Reed Electronics Group

EDN

VOICE OF THE ENGINEER

Oct 13
Issue 21/2005
www.edn.com

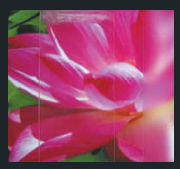

NAND versus NOR

WHICH FLASH IS BEST FOR BOOTIN' YOUR NEXT SYSTEM? PAGE 41

OPTIMIZED LEARNING IN METRO SWITCHES

ETHERNET EXTENDS INTO AGGREGATION AND CORE NETWORKS

PAGE 69

INSTIGATING A PLATFORM TUG OF WAR:

GRAPHICS VENDORS HUNGER FOR CPU SUPPLIERS' TURF

PAGE 50

INTERLEAVING DC/DC CONVERTERS BOOST EFFICIENCY AND VOLTAGE

PAGE 77

The secrets of successful communications using LVDS

PAGE 89

One Chip, One Solution

Integrated Handheld Power Management

The TPS65800 power management IC extends battery life in portable designs while occupying up to 70% less board space than discrete solutions. Designed for multimedia and communication devices powered by a 1-cell lithium-ion battery, the TPS65800 integrates complete battery management, highly efficient power conversion, including a white LED driver, and an I²C control interface in an 8 x 8 mm² QFN package.

Save Up to 70% Board Space

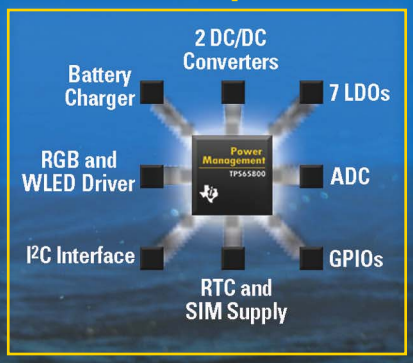

High Performance. Analog. Texas Instruments.

For a datasheet, sample and evaluation module, visit **www.ti.com/tps65800**

Technology for Innovators™

Texas Instruments

Digi-Key is the leader when it comes to breadth of product line... It's that simple. With access to a half-million different part numbers from over 320 quality, name-brand manufacturers, Digi-Key is truly the one-stop source for all your needs.

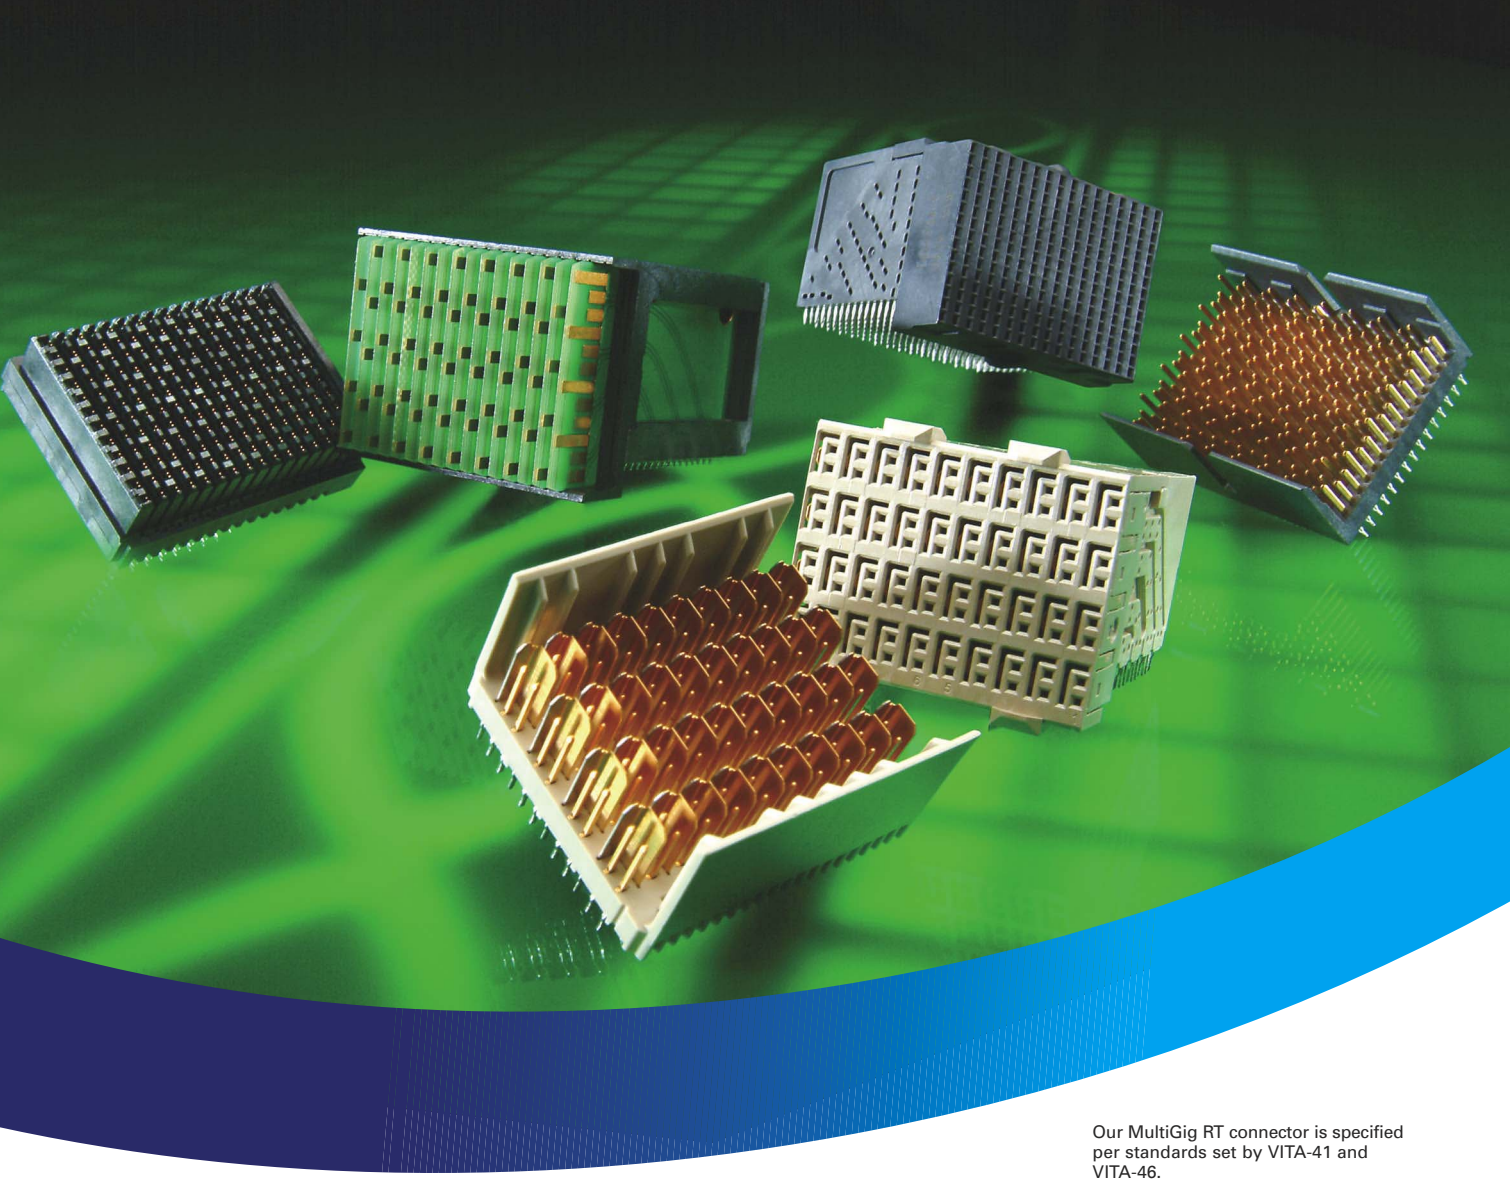

At Tyco Electronics, we have more ways to bring you up to speed.

Getting the high speed performance you need is one thing. Achieving it at the cost you had in mind is quite another. That's why, at Tyco Electronics, we offer more options when it comes to interconnect solutions. Our depth of products allows us to not only meet your precise cost and performance requirements of today, but also help you migrate towards your designs of tomorrow. We even offer extensive simulation and modeling tools that help ensure you get the performance to meet your needs. And with so many options, there's definitely one thing we don't offer: a bias. For more information visit **www.tycoelectronics.com/productinfo** and request our high-speed connector catalog.

Our MultiGig RT connector is specified per standards set by VITA-41 and VITA-46.

The Z-PACK HM-Zd connector is specified by PICMG® consortium as Advanced TCA® standard for Zone 2 Advanced Differential Fabric connector.

Our Z-PACK MAX high performance back plane connector systems are designed for cost pressured, high density applications.

North America 800-522-6752 • Europe +49 6251 1330
• Asia +81-44-844-8292 • www.tycoelectronics.com

a vital part of your world

tyco Electronics

POWER MANAGEMENT

Next Generation Power Modules

Pb-free and 50% Smaller

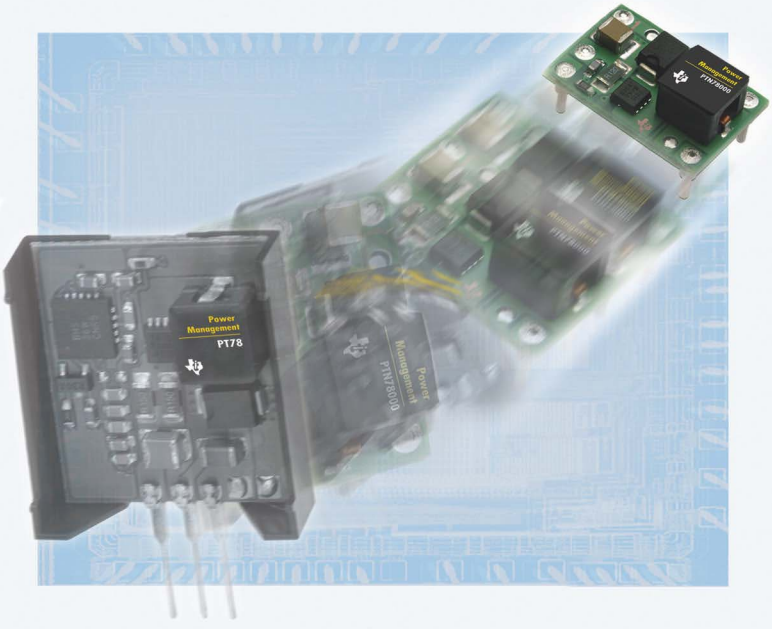

▸Applications

- Industrial controls, HVAC systems
- Test and measurement
- Medical instrumentation
- AC/DC adaptors
- Vehicular, marine and avionic electronics

▸Features

- High efficiency
- Wide input voltage range 7-36V
- Adjustable output voltage
- On/off inhibit
- Current limit/over-temp shutdown
- Operating temp –40°C to 85°C
- Pb-free

The new **PTN78 series** is the next generation of TI's popular PT78 series wide-input wide-output general purpose power module – it's 50% smaller, higher performance and more cost effective. Designed to be Pb-free, RoHS compatible and high-temperature solder process capable, the new PTN78 series also includes adjustable output voltage and on/off inhibit. Choose either positive or negative output voltage. The best just got better!

Device	V_{IN} (V)	I_{OUT}	V_{OUT} (V)
PTN78000W/H	7 to 36	1.5A	2.5 to 12.6/12 to 22
PTN78060W/H	7 to 36	3A	2.5 to 12.6/12 to 22
PTN78020W/H	7 to 36	6A	2.5 to 12.6/12 to 22
PTN78000A	7 to 29	1.5A	–3 to –15
PTN78060A	9 to 29	3A	–3 to –15
PTN78020A	9 to 29	6A	–3 to –15

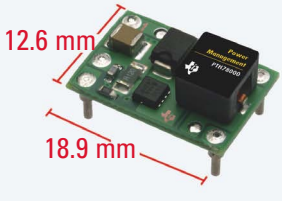

NEW!
Power Management Selection Guide
Data Sheets, Samples, Plug-in Power Selector Guide

www.ti.com/ptn78000 ◦ 800.477.8924, ext. 2355

Technology for Innovators™

Texas Instruments

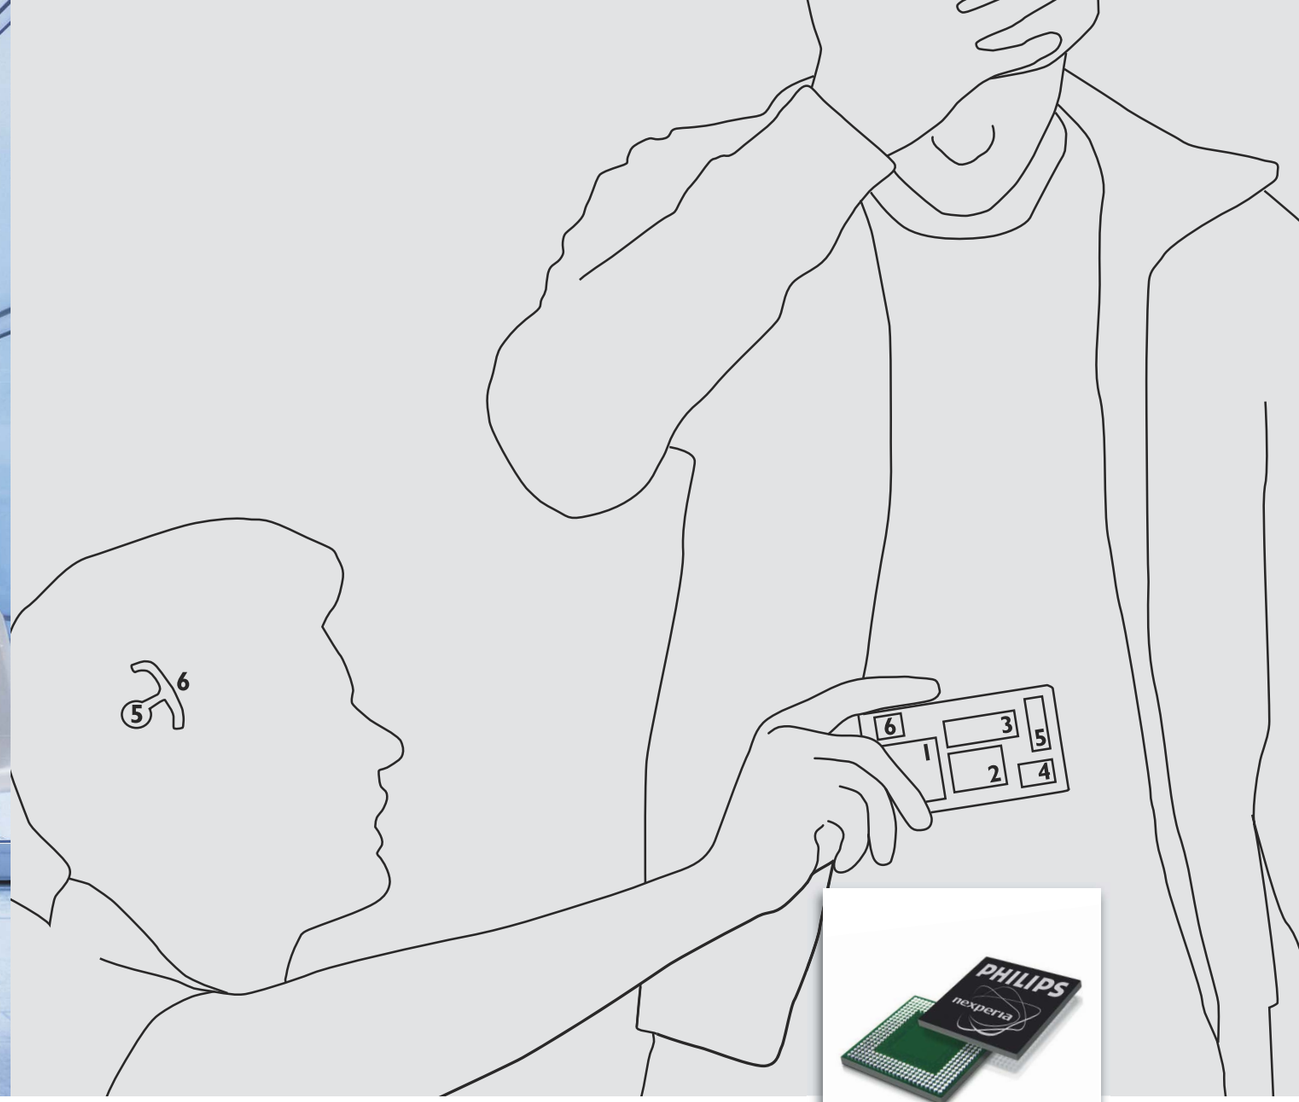

Whatever your next innovative idea is, we'll help you design it.

Philips Nexperia Mobile Solutions. A mobile phone that can buy tickets to the game. A mobile phone that allows you to watch the game when you can't get to the stadium. A mobile phone that takes a picture of your friend coming to terms with his team's loss and shares it with your other friends. People buy their mobile phones today expecting much more than just a mobile phone. That's why you need an adaptable technology solution that can be tailored and changed as fast as your consumer needs. With Philips Nexperia Mobile Solutions, you'll find flexibility in platform design to customize your handset offering for different features and markets, and a total system approach that enables easy porting of various software applications. We address every type of GSM/GPRS/EDGE/UMTS/TD-SCDMA mobile phone, and provide expertise in value-added features such as connectivity, security, broadcast and multimedia. So no matter what's next in mobile phones, Philips will help you design it.

nexperia

www.philips.com/nm

Philips Nexperia Mobile Solutions. To help you design what's next, Philips has complete handset reference designs incorporating the following functions:

1. **Display** – Complete display sub-assemblies and LifePix™ enhancement algorithms bring multimedia content to life.
2. **Cellular** – All the core system elements, including proven Nexperia baseband, power management, RF, PA, and software protocol stack.
3. **Multimedia** – Advanced Nexperia processors, and partnerships with software suppliers for exciting audio/video/gaming features.
4. **Broadcast** – DVB-H for TV-on-Mobile broadcasts, and the smallest FM radios in the industry.
5. **Connectivity** – Bluetooth and WiFi coexistence solutions, and NFC for easy network setup, and mobile ticketing and payment.
6. **Speaker** – Excellent quality speakers and microphones enhance the audio/visual experience.

PHILIPS

sense and simplicity

The first 1.8 V auto-zero instrumentation amp. Best-in-class performance for the lowest price.

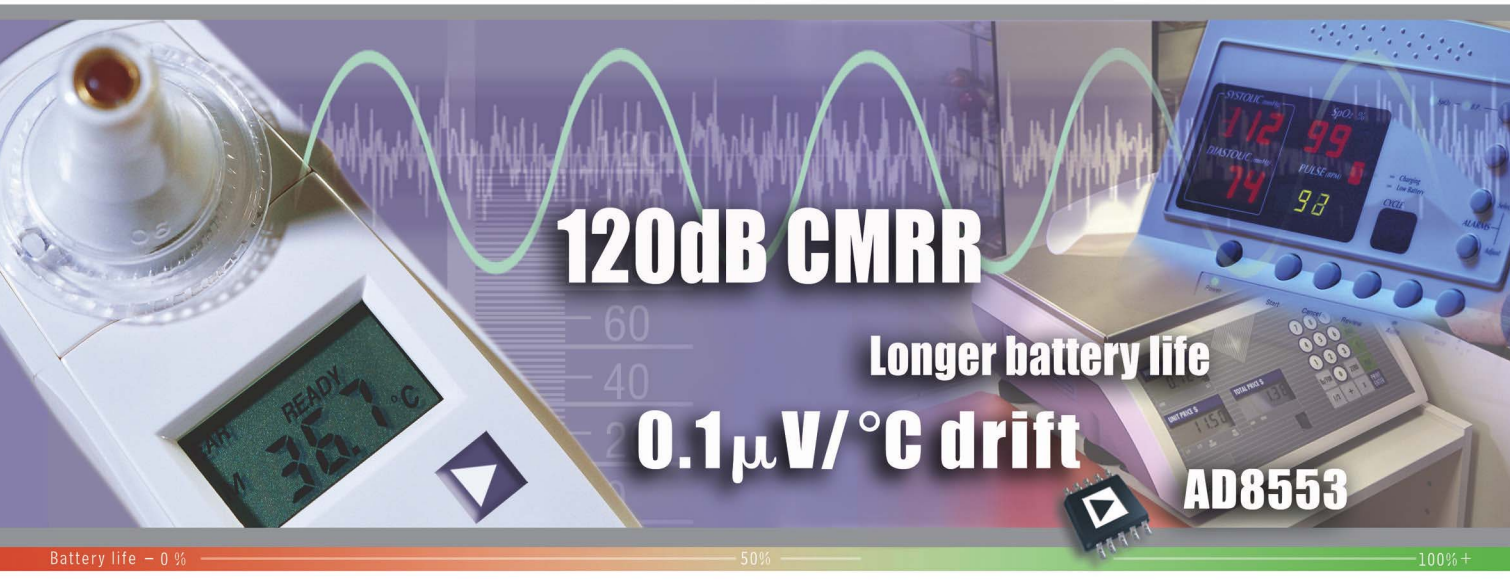

AD8553 Performance…

- High CMRR: 120 dB min
- Low offset drift: 0.1 μV/°C max
- Low offset: 25 μV max
- Low noise: 0.7 μV p-p (0.01 Hz to 10 Hz)
- Pb-free package: 10-lead MSOP
- Price: $1.30/1k quantities

…where it matters

- Pressure sensors
- Strain gages
- Medical instrumentation

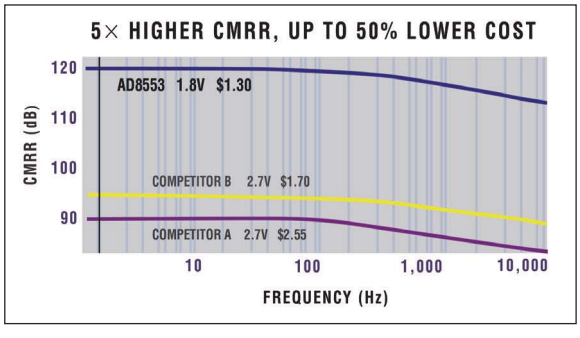

Precision and value for low voltage designs

If you're designing medical or industrial systems, our new AD8553 combines all the in-amp features you wanted for portable applications. ***Extended battery life***—with operation from 5.5 V down to 1.8 V and a shutdown mode that reduces total supply current to less than 4 μA, the AD8553 is the most battery-friendly in-amp in the industry. ***Excellent signal conditioning***—high CMRR, low offset voltage, low drift, and low noise combine to make the AD8553 ideal for precision applications. ***Design simplicity***—forget using discrete solutions that require laser trimming and more board space, or in-amps that require an external op amp to drive the V_{REF} pin. ***Lowest cost***—In addition to better performance and power savings, the AD8553 costs 30% to 50% less than the competition. You can't build or buy better in-amp performance and value for your next low voltage design. To learn more, please visit our website.

40 years ADI 1965 2005 — Innovation • Performance • Excellence

www.analog.com/autozero-inamp

ANALOG DEVICES

THE LEADER IN HIGH PERFORMANCE ANALOG

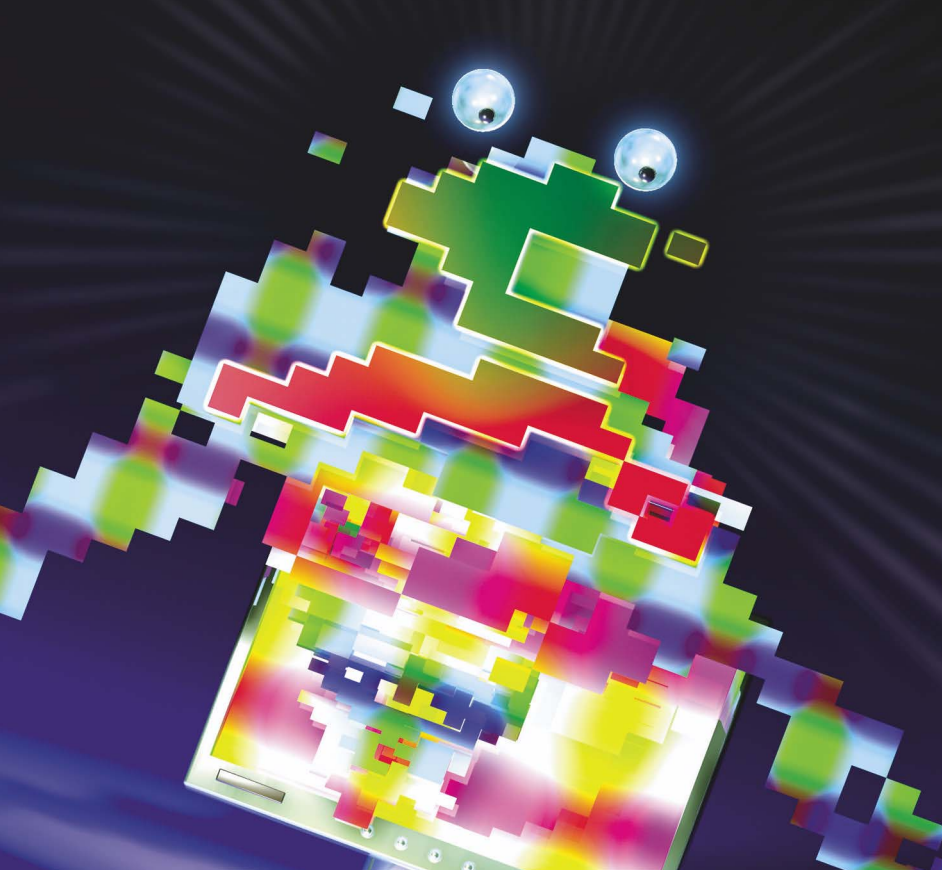

EDN

contents

10.13.05

DESIGNIDEAS

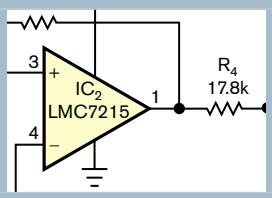

High Speed, Top Performance!

13-Bit, 250-MSPS ADC

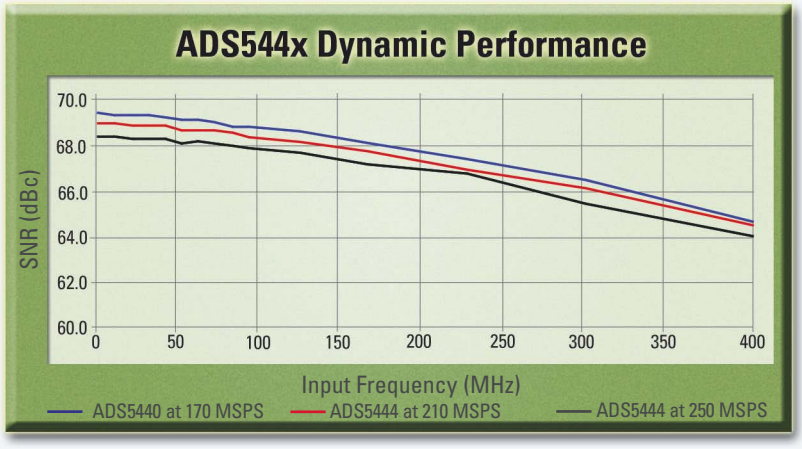

Sample Application: Wideband, High IF DPD Feedback Receiver

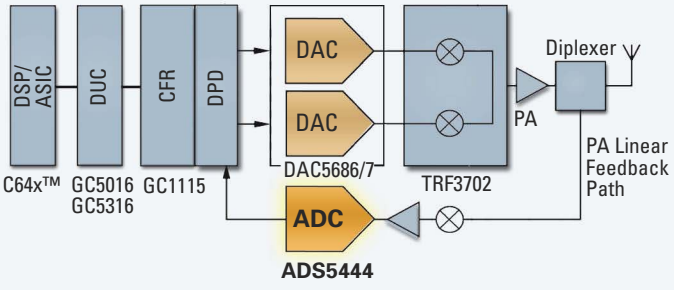

The new **ADS5444** from Texas Instruments sets a new benchmark for high-speed ADCs, providing best-in-class performance at 250 MHz. Look to TI for a complete portfolio of high-speed ADCs, including the recently announced ADS5440 13-bit, 250 MSPS ADC.

Device	Resolution (Bits)	Speed (MSPS)	SNR (dBc)	SFDR (dBc)
ADS5444	13	250	68 at 230 MHz IF	75 at 230 MHz IF
ADS5440	13	210	68 at 230 MHz IF	79 at 230 MHz IF
ADS5500	14	125	69.5 at 100 MHz IF	82 at 100 MHz IF
ADS5424	14	105	74 at 50 MHz IF	93 at 50 MHz IF
ADS5541	14	105	71 at 100 MHz IF	86 at 100 MHz IF
ADS5423	14	80	74 at 50 MHz IF	94 at 50 MHz IF
ADS5520	12	125	68.7 at 100 MHz IF	82 at 100 MHz IF
ADS5521	12	105	69 at 100 MHz IF	86 at 100 MHz IF

▶ Applications

- Software-defined radio
- Base stations:
  - Wideband receiver
  - High IF receiver
  - PA linearization
- Instrumentation
- Test and Measurement

▶ Features

- 100 MHz IF: SNR = 68.7 dBc; SFDR = 73 dBc
- 230 MHz IF: SNR = 68 dBc; SFDR = 75 dBc
- Fully buffered analog inputs
- 2.2 Vpp differential input voltage
- 3.3 V LVDS compatible outputs
- TQFP-80 PowerPAD™ package
- Industrial temperature range –40°C to +85°C
- Price: $95 1k

New!
Amplifier and Data Converter Selection Guide, Datasheets, Samples, Evaluation Modules

www.ti.com/ads5444 ○ 800.477.8924, ext. 2695

Technology for Innovators™

Texas Instruments

contents 10.13.05

pulse

Dilbert 24

24

38

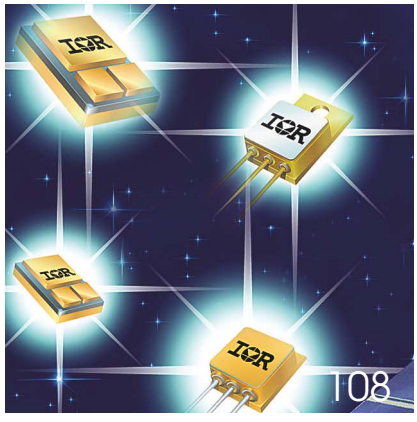

108

DEPARTMENTS & COLUMNS

PRODUCT ROUNDUP

EDN® (ISSN#0012-7515), (GST#123397457, R.B.I. Intl Pub Mail #0280844) is published biweekly, 26 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. Jim Casella, Chief Executive Officer; Stephen Moylan, President, Boston Division. Periodicals postage paid at Littleton, CO 80126 and additional mailing offices. Circulation records are maintained at Reed Business Information, 8878 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 2005 by Reed Elsevier Inc. Rates for nonqualified subscriptions, including all issues: US, $149.90 one year; Canada, $205.90 one year, (includes 7% GST, GST#123397457); Mexico, $195.90 one year; air expedited, $361.90 one year. Except for special issues where price changes are indicated, single copies are available for $10.00 US and $15.00 foreign. Publications Mail Agreement No. 40685520. Return undeliverable Canadian addresses to: Deutsche Post, 4960-2 Walker Road, Windsor ON N9A 6J3. E-mail: subsmail@reedbusiness.com. Please address all subscription mail to EDN®, 8878 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 50, Number 21 (Printed in USA)

Design Today, Prototype Tomorrow

Superior Design Tools and Resources for the Analog Designer

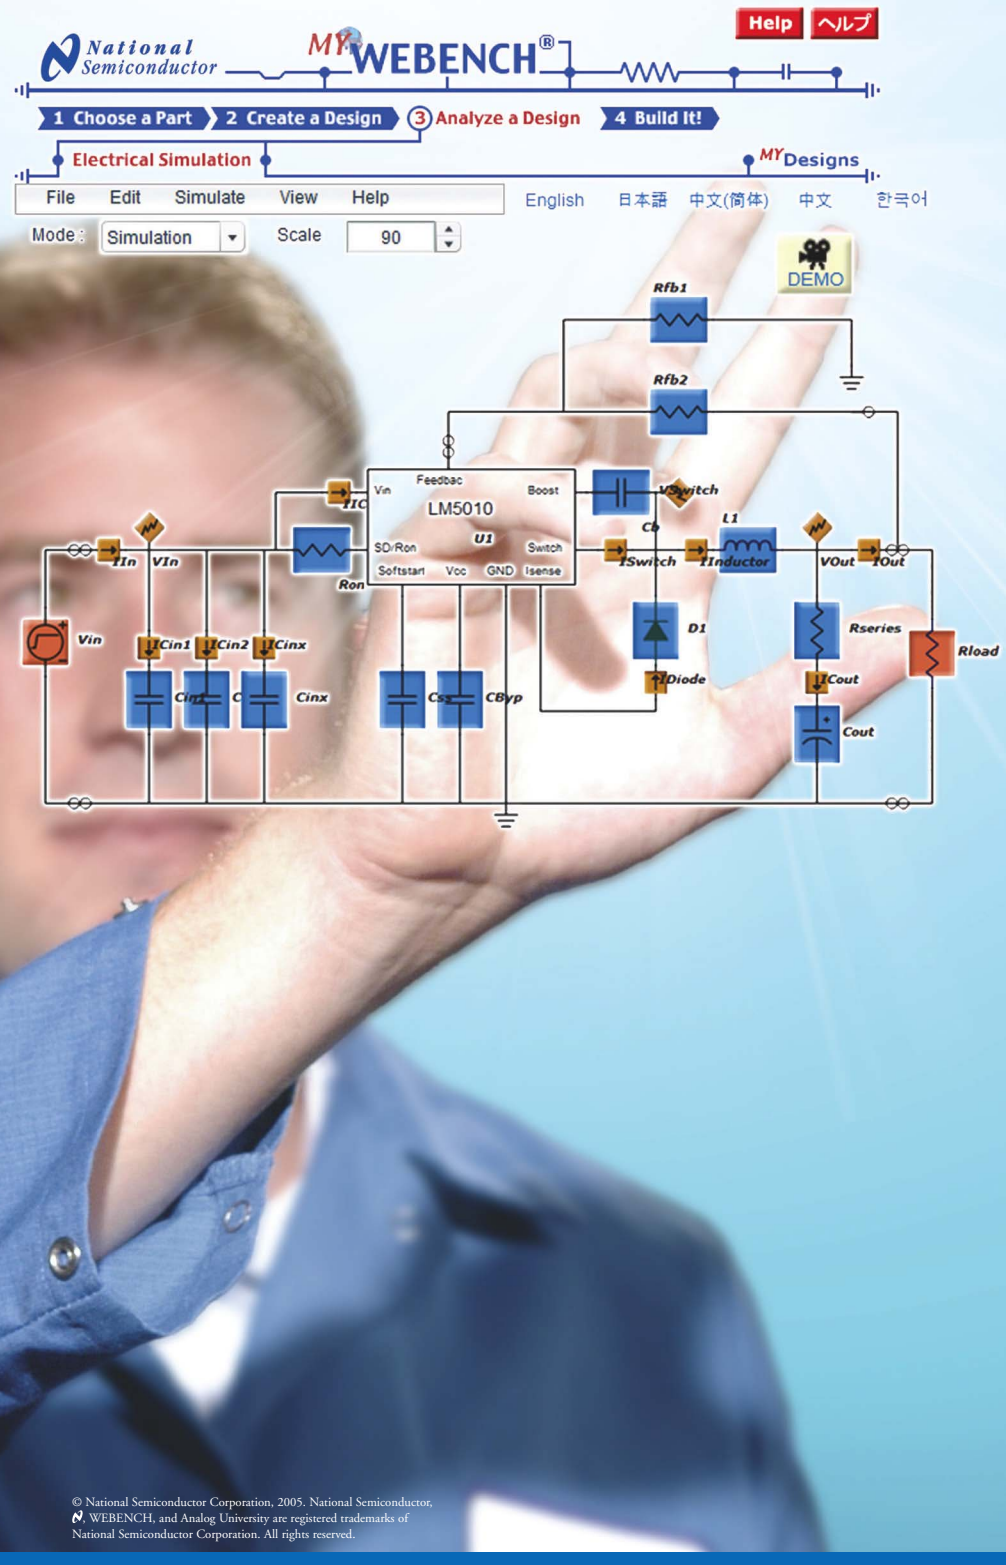

WEBENCH® Online Tools

Get to market faster by accelerating your design process. WEBENCH® tools allow you to select products, design, and receive a custom prototype kit in 24 hours.
webench.national.com

Application Solutions

Over 100 interactive system block diagrams with a list of recommended parts for: automotive, broadcast video, data communications, displays, industrial, medical, and consumer applications.
solutions.national.com

24/7 Online Seminars

View over 50 on-demand design focused seminars for power, amplifiers, audio, data acquisition, interface, and thermal management.
www.national.com/onlineseminar

Analog University®

Get your design right the first time. Log on to Analog University for fast, free information from the pros at National, including in-depth design courses.
analogU.national.com

National Semiconductor
The Sight & Sound of Information

EDN online contents

www.edn.com

ONLINE ONLY

Check out these online-exclusive articles:

Opinion: DFM alliance depends on EDA cooperation

By Michael Santarini, Senior Editor

Don't be surprised if the design-for-manufacturability effort, though noble, takes more time than projected.

→ www.edn.com/article/CA6260497

Bluetooth SIG offers qualification "cookbook"

The online and downloadable tools aim to make the certification of new products a step-by-step affair.

→ www.edn.com/article/CA6260461

More mind-boggling math: Adding and subtracting unsigned BCD

Clive "Max" Maxfield continues his series on binary-coded-decimal math.

→ www.edn.com/article/CA6258681

Meeting the demands of video on demand

As consumers in greater numbers adopt video-on-demand services, cable operators must make sure their networks can rise to the occasion.

→ www.edn.com/article/CA6256544

VOIP chip enables multimedia-rich cordless phones

Broadcom's new chip aims to imbue WiFi-based cordless telephones for home use with features such as video streaming.

→ www.edn.com/article/CA6258645

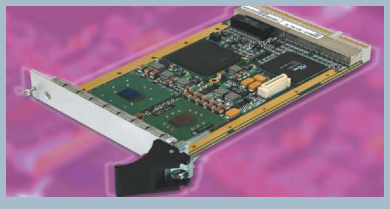

Single-board Pentium M computer targets military use

SBS Technologies says its new 3U CompactPCI-based single-board computer tolerates harsh environments and offers low voltage and low heat.

→ www.edn.com/article/CA6258204

Opinion: When RFI overwhelms RKE

By Bill Schweber, Executive Editor

With wireless systems, engineers must keep in mind that the transmission path they are counting on resides in a crowded swamp.

→ www.edn.com/article/CA6252229

Compression tool shrinks design files by 20×

SoftJin's file-compression program allows users to transfer ever-growing GDSII IC layout files more quickly.

BRIAN'S BRAIN

Is the Xbox a curse for graphics-chip companies? Is Apple's iPod nano a killer of miniature hard disks? Are consumers ready for network-attached storage? Are graphics companies taking a page from PT Barnum (relying on suckers to make a buck)? The answers (or at least opinions) are in Brian's Brain, a blog by *EDN* Senior Technical Editor Brian Dipert. This week, Brian's Brain also contains a raft of information supplementing Brian's graphics-chip cover story in this issue.

→ www.edn.com/briansbrain

EDN Innovation 16th ANNUAL AWARDS

INNOVATION AWARDS: CALL FOR NOMINATIONS

We're now accepting nominations for the 16th annual *EDN* Innovation Awards. Learn how to nominate your company's products and engineers using our easy online process. Nominations close November 23, after which *EDN*'s readers will pick the winners. We'll hand out the awards in early April 2006.

→ www.edn.com/innovation

FROM THE VAULT

Articles and extras from the EDN *archives that relate to this issue's contents.*

Related to …

NAND versus NOR (pg 41):

Hitting their stride

Nonvolatile-memory upstarts draw near established leaders.

→ www.edn.com/article/CA495735

Flash forward to the future

The global digital-consumer trend signals a potential changing of the guard in the flash-memory industry.

→ www.edn.com/article/CA474720

Related to …

See beyond the edge (pg 36):

VNAs and TDRs

Taking the measure of the new millennium.

→ www.edn.com/article/CA633462

Signal-integrity modeling of gigabit backplanes, cables, and connectors using TDR

→ www.edn.com/article/CA231568

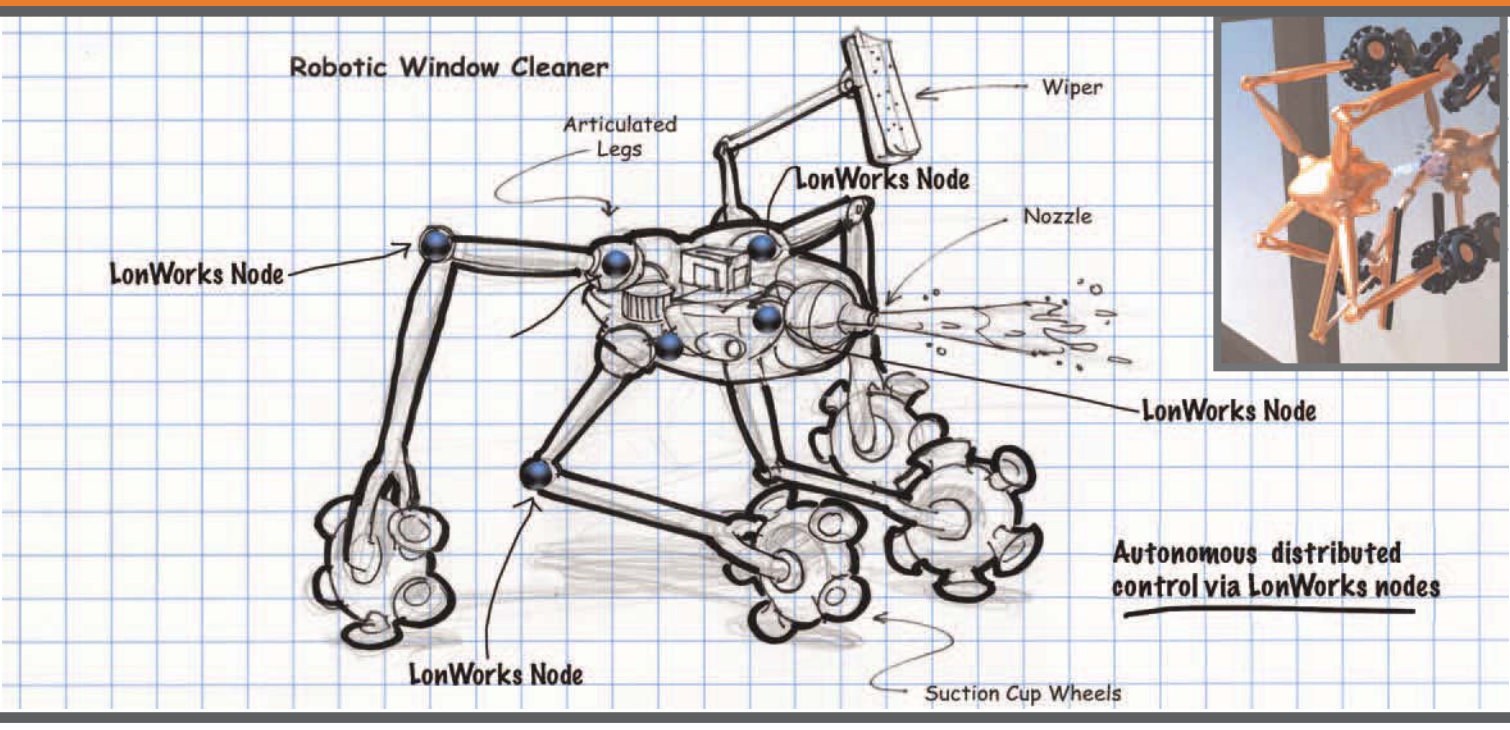

DESIGN WITHOUT LIMITS

What you design is your business. Adding distributed control throughout your application is ours. Our LonWorks® platform is a worldwide standard for networking anything you want to control—pumps, motors, valves, fans—even a remote-controlled window washer. Add networking features to your products and watch them deliver more than you thought possible. Start today with our NodeBuilder® development tool, or ease into your LonWorks development with our new Mini EVK. Call +1 408-938-5200 lto earn more about our development tools, or visit us online at www.echelon.com/embedded.

ECHELON®

ALPS®

Innovation driven by a commitment to excellence.

Striving to create, nurture and accumulate the highest quality electronic components, ALPS continues to forge ahead in an ever-changing marketplace — just like we have for over fifty years.

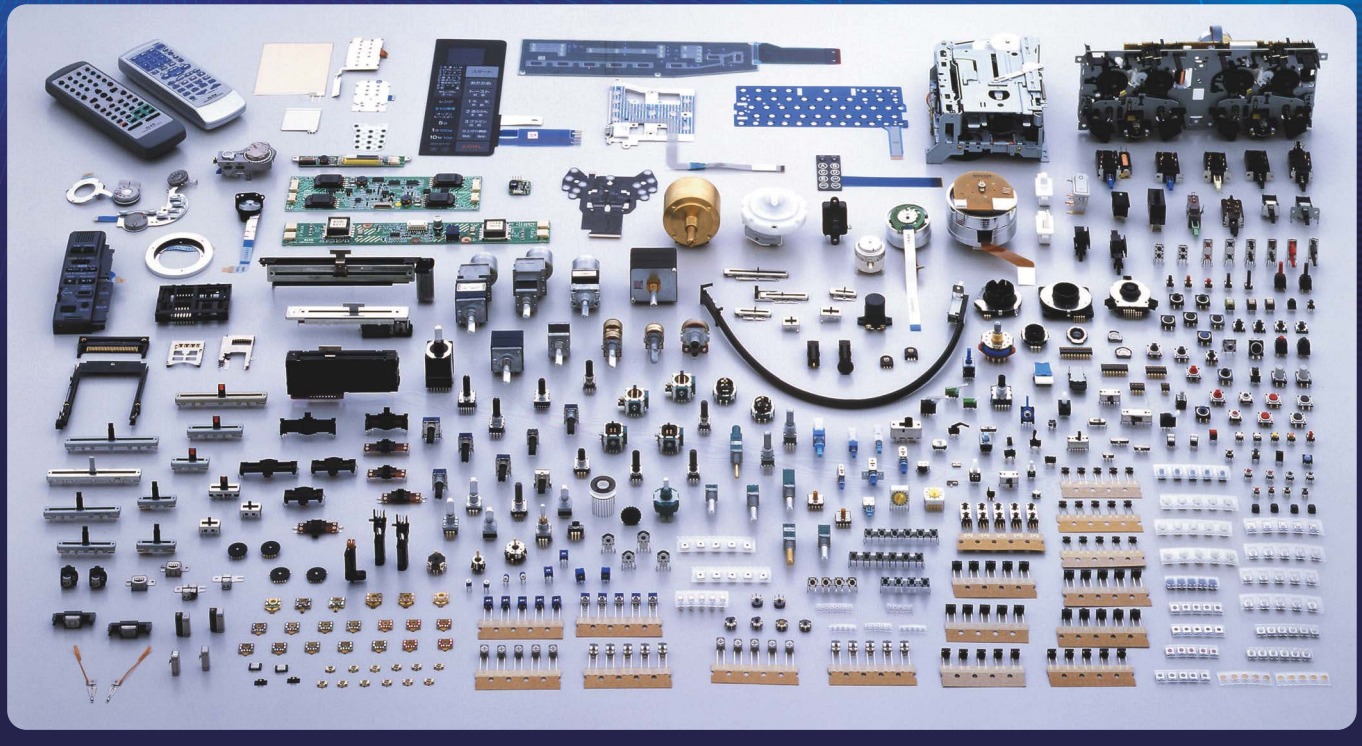

Visit our Web site for more details:

http://www.alps.com

ALPS

ALPS ELECTRIC CO., LTD.

ALPS ELECTRIC (USA), INC. *http://www.alpsusa.com*
910 E. Hamilton Avenue Suite #500 Campbell, CA 95008, U.S.A. Phone +1(408)361-6400 Fax +1(408)226-7301
HEAD OFFICE *http://www.alps.co.jp*
1-7, Yukigaya-otsuka-cho, Ota-ku, Tokyo, Japan, 145-8501 Phone. +81(3)3726-1211 Fax. +81(3)3728-1741

BY JOHN DODGE, EDITOR IN CHIEF

Industry rocks the night away

Music has long been a release for many working in the high-pressure electronics and IT industries. In the '80s, Borland founder Philippe Kahn held jams at the once-huge, now-defunct Comdex computer conference (**Link 1**). The music was so-so, but anyone could play, showing that people working 60 hours a week did have lives. The music is better at the Demo conference, where innovators get to show off new products and ideas to a crowd of influencers and investors (**Link 2**). Long-time *Wall Street Journal* high-tech reporter Don Clark always makes Neil Young proud at this gathering of high tech's glitterati. One year, Shawn Colvin showed up.

The electronics industry now has dozens if not hundreds of bands and musicians to call its own. One noteworthy troupe is Full Disclosure, whose members include rivals from Cadence and Synopsys, along with Gartner Dataquest analyst Gary Smith, who tracks the EDA industry. Full Disclosure plays rhythm and blues and is in its glory every year at DAC.

The electronics industry now has dozens if not hundreds of bands and musicians to call its own.

Another electronics-industry band, Spurious Freedom, aspires to play fundraisers for charities. Formed in 2003, Spurious Freedom takes its name from "spurious-free dynamic range"—a key specification in ADCs. Fame may not be the band's forte, but it hasn't entirely managed to avoid the spotlight. Spurious Freedom was scheduled to compete Oct 7 and 8 at the Rock and Roll Hall of Fame's annual Battle of the Corporate Bands in Cleveland after winning a regional contest in Los Angeles on July 9. The band specializes in '70s and '80s rock tunes.

The band's first show celebrated 100 days of success regarding a bone-marrow transplant for bass guitarist Shawn Eubanks' daughter, who is battling leukemia. These six middle-aged minstrels and one sound engineer hail—or hailed—from Texas Instruments' Analog Division. All engineers, they include guitarist Michael Lanz, vocalist Thomas Armendarez, keyboardist Michael Ashton, drummer Tony Zizzo, rhythm guitarist Lon Mitchell, sound engineer Lenard Milholland, and guitarist Eubanks.

The electronics industry has lately been no stranger to good causes. Another example is a joint campaign by Xilinx, Linear Technologies, and Summit Microelectronics to become a "Colon Cancer Free Zone." The trio is collaborating with the American Cancer Society to encourage colon-cancer screening. And now we have the tragic destruction of the Gulf Coast and New Orleans. I'd love to hear what the electronics industry is doing to help victims in those areas. Write me at john.dodge@reedbusiness.com.EDN

WEB LINKS

1. http://news.com.com/Perspective+A+last+hurrah+for+Comdex/2010-1069_3-965975.html.
2. www.demo.com/demo2/.

MORE AT EDN.COM

To listen to Spurious Freedom, go to www.edn.com/051013ed1.

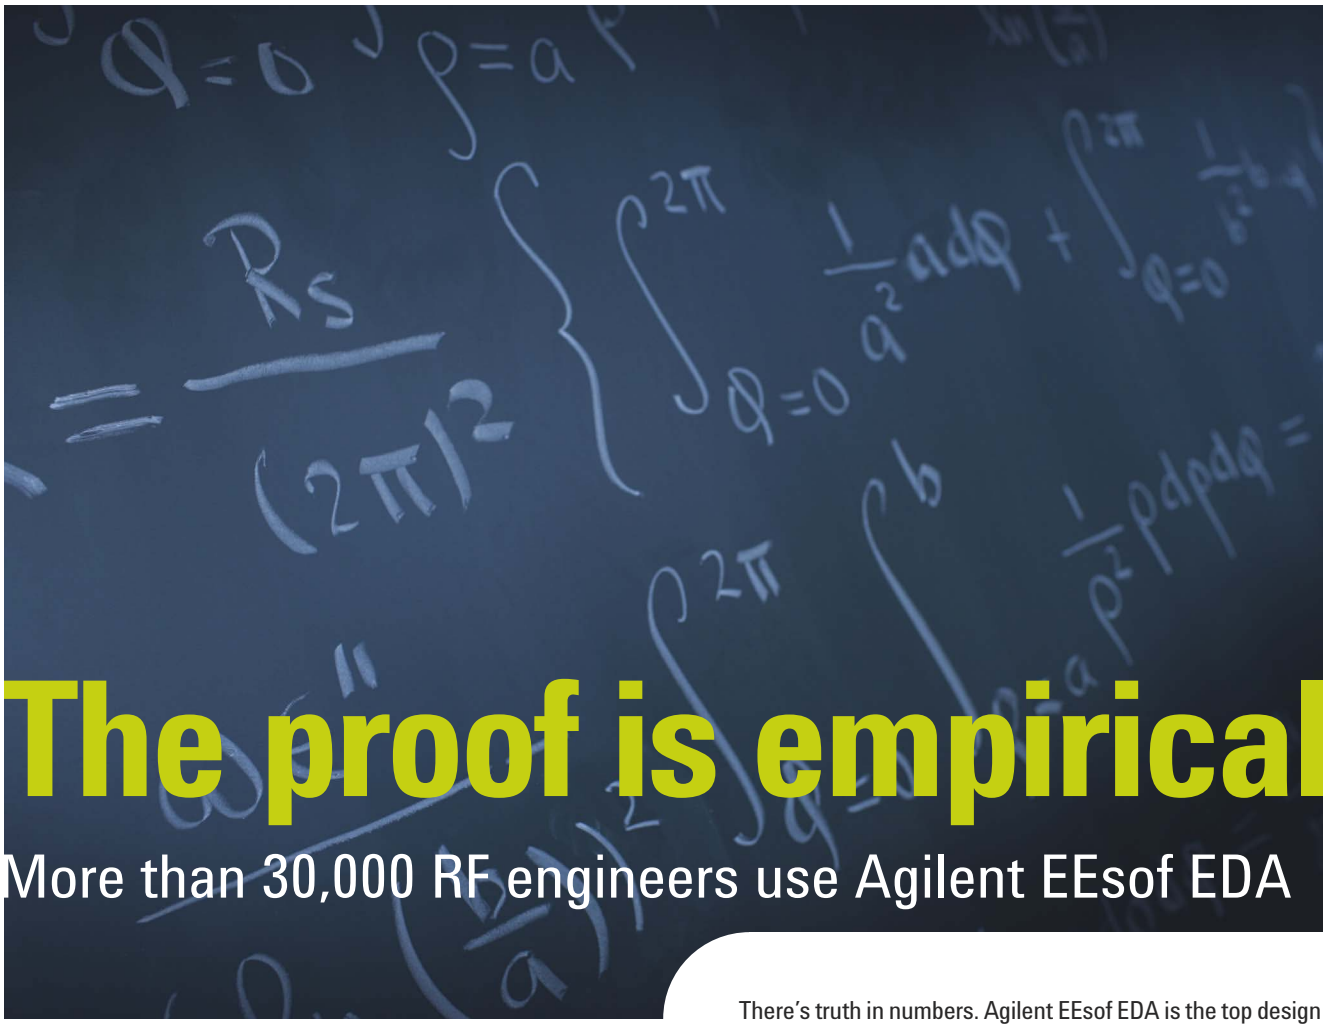

The proof is empirical

More than 30,000 RF engineers use Agilent EEsof EDA

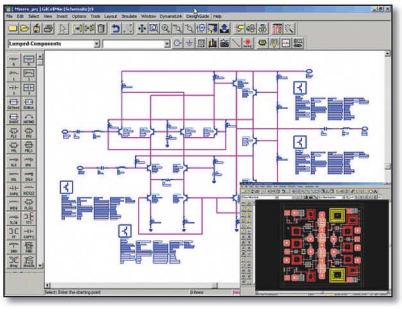

There's truth in numbers. Agilent EEsof EDA is the top design software choice among engineers for the vast the majority of the world's wireless devices.

The fact is derived from software products with undisputed breadth of functionality. Agilent EEsof EDA applications were created for RF engineers by RF engineers. They have the capabilities you need the most. And now with the addition of Eagleware-Elanix you have further ease of use, new synthesis technology, and additional pricing options other products can't touch. Agilent's team of EDA experts releases multiple improvement updates each year. And Agilent has the largest number of technology partners, application examples, and technical articles in the industry.

Work the proof yourself. Agilent EEsof EDA can be customized to fit the precise needs of your projects and budget. Visit www.agilent.com/find/eesof-innovations to find out more.

