team in the packaged-UWB-filter-design competition is Yi-Ming Chen from the National Chung Cheng University (Minhsiung, Taiwan, www.ccu.edu.tw/eng/e-index.php). “[Our team] started from theory and then went to formula,” he says. “But our first model didn’t meet the specifications, so we had to keep finding the problems and ... solutions.”

Most of the students involved with the design competitions used this trial-and-error method as they adjusted their de-



Students and professionals work side by side during the packaged-UWB-filter-design competition at the 2009 IMS (courtesy MTT-S).



Professor Rainer Kronberger looks on as his students Christian Musolff and Andreas Neurberger of the Cologne University of Applied Sciences in Germany present their design to Roger Kaul (seated) during the low-noise-amplifier-design competition (courtesy MTT-S).

signs to fulfill all of the official specifications. Roger Kaul, a retiree of the Army Research Laboratory (www.arl.army.mil) who judged the low-noise-amplifier-design competition, commends the students for the professionalism they displayed in both their work and their attitudes toward the competition. "It's easy for students to make excuses if something doesn't function properly," he says. "Today, no one complained at all. ... This is top-level stuff that they're doing, and the devices they came up with today were very impressive."

Overall, the competitions allow students worldwide to witness today's innovations within the microwave industry. Honorable-mention winners for the low-noise-amp-design competition are Christian Musolff and Andreas Neurberger, undergraduate students at the Cologne University of Applied Sciences (Cologne, Germany, www.international-office.fh-koeln.de/english/). "We did not know that much of the technology we saw here today even existed, so we were quite happy to be involved to see the variety of companies and products available out there," Musolff says.

Conversely, many major companies looked at the students as resources for

future engineering developments, offering their support for the program by donating equipment and personnel for the competitions. "It's important for students to get involved early because they will have to step up when guys like me retire," say Komiak. "They are the future

of the industry, and it's up to [them] to keep it all moving forward." **EDN**

A version of this article appears on the Web site of EDN's sister publication, *Test & Measurement World*: www.tmworld.com/article/CA6666612.

IMS2009 STUDENT-DESIGN-COMPETITION WINNERS

HIGH-EFFICIENCY-POWER-AMPLIFIER DESIGN

Winner: David Yu-Ting Wu, University of Waterloo

Advisor: Slim Boumaiza, PhD

Honorable mention: Paul Saad, Hossein Mashad Nemati, and Mattias Thorsell, Chalmers University of Technology

Advisors: Christian Fager and Kristoffer Andersson

Honorable mention: Junghwan Moon and Jungjoon Kim, Pohang University of Science and Technology (Gyubuk, Korea, www.postech.ac.kr)

Advisor: Bumman Kim

LOW-NOISE-AMPLIFIER DESIGN

Winner: Rohit Gawande, University of

Virginia (Charlottesville, VA, www.virginia.edu)

Advisor: Richard Bradley

Honorable mention: Christian Musolff and Andreas Neurberger, Cologne University of Applied Sciences

Advisor: Rainer Kronberger

PACKAGED-ULTRAWIDEBAND-FILTER DESIGN

First prize: Yi-Ming Chen, National Chung Cheng University

Second prize (tie): Alexander Stark, Technical University of Hamburg-Harburg (Hamburg, Germany, www.tu-harburg.de), and Liang Han, École Polytechnique de Montréal (Montreal, PQ, Canada, www.polymtl.ca)

Techniques for implementing high-performance processor cores

EVERY STEP OF THE INTEGRATION PROCESS GIVES THE DESIGN TEAM OPPORTUNITIES TO MOVE CLOSER TO ITS DESIGN GOALS.

In today's key applications, the performance and power consumption of processor cores are critical issues in the overall success of the design. Yet the processor is often a third-party synthesizable core. Don't despair: You can meet aggressive design goals with a synthesizable core. You can use any of several techniques during integration of processor cores to achieve the best performance and power. You must consider your technology and architectural choices and use some design-optimization techniques. A real-world example of a MIPS (www.mips.com) 24Kc processor core demonstrates the effectiveness of the optimization.

The first major steps of a processor design are the architectural and technology choices. These include processor selection and configuration, along with selection of the process technology, macros, and standard-cell library. The requirements of the application determine these selection criteria. The next step is the implementation, which should yield the best performance and power given the process technology, architecture, and IP (intellectual property). Design teams achieve the best results through the enhancement of standard design flows and added technology, such as augmented cell libraries. For simplicity, this article omits the discussion of techniques that introduce new operational modes into the processor behavior, such as DVFS (dynamic voltage and frequency scaling) and substrate bias. It also does not cover techniques

such as length modulation that rely on manufacturing technology to optimize processor specifications.

ARCHITECTURE AND TECHNOLOGY

The first step is to choose the best processor for the application. Many choices are available from several vendors, including MIPS, ARM (www.arm.com), Tensilica (www.tensilica.com), and ARC (www.arc.com). Also, custom processors are available from vendors such as Freescale Semiconductor (www.freescale.com) and Renesas Technology (www.renesas.com). Your choice depends on many factors, such as application, cost, performance, verification and risk, and infrastructure cost. The first consideration should be the application. Some processors may work best in certain applications. Also, the application may require additional IP that another team has configured to work with a specific processor. You must then consider cost. Because you purchase the IP, the cost of the required IP and license terms are factors when selecting the processor. You must also examine the performance range. Different cores can reach different performance targets. For example, if the application calls for a high-speed clock, you may find only a limited number of possible processor choices. An application may require a small processor, or it might require a specific balance of power, performance, area, and schedule.

You must assess the verification effort and risk to the design.

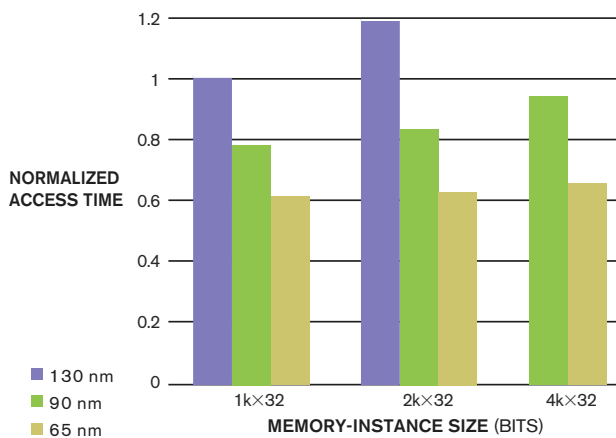


Figure 1 The percentage difference increases as memory size increases for the 130-, 90-, and 65-nm process nodes.




Figure 2 The normalized frequency and power specifications from a commonly used processor core use the general-purpose library and process implementation as the baseline.

Some cores have been in production longer than others, and, depending on the amount of risk a project can afford, using a more mature core might make more sense than using a leading-edge technology. A related question is available infrastructure. The available instruction set, development systems, debugging environment, and current expertise are all factors in choosing a processor. If the design team has experience with a processor or has a developed software infrastructure, that processor may make more sense for the design.

Once you have selected a processor family, you usually need to configure cache memories, as well as any scratchpad or tightly coupled memories. Depending on the size of these memories, they could potentially limit the performance of the end design if the memory paths become critical factors in the design. Some designs organize memories in banks, which dictate the size of the memories the design uses. For example, a 32-kbyte cache could be four banks of 1k×64-bit memories.

Size and organization influence performance. **Figure 1** shows the clock-to-data-output delay of several memories from one compiler. This data comes from data sheets that a memory compiler for a high-performance memory generates. Each group represents a memory size and includes speeds for

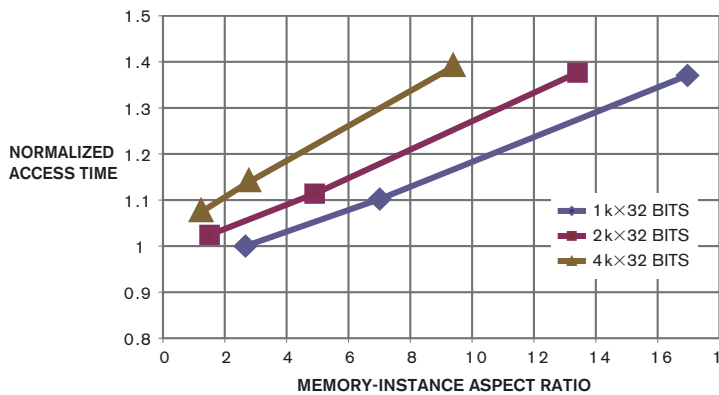


Figure 3 Memory aspect ratio can affect the access time of different memory sizes.

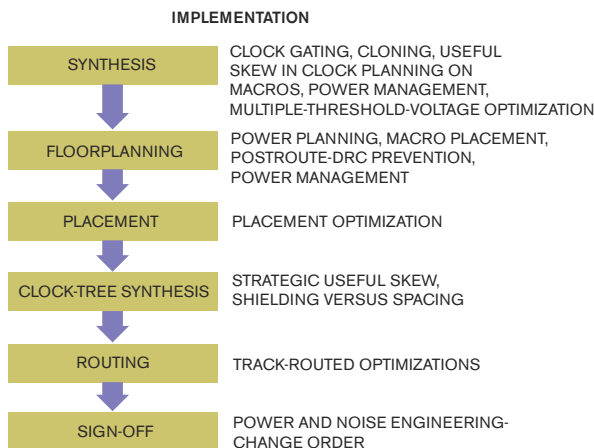


Figure 4 Additional techniques for optimizing processor operation supplement the standard chip-design flow.

the 130-, 90-, and 65-nm process nodes. As the memory gets larger, the delay increases. For example, at the 90-nm node the delay increases 20% from a 1-kbyte memory to a 4-kbyte memory. Typically, the clock-to-memory-data-output path is the critical path, and degradation of 20% could define the final clock frequency. Some soft-processor cores offer options, such as a floating-point unit, instruction-code compression, and clock gating. You can configure these options in the core and allow a tighter implementation for the application.

Selecting the right technology can determine the ultimate performance for an embedded processor core. Trade-offs between cost, power, and performance must happen when you select the target process. For example, before moving to a faster and more expensive process, you should consider advanced optimization techniques if you're nearing your target performance on a given process. Some techniques could yield a 5% improvement, which might be enough to keep the implementation at a lower-cost process. During process selection, you consider such issues as process node or feature size and general-purpose, low-power, triple-gate-oxide, and similar options. You must also consider adjusting certain process parameters, such as threshold voltage and oxide thickness, to meet your requirements. In addition, you can adjust operating voltage to trade off leakage versus performance.

The selection of the standard-cell library also factors into the performance of the processor core. Again, you must consider trade-offs between cost, power, and performance. If the processor core uses a different cell library, you must consider how difficult it would be to integrate the core with the rest of




Figure 5 In the launch side of a successful skew, decreasing the delay, D_1 , of the clock buffer or removing the buffer improves the critical path, S_2 , by that delay delta.




Figure 6 In a skew for capture paths, S_3 improves as the delay, D_2 , increases. The paths using the clock buffer as a launch signal degrade by the same amount, however.




Figure 7 Performance is the most important factor in this implementation of a 90-nm MIPS 24Kc processor with 32-kbyte data and instruction caches.

the chip's fabric. As in process technology, you may need to consider a range of costs, capabilities, and options.

Figure 2 shows a comparison of optimization strategies, libraries, and processes. The normalized frequency and power specifications from a commonly used processor core use the general-purpose library and process implementation as the baseline. You can see the trade-offs between power and frequency; for example, obtaining a frequency boost of 50% results in an increase in the power consumption by more than 50%. Selecting the memories for the processor's caches can also determine the ultimate performance. For example, in the MIPS 24Kc, a memory path's speed dictates the processor's best performance. Again, cost and availability are factors to consider when choosing an IP vendor. Vendors also offer flavors of memories that range in performance and power numbers. Selecting the best fit for the processor could define the final instructions-per-watt count.

Memory compilers offer some input parameters that could trade off footprint versus speed. For example, a shallow multiplexer on the decoder would result in a fast memory but might also yield an aspect ratio that does not fit into the floorplan. Memory aspect ratio can affect the access time of different memory sizes (**Figure 3**). Memories with a higher aspect ratio usually have a wider multiplexer and implementation overhead, resulting in slower accesses. However, the memory area is also an important factor, and you may need to consider trade-offs for the best performance with a workable area and aspect ratio in the floorplan.

IMPLEMENTATION STAGE

Once you select and configure a processor, choose the macros and the process technology, and settle on the standard-cell library, you can begin implementation. Each processor ven-

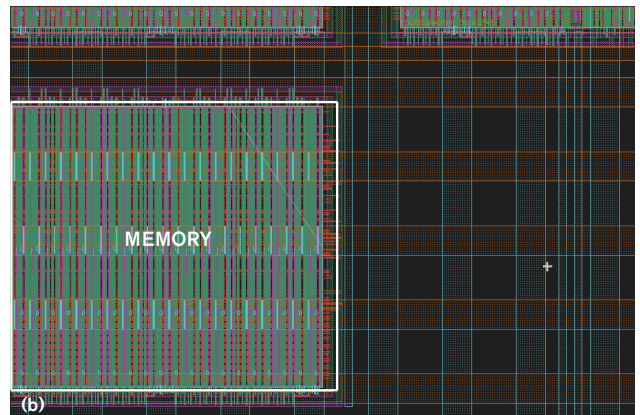
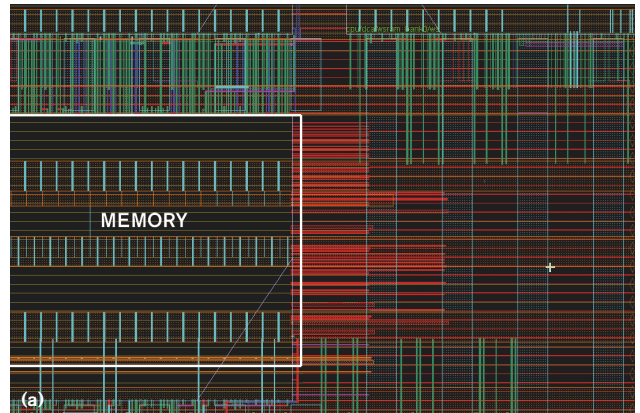


Figure 8 The power wires (shown in red) that connect the memory's pins extend far into the standard-cell area (a). The technique limits the wires to the vicinity of the memory, and the wires don't extend into the cell area (b).

dor provides a reference-design flow that provides results for general-purpose applications. If you need the best result, however, you must supplement certain steps of the flow with additional techniques (**Figure 4**). You can optimize for multiple threshold voltages, implement DFM (design-for-manufacturing) technology, or use other techniques at each stage in the design flow. You can apply library augmentation, such as the CoreMax timing-optimization tool from Open-Silicon (www.open-silicon.com), throughout the design flow. This tool replaces standard digital library cells with optimized cells having fewer logic levels and improved timing performance (**Reference 1**). You can create new cells as full-custom cells or derive them from standard cells and use them throughout the flow as normal standard cells. By using the available standard cells, you can quickly perform the physical implementation of these cells, reducing the verification effort. It is usually best to create these cells for a given process technology and standard-cell library and strategically use them throughout the design flow.

