UBM Electronics

EDN

ELECTRONICS DESIGN NETWORK

MAY 10
Issue 9/2012
www.edn.com

TEARDOWN REVEALS CHEVY VOLT'S ELECTRONIC SECRETS

Page 32

EDN contents

5.10.12

Teardown reveals Chevy Volt's electronic secrets

Generating spatial audio from portable products

Understand and characterize envelope-tracking power amplifiers

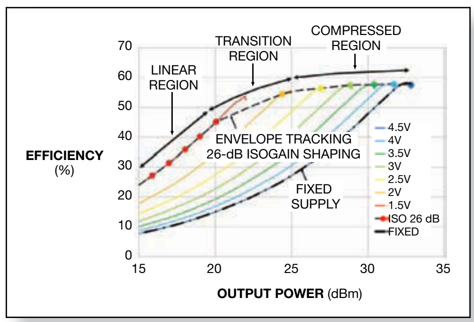

COVER IMAGE: THINKSTOCK/DAVID NICASTRO

DESIGNIDEAS

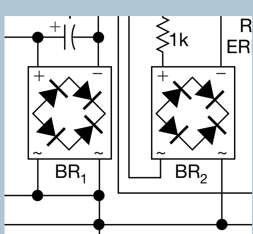

▶ Find out how to submit your own Design Idea: http://bit.ly/DesignIdeasGuide.

The path of least resistance is through our new high efficiency XAL/XFL inductors

Compared to competitive parts, current zips right through our new XAL/XFL inductors.

Their DC resistance is significantly lower than other inductors of the same size. So your batteries last longer and your power supply runs cooler.

They easily handle large peak current, and their soft saturation characteristics often let you use a smaller size part without fear of overloading during transients. Built from a uniquely formulated material, XAL/XFL parts do not have the same thermal aging problems as some competitive parts.

See all the advantages of these new high efficiency inductors. Visit **coilcraft.com/xal.**

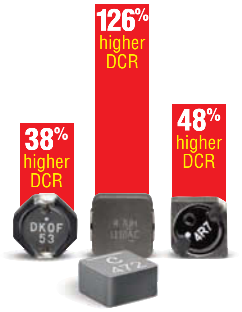

Competitors' 4.7uH inductors have much higher DCR per mm³ than Coilcraft's XAL5030.

Coilcraft®

coilcraftdirect.com
No min order. Next day delivery.

WWW.COILCRAFT.COM

contents 5.10.12

pulse

DEPARTMENTS & COLUMNS

55

58

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 EDN, PO Box 3609, Northbrook, IL 60065-3257. 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 3609, Northbrook, IL 60065-3257. CANADA POST: Publications Mail Agreement 40612608. Return undeliverable Canadian addresses to APC, PO Box 503, RPO West BVR CRE, Rich Hill. ON L4B 4R6.

Volume 57, Number 9 (Printed in USA).

1 GHz never looked this good.

Oscilloscopes Redefined

1 GHz starting at $9,950**

Agilent InfiniiVision X-Series oscilloscopes give you greater signal visibility thanks to patented MegaZoom IV technology. Waveform update rates more than 400x faster than the competition reveal elusive details and even infrequent events. Now get the power of InfiniiVision X-Series up to 1 GHz at amazingly affordable prices.

InfiniiVision Oscilloscopes	2000 X-Series	3000 X-Series*
Bandwidth* (MHz)	70, 100, 200	100, 200, 350, 500, 1 GHz
Update Rate (waveforms/sec)	>50,000	>1,000,000
Price** (starting at)	$1,230	$2,810

*Existing 3000 X-Series oscilloscopes can now also be upgraded to 1 GHz
**Prices are in USD and are subject to change

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

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

FREE WaveGen function/arbitrary waveform generator and digital voltmeter
www.testequity.com/Agilent_Scope

Anticipate ___*Accelerate* ___*Achieve*

Agilent Technologies

EDN online

JOIN THE CONVERSATION

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

DANIEL VASCONCELLOS

In response to "Smack attack," a Tales from the Cube column by Specialty Concepts' Terry Staler, http://bit.ly/JMNngU, Adam Wheeler comments:

"After 30 years in service/field engineering business, I've added the phrase 'percussive maintenance' to my professional vocabulary. This [technique] has saved many a trip out of town by simply talking the end user through the process of banging on the equipment. If it starts working, I tell them to leave it on; I'll overnight them another. I've saved our company thousands of dollars by using this 'tool.'"

In response to "Use a transistor as a heater," a Design Idea by REC Johnson, B Lora Narayana, and Devender Sundi of the Center for Cellular and Molecular Biology, http://bit.ly/IEfp1k, Alan Stummer comments:

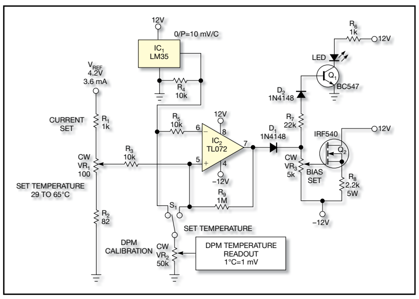

"We used this trick in telecom to control the temperature of fiber, to tune its wavelength. A small chunk of aluminum was machined for the fiber, the transistor, and a thermistor. A small dedicated micro[controller] running a PID algorithm adjusted the pulse width of the transistor at a fixed current. Temperatures were controlled to within a fraction of a degree of the target temperature."

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

WHY MR16 LIGHTS ARE ATTRACTIVE LED-LIGHTING PRODUCTS

Although the tiny lights have their own technical-design challenges, there are three good reasons that Lighting Science Group, the largest producer of LED lights in North America, has jumped into the MR16 market.

http://bit.ly/HWzIV0

PHOTOS: DISCOVERY TAKES VICTORY LAP BEFORE LANDING AT SMITHSONIAN

After logging 39 flights and nearly a year orbiting the Earth, Space Shuttle Discovery ended its epic career with an around-the-Beltway tour of the nation's capital before landing at its final destination near Dulles International Airport in April.

http://bit.ly/Ilo4Fu

ENGINEERING COMMUNITY

Opportunities to get involved and show your smarts

Austin Mini Maker Faire

Austin Mini Maker Faire

Everything is bigger in Texas … even "mini" things. Be sure to check out the Texas-sized do-it-yourself fun and learning at the Austin Mini Maker Faire, a one-day, family-friendly festival of invention, creativity, and resourcefulness and a celebration of the grass-roots maker movement, May 12, 10 am to 6 pm, Pine Street Station, East Fifth and Waller streets, Austin, TX.

http://austinmakerfaire.com

BY BILL SCHWEBER, CONTRIBUTING TECHNICAL EDITOR

Do new technologies ease small-scale product innovation?

The two extremes of product volume have their own attributes. If you're designing and building a high-volume product, you can justify manufacturing tooling and test fixtures—whether at your facility or at a contract assembly house. At the other end of the volume spectrum, if you are producing only a few units per month, or doing semi-custom or full-custom work, you usually must perform many aspects of the manufacture using manual techniques.

But what about those projects with low to moderate volume of approximately 10 to 50 units per month? They are often caught in the small-scale, in-between zone: too few to afford serious tooling and fixturing but too many to build by hand.

I thought about this question when I saw a leading-edge oscilloscope from Agilent, which is certainly not going to have volume runs comparable to those for a smartphone (**Reference 1**). The scope's analog front-end circuitry and assembly include ICs mounted on a custom-milled waveguide subassembly (**photo**). Although this subassembly is clearly sophisticated, basic milling is generally no longer as costly or difficult as it once was. The combination of PC-based CAD (computer-aided design), FEA (finite-element-analysis), CAM (computer-aided-manufacturing), and CNC (computer-numerical-controlled) machining centers makes it easier to design, set up, and make such components.

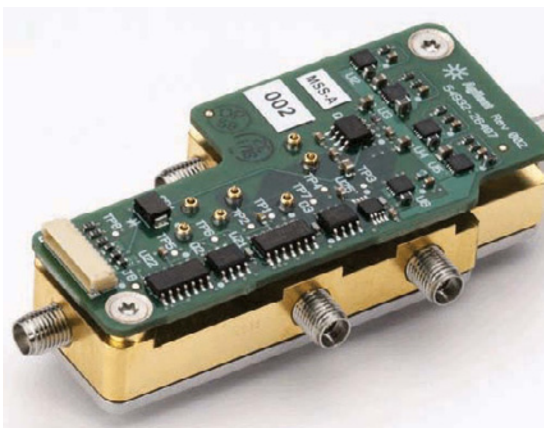

Basic milling is generally no longer as costly or difficult as it once was.

Machining is not the only technique that has changed radically. Using a variety of high-end plastics, sintered powdered metal, and other materials, along with CAD/CAM software, rapid prototyping lets you build both prototypes and modest production runs, with virtually no tooling cost or lag time. CFD (computational-fluid-dynamics) tools let you get a sense of your design's thermal situation to determine whether it needs a fan, a heat sink, or lower-power components.

For the PCB, you can use modeling tools and software to prepare the layout and then get a batch of boards made outside in 24 to 48 hours. Alternatively, you can make them on demand using a machine, such as one from LPKF. Again, these approaches require little or no tooling or setup time. There's still the problem of loading the boards before soldering; we need a way to accomplish that task other than with manual methods or with a special setup. As for parts, there are distributors that can ship prototypes and modest volumes of components to you off the shelf in 24 to 48 hours.

If you step back and look at the tools, tooling, components, and processes it takes to develop and produce a lower-volume product, you'll see that these developments have changed things for the better. You can then market your product directly through the Web, avoiding the need for a more formal channel of distribution until you get some customers and traction.

Does this development mean that we'll see more of those clever, small-scale innovative electronic and electromechanical products coming from "garage" engineers? Will the upfront cost, time, and effort barriers of trying out an idea in the market decrease, or will the inherent marketing challenges and countless regulatory aspects counter and overwhelm the benefits of these advances?

What do you think? Will the lone innovator or lower-volume project team find things better than, about the same as, or worse than they could be? EDN

REFERENCE

1 Schweber, Bill, "Real-time scope punches to 63-GHz true analog bandwidth, on two channels," *EE Times*, April 11, 2012, http://bit.ly/HLkIr8.

Bill Schweber is the editor of Planet Analog *and* Power Management Designline, *both on the Web site of* EE Times, *a sister publication of* EDN. *Contact him at bill.schweber@ubm.com, or comment directly on this column at www.edn.com/120510ed.com.*

EDN

SENIOR VICE PRESIDENT, UBM ELECTRONICS
David Blaza
1-415-947-6929;
david.blaza@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
Consumer
1-516-474-9568;
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

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

SENIOR TECHNICAL EDITOR
Steve Taranovich
Analog, Systems Design
1-631-413-1834;
steve.taranovich@ubm.com

DESIGN IDEAS CONTRIBUTING EDITOR
Glen Chenier
edndesignideas@ubm.com

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

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

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
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
Huang Hua
Operations General Manager
huang.hua@ednasia.com
Grace Wu
Associate Publisher
grace.wu@ednasia.com
Vivek Nanda, Executive Editor
vnanda@globalsources.com

EDN CHINA
Huang Hua
Operations General Manager
huang.hua@ednchina.com
Grace Wu
Associate Publisher
grace.wu@ednasia.com
Jeff Lu, Executive Editor
jeff.lu@ednchina.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, UBM Canon (Publishing), UBM Channel, and UBM Design Central
Kathy Astromoff, Chief Executive Officer, UBM Electronics
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

EDITED BY FRAN GRANVILLE

pulse

INNOVATIONS & INNOVATORS

Demultiplexer extracts the clock from data

The results of jitter and bit-error-rate measurements and eye-diagram evaluations are only as good as the equipment you use, whether an oscilloscope or a bit-error-rate tester, and require you to minimize test-equipment-caused jitter errors. Thus, you need a stable, low-noise clock. Serial data streams embed the clock in the data, making it necessary to extract the clock for measurements.

Another option is to use Agilent's new N4877A electrical clock-data-recovery and demultiplexer, which comes in two flavors. The $50,000 Option 232 allows for clock recovery from 50 Mbps to 32 Gbps, providing the speed necessary for analyzing communications signals, such as 100-Gbps Ethernet, InfiniBand extended data rate, and 32× Fibre Channel. The $32,000 Option 216 reaches 16.5 Gbps, which is enough for PCI Express and other lower-bit-rate serial links.

You can use either instrument to trigger real-time or sampling oscilloscopes or to synchronize a bit-error-rate tester to the incoming data stream. The recovered clock's residual jitter is as low as 100 fsec. The instrument also produces a derived, lower-speed clock and has a tunable loop bandwidth and adjustable peaking, necessary for meeting the requirements of several serial-communications standards.

The N4877A can work with Agilent's N1075A optical coupler/converter, which retains the mode structure of the optical signal in the fiber. Both instruments are available separately or as a package. The N4877A-016 and N4877A-032 sell for $46,000 and $64,000, respectively; the N4877A M14 and S32 each sell for $18,000.

—by Martin Rowe

▷**Agilent Technologies**, www.agilent.com/find/N4877A.

TALKBACK

"Stock ac fan, full of cheap toner; hand vendor coveralls. … Remove vermin bones and jewelry from fuser."

—Engineer Steve Nordquist, in *EDN*'s Talkback section, at http://bit.ly/HARFfp. Add your comments.

Agilent's new N4877A electrical clock-data-recovery and demultiplexer comes in two versions—one providing clock recovery from 50 Mbps to 32 Gbps and one reaching 16.5 Gbps.

Chip sets support 10.3-Gbps Thunderbolt-interface standard

If you're unfamiliar with Thunderbolt technology, you probably should get started. Intel developed the technology, which it markets with Apple. It enables fast data transfers between a PC and a peripheral or a display device, combining PCIe and DisplayPort technology for simultaneous bidirectional transfers at 10.3 Gbps over a single cable.

Such speed comes at a price, however. In addition to the Intel-supplied host processor, graphics-processing unit, and platform-controller hub, the technology requires some dedicated interface management, support components, and carefully built cable using 40-gauge AWG to make it happen. It also demands a significant amount of power and management.

> Thunderbolt combines PCIe and DisplayPort technology, enabling fast data transfers between PCs and peripherals.

A sextet of ICs from Texas Instruments strives to achieve these goals. The family includes the TPS22980 power-load switch, which resides on both the host and the device side, facilitating the delivery and receipt of power to both the active cable and the connected device. You'll also find the LM3017 boost-and-battery disconnect and HD3SS001 FET switch, which work in tandem on the host to connect to the cable. The TPS22985 power-load switch, DS100TB211 signal-conditioning retimer with clock and data recovery, and LMZ10501 Simple Switcher nanomodule reside at the cable. Prices for the components range from 65 cents to $3.95 (1000).

The 3.3 to 18V TPS22980 power-load switch saves board space and simplifies system design by providing a dual-voltage switch with an adjustable current limit that prevents damage to the connected device. It includes high-voltage discharge before low-voltage connection to protect components from overvoltage exposure. The device comes in a 4×4-mm QFN package and sells for $1.10 (1000). The TPS22985 power-selection device manages all power delivery through a Thunderbolt-technology cable, simplifying cable designs. Targeting use with the DS100TB211 for cable implementation, the device minimizes the need for external circuitry and saves space in the active-cable form factor. Measuring 1.6×1.6 mm, the device sells for 65 cents.

The dual-lane DS100TB211 signal-conditioning retimer with clock and data recovery has adaptive four-stage equalization, which enables the use of thin, inexpensive, 40-gauge AWG cable. It also integrates clock synthesis and power filtering, reducing bill-of-materials cost and the number of necessary PCB layers. On-chip cable diagnostics lower assembly time and cost. In a 5×5-mm QFN package, the device sells for $3.95. The HD3SS0001 data-source-selection switch integrates multiple discrete components, saving board space. Its pinout optimizes routing from signal sources to the connector. The device comes in 3.5×5.5-mm TQFN package and sells for 90 cents.

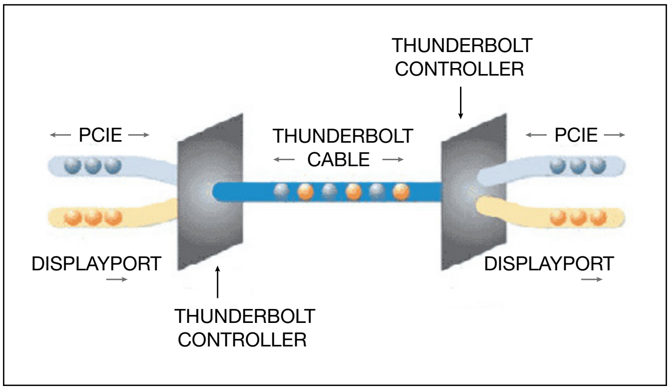

The Thunderbolt technology looks deceptively simple.

The LM3017 boost controller features true shutdown, which protects the battery from excessive current draw in a short-circuit condition. It meets both broadband-output-noise and ripple requirements with less than 50 mV p-p. The device features output power of as much as 30W for driving two Thunderbolt ports at 15W each and a ±1% reference voltage, which provides accurate output to the power load switch. Available in a 2.4×2.7-mm QFN package, the device sells for 97 cents.

The 1A LMZ10501 Simple Switcher nanomodule integrates an inductor, saving board space and easing design. It features efficiency as high as 97%, reducing system-heat generation, and has an adjustable output voltage, providing design flexibility. The 3×2.5-mm SE08A-packaged device sells for $1.80.

—by Bill Schweber

▷**Texas Instruments**, www.ti.com/thunderbolt-pr.

DILBERT By Scott Adams

Digital-LED controller tackles dimmer-compatibility challenges

Cirrus Logic's initial entry into the LED-lighting market indicates that the company believes the key to success in this highly competitive sector lies in achieving nearly 100% compatibility with the installed base of old-fashioned TRIAC dimming switches. The company has tested its new CS161X LED-controller family, which uses Cirrus' new digital TruDim technology, with more than 200 dimmer switches worldwide and claims that the controller achieves 97% compatibility. According to Cirrus, the closest competitive LED controller performs at 71%, and the average among today's LED controllers is 50% compatibility.

Cirrus Logic claims that its CS161X LED-controller family, using the company's new digital TruDim technology, is compatible with 97% of 200 TRIAC-based dimmer switches worldwide.

The CS161X's digital intelligence allows the controller to identify the type of dimmer in use and adapt its dimmer-compatibility algorithm to provide smooth dimming that mimics the response of an incandescent bulb, a response that analog controllers cannot achieve. Among the most commonly used LED controllers today, only a family from iWatt is digital, but the device is a state-machine implementation. Cirrus' design is a fully programmable digital controller.