EAGLEWARE ELANIX

Eagleware-Elanix is now part of Agilent EEsof EDA.

www.agilent.com/find/eesof-innovations

Agilent Technologies

EDN

PRESIDENT, BOSTON DIVISION/ PUBLISHING DIRECTOR, *EDN* WORLDWIDE
Stephen Moylan, 1-781-734-8431; fax: 1-781-290-3431; smoylan@reedbusiness.com

EDITOR IN CHIEF
John Dodge, 1-781-734-8437; fax: 1-781-290-3437; john.dodge@reedbusiness.com

EDITOR AT LARGE
Maury Wright, 1-858-748-6785
mgwright@edn.com

EXECUTIVE EDITOR
Bill Schweber, 1-781-734-8447; fax: 1-781-290-3447; bschweber@edn.com

MANAGING EDITOR
Kasey Clark, 1-781-734-8436; fax: 1-781-290-3436; kase@reedbusiness.com

EXECUTIVE EDITOR, ONLINE
Matthew Miller, 1-781-734-8446; fax: 1-781-290-3446; mdmiller@reedbusiness.com

SENIOR ART DIRECTOR
Mike O'Leary, 1-781-734-8307; fax: 1-781-290-3307; moleary@reedbusiness.com

EMBEDDED SYSTEMS
Warren Webb, Technical Editor
1-858-513-3713; fax: 1-858-486-3646
wwebb@edn.com

ANALOG/COMMUNICATIONS, DISCRETE SEMICONDUCTORS
Joshua Israelsohn, Technical Editor
1-781-734-8441; fax: 1-781-290-3441
jisraelsohn@edn.com

EDA, MEMORY, PROGRAMMABLE LOGIC
Michael Santarini, Senior Editor
1-408-345-4424
michael.santarini@reedbusiness.com

MICROPROCESSORS, DSPs, TOOLS
Robert Cravotta, Technical Editor
1-661-296-5096; fax: 1-781-734-8070
rcravotta@edn.com

MASS STORAGE, MULTIMEDIA, PERIPHERALS, AND PC-CORE LOGIC
Brian Dipert, Senior Technical Editor;
1-916-760-0159; fax: 1-781-734-8070;
bdipert@edn.com

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

DESIGN IDEAS EDITOR
Brad Thompson
edndesignideas@reedbusiness.com

SENIOR ASSOCIATE EDITOR
Frances T Granville, 1-781-734-8439;
fax: 1-781-290-3439;
f.granville@reedbusiness.com

ASSOCIATE EDITOR
Maura Hadro Butler, 1-908-928-1403;
mbutler@reedbusiness.com

WEB/CPS PRODUCTION COORDINATOR
Contact for contributed technical articles
Heather Wiggins, 1-781-734-8448; fax: 1-718-290-3448; hwiggins@reedbusiness.com

EDITORIAL AND ART PRODUCTION
Diane Malone, Manager; 1-781-734-8445
fax: 1-781-290-3445
Steve Mahoney, Production Editor 1-781-734-8442; fax: 1-781-290-3442
Adam Odoardi, Prepress Manager
1-781-734-8325; fax: 1-781-290-3325

NEWS EDITOR
Jeff Berman, 1-781-734-8449; fax: 1-781-290-3449; jeff.berman@reedbusiness.com

CONTRIBUTING TECHNICAL EDITOR
Dan Strassberg, strassbergedn@att.net

COLUMNISTS
Ron Mancini; Howard Johnson, PhD;
Bonnie Baker

PRODUCTION
Dorothy Buchholz, Group Production Director
1-781-734-8329
Kelly Brashears, Production Manager
1-781-734-8328; fax: 1-781-734-8086
Linda Lepordo, Production Manager
1-781-734-8332; fax: 1-781-734-8086
Pam Boord, Advertising Art
1-781-734-8313; fax: 1-781-290-3313

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
Raymond Wong, Managing Director/ Publishing Director
raymond.wong@rbi-asia.com
Kirtimaya Varma, Editor in Chief
kirti.varma@rbi-asia.com

EDN CHINA
William Zhang, Publisher and Editorial Director
wmzhang@idg-rbi.com.cn
John Mu, Executive Editor
johnmu@idg-rbi.com.cn

EDN JAPAN
Katsuya Watanabe, Publisher
k.watanabe@reedbusiness.jp
Kenji Tsuda, Editorial Director and Editor in Chief
tsuda@reedbusiness.jp
Takatsuna Mamoto, Deputy Editor in Chief
t.mamoto@reedbusiness.jp

EDN. 225 Wyman St, Waltham, MA 02451. **www.edn.com**. Phone 1-781-734-8000; fax: 1-781-734-8070. **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

Vol. III, Issue 10

ANALOG edge℠

DESIGN idea: Simplified Power Supply Design ▶

Featured Products

Modular Current Sharing Controller

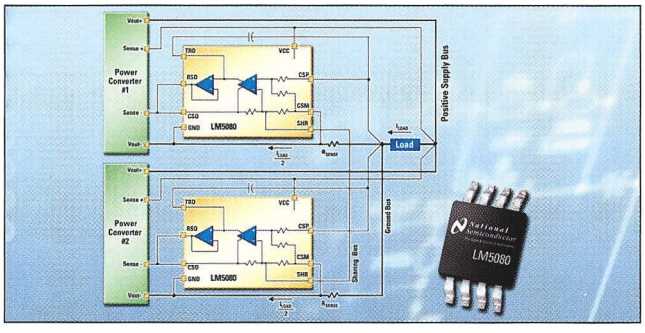

The LM5080 is a simple and cost-effective load-share controller that provides all of the functions required to balance the currents delivered from multiple power converters operated in parallel. The LM5080 implements an average program method of active load share control which adjusts the output voltage of individual power stages either up or down to deliver nearly equal currents to a common load. The average program method improves stability and reduces the output voltage tolerance when compared to other common load sharing methods. The LM5080 supports two common applications for load share controllers: external control in which the load share circuit balances currents between separate power modules (bricks), and internal control where the load share circuit is integrated into the voltage regulation loop of each power converter module or circuit.

Features

- Single-wire star link current share bus
- No precision external resistors necessary
- 3V to 15V bias voltage range
- Adaptable for high or low side current sensing
- Flexible architecture allows 4 modes of operation:
  - Negative remote sense adjustment
  - Positive remote sense adjustment
  - Trim or reference adjustment
  - Feedback divider adjustment

The LM5080 is available in MSOP-8 packaging and is ideal for use in consumer electronics, industrial test equipment, data communications systems, automotive power systems, distributed power systems, and battery-powered applications.

www.national.com/pf/LM/LM5080.html

First 7V to 75V Input, 2.5A Buck Regulator

The LM5005 high-voltage switching regulator features all of the functions necessary to implement an efficient high-voltage buck regulator using a minimum of external components. This easy-to-use regulator includes a 75V N-Channel buck switch with an output current capability of 2.5A. The regulator control method is based upon current-mode control utilizing an emulated current ramp. Current-mode control provides inherent line feed-forward, cycle-by-cycle current limiting, and ease of loop compensation. The use of an emulated control ramp reduces noise sensitivity of the pulse-width modulation circuit, allowing reliable control of very small duty cycles necessary in high input voltage applications. The operating frequency is programmable from 50 kHz to 500 kHz.

Features

- Integrated 75V power MOSFET supports load currents up to 2.5A
- Adjustable output voltage from 1.225V
- Unique, easy-to-use emulated peak current mode control topology enables high frequency operation at V_{IN} up to 75V
- Programmable switching frequency with bi-directional synchronization capability simplifies system design
- Highly integrated, high-speed, full-feature PWM regulator reduces overall solution size

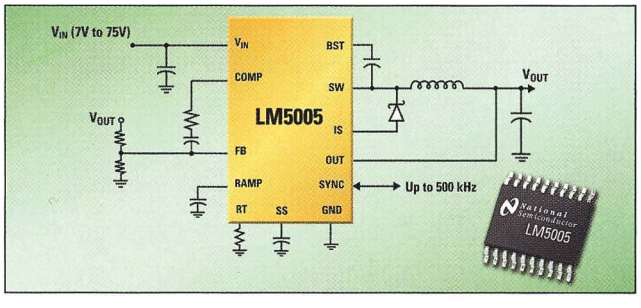

The LM5005 is available in a power enhanced TSSOP-20 package featuring an exposed die attach pad to aid thermal dissipation. It is ideal for use in consumer electronics, telecommunications, data communications systems, automotive power systems, and distributed power applications.

www.national.com/pf/LM/LM5005.html

National Semiconductor
The Sight & Sound of Information

DESIGN idea

L. Haachitaba Mweene, Ph.D., Senior Applications Engineer
National Semiconductor

PWM DC-DC Controllers with Built-In Start-Up Regulators Simplify Switching Power Supply Design

Although PWM DC-DC switching power converters are based on simple topologies, making practical power supplies out of them requires the addition of various functions such as start-up bias, soft-start, switch driving, regulation, short circuit protection, over-voltage protection, over-temperature protection, etc. Today, most of these functions are usually implemented within a compact DC-DC PWM controller integrated circuit.

However, the problem of starting the DC-DC converter in telecom and other high-voltage applications (i.e., where the input voltage exceeds about 15V) is often not addressed. The controller requires a bias supply voltage to run from so that it can produce gate drive pulses and other required signals. But at turn-on the only voltage available is the input voltage, which, if it is greater than 15V, is typically too high to be used as the bias and gate drive supply voltage. It is therefore necessary to lower the input voltage to 15V or below to start-up the power supply. Once the supply is running, the output voltage or a voltage off of a transformer or inductor winding can be used to provide the bias supply for the IC.

But most DC-DC controllers are designed without start-up circuitry, and the power supply designer is expected to add a separate start-up circuit and a bias supply to them (***Figure 1a***).This improves the versatility of the PWM controllers, allowing them to be operated with a wider input voltage range, but the extra start-up circuitry they require increases in the complexity and size of the power supply.

National has solved this problem for the designer in its LM50xx family of 8V to 100V PWM controllers (which includes the LM5020, LM5025, LM5030, and others) by integrating a high-voltage start-up circuit within the IC (***Figure 1b***). This is achieved by fabricating the controller using a 100V process.

The high-input voltage can then be directly applied to the V_{IN} pin of the controller, which is the input to an internal linear voltage regulator. This regulator produces a voltage V_{CC} of about 8V that is used to provide start-up power to the controller. The V_{CC} voltage of the linear regulator is made externally accessible at the V_{CC} pin, for several reasons.

One reason is that the V_{CC} pin is the connection point in the linear regulator for an external output capacitor that keeps the V_{CC} voltage clean.

Another reason is that V_{CC} can serve as the power source for other low-voltage ICs in the circuit such as op amps, logic, and gate drivers.

The V_{CC} pin can also be used to reduce the power dissipation in the controller and to increase the efficiency of the power supply. LM50xx controllers will operate indefinitely off of the input voltage and the internally generated V_{CC}

Figure 1: Two Ways of Implementing a Start-Up Circuit for a PWM Controller

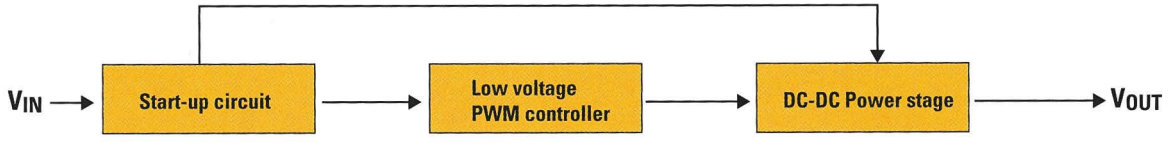

Figure 1a: Power supply built around PWM controller with discrete start-up circuit

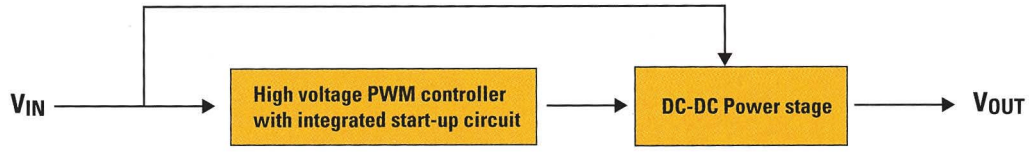

Figure 1b: Power supply built around PWM controller with integrated start-up circuit

voltage (as illustrated in *Figure 2a* in a flyback converter). But this leads to increased power dissipation in the IC. This dissipation is

$$PD=(V_{IN}-V_{CC})I_S$$

where I_S, the supply current of the controller is the sum of the controller quiescent current, and I_G, the frequency-dependent gate drive current. This MOSFET current is given by

$$I_G=Q_{FS}f_S$$

where Q_{GS} is the total gate charge of the MOSFET at a gate voltage of V_{CC} and f_S is the switching frequency.

P_D can be excessively large for the controller at high-input voltages, high-switching frequency, and when the IC is driving a large MOSFET that requires a significant gate-drive current.

LM50xx controllers are designed such that this power dissipation can be circumvented. In all switching power supply topologies, it is easy to derive a bias voltage from a transformer or inductor winding once the power supply has started running. In LM50xx controllers, this voltage (once available), can be applied directly to the V_{CC} pin to provide power for the IC, and can also be used to power other parts of the system. In all LM50xx controllers, if this applied voltage is greater than the 8V output of the internal regulator, the regulator shuts down, eliminating the power dissipation just described. This can lead to an efficiency improvement of 1% or more if the bias supply is properly designed. (*Figure 2b*)

Nevertheless, in lower power systems with lower input voltages it is often advantageous to dispense with the bias supply and to run the supply off of the internal linear regulator. This simplifies the supply and reduces its cost and is often used with LM50xx controllers for output power levels of up to about 30W.

Figure 2: Three Ways of Powering the LM50xx Family of PWM Controllers

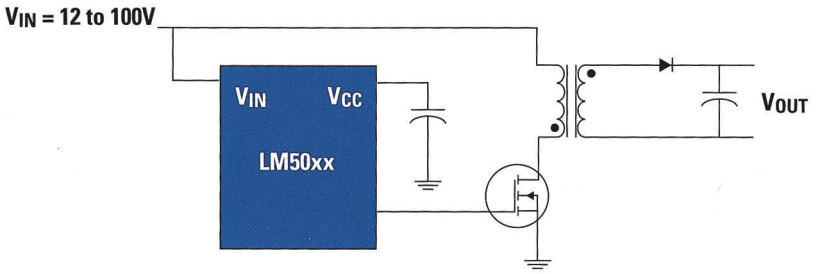

Figure 2a: PWM IC running off a high input voltage in a low power system

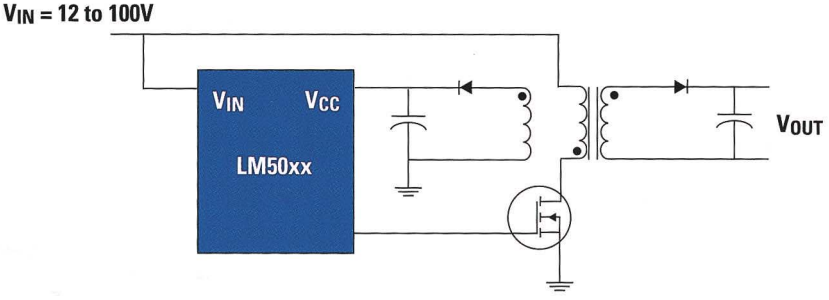

Figure 2b: PWM IC running off a bias supply in a high-power system

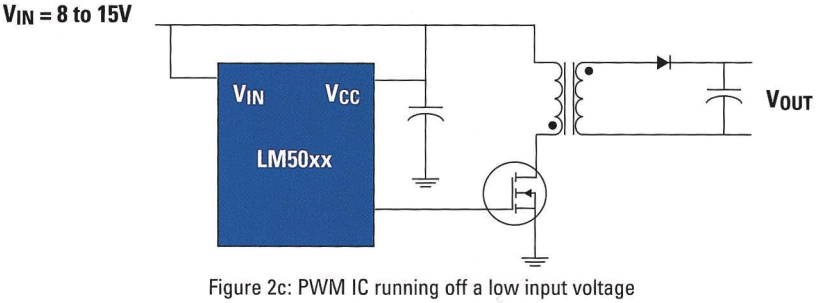

Figure 2c: PWM IC running off a low input voltage

If the input voltage lies between 8 and 15V, the LM50xx controllers are particularly easy to power. The V_{IN} and V_{CC} pins can simply be tied together and to the input voltage, which then directly powers the controller. (*Figure 2c*). If the voltage exceeds 15V the two pins cannot be tied together, and the input voltage has to be higher than about 12V in order for the controller to start up.

In summary, National Semiconductor's high-voltage PWM controllers with integrated start-up regulators allow the power supply designer to reduce circuit complexity, solution size, component costs, design time, and to increase circuit reliability. ■

Featured Products

Current-Mode Controller for Forward Converters with Active-Clamp Reset

The LM5026 PWM controller contains all of the features necessary to implement power converters utilizing the active clamp/reset technique with current-mode control. With the active-clamp technique, higher efficiencies and greater power densities can be realized compared with conventional catch winding or RDC clamp/reset techniques. The device can be configured to control either a P-Channel or N-Channel clamp switch. The main gate driver features a compound configuration, consisting of both MOS and Bipolar devices, providing superior gate drive characteristics. Additional features include line under-voltage lockout, cycle-by-cycle current limit, PWM slope compensation, soft-start, 1 MHz capable oscillator with synchronization input/output capability, precision reference, and thermal shutdown.

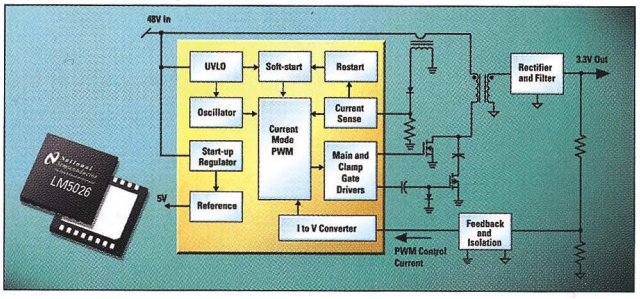

Features

- Wide range (8V to 100V) start-up bias regulator
- Two high-speed power MOSFET drivers: 3A main output driver and 1A clamp driver
- User-programmable maximum duty-cycle and UVLO hysteresis thresholds
- User-programmable gate driver overlap and dead-time
- Versatile dual-mode over-current protection with hiccup mode delay timer

The LM5026 is available in TSSOP-16 or thermally enhanced LLP-16 packaging and is ideal for use in telecommunications power systems, +42V automotive power systems, -48V distributed power systems, industrial power supplies, and multi-output power supplies.

www.national.com/pf/LM/LM5026.html

Industry's First 100V Dual Interleaved Active Clamp Current-Mode Controllers

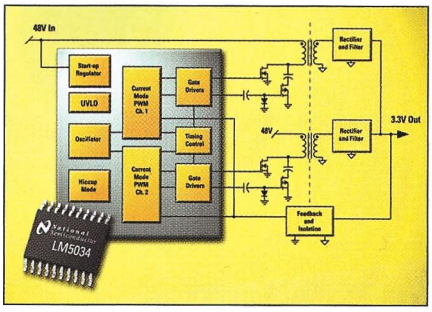

The LM5032 and LM5034 are flexible controllers that can be configured to control either two independently regulated outputs or a single, high-current output from two primary power stages. In the first case, the two PWM channels operate 180 degrees out of phase with one another, or are interleaved, which reduces the input ripple current. In the single-output configuration, the interleaving also reduces ripple current in the output filter capacitor. The LM5032 controller can be used for designing dual-interleaved boost, flyback or standard forward converters. The LM5034 controller is specifically designed for interleaved forward converters with active clamp transformer reset.

Features

- Two independent current-mode controllers
- Interleaved single or dual output operation
- Compound 2.5A main FET gate drivers
- Active clamp FET gate drivers
- Integrated 100V start-up regulator
- Up to 1 MHz switching frequency programmed by a single resistor
- Programmable maximum duty cycle
- Adjustable soft-start and input undervoltage sensing
- Adjustable deadtime between main and active clamp gate drivers

The LM5032/34 are available in TSSOP-16 (LM5032) and TSSOP-20 (LM5034) packaging and are ideal for use in telecom infrastructure, networking, industrial, and automotive power supplies.

www.national.com/pf/LM/LM5032.html
www.national.com/pf/LM/LM5034.html

Adobe® Acrobat® 7.0

In unison, in step. Share critical information more securely. Combine multiple file types into a single, searchable document. View and respond to project feedback seamlessly. And get everyone on the same page. Acrobat 7.0. Try it for free at adobe.com/collaborate and see how much more your team can do. Better by Adobe.™

Adobe

LEADING inside

A series of engineering insights
by Analog Devices.

Specifying A/D Converters: Considerations for IF-Sampling Applications

Choosing the ADC with the highest resolution or sampling speed is often not enough to satisfy the performance demands presented by IF-sampling architectures. ADCs for IF-sampling applications must support high input frequencies while also maintaining adequate SNR, SFDR, and SINAD performance. These features enable designers to eliminate one or more mixing stages and simplify filtering, thereby reducing cost and helping to meet end-system objectives.

Wideband signals having complex modulation—such as those used in many wireless communications, instrumentation, and radar systems—can exhibit time-varying bursts and transients. Furthermore, the data carried by these signals is often spread over multiple channels.

The ADC for these types of architectures must have sufficient input bandwidth to adequately capture and digitize this data. The ADC's dynamic range must also be high enough to detect small signals in the presence of blockers or other large signals in the bandwidth of interest.

Dynamic Range and Noise Requirements

In wideband CDMA systems having a base data rate of 3.84 MHz, data converter clock rates of 16×, 20×, 24×, and 32× are viable. A data converter running at 92.16 MSPS provides good noise performance, and 16-bit ADCs that sample at 100 MSPS are available today. If lower sampling rates are used, the SNR required increases by 1 dB for 76.8 MSPS and 2 dB for 61.44 MSPS.

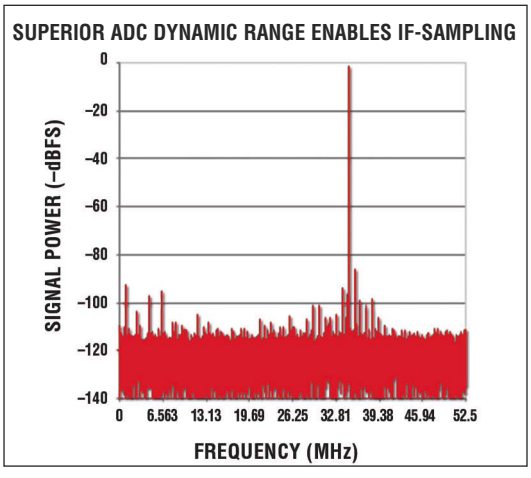

32k point single-tone FFT/ADC: 105 MSPS, 70.3 MHz A_{IN}

The receiver conversion gain and noise figure (NF) sets the ADC's required SNR. At the antenna, the noise spectral density is –174 dBm/Hz, or that of thermal noise. For a conversion gain of 40 dB and a noise figure of 3 dB, the noise spectral density (NSD) at the ADC input will be –131 dBm/Hz (–174 + 40 + 3). If the ADC noise floor is 10 dB below that of the front end noise, it will contribute about 0.1 dB to the overall NF of the receiver. Therefore, a maximum ADC noise floor of –141 dBm/Hz is desirable.

For IF-sampling applications, the total noise of the ADC can be determined by simple integration. For example, a 10 MHz bandwidth signal would have total noise of –71 dBm. This is calculated by adding the effect of the 10 MHz bandwidth [10 log (10 MHz) = 70 dB] to the 1 Hz noise floor of –141 dBm. If the full scale range of the ADC is 4 dBm, the required minimum full-scale SNR for the ADC is then 75 dB.

Selecting the Optimum A/D Converter

What types of ADCs meet the needs of IF-sampling architectures? Typically, they require ADCs with 14 bits to 16 bits of resolution that deliver superior SNR at high input frequencies. Advances in high speed ADC technology offer improved SNR, low additive jitter, higher sampling rates, and increased input frequency capability. These features enable engineers to design more efficient base stations, radar, and measurement equipment.

Additionally, the instrumentation used to validate communications systems must meet even tighter specifications, so as not to mask or distort the end-product's actual performance. These systems allow designers to accurately characterize signals of interest with minimum added distortion from the data converter. The AD9446 16-bit, 100 MSPS ADC from Analog Devices is an example of wideband converter technology that is targeted for IF-sampling applications in communications instrumentation. With a 70 MHz analog input and 100 MSPS sampling rate, the AD9446 provides a spurious-free dynamic range of 83 dB; it provides 82 dB of SFDR with a 100 MHz analog input. For more information on the AD9446 and other data converters for IF-sampling applications, please visit *www.analog.com/PerformanceADCs.*

Author Profile: ***Joanne Mistler*** *is a marketing engineer with Analog Devices' High Speed Converter Group in Wilmington, MA. She has 22 years of RF/MW experience, focusing on low noise synthesizer design, digital communications, and test and measurement applications.*

Embedded modem eases data-communications woes

Targeting applications in remote-monitoring systems, point-of-sale terminals, home-security networks, medical devices, and backup-communication systems, Radicom Research recently announced a low-power, low-cost modem. Standard AT commands control the Half-Inch Modem, which measures just 1×1×0.3 in. with a −40 to +85°C operating temperature. The modem includes a built-in data pump, modem controller, and onboard international DAA (data-access arrangement). With power consumption as low as 18 mA in sleep mode, the modem delivers data rates as high as 56 kbps. Alex Tsau, vice president of operations at Radicom, says, "The Half-Inch Modem simplifies the design process for data-communications functions, allowing designers to add connectivity to their applications in a very small space."

The new Half-Inch Modem from Radicom Research offers drop-in data communications for embedded-system designs.

The self-contained modem requires only a serial TTL interface and phone-line access to provide data, fax, and voice operation. Handset-interrupt and connection-detection features allow the modem to share a phone line with other equipment, eliminating the cost of a dedicated line. The Half-Inch Modem is available now, and prices begin at $21 (1000).

—by Warren Webb

▷**Radicom Research Inc**, www.radi.com.

Get your wall power ac/dc or ac/ac

The ubiquitous wall wart usually provides low-voltage, isolated dc output from the ac mains, but some applications need low-voltage ac, either for circuit operation or because the device it supplies will perform its own local or multivoltage regulation. A new series of Class 2 plug-in units from Foster Transformer provides outputs ranging from 9 to 24V ac, depending on model. The UL-listed transformers are rated for 10 to 50 VA; prices for a basic 10-VA model start at $3.10 (1000).**—by Bill Schweber**

▷**Foster Transformer**, www.foster-transformer.com.

RF connector links boards, uses no cable

Getting dc power or digital signals from one board to an adjacent one is a challenge, but getting RF signals to make the leap is even more difficult. A compression coaxial connector from Tyco Electronics can make the task easier. The Blindmate board-to-board connector has one side that you solder onto the base pc board, and the other side mates by pressing against a tin- or gold-plated target pad on the adjacent target board, using internal spring pressure. Using the one-piece connector is far simpler than using cable-based interfaces or two separate board connectors with a spacer wedged between them.

The connector provides a 50Ω path and operates to 6 GHz with return loss of at least 20 dB to speeds as high as 2.1 GHz. Tyco offers models for 6.65-, 10-, and 14-mm board spacing, and the connector tolerates both axial (interboard) and radial misalignments: ±1 mm for the axial misalignment for the 10- and 14-mm versions and ±0.3 mm for the shortest version. The tolerance for radial misalignment is ±0.8 mm for all three versions. The Blindmate connector sells for $5.80 (10,000).**—by Bill Schweber**

▷**Tyco Electronics**, www.tycoelectronics.com.

Making an RF path between parallel boards is simpler when you eliminate a cable assembly and instead use a spring-loaded compression connector; the Tyco Blindmate operates as fast as 6 GHz and comes in three lengths.

Desktop platform supports voice over WiFi

Azimuth Systems' two-module W-Series desktop WiFi (wireless-fidelity)-test platform enables software designers to test station roaming, performance, function, and range. The system comes in a portable, three-slot chassis and includes a custom-built STM-501 module. According to an Azimuth spokesman, the STM-501 is the first module of its type to provide the multipurpose-testing capabilities that benchtop-WiFi-test systems require. The W-Series software now provides enhanced support for voice over WiFi, including voice-quality measurements to ensure clear and repeatable voice transmissions and a full software-test suite for comprehensive WiFi-phone analysis. "By adding the voice-over-WiFi tests, Azimuth has become the first—and, currently, the only—company to offer fully automated testing of WiFi phones' roaming parameters," says Ray Cronin, the company's chief executive officer.

In addition, the company has added programmable extensions to its testMAC (media-access-control) layer. These extensions permit packet-by-packet traffic generation and analysis, allowing vendors to stress-test their client-to-access-point communication. "In the past, system faults could disable an access point and compromise an entire wireless network," says Cronin. "With this addition, the testMAC can isolate problems, analyze their origins, and allow engineers to correct any anomalies."

The company now also offers a suite of scripts that implement the WiFi Alliance's (www.wi-fi.org) interoperability tests for precertification testing. The battery of tests assures vendors that their products meet the Alliance's standards. Prices for the desktop-engineering system range from $40,000 to $50,000. Software licenses for the voice-over-WiFi- and WiFi-certification enhancements cost $4500 to $9000.

—by Dan Strassberg

▷**Azimuth Systems**, www.azimuthsystems.com.

The W-Series modular desktop WiFi-test system is the first to offer fully automated testing of WiFi phones' roaming parameters.

FEEDBACK LOOP

"I theorize that overly high prices encourage pirating. Send a message to the manufacturers and vendors by not buying their products."

Joseph Travis, in *EDN*'s Feedback Loop at www.edn.com/article/CA633438. Add your comments.

DILBERT *By Scott Adams*

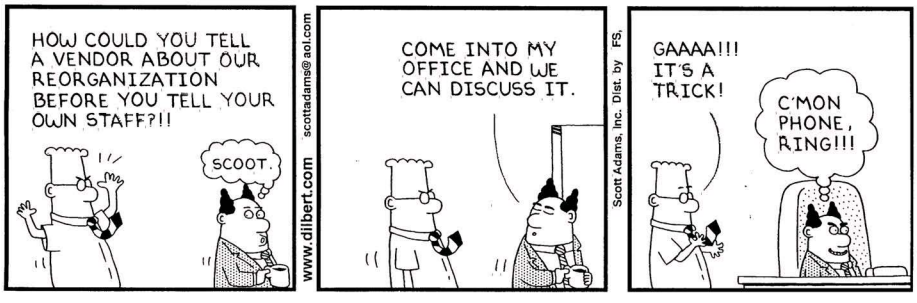

Line terminator says "*hasta la vista*" to reflections

10.13.05

Squeezing an array of Schottky diodes into a single package, the QDN001/002/003 series of diode-protection networks from TT Electronics IRC provides line termination and ESD protection. According to Debasis Roy, PhD, director of IRC's Advanced Film Division Thin Film Business Unit, the devices require no impedance matching, because they terminate load impedances and clamp high- and low-state reflections and noise. The terminators target use with high-speed data lines, especially PCI and SDRAM buses, and reduce overshoot and undershoot with near-zero power dissipation.

The QDN001 includes 36 diodes featuring clamping current of ±50 mA, operating voltage of −0.3 to +7V, and diode forward-voltage drop of 0.5V at 10 mA to 0.8V at 50 mA. The QDN002 series incorporates 17 diodes for ESD protection, with a 15-kV human-body-model rating and 8-kV ESD contact protection; the QDN003 has 18 Schottky diodes and similar ratings.**—by Bill Schweber**

▶**TT Electronics IRC**, www.irctt.com.

The QDN00x series of terminators clamps high-speed reflections, undershoots, and overshoots and protects against ESD.

PC/104 Embedded Machine to Machine Connectivity

Keep in touch with your industrial application by using WinSystems' wired and wireless modular PC/104 modules. Small, rugged, and operational over extended temperature ranges, they are perfect for transportation, pipeline, instrumentation, MIL/COTS, and homeland security applications.

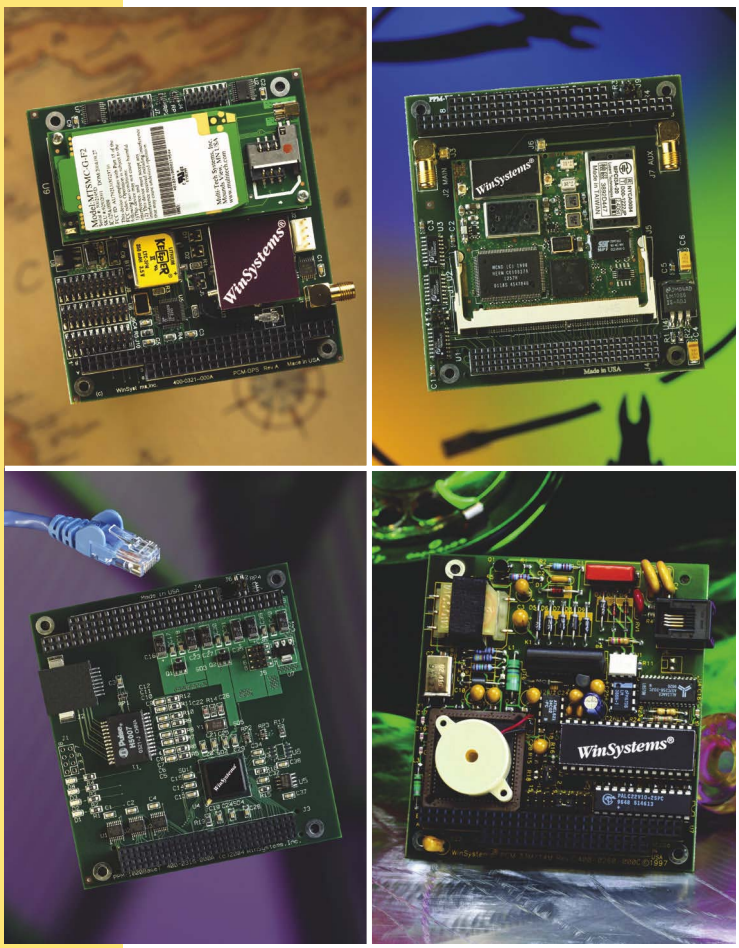

- Wireless
  - GSM/CDMA Cellular
  - 802.11 a/b/g Networking
- Ethernet
  - 10/100 Mbps Ethernet
  - Gigabit Ethernet
- Modem (POTS)
- USB 2.0
- RS-232/422/485 Serial I/O
- Other I/O Expansion
  - A/D, D/A
  - Digital
  - GPS
- Embedded Platforms Supported
  - PC/104 – 3.6 x 3.8 inches
  - EPIC – 4.5 x 6.5 inches
  - STD Bus – 4.5 x 6.5 inches
  - EBX – 5.75 x 8.00 inches
- O/S Support
  - Windows® XPe and CE
  - Linux 2.6
  - x86-compatible RTOS
- Extended Temperature Operation

Off-the-shelf delivery, knowledgeable technical support, and long-term availability makes these products the right choice for your application.

epic™ Microsoft® Windows® Embedded Partner VDC Platinum Vendor Extended Temperature™

Call 817-274-7553 or
Visit www.winsystems.com

Ask about our 30-day product evaluation!

WinSystems®

715 Stadium Drive • Arlington, Texas 76011
Phone 817-274-7553 • FAX 817-548-1358
E-mail: info@winsystems.com

Processor integrates more peripheral support

Atmel's AT91SAM7X128 and AT91SAM7X256 ARM7-based, flash microcontrollers include peripheral support for 10/100 Ethernet; CANs (controller-area networks); full-speed, 12-Mbps USB 2.0 devices; and high-speed AES/3DES (Advanced Encryption Standard/Triple Data Encryption Standard) encryption in a single device. The 50-MIPS microcontrollers offer either 32 or 64 kbytes of SRAM, along with 128 or 256 kbytes of 25-nsec flash memory that supports deterministic memory access. A peripheral DMA controller connects each peripheral directly to on-chip memory, enabling high-throughput data transfers without any processor overhead. Additional peripherals include a 10-bit ADC, an eight-level priority-interrupt controller, an SPI, an SSC (synchronous serial controller), a TWI (two-way interface), UARTs, and supervisory functions. The embedded encryption engine, with the peripheral DMA controller, can encrypt or decrypt data at a rate of 80 Mbps for AES, 32.8 Mbps for DES, and 20 Mbps for 3DES.

Software-development support includes compilers, linkers, and debuggers from IAR Systems (www.iar.com), Green Hills Software (www.ghs.com), and Keil Software (www.keil.com). CMX Systems (www.cmx.com), Micrium (www.micrium.com), and FreeRTOS (www.freertos.com) offer RTOSs for the SAM7X processors. TCP/IP stacks that are available now include open-source uIP/lwIP, Micrium's royalty-free µC/TCP-IP, and CMX Systems' royalty-free MicroNet. The AT91SAM7X128 is available now in a green, 100-lead LQFP for $7.20 (10,000); the AT91SAM7X256 is available in a green, 100-lead LQFP for $8.65 (10,000). The AT91SAM7X-EK evaluation kit is available from Atmel for $250.**—by Robert Cravotta**

▷**Atmel**, www.atmel.com.

EPIC computer pushes temperature extremes

With potential application in medical devices, security equipment, industrial machinery, aerospace projects, and transportation systems, VersaLogic's latest EPIC (embedded-platform-for-industrial-computing)-format single-board computer operates at −40 to +85°C. The company based the Gecko design on the newer AMD GX-500 processor, which offers 500-MHz-equivalent performance and draws only about 1.5W. The low power consumption results in minimal heat dissipation, eliminating the use of an onboard fan.

The module features as much as 512 Mbytes of DDR RAM, integrated video with analog and LVDS flat-panel outputs, a stereo-I/O line, 10/100-Mbps Ethernet, analog and digital I/O, four USB ports, four communications ports, LPT and IDE (integrated-development-environment) interfaces, and a CompactFlash socket. The PC/104-Plus site accommodates both PC/104 and PC/104-Plus modules for system expansion. The Gecko includes safety features, such as transient-voltage-suppression devices for ESD protection, self-resetting fuses for user I/O, and a watchdog timer for hardware-level application control. Prices start at $673 (OEM quantities).

The new Gecko single-board computer in the EPIC form factor fits space-constrained applications with restricted cooling.

—by Warren Webb

▶**VersaLogic Corp**, **www.versalogic.com**.

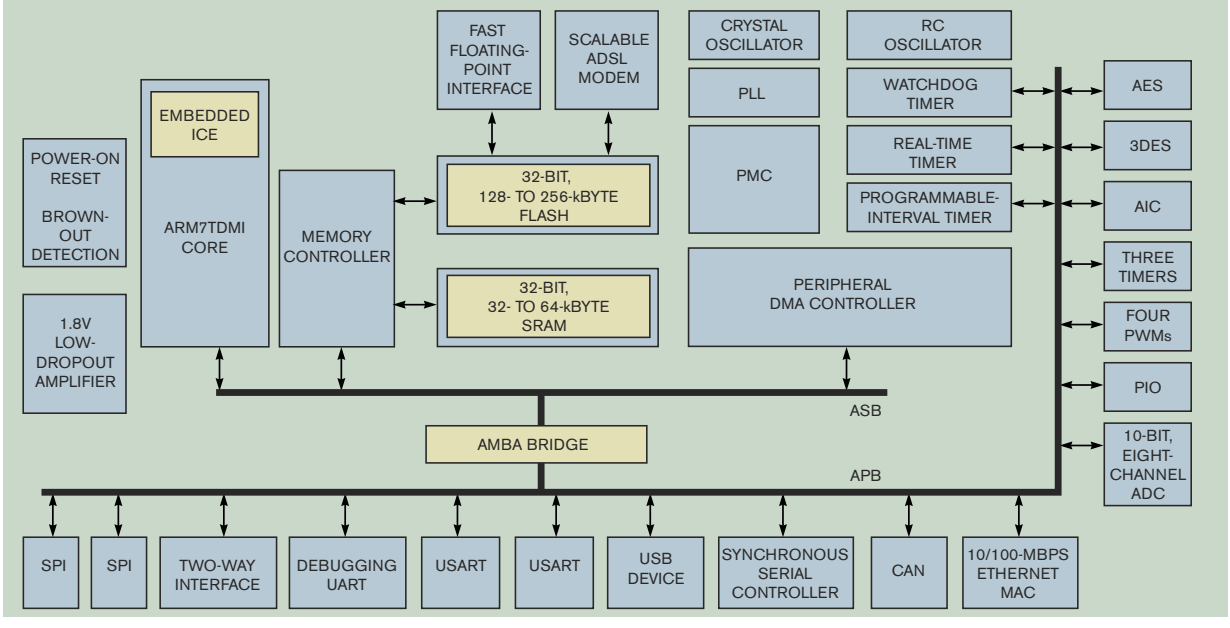

The AT91SAM7X provides an extended peripheral set that supports the ARM7 processor core.

QUESTA GIVES YOU THE POWER OF INTEGRATED VERIFICATION AND PUTS YOU IN CONTROL

Questa™ Verification Platform

You could get carried away with the confidence our verification solution inspires, but try not to let it go to your head. Mentor Graphics® Questa platform is the industry's standards-based, single-kernel verification solution that integrates HDL simulation, assertion-based verification, testbench automation, coverage-driven verification and transaction level-modeling. Questa's unparalleled language support, including SystemVerilog, is critical for the advanced verification of today's most complex designs. But remember, just because you can now tackle a faster way to reach critical verification closure than other methods, it doesn't mean you're invincible. Raise your verification confidence to new levels. Visit our website at **www.mentor.com/products/fv** or call **800.547.3000.**

DESIGN FOR MANUFACTURING + INTEGRATED SYSTEM DESIGN
ELECTRONIC SYSTEM LEVEL DESIGN + FUNCTIONAL VERIFICATION

Mentor Graphics®

THE EDA TECHNOLOGY LEADER

Package provides IDE for heterogeneous distributed systems

According to John Pasquarette, National Instruments' director of software marketing, LabView 8, the latest release of LabView, is the most extensive revision yet. NI developers have been working since 1998 on some of the new features. The company's goal for the product is to provide an IDE (integrated development environment) for heterogeneous distributed systems for design, design verification, test, automation, and control. Such systems can take many forms, some of which incorporate subsystems of multiple types, including conventional PCs; intelligent modular instruments, such as PXI (PCI extensions for instrumentation); distributed I/O, such as NI's Compact FieldPoint and Compact RIO (reconfigurable I/O); PDAs; and custom-designed system components based on FPGAs. Like earlier versions of LabView, this release supports hardware from a broad spectrum of manufacturers, including NI's competitors.

To bring the many disparate elements under the LabView umbrella, the new release focuses on its ability to measure within minutes after installation, thanks to newly enhanced LabView Express features; streamlining application and device management by means of the LabView Project software module; and designing, distributing, and synchronizing intelligent devices and systems. At the heart of these functions are features either that are new to this release or that constitute major improvements over their counterparts in earlier releases.

LabView 8 includes extensive capabilities for locating hardware and software resources that you may want to incorporate into an application. Examples are networked instruments and driver-software modules that reside on networked computers. LabView Project searches the network for these modules, allows you to select them for use in your project, and tracks the availability of resources that you share with another developer or project. Whereas the term "LabView Project" may make you think of the Microsoft Project (www.microsoft.com) project-management tool, the two have little in common. Although it can show you whether another developer is modifying a software module that you wish to use, LabView Project provides no configuration-control functions. It does, however, work with most popular configuration-control packages.

NI offers LabView 8 in a low-cost student version; prices for base, full, and professional versions start at $995. Add-on components include a PDA module, which costs $995, and FPGA, real-time, and data-logging/supervisory-control modules, each of which costs $1995.

—by Dan Strassberg

▷**National Instruments**, www.ni.com/labview, ftp://ftp.ni.com/pub/newsimages/Constellation/LabView_8_%20Distributed_White_Paper.pdf.

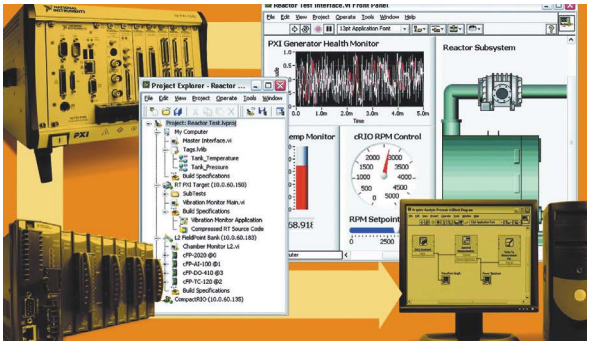

LabView 8 includes LabView Project, whose Explorer window (center left) displays the status of hardware and software resources that your project can use.

FEEDBACK LOOP

"That's what makes Design Ideas so useful; they are starting points for a design or useful concept that can be molded into an individual solution for a particular design."

Steve Hageman, in *EDN*'s Feedback Loop at www.edn.com/article/CA379888. Add your comments.

SMT plasma arrester features low, stable capacitance

A trio of Greentube gas-plasma arresters (similar to gas-discharge tubes) from Littelfuse targets use as surge protection for broadband circuits. The devices have maximum capacitance of just 1.5 pF, along with low insertion loss, and these values remain constant over a range of voltages and temperatures, a critical factor for gigahertz-range applications.

The SL0902 devices, with a 5×5-mm footprint, are available in 90, 230, and 350V versions. The protectors generally work in conjunction with a fuse to provide both surge and fusing protection in applications such as telephone-company and DSL lines, satellite boxes, and CATV interfaces. Their impulse discharge-current rating, per industry-standard test conditions, shows that they can divert a 2.5-kA pulse without destruction.

The 90V version has a dc break-over rating of 72 to 108V at 2 kV/sec with a maximum break-over voltage of 400V at 100V/µsec and 600V at 1 kV/µsec. The 230 and 350V units have dc break-over ratings of 184 to 276V and 400 to 500V, respectively; maximum break-over voltages differ somewhat. All models have an ac discharge current of 2.5A. The SL0902 devices sell for less than 20 cents (OEM quantities).

—by Bill Schweber

▶**Littelfuse Inc**, www.littelfuse.com.

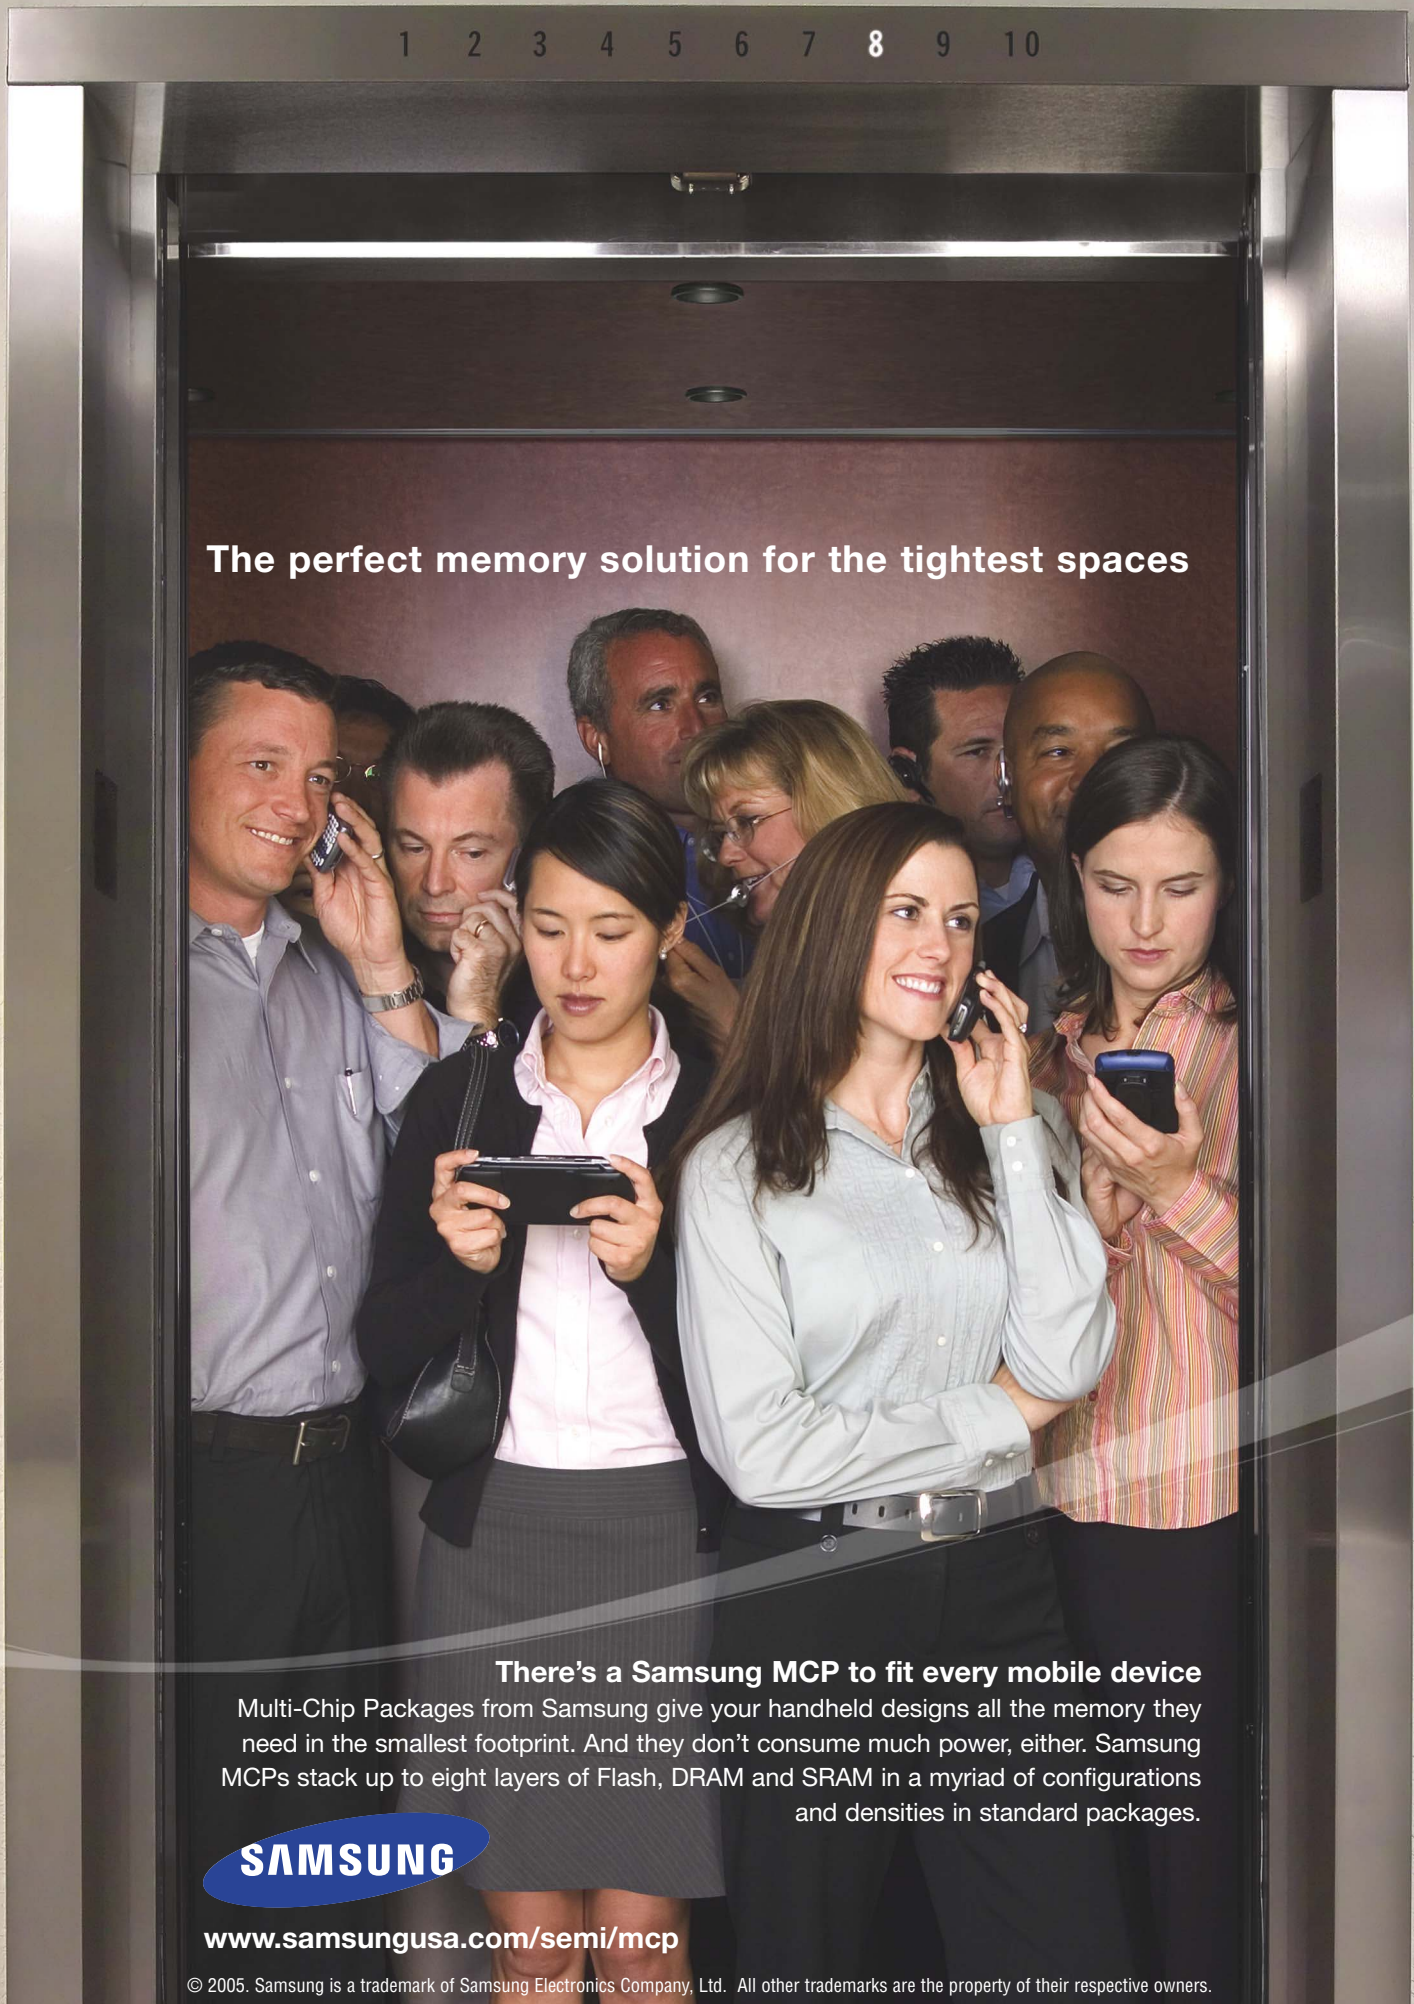
1 2 3 4 5 6 7 8 9 10
The perfect memory solution for the tightest spaces
There's a Samsung MCP to fit every mobile device
Multi-Chip Packages from Samsung give your handheld designs all the memory they need in the smallest footprint. And they don't consume much power, either. Samsung MCPs stack up to eight layers of Flash, DRAM and SRAM in a myriad of configurations and densities in standard packages.
SAMSUNG
www.samsungusa.com/semi/mcp
© 2005. Samsung is a trademark of Samsung Electronics Company, Ltd. All other trademarks are the property of their respective owners.

RESEARCH UPDATE

BY BILL SCHWEBER

Coating may dispel condensation fog

Researchers at the Massachusetts Institute of Technology have developed a coating, made from layers of silica nanoparticles and a special polymer, which may counter droplet accumulation on windows, goggles, and lenses. Normally, condensed droplets scatter light in random patterns, which makes the material translucent and foggy. The coating is superhydrophilic (water-loving) and attracts the water droplets; in turn, the droplets flatten and merge into a uniform, transparent film or sheet on top of the underlying surface.

The photo shows how researchers coated one glass slide (left) with antifogging coating and left the other one uncoated (right). They then placed the glass slides in a freezer, brought them out into humid air, and positioned them over a photo. To learn more, go to http://web.mit.edu/newsoffice/2005/fog.html.

▷**Massachusetts Institute of Technology**, www.mit.edu.

Microstrip line uses layered dielectrics for front-side antenna patch

By using multiple layers of dielectric material on a pc board, engineers at NASA's Jet Propulsion Laboratory have developed a microstrip-line patch antenna that mounts on the front of a pc board and is easy to electrically couple to the electronics. The thin, inexpensive design uses a ground plane on the underside of the high-permittivity pc-board substrate and a low-permittivity-dielectric spacer layer on the top side; the microstrip patch radiator caps all sides.

Designers can use the pc-board area under the ground plane for conventional circuitry. Because there is no feed circuitry behind the ground plane, the patch antenna needs no through hole in the ground plane, minimizing RF leakage. For more, go to the September 2005 issue of *Tech Briefs* at www.techbriefs.com.

▷**National Aeronautics and Space Administration Jet Propulsion Laboratory**, www.jpl.nasa.gov.

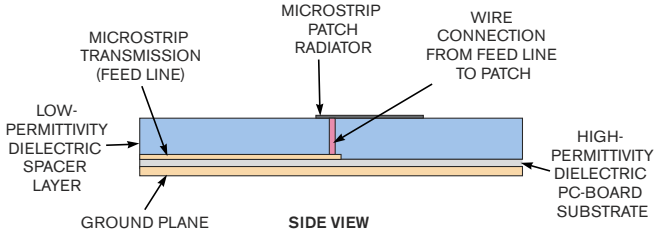

GaN FET produces 174W at 6 GHz

10.13.05

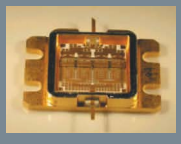

For satellite and cellular base stations, a GaN (gallium-nitride) power FET from Toshiba Corp surpasses the output of GaAs (gallium-arsenide) devices, improving power density by a factor of eight. Currently available GaAs devices are rated at about 90W/6 GHz and 30W/14 GHz, but GaN has higher saturation electron velocity, dielectric breakdown voltage, and operating-temperature range than GaAs. These factors are important for achieving higher power gigahertz-range operation.

The device uses an epitaxial-layer structure, with optimized FET layout and dimensioning, plus a new surface-treatment process for low contact resistance and low gate-leakage current. Toshiba's production also requires a modified version of the conventional stepper process, which is better for mass-production than the electron-beam lithography process that C-band GaN devices commonly use. A GaN power-FET chip measures 2.92×0.71 mm, and the outer cavity of a packaged device with four chips, such as the one in the photo, measures 24.5×17.4 mm. For more, go to www.toshiba.co.jp/about/press/2005_09/pr1201.htm.

▶Toshiba Corp, www.toshiba.co.jp.

Fuel cell beats power-density benchmarks

UltraCell Corp claims that its cells, using a reformed-methanol technology, provide significantly more power density than the conventional "direct-methanol" approach and twice the power density of lithium batteries. Key to the improvement is a microreformer, which the company says generates fuel-cell-ready hydrogen from a highly concentrated methanol solution. The approach requires no water reservoir, cleanup, distinct internal reformer, or water-management system. UltraCell has shipped 32-oz prototypes to the military. Those devices deliver 45W continuous power and 20W average output over 24 hours. For more, got to www.ultracellpower.com.

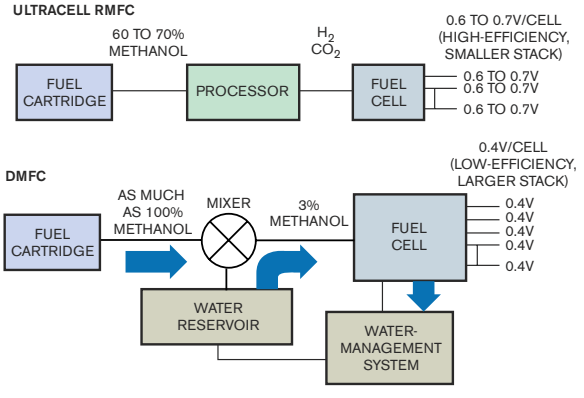

UNLEASH BROADBAND WITH FUJITSU'S WiMAX 802.16-2004 SoC.

