

EDN

VOICE OF THE ENGINEER

JUNE **25**
Issue 12/2009
www.edn.com



Think board to go broad Pg 14

EDN.comment: Can't see the forest ... Pg 6

Baker's Best Pg 16

Peering into ultrasound machines Pg 18

Design Ideas Pg 43

Tales from the Cube Pg 74

ACCELERATING CONSUMERS' NAS ADOPTIONS: ASSESSING YOUR PRODUCT OPTIONS

Page 30

DIAGNOSTIC ULTRASOUND GETS SMALLER, FASTER, AND MORE USEFUL

Page 21

LITTLE-KNOWN FLASH- MEMORY FEATURES PROTECT DATA AND IP

Page 39





COMPLEX DESIGN MADE LOGICAL

Cut development time as much as 50% while adding breakthrough performance, power and cost benefits. The ISE Design Suite 11 enables new domain-specific design for Targeted Design Platforms, allowing you to focus on building design differentiation. Download a free evaluation at www.xilinx.com/ise.



MORE SEMICONDUCTORS THAN EVER*



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

Authorized Distributor for more
than 400 Supplier Partners.

The industry's broadest product selection
available for immediate delivery

www.digikey.com
1.800.344.4539

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



Signal quality is just the beginning.

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

Function/Arbitrary Waveform Generator Series

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



**Add a FREE External Reference
worth up to \$495****

www.agilent.com/find/smartarb

Agilent Authorized Distributors

Allied Electronics

800-433-5700


Davis Instruments

800-358-5525

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




Agilent Technologies



EDN 6.25.09 contents

Accelerating consumers' NAS adoptions: assessing your product options

30 Network storage is a notably bright spot in the otherwise-blah consumer-electronics economy. Carefully select and cull hardware, software, and their jointly implemented features to ensure product success. *by Brian Dipert, Senior Technical Editor*



Diagnostic ultrasound gets smaller, faster, and more useful

21 The signal path in ultrasound machines is a multichannel transmitter-receiver system with blazingly fast data rates. Engineers need to weigh a host of options in designing these complex machines. *by Paul Rako, Technical Editor*

Little-known flash-memory features protect data and IP

39 Features from block locking to encrypted-password-access mechanisms can prevent unintentional disruption, malicious damage, or copying. *by Bill Stafford, Numonyx*

pulse

Dilbert 11

- 10** Scope-based tools support complete debugging and compliance testing of Superspeed USB 3.0
- 10** Analog Devices, Infineon collaborate on air-bag systems
- 11** 1- to 4-GHz, 20G-sample/sec DSOs and MSOs integrate protocol analysis


- 12** Nonvolatile-RAM module swaps batteries for super-capacitors
- 12** Chip allows isolated USB
- 13** POE ICs offer increased power efficiency for PDs
- 13** Tool bundles target FPGAs
- 14** **Voices:** Avnet's Jeff Ittel: Think board to go broad

A SPECIAL EDN SECTION INNOVATORS 55

- 57** Innovation: genius, practice, or luck?
- 59** Application engineers: serving the customer
- 67** Improving on PCB design



DESIGN IDEAS



- 43** Create a swept-sine function in LabView with just one virtual instrument
- 44** Charlieplexing at high duty cycle
- 48** Serial port tests digital circuits
- 51** DAC calibrates 4- to 20-mA output current
- 51** Alarm tells you to close the refrigerator door

The Newest Microcontrollers

8051 Cortex™ ARM®
Z8® Flexis™ AVR® PSoC® Encore!®
Coldfire® PIC® Z80® ST7
HC12 ST6 MSP430 HCS08
HCS12 MCU ST9

New Products from:



STM8S 8-bit Microcontroller
[www.mouser.com/
stm8smcus](http://www.mouser.com/stm8smcus)



TEXAS
INSTRUMENTS
Authorized Distributor


MSP430x5xx Microcontroller
Family
[www.mouser.com/
MSP430F5xx](http://www.mouser.com/MSP430F5xx)



AVR® XMEGA™
Microcontrollers
[www.mouser.com/
atmel_xmega](http://www.mouser.com/atmel_xmega)

The ONLY New Catalog Every 90 Days

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



a tti company

*The Newest Products
For Your Newest Designs*

(800) 346-6873



www.mouser.com

Over A Million Products Online



18



74

DEPARTMENTS & COLUMNS

- 6 **EDN.comment:** Can't see the forest ...
- 16 **Baker's Best:** Is your amplifier offset way out of whack?
- 18 **Prying Eyes:** Peering into ultrasound machines
- 72 **Product Roundup:** Motion, Integrated Circuits
- 74 **Tales from the Cube:** Weather or not: All signs point to "no"

EDN online contents

www.edn.com

ONLINE ONLY

Check out these Web-exclusive articles:

"In-design" physical verification is "on-time" physical verification

It's no longer practical to wait until the end of the physical design of an IC to do physical verification. Physical-design engineers who create chips at the 45-nm node and beyond face a difficult task. The time-tested flows used at previous nodes are no longer viable to maintain productivity at today's advanced nodes.

→ www.edn.com/article/CA6657281

RFIDs power themselves

RFIDs have enabled a significant improvement in many very costly business systems ranging from warehousing to tagging your luggage at the airport. The secure contactless-IC variants have found their way into our credit cards and passports, and their future application almost certainly includes the replacement of the checkout counter at the local grocery chain.

→ www.edn.com/article/CA6655988

MORE EDN INNOVATORS

Obstacles and opportunities

Once you finish reading the special Innovators section in this issue, head to the Web to read our earlier Innovators supplements that high-



light other technologists and their innovations in electronics design.
→ www.edn.com/innovators

ELECTRONIC NEWS TODAY

The latest tech news, delivered every afternoon.


→ www.edn.com/newsletters

EDN® (ISSN#0012-7515), (GST#123397457) is published biweekly, 24 times per year, by Reed Business Information, 8878 Barrons Blvd, Highlands Ranch, CO 80129-2345. Reed Business Information, a division of Reed Elsevier Inc, is located at 360 Park Avenue South, New York, NY 10010. Tad Smith, Chief Executive Officer; Mark Finkelstein, President, Boston Division. Periodicals postage paid at Littleton, CO 80126 and additional mailing offices. Circulation records are maintained at Reed Business Information, 8878 S Barrons Blvd, Highlands Ranch, CO 80129-2345. Telephone (303) 470-4445. POSTMASTER: Send address changes to EDN®, PO Box 7500, Highlands Ranch, CO 80163-7500. EDN® copyright 2009 by Reed Elsevier Inc. Rates for nonqualified subscriptions, including all issues: US, \$179.99 one year; Canada, \$229.99 one year (includes 7% GST, GST#123397457); Mexico, \$229.99 one year; air expedited, \$399.99 one year. Except for special issues where price changes are indicated, single copies are available for \$10 US and \$15 foreign. Publications Agreement No. 40685520. Return undeliverable Canadian addresses to: RCS International, Box 697 STN A, Windsor Ontario N9A 6N4. E-mail: Subsmail@ReedBusiness.com. Please address all subscription mail to EDN®, 8878 S Barrons Blvd, Highlands Ranch, CO 80129-2345. EDN® is a registered trademark of Reed Elsevier Properties Inc, used under license. A Reed Business Information Publication/Volume 54, Number 12 (Printed in USA).

PICO

SURFACEMOUNT
(AND THRU-HOLE)
**Transformers
& Inductors**

Size
does
matter!



from
low-
profile

.19"ht.

- **Audio Transformers**
- **Pulse Transformers**
- **DC-DC Converter Transformers**
- **MultiPlex Data Bus Transformers**
- **Power & EMI Inductors**

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

or send direct for free PICO Catalog
Call Toll Free 800 431-1064
Fax 914-738-8225

E Mail: info@picoelectronics.com

PICO Electronics, Inc.

143 Sparks Ave. Pelham, N.Y. 10803-1837



Delivery - Stock to one week



BY PAUL RAKO, TECHNICAL EDITOR

Can't see the forest ...

Analog design is all about trade-offs, so system-design expertise is vital to any company offering analog chips. For example, Dave Kress, director of technical marketing at Analog Devices, points out that his company is trying to get a better understanding of the entire system design so it can solve design problems that span multiple chips and subcircuits. Meanwhile, Doug Bailey, vice president of marketing

at Power Integrations, says that his company's customers expect help with thermal design and EMI (electromagnetic-interference) characterization, functions that were previously the domain of system engineers.

Texas Instruments is also dedicated to understanding system design to better serve its customers. Matt McKinney, marketing-communications manager, mentions that TI offers reference designs that solve the system-power and analog-signal-path problems of customers using the company's DSPs or embedded microprocessors.

Bob Thomas, technical leader at Cisco Systems, confirms the importance of vendors' having system-level cognizance. "A vendor shouldn't just toss a part on our desk and walk away," says Thomas. "The vendor has to get to know our system and its unique characteristics so [the vendor] can add something special to our designs." Steve Abe, a senior software engineer at the company, mentions that chip manufacturers should help develop the firmware to initialize the chips in Cisco's designs. "The vendors have an intimate understanding of the chips, and we look to them to help us ... keep the chip working in our system environment," he says.



"It's an example of vendors' not looking at the whole receiver-signal chain."

—John Scampini

John Scampini, director of strategic marketing at Maxim Integrated Products, points out that some vendors recommend amplifiers that would limit the SNR (signal-to-noise ratio) of subsequent analog-to-digital stages, meaning that the output-referred noise of the variable-gain amplifiers would be larger than the noise floor of the ADC that it was feeding. "It's an example of vendors' not looking at the whole receiver-signal chain," he says (see "Diagnostic ultrasound gets smaller, faster, and more useful," *EDN*, this issue, pg 21).

Jim Williams, staff scientist at Linear Technology, also discusses the changing role of application engineers (see "Application engineers: serving the customer," *EDN*, this issue, pg 59). Williams observes that modern application engineers are fundamentally systems engineers. "In a well-run analog company, [components are] second- or third-tier parts of the job," he says. "The fundamental responsibility of an analog-application engineer is to understand his customer's problem and provide a solution."

Semiconductor companies are taking an interest in systems design that is essential to continued innovation and success. For example, Qspeed Semiconductor's products are faster than other silicon diodes and cheaper than silicon-carbide diodes. Yet the Qspeed parts are three times more expensive than legacy slow-recovery silicon diodes. If Qspeed thought of itself as just a diode company, this approach would make it tough to sell parts into cost-sensitive applications. The company looks past the two pins of its diodes, however, and sees the total impact of better performance.

Andrew Smith, power-marketing manager at Power Integrations, points out that the company looks at the whole power system, designing controller chips that operate the PFC (power-factor-correction) section of the power system, as well as the primary isolated LC (inductor-capacitor) power stage for generating low-voltage system power.

Ed Lam, vice president of marketing and engineering at Advanced Analogic Technologies, is also looking at the entire system—from the wall plug to the final analog output. The components in this path might be the power in an LCD television, ending in a Class D audio amplifier, or the charging system in a cell phone. "The system is the thing," he says. **EDN**

Contact me at paul.rako@edn.com.