The US Department of Energy and its Energy Star program clearly deem compatibility important: Dimming compatibility was a requirement for the Department's L-Prize LED-bulb competition. It didn't set the bar too high, however. Contest rules stated that the bulbs "must be compatible with at least three widely available residential dimmers."

In addition, dimming switches are most prevalent in the United States, especially California, and relatively uncommon in Asia. Interestingly, the newer, so-called digital dimmers are the most difficult devices for which manufacturers can ensure compatibility. The Maestro digital dimmer from Lutron, for example, can routinely make LED bulbs fail, whereas the old-style dimmers pose less of a challenge. "Digital" in this case refers to the ability of the switches to set and remember dimming levels; they still use TRIACs to chop the ac-wave signal.

The CS161X integrates a continuous-conduction-mode/critical-conduction-mode boost converter, providing power-factor correction and dimmer compatibility with primary-side control, constant output current, and a quasiresonant flyback or buck output stage. It adapts to use with both 100 to 120 and 220 to 240V-ac line voltages. The CS161X is currently in volume production. It is available in a 16-pin SOIC package and sells for 81 cents (100,000).

—by Margery Conner

▷**Cirrus Logic Inc**, www.cirrus.com.

Midpower LED aims at office lighting

Philips Lumileds' first white, midpower LED, the 5630, targets office lighting and other distributed-LED-lighting applications. Midpower LEDs operate at approximately 100 to 150 mA, whereas high-brightness LEDs operate at 350 mA to 1A.

The device comes in a color-correlated temperature range of 2700 to 6500K with an efficacy of 110 lumens/W at 100 mA. It has minimum and typical color-rendering indexes of 80 and 82, respectively; a guaranteed minimum of 26 lumens at 100 mA; and a guaranteed minimum of 30 lumens for higher color-correlated temperatures.

Midrange LEDs fit well in office lighting because office light is usually omnidirectional, such as you'd expect from a fluorescent tube. High-brightness LEDs, on the other hand, yield an intense point of light, which can be distracting; it's just not what people expect from their office lighting. Much of what we expect from lighting goes back to whatever we became used to with the previous generation of technology.


Philips Lumileds' first white, midpower LED, the 5630, has an emitter in the center. The slightly darker yellow is the phosphor, which covers the emitter. Two thin bond wires go to the anode and cathode. The package measures 5.6×3×0.9 mm.

When our ancestors left torchlit caves, they probably briefly complained about the lack of soot. In Japan, which has perhaps embraced LED lighting more quickly than any other culture, the point sources of light from high-brightness LEDs are popular for office lighting, perhaps because it makes a clear statement that the lighting is cutting-edge LED technology. One culture's bug is another's feature.

—by Margery Conner

▷**Philips Lumileds**, www.philipslumileds.com.

Resistor-programmed switching-regulator IC goes rail to rail down to 0V

Designers must often choose whether to use a low-dropout regulator or a switcher and choose whether to use one large regulator or several small ones. With the LTC3600 synchronous buck regulator from Linear Technology, you have another option: a dc/dc regulator that you can easily parallel with other LTC3600s for greater current-output capability and that you can program down to 0V.

This product gives you the choice of locating ganged regulators close together or spreading them—and their inevitable dissipation—around your PCB. The option also allows you to simplify your bill-of-materials costs by using one regulator in many places, both as a solo device and as part of higher-current groups.

You use one resistor to program the IC, which delivers 0 to 15V at currents as high as 1.5A. It includes an internal 50-μA reference current to establish the output voltage, which you can adjust dynamically. It accepts inputs spanning 4 to 15V, making it a good fit for dual-cell lithium-ion designs or fixed-rail 5 and 12V intermediate-bus converters. You can set the output voltage between 0 and 0.5V lower than the input voltage, effectively making the device a rail-to-rail regulator; the 0V setting is useful for powering down individual rails in a system.

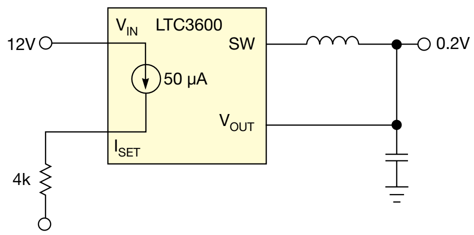

The LTC3600 synchronous buck regulator delivers a resistor-programmable 0 to 15V and 1.5A parallel operation for higher output currents.

Accuracy is ±1% over temperature and is independent of the load, thanks to the laser-trimmed reference. The LTC3600 offers output-voltage tracking or soft-start operation for multirail sequencing through a programmable package pin. The use of internal N-channel power MOSFETs eliminates the need for using separate discrete devices.

Quiescent-supply current is lower than 1 μA, and the regulator provides output regulation and fast transient response independently of the output voltage. Users can set the operational switching frequency at 200 kHz to 4 MHz, which lets them select small, low-cost passive components. Efficiency is approximately 96%.

The LTC3600 is available in a 12-lead, 3×3-mm DFN package and a thermally enhanced MSOP; it sells for $2.80 (1000). Wide-range industrial-grade versions are available for –40 to +125°C operation and sell for $3.22.—**by Bill Schweber**

▷**Linear Technology**, www.linear.com/product/LTC3600.

Intel rolls out first processor for data centers

Intel Corp has designed its first processor built from the ground up for the "green" data centers of the future, claiming a 70% increase in performance for the same energy consumption. The new E5-2600 also features a high-speed, bidirectional ring encircling as many as eight cores per socket and connecting as much as 20 Mbytes of cache, quad-DDR3-memory controllers, and 40 lanes of PCI Express 3 for I/O. The E5 features twin 32-byte-wide ultra-high-speed rings going in opposite directions to encircle eight Sandy Bridge cores and connect them to the cache. The E5 family is Intel's first server-processor family with integrated I/O, rather than using a separate chip, thereby reducing latency by 30% and doubling the bandwidth with PCIe 3. The E5 is also the first Intel server processor to support LAN-on-motherboard by virtue of its integrated 10-GbE LAN.

The processor touts an idle power of 10 to 20% usage. A sophisticated power-management agent puts separate power limits on the device, its cores, memory, and I/O and then intelligently manages them for optimal performance, energy efficiency, or other data-center goals. Using dynamic switching, depending on load conditions and turbo requests, the E5 automatically switches between performance and low-power modes plus a new balance mode that compensates for turbo requests by adjusting the voltage and frequency of other cores.

For instance, if data-center managers decide to clamp power at a certain overall level, then the balance mode adjusts some cores down in voltage and frequency to compensate for the heavy load on a turbo-mode core. The new Turbo 2.0 mode is also smarter on the E5, employing better thermal-management algorithms that keep track of how long a core has been held idle, building up turbo "credits" for use when overclocking is invoked.

Besides voltage and frequency scaling for each core, the new power-management agent also manages energy efficiency in I/O by dynamically reducing its width in response to workload and thermal-management goals. Core power scales from 50 to 95W, which likewise scales memory latency from 118 to 64 nsec; a unicore technique scales cache and ring frequency to match.

The E5's running-average-power-limit architecture adjusts 23 parameters. With Intel's Node- and Datacenter-Manager software, Intel estimates, users can install as many as 40% more servers per rack using E5 processors.—by R Colin Johnson

▶**Intel Corp**, www.intel.com.

05.10.12

Multicore processors tackle human interface

Heterogeneous multicore processors best address the diverse demands of real-time process control, which must nevertheless offer a touch-enabled user interface—from smart appliances to point-of-sale terminals to medical monitors—according to Freescale Semiconductor Inc, which recently introduced its Vybrid microcontroller family.

"Smartphone users now expect their appliances to have smart touch-enabled human-machine interfaces, but microcontrollers that are good at running high-level operating systems, such as Android, have a hard time delivering real-time deterministic control, too," says John Weil, global business manager for industrial microcontrollers at Freescale. "Our new Vybrid family solves that problem with a heterogeneous ARM architecture—using a Cortex-A5 core for the human interface and a Cortex-M5 for real-time control."

The new Vybrid processor is an industry first in its ability to simultaneously run both a high-level Linux/Android OS and a real-time RTOS/MQX OS on its dual heterogeneous ARM Cortex-A5 and -M4 cores, respectively. Separating the two functions but allowing development from one tool platform, Freescale claims, streamlines applications in cost, time to market, and energy consumption.

The Cortex-A5 is ARM's answer to criticisms it received about its higher-performance A9, offering 14.4 Dhrystone MIPS/mW compared with 8 DMIPS/mW for the A9, an 80% savings in energy consumption for only a 56% reduction in performance—that is, 1.6 DMIPS/MHz for the A5 compared with 2.5 DMIPS/MHz for the A9. Plus, the M4 core handles all of the real-time interrupts and other mission-critical functions, offloading these tasks from the A5, making it a better fit for its targeted consumer, automotive, industrial, and medical applications.

Freescale's new Vybrid microcontroller family can run simultaneously a high-level Linux/Android OS and a real-time RTOS/MQX on dual heterogeneous ARM Cortex-A5 and -M4 cores, respectively.

Freescale has several design wins, including two megacorporations, for its new Vybrid processor, which will become available for sampling during the next quarter and should begin volume production in the third quarter. Current applications from OEMs late in 2012 or early in 2013 include touchscreen-enabled point-of-sale terminals, medical monitors and dispensers, and white goods. A typical division of labor between the A5 and the M4 cores includes running a touch-enabled Android user interface as the human-machine interface on the A5 and the security-and-safety algorithms for guaranteed drug delivery for medical applications, as well as the servo-control and swipe-and-read algorithms for industrial and point-of-sale applications, respectively, on the M4 cores.

Hardware support for a semaphore-based message-passing system handles communications between the two cores. All of the usual peripherals for both application processors and real-time-control processors are also on-chip. These peripherals include flex timers; watchdogs; clocks; low- and high-frequency oscillators; interrupt routers; ADCs; DACs; phase-locked loops; debug-and-trace capabilities; direct-memory access; power management; a cryptography module; tamper detection; UART, CAN, SPI, Ethernet, and USB interfaces; and boot-ROM, SRAM, and both flash- and DDR-memory controllers.

—by R Colin Johnson

▷**Freescale Semiconductor**, www.freescale.com.

Dual power booster adds serious muscle to op-amp outputs

In the area of milliwatt dissipation, microwatt drain, and single-volt drive, doing almost anything useful in the real world—that is, driving substantial loads or moving physical things—often requires voltages in the two- or three-digit range and currents measured in amps. The PB63 dual-channel booster amplifier from Cirrus Logic's Apex Precision Power group targets those needs, amplifying a small signal from common, general-purpose op amps for operation at voltages as high as ±75V, or 150V, and ±2A. With two channels in one package, the device suits diverse applications such as power supplies, piezoelectric transducers, electric- and magnetic-field excitation, scanning electron microscopes, motion and motor control, flat-panel-display pattern generators, automated-test equipment, wide-format ink-jet printers, programmable power supplies, medical and surgical instruments, and industrial-audio installations.

The PB63 booster amplifier comes in a 12-lead power SIP.

The device features a unipolar output-voltage swing of 40 to 150V or a split bipolar swing, a slew rate as fast as 1000V/μsec, output current as high as 2A, a 1-MHz power bandwidth, and a quiescent current of 20 mA. The PB63 comes in a 12-pin, 31×30-mm power SIP offering a footprint of 1.5 in.^{2} (3.8 cm^2). The price is slightly less than $129 (1000). Additional support includes an evaluation board with a heat sink and associated hardware and several application notes.

—by Bill Schweber

▷**Cirrus Logic Inc**, www.cirrus.com.

3.5A, 6V STEP-DOWN CONVERTER NEEDS ONLY 40-μA QUIESCENT CURRENT

Targeting handheld and portable devices requiring only one single-cell lithium-ion battery, the MP2130 synchronous step-down converter from Monolithic Power Systems features a quiescent current of only 40 μA. The 3.5A, 6V, 1.2-MHz device fits into a 2×2-mm QFN package. It achieves a 96% efficiency at a load current of 1A and an output voltage of 3.3V, and it still maintains an efficiency of 85% at a load current of 0.01A. You can regulate the output voltage as low as 0.6V. The constant-on-time-control approach provides fast transient response, high light-load efficiency, and fast loop stabilization. Fault condition protection includes cycle-by-cycle current limiting and thermal shutdown and hiccup-mode short circuit. The product also comes with a power-good indicator for output-voltage monitoring and easier power sequencing. It sells for $2.65 (1000).—by Margery Conner

▶Monolithic Power Systems, www.monolithicpower.com.

The 3.5A, 6V, 1.2-MHz MP2130 converter fits into a 2×2-mm QFN package and requires a quiescent current of only 40 μA.

BY HOWARD JOHNSON, PhD

10 measurements defining signal integrity

The following 10 measurements define signal integrity. Master them, and you will become a guru of the art. Work on one measurement every quarter or every year until you fully grasp the relationship among circuit theory, simulation, and measurement. All of these problems harbor subtle difficulties, as well as sparkling gems of insight. If you already know the answers, I hope this outline inspires you to teach others.

Step response: Drive a capacitor with a 5Ω step source having a rise time of 1 nsec. Look at the step-response waveform in the time domain. Assume a basic resistance/inductance/capacitance-series electrical-circuit model for the capacitor and extract the circuit parameters. Plot the curve of impedance magnitude versus frequency that your model predicts. Using a sine-wave source, or S-parameter test set, measure the actual impedance versus frequency and see whether it matches your expectations (**Reference 1**).

Characteristic impedance: Separately measure the inductance and the capacitance of a 1-foot section of coaxial cable at 10 MHz. Estimate the characteristic impedance as the square root of the ratio of inductance to capacitance. Using a time-domain reflectometer, measure the input impedance of a longer section of the same type of cable. Determine whether the two values match (**Reference 2**).

Dispersion: Terminate a 100-foot-long RG-58 coaxial cable. Using a rise time of 1 nsec, measure its signal delay. Then quadruple the cable's length and observe whether the delay quadruples. Learn about signal dispersion and the skin effect (**Reference 3**).

Reflections: String together two 10-foot sections of RG-58 cable. Terminate the endpoint. In the middle, connect a 100-pF capacitor from signal to shield. At the head end of the structure, using a 10-nsec rise time, observe the reflections that the reactive load in the middle generates. Now, imagine a 2-in. PCB trace with a 1-pF load, operating at a rise time of 100 psec. Do you expect the same behavior (**Reference 4**)?

Ringing: Inject a 1V p-p signal from a 50Ω source into a 10-foot coaxial cable with no termination. Use a rise and fall time of 100 nsec. Determine the largest signal you can create at the endpoint. Now make a T configuration, with 10 feet leading to the central T and a pair of 10-foot sections branching off from that section. Add a 39-pF load at just one of the branch endpoints. What is the largest signal you can now make, and what does this signal tell you about the unbalanced T configuration (**Reference 5**)?

Crosstalk: Select two adjacent traces on a working PCB. Inject a step waveform from a 50Ω source into the first trace, making sure to terminate its far end. On the adjacent trace, compare the measured crosstalk at the trace's two endpoints, with both terminated. Is the crosstalk the same at the two ends (**Reference 6**)?

Loading: On a working PCB, disconnect a driver from its load. Record the output waveform with no load, with 50Ω to V_{CC}, and with 50Ω to ground. Do any of the loads change the driver's switching speed and, if so, why (**Reference 7**)?

Jitter: Capture a waveform showing thousands of edges exiting a reference oscillator. Put the waveform into a math-processing spreadsheet, such as Mathcad, Matlab, or Mathematica, and make a list of the exact time of arrival of each rising edge, interpolating between samples. Make one vector showing the delay, from edge to edge, of successive pulses. Make another vector showing the delay from each edge to the edge 100 cycles later. Plot histograms of the two vectors and decide which has the largest standard deviation (**references 8** and **9**).

Simultaneous-switching-output noise: On a large BGA device, find an output that is programmed to stay at logic zero throughout this experiment and route that output to your scope. Activate the rest of the device and observe the crosstalk at the signal under test. What can you do in the BGA device to maximize the crosstalk? Do all of the pins exhibit the same level of crosstalk (**Reference 10**)?

Power-supply noise: On a working PCB, strip off one bypass capacitor and use its mounting pads to connect a small coaxial probe. Observe the power-system noise at that location. Next, start disconnecting bypass capacitors. What happens to the power-system noise as you remove the capacitors, and what aspects of the noise change? Are those changes the same at other locations (**Reference 11**)? EDN

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

See the online version of this column at www.edn.com/120510hj for a list of the references cited and links to the original content.