The first step is to integrate the configured HDL (hardware-description language) from the processor vendor's kit with the rest of the design. During integration, you can add global power-management modes in the logic surrounding the processor, for example, to allow the application to control the processor's power consumption based on performance demand. Various




Figure 9 With further analysis of the slack distribution, the optimization moves a large number of paths to achieve a better result.

options are available for implementing the datapath elements of the design, depending on the synthesis tool. For example, you could use Synopsys (www.synopsys.com) DesignWare to improve implementation for arithmetic elements in the design. Optimization at the RTL (register-transfer level) can produce good results and feedback on whether the performance targets you set for the design are achievable. Usually, you margin the performance at this early stage to allow for effects you are not yet modeling, such as noise and parasitic delays.

You can use augmented library cells in limited numbers during synthesis. Logic optimization during synthesis consists of mapping and restructuring, and these steps could benefit from the additional variety of complex gates. Using the augmented library technique makes available a richer range of drive and load characteristics for a given complex gate, resulting in better design performance. At this stage, you also model the clock network and investigate intentionally applying skew to some blocks. Critical memory paths, for example, can benefit from skewing their clocks. By intentionally performing this useful clock skew early in the flow, you allow the synthesis tool to recognize and optimize the critical paths instead of seeing only a memory path with little or no logic.

Processor-core vendors provide recommended floorplans. You still must optimize these floorplans for each implementation, however. You may need to change the aspect ratio and fit of the processor into the full-chip floorplan in a different way from that of the vendor's recommendation. You must consider how the macros interface with the power grid, the location of pins connecting the processor to the rest of the chip, power requirements, and other factors. Consider how the floorplan will affect downstream-flow steps. The power supply to the macros should not interfere with the standard-cell rows, and you may

need additional power rails to connect decoupling capacitors and substrate ties. If you can avoid a cleanup step by properly using floorplanning, the design will be able to reach a better outcome.

Placement optimization no longer consists of simply placing a netlist; it also includes sizing and buffering. The vendor supplies a placement flow that includes some optimization commands, but more optimization techniques are available at this step, and the use of the additional cells from the augmented library at this stage can yield better results. Applying optimization to a critical range of paths instead of only the most critical path can prevent noncritical paths from becoming critical later in the flow. The critical range is by definition the group of critical endpoints that fall within a certain delta of the critical slack.

You can synthesize the clock tree once you have completed placement. Two popular clock-tree-synthesis strategies are balanced-tree and skewed-tree synthesis. The balanced-tree approach attempts to minimize the delta or skew between endpoints, and the skewed-tree approach attempts to skew the endpoints in a way that minimizes the design's timing slack. Both methods support any clock skews that you apply before clock-tree synthesis. For example, if you skew the clocks for the memories during the logic synthesis and placement steps, clock-tree synthesis preserves these skews.

You can obtain good results using a balanced-tree approach along with performing some skews on the memories before clock-tree synthesis. You follow this step with the strategic application of useful skews after clock-tree synthesis to adjust any difficult paths. The skewed-tree method, even if you constrain its use, may result in the addition of too many clock buffers, which congests the design. Another choice in clock-tree synthesis is how to isolate the clock routes from the other

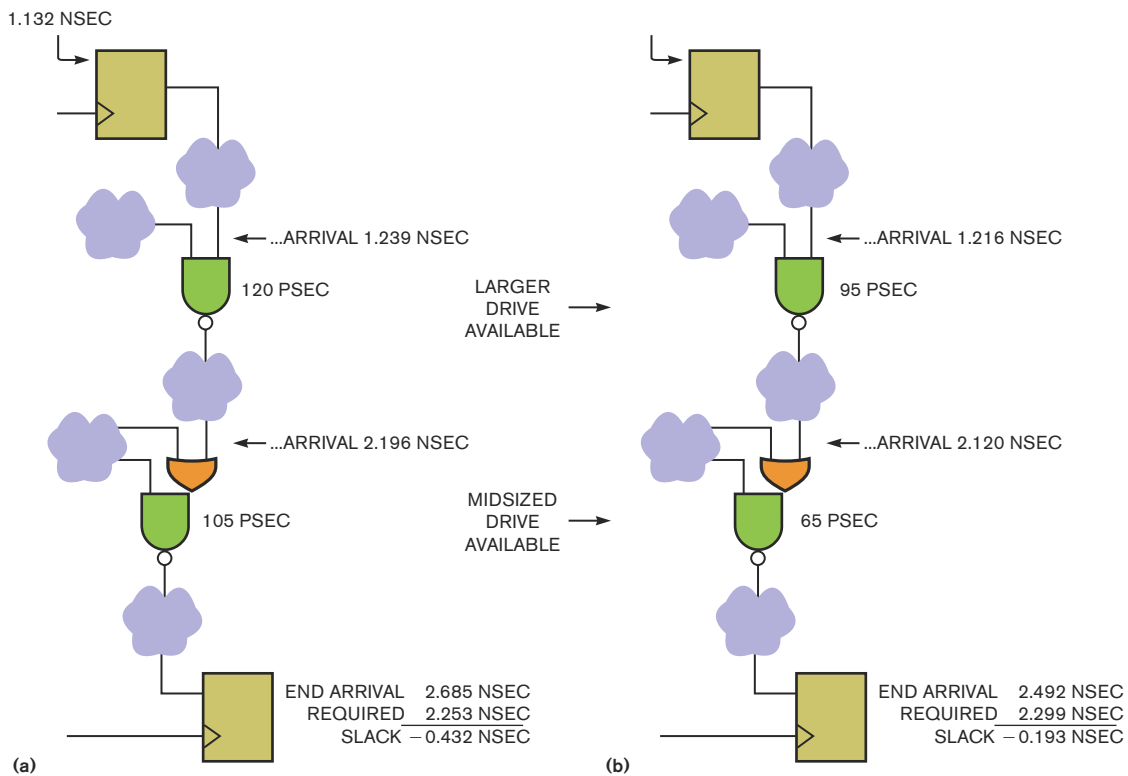


Figure 10 These examples of a critical path before (a) and after applying the augmented library (b) illustrate how the use of augmented library cells can improve path timing.

signal routes. Clock signals are noisy because they switch during every cycle. To minimize the noise effects, you must shield the clock routes or at least space them away from other routes. The processor-implementation flow should include an option for choosing between shielding and spacing of the clock routes to minimize noise coupling to nearby signal lines.

Figure 5 shows the launch side of a successful skew. If you decrease the delay, D_1 , of the clock buffer or if you remove the buffer, the critical path, S_2 , improves by that delay delta. This approach also degrades the timing on all the capture paths sharing the clock buffer, S_1 . These capture paths must not be critical to allow this transformation. **Figure 6** shows the capture side of the critical path, S_3 , which improves as the delay, D_2 , increases or if you add a buffer. The paths using the clock buffer as a launch signal degrade by the same amount, however.

After completing placement and clock-tree synthesis, you can route the design. The router should be aware of signal-integrity effects and timing and design rules. Most flows also include a round of postroute optimizations. Further optimization may be necessary after routing to close timing. The use of an augmented library at this step helps to more efficiently close timing by providing the proper size of cells. You can also apply additional useful clock skew by removing, sizing, and adding buffers to the clock tree.

EXAMPLE PROCESSOR

Performance is the most important factor in this implementation of a 90-nm MIPS 24Kc processor with 32-kbyte data

and instruction caches (**Figure 7**). The design optimization and trade-offs maintain the best performance without unreasonable power consumption. The design uses high-performance memories and standard cells. The design implementation begins with a recommended floorplan from MIPS, using the supplied design flow for Magma Design Automation's (www.magma-da.com) BlastFusion. The cache memories reside in banks along the perimeter of the block. The floorplan needs a denser power grid to meet IR (current/resistance)-drop requirements for the project, so the memory placement requires adjustment to prevent trimming of the power grid. The memories connect directly to the grid only if the power rails align properly, so adjust the memories with this consideration in mind.

To minimize the number of power-DRC (design-rule-check) issues, the design uses a series of rails and small power rings around the memory area. This approach allows BlastFusion's power router to connect the memory's tie-high/tie-low pins without creating metal wires in the standard-cell rows. This extra step prevents short circuits between the cells and the power wires in the standard-cell rows. The power wires that connect the memory's pins extend far into the standard-cell area (**Figure 8a**). The technique limits the wires to the vicinity of the memory, and the wires don't extend into the cell area (**Figure 8b**).

Constructing the clock tree with useful skew causes the appearance of many clock buffers and inverters, so you should try using a balanced tree instead. Apply small amounts of skew to the memory clocks to balance critical paths in the design. Magma's clock-tree synthesis produces a more manage-

able number of clock-tree cells. Once you synthesize the clock tree, you can complete routing using the BlastFusion flow, which the processor-implementation kit includes.

At this point, you can further optimize each major step, including “fix clock” and “run route track.” These improvements carry over to the next step. Running more optimization at early stages saves time when closing the design’s timing after the routing step “fix wire.” Also, applying optimizations after turning on on-chip-variation margins allows the optimizer to recover from the change in the design’s slacks before moving on to the next step in the design flow. These optimization techniques use built-in Magma commands, such as “run place timing—size” and “run timing optimization,” along with custom scripts that employ gate sizing and useful skew transforms.

New cells in the standard-cell library use Open-Silicon’s CoreMax timing-optimization flow. These new cells provide a richer selection for optimization and find use in placement optimization and postroute-timing closure. Combining all these techniques yields a baseline frequency of 518 MHz versus an optimized frequency of 550 MHz. At the baseline frequency, there is 100 psec of jitter; 8% on-chip variation; and worst-case-corner, worst-case-parasitics, and signal-integrity issues. At the ideal baseline frequency of 562 MHz and ideal optimized frequency of 588 MHz, there is no jitter; no on-chip variation; and no worst-case-corner, worst-case-parasitics, or signal-integrity issues.

With further analysis of the slack distribution, the optimization moves a large number of paths to achieve a better result (Figure 9). The approach uses thousands of endpoints to move the worst slack. You must use automation to evaluate and improve these paths. Also, the use of augmented cells makes a significant difference in optimization (Figure 10).

Applications that use embedded processors can benefit from even the smallest performance boost. This article presents some techniques that squeeze the best performance from a design at

Go to www.edn.com/ms4334 and click on Feedback Loop to post a comment on this article.

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

a given process node. Important factors that determine performance include the selection of the processor, the process technology, the choice of standard-cell library, and the configuration and size of the memories. You must make trade-offs between performance and power when choosing optimization variables.

You can use several techniques during design implementation to optimize performance. These techniques include using custom optimization tools and extra scripting in the standard design flow. During clock-tree synthesis, you can implement the clock tree as a balanced tree or as a skewed tree. You can use targeted application of useful skew to optimize the clock tree to minimize critical paths. Using these techniques can result in better performance without moving to a more expensive process node or library. **EDN**

REFERENCE

Roy, Rob, Debashis Bhattacharya, and Vamsi Boppana, “Transistor-Level Optimization of Digital Designs with Flex Cells,” *IEEE Computer*, Volume 38, Issue 2, February 2005, pg 53, www2.computer.org/portal/web/csdl/doi/10.1109/MC.2005.74.


AUTHORS’ BIOGRAPHIES

Gregory Schulte is a software architect at Magma Design Automation. He previously was employed at Open-Silicon, where he worked to bring new technologies to enable state-of-the-art design implementation.

Vamsi Boppana was the senior director of technology for Open-Silicon Inc, where he focused on next-generation technology. He has authored or co-authored more than 40 technical articles and papers and has six granted patents. He has served on the program committee and as a session chair for several computer-aided design and test conferences, including the International Conference on Computer-Aided Design. Boppana has a bachelor’s degree in computer science and engineering from the Indian Institute of Technology (Kharagpur, India) and master’s and doctorate degrees in electrical and computer engineering from the University of Illinois—Urbana-Champaign.

PICO

New High Voltage Hi Power Regulated DC-DC Converters



High Voltage, Isolated Outputs
100-500 VDC

Output Voltages from 5VDC to 500VDC
High Power: to 50 Watts,
Efficiency to 90%

Miniaturized Size Package:
2.5" X 1.55" X 0.50"

Two Standard Wide Input Ranges

Safe: Short Circuit, Over/Under Voltage, and Over Temperature Protected

Options Available: Expanded Operating Temperature, -55 to +85°C Environmental Screening, Selected from MIL Std. 883

Ruggedized for Operation in Harsh Environments

External Bias Control: For Charge Pump Applications

Custom Modules: Available to Optimize Your Designs, Special Input or Output Voltages Available

PICO’s QP Series compliments our 650 plus existing standard High Voltage Modules. Isolated, Regulated, Programmable, COTS and Custom Modules available to 10,000 Vdc and other High Voltage to 150 Watts!

www.picoelectronics.com

E-Mail: info@picoelectronics.com

send direct for free PICO Catalog
Call Toll Free 800-431-1064

PICO Electronics, Inc.
143 Sparks Ave, Pelham, NY 10803-1837

DRAM technology for SOC designers and—maybe—their customers

AN UNDERSTANDING OF DRAM TECHNOLOGY HAS BECOME CRITICAL TO ANYONE DESIGNING CONSUMER-ELECTRONICS SOCs.

Consumers have come to expect—and even demand—the full benefits of the convergence between computing, communications, and digital-media technologies that we’ve all been predicting for the past 15 years. OEMs therefore produce a bewildering array of advanced electronic systems in a variety of form factors, feature sets, and prices to fulfill consumer desires. SOC (system-on-chip) devices enable much of this convergence by integrating increasing amounts of performance at acceptably low costs.

Although many differences exist between the processing requirements of the computing, communications, and digital-media components in converged SOCs, one commonality is the reliance upon external memories, particularly DRAMs (dynamic random-access memories), to provide high-bandwidth stor-

age for data that the SOC processes. For many consumer SOCs, the performance and cost of the overall system depend upon the efficiency of the communications between the SOC and the attached DRAM. This dependence is particularly strong in SOCs that process high-definition digital-video streams, such as HDTVs (high-definition televisions), digital-cable or satellite STBs (set-top boxes), DSCs (digital still cameras) and camcorders, and advanced multimedia and smart mobile telephones.

