UBM Electronics

EDN®

VOICE OF THE ENGINEER

OCT 6
Issue 19/2011
www.edn.com

TIME CAREFULLY THE ADOPTION OF YOUR NEXT-GENERATION PROCESSOR

Page 30

OPTIMIZE ENERGY-EFFICIENT ETHERNET PERFORMANCE IN BUNDLED LINKS

Page 23

Hello future.

Oscilloscopes Redefined

Starting at
$1,230*

	Agilent 2000 X-Series (MSO and DSO)	Popular Competitor's Series (DSO)	Agilent 3000 X-Series (MSO and DSO)	Popular Competitor's Series (MSO and DSO)
Bandwidth (MHz)	70, 100, 200	50, 70, 100, 200	100, 200, 350, 500	100, 200
Max sample rate	2 GSa/s	2 GSa/s	4 GSa/s	1 GSa/s
Max memory depth	100 kpts	2.5 kpts	4 Mpts	1 Mpt
Max update rate (waveforms/sec)	50,000	200**	1,000,000	5,000
Fully upgradable	Yes	No	Yes	No
Function Generator	Yes	No	Yes	No
Notes:	**Refer to Agilent Pub 5989-7885EN for update rate measurements Data for competitive oscilloscopes from competitor publications 3GW-25645-0 and 3GW-22048-1			

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

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

See the difference today.
www.testequity.com/Agilent_Scope

*All prices are in USD and subject to change

Agilent Technologies

EDN contents

10.6.11

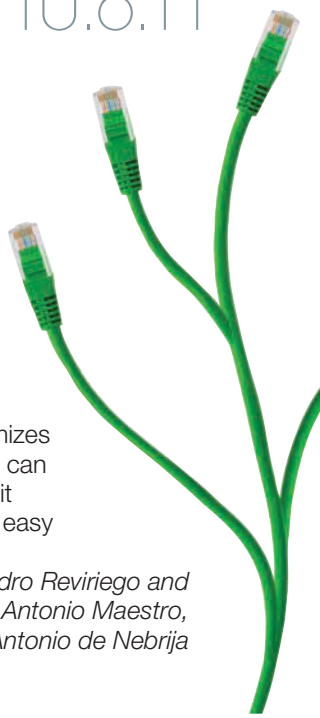

COVER: IMAGE: SPARK STUDIO/GETTYIMAGES

DESIGNIDEAS

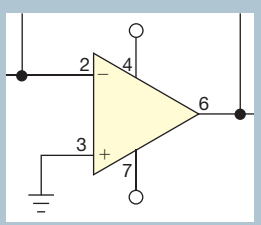

A True Low Power Platform for 8-/16-bit Applications

As low as 66µA per MHz and
0.57µA for RTC and LVD[1]

LOW POWER

RL78
'The Smart Efficiency'

SCALABLE

From 2KB to 512KB
of memory and
20- to 128-pin options

EFFICIENT

1.27 DMIPS per MHz and
1.6V to 5.5V operation

Applications:

- Industrial and Building Automation
- Medical
- Small Appliances
- White Goods
- Metering
- Consumer Electronics

RENESAS RL78

The RL78 Family of microcontrollers combine advanced low power technology, outstanding performance, and the broadest line-up in its class for the most demanding 8- and 16-bit embedded applications.

System Cost Saving

- Up to 8KB of Data Flash for data storage
- 32MHz internal oscillator
- Internal voltage reference
- Temperature sensor

Sophisticated Peripherals

- RTC with full calendar and alarm function
- Multifunction 16-bit timers
- Configurable communication blocks

Advanced Safety

- CRC, RAM parity-check
- Clock fail and frequency detection
- ADC self-test
- Memory protection

Register today to qualify for a

Free RL78 Renesas Demonstration Kit

www.renesas.com/RPBRL78G13/a

[1]RTC: Real-Time Clock LVD: Low Voltage Detect

RENESAS

contents 10.6.11

pulse

DEPARTMENTS & COLUMNS

EDN® (ISSN# 0012-7515) is published semimonthly by UBM Electronics, 600 Community Drive, Manhasset, NY 11030-3825. Periodicals postage paid at Manhasset, NY, and at additional mailing offices. SUBSCRIPTIONS—Free to qualified subscribers as defined on the subscription card. Rates for nonqualified subscriptions, including all issues: US, $150 one year; $250 two years; $300 three years. Except for special issues where price changes are indicated, single copies are available for $10 US and $15 foreign. For telephone inquiries regarding subscriptions, call 847-559-7597. E-mail: edn@omeda.com. CHANGE OF ADDRESS—Notices should be sent promptly to PO Box 47461, Plymouth, MN 55447. Please provide old mailing label as well as new address. Allow two months for change. NOTICE—Every precaution is taken to ensure accuracy of content; however, the publishers cannot accept responsibility for the correctness of the information supplied or advertised or for any opinion expressed herein. POSTMASTER—Send address changes to EDN, PO Box 47461, Plymouth, MN 55447. CANADA POST: Publications Mail Agreement 40612608. Return undeliverable Canadian addresses to BleuChip International, PO Box 25542, London, ON N6C 6B2. *Volume 56, Number 19 (Printed in USA).*

EDN online

JOIN THE CONVERSATION

Comments, thoughts, and opinions shared by *EDN*'s community

In response to Paul Rako's tear-down, "The Tektronix 1101 oscilloscope-probe power supply," at http://bit.ly/p3OukN, David Lockman commented:

"To paraphrase Santayana, those who cannot learn from history are doomed to repeat it. This [tear-down] is an opportunity to learn how to solve design problems, and perhaps save yourself a little time and effort when you are called upon to solve a similar problem."

In response to "Replaceable lithium-ion battery helps keep Droid Bionic out of the landfill," posted in Margery Conner's PowerSource blog at http://bit.ly/nxo7Rg, Netteligent commented:

"We produce so much waste with our electronic devices. I would propose [that] vendors consolidate and standardize their battery and charger into three models each. As volume increases, cost [is] much less. Interoperate among vendors: cheaper and convenient. Greener—much less pollution [in] landfills."

In response to "An engineer walks into a bar … ," a Tales from the Cube column on engineers and social situations, by William M Grissom, Aerojet, at http://bit.ly/pla9ap, Michael Blazer commented:

"In times like these, just do a little meditation. Find a comfortable chair, put your feet up on the bench, and chant the EE mantra. Repeat after me: Volt … Amp ... Ooooooohhhhmmmmmm …."

EDN *invites all of its readers to constructively and creatively comment on our content. You'll find the opportunity to do so at the bottom of each article and blog post. To review current comment threads on EDN.com, visit http://bit.ly/EDN_Talkback.*

CONTENT

Can't-miss content on EDN.com

TOP ENGINEERING SCHOOLS' RANKINGS

With prices high all around, where do kids these days get the most bang for their tuition bucks: at smaller engineering schools or at larger, well-known schools?

http://bit.ly/mZjEaX

HUNTING NOISE IN MIXED SYSTEMS

At some point, most of us need to find and remove noise and spurious signals. For some of us, the source of noise could be a neighbor's dog barking at 3 a.m. For those of us in electrical engineering, however, noise most likely comes from switching power supplies, other parts of the system, and external sources. Find out how to track down and eliminate these noise sources.

http://bit.ly/oROA7s

ENGINEERING COMMUNITY

Opportunities to get involved and show your smarts

DESIGN IDEAS WANTED

Want to see your work featured in *EDN*? Submit your own Design Idea, a short, concise article that helps solve problems or shows innovative ways to accomplish design tasks. To learn what makes a good Design Idea entry and how to submit it for consideration, visit **http://bit.ly/DesignIdeasGuide**.

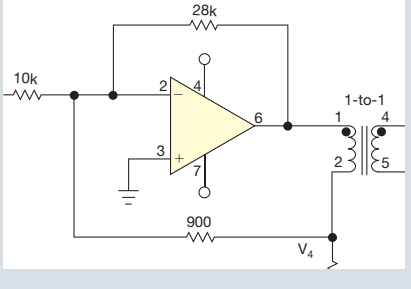

EDN

BRAND DIRECTOR, *EDN*
Jim Dempsey
1-440-333-3040;
jim.dempsey@ubm.com

DIRECTOR OF CONTENT, *EDN* AND DESIGNLINES
Patrick Mannion
1-631-543-0445;
patrick.mannion@ubm.com

EXECUTIVE EDITOR, *EDN* AND DESIGNLINES
Rich Pell
1-516-742-0027;
rich.pell@ubm.com

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

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

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

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

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

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

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

CONTRIBUTING TECHNICAL EDITORS
Dan Strassberg, strassbergedn@att.net
Stephen Taranovich, staranovich@yahoo.com
Brian Bailey, brian_bailey@acm.org
Robert Cravotta, robert.cravotta@embeddedinsights.com

VICE PRESIDENT/DESIGN DIRECTOR
Gene Fedele

CREATIVE DIRECTOR
David Nicastro

ART DIRECTOR
Giulia Fini-Gulotta

PRODUCTION
Adeline Cannone, Production Manager
Laura Alvino, Production Artist
Yoshihide Hohokabe, Production Artist
Diane Malone, Production Artist

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

EDN ASIA
Grace Wu, Associate Publisher
grace.wu@ednasia.com
Vivek Nanda, Executive Editor
vnanda@globalsources.com

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

EDN JAPAN
Masaya Ishida, Publisher
mishida@mx.itmedia.co.jp
Makoto Nishisaka, Editor
mnishisa@mx.itmedia.co.jp

UBM ELECTRONICS MANAGEMENT TEAM
Paul Miller, Chief Executive Officer, UBM Electronics and UBM Canon (Publishing)
Brent Pearson, Chief Information Officer
David Blaza, Senior Vice President
Karen Field, Senior Vice President, Content
Jean-Marie Enjuto, Vice President, Finance
Barbara Couchois, Vice President, Partner Services and Operations
Felicia Hamerman, Vice President, Marketing
Amandeep Sandhu, Director of Audience Engagement and Analytics

UBM Electronics

For a complete list of editorial contacts, see http://ubmelectronics.com/editorial-contacts

BY PAUL RAKO, TECHNICAL EDITOR

Computer History Museum honors Jim Williams and Bob Pease

The Computer History Museum in Mountain View, CA (**photo**, right), is building an "Engineers at Work" exhibit, which it plans to open on Oct 15. The exhibit will honor the late "analog gurus" Jim Williams and Bob Pease, who both died this year. As part of the exhibit, Dag Spicer, the senior curator at the museum, will arrange and display Williams' famously cluttered workbench (**photo**, below left). Jon Plutte, the media producer at the museum, notes that the exhibit will "celebrate the engineers of today and inspire the engineers of tomorrow."

EDN will also have a role in the museum's exhibit. John Hamburger, the director of marketing communications at Linear Technology, selected seven issues of *EDN*, which will be on display. The issues include a sampling of the many articles Jim Williams wrote for *EDN* over a span of more than three decades. These articles and more are available at www.edn.com/jimwilliams. As part of the video assets of the exhibit, the museum officials interviewed many of the people Williams worked with at Linear Technology, and they also spoke with me.

I asked Tim Regan, an application manager at the company, how Linear could possibly move Williams' bench 11 miles without disturbing the pile of work on the bench. "The same way Williams moved the bench seven years ago," he replied with a smile.

Apparently, when he faced moving the bench between two buildings, Williams took plastic-film shipping wrap and simply wrapped the whole bench with his equipment and prototype boards in place. When he arrived at his new lab, he just cut the shipping film off the bench. The staff at the Computer History Museum was nice enough to let me snap a picture of the Williams bench as they received it in the shipping wrap (**photo**, below right), which warmed my heart because it shows how Williams is still solving engineering problems even after his death. EDN

Contact me at paul.rako@ubm.com.

IMAGES: BOTTOM LEFT, PHOTO COURTESY FRAN HOFFART; TOP RIGHT, PHOTO COURTESY COMPUTER HISTORY MUSEUM

The New Future Electronics

Putting the Future into Today's Business

With 42 years of service excellence and proven experience, Future Electronics is bringing the best of its past success to take a creative new look at growing your business for the future. We've got a refreshed outlook and an invigorated spirit to support your business from the ground up.

The only thing we're changing is for the better.

FUTURE
ELECTRONICS

www.FutureElectronics.com **1-800-FUTURE-1**


IN THE EYE OF
EVERY ENGINEER...
...THERE IS A NEW EMBEDDED SOFTWARE SOLUTION.™
One that is integrated and optimized – simplifying what was once a complex problem.
Don't miss the unveiling at ARM TechCon™ 2011
Register today – www.armtechcon.com
Brought to you by
AVNET electronics marketing
ARM

EDITED BY FRAN GRANVILLE

pulse

INNOVATIONS & INNOVATORS

Adapter protects scope from high voltage, chases power-line EMI

AC power lines carry transient spikes and high-frequency EMI (electromagnetic-interference) noise, which can affect and even damage attached equipment. However, measuring the noise can be a hazard in itself because most test equipment, such as your scope, cannot probe high voltages. To address that problem, OnFilter's EMI adapter allows you to safely plug in your scope or spectrum analyzer to power lines and to display waveforms on your equipment. It allows measurements of noise on both differential and common-mode power lines without attenuation, rejects 50- and 60-Hz noise, and provides isolation from the power line. You can use it with any oscilloscope or spectrum analyzer. A protective circuit limits the peak signal at the output to 15V. The adapter works at voltages as high as 250V and on any outlet with a plug adapter. It comes with BNC cable, a 50Ω BNC terminator, and a BNC T adapter, which allows the use of the 1-MΩ input of your oscilloscope for large-dynamic-range measurements. It sells for $389.

OnFilter's EMI adapter allows you to safely plug in your scope or spectrum analyzer to power lines and to display waveforms on your equipment.

—by Margery Conner

▷**OnFilter**, www.onfilter.com.

"When people need a problem solved, who do they call? An engineer. When people need to make money quickly, who do they call? An engineer. Or they lie, cheat, steal, hire marketing people, talk to lawyers, or sell the company."

—Peter L, in *EDN*'s Talkback section, at http://bit.ly/nmdvY4. Add your comments.

Microcontrollers target cost-sensitive applications

Freescale recently expanded its portfolio of 32-bit Kinetis microcontrollers, which it based on the ARM (www.arm.com) Cortex-M4 processor, with 60 new entry-level devices that should become available for sampling in early 2012. The K10 and K20 family members will carry price tags of less than $1 each, making them well suited for cost-sensitive industrial and consumer appliances, including I/O modules for factory automation, portable health-care instruments, USB (Universal Serial Bus) microphones, gaming headsets, and smart-grid applications.

According to the company, the new devices are the lowest-power Kinetis microcontrollers to date, consuming only 630 µA at a CPU speed of 4 MHz in low-power run mode, which is equivalent to 160 µA/MHz. The lowest-power mode consumes only 40 nA, helping to extend battery life.

The devices offer 32 to 128 kbytes of flash memory in packages ranging from 32-pin, 5×5-mm QFNs to 64-pin LQFPs. Freescale built the devices in its 90-nm, thin-film-storage, low-leakage flash technology. They include as much as 32 kbytes of optional FlexMemory or 2 kbytes of EEPROM, low-power run- and stop-mode currents, and fast wake-up times. Peripheral options include a 16-bit ADC; full-speed USB On-The-Go controller with a complementary software stack; a low-power touch-sensing interface; and several general timing, communication, and control peripherals.

Support for the devices is available from a selection of easy-to-use 32-bit tools, including a package of Freescale's MQX RTOS and software and CodeWarrior for microcontrollers featuring Processor Expert software. Kinetis support from the ARM ecosystem will continue to thrive with support for the K10 and K20 families from IAR Systems (www.iar.com), Keil (www.keil.com), Green Hills Software (www.ghs.com), and Atollic (www.atollic.com); open-source GNU tools; and more. An evaluation and demonstration kit employing the Freescale Tower System—a modular development platform targeting reduced development costs and time to market—should debut in the first quarter of 2012.

Freescale expects to begin providing alpha samples of the 50-MHz K10 and K20 devices to selected customers in the fourth quarter of this year, and the devices should become broadly available in early 2012. Prices start at 99 cents (10,000).

—by Ismini Scouras

▷**Freescale**, www.freescale.com.

BMW's laser headlights: the future of solid-state lighting?

BMW recently announced that it will include laser headlights in one of its high-end cars within the next few years. Laser diodes are attractive light sources for headlights for several reasons. Laser light appears from a point source, whereas halogen, HID (high-intensity-discharge), and HB-LED (high-brightness-light-emitting-diode) lights use Lambertian reflectance, in which light is scattered so that the apparent brightness of the surface to an observer is the same regardless of the observer's angle of view.

Laser light is also highly collimated, traveling in a narrow beam, making for easy manipulation by a headlight lens because the optics are simpler for laser-diode light. Most important, it delivers a lot more lumens per watt than, for example, an HB LED: Whereas an HB LED might produce 100 lm/W, a laser diode can produce 170 lm/W, according to BMW.

Safety and concerns about eye damage always follow any discussion of lasers. However, according to Sevugan Nagappan, marketing manager for Osram's (www.osram.com) infrared-business unit, which is also responsible for laser diodes, lasers are common parts of everyday electronics, such as optical-storage drives. He points to Microsoft's (www.microsoft.com) Kinect-game controller, which also uses a laser diode as an illumination source for the gesture-recognition system. "The Kinect's optics diffuse the laser light, which makes it safe for the human eye," he says. Similarly, the sealed-lensing system for an automobile's headlights ensures that the intense laser beam never directly illuminates your eye.

BMW is promising laser-diode-based headlights for the production version of its i8 concept car.

"The first use for blue-laser diodes and still their biggest application are in the Blu-ray optical drives," says Nagappan. "At 405 nm, these diodes are actually in the UV [ultraviolet] range." Development of laser diodes for other applications has only fairly recently grown—for example, as light sources for projectors. "The Casio [www.casio.com] laser-LED hybrid projector uses a blue-laser diode with a longer wavelength than the UV-laser diodes used for data storage, as well as much higher output power," he adds. "In the last three years, people have been developing high-output laser diodes for the projection market." According to Nagappan, you can add a phosphor to high-power blue-laser diodes; the combination yields white light. "That's what BMW is doing," he says.

So, laser diodes are just a little lower on the maturity curve than white HB LEDs, which are currently the darlings of the lighting world. Laser-diode emitters have higher system efficiency and are easier to control through lensing. According to Nagappan, nothing in the laser-diode-manufacturing process makes them inherently more expensive than HB LEDs. "There's no reason that laser diodes cannot be competitive with LEDs, but you have to give it time for that to happen," he says. "The same thing happened with LEDs 15 years ago; the market was a lot smaller, and they were a lot more expensive than traditional light sources. You're going to see the same thing again; lasers are the next-generation light source. Will that be in all applications or in selected ones? We just don't know at this time."

> "Laser diodes are just a little lower on the maturity curve than white HB LEDs, which are currently the darlings of the lighting world. Laser-diode emitters have higher system efficiency and are easier to control.

According to Nagappan, only Osram and Nichia (www.nichia.com) now offer both LEDs and laser diodes. He does not disclose the source of the blue-laser diode in the BMW prototype.

—by Margery Conner

▷**BMW**, www.bmw.com.