More Q. Less Cu

29

Copper
63.546

These tiny new air core inductors have the highest Q and current handling in the smallest footprint.

Coilcraft's new SQ air core inductors have unmatched Q factors: most are above 200 in the 1-2 GHz range! That's 3 times higher than comparably sized 0805 chip coils.




Q factors are 3X higher than standard chip inductors

And with their extremely low DCR, they can handle 4 to 8 times more current: up to 4.4 Arms.

SQ air core inductors are perfect for your LC filter and RF impedance matching applications. They come in 15 values ranging from 6 to 27.3 nH, all with 5% tolerance.

These coils are significantly smaller than existing air core inductors. We reduced the footprint by using close-wound construction and keeping the leads close to the body. The square shape cuts the height to as low as 1.5 mm and creates flat top and bottom sur-



faces for easy automated handling and stable mounting.



The square shape and narrow footprint reduce board space by 60-75% over conventional air core inductors.

See how the ultra-high Q and current handling of Coilcraft's new SQ air core inductors can maximize the performance of


your next design. For complete specifications and free evaluation samples, visit www.coilcraft.com/sq

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

Coilcraft

www.coilcraft.com 800/322-2645

Microsoft®



it's everybody's business
to get answers as **FAST**
as yesterday.

Does your enterprise software help give you the insight
to keep up with the pace of demand?

ask for *people_ready* enterprise high tech solutions

microsoft.com/hightech

EDN

PRESIDENT, BOSTON DIVISION, REED BUSINESS INFORMATION

Mark Finkelstein, mark.finkelstein@reedbusiness.com
1-781-734-8431

PUBLISHER, EDN WORLDWIDE

Russell E Pratt, 1-781-734-8417;
rpratt@reedbusiness.com

ASSOCIATE PUBLISHER, EDN WORLDWIDE

Judy Hayes, 1-925-736-7617;
judy.hayes@reedbusiness.com

VICE PRESIDENT, EDITORIAL DIRECTOR

Karen Field, 1-781-734-8188;
kfield@reedbusiness.com

EDITOR-IN-CHIEF, EDN WORLDWIDE

Rick Nelson, 1-781-734-8418;
rnelson@reedbusiness.com

EXECUTIVE EDITOR

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

MANAGING EDITOR

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

SENIOR ART DIRECTOR

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

ANALOG

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

MASS STORAGE, MULTIMEDIA, PCs, AND PERIPHERALS

Brian Dipert, Senior Technical Editor
1-916-760-0159;
fax: 1-303-265-3187;
bdipert@edn.com

MICROPROCESSORS, DSPs, AND TOOLS

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

NEWS

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

POWER SOURCES, ONLINE INITIATIVES

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

DESIGN IDEAS EDITOR

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

SENIOR ASSOCIATE EDITOR

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

EDITORIAL/WEB PRODUCTION

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

CONSULTING EDITOR

Jim Williams, Staff Scientist,
Linear Technology

CONTRIBUTING TECHNICAL EDITORS

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

COLUMNISTS

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

PRODUCTION

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

EDN EUROPE

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

EDN ASIA

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

EDN CHINA

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

EDN JAPAN

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



MAXIMUM

PRECISION

From a full spectrum of pin receptacles

Regardless of whether your components fly with the Blue Angels or control industrial robots, good design should never be limited by a lack of options. Mill-Max Mfg. Corp. offers hundreds of high-precision receptacle styles with an unprecedented range of options and features:

- 35 pre-tooled contact styles.
- A full selection of 3-, 4- and 6-finger contact designs.
- .008" to .102" (0,20 - 2,59 mm) pin acceptance.
- Accepts round, square and rectangular component leads.
- Styles for soldering, swaging or press-fitting.



Large Variety of Styles



Knock-out Bottom OFP®



Available on Tape and Reel

Mill-Max receptacles. Plug into the simple, reliable connectivity solution.

To view our Design Guide, new product offerings and order free samples, visit

www.mill-max.com/EDN585



 Reed Electronics Group

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

pulse

INNOVATIONS & INNOVATORS

Scope-based tools support complete debugging and compliance testing of Superspeed USB 3.0

Tektronix's new DPO (digital-phosphor-oscilloscope)/DSA (digital-signal-analyzer) 7000B family is now available with a USB-TX option, which, with the press of a button, allows you to automatically and rapidly validate USB (Universal Serial Bus) 3.0 Superspeed transmitters. These devices transmit data more than 10 times as fast as high-speed devices that conform to the previous (2.0) version of the USB standard. Tektronix has added the option—a comprehensive tool



With the addition of the new USB-TX option, one button press is all it takes for 7000B oscilloscopes to begin extensive validation testing on USB 3.0 Superspeed transmitters.

set for characterizing, debugging, and automating compliance testing of Superspeed devices—to its TekExpress oscilloscope-based serial-data-test framework, which also supports several other ultra-high-speed serial buses. In addition, the manufacturer has introduced a full set of USB 3.0 fixtures that enable you to perform transmitter, receiver, and cable tests.

The increased bandwidth of the Superspeed mode brings with it critical signal-fidelity challenges. Tektronix points out that, whereas other industry offerings provide only normative measurements in accordance with the USB-IF (Implementers' Forum) electrical-test specification, the USB-TX option supports all measurements, including normative and informative tests, such as SSC (spread-spectrum clocking), slew, and voltage levels. A unique plug-style fixture enables probing close to the silicon to provide a true representation of the signal. The list price for the USB-TX option is \$5000. Prices for the DPO/DSA 7000B scopes begin at \$53,300.

—by Dan Strassberg

► Tektronix Inc, www.tektronix.com.

FEEDBACK LOOP

“Brilliant anecdote. ... A classic case of everything is possible but not always feasible.”

—EDN reader and engineer “Harry,” in EDN’s Feedback Loop, at www.edn.com/article/CA6651606. Add your comments.

Analog Devices, Infineon collaborate on air-bag systems

Analog Devices and Infineon Technologies have collaborated to produce next-generation automotive-air-bag systems. In the immediate future, the two companies will produce a hardware platform using parts from each company’s current offerings. Designers will be able to adopt this platform to begin development of code using a well-defined software interface. Later, the two companies will align their plans to ensure interoperability.

Analog Devices brings its MEMS (microelectromechanical-system)-air-bag-sensor technology to the mix, and Infineon contributes most of the rest of the chip set, including sensor interfaces, microcontrollers, power-management chips, CAN (controller-area-network) and LIN (local-interconnect-network) transceivers, pressure sensors, and air-bag-squib-firing

drivers. Both companies are now offering the results of the collaboration and continuing to sell their product lines. A typical product-development cycle in this market is approximately two years, according to the companies. They claim that starting with their hardware platform could save six months of development time.

Analog Devices has shipped more than 500 million accelerometers and gyroscopes in more than 15 years and is the largest supplier of MEMS-based inertial sensors to the automotive industry. Infineon has supplied more than 600 million air-bag components in the last 15 years.—by Graham Prophet

► Analog Devices, www.analog.com.


► Infineon, www.infineon.com.

1- to 4-GHz, 20G-sample/sec DSOs and MSOs integrate protocol analysis

Agilent Technologies has expanded its MSO (mixed-signal-oscilloscope) and DSO (digital-storage-oscilloscope) portfolios with the addition of the six-model Infiniium 9000 series. All of the units offer true analog bandwidths of 1, 2.5, or 4 GHz on four channels and optional built-in logic and protocol analysis on 16 digital channels. All analog channels offer both 50Ω and 1-MΩ inputs and acquire a maximum of 20G samples/sec in the two-channel mode. In the

four-channel mode, the analog channels offer a standard memory depth of 10M samples, with options for as many as 1G samples in the two-channel mode. On any of the three analog-only models, you can in minutes enable the logic channels without returning the scopes to the factory or recalibrating them.

All models come in easy-to-carry, 9-in.-deep, 25-lb packages and feature screens that, according to the company, are the industry's largest: 15-in., 1024×768-pixel, touch-sen-

 The 15-in. displays offer 57% greater display area than do 12-in. displays.

sitive XGA (extended-graphics-array) LCDs.

The units are also the first scopes to integrate both logic-analysis channels and a protocol-analyzer viewer for buses such as PCIe (peripheral-component-interconnect express) and USB (Universal Serial Bus), according to the company. These features simplify setting up and using the instruments to test complex designs, encouraging greater instrument usage and increased return on investment.

Regardless of the scope's analog bandwidth, the integrated digital channels acquire digital signals at rates as high as 2G samples/sec and trigger on a rich set of mixed analog/digital conditions for accurate analysis of timing relationships among control signals and data buses. A protocol viewer for PCIe and USB lets you extend the scope's debugging and testing capabilities to obtain rapid insights into the behav-


ior of protocol-rich serial buses without finding and connecting additional test equipment.

The digital inputs accept many Agilent logic-analyzer accessories. In addition, general-purpose measurement and analysis features allow you to customize the instruments by adding software. To provide insight into common serial buses, FPGAs, and RF measurements, 25 optional application packages allow for rapid verification of a unit under test's compliance with industry standards.

Debugging software includes protocol triggering and viewing for PCIe and USB; serial decoding and triggering for I²C (inter-integrated-circuit), SPI (serial-peripheral-interface), CAN (controller-area-network), RS-232, and UART (universal-asynchronous-receiver/transmitter) interfaces; rapid core-assisted debugging of designs with Xilinx (www.xilinx.com) and Altera (www.altera.com) FPGAs; and an InfiniiScan capability for triggering on infrequent signal anomalies that meet visual criteria. Compliance-testing features include communication-mask testing and compatibility with DDR-1, -2, and -3 interconnects; Ethernet; and USB 2.0.


The 15-in. displays offer 57% greater display area than do 12-in. displays. The larger displays provide the space users need to comfortably view large numbers of digital and serial signals in addition to analog traces. Infiniium 9000 series base prices range from \$19,900 for a 1-GHz-bandwidth DSO to \$41,500 for a 4-GHz MSO.

—by Dan Strassberg
 > Agilent Technologies,
www.agilent.com/find/9000.



The 20G-sample/sec Infiniium 9000 scope's 15-in., touch-sensitive XGA LCD proves helpful in mixed-signal debugging, letting you view 16 logic-analysis channels and four analog waveforms.


DILBERT By Scott Adams



Nonvolatile-RAM module swaps batteries for supercapacitors

Systems often need some form of nonvolatile RAM to protect against memory losses after a power interruption or loss. Battery-backed RAM has been the most common form of high-density nonvolatile memory, but it is bulky, involves long charging times for the battery, has a relatively low-temperature operating environment that requires low-temperature operation, and has a limited number of cycles.

AgigaTech's new AgigaRAM nonvolatile-system technology addresses these issues by combining synchronous DRAM with flash memory. The approach provides memory capacity of 4 Mbytes to



The Bali version of the AgigaRAM nonvolatile-memory system combines synchronous DRAM with flash memory. If the system loses power, the module automatically saves the SDRAM data to NAND-flash memory using the energy it stores in its internal supercapacitors.

2 Gbytes with transfer rates equivalent to those of DRAMs. The AgigaRAM relies on its

SDRAM during normal operation, but, if the system loses power, the device automati-

cally saves the SDRAM data to NAND-flash memory using the energy it stores in its internal supercapacitors.