Classic computer-architecture texts describe a computer as comprising three classes of hardware: processing, memory, and I/O (input/output) devices. In consumer-electronics equipment, you find hierarchies of these same components. Although the phrase “system on chip” implies the integration of all a system’s functions into one device, SOCs typically implement most of the user-visible processing portions of the system but rely on ex-

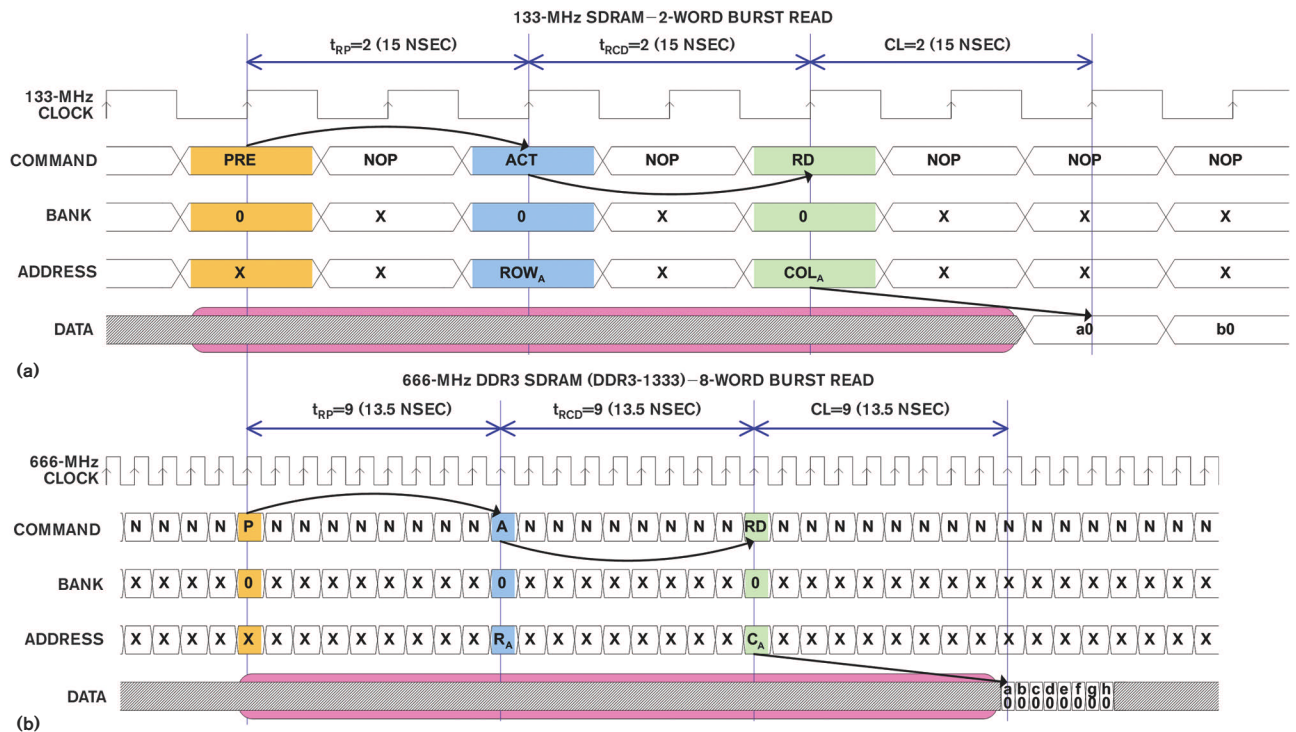


Figure 1 Sample timing diagrams from synchronous DRAMs show differences in various vintages: (a) a 2-word burst read for a 133-MHz SDRAM and (b) an 8-word burst read for a 666-MHz DDR3 SDRAM.

word—to or from the sense amplifiers, using a column address. Because accessing the sense amplifiers is much faster than accessing a new row of cells, the latency of column accesses within a page is much quicker than opening a new page. Before the chip can access another page in the same array, it must precharge the array, or close the page—that is, load the data values stored in the sense amplifiers back into the cells of the open page. While the chip opens and closes a page, the system can access no data from pages in that array at the DRAM interface, which results in a loss of data throughput.

Most modern DRAMs contain banks, each of which contains multiple such arrays. This organization allows data access to one bank to occur simultaneously with page-closing and -opening operations on other banks. In this way, the DRAM can hide the throughput penalties of page operations by exploiting bank-level parallelism in the access patterns from the system.

In the 1990s, vendors introduced SDRAM (synchronous DRAM) with clocked interfaces that enabled higher data rates than previous asynchronous DRAMs. In recent years, this trend has continued with DDR (double-data-rate) SDRAM devices that send data on both edges of the clock, enabling even higher data rates. SDRAM architectures support the concept of burst-mode accesses, in which a column access results in the transfer of several interface words of data. The SOC's memory controller programs the burst-length value into the SDRAM's mode register to choose between burst lengths of 1, 2, 4, or 8 words. Burst-mode accesses are important because they free up the DRAM command interface to transmit precharge and activate commands for other banks while one bank sends or receives data.

Figure 1a shows a timing diagram for a burst-read access to a 133-MHz SDRAM. This access causes a page miss in Bank 0, so the memory controller issues a precharge command to close the open page. After the number of clock cycles necessary to cover t_{RP} —the minimum time between the precharge command that closes a page and the activate command that opens a page in the same bank—the controller activates the new page at Row 8 in Bank 0. After at least t_{RCD} —the minimum delay between the activate command, which opens a row in the bank, and the read- or write-column command, which moves data in the same bank—the desired page opens, so the controller issues a burst-read command to Column A in Bank 0. Two cycles later, the SDRAM returns the requested data, which is two words, a0 and b0, because the SDRAM's burst-length value is 2. Bank 0 is busy closing and opening pages during the entire period **Figure 1a** depicts in pink. You lose these data cycles unless the controller can schedule commands to other banks that use the data bus during this period.

Achieving the lowest cost per bit of storage requires using relatively large arrays of bit cells with resultantly slow array circuitry. The fundamental operating frequencies of SDRAM arrays have not kept pace with technology scaling. Instead, improvements in DRAM bandwidth result from the prefetch architecture that DDR SDRAMs introduced. The prefetch architecture

TABLE 1 SDRAM PREFETCH VALUES

Architecture	Prefetch degree	Efficient/available burst-length values
SDRAM	1	1, 2, 4
DDR1 SDRAM	2	2, 4, 8
DDR2 SDRAM	4	4, 8
DDR3 SDRAM	8	8

relies on choosing DRAM-interface words that are narrower than the DRAM-array words. If the ratio of the array-word size to the interface-word size is N, then interface words can operate at an effective speed N times that of the array without loss of throughput. However, the minimum efficient access size for the DRAM array is the array word, so the minimum efficient interface burst length becomes N.

The degree of prefetch is the major architectural differentiator in SDRAM technologies. **Table 1** shows the current prefetch values for SDRAMs. The minimum efficient access size for DDR3 devices is a burst of 8 interface words. Unfortunately, this size often proves troublesome for consumer SOCs.

Figure 1b shows the timing differences between the 133-MHz SDRAM device and a recently announced 666-MHz DDR3 SDRAM. Although the latency of a page miss at 40.5 nsec has improved by only 10% over the other device's 45 nsec, the peak data throughput has increased by a factor of 10: 1333 versus 133 Mbps/pin. Note that, even though the burst length is now 8, the entire burst of DDR-read data takes less time—6 versus 7.5 nsec—than a single word in the SDRAM case. Note also that the time when Bank 0 is busy is nearly as long, and it will thus take more DDR bursts to cover this busy time than the SDRAM requires.

CHALLENGES IN SOC DRAMs

Because the operating rates of the DRAM arrays have increased more slowly than those of the interface data, each of the DRAM-bank operations has latency measuring many more interface data cycles for DDR3 than did earlier SDRAM devices. To keep the DRAM-data interfaces fully occupied for high efficiency, the SOC's DRAM controller must overlap precharging and activation of pages in other banks while accessing data based on burst commands to an open page in a first bank. This requirement means that the DRAM controller manages a pipeline of DRAM commands, and this pipeline becomes deeper as DRAM technology advances.

Figure 2 shows a sample schedule that a DRAM controller could implement to maintain high throughput during Bank 0's page miss. The schedule relies on cycling through five DRAM banks, assuming a page miss every other burst. This situation allows consecutive 8-word reads to column addresses CA and CI in the same page. To keep the schedule full and therefore minimize the efficiency losses that may result from page closing and opening, the DRAM controller must manage 10 column commands plus five precharge and five activate commands in a pipeline. For the SOC designer, managing such deep memory pipelines requires more complexity in the DRAM controller and the on-chip interconnect, as well as implementation of scheduling logic that feeds the controller.

Although pipelining the DRAM system maintains high efficiency and throughput, it also increases latencies for processors and other initiators on the SOC because a new request must normally wait for the pipeline of requests ahead of it to drain before receiving service. Thus, DRAMs' high efficiency and low latency are in conflict, which is particularly acute in

SOCs, in which the large number of processors and other initiators have varying burst characteristics, throughput requirements, and latency sensitivities. This situation forces the designers of many consumer SOC, who must optimize for highest DRAM efficiency, to implement complex arbitration and scheduling techniques to balance the efficiency-versus-latency trade-offs for different classes of initiators.

For instance, an LCD controller can naturally fetch an entire scan line's worth of pixels that it will display. This traffic pattern looks like a long incrementing burst, which is nearly ideal for maximizing DRAM efficiency because it reads most or all of the bits in one page before moving on to the next one. However, making a general-purpose CPU that is serving a cache miss wait behind such a large sequence of DRAM requests substantially reduces CPU performance and could even prevent the CPU from serving a critical interrupt in a timely manner. Thus, SOC designers would normally break up the scan-line fetch into smaller bursts so that latency-sensitive CPU requests can interleave between these bursts.

When the interleaved CPU request targets a different DRAM bank from the scan line's bank, the DRAM controller can schedule the requests without losing efficiency. If the CPU targets the same bank, however, the CPU's request requires closing the scan line's page so that the DRAM can open the CPU's page. After the DRAM services the CPU request, the chip must close the CPU's page so that it can reopen the scan line's page. If the CPU has another cache miss on the same page, this process will recur and waste substantial DRAM efficiency. This situation, page thrashing, can be so inefficient that the CPU would have achieved higher performance by waiting until the scan-line accesses had finished—that is, not interleaving in the first place.

Because the DRAM system must deliver the combined throughput requirements of several initiators and because higher-efficiency DRAM requires longer bursts, it follows that the DRAM system normally transfers data at higher peak bandwidth than most initiators—with the notable exception of the CPU—require. Most initiators therefore include FIFO (first-in/first-out) buffers so that the DRAM can efficiently service their requests into or out of the FIFO buffer, allowing the initiator to move data at lower bandwidth on the other side of the FIFO buffer. For those initiators requiring guaranteed throughput from DRAM, these FIFO buffers provide latency tolerance by covering the communication requirements of the initiator for the time it takes the initiator to drain or fill the FIFO. Deeper buffers make the initiator more latency-tolerant, giving SOC designers more flexibility in scheduling the initiator's traffic to DRAM at the cost of additional buffering area. For reads, such an architecture relies on the initiator's providing burst requests far in advance of needing the data; this approach is simply another form of pipelining.

The prefetch architecture of DDR3 raises a new challenge for consumer SOC. As the minimum efficient burst lengths of DRAMs increase and the high bandwidth requirements of SOC force wide DRAM words, the minimum efficient DRAM burst can reach 64 bytes or higher. Initiators are accessing data structures of various sizes and access patterns. Of par-

[Go to www.edn.com/ms4325](http://www.edn.com/ms4325) and click on Feedback Loop to post a comment on this article.

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

ticular interest are CPU-cache lines and MPEG macroblocks, each of which is relatively small and has effectively random access patterns. Initiators' fetches that are smaller than the DRAM burst waste DRAM-data transfers, and the system operates less efficiently. Today, this access-granularity problem is particularly acute in HDTV and STB applications, but it is also becoming common in other video-capable SOC. SOC designers often address the access-granularity problem by reducing the size of the DRAM burst by either rebuilding

their designs to use less DRAM bandwidth—for instance, by adding on-chip memory—or splitting the DRAM system into multiple independent channels, in which each channel is narrower and thus operates at lower burst size.

A final challenge facing SOC designers is the variety of operating modes in consumer-electronics devices. Cost concerns preclude optimizing for the sum of the worst-case bandwidth requirements of each initiator across all modes, so the designer must instead independently consider each key operating mode. In each mode, the various initiators have different performance requirements, and optimum scheduling and buffering choices often differ across modes. Because a single SOC must support these modes, the designer must eventually choose an architecture and design parameters that cover these needs. Sometimes, designers leverage programmable scheduling and other features to allow runtime optimization of the SOC for the different modes.

Many challenges face designers of consumer SOC when they are considering the implications of external-DRAM systems. Designers must select DRAM channel counts, burst lengths, arbitration and scheduling policies, and FIFO-buffer depths to carefully optimize their design across a range of key operating modes of the end applications. Although the underlying DRAM technology continues to deliver both higher capacity and peak bandwidth, these features bring with them additional latency and complexity. The wide variety of initiators and associated traffic requirements of the SOC, along with the extreme cost sensitivities of consumer markets, make these challenges more difficult. These challenges force designers to operate DRAMs at the highest achievable efficiency. Commercial products to address these problems are available from Sonics and other manufacturers, and you should consider them when embarking on new consumer-SOC development. **EDN**

AUTHOR'S BIOGRAPHY



Drew E. Wingard, PhD, co-founded Sonics in September 1996 and has been chief technical officer and secretary since March 1997. He is also a member of the board of directors. Before co-founding Sonics, Wingard led the development of advanced circuit and CAD methodology for MicroUnity Systems Engineering Inc. He also co-founded and worked at Pomegranate Technology from 1992 to 1994. Since December 2001, Wingard has served as secretary and a director of the OCP-IP (Open Core Protocol-International Partnership), a nonprofit trade organization. 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).

First-pass success in silicon packaging

PACKAGE ISSUES INTERACT WITH ALL ASPECTS OF THE DESIGN FLOW, FROM CHIP ARCHITECTURE TO MANUFACTURING DECISIONS.

There are few aspects of a semiconductor program that can be more detrimental than a re-spin. This re-spin may be of a die, a package, a PCB (printed-circuit board), or another portion of a project. When you consider the costs of such an event, you must include labor, software tools, masks and other manufacturing tooling, lost-opportunity costs, and time to market. These costs can be crippling to a program or even to a company, particularly if the delay also means missing a window of opportunity for a certain design win.