With a MAC-to-PHY implementation based on the IEEE 802.16-2004 broadband wireless access standard, the Fujitsu WiMAX SoC, MB87M3400, offers a cost-effective solution for both subscriber station and base station applications. This highly integrated SoC implements MAC, PHY, radio control and all the necessary analog circuits for the appropriate 2 to 11GHz licensed or license-exempt bands. The Fujitsu WiMAX SoC fully complies with the IEEE 802.16-2004 standard using an OFDM PHY.

ADVANCED FEATURES

- 256 OFDM PHY with 64QAM, 16QAM, QPSK and BPSK modulation
- Uplink subchannelization
- Flexible baseband interface with integrated high-performance ADC and DAC
- Security implementation using DES, CCM encryption/decryption
- Rich set of integrated peripheral and RF control

For more information, call **(800) 866-8608** or visit **http://us.fujitsu.com/micro/wimax**

FUJITSU

THE POSSIBILITIES ARE INFINITE

GLOBAL DESIGNER

Low-cost handsets gain traction in India

Burgeoning demand from mobile-service operators for ultralow-cost handsets to service price-sensitive emerging markets, such as India, is motivating semiconductor vendors to develop single-chip-handset offerings. For example, Texas Instruments recently demonstrated the first less-than-$40 handsets built in India from concept to design to production. BPL Ltd (www.bplmobile.com) and Quasar Innovations (www.quasarinnovations.com) based the Primus GSM phones that they developed on TI's TCS chip set. These handsets incorporate only basic voice and short-message service—no cameras, color screens, or MP3 players, instead emphasizing low cost and an intuitive user interface.

"We expect the industry will be selling an ultralow-cost mobile phone for less than $20 by the start of 2006. These costs include electronic components, connectors, the pc board, the casing with display and keypad, software, and the battery," says Horst Patch, vice president of the Communication Business Group at Infineon Technologies.

Ultralow-cost handsets will enable telephony in parts of the world in which land-line service is unavailable or unreliable, says Douglas Grant, business-development manager with Analog Devices' RF and Wireless Systems Group. Grant adds that chip-set suppliers pursuing the ultralow-cost-handset market need to provide robust reference platforms and even complete turnkey designs with preloaded software features to simplify the development of the final handset product.

"In phone design, there are a few major building blocks—among them: digital baseband, analog baseband, RF section, and memory, which comprise the bulk of the cost drivers for these handsets," explains Harish M, general manager of business development at TI India. An ultralow-cost handset should have no more than 100 components, compared with 600 or more for a high-end feature phone or smart phone.

Grant says that manufacturers can reduce BOM (bill-of-materials) costs by integrating core handset functions, such as digital and analog processing, power management, and a radio transceiver, into one chip. But this approach is risky, because it requires a high-performance wafer process, and may be unsuited for consistent production.

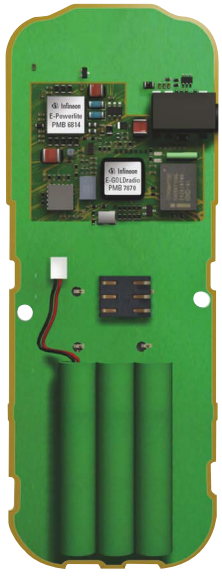

Infineon's new, low-cost E-Goldradio handset combines the baseband processor and the RF transceiver in a 9×9-mm-footprint chip.

"A lower risk approach is to combine functions into a system-in-package design, which saves board area and manufacturing cost but retains the performance of the radio and minimizes risk because vendors can easily port software to such a device," says Grant.

Infineon's new, low-cost E-Goldradio handset combines the baseband processor and the RF transceiver in a 9×9-mm-footprint chip. It enables manufacturers to implement the baseband and RF functions on less than 4 cm^2 of board space—about 30% less than two-chip offerings require. It also cuts the BOM cost by about 30% by eliminating external components, including capacitors and discrete components, that a two-chip approach requires. And a low-component-count design and simplified layout enable the use of cheap, four-layer pc-board technology and single-side mounting, further reducing manufacturing costs. TI also offers the similar single-chip DRP (digital-radio processor).

—by Chitra Giridhar, *EDN Asia*

▷**Analog Devices**, www.analog.com.
▷**Infineon**, www.infineon.com.
▷**Texas Instruments**, www.ti.com.

Japanese-automotive software goes horizontal

Toyota Motors is moving to a horizontal and cooperative model for software development to meet the increasing demand for more sophisticated and widespread automotive electronics. Traditionally, Japanese designers in the automotive industry have developed software for single applications instead of using standardized software platforms upon which vendors could build myriad applications, according to Koichi Tanigawa, general manager of the development department at the integrated-system-engineering division of the vehicle-engineering group of Toyota. He recently spoke at the Freescale (www.freescale.com) Technology Forum Japan 2005 in Tokyo.

With the increased use of microcontrollers and sensors in such automotive applications as vehicle control, safety, and pollution control, automotive electronics have become hot markets. Fulfilling the needs of those applications requires a lot more software development, he says.

In the Japanese-automotive industry, suppliers and automotive manufacturers have a vertical-integration relationship, in which conventional and limited suppliers deliver components to automotive manufacturers. This approach makes suppliers a type of subsidiary of the automakers in vertical-integration systems. For example, Toyota has developed its own dedicated software. But, as the amount of the software increases, it is using multiple developers to develop individual layers, including the operating system, middleware, and applications.**—by Takatsuna Mamoto, *EDN Japan***

▶**Toyota Motor Corp**, www.toyota.com/jp.

10.13.05

Intersil Battery Charger ICs

Intersil High Performance Analog

Have **YOU** Seen the World's Smallest Battery Charger IC?

Not only is the 2mm x 3mm ISL6294 the industry's smallest, but this fully integrated, single-cell Li-Ion / Li-Polymer battery charger IC can handle input voltages up to 28V, eliminating the need for an over-voltage protection circuit.

The constant current I_{REF} is set with the external resistor R_{IREF}. The constant voltage is fixed at 4.2V.

End-of-charge (EOC) current indicated through the CHG pin (which can be interfaced to a micro processor), but the charger continues to output the 4.2V.

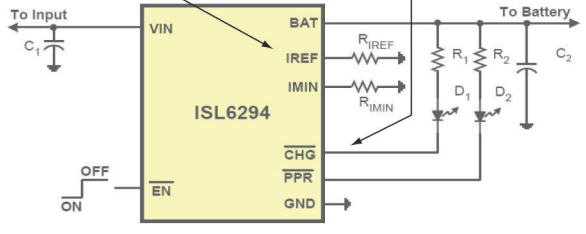

If the battery voltage is below 2.6V the ISL6294 charges the battery with a trickle current of one-tenth of I_{REF}. When the battery voltage reaches 4.2V, the charger enters a CV mode and regulates to fully charge battery without the risk of over charge.

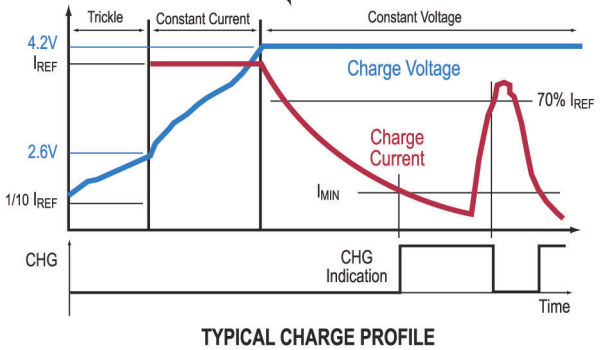

TYPICAL CHARGE PROFILE

Key Features:

- 2mm x 3mm 8 Ld DFN package
- 28V maximum input voltage
- Programmable end-of-charge current with status interfaced to a micro device through CHG pin
- Thermaguard™ charge current thermal foldback for thermal protection
- No external blocking diode required
- Integrated pass element and current sensor
- 1% voltage accuracy
- Trickle charge for fully discharged batteries
- Less than 1µA leakage current off the battery when no input power attached or charger disabled
- Input over-voltage protection
- End-of-charge indication with large hysteresis to prevent unwanted re-charge

Datasheet, free samples, and more information available at **www.intersil.com/edn**

Intersil – Switching Regulators for precise power delivery.

intersil®
HIGH PERFORMANCE ANALOG

Power Management for FPGAs

High Performance Analog Solutions

Power management and signal chains demand a system approach – that's why Avnet Electronics Marketing and National Semiconductor together are such a powerful combination.

- Power and signal chain products that blend maximum performance with overall efficiency
- Collaborative engineering support from conceptual design through volume manufacturing
- System level expertise to assist in determining the right architecture and products for your application
- Factory-trained Field Application Engineers well versed in National's technology and tools
- Broad portfolio of devices to compliment FPGA-based designs
- Solutions sample kits to assist you in your designs

Whether you're working on a portable, handheld, security, medical or other analog design, let Avnet EM help you solve your power and signal design challenges and get your application up and running.

For sample kits, web seminars and design tools, go to **avnet.national.com**

National Semiconductor
The Sight & Sound of Information

AVNET electronics marketing

Enabling success from the center of technology™

Avnet Green Initiative

Supplier Authorized Distributor

1 800 332 8638

FPGAs in High Performance Systems

FPGAs enable developers to accelerate high performance applications by integrating I/O acquisition and processing functions into programmable logic. While system power is primarily determined by how logic is utilized and how fast it is clocked, the selection of power components has a substantial impact on the efficiency of power supply systems and, as a consequence, on overall power consumption.

To read this case study, go to **avnet.national.com**

POWER MANAGEMENT DESIGN GUIDE FOR XILINX® FPGAS

The National Power Management Design Guide features parametric tables, sample designs and step-by-step directions, making it a fast, accurate source for choosing the best National power supply solutions for your Xilinx FPGA-based applications.

The recommended power supply solutions feature **voltage regulators**, **voltage supervisors** and **voltage references**.

The guide also details available design tools, including the National Dual High-Current Power Supply Kit developed by Avnet (part # ADS-NSC-XP).

Learn more at avnet.national.com

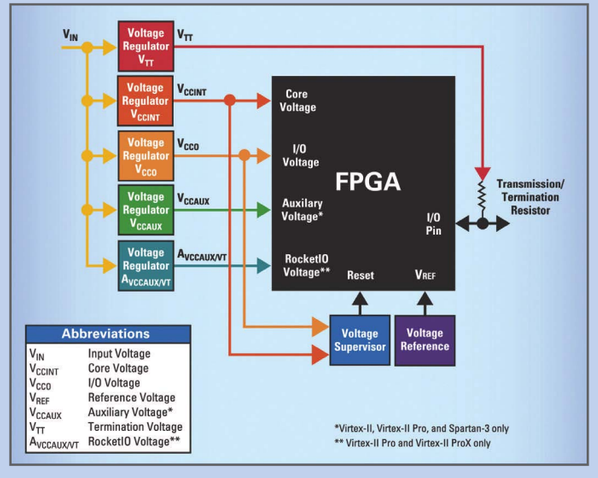

Reigning in Power Consumption for FPGA-Based Applications

Through an understanding of the interdependencies of power systems, developers can include power considerations in the early stages of design that will simplify power supply design and maximize the power efficiency of the final architecture. Industry expert Nicholas Cravotta explores the key issues behind managing power in FPGA-based high performance systems, including:

- How to efficiently and cost-effectively supply all of the different voltages an FPGA requires
- Designing flexible power supply systems that can accommodate the power differences between FPGA devices – even those within a single product family – in order to enable efficient scaling of device architectures without a power subsystem redesign
- How to increase power utilization through power sequencing, limiting the rate of current change, and synchronization mechanisms
- Understanding the impact of temperature on power utilization and how to manage and reduce power dissipation to result in overall higher system reliability

For additional application solutions, visit avnet.national.com

BY HOWARD JOHNSON, PhD

See beyond the edge

To measure the characteristic impedance of a pc-board trace, most engineers use a TDR (time-domain reflectometer). Connect the instrument to a long, unterminated trace, and it blasts into the trace one very quick, precise rising edge. By analyzing the signal that reflects back from the trace, you can deduce the trace impedance. Some instruments provide averaging capabilities to help reduce the noise floor during particularly fast measurements.

In the usual setup, a TDR instrument uses only the first few nanoseconds of your reflected signal. After the initial step edge propagates to the far end of the unterminated trace, it bounces and returns to the instrument, corrupting further readings of characteristic impedance. **Figure 1** illustrates the typical result. The top curve (TDR, in red) illustrates the signal you typically observe at the front end of a pc-board trace. This plot displays the initial TDR step (first edge) and the signal that reflects from the far end of the trace (second edge).

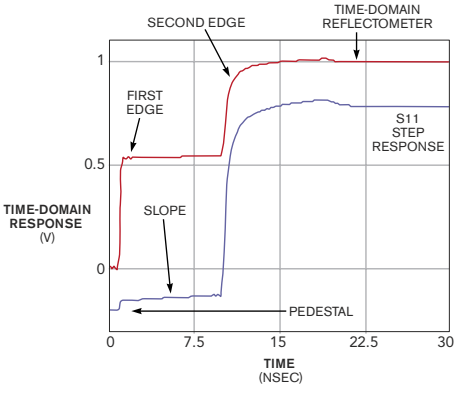

Figure 1 Engineers normally consider data they receive after the second edge as unusable.

You calculate the lower plot (blue, offset below for visual clarity) from the TDR plot. It is the step response of the S-parameter function, S11. In the frequency domain, S11(f) equals $(2\times(j2\pi f)\times\mathrm{TDR}(f)-1)$. The S11 step response shows only the signal reflected from the trace, in the absence of the outgoing signal. (Subtracting unity accomplishes this reduction.)

The S11 step response displays a pedestal, from whose amplitude you may deduce the effective trace impedance over a scale of time of 1 nsec or so. After the pedestal, this S11 step response also displays a gentle upward tilt. This tilt is the hallmark of a trace marred by significant amounts of skin-effect loss.

Engineers usually consider the second edge the end of usable data in a TDR waveform. Even though the latter stages of the waveform contain a wealth of information about trace loss and impedance, these details are hidden from view—unless you learn to see beyond that second edge.

Here is the catch: You must make two measurements, not one. Make the first measurement as usual, with the trace open-circuited at the far end. Make the second measurement with the trace shorted to ground at the far end. Now, convert both measurements to the frequency domain using an FFT.

If you are unfamiliar with FFT calculations, look up the "FFT windowing functions" in your instrument's help screens. The time-domain window exists to chop off the truly unusable negative edge of your pulse generator in a graceful way without inducing other undesirable side effects, such as wiggles in the frequency-domain output. It sometimes helps to differentiate the waveform before windowing and then patch it up later in frequency-domain form.

Now, convert each of your TDR results to S11 form using $S11(f)=(2\times(j2\pi f)\times\mathrm{TDR}(f)-1)$.

From the two S11 functions now in your possession and from knowledge of the source impedance, $Z_S(f)$, of your TDR tester (usually 50Ω), you may now calculate the characteristic impedance, $Z_C(f)$, of your pc-board trace. This calculation works at extended frequencies corresponding to the full length of the waveform you capture and is not limited by the reflection time of your TDR-test coupon:

$$Z_C(f)=Z_S(f)\sqrt{\left(\frac{1+S11_{OPEN}}{1-S11_{OPEN}}\right)\left(\frac{1+S11_{SHORT}}{1-S11_{SHORT}}\right)}.$$

This clever frequency-domain technique derives from procedures that engineers commonly use to calibrate the SMA cables you use with a network analyzer. EDN

MORE AT EDN.COM

Go to www.edn.com/051013hj and click on "Feedback Loop" to post a comment on this column.

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

Introducing the new Spartan-3E family — the world's lowest-cost FPGAs

Priced to go.

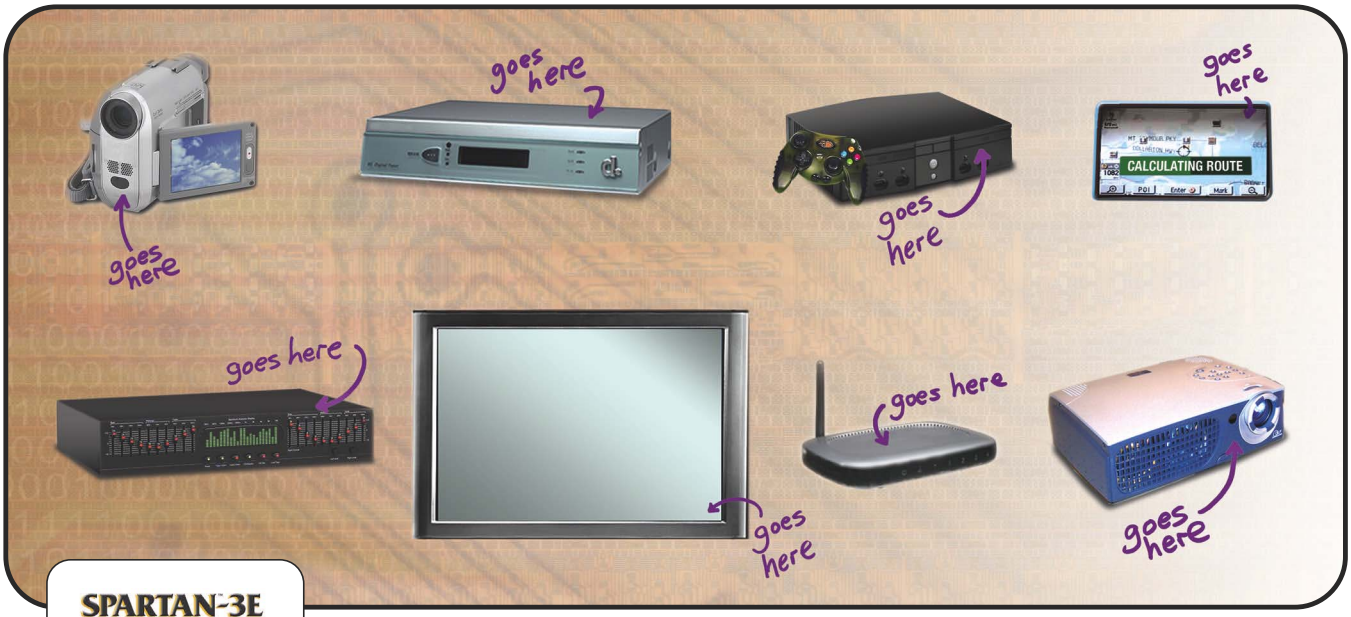

SPARTAN-3E FPGA

The industry's first 100,000-gate FPGA for only $2.00*

Spartan-3E Platform FPGAs offer an amazing feature set for just $2.00! You get 100K gates, embedded multipliers for high-performance/low-cost DSP, plenty of RAM, digital clock managers, and all the I/O support you need. All this in a production-proven 90nm FPGA with a density range up to 1.6 million gates.

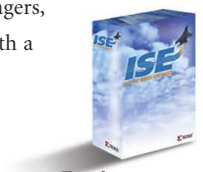

Easiest to use Software

Perfect for digital consumer apps and much more!

With the Spartan-3E series, we've reduced the previous unit cost benchmark by over 30%. Optimized for gate-centric designs, and offering the lowest cost per logic cell in the industry, Spartan-3E FPGAs make it easy to replace your ASIC with a more flexible, faster-to-market solution. Compare the value for yourself . . . and get going on your latest design!

Over 5 million Spartan-3 devices shipped!

MAKE IT YOUR ASIC

XILINX®
The Programmable Logic Company℠

For more information visit
www.xilinx.com/spartan3e

* Pricing for 500K units, second half of 2006

Earth Friendly

Pb-free devices available now

Coping with changes of employment

A few weeks ago, Boeing completed the sale of the Rocketdyne operation, where I work, to Pratt & Whitney. P&W (part of UTC) is the ninth company I've worked for in the last (almost) 40 years. I carefully picked my first job in 1966, interviewing with and comparing dozens of companies. I chose RCA—mainly because it had so many knowledgeable engineers, an opinion that only got stronger during the 15 years I worked there. The eight employment changes I've experienced include one resignation and two incidents in which I managed to lay myself off. Both were plant closures. I was sold, along with the furniture and facilities, on the other five occasions.

It will take a while before the dust settles and we really understand our situations with our new ownership. But, based on past experience, I expect to benefit from the change. Each of the previous changes has exposed me to product types and design challenges I'd never before encountered, resulting in a relatively wide range of engineering experience that has worked to my benefit. It keeps life interesting. But you need to have the right attitude to go through these changes without damage.

The first time a company leaves you, or asks you to leave, you might find that your feelings are hurt. After all, you carefully chose this company, identified with it, and intended to stay indefinitely. In the '50s and '60s, it was not at all unusual to encounter engineers who were about to retire after working 45 years for the same company. That situation is becoming unusual, at least in the private sector. Acquisitions, mergers, bankruptcy, plant closings, and relocations are just some of the reasons that you might find yourself working for a new company.

First impressions of new ownership are frequently negative. The benefits, policies, procedures, and organizations might differ. Some of your favorite managers might be pushed aside—or out. You can always pick up your stuff and leave, of course, but that choice might not be the best. And sulking won't help. I've always given the new situation a year to sort itself out, and, with one exception, I've been satisfied with the new company. You should enjoy your work, as much as is possible. If a company makes every task an onerous ordeal, it's time to move on. There's nothing wrong with being skeptical of new ownership, but you should try to avoid being cynical. If you can be patient for a year, keep an open mind, and be flexible enough to accommodate some changes, things are likely to work out.

When you separate from a company through a layoff or resignation, try to stay in touch with the best and brightest that you worked with there. Similarly, if some of them leave and you stay, get a home e-mail address or phone number of those you hold in high esteem and follow up every six months or so with a call or note. Nowadays, they call it networking. I always thought it was just common sense. Maybe I was just lucky, but I've never missed a day's pay or collected unemployment, largely because I knew key people in other companies who would gladly recommend me to their management.

If you're an engineer in the private sector, it's likely that you have experienced or will experience a change of employment. Try not to stress out about it or make decisions without sufficient data. Treat it like a design assignment: Examine your requirements and the new environment, look at your options, do the trades, and make an informed decision on how to proceed. That's what I've always done—even going to the trouble of building a spreadsheet on occasion. Give it a try—and take your time. **EDN**

Charles Clark is a Technical Fellow with The Boeing Company and is a member of EDN*'s Editorial Advisory Board. Like Clark, you can share your Tale from the Cube. Contact mgwright@edn.com.*

ILLUSTRATION BY DANIEL VASCONCELLOS

CAN YOUR BRICK DO THIS?

IR's DC Bus Chipset Enables Bricks With $48V_{in}$, $9.6V_{out}$, 330W at 97% Efficiency

Full-Bridge Bus Converter Chipset (IRF6646, IR2086S, IRF6635)
$48V_{in}$ / $9.6V_{out}$ / 175kHz

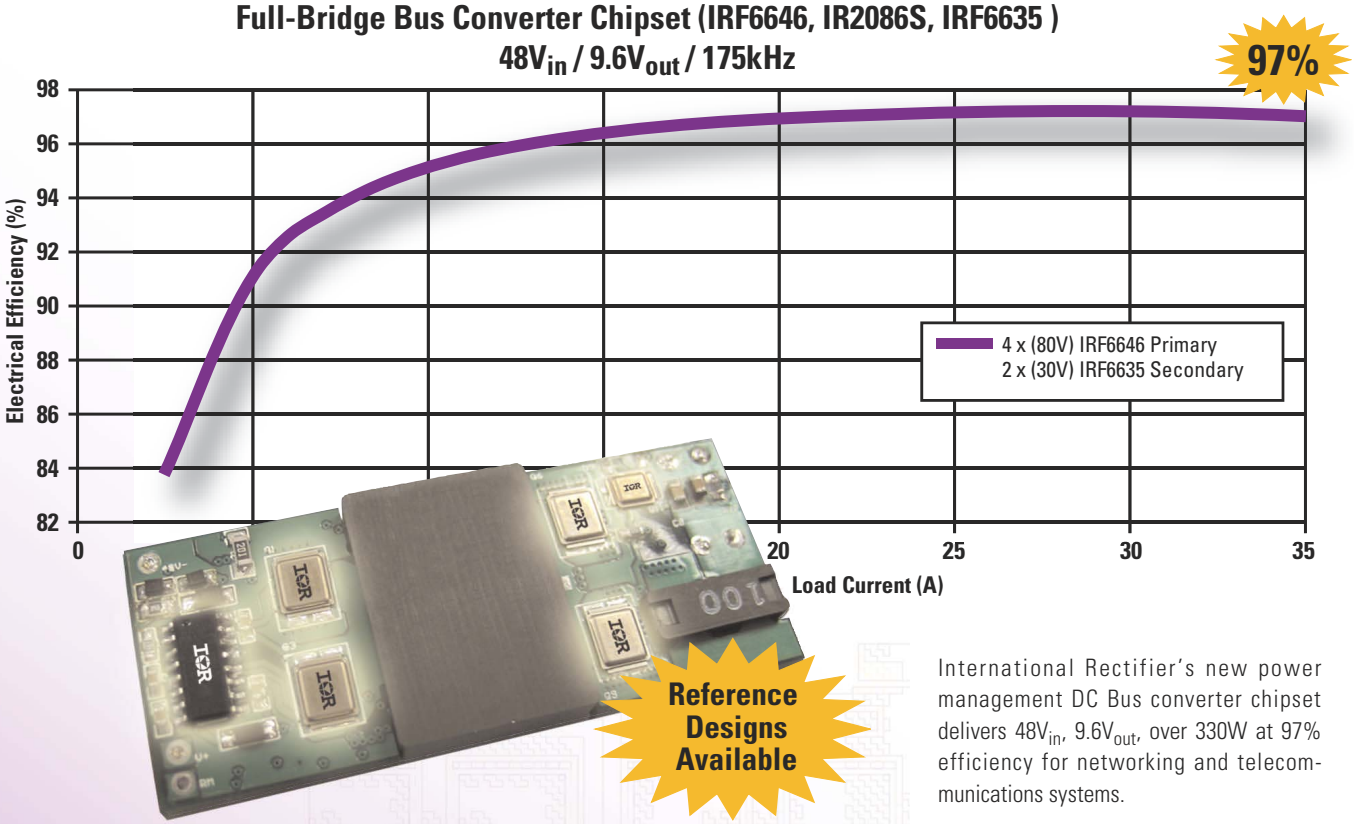

Reference Designs Available

New DC-DC Chipset Solution from International Rectifier

DirectFET™ MOSFETs					
Part #	Package	V_{DSS}	$R_{DS(on)}$ max @ V_{GS}=10V	Q_G Typical	Q_{GD} Typical
IRF6644	Medium can	100V	13mΩ	35nC	11.5nC
IRF6655	Small can	100V	62mΩ	8.7nC	2.8nC
★ IRF6646	Medium can	80V	9.5mΩ	36nC	12nC
IRF6613	Medium can	40V	3.4mΩ	42nC	12.6nC
IRF6614	Small can	40V	8.3mΩ	19nC	6.0nC
★ IRF6635	Medium can	30V	1.8mΩ	47nC	17nC

Control IC			
Part #	Package	Voltage Rating	Description
IR2085S	SO-8	100V	Primary-side half-bridge control IC, fixed 50% duty cycle, self-oscillating
★ IR2086S	SO-16	100V	Primary-side full-bridge control IC, fixed 50% duty cycle, self-oscillating

★ – IR product featured in Full-Bridge Bus Converter reference design above

International Rectifier's new power management DC Bus converter chipset delivers $48V_{in}$, $9.6V_{out}$, over 330W at 97% efficiency for networking and telecommunications systems.

- Single DirectFET MOSFET can replace two SO-8 devices
- Reduces component count by >45%
- Achieves up to 1.5% better efficiency compared to industry standard quarter brick form factors
- Reduces temperatures up to 40°C
- Board space reduced by 29% vs. quarter brick form factors
- Industry Best $R_{DS(on)}$ for IRF6644 (13mΩ), IRF6646 (9.5mΩ), IRF6635 (1.8mΩ)

Design Services available at www.mypower.com

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

International
IOR Rectifier
THE POWER MANAGEMENT LEADER

POWER, PERFORMANCE, VALUE

Precision ICs

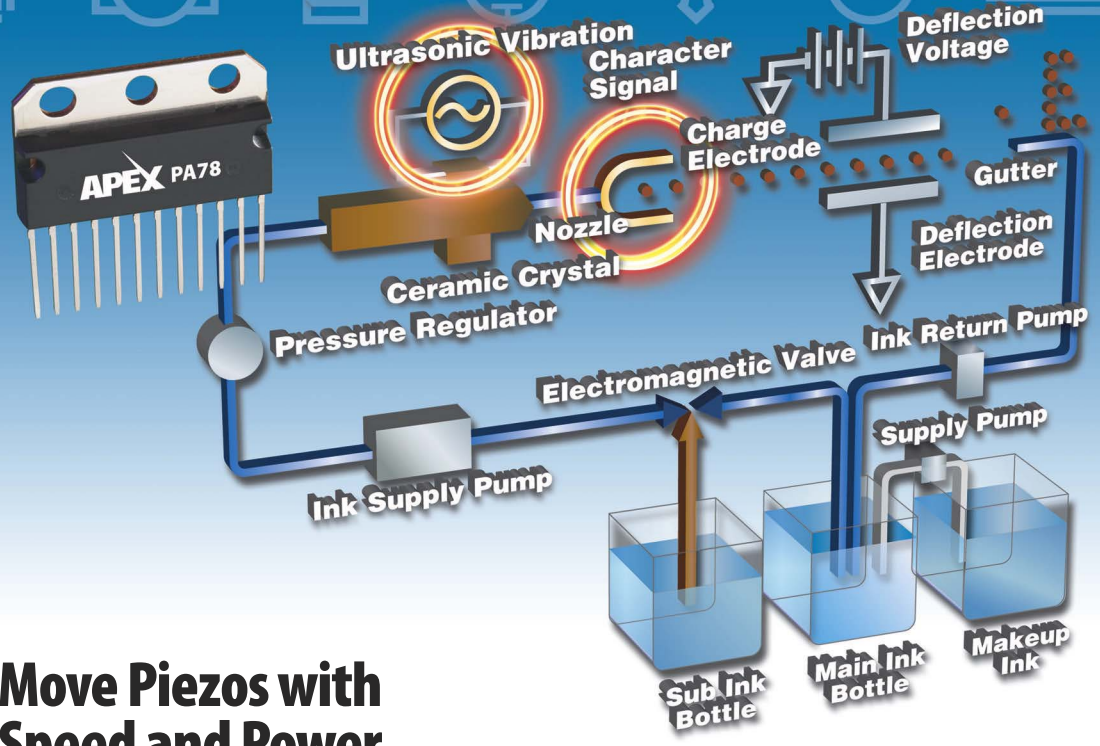

Move Piezos with Speed and Power.

350V/µs, 350V Power Amplifier IC Delivers on Price too!

Apex's new Precision power amplifier IC offers a best in class combination of voltage supply and slew rate while running at a cool standby current less than 1mA

The new Apex PA78 is the flagship model in a new series of Precision ICs that combine new levels of power amplifier performance for speed, voltage supply, quiescent current and price. For applications that demand high voltage drive power that's equally responsive, the PA69, PA86 and PA78 offer you a multitude of performance and per unit cost combinations to choose from.

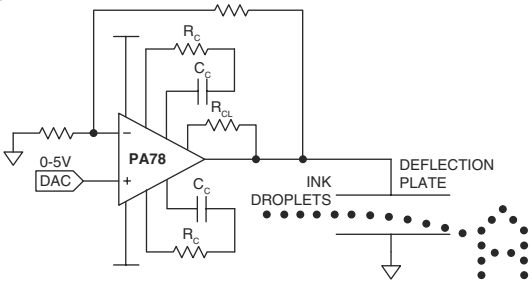

PA78 Power Amplifier for Piezo Actuation and Deflection

Model	Output Voltage V	Output Current mA	Slew Rate V/µs	Power Bandwidth kHz	Production Volume Pricing 10k pcs USD
PA69	±100	50	200	200	$10.50
PA69A	±100	75	250	200	$13.60
PA86	±100	100	350	300	$18.15
PA86A	±125	150	350	300	$21.60
PA78	±175	150	350	200	$24.85

Visit us online today at www.apexmicrotech.com/EDN or call 1-800-862-1015.

- Data sheets
- Design tools
- Application notes
- Spice models
- Eval kits
- Request v12 data book

POWER, PERFORMANCE, VALUE

Precision ICs

APEX MICROTECHNOLOGY

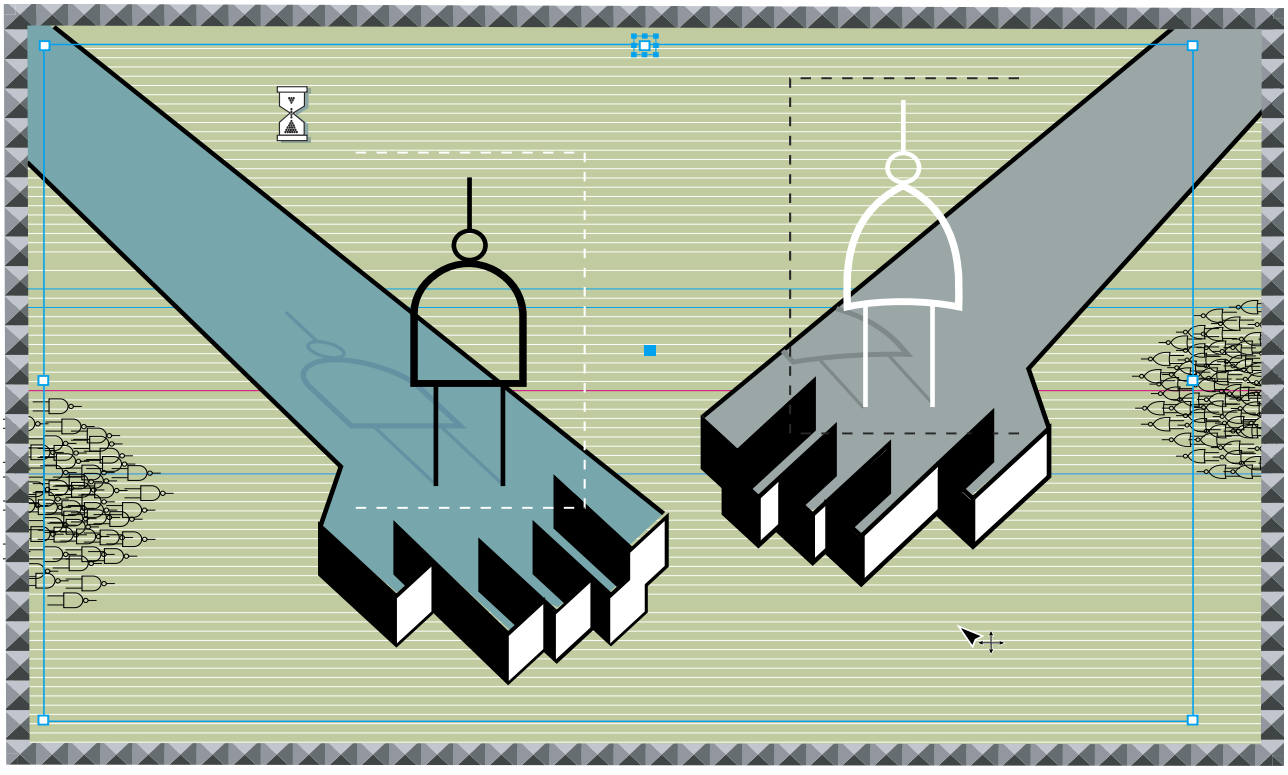

BY MICHAEL SANTARINI • SENIOR EDITOR

NAND versus NOR

WHICH FLASH IS BEST FOR BOOTIN' YOUR NEXT SYSTEM?

If you are a handset-chip architect, you have more choices than ever when it comes to picking a memory architecture for your next project. Users can go with tried-and-true methods using NOR for system booting or try their hands at designing a new architecture that boots with two of today's hybridized flash chips: Samsung's OneNAND and M-Systems' mDOC (mobile disk on chip).

The hybridized model promises to eliminate the pricey NOR device for high-end-system booting and to handle storage, too. In demand-paging architectures, it even promises to reduce the amount of RAM needed, thus reducing overall system power and cost.

But opponents say that implementing hybrid architectures is complex and error-prone. Intel, the current leader in the traditional NOR market, claims a system can make only so many "reads" from a NAND before losing data-storage integrity, which can ultimately lead to system failures, especially in demand-paging systems.

Experts say there are pluses and minuses to implementing any of the flash architectures, so users have to find the right balance of target market and user, features, unit cost, and design cost for their next designs.

Today, designers can employ four primary flash-memory architectures: the traditional XIP (execute-in-place) model, the shadow model, the store-and-download model with NAND, and the newer store-and-download model with hybridized NAND.

THE XIP MODEL

The XIP model consists of NOR memory and volatile memory, likely DRAM (**Figure 1**). In the XIP model, the NOR memory executes code, and the volatile memory accounts for constantly changing system elements, such as variables, stack, and heat. In the XIP model, the NOR can also provide data and code storage as well.

"The advantage of the XIP model is simplicity, but the disadvantage is its dismal write speed," says Samsung associate director for flash marketing Don Barnetson. "It takes five seconds to write a new phone number into my cell phone, because NOR has dismal write performance. If you are taking pictures or downloading movies, NOR simply doesn't have the bandwidth to keep up. At best, XIP can move 100 Mbytes of data over the bus."

Advantages of the XIP model are that it is simple and well-understood across multiple system disciplines. So although the cost of NOR is high compared with that of NAND, designing it into a system is easier, thus it typically helps get designs to market faster.

THE SHADOW MODEL

To get around the storage and write-speed issues of the NOR device, some designers—especially those working on

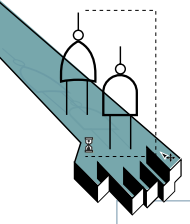

AT A GLANCE

- NOR devices are 30 to 40% more expensive than hybrid devices.
- The NOR architecture is easy to use but expensive.
- Hybrid flash devices offer improved write and read performance.
- Hybrid Flash devices require at least single-bit error correction and detection.

higher end phones—employ the shadow model (**Figure 2**). In this model, users boot a system with NOR and use a NAND for storage (or, in rare cases, a mini hard drive for storage); the volatile memory handles all of the execution.

"On power up, you execute under NOR and almost immediately shadow most of the operating system over to the DRAM, so your operating execution happens back and forth between the application processor and the DRAM," says Barnetson. "The reason being is the DRAM is an order of magnitude faster."

But although the shadow model offers seemingly the best of both sides of the flash coin—easy boot-up and faster read and write with better storage density—it is also expensive in that you are using a relatively pricey NOR only to boot up your system. The architecture is also a bit more complex, which means it consumes more design time or floor costs. The designs also tend to be power-hungry, because the volatile memory is constantly active. You can, however, offer customers richer features.

NAND STORE AND DOWNLOAD

To overcome the space issue, which is a huge factor in handheld devices, some design teams employ a store-and-download architecture (**Figure 3**). This model has no NOR, but there is an OTP (one-time-programmable) or ROM core designed into the main application processor. The processor loads information into the volatile memory (likely DRAM), which accesses a NAND core for data storage.

The architecture is a bit more complex and requires more initial engineering or floor costs, but ultimately, the unit cost of the system is less expensive. The main difficulty of the model is that users must employ extensive error-correction and -detection coding because NAND is typically less reliable. Storing and downloading designs tend to require more power, as the RAM takes a more active role.

HYBRID STORE AND DOWNLOAD

To get around the shortcomings of the traditional memory models, memory vendors are now offering hybridized devices. M-Systems and Samsung are the first two companies to offer such products. M-Systems introduced its MLC (multilevel-cell) mDOC EFD (embedded-flash drive) in 2000, and Samsung followed with its SLC (single-level-cell) OneNAND in the second half of 2004.

Both offerings mix SRAM, control logic, and NAND to basically create a best-of-both-worlds device that is supposed to look to the designer like a NOR Flash. But, OneNAND, for example, performs reads much faster than a standard NAND and at the same speed as a NOR device. It also offers better write performance than NOR devices (**Table 1**). M-Systems' mDOC offers faster write performance than NOR but slower read speed and comes in SLC and MLC configurations (**Reference 1**). The hybridized version one-ups the traditional store-and-download model by eliminating the OTP block (**Figure 4**).

OneNAND has a 16-bit NAND interface instead of the standard 8-bit multiplexed interface in storage NAND. OneNAND has 1 kbyte of boot RAM. "Effectively [it copies] the 1 kbyte from the NAND into the SRAM block, and then the processor executes from that SRAM," says Barnetson. "It gives it 1 kbyte, which is just enough for the processor to instruct the OneNAND to load the next boot loader into two other SRAM locations Samsung calls data RAM. The OneNAND has a total 4 kbytes of data RAM."

Barnetson says the second boot loader is much larger; you use it to configure the application processor and begin streaming code into the DRAM. "Depending on the system, which might have one more boot loader, what you end up with is enough of your operating system in the DRAM to fire up your OS," he says. "You don't need the OTP block."

Hybrids require less error-correction and -detection coding than store-and-download models with standard NAND. Most hybrids require single-bit error-detection and -correction coding, but vendors recommend double to be safe. John Nation, marketing manager at Spansion, says that NOR vendors could

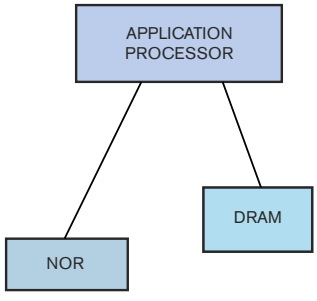

Figure 1 Advantages of the XIP (execute-in-place) model include a long legacy of implementation, thus the ability to take advantage of existing software and staff knowledge. Disadvantages include its price and write speed. NOR is 30 to 40% more expensive than NAND. A write speed of 150 kbytes is slow for smartphone features.

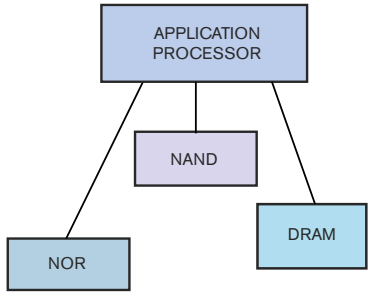

Figure 2 The shadow model offers the ability to use existing software for code and to use NAND for data. Users benefit from good performance from the NOR on read and from the NAND on write. Shortcomings of this model are primarily cost and space (two chips instead of one), and the NOR is still expensive.

Micrel Ethernet Solutions:
We've Got Your Processor Covered

Micrel's Family of Ethernet Solutions Addresses All Your Networking Interface Needs

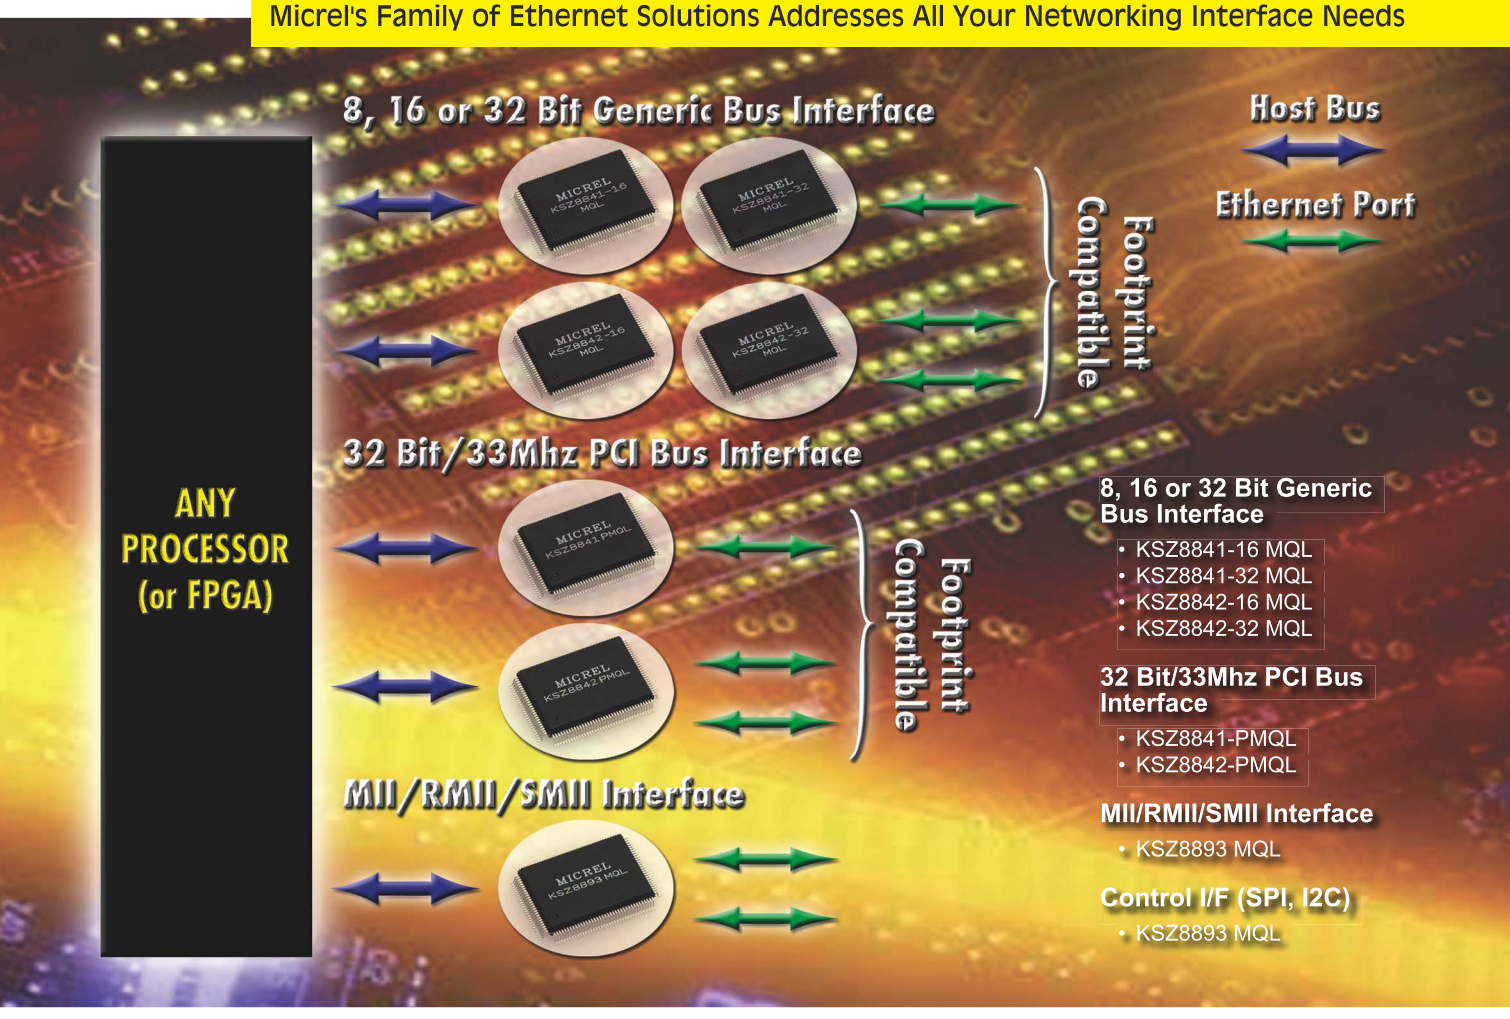

The road to network connectivity is never smooth but Micrel has you covered. Whether you need networking via an 8, 16 or 32-bit generic-bus, PCI-bus, MII, RMII, or SMII host interfaces, Micrel has the answer in easy to install single and dual-port Ethernet solutions. The devices address the growing need for streamlined networking connectivity in IP-Set Top Boxes, VoIP phones, Network Printers, Industrial Controls and networked Game Console applications, to name but a few. The dual port devices have the lowest latency (sub 310nS) in the industry and are ideal for daisy-chaining applications, or simply as two port switches to connect to voice, video and data.

All of the ICs incorporate HP Auto-MDIX to take the guess-work out of whether your device is connected using straight or crossover cables. In addition, Micrel's LinkMD™ cable diagnostics function not only determines the length of the cable and the distance to fault, but also diagnoses common cabling faults such as open and short circuits. These features reduce the need for costly customer calls and IT service requests. Along with Micrel's trademark high reliability, outstanding performance, and low power consumption, the KSZ88xx family offers ideal solutions for applications that require compact, cost effective, RoHS compliant networking connections.

For more information, contact your local Micrel sales representative or visit us at www.micrel.com/ad/ks88xx.
Literature: 1 (408) 435-2452 • Information: 1 (408) 944-0800

MICREL®
www.micrel.com

offer error detection and correction with the devices, but vendors test NOR devices so thoroughly that they don't require that extra step. "Users don't have to spend time worrying about ECC [error-correction coding] with a NOR," says Nation.

But Blain Phelps, director of worldwide marketing at M-Systems, says that error correction and detection of hybrids, or what M-Systems calls EFDs, are well-understood. "EFDs have been around for a while," says Phelps. "Ignorance is no longer an excuse. You either go with one or the other. We know how to work with both of them and integrate both of them. It gets back to cost, performance, and use model."

OneNAND's interface can read and write at 108 Mbytes/sec. Samsung claims that although the NAND hybrids can replace the stand-alone NOR and OTP block in a design, their read and write speed makes them especially well-suited for advanced operating systems using demand paging, which can ultimately reduce the amount of DRAM a system requires. Typically, other flash modes simultaneously load the entire operating system and application.

"In demand paging, you only bring into DRAM the things you need," says Barnetson. "You don't have to load the entire OS and all the applications simultaneously. Because the interface of OneNAND is so fast, if you use it in demand page, that likely means you don't need as much DRAM." Less DRAM, of course, cuts down on the power budget, form factor, and cost.

Not all vendors share the same view. Ed Doller, chief technology officer of Intel's flash group, says there is a huge difference between using NAND in a code environment and using it as a storage medium. "Demand-page systems are extremely complex to design, manage, and test," says Doller. "The minute you start burdening a cellular phone with a demand-paging system, your engineering validation and resources [increase] dramatically."

Doller says he is beginning to see that purchasing departments are starting to apply more pressure on system architects to use the NAND architectures. "None of those architects will say there is a technical reason for switching to NAND in a demand-page, store, and download architecture," says Doller. "Most designers despise it. They would rather be working on phone features than debugging architecture ... The minute you are late to market because you are debugging a demand-paged OS versus getting revenue for those phones, you have a huge problem."

The unit cost for a hybridized NAND is 30 to 40% cheaper than NOR at the same density; stand-alone NAND is slightly less expensive than hybrid NAND.

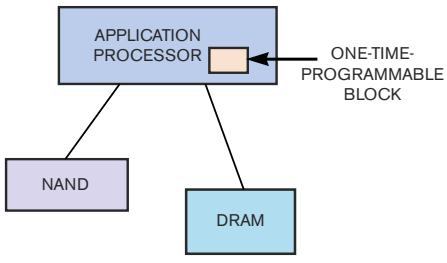

Figure 3 Pluses of the store-and-download model include performance, cost, and simplicity. The architecture has a better write speed than NOR and reduces energy consumption during write. Minuses include slower read speed than NOR, a requirement for users to adopt new software infrastructure, and a processor that can boot from NAND.

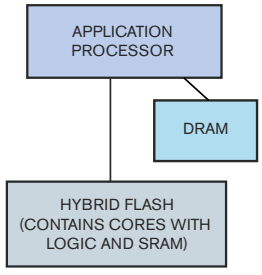

Figure 4 The hybridized model offers improved performance and lower cost than NOR at comparable densities. The device has equal read speed to NOR and better write speed than NOR. It also has reduced energy consumption in writes and doesn't require a processor with the NAND controller. The drawback is that it requires a new software infrastructure.

OTHER METHODS/DEVICES?

Although the hybridized-NAND market currently contains only two players, a lot of other players are eyeing it closely. To date, Spansion is the only traditional NOR vendor actively working on a response (nonverbal, that is). The company is currently developing ORNAND, which it bases on its MirrorBit technology and which promises faster write speeds than NOR and faster read speeds than NAND. Spansion expects to ship its first ORNAND devices to customers in 2006.

Meanwhile, Infineon Technologies, which makes NAND, has yet to jump into the market and in part is expecting the market may come to it. Indeed, Eugene Chang, a senior marketing man-

TABLE 1 PERFORMANCE COMPARISON BY SPECIFICATION

	NOR multilevel cell (Mbytes/sec)	NAND 90-nm single-level cell (×8, large block) (Mbytes/sec)	Samsung OneNAND 90 nm (Mbytes/sec)	M-Systems mDOC 90-nm H1 (Mbytes/sec)
Read	108	16.2	108	9.5 (sustained)
Write	0.14	6.8	8.2	3.9 (sustained)
Erase (single)	0.11	64	64	NA
Erase (multiple)	0.11	NA	2	NA

Notes:
Samsung OneNAND and M-Systems' mDOC hybrid devices offer an alternative to NOR for booting systems. OneNAND is a system on chip, and mDOC is a multichip module.
NOR data courtesy of Samsung.

Take 10/100 Mbps Ethernet to the EXTREME!

DP83848 PHYTER® 10/100 Mbps Application-Specific Single Port Physical Layer for Commercial, Industrial, and Extreme Temperature Applications

Features

- IEEE 802.3u compliance
  – *Guaranteed by Test*
- UNH compliant
- Lowest latency for real-time operation
- Integrity utility for system monitoring and maintenance
- 4 kV ESD protection (HBM)
- 25 MHz clock output
- Low power <270 mW
- Auto-MDIX
- Selectable MII/RMII interface
- Available in LQFP-48 packaging

AVAILABLE LEAD-FREE

Ideal for embedded computers, high-end peripherals, industrial controls, building/factory automation, transportation, test equipment, and wireless basestations

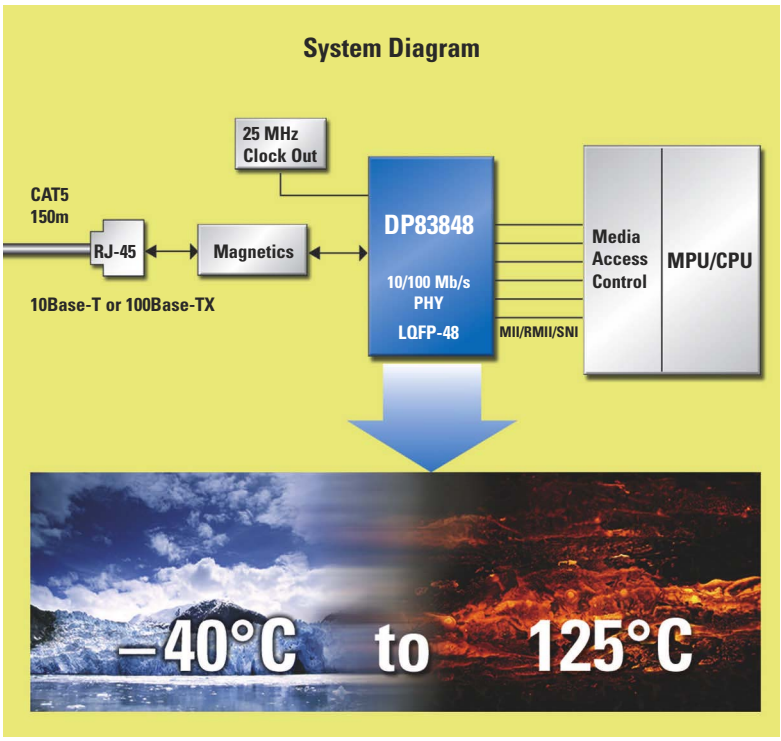

Unique Specifications

Product ID	Grade	Temp. Range (°C)	Cable Reach (m)	JTAG
DP83848C	Commercial	0 to 70	137	–
DP83848I	Industrial	-40 to 85	150	1149.1
DP83848YB	Extreme	-40 to 125	150	1149.1

For samples, evaluation boards, datasheets and online design tools, visit us today at:

ethernet.national.com

Or call 1-800-272-9959

National Semiconductor
The Sight & Sound of Information

TABLE 2 WORLDWIDE-MEMORY-REVENUE HISTORY AND FORECAST FROM 2000 TO 2010 (MILLIONS OF DOLLARS)

	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	CAGR 2004 to 2010 (%)
NAND flash	1563	1378.2	2363.8	4131	7007	10,134.5	12,996.2	11,662.8	14,365.7	14,420.8	15,077.9	13.6
NOR flash	10,141.1	6873.2	5820.1	6583	8429	7121.4	6871.5	7191.5	8055.8	6335.8	6324.5	−4.7

Notes: Data courtesy of Gartner Dataquest (August 2005). CAGR: compound annual-growth rate.

ager for Infineon, believes that hybridized devices will have a short shelf life, because some IP (intellectual-property) vendors will figure out a way to create a controller core that allows users to boot directly from a standard NAND, which is cheaper than a hybridized NAND.

Chang argues that the controller core, unlike an OTP device, would actually facilitate some of the bit switching to speed the NAND read performance, not just wake up the NAND. "It poses a great opportunity for an IP company," says Chang. "If you can create a controller that can be dropped into an ASIC or processor and package it with firmware and a pure NAND interface, [you] could probably make some serious money."

Lane Mason, Denali Software's memory market analyst, notes that Denali offers a OneNAND controller, but he declines to say whether the company is working on a controller to boot from a standard NAND.

Intel, which currently holds a slight lead over Spansion in NOR-market share, is closely watching the NAND market, especially its movement into NOR territory. "Right now, we look at the NAND market and hockey-stick projections and ... at various options to service that market," says Doller. "Over the long term, that option is next-generation memory technology. We are evaluating what makes sense for Intel from a NAND perspective. There hasn't been a strategy meeting here in the last few years about whether or not it makes sense for Intel to be in that business. We continue to evaluate it."