The AgigaRAM product encompasses the Bali and Capri product families. Bali comes in 4- to 64-Mbyte densities, incorporates 100-MHz SDRAM with 200-Mbyte/sec peak transfers, requires a 3.3V power supply for 4 to 32 Mbytes or 5V for 64 Mbytes, and operates over a 0 to 70°C operating range. It comes in a 200-pin SO-DIMM package or a mezzanine card and sells for approximately \$4 per megabyte. The Capri comes in 256-Mbyte to 2-Gbyte capacities, uses a high-speed DDR-800 interface, and sells for less than \$1 per megabyte.

—by Margery Conner

► **Agiga Tech**, www.agigatech.com

CHIP ALLOWS ISOLATED USB

Analog Devices' new ADuM4160 bidirectional isolator provides UL (Underwriters Laboratories) 1577-rated, 5-kV isolation for 1.5- or 12-Mbps USB (Universal Serial Bus)-data lines. The isolator has propagation delay comparable with that of a standard hub and cable. It operates at 250V rms for IEC (International Electrotechnical Commission) 60602-1 CSA (Canadian Standards Association), 600V rms under IEC 60950-1 CSA, and 846V peak for VDE (Verband der Elektrotechnik, Elektronik und Informationstechnik) certification. IEC 60601-1 medical-safety approvals are pending.

The device operates from the 5V USB power supply



You can create a fully isolated USB 2.0 interface by combining the ADuM4160 data isolator with the ADuM5000 dc/dc converter.

or 3.3V power and provides isolated control of the pullup resistor, allowing the peripheral to control connection timing. The maximum upstream-supply

current is 7 mA at 1.5-Mbps data rates and 8 mA at 12-Mbps rates; idle current is 2.5 mA. The unit operates over a -40 to +105°C temperature range.

You can implement a fully isolated USB interface by combining the ADuM4160 data isolator with the ADuM5000 isolated dc/dc-converter IC. Both parts use Analog Devices' iCoupler isolation technology, which employs on-chip transformers made during die metallization, along with a Kapton polyimide film, to provide the voltage isolation. The high-quality isolation makes the parts suitable for medical, scientific, and industrial applications.

The ADuM4160 is available in a 16-pin, wide-body SOIC with a suggested retail price of \$4.89 (1000). ADuM4160EBZ evaluation modules are also available.—by Paul Rako
► **Analog Devices**, www.analog.com

POE ICs offer increased power efficiency for PDs

Akros Silicon cannily plays the efficiency card—for both energy and space—with its new AS1854 POE (power-over-Ethernet) PD (powered-device) IC. Akros claims that the integration of parts into the 1854 PD-controller platform saves more than 75% in board space and 25% in BOM (bill-of-materials) costs. The chip enables higher power efficiency with its GreenEdge synchronous rectification, which the company claims can yield greater-than-8% improvement in efficiency in the 6 to 8W “sweet-spot” range for VOIP (voice-over-Internet Protocol) phones, currently the most common application for POE. GreenEdge digital-isolation technology also provides 2-kV on-chip isolation, eliminating the need for optoisolators that reduce system reliability and add cost and

space. PDs require isolation between the power-management circuitry for the system devices, such as VOIP-phone electronics and display, and the primary-side electronics and sensor-control circuits.

POE systems require “master” PSE (power-sourcing equipment) that provides power to networked equipment through an Ethernet cable. For an IEEE 802.3at draft version, the PSE must allocate as much as 25.5W per PD and 4.5W of cable loss. The AS1854’s PSE can determine the PD’s power usage and cable losses. In this way, you need not rely on a worst-case power budget. Budgetary pricing for the AS1854 is \$4.69.

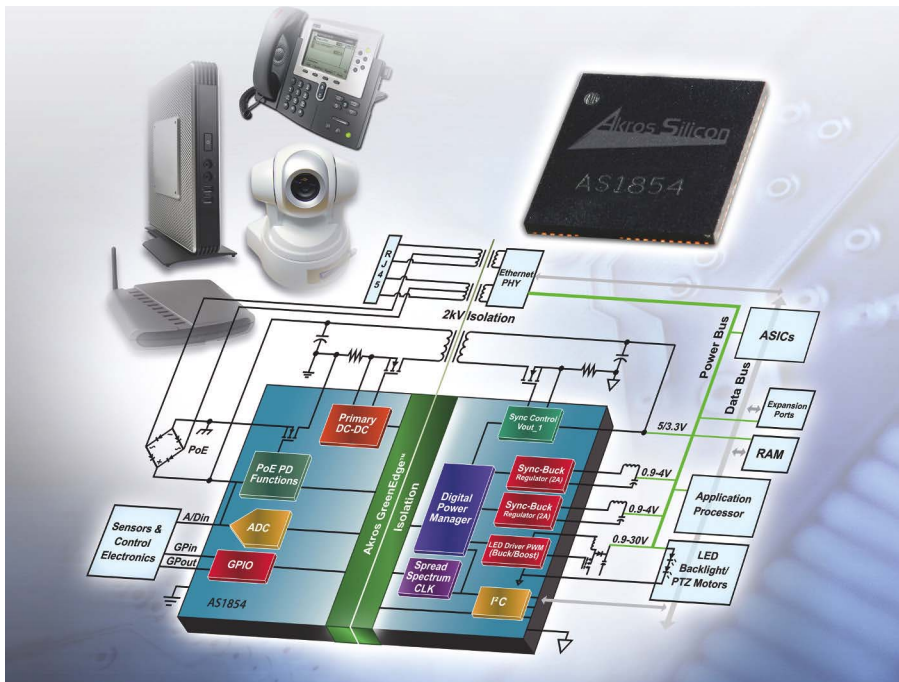
Texas Instruments has also introduced a family of three devices supporting POE PDs. The TPS2375X relies on external optoelectronic isolators

rather than on-chip isolation. The TPS23754 supports 13 or 26W PD applications and provides dc/dc-converter topologies that can achieve more-than-90% power-conversion efficiency. The TPS23756 enables PDs to accept auxiliary input voltages as low as 12V, allowing the use of popular and inexpensive 12V wall adapters. The TPS23757 supports less-than-13W systems that require the higher-efficiency dc/dc topologies. The TPS23754 is available now, and the TPS23756 will be available in the third quarter. Each comes in an HTSSOP-20 package for \$1.90 (1000). The TPS23757, also available in the third quarter, will sell for \$1.65 (1000).

—by Margery Conner

► **Akros Silicon**, www.akros-silicon.com.

► **Texas Instruments**, www.ti.com.



The AS1854 POE platform chip from Akros Silicon includes on-chip 2-kV isolation.

TOOL BUNDLES TARGET FPGAs

Xilinx this year introduced two families of Virtex-6 and Spartan-6 FPGAs that enable targeted-design platforms. The company is now adding domain-specific-tool support to the targeted-platform concept. Xilinx’ ISE (integrated software environment) Design Suite 11.1 offers interoperable domain-specific design flows and tool configurations for logic, DSP, embedded processing, and system-level design. Each of the four editions of the design-tool suite includes a bundle of tools that will match different FPGA users’ needs.

For example, the logic edition—for classic logic and connectivity-function design—includes the ISE simulator, the PlanAhead design-and-analysis tool, the ChipScope debugging- and serial-I/O-tool kit, and a catalog of IP (intellectual property). Similarly, the DSP edition includes Xilinx’s System Generator for DSP, plus related tools and DSP-specific IP. The embedded edition supports both hardware and software designers, and its IP includes the MicroBlaze soft processor. Xilinx licenses the packages using the Flexnet system, and all packages come configured for global use over a network. For more information on these products, go to www.edn.com/090625pa.

—by Graham Prophet

► **Xilinx**, www.xilinx.com.

VOICES

Avnet's Jeff Ittel: Think board to go broad

Jeff Ittel serves as senior vice president of business development and marketing at Avnet Electronics Marketing Americas, a division of components distributor Avnet Inc (www.avnet.com). Previously, Ittel served for two years as president of the former Avnet Cilicon division of Avnet Electronics Marketing, where he and his management team led successful efforts throughout a previous industry downturn. Ittel spoke with *EDN* about the changing role of distributors and the support they can offer to board-level designing engineers and suppliers. Excerpts of that interview follow. For the full interview, go to www.edn.com090625pb.

An Avnet exec once told me that “designing across the board” was the biggest opportunity for distributors. Do you agree?

A Yes, I do. As a distributor, we see across a whole technology or commodity, so we've developed some expertise in which products might optimize certain applications in certain uses. Probably one of the biggest things we see more and more now is that we are not just asked to go in and design in a chip. Customers want [to know] how everything else works together. We've been asked and are required to provide more overall solutions. That's all from a design point. From a supply-chain point, if the product is designed-in off our line card, it's much easier to support from our vantage point, whichever way that customer wants to be supported in the supply chain or how they set it up. Supply chains are getting more and more complex, and your design decisions affect your supply-chain decisions more than ever these days.

Why is that?

A You've got to have product that is going to be around for a while for manufacturability—for lifetime. The guys that are designing in one room and building in another room are becoming fewer and farther between. They are outsourcing their manufacturing. It's easy if they are outsourcing in the same city to a customer we already call on; it's a little more difficult when they are outsourcing out of their state or region. More and more is getting outsourced overseas. And most [customers] have a combination of all of these [scenarios]. They want to be able to pick someone who can support the supply chain and move wherever they want to go and keep the support ongoing.

Can you share an example of designing across the full board?

A Lighting is a good example of a solution that is much more than having a neat, high-tech LED that is the brightest one on the mar-



ket. You've got to have drivers, you've got to have thermal management, and you've got to have optics on virtually every single application. To be able to deliver the product and our ability to have expertise in each one of those areas, coupled with leading suppliers in each one of those areas, allows us to bring a solution to an engineer that is more than just the latest high-brightness LED. ... It's an exciting area that requires system-level knowledge.

There's now a \$10 million stimulus prize for the first company or engineer that can design a screw-in LED-based replacement for the standard Edison 60W light bulb. The new bulb would have to be dimmable and include LEDs, as well as driver circuitry, and have efficacy greater than 90 lumens per watt. This is no easy feat.

A It's a great recognition just for the future of LEDs for somebody to put that much money out there. It says that this is a real technology, and we encourage you to try to adapt to it. For us, it's also very clear that, even in your description, it's not as simple as just plugging in an LED. It is much more complex. It is not just a component sale. It is a system sell to be able to deliver the end state that people are looking for there.

Again, it's not only a latest-technology, high-end LED, but drivers, thermal management, optics, everything involved to be able to deliver the concept there. For us, that's a really nice testimonial to the fact that system knowledge and innovative technologies are going to be two key factors, and those are two things we think we bring to the market.

While we are on the topic of power-saving LEDs, why is “green” more than a buzzword for components distributors?