Historically, design teams completed silicon design, packaging, and PCB layout in series, with minimal interaction between these various entities. As technology nodes shrink, data rates accelerate, signal density increases, and systems become more densely populated, the boundaries between these entities quickly blur. Power considerations play heavily into leakage, standby time, environmental impacts, and reliability. Therefore, with the possible exception of simple devices, any modern packaging engagement should begin at the floorplanning and IP (intellectual-property)-selection stage and extend all the way through the planning of system-level integration (Figure 1).

PLANNING STAGES

During floorplanning and IP selection, designers require packaging input for considering the signal- and power-integrity requirements and the exit-routing strategy. This early stage is the most cost-effective time for packaging to supply feedback to the planning team. For example, a given IP element may have a bump pattern that requires an expensive packaging technology, but, with a modification of the selection, a more cost-effective approach may be possible. You can find another example of feedback in a core-limited design: Signals near a higher-speed SERDES (serializer/deserializer) should move farther away from the SERDES to provide for better isolation, shorter wire bonds, and a simplified power-delivery structure.

The package technology should heavily influence the pad ring. For example, a wire-bond package with higher-inductance connections may require more power-supply buffers on the die (Figure 2). Another example is that for a flip-chip

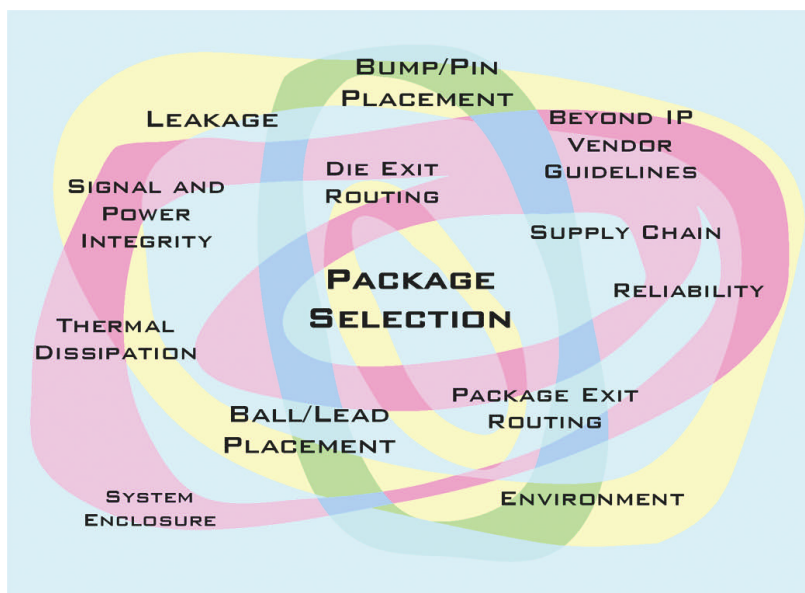


Figure 1 Many disparate technical considerations go into package selection.

design, in which it may be desirable to have a repeatable pattern of signals and supplies (Figure 3). This pattern allows a repeatable redistribution layer on the wafer. You can model the repeated pattern once to represent the other sections of the chip with a similar pattern.

Few details in the planning stage have a higher influence on signal integrity than bump, pad, and ball placement. Mistakes at this stage are difficult to counter in later portions of the design phase. The bump or bond-pad pattern on the die provides the most influence on the power-delivery structure in the package. For example, the order of the bumps from the edge of the die inward has a close relationship with the layer assignments on the package substrate. A bump-assignment pattern that physical designers believe gives them the best power delivery may in fact pose a large power-delivery problem. Knowing how to route the package gives packaging engineers the best insight about how to arrange these bumps on the die and can influence power deliver, thereby saving a significant amount of time on later design fixes.

A PCB's exit-routing strategy requires a similar train of thought. Packaging engineers must truly understand the PCB's routing challenges and the methods available to them when they select the best possible ball pattern for a device. These

challenges are worse for ASIC suppliers, which may sell into markets and applications in which the supplier doesn't know all of the PCB information.

Designers must also consider the PCB stackup when contemplating an exit-route strategy. The placement of these signals and supplies has a large influence on power delivery and return currents. In most cases, designers must present higher-speed differential pairs in a manner that allows the nets to exit in a differential format on the same layer. Some designs require striplines in internal layers, broadside-coupled pairs, or microstrips on the outer layers, and designers must place these balls in a manner that accommodates these strategies. Failure to consider these options may force users to adopt more expensive PCB technologies with blind, buried, or partially drilled vias to meet the signal requirements with an inadequate ball pattern, whereas, with a better ball pattern, they could use a more conventional laminate technology that allows significant system-cost savings.

This same thought process works for package sizing. It is not uncommon for a less expensive package to require a more expensive PCB. One familiar example occurs when designers reduce pitch size so that the PCBs will fit into smaller packages. The tighter pitch may force tighter via and trace rules to provide an exit route for a package. A less understood example is when designers shrink a package by driving signal nets deeper into a package. Having signals closer to the center of a package may affect power delivery to the core logic and drive up the number of layers required for providing an exit route for the package. Either way, these examples show how a shortsighted attempt to reduce package costs ends up increasing the total system costs.

THE THERMAL MISTAKE

The design team should consider thermal characteristics of a device in a system early in a project. Normally, the team has insufficient data on the package design, die, PCB, and system enclosure to generate an accurate number, but some approaches work for crossing off thermal issues in a project. Engineers must also deal with many pitfalls when making this assessment. Many engineers make the critical mistake of using JEDEC (www.jedec.org) terms such as “ θ -JA” (thermal resistance from junction to ambient) to determine how their part will work without a heat sink or “ θ -JC” (thermal resistance from junction to case) to see how a part will work with a heat sink. JEDEC created these terms for comparing one package with another, and the terms rarely match those for a real application. Designers have for years misused the JEDEC JESD51 specification, which governs this thermal terminology, to deter-

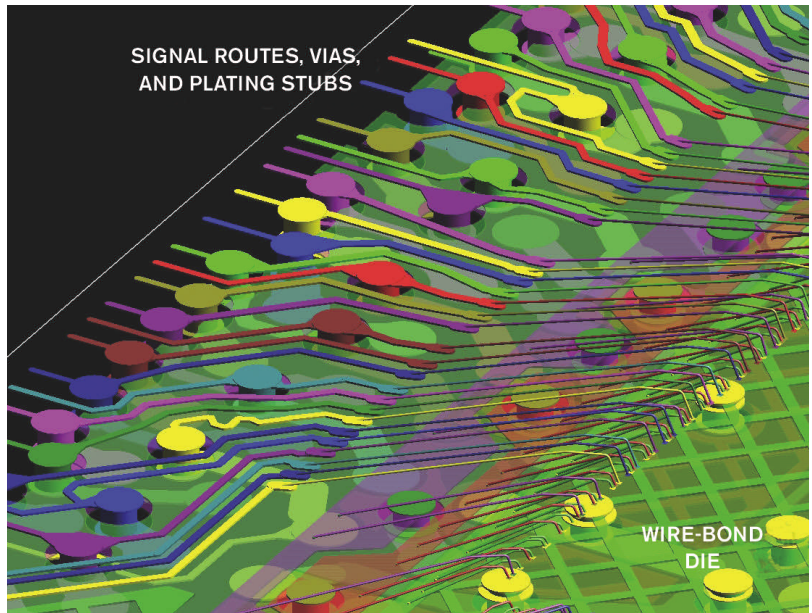


Figure 2 The high inductance of wire-bond exit routing should influence downstream decisions about the die's power grid.

mine the maximum power of devices in non-JEDEC environments, and this trend is unfortunately still going strong today.

Another common mistake is to take thermal performance data from a generic data sheet on a package. This generic data works for only one environment, die size, via count, and ball pattern, and you can use it to get a general idea of thermal performance. Be sure to use this information for only its intended purpose of general performance.

The most efficient approach would be to better understand the environment in which your device will need to perform. Determine whether your device will have a heat sink, airflow, an enclosure that touches the package or PCB, or other such heat-conducting factors. Also determine whether your PCB is more conductive than the standard JEDEC thermal board. Compare these factors with the available JEDEC analysis setup as a guide and then determine which factors are more conservative or more aggressive than those

FEW DETAILS IN THE PLANNING STAGE HAVE A HIGHER INFLUENCE ON SIGNAL INTEGRITY THAN BUMP, PAD, AND BALL PLACEMENT.

in your own environment. Most applications are more conservative than those in the JEDEC environment. If your calculated maximum junction temperature has sufficient head room compared with the specified maximum junction and your environment is more conservative than the JEDEC environment, then this approach will usually suffice. Otherwise, a thermal simulation would be a wise option to consider before getting too far along in your project. Some packaging

foundries, IDMs (integrated-device manufacturers), and value-chain producers perform this service free for customers, whereas others charge a nominal fee. Value-chain producers combine the high volume of design activity of a fabless ASIC company with more direct responsibility for customers' total

cost of ownership. Industry consultants can also help with this activity.

ELECTRICAL CONSIDERATIONS

You should use electrical simulation as a tool to verify what you expect to see in a package design. You should not design a package first and then just simulate to see whether it will work. Using simulation for package design considerably increases development time due to the iterations necessary to sufficiently tweak a design to make it function. It is important to have proper knowledge, guidance, or both at this stage.

If you send a design to a package supplier or a design house that is experienced in design but not in signal integrity, expect issues as your design moves to higher data rates. Consequently, using a group or a company that is separate from the package-design group to perform the simulation introduces a significant delay while the simulation group repeatedly feeds back changes to the design group. You must better integrate this process for higher-speed designs. The best approach is to use a team who knows the intricacies of high-speed design and has simulated various configurations before beginning a design. Such a team can then simulate the packaged device either stand-alone or in system configurations to confirm what you expect. An unexpected issue, such as resonance, could occur in a simulation, but such an event should be the exception and not the rule.

It is also important to understand which simulations for a device make sense in each environment. You cannot simulate every environment for a device and still remain within budget. The approach should be to first understand where the risks are and then identify the appropriate simulation activity, if it is necessary at all. Software tools can be expensive, as are the experienced signal-integrity engineers who can effectively operate them, so, depending on the throughput of the overall design organization, the team may not have access to these costly resources on a full-time basis. It is more important to have the right people and the right tools at your disposal for this activity when you need them.

KNOW YOUR RELIABILITY REQUIREMENTS

You can make trade-offs to improve reliability in one application at the expense of another. One example occurs in the use of lead-free solder-ball alloys. Some alloys perform well with thermal cycling but poorly in drop testing. Others exhibit just the opposite behavior. If your application is for handheld systems, you would tune your materials choices differently from how you would tune them for network applications.

Supply-chain selection is generally an afterthought in a semiconductor program. It is a mistake to assume that all offshore subcontractors are generally the same and that they will effectively handle a project regardless of volume forecasts and relationships. It is important to carefully select a supply strategy. Corporations with high and repeatable volume forecasts have a generally easier time working with offshore suppliers. Most companies do not fall into this category and instead face pricing issues, customer-service and technology-support prob-




Figure 3 Flip-chip technology can make possible repeatable routing patterns on the redistribution layer.

lems, and steep learning curves. Most offshore suppliers are not looking to bring in whatever business they can. To maximize profit, their strategy has instead been to bring in the business that best aligns with their own strategy. This situation means that offshore suppliers prefer to work with customers that are technically self-sufficient, can forecast well, and serve growing markets.

Smaller tier-two and tier-three suppliers may better serve lower-volume customers that require more support and programs that do not easily fit the cookie-cutter model of an offshore-assembly supplier. Many of these suppliers are in North America and Europe. These smaller companies do not specialize in high volume and low cost but do provide the support and customization that certain projects require.

Assembly subcontractors in fields such as thermal-analysis, PCB-routing, and signal-integrity engineering have varying degrees of proficiency. Many differences exist between wafer foundries, assembly suppliers, and contract manufacturers at the board and system levels. Even IDMs that generally have these capabilities in-house also experience trouble bridging this gap due to their own organizations' barriers. Bridging this gap has historically been the responsibility of the company contracting with these various entities, but it may be difficult to find resources with know-how in manufacturing.

WHAT ARE THE OPTIONS?

With all of these crucial elements to consider, various questions come into play. You must select the right supply-chain partner or partners for your devices, and you must determine which technology best suits your needs. You also need to determine which factors affect pricing and how you should trade off these factors against performance and flexibility. You must know which level of simulation is sufficient to ensure performance and where you can get the information you need to select the best package for your product. Some large companies have large, seasoned teams who may be able to address these and other questions, but these companies would need a significant throughput of products to keep them working. Larger IDMs have been able to fill in as necessary but are usu-

ally limited to their captive fab, assembly, or test capabilities.

Using consultants as necessary to fill in gaps can be helpful, especially when engaging in newer technologies or with any technology outside a client's comfort zone. However, customers generally base consultants' compensation on the number of hours they spend on a project, so it is always within the best interest of the consultants to spend as much time on a program as possible.

Their bias generally helps to improve the program's chances of functioning properly but generally does not help with maintaining a reasonable budget or program schedule. In addition, consultants may not have enough contacts and relationships with various necessary suppliers because they generally work independently and have limited resources.

Outsourcing activities to different companies has several significant drawbacks, as well. If one company works on the design of a laminate and another performs signal-integrity simulation, then multiple issues can arise from both sides. Say, for example, that a laminate has signaling and power requirements that would require a company to perform analysis. In that case, it would be a mistake to use another company to handle the design because the laminate almost always requires modification based on input from the simulation company. This cycle of feedback and the changes that result lengthen the development time of a packaged device.

The value-chain producer is emerging as an alternative to

Go to www.edn.com/ms4333 and click on Feedback Loop to post a comment on this article.

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

in-house expertise, consultants, and outsourcing to multiple contractors. Due to their unique business model, most value-chain producers profit only when a program enters production. Hence, they tend to find the most cost-effective and timely manner for getting the product out the door. First-time-right approaches become paramount, as any program delays negatively affect the value-chain producers' finances. With this in mind and benefiting from their relatively large design volumes,

such vendors tend to have extensive capabilities in design, extraction, simulation, and manufacturing. Furthermore, they are better able to shift resources from program to program than are other business models, due to having larger pools of more highly skilled employees and by using their existing production and supplier relationships. They can also spread these supplier relationships across a wide industry ecosystem, including the more-difficult-to-reach second- and third-tier suppliers, which can handle specialized jobs and lower volumes.


Understanding the technical and supply-chain considerations regarding packaging engagements is a critical step toward successful implementation. You can address and mitigate many of the issues at the start of a program with appropriate planning and strategic partnerships. **EDN**

AUTHOR'S BIOGRAPHY

Javier DeLaCruz is director of semiconductor packaging at eSilicon Corp.