HAVING A STRONG PLAY

So although debate continues over which memory architecture is best for your system, there is little doubt that hybridized-NAND devices are indeed making an impact in the high-end mobile-handset market. According to Neta Yacoby, worldwide marketing manager for mobile at M-Systems, six of the seven largest mobile-phone manufacturers in the world use mDOC.

"We see convergence in the mobile-phone market," says Yacoby. "More applications ultimately mean there will be a growing need for more storage, driving the replacement of NOR with NAND."

Meanwhile, Samsung claims to be shipping 2 million OneNAND devices a month. "Two million per month is a very small percentage [compared with NOR], so we have a huge opportunity ahead of us," says Barnetson.

Analysts are still trying to figure out whether hybrids will add more fuel to the blazing NAND market. This year, NAND revenue for the first time in its history surpassed NOR revenue (**Table 2**). Joseph Unsworth, senior analyst for semiconductor memories at research company Gartner, says that most of that revenue growth came from new markets, such as USB sticks and digital cameras, but some growth—including emerging MP3 players such as Apple Nano—is coming at the expense of the hard-disk-drive market. And growth will continue as NOR declines.

"The NOR market losing value is due in large part to Intel cutting prices on NOR," says Unsworth. He notes that Intel a few years ago irked customers by raising the price of NOR. The high price allowed Spansion, which was late to market with MirrorBit, to grab market share from Intel, which then responded by dropping the price of NOR. That move, says Unsworth, has been the main reason for NOR revenue's decline and NAND revenue passing it for the first time.

"A NOR-market recovery, if indeed it ever recovers, will be further slowed by hybridized vendors cutting into the high-end NOR market. Certainly one of Intel's biggest competitors today is the NAND flash," says Unsworth. "Samsung and M-Systems are competing for their business. Intel has to keep NOR flash pricing aggressive ... to make a convincing argument to keep NOR flash in handsets."

Unsworth is in the process of releasing Gartner's first study on the hybrid flash market and its impact on NOR. He was unable to share any data but did note that future dominant memory systems will depend largely on when full-featured phones become mainstream, whether NOR vendors can keep up with the density requirements for those phones and still turn a profit, and whether the inevitable use of removable storage will usurp embedded data storage.

"In the next few years most of these phones will be able to play music," says Unsworth. "That means that removable slot will become imperative to help people move songs and photos between phones and personal computers and printers and to move the songs to their next cell phones."

Doller notes that the Sony W800 Walkman, one of the more advanced phones, uses NOR for code execution, employs a small amount of embedded NAND, and includes an additional card slot for data storage. "As the 512-Mbit device hits the street, I'd expect the next generation will simply replace that embedded NAND with the NOR but leave that card slot," says Doller.

Doller believes that the market will sway back and forth between using NOR and NAND architectures to boot systems. "Once the NOR gets into a particular lithography, because our die sizes are fundamentally smaller, it ultimately becomes a cheaper solution in terms of floor cost," he says. "In time, the main questions will

8/10/12-Bit Pin- and Function-Compatible DACs

Select the Resolution Your Application Demands

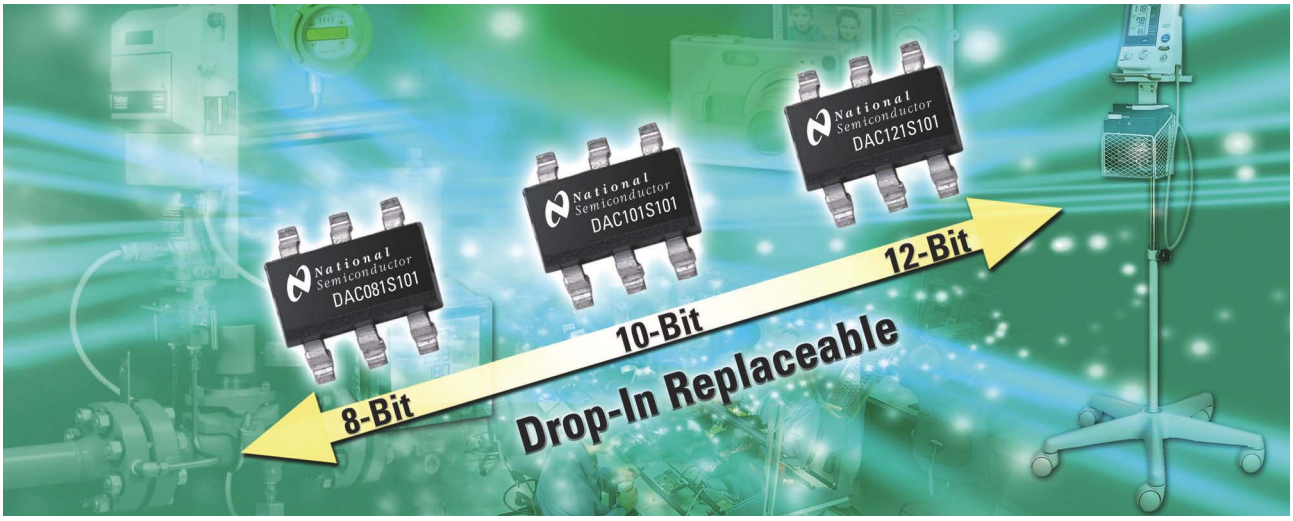

Part Number	Resolution (bits)	Setting Time (µSec)	INL (LSB)	DNL (LSB)	Power Consumption mW (at 3.6V, 20 MHz)
DAC121S101	12	10	±2.6	+0.25, -0.15	0.57
DAC101S101	10	8	±0.6	+0.15, -0.05	0.55
DAC08S101	8	5	±0.16	±.04	0.55

Features

- Guaranteed monotonicity
- Low-power operation
- Wide power supply range (+2.7V to +5.5V)
- Power-on reset to zero volts output
- Rail-to-rail voltage output
- SPI interface
- Power down feature
- ThinSOT and MSOP packages

AVAILABLE LEAD-FREE

Ideal for battery-powered instruments, digital gain and offset adjustment, programmable voltage and current sources, and programmable attenuators

For free samples, evaluation boards, datasheets and online design tools, visit us today at:

national.com/adc

Or call 1-800-272-9959

National Semiconductor

The Sight & Sound of Information

be: What is that lowest floor cost, and at what density is [it] occurring?"

"There is a suitable architecture for store-and-download phones, and I think some of the high end and more so the mainstream phones are going to stick with the NOR flash because companies like Intel and Spansion will have more aggressive NOR pricing, and it is going to be accompanied by a removeable slot," says Unsworth. "That way, if the consumer wants more storage, they can buy it."

Most believe that mobile-phone-service carriers will at first be reluctant to supplement the cost of phones with removable storage, because the carriers currently make extra revenue from charging customers to store data on their networks. Record companies may also be unhappy about the situation, because they want to charge customers to download a song for each device rather than have them download a song only once onto a storage medium that can play on multiple devices. All agree that over time, the obvious consumer preference for removable-data storage—ultimately promising the ability to move stored music, photos, and video from their mobile handsets to other devices, such as PCs, printers, and new mobile handsets—will likely prevail.

THE CUTOFF AT WHICH NAND STARTS TO BECOME MORE FEASIBLE THAN NOR-BASED ARCHITECTURES IS THE 256-MBIT DENSITY, BUT NOT EVERYONE AGREES WITH THAT ASSERTION.

Today, the cutoff at which NAND starts to become more feasible than NOR-based architectures is the 256-Mbit density, but not everyone agrees with that assertion. Samsung, for example, doesn't offer OneNAND with less than 128 Mbits, so it claims 128 Mbits is the point at which a hybridized architecture becomes more feasible; Intel claims that 512 Mbits, the biggest NOR the company offers, is the inflection point.

Doller notes that Intel has to stay at the top of its game. "We clearly need to make sure we move our lithography as fast as humanly possible in that our densities keep up with the sweet spot of the market, because if they don't, we've got a business imperative," says Doller. "We are in a great path for meeting that requirement. It doesn't mean NAND is not going to make its way into handsets. I do fundamentally believe there will be data growth that will outpace at the high end the requirements of a NOR-only solution."

Doller says that ultimately, the industry will require a new flash device. He notes that Intel is in the early stages of a study indicating that NAND has an inherent flaw that makes it ill-suited for booting systems, especially in demand-paging architectures. "NAND inherently has a limit to the amount of times you can read it before you have disturb issues where you start to change the data in the device," says Doller. "I fundamentally believe that the demand-page architecture out of NAND will be problematic for that reason."

According to Doller, Intel took one OEM's platform, engineered around it, and collected data that shows that this particular platform "could be problematic." The real question is whether this problem is isolated, says Doller. "We're currently looking at a completely different OEM handset to try to ascertain if it is a problem there, also." He believes the study will at least show that the number of reads performed in demand paging exceeds the number of reads that testers spec for NAND-based parts. EDN

REFERENCE

1 www.m-sys.com/NR/rdonlyres/8E18535D-7A08-4E2B-8C0A-223C5F27F618/0/Alternatives_to_Using_NAND_FLASH.pdf.

FOR MORE INFORMATION

Apple www.apple.com	**M-Systems** www.m-sys.com
Denali Software www.denali.com	**Samsung** www.samsung.com
Infineon Technologies www.infineon.com	**Sony** www.sony.com
Intel www.intel.com	**Spansion** www.spansion.com

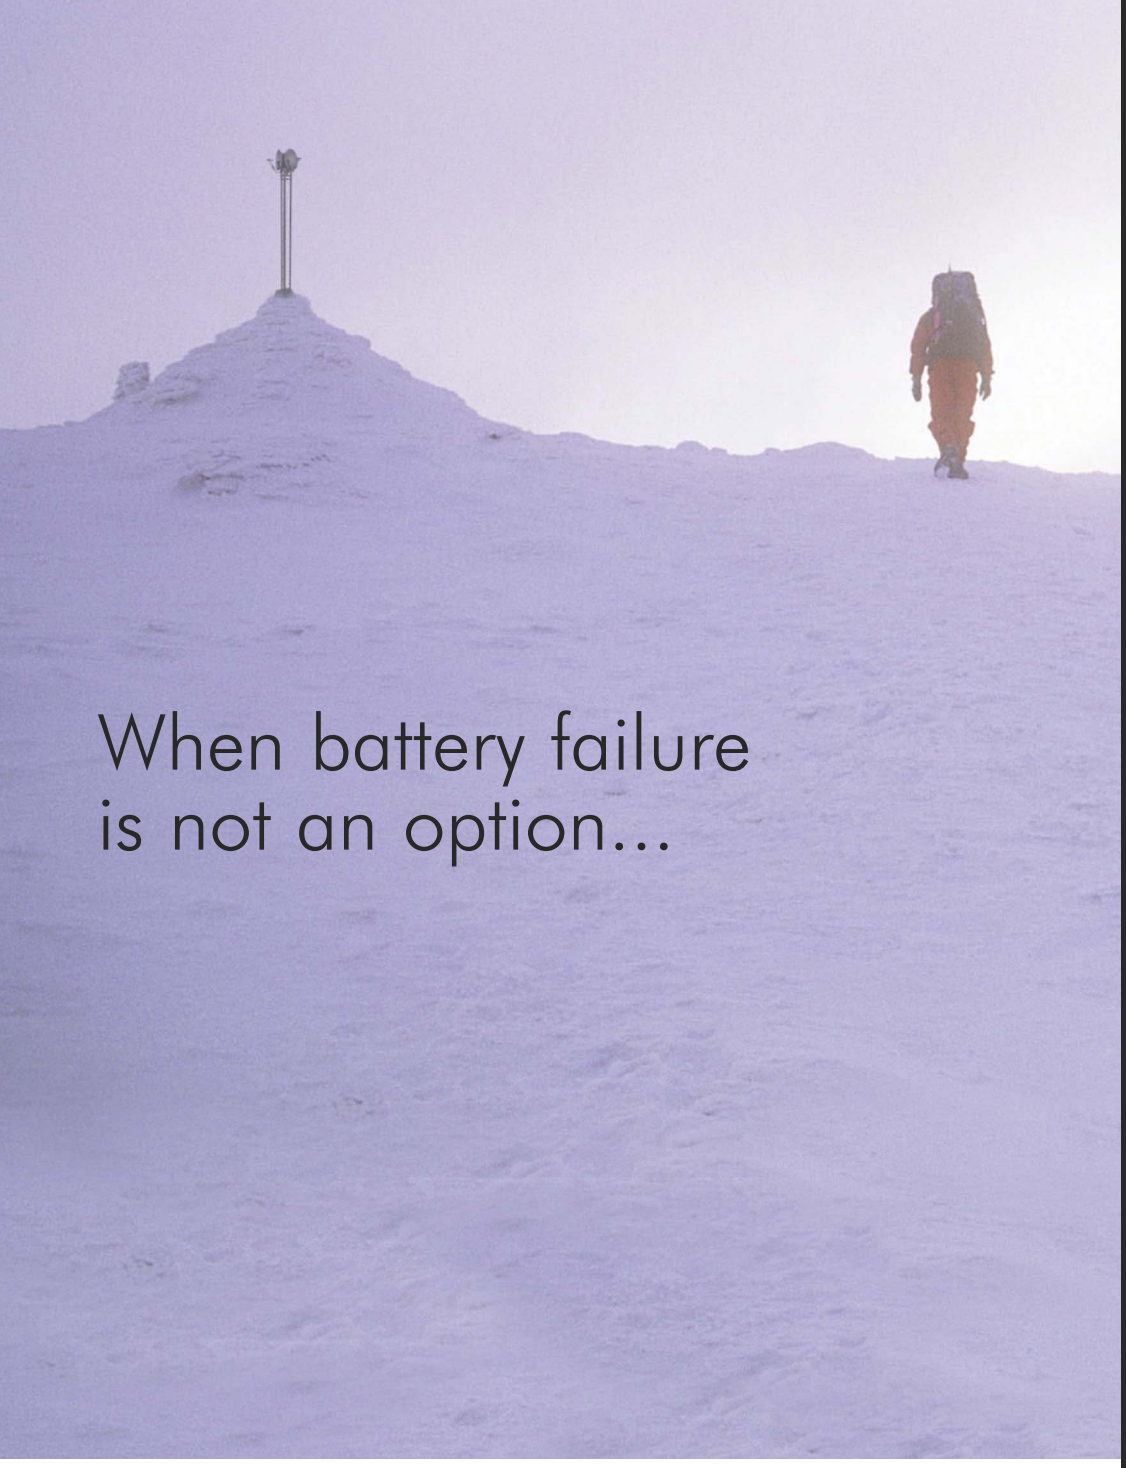

PROVEN
20
YEARS

When battery failure is not an option...

Remote locations and extreme environments are simply no place for unproven battery manufacturers. When battery failure is not an option, your only option is Tadiran, the only *proven* 20-year battery. So don't take chances, take Tadiran.

TADIRAN
BATTERIES

2 Seaview Blvd.
Port Washington, NY 11050
1-800-537-1368
(516) 621-4980
Fax: (516) 621-4517

www.tadiranbat.com

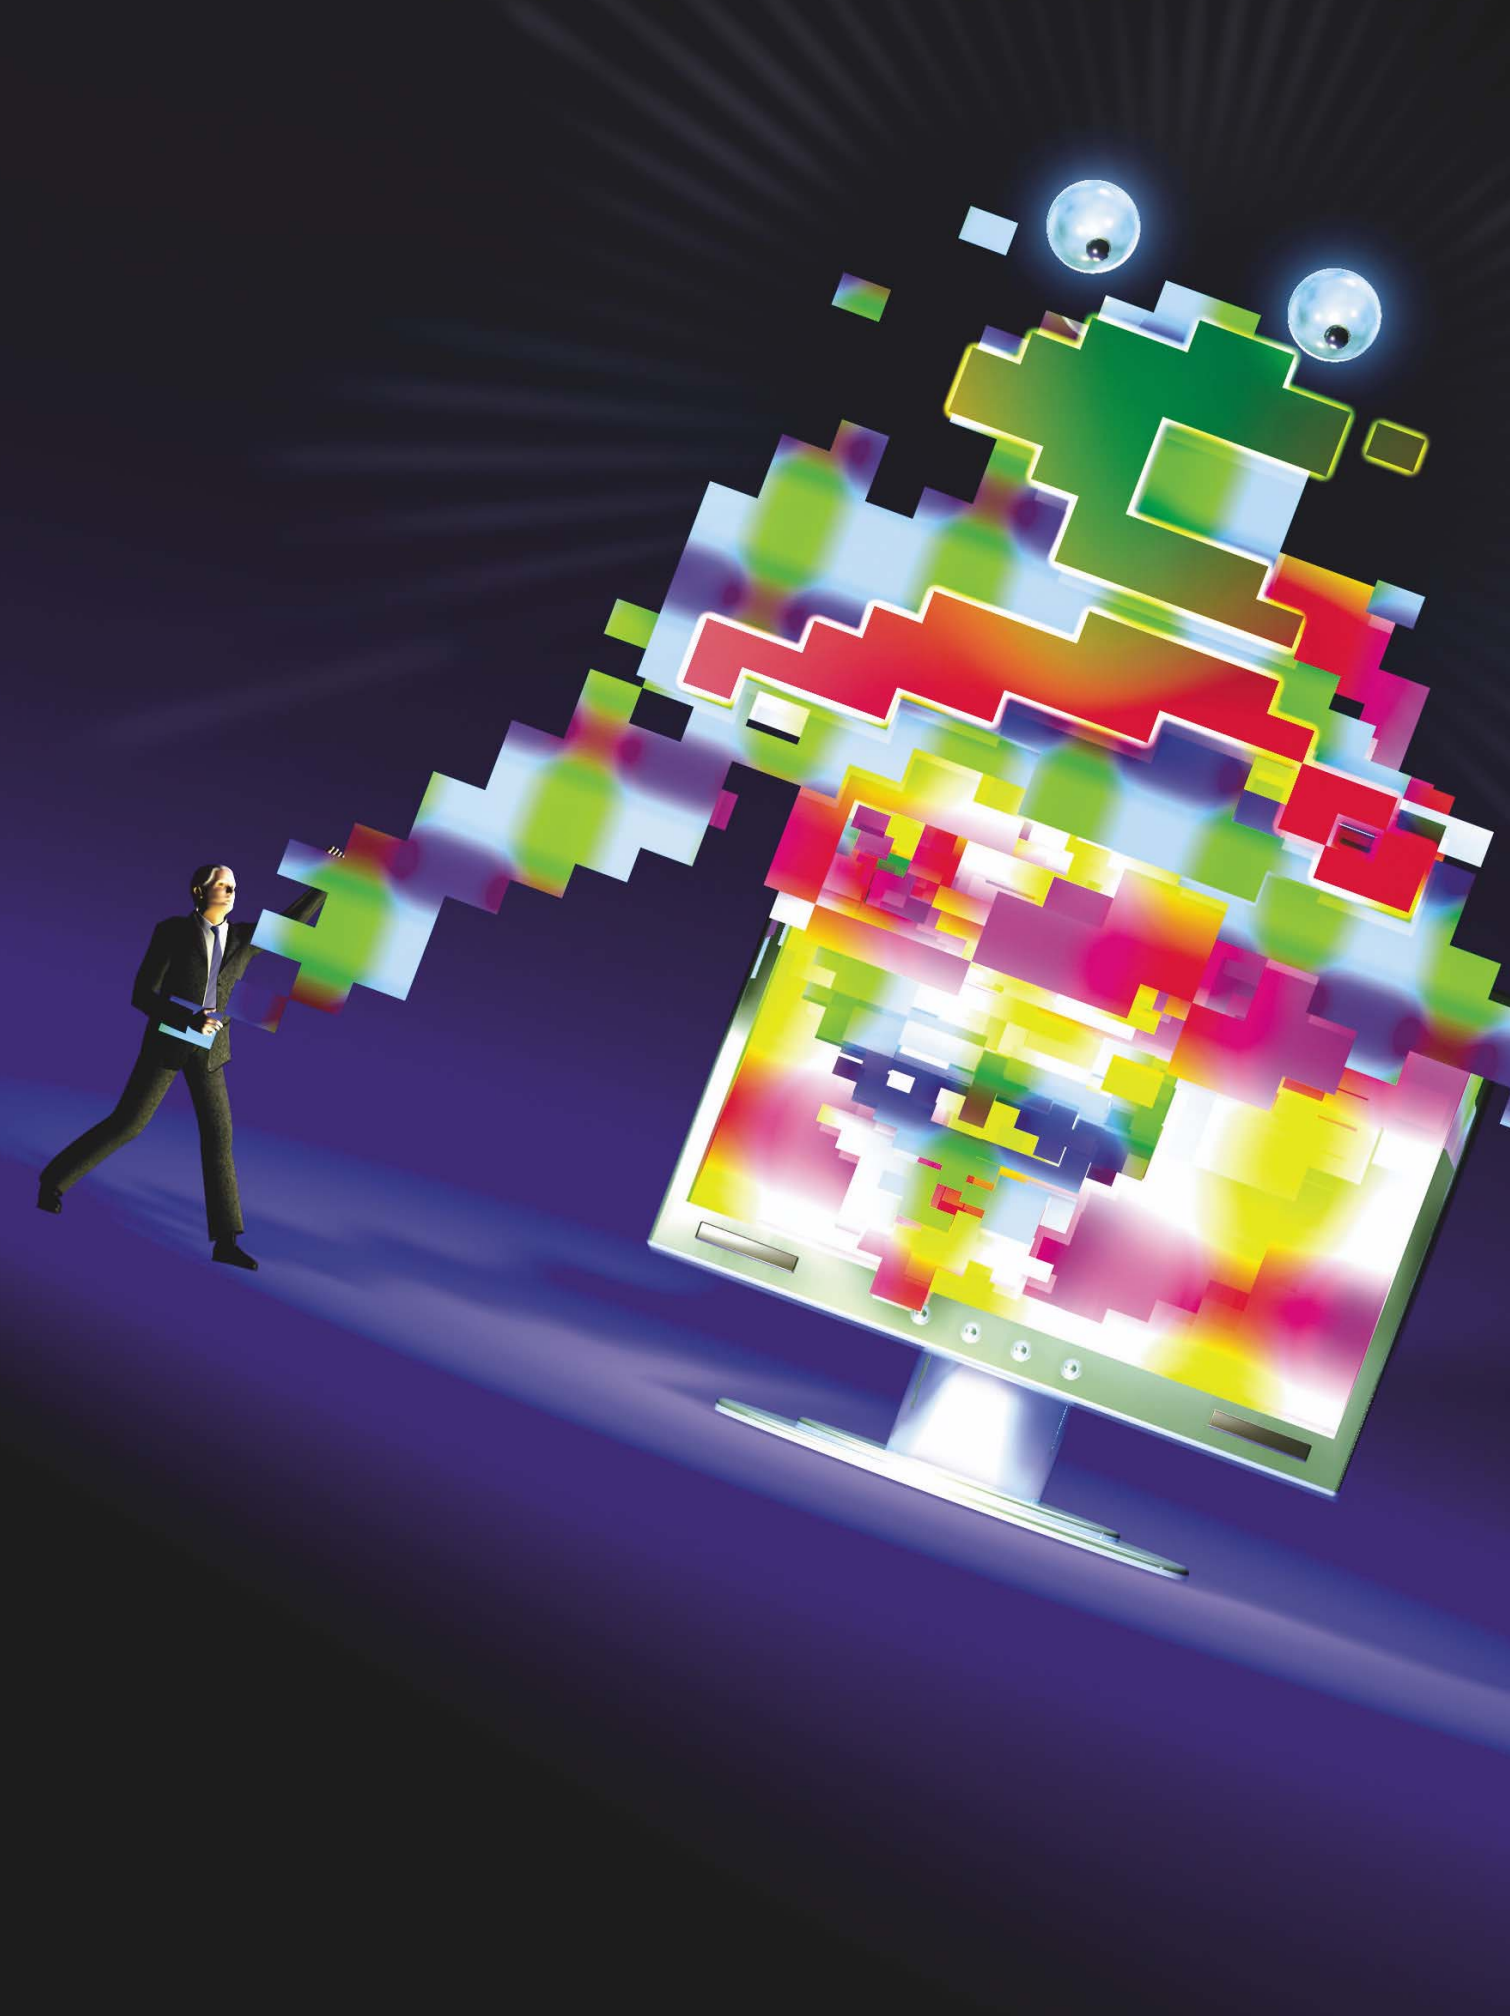

BY BRIAN DIPERT • SENIOR TECHNICAL EDITOR

INSTIGATING A PLATFORM TUG OF WAR:

Graphics vendors hunger for CPU suppliers' turf

PROCESSOR VENDORS' FISCAL FORTUNES HINGE ON THE ABILITY TO SELL INCREASINGLY POWERFUL CHIPS. GRAPHICS-PROCESSOR VENDORS, TOO, DEPEND ON UPGRADES TO PAVE A PATH TO CONTINUED SUCCESS. THESE UPGRADES WILL INCREASINGLY RELY ON STEALING FUNCTIONS THAT THE CPU ONCE HANDLED.

In his keynote at mid-July's IEEE Hot Chips conference (www.hotchips.org), Nvidia's chief scientist David Kirk made a startling public confession, validating a trend that graphics-industry observers had long suggested was under way. Kirk stated: "As pixel/vertex/triangle growth slows and plateaus ... "(**Reference 1**). Translation: The increasingly complex treadmill of 3-D application GUIs (graphical user interfaces) is losing steam, although Kirk was quick to point out that other performance-consuming graphics factors, such as the number of color samples per pixel (antialiasing) and the number of calculations per pixel and per vertex, will take up some of the slack.

What factors are driving the graphics-complexity slowdown? With the exception of Windows XP Media Center Edition-based systems and like-minded computers tethered to HDTV displays in living rooms and dens, most Macs and PCs drive monitors with 17-in. and smaller diagonal viewing dimensions (**references 2** and **3**). Due in part to the raster-based approach of implementing icons, cursors, text, and other GDI (graphics-display-interface) elements available in current Windows operating systems, you rarely encounter resolutions beyond 1024×768 pixels. (The Windows OS, unlike the more advanced vector-based-graphics approach in Mac OS X, doesn't enable device-independent high-quality display scaling.) Multidisplay-PC configurations haven't widely caught on. And, although incremental quality improvements require an exponentially greater number of transistors running at increasingly higher frequencies, as well as exponentially more complex software driving those transistors, those improvements are less and less noticeable to users. This scenario is particularly true with so-called first-person shooters and other fast-action genres that dominate gaming.

A pessimist might interpret Kirk's words as a foretelling of ill fortunes for graphics vendors. An optimist, however, might consider that, during that same presentation, Kirk forecasts demand growth with-

AT A GLANCE

- CPU and GPU (graphics-processing-unit) evolutions are bringing the two system building blocks into more direct competition.
- The migration from the fast unidirectional bandwidth of AGP (accelerated-graphics port) to the even speedier and bidirectional-transfer capabilities of PCI Express is a key element of the system transformation.
- Windows-based systems currently lead their Mac OS-powered peers in hardware-assisted video-decoding and -encoding capabilities.
- Image editing, both still and video, is another area in which GPUs can lend a helping hand. In the Mac-OS world, they already do.
- In the future, GPUs will likely not only render, but also animate polygon-based characters. It's less clear though, which chip will process the physics-based interactions *between* characters.

out bound for the number of general-purpose, programmable, 32-bit Gflops per pixel and, more generally, that GPUs (graphics-processing units) will become general-purpose parallel processors. Nvidia's prediction isn't merely a delusional pipe dream; recent historical developments and current platform-architecture trends both support the premise. Although the display and operating-system limitations constitute part of the reason for the graphics slowdown, another key factor is that GPU's can now process polygon data and other graphics-related traffic faster than the CPU can provide it (**Figure 1**). If the GPU can in the future operate at higher levels of scene abstraction, it can dodge the CPU bottleneck. David Blythe, Microsoft's software architect for Windows Graphics & Gaming Technologies, strongly supports such an approach. In his presentation at July's DirectX Meltdown 2005 conference (www.microsoft.com/directx), he flatly stated that "games are CPU-limited" and that multicore CPUs are "not a panacea" (**Reference 4**). In addition to making games multithreaded, Blythe exhorted his audience to "offload to the GPU," citing five reasons that this shift now makes sense:

- Shader models are increasingly expressive;
- memory datapaths support iterative calculations;
- better data transfers to and from the CPU are now possible;
- data-amplification support exists, for example, with geometry shaders; and
- high-level shader languages are evolving to support new capabilities and abstractions.

GPUs are transforming into more generic coprocessors, with the evolution of APIs (application-programming interfaces) that enable operating systems and applications to tap into their capabilities setting the pace. As this transformation continues, GPUs will potentially also be able to wrest away other functions that you currently implement in software running on CPUs. However, CPU suppliers are highly motivated to keep their own upgrade treadmills smoothly running and won't allow the GPU vendors' aspirations to go unchallenged (see **sidebar**, "The CPU perspective"). Other specialty-function processors, too, are hoping for their turns in the limelight and view both the CPU and the GPU as competitors. And this tug of war isn't restricted to PCs; it will potentially play out in *any* system that includes a display.

OPENING DOORS TO CHANGE

Historically, CPUs and "graphics accelerators," as they were originally called, were quite different devices and, as such, were symbiotic partners. The CPU is software-driven and is therefore infinitely flexible in the kinds of functions it can perform. CPUs also, in the words of Computer Science Professor Emeritus H Norton Riley, historically implement a "model, deeply rooted in the von Neumann tradition ... [which] sees a program in terms of an orderly execution of instructions as set forth by the program. The programmer defines the order in which operations will take place, and the program counter follows this order as the control executes the instructions" (**Reference 5**). Graphics accelerators, in contrast, have long simultaneously operated on multiple pieces of information. (As Nvidia's Vice President of Technical Marketing Tony Tamasi put it in a recent presentation: "Graphics is embarrassingly parallel.") Historically, however, graphics accelerators were hard-wired state machines that took in graphics primitives and spat out rendered pixels (**Reference 6**).

Both CPUs and GPUs have evolved in recent years, however, and in directions that diminish their distinctions and redefine their relationship, pointing them toward a colli-

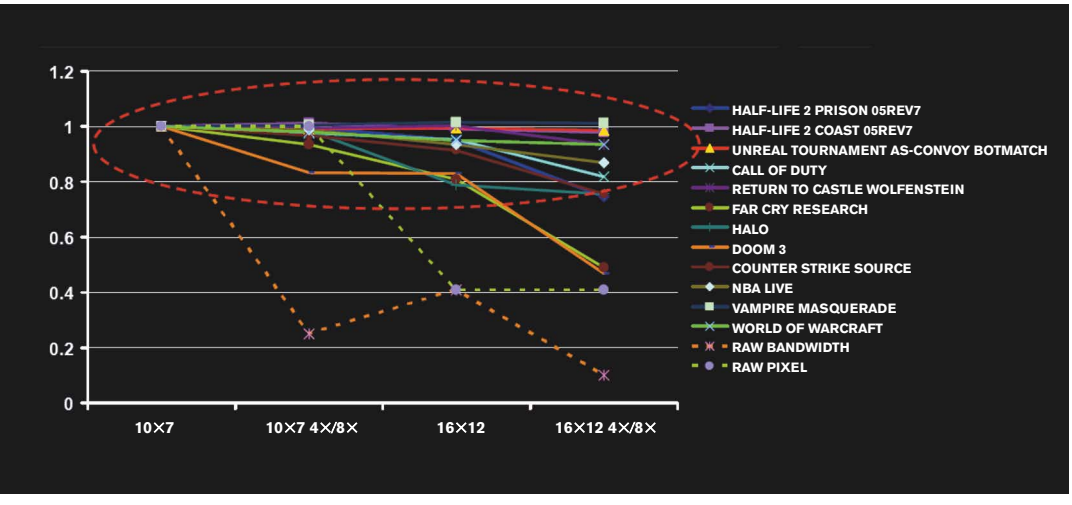

Figure 1 Minor to nonexistent decreases in frame rates with increasing graphics complexity suggest that the CPU, not the GPU, is the performance bottleneck (courtesy Nvidia).

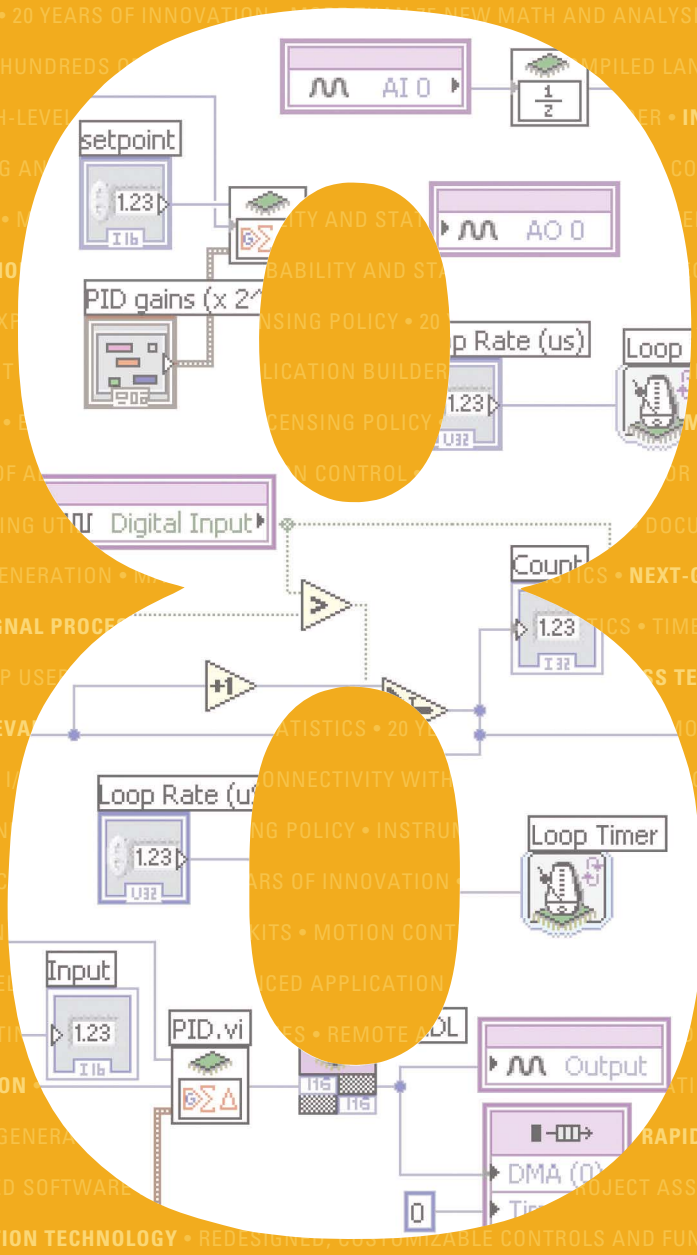

LabVIEW 8 – Discover Distributed Intelligence

National Instruments LabVIEW 8 delivers distributed intelligence – a powerful suite of new capabilities for you to easily design, distribute, and synchronize intelligent devices and systems. Whether you are testing electronics, mechanical structures, or dynamic systems, NI LabVIEW 8 equips you to tackle the most difficult test challenges with the ease and productivity of graphical development. With more than 100 new features combined with the open, project-based LabVIEW environment, connectivity to thousands of measurement devices, and scalable development tools for projects ranging from benchtop measurements to ATE systems, LabVIEW 8 is the intelligent solution for your application.

LabVIEW

ni.com/labview

(800) 453 6202

NATIONAL INSTRUMENTS™

sion course. Beginning with SIMD (single-instruction, multiple-data) instruction sets, such as Intel's MMX (multimedia extensions) and subsequent iterations of SSE (streaming SIMD extensions), AMD's 3DNow!, and the PowerPC's AltiVec, CPUs could simultaneously apply a common instruction to multiple pieces of data (**Reference 7**). Superscalar CPUs, which could concurrently process multiple independent instructions, came next. Intel's HyperThreading feature took parallelism to the next level, enabling limited parallel execution of multiple instruction threads, and the multicore CPUs now emerging from numerous suppliers extrapolate this capability in a more generic form (**Figure 2**). And what of graphics accelerators? Consider this quote from Kirk's introduction to the seminal graphics reference guide *GPU Gems*: "We have entered the era of programmable GPUs. The graphics-hardware pipeline, which had not previously changed significantly in 20 years, was broken down to its component, hard-wired elements and rebuilt out of programmable, parallel-pipelined processors. In a hard-wired pipeline, triangle vertices are transformed and lit; triangles are rasterized; and pixels are shaded with diffuse lighting, specular exponentation, fog blending, and frame-buffer blending."

Kirk continues, "In a programmable pipeline, each of these operations is abstracted to its component memory accesses and mathematical operations. A programmer can still write a program that calculates the same results as a hard-wired pipeline ... but the opportunity presented is so much greater" (**Reference 8**). Another critical piece of the platform transformation, beyond the evolution of the CPU and GPU, is the enhancement of the bus that interconnects them. Multiple contending peripherals shared the bandwidth of PCI, and, in today's terms, it was as slow as molasses—32 bits and 33 MHz—for 1.05-Gbps peak unidirectional bandwidth. The graphics-tailored AGP (accelerated-graphics-port) bus's highest speed proliferation, the 8× variant, delivered 16.8-Gbps peak bandwidth but in only *one* direction—from the CPU to the GPU. Data flowing from the GPU back to the CPU traversed the AGP bus at much slower, 1× AGP (2.1 Gbps, or

THE CPU PERSPECTIVE

Talk to PC-microprocessor vendors AMD and Intel about the looming GPU (graphics-processing-unit) competitive threat, and you get somewhat divergent perspectives. "I've seen the same sort of thing in servers with TCP/IP-offload hardware," said AMD's Steve Demski, product manager for the microprocessor business unit's server and workstation business segment, at the early-July Siggraph conference in Los Angeles. "I think it's a good thing. Anything that can be done better, faster, and cheaper in the GPU ... that's fine. That just frees up the general-purpose CPU to do different things."

But would a transfer of power for some functions from the CPU to the GPU negatively impact AMD's CPU business in the future? "Look at the operating system," Demski responded. "Operating systems are multithreaded today. They will run better on a dual-core CPU." He pointed out that you will still have multiple applications: background tasks, pop-up blockers, virus protection, and more. "There are plenty of applications wanting that CPU, wanting even the dual-core CPU. If you can offload some of the more graphics-intensive things to the GPU, you're just going to speed up the overall system," he added.

"Our company focuses on what the customer needs and what the end-user experience is, and if that means working with better graphics cards and graphics-card companies, we will work with all of them," chimed in AMD public-relations manager Scott Carroll. "We don't compete with those guys; we compete with Intel," he added. Carroll pointed out that AMD removes the system bottlenecks as much as possible–for example, as the company has done with the Opteron direct-connect architecture.

What's Intel's opinion on the role of the GPU? Regarding Intel's support for video formats beyond MPEG-2 in core-logic chip sets, Patrick Smith, a graphics architect for the company, says, "At some point, we'll need to implement more and more of the decode architecture, primarily to support emerging usage models." By "emerging usage models," Smith is referring to the VIIV-branded digital home, for example, which combines a central hub, the digital server, with comparatively "dumb" media-playback devices scattered around the home. How does Intel reconcile the need for increasing amounts of dedicated video- decoding circuitry with the increasingly powerful CPUs on Intel's road map? "CPU horsepower alone isn't enough, no matter how many cores you throw at the problem," he says.

Smith is less sanguine, though, about the need for a dedicated physics processor. "Where we're doing physics today, in the CPU, there's ample compute power. I'm not sure we need to offload to yet another device and burden the system with even more cost," he suggests. He also points out the thermal issues of yet another high-powered processor in the system. The bottom line is that mainstream applications don't currently push the system and CPU hard enough to justify a hardware hand-off, according to Smith. He further admits that the graphics-chip companies are simultaneously competitors and partners; as partners, they were instrumental in co-defining PCI Express, for example. In contrasting AMD's and Intel's perspectives, keep in mind that Intel offers graphics cores in some of its core-logic chip sets, whereas AMD relies on semiconductor partners for both core logic and graphics support.

twice that of PCI) peak speeds, and upstream AGP traffic also underwent snooping for coherency and used PCI semantics.

With PCI Express Version 1, each four-signal "lane" supports simultaneous and bidirectional, 2-Gbps data transfers (2.5-Gbyte/sec raw bandwidth minus 8/10-bit-encoding overhead). In other words, it supports an aggregate sum of 4 Gbps of peak bandwidth. Common PCI Express implementations in today's PCs devote a 16-lane connection to the graphics subsystem, thereby delivering a truly staggering amount of bandwidth between the CPU and the GPU. In the past, it might have made no sense to send information to the GPU for intermediate processing, because, even though the GPU might process the data faster than the CPU could, passing the results back to the CPU would incur unacceptable latency. The move to PCI Express clearly relieves this AGP upstream bottleneck. And at the late-July Intel Developer Forum (www.intel.com/idf), presenters paved a path for the upcoming Version 2 PCI Express specification, which doubles the per-lane bandwidth yet again.

THE FIRST SHIFT

When DVD-ROM drives—primarily for playing video DVDs—began appearing in PCs a few years ago, GPU vendors saw their first solid opportunity to break out of the graphics-only box. The appeal was particularly attractive with laptops, whose CPUs were comparatively underpowered versus desktop systems. And, because they are battery-powered, laptops were comparatively more concerned with energy consumption. Hard-wired MPEG-2 decoding is more energy-efficient than the software-centric approach. Thus, beginning with the concluding color-space-conversion step and later broadening to earlier stage operations, such as iDCT (inverse discrete-cosine transform) and motion compensation, GPUs took over most of the 480-line-resolution DVD-decoding burden. (Nowadays, this takeover includes 720- and 1080-line-resolution HDTV, as well.) The desktop-versus-laptop differentiation continues to this day. The graphics core inside Intel's 945G desktop core-logic chip set, for example, handles color-space conversion,

Although the GPU in the future may absorb functions that currently take place on the CPU, could the opposite scenario play out? Could the CPU retake tasks that the GPU currently handles? Looking at the big picture, consider first the significant horsepower upgrades on both AMD's and Intel's published CPU road maps for the remainder of this decade. The processor vendors can unleash significantly larger marketing budgets than the GPU suppliers to convince customers to buy those CPUs.

Next, focus your attention on video encoding. Although the first few generations of Windows MCE (Media Center Edition)-based PCs included powerful CPUs and GPUs, they also contained dedicated MPEG-2-encoder chips to handle the PVR (personal-video-recorder) function. Microsoft wanted to ensure that, no matter what else was taking place on an MCE system at a given time, the consumer wouldn't experience dropped video frames or other recorded television blunders. In mid-August, however, CyberLink announced that its MPEG-2 software-encoder plug-in for Windows XP MCE 2005 had obtained approval from Microsoft and claimed that it drastically reduces costs for tuner-card manufacturers by avoiding reliance on hardware-chip sets when recording TV content with MPEG-2-video and -audio quality (Reference A).

Finally, look at graphics. Remember that, before the advent in 1996 of 3dfx's Voodoo 1 chip and the all-important *Quake* game that took advantage of it , 3-D-graphics acceleration was limited to a narrow niche of high-end workstations, along with specialized visualization applications, such as flight simulators. In the small amount of mainstream 3-D software that existed at the time, software running on the CPU fully rendered graphics primitives to pixels before their subsequent hand-off to the 2-D-graphics chip or, if the CPU also handled 2-D rendering, the RAMDAC. Now consider that, whereas the graphics core in Intel's current 945G chip set is shader-based, it hardware-accelerates only *pixel-shader* operations. (Engineers also sometimes refer to these operation as *fragment-shader* operations.) Vertex-shader code is software-emulated on the CPU. If 3-D doesn't expand beyond its rabid but minuscule gaming niche and if Kirk's predictions of slowing pixel, vertex, and triangle growth in that niche come to pass, will the processing burden shift back to the CPU, fueled by unrelenting bill-of-materials cost-reduction pressure?

This forecast may seem outlandish at first glance. But consider that, although neither Nvidia nor Sony comments on the rumor, industry insiders believe that including Nvidia's GPU (likely, a kissing cousin of the GeForce 7800 GTX) to Sony's Playstation 3 was a late-stage-development decision. According to industry gossip, as Sony originally structured the Playstation 3, it would completely handle graphics operations with its Cell processor. If the tales are true, then Sony's engineers were too optimistic this time around. But the fact that they even seriously *considered* a CPU-only approach to graphics says a lot about how the CPU-versus-GPU tug of war may evolve over the next few years.

REFERENCE

A www.gocyberlink.com/eng/press_room/view_886.html.

iDCT, and motion-compensation duties, and the mobile-tuned 945GM devotes extra transistors to tackle the variable-length lossless-decoding task.

Initially, DVD-playback software was forced, out of necessity, to interrogate the graphics subsystem and find out what GPU it contained and subsequently to include numerous GPU-specific routines that reflected each chip's hardware-acceleration capabilities. For example, early Nvidia chips had fewer MPEG-2 features than their ATI Technologies counterparts. (However, this gap has closed in recent years. Nvidia GeForce 6xxx and 7xxx chips, for example, contain three dedicated video engines for MPEG-2 decoding, motion estimation, and video processing. They also take advantage of the chips' shader processors for video functions.) Integrated graphics cores within core-logic chip sets also tend to have fewer features than stand-alone GPUs. This development and maintenance quagmire eased when Microsoft unveiled its DirectX VA (video-acceleration) API in late 2000. API support for video-decoding functions is one area in which Microsoft's operating systems continue to have a solid lead over Apple's Mac OS X. (This lead is surprising, too, given the Mac's historic strength in multimedia applications.) The only Mac OS application that currently taps into the video-decoding features of systems' graphics chips is Apple's own DVD Player program. This scenario is the reason that, for example, high-resolution MPEG-2 and -4 playback at full frame sizes and without dropping frames requires a high-end, dual-CPU G5 Power Mac Apple system. On the other hand, almost any Windows-based system for sale today can smoothly display high-resolution MPEG-2 streams without breaking a sweat.

With the last few generations of ATI, Nvidia, and competitors' GPUs, and with their evolution beyond hard-wired video-decoding pipelines to more generic shader-based processors, video-format support has broadened beyond MPEG-2. First is WMV9 (Windows Media Video 9), or, in SMPTE (Society of Motion Picture and Television Engineers) terminol-

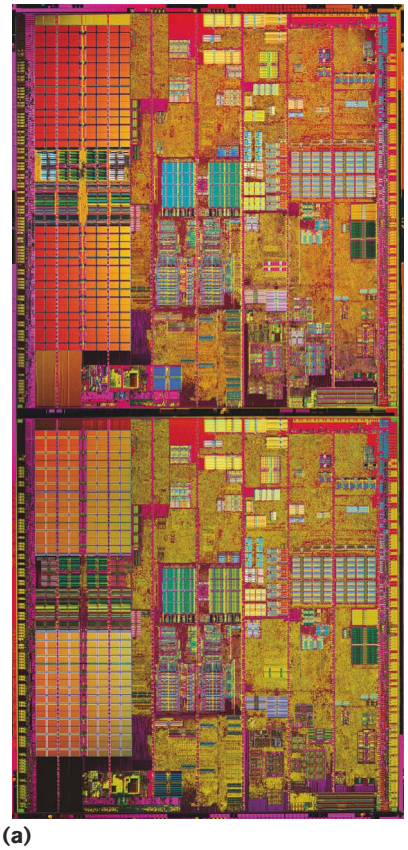

(a)

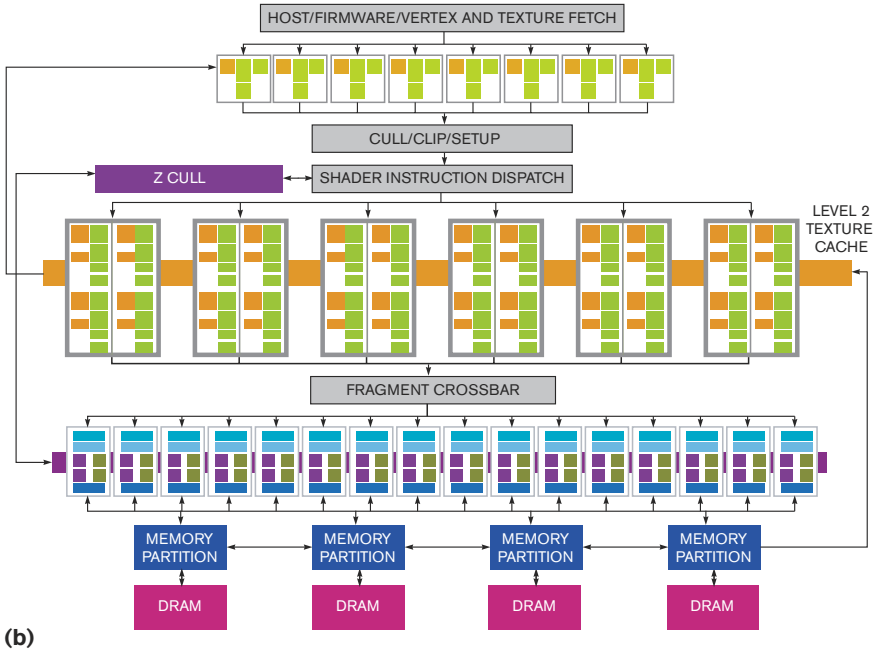

(b)

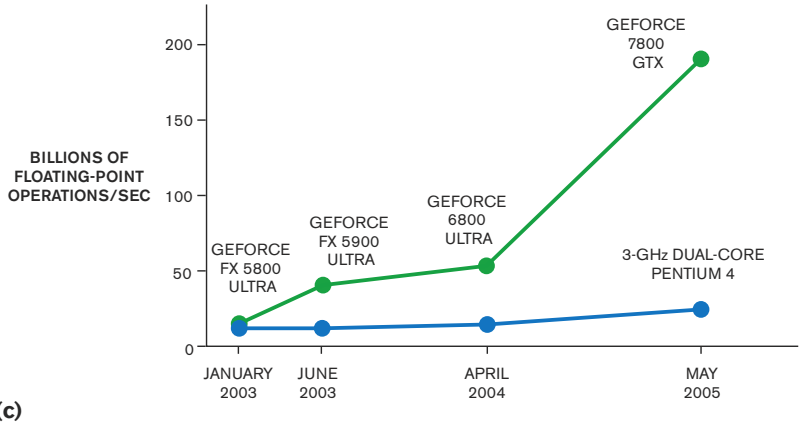

(c)

Figure 2 Single-core-to-multicore-CPU evolutions (a, courtesy Intel), paced by graphics-chip transformations from hard-wired state machines to software-programmable processors (b, courtesy Nvidia) bring the formerly symbiotic partners into more direct competition. Because a GPU need not to devote a high percentage of its die area to cache and control circuitry, it's more computationally efficient than its CPU counterpart (c, courtesy Nvidia).

Making**Wireless**—
Making**3G**

Thanks to TI 3G technology, you helped Susan conduct a meeting in her bathrobe.

TI Wireless Technology for Videoconferencing – Delivering All the Promise of 3G

With industry-leading 3G technology from TI, the wireless possibilities are endless. No wonder six of the seven leading 3G handset manufacturers use our technology in over 60 3G handsets on the market today. Open and flexible hardware and software solutions from TI give you leading-edge wireless technology and the multimedia performance you need to deliver movies, live digital TV, videoconferencing, gaming and even the purchase of a favorite song on a handset. **Put TI at the heart of your next 3G handsets. For more information on TI's 3G OMAP™ and OMAP-Vox™ solutions, visit www.ti.com/3g**

Technology for Innovators™

Texas Instruments

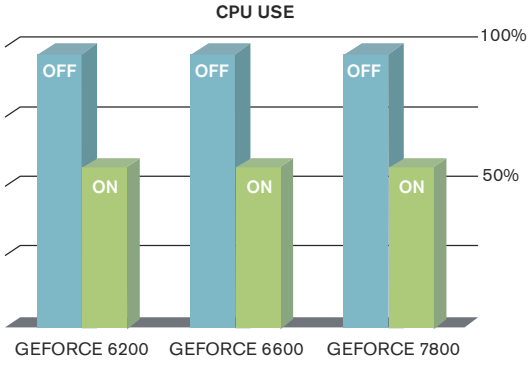

Figure 3 Latest generation GPUs have broadened their video-decoding support beyond MPEG-2, and harnessing their hardware-acceleration capabilities leads to dramatically decreased CPU loading (courtesy Nvidia).

ogy, VC-1. The latest iteration of DirectX VA, with supportive silicon and drivers, dramatically reduces the CPU burden when playing back a high-resolution WMV9 clip (**Figure 3**). The most recent format to receive graphics-vendor attention is MPEG-4 AVC (advanced video coding, or MPEG-4 Part 10, and H.264). ATI's late-September Avivo announcement brands the capability of the company's upcoming Radeon X1300, X1600, and X1800 GPUs to both decode and encode video in a GPU-accelerated manner in MPEG-4 (up to AVC), WMV9, MPEG-2, and DivX formats. Video-encoding acceleration employs the DirectShow API, just as MPEG-2-encoder chips' drivers do. Nvidia, as part of its GeForce 7800 GTX introduction in June, publicly stated that it hoped to have MPEG-4 AVC support in place in time for this year's holiday buying season, in partnership with companies such as CyberLink and InterVideo.

At the Spring 2004 Intel Developer Forum, Pinnacle Systems (now a division of Avid) co-delivered two presentations—one with ATI and the other with ATI and Intel—that highlighted key advantages of PCI Express over AGP in high-definition video-editing applications (**Figure 4**). A typical scenario involves the editing and merging, on the GPU, of multiple-source video streams. (These video streams are both compressed and uncompressed; the uncompressed streams require as much as 250 Mbytes/sec of bus bandwidth per stream.) The video streams reside on the system hard drive, in main memory, and on a connected high-definition video camcorder. The GPU sends back the resultant final product to the CPU for archiving. AGP's limited upstream bandwidth is a bottleneck—one that PCI Express removes—in this final step. GPU-accelerated video encoding is attractive in such a scenario. It's also attractive in PVR (personal-video-recorder) applications and when you're transcoding and streaming video over the LAN (local-area network) or WAN (wide-area network) to a network appliance that doesn't support the source video's attributes. DirectX VA is currently a video-decoding-centric API. However, Microsoft's Blythe says, "There will definitely be future support [for API-enabled video-encoding.]"

THE NEXT FRONTIER

Simplistically speaking, GPUs have been processing images as long as they've been handling MPEG-2 decoding. Proper playback of DVDs requires that the GPU scale the 480-line video (standard or wide-screen aspect ratio) to the resolution and dimensions of the output device. It also requires that the GPU deinterlace the video to match the progressive-scan characteristics of the display. Initially, deinterlacing employed relatively crude "bob" and "weave" algorithms, but, as the capabilities of both CPUs and GPUs have improved over time, the deinterlacing algorithms have become increasingly sophisticated (**Reference 9**). With the latest generation GeForce 7800 GTX GPU and its shaders' performance potential, for example, Nvidia now claims to tackle spatial-temporal deinterlacing of high-definition MPEG-2 sources, such as high-definition video and HDTV. Both ATI's and Nvidia's product literature mentions other image-processing functions the chips can perform, such as postdecoding deblocking and suppression of other lossy-compression artifacts, which are particularly attractive with low-bit-rate streaming video; random broadcast-noise removal; color intensification; and more.

Creative Labs, in May 2003, introduced a visionary graphics card based on subsidiary 3Dlabs' VP500SE GPU: the Graphics Blaster Picture Perfect, a card that—as is usually the case with visionary products—hasn't been particularly successful in the marketplace. According to the vendor, the card includes a suite of software from ArcSoft. The suite includes PhotoImpression, which offers a variety of advanced image-editing features and special-effects filters running directly on the Graphics Blaster Picture Perfect. It also includes Panorama Maker, which combines horizontal, vertical, or tiled sets of images to create panoramic photos. Panorama Maker uses the VPU (Virtual Processing Unit) on the Graphics Blaster Picture Perfect to increase the stitching speed. Finally, the suite incorporates Pho-

1. HD CAMERA DOWNLOADS VIDEO TO SYSTEM MEMORY, THEN WRITES TO THE HARD DRIVE.

2. PROCESSOR READS VIDEO STREAM FROM SYSTEM MEMORY, DECODES IT, AND WRITES IT TO VIDEO MEMORY.

3. VIDEO CARD DISPLAYS VIDEO.

4. SYSTEM DECODES MULTIPLE VIDEO STREAMS RESIDING ON THE HARD DRIVE AND SENDS THEM TO THE DISPLAY.

5. GRAPHICS CARD EDITS AND SENDS EDITS BACK TO SYSTEM MEMORY.

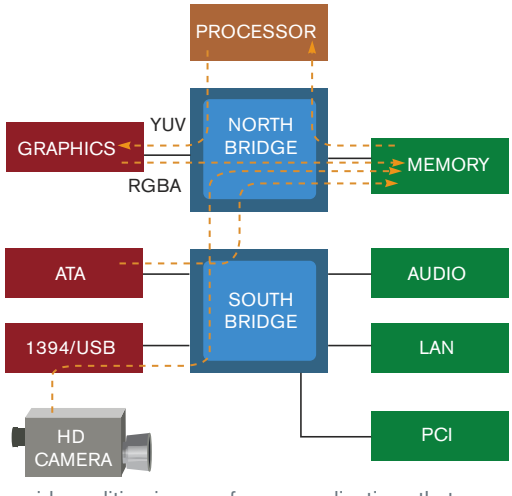

Figure 4 High-resolution, multistream video editing is one of many applications that showcases the advantages of the high upstream bandwidth of PCI Express versus the AGP predecessor (courtesy Pinnacle Systems, a division of Avid).