High Speed / Low Cost
Development, Test & Burn-In Sockets
Compression, SMT, &Thru-Hole PCB Mounting
Lead pitches as low as 0.4mm, up to 125,000 insertions.
Quick On/Off Lid
Easy Screw - Lid removed to show detail.
Easy Knob configuration
Multi-Cavity Sockets
Significantly reduce your socket & labor costs with these multiple IC test and burn-in solutions.
Emulation Technology, Inc
Custom Test Fixtures
Fully automated or manually operated solutions to test any lead pitch & IC package
Quick Turn
Low NRE
ET®
EMULATION TECHNOLOGY, INC
1-800-232-7837
www.emulation.com

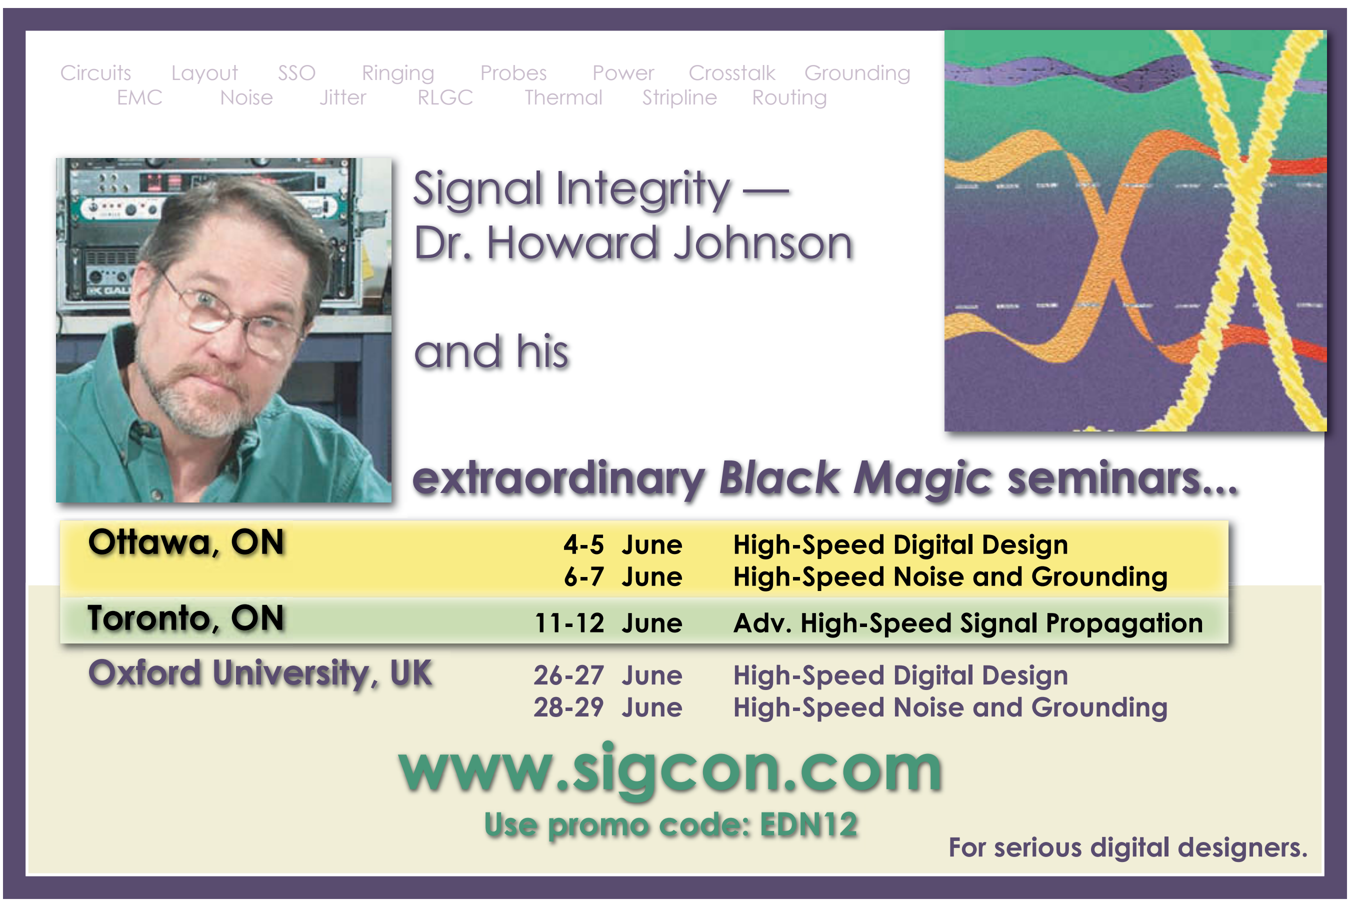

Circuits Layout SSO Ringing Probes Power Crosstalk Grounding
EMC Noise Jitter RLGC Thermal Stripline Routing
Signal Integrity —
Dr. Howard Johnson
and his
extraordinary *Black Magic* seminars...
Ottawa, ON 4-5 June High-Speed Digital Design
6-7 June High-Speed Noise and Grounding
Toronto, ON 11-12 June Adv. High-Speed Signal Propagation
Oxford University, UK 26-27 June High-Speed Digital Design
28-29 June High-Speed Noise and Grounding
www.sigcon.com
Use promo code: EDN12
For serious digital designers.

BY PALLAB CHATTERJEE

Drive interfaces require trade-offs

As the amount of data in networks continues to grow at an exponential rate, higher-performance storage at both the client side and the enterprise server is becoming a necessity. The client side must balance the performance with power and space constraints. The server side must balance capacity and throughput. These capacity issues drive a power constraint on a multi-unit rather than a single-unit basis.

Disk drives have moved from power-inefficient and high-pin-count parallel interfaces to low-power, high-speed serial connections with low pin counts. The older, high-power interfaces, including PATA, SCSI, and PCI, had multiple-data-pin, wide-cable interfaces. The new interfaces, including SATA, SAS, and PCIe, all have high data rates with only a few active pins, thanks to advancements in the use of serializers/deserializers on the interface (**Figure 1**).

Space and power/thermal issues for capacity and storage limit the client side. As a result, the new serial interfaces feature rotating media as the dominant high-capacity technology. For newer, power-conscious and application-optimized systems, solid-state drives are making inroads, but they have smaller capacities due to their cost. The form factor for this client-side storage is a 2.5-in. or smaller drive. Bandwidths of approximately 3 Gbps for SATA II are now moving to 6 Gbps for SATA III for the primary interface, including the 5-Gbps USB 3 interface, which is generally SATA II or III devices with a protocol converter and a buffer.

With the arrival of the Internet of Things and the rise in embedded computing, the storage needs of new embedded clients have also changed. The embedded-system world has several storage formats to use. The EMMC (embedded multimedia card) is the most common format for cell phones, tablets, and other microcontroller-directed processing environments featuring directly mapped storage and I/Os. The developers of this interface designed it to employ flash memory and wear-leveling with an internal memory-controller interface.

According to Greenliant, a premier supplier of embedded solid-state drives and NAND-flash controllers, extended-temperature EMMC products, in addition to traditionally strong SATA and PATA storage solutions, have been gathering interest from industrial-, automotive-, and networking-system designers. The company is also seeing requests for and investigation of PCIe as a high-throughput interface for these applications, but the use of PCIe has not yet materialized for the mainstream embedded-system market. SATA, PATA, USB, and other legacy interfaces are still dominant in the embedded systems requiring high reliability and new, smaller form factors, such as Greenliant's NANDrive, which offers an integrated controller, flash memory and file system, and ECC (error-correcting code), all in a 14×24×1.95-mm size.

Just as in the industrial embedded-system market, enterprise storage focuses on an extended range of operating temperature as well as ECC, wear-leveling, data integrity, and a high BER (bit-error rate). These features fill the needs for speed, capacity, and power for consumer requirements. The enterprise market splits between using long sequential read/writes and short random read/writes. Unlike consumer storage, these enterprise storage-area-network and direct-attached-storage products may scale into the petabyte, or 10^{15}-byte, and exabyte, or 10^{18}-byte, level. These levels require fast, large data, but the data must also be correct. Enterprise storage must push these known interfaces and protocols beyond the historic one in 10^{16} BER. Some of the early high-speed controllers exhibited BER reduction to one in 10^{10}. This rate required a 6-Gbps interface, and errors occurred every second.

The enterprise market is also trying to deliver competitive densities with consumer solid-state drives by implementing multilevel-cell and triple-level-cell flash architectures. Density matters, but getting the correct data is the key. EDN

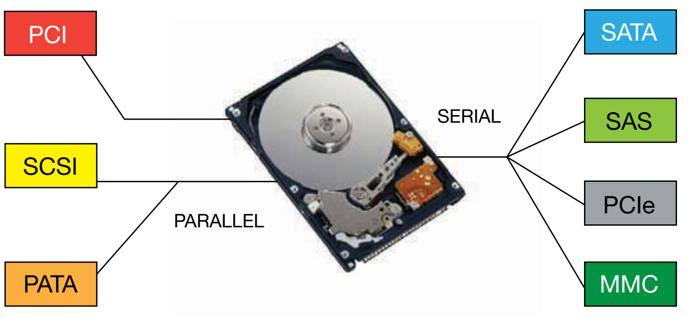

Figure 1 The older, high-power interfaces, including PATA, SCSI, and PCI, had multiple-data-pin, wide-cable interfaces. The new interfaces, including SATA, SAS, and PCIe, all have high data rates with only a few active pins.

Pallab Chatterjee has been an independent design consultant since 1985.

GENERATING SPATIAL AUDIO FROM PORTABLE PRODUCTS

BY KEN BOYCE • TEXAS INSTRUMENTS

THE EVER-SHRINKING MECHANICAL COMPONENT OF PORTABLE DEVICES POSES MAJOR CHALLENGES FOR STEREO AUDIO PLAYBACK.

The sound stage generated by a stereophonic audio system is typically restricted by the physical location of the speakers, while the sound events perceived by the listener are limited within the span of the two speakers. In small stereo speaker systems, such as those in portable devices, the perceived stereophonic sound stage becomes very limited, almost monophonic. To overcome this limitation, you can use spatial audio sound-generation techniques to expand the stereo sound stage, achieve better crosstalk cancellation, and enhance certain spatial cues.

IMAGE: SHUTTERSTOCK

SPATIAL AUDIO

Naturally occurring sound in the space around you is inherently spatial. Sound sources reside at a point in a small region of the total space, although some sources, such as earthquakes and landslides, emanate from a broader area. The sound scatters off various objects in the environment. You hear the direct and scattered sound binaurally with your ears and then, after some human auditory processing, finally recognize the sound. You make decisions about the processed sounds, characterizing them with labels such as direction, location, loudness, ambience, quality, near, far, tone, fat, and thin. In the case of two or more sound sources, you also determine the characteristic mix of each sound source relative to the total sound your ears are receiving.

Spatial audio is a term applied to the reproduction of sound by electronic or mechanical methods that attempts to artificially re-create the real-world listening experience of sounds. Alternatively, it attempts to artificially alter reproduced sounds to create a perceived spatial environment that may not have originally existed.

The principle of spatial audio reproduction is simple: If the reproduced sound waves arriving binaurally at your eardrums are identical to those of the real audio source at a particular position, you will perceive the reproduced sound as coming from a source at that position. It does not matter whether the sound source originated at some other position. The sound data arriving at your ears is what the brain processes to ultimately characterize all aspects of the sound.

AT A GLANCE

- The number of speakers and the spacing between them limit the sound stage of portable stereo sound systems.
- Spatial audio attempts to artificially re-create the real-world listening experience of sounds or to create a perceived spatial environment that may not have originally existed.
- Using HRTF (head-related-transfer-function) information, you can synthesize binaural sounds that appear to originate from any point in space around a listener.
- Acoustic beam-forming directs audio waves in a direction other than a loudspeaker's typical radiation pattern.
- To deliver spatial stereo sound over a speaker array requires an effective crosstalk-cancellation algorithm.
- Many methods can deliver spatial audio, but most are unsuitable for any reasonable array size on a small portable device.

HRTFs

An HRTF (head-related transfer function) of the ear describes how you receive and process sound from a point in space. The frequency-response HRTF describes how the human body scatters acoustic signals and how the pinna, or external ear, and the ear canal filter these signals before the sounds reach the eardrum. The circularly asymmetric external ears, or pinnae, form specially shaped antennas, which cause location- and frequency-dependent filtering of the sound reaching the eardrums, especially at higher frequencies.

HRTFs are the Fourier transforms of static measurements of the left- and right-ear impulse responses, or head-related impulse responses, of sound the ears receive from different distances and directions. The ILD (interau-

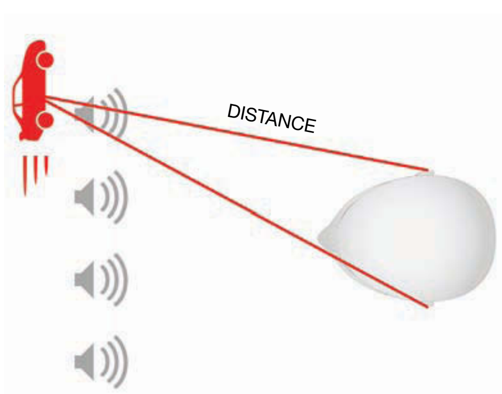

Figure 1 A head-related transfer function localizes an audio source in space by considering that sound waves enter the ears at different times and with varying intensities due to the difference in distance between ears.

ral level difference) and ITD (interaural time difference) are derived through the sounds each ear hears (**Figure 1**).

Each person's HRTF differs because of sometimes significant differences in hearing capability and physical characteristics. However, several HRTF-measurement databases use broad classifications, such as male or female and young or old, and generally find use in consumer audio applications, which require HRTFs. Such measurements can take a lot of time because humans cannot hold their heads in a fixed position for long periods, resulting in imprecise data. For this reason, some creators of HRTF databases take measurements only from dummy heads that they model on an average human head and ears to avoid head-movement errors.

Your ears can locate sound directions in three dimensions—front/back, above/below, and either side—to an angular resolution of approximately 3°, and you can also estimate distances because your brain, inner ear, and external ear use monaural cues derived from one ear and by comparing binaural, or difference, cues received at both ears. In the natural environment, individuals have learned the accuracy of their sound-location ability—their HRTF data—through trial and error and lifelong experience and have effectively compensated for their body shape and composition.

It is possible to synthesize binaural sounds that appear to originate from any point in space around the listener by applying the appropriate filters to existing audio signals and combining the sound with HRTF information, resulting in left- and right-channel sound specifically designed for each ear. Similar to left- and right-sound separation that headphone users experience, each ear hears only what it should hear. By contrast, a stereo audio signal playing through headphones appears to emanate in an area restricted to a line between the ears. This difference between ordinary stereo and spatial audio gives rise to 3-D, or virtual-surround, sound.

AUDIO CROSSTALK

Creating a spatial audio effect is more difficult when using two loudspeakers at some distance from the listener because both ears can hear the sound from any one channel, causing audio crosstalk. You can achieve crosstalk cancellation by using destructive-wave interference to cancel unwanted signals. Antiwaves, or cancellation waves, sent to the right ear cancel unwanted left-channel audio signals (**Figure 2**). The same thing happens to unwanted right-channel signals at the left ear. The result is distinct right- and left-channel sound-enhancement areas that promote an elevated sense of audio placement in 3-D space.

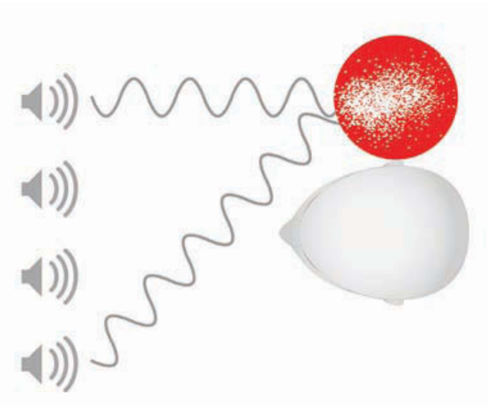

Figure 2 Crosstalk cancellation uses destructive-wave interference to cancel unwanted signals.

This additional processing of the audio in each channel is necessary to eliminate or reduce the effects of crosstalk and to take into account the possible effects on the sound by the listener's position, including the angle and distance from loudspeakers; ear sensitivity and shape due to age, sex, or ethnicity; head and torso size and mass; and the localized physical environment, including the presence or absence of reflective or absorptive materials. All of these factors determine how or whether a listener can accurately identify where a sound has originated. For these reasons, spatial-audio-creation techniques for loudspeaker systems must also include acoustic beam-forming.

ACOUSTIC BEAM-FORMING

Acoustic beam-forming in loudspeaker audio reproduction refers to the ability to direct audio waves in a direction other than the loudspeaker's typical radiation pattern. An ideal loudspeaker acts as a piston-type source set in an infinite baffle and has a radiation pattern that depends on the frequency being reproduced relative to the speaker's disk radius. The degree of acoustic sound beaming relates to the ratio of the radius of the loudspeaker's piston to the wavelength of the sound. At low frequencies, the loudspeaker's sound spreads out evenly in all directions in front of the speaker. For this reason, you can set a subwoofer at almost any location in the front of a room and hear it equally well from anywhere else in the room.

As the frequency increases, the radiation pattern focuses more in front of the speaker, increasingly becoming a narrowing cone-shaped pattern around an axis perpendicular to the loudspeaker's face. The sound-pressure level is strongest within the cone pattern and drops off rapidly outside the pattern. You can test this effect by listening to high frequencies from high-fidelity loudspeakers while moving from side to side.

Also, real loudspeakers reside in relatively small, finite boxes—not infinite baffles. The edges of the box cause diffraction of the sound waves, resulting in a more complex radiation pattern. If you place the loudspeaker in an open-air setting, you would correctly hear all of the reproduced frequencies only if you listened at a position that was on the axis perpendicular to the front face of the speaker.

If the loudspeaker is placed in a room, the listening position is less critical because the walls and furnishings also reflect the sound at various angles. Although reflections make it easier to hear from any position, the reflections arrive at different times and intensities from the original signal and result in sound that lacks clarity.

Acoustic beam-forming in the extreme case attempts to direct the sound energy emanating from the speaker to an angular position within the room environment. Acoustic beam-forming with one speaker is difficult, so typical applications use two or more speakers. Using multiple speakers also allows the use of constructive and destructive combinations of sound-wave energy to create certain directional patterns.

In an array of multiple speakers, the size and shape of the array also make certain directional patterns possible. The array is usually physically linear but may also be a 2-D array, such as curvilinear, planar, circular, or combinations of these types (**Figure 3**). As a general rule, you can more easily accomplish and discern spatial effects at higher frequencies. However, a large array and speakers with good low-frequency response allow better directional control over low frequencies.

The size and shape of the array partially determine the techniques for achieving the desired spatial effect. The other factor is the purpose for which the array is being built. For example, a picture-in-picture feature allows two viewers to watch two TV channels, but you would need another technique if you wanted to simulate a 5.1-channel surround-sound environment for several people seated in front of the TV. As a practical matter, a spatial-audio system may have to support both methods and possibly several others for consumers to have pleasurable listening experiences, regardless of the TV's program mode.

BEAM-FORMING TECHNIQUES

Mechanical beam-formers rely on the physical sizes and positions of the speakers to produce desired spatial effects, whereas electronic beam-formers rely on DSPs to process signals before providing audio signals to the speakers. You can combine these methods, depending on the application.

The creators of electronic beam-forming initially developed it for radar applications. Its application to audio first appeared in microphone arrays for speech and audio capture. The abundant applications in this area have led to many years of innovations in audio-beam-forming algorithms.

The basic idea of microphone-array beam-forming is to individually adjust the phase and amplitude of the received signal of each array element so that the combined output can achieve a maximum signal-to-noise ratio in certain directions. The concept is similar to extracting the desired signal in the frequency domain by bandpass filtering. With microphone-array beam-forming, however, this task takes place in the spatial domain, and the passband can be considered as a range of directions. Many well-documented beam-forming techniques exist, and the selection of

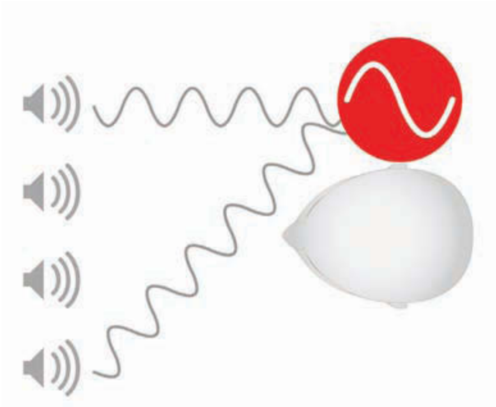

Figure 3 In speaker-array beam-forming, sound waves constructively interfere at desired locations in space.

certain techniques usually depends on the requirements and constraints of the application.

Although audio beam-forming finds wide use in capturing audio signals, its application in audio playback is relatively limited. A major reason is that stereo systems have been delivering relatively good performance, and it was unnecessary to use more than two speakers for many applications. However, the ever-shrinking mechanical component in portable devices now poses major challenges for stereo audio playback. For example, output volume-level loss due to speaker-size reduction and diminished stereo-sound image due to narrow speaker spacing are both problematic.