DILBERT By Scott Adams

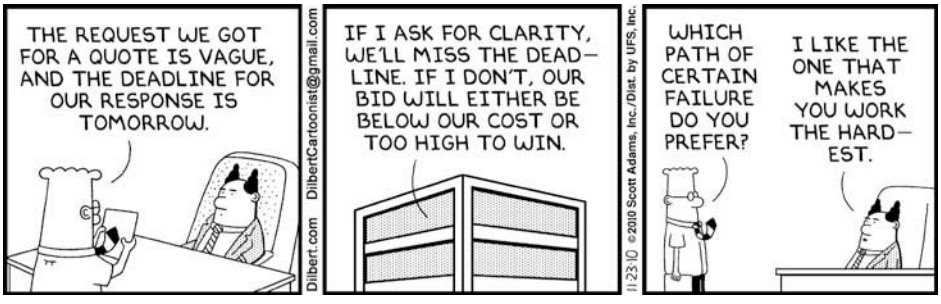

Your question:
Why full bandwidth up to 1 mV/div?

Our answer: For high signal fidelity and measurement accuracy on fast, low-amplitude signals.
The R&S®RTO oscilloscopes offer high vertical input sensitivity of up to 1 mV/div. They always operate at high measurement accuracy because the full bandwidth of the oscilloscope can be used at even the lowest sensitivity levels. Furthermore, the R&S®RTO oscilloscopes do not use software-based zooming. They show the signal's real test points at all sensitivity levels.

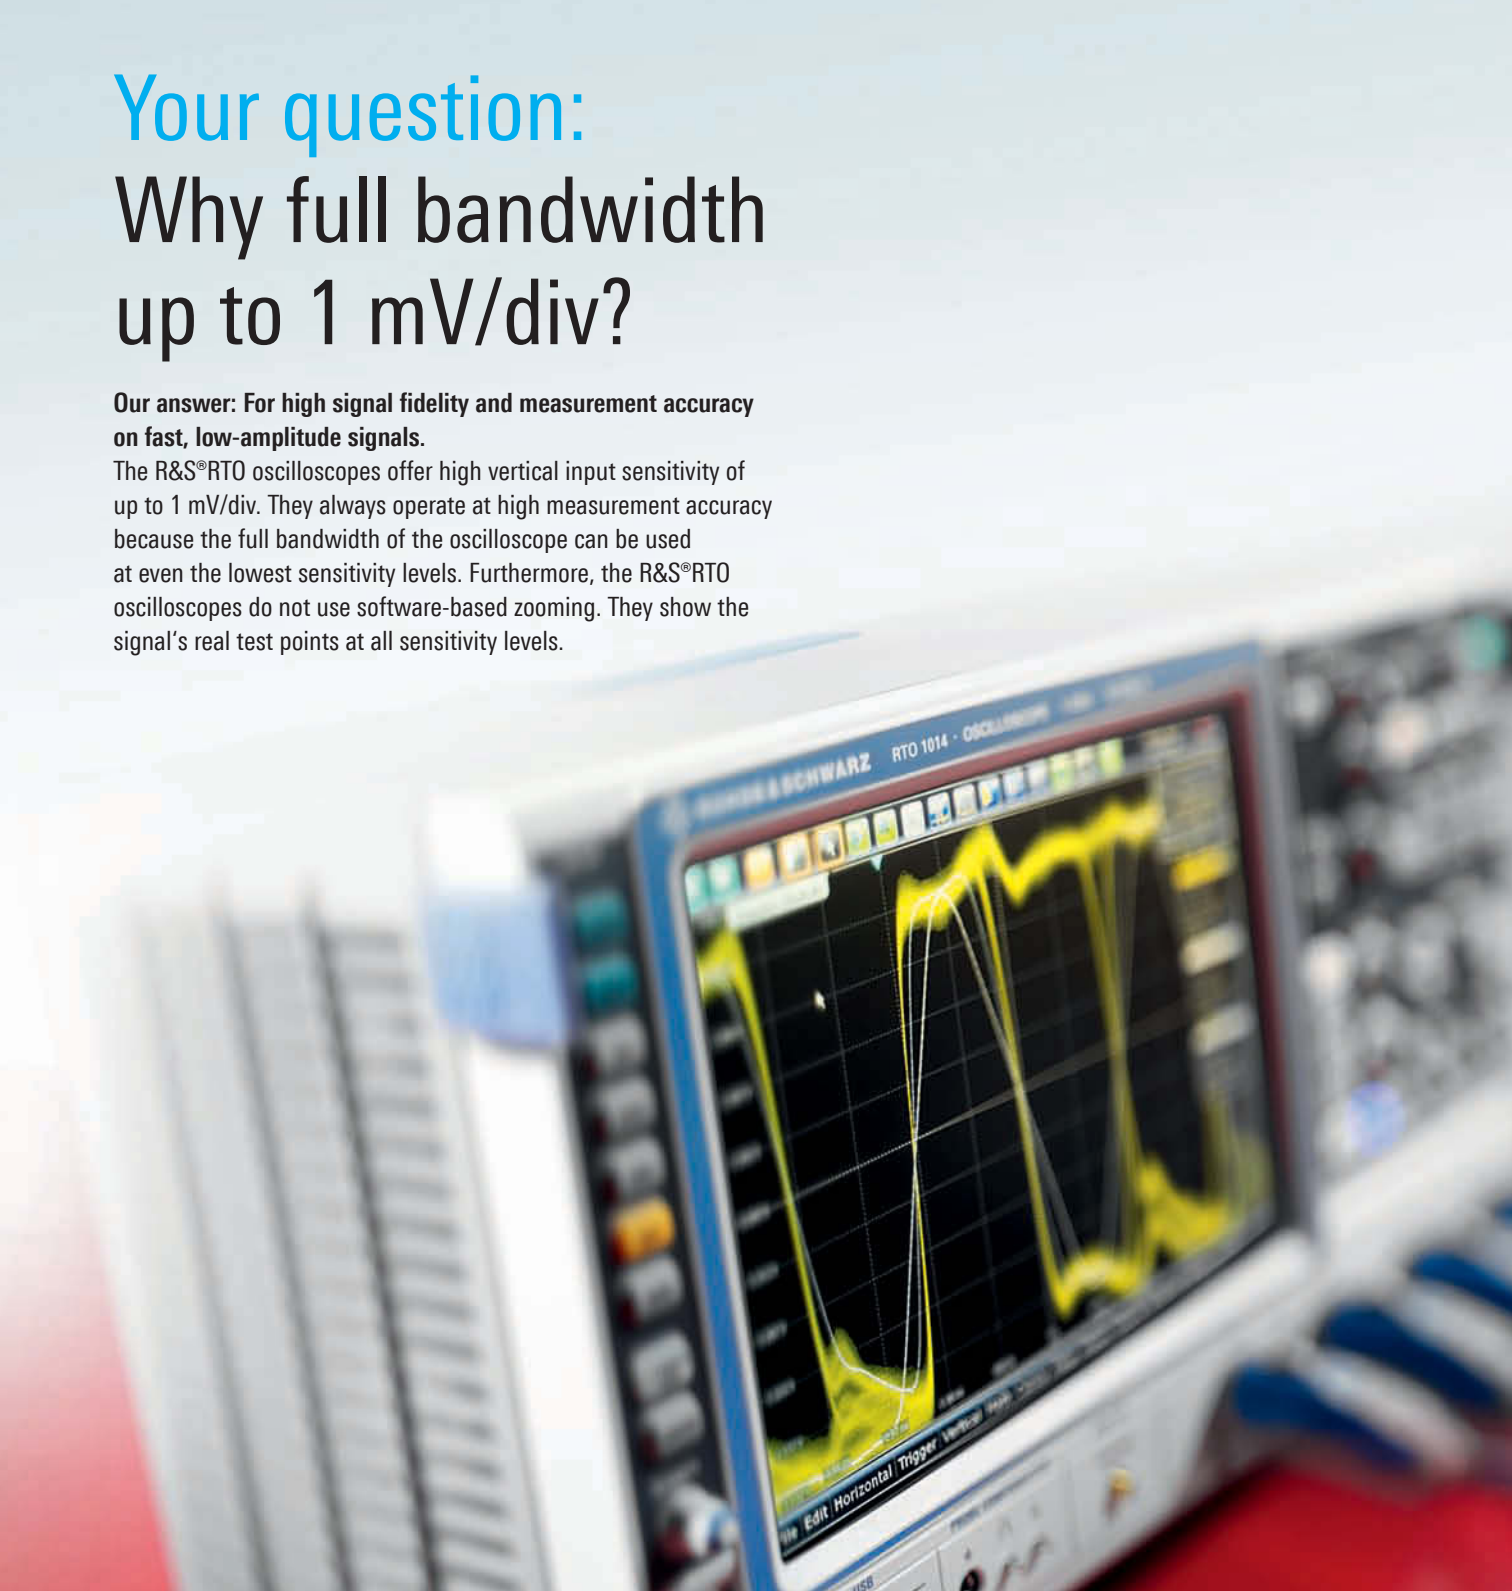

See for yourself
www.scope-of-the-art.com/ad/faq-bw/edn

ROHDE&SCHWARZ

Polyphase energy-metering IC monitors harmonic frequencies

Analog Devices' new three-phase ADE7880 energy-metering IC for polyphase meters employs the company's harmonic-analysis technology for advanced power-quality monitoring. The ADE7880's harmonic analysis produces magnitude and phase information, including current rms (root/mean/square); voltage rms; active, reactive, and apparent powers; and power-factor, harmonic-distortion, and THD+N (total-harmonic-distortion-plus-noise) calculations.

The ADE7880 delivers better than 1% accuracy over a dynamic range of 2000 to 1 to the 63rd harmonic. It delivers active- and reactive-energy accuracy better than 0.2% over a dynamic range of 5000 to 1 and accuracy better than 0.1% over a dynamic range of 1000 to 1 with gain calibration only. Using the company's proprietary ARTM (adaptive-real-time-monitoring) harmonic-analysis technology, the ADE7880 eliminates the need for custom development of sophisticated digital-signal processing and reduces the demands on the system microcontroller. ARTM automatically tracks changes in the fundamental frequency, eliminating any effects of frequency drift on the accuracy of the measurements.

The device is pin-compatible with the ADE78xx high-accuracy energy-metering ICs for three-phase configurations and incorporates seven sigma-delta ADCs, a digital integrator, a voltage reference, and all metrology signal processing. It is suitable for Class 0.2 meters and sells for $9.47 (1000). An evaluation board is available for $390.

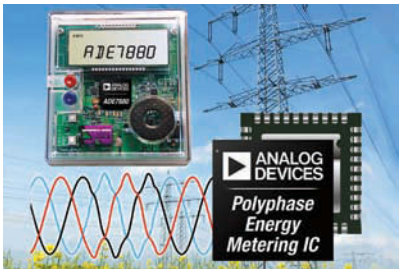

The three-phase ADE7880 energy-metering IC for polyphase meters employs the company's harmonic-analysis technology for advanced power-quality monitoring.

—by Ismini Scouras

▷**Analog Devices**, www.analog.com.

Renesas, x86 take back seat to ARM in 32- and 64-bit microcontrollers, embedded microprocessors

Given the accelerating number of ARM-based design wins, particularly in mobile applications, it's no surprise that a new Semicast (www.semicast.net) report marks ARM's dominance, just days after Kontron (http://us.kontron.com) announced that it was adopting ARM across the board. What is surprising, however, is the report's other conclusion: Freescale (www.freescale.com) has overtaken Renesas (http://am.renesas.com). In Semicast's 2011 edition of its study of 32- and 64-bit microcontrollers, embedded microprocessors, and DSPs, ARM surpassed the x64/x86 and Power Architecture; over 2010 to 2016, it should enjoy the highest growth, with x64/x86 remaining the leading challenger (**Reference 1**). Semicast expects revenue for 32- and 64-bit microcontrollers and embedded microprocessors to grow from $8.3 billion in 2010 to $19.8 billion in 2016, a compound annual growth rate of 15.6%.

Kontron recently announced plans to augment its x86 lineup with a broad range of boards, single-board computers, tablets, and industrial-PC units using the ARM architecture. Kontron's move is all about power and cost; the company claims that it can now easily port current and new products between RISC (reduced-instruction-set-computer) and CISC (complex-instruction-set-computer) architectures. The processor architecture is becoming less relevant as a decision criterion. Instead, price, power consumption, and performance per watt are now among the most important factors, according to Dirk Finstel, chief technology officer of Kontron. Kontron plans to develop ARM SOC (system-on-chip)-based computer modules with extensive support, including services such as customization, driver adoption, and porting of applications.

In Semicast's report, the writers lay out the competitive field. In 32-bit microcontrollers, the ARM architecture primarily competes against products employing Coldfire, Power Architecture, SuperH, and V850, with automotive and industrial the key markets in 2010. In 32/64-bit embedded microprocessors, the ARM architecture held a market share of less than 10% in 2010, with x64/x86 the leading architecture. ARM licensees have also started to compete in the market for 32/64-bit embedded microprocessors with products using the A8 and A9 cores but face strong competition from established suppliers such as AMD (www.amd.com), AppliedMicro (www.apm.com), Cavium (www.cavium.com), Freescale, Intel (www.intel.com), and PMC-Sierra (www.pmc-sierra.com).

Microsoft's (www.microsoft.com) announcement that both the x64/x86 and ARM architectures will support Windows 8 may have only a minor impact on this market until the late stages of the forecast period but will have significant implications over the long term. "Semicast judges that this announcement will weaken the stranglehold that the x64/x86 architecture has on applications that historically have been developed to be run on Microsoft Windows, thus opening up yet more designs wins to run on ARM," says Colin Barnden, principal analyst at Semicast Research and an author of the study.

Semicast deems Freescale the leading supplier of 32- and 64-bit microcontrollers and embedded microprocessors in 2010, ahead of Intel and Renesas. Freescale's offerings in this market include ARM, Coldfire, and Power Architecture products, whereas Renesas supports SuperH, V850, H8SX, M32, R32, and RX.

So, if you've opted for ARM, go forth and design in peace. You're not alone. If you've done so, please let us know your thoughts.

—by Patrick Mannion

▷**ARM**, www.arm.com.

REFERENCE

1 "32/64-bit Microcontrollers, Embedded Microprocessors & General Purpose DSPs—2011 Edition," Semicast, 2011, http://bit.ly/o7jwme.

Your question:
Why a configurable diagram display?

Our answer: For a well-structured, accurate display of multiple waveforms.
The Rohde&Schwarz SmartGrid function helps users arrange multiple diagrams on the screen. Individual waveforms can be displayed in a clear, well-structured manner. The A/D converter range is optimally used for highest accuracy waveforms.

See for yourself
www.scope-of-the-art.com/ad/faq-dis/edn

ROHDE&SCHWARZ

Gel electrolyte promises better battery

Researchers at the University of Leeds have invented a polymer gel that they say yields cheaper, more flexible lithium-ion batteries without compromising performance. The university licensed the technology to Polystor Energy Corp (www.polystorenergy.com), which is attempting to commercialize the cells for portable consumer-electronics applications.

Professor Ian Ward, the research physicist who developed the technology, says that it could replace the liquid electrolytes that rechargeable lithium-ion cells currently use. Traditional lithium-ion batteries require cells, which contain a porous polymer-film separator plus a liquid chemical filter, allowing lithium ions carrying charge to flow between the two electrodes and acting as a barrier between the electrodes to prevent short-circuiting.

This thin, flexible lithium-ion battery uses a gel electrolyte rather than the liquid electrolyte of traditional cells. It uses no separator and extends the life of portable electronics (courtesy University of Leeds).

The new automated process makes the polymer gel into a thin, flexible film, eliminating the need for a separator and making for a lighter pack. The researchers have also developed a patented manufacturing process that extrudes the gel between an anode and a cathode at speeds as high as 10m/minute to create a nanometers-thick battery. You can cut the resultant battery film to any size; the lamination process seals the electrodes.

"The polymer gel looks like a solid film, but it actually contains about 70% liquid electrolyte," says Ward. "It's made using the same principles as making a jelly: You add lots of hot water to gelatin—in this case, a polymer and electrolyte mix—and, as it cools, it sets to form a solid but flexible mass." Designers can shape and bend the flexible film to fit the geometries of virtually any device. **—by Peter Clarke**

▷**University of Leeds**, www.leeds.ac.uk.

10.06.11

Intel describes 22-nm Ivy Bridge CPUs

Intel publicly disclosed the Ivy Bridge, the company's first 22-nm processor family, at the recent Intel Developer Forum in San Francisco. The chip boasts significant improvements in onboard graphics and overall performance and a handful of stepwise advances in power management and security.

Rivals AMD (www.amd.com) and Nvidia (www.nvidia.com) have long criticized Intel for the relatively low-performance graphics cores the company uses in its chip sets and processors, but the gap between Intel and its rivals is closing quickly, according to Tom Piazza, an Intel fellow who helped design the graphics cores. "We have our foot on the gas [with the Ivy Bridge]. In a couple of years, the gap may be on the competitors' side."

Ivy Bridge graphics support Microsoft's (www.microsoft.com) Direct X 11 graphics APIs (application-programming interfaces), an area in which Intel parts once lagged by a generation. The devices also add support for three simultaneous displays and an L3 cache. Graphics features, among the few details that Intel has revealed, support 32 times more scatter/gather operations than the current Sandy Bridge chips. By increasing thread counts and moving to issuing multiple threads in parallel, the cores now support twice the instructions per cycle.

The current Sandy Bridge chips can simultaneously decode in hardware as many as 20 high-definition video streams. They also boost real-time transcoding speeds fivefold over the previous Westmere generation. With the Ivy Bridge, Intel says, it has enhanced video performance in at least two ways. Designers boost media-sampler throughput for better scaling and filtering and add new color and contrast enhancements to the pixel processing at the back end of the process.

The media blocks also add support for encoding using the multiview codec, key to support for stereo 3-D. Overall, the 22-nm process with its trigate FinFET transistors delivers twice the performance or half the power of 32-nm Sandy Bridge chips, says Varghese George, a senior principal engineer at the company.

Intel also adds to the chip a digital random-number generator that meets key ANSI (American National Standards Institute) and FIPS (Federal Information Processing Standard) security standards. It also adds the ability to prevent security attacks based on a process requesting an escalation of privileges.

For power management, Ivy Bridge can shut off I/O power to DDR memory in deep-sleep states. It can also automatically route threads to the most power-efficient core and optimize voltage use to the most optimal level. Ivy Bridge supports both DDR3 memory and the new DDR3L low-power chips. It also allows overclocking in 200-MHz increments without resetting the BIOS, including overclocking memory to as much as 2800 million transfers/sec.

—by Rick Merritt

▶**Intel Corp**, www.intel.com.

Your question:
Why hardware-accelerated histogram analysis?

Our answer: For fast, statistically conclusive results.
With their acquisition rate of one million waveforms per second — the highest rate available on the market — the R&S®RTO oscilloscopes find signal faults quickly. An ASIC in the R&S®RTO oscilloscopes employs 20-fold parallel signal processing which ensures high acquisition rates, even for complex signal analysis. The results are available quickly and are based on a large number of waveforms, providing statistically conclusive information.