SEARCH
FIND
BUY

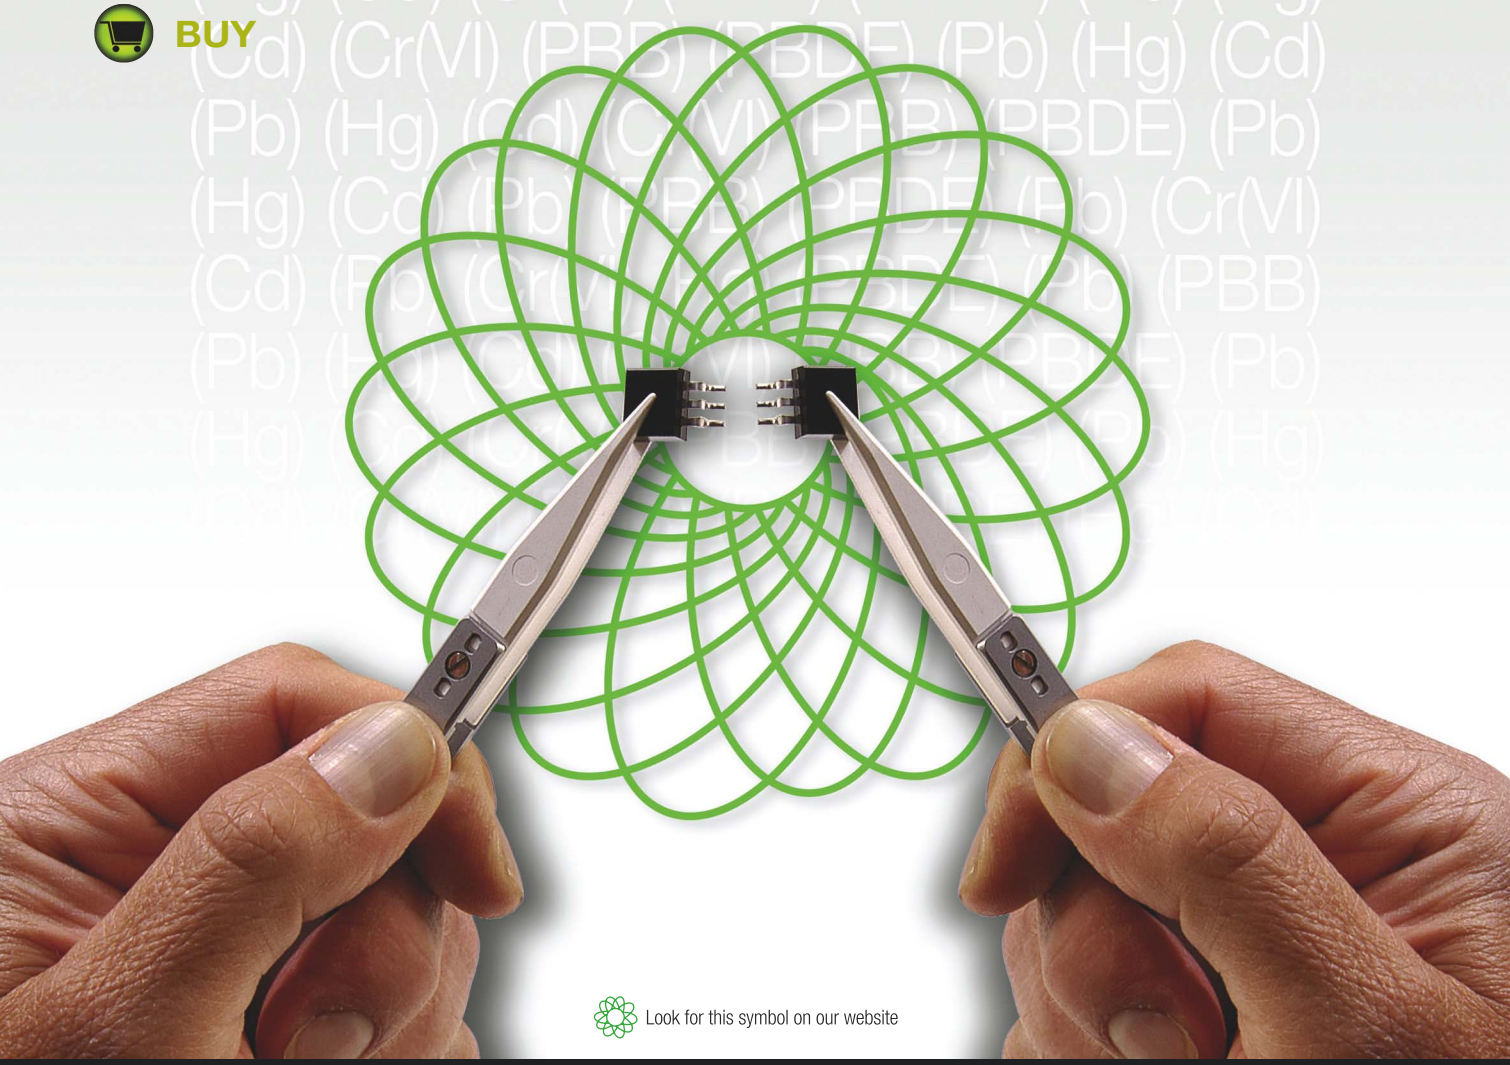

Look for this symbol on our website

Which component is RoHS compliant?

RoHS COMPLIANT

ELECTRICAL & CIRCUIT PROTECTION
INTERCONNECT PRODUCTS/ WIRE & CABLE
LAMPS & DISPLAYS
OPTOELECTRONICS
PASSIVES
POWER PRODUCTS
SEMICONDUCTORS & PROTOTYPING
TEST & MEASUREMENT
TOOLS & PRODUCTION SUPPLIES

Supplier Authorized Distributor

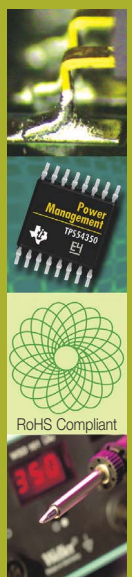

Your #1 source for compliant products makes it easy.

- Each compliant product on our website is flagged with our RoHS symbol.
- Our part number changes for any product that becomes compliant.
- We are continuously adding compliant products online; visit often.
- Our new detailed product ID label includes compliance information.

Find RoHS compliant products at www.newarkinone.com
Visit www.newarkinone.com/rohs for exclusive RoHS information.

NEWARK inone℠

(a)

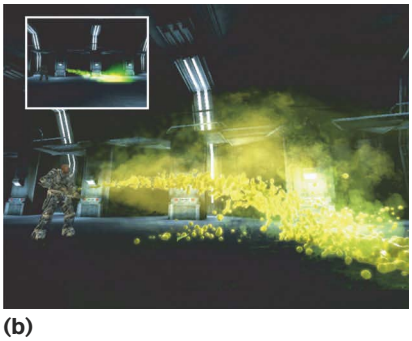

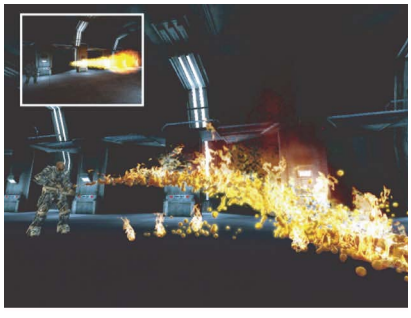

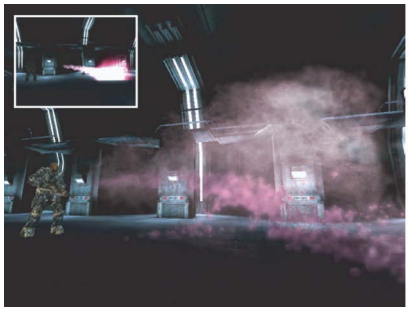

(b)

Figure 5 Although video-clip-based short cuts may still deliver convincing results (a, courtesy Nvidia), physics-processor advocates claim that a full-blown particle-based approach (b, courtesy Ageia) is the optimum option from realism and flexibility standpoints. (Inset graphics show the less lifelike short-cut results.)

toPrinter, which allows multiple image printing on one page and multiple-page printing at one time (**Reference 10**).

The Graphics Blaster Picture Perfect lacks the support of a leading image-editing program, such as Adobe Photoshop. More generally, only the customized ArcSoft programs can access Graphics Blaster's acceleration features. Broader industry adoption requires API support within the Windows operating system, which Microsoft has yet to deliver. With regard to API support for still- and video-image-editing acceleration on the GPU,

Blythe says, "Those are big, big, interesting things to us. And they seem like pretty natural things. Once you've got a programmable pixel-processing pipeline, video and still images are just blocks of pixels, as well, so why not? The way that I think of those is that they're just another stage, at least in the video-decode path. After I produce these images that have been deblocked, scaled, color-corrected, and all that kind of stuff ... well, I can do more processing on them, and it's just a buffer of pixels that, for all practical purposes, you can think of as being like a texture map for 3-D graphics. And you can do an arbitrary amount of pixel-shader processing on it."

Unlike Apple, which provides the Core Image library, Microsoft so far has no comparable offering, according to Blythe. "We've been focusing on the high-level programming-language capability to be able to express these kinds of operations in the programmable shading languages, without providing any sort of packaged set of operations. But we're definitely looking at those sorts of technologies. And we see this as an opportunity to broaden the use of the GPU for other kinds of applications beyond video decode and playing games." As Blythe's comments point out, although Apple may lack Microsoft's third-party-accessible API for video decoding on GPUs, Apple *has* rolled out an image-editing API, Core Image, in its latest OS 10.4 Tiger operating-system release, which it began shipping in late April. Approximately 30 applications already support Core Image, according to Wiley Hodges, Apple's senior product-line manager for developer products. Apple also offers the Core Video API. "In Tiger, Core Video is basically Core Image applied to a sequence. So, there's not really a fundamental difference," explains Hodges.

According to the Core Image portion of Apple's Web site, the company bundles 100 Core Image units with Tiger, along with Core Image Funhouse, a front-end demonstration program. Hardware acceleration requires a pixel-shader-based GPU, but the Web site explains that for computers without a programmable GPU, Core Image dynamically optimizes for the CPU, automatically tuning for Velocity Engine—that is, AltiVec support—and multiple processors as appropriate (**Reference 11**). Says Hodges, "For a developer, Core Image allows someone who might not be particularly expert in the field of image processing to add that capability into an application that might otherwise be too expensive or difficult to take advantage of it. For people who do know what they're doing, it's a great way to provide access to a standard set of effects and filters." According to Hodges, effects are written in a high-level language that's based on a subset of the OpenGL shader language, and Apple will provide just-in-time compilation of the effects for the target hardware on which an application is executing.

"Effects can run in the CPU or the GPU," Hodges says, "and one of the interesting things is that, while it's true that GPUs have enormous amounts of paral-

lel capability, they are not always good at every possible kind of effect or transformation. And so we can look at what something requires and execute it where it's going to most optimally execute." When doing Core Image compilation, does Apple take into account 24-bit floating-point support in the ATI GPUs now in Macs versus the 32-bit floating-point capabilities in the last three generations of Nvidia chips (5xxx, 6xxx, and 7xxx)? "We do [take them into account], but there's no effect on the ultimate end result," says Hodges. "It means that, for the same level of quality, you might not get the same performance on ATI as on Nvidia. We definitely see better performance of complex Core Image effects on the high-end Nvidia cards right now."

Even in the absence of Windows API support, several video-software developers have concluded that the benefits of GPU acceleration are compelling enough that they've taken the early-adopter plunge and implemented some degree of support. According to Giles Baker, senior product manager at Adobe, "Taking advantage of the increasing power of GPU processors is a key area of focus for Adobe's image-editing and video products, both now and in the future." He claims that, as the power of GPUs continues to increase more rapidly than host CPUs, more and more of Adobe's users are simply upgrading their graphics cards to increase the overall performance of their systems. Doing so provides a more cost-effective way of keeping systems up to date.

"In Adobe Premiere Pro [and Premiere Elements], we use the GPU directly to provide a number of accelerated effects that are available only if you have a capable graphics card," says Baker. Fortunately, he adds, most cards available today support these effects, so most people benefit from the features. "As HD production becomes more and more widespread, the GPU is a key technology that we can use to help move the enormous amount of data that is needed to create HD-resolution video content. As operating systems evolve, they provide lower level access to the GPU, which we can leverage to achieve levels of performance that were previously available only when using dedicated hardware," he says. With regard to GPU support in Photoshop, Baker comments that none of the plug-ins that ship with Photoshop is currently GPU-accelerated. However, he adds, "In the future, we see a real opportunity to take advantage of GPU acceleration in all our products." Feel free to read between the lines.

Video editing doesn't just involve video, of course. Videographers also often add 2-D- and 3-D-graphics-based scene transitions, text, and other "eye candy" to a clip before burning it to DVD or another archival-and-playback format. Because the graphics subsystem finds use in displaying graphics information on the computer monitor, it's logical to employ that same GPU to speed the rendering of graphical effects. These effects merge with the video in system RAM and on the hard-disk drive, so they benefit from the high upstream bandwidth that PCI Express delivers. According to Adobe's Baker, Adobe After Effects 6.5 uses OpenGL to accelerate on-screen playback of motion-graphics projects during production. "By offloading some of the processing directly to the GPU, we gain more responsive performance with less waiting. This leads to a more creative design experience that encourages experimentation," he says.

Baker continues, "Since Adobe products are all about creating content productively, this is a huge advantage for our users. OpenGL is an evolving standard, and we plan to take advantage of new capabilities in OpenGL as the technology develops." The Mac OS counterpart of After Effects, Apple's Motion, also harnesses the GPU by means of OpenGL. According to Apple's Web site, Motion is the first motion-graphics software with GPU-accelerated, 32-bit floating-point rendering for true film quality. The documentation further says that this 32-bit floating-point rendering produces fine color accuracy, eliminates banding artifacts, and even improves quality when rendering to 8-bit formats. Finally, it claims that you get great detail, quality, and range of color that automatically scales with new generations of GPUs (**Reference 12**). Hodges points out and the Apple Web site confirms that installing Motion 2, the latest version of the software, *requires* a shader-based GPU.

GAMES BEYOND GRAPHICS

Beyond rendering and image processing, what's the next application frontier in

Design to a higher standard with nonvolatile FRAM memory

Nonvolatile FRAM Memory

- Industry standard interfaces and packages
- High endurance
- NoDelay™ reads and writes
- Low power consumption
- Industrial temperature range
- Densities from 4- to 256-Kbit

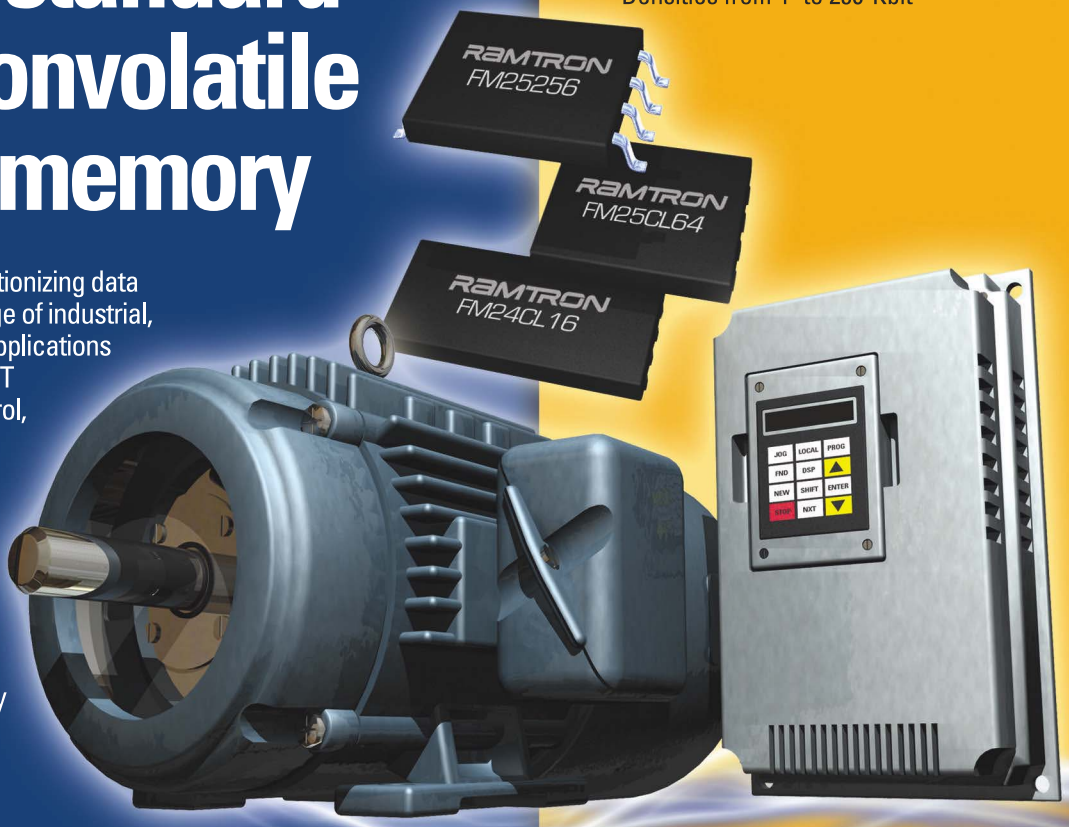

FRAM memory is revolutionizing data collection in a wide range of industrial, scientific and medical applications such as motor control, CT scanners, elevator control, and health monitors.

With virtually unlimited write endurance, nanosecond-class read/write speeds and industrial temperature range operation, it's no surprise that the industry is embracing FRAM memory.

It's not just industrial applications that are leveraging FRAM features. From LCD TVs and automobiles to embedded systems and office equipment, FRAM is proving that it is the best choice in high-performance nonvolatile memory.

Find out how FRAM memory can raise the standard in your next design by calling a Ramtron technical engineer at

1-800-545-3726

or simply email framinfo@ramtron.com.

RAMTRON
International Corporation

www.ramtron.com

When Metro Networks demand bandwidth…

PMC delivers Direct.

www.pmc-sierra.com/metro

To get to market reliably, effectively and on time with equipment that meets the demands of today's bandwidth hungry networks, PMC's metro networking solutions deliver what it takes to get you there direct, from concept to market, and meet the challenge of emerging Triple Play, VoIP, and Wireless services.

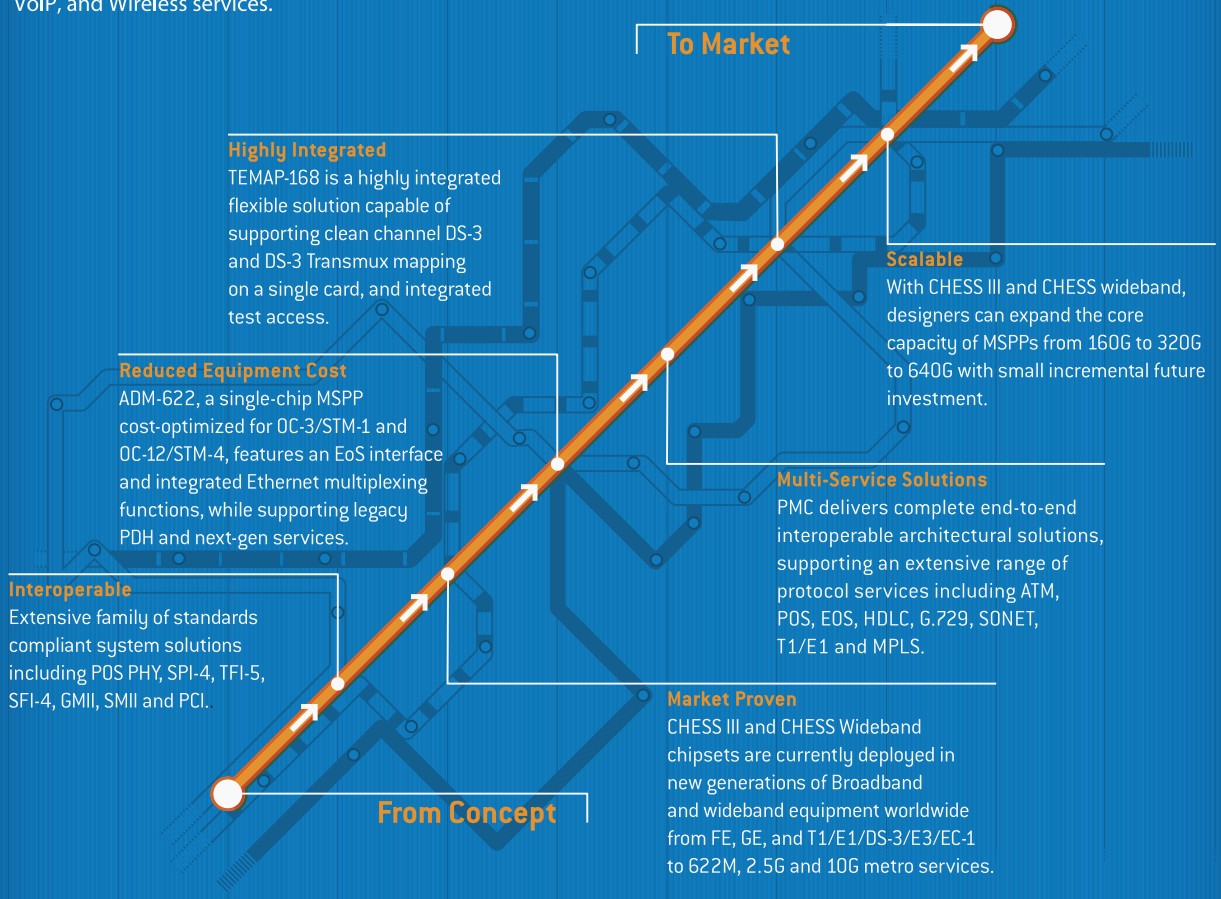

Highly Integrated
TEMAP-168 is a highly integrated flexible solution capable of supporting clean channel DS-3 and DS-3 Transmux mapping on a single card, and integrated test access.

Reduced Equipment Cost
ADM-622, a single-chip MSPP cost-optimized for OC-3/STM-1 and OC-12/STM-4, features an EoS interface and integrated Ethernet multiplexing functions, while supporting legacy PDH and next-gen services.

Interoperable
Extensive family of standards compliant system solutions including POS PHY, SPI-4, TFI-5, SFI-4, GMII, SMII and PCI.

Scalable
With CHESS III and CHESS wideband, designers can expand the core capacity of MSPPs from 160G to 320G to 640G with small incremental future investment.

Multi-Service Solutions
PMC delivers complete end-to-end interoperable architectural solutions, supporting an extensive range of protocol services including ATM, POS, EOS, HDLC, G.729, SONET, T1/E1 and MPLS.

Market Proven
CHESS III and CHESS Wideband chipsets are currently deployed in new generations of Broadband and wideband equipment worldwide from FE, GE, and T1/E1/DS-3/E3/EC-1 to 622M, 2.5G and 10G metro services.

Visit www.pmc-sierra.com/metro to access a wide range of technical documents, including reference design collateral, white papers and webinar materials.

PMC
PMC-SIERRA
Thinking you can build on.

which the GPU might be able to wrest control away from the CPU, or where might yet another processor with tailored functions emerge? Blythe's Meltdown presentation points developers in three GPU-leveraging directions: skinning, noting that simple 1-to-4-bone, GPU-based skinning is already pervasive; morphing, such as GPU-calculated facial animation; and simulation of particle systems and fluids. Echoing comments Kirk made in his Hot Chips keynote speech, Blythe says, "The place where we're going, again staying close to the games side of things, is that we've got this way of putting pixels on the screen, and, over time, we got to a point where I could get enough pixels on the screen. So I then wanted to start increasing the quality of the pixels." As a result, Blythe asserts, Microsoft's next areas of focus were finding ways, for example, to better shade the pixels, have better lighting models, or do better texture-mapping operations. Now, the company is also starting to work toward getting better geometric models for characters; better animations; and better special effects, such as water, fire, and foliage. "Fundamentally, what I'm trying to do is get better looking visuals on the screen. And those are the places where the 'big bang,' or the most cost-effective improvements are going to happen," he says.

Microsoft is also looking at options for increasing the quality of the geometric complexity of characters so that they, for example, have better silhouette edges or cast better shadows. After image-quality improvements, Blythe claims that the next set of factors that are going to contribute to the quality of the end-user experience encompasses the ability to do better kinds of physical simulations, such as creating better looking fluids, and the ability to do destructible environments in which there's a true level of realism to how the objects interact. These tasks require untraditional graphics computations, such as solving both ordinary and partial differential equations and linear systems. Thus, says Blythe, it makes sense for Microsoft to look at where best to solve these kinds of problems. "If I can express them as data-parallel kinds of problems, then the GPU starts to look interesting as a place to do them. If they have more of this sort of squirrelly control-level parallelism to them, then the CPU might be the most appropriate place. And then there's the physics processor. I can't say much about it, because there's a question of how close is it to a CPU and how close is it to a GPU in terms of the kinds of processing," says Blythe.

The *GPU Gems* book series is chockfull of ideas for simplifying the rendering, animation, and elemental interaction of objects. Such objects might include trees and their branches and leaves; windblown blades of grass in a field; a head of flowing hair; water and other fluids; fire; and atmospheric effects, such as fog and smoke. The word "simplifying" is critical here; rendering smoke by separately generating each particle and calculating its interactions with all other present particles, for example, is prohibitively expensive even for today's leading-edge CPUs and GPUs. So, visually comparable, but arithmetically easier, approaches are necessary. Chapter 6 of *GPU Gems*, for example, discusses a fire effect that the Nvidia team developed for a demo called Vulcan at the launch of the GeForce FX 5900 Ultra GPU (**Figure 5**). "When we started working on the demo, we first tried two solutions that looked promising: fully procedural flames and screen-space, 2-D, distortion-based flames," writes chapter author Hubert Nguyen.

Nguyen states that the fully procedural approach consumed little memory yet created an appealing flame effect. However, to produce well-defined flames, the

MORE AT EDN.COM

- Go to www.edn.com/051013cs to find this article's associated vendor box. At this URL, you can also post a comment on this article.
- At the Brian's Brain blog, www.edn.com/briansbrain, learn more about the topics discussed here, post comments and questions, and peruse and respond to the postings that other *EDN* readers make. Visit the "Instigating a platform tug of war" entries to check out the following postings:
  - Chips for chips?
  - General-purpose GPUs: going beyond graphics
  - Interesting lit and other bits
  - Non-PC opportunities
  - What about audio?

BigaHurts

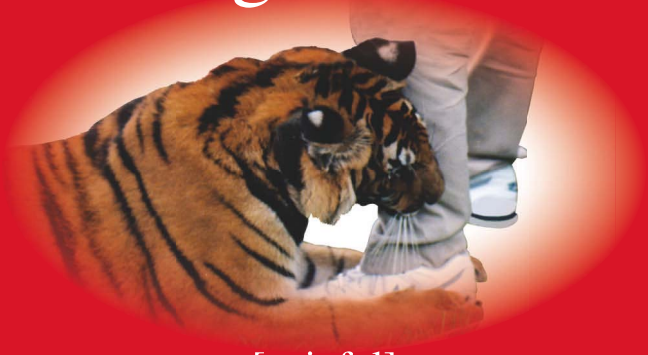

[painful]

GigaHertz

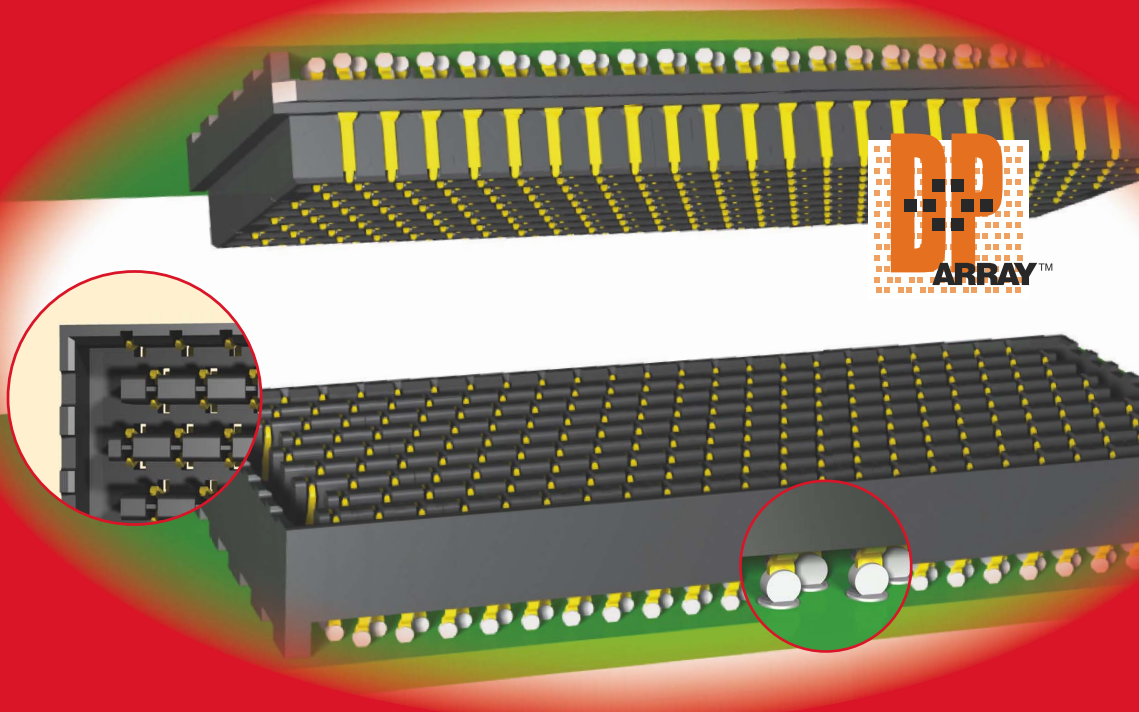

[easy]

[easy] to maximize bandwidth
with over a terabit per connector (8 Gbps per pair x 168 usable pairs).

[easy] to minimize board space
with staggered pin layout that reduces grounds required.

[easy] to route signals
with Final Inch® PCB design tools that save design, development and validation time.

[easy] to use
with solder crimped tails, low insertion/extraction forces.

[easy] to find out more
about high speed, high density interconnects at **www.dparray.com**
or **www.samtec.com/easy**

www.samtec.com
SUDDEN SERVICE

demo had to display thousands of particles, and processing all of those vertices and pixels put a heavy load on the CPU and the GPU. The 2-D distortion-based flames used a GPU-generated perturbation function that altered a flame shape to give it a realistic motion. The distortion involved making several render-to-texture passes and shifting 2-D texture coordinates. Although it consumed more memory than the particle-system technique—because the demo had to allocate several render targets—the chapter claims that the effect is perfect for creating candlelike flames. However, the screen-aligned nature of the effect made it sensitive to motion in general and to the camera viewing angle. (Top and bottom views required constraints, and moving toward and away from the camera sometimes didn't work well in 2-D.) Integrating smoke was also a challenge, according to the book.

"Both procedural techniques have strong advantages, but they didn't meet our goal of creating a believable raging fire with smoke in a real-time, user-controllable environment," writes Nguyen. Thus, the team turned to video-textured sprites—that is, video-based footage—to make the fire more realistic. "Although full procedural and physically based flame generation is clearly the wave of the future, some cutting-edge movies, such as *The Lord of the Rings*, still use special effects composed of sprite-based flames" (**Reference 13**). In reference to Nguyen's wave-of-the-future comment, CPUs and GPUs will inevitably evolve to better manage the complex processing necessary for full-blown, real-time particle animation as well as for other computationally intensive operations, such as real-time ray-tracing-based lighting. For example, Microsoft revealed during its Hot Chips 2005 presentation that it had added extensions to the variant of DirectX 9 running on the upcoming Xbox 360 game console to support the execution of particle-physics calculations on the system's ATI-designed GPU.

However, officials at companies such as Ageia, with its upcoming PhysX processor, believe that a compelling need will still exist for a dedicated physics processor. One example would be a game character that needs not just to be present alongside, but also to interact with and deform a particle-based object—for example, when walking through a fog bank. "While dual-core machines can demonstrate physics effects impossible on single-CPU systems, the PhysX processor brings a realism and quality of effect impossible in software alone," says Suneil Mishra, Ageia's director of software-product marketing. "While dual-core processors can potentially handle hundreds of real-time objects interacting as opposed to dozens with a single-core CPU, Ageia's PhysX processor offers tens of thousands of fluid particles and rigid-body objects concurrently. The leap in performance and quality enables a completely different level of realism and immersion for gamers, both via effects and game-play physics."

Mishra points out that GPUs have focused recently on mocking physical effects, such as fire, water, hair, or cloth, with clever visual imagery that has limited dynamic interaction or motion. He further points out that, although convincing, these cases are limited, as are other GPU forays into coding small parts of the physics-simulation pipeline. Ageia's physics processor, he claims, handles thousands of real physical objects interacting throughout an environment, allowing gamers to experience true physical reality across and throughout dynamic game levels. Here, things not only look right, but also act and feel right. "While we fully expect GPUs to continue to grow in performance, allowing them to take on more of the most basic physics operations, support for advanced physics effects and interaction will continue to remain beyond GPUs, simply due to their chip architecture," says Mishra. Ageia's PhysX processor aims at accelerating the needs of physics-simulation algorithms. The memory and floating-point-bandwidth requirements alone cause even the most programmable GPU to fail under the load of tens of thousands of interactive physics objects, according to Mishra. EDN

Design with Flexibility

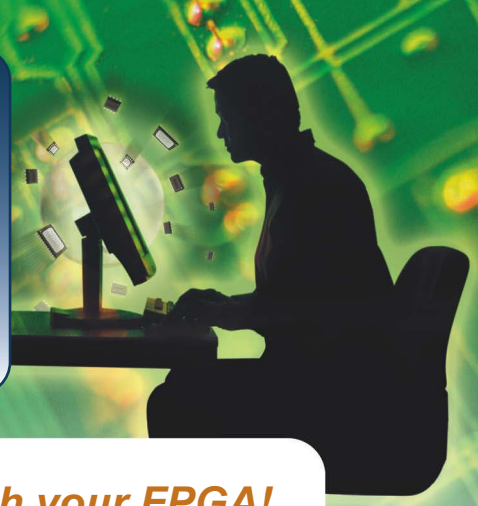

Now your power solution can migrate with your FPGA!

Pin-Compatible 3A, 2A and 1A Integrated Synchronous Bucks

MPS provides worry-free board design and reduces cycle time with the MP2305/7/9 series of pin compatible synchronous step-down regulators. Now, both your custom logic design and power supply have a simple migration path.

3A  ***MP2307***

2A  ***MP2305***

1A  ***MP2309***

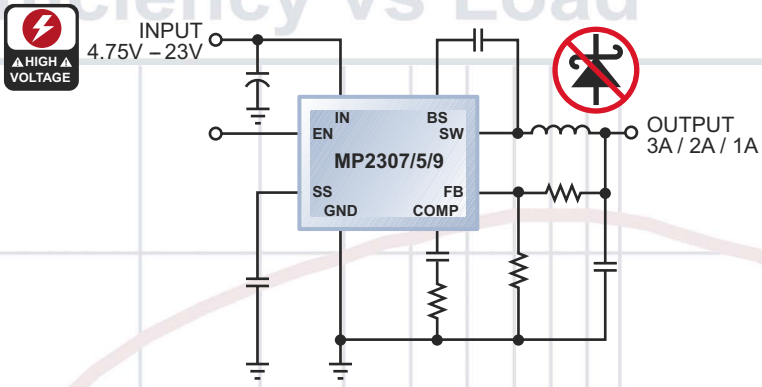

- ***4.75V - 23V Input***
- ***Common PCB Layout***
- ***Over 90% Efficiency***
- ***Output as Low as 0.9V***
- ***Integrated Switches***
- ***All Ceramic Cap Design***
- ***Small SO8 Package***

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

MPS®

The Future of Analog IC Technology®

www.monolithicpower.com

Smart solutions
for comfort and safety

Saia-Burgess manufactures solenoids, motors, switches, electronic controllers and subsystems. Together with our customers we develop value-based solutions for countless industrial applications. Our innovative, experienced design team and TS-16949 registered manufacturing facilities support our commitment to quality in developing world class solenoid, switch and motor products. Call us today.

www.saia-burgess-usa.com/hvac

Saia-Burgess Inc.
Technical Assistance 1-800-429-0365
Commercial Assistance 1-800-998-2298

saia-burgess
Smart solutions for comfort and safety

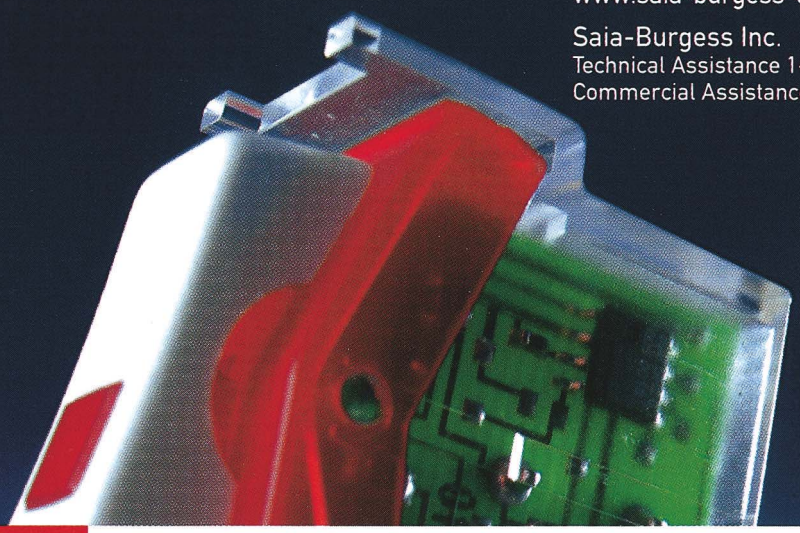

- UAG Stepper Motor
- 3600 Rocker Switch
- Ledex C8 Open Frame Solenoid
- G3 Snap Action Switch
- KOP.K Electronic Timer

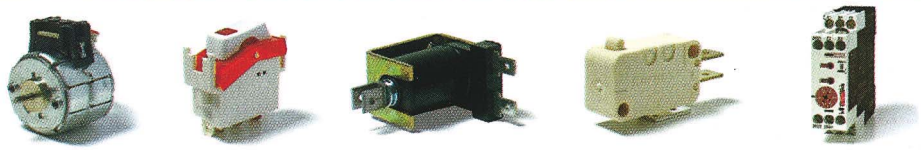

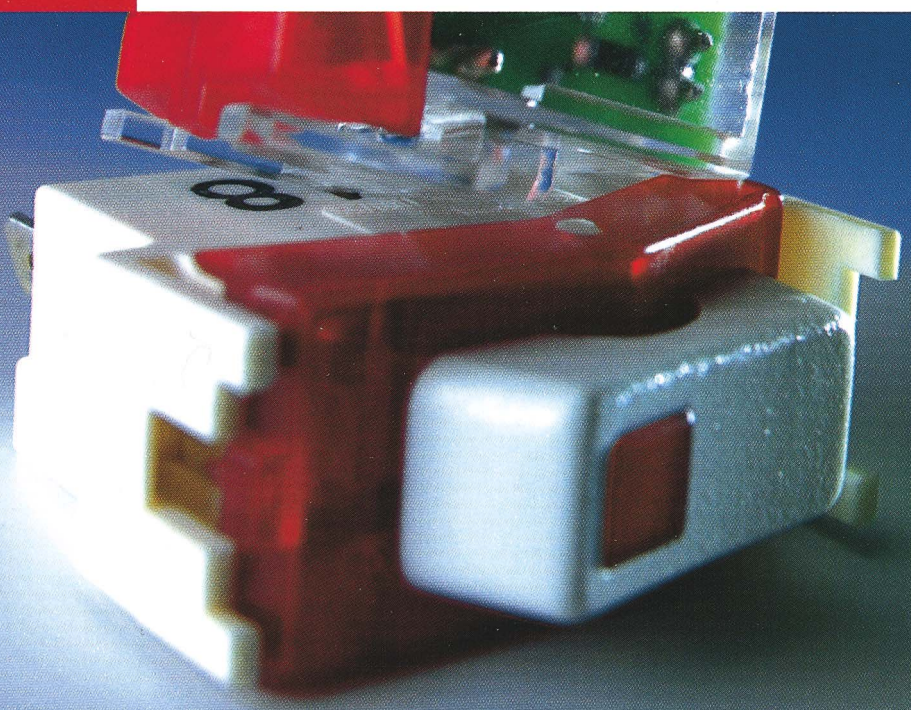

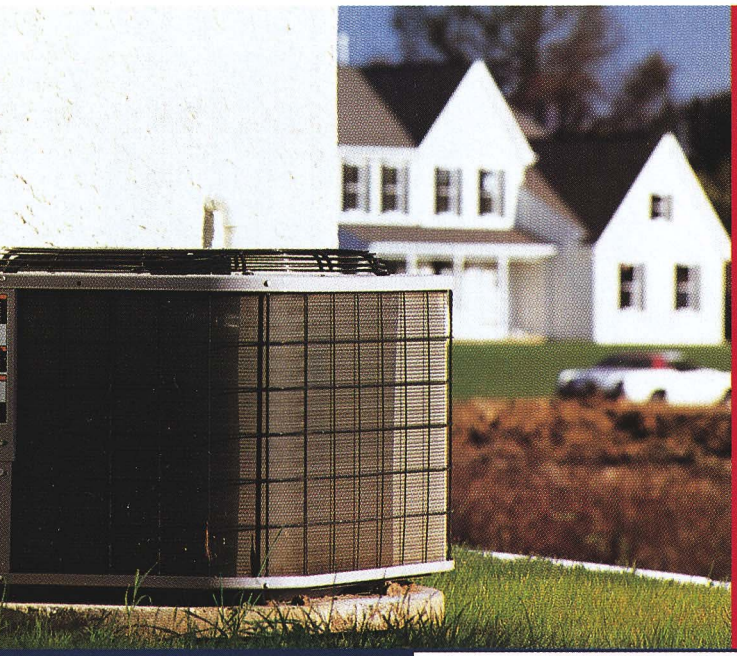

Saia-Burgess products — simply everywhere

Saia-Burgess products are found simply everywhere in the industrial and residential dampers and air management industry. From the largest central HVAC systems to the smallest residential pellet stoves, Saia-Burgess products are used by engineers designing the most reliable and highest quality products for air management applications. Our innovative design team and TS-16949 registered facilities are sure to add value to your state-of-the-art design. Call us to discuss your solenoid, switch and motor requirements.

saia-burgess

Smart solutions for comfort and safety

- UAG Stepper Motor
- STA Tubular Solenoid
- Ledex BTA Solenoid
- G3 Snap Action Switch
- KOP.K Electronic Timer

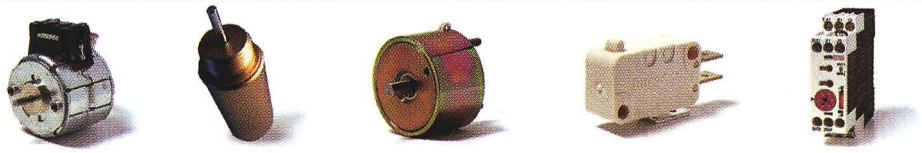

Industrial and Residential

- Dampers
- Regulators
- Ventilators
- Cooling fans
- Valve control
- Air management
- Compressor controls
- Pellet stoves

Looking for more information on Saia-Burgess solenoid, switch and motor products? Visit our website for performance and mechanical specifications on all our products:
www.saia-burgess-usa.com/hvac

Saia-Burgess Inc.
Technical Assistance 1-800-429-0365
Commercial Assistance 1-800-998-2298

Smart solutions for comfort and safety

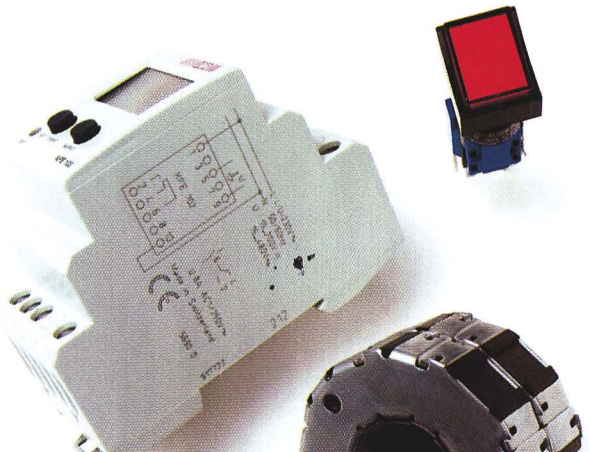

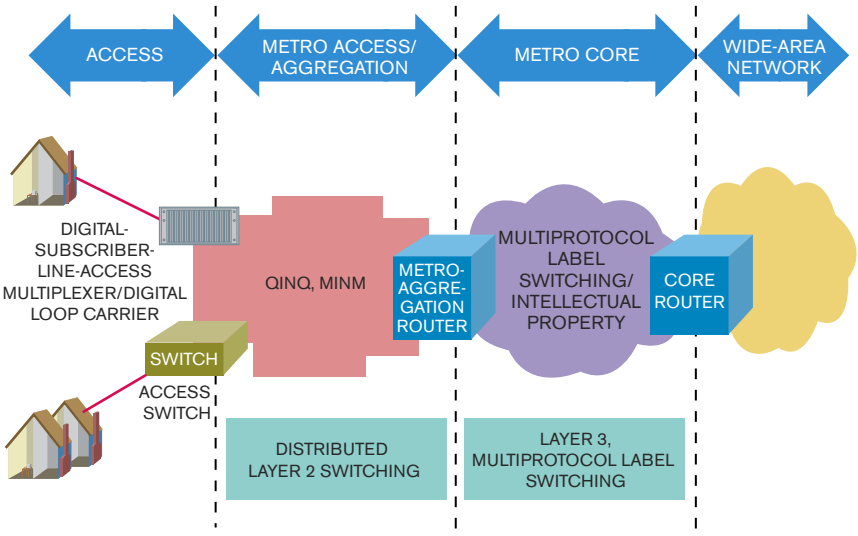

Figure 1 Most current broadband deployments use a two-tier metro-aggregation-network and metro-core-network architecture.

TECH TRENDS BY GOPAL GARG, CYPRESS SEMICONDUCTOR, AND R THIRUMURTHY, MIDAS COMMUNICATION TECHNOLOGIES

Optimized learning in metro switches

ETHERNET EXTENDS INTO AGGREGATION AND CORE NETWORKS.

Ethernet has emerged as an alternative access technology to SONET/SDH (synchronous optical network/synchronous digital hierarchy) for broadband deployments; its low cost and simple provisioning mechanism mainly drive this emergence. Most new broadband deployments across the globe use some form of Ethernet for their access, aggregation, and core networks. Next-generation SONET/SDH and RPR (resilient-packet ring, IEEE 802.17) have given greater impetus to the deployment of Ethernet services and extended Ethernet technology into aggregation and core networks.

Most current broadband deployments use a two-tier metro-aggregation-network and metro-core-network architecture (**Figure 1**). In the metro-aggregation-network portion, Layer 2 Ethernet switching aggregates traffic from access devices, such as DSLAM (digital-subscriber-line-access multiplexer), MSPP (multiservice provisioning platform), and access switches. This Layer 2 traffic terminates in the metro aggregation router. In some cases, the BRAS (broadband remote-access server) collocates with (or is built into) the metro-aggregation router. The BRAS terminates subscriber traffic, such as PPPoX (Point-to-Point Protocol over Ethernet/ATM), VLAN (virtual local-area network), or other traffic, and applies policing and QOS (quality of service) per the subscriber profile.

The metro-aggregation router is the point at which normal residential Layer 2 traffic terminates and premium services, such as Layer 3 VPN (virtual private network), VPLS (virtual private-LAN service), EPLAN (Ethernet private LAN), and EPL (Ethernet private line), map to the metro-core layer as MPLS LSP (multiprotocol label-switching label-switched path), Layer 2/3 VPN, for example.

SCALABILITY

The aforementioned network architecture relies heavily on Ethernet-switching technology in the metro-access/aggregation network. However, Ethernet is a technology targeting the enterprise environment, with scalability, resilience, and other key network characteristics tuned for enterprise needs rather than for more demanding service-provider requirements. Over the past few years, the industry has introduced new protocols and management features to enable Ethernet to perform according to telephone-company standards. With respect to the scalability of metro Ethernet switches, two major issues developers must address are VLAN limitations and FDB (filtering-database) table size.

VLAN

One advantage of employing Ethernet in the enterprise domain is the ability to logically partition distinct user groups over the same physical network through a VLAN. This capability of Ethernet extends into the metro domain with user groups becoming individual subscribers or companies. However, the IEEE 802.1Q standard defines an address space of only 4096 available tags. With companies of-

AT A GLANCE

- Ethernet targets the enterprise environment and metro-area switches.
- New protocols and features enable Ethernet to comply with telephone-company standards.
- Major issues are virtual-local-area-network limits and filtering-database-table size.
- Adaptive, optimized-learning schemes can enhance performance.

fering subscribers multiple services, having only 4096 VLANs becomes a serious limitation. Providers can address this VLAN-scalability problem by increasing the VLAN available space through double-tagging (the IEEE 802.1ad draft, or QinQ) or VLAN stacking. This scheme essentially tags the packet with an outer VLAN tag, thereby expanding the address space to 4096×4096 unique subscriber/service-identification tags.

MAC-ADDRESS TABLE

The other scalability problem is the size of the MAC (medium-access-control)-address-learning table in switches. Because traffic in the aggregation network is Layer 2-switched, the Ethernet switches need to have a large MAC-address-learning (FDB) table that is proportional to the number of subscribers connected throughout the Layer 2 network. Proprietary methods, such as MinM (MAC in MAC), which encapsulate the subscriber Ethernet packet within the Ethernet header of the switch (using the switch's MAC address as the source), can address issues resulting from the table's large size. With MinM, intermediate metro switches need to learn only the switch addresses and not the actual subscriber MAC address. Only the switch that encapsulates the Ethernet packet needs to learn the MAC address of the subscribers that directly connect on its access links.

OPTIMIZED LEARNING SCHEME

Most of the traffic-flow patterns in the access/aggregation network are P2P (point-to-point) networks, such as networks in which the subscriber's Internet traffic terminates in the BRAS, the ASP (application-service-provider) server, or another area, or P2M (point-to-multipoint) networks, such as networks for multicast-video service, VPLS, and other services. In both cases, if you use QinQ, then the inside VLAN identifier identifies and classifies in the end nodes and is unseen by the intermediate switching nodes. Therefore, this discussion considers only the outside VLAN ID.

Consider a P2P case in which the traffic flows from the source to the edge switch/router and vice versa. This flow uniquely identifies with a VLAN ID. If the source node and the destination node establish a path, then in each intermediate switch, the path enters through a particular interface and exits through another specific interface. In other words, packets entering through one interface of the path will have exactly one interface through which they exit. This scenario can avoid the MAC-address learning for that flow in the transit switches.

In **Figure 2**, six switches (N1 through N6) connect in a partial mesh topology. The bold lines indicate the paths that loop-detection mechanisms, such as STP/RSTP (Spanning Tree Protocol IEEE 802.1d/Rapid STP, IEEE 802.1w), select for forwarding traffic. The P2P flow between nodes A1 and A2 passes through N1, N2, and N3. In this case, only nodes N1 and N3 need to learn the MAC addresses within the flow, and, if the ingress and egress interface of the flow is known, node N2 can pass through this flow without looking up the address.

Similarly, in the case of a multipoint flow (typically a VPLS), only certain nodes need to actually learn the MAC address; other nodes in the path can just forward the traffic. This situation is possible if you can create a tree between the nodes for each multipoint flow (identified with the VLAN ID). As **Figure 3** shows, assume that B1, B2, and B3 belong to a VPLS; the red lines in the **figure** indicate the subtree connection between the nodes for this flow. In this case, only the switch nodes N1, N3, and N5 need to learn the MAC address for this multipoint flow. Nodes N2 and N4 can forward the traffic without learning the MAC address.

To summarize, only endpoints (for example, source and edge switches) and certain intermediate nodes should handle MAC-address learning for a P2P or P2M flow. All other intermediate nodes can blindly forward the packet from one interface to another without having to learn or look up the destination MAC address. The above optimization is possible if you can establish a path/tree for each flow (VLAN ID) in the intermediate switch.

With the aforementioned scheme, the switching nodes in the network need to learn the MAC addresses only for certain flows. This stipulation significantly reduces the FDB-table size for each node, thereby enhancing the scalability of Ethernet switching in the access/aggregation network.

The main challenges in implementing this scheme include:

- deriving a path/tree for P2P flows, multipoint flows (QinQ ID), or both; and
- identifying the flows each node must learn.

Establishing a path or tree for each flow requires a mechanism that selects the membership ports for each flow in the node. Because you use the VLAN for flow identification, you can use the GVRP (GARP VLAN Registration Protocol) for deriving the port membership. (GARP stands for Generic Attribute Registration Protocol.) IEEE 802.1p and 802.1Q define GVRP as providing a mechanism for dynamic maintenance of VLAN-registration entries for each VLAN and for propagating the information to other nodes. This information allows nodes to dynamically establish and update their knowledge of the set of VLANs that are active and the ports through which you can reach them.

Considering **Figure 2**, to provide P2P connectivity between A1 and A2, the static VLAN (say, V) entry is configured in port 2 of node N3 for user A2. Then:

- The GVRP registers the VLAN on

Ready!

Design with Quartus® II development tools.

QUARTUS® II

Set!

Prototype with Stratix® II FPGAs.

Go to Market!

Ramp to production with Stratix II FPGAs or migrate seamlessly to HardCopy® II structured ASICs.

The quickest & easiest path to production.

ALTERA®

The Programmable Solutions Company®

www.altera.com

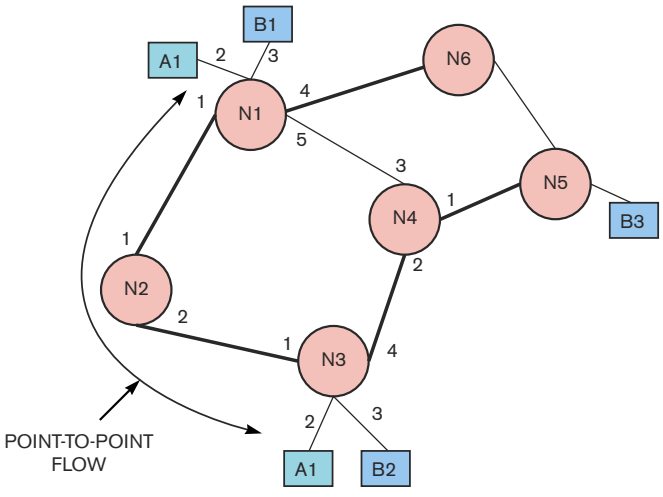

Figure 2 Most of the traffic-flow patterns in the access/aggregation network are P2P (point-to-point) networks.

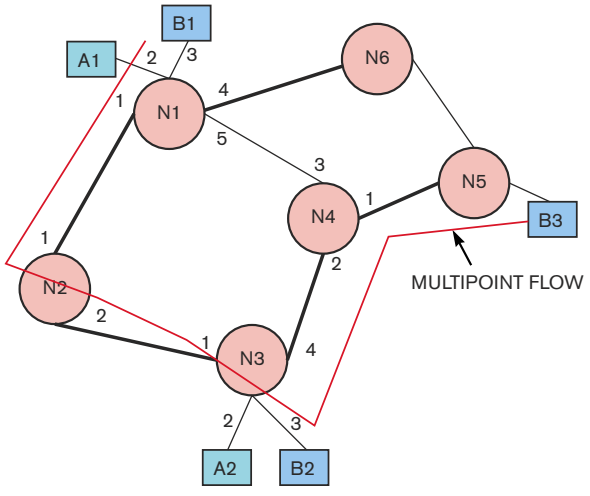

Figure 3 In a multipoint flow (typically a virtual private-LAN service), only certain nodes need to actually learn the MAC address.

the port connected to A2 and propagates it on all other ports.

- Node N2 receives the registration and registers the VLAN on the port (port 2) received, then propagates the registration on port 1.
- N1 receives the registration in port 1 and includes port 1 as a member of VLAN V, then propagates the registration on all other ports.

In the above sequence, port 2 of N2 and port 1 of N1 have become members of the VLAN V (flow).

Because A1 belongs to the same flow, the same VLAN entry will be statically configured in N1 on port 2. Hence:

- N1 propagates the VLAN membership on ports 1, 3, and 4.
- N2 receives the registration message on port 1 and forwards it to port 2.
- N3 receives the registration on port 1 and includes port 1 as a member of VLAN V.

This sequence includes ports 1 and 2 of N1, ports 1 and 2 of N2, and ports 1 and 2 of N3 as members of VLAN V, creating a bidirectional path. You can easily extend this procedure to create a tree for a multipoint flow.

The GVRP runs within the context of the spanning tree created through RSTP/STP. This step ensures that the VLAN membership updates only on the primary path that the RSTP/STP creates. When the primary path breaks, GVRP updates VLAN membership on the nodes in the secondary path. The system automatically recalculates the path or tree created on link or node failure.

FDB LEARNING RULE

Once you establish the tree/path, a simple learning rule enables implementation of the above scheme. In the switch, MAC-address learning occurs selectively for each VLAN; in other words, for each VLAN, the following rules enable or disable learning.

For a particular VLAN, learning is enabled if:

1. one of the members of the VLAN is an access port, or

2. the number of members of the VLAN is greater than two.

Multipoint flows require the first rule to prevent flooding of downstream traffic on the access ports and for access control of P2P flow, such as EIA (Ethernet Internet access). The second rule enables learning only if the number of member ports for any particular VLAN is greater than two and disables learning for VLANs whose members are just two (intermediate switches on the VLAN path). In metro switches, for most transit VLANs, the members for the VLAN are two; hence, learning is disabled for those VLANs, and one interface can efficiently forward the packet to another.

REQUIREMENT IN SWITCHES

Implementing the above scheme in a switch requires minor modifications in the protocol stack and switching fabric: a capability in the switching fabric to enable/disable learning for each VLAN and a control-plane implementation of the above-mentioned rules. Most switching fabrics available today provide the capability to enable/disable learning for each VLAN indirectly, through filtering rules, forwarding rules, or other mechanisms.

The control-plane implementation requires only a simple software modification that can be a special configuration for GVRP. GVRP switches off learning for a VLAN depending on the number of members in that VLAN, along with the rules mentioned in the previous section.