Flat-panel TVs are trending toward thin enclosures that severely limit loudspeaker-cone excursion, which in turn diminishes output-volume levels and audio quality. One way to overcome these limitations is to use an array of small speakers to increase the overall volume and to render a more desirable sound field using audio-beam-forming techniques. Many other methods, such as WFS (wave-field synthesis) and ambisonics, can deliver spatial audio over a speaker-array system, but they typically require dozens to hundreds of speakers and a large amount of space. Therefore, they typically find use in speaker systems in theaters and sound rooms but are unsuitable for small to midsized speaker arrays.

You can readily adapt beam-forming techniques for microphone-array applications to speaker-array applications because the playback is basically a reverse of the capturing process. However, full-bandwidth audio playback requires wider bandwidth than do typical speech applications and therefore requires special considerations in algorithm choice and array setup.

To deliver spatial stereo sound over a speaker array requires an effective crosstalk-cancellation algorithm. Additionally, speaker-array algorithms must minimize the distortion of audio playback as much as possible, including artifacts and coloration. The challenge has been to develop techniques suitable for small arrays, which yield subjectively better results than using stereo alone, and apply those techniques in a manner that uses several speakers without requiring complex algorithmic-programming skills.

SPATIAL-ARRAY AMPLIFIER

Texas Instruments has addressed this challenge with the LM48901, the first in a series of spatial-array audio amplifiers for implementing loudspeaker arrays to produce an immersive audio experience for space-constrained applications. The four-channel, Class D LM48901 audio IC implements distributed, electronic beam-forming algorithms plus sound placement using HRTF data to create beam-formed loudspeaker arrays.

You can use one LM48901 in two-, three-, or four-speaker applications. Daisy-chaining several LM48901 ICs enables the design of eight-, 12-, and 16-speaker arrays that deliver a wider and more exciting audio experience. Spatial-array-technology videos showing more information about the product and the Web-based design tool are available at www.ti.com/spatial-pr. **EDN**

AUTHOR'S BIOGRAPHY

Kenneth Boyce was formerly an audio technologist for Texas Instruments' Silicon Valley Labs. He previously served as marketing and technology director for National Semiconductor's audio-products group and as a contributing member to the MIPI SLIMbus specification. Before joining National, Boyce was director of Oak Technology's audio and communications division. He holds a bachelor's degree in electronics from West Virginia University (Morgantown, VA).

TEARDOWN REVEALS CHEVY VOLT'S ELECTRONIC SECRETS

BY RICK DEMEIS • AUTOMOTIVE DESIGNLINE

Engineers sometimes get assignments that are not only challenging learning experiences but also just plain fun. That scenario happened when John Scott-Thomas, UBM TechInsights' product-marketing manager, and Al Steier, Munro & Associates' senior associate and "design prophet," recently took apart a Chevy Volt plug-in hybrid car to see what makes it tick, whir, and hum—and how its designers put together all of the technology in the car. Over the three days it took to creatively disassemble the Volt, they learned many things about the vehicle (see **sidebar** "Anatomy of an automotive teardown" and **references** 1 and 2).

IMAGE: COMPOSITE/SHUTTERSTOCK

BEYOND BATTERY CHEMISTRY AND ELECTRIC PROPULSION AND CONTROL, THE CHEVY VOLT ENHANCED-RANGE ELECTRIC VEHICLE BUILDS IN FLEXIBILITY, RUGGEDNESS, AND DIAGNOSTICS—WITH ATTENTION TO QUALITY CONSTRUCTION.

BATTERY PACK

The 288-cell Volt lithium-ion battery pack comprises four modules in a T shape that fits below the rear seat and in the "tunnel" between the front seats. Bus bars connect the four modules, and a service-disconnect bar connects to the pack contacts (**Figure 1**). The pack physically divides into plastic-encased slices, or blades, each of which includes two cells. A cooling fin carrying five channels of coolant separates the two cells. Electrically, groups of three cells connect in parallel, and 96 of these groups are in series so that the 288 cells produce 360V with a capacity of 16 kWhr. To prolong battery life, the battery never fully charges or discharges, so it uses only the "middle" 9.4 kWhr of battery energy.

LG Chem manufactures the battery, which uses lithium-manganese-spinel chemistry, but GM has licensed battery cobalt chemistry from the US Argonne National Lab, indicating that a switch to a nickel-manganese-cobalt battery could be in the offing. The battery-cooling-fluid circuit is one of four in the Volt, each with its own controller and radiator module. The other three loops are for the internal combustion engine, the two electric motor/generator inverters, and the power-line plug-in charger's power converter.

When the battery is operating at lower than the optimum operating temperature, the fluid heats the battery to operating conditions and then cools it to avoid overtemperature. Even if the car is not operating, the control electronics activate the coolant loop to avoid overheating the battery during hot weather or overcooling in cold weather. Thus, keeping the Volt on its external charger when the car is not in use avoids draining the battery under such conditions.

AT A GLANCE

- Numerous inspectors' marks throughout the vehicle provide evidence of quality checks during the manufacture of the Volt.
- The layout of the electronics allows for easy integration of new modules and ICs.
- Ruggedized features, such as potting and taping, allow for the demanding automotive environment. Extensive software diagnostics ensure the vehicle's safety and reliability.

The battery-pack coolant loop connects using hose clamps, indicating that the car is a limited-production vehicle. Higher production volumes would allow use of brazed joints. The bolts clamping together the pack each have three inspectors' paint marks, showing that the assembly is carefully inspected to ensure quality and function for this $8000 component that is at the heart of the Volt.

CONTROL AND MONITORING

The complex Volt battery pack, as the teardown revealed, has equally sophisticated control and monitoring, which are typical of the entire car. Scott-Thomas observes that 40% of the value of the vehicle is in its electronics, typified by the nearly 100 onboard microcontrollers. Nearly 10 million lines of code control this electronic suite—more code than it takes to control the Boeing 787 Dreamliner, at 8 million lines.

As for the battery pack itself, Scott-Thomas notes that long battery life is a key objective. Toward this end, the manufacturer regulates pack temperature to within 2°F and balances cell charge between cells so that each ages at the same rate. The control software also factors in differences in manufacturing and other variables in aging.

For example, the controllers monitor voltage on each cell during charging. To ensure the same maximum charge on each, if one cell reaches capacity early, a resistive shunt across the cell connects to prevent it from being overcharged while the other cells come up to full charge. "The level of control and software is hard to appreciate," says Scott-Thomas. The car's controllers monitor the battery pack's voltage and temperature with 500 diagnostics 10 times every second, with control activity even when the car is at rest.

THE USE OF HOSE CLAMPS INDICATES THAT THE CAR IS A LIMITED-PRODUCTION VEHICLE.

The battery-interface and monitoring module mounts atop the pack's front. This unit has four orange monitoring PCBs, indicating high voltage—one for each pack section (**Figure 2**). Freescale, LG Chem, and STMicroelectronics chips populate these PCBs; the LG Chem and STMicro chips use bipolar CMOS DMOS (diffused-metal-oxide-semiconductor) technology. Midvoltage boards are blue, and low-voltage PCBs are green. Quality checks are in place throughout the manufacturing process; each cell connector bears multiple inspectors' marks.

Getting the battery electronics correct is difficult; the system must measure within several millivolts at the top of each cell, while a cell can be offset from

Figure 1 The 375-lb (170-kg) lithium-ion battery pack is the heart of the Volt. The vehicle's systems and software maintain the battery pack's health for a long service life (courtesy Munro & Associates).

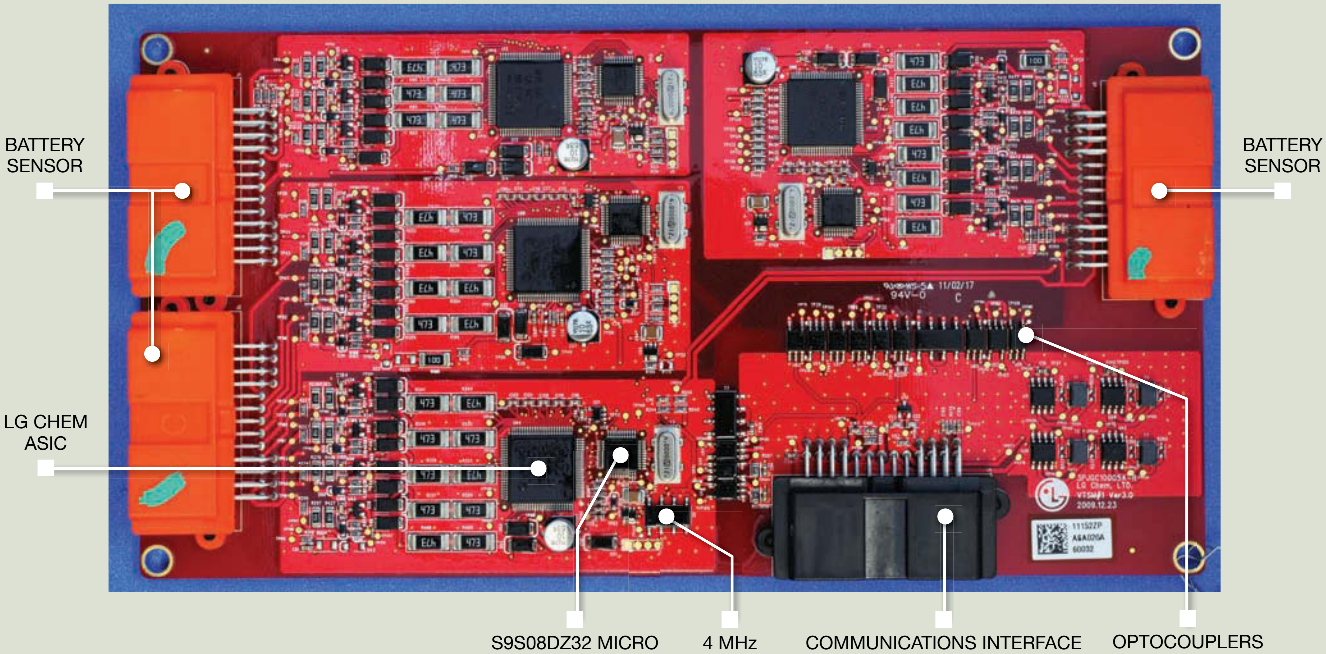

Figure 2 On the battery-interface/battery-monitor PCB, sensors on each cell monitor temperature and voltage. Their data routes in clusters, in which readouts for 10 cells are on one circuit that digitizes to a microcontroller. An optoelectric coupler feeds a common bus to the main controller in the inverter module (courtesy Munro & Associates).

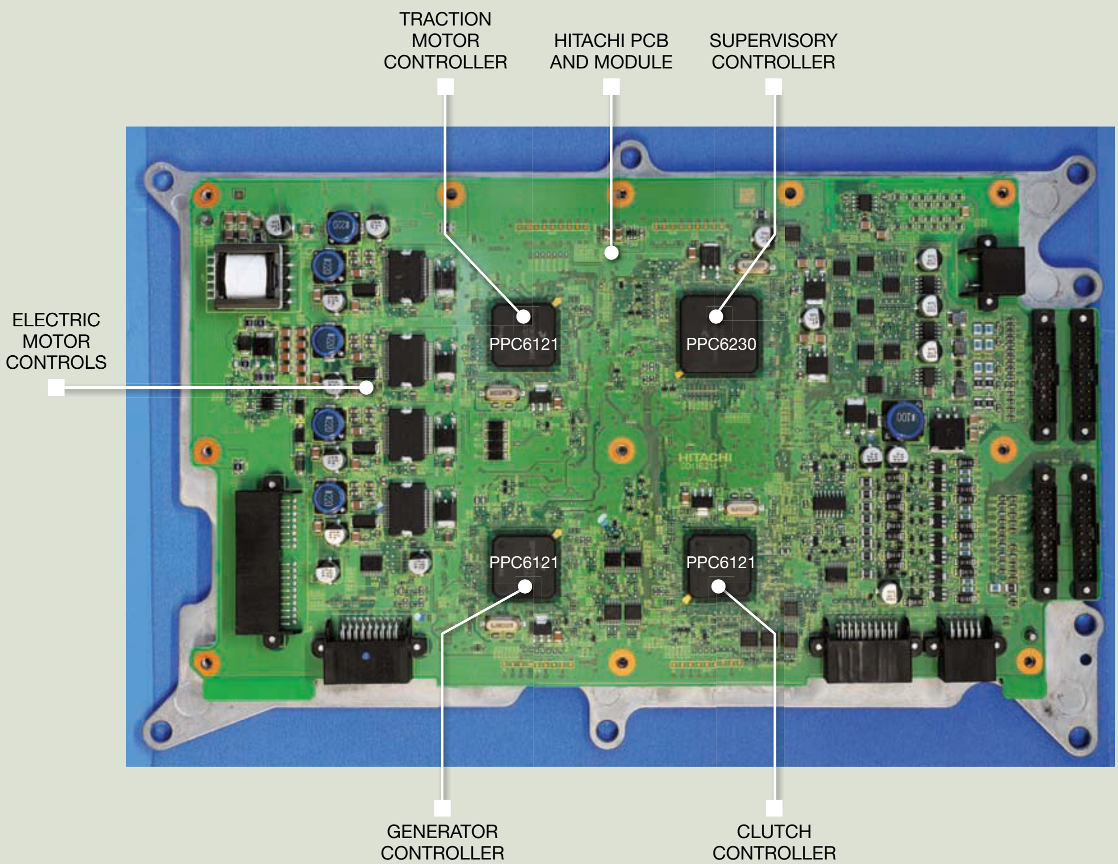

Figure 3 The motor/generator inverter module contains the brains of the Volt power train, in which one supervisory microcontroller and three others define the operating state of the drive and the regenerative-braking system (courtesy Munro & Associates).

ground by hundreds of volts. This task requires attention to PCB layout, trace design, ground planes, and voltage-isolation techniques. Scott-Thomas observes that the car's design is a work in progress, with flexibility and modularity to easily introduce new cells, battery packs, electronics, and controls.

The teardown team discovered one unexpected battery-related module in the Volt. In addition to the standard onboard-diagnostics port under the driver's side dashboard, the team found a sealed and potted module under the front passenger seat. This module stores battery- and hybrid-operation diagnostic codes and has a connection for an appropriate cable for a technician to access them.

CHARGING SYSTEM

Besides using regenerative braking, the battery pack stores energy by charging from the power grid using the supplied 110V charger or an optional 220V charging station—for faster recharging—installed by a licensed electrician (**Reference 3**). Lear Corp manufactures the 110V home charger whose power electronics and software are sophisticated enough that charging does not occur if the user plugs it into an insufficiently grounded circuit. The charger's relays and monitoring-electronics board communicate with the battery pack and onboard monitoring systems. As noted previously, the onboard system for changing ac power mains into dc power to charge the battery has its own cooling loop.

The charger plugs into a standardized receptacle behind a door on the left front fender. According to Scott-Thomas, dismantling this interface unit reveals GM's attention to design

ANATOMY OF AN AUTOMOTIVE TEARDOWN

Before beginning a teardown for Munro & Associates, Al Steier, senior associate at the company, reads all of the available information he can find on the car. For hybrid and electric vehicles, a basic step is to locate the service disconnect that makes the high-voltage lines safe when he removes it and then secure it in his toolbox. He takes photos, from all possible angles, of components before and after their removal for documentation and as part of the effort to fathom materials and manufacturing processes. Steier determines components and their makers at the PCB level. Capping off ICs and ASICs shows memory capacity if he cannot get this information from the components' published data.

The Volt came to the teardown just like any other new car from a dealer—with a full tank of gas. The teardown team left systems on to drain the lithium-ion battery before taking apart the car, but the vehicle software started the gas engine to prevent deep discharge of the battery. The teardown team decided to drain the tank and then leave on the lights, radio, and other systems, draining the 12V battery. The high-voltage system recharged the battery overnight—the system software did not allow the battery to completely discharge—but only to a level at which the car could travel 35 miles. A company specializing in electric vehicles then drained the battery using a power resistor across the terminals.

Steier had previously taken apart a Toyota Prius hybrid and found some differences between the Prius and the Volt (Reference A). For example, as a plug-in vehicle, the Volt has an extra inverter module for charging. The Volt uses a lithium-ion battery, whereas the Prius uses a nickel-metal-hydride battery. For hybrid thermal management, the plug-in Volt uses liquid cooling, whereas the Prius features air cooling. Further, in supplying the electronic components, the Volt seems to have a more diverse supplier base; the Prius primarily uses Toyota technology.

REFERENCE

A Mannion, Patrick, "Video: Tear down that Prius!" *EE Times*, April 26, 2007, http://bit.ly/HiDFoq.

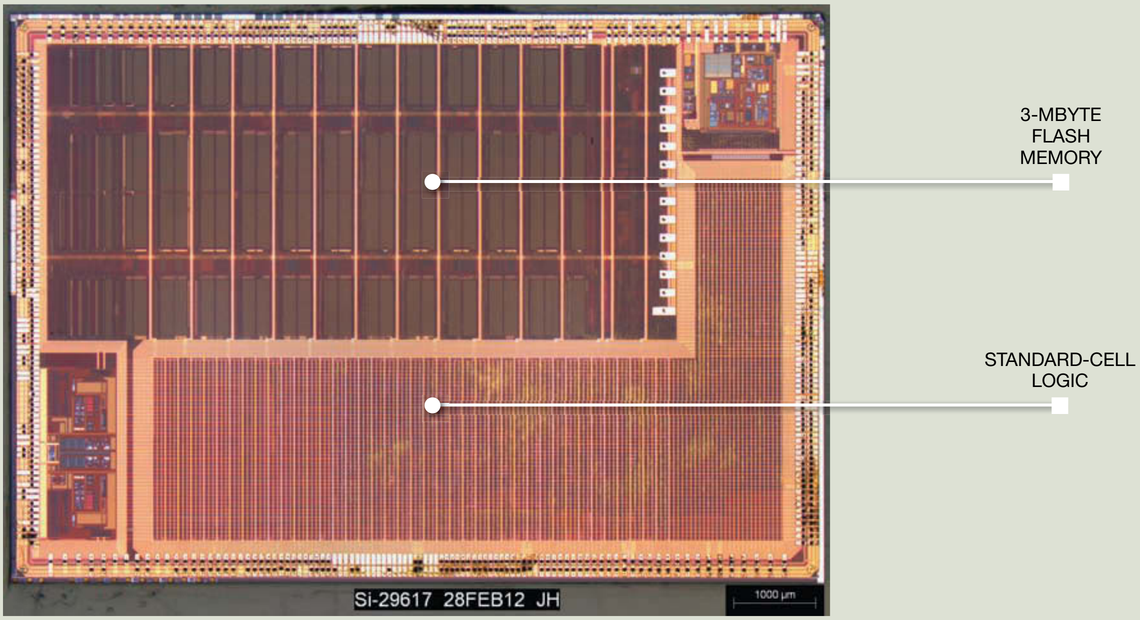

Figure 4 The inverter's supervisory controller is the center of the Volt's hybrid architecture and determines current power-train state (courtesy Munro & Associates).