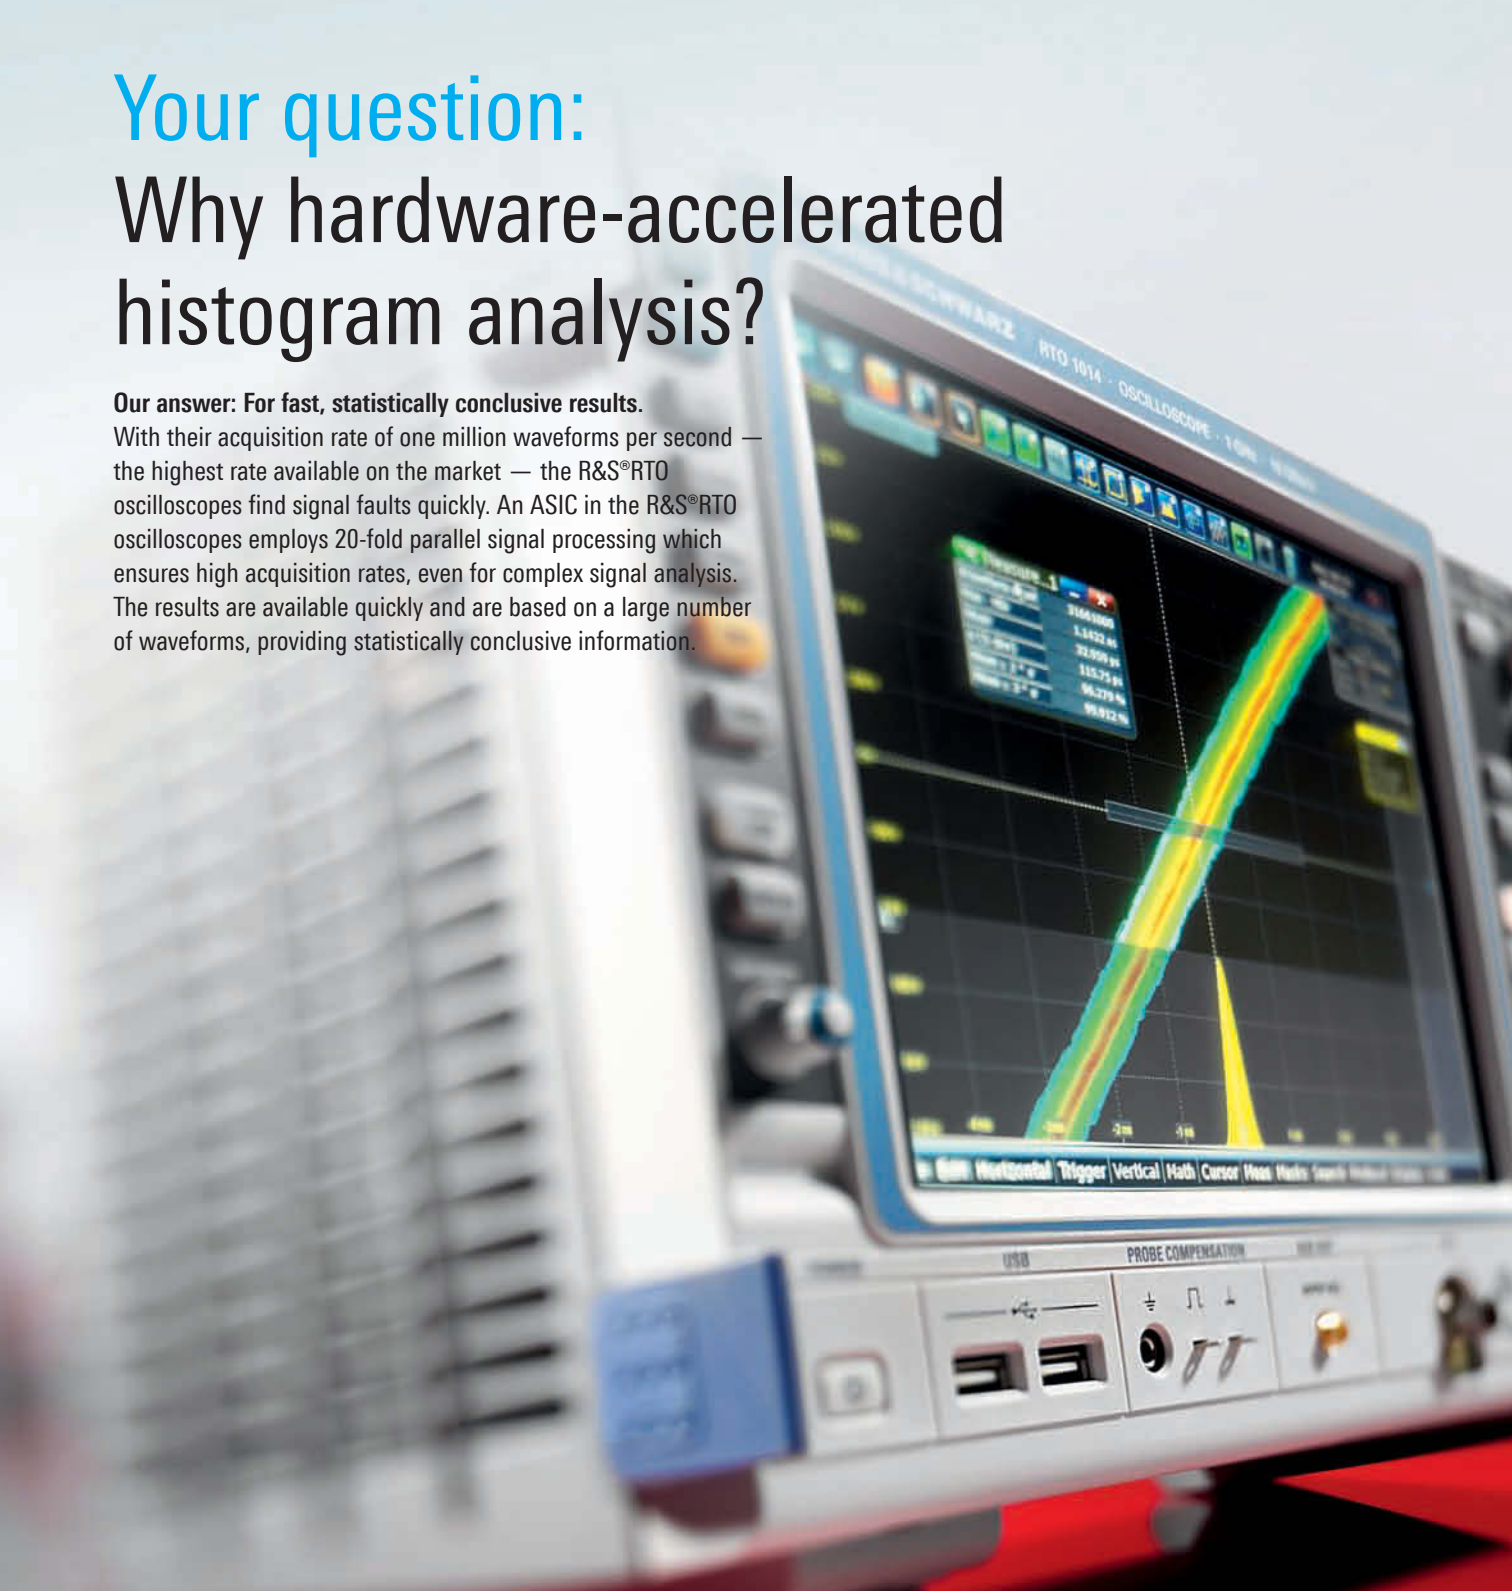

See for yourself
www.scope-of-the-art.com/ad/faq-histo/edn

ROHDE&SCHWARZ

BY HOWARD JOHNSON, PhD

Wafer-probe launch

Successful high-speed projects begin with a series of test cards. The cards test whatever new package layout, PCB (printed-circuit-board)-layer stack, and connector features the product will use. They interface directly with either a VNA (vector network analyzer) or a TDR (time-domain reflectometer). Test cards, implemented early in the development schedule, can root out difficulties with impedance mismatch, signal loss, and crosstalk that might otherwise not show up until near the end of the production cycle.

The usual approach brings coaxial cables from the VNA to the test card using SMA (surface-mount-assembly) connectors, which can achieve a bandwidth of 18 GHz if laid out perfectly on your PCB. According to the rule of thumb that your measurement equipment should have a bandwidth three times greater than the frequencies you are trying to measure, properly configured SMA connectors should be able to accurately test layout features at frequencies as high as 6 GHz. That bandwidth would be adequate if you were working on a serial link at, say, 10 Gbps, for which the maximum frequency of the 101010 repeating pattern equals 5 GHz.

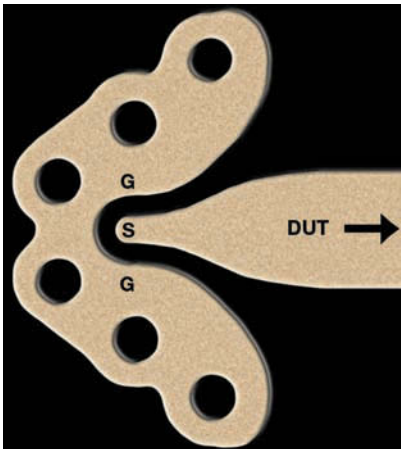

Figure 1 A compact launch zone minimizes impedance disruption.

For the latest generation of 28-Gbps serial links, however, the SMA runs out of gas. Although manufacturers of oscilloscopes and VNAs have developed "super-SMA" connectors that have considerably higher bandwidth, the generic SMA-type connectors you are likely to use on your PCB, in combination with your less-than-optimal layout, do not perform well at 28 Gbps.

A better approach connects your VNA to the test card using a high-performance microwave wafer probe. **Figure 1** illustrates a PCB layout that Tibor Lapohos, PhD, developed for use with the GGB Industries (www.ggb.com) model 110H Picoprobe (**Reference 1** and **Figure 2**). The Picoprobe sits at the end of a tiny, 1-mm coaxial cable. It touches down onto the layout in a GSG (ground-signal-ground) arrangement. The pin spacing can be as small as 50 microns, or 2 mils.

The butterfly-shaped region on the left in **Figure 1** is a ground land, which incorporates six ground vias. The signal trace carries the signal power off to the right, toward the DUT (device under test), somewhere off-screen. The combination of the Picoprobe and an optimized launch layout replaces the SMA-style connector, providing a higher overall rated bandwidth.

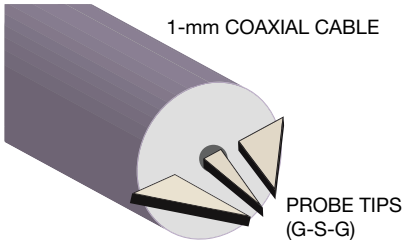

Figure 2 This conceptual view of a high-performance microwave wafer probe shows three points of contact.

Lapohos concentrates the launch zone in one small area, minimizing disruptions due to the sudden change from the coaxial geometry within the Picoprobe to the microstrip configuration on the PCB. After that change, the geometry fans out to full PCB-microstrip width.

The ground land surrounds the signal trace near the launch point. Lapohos carefully sculpted the ground land to provide enough distributed capacitance near the launch region to maintain a characteristic impedance of 50Ω on the skinny signal trace at that point. As the signal moves to the right, the grounded metal on either side forms a co-planar-waveguide structure. As the signal trace widens, the ground land peels away, leaving a simple microstrip. This layout maintains a consistent 50Ω characteristic impedance at all points. EDN

REFERENCE

1 Lapohos, Tibor, PhD, "Microwave Probe Pad Design: The First Step," BreconRidge Manufacturing Solutions Corp, Ansoft Partners in Design Workshop, 2004, Presentation 13, http://bit.ly/nQZvl2.

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.

SET YOUR DESIGNERS FREE WITH CALIBRE® AUTOMATIC WAIVERS | It keeps track of waivers and eliminates them from SoC verification results, saving hours of wasted effort by you and your ecosystem partners. Now your designers can do more productive things. Take a test drive at: http://go.mentor.com/autowaivers.

Mentor Graphics®
THE EDA TECHNOLOGY LEADER

Get in.

Buckle up.

Come Along for the Ride on the Ultimate Electronics Road Trip

DRIVEFORINNOVATION

12 months. 10,000 miles. Many technology stops.

Join us on the ultimate electronics road trip. Learn new technologies and supply chain options, access tools, industry icons and innovators along the way.

AVNET® Express

Enter for your chance to win an electric car and other prizes.

driveforinnovation.com

OPTIMIZE ENERGY-EFFICIENT ETHERNET PERFORMANCE IN BUNDLED LINKS

BY PEDRO REVIRIEGO AND JUAN ANTONIO MAESTRO
UNIVERSITY ANTONIO DE NEBRIJA

The IEEE recently approved the 802.3az EEE (energy-efficient Ethernet) standard, which promises significant energy savings to adopters in the coming years. Although industry participants have recently done several studies on reducing the energy consumption in bundled links in Ethernet, these studies do not consider the use of EEE. This article describes a technique that optimizes the use of EEE in bundled links, offers better response than do previous approaches, and minimizes frame delays and loss. The technique also allows you to independently implement each link's endpoint for easy deployment. When traffic is asymmetric, the technique significantly improves energy efficiency.

IMAGE: THINKSTOCK

The EEE standard targets more than 4 TWhrs per year in energy savings (**Reference 1**). It specifies methods of reducing energy consumption in Ethernet devices by defining low-power modes (**Reference 2**). A transceiver that has no frames to transmit can enter a low-power mode. When new frames arrive, the transceiver returns to active mode within a few microseconds, enabling energy savings that are almost transparent to upper protocol layers. The energy savings in a given link directly relate to the time the transceiver spends in the low-power mode. This time in turn depends on the traffic load, and, due to mode-transition overheads for loads with more than a few percentage points, the savings greatly decrease (**Reference 3**), meaning that EEE provides limited savings when the links operate at loads larger than 10%.

It is common in Ethernet to bundle links to provide larger capacity, and this feature has motivated a number of studies that aim to optimize the energy consumption on bundled Ethernet links (**references 4**, **5**, and **6**). Those methods dynamically adapt the number of active links in a bundle to the traffic load, meaning that you can save a significant amount of energy when you use lightly loaded bundles. However, those methods do not consider the use of the EEE standard.

AT A GLANCE

- This technique optimizes the use of EEE (energy-efficient Ethernet) in bundled links, offers better response than do previous approaches, and minimizes frame delays and loss.
- The technique also allows you to independently implement each link's endpoint for easy deployment.
- The EEE standard targets more than 4 TWhrs per year in energy savings.
- In the link-aggregation implementations that use EEE, each of the links that composes the aggregation receives the same load and, depending on whether it has frames to transmit, enters or exits the active state.

LINK AGGREGATION

The IEEE standardized link aggregation in Ethernet in IEEE 802.3ad, which the organization later renamed IEEE 802.1ax for consistency with other 802.3 standards (**Reference 7**). The standard enables the bundling of multiple Ethernet links into one logical link: the LAG (link-aggregation group). IEEE 802.3 stipulates that link aggregation must take place above the MAC (media-access-control) layer, where entire frames are sent to the MAC layer of one of the aggregated links. Link aggregation provides several advantages (**Reference 8**). For example, the aggregated link's capacity is the sum of the capacities of the links, enabling increases in link capacity that are smaller than the usual 10-times factor for different Ethernet technologies. Aggregated links also can be used to provide larger capacity than the capacity that the PHY (physical) layer provides. Link aggregation also increases

the link's availability because the aggregated link would be unavailable only if all links were unavailable.

Link aggregation can cause problems with arbitrary frame distribution, such as frame reordering, however. Frame reordering can occur when a long frame that arrives first transmits over one link and a second short frame that arrives shortly after the first one transmits over a different link, thus causing it to arrive first. Frame reordering can cause problems to higher-layer protocols, and the developers of the IEEE 802.3ad standard therefore designed it to avoid frame reordering. Frame reordering is an issue only in frames that belong to the same conversation. Therefore, algorithms that distribute frames across links for transmission ensure that frames that belong to a conversation transmit over the same link and preserve frame order. These algorithms ensure that those frames remain in the correct order but limit the maximum bandwidth to a link that a conversation can achieve.

Figure 1 shows the link-aggregation sublayer. The frame-distribution block assigns frames to links for transmission and must ensure that all frames from a conversation transmit over the same link. The frame-collection block receives frames from different ports and passes them to the MAC client. The frame-collection block must ensure only that frames from the same port pass in order to the MAC client. This design makes the frame-collection block independent of the algorithm in the frame-distribution block and avoids buffering and reordering in the collector. This design also avoids fragmentation and reassembly.

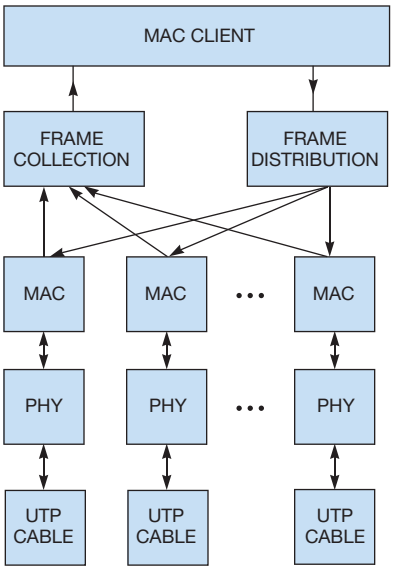

Figure 1 The frame-distribution block assigns frames to links for transmission and must ensure that all frames from a conversation transmit over the same link. The frame-collection block receives frames from different ports and passes them to the MAC client.

You can use other alternatives for link aggregation. For example, links can be aggregated at the PHY layer (**Reference 9**). In this method, frame transmission is accomplished by fragmenting the frames in segments that can then be distributed among the links for transmission. At reception, the technology reassembles the fragments to reconstruct the original frame. WAN (wide-area-network) aggregation widely uses this strategy. In this case, fragments of a conversation transmit over all the links; a single conversation can thus use all the capacity of the aggregated link. The IEEE 802.3ad standard did not select this type of design, however, because it complicates the frame-distribution and -collection processes.

To efficiently use the capacity in the aggregated link, the distribution algorithms try to assign conversations to links in such a way that the load in each of the links is similar. This approach enables high link usage when many conversations are taking place. You can define conversations at different levels, depending on the network topology and traffic. Annex A of the standard discusses different options, such as using the Ethernet source or destination address to identify a conversation. In some cases, it is more convenient to use higher-layer protocol information to define a conversation. For example, some methods define a conversation as a TCP (Transmission Control Protocol) connection. The standard also provides mechanisms for reassigning conversations to different links and ensures frame order. This feature is useful for dynamically balancing the load of the links.

When the traffic load is low and all of the aggregated links are active, link aggregation can waste a significant

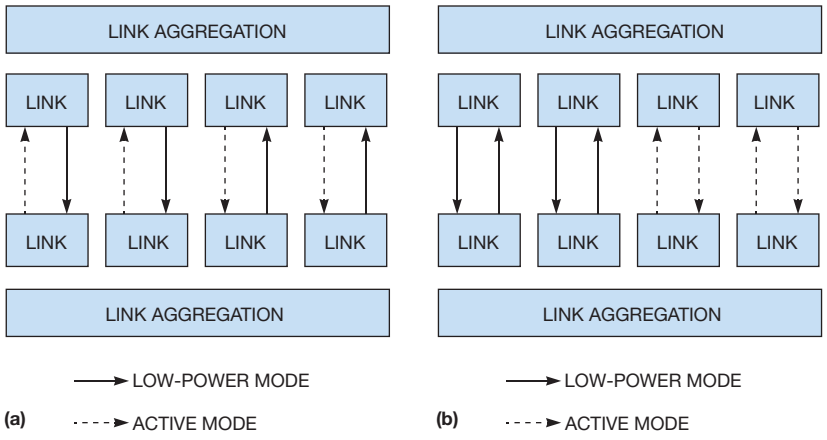

Figure 2 You can manage links independently when they are in low-power mode (a); otherwise, link directions enter low-power mode together (b).

amount of energy because energy consumption in PHY-layer devices is roughly independent of the traffic load once the link is active. One method proposes that the number of active links should adapt to the traffic load so that only a few links remain active when there is little traffic. Transitions between active and standby modes of the links require coordination between the link partners using the LACP (Link Aggregation Control Protocol). Transitions in this case would require hundreds of milliseconds because the link must also re-establish itself, which also requires a significant amount of time (**Reference 10**). The LACP manages links as a whole—that is, both link directions are active or idle, but they cannot be active in one direction and idle in the other. This approach can significantly reduce energy consumption in aggregated links but does not include coordination with the EEE.