A When you look at the way industries have been forced to go with ROHS [restriction-of-hazardous-substances] directives and everything, the new technologies that come out are more energy-efficient; they are more environmentally sustainable and cost-effective. When you look at that [scenario], ... everybody wins. There's less throw-away [product], and product life is lengthened. Our customers' customers, especially if they work with government, are telling us they are required to start designing more green, and that [trend] is not just smaller, faster, and cheaper; it's all the different things I just mentioned: energy efficient, sustainable, cost-effective, that type of thing. Customers really have to keep up on it. There's a lot of new technology coming out. They can stick with what they know and try to figure this all out, or they can reach out to people who have maybe more marketing insights and technical training and that can offer support. That's an area we like to think is one of our places in the whole supply and design chain.

—interview conducted and edited by Suzanne Deffree

Rarely Asked Questions

Strange stories from the call logs of Analog Devices

What shall we do with an unused op-amp?

(Sung to the tune of "What shall we do with the drunken sailor?")

Q. In the RAQ on op-amps as comparators you mentioned an unused op-amp in a quad as being possible justification for using it as a comparator. What do you do if you've an unused op-amp and don't need a comparator?

A. This is a trickier than it looks. If an op-amp is overdriven, the output stage will saturate at one of the supply rails, and the op-amp will consume excess power. Many common configurations of an unused op-amp will overdrive it.

If the terminals are all left unconnected, there is a real risk that stray electrostatic fields will cause an input to go outside the supply rails. This can cause latch-up and destroy the whole chip. Even if latch-up does not happen, a dc field may cause saturation and power waste. In addition, the amplifier may amplify an ac field and, if overdriven, will heavily modulate its own supply current and cause crosstalk to other amplifier(s) on the chip.

Some users connect one input to the positive supply and the other input to the negative supply. This again saturates the output and wastes power; it may also exceed the differential input voltage rating and damage the device. Even if damage does not occur, some input stages draw several tens of milliamps under these conditions, wasting even more power.

Grounding both inputs, or shorting them together at some other potential, also causes the output stage to saturate, since the offset voltage of an op-amp is never exactly zero; shorting them together and not biasing them has the same latch-up risks that we have already mentioned.



What we should do is connect the device as a follower (output to inverting input) and connect the non-inverting input to a potential somewhere between the supply rails. With a dual-supply system, ground is ideal, but connecting to the positive or negative supply of a single supply system will cause saturation and the resulting power waste if the offset voltage has the wrong polarity. The "potential somewhere between the supply rails" may be any point in the circuit with a suitable potential, since the loading caused by the op-amp input is minimal. For diagrams see the linked article.


Or you might use it as a buffer amplifier in a part of your system that does not need one but might perform slightly better if it had one.

*What shall we do with the unused op-amp? (X3)
Early in the morning.*

*Hook-up as a buffer with a dc input, (X3)
Early in the morning.*

**To Learn More About
Op-Amps as Comparators**

<http://designnews.hotims.com/23107-101>



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

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

For Analog Devices' Technical Support, Call 800-AnalogD

SPONSORED BY





BY BONNIE BAKER



Is your amplifier offset way out of whack?

Have you ever spent a great deal of time selecting the perfect operational amplifier for your circuit, only to find that the offset voltage is wrong at the manufacturer's bench-specified input? What if you find that it is more than 10 times higher than specification in your application circuit? Do you send the chip in for failure analysis or just toss the chip out and have another look at your list of amplifiers? As an alternative, I suggest that you try to explain the offset error by re-examining your amplifier's specifications.

If you are using your amplifier as the key component in a transimpedance amplifier, an analog filter, a sample-and-hold circuit, an integrator, a capacitance transducer, or any other circuit with high-impedance components around your amplifier, you might find that the amplifier's input-bias current creates an offset-voltage error through the resistors in your circuit.

In the bipolar-amplifier days, the term "input-bias current" was an accurate descriptor, and it still is. A bipolar amplifier's input-bias current is the

same as the base current of the NPN or PNP transistors at the input of the amplifier. The magnitude of the bipolar amplifier's input-bias current ranges from a few nanoamperes for low-power devices to hundreds of nanoamperes for higher-power devices.

The term "input-bias current" loses its meaning when you look at JFET or CMOS input amplifiers. With these types of amplifiers, the current sinking or sourcing from the amplifier's input pins is actually the leakage current from the input-ESD (electrostatic-dis-

charge) cells (Figure 1). A more accurate descriptor for this current error is "input-leakage current." The magnitude of leakage current with JFET or CMOS amplifiers is less than 1 pA at 25°C. This specification is independent of the common-mode voltage and the magnitude-amplifier power. Almost all amplifiers have ESD cells for protection from an ESD event, but you will never see ESD-leakage current in bipolar amplifiers. The input-bias current swamps out the picoampere leakage current from the ESD cells.

Input-bias and input-leakage current can change over temperature. However, depending on the operational-amplifier design, the bipolar input-bias current can be fairly stable. The JFET and CMOS input amplifiers may not be, however. Because the leakage current is from the reverse-biased ESD diodes, the leakage current increases approximately two times per 10°C change.

In ensuring that the input-leakage current remains low with JFET and CMOS amplifiers, you must understand the impact of your PCB (printed-circuit board) on the picoampere levels of current. For instance, a small amount of dust, oil, or water molecules can increase leakage current and masquerade as input-bias current. The good news is that, if you exercise special care, you can build a PCB that will adhere to a 1-pA performance specification.

The most effective way you can reduce or minimize the effects of input-bias or input-leakage current is to check your circuit configurations. As you examine your circuits, look at the voltage characteristics of each node and make sure that you understand the impact of all of the current paths in your circuit. **EDN**

Bonnie Baker is a senior applications engineer at Texas Instruments and author of A Baker's Dozen: Real Analog Solutions for Digital Designers. You can reach her at bonnie@ti.com.

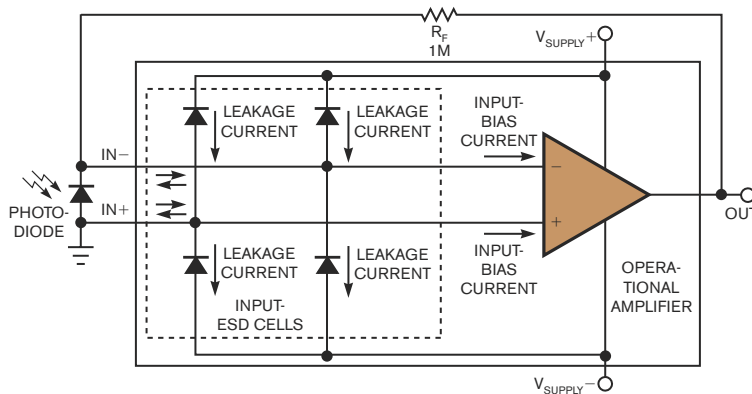


Figure 1 Input-bias or -leakage current creates a voltage drop across R_f .

TAPE-OUT COMES A LOT FASTER WITH CALIBRE nmDRC.



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

Calibre® nmDRC | There's nothing like having an advantage when you're racing to market. That's exactly what you get with Calibre nmDRC. Widely recognized as the world's most popular physical verification solution, Calibre's hyperscaling architecture produces the fastest run times available. To accelerate things even more, Calibre nmDRC adds incremental verification and a dynamic results-viewing and debugging environment. Designers can check, fix and re-verify DRC violations in parallel, dramatically reducing total cycle time. Start out and stay ahead of the pack. Go to mentor.com/go/calibre_nmDRC or call us at 800.547.3000.

**Mentor
Graphics®**
THE EDA TECHNOLOGY LEADER

+ Read an expanded version of this article at www.edn.com/090625pry.

+ Go to www.edn.com/pryingeyes for past Prying Eyes write-ups.

Peering into ultrasound machines

Ultrasound finds use in noninvasive imaging in cardiac, obstetric, gynecologic, and other diagnostic areas. What's inside this tool that is playing an expanding role in today's medical world?

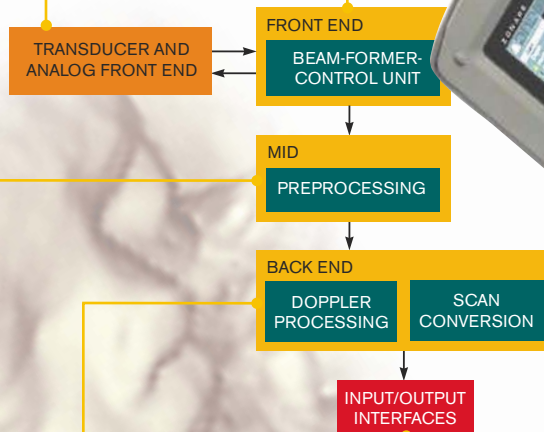
The systems often operate in the 2- to 20-MHz frequency range. An ultrasound system transmits a phased array of sound waves through a linear transducer so that the waves constructively combine at a focal point. As the generated sound waves propagate toward the focal point, they undergo a slight change in direction and produce a reflected sound wave each time they cross through matter of different densities. A variable controlled amplifier scales the reflected sound waves before an ADC samples and passes the data to the front-end processing of the system.

The front-end processing controls and performs beam-forming, which involves steering and focusing the phased array of sound waves. Steering involves sweeping the angle and direction of the beam to focal points with a precomputed depth. Focusing involves exciting multiple piezoelectric elements in the transducer with precisely time-delayed pulses so that the sound waves converge at each focal point along a scan line. The receiver's beam-forming ability detects the time delay, phase, and amplitude of the reflected sound wave at each focal-point location to reconstruct the flight path of the waves using a delay and sum algorithm to support creating an ultrasound image.

The midprocessing end lacks a common definition, but it can include filter, detection, and compression processing on a scan line of beam-formed data. The filter processing is typically bandpass filtering to reduce noise. The detection processing extracts a signal from the envelope of the signal. A complex rotator demodulates the signal in baseband; lowpass filtering then eliminate side lobes. The system may perform additional lowpass filtering with decimation or interpolation before presenting this data for back-end processing.

The back-end processing focuses on forming quality images for display from the received data. A scan conversion interpolates raw-data coordinates to displayed-data coordinates. The raw data can be in Cartesian coordinates for linear probes or polar coordinates for curvilinear or phased-array probes. Frame-smoothing techniques reduce noise without blurring the image's edges. Edge-detection techniques can help remove this blurring. Doppler processing focuses on measuring and displaying shifts and motion of structures, such as blood flow, in the data. A CW (continuous-wave)-Doppler system is analog and is highly sensitive and selective so that it can estimate velocities.

Display-mode-dependent processing is vendor-specific and plays a critical role in producing images on the display. This type of processing can involve the combination of overlays, such as color Doppler and color-flow imaging. It can also include support for the user interface, including menus, help, and display options. Images courtesy Zonare Medical Systems (www.zonare.com); for more information about ultrasound, see "Diagnostic ultrasound gets smaller, faster, and more useful," *EDN*, this issue, pg 21.



Falas MATLAB?