TCAM-BASED ARCHITECTURE

You can most efficiently implement the route look-up mechanism in a network processor and TCAM (ternary-content-addressable-memory)-based architecture. You can implement the control function in the forwarding plane; in this way, when the system encounters an ingress/egress interface failure, it can detect the alternate interface for a flow and transmit the packet with little or no dis-

CIRRUS LOGIC®

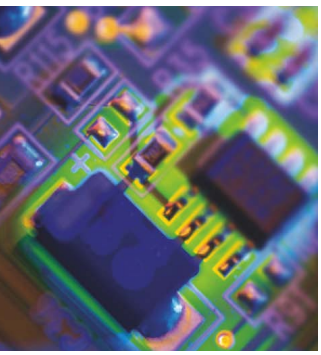

Highly Integrated Stereo CODEC Ideal for Portable, Low-Power Digital Audio Applications

Free Sample

Register for a free sample of the CS42L51 today.
Plus, get the datasheet, application notes and more.

www.cirrus.com/CS42L51

- Ground-centered outputs reduce bill-of-material costs and save space—eliminating the need for large, expensive DC-blocking capacitors.
- Consumes less than 12 mW @ 1.8 V power supply.
- Available in a 32-QFN lead-free assembly.

Priced from $3.95 (US) in 1K qty.

CS42L51—New IC Minimizes Player Component Count, Size & Cost

The CS42L51 delivers excellent audio performance while providing portable device developers with an innovative solution that maximizes battery life while minimizing end-product form factors. The CS42L51's key distinction is that it operates from a single, 1.8 V power supply for low-power consumption, yet delivers plenty of output power, 46mW, into stereo 16 Ω headphones for a great listening experience. This highly integrated 24-bit, 96-kHz stereo CODEC is based on a multi-bit Delta-Sigma architecture which allows infinite sample-rate adjustment between 4 kHz and 100 kHz.

By generating its own on-chip negative power supply voltage, the CS42L51 provides ground-centered outputs, which eliminates costly, space-consuming DC blocking capacitors required by some competing solutions. The elimination of these capacitors not only reduces the implementation cost but it also greatly simplifies product design and saves valuable circuit board space.

This next-generation CODEC is an ideal choice for any portable audio product developer. For complete specifications and to register for a free sample, visit our web site today.

Applications

- Hard-disc drive & Flash-based portable audio players
- Personal media players
- Mini-disc players & recorders
- PDA's & smart phones
- Digital cameras & camcorders
- Digital voice recorders
- Wireless headsets
- Guitar effects pedals
- Portable audio recording systems
- Portable gaming systems

www.cirrus.com/CS42L51

ruption. You can easily accomplish the mapping of a failed interface to the flow (VLAN) if you maintain the VLAN table in TCAM. You can optimize this type of implementation with an extra procedure to provide fast convergence for flows.

Such an implementation is straightforward in metro switches and optimizes learning without affecting any other protocol behavior of the switch, enabling switches with optimized learning to interoperate seamlessly with other switches. The scheme works effectively in metro-access/aggregation networks with any topology, including ring or mesh. Because the technique is based on GVRP, the optimization procedure can work with VLAN STP and MSTP (multiple STP).

SCALABILITY

Without optimized-learning mechanisms, intermediate switches in the access/aggregation network must learn the source MAC address of all traffic that flows through it. Hence, if N is the number of switches in the Layer 2 access/aggregation network, and M is the number of stations connected in each node, then the FDB-table size that each node requires is $O(M \times N)$.

Using an optimized-learning approach, switches learn the MAC addresses of the nodes in the access interfaces and for some multipoint transit flows for which the switch has more than two ports as members. Hence, if f is the number of flows with more than two ports as members, and K is the number of nodes for each flow, then each switch requires an FDB-table size of $O(N+f \times K) \rightarrow O(N)$.

Few flows require address learning, but if you ignore the ones that do, the required FDB-table size is $O(N)$. Hence, you can significantly reduce the large FDB-table-size requirement in metro Ethernet switches. The aforementioned optimization procedure reduces the MAC-address learning-table size from $O(M \times N)$ to $O(N)$.

The optimized-learning scheme that this article describes reduces the MAC-address learning-table size necessary in metro Ethernet switches. You can implement this scheme on most common switch fabrics and network-processor-based architectures. The implementation of optimized learning reduces switch complexity as well as the required capital expenditure. The technique simplifies network management and troubleshooting, because learning occurs only at the node level, and traffic tunnels through the rest of the network. EDN

AUTHORS' BIOGRAPHIES

Gopal Garg is a senior technology-marketing director at Cypress Semiconductor, where he markets networking semiconductor products. He holds a BS from Birla Institute of Technology and Science (Pilani, India) and an MBA from PU India.

R Thirumurthy is the head of R&D at Midas Communication Technologies, where he is responsible for broadband-product architecture, design, and product positioning. He holds a BE in computer science from College of Engg, Guindy (Chennai, India), and an MS in computer science from the Indian Institute of Technology, Madras.


90% CUT IN EMISSIONS
AND 50% CUT IN DEVELOPMENT TIME.
THAT'S MODEL-BASED DESIGN.
To meet a tough performance target, the engineering team at Nissan used dynamic system models instead of paper specifications.
The result: 50% time savings, the first car certified to meet California's Partial Zero Emissions Vehicle standard, and a U.S. EPA award.
To learn more, visit mathworks.com/mbd
MATLAB®
&SIMULINK®
The MathWorks
Accelerating the pace of engineering and science
©2005 The MathWorks, Inc.

THE DATA DETECTIVE

Use Sound Judgment with Dynamic Range

Almost everyone has heard the difference in the quality of sound produced by movie-theater system and that produced by home-theatre audio equipment. The difference can arise from audio-signal distortion that human ears can detect at low levels. Thus, rigorous testing must occur during product design and during manufacturing to ensure companies ship good MP-3 players, cell phones, sound cards, and so on. Testing involves measuring many characteristics, such as signal-to-noise ratio, total harmonic distortion, dynamic range, and intermodulation distortion.

Although distortion may seem like a qualitative criterion, instruments can readily measure energy at harmonic frequencies that we hear as distorted sound. To help quantify distortion, engineers use instruments that detect low-amplitude signals produced by combinations of standard test signals in audio electronics. Intermodulation distortion (IMD) measurements, for example, compute the ratio of the root-mean-square (rms) value of two test signals, f_1 and f_2, to the rms value of "products" that arise from mixing or modulating the two test signals. Signal

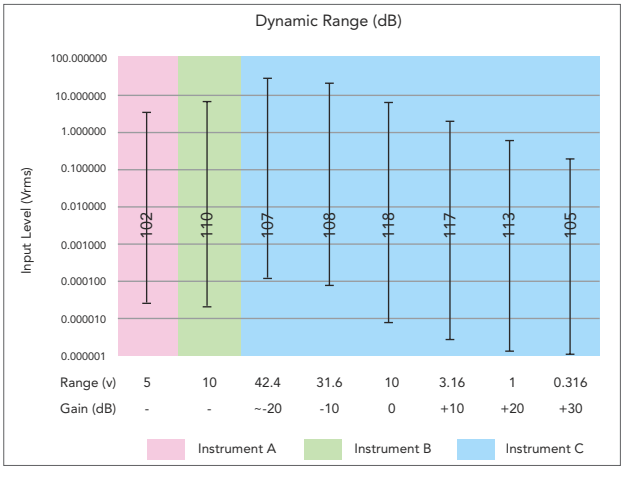

The plot above shows the range of measurement voltages and the corresponding dynamic range for three instruments. Instrument C offers six gain settings. Note the dynamic range varies from one gain setting to the next.

mixing takes place in non-linear circuits and components that exist to some extent in all electronics. Second-order IMD frequencies exist at f_2+f_1, f_2-f_1, $2f_2$, $2f_1$. Third-order IMD frequencies exist at $2f_1+f_2$, $2f_1-f_2$, f_1+2f_2, and f_1-2f_2. Good design and testing techniques ensure electronic equipment produces harmonics at low levels, which minimizes distortion.

Often, the second-order IMD frequencies exist some distance from f_1 and f_2, but two third-order IMD frequencies occur close to f_1 and f_2. Thus, filtering can remove second-order IMD components but not third-order IMD frequencies, which exist too close to f_1 and f_2. The change in amplitude of the test signals f_1 and f_2 also affects the amplitude of the IMD-generated signals. An increase of 5 dB in a fundamental frequency increases the amplitude of its second harmonic by 10 dB and the amplitude of its third harmonic by 15 dB. To ensure IMD frequencies will not distort audio signals, engineers require audio test equipment that can accurately measure IMD frequencies over as large a dynamic range as possible.

Dynamic range expresses how well an instrument can detect small signals in the presence or large signals. Because an IMD measurement requires simultaneous measurement of stimulus signals and their low-amplitude harmonics, instruments require a wide dynamic range. The comparison of the largest measurable signal to the smallest detectable signal in a ratio, expressed in decibels (dB), provides a dynamic-range value. The ratio does not specify that an instrument can measure a specific signal. In many cases, dynamic range will equal an instrument's signal-to-noise ratio because the smallest detectable signal amounts to noise. Audio specifications may refer to an A-weighted dynamic range, which involves measurements on filtered signals between 20 Hz and 20 kHz, the range of human hearing.

Most instruments provide adjustable signal-input settings, say from ±10V down to ±10 mV. Each setting specifies the maximum signal the instrument can work with, and each setting provides its own dynamic range. Exercise care when you determine the gain for a measurement. Say you have a 0.5 Vrms signal to measure and you can choose either a 10-V input range with a 118 dB dynamic range, or a 1-V input range with a 113 dB dynamic range. The higher dynamic range of the 10-V setting seems attractive, but do the math and you'll find that range lets you measure only down to about 12 µV. The noise floor of the instrument may hide low-amplitude harmonics and distortion produced by a device under test. On the other hand, the 1-V range with a 113 dB dynamic range lets you measure signals as low as 2 µV.

Distortion Contortion

When Bonnie substituted a new amplifier in a prototype audio circuit, several of her colleagues said they detected distortion. Bonnie wants to determine which harmonics dominate the signal and how much distortion they cause.

How should Bonnie test her circuit to get quantitative distortion results?

http://rbi.ims.ca/4398-501

Go to http://rbi.ims.ca/4398-501 to solve the challenge!

BY JOHN BETTEN AND ROBERT KOLLMAN • TEXAS INSTRUMENTS

Interleaving dc/dc converters boost efficiency and voltage

THOUGH SOMEWHAT MORE COMPLEX THAN SINGLE-PHASE DESIGNS, INTERLEAVED-BOOST CONVERTERS RUN COOLER, OCCUPY LESS SPACE, AND CAN COST LESS.

Boost power supplies are popular for creating higher dc voltages from low-voltage inputs. As the power demands from these supplies increase, however, a single power stage may be insufficient. This article presents an interleaved-boost approach, which, compared with a single-boost converter, both analytically and empirically provides efficiency, size, and cost advantages. The article also compares test results of 250W, single-phase and interleaved-boost power supplies. The interleaved-boost topology provides superior performance, albeit with increased complexity.

The motivation for the work this article describes was the selection of a power-supply topology for an ink-jet printer's solenoid driver. The input voltage was 12V dc; the required output was 37V dc at 7A. The supply's input current could exceed 20A. It was initially unclear whether a single power stage or a multiphase stage was most appropriate. As in buck regulators, currents could be high enough to make duplicate power stages desirable; they spread the dissipated heat around the pc board and reduce the stress on many of the circuit components. The work discussed included evaluation of single- and two-phase-boost topologies.

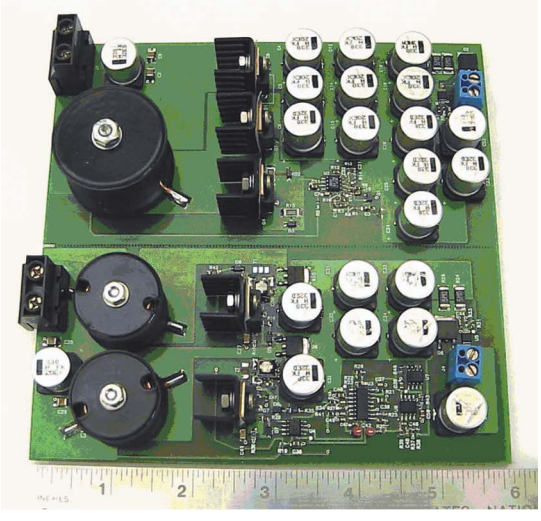

Figure 1 A converter with a single-phase topology (top) is larger than a converter with a two-phase topology (bottom) that produces the same output voltage and current.

Table 1 presents the power-supply requirements. To maintain the desired output voltage within a small margin, this supply is subject to substantial current surges as the solenoid energizes and de-energizes. In addition, high efficiency is important for maintaining an acceptable temperature rise. The 37V, 7A output delivers more than 250W to the load. Even with an efficiency of 91%, the power supply dissipates 25W, requiring multiple heat sinks. Moreover, the supply's size and cost were important, although no specific requirements were provided.

Figure 1 compares two power supplies. The top supply is the single-phase design with a single input inductor. The circuit below it is the two-phase design. The single-phase design requires approximately 18 in.^{2} of pc-board area, whereas the interleaved design requires 14 in.^{2} The largest differences in area between the two approaches are in the inductors, output capacitors, and heat sinks. The maximum height of the interleaved inductors is also less than that of those in the single-phase design.

Figure 2 shows the schematics of the single-phase- and interleaved-boost converters. In the single-phase design, applying a

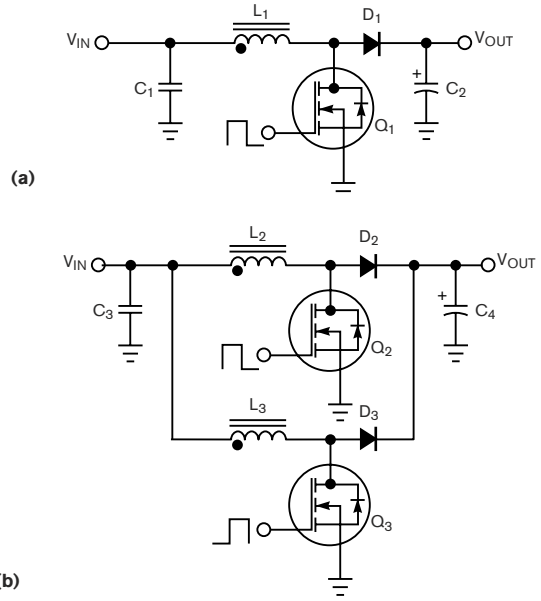

Figure 2 The single-phase supply (a) is simpler than the interleaved-boost supply (b).

gate voltage to FET Q_1 pulls the drain potential to ground, applying the input voltage across inductor L_1 and causing current to ramp up. During the ramp time, output capacitor C_2 must alone supply the load current. When Q_1 turns off, the voltage across L_1 reverses polarity to maintain current flow. This scenario forces the switch node more positive than the input voltage and forward-biases diode D_1, charging output capacitor C_2 and supplying the output current. For each of the two switching states, the inductor's volt-microsecond product must balance. That is, $d/f_S \times V_{IN} = (1-d)/f_{IN} \times (V_{OUT} - V_{IN})$, yielding the relation $V_{OUT} = V_{IN}/(1-d)$, where d is the duty ratio, f_S is the switching frequency, V_{IN} is the input voltage, and V_{OUT} is the output voltage. This expression is valid in CCM (continuous-conduction mode), in which the inductor current remains positive at all times.

Each phase of the interleaved-boost converter (**Figure 2**) works in the same way that this single-phase-boost converter does. The two power stages operate 180° out of phase, canceling the ripple current in the input and the output capacitors. The interleaved-boost approach uses forced current-sharing between the power stages to equalize the power that the stages deliver. Without this feature, one power stage could deliver substantially more power than the other, which would defeat the ripple cancellation.

DESIGN ANALYSIS

Figure 3 shows how interleaving benefits input-capacitor ripple-current cancellation. The two power stages operating 180° out of phase provide a two-to-one reduction in peak-to-peak ripple current. Because the interleaved-boost converter's combined input-ripple current equals that of the single-phase converter, the two-phase design's individual-phase ripple currents can each be twice as large as that of the single-phase design. The individual interleaved power stages operate at the same frequency as the single-phase design, 100 kHz, but the effective input- and output-ripple frequency is 200 kHz. The interleaved-design calculation used a frequency of 100 kHz and twice the ripple current of the single-phase design, yielding half the inductance. Because the two-phase design's effective input-capacitor ripple current was the same as that of the single-phase design, the two designs used an equal number of input capacitors. Ripple cancellation allows a choice of which components to reduce in number. Using two inductors, each having the same value as that in the single-phase design, halves the input-capacitance requirements. In a boost design, however, the inductor requirements are generally more critical than those of the input capacitors.

Interleaving benefits the output capacitors in about the same way as it affects the input capacitors. **Figure 4** shows the sin-

TABLE 1 ELECTRICAL SPECIFICATIONS OF AN INK-JET-PRINTER POWER SUPPLY

Parameter	Specification
Input voltage	10.8 to 13.2V
Output voltage	37V
Output current	7A maximum
Load step	1 to 7A
Efficiency	91% minimum

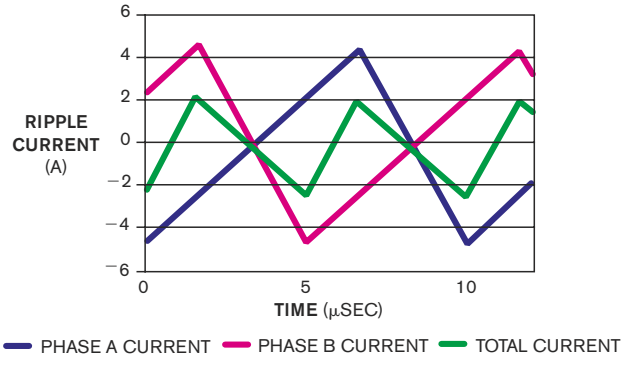

Figure 3 Because the two-phase converter's input-ripple currents are in phase opposition, the total input-ripple current is smaller than that of either phase.

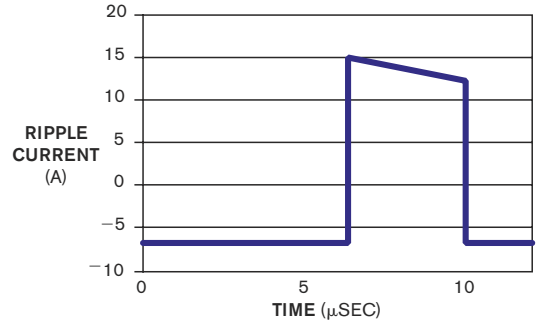

Figure 4 In the single-phase converter, the output-filter capacitor supplies all of the output current during the FET's on-time. During the off-time, however, a current of $I_{OUT} \times d/(1-d)$, or 14A, flows into the capacitor to recharge it.

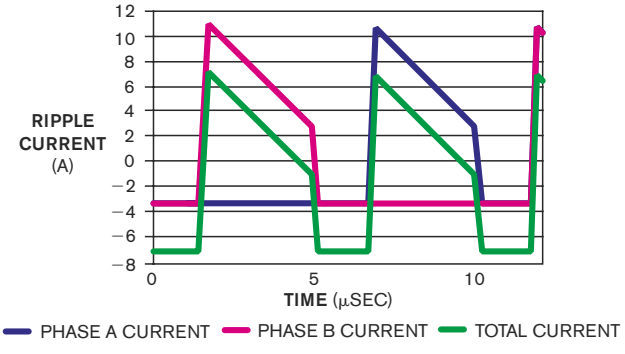

Figure 5 Interleaving reduces the two-phase converter's output-capacitor ripple current.

TABLE 2 SINGLE-PHASE VERSUS INTERLEAVED INDUCTOR

	Single phase	Interleaved
No. of inductors	One	Two
Inductance specification (mH)	27	10
Inductor-current specification (A)	21	10.5
Total LI² specification (mH×A²)	11,907	2205
Height (in.)	1.6	1.1
Total actual volume (in.³)	2	1.8
LI²/volume (actual)	6075	1225
Actual resistance (V)	0.008	0.006
Total loss (W)	5.8	1.2

A Purely Technical Experience

Yours for the Asking!

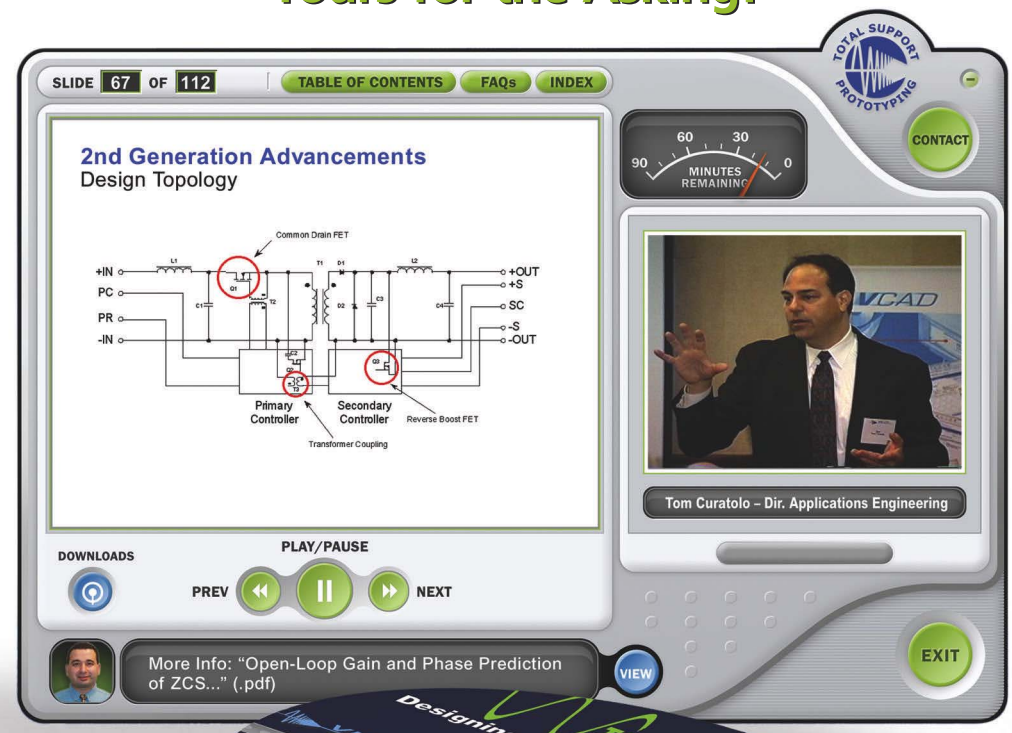

TOPICS COVERED

- Zero-current/Zero-voltage Switching Topology
- 2nd Generation Advancements
- Vicor Design Assistance Computer
- Electrical Performance
- Thermal Considerations
- Control Functions
- Unique Paralleling Methods
- Electrical Safety Concepts
- Module Do's and Don'ts
- Mounting Techniques
- Application Examples
- EMI Considerations

"DESIGNING WITH COMPONENT POWER MODULES"

A Technical Seminar on CD

Presented by leading power applications expert, Tom Curatolo, VICOR's CD Seminar is a purely technical 90-minute experience that you can run at your own pace. The CD provides power system designers with an in-depth technical review of issues including electrical, mechanical, and thermal design; EMI, and applications. The information is presented in a synchronized format and includes printed and video back-up materials which are downloadable.

YOURS FOR THE ASKING!

For your complimentary copy of "Designing with Component Power Modules" go to vicorpower.com/cd or contact us at **800-735-6200.**

This is the first in a series of Tech Seminar CDs and by registering now you will be automatically informed as new ones become available.

VICOR

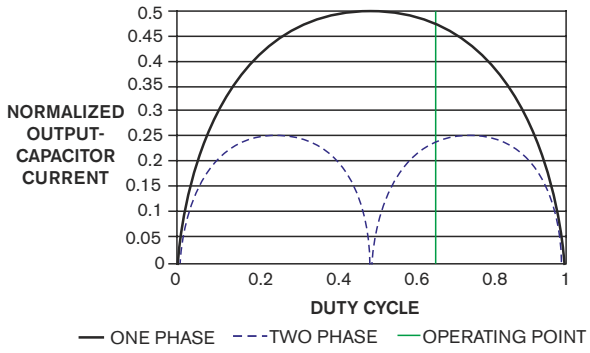

Figure 6 This interleaved-boost converter's output-capacitor ripple current is only half as great as that of a single-phase converter.

gle-phase output-capacitor ripple current. In this design, this waveform's rms current is approximately $I_{PP}\times(d\times(1-d))$, or 10A rms. The inductor slope, which appears at the top of the waveform, does not significantly add to the total rms current. This capacitor supplies all of the output current during the FET's on-time. During the off-time, however, a current of $I_{OUT}\times d/(1-d)$, or 14A, flows into the capacitor to recharge it. In designs that use aluminum-electrolytic output capacitors, capacitor ripple-current ratings determine the required number of capacitors.

Figure 5 shows the interleaved-boost converter's individual and combined output-capacitor currents. Not counting the inductor slope, the phase A and B currents' peak-to-peak amplitudes are half those of the single-phase design because the duty cycle of the current flowing into the output capacitors is twice that of the single-phase design. In **Figure 5**, the rms value of the combined or total waveform is 5A, allowing half the number of output capacitors to maintain a ripple voltage no greater than that of the single-phase design.

Figure 6 shows the ripple-current cancellation that you can obtain at various duty cycles. The vertical line indicates the operational duty cycle and shows the interleaved-boost circuit's two-to-one rms-current reduction compared with that of the single-phase circuit. A 50% duty cycle can provide perfect cancellation.

Figures 7 and **8** show the completed single-phase- and interleaved-boost-converter designs. In the single-phase design, a UCC38C43 PWM (pulse-width-modulation) controller operating in voltage mode drives a pair of MOSFETs. Because the boost converter offers no way to limit the output current in the event of a short circuit, a TPS2490 hot-swap circuit with overcurrent protection was added during testing to halt current flow during overcurrent faults. **Figure 8** illustrates the interleaved design using a UCC38220 dual-interleaved PWM controller.

TABLE 3 INTERLEAVED VERSUS SINGLE-PHASE DESIGN

	Interleaved design	Single-phase design
Circuit area (in.^{2})	14	18
Height (in.)	1.2	1.6
Full-load efficiency (%)	93.8	91.6
Full-load loss (W)	16	23
No. of power components	15	19
No. of heat sinks	Two	Three
No. of control components	53	30

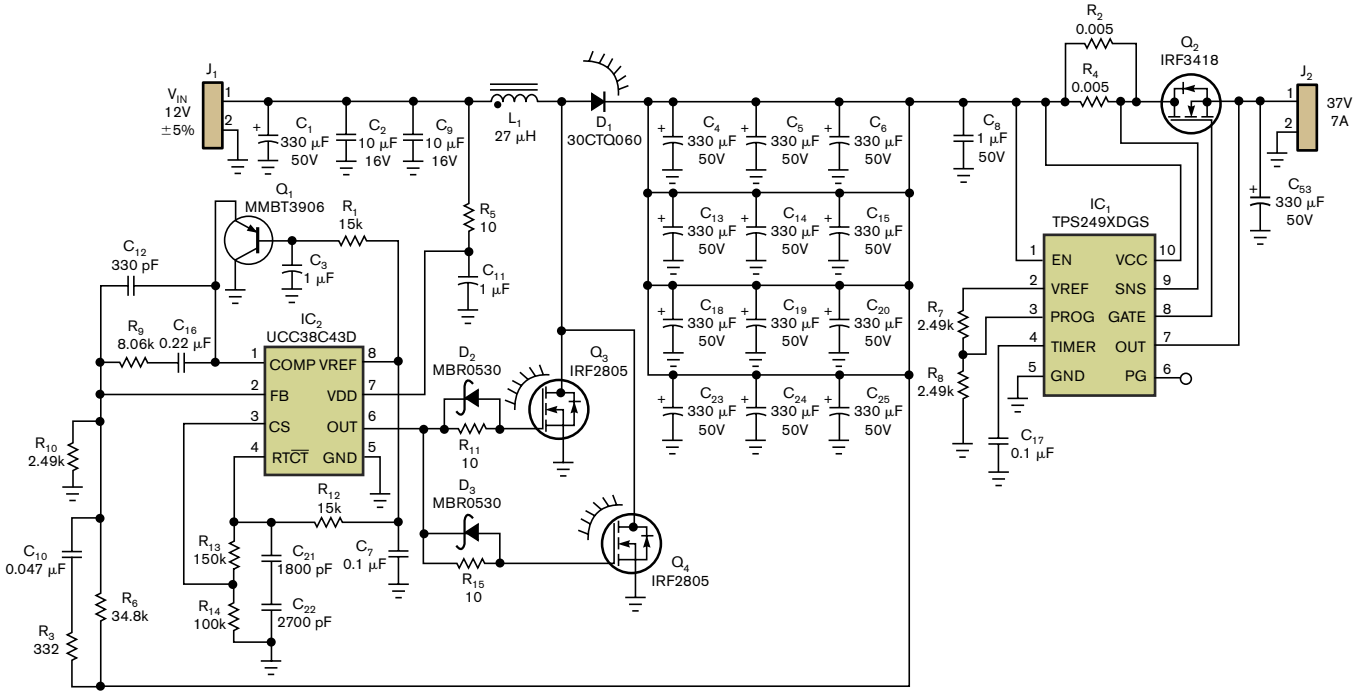

Figure 7 This single-phase-boost converter uses only two ICs, one FET, and one inductor, but it includes 13 large aluminum-electrolytic capacitors.

Intersil Video Products

Intersil High Performance Analog

Guess How Many Chips You Need for a Triple, 16x5 RGB Video Crosspoint

You win. Intersil's EL4544, with video sync extraction and pixel-by-pixel overlay for on-screen display, is the industry's first fully-integrated 16x5 RGB Video Crosspoint Solution.

Intersil's EL4544 is a 300MHz fully buffered RGB video crosspoint switch. With crosstalk rejection of 70dB and low offset enabled by an auto-calibration mode, the EL4544 provides razor sharp video performance.

Switch up to 16 SXGA video signals

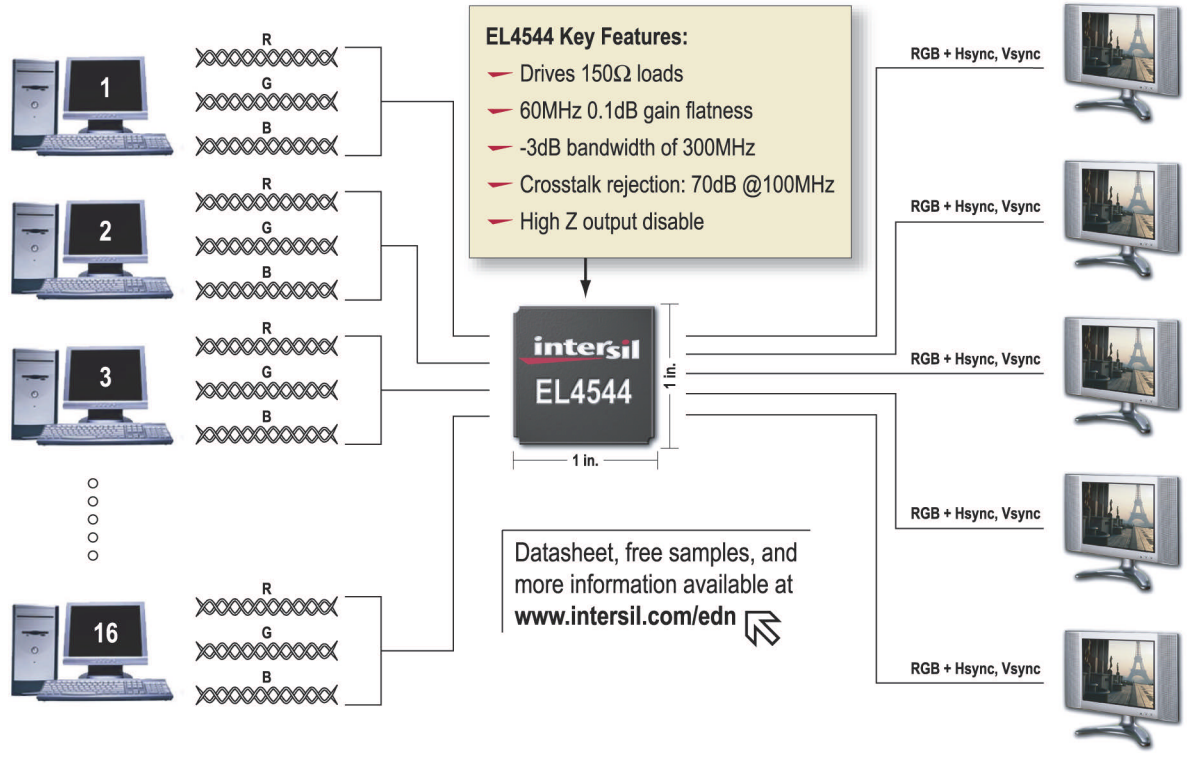

Intersil – Amplify your performance with advanced signal processing.

intersil®
HIGH PERFORMANCE ANALOG

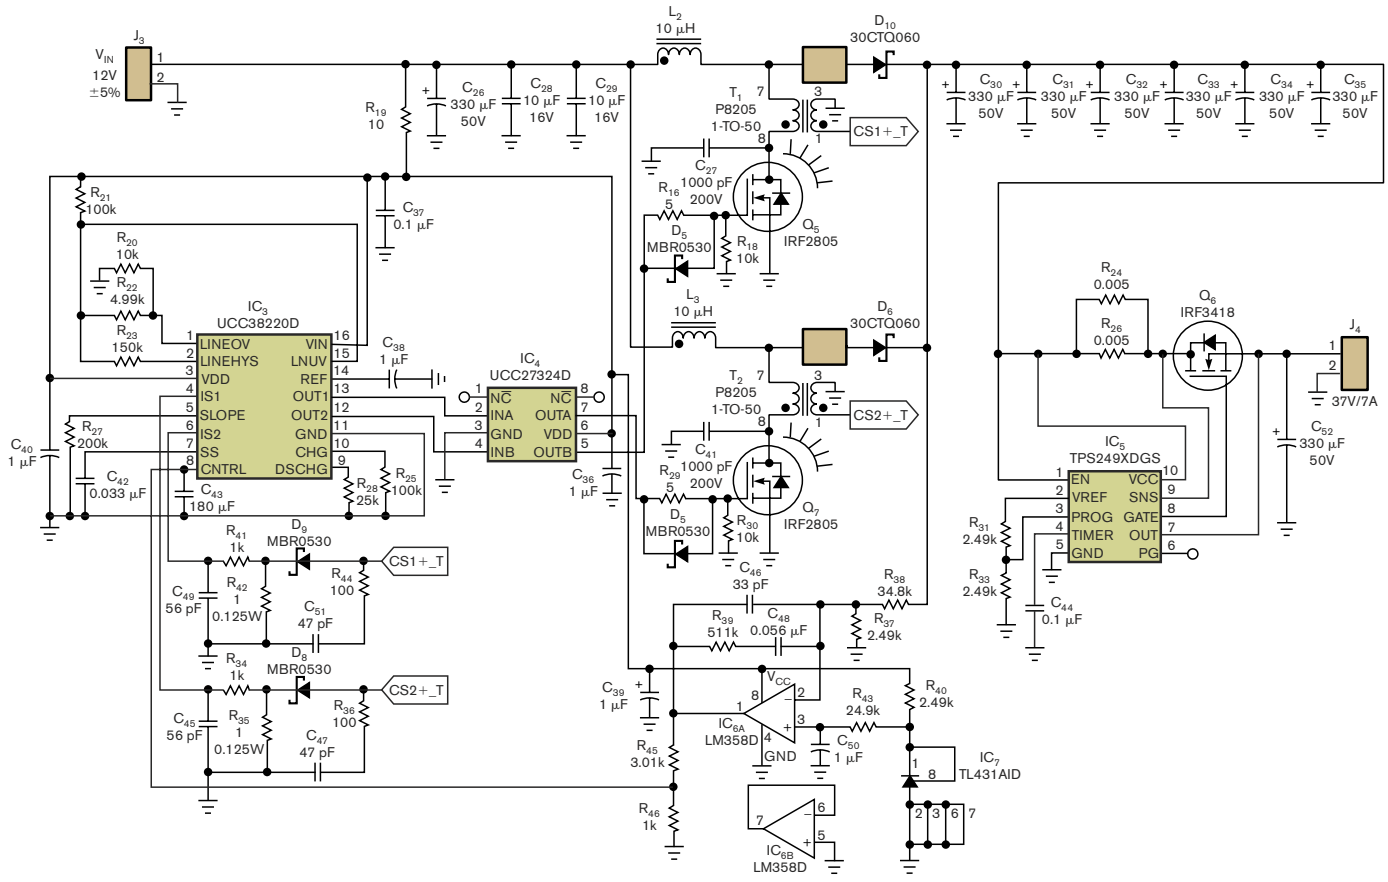

Figure 8 This interleaved design halves the output-capacitor requirements.

Low-cost transformers in the drain leads of Q_5 and Q_7 sense the FET current. The controller forces equal current in the two phases. Reduced current in the rectifiers eliminates the need for heat sinks and lowers assembly costs.

EXPERIMENTAL RESULTS

The above scenario compared the two designs for efficiency, input- and output-ripple voltage, and transient loading. In most situations, the two-phase approach exhibited better performance than did the single-phase approach. **Figure 9** compares the efficiency of the two approaches. Both meet the target 91% efficiency; however, at full load, the two-phase approach's efficiency is more than 2% better. Although this improvement may sound insubstantial, the difference in losses between the two supplies is significant. The single-phase design dissipates 23W, whereas the two-phase approach dissipates only 16W, significantly affecting the choice of heat sink and the thermal design.

The single-phase curve's early maximum and rapid decline indicate that the design has significant conduction losses. The big differences between the two designs are the losses in the inductor, boost diode, output capacitors, and pc board. **Table 2** compares the inductor requirements and designs' performance. The two-phase approach uses significantly less inductance than does the single-phase approach, and each inductor carries half the current. Energy-storage requirements and temperature rise determine an inductor's volume. The formula $\frac{1}{2} \times L \times I^2$ determines energy storage. **Table 2** shows that the energy storage of the single-phase design is five times that of the two-phase approach. Therefore, if the temperature rise of the inductors had been equal, the single-phase inductor would have been five times as large.

Rather than keeping the energy density equal, the designers chose to allow a higher temperature rise in the single-phase design, sacrificing some efficiency by using an inductor with higher losses; consequently, losses in the single-phase design are about 5W higher. Output capacitors accounted for about 1W of the power-loss difference. Ripple current in each of the output capacitors produced about 100 mW of dissipation, and the single-phase approach needed approximately six more capacitors than did the two-phase design. The two-phase approach required the use of two diodes in the power stage, with each carrying half of the total current. Consequently, the diodes exhibited a lower voltage drop, resulting in approximately 1W less loss.

Figure 10 shows the input- and output-voltage-ripple measurements. **Figure 10a** is the single-phase converter, and **Figure 10b** is the interleaved converter. The upper traces, which show the output-ripple voltage, illustrate several key points. The inductor current flowing through the output capacitor's ESR (equivalent series resistance) mainly determines ripple voltage. The traces in **Figure 10b** show the higher frequency ripple that

Intersil Switching Regulators

Intersil High Performance Analog

What's Black, White, and Cool All Over?

Maintaining IC temperatures doesn't have to be a riddle. Intersil's new high current Integrated FET Regulators have the industry's only true Thermal Protection with auto shut down at 135°C.

Intersil's EL7554 and EL7566 DC-DC buck regulators with internal CMOS power FETs operate from 3V-to-6V input voltage and are capable of up to 96% efficiency. But what's really cool about these devices is ground breaking features like built-in Thermal Protection and Voltage Margining for actual in-circuit performance validation.

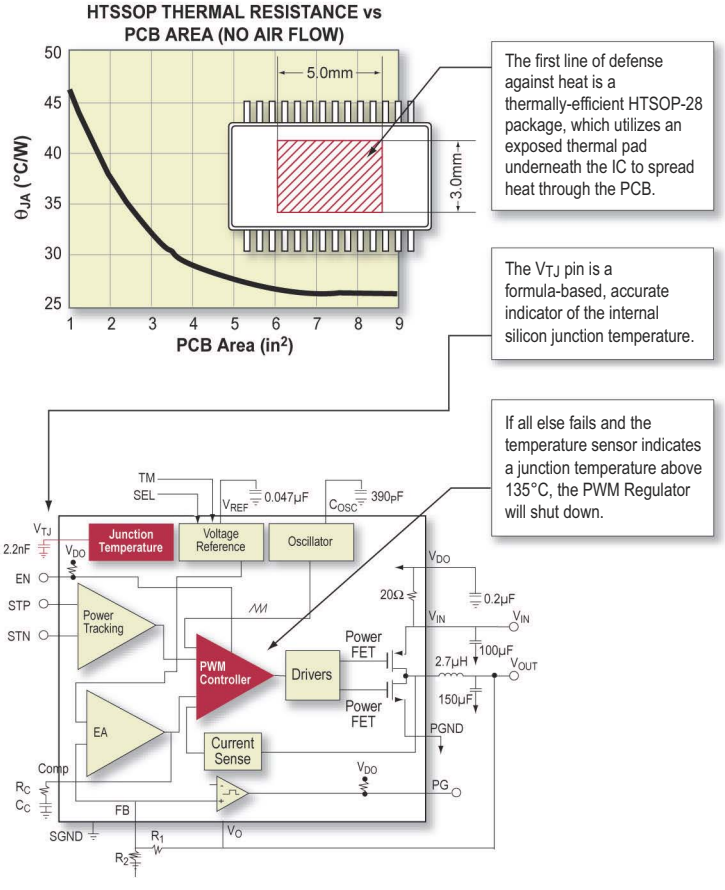

Key Features:

- 4A (EL7554) and 6A (EL7566) continuous output current
- Up to 96% efficiency
- Built-in 5% voltage margining
- 3V-to-6V input voltage
- 0.58 in² (EL7554) and 0.72 in² (EL7566) footprint with components on one side of PCB
- Adjustable switching frequency to 1MHz

For more information and samples, go to **www.intersil.com/edn**

Easy-to-use simulation tool also available. Modify switching frequency, voltage ripple, ambient temperature and view schematics waveforms, efficiency graphs and complete BOM with Gerber layout.
www.intersil.com/iSim

Intersil – Switching Regulators for precise power delivery.

intersil®
HIGH PERFORMANCE ANALOG

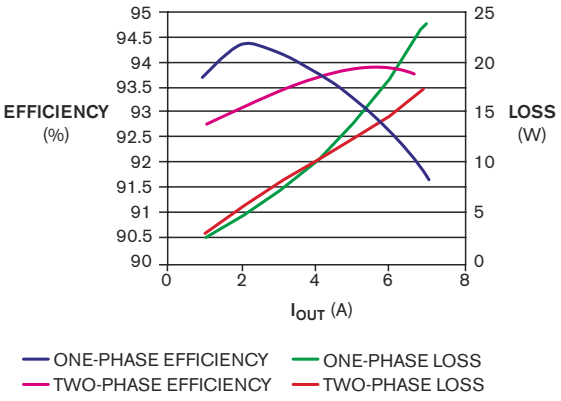

Figure 9 The two-phase converter is more efficient than the single-phase unit at output currents greater than 4A.

the interleaved approach achieves. In **Figure 10a**, the top of the ripple is nearly flat because of the large value of the boost inductor. In **Figure 10b**, the slope is significant because of the large change in inductor current during the power-switch off-time. The lower traces also show the input-ripple voltage's higher frequency with the two-phase approach.

Just as with buck regulators, interleaved-boost regulators can provide performance benefits over single-phase designs. **Table 3** compares the completed single-phase-boost design with the interleaved-boost approach. The interleaved-boost circuit is smaller, shorter, and more efficient. The fact that it has fewer output capacitors is largely due to lower output-ripple current, which results in lower cost and lower power dissipation. This circuit also significantly reduces the energy-storage requirement of the combined input inductors, thus reducing the magnetic volumes, heights, and dissipations. The multiphase approach reduced the overall power dissipation by 30% and spread that dissipation over a larger board area, allowing better thermal management. The main drawback of the multiphase approach is added circuit complexity, requiring measurement and balancing of each phase current as the larger number of control components illustrates. EDN

MORE AT EDN.COM

Go to www.edn.com/ms4136 and click on Feedback Loop to post a comment on this article.

AUTHORS' BIOGRAPHIES

John Betten is an application engineer and a member of the group technical staff at Texas Instruments (Dallas). He has 20 years of design experience in ac/dc- and dc/dc-power conversion, has published more than 20 articles, and has received one patent. He received a bachelor's degree in electrical engineering from the University of Pitts-

Intersil Digital Potentiometers

High Performance Analog

Huge Digital Control in a Tiny Package

Intersil's unleashes the X93154/55/56, the world's smallest and lowest cost Non-Volatile Digital Potentiometer.

The X93154, X93155 and X93156 addresses new market needs for high volume and space constrained applications such as portable or personal communications devices. The integration of non-volatile EEPROM for the wiper position provides design advantages including lower programming current and the elimination of additional high voltage supplies required by one-time programmable products.

All this functionality in tiny 2mm x 2.5mm TDFN package

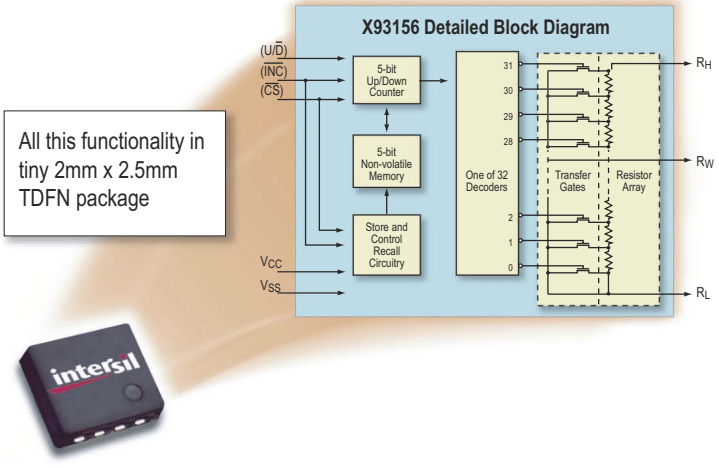

Features

- 3-wire up/down interface
- 32 wiper tap points. Wiper position stored in non-volatile memory and recalled on power-up
- Low power CMOS, with V_{CC} of 2.7V to 5.5V, active current of 250µA max, and standby current of 1µA max
- High reliability with endurance 200,000 data changes per bit and register data retention of 100 years
- Available in 8-lead MSOP and TDFN packages
- Pb-free and RoHS compliant packaging available

Key Parameters

Description	Conditions	MIN	TYP	MAX	Unit
Supply Voltage	X93154	2.7	3	3.3	V
	X93155	4.5	5	5.5	V
	X93156	2.7	-	5.5	V
End-to-end Resistence		35	50	65	kΩ
R_H, R_L Terminal Voltages		0	-	V_{CC}	V
Power Rating	R_{TOTAL} = 50 KΩ	-	-	1	Mw
Noise	Ref: 1kHz	-	-120	-	dBV
Wiper Resistance	X93156	-	-	1100	Ω
Wiper Current		-	-	0.6	mA
Resolution		-	3	-	%
Temperature (Industrial)		-40°C	-	+85°C	C

Datasheet, free samples, and more information available at **www.intersil.com/edn**

Intersil – Amplify your performance with advanced signal processing.

intersil®
HIGH PERFORMANCE ANALOG

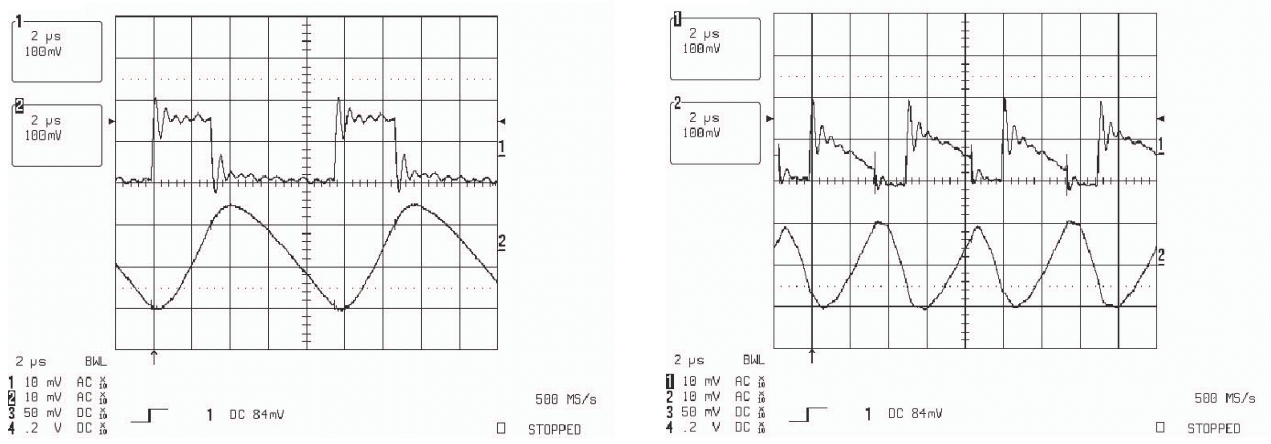

(a) (b)

Figure 10 The output-ripple voltage (upper traces in (a) and (b)) illustrates several points. For example, the higher frequency of the interleaved converter's ripple current eases filtering.

burgh in 1985 and is a member of IEEE. You can reach him at j-betten@ti.com.

Robert Kollman is a distinguished member of the technical staff at Texas Instruments and is a nationally recognized power-supply expert with more than 30 years of power-electronics experience. He is currently a power-management applications manager at TI (Dallas). He has authored more than 40 papers in the field. He holds a bachelor's degree in electrical engineering from Texas A&M University (College Station) and a master's degree in electrical engineering from Southern Methodist University (Dallas). Contact him at r-kollman@ti.com.

Support Across The Board.™

Bringing Products to Life.

At Avnet Electronics Marketing (EM), support across the board is much more than a tagline for us. From initial design through end of life – we are deeply committed to driving maximum efficiency throughout the product lifecycle. Take Battery Technology Inc. (BTI) for example.

Battery Tech – The Challenge

BTI technology keeps more than half a million laptops alive with its line of batteries. When it came time for BTI to recharge its product designs, it needed a product engineering solution that reduced the number of components on its board, and lowered overall costs.

Avnet EM and Atmel – The Solution

Avnet's MCU specialist introduced a new product solution involving Atmel's 8-Bit Flash memory based AVR microcontroller that solved BTI's challenge quickly and efficiently. Today, BTI utilizes Avnet's point of use replenishment system (POURS) program to ensure the proper flow of components into the manufacturing line, as it readies these new products for volume production.

BTI has also charged Avnet and Atmel to move its existing designs to Atmel's AVR platform – the industry's leading flash-based microcontroller. It's no shock to see why – with Avnet and Atmel's focused energy, BTI found total support across the board.

For additional application solutions and to download the BTI case study, visit: www.em.avnet.com/atmel/satb

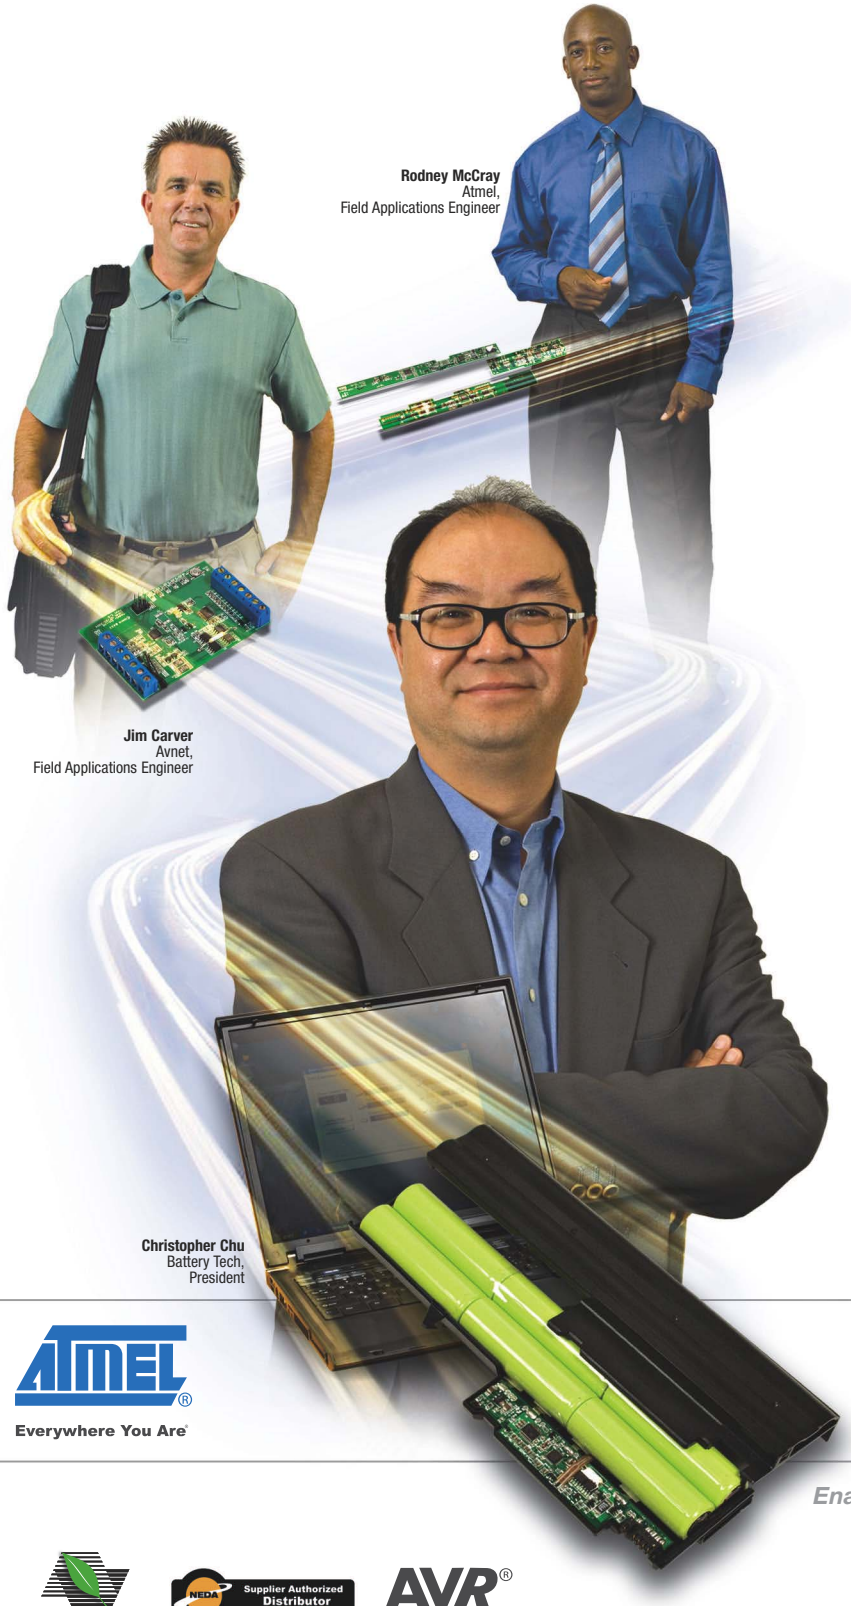

ATMEL® Everywhere You Are®

AVNET® electronics marketing

Enabling success from the center of technology™

Avnet Green Initiative

NEDA Supplier Authorized Distributor

AVR®

1 800 408 8353
www.em.avnet.com

PORTABLE AUDIO/VIDEO SOLUTIONS TO GET YOU TO MARKET FASTER.

POWER CONVERSION
Smallest, highly efficient 3-MHz DC/DC converters extend battery life and save space in portable applications. (TPS62300)

AUDIO POWER AMPLIFIER
Smallest, high efficiency stereo Class-D APA enables longer battery life. (TPA2012D2)

AUDIO CODEC
Highly integrated, low-power stereo codecs ideal for a variety of portable audio equipments. (TLV320AIC33)

INTERFACE
Highly configurable, low-power USB 2.0 compliant solutions provide world-class throughput. (TUSB6250)

DSP
TI DSP-based digital media processors come with razor-sharp reflexes, making them easily programmable to adapt to new standards:
- Video: H.26x, VC-1, MPEG-2, MPEG-4, Microsoft® WMV9
- Imaging: JPEG, MJPEG, JPEG2000
- Audio: MP3, AAC, Microsoft WMA9, G.7xx, Dolby® Digital, Pro Logic® II, DTS®-ES, DD-EX, THX®
- HD Radio™: Eureka 147 DAB, Digital Radio Mondiale (DRM)

Drive your portable audio and video innovations to market faster with TI solutions designed for your entire signal chain. With industry-leading DSP-based digital media processors, high-performance analog ICs, logic and application software, TI helps you get your product to the consumer fast. If you are looking for the easiest to use, most reliable, scalable and power-efficient solutions for portable consumer electronics, then look no further than TI.

For technical information including application notes, visit www.ti.com/tiportable

Technology for Innovators™

Texas Instruments

BY JIM DIETZ AND RICHARD HUBBARD • TEXAS INSTRUMENTS

The secrets of successful communications using LVDS

RELIABLE SINGLE- AND MIXED-TECHNOLOGY LVDS DESIGNS REQUIRE ATTENTION TO VOLTAGE LEVELS, NOISE MARGINS, AND DRIVE LEVELS.

Many choices now exist for LVDS (low-voltage differential-signaling) devices. Versions of these devices often coexist in the same system, which creates interoperability concerns. Among the available and practical versions, designers can choose from LVDS, bus-based LVDS, and M-LVDS (multipoint LVDS). Can these devices work in the same system? What issues should designers address when trying to mix these technologies in a system? What limitations do the combination of similar but different devices impose?

Each LVDS technology has strengths and limitations. When you compare them with single-ended signaling, LVDS technology's strengths include lower power, higher speed, and lower EMI. Technological limitations include the topologies they support, the number of nodes they allow, drive capability, and standard compliance. Standards minimize integration concerns; for example, the TIA/EIA-644A standard specifies the performance of LVDS devices for point-to-point and multidrop applications.

The TIA/EIA-899 standard for M-LVDS specifies requirements for multipoint devices. When developing a system from scratch, a homogeneous system is the best choice. The reality is, however, that systems often contain modules from various vendors, each complying with variants of the LVDS physical layers. The following guidelines will identify potential pitfalls and help designers to avoid them when integrating LVDS-I/O types. Examining M-LVDS devices demonstrates the range of I/O levels that this technology supports.