CMOS Logic Distribution Amplifier

- High input impedance with hysteresis
- High current outputs
- Fast transition times
- Small insertion delay
- Low channel-to-channel timing skew
- \$950 (US list)



The FS730/3 CMOS Logic Distribution Amplifier is used to distribute CMOS level logic pulses. Each amplifier has one input and seven outputs. All inputs and outputs are logic levels. The 50 Ω outputs will reverse terminate reflected pulses when driving unterminated lines. All of the outputs are driven by

separate drivers to provide high isolation, along with fast transition times and very low overshoot. Both bench-top and rack-mount configurations are available.

SRS also offers state-of-the-art solutions for distributing 10 MHz, Broadband (50 Ω and 75 Ω), and SDI digital video. Please see our website for details.



Stanford Research Systems

• www.thinkSRS.com


• 408-744-9040

Industrial Precision Op Amps


Out with the old...




...in with the new



www.linear.com/circuits




Tiny Package




Rail-to-Rail
Outputs/Inputs




Minimal Drift



Low Noise



Low Bias Current



Wide Supply
Range

Give Your Industrial Design a Facelift

Our precision amplifiers provide the ideal combination of specs for the industrial market. We have improved on industry standard amplifiers with wider supply ranges, higher precision, lower noise and supply current, along with rail-to-rail outputs. These amplifiers minimize board space, extend battery life, and improve system performance, driving your product's performance to the next level.

Precision Op Amps

Industry Standard Amplifiers	Description	Updated Amplifiers	New Features
LT*1078/9 LT2078/9	Precision, Micropower, Single Supply	LTC*6078/9*	Rail-to-Rail, Lower Noise, Faster, 125°C Specified
LT1008 LT1012 LT1097	Precision, Low Noise, C-Load Stable	LT1880 LT6010	Rail-to-Rail, SOT-23
LT1112 LT1114	Low Power, Matching Specs, C-Load Stable	LT1881/2 LT6011/2/4*	Rail-to-Rail, 3mm x 3mm DFN
LT1494/5/6	Ultralow Power, Rail-to-Rail, Precision	LT6003/4/5*	Lower Supply Range, Smaller Packages
LT1055/6/7/8 LT1169	Picoamp Input Bias Current	LTC6240/1/2/4* LTC6084/5* LTC6087/8*	Rail-to-Rail, Lower Power, Smaller Packages, Faster, 125°C Specified
LT1013/4	Low Offset, Low Offset Drift, Low Power	LT1490A LT1491A	Rail-to-Rail, Rugged, 125°C Specified, 3mm x 3mm DFN
LT1028/LT1128	Low Noise, Low Drift	LT6200/1* LT6230/1/2	Lower Power, Faster, Rail-to-Rail, SOT-23
LT1007 LT1037	Extremely Low Noise	LT1677/8/9	Rail-to-Rail
LT1124/5/6/7	Low Noise, Low 1/f Corner, Precision	LT6202/3/4* LT6233/4/5	Lower Power, Faster, Rail-to-Rail, SOT-23
LTC1050/1/2/3	Zero Drift	LTC2050/1/2/4/5*	Smaller Packages, 125°C Specified

*Maximum supply voltage is lower than predecessor.

Info & Free Samples

www.linear.com/precisionopamps

1-800-4-LINEAR



Free Industrial
Signal Chain
Brochure

www.linear.com/indsolutions

LTC, LTC, LT, LIM and Over-the-Top are registered trademarks of Linear Technology Corporation. All other trademarks are the property of their respective owners.



designideas

READERS SOLVE DESIGN PROBLEMS

High-speed op amp enables IR-proximity sensing

Arpit Mehta, Maxim Integrated Products Inc, Sunnyvale, CA

IR (infrared)-proximity sensors can sense the presence of an object, its distance from a reference, or both. Applications include speed detection, sensing of the hand in automatic faucets, automatic counting or detection of objects on conveyer belts, and paper-edge detection in printers. The latest-generation smartphones, for example, can turn off the LCD touchscreen to prevent the accidental activation of buttons when you press the screen against your chin or your ear.

To sense an object, a proximity sensor transmits IR pulses toward the object and then “listens” to detect any pulses that reflect back. An IR LED transmits the IR signals, and an IR photodiode detects the reflected signal. The strength of this reflected signal is inversely proportional to the distance of the object from the IR transceiver. Because the reflected IR signal is stronger when the object is close, you can calibrate the output of the photodiode detector to determine the exact trigger distance of an object. The trigger distance indicates the threshold for making a decision on whether an object is present.

The photodiode detects IR not only that the object reflects, but also from the ambient conditions. You must filter out this IR noise to prevent false detections. A common method is to

modulate the LED’s IR signal with a convenient frequency and then detect only the IR with that modulation, which identifies it as a reflection from the object.

This Design Idea describes an IR-proximity sensor with simple transmitter and receiver sections (Figure 1). The transmitter consists of an Everlight (www.everlight.com) 940-nm IR11-21C IR LED, which turns on and off using a 10-kHz oscillator frequency. By varying the LED’s current, you control the level of transmitted power and, hence, the detection range. To save power, the transmitting pulses have a typical duty cycle of only 10%.

The receiver circuit demodulates and amplifies the IR signals that the Everlight PD15-22C photodiode de-

DIs Inside

44 Set your lights to music

48 Current limiter allows large USB bypass capacitance

50 High-speed pulse modulator retains signal envelope

► What are your design problems and solutions? Publish them here and receive \$150! Send your Design Ideas to edndesignideas@reedbusiness.com.

► To see all of EDN’s Design Ideas, visit www.edn.com/designideas.

fects; the photodiode’s peak sensitivity occurs at 940 nm. The photodiode output ac couples to the op amp’s noninverting input. This coupling allows the 10-kHz signal to pass, but the coupling capacitor sets a 300-Hz cutoff frequency that prevents dc noise and background IR from reaching the amplifier.

Low noise, high bandwidth, and rail-to-rail-I/O capability make the op amp a good choice for demodulation and amplification in this circuit. In addition, its RF immunity prevents the annoying 217-Hz audio buzz that you commonly find in GSM (global-system-for-mobile)-communications cell phones. For the IR receiver, the op amp acts as a gain-of-100, second-order bandpass fil-

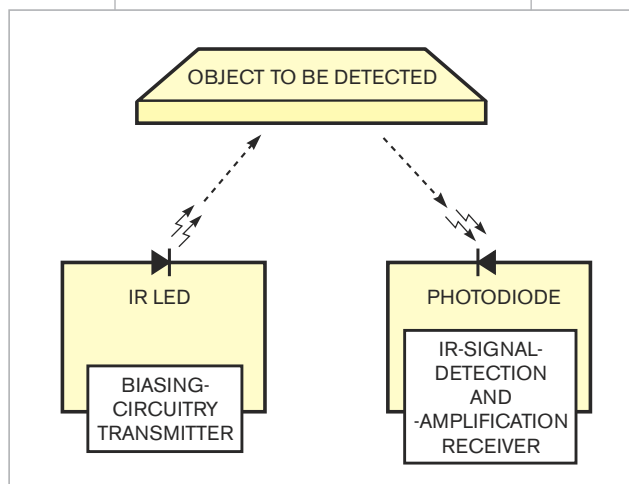


Figure 1 An IR-proximity sensor detects an object by receiving reflected light.

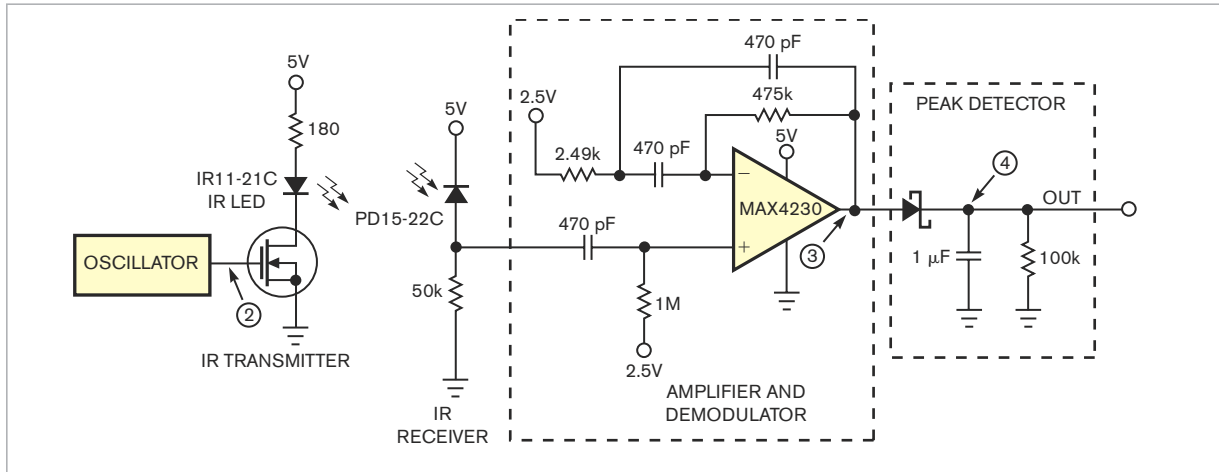


Figure 2 An IR transceiver detects the presence of an object and provides an approximate distance from the transceiver.

ter with a center frequency of 10 kHz. Thus, the op amp amplifies the incoming IR signals and demodulates them with a bandpass filter.

With no input IR signal present, the op amp is biased at 2.5V. With a 10-kHz IR signal incident, its output varies around 2.5V with a dynamic range of 5V. The output drives a simple diode detector, which rectifies the 10-kHz signal and provides a dc signal proportional to its amplitude. This analog-output signal is proportional to the distance of the object from the IR transmitter. You can use it as is or feed it to an ADC for further processing.

Figure 2 shows circuit operation at three nodes for objects at 1.2 and 1.4 in. from the IR transceiver. The circled numbers in **Figure 2** refer to the oscilloscope traces in **Figure 3**.**EDN**

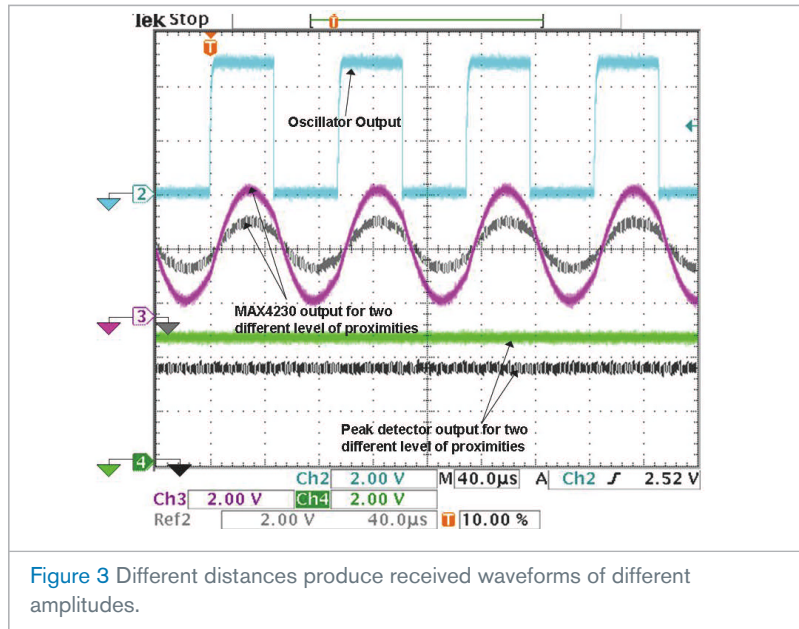


Figure 3 Different distances produce received waveforms of different amplitudes.

Set your lights to music

Hanif Saeed, Maxim Integrated Products Inc, Sunnyvale, CA


As one of many ways you can implement a light show, the circuit in this Design Idea selectively activates various subsets in a group of six strings of lights, causing them to flash on and off according to the level and tempo of music you are playing. The stand-alone circuit requires no microcontroller, no software, and no trimming (**Figure 1**). You apply the

audio signal you want to display to IC₁, a 12-bit ADC. The signal ranges from 0 to 2.048V, causing the first string of lights to come on at 2 mV. Although the circuit controls six ac outlets, you can expand it to control 12 outlets.

A short positive pulse at the CNVST pin of IC₁ triggers it to initiate a conversion, which the SCLK signal clocks.

Its output (DOUT), which the rising edges of SCLK clock, comprises four leading zeros followed by the 12-bit conversion result, MSB (most-significant bit) first. Thus, one conversion requires 16 clock pulses at SCLK.

A vertical stack of six switched outlets, in which the top outlet represents the MSB, powers the display. You might, for example, plug a separate string of lights into each outlet. During operation, the circuit scans each conversion result as it is generated (MSB



DESIGN NOTES

Triple Output DC/DC μ Module[®] Regulator in 15mm \times 15mm \times 2.8mm Surface Mount Package Replaces Up to 30 Discrete Components Design Note 469

Eddie Beville and Alan Chern

Introduction

When space and design-time are tight in multivoltage systems, the solution is a multioutput DC/DC regulator IC. For more space and time constraint systems, a better solution is an already-fabricated compact multioutput DC/DC system that includes not only the regulator ICs but the supporting components such as the inductors, compensation circuits, capacitors and resistors.

Dual Switching 4A and 1.5A VLDO[™] Regulators

The LTM[®]4615 offers three separate power supply regulators in a 15mm \times 15mm \times 2.8mm LGA surface mount package: two switching DC/DC regulators and one very low dropout VLDO linear regulator (Figure 1). MOSFETs, inductors, and other support components are all built in. Each power supply can be powered individually or together, to form a single input, three output design. Moreover, for an otherwise complex triple output circuit design, the task is eased to designing with only one device while the layout is as simple as copying and pasting the LTM4615's package layout. One LTM4615 replaces up to 30 discrete components when compared to a triple-output high efficiency DC/DC circuit.

The two switching regulators, operating at a 1.25MHz switching frequency, accept input voltages between 2.35V

to 5.5V and each delivers a resistor-set output voltage of 0.8V to 5V at 4A of continuous current (5A peak). The output voltages can track each other or another voltage source. Other features include, low output voltage ripple and low thermal dissipation.

The VLDO regulator input voltage (1.14V to 3.5V) is capable of up to 1.5A of output current with an adjustable output range of 0.4V to 2.6V, also via a resistor. The VLDO regulator has a low voltage dropout of 200mV at maximum load. The regulator can be used independently, or in conjunction with either of the two switching regulators to create a high efficiency, low noise, large-ratio step-down supply—simply tie one of the switching regulator's outputs to the input of the VLDO regulator.

Multiple Low Noise Outputs

The LTM4615 is capable of operating with all three regulators at full load while maintaining optimum efficiency. A typical LTM4615 design (Figure 2) for a 3.3V input to three outputs has the VLDO input driven by V_{OUT2}. The efficiency of this design is shown in Figure 3.