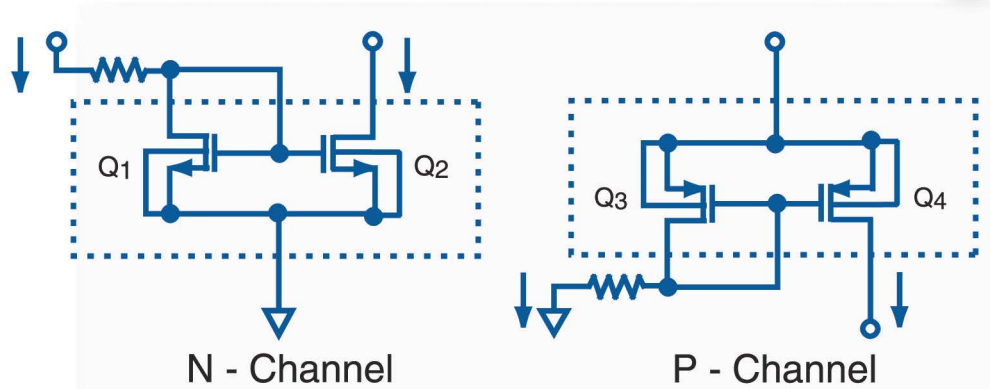
Over one million people around
the world speak MATLAB.
Engineers and scientists in every field
from aerospace and semiconductors
to biotech, financial services,
and earth and ocean sciences
use it to express their ideas.
Do you speak MATLAB?

*Electrons gain 42 billion
electron volts in a plasma
wakefield accelerator.
Provided by Stanford Linear
Accelerator Center.*

Read more at mathworks.com/ltc

MATLAB[®]
The language of technical computing

Small Signal Analog



Zero Threshold™ / NanoPower MOSFET ARRAYS

Enhancement • Matched Pair • Depletion • Complementary

ALD EPAD® Enabled MOSFET Arrays

convincingly demonstrate another innovative technology breakthrough for small signal analog circuit designs. These unique devices can be operated at *extremely* low voltage and current levels, ideal for applications in many battery or alternative-energy powered systems, including differential amplifiers, current mirrors, current sources, synchronous rectifiers, RF oscillators and normally closed switches.

Other innovative ALD precision products

- High efficiency Energy Harvesting™ modules
- Low charge injection analog switches
- Low drift CMOS timers
- Voltage comparators with high current drivers
- Ultra low input signal rail to rail CMOS op amps
- High accuracy dual-slope integrating analog and digital processors
- $\pm 5 \frac{1}{2}$ digit DVM boards

Performance features include:

- High Input impedance & DC current gain
- Low operating voltages from 0.03V to 10V
- Threshold offset voltage 1-2 mV typical
- Low input capacitance C_{ISS} 2.5 pF typical
- nA to mA output currents



In
stock
at

MOUSER
ELECTRONICS

a tti company


www.mouser.com

For more information, go direct:
www.aldinc.com/ssa81.htm



ADVANCED
LINEAR
DEVICES, INC.

415 Tasman Drive, Sunnyvale, CA 94089-1706
Tel: 408-747-1155 Fax: 408-747-1286
www.aldinc.com



BY PAUL RAKO • TECHNICAL EDITOR

DIAGNOSTIC ULTRASOUND GETS SMALLER, FASTER, AND MORE USEFUL

THE SIGNAL PATH IN ULTRASOUND MACHINES IS A MULTICHANNEL TRANSMITTER-RECEIVER SYSTEM WITH BLAZINGLY FAST DATA RATES. ENGINEERS NEED TO WEIGH A HOST OF OPTIONS IN DESIGNING THESE COMPLEX MACHINES.

Developers of diagnostic-ultrasound machines must carefully consider these devices' design and their intended applications, making trade-offs among such factors as SNR (signal-to-noise ratio), channel count, selection of ADCs, Doppler versus conventional technology, PW (pulsed-wave) versus CW (continuous-wave) approaches, power consumption, cable selection, and cost. With advancements in electronics, smaller and faster machines are emerging that are applicable in a variety of new applications. Portable and 3- and 4-D machines are also making inroads in this expanding field.

The diagnostic-ultrasound technique creates images of organs and measures blood flow within patients' bodies (Figure 1). The technology is now also finding use in therapeutic applications, such as targeted drug delivery. In this scenario, the machine emits a pulse waveform that couples with microspheres in a patient's blood; the pulses break these

spheres, accurately delivering drugs to their intended targets. Other applications include chemopotential, which helps chemotherapy drugs find and destroy cancerous tumors, and veterinary medicine, in which patients cannot verbally describe symptoms.

You can divide diagnostic ultrasound into two broad categories: convention-

This output of a CW-Doppler system (above left) does not show an image of your body. It represents blood flow over time. You can measure maximum velocity and observe the negative periods of blood flow as arterial valves close. A cardiogram signal is superimposed below the ultrasound signal (courtesy Analog Devices). By using a 2-D transducer array (above right) or wobbling a linear array with a stepper motor, you can derive a 3-D ultrasonic image (courtesy Texas Instruments).



al and Doppler. Conventional ultrasound, also known as B (body) mode or B (beam) scan, represents the classic use of technology: to peer into organs of the body, perhaps to determine the sex of a fetus. Doppler ultrasound relies on the Doppler effect, which takes its name from Austrian mathematician Christian Doppler, who proposed it in 1842. The Doppler effect describes the change in the frequency of sound waves for an observer relative to the source of the waves. You can imagine this effect when you think of the sound an ambulance siren makes as it approaches you, when it is right next to you, and as it moves away from you. PW, or color, Doppler encodes velocity information as colors on a display. CW, or spectral, Doppler broadcasts a continuous wave of sound into the body and measures the phase shift of the returning signals. This technique yields only velocity information, with no profile of where the velocities occur.

All ultrasound systems are able to use beam-forming and beam-steering approaches—signal-processing techniques for directional-signal transmission or reception. Physicians can employ PW-Doppler technology to map a patient's interior organs and blood flow, for example. CW systems must be aimed at a certain point because they have no way of knowing where the blood is flowing.

AT A GLANCE

Physicians use ultrasound therapeutically for tasks such as breaking up kidney stones and diagnostically for peering into patients' bodies.

You can divide diagnostic ultrasound into conventional B (body)-wave, PW (pulsed-wave) Doppler, and CW (continuous-wave) Doppler.

The transmitting path applies digital pulses or DACs to high-voltage amplifiers.

The receiving signal path is similar to that of a phased-array radar system.

Spatial resolution is a function of the channel count and the SNR (signal-to-noise ratio) of the signal path.

For this reason, physicians typically use CW Doppler as an add-on to a conventional ultrasound-signal path. The physician can find an artery or a vein with conventional ultrasound and then closely examine its blood flow with the high resolution of CW techniques.

Doppler systems employing beam-steering use linear arrays of piezoelectric transducers that steer the beams and provide for increased spatial resolution. By delaying the pulses on the center of a linear array, they can focus the trans-

mitting pulse on one point in the body (Reference 1). Applying delay to the receiving side makes the reflected waves adhere to that point.

Ultrasound machines can have 16 to 512 channels; those with fewer channels may be portable, whereas you may need wheeled carts to move larger machines with many channels. Having many signal channels affects not only the cost and size but also the power consumption of these machines.

Beware of manufacturers' channel-count claims, however, because they confuse some would-be purchasers. Some manufacturers use signal-processing techniques to increase the apparent channel count and then claim double the number of analog channels in the front end. Others multiplex the analog channels across a wider-channel-count transducer. For example, they may switch a 16-channel-wide analog-input section across a 64-channel transducer, perhaps moving one element at a time to "paint" an entire beam-formed area across the transducer length. Some manufacturers also provide more transmitters than receivers, an approach that can provide acceptable but less accurate results than using one analog channel for each transducer element. To increase spatial resolution, you can use more analog channels than elements and focus between elements, according to Lee D Dunbar, vice president of market innovation at SonoSite, a maker of portable ultrasound equipment.

A recent development in this market is 3-D ultrasound, in which the transducer is a 2-D, rather than a linear, array. This approach allows the system software to create a 3-D view into a patient's body. (For more on 3-D ultrasound, see "Peering into ultrasound machines," *EDN*, this issue, pg 18.) A 4-D ultrasound adds a time dimension to a 3-D image. It can show a moving image of organs and structures inside the body.

Ultrasound images also rely on the SNR of the analog-signal path. If there is less dynamic range in the signal chain, the ultrasound system cannot discriminate among the reflected signals. The analog chips' high noise floor would mean that conventional ultrasound could not peer as far into the body and would not resolve small features, whereas CW-Doppler ultrasound would not



Figure 1 Portable ultrasound machines allow doctors to get early information about trauma and emergency conditions. US Army medics donated this SonoSite machine to the Hawija Hospital in Iraq. Major Charles Buck of the 25th Infantry Division shows an Iraqi physician how to use the machine (courtesy Sergeant Sean Kimmons, US Army).

pick up slow or subtle variations in blood flow. By trading off channel count versus SNR, however, designers can provide the system software with enough information to make accurate images.

As in most analog-signal paths, the first amplifier in the receiver path controls SNR. You can increase SNR by broadcasting a larger pulse into the body, but you must stay within the safe limits of energy. The ADC in ultrasound systems must provide enough SNR to accommodate the amplifier chain in front of it. In modern systems, this task requires a 12-bit ADC for the pulsed path and at least an 18-bit unit for the CW path. Similarly, you can increase SNR in an analog-signal chain by burning more power in the front-end amplifiers, but this approach affects battery-power consumption. Equally daunting, if you build a system with hundreds of channels, you need the digital processing and software to do something with all the information. This approach also consumes power. You can use a “slice-count” technique—that is, increase the number of channels—to improve not only the resolution but also the speed of imaging and the volume of tissue you can examine.

A medical-ultrasound-signal path looks similar to that of a phased-array radar installation, except that the ultrasound-signal path operates at much lower excitation frequencies. Unlike a radar receiver, a modern ultrasound unit may be a portable device that fits into the palm of your hand and runs on batteries. These machines transmit a 2- to 17-MHz ultrasound frequency into a patient’s body; the lower the frequency, the deeper the machine can peer into the body. The round-trip attenuation of an ultrasonic signal in the body is 1.4 dB/cm-MHz. Doppler systems can discriminate the velocity of only those particles that are smaller than the wavelength of the excitation

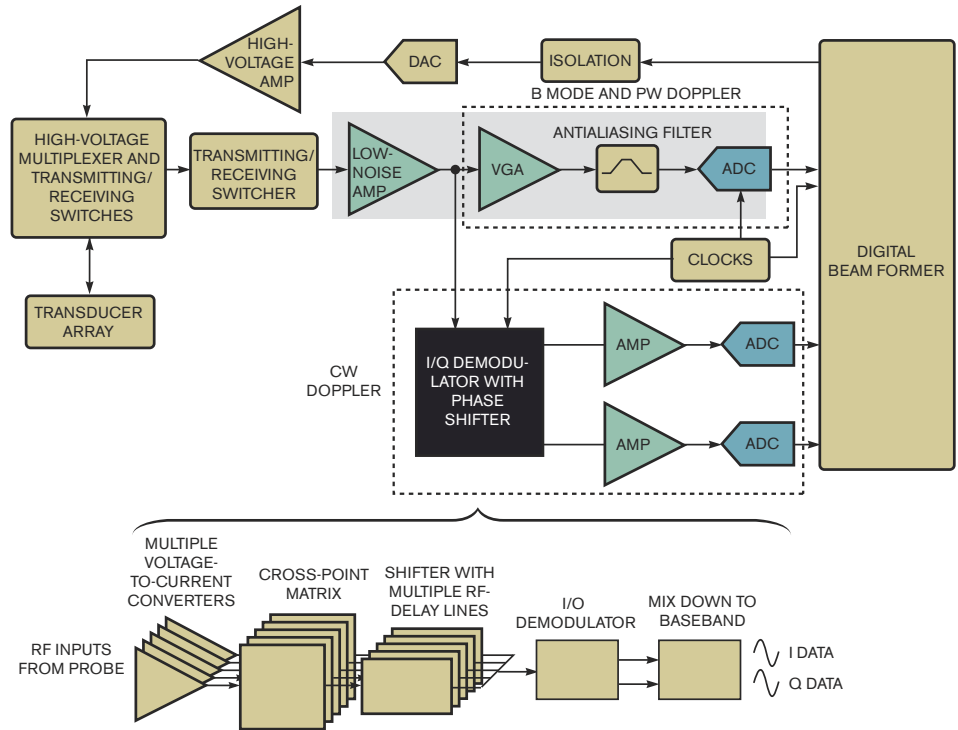


Figure 2 The analog front end of an ultrasound machine is a complex transmitter-receiver system similar to those of phased-array radar systems (courtesy Analog Devices).

frequency. Ultrasound cannot travel through air, so physicians typically apply a gel to the skin where they are using the transducer. Specialized transducers work inside the body to provide closer examination of the esophageal, reproductive, or digestive tracts. The ultrasound signal does not penetrate air-filled cavities in the body, including lungs and intestines; it can transmit images only from organs and from measuring blood flow.