EEE IN BUNDLED LINKS

Once EEE reaches widespread adoption, enabling EEE on each of the aggregated links would seem to ensure energy efficiency when there is low link usage because the links would usually be in low-power mode. This idea would make unnecessary those schemes that adapt the number of active links to the traffic load. EEE provides significant savings, however, only when the load is low. EEE provides limited savings for links in a bundle with only a 15% load, for example, whereas it could achieve a significant amount of savings if the traffic concentration occurs on only some of the links of the bundle. Therefore, even when the industry adopts EEE, it may be useful to study additional energy-efficient policies for aggregated links that operate at those loads.

Another method enables all links in the LAG for EEE and assigns no conversations to links in low-power mode, meaning that no frames transmit on those links. For most EEE PHY and MAC layers, you can thus independently manage the state of link directions; one direction can be in low-power mode, and the other can be in active mode. This approach simplifies the selection of which links to place in standby mode because the decision can be made locally with no coordination with the other end to ensure that the links in low-power mode are the same. **Figure 2** illustrates two possible configurations. When transmission takes place only in one direction on each link, there is no need to perform near-end-crosstalk and echo cancellation in the receivers (**Reference 11**), providing significant energy savings for the scheme when it is implemented with smart management of the PHY devices.

The local frame-distribution block monitors the traffic load and decides the number of links that have assigned conversations. If traffic decreases, the block can reassign the conversations to fewer links to reduce energy consumption. Conversely, if traffic increases, the block can use more links for the conversations. These reassignments can be performed in the same way that conversations are dynamically reassigned to balance loads. **Figure 3** shows an algorithm that requires only the reassignment of conversations to reduce the number of active links. Those links with no assigned conversations automatically enter and remain in the EEE low-power mode until the load increases. At that point, the approach assigns some conversations to the link, and frame transmission can start immediately by entering the EEE active mode.

The main parameter of the algorithm is the load threshold that determines a change in the number of links. Implementations should define a hysteresis mechanism to avoid frequent changes, such as the use of two thresholds for increasing and decreasing the number of active links. However, a single threshold is sufficient. If this threshold is large—for example, a load approaching 100% of the link's capacity—then the latency of the link due to large queue occupancy would increase, and the same situation would occur with the probability of discarded frames due to queue overflow. If the threshold is small, links will be lightly loaded, reducing energy efficiency. Therefore, a compromise value—in this case, 80% of the link's capacity—has been selected to illustrate this approach.

This approach, like traditional frame-distribution schemes, balances the load among the active links—that is, those that have conversations assigned to them—ensuring that conversations receive a similar service regardless of their link assignment. The approach effectively combines EEE with link aggregation, in most cases resulting in larger energy savings than those obtained using EEE alone.

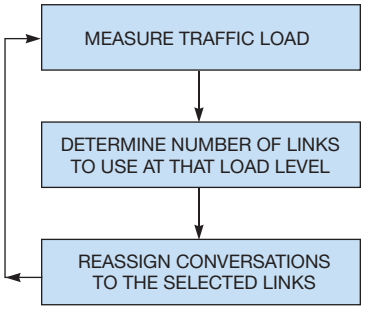

Figure 3 This algorithm combines link bundling with EEE.

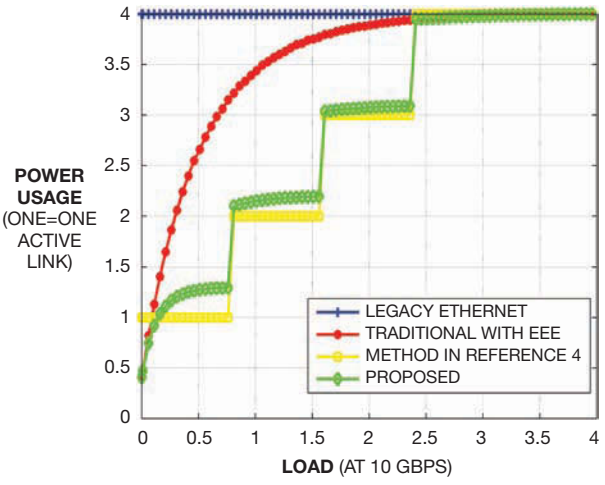

Figure 4 The link-aggregation algorithm compares energy efficiency for four technologies.

You can also implement the approach independently of the link partner; a manufacturer that implements this algorithm achieves some energy savings even if the other end does not implement the algorithm. Further, you can independently manage link directions, enabling configurations in which a different number of links are active in each direction. This feature is useful for links that have an asymmetric load, such as when one direction carries mostly data frames and the other carries acknowledgments. In those cases, using fewer links in the acknowledgment direction can save a significant amount of energy. Other studies have proposed energy-optimization techniques employing the asymmetric operation of Ethernet links (**Reference 12**). This approach also allows quick changes in the number of links because the approach makes all decisions locally with no need to exchange information with the other end and because the transition from low-power to active mode occurs quickly, improving the responsiveness of the system when the load changes, avoiding frame delay or loss.

In the link-aggregation implementations that use EEE, each of the links that composes the aggregation receives the same load and, depending on whether it has frames to transmit, enters or exits the active state. Therefore, the power consumption equals the number of links in the aggregation times the power consumption for a link at that load level. In the alternative approach, the power consumption would directly depend on the number of active links.

PERFORMANCE EVALUATION

A number of simulation experiments using 1250-byte frames estimate the potential energy savings of this method. The experiments also simulate other frame lengths and produce similar results. Frames arrive following a Poisson distribution—a discrete probability distribution that expresses the probability of a given number of events occurring in a fixed interval of time, interval of space, or both if these events occur with a known average rate and independently of the time since the last event. Frame arrivals deviate from a Poisson model in LAN traffic, so this simulation provides only a rough approximation. However, for LAGs that aggregate traffic from many sources, the Poisson assumption can be valid for short intervals, which are the time scales relevant for this approach (**Reference 13**).

Assume the existence of independently managed link directions and assume that the MAC and PHY layers may enter the low-power mode asymmetrically. Also assume that power consumption in the low-power mode is 10% of the power in the active mode and that power during state transitions is the same as that in the active state. The duration of the state transitions complies with the IEEE 802.3az standard. Other methods assume that the links consume no power in standby mode because those links are effectively idle rather than in the low-power mode in EEE that keeps some elements active and schedules refresh periods.

Figure 4 compares energy-consumption results for different loads and technologies. In this case, the aggregation

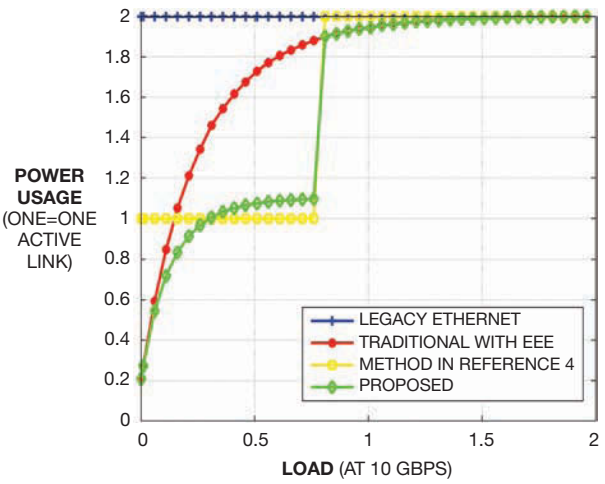

Figure 5 The link-aggregation algorithms compare energy efficiency for two links.

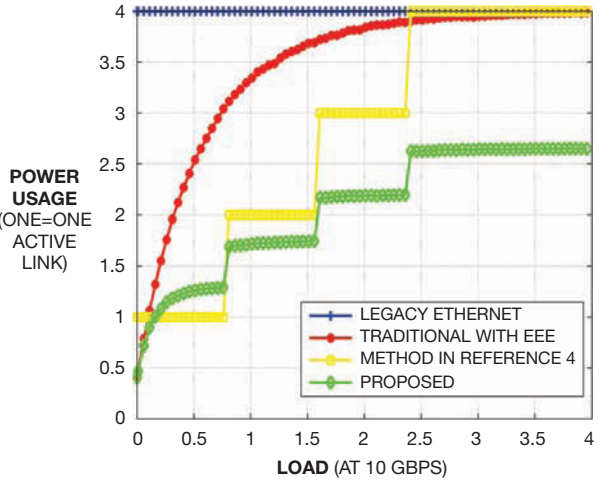

Figure 6 On a link with asymmetric traffic, the proposed link-aggregation algorithms outperform those of the traditional algorithms.

algorithm comprises four 10BaseT links. Experiments yield similar results for 100BaseTX and 1000BaseT links. The next experiment considers aggregations with two links (**Figure 5**). The overall shape of the energy consumption versus load is similar for the other number of links. Additional experiments use different frame sizes, and the results are similar: Most of the savings come from having some links in low-power mode, depending on the load of the LAG and not on the frame size.

The results so far focus on one link direction, but traffic is sometimes

asymmetric, with one direction having a significant load and the other having a much smaller one. In those cases, the technique described in this article enables a different number of links to be active on each direction. Simulations can evaluate the potential savings in links with asymmetric traffic with 1250-byte frames in one direction and the same number of 125-byte frames in the other direction (**Figure 6**). In this case, the simulation estimates the consumption in each direction and computes the average to provide a normalized result. The experiments show that the proposed algorithm obtains additional energy savings for a wide range of link loads compared with the use of EEE alone. EDN

REFERENCES

1 Christensen, K; Pedro Reviriego; B Nordman; M Bennett; M Mostowfi; and Juan Antonio Maestro, "IEEE 802.3az: The Road to Energy Efficient Ethernet," *IEEE Communications Magazine*, Volume 48 , Issue 11, November 2010, pg 50, http://bit.ly/qwQs2A.

2 "P802.3az Energy Efficient Ethernet Standard," IEEE, September 2010, http://bit.ly/r0DHQC.

3 Reviriego, Pedro; JA Hernandez; D Larrabeiti; and Juan Antonio Maestro, "Performance evaluation of energy efficient Ethernet," *IEEE Communication Letters*, Volume 13, No. 9, September 2009, pg 697, http://bit.ly/orxFVc.

4 Koenen, David J, and Mike Chuang, "Method for Ethernet Power Savings on Link Aggregated Groups," United States Patent Application 20080304519, Dec 11, 2008, http://bit.ly/nVzIlz.

5 Fukuda, Y; T Ikenaga; H Tamura; M Uchida; K Kawahara; and Y Oie, "Performance evaluation of power saving scheme with dynamic transmission capacity control," *Proceedings of Globecom 2009*, December 2009, http://bit.ly/nUFlpS.

6 Imaizumi, H; T Nagata; G Kunito; K Yamazaki; and H Morikawa, "Power saving mechanism based on simple moving average for 802.3ad link aggregation," *Proceedings of Globecom 2009*, December 2009, http://bit.ly/nHslBj.

7 Frazier, Howard, "IEEE 802.3ad Link Aggregation: what it is and what it is not," IEEE, April 17, 2007, http://bit.ly/rbFgoh.

8 Seifert, Rich, and James Edwards, *The All-New Switch Book: The Complete Guide to LAN Switching Technology*, John Wiley and Sons, 2008, ISBN 0470287152.

9 Frazier, Howard M, "Aggregation at the Physical Layer," *IEEE Communications Magazine*, Volume 46, Issue 2, 2008, Jan 23, 2008, pg S12, http://bit.ly/q4k7w5.

10 Gunaratne, C; K Christensen; B Nordman; and S Suen, "Reducing the energy consumption of Ethernet with adaptive link rate (ALR)," *IEEE Transactions on Computers*, Volume 57, No. 4, pg 448, April 2008, http://bit.ly/oSxVeQ.

11 Muller, Shimon; Brad Booth; Alan Flatman; George Zimmerman; and Sailesh Rao, "IEEE 802 10GBASET Tutorial," IEEE, November 2003, http://bit.ly/pE0jsW.

12 Diab, Wael William, and Kevin Brown, "Method And System For Asymmetric Operation In A Network Node In An Energy Efficient Ethernet Network," United States Patent Application 20090154465, http://bit.ly/o0JFML.

13 Karagiannis, T; M Molle; M Faloutsos; and A Broido, "A non-stationary Poisson view of Internet traffic," *Proceedings of IEEE Infocom*, March 2004, pg 1558, http://bit.ly/pkXNRM.

AUTHORS' BIOGRAPHIES

Pedro Reviriego is an associate professor in computer science at the University Antonio de Nebrija (Madrid, Spain). His research interests include fault-tolerant systems, signal processing, and computer networks. Reviriego has a doctorate in telecommunications engineering from the Polytechnic University (Madrid, Spain). He has participated in the IEEE 802.3 standardization for 10GBaseT and worked on the design of 1000BaseT transceivers.

Juan Antonio Maestro manages the computer-architecture and -technology group at the University Antonio de Nebrija (Madrid, Spain). His research interests include signal processing and real-time systems, fault tolerance, and reliability. Maestro has a doctorate in computer science from Complutense University of Madrid (Madrid, Spain).

TIME CAREFULLY THE ADOPTION OF YOUR NEXT-GENERATION PROCESSOR

CAN YOU MAXIMIZE FIRST-TO-MARKET OPPORTUNITIES AND MINIMIZE THE EXTRA RISKS OF EARLY ADOPTION?

BY ROBERT CRAVOTTA • EMBEDDED INSIGHTS INC

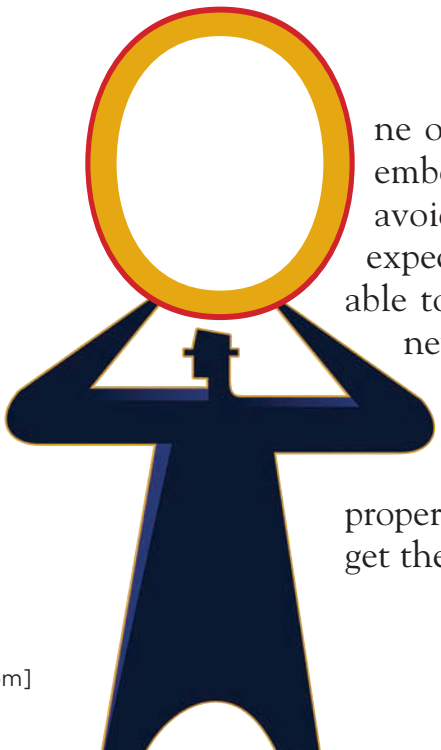

IMAGE: SPARK STUDIO/GETTYIMAGES

One of the first expressions I learned as a young embedded-system developer was "Smart money avoids Version 1.0." Back then, there was little expectation that even a compiler would be available to developers soon after the availability of a new processor. Working with a new processor or development tool meant your own development team would experience some of the growing pains of determining what worked properly and what work-arounds were necessary to get the system to do what it needed to do.

Suppliers use major releases of available products to signal to users that the new features and capabilities are built on top of a proven product line; this helps mitigate some of the challenges of launching completely new and unproven products. Suppliers often mark a major release of their product lines by increasing the major version number, such as to 2.0, or by launching a new family name within the product line. Major releases or new-product families can span and employ multiple already-available product lines. For example, Texas Instruments' recent launch of the Hercules safety-microcontroller family spans the TMS570 and RM4x product lines. TI based the new microcontrollers on proven processor families but packages new or different peripherals and features in a way that is appropriate to a vertical market—in this case, safety applications.

Today's newest processors are significantly more complicated than their predecessors. They may include new peripherals, application-specific accelerators, and additional processor cores that may not share the same architecture as the main processor core. New features can mean that the developer community may have little or no direct experience with the parts of the system and how to best use the new devices' features. To sufficiently minimize developers' learning curves for using the new processor and encapsulate the complexity of these systems, the supplier must provide appropriate development tools, drivers, boards, and support IP (intellectual property) when the company launches the processor.

Unfortunately, the wisdom contained in the smart-money expression also applies to major releases of products that have been in the market as production systems because customers have not yet robustly applied the new features and capabilities they contain to an exhaustive set of real-world use cases. A development team that adopts a new product or major update may experience the pain, cost, and risk of changes to their project schedules and engineering resources. As a result, some development teams enforce a policy of avoiding using the latest one or two versions of a component in their designs.

The intent of this type of policy is not to avoid using the new capabilities that the latest version supports but to avoid the risk of having to rework portions of a design because of changes in how to access and manage a new capability as the system undergoes final engineering tweaks and adjustments to deliver optimum performance and manufacturing yields before full production availability. This type of policy relies on the supplier's version nomenclature and production-availability date to signal when the volatility of the system is stable enough that additional engineering tweaks are unlikely to cause rework by the development team. Although this policy is simple to follow, it can potentially cause a development team to miss out on substantial first-to-market benefits by adopting a new processor or tool months after the adoption sweet spot—that point in a new processor's or development tool's early life cycle that provides the most opportunity for a design to reach the market first but incurring the lowest additional risks to the development effort (**Figure 1**).

AT A GLANCE

- Suppliers use major releases to signal when new capabilities are built on top of proven products.
- A development team that adopts a new product or major update may experience the pain, cost, and risk of changes to their project schedules and engineering resources.
- Early adopters gain the largest first-to-market-window opportunity, but they may expend more engineering resources than developers who adopt the processor later.
- Working directly with an FAE (field-application engineer) usually provides the most reliable, accurate, and timely information about a new processor.
- Community-support forums provide avenues for finding timely information about using a new processor.

ADOPTION SWEET SPOT

The first users of a new processor, the early adopters, gain the largest first-to-market-window opportunity, but they may expend more engineering resources than developers who adopt the processor after it has reached general or production availability months later. As a processor proceeds through the normal life cycle after first availability, the number and severity of risks that users of the new processor are subject to decrease (see **sidebar** "Sources of risk").

General adopters adopt a new processor after the silicon and development tools have matured sufficiently and after the formation of a pool of people who have experience working with the new processor. General adopters expect that a stable set of development tools, IP blocks or software, and expert support staff is available to support the new processor so that they can focus their engineering resources on completing their design in the shortest time without

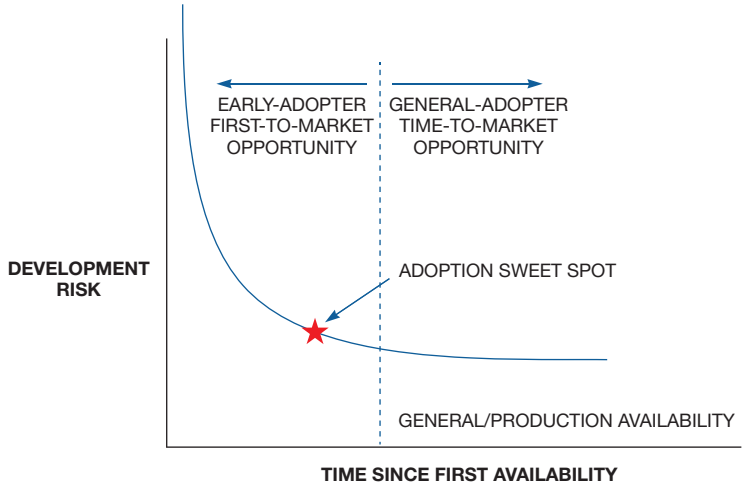

Figure 1 The adoption sweet spot provides the most opportunity for a first-to-market design and incurs the lowest additional development risk (courtesy Embedded Insights).

necessarily requiring heroic engineering efforts outside their field of expertise. In other words, accessing the new features is fairly straightforward and well-defined, such as using available drivers, libraries, or API (application-programming interfaces).