Figure 1. Three DC/DC Circuits in One Package

LT, LT, LTC, LTM, Linear Technology, the Linear logo and μ Module are registered trademarks, and VLDO is a trademark of Linear Technology Corporation. All other trademarks are the property of their respective owners.

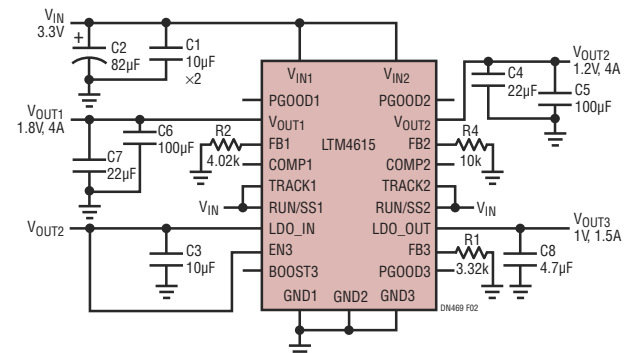


Figure 2. Triple Output LTM4615: 3.3V Input, 1.8V (4A), 1.2V (4A), 1.0V (1.5A)




Figure 3. Efficiency of the Circuit in Figure 2, 1.8V, 1.2V and 1.0V (VLDO)




Figure 4. Low Output Voltage Ripple (3.3V Input)




Figure 5. Top View Thermal Imaging of the Unit at Full Load in Ambient Temperature with No Airflow. Even Temperatures (Cursors 1 and 3) Indicate Balanced Thermal Conductivity Between the Two Switching Regulators. 3.3V Input, 1.8V (4A) and 1.2V (4A).




Figure 6. Output Voltage Tracking Design V_{OUT2} (1.2V) Tracks V_{OUT1} (1.8V)

The LTM4615 comes prepackaged with ceramic capacitors and additional output capacitors are only needed under full 4A load and if the input source impedance is compromised by long inductive leads or traces.

The VLDO regulator provides a particularly low noise 1.0V supply as it is driven by the output of the 1.2V switching regulator (V_{OUT2}). The low output voltage ripple for all three outputs is shown in Figure 4.

Thermally Enhanced Packaging

The LGA packaging allows heatsinking from both the top and bottom. This design utilizes the PCB copper layout to draw heat away from the part and into the board. Additionally, a heat sink can be placed on top of the device, such as a metal chassis, to promote thermal conductivity. Thermal dissipation is well balanced between the two switching regulators (Figure 5).

Output Voltage Tracking

A tracking design (Figure 6) and output (Figure 7) can be programmed using the TRACK1 and TRACK2 pins. Divide down the master regulator's output with an external resistor divider that is the same as the slave regulator's feedback divider on the slave's TRACK pin for coincident tracking.




Figure 7. Start-Up Voltage for Figure 5 Circuit V_{OUT1} (1.8V) Coincidentally Tracks V_{OUT2} (1.2V) for Coincident Tracking

Data Sheet Download

www.linear.com

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

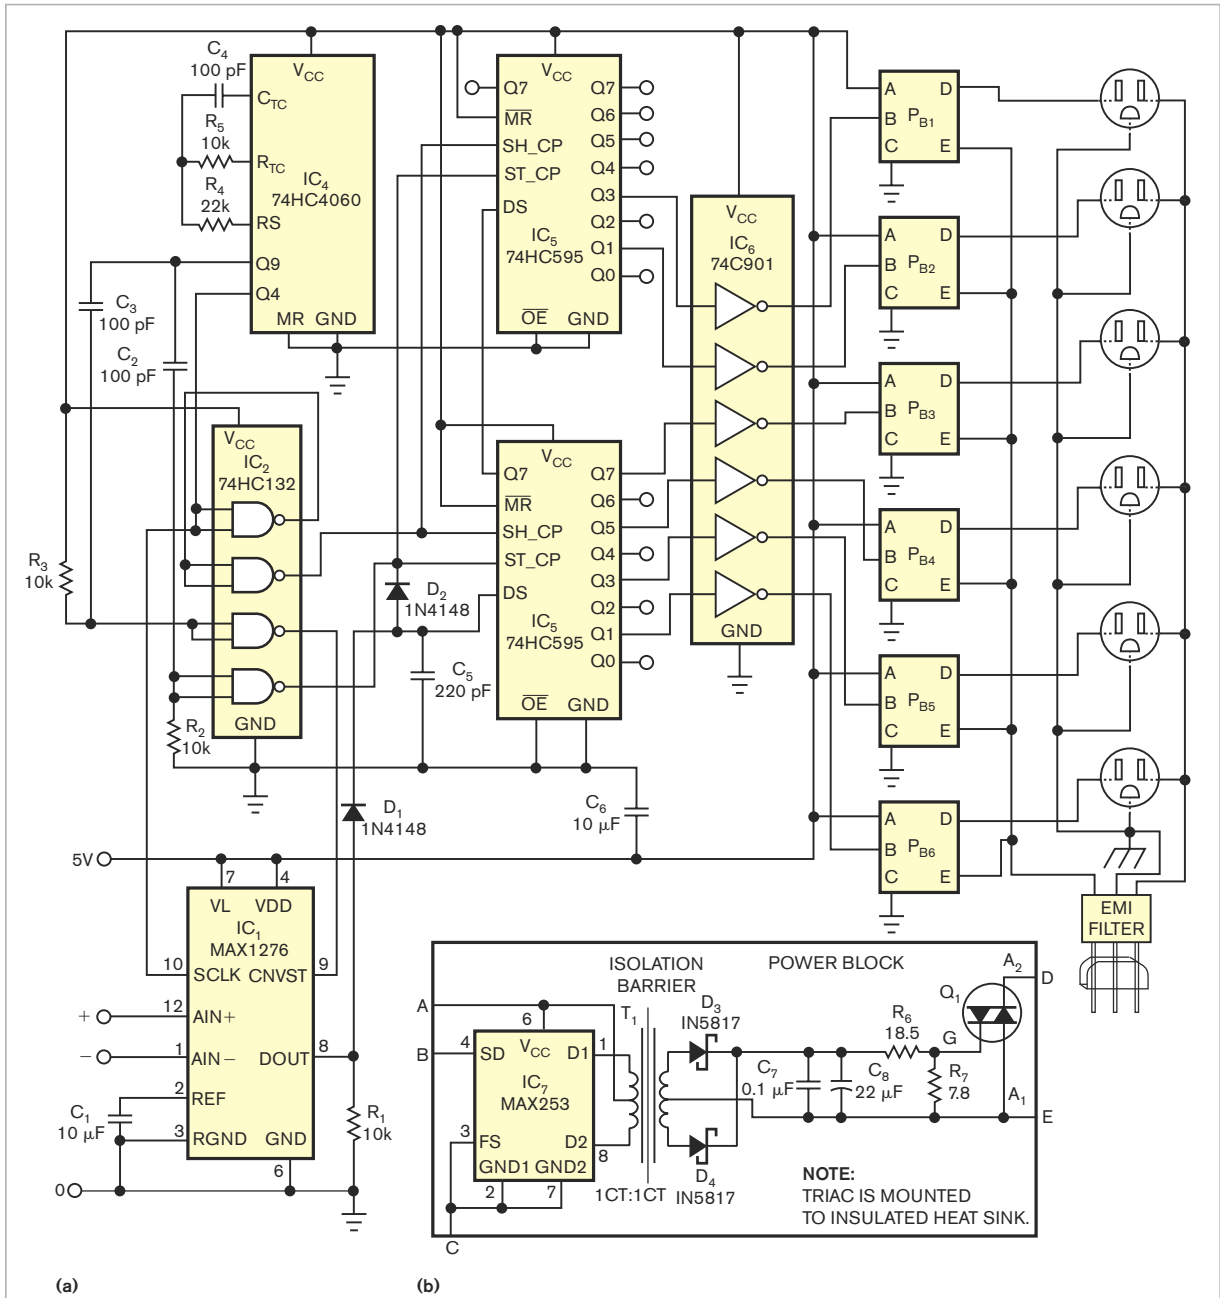


Figure 1 Driven by a 0 to 2.048V music signal at pins 12 and 1 of IC₁, this circuit activates the six ac outlets according to the music amplitude, in a logarithmic thermometer-code format (a). The power block (b) represents each of the power blocks, PB₁ through PB₆.

first, as described previously) and notes the first bit to assume a value of one. It then turns on the corresponding outlet and all those below it in the stack. The result is a logarithmic column, in which the change of input voltage necessary to move the column one step up or

down (a 12.04-dB increment) is either quadruple or one-fourth the immediate value. Although the number of steps available equals the ADC's resolution of 12 bits, this circuit uses only every other one to drive the six outlets.

At DOUT, the first output bit with

a value of one charges C₅ through D₁ to the logic-one level. The voltage on C₅ connects to the data input (DS) of the first of two cascaded 74HC595 ICs, which together form a 16-bit shift register. The signal that clocks the ADC, slightly delayed, also clocks this shift

register through the NAND gates in IC₂ and thereby inserts into the shift register the value present at its input. At the end of a conversion, the voltage stored on C₅ forces to one all the bits following the first one that exhibits a value of one.


At the completion of each conversion, a negative pulse applied to the ST_CP inputs of both 74HC595 ICs transfers these shift-register contents to a parallel-output register, IC₆. The same pulse discharges the storage capacitor through diode D₂, leaving the circuit ready for the next conversion scan. The parallel-register outputs then serve as drivers for the 12-bit logarithmic column, with the MSB driving the top outlet.

IC₄, a 74HC4060, serves as a clock and timing-sequence generator, and IC₇, a 74HC132, provides some necessary glue logic. For each connected 74HC595 output, the signal, which IC₆ inverts, activates the corresponding MAX253 transformer driver, IC₇, in one of the six power blocks. A 1-to-1 transformer isolates this driver signal, which then triggers solid-state TRIAC (triode for alternating current), Q₁, to its on state. For the component values in the figure, the circuit has a display-sampling rate of about 2.5 kHz and uses the 12th, 10th, eighth, sixth, fourth, and second bits to control the six outlets. The resulting light show adds an extra dazzle to the music you are playing.

This circuit operates at lethal voltages and requires proper handling. Note that the transformer must withstand a line level of 120V ac. It operates with incandescent light bulbs; you should not use any other type of light bulb. Even though the outlets are standard 120V-ac outputs for use with commercial incandescent lights, fast switching in the TRIACs makes them unsuitable for driving other types of loads, such as appliances, electronics, or ac adapters. Transformer T₁ is a TGM-350NA from Halo Electronics Inc (www.halo-electronics.com), and TRIAC Q₁ is a T1235-T from STMicroelectronics (www.st.com). For a video of this circuit in action, go to www.edn.com/090806dia. **EDN**

Current limiter allows large USB bypass capacitance

Daniel Morris, Group IV Technology, Renton, WA

 The USB (Universal Serial Bus) specification requires a connected USB device to present a load to the host or hub of no greater than 10 μF in parallel with 44Ω, including the effects of any bypass capacitance visible through the device's voltage regulator. This limit avoids excessive volt-

age drop at the device as inrush current charges its capacitance. Occasionally, a bus-powered device needs more than 10-μF bypass capacitance to provide an adequate reservoir for current spikes. The circuit in this Design Idea repurposes a Linear Technology (www.linear.com) LTC6102 precision cur-

rent-sense amplifier, IC₁, to limit inrush current below the specified maximum, allowing the device to use more capacitance when necessary.

The LTC6102 usually translates the voltage across a current-sense resistor to a larger ground-referenced voltage in an output resistor. The part features an amplifier with low offset voltage, letting you use low-value sense resistors. In the usual circuit configuration, output current flows through an onboard FET whose source connection connects to a force pin separate from the amplifier input pin to minimize errors across trace and pin resistances.

This circuit grounds the LTC6102's output pin and uses the onboard FET as a source follower to drive the gate of an external current-limiting FET (Figure 1). The feedback loop around the LTC6102 maintains equal voltages at the positive and negative inputs of the amplifier, pins 8 and 1 of IC₁. Resistor divider R₂/R₄ sets the positive input of the amplifier, IC₁'s Pin


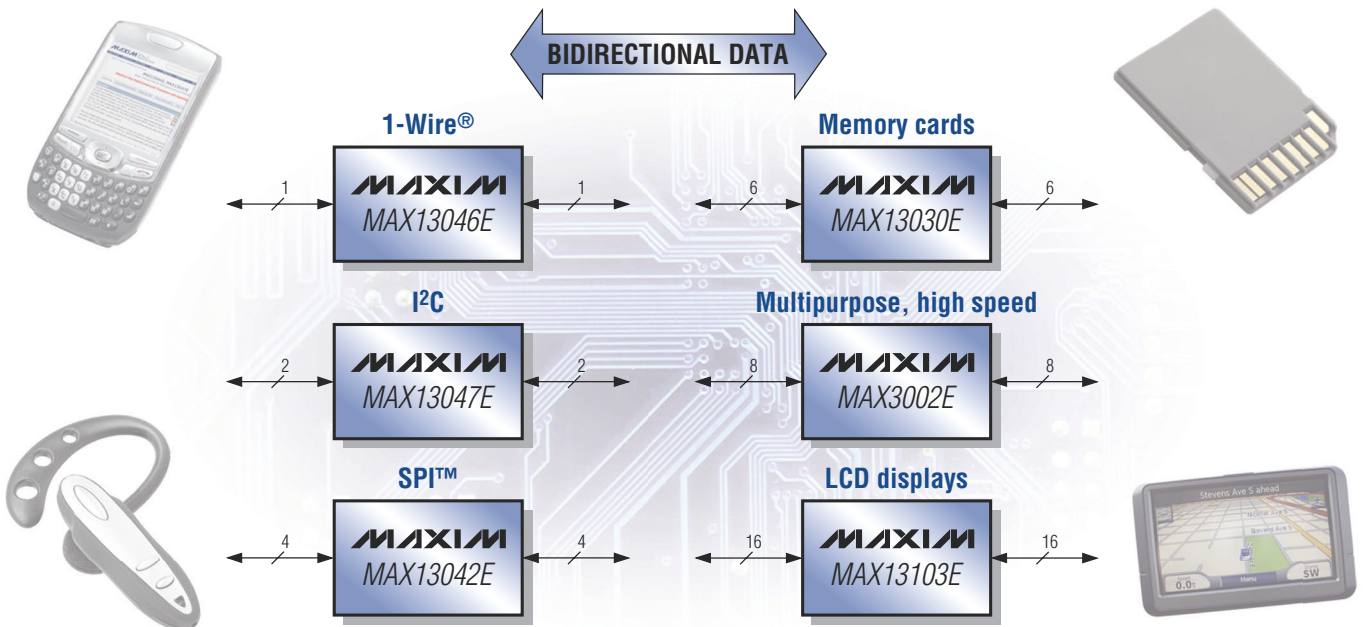


Figure 1 This circuit limits USB-device current both at connection and after configuration.