A medical-ultrasound system is in many ways similar to other communication systems (Figure 2). The presence of multiple transmitters and receivers gives

ultrasound a similarity to MIMO (multiple-input/multiple-output) wireless systems. The same ADC might find use in either system, according to Allan Evans, vice president of marketing for fabless-semiconductor company Samplify Systems. The transmitting section emits a pulse in the same way that a sonar or a radar system does. The input section is a receiver similar to a radar receiver, except that it works at the speed of sound instead of the speed of light. The input path of a CW-Doppler-ultrasound system has I (in-phase) and Q (quadrature) demodulation that is familiar to the designers of base stations and cellular phones.

The transducers usually comprise quartz-piezoelectric materials. Researchers are also developing MEMS (microelectromechanical-system) CMUTs (capacitive-micromachined-ultrasonic transducers) to create and receive the sound pulses. One advantage of these devices is that they can make repetitive linear and 2-D arrays. “These CMUTs [eliminate] the whole interconnect problem,” says




Figure 3 The transmitted pulse from an ultrasound system affects the quality of the image the software produces (courtesy Texas Instruments).



John Scampini, director of strategic marketing at Maxim Integrated Products. In addition, designers can include some circuitry on the transducer itself because it is a silicon die. Unfortunately, these devices are currently in the research phase.

In general, most ultrasound machines have no signal-processing abilities in their transducer heads. Instead, all of the channels connect through coaxial cables that must withstand 150V transmitting pulses and must have low stray capacitance to carry the received signals without undue attenuation. Gore, a participant in the ultrasound-imaging-equipment area and a range of other consumer and electronics markets, makes a ribbon-based cable that is electrically equivalent to but thinner than 75Ω coaxial cable. This thinner cable targets use in sonographic ultrasound whose operators experience repetitive-stress injuries from manipulating heavy cables. Gore's thinner and more flexible cables solve this problem for the operators of these machines. Also, machines that create virtual channels by multiplexing 16 analog channels across a 64-element transducer may include multiplexing circuitry in their transducer heads—a benefit because the cable needs to carry only 16 analog channels and can be four times thinner and one-fourth the cost of a 64-channel cable.

Ultrasound transmitting side of an ultrasound system needs to deliver high-voltage pulses to the transducer elements. The ultrasound machine delays the pulses at the center of the linear transducer so that the sound waves converge on a point inside the body. For this reason, each channel needs its own pulser and perhaps a DAC. The voltages are typically 70V or higher, and peak output currents reach $\pm 2A$. Frequencies range from 2 to 17 MHz, and burst times are 1 μsec to several microseconds. The waveshapes are often just simple digital pulses, but more advanced machines include DACs to tailor the pulse shapes. In this case, the transmitter-driver chip is not a simple MOSFET array but a high-voltage amplifier. These DAC-generated pulses can be useful in therapeutics, in ultrasounds using contrast agents, and in chemopotential. The quality of the burst waveshape also factors into the image quality. Texas In-




Figure 4 Rather than make the analog front end on one die, Texas Instruments combines a silicon-germanium die and a CMOS die into one package.

struments, for example, ensures that its TX734 pulser transmitter has no undue ringing as the burst ends (Figure 3). In CW-Doppler applications, you use one-half of the transducer array as transmitters and the other half as receivers.

The receiver side of a medical-ultrasound system shares many characteristics of communications receivers. Because a 150V transmitting pulse would overload any amplifier system, the first block is a switch that protects the amplifiers from the transmitting pulse and then quickly switches in the transducer so that it can receive the sonic echoes. These switches are distinct from systems that may be multiplexing 16 analog channels into a 64-element transducer. All ultrasound systems need the receiver switches to prevent overloading in the signal path. The switches are often current-fed diode bridges, much like the front end of a sampling oscilloscope (Reference 2).

The first gain block is a fixed-gain, low-noise amplifier that typically runs at 19 dB and often has an active-termination network to minimize noise. The amplifier's noise limits the SNR of the

entire signal path. In addition, the amplifier must have a sufficient power supply so that there is no clipping of the signal. Input signals can be 0.5V p-p from hard reflections, such as the bone in the patient's body that is closest to the transducer. Furthermore, should the inputs overload, the low-noise amp must quickly and predictably exit saturation.

Manufacturers frequently use a silicon-germanium or a CMOS process for low-noise amps, but silicon germanium provides better linearity and a lower noise figure (Reference 3). You can operate amplifiers employing this process at a higher supply voltage, ensuring that you do not overload the amplifier. Another benefit of silicon germanium is that it yields a low flicker-noise corner, an important criterion in CW-Doppler ultrasound. Even though the flicker noise is worse at dc, it still modulates the signal when it passes through the I- and Q-demodulator sections. Thus, the flicker noise results in a broader "skirt," the widening of the base of the frequency-domain-receiver spectrum. This phase noise obliterates small Doppler signals and prevents the display of slow and subtle changes in blood flow.

On the other hand, using large CMOS transistors reduces voltage noise. Analog Devices, for example, claims that its AD9272 low-noise amplifier has lower terminated noise than any other analog front end. CMOS always excels at current-noise specifications because the input-bias currents are so low to begin with. "CMOS has the advantage of having very low input-noise current," says Corey Peterson, project leader of the AD927x-product line at Analog Devices. "For higher probe impedance, the in-



Figure 5 Ultrasound systems are becoming available in handheld form factors (courtesy Signostics Medical).

put-noise current can be as significant as the input-noise voltage.” He also notes that the company’s triple-well process prevents digital-substrate currents from interfering with the analog signals and that using large input transistors can reduce flicker noise to acceptable levels. Although the higher flicker-noise corner of CMOS is unavoidable, it is often acceptable if you do not use CW mode and if the machine is a low-performance portable model.


After passing through the low-noise amplifier, the signal splits off between a CW-Doppler path and a B-mode path and conventional color-Doppler path. The conventional path starts with a PGA (programmable-gain amplifier) that typically adds 40 dB of gain. The PGA varies dynamically as it receives each pulse; thus, the signals that reflect from deeper in the body receive more gain. This approach allows a channel SNR of 110 dB and extends the effective 70-dB dynamic range of a 12-bit ADC. As with a low-noise amp, bandwidths for PGAs are approximately 20 MHz. One vital characteristic of PGAs is recovery from overload. You can expect hard tissue or bone to reflect a large signal that overloads the PGA. The system can sense this signal and reduce the gain at that point in the receiving time, but the faster the PGA recovers from overload, the sooner the system will have accurate information to process. Overload can also create frequency modulations that are indistinguishable from blood flow in a Doppler system.

Amplifier and ADC vendors measure overload recovery differently, and even ADC vendors differ among themselves in how they evaluate overload recovery. This measurement is not so much about specmanship but about the probabilistic nature of the overload regimens the system must undergo. “You have to work with customers carefully and closely to translate what is a system phenomenon of overload into a spec for a part,” says Scott Pavlik, a marketing manager in the health-care segment at Analog Devices.

An antialiasing SAW (surface-acoustic-wave) or a passive filter follows the PGA. Although the 20-MHz bandwidth is high for active filters, amplifiers are continuously pushing this boundary. It is common for the ultrasound ADC to op-

erate at 50M samples/sec so that the antialiasing filters have sufficient roll-off and flatness to prevent higher-frequency harmonics from mixing back into the signal and reducing quality. With active filtration, however, you also must speed up overload recovery. Passive networks, on the other hand, need no recovery period. You should characterize your SAW devices for overload, but they also recover nearly instantaneously.

The analog signal ends at the ADC, which is typically a 12-bit, 50M-sample/sec pipeline device. Older systems use 8-bit converters for B-mode scans, but the



requirements of pulsed Doppler have driven bit counts to 12. Another factor to consider is harmonic-imaging mode. Compressed tissue under an ultrasound pulse reflects back second harmonics, and a 12-bit converter allows you to sense these harmonics and compensate for the tissue compression, allowing for better resolution in the scan. Although SAR (successive-approximation-register) converters provide better SNR figures, pipelined converters are acceptable. The VGA rather than the ADC determines the SNR of the signal chain. How quickly the converter recovers from overload conditions may be more important than SNR, however.

Ultrasound systems can alternatively use CTSD (continuous-time-sigma-delta) converters, which are available from National Semiconductor and Analog Devices. Systems using this topology have low power consumption for their speed and SNR. CTSD converters also require no antialiasing filter in

front of the converter because the loop filters in the converters inherently prevent aliasing. With CTSD converters, the internal loop amplifiers control the bandwidth, so you cannot undersample them. As a result, a 50M-sample/sec CTSD converter can operate only at that sampling rate; it doesn’t work at 25M or even 40M samples/sec. Another caveat has to do with overload recovery. Because pipelined converters carry their samples in separate sections of a pipeline, they inherently recover well from overload—often during one sample. CTSD-ADC designers must tack overload recovery with clamping networks onto the internal integrator, an approach that typically reduces SNR and increases power consumption while the part is in an overloaded state. This situation is not a serious problem because the ADC typically does not remain in saturation for more than a few cycles.

Because dozens of channels of 12-bit data would use up all of the pins on even the largest FPGA, the converters use serial outputs, such as LVDS (low-voltage-differential signaling). Parts that have eight ADC channels in one package and multiple channels on one pair of differential outputs further reduce the pin count in beam-former FPGAs. Be aware that SNR specs can be deceiving. “An ADC can have very poor near-carrier SNR and still having excellent full-Nyquist-band SNR specifications,” says Maxim’s Scampini, who also warns that flicker noise in the ADC reference can cause poor near-carrier SNR.