Table 1 highlights the key device parameters for the most common classes of LVDS devices. The TIA/EIA-644A standard and its predecessor, TIA/EIA-644, define the requirements for point-to-point (one driver, one receiver) and multidrop (one driver, multiple receivers) devices. Designers can connect as many as 32 receivers in a TIA/EIA-644A bus. Drive current is 3.5 mA, which is enough for single-termination applications but insufficient for double-terminated designs.

Bus-based LVDS increases the drive-current strength and preserves most of the features of the TIA/EIA-644A standard. These technologies share the same receiver common-mode range and receiver threshold as TIA/EIA-644A. They are driver enhancements of TIA/EIA-644A. M-LVDS provides a full complement of true multipoint features. Its drive capability of 11 mA supports double termination or heavily loaded backplanes. The receiver threshold is half that of the other technologies, thus providing more sensitivity. The receiver's ground-potential offset, relative to the driver, is twice that of other technologies. M-LVDS provides a superset of features of the other bus-based LVDS technologies and complies with an industry standard.

TRANSMITTER AND RECEIVER SPECIFICATIONS

In the driver-parameter section in **Table 1**, the test load is the impedance that the test circuit uses for measuring and reporting data-sheet specifications. The output differential voltage, V_{OD}, is associated with this test load. The driver-output current is the derived load current that the device sources to achieve the indicated output-differential-voltage value. Designers can more meaningfully compare the technologies by normalizing driver strength, using a common test impedance. The normalized output-differential-voltage value assumes that each transmitter acts as an ideal current source and that this current source drives a 100Ω load. None of the technologies in the **table** truly respond like an ideal current source over a wide range of loads, but the assumption is fairly accurate for each technology and allows meaningful comparisons.

Table 1 also shows that all classes of signaling, except M-LVDS, have a receiver threshold of 100 mV. M-LVDS receivers are twice as sensitive as the other technologies, providing improved noise margin over the other LVDS receivers. Low-voltage-signaling technologies generally find use in short-distance data transmission. A ground-potential difference of ±1V

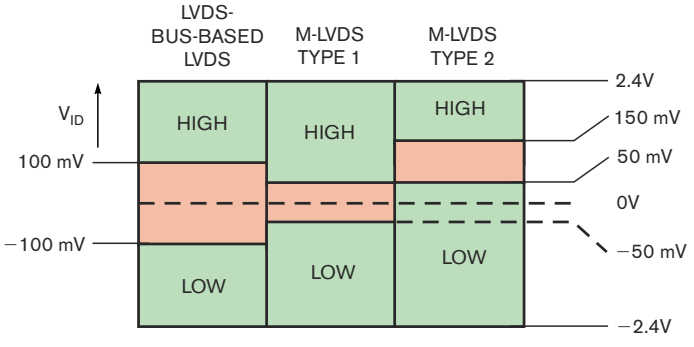

Figure 1 A comparison of receiver threshold voltages shows the differences among various technologies.

TABLE 1 COMPARISON OF MAJOR LVDS-FAMILY SPECS

Driver parameters	LVDS TIA/EIA-644A	Bus-based LVDS Nonstandard	M-LVDS TIA/EIA-899
Test load (Ω)	100	27 to 50	50
Output differential voltage (V_{OD}) (mV)	350	350	565
Driver output current (I_{OD}) (mA)	3.5	7 to 11.1	11.3
V_{OD} normalized (100Ω)	350 mV	700 mV to 1.11V	1.13V
Steady-state-output common-mode voltage ($V_{OC(SS)}$) (V)	1.2	1.2 to 1.3	1
Receiver parameters			
Input threshold voltage (VI_{TH}) (mV)	100	100	50
Common-mode voltage (V_{ICM}) (V)	0 to 2.4	0 to 2.4	−1 to 3.4
Voltage-potential difference (V_{GPD}) (V)	1	1	2
Fail-safe operation	Nonstandard	Nonstandard	Type 2-standard compliant
Supported architecture			
	Point-to-point, multidrop	Multipoint	True multipoint

had been the required performance specification until the release of the M-LVDS standard. M-LVDS doubles the allowable ground shift between drivers and receivers and ensures that receiver dynamic range is wide enough to handle multipoint applications, in which enabling and disabling of drivers is common. Manufacturers incorporate various fail-safe operation into receivers, so designers must pay attention to the fail-safe each device incorporates. The M-LVDS standard clearly specifies fail-safe operation.

M-LVDS provides many additional features to support true multipoint operation. A common application of M-LVDS devices is in multislot backplanes. Such systems have numerous impedance mismatches due to transmission-path stubs at the backplane connectors and line cards. To minimize the reflection from these stubs, M-LVDS specifies a controlled rise time, setting a minimum allowable transition time of 1 nsec. This edge rate limits the maximum achievable signaling rate, but it places no real restriction on applications because multipoint signaling is generally limited to 200 to 400 Mbps.

M-LVDS bus drivers never drive voltages greater than 2.4V, even under conditions of driver contention. This voltage limitation, coupled with the specification for receivers to operate over a wide range, ensures that a homogeneous M-LVDS system will operate and that the receivers will always be able to determine the correct bus state.

LVDS and bus-based LVDS all require receiver operation over a 0 to 2.4V range. This requirement effectively allows for a voltage-ground-potential difference of 1V. M-LVDS requires a range of −1V to +3.4V, which allows for a greater offset in ground potential. This increased common-mode range ensures that an M-LVDS receiver can accept LVDS, bus-based LVDS, and M-LVDS signals, thus positioning M-LVDS as a flexible receiver technology when designing mixed-use systems.

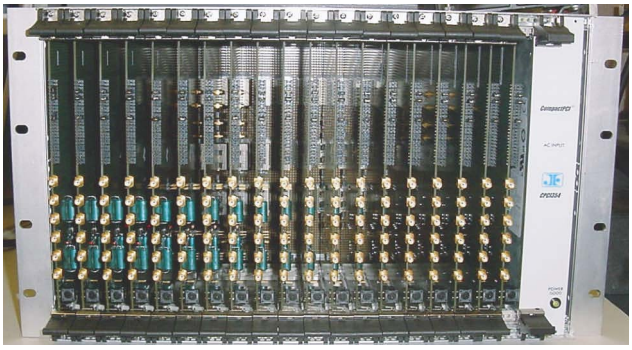

Figure 2 A full-scale hardware system confirms the results of analysis and predicted interoperability performance.

Another feature of M-LVDS, fail-safe operation, refers to the response of receivers under certain fault conditions or when drivers are inactive. Until TIA/EIA-899 emerged, many ways existed to accomplish fail-safe operation for LVDS technologies. Some of these techniques rely on external circuits to provide known outputs, whereas others are integrated approaches that force outputs to a known state. These methods are sometimes not interchangeable; thus, designers must heed how the data sheet specifies fail-safe operation. Bus-based LVDS technologies are enhanced driver specifications, and they address fail-safe operation in the same manner that LVDS does. In other words, no standard exists for fail-safe operation. Closely examine data sheets because this parameter can change from device to device.

The TIA/EIA-899 standard defines the fail-safe requirement, which must function over the full common-mode range of the driver. The standard identifies 50-mV-threshold Type 1 receivers and 100-mV-offset-threshold Type 2 receivers, which detect open-circuit and idle-bus conditions. Type 1 receivers are similar to LVDS receivers but with a more sensitive threshold range. Type 2 receivers provide standardized fail-safe operation by requiring an offset threshold. Type 2 receivers "see" signals lower than 50 mV as low, whereas they see those higher than 150 mV as high (**Figure 1**).

DRIVER CONTENTION

What happens when more than one driver is active on a bus at once, and how does each technology deal with driver contention? A designer hopes that, at a minimum, no damage occurs and that the bus voltage stays within some limit. Driver contention does not occur in point-to-point or multidrop systems; hence, TIA/EIA-644A does not cover it. The nonstandard, bus-based LVDS technologies also do not cover this issue. However, M-LVDS, a comprehensive, multipoint standard, addresses driver contention. M-LVDS drivers monitor the bus

POWERpuzzler #7

sponsored by POWER INTEGRATIONS

Stable Modes and Bodes

by Peter Vaughan
Manager of Product Applications
Power Integrations

Take a break from your daily routine and test your power supply design knowledge by trying your hand at answering the three questions below regarding stability of switched mode power supplies. Then check your answers at ***www.powerint.com/puzzler7*** and enter for a chance to win a new Apple iPod Mini.

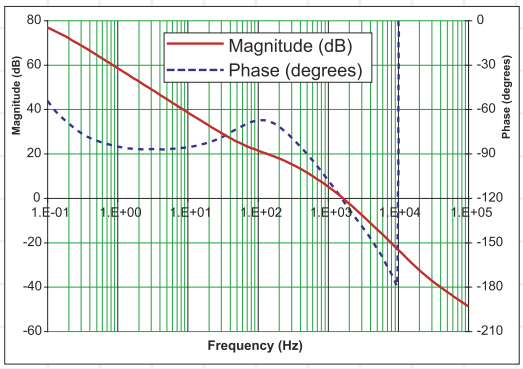

Figure 1A

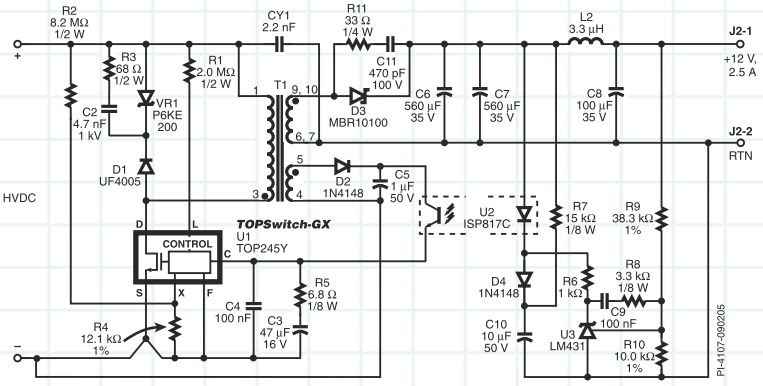

Figure 1B

Question 1: beginner

The graph in Figure 1A shows a gain / phase bode plot as measured during the routine development of a switched mode power supply using a *TOPSwitch*®*-GX* as shown in Figure 1B (See EPR-34 at www.powerint.com/appcircuits.htm)

a) Identify the gain crossover frequency and the associated phase margin.
b) How does setting the gain crossover frequency affect the power supply performance?
c) What is the importance of the phase margin?

Question 2 : advanced

Waveforms in Figures 2A and 2B show a step load change (top waveform) in two different power supplies, causing the duty cycle to exceed 50%. The output response is shown in the lower waveform.

Which supply uses current mode control and which one uses voltage mode control?

What steps need to be taken to avoid the instability shown in Figure 2A?

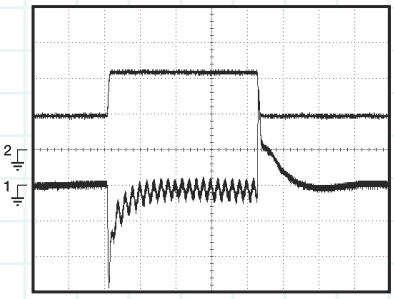

Figure 2A

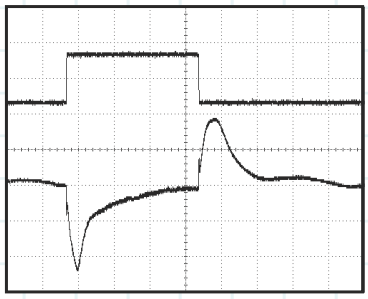

Figure 2B

Question 3 : expert

When choosing between current mode control and voltage mode control, the designer usually has to consider the relative tradeoffs between each approach and make compromises to achieve the best overall solution.

What are some of the relative pros and cons of current mode control and voltage mode control and why does the *TOPSwitch-GX* provide the best of both worlds?

The answers to these questions can be found at ***www.powerint.com/puzzler7****. Check out how well you did and enter to win an Apple iPod Mini!*

TABLE 2 COMPARISON OF OUTPUT-DIFFERENTIAL VOLTAGE IN LVDS TECHNOLOGIES

Driver parameters	LVDS TIA/EIA-644A	Bus-based LVDS Nonstandard	M-LVDS TIA/EIA-899
Test load (Ω)	100	27 to 50	50
Minimum output-differential voltage (mV)	247	200 to 247	480
Minimum output-differential voltage normalized to 100Ω (mV)	247	494 to 741	960
Minimum output-differential voltage normalized to 40Ω (mV)	100	200 to 300	400

voltages and control the output current, such that the bus does not exceed 2.4V. TIA/EIA-899 also requires that disabled devices and receivers do not impact the bus in a manner that causes it to exceed the 2.4V limit. Considering these provisions, a homogeneous M-LVDS system does not see a bus voltage that exceeds 2.4V.

The driver's output voltage, receiver's threshold voltage, and receiver's common-mode voltage are important in addressing interoperability between classes of LVDS devices. Referring to **Table 1**, normalized-load-voltage technologies vary by a factor of almost three. This difference shows how mixing technologies in a point-to-point and multidrop system can become complex. Whether the receiver can handle these increased load-voltage levels depends on the receiver's common-mode-voltage range and maximum differential-input voltage. From the **table**, you can deduce that the M-LVDS receivers provide the widest common-mode-voltage range and that they can provide the greatest margin when interfacing to other drivers. This fact still does not provide the complete picture, though. Noise margin is a key concern to draw conclusions relating to interoperability.

Noise margin for differential-bus architectures is the minimum driver-differential-output voltage minus the maximum receiver-input threshold. The current-source assumption of **Table 1** generates the normalized voltages in **Table 2**, which provides the minimum differential-output voltage across data-sheet loads, normalized with 100 and 40Ω test loads. The 40Ω test-load value derives from realistic expectations for a multipoint-backplane system.

The standards fully define noise margin for standards-compliant devices. The TIA/EIA-644A standard provides for a minimum output-differential voltage of 247 mV and a maximum threshold voltage of 100 mV for a noise margin of 147 mV in a homogeneous system. The other bus-based-LVDS technologies do not comply with standards, so designers need to carefully study each bus-based-LVDS data sheet to calculate the noise margin. For M-LVDS, the noise margin is 480−50 mV, or 430 mV. It is important to understand the load conditions. **Table 2** shows how output differential voltage can vary while the load changes. The **table** also raises questions concerning the noise margin that exists when the system uses mixed technologies.

M-LVDS DEMONSTRATION SYSTEM

Figure 2 shows a multipoint-M-LVDS demonstration system, illustrating the performance of M-LVDS in a realistic environment. The 21-card demonstration system has a 0.8-in. pitch between cards. The backplane traces have 130Ω differential impedance. The connectors, devices, and stubs couple with tight card pitch to reduce the 130Ω backplane impedance to an effective impedance of approximately 40Ω. The demo uses 40Ω for a multipoint backplane and 100Ω for point-to-point comparison.

Table 3 provides calculated differential noise margin for a homogeneous system and various mixed technologies, using M-LVDS for normalization. Even though LVDS has less noise margin, it can drive any of the other technology receivers in a 100Ω environment. Designers usually select LVDS for speed in a point-to-point system, and they choose M-LVDS if greater noise margin rather than speed is the relevant issue.

In a homogeneous system, M-LVDS provides the most noise margin for either a 40 or a 100Ω architecture. This noise margin often provides a level of comfort for driving a signal through trace, connector, cable, and backplane, even in a point-to-point system. A true multipoint system creates loads that need more noise margin, which necessitates a technology such as M-LVDS. LVDS targets use in 100Ω environments and thus is unsuitable for loads greater than those that the TIA/EIA-644A standard specifies.

The last area of concern for mixing technologies involves the receiver common mode. In most instances, the allowable ground-potential difference between driver and receiver for

TABLE 3 CALCULATED DIFFERENTIAL-NOISE MARGINS FOR SINGLE- AND MIXED-TECHNOLOGY

Noise-margin parameters	Driver LVDS / Receiver LVDS	Driver Bus-based LVDS / Receiver Bus-based LVDS	Driver M-LVDS / Receiver M-LVDS	Driver LVDS / Receiver M-LVDS	Driver Bus-based LVDS / Receiver M-LVDS	Driver M-LVDS / Receiver LVDS	Driver M-LVDS / Receiver Bus-based LVDS
Input threshold voltage (mV)	100	100	50	50	50	100	100
Differential voltage (V)	247	200 to 247	480	247	200 to 247	480	480
Minimum output-differential voltage normalized to a 100Ω load (mV)	247	494 to 741	960	247	494 to 741	960	960
Minimum output-differential voltage normalized to a 40Ω load (mV)	100	200 to 300	400	100	200 to 300	400	400
Noise margin with a 100Ω load (mV)	147	394 to 641	910	197	444 to 691	860	860
Noise margin with a 40Ω load (mV)	0	100 to 200	350	50	150 to 250	300	300

16-Bit, 130Msps ADC

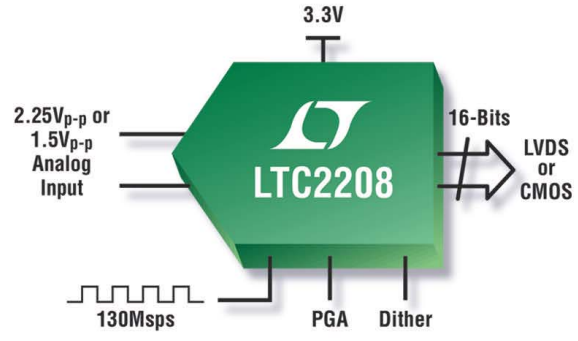

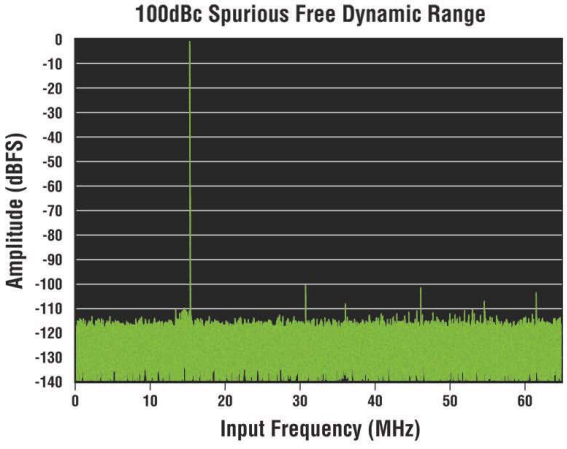

LTC2208 with Internal Dither Achieves 100dB SFDR from a Single 3.3V Supply

Linear Technology extends its leadership in high-speed ADCs with the LTC®2208, the fastest 16-bit ADC. The LTC2208 delivers exceptional dynamic range performance for the most demanding wideband, low noise, signal acquisition applications. This high performance family offers two unique performance enhancements: an internal transparent dither circuit that dramatically improves SFDR for low level input signals and a digital output randomizer that suppresses the effects of coupling from the ADC outputs.

Features

- 16-bits, 130Msps, 1250mW
- 78dB SNR, 100dB SFDR
- PGA Front End, 2.25V_{p-p} or 1.5V_{p-p} Input Ranges
- Internal Transparent Dither
- Digital Output Randomizer
- LVDS or CMOS Outputs
- Single 3.3V Supply
- Clock Duty Cycle Stabilizer
- 9mm x 9mm QFN Package

16-Bit ADC Family

Part No.	Speed	SNR	Power
LTC2208	130Msps	78dB	1250mW
LTC2207	105Msps	78dB	850mW
LTC2206	80Msps	78dB	650mW
LTC2205	65Msps	79dB	450mW
LTC2204	40Msps	79dB	350mW
LTC2203	25Msps	81dB	220mW
LTC2202	10Msps	81dB	150mW

Info & Online Store

www.linear.com/2208

Literature: 1-800-4-LINEAR

Support: 408-432-1900

LINEAR TECHNOLOGY

LVDS and nonstandard-bus-based technologies is ±1V. Although this value is approximately correct, the designer needs to understand the allowable effects that the receiver can withstand under extremes of specification limits. For LVDS, the driver has a common-mode-output range of 1.125 to 1.375V, with 1.2V as the typical value. LVDS receivers accept an input of 0 to 2.4V, which means that the receiver can have a ground shift that can be approximately 1V. This shift value may actually be less if the driver provides the maximum output differential voltage—such as 450 mV at 1.375—at the high end of the common-mode range. This requirement means that the allowed ground shift is only 800 mV. Some vendors offer receivers, such as the SN65LVDS33D, with larger input ranges—in this case, −4 to +5V. The TIA/EIA-899 specification for M-LVDS requires an input-voltage range of −1 to +3.4V.

LVDS is an appropriate technology for point-to-point- and multipoint-system technologies. LVDS and M-LVDS interfaces are based on standards, whereas the other bus-based-LVDS technologies are not. LVDS technologies offer speed improvements, lower power, and better EMI than older, single-ended- bus technologies. A homogeneous system is the best design option, but the interface technologies are sometimes mixed. In such cases, designers should never exceed the data-sheet limits, must take noise margins into account, and must understand the common-mode range of the application. M-LVDS provides a new type of LVDS that can either increase the noise margin over other options or provide a true multipoint approach when a design requires it. Fail-safe operation, driver contention, and compliance are all available with M-LVDS implementations. Each technology has its place, and interoperability is possible if designers take the appropriate steps. EDN

MORE AT EDN.COM

We encourage your comments! Go to www.edn.com/ms4146 and click on Feedback Loop to post a comment on this article.

AUTHORS' BIOGRAPHIES

Jim Dietz is systems-engineering manager at Texas Instruments, where he provides strategic-marketing and applications support for TI's high-speed-interconnect portfolio. He has a master's degree in systems engineering from Texas Tech University (Lubbock, TX), a master's degree in international management studies from the University of Texas—Dallas, and a bachelor's degree in electrical engineering from Clarkson University (Potsdam, NY).

Richard Hubbard is a new-product-development/characterization manager at Texas Instruments, where he manages LVDS/M-LVDS/CML (current-mode logic)/LVPECL (low-voltage positive-emitter-coupled-logic)-device development, including buffers, translators, crosspoint switches, and serializers/deserializers. He has a master's degree in business administration from the University of Texas—Austin and bachelor's degrees in electrical engineering and liberal studies from the University of Central Florida (Orlando, FL). His other interests include raising his five-year-old son, painting, reading, bowling, and pursuing higher education opportunities.

Cut Switcher EMI >20dB

Fundamental and Harmonic Frequencies

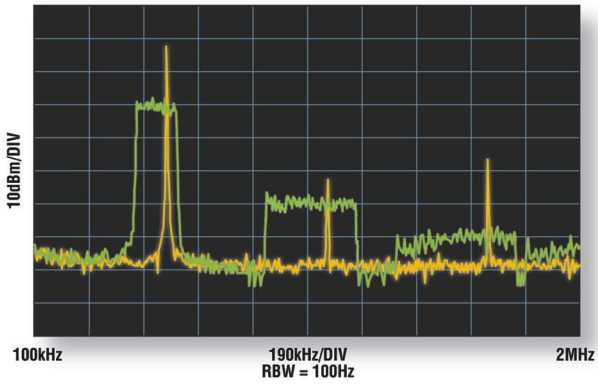

Fundamental Frequency

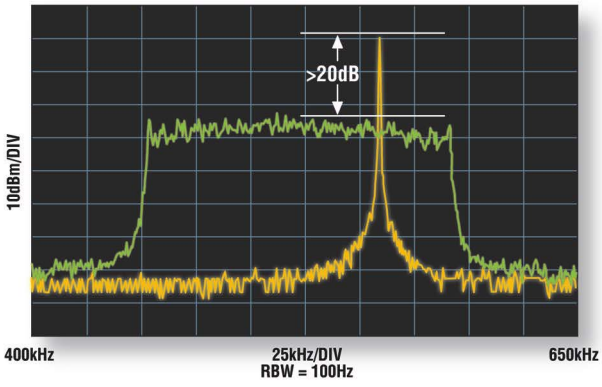

Spread Spectrum Enabled

Spread Spectrum Disabled

Eliminate Noise Problems with On-Chip Spread Spectrum

Built-in spread spectrum modulation in Linear Technology's new family of DC/DC regulators reduces peak radiated and conducted energy at any frequency by at least 20dB. This improvement significantly suppresses peak electromagnetic interference (EMI), reducing the need for metallic shielding or expensive filtering. Many of our future switching regulators will incorporate this technology – see the table below for our current offering.

DC/DC Converters with Spread Spectrum Modulation

Part Number	Step-Down Switching Regulators	V_{IN} (V)	V_{OUT} (V)	I_{OUT} (A)	Tracking	PLL	Fixed Frequency (kHz)	Spread Spectrum Frequency Range (When Enabled) (kHz)	Package
LTC®3736-1	2-Phase, Dual Synchronous Controller	2.75 to 9.8	1: 0.6 to V_{IN} 2: 0.6 to V_{IN}	5	✓	-	300, 550 or 750	450 to 580	QFN-24 SSOP-24
LTC3808	Synchronous Controller	2.75 to 9.8	0.6 to V_{IN}	5	✓	✓	300, 550 or 750	460 to 635	DFN-14 SSOP-16
LTC3776	DDR Memory Dual Controller	2.75 to 9.8	1: 0.6 to V_{IN} 2: V_{DDQ}/2	5	✓	✓	300, 550 or 750	450 to 580	QFN-24 SSOP-24
LTC3809	Synchronous Controller	2.75 to 9.8	0.6 to V_{IN}	5	-	✓	300, 550 or 750	450 to 580	DFN-10 MSOP-10
LTC3252	Inductorless, Dual, 2-Phase	2.7 to 5.5	1: 0.9 to 1.6 2: 0.9 to 1.6	0.25	-	-	-	1,000 to 1,600	DFN-12
LTC3251	Inductorless, 2-Phase	2.7 to 5.5	0.9 to 1.6	0.5	-	-	1,600	1,000 to 1,600	MSOP-10
LTC3445	Monolithic, I²C, Triple Output	2.5 to 5.5	1: 0.85 to 1.55 2, 3: ≥ 0.3	1: 0.6 2,3: 0.05	✓	-	1,500	Adjustable Spread 0% to 22.4%	QFN-24
LTC3415	Monolithic, PolyPhase, Stackable	2.5 to 5.5	0.6 to V_{IN}	7 x n	✓	✓	2,000	1,000 to 3,000	QFN-38

Info & Online Store

www.linear.com

Literature: 1-800-4-LINEAR

Support: 408-432-1900

Free SwitcherCAD Download

LINEAR TECHNOLOGY

60V to 100V_{IN} Regulators– No Compromise

Switching

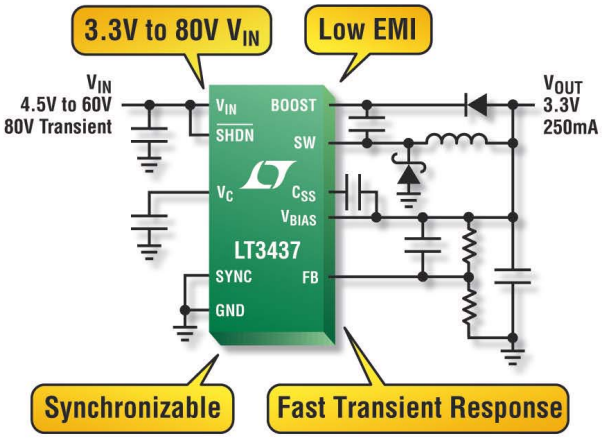

Linear

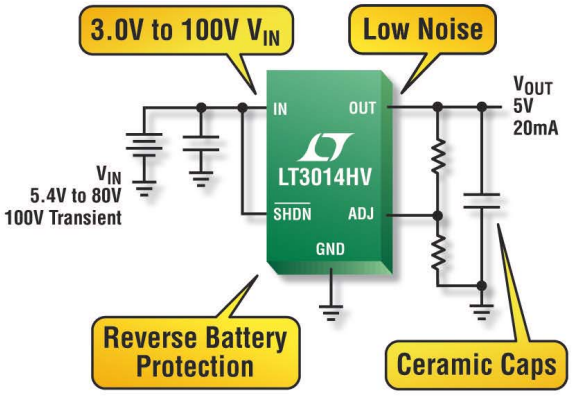

Rugged, Wide V_{IN} Range, Low I_Q & Low EMI

Linear's growing line of 60V to 100V input voltage capable regulators enables simpler DC/DC converter designs by eliminating the need for transient protection. The LT®3437 is a 500mA, 60V input capable monolithic step-down switching regulator in a tiny 3mm x 3mm DFN package. Its current mode topology provides fast transient response and excellent loop stability. The LT3014HV 20mA-rated linear regulator can take 100V input transients and has only 7µA of quiescent current, reverse battery protection and thermal shutdown. Our high voltage linear and switching regulators are well suited to telecom, automotive and industrial applications.

High Voltage Regulators

Part Number	Device Architecture	V_{IN} Range	I_{SW} (A)	Frequency	I_Q	Package
LT3010	High Voltage LDO	3.0V to 80V	0.05†	N/A	30µA	MS8E
LT3014/HV	High Voltage LDO	3.0V to 80V/100V**	0.02†	N/A	7µA	DFN, ThinSOT™
LT3012/13	High Voltage LDO	4.0V to 80V	0.25†	N/A	55/65µA	DFN, TSSOP-16E
LT3433	Buck-Boost Regulator	4V to 60V	0.50	200kHz	100µA	TSSOP-16E
LT3437	Step-Down Regulator	3.3V to 60V/80V**	0.50	200kHz	100µA	DFN, TSSOP-16E
LT1976/77	Step-Down Regulator	3.3V to 60V	1.50	200/500kHz	100µA	TSSOP-16E
LT3434/35	Step-Down Regulator	3.3V to 60V	3.00	200/500kHz	100µA	TSSOP-16E
LT3800/LT3724	Step-Down Controller	4.0V to 60V	10.00*	200kHz	100µA	TSSOP-16E
LTC3703/-5	Synch. Step-Down Controller	4.1V to 100V	20.00*	600kHz	1.5mA	SSOP-16E

*Depends on MOSFET Selection, **Transient Capable, †I_{OUT} for LDO

Info & Online Store

www.linear.com
Literature: 1-800-4-LINEAR
Support: 408-432-1900

Free SwitcherCAD Download

LINEAR TECHNOLOGY

LINEAR TECHNOLOGY

DESIGN NOTES

ThinSOT Micropower Buck Regulator Has Low Output Ripple

Design Note 373

Keith Szolusha

Introduction

High voltage monolithic step-down converters simplify circuit design and save space by integrating the high-side power switch into the device. In most cases, the switch is an n-type transistor (NMOS or NPN) with a boot-strapped drive stage, requiring an external boost diode and capacitor as well as the main catch diode, complicating the applications circuit.

The LT®3470 is a 40V step-down converter with the power switch, catch diode and boost diode integrated in a tiny ThinSOT™ package. The boosted NPN power stage provides high voltage capability, high power density and high switching speed without the cost and space of external diodes.

The LT3470 accepts an input voltage from 4V to 40V and delivers up to 200mA to load. Micropower bias current and Burst Mode® operation enable it to consume merely 26µA with no load and a 12V input. Hysteretic current mode control and single-cycle bursts result in very low output ripple and stable operation with small ceramic capacitors. The combination of small circuit size, low quiescent current and 40V input makes the LT3470 ideal for automotive and industrial applications.

Current Mode Control

The LT3470 uses a hysteretic current control scheme in conjunction with Burst Mode operation to provide low output ripple and low quiescent current while using a tiny inductor and ceramic capacitors. The switch turns on until the current ramps up to the level of the top current comparator, then turns off and the inductor current ramps down through the catch diode until the bottom current comparator trips and the minimum off-time has been met.

In continuous mode, the difference between the top and bottom current comparator levels is about 150mA. Since the switch only turns on when the catch diode current falls below threshold, switching frequency decreases, keeping switch current under control during start-up or short-circuit conditions.

If the load is light, the IC alternates between micropower and switching states to keep the output in regulation (Figure 3a). Hysteretic mode allows the IC to provide single switch-cycle bursts for the lowest possible light-load output voltage ripple (<20mV peak-to-peak from 12V to 3.3V at zero load.) During continuous switching mode (Figure 3b) at higher current levels, the output voltage ripple is even smaller (<10mV peak-to-peak).

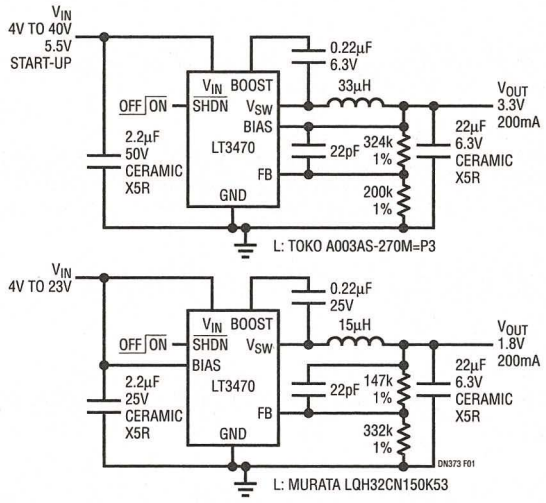

Figure 1. Typical 3.3V and 1.8V Output Applications

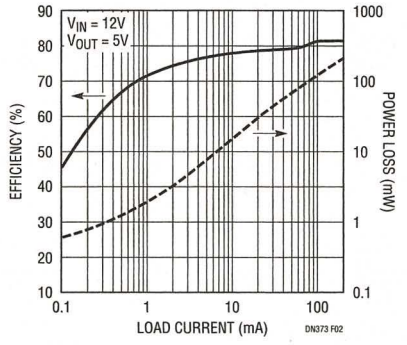

Figure 2. Efficiency and Power Loss vs Load Current

Design Flexibility with Integrated Boost Diode

A high side NPN power switch in a buck regulator design needs a driver voltage that is at least a few volts higher than the switch or input voltage. When there are no other high voltage lines available, a bootstrapping method of providing several volts of boost to the IC is required. When there is at least 2.5V on the output, the boost voltage can be most efficiently derived from the output. If the output voltage is too low, 1.8V for example, the boost voltage must be derived from the input.

Integration of the high side bootstrapping boost diode into the IC does not limit boost diode flexibility. Boost diode flexibility such as the ability to connect to various sources and/or the inclusion of a Zener blocking diode is needed for both high and low output voltages with and without wide input voltage ranges. The anode of the boost diode can be connected to different sources via the BIAS pin. In most cases, this is a simple connection to either the input, when the output voltage is below 2.5V, or the output, for output voltages above 2.5V. Additional Zener diode voltage drop in the boost diode path or a transistor bias supply as shown in Figure 4 protects the IC from BOOST pin overvoltage when there is a wide input voltage range.

Conclusion

The LT3470 is a wide input voltage range, hysteretic mode, fully integrated monolithic 300mA step-down DC/DC converter. The onboard high side NPN power switch, Schottky boost diode, and Schottky catch diode combined with the small ThinSOT package and high 40V input voltage make this a simple and versatile IC to use for many step-down applications with less than 200mA load current.

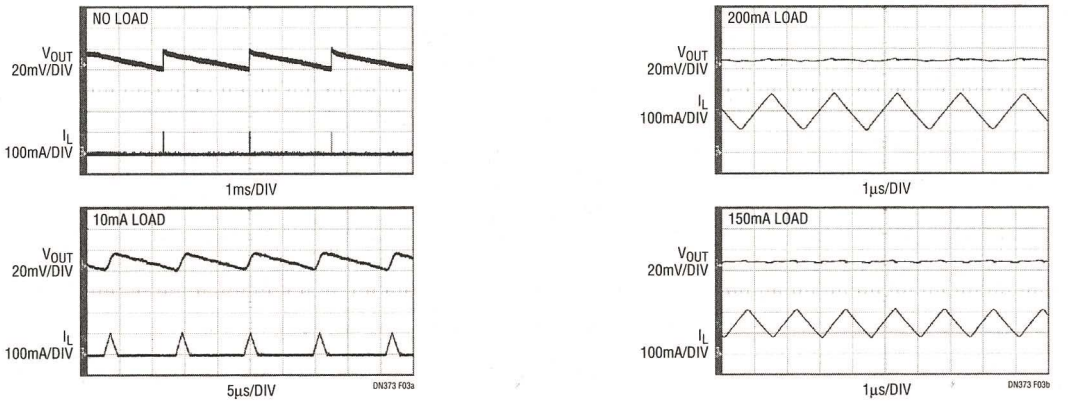

(3a) (3b)

Figure 3. (3a) Burst Mode Operation—Single Pulse Burst Mode Operation Has Only 20mV$_{P-P}$ Ripple. (3b) Continuous Operation—Extremely Low Output Voltage Ripple

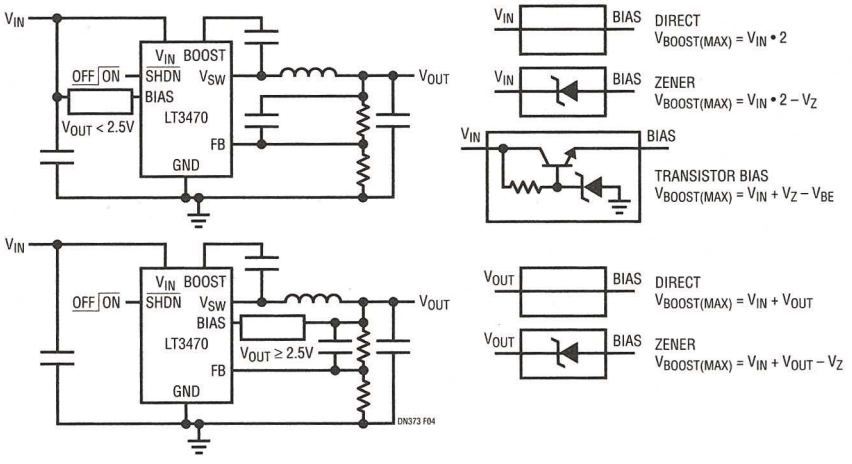

Figure 4. BIAS and BOOST Pin Connection Variations Provide Input and Output Voltage Range Flexibility

Data Sheet Download

http://www.linear.com

For applications help,
call (408) 432-1900, Ext. 2759

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

dn373f LT/TP 1005 305K • PRINTED IN THE USA

EDITED BY BRAD THOMPSON
AND FRAN GRANVILLE

designideas
READERS SOLVE DESIGN PROBLEMS

Dither a power converter's operating frequency to reduce peak emissions

Bob Bell and Grant Smith, National Semiconductor, Phoenix, AZ

Designers of dc/dc switching power converters face the challenge of controlling EMI (electromagnetic-interference) emissions produced during normal operation. If large enough, these emissions conduct through power lines or radiate to other assemblies within a system and can compromise a system's performance. Emission peaks typically occur at the converter's fundamental switching frequency and gradually reduce in amplitude at each higher order harmonic, with most of the emitted energy confining itself to the fundamental and lower order harmonics. Modulating, or dithering, the power converter's operating frequency can reduce the peak emissions by spreading EMI over a band of frequencies.

Most modern PWM controllers use an external resistor to set the operating frequency, which typically increases with decreasing resistor values. For example, the LM5020's internal oscillator delivers a regulated 2V at its programming pin (RT), and a programming resistor connected to RT sets the current that RT delivers. The oscillator also delivers a proportional current into an internal timing capacitor (**Reference 1**). The period of the timing capacitor's ramping voltage determines the oscillator's frequency.

The external dithering circuit in **Figure 1** comprises a simple stand-alone comparator-based oscillator configured to operate at approximately 800 Hz. The output state of comparator IC_2 goes high upon power-up. R_1, R_2, and R_3 set the comparator's positive input, which initially rests at 2.9V. The voltage at capacitor C_3 ramps up toward the positive threshold.

When the voltage at the comparator's negative input reaches the positive-threshold voltage, the comparator's output switches low, which also lowers the threshold at the comparator's positive input to 2.1V. The voltage at capacitor C_3 then ramps down toward the new threshold, and, when it reaches the lower threshold voltage, the cycle repeats. The voltage across C_3 approximates a triangular wave with a minimum voltage of 2.1V and a maximum of 2.9V.

To dither the LM5020 controller's PWM-oscillator base frequency, the triangular wave generated by IC_2 modulates the current from the controller's RT pin. Resistor R_5 sets the percentage of modulation dither. The right side of R_5 is fixed at the RT pin's regulated potential of 2V, and the low-frequency triangle wave coupled from IC_2 through capacitor C_2 drives R_5's left side. For R_5 with a value of 64.9 kΩ, the peak-to-peak current through resistor R_5 is approximately 12 µA. With the dither circuit disconnected, the steady-state current that RT sources is approximately 121 µA, and the 12-µA p-p dither current thus represents 10% total modulation.

An LM5020-controlled PWM flyback dc/dc converter, IC_1, evaluates the dither circuit's effectiveness. The circuit's fundamental operating frequency is 250 kHz, which the controller's R_T

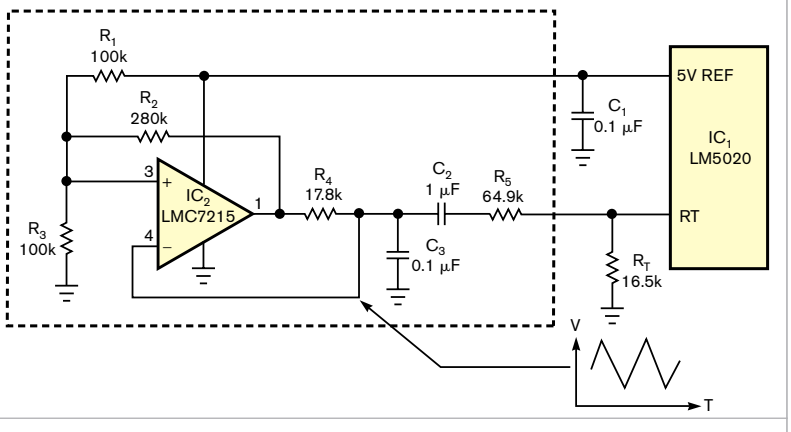

Figure 1 A low-frequency triangular-wave oscillator modulates a PWM controller's oscillator frequency.

DIs Inside

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

resistor sets. The red trace of **Figure 2** shows the conducted emissions on the circuit's positive input-power line without the dither circuit in operation. The peak emissions are narrowly confined around the fundamental oscillator frequency of 250 kHz with a measured amplitude at the fundamental frequency of −24 dB.

Connecting the dither circuit to the controller's RT input produces the blue trace of **Figure 2**. Conducted emissions around the fundamental frequency now disperse around the fundamental frequency with maximum amplitude reduced by −34 to +10 dB. EDN

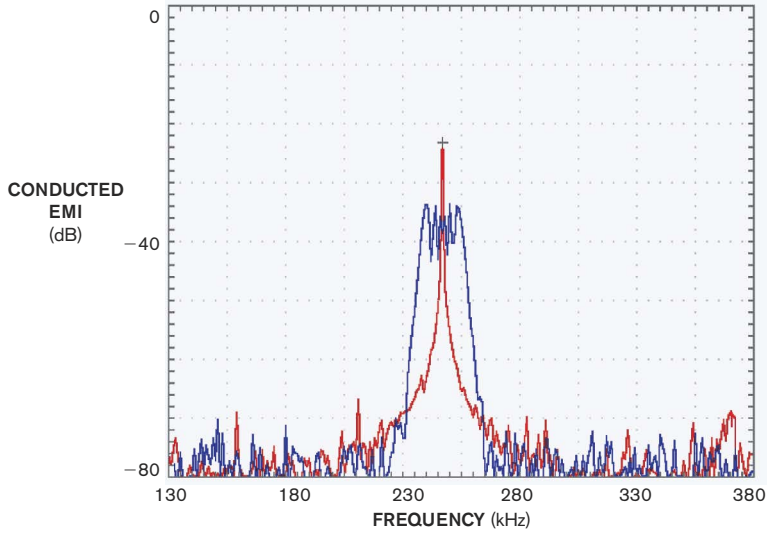

Figure 2 Measured at the dc/dc converter's input, peak conducted emissions without dither (red trace) decrease by 10 dB with dithering applied (blue trace).

REFERENCE

1 LM5020 data sheet, National Semiconductor, www.national.com/pf/LM/LM5020.html#Datasheet.

Single-port pin drives dual LED

Tom Gay, Dornstadt, Germany

Most current microcontrollers offer I/O ports that can change their functions during program execution. As outputs, the circuits can sink and source reasonably large amounts of current. This Design Idea shows three alternative methods for driving a two-pin, two-color LED from a single I/O pin. **Figure 1** illustrates one possible approach that uses external inverter IC_1 to drive D_1, a red/green bidirectional LED. A logic-high output on the port pin forces current through the green (upper) LED and pulls the inverter's input high, which drives the inverter's output low and sinks current from the green LED. A logic-low output on the port pin raises the inverter's output high, delivering current to the red (lower) LED; the microcontroller's output sinks current from the red LED.

To turn off both LEDs, you reconfigure the microcontroller's port pin from output to input or switch the pin to tri-state mode, either of which prevents the microcontroller's port pin from sourcing or sinking current. This circuit's primary disadvantage is that it yields no control over each LED's brightness; instead, resistor R_5 determines forward current for both LEDs.

Figure 2 presents an approach that also involves a major disadvantage. Zener diodes D_3 and D_4 and resistors R_3 and R_4 form a low-impedance voltage divider that applies $V_{CC}/2V$ to one end of LED D_5. The value of V_{CC} drives the selection of the zener diodes' voltage, V_Z, with lower voltage zener diodes allowing more LED current and higher voltage ones limiting maximum LED current. Given that the microcontroller's outputs can deliver rail-to-rail voltages, the difference between V_{CC} and V_Z limits maximum forward current for both LEDs. For example, if V_{CC} is 5V and V_Z is 3V, the forward voltage across either LED is less than 2V. Once a designer selects the zener-diode voltage, only small variations in V_{CC} can occur; otherwise, the LEDs' brightness would fluctuate.

Using discrete components, another circuit offers an inexpensive approach that avoids the other circuits' disadvantages (**Figure 3**). When the microcontroller's output port goes high, current flows through the green (upper) LED, R_2, D_2, and FET Q_2, which the port's high level turns on. When the microcontroller's output port goes low, transistor Q_1 turns on and delivers current to the port pin through R_2 and the red (lower) LED. The circuit operates symmetrically because silicon diode D_2's forward-voltage drop is present

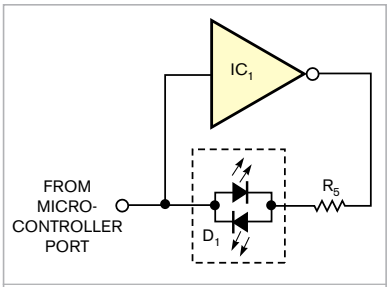

Figure 1 An inverter can drive a bidirectional, two-color LED but applies the same amount of current to both LEDs.

Intersil Interface Products

Intersil High Performance Analog

Want a More Dependable Bus?

Try a better Transceiver. Intersil's new ISL8308XE 5V Fractional (1/8) Unit Load, RS-485/RS-422 Transceivers incorporate "Hot Plug" functionality to keep your bus from crashing during power-up and power-down.

That's not all. These devices feature 15kV ESD Protection; "Full Fail-Safe" design to ensure a high Rx output if Rx inputs are floating, shorted, or terminated but undriven; and low bus currents to allow up to 256 transceivers on the network without violating the RS-485 network specification's 32 unit load maximum without using repeaters.

HOT PLUG PERFORMANCE (ISL83080E) vs DEVICE WITHOUT HOT PLUG CIRCUITRY (ISL83086E)

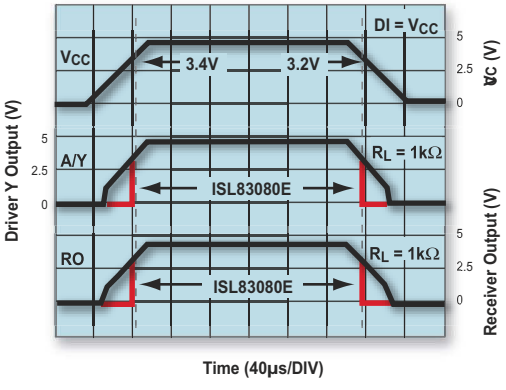

Key Features:

- True 1/8 Unit Load allows up to 256 devices on the bus
- Hot Plug circuitry to maintain three-state Tx and Rx outputs during power-up and power-down
- Full Fail-Safe (open, short, terminated and floating) receivers
- ±15kV HBM ESD Protection on RS-485 I/O pins and Class 3 ESD Protection on all pins
- Available in Pb-Free and small MSOP packages

Datasheet and more info available at **www.intersil.com/edn**

5V, High ESD, Fractional (1/8) Unit Load RS-485/RS-422 Key Specifications

Device	# of Tx/ # of Rx	Devices Allowed on Bus	Half/ Full Duplex	High ESD?	Hot Plug?	Data Rate (Mbps)	Slew Rate Limited?	Tx/Rx Enable?	ICC EN / DIS (µA)	SHDN ICC (µA)	VCC Range (+V)	Pkg.
ISL83080E	1 / 1	256	Full	Yes	Yes	0.115	Yes	Yes	530 / 530	0.07	4.5 to 5.5	14 Ld SOIC
ISL83082E	1 / 1	256	Half	Yes	Yes	0.115	Yes	Yes	560 / 530	0.07	4.5 to 5.5	8 Ld MSOP 8 Ld SOIC
ISL83083E	1 / 1	256	Full	Yes	Yes	0.5	Yes	Yes	530 / 530	0.07	4.5 to 5.5	14 Ld SOIC
ISL83085E	1 / 1	256	Half	Yes	Yes	0.5	Yes	Yes	560 / 530	0.07	4.5 to 5.5	8 Ld MSOP 8 Ld SOIC
ISL83086E	1 / 1	256	Full	Yes	No	10	No	Yes	530 / 530	0.07	4.5 to 5.5	14 Ld SOIC
ISL83088E	1 / 1	256	Half	Yes	No	10	No	Yes	560 / 530	0.07	4.5 to 5.5	8 Ld MSOP 8 Ld SOIC

Intersil – Amplify your performance with advanced signal processing.

intersil®
HIGH PERFORMANCE ANALOG

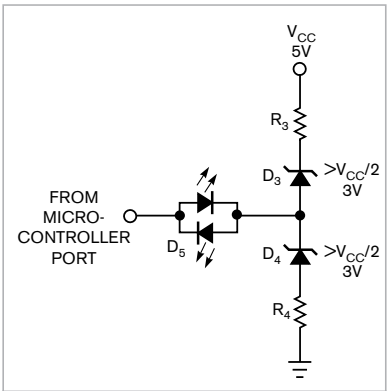

Figure 2 A zener-diode and resistor-voltage divider exhibits sensitivity to supply-voltage variations and allows LED brightness to vary.

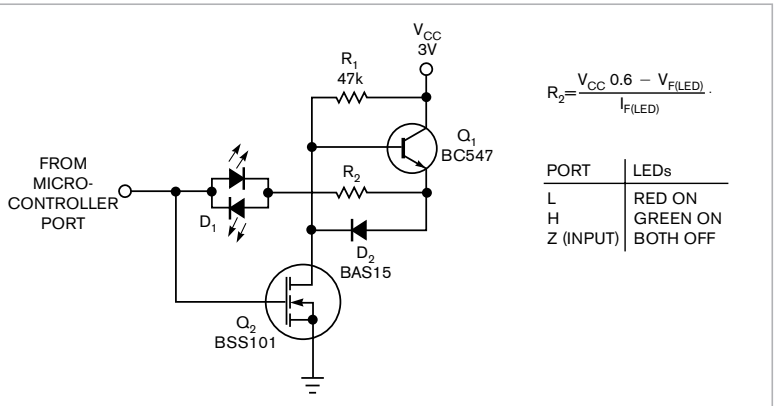

Figure 3 This circuit provides immunity to supply-voltage fluctuations and more uniform LED brightness.

regardless of whether the microcontroller's port pin goes high or low. V_{CC} may vary during operation but must remain higher than 3V.

You can individually adjust the LEDs' currents to equalize brightness or compensate for a difference between the microcontroller's power-supply voltage and the LED-driver circuit's V_{CC}. Replace R_2 with two resistors connected in series between Q_1's emitter and D_2's anode. Connect the midpoint of the two resistors to the LEDs.

With the microcontroller's port pin configured as an "input with pullup," the port delivers a small current to the green LED. However, pullup-resistor values of 22 kΩ or higher do not cause misleading light output from LEDs in the off-state. When the input signal from the port pin floats—that is, with V_{CC} at 5V and the port configured as an input with no pullup resistor—the circuit draws no additional current, and the quiescent current, which R_1 determines, averages less than 100 μA. EDN

Network linearizes dc/dc converter's current-limit characteristics

John Guy and Lance Yang, Maxim Integrated Products Inc, Sunnyvale, CA

Recently announced versions of integrated step-down dc/dc converters have eliminated the requirement for a high-side current-sense resistor by sampling the voltage drop across an external, low-side, MOSFET synchronous rectifier. This topology eliminates the sense resistor's cost and pc-board-space requirement and also provides a modest increase in circuit efficiency. However, the MOSFET's highly temperature-dependent on-resistance dominates the current-limit value. Fortunately, certain newer dc/dc converters, such as Maxim's MAX1714, allow external adjustment of the current-limit threshold. The circuit in **Figure 1** shows how a thermistor applies temperature compensation to the circuit's output-current limit.

The MAX1714's linear current-limit (I_{LIM}) input range at Pin 6 of IC_1 spans 0.5 to 2V, which corresponds to current-limit thresholds of 50 to 200 mV, respectively. For the default current-limit setting, 100 mV, the circuit imposes a 7.5A current limit at 25°C. However, **Figure 2** shows that the current limit varies from 9A at −40°C to 6A at 85°C. To design the temperature-compensation network, begin by breadboarding the circuit and using an external power supply to vary the MAX1714's current-limit input voltage such that the output-current-limit value remains constant. You repeat the

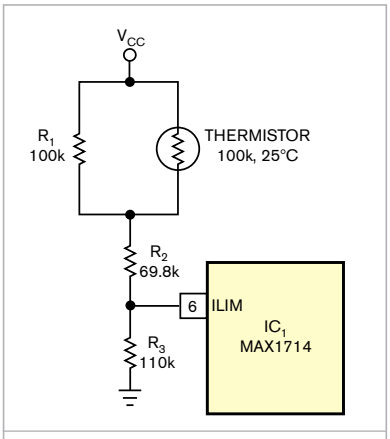

Figure 1 A thermistor-resistor network provides temperature compensation for a dc/dc converter's current-limit input, I_{LIM}.

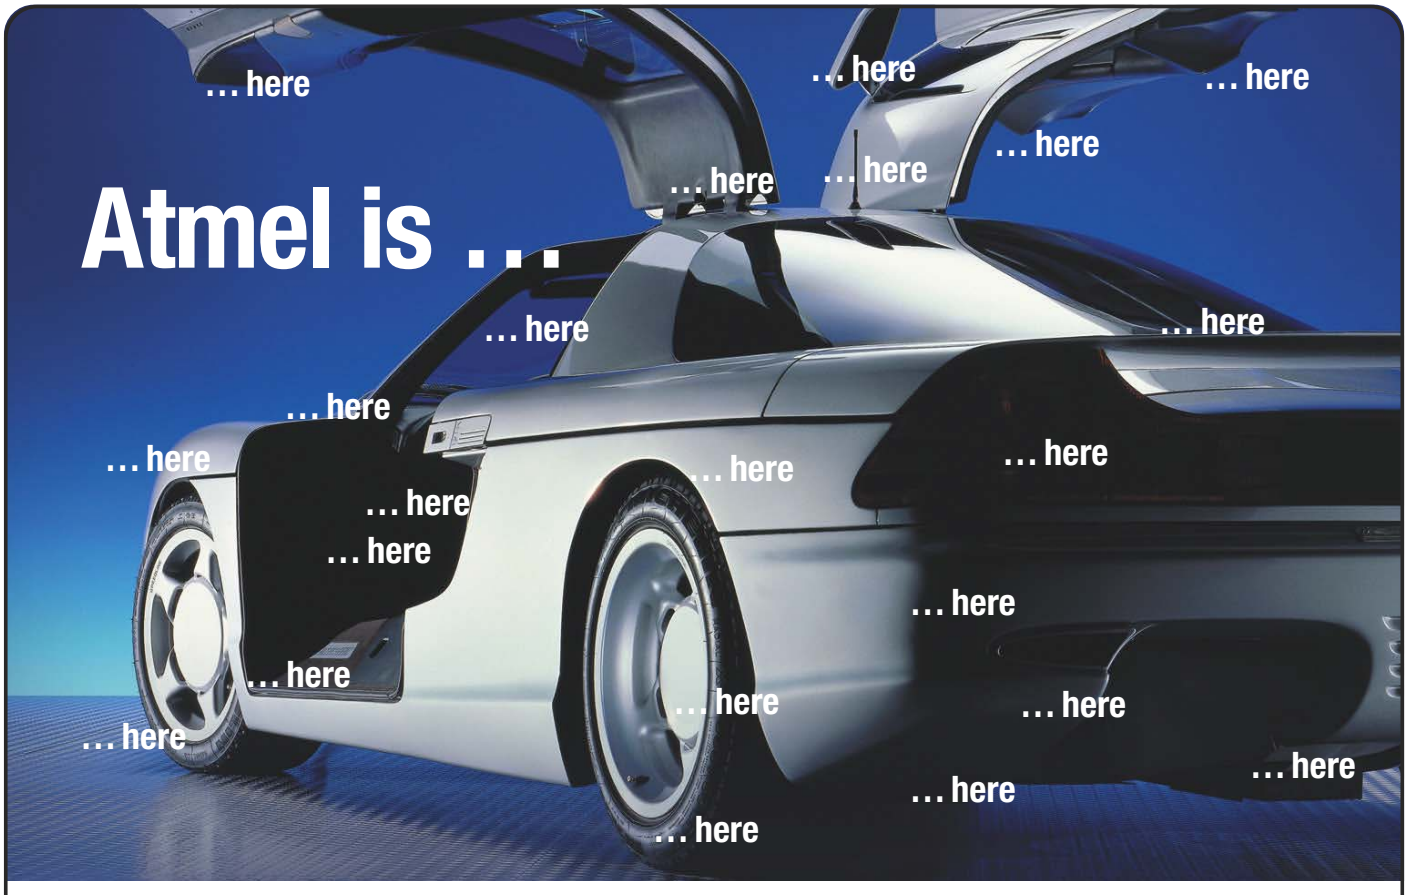

… Everywhere in Automotive

Take the easy road ahead to successful automotive designs with Atmel's dedicated IC solutions.