Two previous hands-on projects illustrate some of the possible barriers between expectations and reality for early adopters (**references 1** and **2**). Both projects involved a recently launched system that comprised an ARM core and a DSP core within the same device. Some of the challenges included learning how to configure and use the system and the software at the same time that the engineering-support people were learning it, using drivers that were "hot off the press" to meet the project deadline, and having only one choice for the host-workstation configuration because the alternative configurations would not become available for several more months.

A key lesson from both of these projects is that early adopters must be aware that the ability to get to market before your competitors requires risk-taking in working with components that might

SOURCES OF RISK

As contemporary processor offerings pack more features and on-chip resources into single packages, it becomes increasingly more difficult and time-consuming for a developer to learn and understand how all of the parts of the entire system operate together. An effective way to improve the productivity of developers is to use operating systems, BIOS, drivers, software libraries, development kits, and feature-rich software-development-tool suites. However, these productivity tools take time—typically, a few months—to build and test against the actual processor when it first becomes available.

Early adopters might need to work with the processor before any of these productivity tools target the processor core they are using. One work-around for this problem is to use a simulator or an available processor that is similar to the new one; the corresponding risk is that the new processor implements some features in a subtly different way that will require the early adopter to rework a portion of the system to accommodate those differences once they are identified.

Development-tool providers are obvious candidates for early adoption of new processors because they supply the tools that general adopters must rely on. Development-tool providers support multiple versions of their tools for each processor target because their customer base does not automatically adopt the latest version of the tools. Moving to the latest revision may involve gaining access to a new capability that earlier versions do not support. In any case, a business-case scenario justifies the risk of moving to the latest version. Early-adopter issues may differ in the details, but they are similar whether you are an FPGA-, a board-level-, or a software-tool provider.

Rodger Hosking, vice president of Pentek, shares his company's experience with supporting new FPGAs. The FPGA-development tools may undergo modification because of hardware changes for as much as a year of production because the manufacturer is still "cooking" the silicon. During this time, the hardware vendor is still tweaking the system IP (intellectual property) to improve the performance of blocks in the system. Developers considering the purchase of a new FPGA frequently ask whether the vendor has tested it at a certain speed and what tools are available to support that level of performance. The answers to these questions are relevant because new silicon sometimes does not immediately support every function at the expected top clock rate. The details of the answers to these two questions provide prospective adopters with a sense of the maturity, stability, and engineering effort they will need to bring to a project if they adopt that FPGA at that moment.

Jason Kridner, a BeagleBoard expert at Texas Instruments, weighs in from the perspective of supporting a board-level resource for the BeagleBoard community. Important areas affecting the stability of a new board include interconnect and interprocessor communication. It takes weeks for the system to become stable after a change in ISA (instruction-set architecture) or in specialized processing engines in the target processor that affect the timing of the system. The support software, including the supplied kernel, drivers, and operating system, encapsulates such changes. Examples of such changes in the software could be to enable a driver to support a specific DMA (direct-memory-access) mode or to activate a vector processor. Another type of change that can occur and causes rework and delay is to substitute a component on the board with a similar part, such as a memory chip; the substitution can cause a ripple of changes in the system that can take a few weeks for testing and deployment.

Robert Day, vice president of marketing at LynuxWorks, sheds light on a different set of risks facing software-tool providers. It's easier to provide early support for some processors, such as those from Intel, because the embedded variants employ the PC or the server architectures that precede by several months the availability of the embedded parts. In contrast, a significantly longer lag occurs between when the IP core for ARM processors debuts and when actual silicon is first available. The development tools for the embedded processors come from a base that focuses on embedded-system constraints.

However, providing software-development tools for embedded processors involves validating what features the end device includes. This issue is important because each instantiation of an ARM processor by each licensee can expose dissimilarity between expected and missing features. An example with an ARM core involves a Cortex-A8 port that does not implement ARM's TrustZone technology, a systemwide approach to security. An example with an Intel processor involves an Atom core that lacks virtualization support. An instantiation of a processor core sometimes lacks all of the possible specified features.

be undergoing build and test cycles; building components in-house to meet schedule constraints, even though those components will be readily available to later adopters; or both of these scenarios. The payoff is that early adopters may be able to release their products with a new and important differentiating capability many months before their competitors.

Early adopters rarely shoulder all of the development risks, however (**Reference 3**). Rather, the processor vendor often enters into a strategic alliance with the early adopter that involves early access to the silicon, software, tools, engineering team, and dedicated FAEs (field-application engineers). Obvious examples of early adopters include tool providers that will support the new processor with their tools as well as strategic-partner companies that will be the first primary users of the new processor. The strategic-partner company is often key in specifying and testing the processor so that it best meets the requirements of the targeted application.

FIND THE KNEE

It is not feasible for everyone to enter an early-adopter partnership with a processor vendor. So, must smaller design houses resort to adopting a processor after it has reached general availability, or can they find the "knee," or sweet spot, on the early-adopter side of the risk-and-adoption curve that gives them a competitive advantage? A number of engineering and marketing resources, such as FAEs, data sheets, and application notes, are available from the processor vendor to help the public engineering community determine when they should adopt a processor. It is important, however, to realize that each resource has a different timeliness and reliability, depending on the company's procedures for capturing and disseminating engineering changes (**Figure 2**).

A substantial time interval usually occurs between when someone first identifies a bug and when an informal fix is available; more time elapses during the approval and implementation of a formal fix, the updating of the documentation, the release of the fix, the publication of the updated documents, and the awareness of interested developers that something about the new processor has changed.

Working directly with an FAE usually provides the most reliable, accurate, and timely information about a new processor because the FAE can obtain informal information directly from the engineering group and bypass the delays of days to weeks that can be involved in waiting for the results of engineering-change-review boards. The FAE provides the expert judgment to know when to share an engineering change that the vendor is considering but has not yet formally adopted.

Community-support forums provide avenues for finding timely information about using a new processor. Forums typically connect developer peers with each other and engineers from the processor supplier. In some cases, the original developers of the processor may participate in the forums. Active forums can provide a place for developers to ask questions and to read or participate in conversations regarding how to address unexpected issues when operating in a scenario with a processor. Forums provide appropriate means of finding approved-formal-fix information; because they are static and persistent resources, however, they are unlikely to be good resources for finding informal-fix information. Examples of developer community forums include Atmel's AVR Freaks and Texas Instruments' e2e forums. Offering a community forum requires active management to ensure that the material in the forum is relevant and accurate. Without active monitoring, the "signal-to-noise ratio" can suffer from spam posts as well as posts that are incorrect or even misleading—intentionally or otherwise.

A knowledge base—a repository of relevant documents to disseminate information to the developer community—can complement the technical-support group. It may contain troubleshooting information, application notes, articles, white papers, data sheets, errata, bug lists, sample programs, user manuals, hardware-support documents, downloadable software, archived software, and answers to frequently asked questions. Running a knowledge base does not require the same level of active management that a forum needs because it is less interactive than a forum.

Participants sometimes use IRC (Internet-relay-chat) channels to engage in active conversations in active developer communities, such as at BeagleBoard. These channels provide a faster way to ask peers questions and engage in back-

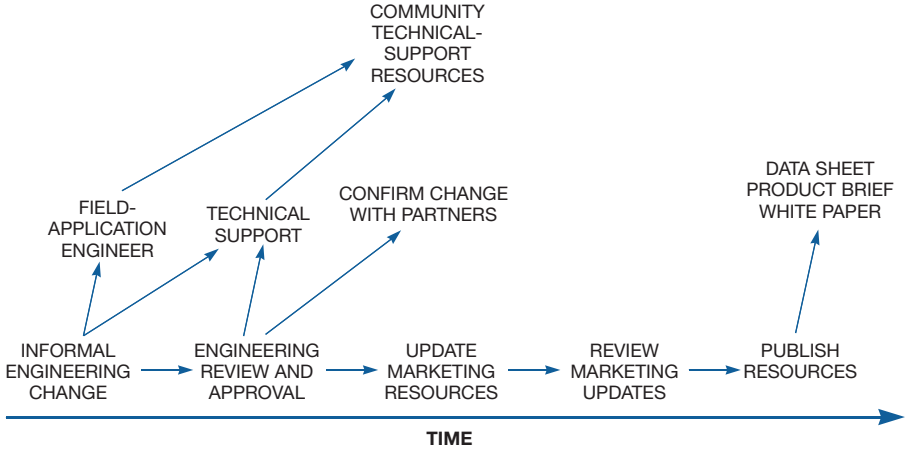

Figure 2 Engineering changes pass through a number of groups and cause a lag between when a change is made and when it is incorporated into the various information channels open to developers (courtesy Embedded Insights).

FOR MORE INFORMATION

ARM
www.arm.com

Atmel
www.atmel.com

BeagleBoard
www.beagleboard.org

Intel
www.intel.com

LynuxWorks
www.lynuxworks.com

Pentek
www.pentek.com

Texas Instruments
www.ti.com

Xilinx
www.xilinx.com

and-forth conversations; however, users of these channels should check the reliability of any information they see with their FAEs and suppliers because these conversations lack full moderation.

Many companies offer an e-mail service to notify developers whenever there is an update to the documentation pertaining to a processor family or a development tool. This service provides developers with reliable information, but relying on e-mail notification can mean that you find out about the latest information weeks later than you could have through other resources. E-mail notification is ideal for those developers who are interested in knowing what changes are happening but are not looking to immediately make a selection.

Adopting a processor too early can cause you to accept more risk than you are prepared to handle. Likewise, adopting a processor past the sweet spot means that you may miss out on a first-to-market opportunity. Hitting the adoption sweet spot for a new processor requires diligence and active effort. Numerous resources keep track of that effort, which together can provide you with a sense of where a processor is on the risk-and-adoption curve. EDN

REFERENCES

1 Cravotta, Robert, "Forge ahead?" *EDN*, March 6, 2003, pg 51, http://bit.ly/oW1f5T.

2 Cravotta, Robert, "Embedded-processing platforms: a second look," *EDN*, Nov 9, 2006, pg 54, http://bit.ly/n73rl0.

3 Cravotta, Robert, "Welcome to the jungle: early adoption in embedded development," *EDN*, Oct 30, 2003, pg 38, http://bit.ly/ncdyFj.

AUTHOR'S BIOGRAPHY

Robert Cravotta is principal analyst and co-founder of Embedded Insights. Previously, he worked as a technical editor at EDN, *covering embedded processors. He also worked in aerospace on electronics and controls for pathfinding projects, such as fully autonomous vehicles, spacecraft and aircraft power-management systems, Space Shuttle payloads, and building-automation systems. Cravotta received a master's degree in engineering management from California State University—Northridge and a bachelor's degree in computer-science engineering from the University of California—Los Angeles.*

Unlock the potential of ARM-Powered® designs.

Optimize power consumption, leverage advanced process technology, and gain efficiency in your next embedded design.

200+ speakers, 150+ technical sessions, 100+ ARM IP experts, a single architecture—join the community defining the future.

ARM® TechCon™ 2011 is your "can't miss" opportunity to unlock the proverbial Rubik's cube of design challenges.

- **Immerse** yourself in the latest technologies for chip design utilizing advanced manufacturing technology.
- **Connect** with experts and learn industry-leading strategies and methodologies for optimizing your ARM IP-based design.
- **Learn** proven techniques for accelerating your design with leading operating systems and software.

Advance your knowledge of the best-in-class techniques for designing with the leading embedded microprocessor architecture—ARM. It's simply a smarter way to design everything from sensors to servers. **The design possibilities are endless, but the seats are not. Go to ARMtechcon.com and register today.**

Chip Design Day
October 25

Software & Systems Design Days
October 26-27

The Santa Clara Convention Center
Santa Clara, CA

2011 ARM® TechCon™

UBM Electronics

ARMtechcon.com

Join the community defining the future

BY DWIGHT BYRD AND THOMAS KUGELSTADT • TEXAS INSTRUMENTS

Understanding and comparing the differences in ESD testing

DIFFERENT TESTING STANDARDS FOR ESD EXIST, AND WHICH ONE YOU USE AFFECTS YOUR DESIGN CHALLENGE.

ESD (electrostatic discharge)—the sudden and momentary electric current that flows between two objects at different electrical potentials—causes equipment failure and network downtime, thus causing production losses of multiple billions of dollars annually. From portable consumer electronics to industrial-automation, process-control systems, and military and aerospace applications, every electronics manufacturer must consider ESD during equipment design. Myriad testing standards exist for addressing the range of technical requirements of the various industrial segments.

To help you select the correct testing standard for a design, you need to understand the main ESD standards and the differences between device- and system-level testing. ESD protection includes a range of protection schemes, the most common of which are steering-diode arrays, TVS (transient-voltage-suppressor) diodes, and zener diodes. No matter which protection scheme you select, you must perform a final EMI (electromagnetic-interference) test and a test of the protection circuit itself.

HBM TESTING

The HBM (human-body-model) device-level test is the most common model for ESD testing. It is used to characterize the susceptibility of an electronic component to ESD damage. The test simulates an electrical discharge of a human onto an electronic component, which could occur if a human has built up residual charge—for example, by dragging his feet, in socks, across a carpet and then touching an electronic device. The failure modes for the HBM testing of ICs typically comprise junction damage, metal penetration, melting of metal layers, contact spiking, and damage to the gate oxides.

You set up the test procedure by applying a high-voltage supply in series with a 1-MΩ resistor and a 100-pF capacitor. After the capacitor is fully charged, a switch is used to remove it from the high-voltage supply and series resistor and to apply it in series with a 1.5-kΩ resistor and the DUT (device under test). The voltage thus fully dissipates through the resistor and the DUT (**Figure 1**). Values for the high-voltage supply range, according to the test level, from 0.5 to 15 kV.

Figure 2 shows a typical oscilloscope readout with an initial current spike as large as 1.4 to 1.5A when the capacitor starts discharging and the ramp-down until it asymptotically approaches 0A at approximately 500 nsec. The DUT can experience a maximum power of 22.5 kW at a single discharge event on a traditional HBM. Keep in mind that power equals the current times the voltage.

MM TESTING

The MM (machine-model) device-level test, which emerged in the 1990s, is now less common than the HBM test. Industrial-automation-manufacturing sites became increasingly popular in the '90s to increase output. These machines

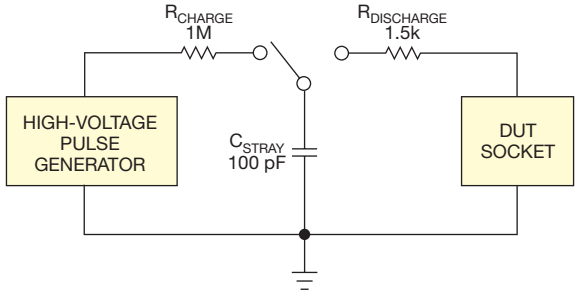

Figure 1 The HBM device-level test is the most common model for ESD testing. You use it to characterize the susceptibility of an electronic component to ESD damage.

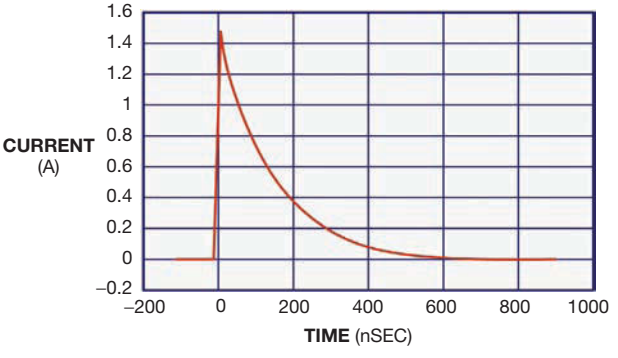

Figure 2 A typical oscilloscope readout shows an initial current spike as large as 1.4 to 1.5A when the capacitor starts discharging and the ramp-down until it asymptotically approaches 0A at approximately 500 nsec.

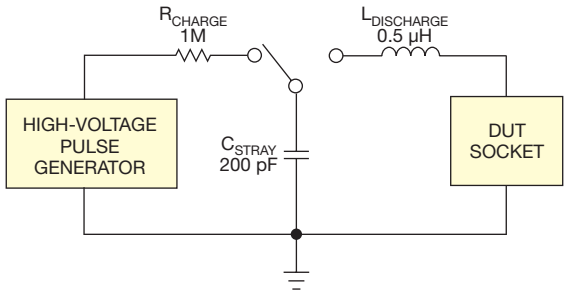

Figure 3 You set up the test procedure for MM testing with a high-voltage supply in series with a resistor and a 200-pF capacitor.

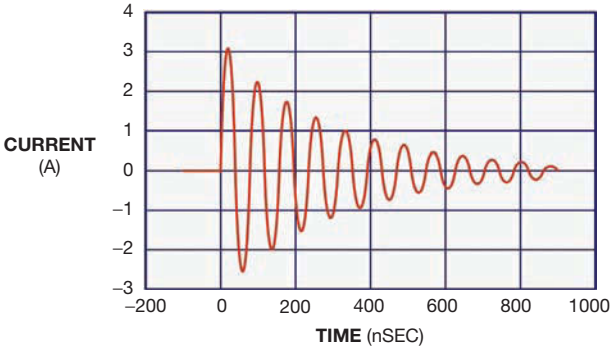

Figure 4 An oscilloscope measurement of current over time shows that the RLC circuit creates an alternating current.

become electrically charged after turn-on and discharge into an electronic component after making contact. Thus, MM tests became a model for testing this type of ESD event. Failure modes in MM testing are similar to those in HBM testing. These failure modes include junction damage, melting metal layers, and gate-oxide damage.

You set up the test procedure for MM testing with a high-voltage supply in series with a resistor and a 200-pF capacitor. After the capacitor fully charges, a switch is used to remove it from the high-voltage supply and series resistor and then apply it in series to a 0.5-µH inductor and the DUT. The inductor with the capacitor voltage dissipates through the DUT (**Figure 3**). Traditional values for the high-voltage supply can vary, but the most common range is 50 to 400V.

When looking at an oscilloscope measurement of current over time, you can see that the RLC (resistance/inductance/capacitance) circuit creates an alternating current (**Figure 4**). The current reaches approximately ±3A, which is about four times higher than the HBM's peak-to-peak current amplitude. Furthermore, the dissipation is much longer for the MM test because it is still asymptotically approaching 0A at 900 nsec (**Figure 4**). The DUT experiences a maximum power dissipation of approximately 1.2 kW during an MM discharge event.

MM testing requires that you test each pin on the DUT to its standard. The electronic chip is mounted on a specially designed load board that interfaces with an automated ESD tester. You ground the other pins on the board and then individually test each pin. You continue this procedure until all pins have been tested.