The other analog path in a medical ultrasound front end is for CW Doppler. This path needs greater SNR than the 110 dB that the B-mode and PW-Doppler paths provide. “Bear in mind that the transmission signal is continuous and is also being received,” says Samplify’s Evans. “You can also imagine that the refraction from a blood vessel wall would be far larger than the Doppler scattering of the blood carried inside the vein. In the CW Doppler, you are trying to maintain a dynamic range of 154 dBm [decibels referenced to milliwatts]/Hz. The thermal noise floor of 50Ω is –174 dBm/Hz.”

Because the CW-Doppler technique does not use pulses, you cannot disregard the large signals by blanking them out




or attenuating them for a certain time. Instead, the technology picks off the ultrasound signal after the low-noise-amp section and sends them to a demodulator section that performs analog-signal processing. Doppler signals represent changes in phase, so the I and Q demodulation that advanced-communications techniques use is also applicable to mixing down the 2-MHz CW signal to baseband. This demodulation allows the system to discriminate between blood flow toward the sensor and blood flow away from the sensor. The I and Q components then go to a higher-resolution, slower-sampling ADC. Demodulation has moved down the signal from the frequency of the transmitting pulse to the frequency of blood flow in your veins. Suitable ADCs are 16- or 18-bit devices sampling at 150k or 100k samples/sec, respectively. The I and Q components then move to the spectral-processing subsection in the beam-former FPGA.

INTEGRATION TRADE-OFFS

A fundamental choice you will have to make in building an ultrasound sys-

A FUNDAMENTAL CHOICE IN BUILDING AN ULTRASOUND SYSTEM IS HOW MUCH INTEGRATION TO USE.



tem is how much integration to use. Analog Devices this year won an *EDN* Innovation Award for its latest ultrasound front end (Reference 4). The company's eight-channel AD9276 includes the low-noise amp, the PGA, antialiasing filters, and an ADC—all on one CMOS-silicon die. Leveraging Analog Devices' expertise in communications, the part also integrates the I and Q demodulators. To make an analog

front end, you would need to add only the transmitting functions, the input switches, and the amplifiers and ADCs for the CW-Doppler function. Texas Instruments' AFE5804 analog front end integrates the amplifiers and ADC in one package but uses a low-noise amp and VGA on a silicon-germanium process and mounts it as a separate die into the same package as the CMOS ADC (Figure 4). The company also offers the AFE5851, which integrates the VGA and ADC but omits the low-noise amplifier, so you can use an external part.

"Customers want in the long term to put the electronics into the transducer," says Veronica Marques, strategic-marketing manager for the medical-business unit at Texas Instruments. "If customers want to integrate the low-noise amp into the transducer, they can still use a TI integrated analog front end," she says. Choosing the integrated front end locks you into one vendor, making it impossible to upgrade separate blocks, such as the PGA or the VGA.

Maxim, on the other hand, offers high-performance, silicon-germanium,

\$51^{For 3} PCBs
FREE Layout & Schematic Software!

- 01 DOWNLOAD our free CAD software
- 02 DESIGN your two or four layer PC board
- 03 SEND us your design with just a click
- 04 RECEIVE top quality boards in just days

expresspcb.com

YES.
We Can Even Make a BEACH UMBRELLA a Hot Spot!

- Design Services
- Certification
- Modules

LS RESEARCH, LLC
 WIRELESS PRODUCT DEVELOPMENT

The leader in wireless product development

www.lsrdesign.com

low-noise amps and then lets you choose PGAs and ADCs from the company or from a competitor. The MAX2038 integrates an amplifier with the quadrature mixer. The MAX2078 integrates the amplifier, VGA, filters, and CW-Doppler mixers across eight channels. If you want to differentiate your product by its software or user interface, using an integrated analog front end makes sense, especially if your design has space and power constraints. SonoSite, on the other hand, uses a proprietary analog front end because the results of the cost and design time show up in the displayed images. All of the company's analog front ends are custom-designed. "The chip makers are interested in leveraging our technologies," says the company's Dunbar. "But they still have a way to go."

Another integration factor involves the digital-data side of the front end. Samplify Systems, for example, which also this year won an *EDN* Innovation Award, does not integrate the low-noise amp and PGA into the ADC converter because high-voltage or low-noise pro-

FOR MORE INFORMATION

Analog Devices
www.analog.com

Gore
www.gore.com/en_xx/
products/medical/oem/
imaging

Maxim Integrated Products
www.maxim-ic.com

National Semiconductor
www.national.com

Samplify Systems
www.samplify.com

Signostics Medical
www.signostics
medical.com

SonoSite
www.sonosite.com

Tektronix
www.tek.com

Texas Instruments
www.ti.com

TSMC
www.tsmc.com

25th Infantry Division
www.25idl.army.mil

UMC
www.umc.com

cesses are better suited to those blocks (Reference 5). Samplify's Evans notes that a low-power, 12-bit, 50M-sample/sec converter will work across several products—from cart-carried systems to portables. "You can then mix and match low-noise-amp/VGA front ends, depending on your power/performance trade-offs," he says. Samplify's 16-channel ADCs have built-in lossless and lossy data compression, reducing the cost, the complexity, and the data rates to the beam-forming chip. The reduced data rates have lower EMI

(electromagnetic interference), and the compression routines also reduce frequency peaks and keep the EMI spectrum more like pink noise. In addition to the chips, the company can provide you the IP (intellectual property) for your FPGA or DSP to decompress the signal. One interesting development has been the use of lossy compression, which provides medical images that are indistinguishable from those from non-compressed systems.

Other factors you should consider when using integrated chips are their availability and their obsolescence. Due to the fine-line CMOS content in a multichannel ADC system, many of the analog-semiconductor vendors farm the fabrication out to TSMC (Taiwan Semiconductor Manufacturing Co), UMC (United Microelectronics Corp), or other large digital-CMOS-fab vendors. This approach reduces your chances of ending up on an allocation list when the boom in semiconductor demand inevitably comes. Equally concerning is obsolescence. It is not beneficial to design in a proprietary chip for

View Online Now at
vicorpower.com/promotions/powertech_series

The screenshot shows the 'VICOR POWER TECHTORIALS' website interface. It features a 'Choose a Power Techtorial' dropdown menu, a 'Related Documents' list with links to various technical articles, and a 'Related Links' section. A video player is embedded, showing a man in a blue shirt holding a component. Below the video is a 'Now Playing:' section with the title 'Electromagnetic Interference: How to Get the Lowest Noise' and a brief description of the content.

NEW

**Electromagnetic Interference:
How to Get the Lowest Noise**

NEW

Thermal & Mechanical Considerations

NEW

Input Overvoltage Protection

Improving Output Filtering

Vicor DC-DC Converter Theory Overview


VICOR POWER TECHTORIALS™

Get Practical Answers To Your
Most Important Power Questions
View Online NOW!

Vicor's PowerTechtorial Series concentrates on important, real-world technical issues in power system design.

Questions, posed by power designers worldwide, are answered through concise, expert instruction. Each techtorial, on vicorpower.com or on CD, is accompanied by relevant technical articles and application notes.

Gain access to view an ever-growing number of Vicor PowerTechtorials. You'll be informed as each NEW techtorial becomes available. As always, technical assistance with your power system design is there, worldwide, from Vicor Applications Engineering at 800-927-9474.





cost, space, and power considerations if the vendor plans to stop producing the chip. Semiconductor vendors know that the medical-device market differs from the consumer-electronics market. Customers for medical devices must get FDA (Food and Drug Administration) approvals and expect chip runtimes of a decade or more. Although it is more difficult to win a medical-ultrasound socket, once a company gets that socket, the part can flourish for years. Even in a recession, people still become sick. Therefore, all the analog-semiconductor companies keep their proprietary chips in production for as long as customers need them.

Remember that the architecture of an ultrasound system involves an analog trade-off (Figure 5). You must decide whether CW-Doppler capability is important, given that it requires so many dedicated analog circuits. You must weigh multiplexing 16 channels into 64 against the performance this approach will yield. You must trade off silicon germanium versus CMOS, quartz versus MEMS, and DSPs versus FPGAs. Design

⊕ Go to www.edn.com/090625df and click on Feedback Loop to post a comment on this article.

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

cycles have decreased from five years to less than two years. A product that finds use in aircraft in a battlefield will need approvals from the FDA, Federal Communications Commission, CE (Conformité Européenne), FAA (Federal Aviation Administration), and Department of Defense. You may have to design systems that work in a decompression chamber to evaluate Navy SEAL (sea/air/land) special forces or Air Force personnel at altitudes of 30,000 feet.

You can understand why engineers devote their entire careers to designing analog front ends for medical-ultrasound systems. "There are a lot of people that get into it and then get addicted to it," says SonoSite's Dunbar. "All of us are addicted to ultrasound. You are building something that helps somebody." **EDN**

REFERENCES


- 1 Brunner, Eberhard, "How Ultrasound System Considerations Influence Front-End Component Choice," *Analog Dialogue*, Volume 36, No. 3, May-July, 2002, www.analog.com/library/analogDialogue/archives/36-03/ultrasound.
- 2 "Sampling Notes," Tektronix, 1964, www2.tek.com/cmswpt/tidownload.lotr?ct=TI&cs=afs&ci=14824&lc=EN.
- 3 Rako, Paul, "Silicon germanium: fast, quiet, and powerful," *EDN*, Sept 18, 2008, pg 27, www.edn.com/article/CA6594090.
- 4 "AD9272 eight-channel ultrasound receiver," *EDN*, March 30, 2009, www.edn.com/article/CA6631763.
- 5 "AM1610/05/00 12-bit ADCs," *EDN*, March 30, 2009, www.edn.com/article/CA6631770.

You can reach Technical Editor Paul Rako at 1-408-745-1994 and paul.rako@edn.com.



NEW DATEL Panel Meters

We're METER MAD at Murata Power Solutions!



Our design engineers have been working like crazy to bring you the very best in digital panel meter technology. Here are just a few of their latest designs!

DMS-20RM Series

True-rms AC Voltmeters with Isolated Inputs

- AC mains or DC powered
- 2000V isolation
- True-rms AC measurements of complex waveforms
- Large, easy-to-read red display



DMS-20PC Series

Low-Power Blue and Green LED Voltmeters

- Only available from Murata Power Solutions
- Draws 12mA from 5V supply or <0.5mA in "standby" mode
- 12-pin DIP package



ACA-20RM-ALM Series

True-rms AC Ammeters with Alarm Function

- User-settable alarm function, factory or field adjustable
- Flashing display indicates overcurrent fault conditions




For specifications or a sample, visit our website: www.murata-ps.com/dpm

muRata 
Murata Power Solutions

11 Cabot Boulevard, Mansfield, MA 02048-1151, USA Tel: 800 233 2765 Fax: 508 339 6356

Why is Agilent the fastest growing oscilloscope manufacturer?



Because we listen to you.