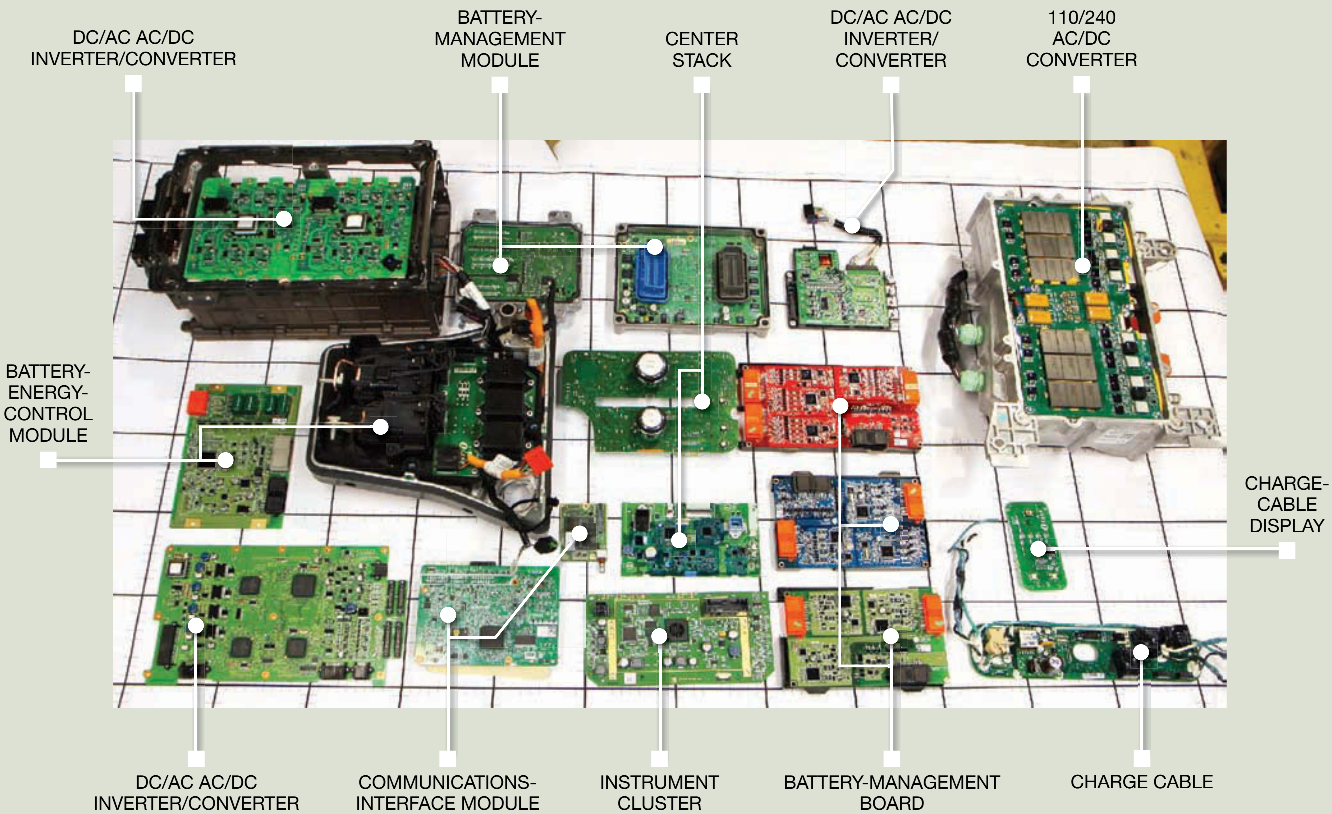

Figure 5 Of the 18 electronic modules in the Volt, roughly three-fourths handle hybrid-power-train functions (courtesy Munro & Associates).

DRIVE FOR INNOVATION
DESTINATION: AVNET EXPRESS

The Chevy Volt teardown took place during the months leading up to Design West 2012 as part of an Avnet/UBM Electronics program to highlight innovation.
For more on this program, go to www.driveforinnovation.com.

detail. The manufacturer taped and foam-isolated the high-voltage components—that is, the capacitors and the common-mode choke—in this high-vibration environment for robustness and protection, and windings are robust, stable, and mechanically redundant.

Steier found one puzzling feature with the charger configuration: Although the charger receptacle is on the left fender, the charging inverter it feeds is under the right headlight. Likewise, the gas-engine controller is on the left side, whereas the engine is on the right. Such an arrangement has greater wiring weight than if the manufacturer had reversed this configuration.

BRAINS OF THE VOLT

On the electric motor and generator housing, which looks like the transmission case of a gas-engine car, a liquid-cooled inverter module feeds battery power into the traction motor. The high-voltage orange cables leading to this module have disconnects with relays for safety; Steier notes that the module cover itself is also a safety-circuit disconnect. Inside is what Scott-Thomas says is the closest thing to a central brain in the car (**Figure 3**).

The Hitachi PCB includes four 32-bit Freescale Qorivva microcontrollers. Scott-Thomas first noticed the large amount of real estate available, noting that it allows future modifications, either by changing or adding circuits. One of the four controllers functions as the supervisor, using inputs, including vehicle and wheel speeds; acceleration, or throttle; braking; and battery state, to decide which state is the most efficient (**Figure 4**). Decisions to be made could include, for instance, choosing which combination of outputs from the traction motor and combustion-engine-generator motor to use, when to activate regenerative braking, and to what extent to recover energy. The supervisory controller is the largest of the four microcontrollers, with 3 Mbytes of flash memory, taking half the area of the die. The controller also endeavors to run the electric motors at lower rotation rates for more efficiency. The other three Freescale microcontrollers control the traction motor; the combustion-engine-driven generator; and the clutched planetary gear set, which the IC engine can engage if necessary.

OTHER ELECTRONICS

The rest of the Volt's electronics, unless they connect with the hybrid-drive system, are fairly conventional state-of-the-art automotives (**Figure 5**). An air-cooled dc/dc converter, with PCBs from TDK and a Renesas microcontroller, takes the place of an alternator to provide 12V for running standard auto systems, such as doors, lights, navigation, and audio, and to charge the auxiliary 12V battery. Taking apart the center stack uncovered a communications-module PCB by LG hosting a Freescale memory controller with Spansion flash. These infotainment boards are sparse, according to Scott-Thomas, who notes there is a lot of space but not that much required processing power; thus, the electronics combine several functions onto single chips. In addition, adequate space between the resistive-touch switches on the front of the panel helps prevent the driver from incorrectly selecting the wrong function.

As customers gain experience, it will be interesting to see how—and how fast—this plug-in hybrid platform evolves in the coming years. EDN

REFERENCES

1 DeMeis, Rick, "Innovation revelation: engineering the Chevy Volt," *EE Times*, May 13, 2011, http://bit.ly/HJSyOx.
2 DeMeis, Rick, "We tear down a Chevy Volt," *EE Times*, March 12, 2012, http://bit.ly/HOpY0Y.
3 DeMeis, Rick: "Driving impressions: Chevy Volt," *EE Times*, Aug 2, 2011, http://bit.ly/tB3W3U.

You can reach Designline Editor Rick DeMeis at jldmld@comcast.net. Read more at www.eetimes.com/design/automotive-design.

BY GERARD WIMPENNY • NUJIRA LTD

Understand and characterize envelope-tracking power amplifiers

IN CONTRAST TO FIXED-SUPPLY POWER AMPS, THE PERFORMANCE OF AN ENVELOPE-TRACKING POWER AMPLIFIER IS NOT SELF-CONTAINED AND REQUIRES SUBSTANTIALLY MORE DATA TO PREDICT SYSTEM PERFORMANCE.

The process of designing traditional fixed-supply power amplifiers has been well-established for many years. Well-defined metrics for performance assessment exist, and the amplifier designer's job is to design a power amp with the best set of performance metrics. This task is far from simple, but designers at least understand the well-established assessment criteria. For envelope-tracking power amplifiers, the situation is more complex and requires the use of more sophisticated characterization techniques.

The objective of envelope tracking is to improve the efficiency of power amps carrying high peak- to average-power-ratio signals. The drive to achieve high data throughput within limited spectrum resources requires the use of linear modulation with high peak to average power. Unfortunately, traditional fixed-supply power amplifiers operating under these conditions have low efficiency. You can improve the efficiency of an envelope-tracking power amplifier by varying the amplifier's supply voltage in synchronism with the envelope of the RF signal. The power amplifier's fundamental output characteristics—power, efficiency, gain, and phase—now depend on two control inputs, RF input power and supply voltage, and can be represented as 3-D surfaces.

A typical envelope-tracking system dynamically adjusts the supply voltage to track the RF envelope at high instantaneous power. In this case, the power amplifier operates with high efficiency in compression. The instantaneous supply voltage primarily determines the amplifier's output characteristics. Conversely, when the instantaneous RF power is low, the supply voltage remains substantially constant, and the instantaneous input power in the linear region primarily determines the power amplifier's output characteristics. A transition region in which both supply voltage and input power influence the output characteristics exists between these two extremes (**Figure 1**).

ENVELOPE-TRACKING LINEARITY

You can construct a simple quasistatic—that is, memoryless—behavioral model of a power amplifier if you know

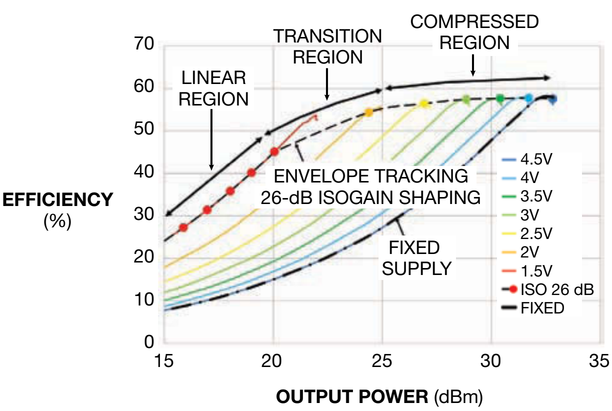

Figure 1 When the instantaneous RF power is low, the supply voltage remains substantially constant, and the instantaneous input power in the linear region primarily determines the power amplifier's output characteristics.

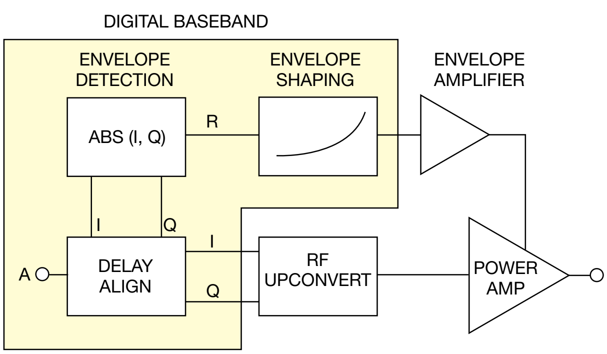

Figure 2 The mapping between the instantaneous RF envelope and the applied supply voltage profoundly influences these characteristics, along with other key power-amp metrics, such as power and efficiency. In an envelope-tracking system, the contents of a shaping table in the envelope path determine this mapping.

its AM (amplitude-modulation)/AM and AM/PM (phase-modulation) characteristics. The mapping between the instantaneous RF envelope and the applied supply voltage profoundly influences these characteristics, along with other key power-amp metrics, such as power and efficiency. In an envelope-tracking system, the contents of a shaping table in the envelope path determine this mapping (**Figure 2**).

To achieve "ISOgain" shaping, the mapping between RF envelope and supply voltage is chosen to achieve a particular constant power-amplifier gain (**Figure 3**). With this mapping, the envelope-tracking amplifier system achieves low AM/AM distortion despite operating in compression over much of the envelope cycle (**Figure 4**). The **figure** also shows the equivalent trajectory for fixed-supply operation; from this trajectory, it is apparent that you can use envelope tracking to linearize a power amplifier, reducing adjacent-channel power ratio and error-vector magnitude.

The system trade-off of using the shaping table to linearize the power amplifier is a small loss of efficiency for a substantial improvement in linearity (compare **figures 1** and **5** and **figures 4** and **6**). The choice of shaping function also has a strong influence on the bandwidth requirement of the envelope path. A smooth transition between the linear and the compressed regions results in a lower bandwidth requirement for the envelope amplifier for a 1 to 2% loss in system efficiency.

When designing a fixed-supply linear power amplifier, you must pay a great deal of attention to achieving adequate linearity characteristics at maximum output power. Many factors, including fundamental technology characteristics, biasing, and RF matching, influence the linearity, and it is up to the designer to achieve the best trade-off between efficiency and linearity. For an envelope-tracking power amplifier, however, the linearity in the compressed region is no longer a self-contained power-amplifier parameter. The amplifier still must be linear in the low-power, low-voltage region. At

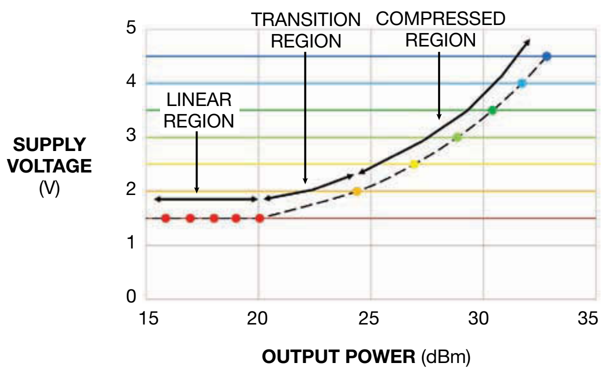

Figure 3 To achieve "ISOgain" shaping, the mapping between RF envelope and supply voltage is chosen to achieve a particular constant power-amplifier gain.

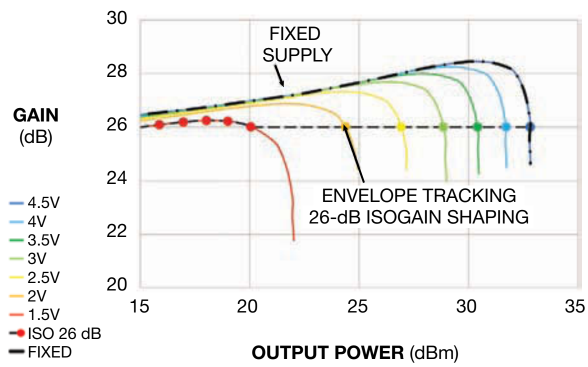

Figure 4 The envelope-tracking amplifier system achieves low AM/AM distortion despite operating in compression over much of the envelope cycle.

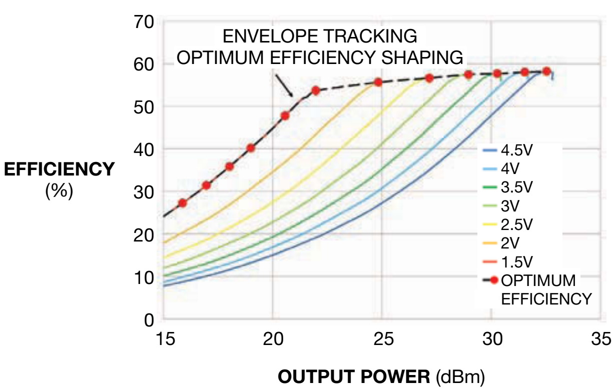

Figure 5 The system trade-off of using the shaping table to linearize the power amplifier is a loss of efficiency for a substantial improvement in linearity. A smooth transition between the linear and the compressed regions results in a lower bandwidth. See also figures 1, 4, and 6.

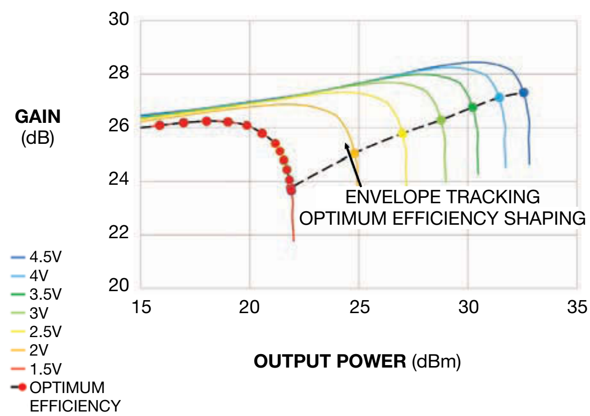

Figure 6 Choosing a shaping table for optimum efficiency may introduce AM/AM nonlinearities in the power amplifier.

higher powers, however, there is no AM-linearity constraint, and developers can design the power amplifier for optimum envelope-tracking efficiency without regard to AM linearity. Unlike with AM distortion, the envelope shaping table does not directly control phase distortion. However, many power amplifiers show reduced PM distortion when operating in envelope-tracking mode.

As a result of this self-linearization, you can push harder into compression at signal peaks with an envelope-tracking system than with a fixed-supply amplifier, allowing increased output power for given linearity. **Figure 7** shows measured adjacent-channel leakage ratio and error-vector-magnitude performance for a power amplifier operating in fixed-supply and envelope-tracking modes. In this example, the amplifier's output power for −40-dBc adjacent-channel leakage ratio is 2 dB higher in envelope-tracking than in fixed-supply mode.

CHARACTERIZATION TECHNIQUES

You cannot measure the stand-alone performance of envelope-tracking power amps without first defining the shaping table. This definition requires measurement of the power amplifier's fundamental characteristics—output power, efficiency, gain, and phase—over the full range of supply voltage and input power.