CDM TESTING

The CDM (charged-device-model) device-level testing procedure is a simulation for situations that often happen in automated-manufacturing environments in which machines often remain on indefinitely, causing the electronic ICs to electrically charge over time. When the part comes into contact with a grounded conductor, the built-up residual capacitance discharges. For the CDM test, the DUT is placed on its back facing upward on a testing board.

Separate the metal field plate and the DUT with an insulating material, which acts as a capacitor between the two objects. You then connect the metal field plate to a high-voltage supply and increase its voltage to the required CDM-test-voltage level. A probe then approaches the pin under test where an ESD event occurs. Monitoring the ground connection of the pin under test verifies this action. Repeat this test on each pin of the DUT for three positive and three negative pulses. The result is six total discharges per pin (**Figure 5**).

Figure 6 indicates that the CDM discharge takes place over 2 nsec at most, which makes it difficult to test and to model. This test results in a current of 5 to 6A discharging in less than 1 nsec. The current dissipates within 5 nsec, making this part of the test succinct but volatile. Due to this fast transient, the failure modes typically seen in CDM tests are gate-oxide damage, charge trapping, and junction damage. **Figure 6** shows the current waveform during a CDM test.

The HBM, MM, and CDM are the most common ESD device-level testing procedures for electronic components. **Table 1** summarizes their similarities and differences.

ESD IMMUNITY

The system-level ESD-immunity test simulates the ESD of a human onto an electronic component (**Figure 7a**). Electrostatic charge on a human can develop in low relative humidity, on low-conductivity carpets, and on vinyl garments. To simulate a discharge event, an ESD generator applies ESD

TABLE 1 COMPARISON OF HBM, MM, AND CDM TESTING

Model	HBM	MM	CDM
Test levels	2, 4, 8, and 15 kV	100, 150, and 200V	250, 500, 750, and 1000V
Peak current (A)	1.5	±3	5 to 6
Pulse width (nsec)	Approximately 150	Approximately 80	Approximately 1
Rise time	2 to 10 nsec	Approximately 1 nsec	Less than 400 psec
Typical ESD failures	Junction damage, metal penetration, metal melting, contact spiking, and gate-oxide damage	Junction damage, metal melting, and gate-oxide damage	Gate-oxide damage, charge trapping, and junction damage

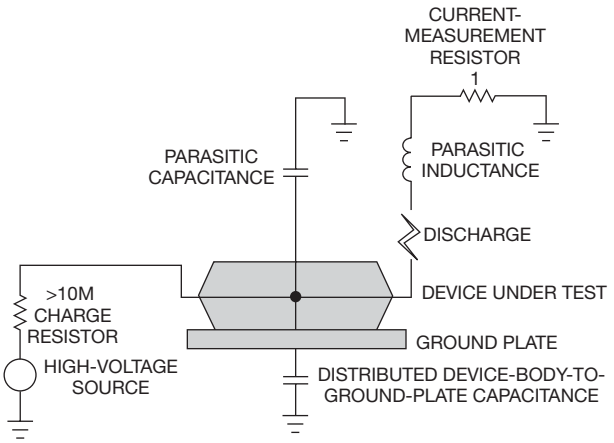

Figure 5 Separate the metal field plate and the DUT with an insulating material, which acts as a capacitor between the two objects. You then connect the metal field plate to a high-voltage supply and increase its voltage to the required CDM-test-voltage level.

pulses to the EUT (equipment under test) in two ways. The first is through contact discharge, or direct contact with the EUT, in which something makes physical contact with the EUT. The second is through air-gap discharge, or indirect contact with the EUT, in which the discharge occurs through the air. The IEC (International Electrotechnical Commission) defines this test in the IEC61000-4-2 ESD-immunity-test specification.

Characteristics for this test are a rise time of less than 10 nsec and a pulse width of approximately 100 nsec, indicating a low-energy, static pulse. The ESD-immunity test requires that you administer at least 10 discharges of both positive and negative polarity at 1-sec intervals. Thus, you test the EUT at least 20 times for the ESD-immunity system-level specification (**Figure 7b**).

Figure 8 shows the differences between device- and system-level testing standards. The IEC ESD test, which many consider the gold standard for component testing, typically has an eight-times-higher testing voltage than CDM and 20-times-higher peak-current testing than HBM.

EFT IMMUNITY

The system-level-testing standard of IEC61000-4-4 is the EFT (electrical-fast-transient) immunity-testing model (**Figure 9a**). The EFT-, or burst-, immunity test simulates transients that can happen in everyday environments due to switching off inductive loads, relay-contact bounce, and the operation of dc or universal motors. This test is performed on all power, signal, and earth wires. A burst is the sequence of pulses with a finite duration. In the EFT-immunity test, a burst generator produces a sequence of test pulses that attenuate to 50% of their peak values in less than 100 nsec. The next adjacent pulse typically occurs 1 µsec later. A burst typically lasts for 15 msec, and the burst period, the time from one burst's start to the next burst's start, is 300 msec. This cycle repeats for 10 sec, after which there is no testing for 10 seconds. This scenario represents one test cycle. The test cycle must repeat six times, taking 110 sec. The significance of the EFT-immunity test is its short pulse rise times, high repetition rates, and low energy content.

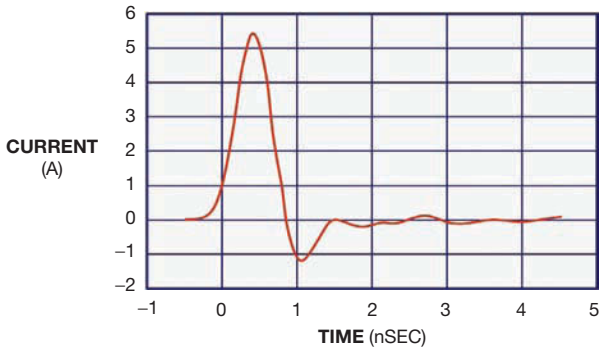

Figure 6 CDM discharge takes place over 2 nsec at most, which makes it difficult to test and to model. This test results in a current of 5 to 6A discharging in less than 1 nsec.

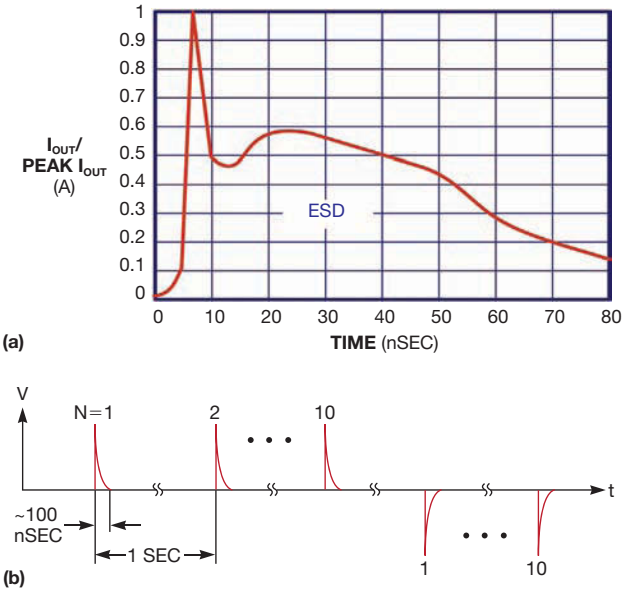

Figure 7 The system-level ESD-immunity test simulates the ESD of a human onto an electronic component (a). You test the EUT at least 20 times for the ESD-immunity system-level specification (b).

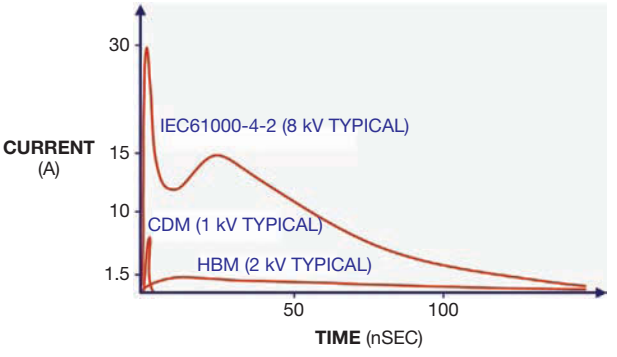

Figure 8 The IEC ESD test typically has an eight-times-higher testing voltage than CDM and a 20-times-higher testing peak current than HBM.

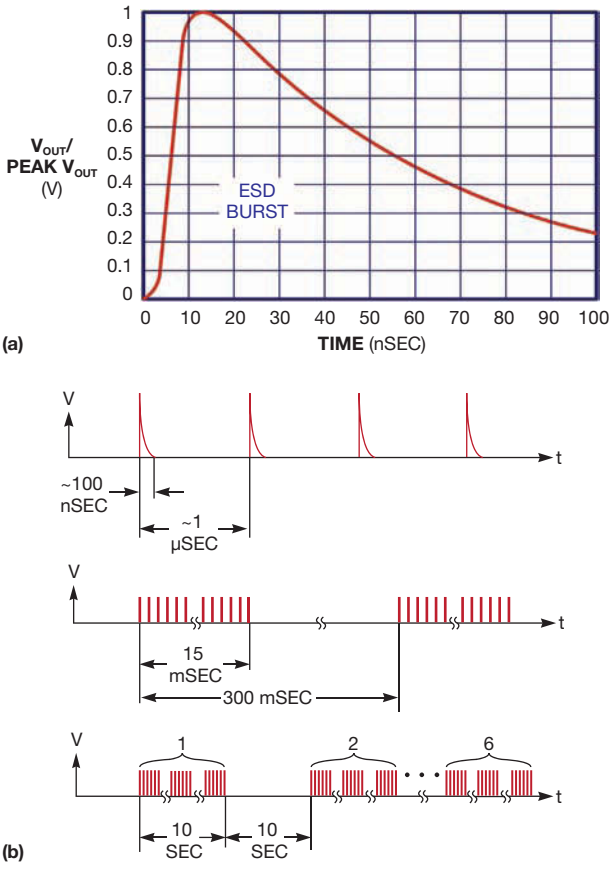

Figure 9 The system-level-testing standard of IEC61000-4-4 is the EFT-immunity-testing model (a). Industrial-grade cabling with internal shielding can produce great results to the DUT by drastically attenuating the coupling of EFT energy into the conductors (b).

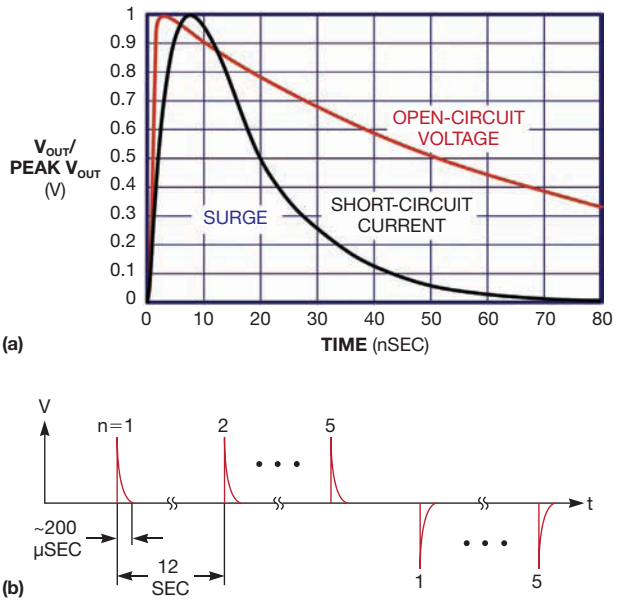

Figure 10 The surge-immunity, or lightning, test, IEC61000-4-5 represents the most severe transient-immunity test in current and duration (a). A common procedure is to shorten the pause intervals to 12 sec, thus reducing total test time to less than two minutes. This approach intensifies the surge impact but significantly reduces test cost (b).

Although the fast rise time and the low energy content of an EFT are somewhat similar to those of an ESD pulse, the number of pulses per test cycle is not. Assuming a 1-µsec interval between one pulse front and the next, a 15-msec EFT burst contains at least 15,000 pulses. Multiplying the number of bursts within a 10-sec window yields 10 sec/300 msec=33.3 bursts and 500,000 pulses per 10-sec window. Thus, the application of six 10-sec windows with a 10-sec pause interval results in 3 million pulses within 110 sec.

Because EFT testing involves no direct contact of conductors but instead the indirect application through a capacitive clamp, proper, industrial-grade cabling with internal shielding can produce great results to the DUT by drastically attenuating the coupling of EFT energy into the conductors (**Figure 9b**).

SURGE IMMUNITY

The surge-immunity, or lightning, test, IEC61000-4-5, represents the most severe transient-immunity test in current and duration (**Figure 10a**). However, testers often employ it on signal and power lines longer than 30m. The surge-immunity test simulates switching transients due to direct lightning strikes; induced voltages and currents due to indirect strikes; or switching the power systems, including load changes and short circuits.

The test specifies the surge generator's output waveforms for open- and short-circuit conditions. The ratio of the open circuit's peak voltage to the short circuit's peak current is the generator's output impedance. High current due to low generator impedance and pulse duration approximately 1000 times longer than the ESD- and EFT-immunity tests characterize this test, indicating a high-energy pulse.

TABLE 2 COMPARISON OF SYSTEM-LEVEL TESTS

Immunity test	Standard	Lines tested	Voltage (kV)
ESD	IEC61000-4-2 air gap	Power, signal	±15
	IEC61000-4-2 contact	Power, signal	±8
EFT/burst	IEC61000-4-4	Power	±4
		Signal	±2
Surge/lightning	IEC61000-4-5 (8 and 20 µsec, 42Ω, 0.5 µF)	Signal	±0.5
	IEC61000-4-5 (8 and 20 µsec, 2Ω, 18 µF)	Power	±1

This test requires five positive- and five negative-surge pulses with a time interval between successive pulses of one minute or less. A common procedure is to shorten the pause intervals to 12 sec, thus reducing total test time to less than two minutes. Although this approach intensifies the surge impact due to the protection circuits' reduced recovery time between pulses, it contributes to a significant reduction in test cost (**Figure 10b**).

SYSTEM-LEVEL TESTING

The IEC compiles the system-level-testing standards according to IEC61000-4. This family of standards includes approximately 25 system-level-testing specifications for transient-immunity testing: IEC61000-4-2 for ESD, IEC61000-4-4 for EFT, and IEC61000-4-5 for lightning. **Table 2** compares these tests.

Today's rising demands for system-level testing renders inadequate device-level testing at the low voltage/current levels of HBM, MM, and CDM. A strong distinction exists between system ESD and burst/surge-level testing between consumer products and industrial equipment and systems, however. In consumer designs, ESD testing assumes a high priority due to the increased probability of human contact with electronic components through cable connectors.

In strong contrast, industrial designers rate the burst- and surge-immunity tests higher than ESD testing. In this case, the daily bombardment of electrical transients due to electric motors and other inductive switching loads poses far greater risks to the system than ESD, whereas human contact occurs only during system installation and maintenance and even then only when the operator is wearing ESD-protection gear. For more information about ESD and testing, visit www.ti.com/esd-ca. EDN

ACKNOWLEDGMENT

This article originally appeared on EDN's *sister site,* Planet Analog Designline, *at http://bit.ly/omxWgn.*

AUTHORS' BIOGRAPHIES

Dwight Byrd is a product-marketing engineer in the precision-analog group at Texas Instruments, where he is responsible for ADCs. Byrd was previously part of TI's interface group, where he was responsible for ESD/EMI products, I²C peripheral devices, signal switches, and voltage-level translators. He received a bachelor's degree in electrical engineering from Texas A&M University (College Station, TX).

Thomas Kugelstadt is an application manager at Texas Instruments, where he is responsible for defining new, high-performance analog products and developing complete systems that detect and condition low-level analog signals in industrial systems. During his 22 years with TI, he has worked in various international application positions in Europe, Asia, and the United States. Kugelstadt is a graduate engineer from the Frankfurt University of Applied Science (Frankfurt, Germany).

Isolated LED Current Control with Active PFC

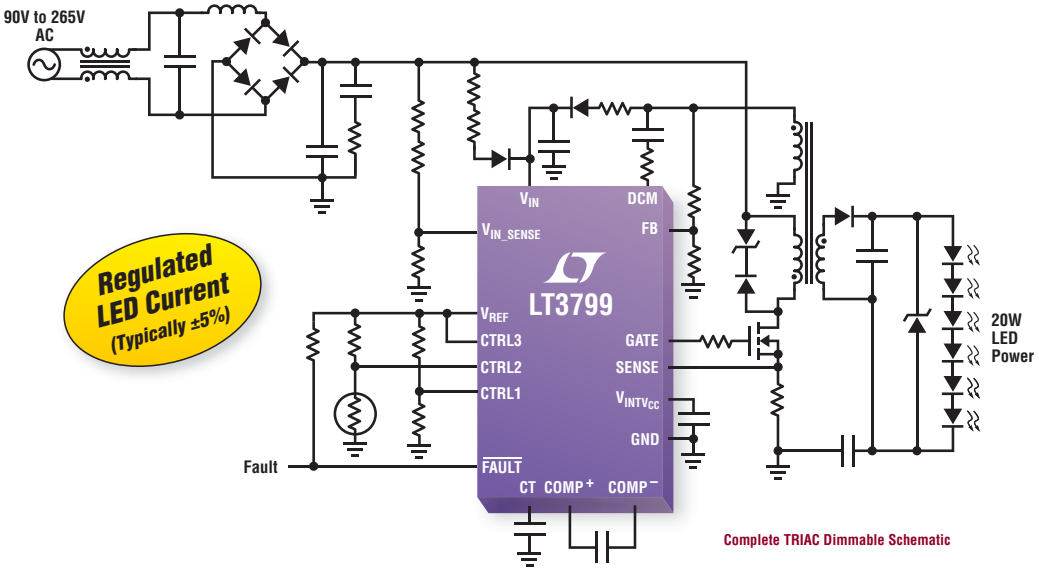

Complete TRIAC Dimmable Schematic

TRIAC Dimmable LED Driver Needs No Opto-coupler

Our LT®3799 isolated LED controller with active power factor correction (PFC) is specifically designed for driving LEDs over a wide input range of 24V to 480V+. It is ideal for LED applications requiring 4W to over 100W of LED power and is compatible with standard TRIAC in-wall dimmers. The LT3799's unique current sensing scheme delivers a well regulated current to the secondary side with no opto-coupler, enabling it to provide ±5% LED current accuracy. It also offers low harmonic distortion while delivering efficiencies as high as 90%. Open and short LED protection ensures long term reliability and a simple, compact solution footprint addresses a wide range of applications.