From body electronics, chassis, automotive security and safety, car infotainment to power train – Atmel, the automotive specialist, has more than 20 years of automotive electronics design expertise and is a market leader in numerous areas.

Our certified devices support key application areas from flashers, drivers, door modules to lighting, sensor acquisition, and networking, just to name a few. But that's not all! If you're looking for complete system solutions such as tire-pressure monitoring, car access, GPS systems and car radio, we're right there with you too.

Adding our high-performance microcontrollers, including AVR® and ARM®, and memories, Atmel can support your toughest requirements. We can also help you as a development partner with a broad range of dedicated ASICs based on Intellectual Property (IP).

When it comes to technology, we've really got you covered. With dedicated technologies and qualified fabs, our high-end products meet the strict automotive quality demands.

So whether ASSP or ASIC, Atmel is by your side every step of the way. Design-ins are supported by a full range of kits, software and other tools. If you are looking to improve your next automotive design, check out why Atmel is Everywhere in Automotive.

Check out Atmel's Automotive products today at www.atmel.com/ad/automotive

ASSPs (Chassis, Body Electronic, Security, Safety, Car Infotainment, Safety, Power Train)	**Microcontrollers** (ARM®, AVR®, MARC4, 80C51)
ASICs	**Networking** (CAN, VAN, LIN, FlexRay™)
Memories	**Sensors Acquisition**

Our fabs are qualified:
ISO 9001:2000
ISO TS 16949:2002
ISO 14001:1996

Everywhere You Are®

measurements at 10°C intervals over the circuit's operating-temperature range.

To compensate for IC_1's temperature variation, you can select from among several possible resistor-thermistor-network topologies. First, you need to select a suitable thermistor and characterize its resistance-versus-temperature variation. Because the MAX1714's current-limit input pin feeds a relatively high input-impedance voltage-follower stage, this thermistor requires a high nominal resistance of 100 kΩ. Resistance-versus-temperature characteristics of inexpensive thermistors exhibit considerable nonlinearity, but one relatively simple approach to linearization involves paralleling the thermistor with a fixed resistor equal to the thermistor's nominal resistance (**Reference 1**). In the network of **Figure 1**, R_1 linearizes the thermistor, and R_2 and R_3, respectively, set the slope and intercept of the current-limit-voltage-versus-temperature-characteristic curve.

To arrive at optimal values for R_2 and R_3, we prepared a spreadsheet incorporating the original current-limit-voltage-versus-temperature data and added columns for each of the network's resistors, plus the thermistor specification sheet's resistance-versus-temperature data. While observing the circuit's temperature-versus-voltage transfer function, we varied the spreadsheet's values for R_2 and R_3 until the transfer function best approximated the measured current-limit-voltage-versus-temperature data. Finally, we constructed the circuit and tested it over the temperature range and noted that it yielded a reasonably flat response.

The curvature of the corrected output characteristic of **Figure 2** (red trace) is intrinsic to the thermistor. Though not perfectly flat, the corrected curve represents a great improvement over the original (black trace) and is sufficient to meet the original design goal. You can achieve more precise compensation by selecting a different thermistor or by incorporating multiple thermistors. EDN

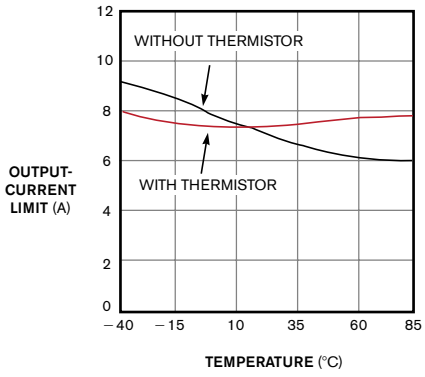

Figure 2 Before (black trace) and after (red trace) current-limit-versus-temperature characteristics show the performance enhancement that the circuit in Figure 1 provides.

REFERENCE

1 Horowitz, Paul and Winfield Hill, *The Art of Electronics*, ISBN 0 521 37095 7, Cambridge University Press, New York, 1980.

Add a Schmitt-trigger function to CPLDs, FPGAs, and applications

Stephan Roche, Santa Rosa, CA

Thanks to its internal hysteresis, the highly useful Schmitt-trigger circuit accepts a low-slew-rate input signal and produces a clean, glitch-free output transition. Unfortunately, user-programmable logic devices, such as CPLDs and FPGAs, generally offer no direct method of synthesizing Schmitt-trigger gates and buffers. This Design Idea shows how a few external components and some VHDL code can implement a Schmitt trigger and put it to work in several useful applications.

To create an equivalent of the basic Schmitt-trigger buffer, you use two external resistors to create positive feedback around a buffer (**Figure 1a** and **b**). You can also use four external resistors to set two threshold levels around an R-S flip-flop (**Figure 1c**).

The following equations, respectively, describe the basic Schmitt trigger's positive- and negative-threshold levels:

$$V_+ = \frac{R_1}{R_2} V_{CC} + V_{TH}\left(1-\frac{R_1}{R_2}\right);$$

$$V_- = V_{TH}\left(1-\frac{R_1}{R_2}\right).$$

In these **equations**, V_{TH} represents the input-voltage threshold of the CPLD/FPGA device, and V_{CC} is its power-supply voltage.

Based on the equivalent Schmitt-trigger circuit in **Figure 1b**, the low-cost resistance-capacitance oscillator in **Figure 2** requires four external passive components. Resistor R and capacitor C set the circuit's oscillation frequency. Note that the resistance values of R_1

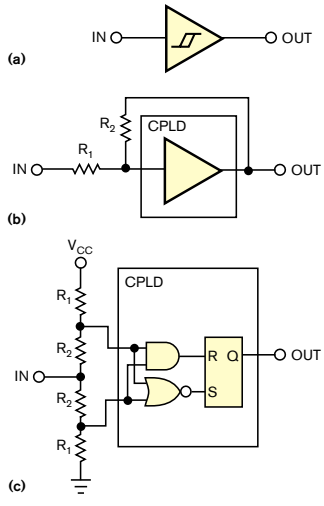

Figure 1 Use a portion of a programmable-logic device or gate array to implement a Schmitt-trigger buffer (a) by adding either two (b) or four external resistors (c).

INDUSTRY'S SMALLEST DUAL PCI EXPRESS HOT-PLUG CONTROLLER

Hot Swaps All Three Voltage Rails for Two PCI Express™ Cards

BEATS THE COMPETITION THREE-WAYS

1 REDUCES COST
- REQUIRES 21 FEWER EXTERNAL COMPONENTS
- USES LOWER COST n-CHANNEL MOSFETs vs. THE COMPETITION'S p-CHANNEL MOSFETs
- OPTIONAL EXTERNAL SMBus™ INTERFACE WITH I/O EXPANDER

2 SUPERIOR RELIABILITY AND FLEXIBILITY
- 7.5x LARGER GATE PULL-DOWN CURRENT FOR FASTER FAULT RESPONSE
- 30% LOWER MOSFET RDSON RESISTANCE IMPROVES EFFICIENCY AND REDUCES POWER DISSIPATION
- CHOICE OF LATCHED OR AUTO-RETRY FAULT MANAGEMENT vs. LATCHED FROM COMPETITION

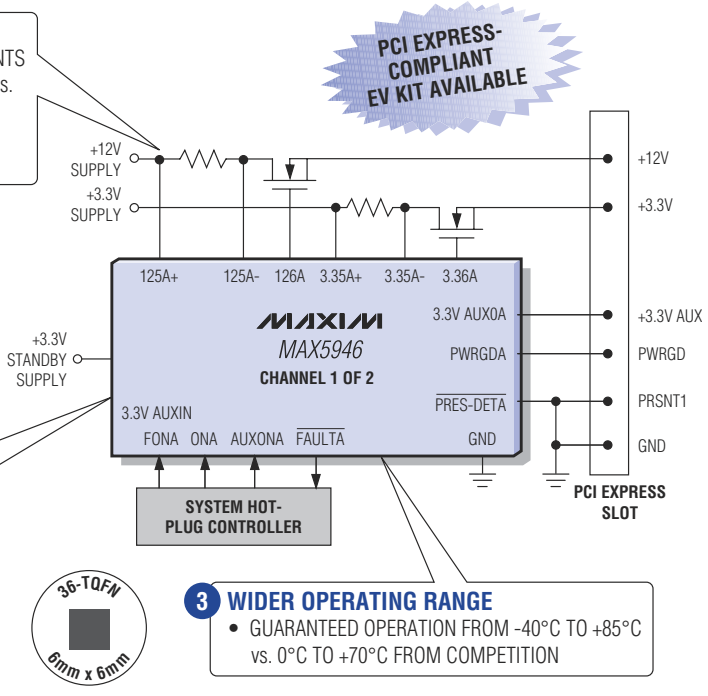

3 WIDER OPERATING RANGE
- GUARANTEED OPERATION FROM -40°C TO +85°C vs. 0°C TO +70°C FROM COMPETITION

- ♦ Hot Swaps +12V, +3.3V, and +3.3V Auxiliary Voltage Rails for Two Slots
- ♦ Integrated 0.24Ω Power MOSFETs for +3.3V Auxiliary Supply Rails
- ♦ Independent ON/OFF Control for Each Slot
- ♦ Active Current Limit Provides Overcurrent and Short-Circuit Protection for Each Slot
- ♦ Contact Factory for Competitive Pricing

PCI Express is a trademark of PCI-SIG.
SMBus is a trademark of Intel Corporation.

FOR ALL YOUR HOT-SWAP NEEDS, VISIT www.maxim-ic.com/HotSwap

MAXIM

www.maxim-ic.com

FREE* Power Supplies Design Guide—*Sent Within 24 Hours!

CALL TOLL FREE 1-800-998-8800 (6:00 a.m.–6:00 p.m. PT)
For a Design Guide or Free Sample

MAXIM/DALLAS DIRECT! DISTRIBUTION 1-888-MAXIM-IC
ARROW ELECTRONICS, INC. 1-800-777-2776
AVNET electronics marketing 1-800-332-8638

Distributed by Maxim/Dallas Direct!, Arrow, Avnet Electronics Marketing, Digi-Key, and Newark.

and R_2 must be larger than that of R. **Listings 1** and **2** contain the circuit's VHDL implementation and RTL architecture, respectively.

In **Figure 3**, an open-collector buffer provides the trigger for the basic Schmitt-trigger-retriggerable monostable circuit by discharging timing capacitor C. The circuit's output pulse width approximately equals the time constant RC. **Listing 3** shows the VHDL implementation and RTL architecture, respectively.

You can convert the retriggerable monostable into the nonretriggerable monostable in **Figure 4** by using an open-collector NAND gate to discharge timing capacitor C. As long as the circuit's output remains high during the timing interval, the system locks out external triggers. As in the previous circuit, the output pulse width approximately equals the time constant RC. **Listing 4** contains the VHDL and RTL codes.

You can use the basic CPLD buffer-with-feedback circuit to provide hysteresis for a contact-debouncing circuit. In **Figure 5**, resistor R_4 provides contact-cleaning current, and R_3 and C form a low-pass filter to reduce noise that contact bounce generates. Component values vary depending on the application. EDN

```
Entity Oscillator is
  Port (
        A : in std_logic;
        B : in std_logic;
        OUT : out std_logic
  );
end Oscillator;



architecture RTL of Oscillator is
begin
  A <= B;
  OUT <= not A;
end RTL;



Entity Monostable is
  Port (
        A : in std_logic;
        B : in std_logic;
        Trigger : in std_logic;
        C : out std_logic;
        OUT : out std_logic
  );
end Monostable;
architecture RTL of Monostable is
begin
  A <= B;
  OUT <= not A;
  C <= '0' when Trigger= '1' else 'Z';
end RTL;



  Port (
        A : in std_logic;
        B : in std_logic;
        Trigger : in std_logic;
        C : out std_logic;

          OUT : out std_logic
   );
 end Monostable;
 architecture RTL of Monostable is
 begin
   A <= B;
   OUT <= not A;
   C <= '0' when Trigger= '1' and A='0' else 'Z';
 end RTL;
```

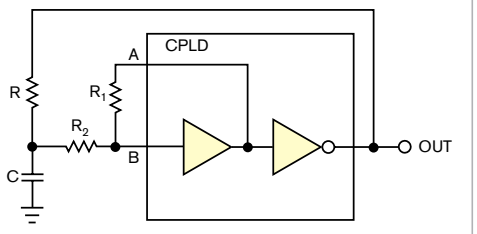

Figure 2 Add a resistor and capacitor to a basic Schmitt-trigger buffer to form a low-cost oscillator.

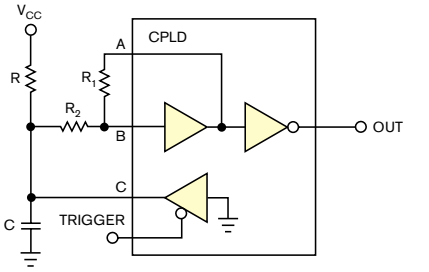

Figure 3 An active-low buffer holds timing capacitor C discharged in this version of a retriggerable monostable multivibrator based on the Schmitt-trigger buffer.

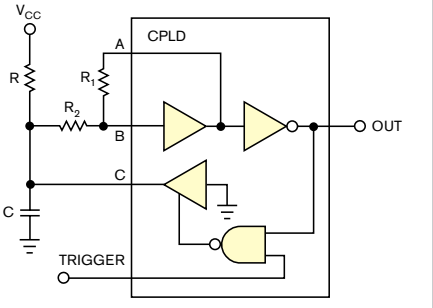

Figure 4 A NAND gate locks out trigger pulses, forming a nonretriggerable, monostable circuit.

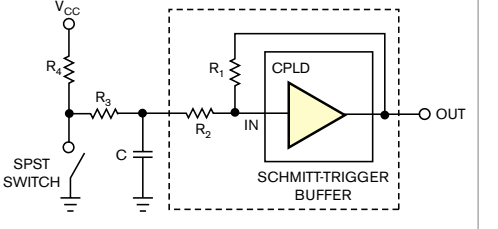

Figure 5 Use a Schmitt trigger to debounce switch contacts.

WORLD'S SMALLEST 3.5- AND 4.5-DIGIT PANEL-METER ICs WITH INTEGRATED LED DISPLAY DRIVER

Integrating and Autozeroing Capacitor Not Required

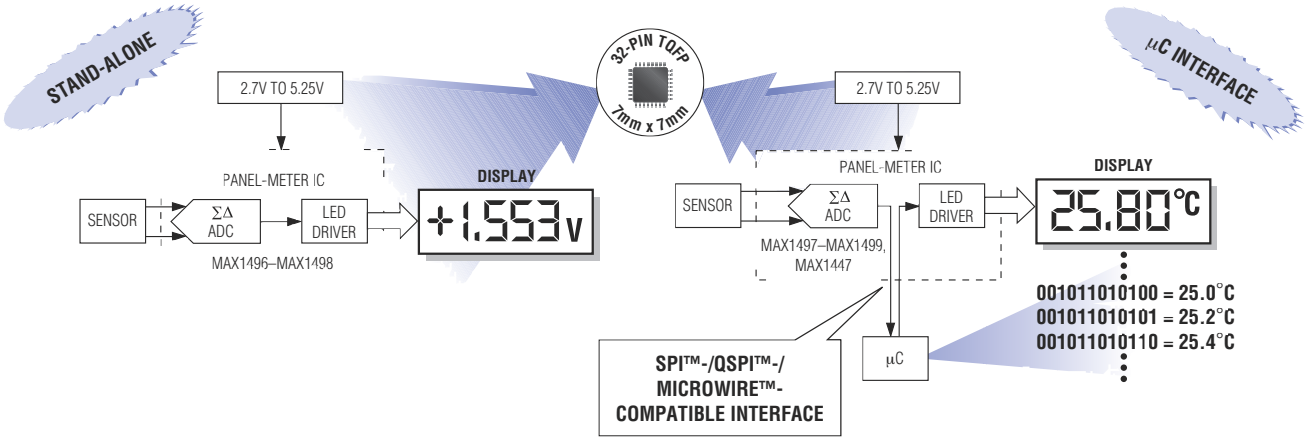

- ♦ Integrated Panel-Meter Solution
  - ♦ On-Chip Oscillator, Reference Input Buffer, LED Driver
- ♦ Programmable LED Drive Current
- ♦ Peak and Low-Battery Detection
- ♦ Hold Function

Part	Display	Digits	Resolution (Counts)	Interface	Pin-Package	Price† ($)
LED Display Panel-Meter ICs						
MAX1496	LED	3.5	±1999	Stand-alone	28-SSOP, PDIP	3.96
MAX1497	LED	3.5	±1999	µC	28-SSOP, PDIP	4.91
MAX1498/1447	LED	4.5	±19999	Stand-alone	32-TQFP	6.18
MAX1499	LED	4.5	±19999	µC	32-TQFP	7.13
LCD Display Panel-Meter ICs						
MAX1491	Triplex LCD	3.5	±1999	Stand-alone	28-SSOP, PDIP	3.96
MAX1492	Triplex LCD	3.5	±1999	µC	28-SSOP, PDIP	4.91
MAX1493/95	Triplex LCD	4.5	±19999	Stand-alone	32-TQFP	5.62
MAX1494	Triplex LCD	4.5	±19999	µC	32-TQFP	6.57

SPI and QSPI are trademarks of Motorola, Inc. MICROWIRE is a trademark of National Semiconductor Corp.
†1000-up recommended resale. Prices provided are for design guidance and are FOB USA. International prices will differ due to local duties, taxes, and exchange rates. Not all packages are offered in 1k increments, and some may require minimum order quantities.

For More Information, Visit www.maxim-ic.com/solutions/panel_meter/

MAXIM

www.maxim-ic.com

FREE* A/D Converters Design Guide—*Sent Within 24 Hours!

CALL TOLL FREE 1-800-998-8800 (6:00 a.m.–6:00 p.m. PT)
For a Design Guide or Free Sample

MAXIM/DALLAS DIRECT! DISTRIBUTION 1-888-MAXIM-IC

ARROW ELECTRONICS, INC. 1-800-777-2776

AVNET electronics marketing 1-800-332-8638

Distributed by Maxim/Dallas Direct!, Arrow, Avnet Electronics Marketing, Digi-Key, and Newark.

productroundup

DISCRETE SEMICONDUCTORS

Midvoltage rad-hard MOSFETs have improved device on-resistance

The R6 line of HiRel 100, 150, 200, and 250V radiation-hardened MOSFETs target hybrid military PRF-38534 Class K modules for satellite applications, such as low-, middle-, and Geostationary-earth orbit and deep-space missions. The product provides 100 to 250V, 0.010 to 0.130Ω on-resistance, a 16 to 63A maximum current, and a 35- to 240-nC Q_C range. These devices come in TO-254 and TO-257 and tablet packages. Prices for the R6 line begin at $290.63 each (1000).
International Rectifier, www.irf.com

Two high-current SCRs and three triac series now available

Two series of high-current SCRs (silicon-controlled rectifiers) with 40 and 50A ratings come in nonisolated packaging. The CYNB40 and CYNB55 feature 400 to 1000V blockage-voltage ratings, and the CYNB55 has a 700A rating at 60 Hz, full cycle. Three triacs are also available with 30A ratings in isolated packaging and 35A in nonisolated packaging. The products cost $1.25 (100).

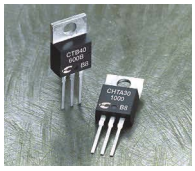

Crydom Corp, www.crydom.com

Small, low-voltage device features lightning protection

Protecting as many as four low-voltage interfaces against ESD, the 3.3V RClamp3304N TVS (transient-voltage-suppression)-protection device features protection levels that meet IEC 61000-4-2 ESD specifications at 15V on contact and 25 kV on air. With a 5-pF capacitance for high-speed interfaces, the device also provides 25A lightning protection at 8/20 µsec, as well as a low leakage current. Measuring 2.6×2.6×0.6 mm, the device comes in a leadless SLP packaging and complies with ROHS (reduction-of-hazardous-substances) and WEEE (waste-from-electrical/electronics-equipment) directives. The RClamp3304N costs 65 cents (1000).
Semtech, www.semtech.com

Low-leakage Schottky diode raises efficiency

With reverse-current ratings of 30V at 1 µA, the ZLLS350 Schottky barrier diode has a 380-mV forward voltage for a 30-mA forward current. The device also handles a 380-mA continuous forward current. With a 1.7×0.9 footprint in an SOD-523 package, the ZLLS350 costs $0.125 (10,000).
Zetex Semiconductors, www.zetex.com

EMI-filter/ESD-protection device has low leakage current

This six-line EM6D-100L lowpass-filter array integrates TVS (transient-voltage-suppressing) diodes, suppresses undesirable EMI/RFI signals, and provides ESD protection for high-speed data interfaces. A cutoff frequency of 150 MHz provides EMI/RFI attenuation of better than 25 dB in the 800-MHz to 3-GHz bandwidth. Meeting immunity-test standard IEC 61000-4-2 at all levels, the device features a leakage current of less than 100 µA. The EM6D-100L comes in a DFN-12 package and costs 50 cents (12,000).
ProTek Devices, www.protekdevices.com

MOSFET arrays support nanowatt-power circuit operation

This family of quad/dual N-channel, matched-pair MOSFET arrays achieves a 0.2V threshold voltage with ±0.02V tolerances and ±10-mV offset. The ALD110802/ALD110902, ALD110808/ALD110908, and ALD110814/ALD110914 devices have gate-threshold

You need this clock generator

CG635 – Precise, low jitter clocks from DC to 2.05 GHz

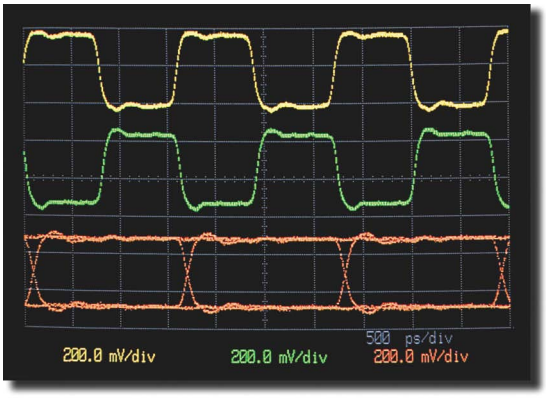

Clock and PRBS signals at 622.08 MHz

Plot shows complementary clock and PRBS (opt. 1) outputs at 622.08 Mb/s with LVDS levels. Traces have transition times of 80 ps and jitter less than 1 ps (rms).

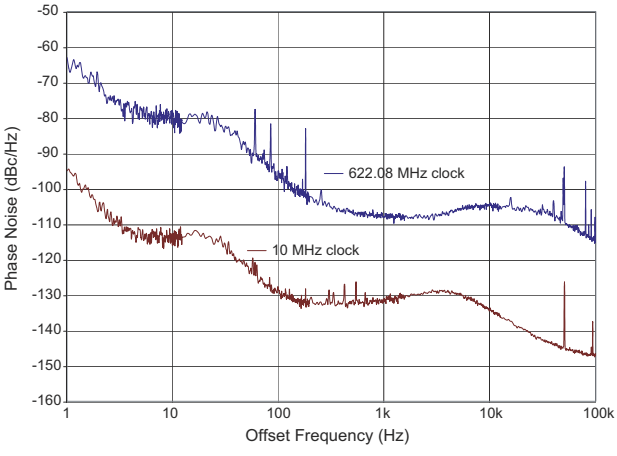

Phase noise for 10 MHz and 622.08 MHz outputs

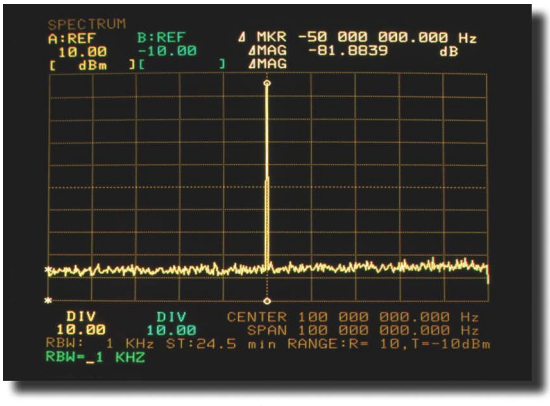

RF Spectrum of a 100 MHz clock

Graph shows a 100 MHz span around a 100 MHz clock. Only two features are present: the clock at 100 MHz, and the spectrum analyzer's noise floor (around –82 dBc).

- **Square wave clocks from DC to 2.05 GHz**
- **Random jitter <1 ps (rms)**
- **80 ps rise and fall times**
- **16-digit frequency resolution**
- **CMOS, LVDS, ECL, PECL, RS-485**
- **Phase adjustment and time modulation**

The CG635 Synthesized Clock Generator provides square wave clocks between DC and 2.05 GHz that are clean, fast and accurate. With jitter less than 1 ps, transition times of 80 ps, and 16 digits of frequency resolution, the CG635 will meet your most critical clock requirements.

The instrument can provide clocks at virtually any logic level via coax or twisted pairs. The outputs have less jitter than any pulse generator you can buy, with phase noise that rivals RF synthesizers costing ten times more.

Optional OCXO and rubidium timebases improve frequency stability by 100× and 10,000× over the standard crystal timebase. And an optional PRBS helps you evaluate high-speed serial data paths.

Whether you are trying to lower the noise floor of an ADC, increase SFDR of a fast DAC, or squash the bit error rate in a SerDes, the CG635 is the tool you need to get the job done.

SRS Stanford Research Systems, Inc.

Phone: (408)744-9040
www.thinkSRS.com

voltages of 0.2V±0.02V at 1 µA, 0.8V±0.02V at 1 µA, and 1.4V±0.04V, respectively. An A-grade version, the ALD110808A/ALD110908A, is also available with a 0.8V±0.01V at 1 µA. Available in PDIP and SOIC packages, the devices cost 67 cents (100).

Advanced Linear Devices, www.aldinc.com

Low-profile ESD-suppression device meets protection standards

Targeting portable devices, the µESD dual series of ESD-protection diodes clamps 30-kV ESD transients—per the IEC61000-4-2 standard—in less than 1 nsec. The devices

clamp ESD pulses to less than 7V and offer a leakage of 0.05 mA with a capacitance of 35 pF. Available in SOT-723 packages, the µESD3.3D and µESD5.0D cost 7 cents each (10,000).

On Semiconductor, www.onsemi.com

EMBEDDED SYSTEMS

Stand-alone radio features RS-232 or USB connectivity

Allowing quick connections to RS-232- or USB-enabled devices, the 2.4-GHz Xbee-Pro stand-alone radio is compatible with networks operating on Xbee and Xbee-Pro technology. The radio is made of an aluminum housing and a 2.1-dBi dipole antenna. Its peripherals include a 100-mW RF module with a transmitting range of 0.9 miles in line-of-sight conditions, 250-kbps communication, and ZigBee-ready capability; it accepts compliant networking-protocol upgrades. The RS-232 and USB radios cost $99 each.

MaxStream, www.maxstream.net

Low-cost development-tool suite adds upgrades

Targeting development of software applications and systems for the Java 2 Enterprise Edition and other environments, Version 4.0 of the MyEclipse Enterprise Workbench enterprise-class platform and tool suite supports 25 application-server connectors, including WebLogic 9.0 and WebSphere 6.0. Extended features include MyJSF (Java Server Faces) developer with support for Sun JSF RI 1.1 and MyFaces 1.0.9 JSF implementations, Faces Configuration Designer, integration with MyEclipse Hot-Sync Deployment, and WAR Expert Tools. The upgrade also features Jakarta Tapestry; the Tapestry-aware template editor; Tapestry project validations; and the Integration Enhanced IDE Oracle Con-

Triple-play. IMS. Converged services. **Profit?**

profitable

OPEN MODULAR SOLUTIONS

ACCESS

EDGE

CORE

TRANSPORT

DATA CENTER

DESIGN AND DEPLOY

Your new IMS infrastructure applications for the next generation wireless network using Kontron ATCA / AMC modular solutions.

Kontron simply takes the worry – and the expense – out of building complex IMS communication platforms for **next generation 3G wireless networks.** Whatever the application, your project is designed and deployed in a hearbeat with fully integrated, open standard modular solutions that are application-ready, right off the shelf. That means reduced development costs for you, and tremendous "swap in - swap out" service flexibility for your carrier customers. It's a very **profitable** win-win go-to-market strategy for everything from data and signaling platforms to IP streaming multimedia applications for video-on-demand, real-time voice and video telephony. It's so simple. It's the way of the wireless future. **Open.**

> Go Open Standards > Go Kontron > Ask for an Eval today >

kontron.com/goATCA

1-888-526-ATCA

EMEA: +49 0 800 7253756
ASIA: +886 2 2910 3532
sales@us.kontron.com

AdvancedTCA™ Intel® Communications Alliance Premier Member GOLD solid. montavista

Computer On Modules | **Blades & Mezzanines** | CPU Boards | **Systems** | Mobile Rugged | Custom Solutions

kontron

nector, including a procedure-runner tool. MyEclipse 4.0 costs $29.95.
Genuitec LLC, www.genuitec.com

Single-board computer features a Pentium 4 CPU

Targeting industrial applications, the ATX-807 single-board computer Mini ITX form factor suits compact chassis. It includes a 90-nm Pentium 4 Prescott CPU with selectable 400-, 533-, or 800-MHz front-side-bus speed and Intel VRD 10.1 compliance for support of upgraded advanced processors. The Pentium 4 processor has hyperthreading technology, two DIMM sockets supporting 2 Gbytes of DDR memory, and onboard dual Intel Ethernet controllers allowing Gigabit and 10/100-Mbit speeds. The unit's RJ45 connector has a built-in LAN LED, which displays speed and link activity, as well as an available external connector, displaying the LAN status on the front panel of the computer chassis. The package, including a support hardware monitor and watchdog timers to alert administrators of abnormal operation, costs $380.
Arista Corp, www.aristaipc.com

PMC card features low- or ultralow-voltage processors

The PSL09 PrPMC high-performance-processor module includes either a 1.4-GHz Intel Pentium M processor or an ultralow-voltage Intel Celeron M processor. The module includes an Intel 855GME graphics-memory-controller hub and an Intel 6300ESB I/O-controller-hub chip set. The device operates at −40 to +55°C and supports Linux, VxWorks, and Windows XP. The PSL09 PrPMC costs $2140.
SBS Technologies, www.sbs.com

MICRO-PROCESSORS

Microcomputer has low power consumption

Suited for watches and clocks, the ultralow-power S1C63708 microcomputer features a 4-bit CMOS core. The device integrates an 8192-word×13-bit ROM, a 1024-word×4-bit RAM, a serial interface, a programmable PWM timer, and a sound generator, as well as a built-in stepping-motor driver and LCD driver.
Seiko Epson Corp, www.epson.co.jp/e/

Evaluation kit includes Linux evaluation board

This evaluation kit for the PowerPC 440EP processor includes the Yosemite evaluation board with a Linux configuration but is flexible enough to

Intersil Switching Regulators

Intersil High Performance Analog

Running Out Of Room? We Suggest Fewer Chips.

Reduce your board size and increase efficiency with Intersil's multiple output DC/DC Switching Regulators. Get up to 20A of output current and four regulated outputs from a single V_{IN} ranging from 3.3V to 24V.

Design has never been so easy, with a variety of combinations of linears and PWMs, Most devices are available in thin scale packages to maximize real estate.

Intersil's Family of Multiple Output DC-DC Solutions

# of Outputs	Device	Regulators PWMs	Regulators Linears	Int. FETs	V_{IN} (V)	I_{OUT} (max) (A)	Package
4	ISL6521	1	3		5	20	SOIC-16
3	ISL6455	1	2	✓	3.3	0.6	QFN-24
	ISL6455A	1	2	✓	5	0.6	QFN-24
	ISL6537	2	2 + Ref		2.5, 12	20	QFN-28
	ISL6532A	1	2		5, 12	20	QFN-28
	ISL6441	2	1		4.5 to 24	6	QFN-28
	ISL6443	2	1		4.5 to 24	10	QFN-28
2	ISL6227	2	0		4.5 to 24	16	SSOP-28
	ISL6440	2	0		4.5 to 24	10	QSOP-24
	ISL6539	2	0		5 to 15	8	SSOP-28
	ISL6530/1	2	Ref		5	1	SOIC-24, QFN-32
	ISL6528	1	1		3.3, 5	15	SOIC-8
	ISL6529	1	1		3.3 to 5, 12	15	SOIC-14, QFN-16

Datasheet, free samples, and more information available at **www.intersil.com/edn**

Intersil – Switching Regulators for precise power delivery.

intersil®

HIGH PERFORMANCE ANALOG

load and run other PowerPC-compatible embedded operating systems. The 5×7-in.-form-factor board features an AMCC 440EP processor with a 533-MHz clock frequency. Additional peripherals include 256 Mbytes of SDRAM, 32 Mbytes of flash, two 10/100 Ethernet ports; a USB 2.0 port and two USB 1.0 ports; JTAG, trace, and PCI host connectors; a 2.6 Linux kernel in flash; boot firmware in flash; and Kozio's board-diagnostics suite.
Applied Micro Circuits Corp, www.amcc.com

IPsec software encrypts and safeguards data

Integrating with the IPv4 (Internet Protocol Version 4) layer and IPv6 layer of the Nucleus Net TCP/IP stack, the Nucleus IPsec (Internet Protocol-security) software includes protection mechanisms, such as data-origin authentication, data integrity, data confidentiality, antireplay protection, and limited traffic-flow confidentiality. The software conforms with the IETF (Internet Engineering Task Force) IPsec specifications and provides support for the IKE (Internet-key-exchange) protocol. A MIB (management-information database) allows for remote configuration and monitoring of the module. The software comes in source-code format; a license costs $19,995 without royalty fees.
Accelerated Technology, www.acceleratedtechnology.com

Software supports all Intel Itanium 2-processor-based platforms

Supporting the Intel Itanium 2-processor-based computing platform, QuickTransit software allows applications for one processor and operating system to run on another processor and operating system without source-code or binary changes. The software is available as a component of the Silicon Graphics Prism family of visualization systems.
Transitive Corp, www.transitive.com

CAN software integrates support tools

CANopenRT integrates device-driver support and includes all functions necessary for implementing slave or simple master devices in accordance with CANopen specification DS-301, Version 4.02. Supported processors include Analog Devices' Blackfin BF536, BF537, and BF534; Freescale's Coldfire MCF5214/6X, MCF523X, MCF528X, and MCF548X; and Infineon's C16X. Tested with CANopen conformance-test software, CANopenRT and CAN drivers are available in C source code. Licenses cost $1900 for basic CAN drivers.
Quadros Systems, www.quadros.com

© David Corio

"IEEE is a magnet bringing engineers together. Engineering is a language that transcends bureaucracy and nationalism. It gives to us all a common language and understanding."

Ya-Qin Zhang, IEEE Fellow
Managing Director, Microsoft Research Asia

IEEE

Build relationships in a worldwide community of innovators. Join IEEE.
www.ieee.org/member

EDN productmart

This advertising is for new and current products.

EDN ADVERTISER INDEX

This index is provided as an additional service. The publisher does not assume any liability for errors or omissions. For immediate information on products and services, go to Reader Service under Tools & Services at www.edn.com.

EDN literature link

"The engineering professional's link to technical literature"

RAPIDLY CREATE TEST, MEASUREMENT AND CONTROL SYSTEMS WITH LABVIEW

NI LabVIEW is the graphical development environment for creating flexible and scalable test, measurement, and control applications, rapidly and at minimal cost. With LabVIEW, engineers and scientists interface with real-world signals, analyze data for meaningful information, and share results and applications. Take a FREE online guided tour. Visit ni.com/info and enter **ebpk4g**

National Instruments
Tel: (800) 453-6202 (U.S. and Canada) (512) 794-0100
Fax:(512) 683-9300 Email: info@ni.com

TEST & MEASUREMENT ACCESSORIES

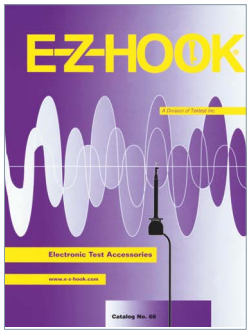

Broad line of over 12,000 high quality test and measurement products that solve application problems in the electrical, electronic, industrial, automotive and medical fields. Product line includes test leads, adaptors, connectors, DMM kits, probes, nail clips, cable assemblies, and coaxial test accessories. Custom manufacturing is available. Selling direct, s ince 1956.

E-Z-HookW
800-995-HOOK (4665)
Fax: 626-446-0972
Web: www.e-z-hook.com
Email: info@e-z-hook.com

BATTERIES SHORT FORM CATALOG WITH TECHNICAL DATA ON CD-ROM

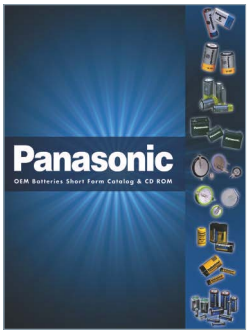

Panasonic's broad line of battery sizes and chemistries is condensed for review in this new 16-page, 4-color catalog with technical data and co mpany overview on enclosed CD-ROM. Batteries covered include NiCd, NiMH, Lithium, Li-ion, VRLA and Alkaline. Easy-to-use tables give technical specifications for model numbers.

Toll Free: 1-877-726-2228
oembatteries@us.panasonic.com
www.panasonic.com/batteries

EXPERTISE IN DFM / DFT

- Reduce production costs
- Improve product quality
- Speed time to market
- Simplify your design process

That's what can be achieved through DFM / DFT methods and procedures.
For more information visit
www. mfg.com/request/

Electronic Technologies International, Inc.

MILL-MAX'S DESIGN GUIDE NOW AVAILABLE ON CD ROM

Mill-Max, is the leading US manufacturer of precision-machined interconnect components. CD includes our most up-to-date:

- Catalog PDF Pages
- New Product Information
- Plant Tour
- Links to our Website

Visit www.mill-max.com/respond
Response Code: EDN555

EMBEDDED HARDWARE SOLUTIONS

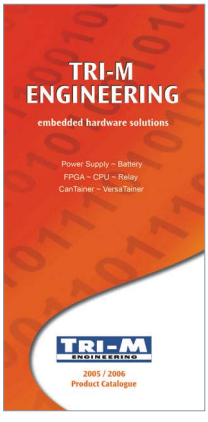

Tri-M Engineering continues to maintain a leading role in the creation of hardware solutions for Embedded Systems. From our Port Coquitlam facility, Tri-M specializes in the design and production of DC to DC Power Supplies, Relay Boards, CPU Boards and Hostile Environment Enclosures, on time and on budget.

Tri-M Engineering
1407 Kebet Way, Unit 100, Port Coquitlam
British Columbia, Canada, V3C 6L3
www.tri-m.com e-mail: info@tri-m.com
Toll Free in Canada and USA 1-800-665-5600
Phone: 604-945-9565 Fax: 604-945-9566
call or e-mail to request a catalogue

GRAPHIC OVERLAYS & MEMBRANE KEYPADS

Amazing design options are now possible with NFi's unique full color applied graphic products. With our Design Guide, you'll discover special features to enhance the function and appearance of your products! Receive samples from an ISO Certified company:

- Filtered LED/LCD Windows
- Serial Numbers • UL/CSA Markings
- Urethane Doming • Keypad Embossing
- Laser Cutting • Prototyping

NFi – Nameplates for Industry, Inc.
Ph: 800-999-8900 Fax: 508-995-0099
Email: info@nfi.name www.nfi.name

SURGE PROTECTION

- COAXIAL Surge Suppressors for PCS, GPS, RF Equipment
- AC Protector UL1449
- DC/DATA Line Protectors

- Surge Arrester Gas Tubes
- Voltage from 75V To 2500V
- Available in Surface Mount

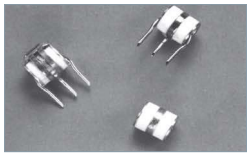

CITEL, Inc.
800-248-3548 FAX: 305-621-0766
www.citelprotection.com

FIBERGLASS LAMINATED EPOXY 1557C READY-TO-USE SHAPES

Design data pamphlet features materials, properties, and tolerances for glass epoxy components. It shows designers how to specify from open stock tools, their critical requirements for potting forms, bobbins, coil forms, structurals, and circuit board manufacturing aids. It is ideal for any electronics manufacturer looking for high reliability glass epoxy shapes.

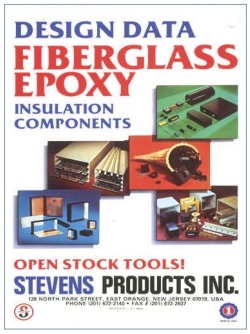

Stevens Products, Inc.
128 N.Park St., E. Orange, NJ 07017
973-672-2140
FAX: 973-672-2837

AC POWER DISTRIBUTION & CONTROL SYSTEMS

Product is rack or strip mounted with: Circuit Breaker Protection, EMI/RFI Filtering, Spike/Surge Protection, Remote Power On/Off, Remote Reboot, Emergency Shutdown, Redundant Power, Sequential Power Up/Down and power regulation Agency systems. ISO 9001:2000 Certified.

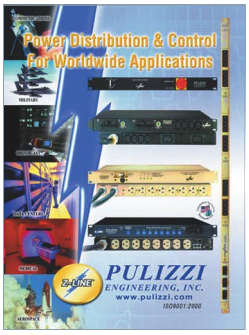

Pulizzi Engineering, Inc.
3200 S. Susan St., Santa Ana, CA 92704-6839
Sales: 605-334-8999 or 800-870-2248
Fax: 605-334-2611 E-mail: sales@pulizzi.com
Web Site: http://www.pulizzi.com

DESIGN GUIDE FOR OPTICAL ENCODERS

This Guide provides a comprehensive summary of optical encoders, accessories and interface electroNics. Included are design examples, absolute and incremental encoders, shafted and hollow-sectionsincludehazardousenvironments. Call 1-800-ENCODER or visit our web site at www.beiied.com.

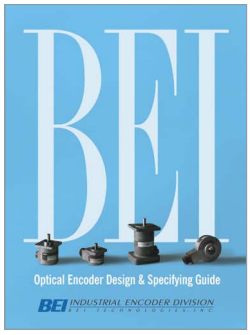

SINGLE BOARD COMPUTERS

PC/104-*Plus*, EPIC, and EBX embedded computers feature processors from 586 to Pentium M. Ideal for OEM applications. Benefits include extended temperature operation, industrial I/O, 5-year product availability guarantee and outstanding warranties. Voted a "Platinum" embedded board vendor (VDC surveys, 2002-2005). Request a Free 2005/2006 catalog.

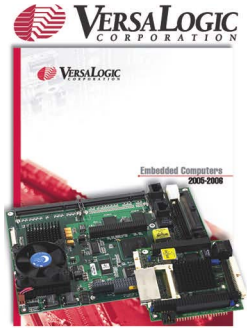

VersaLogic Corp.
Toll-free: (800) 824-3163
E-mail: info@VersaLogic.com
Web site: www.VersaLogic.com

MULTI-AXIS, MULTI-MOTOR MOTION CONTROL ICS

PMD, the leader in IC-based motion, introduces the Magellan™ Family of Motion Processors. Available in 1- 4-axis versions these flexible, programmable chips control DC brush, brushless DC, microstepping and pulse & direction motors. Magellan was specifically designed to lower your production cost, improve your performance and simplify your design. Contact us at 781-674-9860, or visit www.pmdcorp.com for more information.

NEW THYRISTOR DEVICE DATABOOK

Central's new 2005 Thyristor Databook features SCRs and TRIACs in leaded and SMD packages. This databook includes an extensive cross reference, selection guide, detailed datasheets and package specifications. Also featured is Central's new Extended Range of SMD Thyristors. To Order: www.centralsemi.com/ catalog

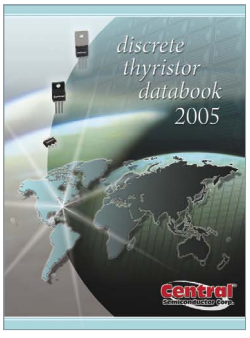

Central Semiconductor Corp.
145 Adams Avenue
Hauppauge, New York 11788 USA
631.435.1110 Fax: 631.435.1824
Email: salesinquiry@centralsemi.com
Web: www.centralsemi.com

AUTOMATIONDIRECT 2005 CATALOG

AutomationDirect's 2005 free catalog offers 5500 industrial automation products for a wide range of applications. The expanded 1750-page volume features product descriptions, technical data, photos, and price lists for PLCs, AC drives/motors, operator interfaces, sensors, encoders, and more.

For more information, visit http://www.automationdirect.com or call 800-633-0405.

AutomationDirect

***EDN* BUSINESS STAFF**

PRESIDENT, BOSTON DIVISION/PUBLISHING DIRECTOR, *EDN* WORLDWIDE
Stephen Moylan, smoylan@reedbusiness.com; 1-781-734-8431; fax: 1-781-290-3431

ASSOCIATE PUBLISHER, *EDN* WORLDWIDE
John Schirmer, jschirmer@reedbusiness.com; 1-408-345-4402; fax: 1-408-345-4400

OFFICE MANAGER
Rose Murphy, MurphyS@reedbusiness.com; 1-781-734-8457; fax: 1-781-290-3457

NORTHERN CA/ SILICON VALLEY
Patti Sellman
patti.sellman@reedbusiness.com
1-408-345-4439, fax: 1-408-345-4400

NORTHERN CA/SILICON VALLEY ID/NV/OR/WA/BRITISH COLUMBIA
Alan Robinson
aarobinson@reedbusiness.com
1-408-345-4450
fax: 1-408-345-4400

EASTERN CANADA/GA/IA/IL/ IN/KS/KY/MI/MN/MO/NC/ ND/NE/OH/WESTERN PA/ SC/SD/TN/VA/WI
Jim Dempsey
jim.dempsey@reedbusiness.com
1-440-333-3040
fax: 1-440-333-3044

CT/DE/MA/MD/ME/NJ/NY/RI/ EASTERN PA/VT
Joe McCabe
jmccabe@reedbusiness.com
1-781-734-8433
fax: 1-781-290-3433

SOUTHERN CA/AZ/CO/NM/UT
Leah Vickers
ljvickers@reedbusiness.com
1-562-598-9347
fax: 1-214-220-5420

AL/AR/FL/LA/MO/OK/TX
Roger Buckley
rbuckley@reedbusiness.com
1-972-216-5104
fax: 1-972-216-5105

DIRECTOR OF ONLINE SALES
Barbara Couchois
barbara.couchois@reedbusiness.com
1-408-345-4436

ADVERTISING SERVICE COORDINATOR
Nancy McDermott
nmcdermott@reedbusiness.com
1-781-734-8130
fax: 1-781-734-8086

DISPLAY ADVERTISING, SUPPLEMENTS, PRODUCT MART, INFO CARDS, LIT LINK, AND RECRUITMENT
Dara Juknavorian
dara.juknavorian@reedbusiness.com
1-781-734-8343
fax: 1-781-290-3343
Judy Keseberg
judy.keseberg@reedbusiness.com
1-781-734-8547
fax: 1-781-734-3547

WEB
Cathy Baldacchini, Director
cbaldacchini@reedbusiness.com
Jessica Mason
jmason@reedbusiness.com

VP OF INSIDE SALES
Jason Miner
jason.miner@reedbusiness.com
1-781-734-8545
fax: 1-781-290-3545

VP OF INTERNATIONAL SALES
Mike Hancock
mike.hancock@rbp.co.uk
+44 208-652-8248
fax: +44 208-652-8249

UK/NORWAY/SWEDEN/ DENMARK/FINLAND/ NETHERLANDS/BELGIUM/ LUXEMBOURG
John Waddell
jwadds@compuserve.com
+44 208-312-4696
fax: +44 208-312-1078

AUSTRIA/GERMANY
Adela Ploner
adela@ploner.de
+49 8131 366 99 20
fax: +49 8131 366 99 29

ISRAEL
Asa Talbar
talbar@inter.net.il
+972-3-5629565
fax: +972-3-5629567

ITALY
Roberto Laureri
media@laureriassociates.it
+39 02-236-2500
fax: +39 02-236-4411

SWITZERLAND
Gino Baretella
baretella@exportwerbung.ch
+41 1880-3545
fax: +41 1880-3546

FRANCE/PORTUGAL/SPAIN
Alain Faure
Alain.Faure@wanadoo.fr
+01 53 21 88 03
fax: +01 53 21 88 01

JAPAN
Toshiyuki Uematsu
t.uematsu@reedbusiness.jp
+81 3-5775-6057

SOUTH KOREA
Andy Kim
andy.kim@rbi-asia.com
+822 6363 3038
fax: +822 6363 3034

SINGAPORE, MALAYSIA, THAILAND
Chen Wai Chun
waichun.chen@rbi-asia.com
+65 6780 4533
fax: +65 6787 5550

TAIWAN
Charles Yang
+886 4 2322 3633
fax: +886 4 2322 3646

AUSTRALIA
David Kelly
david.kelly@rbi.com.au
+61 2-9422-2630
fax: +61 2-9422-8657

HONG KONG
Simon Lee
simonlee@rbi-asia.com.hk
+852 2965-1526
Dolf Chow
dolfchow@rbi-asia.com.hk
+852 2965-1531

MARKETING MANAGER, *EDN* WORLDWIDE
Wendy Lizotte
wlizotte@reedbusiness.com
1-781-734-8451
fax: 1-781-290-3451

DIRECTOR OF CUSTOM PUBLISHING
Cindy Fitzpatrick
cfitzpatrick@reedbusiness.com
1-781-734-8438
fax: 1-781-290-3438

CPS/WEB PRODUCTION COORDINATOR
Heather Wiggins
hwiggins@reedbusiness.com
1-781-734-8448
fax: 1-781-290-3448

ADMINISTRATION
John Blanchard
Vice President of Manufacturing
Norm Graf, Creative Director
Gloria Middlebrooks
Graphic Production Director
Dorothy Buchholz
Group Production Director

DESIGN
Dan Guidera
Senior Art Director/Illustration

RESEARCH DIRECTOR
Rhonda McGee
rmcgee@reedbusiness.com
1-781-734-8264,
fax: 1-781-290-3264

CIRCULATION MANAGER
Jeff Rovner
jrovner@reedbusiness.com
1-303-470-4477

REED BUSINESS INFORMATION
Jim Casella
Chief Executive Officer
Stephen Moylan
President, Boston Division
John Poulin
Senior Vice President, Finance
Sean T Keaveny
Vice President,
Finance, Boston Division

MOVERS and SHAKERS 2005

Sixth Edition

LEADERS

CUSTOMERS

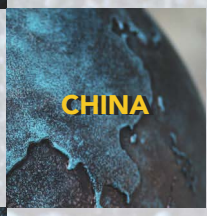

INNOVATORS

OEMs/ODMs

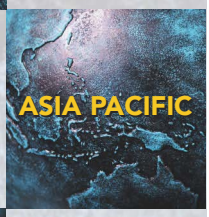

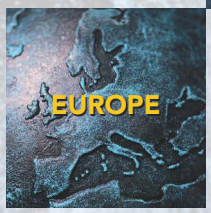

CEOs

Reed Electronics Group Presents the Leaders in Electronics.

Sponsor partners thus far:

Reed Electronics Group is the largest and most diverse organization in the world devoted to gathering and distributing information for the global electronics industry. Reed Electronics Group is part of Reed Elsevier plc group (NYSE: RUK and ENL), a world-leading publisher, trade-show producer, and information provider in print and online, with over 38,000 employees worldwide, operating in the science and medical, legal, education, and business-to-business industry sectors.

www.reed-electronics.com/moversandshakers

COMING NOVEMBER 2005

realitycheck

YESTERDAY'S HYPE MEETS TODAY'S REALITY

STATS First cell-phone call: April 3, 1973 / First commercial cell phone: Motorola DynaTRAC 8000x

Cell phones shrink, their market explodes

When Martin Cooper of Motorola made the first cell-phone call from a street in New York to his rival at Bell Labs on April 3, 1973, he used a bricklike prototype that weighed 30 oz and implemented the AMPS (analog-mobile-phone-system) format. The first commercial cell phone, the Motorola DynaTRAC 8000X (see photo) hit the market 10 years later; it weighed 28 oz, offered 30 to 60 minutes of talk time and eight hours of standby time, and retailed for $3995. By 1990, there were about a million subscribers in the United States; by 1994, that number had reached 24 million, and by 2004, it had reached 190 million, according to CTIA Wireless (www.ctiawireless.com).

But cell-phone systems need much more than a handset: They need base stations that link to each other and to the land-line system, through the Mobile Telephone Switching Office, plus a complex network-management infrastructure that ensures proper cell hand-off and billing. In 1994, there were 18,000 cell sites, compared with more than 175,000 now.—by Bill Schweber

Cutting-Edge Technologies for Wireless Solutions

AEROCOMM
MICROCHIP
NEC CEL
Chipcon
PARALLAX INC
DLP Design
RABBIT SEMICONDUCTOR
FAIRCHILD SEMICONDUCTOR
Radiotronix
freescale semiconductor
SILICON LABORATORIES
JRC NJR CORPORATION
VISHAY Semiconductor
LANTRONIX
ZiLOG

Delivering the latest with the greatest.

For the latest cutting-edge technologies for all your wireless applications, turn to Mouser. We are the catalog distributor focused on the needs of design engineers. We deliver the latest products with the great service and support you deserve.

SEMICONDUCTORS | PASSIVES | INTERCONNECTS | POWER | ELECTROMECHANICAL | TEST, TOOLS & SUPPLIES

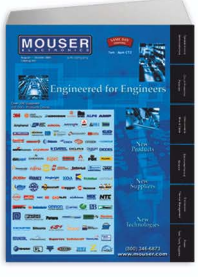

Supplier Authorized Distributor

New Products
New Technologies
New Suppliers
New Catalog Every 90 Days!

MOUSER ELECTRONICS
a tti company

(800) 346-6873
www.mouser.com

Motor Control

Solutions

powerSPIN™ Monolithic Motor Control ICs: simple-to-use, open, scalable single-chip platform that fits into a wide spectrum of motion-control-based applications

Innovative products for multi-segment application systems

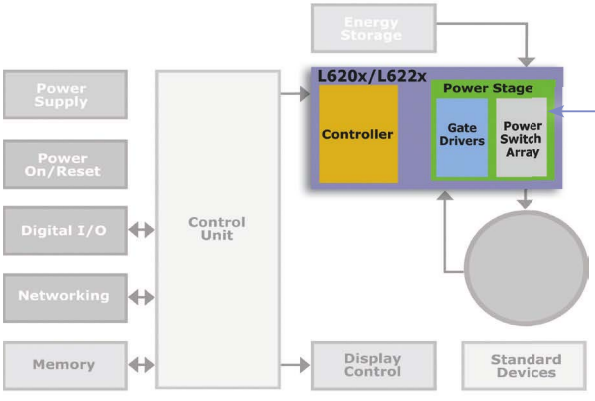

ST's global solutions make designing a flexible motor controller easier than ever with a portfolio of systems-on-chip, driver ICs, communication ICs and a wide range of advanced discrete devices. The powerSPIN™ family for example, combines control logic and power stages on the same chip, reducing the CPU load and making motor control design so much easier for all DC, stepper and brushless motor applications.

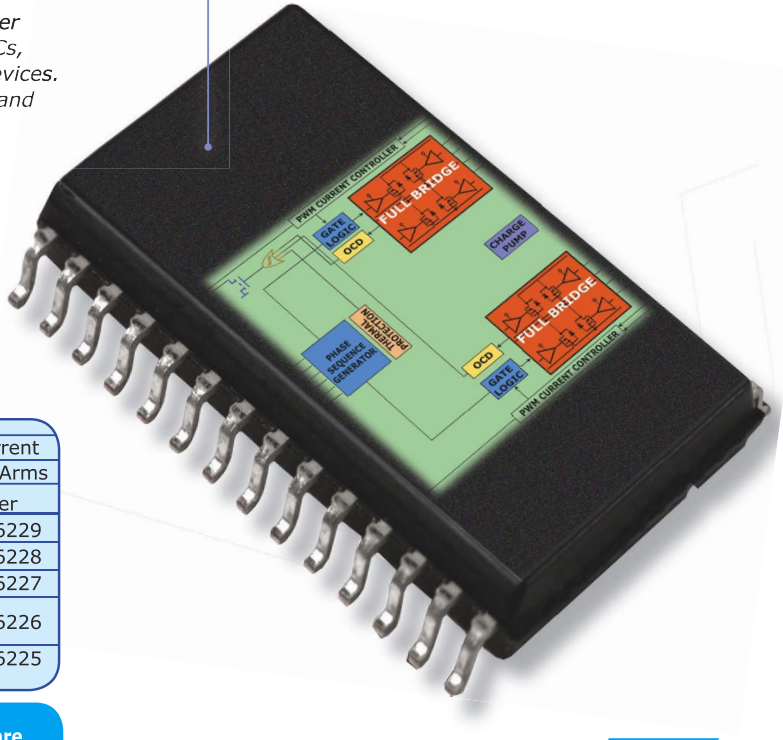

powerSPIN Key Features

- 60V-rated DMOS power stages
- $R_{DS(on)}$ = 0.3Ω for 2.8Arms L620x series
- $R_{DS(on)}$ = 0.73Ω for 1.4Arms L622x series
- Extensive built-in diagnostics

powerSPIN Family

Function	Output Current 2.8Arms	Output Current 1.4Arms
	Part Number	
Three-Phase brushless DC motor driver	L6235	L6229
Bipolar stepper motor driver with translator	L6208	L6228
Dual full bridge with twin PWM current control	L6207	L6227
Dual full bridge with user-adjustable overcurrent limit	L6206	L6226
Dual full bridge with fixed overcurrent limit	L6205	L6225

For datasheets, reference designs, application boards, software and application notes visit www.st.com/powerspin

Solutions @

™ powerSPIN is a trademark of STMicroelectronics.

DEVICE AVAILABLE FROM THESE DISTRIBUTORS: ARROW • AVNET • DIGI-KEY • FUTURE • MOUSER • NU-HORIZONS