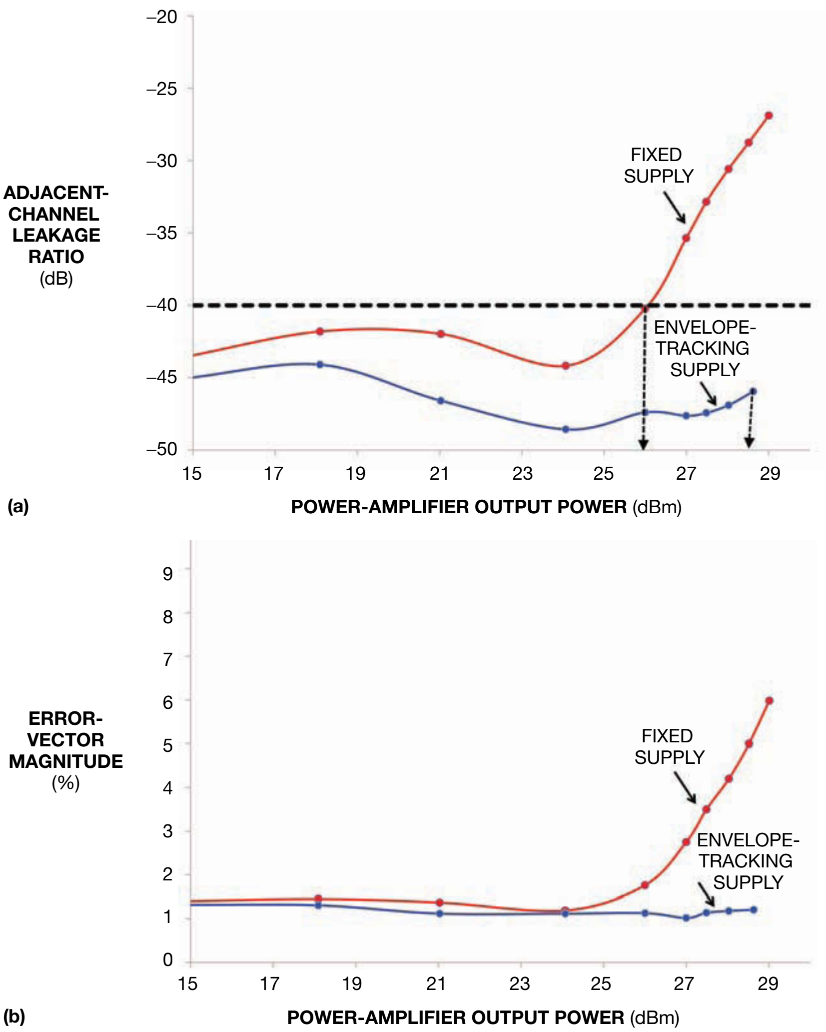

Figure 7 The amplifier's output power for –40-dBc adjacent-channel leakage ratio is 2 dB higher in envelope-tracking than in fixed-supply mode (a). Error-vector-magnitude performance is also shown (b).

In principle, you could perform this characterization using a continuous-wave network analyzer and a variable dc supply, but results are typically poor due to thermal effects, ranging errors, and drift in phase measurements. It is also too slow to allow the use of load-pull techniques. An alternative approach is to use a pulse characterization using standard automatic-test equipment. This approach avoids the need for a high-bandwidth, low-impedance supply and is sufficiently fast for load pull to be viable. The approach makes it difficult to make accurate phase measurements, however. A third approach is to use real waveforms and to vary the shaping table to allow the measurement of all combinations

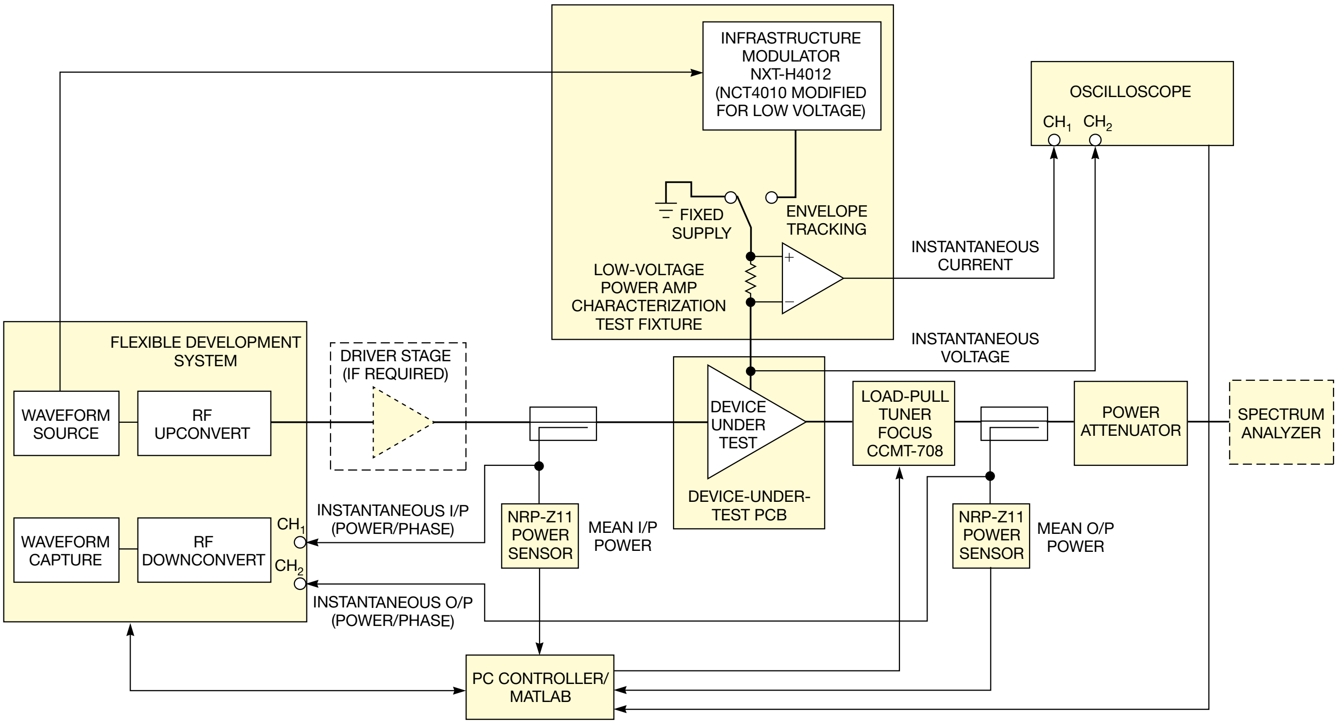

Figure 8 An automated test-and-measurement configuration using an envelope-tracking supply modulator allows accurate capture and measurement of instantaneous power-amplifier efficiency, gain, and phase across all combinations of input power and supply voltage under dynamic supply-modulation conditions.

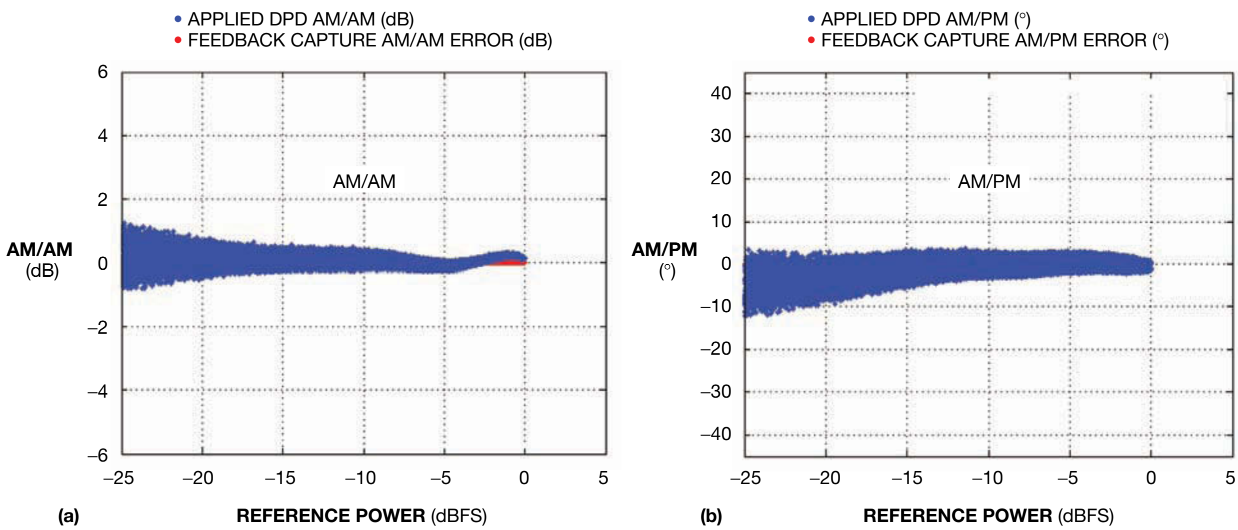

Figure 9 You can use the same hardware for both the power-amplifier device-level characterization and the direct verification of power-amp system performance using a defined shaping table to capture AM/AM response (a) and AM/PM response (b).

of input power and supply voltage. This approach requires a supply modulator but is fast, allows you to gather accurate phase information, and can also characterize memory effects (**Figure 8**).

You can use a basic envelope-tracking power-amp characterization to create a quasistatic data model of the power amplifier. This model can have output power, phase, and efficiency as outputs and input power and supply voltage as inputs. Once the shaping table is defined, you can use the model to predict the amplifier's performance parameters, such as adjacent-channel power ratio, error-vector magnitude, and efficiency for standard test waveforms.

You can use the same hardware for both the power-amplifier device-level characterization and the direct verification of power-amp system performance using a defined shaping table (**Figure 9**). For higher-bandwidth waveforms, the amplifier's memory effects can be a significant source of nonlinearity. The power amplifier's output parameters, including AM, PM, and efficiency, now depend on time—that is, the signal history—along with instantaneous input power and supply voltage. Memory effects show up in the amplifier's characterization as a broadening of the AM/AM and AM/PM characteristics and can result from electrical time constants in input or output bias circuits, thermal time constants associated with local die heating, or technology-specific charge-storage effects.

INCREASING EFFICIENCY

The statistics of typical high peak- to average-power-ratio signals are such that an envelope-tracking power amplifier typically spends most of its time operating with relatively low supply voltage, with only occasional high-voltage excursions on high-power peaks. It makes sense, therefore, to optimize the amplifier's matching to achieve the best efficiency with the target peak- to average-power-ratio signals rather than simply designing for best efficiency at peak power and maximum supply voltage, as would be the case for a fixed-supply power amplifier. Designers can alter the amplifier's matching to increase efficiency around the peak of the signal's probability-density function, even if this necessitates a slight compromise in the peak power efficiency, as the following **equation** shows (**Figure 10**):

$$\text{EFFICIENCY (A)}=\frac{\int_0^{V_{MAX}} \text{PROBABILITY (V)}\times\text{POWER (V)}}{\int_0^{V_{MAX}} \dfrac{\text{PROBABILITY (V)}\times\text{POWER (V)}}{\text{EFFICIENCY (V)}}}.$$

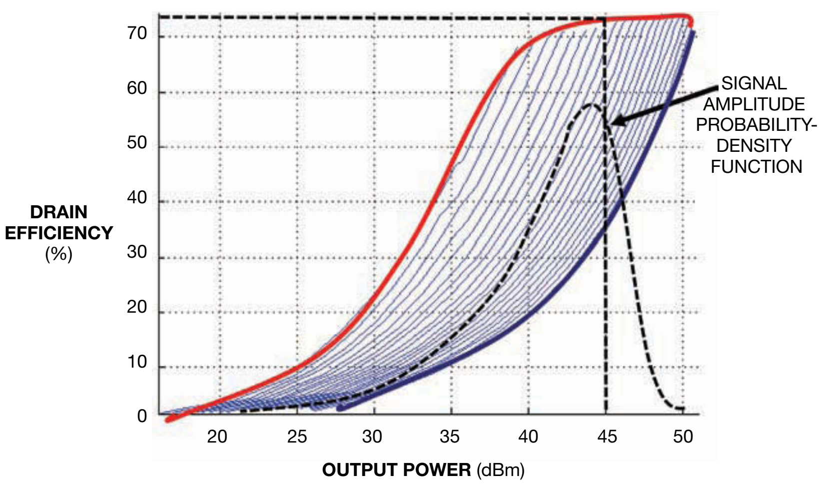

Figure 10 Designers can alter the amplifier's matching to increase efficiency around the peak of the signal's probability-density function, even if this necessitates a slight compromise in the peak power efficiency.

To fully optimize the efficiency of an envelope-tracking power amplifier, you can extend the device's characterization to include sweeping the load impedance—using fundamental or harmonic load pull—along with the input power and the supply voltage. This characterization produces a large set of data, and tools, such as Matlab, can be used to automate the analysis of this data to predict the average efficiency when operating with a specific set of envelope-tracking parameters. Using this characterization method, you can predict how an amplifier's average efficiency varies with shaping function, output-voltage swing, backoff from maximum power, and waveform statistics when operating in envelope-tracking mode (**Figure 11**).

PARAMETER VARIATION SENSITIVITY

You might expect the performance of envelope-tracking power amplifiers over temperature to be worse than that of their fixed-supply counterparts. The reverse situation is true, however. An envelope-tracking power amplifier's performance is more sensitive than that of a fixed-supply amp to changes in the supply-voltage characteristics than to changes in gain of the RF chain driving the power amplifier. Because you can better control the characteristics of the supply voltage over temperature than the variation of the RF gain, little variation in linearity occurs for extreme temperature variations (**Figure 12**).

In a handset environment, the power amp receives an uncontrolled load impedance due to reflections from nearby objects, which can result in the amplifier's having to work into load mismatches with a VSWR (voltage-standing-wave ratio) as high as 3-to-1. The envelope-tracking power amplifier's self-linearization principle also applies under high-VSWR conditions, and this operation can result in significantly better adjacent-channel power ratio and error-vector-magnitude performance than that of an amplifier operating

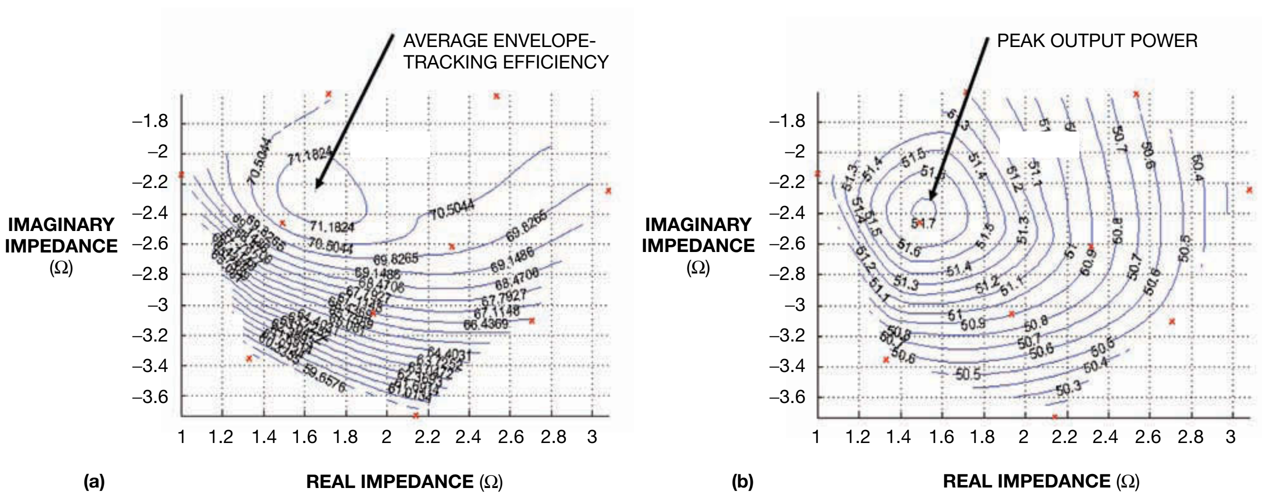

Figure 11 You can predict how an amplifier's average efficiency varies with shaping function, output-voltage swing, backoff from maximum power, and waveform statistics when operating in envelope-tracking mode (a); also shown is peak output power (b).

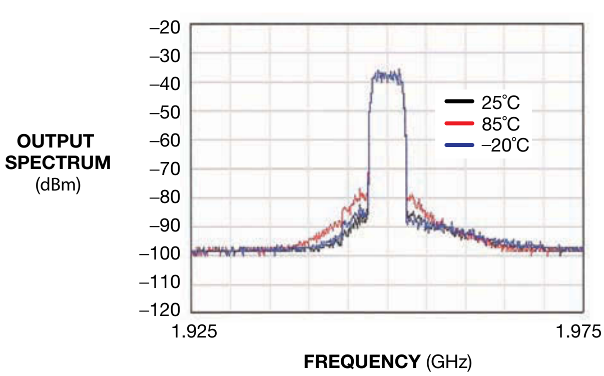

Figure 12 Because you can better control the characteristics of the supply voltage over temperature than the variation of the RF gain, little variation in linearity occurs for extreme temperature variations.

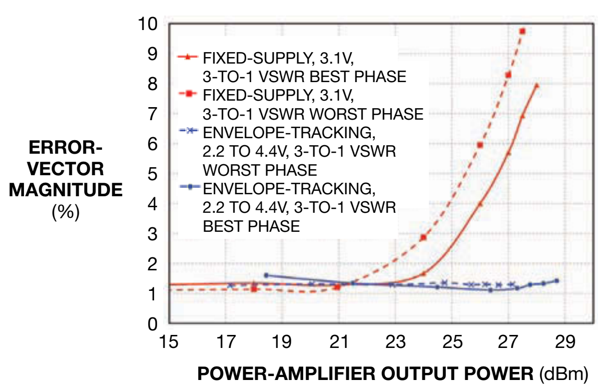

Figure 13 The envelope-tracking power amplifier's self-linearization principle also applies under high-VSWR conditions, and this operation can result in significantly better adjacent-channel power ratio and error-vector-magnitude performance than that of an amplifier operating in fixed-supply mode.

in fixed-supply mode (**Figure 13**).

The system-efficiency benefit of operating a power amplifier in envelope-tracking mode is well-known. However, it also offers other useful system benefits, such as increased output power, improved operation into mismatched loads, and insensitivity to temperature variations. In contrast to fixed-supply power amplifiers, the performance of an envelope-tracking power amplifier requires the gathering of substantially more data to predict system performance and the use of a test environment that allows sweeping of the supply voltage and the input power. A key aspect is the definition of the shaping table, which defines the relationship between supply voltage and RF power. Once you define the shaping function, you can directly measure efficiency and linearity using an appropriate system-characterization bench. EDN

A KEY ASPECT IS THE DEFINITION OF THE SHAPING TABLE. ONCE YOU DEFINE THE SHAPING FUNCTION, YOU CAN MEASURE EFFICIENCY AND LINEARITY USING A SYSTEM-CHARACTERIZATION BENCH.

ACKNOWLEDGMENT

This article originally appeared on EDN's *sister site*, RF and Microwave Designline, *http://bit.ly/w7gTef*.

AUTHOR'S BIOGRAPHY

Gerard Wimpenny is the chief technology officer of Nujira Ltd and a member of the OpenET Alliance. He has more than 20 years of RF and signal-processing experience and has been responsible for strategic R&D, design-process definition, and top-level technical support for business-development activities. Wimpenny has been instrumental in delivering numerous wireless products to semiconductor vendors and infrastructure and handset manufacturers. He has a master's degree from Cambridge University (Cambridge, UK).