Maxim's bidirectional level translators seamlessly interface between ICs

Level translate without a direction pin and save board space



Low-voltage operation

- $V_L = 1.2V$ to $5.5V$, $V_{CC} = 1.65V$ to $5.5V$

Save board space compared to discretes

- Small UCSP™, μ DFN, and UTQFN packages

Save processor overhead

- Bidirectional operation with no direction pin

Robust protection

- $\pm 15kV$ ESD protection* (V_{CC} pins)

Save power

- Ultra-low, $1\mu A$ supply current in shutdown mode

Part	I/O Channels	V_L Supply (V)	V_{CC} Supply (V)	Data Rate (Mbps, max)	I/O V_L Shutdown State	I/O V_{CC} Shutdown State	Package (mm x mm)
MAX13046E	1	1.1 to 3.6	1.65 to 5.5	8	High-Z	High-Z	6- μ DFN (1 x 1.5)
MAX13047E	2	1.1 to 3.6	1.65 to 5.5	8			10-UTQFN (1.4 x 1.8)
MAX13042E	4	1.62 to 3.2	2.2 to 3.6	100			12-UCSP (1.5 x 2.1)
MAX13030E	6	1.62 to 3.2	2.2 to 3.6	100		16-UCSP (2 x 2)	
MAX3002E	8	1.2 to 5.5	1.65 to 5.5	20		6k Ω to GND	20-UCSP (2 x 2.5)
MAX13103E	16	1.2 to 5.5	1.65 to 5.5	20		High-Z	36-UCSP (3 x 3)

1-Wire is a registered trademark and UCSP is a trademark of Maxim Integrated Products, Inc.

SPI is a trademark of Motorola, Inc.

*Human Body Model.

www.maxim-ic.com/MAX13046-info



www.maxim-ic.com/shop



www.avnet.com



For free samples or technical support, visit our website.

Innovation Delivered is a trademark of Maxim Integrated Products, Inc. Maxim is a registered trademark of Maxim Integrated Products, Inc. © 2009 Maxim Integrated Products, Inc. All rights reserved.

8, approximately 2 mV below the 5V USB-voltage rail. With Q_1 initially off at device connection, the negative amplifier input, IC_1 's Pin 1, is higher than the positive input, causing the amplifier's output to go low. As the amplifier's output drops, the onboard FET follows, pulling the gate of Q_1 low and turning it on. Current increases in Q_1 until the voltage drop across sense resistor R_1 matches the drop across resistor R_2 .

Resistor R_3 and capacitor C_2 com-

pensate the feedback loop against oscillation and slow the turn-on of Q_1 , preventing an initial current spike when the device connects to the bus. Capacitor C_3 bypasses a regulator on IC_1 . Resistor R_7 meets the allowed maximum 1-mA current through the FET on IC_1 . Q_1 turns on at a gate voltage low enough that it does not exceed the input range of 4V positive voltage to IC_1 's Pin 7 to Pin 2.

Instead of the large capacitive load of C_1 , the circuit presents a resis-

tive load to the USB host equal to $R_1(R_2+R_4)/R_4=49.8\Omega$, lighter than the 44Ω maximum requirement. After C_1 charges, the circuit continues to limit current below the 100-mA maximum permitted to a low-power USB device. Upon configuration, the device can raise the current limit to the 500-mA maximum permitted to a high-power device by turning on FET Q_2 to place R_3 in parallel with R_4 , increasing the voltage maintained across sense resistor R_1 .**EDN**

High-speed pulse modulator retains signal envelope

Marián Štofka, Slovak University of Technology, Bratislava, Slovakia

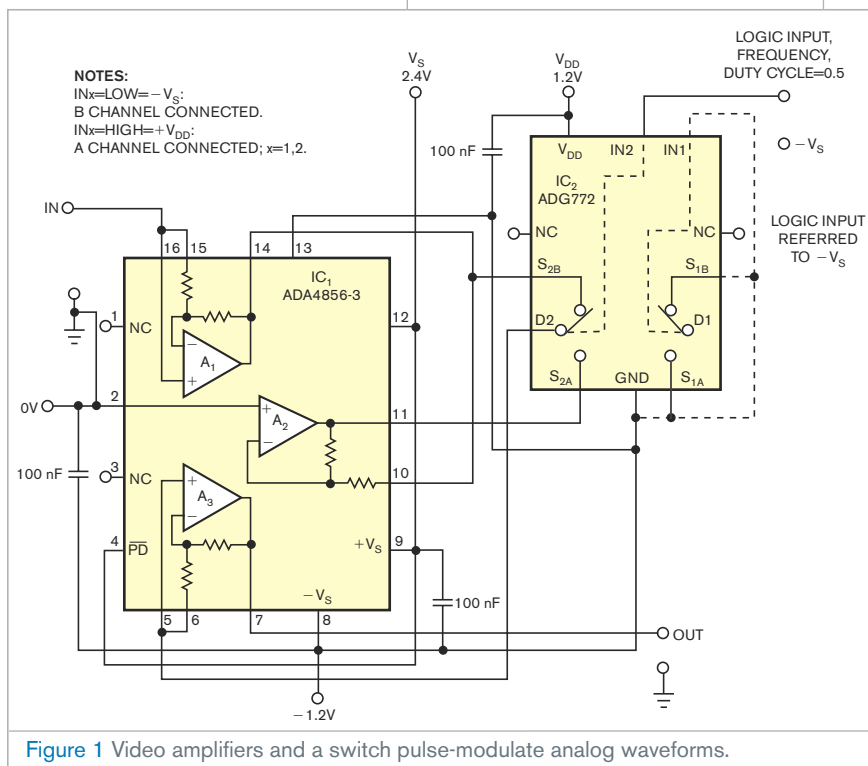
The circuit in **Figure 1** enables you to convert an arbitrary, relatively slowly varying voltage waveform to a new waveform in which the instantaneous values of the original waveform alternate with positive and negative signs. The new waveform retains information about the original waveform,

and its mean value approaches zero. This situation holds true for any input waveform, even a dc voltage. The nearly zero dc component of the output of the circuit in conjunction with the up-conversion of the frequency band lets the modulated waveform pass easily through a transformer (**Figure 2**).

The circuit uses just one and one-half ICs from Analog Devices (www.analog.com). IC_1 is a triple video amplifier, the ADA4856-3 with a gain of two (**Reference 1**). Amplifier A_1 acts as a voltage follower, which gives a maximally smooth and flat frequency response. Amplifier A_2 acts as an inverter, having a voltage gain of negative one, and A_3 serves as an impedance converter with a voltage gain of one.

IC_2 , an ADG772 high-speed 2-to-1 multiplexer (**Reference 2**), alternately switches the outputs of A_1 and of A_2 to the input of A_3 . You must keep the duty cycle of IC_2 's logic-control signal, IN2, close to 0.5 to ensure the "zero" mean value of the output voltage, even at a nonzero input voltage. At a modulation rate, or the frequency of the logic-control signal, of approximately 6 MHz, the output voltage's dc component shifts negligibly only from the low-frequency mean-offset voltage of the circuit, which is less than 4 mV.

Experiments have confirmed this value for an input voltage of 0V and for the precise reference dc voltage of 0.8188V. At a frequency of 60 MHz, the dc component of the output voltage remains at about 4 mV for an input voltage of 0V and rises to approximately 175 mV for an input voltage of 0.8188V. This result is still remarkable because the ADG772 is a BBM (break-before-make) type of multiplex-







Industry's lowest power, dual-channel, 10-bit, 130Mps ADC has superior dynamic performance

Extensive feature set minimizes external component count

The MAX19517 ADC is a member of Maxim's pin compatible family of ultra-low power, dual-channel ADCs. Its extensive set of features (a partial list is shown below) minimizes external component count and supports a wide range of applications that demand high performance.



Selected key features

- Lowest power: 75mW per channel
- Excellent performance for input frequencies from DC to > 400MHz
- Self-sensing supply regulation
- 1.8V, 2.5V, 3.3V operation
- Wide input common-mode range
- Programmable CMOS termination
- Flexible data formatting
- Programmable data timing
- Programmable clock divider: 1x, 2x, 4x
- Data output test patterns
- 7mm x 7mm, 48-pin TQFN package

Pin-compatible, ultra-low-power, dual-channel ADCs

Part	Resolution (Bits)	Sample Rate (Mps)	Power per Channel (mW)	SNR (dBFS)	SFDR (dBc)
MAX19517	10	130	75	60.0	85
MAX19516		100	57		
MAX19515		65	43		
MAX19507	8	130	75	49.7	69
MAX19506		100	57		
MAX19505		65	43		

www.maxim-ic.com/MAX19517-info



www.maxim-ic.com/shop



www.avnet.com



For free samples or technical support, visit our website.

er/switch. During time interval t_{BBM} , which is typically 5 nsec, both the S_{2A} and S_{2B} switches are temporarily off. Thus, the corresponding switch is on for approximately 8.2 nsec within a half-period of a 60-MHz control signal, yielding an on-state duration of only 3.2 nsec. An eventual 320-psec difference of the turn-on times of switches S_{2A} and S_{2B} would cause a shift in the dc component of 81.88 mV. The corresponding dc components of output voltages for an input voltage of 0V and an input voltage of 0.8188V differ by about 175 mV as a result of the difference in turn-on times of S_{1A} and S_{1B} . You can estimate this difference using the following equation:

$$320 \text{ psec} \times \frac{175 \text{ mV}}{81.88 \text{ mV}} \cong 684 \text{ psec.}$$




Figure 2 The pulse-modulated waveform can pass through a transformer, providing signal isolation. Comparing the input waveform, $V_{IN}(t)$ (a) with the output waveform, $V_{OUT}(t)$ (b) shows that the frequency spectrum of the output waveform upconverts while its dc component becomes zero.

Thus, this application calls for an analog multiplexer having the speed and bandwidth of the ADG772, and it should operate as an MBB (make-before-break) type. At a switching rate of 60 MHz, the channels of such a multiplexer will conduct almost three times

longer, and the difference in turn-on times of the A and B channels will be less significant. To prevent short overloading of amplifiers A_1 and A_2 , you can place SMD resistors of about 20Ω to the outputs of A_1 and A_2 when using an MBB multiplexer. **EDN**


REFERENCES

- 1 "ADA4856-3 Single-Supply, High Speed, Fixed Gain ± 2 , Rail-to-Rail Output Video Amplifier," Analog Devices, 2008 to 2009, www.analog.com/en/audiovideo-products/video-amps/buffers/filters/ada4856-3/products/product.html.
- 2 "ADG772: CMOS Low Power Dual 2:1 Mux/Demux USB 2.0 (480 Mbps)/USB 1.1 (12 Mbps)," Analog Devices, 2007 to 2008, www.analog.com/en/switches/multiplexers/analog-switches/adg772/products/product.html.

High Resolution Linear Hall Encoder

AS5311 – First magnetic encoder to break into submicron range

- ▶ 0.5 μ m/step absolute output
- ▶ 2.0 μ m/step incremental output
- ▶ 650 mm/s linear speed



ae austriamicrosystems

a leap ahead in analog

West Coast (408) 345-1790 · East Coast (919) 676-5292
www.austriamicrosystems.com/ANC

EDN literaturelink

"The engineering professional's link
to technical literature"

VISIT MILL-MAX'S DESIGN GUIDE ONLINE

Search • Download individual product drawings
and specs.

Browse • 237 pages of standard offerings
available with traditional platings, as well as
RoHS compliant platings;

Create • Design in Mill-Max products from
choices of standard and application-specific parts.
For our new Design Guide and to see new
products, go to www.mill-max.com/EDN567

Mill-Max Mfg.

190 Pine Hollow Road

Oyster Bay, NY 11771

Web: www.mill-max.com



ENCLOSURE POWER DISTRIBUTION UNITS (ePDU™)

Eaton® manufactures enclosure Power distribu-
tion units (ePDU™) providing Circuit Breaker
Protection, EMI/RFI Filtering, Spike/Surge
Protection, Remote Power On/Off/Reboot,
(EPO) Emergency Power Off, Redundant Power,
Sequential Power Up/Down. Domestic and inter-
national configurations for OEMs and system
integrators are available. Agency approved
systems. ISO 9001:2000 Certified.

Eaton – (ePDU™)

Tel: 800-870-2248 • Fax: 605-334-4999

pulzisaales@eaton.com • www.pulizzi.com



ePDU
Enclosure Power
Distribution Units

ADVERTISE IN

EDN literaturelink

For more information on how to
place your ad in this section
please contact:

ROSE MURPHY

T: 781-734-8457

E-mail: murphys@reedbusiness.com

To advertise in Literature Link, call Rose Murphy at 781-734-8457.

GHz BGA/QFN Sockets 0.3mm to 1.27mm

Industry's Smallest Footprint

- Up to 500,000 insertions
- Bandwidth to 40 GHz
- 2.5mm per side larger than IC
- Ball Count over 3500,
Body Size 2 - 100mm
- Four different contactor
options
- Optional heatsinking
to 100W
- Four different Lid Options
- <25 mΩ Contact Resistance
throughout life



Ironwood
ELECTRONICS

1-800-404-0204

www.ironwoodelectronics.com

NEW Digital Humidity Sensor

RH&T

Size that inspires!

SHT21



real size (3 x 3 mm DFN)

- High Functional Integration
- Stable in Extreme Environment
- Suitable for Mass Production

Specifications

Range RH	0% to 100%
Range T	-40°C to +120°C
Accuracies RH/T	max. ± 3% RH / ± 0.4°C
Interface	Digital I ² C
Consumption	1.5 µW typ.

SENSIRION
THE SENSOR COMPANY

SENSIRION Inc. Westlake Village
CA 91361, Phone 805-409-4900
www.sensirion.com

humidity | gas flow | differential pressure | liquid flow

Place your ad in
the next

**Product
Mart!**

Call
1-781-734-8457
for rates and
deadline information.

supplychain

LINKING DESIGN AND RESOURCES

Job board targets distributors

NEDA (National Electronic Distributors Association, www.nedassoc.org) has launched an interactive job board, focusing on authorized electronics distribution and the positions—from entry level, to engineer, to executive—that components distributors require.

“For a number of years, the association has considered but never really decided to implement some kind of job posting,” says Robin Gray Jr (photo), executive vice president of NEDA. “There were some concerned when things were good in the industry about providing a vehicle for everybody to poach people from each other. Now that we are in kind of a difficult time, everyone recognizes that that [scenario] is not likely to happen.”