LED Current vs TRIAC Angle

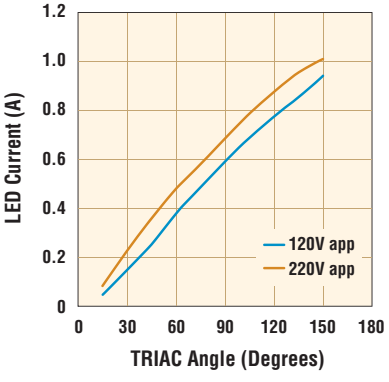

LT3799 Demo Board (25W)

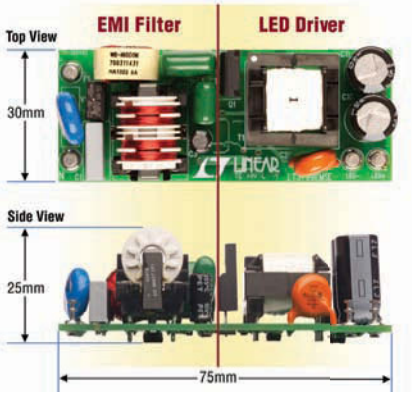

Info & Free Samples

www.linear.com/product/LT3799

1-800-4-LINEAR

LINEAR TECHNOLOGY

EDITED BY PAUL RAKO
AND FRAN GRANVILLE

designideas

READERS SOLVE DESIGN PROBLEMS

Circuit controls inrush current in ac-operated power supplies

Peter Demchenko, Vilnius, Lithuania

Large power supplies that operate from ac wall voltage have large input-filter capacitors. You must limit the inrush current to those capacitors. Otherwise, the supply may trip the ac circuit breaker, or you may damage the rectifier, filer chokes, or PCB (printed-circuit-board) traces. The circuit in this Design Idea inserts a current-limiting resistor in the capacitor's charging path. It senses when the capacitor has charged to a minimum threshold voltage. It then uses a TRIAC (triode alternating current) to short the resistor. Monitoring the capacitor voltage is preferable to monitoring the input current, thus avoiding large load currents during operation that can cause inrush limiting.

This circuit uses a dynistor to detect the minimum threshold voltage across the filter capacitor. Once called a Shockley diode, a dynistor is an asymmetric thyristor with an alternating P+ and N– structure in its anode. In this break-over diode, avalanche current triggers a gateless thyristor. The unidirectional dynistor differs from the bidirectional STMicroelectronics (www.st.com) Trisil, the Bourns (www.bourns.com) TISP (thyristor-surge-protector), and the Littelfuse (www.littelfuse.com) SIDACtor (silicon-diode-alternating-current) devices. Dynistors are somewhat rare, but you can still obtain a DB3 type with a 32V break-over voltage. These devices sense capacitor voltages greater than 38V. You can use PNP/NPN transistor pairs or a low-power thyristor with a zener diode to emulate a dynistor of any voltage.

A dynistor makes a 700-Hz oscillator when the capacitor reaches the threshold voltage. You can easily transmit this ac signal across an isolation boundary comprising a transformer or capacitors. If you are unsure of the isolation rating of your signal transformer, you can capacitively couple the transformer (**Figure 1**). If you use optional capacitors C_2 and C_3, their rating voltage should exceed 800V. Power resistor R_3 limits inrush current and should have a power rating greater than 2 to 10W. Resistors with integrated thermofuses are preferable. Use the ZCD (zero-crossing detector) to synchronize R_3's shorting event with the transition of the ac voltage through 0V. Parts such as Fairchild's (www.fairchildsemi.com) MOC3062M zero-crossing phototriac-driver optocoupler serve this function.

Dynistor D_1, capacitor C_1, resistor R_1, and transformer T_2 comprise an oscillator. It begins working when the value of E exceeds the value of the dynistor's break-over voltage. The oscillator provides current pulses greater than 20 mA, enough to trigger many types of TRIACs and consuming less than 1.5 mA dc. Because the frequency of pulses is approximately 700 Hz, transformer T_2 is small. Resistor R_2 limits the discharge current through dynistor D_1. If the transformer has adequate dc resistance, you can omit R_2. Choose a TRIAC with a gate-trigger current lower than 20 mA. You may not need the

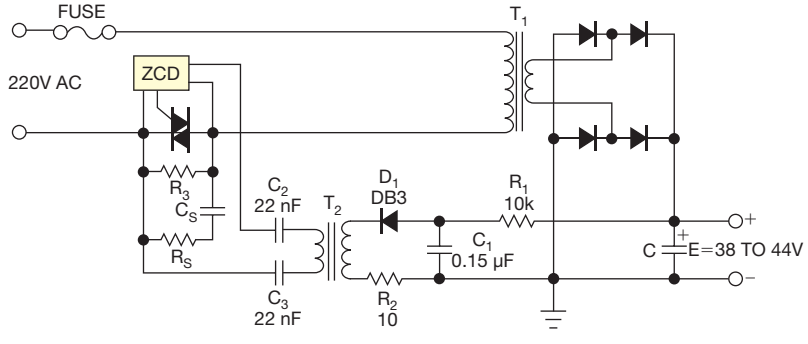

Figure 1 This circuit uses a dynistor to create a 700-Hz oscillator when the voltage on the filter capacitor reaches 32V.

DIs Inside

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

▸Have your own Design Idea to share? Find out how at http://bit.ly/DesignIdeasGuide.

snubber network comprising R_S and C_S if the leakage inductance of T_1 and the inductance of the ac lines are low.

You can adapt the circuit for capacitor voltages greater than 45V (**Figure 2**). Feed the dynistor oscillator through a voltage divider comprising R_4 and R_5. This divider consumes 10 to 20 mA but keeps the oscillator's frequency close to 700 Hz. To avoid dc-current draw, you can use or simulate a higher-voltage dynistor. This circuit dispenses with the isolation transformer and uses capacitors C_2 and C_3. Replacing R_1 in **Figure 1** with R_{11} and R_{12} helps reduce current injection into earth ground and audio interference due to 700-Hz oscillations. EDN

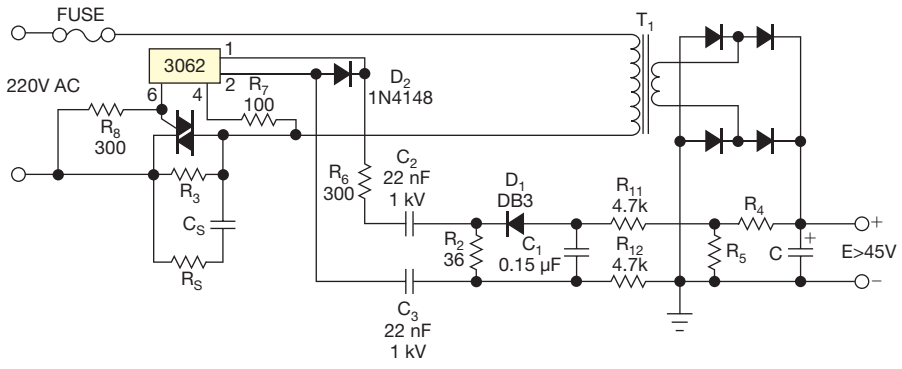

Figure 2 By using the R_4/R_5 voltage divider, you can use the same dynistor to sense capacitor voltages greater than 45V.

Save 3 dB of output power using feedback to set the output impedance

Vic Jordan, Carson City, NV

It's common practice to series-terminate an op amp to match the impedance of the load. This practice causes a 3-dB loss of output power in the termination resistance, however (**Figure 1**). Newer op amps operating on 3 and 5V have limited output swings, meaning that you should avoid using series-build-out resistors. Instead, you can use a series-feedback circuit to set the output impedance. John Wittman, then a senior staff engineer at GTE Lenkurt Electric Co, introduced this technique more than 40 years ago.

With this technique, you add 6 dB of the reciprocal feedback to set the output impedance, obtaining a return loss of more than 30 dB. You add a series-current-sensing resistor, another op amp, and a limiting resistor (**Figure 2**). This example shows a high-side sensor and an unbalanced load. The forward amplifier is designed to have twice the needed gain when unloaded. In this example, the open-circuit gain is 2.7, and the input impedance is 1Ω. The input current is 1A, with an input signal of 1V.

To match the amplifier's 1Ω load, the series-feedback circuitry must divert one-half of the input current from the negative input node of the op amp. The original 1A input current that flows through R_F then decreases to 0.5A, meaning that the output voltage is half of the open-circuit voltage. The output impedance is now 1Ω, and the series feedback is 6 dB, allowing you to match the output impedance to the load and still get almost the full voltage swing from the amplifier. You no longer waste half the output power in a series termination. This example uses a current-sense resistance value that is 3% of the output load, so power loss will be 3%. With careful design, you can lower the loss to less than 1%.

In telecommunications lines, for longitudinal balance, the impedance of both conductors to ground should be the same. Longitudinal balance protects against crosstalk and 60-Hz-induced noise. It is also important at the higher frequencies that DSL (digital-subscriber-line) service uses. Telecom companies commonly use transformers to provide longitudinal balance of 80 to 120 dB. The transformers also isolate transients, such as those due to lightning. You can apply this technique using transformer coupling and low-side current sensing

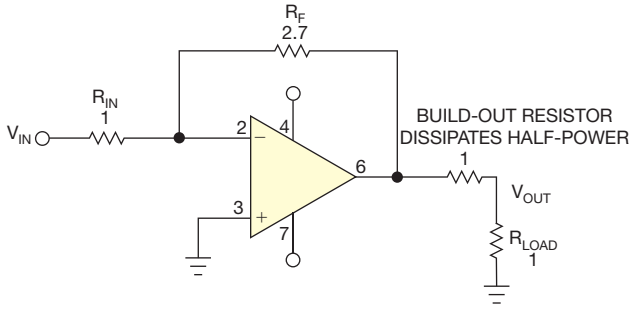

Figure 1 Using a series-termination resistor equal to the load wastes 3 dB of power and reduces the output swing by one-half.

Keep Industrial Controls Online

Power-supply managers with "black box" functionality increase uptime

Failure may not be an option, but it's always a possibility. That's why Maxim's power-supply managers embed a unique "black box" failure-logging function, along with multiple channels of power and thermal management. Measurements are stored in on-chip flash, allowing you to rapidly determine root cause and get your customers back online quickly.

Maximize Safety and Reliability

- Monitor up to six power supplies and four remote temperature sensors
- Prevent impaired equipment from being throttled by implementing a quick system-diagnostic check
- Reduce downtime using root-cause data from the "black box" failure log
- Support standard sequencing, monitoring, margining, and fault-management commands over a PMBus™ interface

PMBus is a trademark of SMIF, Inc.

Watch the product video to learn more.
www.maxim-ic.com/black-box

MAXIM
INNOVATION DELIVERED®

MAXIM DIRECT www.maxim-ic.com/shop
AVNET electronics marketing www.em.avnet.com/maxim
Digi-Key www.digikey.com/maxim
MOUSER ELECTRONICS www.mouser.com/maximic

(**Figure 3**). The design process is still the same, except that just two resistors can provide 6 dB of feedback.

You can formalize the analysis of the circuits using state **equations**. For the circuit of **Figure 2**, because the negative input of an op amp is at virtual ground, you can relate input voltage and current by inspection: $I_{IN}=(V_{IN}-V-)/1\Omega=V_{IN}/1\Omega$. You can derive a second **equation** employing the fact that the negative input of an op amp has high impedance, so the currents at that node must add to 0A.

Sum the currents at V–, except you reference the node currents back to the sense resistor, including the 0.3Ω resistor and the 0.03Ω sense resistor: $0=V_{IN}/1\Omega+V_{OUT}/2.7\Omega+0.37V_{OUT}/R_{LOAD}$. You can express the circuit function in vector and matrix terms: (I)=(**ADMITTANCE**)×(**V**). You can also expand for the appropriate states of current:

$$(I)=\begin{pmatrix}I_1\\I_2\end{pmatrix},\ \begin{matrix}I_1=I_{IN}\\I_2=0\end{matrix}$$

and then expand for the vector expression of voltage:

$$(V)=\begin{pmatrix}V_1\\V_2\end{pmatrix},\ \begin{matrix}V_1=V_{IN}\\V_2=V_{OUT}\end{matrix}$$

Plug these values into the (I)=(**ADMITTANCE**)×(**V**) **equation** and solve for (**V**):

$$(\mathbf{ADMITTANCE})=\begin{pmatrix}1 & 0\\ 1 & \frac{0.37}{R_{LOAD}}+\frac{1}{2.7}\end{pmatrix}.$$

For the circuit of **Figure 2**, the forcing function is I_1; the input current is 1A. Invert the admittance matrix and then multiply the current vector to find the voltage vector. You can use a Hewlett-Packard (www.hp.com) HP-48 calculator to do the hard work. It yields the result that V_{IN} is 1V and calculates V_{OUT} at –1.35V, one-half the unloaded gain of 2.7. You then repeat the analysis for a load resistance of 1000Ω:

$$(\mathbf{ADMITTANCE})=\begin{pmatrix}1 & 0\\ 1 & \frac{0.37}{1000}+\frac{1}{2.7}\end{pmatrix}.$$

Inverting the (**Y**) matrix and multiplying it with the I matrix, with I_1 of 1A, yields an open-load voltage vector, with V_{IN} equal to 1V and V_{OUT} equal to –2.7V, confirming that the design is correct.

Be careful writing your own equations; two dependent equations can easily lead to an incorrect answer. An HP-48 calculator solves them using the "least-squares" method, but it does not check for a determinant condition of zero to warn you of nonindependent equations. You can use the HP-48 to sum two real matrices to form a complex one. This approach comes in handy when you include reactive elements in the circuit models. If you prefer using a computer rather than a paper napkin, you can also use Spice to analyze this circuit.

Three **equations** are used to analyze the circuit of **Figure 3**. You can express the input current as a function of the input resistance: $I_{IN}=(V_{IN}-V-)/R_{IN}=V_{IN}/R_{IN}$. As in the previous example, you sum the currents at the amplifier's negative pin to zero, $0=V_{IN}/R_{IN}+V_{OUT}/28k\Omega+(V_4-V-)/900\Omega$, and then sum the currents at the V_4 node: $0=(V_4-V-)/900\Omega+(V_4-V_{OUT})/R_{LOAD}+V_4/20\Omega$.

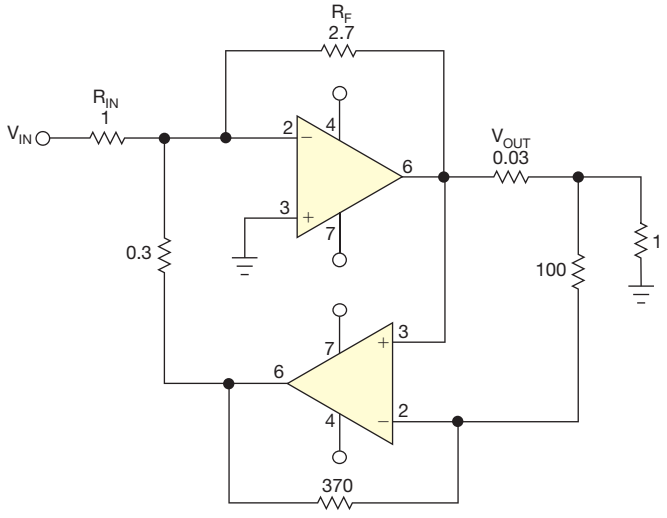

Figure 2 This scheme uses a high-side-current-sense resistor and a second amplifier to set the output impedance to match the load, and it allows almost full output swing.

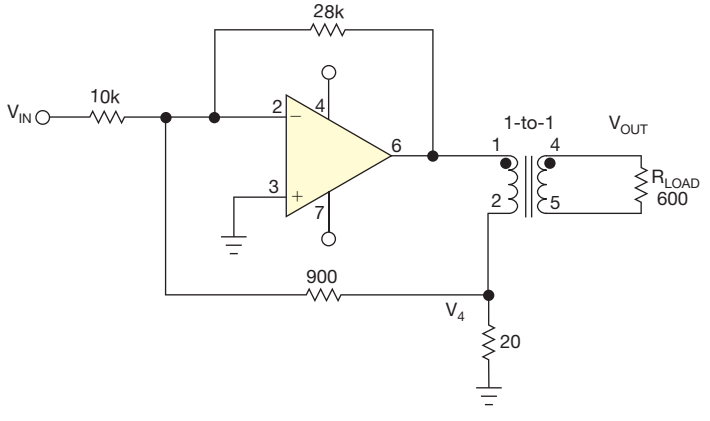

Figure 3 You can also use low-side sensing of output current to match the output impedance when you have a transformer-coupled output.

You express the currents as a vector:

$$(I)=\begin{pmatrix} I_{IN} \\ 0 \\ 0 \end{pmatrix}, \text{ where } (V)=\begin{pmatrix} V_{IN} \\ V_O \\ V_4 \end{pmatrix}.$$

The admittance matrix becomes:

$$(\text{ADMITTANCE})=\begin{pmatrix} \frac{1}{R_{IN}} & 0 & 0 \\ \frac{1}{R_{IN}} & \frac{1}{R_F} & \frac{1}{900\Omega} \\ 0 & \frac{-1}{R_{LOAD}} & \frac{1}{20\Omega}+\frac{1}{900\Omega} \end{pmatrix}.$$

This **equation** determines (**Y**), the admittance matrix.

In this case, the input current should be 100 µA, and the load resistance should be 600Ω. Use an HP-48 calculator to invert the admittance matrix and multiply it by the current matrix. The resulting voltage matrix yields an input voltage of 1V, an output voltage of –1.4V, and a V_4 of –0.05V. Next, set the load to 10,000Ω. Assume that the magnetizing inductance of the transformer is infinite. You then repeat the exercise to check that the output voltage is |2.8V|.

You can match the maximum available signal power from the op amp to the load by changing the transformer turns ratio. Calculate the optimum op-amp signal output impedance to be the peak output swing voltage divided by the maximum peak capability of the op amp. EDN

Optically isolated overcurrent detector works from ac mains

José M Espí, Jaime Castelló, Rafael García-Gil, and Marcos Arasa, University of Valencia, Valencia, Spain

The circuit in this Design Idea detects short-circuit faults or excess-current conditions in line-operated equipment, such as a UPS (uninterruptible power supply), across an isolation boundary. If the load current rises higher than a predetermined level, the circuit generates a 5V pulse to immediately shut down the power inverter. This compact circuit provides a reliable means for protecting power semiconductors and batteries against overcurrent conditions. It features an isolated output and a fast response. The ac mains self-power the sensing network. The circuit requires a 5V power supply to the output-enable circuit on the other side of the isolation boundary. The circuit detects bipolar over-

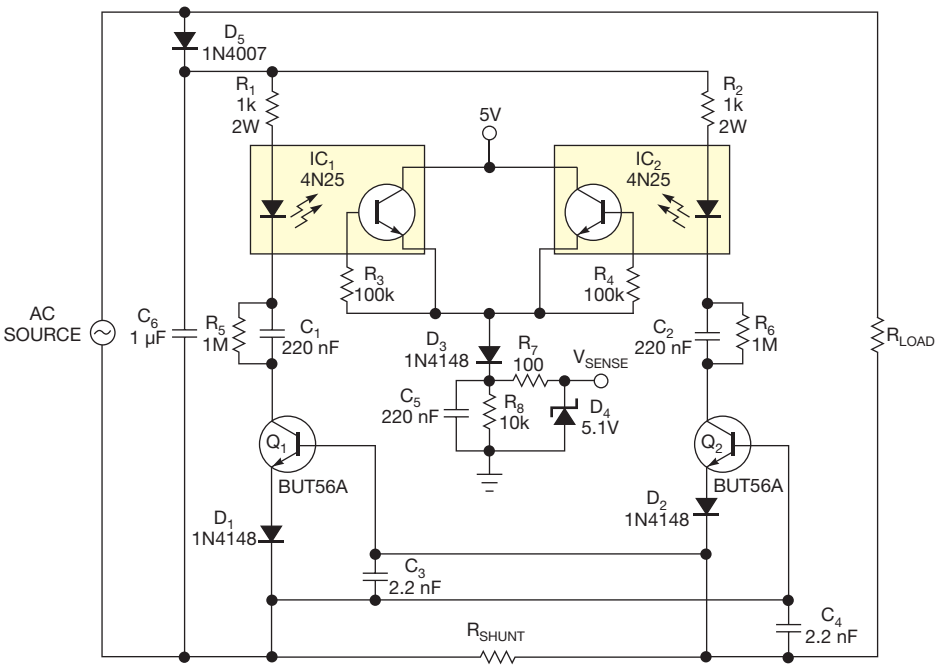

Figure 1 This isolated circuit detects bipolar overloads in a UPS or another ac load.

currents through the load, using shunt resistor R_{SHUNT}'s resistance.