42V, 2μA I_Q Low Dropout Switcher

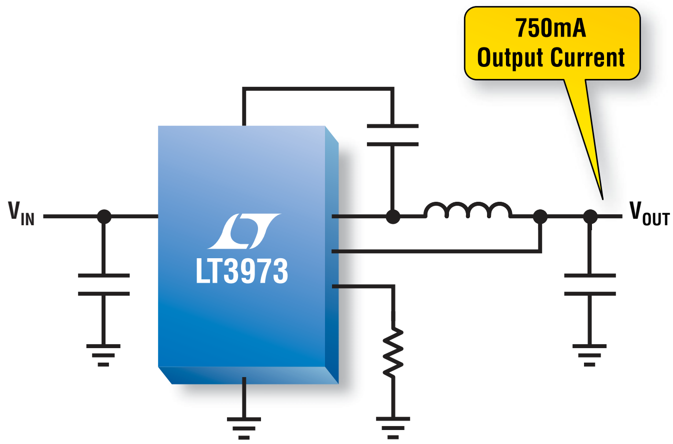

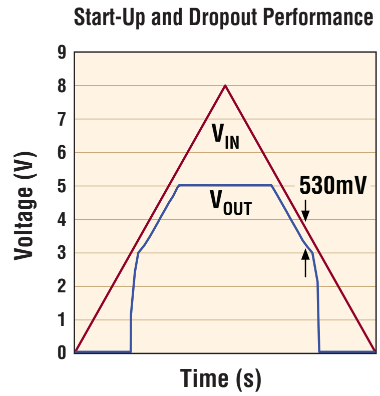

530mV Maximum Dropout

The LT®3973 is the newest member of our growing family of ultralow quiescent current high voltage monolithic buck regulators. It consumes only 1.8μA of quiescent current while regulating an output of 3.3V from a 12V input source. A high efficiency switch is included on-chip along with the catch diode, boost diode and all necessary control and logic circuitry. A minimum dropout voltage of 530mV is maintained when the input voltage drops below the programmed output voltage, providing a regulated output to the downstream load. Its low ripple Burst Mode® operation maintains high efficiencies at low output currents while keeping output ripple below $10mV_{P-P}$.

Features

- Ultralow Quiescent Current: 1.8μA I_Q at $12V_{IN}$ to $3.3V_{OUT}$
- Low Ripple Burst Mode Operation
- Input Voltage Range: 4.2V to 42V
- Integrated Boost and Catch Diodes
- Excellent Start-Up & Dropout Performance
- 750mA Output Current
- Adjustable Switching Frequency: 200kHz to 2.2MHz

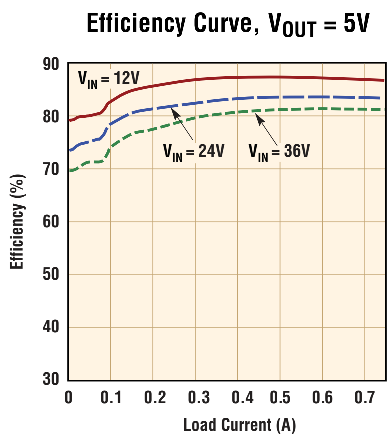

Info & Free Samples

www.linear.com/product/LT3973

1-800-4-LINEAR

LINEAR TECHNOLOGY

Build a digital PLL with three ICs

Dave Allen, Dash Inc, Kansas City, KS

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

The simple circuit in this Design Idea exhibits the basic characteristics of a traditional analog phase-locked loop but has no analog components other than the reference oscillator. Other digital PLLs exist, including those employing an up/down counter, but this one is simpler and more flexible.

The circuit initially found use more than 30 years ago as a clock regenerator in a data separator for a self-clocking code, such as Manchester or biphase, in magnetic recording. It quickly became clear that it has many other applications. The circuit also served as the basis of a servo controller for a tape drive's capstan motor/tachometer. LSI disk/tape-controller chips incorporated both the data separator and the capstan servo controller, with the advantage of having no analog circuitry and no requirements for adjustment. Because it was used in the production of commercially available products so long ago, it is not patentable today and is free for use.

The example in **Figure 1** uses only three ICs to make prototyping quick and the explanation simple. The connections between the 74161 counter outputs and the preset inputs form a rudimentary ROM implementing a look-up table (**Table 1**). The 16XREF should be a square wave or at least not a narrow pulse because you must take into account things that happen on both the leading and the trailing edges and setup times. The INPUT pulse must be long enough to meet the clock pulse-width requirements of the logic family you use for the 7474 D flip-flop.

To test your prototype, make the INPUT approximately one-sixteenth of the 16XREF frequency and watch the output as you slowly vary the INPUT frequency. Use a signal generator that allows fine adjustment of the INPUT to measure lock range slightly above and below one-sixteenth of your XREF source. The dither is equal to the period of the 16XREF's clock, but the output stays locked to the INPUT as you vary the INPUT ±20% or more. You can temporarily disconnect Pin 9 of the counter to see the output slipping past the INPUT when the frequencies are close to each other. Reconnecting Pin 9 demonstrates the locking action. The output is a square wave when the INPUT is exactly one-sixteenth of 16XREF but becomes rectangular as you go above or below the center frequency.

In operation, the counter counts continuously, but each rising edge of the INPUT signal causes a preset pulse

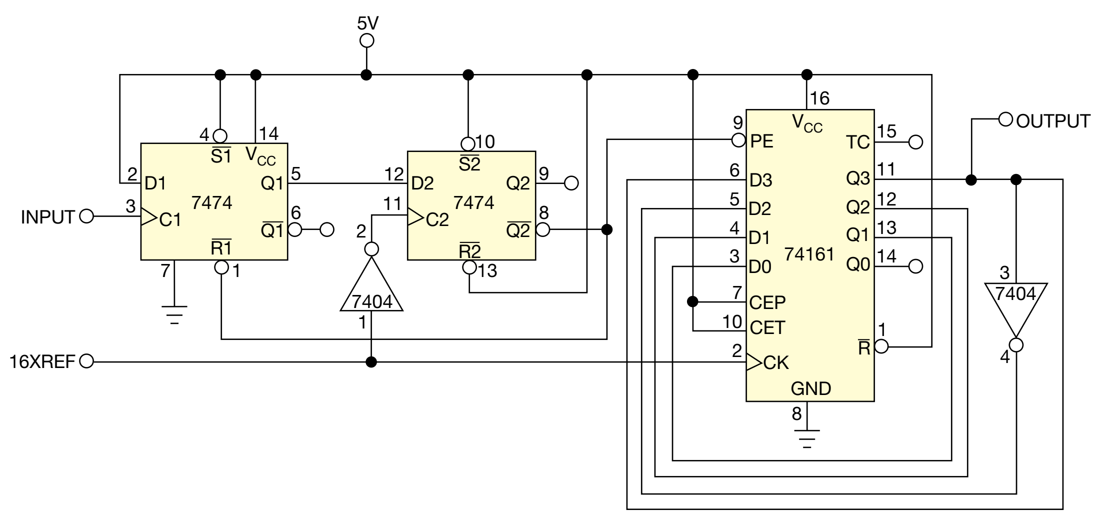

Figure 1 The counter loads its own data input to generate an output locked to the input signal.

at the counter. From **Table 1**'s count and preset values, you can deduce that, whenever the counter gets a preset pulse, the count moves closer to seven or eight. If it is already at seven or eight, it remains there. The servo-loop error signal is the difference, at the moment the preset signal arrives, between the counter's current state and seven or eight. This simple example uses the count value to halve the error signal for the preset.

If the INPUT signal is exactly one-sixteenth of the reference but starts up at 180° out of phase, then the first preset pulse might occur when the counter is at 15. So, the counter presets to 11 and resumes counting from there. At the next preset pulse, the counter is at 10 and presets to nine. The next preset pulse occurs at a count of eight and presets to eight. It next presets to seven; when the next preset pulse comes in while the counter is at six, it again presets to seven and is now synchronized. The preset pulses arrive just before the counter's most-significant bit goes from 0 to 1, which is what the INPUT signal is also doing.

If the INPUT signal is a little slower than one-sixteenth of the reference, the preset pulses arrive after the counter has counted beyond eight—to 12, for

TABLE 1 MINIMAL COUNTER PRESET

Count value	Q3	Q2	Q1	Q0	D3	D2	D1	D0	Preset value
15	1	1	1	1	1	0	1	1	11
14	1	1	1	0	1	0	1	1	11
13	1	1	0	1	1	0	1	0	10
12	1	1	0	0	1	0	1	0	10
11	1	0	1	1	1	0	0	1	9
10	1	0	1	0	1	0	0	1	9
9	1	0	0	1	1	0	0	0	8
8	1	0	0	0	1	0	0	0	8
7	0	1	1	1	0	1	1	1	7
6	0	1	1	0	0	1	1	1	7
5	0	1	0	1	0	1	1	0	6
4	0	1	0	0	0	1	1	0	6
3	0	0	1	1	0	1	0	1	5
2	0	0	1	0	0	1	0	1	5
1	0	0	0	1	0	1	0	0	4
0	0	0	0	0	0	1	0	0	4

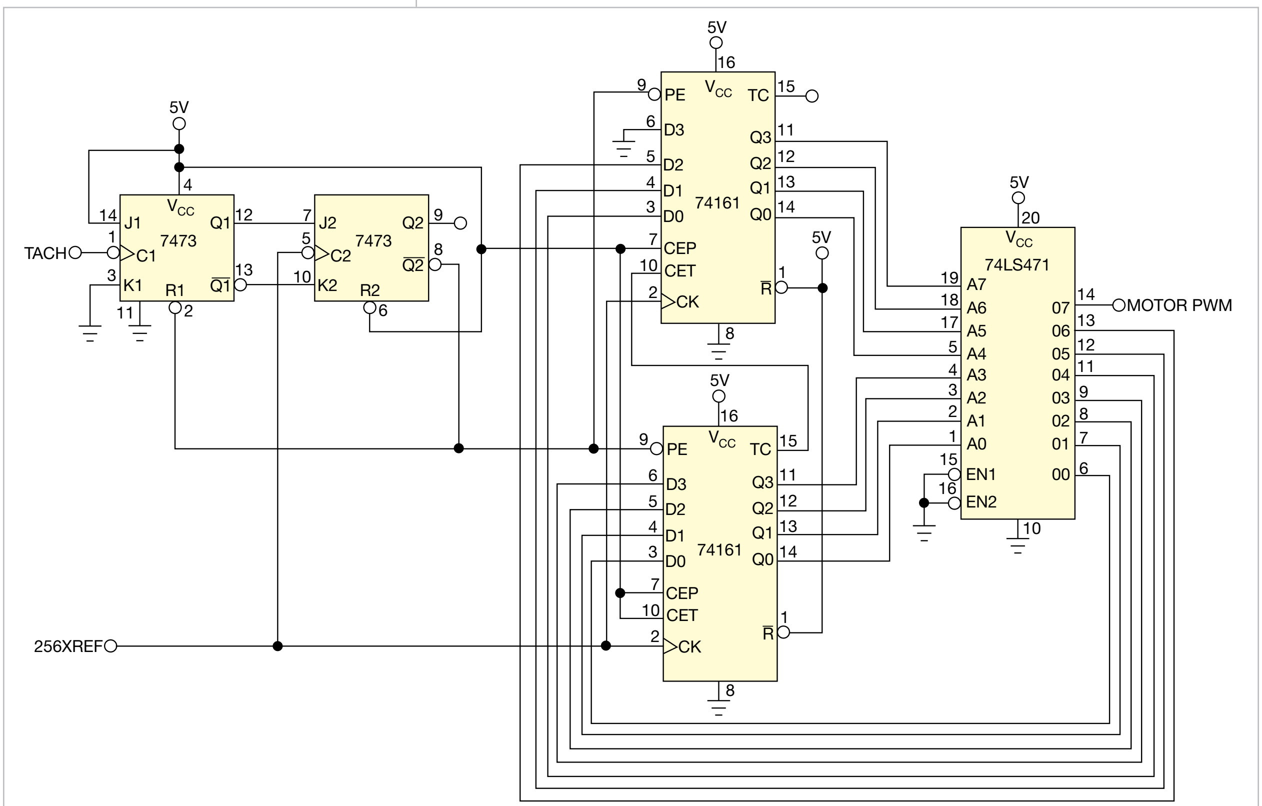

Figure 2 Two counters and a PROM enhance the versatility of the lock function.

example. It gets preset to 10. Because the INPUT signal is too slow, however, the counter again goes to 12 when the next preset pulse arrives. The circuit is still locked, but the MSB stretches to match the slower INPUT signal. The same process maintains lock for an INPUT signal that is faster than one-sixteenth of the reference, but the MSB signal's period shrinks to maintain lock.

An out-of-lock condition occurs if the INPUT is so slow that the counter goes past 15 and wraps around to zero or beyond before the preset occurs. It's likewise out of lock if the counter can't even count to zero before the next preset pulse. The circuit can lock on multiples and submultiples of the 16X reference.

You can tailor the locking characteristics and reduce the dither by adding more counter bits and putting a ROM between the counter's outputs and the preset's inputs (**figures 2** and **3**). By using a PROM, you can, for instance, divide the error by three or by four, which increases the lock range. You can also use a PROM to subtract one or two from the error signal instead of dividing the error by two. This approach dramatically narrows the lock range. You can use additional PROM output lines—that are not presetting the counter—for other functions.

Because the counter's preset pulse occurs at the moment that the error signal is available, you can program some more bits in the PROM and latch the error condition for another application's benefit, such as to indicate an unlocked state or to indicate the INPUT frequency as high, low, or centered relative to the reference. This scheme, with the motor controller, could indicate the motor's load as light, moderate, or heavy. For other applications, you could program a second parallel PROM to generate a sine wave by feeding samples to a DAC or a quasi-sine wave for a power inverter.

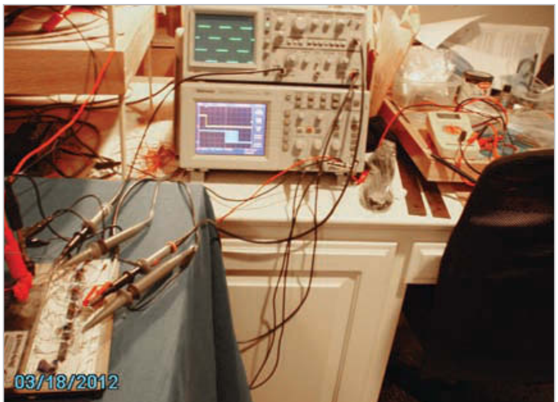

Figure 3 You can trigger from the input signal in the bottom scope's top trace to see the dither on the locked output signal on the bottom scope's bottom trace.

As a capstan-motor servo controller, the tachometer is the INPUT to the PLL, and the motor speed locks to a crystal reference. One output bit from the PROM, tailored for the motor PWM (pulse-width-modulated) signal, enables the servo to better control the duty cycle. The MSB of the preset is 0, forcing the system to work in the lower half of the address space whenever it is locked. This approach frees up the top PROM output line, which becomes the motor-control signal.

Using an 8-bit counter and a 256×8-bit PROM provides lots of room and many options for optimizing the motor's behavior under varying load conditions. The programming of the top PROM line determines where in the counter's cycle the motor PWM signal turns on and off. If the load on the motor is heavy, it slows down, letting the counters count longer and slightly higher before the preset occurs. As the counters count higher, the motor bit stays on longer, increasing the duty cycle of the PWM signal to compensate for the heavy load. The center point of the servo is 63/64, keeping locked operation in the lower half of the address space. The upper half of the address space is therefore in use only during motor start-up, so programming the PROM's motor PWM bit "on" whenever the counter is that high provides extra starting torque.

By programming the PROM, you can control the lock range, or loop gain, to match the load variations; you can tailor the duty cycle to match the motor's torque characteristics; and you can control the start-up torque. EDN

Power-supply decoupler protects your UUT

Raju Baddi, Tata Institute of Fundamental Research, Pune, India

You can install the bench-top accessory in **Figure 1** between a variable-voltage bench power supply and a breadboard or UUT (unit under test) to protect against accidental overvoltage and reverse polarity. It draws its power from the supply and pulses a 5V, two-coil latching relay to interrupt power to the load under abnormal conditions. The latching relay uses permanent magnets to hold the DPDT (double-pole/double-throw) contacts in the position, which the most recent pulse to the associated coil sets. It is nominally rated at 5V for the coils but operates at voltages as low as 3.5V.

In the normal state, the relay passes power from the attached supply at the input voltage to the load at the output voltage through the inductor and holds Q_3's emitter at 0.6V less through BR_1. C_2 charges to 1.2V less than the input voltage. C_1 cannot charge through reverse-biased D_1 or D_2.

Q_1 forms a variable-zener function, which you adjust by using the coarse and fine potentiometers to set the overvoltage-threshold base to the emitter voltage. If the voltage should exceed the threshold—for example, if a user accidentally bumps the power-supply voltage knob—the base-to-emitter voltage increases to the 0.6V necessary to turn on Q_1. This action then turns on Q_2, which, in turn, turns on Q_3 through D_1. Q_3 draws current from C_2 through the appropriate relay coil to open the contact to V_{OUT} and BR_1 and close the contact to the collector of Q_4 and an ac input terminal of BR_2, which lights the error LED. The charge on C_2 allows the relay to fully complete its latching action, even though it has disconnected its own voltage supply.

You should calibrate a dial for the potentiometers or adjust them with the power supply on but with the UUT load unconnected. This approach helps to quickly and easily set the decoupler's overvoltage threshold.

Voltage follower Q_4 charges its emitter capacitor to no more than 4.5V, which the 5.1V zener diode at its base sets. After the overvoltage condition is cleared, momentarily pressing the RST (reset) pushbutton switch discharges Q_4's emitter capacitor into the other relay coil, latching it back into its normal position.

If you accidentally apply reverse polarity to the V_{IN} and ground terminals, Q_3 biases on through D_2. BR_1 applies the normal polarity to C_2 and the relay coil to allow Q_3 to operate the relay for the overvoltage condition. BR_2 lights the LED despite the reversed polarity.

The design includes a 47- to 500-mH inductor and a 1000- to 4700-µF capacitor at the output stage to delay the rise of output voltage and current. This step avoids damage to the recipient circuit during the operating-time delay of the relay. Choose the inductor for sufficient current rating and minimum dc resistance for the load current you anticipate.

The following **equation** calculates the rise of current, I, through an inductor, L, as a function of time, T, when a voltage, V, exists across it: I=(V/L)T. The next **equation** calculates the rise in voltage of a capacitor, C, when a charge, Q, is deposited into it: V=Q/C. You can use these **equations** to calculate the required values of L and C if the relay were to operate in time T. You can obtain the charge by integrating the first **equation** with time limits zero to T. Because the voltage across the capacitor at the output must be within safe limits for the recipient circuit, the charge on the capacitor must be small enough that its voltage hardly changes, meaning that the voltage and the current are more or less constant. In this case, the charge is (I×T)/2 and I=(V/L)T. See the circuit layout online at www.edn.com/120510dia.