Gray notes that the not-



for-profit trade association launched the NEDA Industry Career Center to provide displaced workers with another vehicle for networking or finding employment in the industry, to attract new or soon-to-be new graduates who may be interested in pursuing a career in the electronics industry, and to offer NEDA-member companies an additional resource as they try to help employees they may have to lay off.

The targeted resource for online employment connections is open to both NEDA members and nonmembers

aiming to reach qualified candidates. Employers can post jobs online, search for qualified candidates based on specific job criteria, and create an online résumé agent to e-mail qualified candidates daily.

For job seekers, the NEDA Industry Career Center allows résumé posting, browsing and viewing of available jobs based on their criteria, and creation of a search agent to provide e-mail notifications of jobs that match their criteria. The service is free to job seekers; rates for NEDA-member employers start at \$250 for a job posting and at \$350 for nonmembers.

As of June, postings were off to a slow start. But, as Gray notes, “There is always a need for new blood to continually prime the pump in the system.” Visit the site at <http://careercenter.nedassoc.org>.

GREEN UPDATE

POSSIBLE CHANGES TO TOXIC SUBSTANCES CONTROL ACT ECHO ROHS

A proposed amendment to the 1976 TSCA (Toxic Substances Control Act) has been put before the US House of Representatives that could set regulations in the United States similar to those that the ROHS (restriction-of-hazardous-substances) directive set in the European Union. The EDEE (Environmental Design of Electrical Equipment) Act, bill HR2420, aims to “ensure a uniform federal scheme of regulation of restrictions in the use of certain substances in electrical products and equipment in interstate and foreign commerce and for other purposes.” EDEE cites among its aims better trade and the prevention of potential disparities between state laws regarding restrictions on use of toxic substances in electrical products

and equipment “that could create barriers to interstate commerce, domestic and foreign trade, and distort global competition.”

EDEE states that, after July 1, 2010, electronic-industry manufacturers cannot produce any product that contains a concentration value greater than 0.1% by weight of lead, mercury, hexavalent chromium, PBB (polybrominated biphenyl), and PBDE (polybrominated diphenyl ether) as measured in any homogeneous material the product contains. EDEE lists exemptions that include certain medical equipment, equipment with a voltage rating of 300V or more, and some fixed installations. For more information on EDEE, visit: www.ec-central.org/lead-free/hr_2420.pdf.

OUTLOOK

PC MARKET TO SLUMP 6% IN 2009

Gartner Inc (www.gartner.com) reports that worldwide PC sales will likely return to market growth by the fourth quarter but warns that vendors should brace for continuing difficulties in the meantime. The market-research company projects that worldwide PC shipments will reach 274 million units in 2009, a 6% decline from 2008 shipments of 292 million units. The fourth quarter should see growth, setting the stage for a healthy market recovery in 2010 with unit shipments forecast to increase 10.3%.

Gartner in mid-May forecast a 6.6% unit decline in 2009 and in March projected a 9.2% unit decline for the full year. “PC-unit growth was stronger than we expected in all markets but Eastern Europe in the first quarter of 2009,” says George Shiffler, research director at Gartner. “In particular, consumer shipments were much stronger than we anticipated. However, professional shipments continued to struggle, and much of the growth in consumer units was due to vendors’ and the channel’s restocking inventories rather than an upsurge in demand.” As of a mid-July update, Gartner reports that second-quarter PC shipments declined 5% from the second quarter of 2008. The company had previously reported a nearly 10% year-over-year decline for the June quarter.”

Support Across The Board.TM From Design to DeliveryTM



Now, you can have it all.TM

Faster and easier than ever before. Our commitment to customer service is backed by an extensive product offering combined with our supply chain and design chain services – which can swiftly be tailored to meet your exact needs. We have dedicated employees who have the experience to provide the highest level of customer service with accuracy and efficiency. All of our technical experts are factory certified on the latest technologies, providing you the expertise to move projects forward with speed and confidence.

Avnet offers the best of both worlds: extensive product and supply chain knowledge, and specialized technical skill which translates into faster time to market – and the peace of mind that comes from working with the industry's best. **Avnet is ranked Best-In-Class* for well-informed sales reps, knowledgeable application engineers and our design engineering services – proof that we consistently deliver:**

- > Industry recognized product expertise
- > Specialized technical skills

Ready. Set. Go to Market.TM

Visit the Avnet Design Resource CenterTM at:
www.em.avnet.com/drc



Accelerating Your SuccessTM


1 800 332 8638
www.em.avnet.com



*As rated by Hearst Electronics Group: The Engineer & Supplier Interface Study, 2009.
©Avnet, Inc. 2009. All rights reserved. AVNET is a registered trademark of Avnet, Inc.

productroundup

DISCRETE SEMICONDUCTORS



High-performance MOSFETs implement superjunction technology

▾ Aiming at energy-conversion applications, such as PFC or PWM stages, the high-performance, 600V CoolMOS C6 MOSFET series enables control of switching behavior and high body-diode ruggedness. The technology combines the characteristics of superjunction or compensation devices, such as ultralow area-specific on-resistance, with reduced capacitive switching losses. The IPA60R190C6 comes in a TO-220 full pack with a 190-m Ω on-state resistance and costs 88 cents, and the IPD60R950C6 comes in a DPAK with a 950-m Ω on-state resistance and costs 30 cents (10,000).

Infineon Technologies, www.infineon.com

30V P-channel MOSFET comes in an SO-8 package

▾ Part of the vendor's TrenchFET Fen III P-channel-MOSFET family, the 30V Si7145DP provides a 2.6-m Ω maximum on-resistance at 10V gate drive and a 3.75-m Ω on-resistance at 4.5V. The MOSFET suits use as an adapter switch and for load-switching applications in notebook computers and industrial and general systems. A low on-resistance results in lower conduction losses, enabling adapter switches to switch between the adapter or wall pow-

er or the battery power, always on and drawing current. Available in a PowerPak SO-8, the Si7145DP costs 69 cents (100,000).


Vishay Intertechnology, www.vishay.com

MOSFET series aims at secondary stage of SMPS

▾ Targeting synchronous rectification during the secondary stage of SMPS, the MOSFET series comprises four devices featuring low on-state resis-

tance for reducing power loss. On-state resistance in a second-stage MOSFET can account for 50% of the total power loss. The series suits use in ac-power adapters or ac/dc-power supplies in computing and consumer-electronics applications. You can use the three 75V devices and the 100V device as switch or synchronous MOSFETs in the secondary stage of SMPS. Available in a TO-220SIS package, the 75V TK80A08K3 has a 40A current rating with a 7-m Ω on-state resistance, the 75V TK80A08K3 has an 80A current rating and a 3.6-m Ω on-state resistance, and the 100V TK40A10K has a 40A current rating and an 11.5-m Ω on-state resistance. The 75V, 80A TK80D08K3 has a 3.6-m Ω on-state resistance and comes in a TO-220W package. Prices for the MOSFET series start at 80 cents (high volumes).

Toshiba America Electronic Components, www.toshiba.com/taec



ORing-controller chip replaces Schottky blocking diodes

▾ Targeting control of MOSFETs in high-reliability, N+1 redundant power systems, the XZGD3102 active ORing-controller chip enables shared-power-system designers to replace heat-dissipating blocking diodes with high-efficiency MOSFETs. As an alternative to Schottky blocking diodes with high heat dissipation and 500-mV forward-voltage drop, the controller chip provides typical voltage drops of less than 100 mV,

reducing system-power dissipation. Sinking a 5A peak turn-off current, the gate driver achieves a 160-nsec MOSFET-turn-off time with a 180V input-voltage blocking capacity. The chip uses a 5 to 15V low-current bias rail. Available in an SM8 package, the ZXGD3102 costs 98 cents (1000).
Diodes Inc, www.diodes.com

Dual MOSFET targets synchronous-buck applications

➔ Aiming at notebooks, netbooks, servers, telecom, and other dc/dc designs, the Dual-MOSFET FDMC8200 integrates a high-side control and low-side synchronous

30V N-channel MOSFET, suiting synchronous-buck applications. Features include a high-side, 24-mΩ on-resistance; a low-side, 9.5-mΩ on-resistance; and a 9A current. Available in a 3×3-mm Power33 MLP module, the device costs 50 cents (1000).
Fairchild Semiconductor, www.fairchildsemi.com

COMPUTERS AND PERIPHERALS

43-in. curved display has double-WGXA resolution

➔ The 43-in. CRV43 curved display has a 32-to-10 aspect ratio, providing double-WGXA resolutions. Features include 200-cd/m² brightness, 0.02-msec response, and a 10,000-to-1 contrast ratio. The display provides single-link DVI-D and HDMI 1.3 input connectors, front-panel controls, and on-screen display and software-based GUIs. The CRV43 costs \$7999 with a three-year parts and labor warranty.
NEC Display Solutions, www.necdisplay.com

Wide-screen LCD reduces power consumption and heat generation

➔ The 22-in., wide-screen, ecologically friendly MultiSync E222W LCD uses four-way-adjustable ErgoDesign, featuring 110-mm-high adjust, tilt, swivel, and pivot. The display's ECO mode reduces power consumption and heat generation, and the intelligent-power management and off-timer conserve energy and reduce carbon-dioxide emissions by switching to a low-power state or automatically powering down when the monitor is on but not in use. Features include 1680×1050-pixel native resolution in a 16-to-10 aspect ratio, 250-cd/m² brightness, and 5-msec response time. The MultiSync E222W costs \$269 with a three-year

parts-and-labor warranty.
NEC Display Solutions, www.necdisplay.com

High-performance chassis suits use with liquid-cooled devices

➔ The Obsidian 800D high-performance-chassis series comes in a steel structure with a matte-black coating and a brushed-aluminum faceplate. Supporting five 5.25-in. and six 3.5-in. drives, the chassis are compatible with ATX, Extended ATX, and Micro ATX motherboards. The chassis suits use in high-end systems and in liquid-cooled devices. The Obsidian 800D chassis costs \$299 and has a two-year warranty.
Corsair, www.corsair.com

LED backlit displays include resistive touchpanels

➔ Joining the vendor's LED-backlit-display line, the 5-in. LQ050W1LA0A panel with resistive touch and the 7-in. LQ070Y3DG3A TFT-LCD panel feature integrated LED drivers and a 3.3V-dc supply voltage for TFT-LCD-panel driving and backlight operation. The 5-in., 1024×3×600-dot panel with 262,144 colors uses an LVDS interface; the 7-in. panel displays graphics and text on an 800×3×480-dot panel with 16,194,277 colors using a 24-bit digital-signal interface. The 5-in. panel and the 7-in. panel cost \$169 and \$130, respectively.
Sharp Microelectronics of the Americas, www.sharpsma.com

EDN ADVERTISER INDEX

Company	Page
Agilent Technologies	11
austriamicrosystems AG	52
Avnet Electronics Marketing	55
BuyerZone	C-3
Cirrus Logic Inc	C-2
Digi-Key Corp	1
Eaton Corp	53
International Rectifier Corp	3
Ironwood Electronics	53
Linear Technology Corp	42
	45, 46
Maxim Integrated Products	49
	51
Mentor Graphics	C-4
Mill Max Manufacturing Corp	53

Company	Page
Monolithic Power Systems	17
National Instruments	12
NEC Electronics	C-1
Pico Electronics	4
	33
Samsung Electro-Mechanics	16
Samtec USA	24
Sensirion AG	53
Stanford Research Systems Inc	41
Xilinx Inc	23

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

Dark side of the light



I was trying to get rid of the last nasty bugs in my electronics system. My company builds laser systems, and this system was the latest in a series of small, handheld laser designators for a military customer. I had designed a small control board to operate the laser, and the software was giving me a lot of trouble. I implemented background interrupts for timing-driven control loops, user-interface devices for control and configuration, and a data-communications SPI (serial-peripheral interface) connecting the microcontroller, an FPGA, and a few peripheral devices. Getting everything to behave had been difficult, but I thought I had solved most of the problems and believed I was in the home stretch.

I was wrong. Some intermittent behavior that I had attributed to earlier bugs refused to go away, and I was under pressure to finish the system so we could package and ship it out for a demonstration. I would start the device up, and, within minutes, some odd behavior would begin. An LED that

displayed the operating mode of the controller would suddenly start blinking in a pattern that didn't correspond to any normal state. More alarmingly, the system would sometimes emit laser pulses when it wasn't supposed to. The variety of bad behaviors indicated a problem with the SPI, which was responsible for programming the FPGA registers that we used to configure and operate the laser.

I was resigned to having my electronics technician tack-solder a bunch of leads to various SPI-related signals and connecting them to a digital oscil-

loscope to capture and analyze the SPI traffic. If I could find a corrupt message, I could trace it back to its source code and find the bug. There were a lot of SPI messages routinely zipping back and forth on the SPI bus, so I knew it would be a search for the proverbial needle in a haystack.

The morning after my technician had prepared the board for its date with the oscilloscope, we started to talk about the debugging plan. He told me that he thought he had identified the source of the problem. "It's the lamp we're using on the lab bench," he said. "The light is giving off noise and messing up your software." I was excited to learn he had a theory, but its unlikelihood immediately let me down.

"It's not the light. How could the light be causing the software to get confused?" I said, somewhat scornfully. I turned on the light and shined it at the board, and the LED pattern immediately changed from normal to anomalous. I was stunned. I put a piece of paper between the light and the board, cycled power on the board, and watched the LED blink happily. I removed the paper, and the LED pattern went bad again.

By cutting a hole in the paper about the size of an IC, I was able to selectively direct the light onto each component on the board and soon identified the problem: Light shining on a Bluetooth-interface IC caused the corruption. The Bluetooth chip was connected to the SPI bus for programming, although the connection was unnecessary because the chip's default settings worked fine. I solved the problem by cutting the SPI lead connections to the Bluetooth part, and the controller buzzed along peachily ever after. **EDN**

Edward Sullivan is an engineer at Fibertek Inc (Herndon, VA).

➦ Go to www.edn.com/090806tales and click on Feedback Loop to post a comment on this article.



Looking to buy? With BuyerZone, every big deal is no big deal at all.


BuyerZone is the trusted, unbiased site that brings buyers and sellers of business purchases together to make the purchasing process easier while saving both parties time and money. Credit Card Processing, Building Services, Home Security, Trade Show Displays, Forklifts and Construction Equipment, Business Purchases and more... Let us do the work so you don't have to.

Join the millions who've already saved time and money on many of their 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) 623-5525 or visit
BuyerZoneQuotes.com

BuyerZone
Where Smart Businesses Buy and Sell

A division of
 Reed Business Information.

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