During the positive cycle of the ac source, components R_1, IC_1, R_5, C_1, Q_1, D_1, and C_3 work to detect an overcurrent (**Figure 1**). If the load produces an overcurrent, the voltage drop at the shunt resistor generates a positive voltage on the base of Q_1 that turns it on. Consequently, capacitor C_1, which R_5 initially discharges, receives charge from current across R_1 and IC_1's photodiode. IC_1's photoresistor saturates, charging capacitor C_5 and providing a 5V output at V_{SENSE}, indicating the overcurrent detection. D_1 increases Q_1's base-voltage turn-on threshold to approximately 1.4V. Capacitor C_3 filters any noise or peak voltage that may accidentally turn on the transistor, thus enhancing the system's detection reliability.

A circuit, comprising R_2, IC_2, R_6, C_2, Q_2, D_2, and C_4, for fault detection during the negative cycle of the ac source works similarly. In this case, the direction of current flow through the shunt resistor generates a positive voltage on the base of Q_2. D_5 and C_6 rectify the ac source, providing the isolated supply voltage to polarize the photodiodes. The inverter logic's power supply supplies the 5V source. When the load currents are within limits, the output at V_{SENSE} extinguishes due to the discharge of C_5 through R_8. EDN

Simple circuit helps to protect a vehicle-reverse camera

Aruna Prabath Rubasinghe, University of Moratuwa, Moratuwa, Sri Lanka

The circuit in this Design Idea uses a simple comparator circuit to make a power-on time delay for an automotive rearview camera. Auto manufacturers typically power reverse-view cameras from the reverse-light circuit. In automatic-transmission vehicles, a short power pulse is applied to the camera when you shift through reverse as you go from park to drive, or vice versa. This sudden voltage pulse is bad for the sensitive circuits in the camera and may reduce its lifetime. This Design Idea suggests a simple and cheap method for avoiding this situation.

The input to this circuit connects to the positive and negative terminals of the reverse light (**Figure 1**). The circuit powers the camera using a MOSFET. R_1 and C_1 form a time-delay element (**Reference 1**). When the reverse light turns on, it slowly charges the capacitor through resistor R_1. R_3 and R_4 form a voltage divider, which you use to set 6V on the inverting pin of the comparator. At the instant of power application to the circuit, the comparator output is

low, and the MOSFET is off. Once the voltage of C_1 rises above 6V, the comparator's output becomes high, and the MOSFET turns on. The values of R_1 and C_1 set the time delay to 2.2 sec. You can calculate this time based on the exponential charging of a capacitor using the following **equations**:

$$V_C(t)=V_{IN}\times\left[1-e^{-(t/R_1C_1)}\right],$$

$$6=12\times\left[1-e^{-(t/33\ k\Omega\times100\ \mu F)}\right],$$

and

$$t=2.287\ SEC.$$

You can set a different time delay by changing the value of R_1 or C_1. When you shift the gear lever from the reverse position to any other position, capacitor C_1 discharges within 60 msec through D_1, R_3, and R_4. As you pass through reverse, shifting between park and drive, the camera does not turn on due to the 2-sec delay. EDN

AS YOU PASS THROUGH REVERSE, SHIFTING BETWEEN PARK AND DRIVE, THE CAMERA DOES NOT TURN ON.

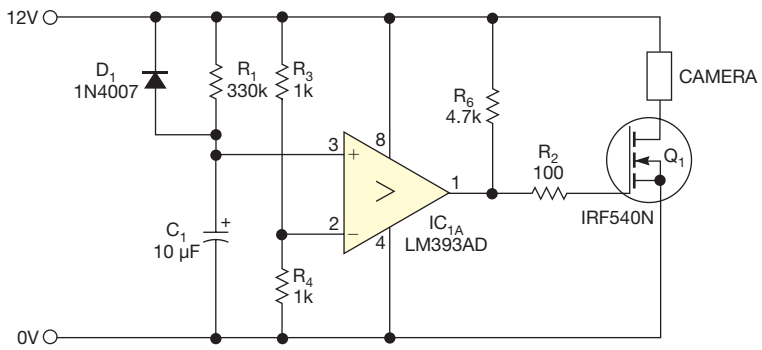

Figure 1 This circuit creates a 2-sec turn-on delay for a car's reverse camera.

REFERENCE

1 "Charging and discharging a capacitor," Kotisivukone, Oct 31, 2007, http://bit.ly/qsmTb6.

supplychain

LINKING DESIGN AND RESOURCES

Smoothing out bumps during supply-chain expansion

Electronics manufacturers are expanding their supply chains for one of two main reasons: They design their products in one region but manufacture them in another, or they want to accelerate growth by expanding organically. In either case, expansion across regions means significant and often costly disruption to the OEM supply chain.

"The biggest challenge we see is expectation," says Ravi Kichloo (**photo**), head of the global business-migration team at Avnet Electronics Marketing (www.em.avnet.com). End customers expect the same level of quality from their manufacturers, no matter where the vendor built the product. OEMs expect consistent pricing and delivery from component suppliers, no matter the location of the manufacturing facility. So companies face a choice: Duplicate their supply chain from scratch or outsource the portion that makes sense. Increasingly, says Kichloo, suppliers and customers are turning to distribution to manage business migration.

Global distributors, such as Arrow Electronics (www.arrow.com), Avnet, Digi-Key (www.digikey.com), Future Electronics (www.futureelectronics.com), Mouser Electronics (www.mouser.com), TTI (www.ttiinc.com), and WPG (www.wpgamericas.com), have invested in the necessary infrastructure to duplicate services across regions. "We already do BOM [bill-of-materials]-analysis, pricing, programming, proximity-warehousing, inventory-pipelining, VMI [vendor-managed-inventory], and buffer programs," says Kichloo. "We have made an effort to support those services globally and maintain the nuances of regional practices. We have the skill, the scope, and the team in place to be able to do this."

Distributors are also helping customers anticipate the types of problems they'll encounter as they make the transition. Component availability and pricing differ from region to region, and transferring IP (intellectual property) for programming purposes puts that IP at risk. "Now that they have multiple locations, they have to look at planning and forecast management, aggregating demand for multiple sites, and new business partners," says Kichloo. "We can make recommendations [on those partners] based on our experience with EMS [electronics-manufacturing-service] providers and contract manufacturers."

Distributors also must consider supplier relationships during a supply-chain transition. Suppliers may lose visibility into customer forecasts as business moves overseas. "We collaborate daily with our suppliers," says Kichloo. Distributors can also help smooth out some of the pricing differences that occur from region to region.

"More suppliers are providing global pricing, but we still have a way to go," he adds.

According to Kichloo,the profile of the global customer is changing. Large manufacturers and OEMs with high-volume, low-mix profiles were the first to expand overseas. Now, second- and third-tier manufacturers—typically, low-volume, high-mix vendors—are also expanding. Because of that complexity, few of these companies are developing their supply chains from the ground up.

"This initiative is important for us and an activity that continues to grow as it becomes more global," says Kichloo. "It's important for suppliers, distributors, and manufacturers to collaborate very closely among themselves. In this way, we can guarantee the transition will be smooth."

—by Barbara Jorgensen, *EBN* Community Editor

This story was originally posted by EBN: *http://bit.ly/oQMD0o.*

OUTLOOK

SMARTPHONES, TABLETS DRIVE MEMS MARKET

The market for MEMS (microelectromechanical-system) devices in consumer electronics and mobile communications devices should reach record growth in 2011, according to IHS iSuppli (www.isuppli.com), which notes that the MEMS blockbuster this year is the three-axis gyroscope.

Consumer- and mobile-MEMS devices make up the largest sector of the market. Revenue should hit $2.25 billion in 2011, up by a record annual-expansion rate of 37%. This rate compares with the previous rate of 27% in 2010, when revenue reached $1.64 billion. Overall, the five-year revenue forecast starting from 2010 calls for growth by a factor of nearly three to $4.54 billion in 2015, equivalent to a compound annual growth rate of 22.5%.

"From the accelerometers and gyroscopes that provide intuitive motion-based displays, to the microphones that allow people to talk on the phone, to the bulk acoustic-wave filters that facilitate wireless-Internet access, MEMS devices provide many of the basic functions that make tablets and smartphones such compelling products," says Jérémie Bouchaud, director and principal analyst for MEMS and sensors at IHS iSuppli.

—by Suzanne Deffree

EE Times Virtual Conference

Upcoming Virtual Conferences

DesignMED virtual event

When: Wed., Oct. 26, 2011 • 11am – 6pm EDT

EE Times virtual event
POWER SOLUTIONS
Smart Energy Technologies and Techniques

When: Thu., Nov. 10, 2011 • 11am – 6pm EDT

2011 ARM TechCon virtual event
JOIN THE COMMUNITY DEFINING THE FUTURE

When: Wed.-Thu., Nov. 16-17, 2011 • 11am – 6pm EDT

On Demand Virtual Conferences

Digi-Key Sensors Virtual Conference
e.ubmelectronics.com/sensors

ESC Silicon Valley Virtual Conference
e.ubmelectronics.com/escsv

System On A Chip
e.ubmelectronics.com/soc

Digi-Key Microcontrollers Virtual Conference
e.ubmelectronics.com/mcu

Industrial Control
e.ubmelectronics.com/industrialcontrol

Multicore: Making Multicore Work for You
e.ubmelectronics.com/multicore

EE Times, the leading resource for design decision makers in the electronics industry brings to you a series of Virtual Conferences. These fully interactive events incorporate online learning, active movement in and out of exhibit booths and sessions, vendor presentations and more. Because the conference is virtual you can experience it from the comfort of your own desk. So you can get right to the industry information and solutions you seek.

Why you should attend:

- Learn from top industry speakers
- Participate in educational sessions in real time
- Easy access to EE Times library of resources
- Interact with experts and vendors at the Virtual Expo Floor
- Find design solutions for your business

For sponsorship information, please contact:
Christian Fahlen, 415-947-6623 or
christian.fahlen@ubm.com

UBM Electronics

productroundup

SENSORS AND TRANSDUCERS

Two key markets—automotive and consumer—are fueling the growth of MEMS (microelectromechanical-system) sensors. Looking at the numbers, it appears as though demand will yield healthy worldwide sales increases in the years ahead. This year's global market for MEMS devices should grow 9.5%, to $7.73 billion, according to a report from market-research company IHS iSuppli (www.isuppli.com).

Catalysts for rising demand in the consumer-electronics and mobile-communications sectors are the use of MEMS sensors, including accelerometers, gyroscopes, magnetometers, pressure sensors, and microphones, in media tablets and smartphones. Revenue for consumer- and mobile-MEMS devices will hit $2.25 billion in 2011, up by an annual expansion rate of 37%, compared with 27%, when revenue reached $1.64 billion, the company says.

Major sensor-device manufacturers are gearing up. STMicroelectronics, for example, has been consistently producing new devices, one of its most recent being a multisensor device that integrates three-axis sensing of linear and angular motion in a 4×5×1-mm package for small consumer electronics (see "ST couples 3-D digital accelerometer, gyroscope in SiP," *EDN*, Sept 8, 2011, http://bit.ly/ny3vov). Analog Devices also recently rolled out a gyroscope that features a patented differential-quad-beam architecture to minimize the influence of linear shock and vibration in automotive applications.

In the automotive segment, global automotive-MEMS-device revenue in 2011 should reach $1.99 billion, up just 4% from $1.91 billion in 2010, compared with last year's 28% expansion. But next year's projections paint a new picture. As vehicle production quickly recovers and countries, including the United States and many in Europe, are enforcing safety-related mandates, the automotive-MEMS-sensor market will return to double-digit growth. According to IHS, global automotive-MEMS-sensor revenue will jump a solid 16% in 2012 to $2.31 billion on its way to $2.93 billion in 2015. **—by Ismini Scouras**

Analog Devices' fail-safe iMEMS gyroscope provides shock, vibration immunity for auto applications

Complying with automotive-safety standards, the ADXRS800 iMEMS gyroscope features a patented differential-quad-beam architecture that minimizes the influence of linear shock and vibration. Sensitivity to linear acceleration is 0.03°/sec/g, vibration rectification is 0.0002°/sec/g², noise-rate density is 0.02°/sec/√Hz at 105°C, and maximum null-offset variation is 3°/sec over temperature and product life. The gyroscope delivers better offset stability without calibration than alternatives. Power consumption is 6 mA under typical conditions. The device comes in a cavity plastic SOIC-16 (Z axis) and an SMT-compatible vertical-mount package (X axis) and operates across an extended temperature range of –40 to +125°C. Prices start at $50 (1000).

Analog Devices, www.analog.com

NXP position sensors boost automotive performance

The KMA210 automotive-position sensor incorporates the vendor's magnetoresistive-sensor chip and a signal-conditioning ASIC the vendor developed using the SOI (silicon-on-insulator) ABCD9-process technology. KMA210 targets use in automotive applications requiring the measurement of a precise mechanical angle—from electronic steering and active suspension to automatic headlight adjustment. The

sensor supports temperatures as high as 160°C, making it well-suited for EGR (exhaust-gas-recirculation) applications. Designing the signal-conditioning ASIC in ABCD9 enhances EMC performance compared with that of previous sensor products with integrated ASICs. The device contains two embedded capacitors in the same package, reducing system cost because it requires no PCBs or external filter components. Other features include contactless angle measurements as large as 180° and overvoltage protection. Prices start at $1.40 each.
NXP Semiconductors, www.nxp.com

Capella's optical sensor eliminates the need for magnetic-lid devices

Targeting use in tablet computers, notebooks, and smartphones, the CM36262 optical-lid-sensor chip contains a patent-pending optical-lid sensor, an ambient-light sensor, and a proximity sensor. The lid sensor makes it easier for users to power down touchscreen products, the ambient-light sensor adjusts screen brightness to extend battery life, and the proximity sensor turns off dialing when a user holds a smartphone in the listening position next to an ear. The CM36262 optically senses the presence of a lid or a fabric cover on touchscreen products, replacing the standard method of using magnets or mechanical switches. The optical method is less expensive and more reliable; it enables users to power down a tablet computer using a fabric or a leather cover without aligning an embedded magnet over a pickup sensor. Prices begin at 79 cents (1000) in an optical-land-grid package that measures 2.35×1.8×1 mm.
Capella Microsystems, www.capellamicro.com

TI reshapes noncontact temperature sensing

The TMP006 single-chip passive IR MEMS temperature sensor enables accurate measurement of device-case temperature. Contactless temperature sensing provides advantages over the current approach of approximating case temperature from the system temperature. Users can also employ the TMP006 to measure temperature outside the device. The TMP006 integrates an on-chip MEMS thermopile sensor, signal conditioning, a 16-bit ADC, a local temperature sensor, and voltage references on a 1.6×1.6-mm chip. It uses 240-µA quiescent current and 1 µA in shutdown mode and supports a temperature range of –40 to +125°C with a typical accuracy of ±0.5°C on the local sensor and ±1°C for the passive IR sensor. It also includes an I²C/SMBus digital interface. The TMP006 sells for $1.50 (1000).
Texas Instruments, www.ti.com

EDN ADVERTISER INDEX

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

The red-cable blues

As a young electronics technician, I got a job working at the end of an assembly line that built a four-channel analog instrumentation recorder. We could build these recorders as AM units to record frequencies as high as 100 kHz or as FM units to record frequencies as low as dc. The problem started when our manager asked the assembly line to build a new run of all-AM units. As these recorders reached the end of the line, it became clear that something was wrong. None of the units could pass the SNR (signal-to-noise-ratio)-specification test, and it was difficult to find playback-amplifier cards that had enough HF equalization to meet the frequency-response requirement.

We initially thought that the run of playback cards was defective. The line chief took all the cards back to the original vendor and asked the company to test them again. Of course, they all passed with flying colors.

We couldn't ship anything because nothing would pass the required tests. The line chief told me to find out what was wrong with this run of AM recorders. I had a reference recorder that would pass all the tests. My plan was to find a different way to test these recorders and compare the results with those of the reference unit. I tested the units with a swept sine wave instead of the fixed signals that the test procedure called for. I then set up the reference recorder and a problem recorder side by side and fed the signal to both units. I probed both recorders at various points throughout the signal path and compared the results on an oscilloscope. After testing the first channel on both recorders, I found no difference between them.

When I moved on to the second channel, however, I found that the first test point on the record amplifier's PCB (printed-circuit board) showed a marked roll-off of the high frequencies. It looked as if the system were lowpass-filtering the input signal. Channels 3 and 4 tested OK. Swapping the record boards didn't change anything. The problem stayed with Channel 2.

I then noticed that each input was wired from a BNC on the front panel to the card-edge connector using a shielded-cable assembly. The cables were different colors—brown, red, orange, and yellow—to identify the corresponding channel. The brown, orange, and yellow cables were all coaxial; their center conductors' insulation was polyurethane foam. The red cable, however, looked more like plain-old PVC (polyvinyl chloride).

The line chief called over the person who had preassembled the cables and asked about the difference. It turned out that the supply room didn't have the red cable that the parts list called for but had some other red cable. The line engineer had approved the use of this cable. Our problems then got worse.

We tried to compensate for the poor HF-record response on Channel 2 by adjusting the playback equalizer to the maximum. This approach caused the playback card to oscillate, and we couldn't detect that oscillation. The oscillation got into everything and was enough to cause the SNR test to fail on all four channels. We rewired the units with the correct cable, and retesting them showed that we had eliminated the problems. The lesson: Although shielded wire and coaxial cable often look similar, they are not the same. EDN

Larry Goga is a retired sound engineer in Albuquerque, NM.

DANIEL VASCONCELLOS

"We're crossing the boundary between the physical and the digital world."

Vida Ilderem, Vice President, Intel Labs, Director, Integrated Platforms Research Lab

EMBEDDED WITH INNOVATION: Intelligent Devices. Smart Connectivity.

Develop intelligent solutions that move markets.

Challenge >

How to meet the growing demand for intelligent connected systems that can transform raw data into smart services.

Solution >

Integrate low-power computing, robust security and flexibility, with always-on connectivity.

Intelligent connectivity is driving unprecedented smart services and business models, and promises to be one of the decade's most disruptive technologies. At Intel, we are reducing development complexity by combining Intel® architecture-based computing, I/O and connectivity in small-footprint modules. Like the Intel® Reference Design for Smart Services Development based on the Intel® Atom™ Processor, a platform that can greatly reduce time and cost for developers of smart services, intelligent connected devices and software solutions.

$35B

Projected annual revenues from M2M services by 2016.[1]

84x44mm

Footprint of the Intel® Reference Design for Smart Services Development.[2]

<5W

DC idle power of the Intel® Reference Design for Smart Services Development.[3]

Maximize your opportunity with the M2M Smart Services Developer Kit based on Intel® architecture. Join the discussion, and see how Intel and our growing M2M ecosystem are creating new possibilities in embedded technology.

READ > WATCH > DISCUSS >>> **intel.com/embedded/innovation/connectivity**

Embedded