The data sheet usually specifies the relay's operating time (**Reference 1**). Alternatively, you can measure it with an oscilloscope or a dual-event timing circuit (**Reference 2**). The circuit uses only one of two sets of contacts. If your design requires a second protection circuit for a negative-supply voltage—with NPN swapped for PNP and the diodes reversed—you can cross-couple the extra contacts in series with the opposite supply so that a fault of either supply would disconnect both.

This Design Idea describes a general approach. You should verify that the various parameters meet the requirements of the power-recipient circuit and make appropriate modifications. EDN

REFERENCES

1 "G6A2," Omron Electronics, http://bit.ly/FQ97wF.

2 Baddi, Raju, "Bi-Event Timer for Physics Lab," January 2012, http://bit.ly/FR3RBR.

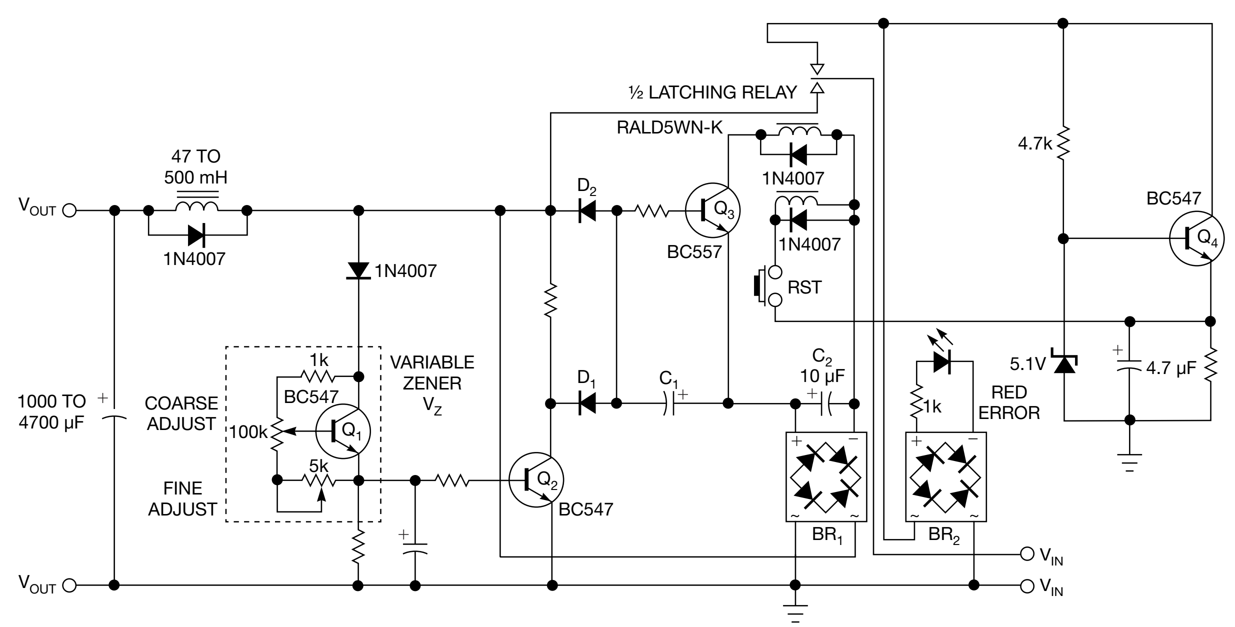

Figure 1 A bistable latching relay is the heart of this protection circuit that guards your load against excessive voltage and wrong polarity.

Originally published in the April 14, 1988, issue of *EDN*

Signal-powered switch connects devices

James Dean, National Research Council, Ottawa, ON, Canada

WANT TO SEE MORE OF THE CLASSICS?

Revisit 50 of the best Design Ideas from the Golden Age of electrical engineering.

http://bit.ly/DesignIdeasClassics

Figure 1's switch can selectively connect a terminal, personal computer, or any other originating device (DTE) to one of four receivers such as modems and printers. Using a thumbwheel switch, you select which device you want to connect. The circuit handles data rates to 19.2k baud without any signal degradation or crosstalk.

Four Maxim DG509A dual 4-channel multiplexers switch the eight critical RS-232C lines. The Maxim multiplexers suit this application because they consume only 20 μA typ, handle ±18V signals, exhibit only 130Ω on-resistance typ, and have input protection.

The unit doesn't require an external power supply because it gets its power from the RS-232C ready lines (DTE pins 4 and 20, DCE pins 5, 6, and 8); the SI7661 converter provides the negative supply. The Schottky power diodes in the dual-diode MBR2045CT packages ensure that the switch draws its power from whichever connected device has the highest output voltage at its RS-232C connector. The low forward-voltage drop of these diodes guarantees that the internal-protection diodes of the switch ICs do not conduct. The circuit's V+ supply is merely one Schottky-diode drop below the highest signal potential.EDN

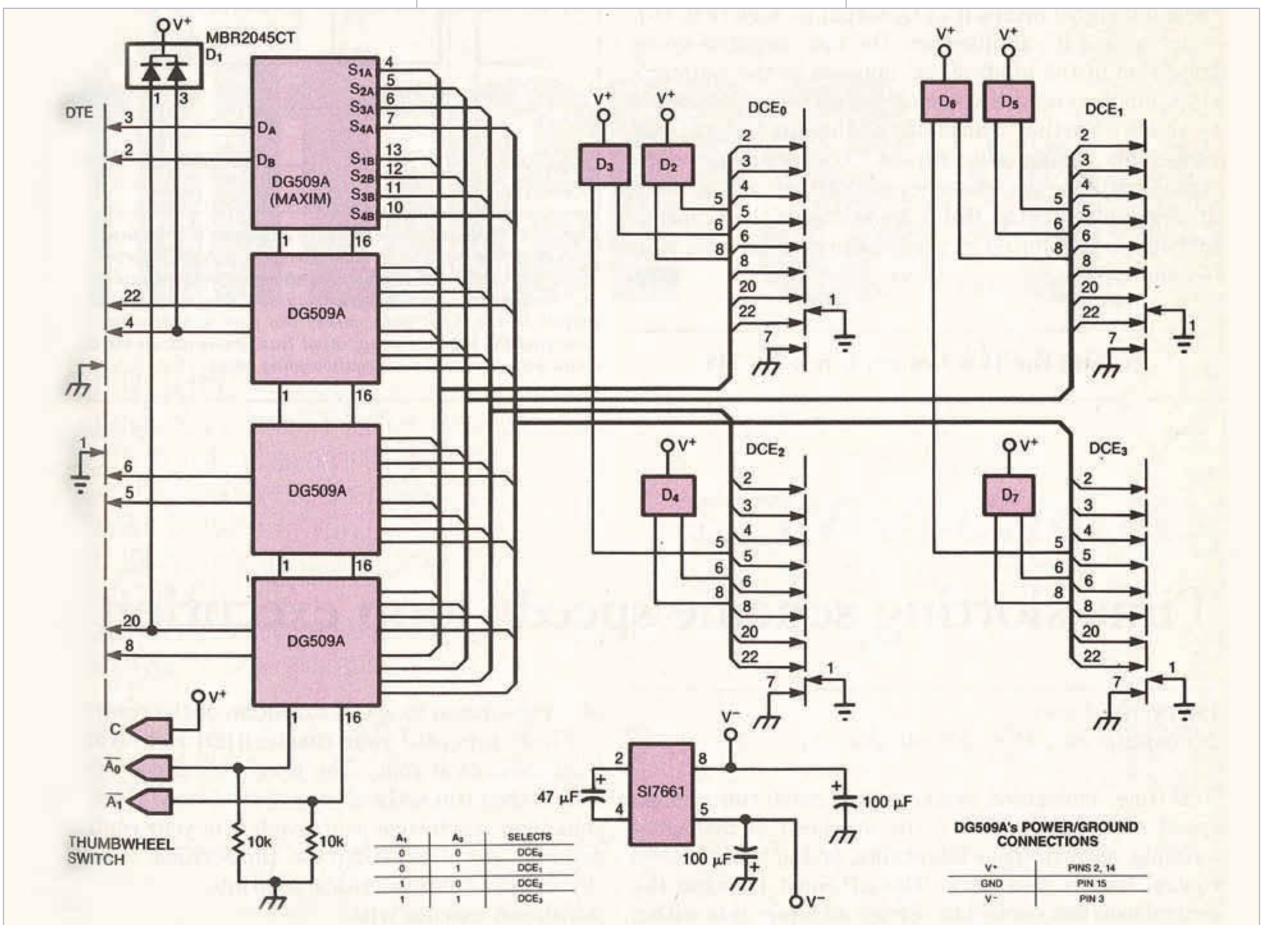

Figure 1 This signal-powered selector uses analog switches to connect one of four RS-232C devices to a terminal or personal computer.

productroundup

CONNECTORS

Harting angled RJ-45 connectors offer options with 45° cable outlet

The RJ industrial 10G angled connector offers insulation-displacement termination for flexible and rigid wires with profiles of AWG 27/7 to AWG 22/1. It allows the connection of cables with a diameter of 4.5 to 8 mm, and incorporates a 45° cable outlet that can be mounted in four directions. The connector offers secure and rapid connection and data-transfer rates as high as 1 and 10 GbE. The rugged and industry-standard IP20 model is also suitable for multiport RJ-45 jacks.

Harting Technology Group, www.harting.com

TE's zQSFP+ connector targets telecommunications

The high-density, high-speed zQSFP+ (z-quad small-form-factor pluggable plus) connector and cage assembly targets telecommunications, data-center, medical, network-interface, and test-and-measurement equipment. The device features an enhanced EMI cage, 1×1 single-port or 1×3 ganged cages with heat-sink and light-pipe options, and staggered press-fit pins. Data rates range from 25 to 40 Gbps, and the connectors meet 100-Gbps Ethernet and 100-Gbps 4X InfiniBand enhanced-data-rate requirements. Other specifications include a mating force of 40N; minimum durability of 250 cycles; an operating temperature range of −20 to +65°C; an electrical-current rating of 0.5A per pin; a maximum voltage rating of 30V dc; and a high-temperature, thermoplastic housing.

TE Connectivity, www.te.com

Phoenix Contact's connectors charge large dc-battery units

These dc-interchangeable connectors charge large dc-battery units in electric commercial vehicles. They are used in city and community transportation systems, transport vehicles in logistics environments, and people movers. The connectors handle currents as high as 400A, voltages as high as 750V, and 10,000 insertion and withdrawal cycles. An integrated data module monitors temperature to prevent overheating. Spring-damped guide bolts compensate for any tolerance when inserting the connector and dampen any vibration when the vehicle moves.

Phoenix Contact, www.phoenixcontact.com

Conec LC fiber-optic connectors comply with ODVA

The ODVA-compliant LC series of fiber-optic connectors finds use in harsh environments, including WiMax, LTE, and remote radio heads using fiber-to-the-antenna connectivity. The connectors use optical-fiber-nonconductive-riser breakout cables for outdoor usage or direct burial. The IP67-rated connectors feature a one-sixth-turn bayonet coupling for fast and secure mating and unmating, even when users are wearing gloves. The devices come in black-plastic and nickel-plated, die-cast housing materials. Prices start at $20 (OEM quantities).

Conec, www.conec.com

Amphenol Helios H4 is UL-certified to 1000V

The Helios H4 photovoltaic connector is UL-certified to 1000V, making it useful in current systems and new ones that require more voltage without an increase in cable size. The UL/TUV-rated connector features the vendor's RADSOK technology, offering higher current ratings and lower contact resistance. The ROHS-compliant connector uses a simple unlocking tool for easy unmating in the field. It is available in 14-gauge AWG at 32A, 12-gauge AWG at 40A, 10-gauge AWG at 44A, and 8-gauge AWG at 65A. Pricing is $3/mated pair.

Amphenol Industrial Global Operations, **www.amphenol-industrial.com**

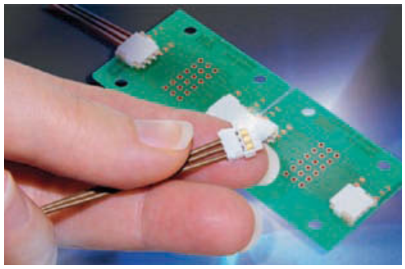

Hirose DF59 connector finds use in LED lighting

The DF59 multifunctional connector finds use in wire-to-board and board-to-board power-supply applications. It features contact spacing of 2 or 4 mm, according to the requirement for increased air or creepage path. The system handles a current of 3A at a wire gauge of AWG 22. The nominal voltage is 100V for 2-mm contact spacing or 230V for 4-mm spacing. The product is available with two, three, or four contacts and meets the requirements of monochrome and RGB lighting. The device has a profile of 2.5 mm, and the three-pin cable connector measures 8×9 mm, avoiding the need for shading.

Hirose Electric Europe BV, **www.hiroseeurope.com**

Molex SlimStack B8 features robust housing

The 0.4-mm-pitch, 2.5-mm-wide SlimStack B8 SMT board-to-board connectors target high-end medical, consumer-electronics, data-communications, and telecom mobile-device applications. They include the vendor's CleanPoint contact with a beveled shape for removing flux and other contaminants. The SlimStack assembly method features a top housing wall with a rugged metal cover that protects the terminal from damage due to forcible angled unmating. Sample prices start at 70 cents (one) from Mouser and other distributors.

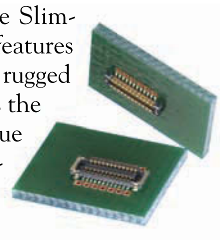

Molex, www.molex.com

EDN ADVERTISER

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

Long shots, short shots, and hip shots

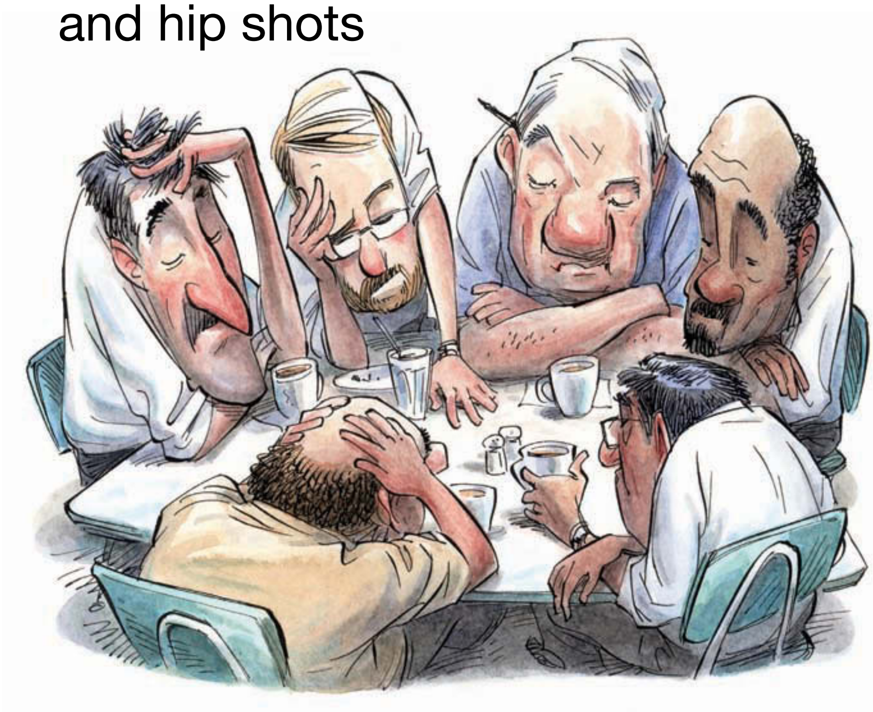

One morning, at an employer where I once worked, there were long faces on the manufacturing technicians and engineers sitting around the cafeteria table. Several weeks earlier, production of our new automatic-test systems had screeched to a halt when a simple self-test for the driver output's path resistance started intermittently failing its high limit by a few ohms. Since then, all of the involved driver cards, connectors, cables, and instrumentation tested OK separately, and we verified that the associated measurement code hadn't changed. The pressure to resume shipments was intense, and we had no clue what the problem was.

Initially, we suspected that it was the fault of the custom zero-insertion-force connectors that tied the driver modules to the outside world. However, a series of experiments ruled out the connectors as the culprits. We needed some new ideas; shooting from the hip wasn't working. We reviewed the status of the investigation and listed what we knew—which wasn't much. Yields on the new system had been fine for a couple of weeks, and, suddenly, nothing passed. Obviously, something had changed, but what?

A second search of engineering-change orders produced a clue: The manufacturing rollout of a software/firmware update matched the date when yield had gone to hell. Although there were no self-test-related changes in the code, the coincidence was too much to ignore. We returned a failing system to the previous revision of code, and the test failures vanished. The new systems still required the new software revision, but a simple code change should get things rolling again.

The initial look at the release was discouraging. There were hundreds of edits in the 300,000-line code base. We reverified that none were even remotely associated with the failing self-test routines. Rolling back these changes to see which had caused the failures could tie up software resources for weeks. We looked for something substantive that we could easily swap out. An FPGA-firmware download to the system's digital pattern sequencer fit the bill. It was a long shot because the sequencer wasn't running during the failing tests, but we were out of short shots.

To our surprise, this change eliminated the problem. We now had a much shorter list of code deltas to examine. Most of the FPGA updates were sequencer enhancements, but one modification enabled a 50-MHz clock that connected to another subsystem. Perhaps this clock was disrupting the measurement. Sure enough, the failures stopped when we again disabled the clock. One of the engineers noticed that holding the flat cable carrying the clock closer to the metal chassis reduced the incidence of failures—suggesting that RF interference was disrupting our analog measurements.

The failing test forced a known current through the driver-output path and a series-connected test resistor. The high side of the floating measurement ADC connected to one end of the driver path, and the low side connected to a sense circuit that provided a buffered version of the ground voltage near the grounded end of the test resistor. Referencing the ADC's low side to the ground near the test resistor decoupled most of the ground-return path from the measurement. The test determined the resistance of the unknown driver path plus the known test resistor.

The low-offset op amps buffering the ground-sense voltage had unity-gain bandwidths much lower than 50 MHz. When receiving the new 50-MHz clock signals on their inputs, they were rectifying instead of amplifying—generating an offset voltage that the system interpreted as additional resistance. Adding a lowpass RC filter on the ground-sense buffer input to attenuate the offending signals fixed the problem. **EDN**

Vance Harwood is a consultant in Loveland, CO.

DANIEL VASCONCELLOS