To build our scopes, Agilent carefully examines the challenges you face. Then we deliver products that solve your problems in imaginative ways. Like the multi-chip module that enables Infiniium's industry-leading signal integrity. And the ASIC that underlies InfiniiVision's patented MegaZoom deep memory giving you the industry's best signal visibility. You'll find innovations like these in each of our scopes — that's why more and more engineers are choosing Agilent over other scope brands.*

Agilent 20 MHz to >90 GHz real-time and sampling scopes

- Handhelds, portables, rack-mounts and benchtops
- 40+ application-specific software packages for exceptional insight
- Innovations that satisfy your toughest demands

See why more and more engineers choose Agilent.
Download our catalog
www.agilent.com/find/scopecatalog



NETWORK STORAGE IS A NOTABLY BRIGHT SPOT IN THE OTHERWISE-BLAH CONSUMER-ELECTRONICS ECONOMY. CAREFULLY SELECT AND CULL HARDWARE, SOFTWARE, AND THEIR JOINTLY IMPLEMENTED FEATURES TO ENSURE PRODUCT SUCCESS.

Accelerating consumers' NAS adoptions:

ASSESSING YOUR PRODUCT

OPTIONS

At least one potentially positive counterpoint—the NAS (networked-attached-storage) server—shines among the abundance of predominantly negative economic news about the technology sector, particularly consumer electronics. People continue to take still and video pictures, listen to music, and download movies—maybe even more so than in the past—because they're now staying home and looking to entertainment as a means of distracting themselves from their recession-related woes. More of them are also now working from home-based offices rather than in the cubicles of times past, when large enterprise servers and IT (information-technology) personnel met and managed their corporate-storage needs. Further, an increasing percentage of their homes contain reasonably robust networking setups, enabling various LAN (local-area-network) clients, such as computers, game consoles, media extenders, and printers, to not only share a common Internet connection but also intercommunicate.

All of these trends suggest the allure of a consolidated nexus in consumers' residences for both professional and personal content that multiple LAN clients could simultaneously access. Ideally, this centralized storage would implement a RAID (redundant array of independent disks), which would protect the NAS from the failure of any one hard-disk drive, and the


NAS would also act as a backup repository for all the computers on the network. Translating this vision into reality, however, requires that home-NAS suppliers deliver an easily justifiable price for the target market; an easy-to-grasp and compelling sales pitch; an easy way for consumers to both integrate the NAS within their networks and subsequently access it from diverse devices; sufficient speed in storage, retrieval, other processing functions, and network bandwidth; and a carefully crafted set of features and cosmetics.

In the more than 12 years that I've been dabbling in home-office NAS, I've seen abundant evidence of both evolution and maturation in the consumer-NAS-product category. Accompanying these trends, both diamonds and lumps of coal have emerged across the dozens of products I've used (see sidebar "Hardware-test beds"). Therefore, this article aims to provide not just a snapshot of current system and silicon-and-software building blocks but also a forecast of how the NAS category might further develop, with the guidance of historical precedents, product capabilities, and customer expectations.

THE NETWORK TETHER

Begin the architecture definition of your next NAS design from the outside, focusing first on its LAN interfaces. Wired Ethernet is the most common LAN-tethering approach—with good reason. Because NAS normally operates in a "headless" fashion—that is, without the need for a keyboard, a mouse, and a display—it can easily locate nearby the router and connect to it over Category 5, 5e, or 6 cable. Wired-Ethernet connections are comparatively robust and speedy. And your customers can leverage some other networking technology by using an external bridge adapter.

However, for aesthetics, operating noise, or other reasons, your customers might instead want to hide the NAS in some





out-of-the-way location, such as a closet. In these cases, you'll want to first ensure that you've educated your customers on the need for consistent NAS access to sufficient supplies of cool, ambient air. Consumers probably won't want to string Ethernet cable around their homes, so they might be willing to pay extra for an integrated alternative network-access technology.

Wi-Fi is probably not only the first approach that would come to mind but also the leading candidate by virtue of its pervasive presence in modern routers. Keep in mind that, even in its latest-generation 802.11n form, it's likely to be a lower-performance approach than wired Ethernet, however. Performance isn't the only reason to focus on 802.11n. Because 802.11g and other wireless predecessors are now mature, they won't provide justification for a substantive price premium.

Speed aside, Wi-Fi is also less reliable than wired Ethernet, due to RF (radio-frequency) interference and other issues, so you'll need to ensure that the NAS recovers from dropped network connections in a user-friendly and data-preserving manner regardless of what operating mode it's in at the time. And the need to provide the NAS with both the WLAN (wireless-LAN) SSID (service-set-identifier) and encryption-key information before it can make the Wi-Fi connection is a challenging setup requirement for a headless-system design. Finally, you need to decide whether to support both the 2.4- and the 5-GHz ISM (industrial/sci-

AT A GLANCE

Only power users can justify paying for blazing-fast network tethers.

Mirrored drives increase a system's size, weight, power consumption, and price, but they also prevent customers' data loss and frustration.

Ensure that your selected operating system and applications have robust features and are interoperable, but hide advanced features from neophyte users.

ARM and x86 appear to be the dominant CPU architectures for consumer-tailored network storage in the future by virtue of their ubiquity and intense industry focus.

Cost-effective and power-thrifty hardware has proved valuable for in-depth hands-on evaluations.

entific/medical)-band options, as well as how elaborate and expensive to make the unit's MIMO (multiple-input/multiple-output)-antenna array (Reference 1).

Because the NAS is ac-powered, thereby requiring a nearby wall outlet no matter where your customer puts the

unit, power-line networking might be a tempting alternative-LAN-interface approach. Not every power outlet is a valid power-line-networking candidate, however, and performance and reliability also vary with time of day and time of year (Reference 2). So, at least for now, ignore the temptation to integrate this feature and stick with an external Ethernet-to-power-line-bridge adapter. The three contending power-line-networking technologies—HomePlug AV, UPA (Universal Powerline Association), and HD-PLC (high-definition power-line communications)—exhibit no serious signs of interest in pursuing interoperability, so if you “bet on the wrong horse,” you'll add cost to your design and gain nothing (Reference 3).

Similarly, I don't currently recommend that you integrate either a HomePNA (Phoneline Networking Alliance) or a MOCA (multimedia-over-coaxial) transceiver. Neither technology is sufficiently mature to be in use by much of your potential customer base. And the need to connect such a NAS to a phone-line- or coaxial-cable-based network tether is too location-restrictive for many homes.

HOW MANY DRIVES?

Unless your target customer is a power user and particularly considering that a notable percentage of LAN clients will likely be accessing the NAS over low-bandwidth Wi-Fi connections, it may be difficult to justify the incremental cost of 1-Gbps GbE (gigabit-Ethernet) LAN transceivers versus conventional 10/100-Mbps alternatives. If a built-in Wi-Fi or 100-Mbps wired-Ethernet bottleneck constrains the NAS transfer-rate speed, there's little rationale for a performance-tailored RAID 0, RAID 5, or similar multidrive-striped-storage architecture behind the network PHY (physical)-layer IC. Conversely, if you believe your target customer will see tangible value in GbE or multistream, bonded-channel 802.11n-networking capabilities, you should seriously consider correspondingly beefing up your design's drive array.

Including more than one drive in your design typically costs more unless you're comparing, say, a leading-edge 2-Tbyte drive with two more mature 1-Tbyte alternatives in a concatenated arrangement (Reference 4). Using a multidrive





Figure 1 Linksys' first stab at a full-blown network-storage appliance for consumers, the NAS200 (a), had insufficient performance potential and, therefore, too few features. A switch from x86 to a Marvell ARM-based CPU substantially boosted the capabilities of follow-on Media Hub devices (b).



design also means that the NAS will need a larger system form factor, generate more heat, and, therefore, have a

greater likelihood of needing to employ a noisy system fan. As such, seriously consider 5400-rpm drives instead of

7200-rpm alternatives. Thanks to dense bit-packing PMR (perpendicular-magnetic-recording) techniques, the slower-

HARDWARE-TEST BEDS

I've been dabbling with network storage as long as I've been with *EDN*. Check out this list of hardware I've tried out over the years, which, given my imperfect memory, may not be comprehensive:

- Addonics Technologies' NASU2;
- ADS Technologies' NAS (network-attached-storage) drive kit;
- Buffalo Technology's LinkStation, LinkStation Pro, and TeraStation;
- D-Link's DNS-323 and DSM-604H;
- Intel's SS-4200E;
- Linksys' NAS200, NMH305 Media Hub, and NSLU2;
- Maxtor's Shared Storage II;
- Netgear's ReadyNAS X6, ReadyNAS NV, ReadyNAS NV+, ReadyNAS NVX, and SC101 Storage Center;
- Toshiba's Magnia SG10;
- Tritton Technologies' T-NAS;
- Via Technologies' Artigo A2000;
- various Windows and Mac OS X-based computers, along with an Apple Power Mac G4 Cube running OS 10.3 Server; and


• Ximeta's NetDisk.

One other recent hardware evaluation begs for more in-depth discussion. I mated a 1-GHz Via Technologies C7 CPU-based EPIA SN mini-ITX motherboard to a Casetronic Travla C137 enclosure (Figure A). I customized the C137 to hold dual 3.5-in. hard-disk drives from both Seagate and Western Digital for mirrored storage. Because many of the NAS systems on the earlier list use modified Linux distributions, I focused this evaluation on Windows Home Server, which Microsoft derived from Windows Server 2003, and FreeNAS, which its open-source developers based on FreeBSD (Berkeley Software Distribution) and which they recently and conveniently upgraded to Version 0.69.1. In its "embedded"-mode option, FreeNAS runs completely from flash memory. Although the EPIA SN offers an integrated CompactFlash card slot, I instead installed FreeNAS on a USB (Universal Serial Bus)-flash drive.

Both operating systems ran smoothly, although I admittedly didn't load them up with abundant add-ons, and they had comprehensive feature sets. Via's core-logic chip set on the EPIA SN supports hardware-accelerated RAID (redundant array of inexpensive disks), a key advantage when you partner it with a power-thrifty but performance-strapped CPU. However, Microsoft's proprietary mirroring approach in Windows Home Server couldn't tap into its features.

Similarly, I couldn't track down BSD drivers that would enable me to use the EPIA SN's core-logic chip set to offload the C7 CPU from handling some or all of the FreeNAS software-RAID algorithms. As with many other enthusiast-driven open-source projects, FreeNAS offers scant and incomplete documentation. The user interface, although feature-rich, is correspondingly complex and unintuitive. Should you decide to use FreeNAS as the foundation of your next NAS design, I'd encourage you to focus some tangible effort in polishing these areas and, per the open-source license, to return your results to the organization so that it can incorporate your improvements.

One other minor frustration involved the EPIA SN BIOS (basic input/output system). My version of the board contained the initial Version 1 firmware image. Via offers a newer Version 2.01 BIOS for downloading from its Web site, but the only corresponding update utility the company provides runs under DOS. I couldn't even execute it in command-line mode under Windows Home Server. As soon as I track down an old copy of DOS, along with a floppy drive to install it, I'll be able to upgrade my BIOS.



(a)



(b)

Figure A Combining the 1-GHz fanless-CPU version of Via Technologies' EPIA SN mini-ITX board (a) with a two-drive customized version of Casetronic's Travla C137 enclosure (b) and hard-disk drives from Seagate and Western Digital forms a stable and robust test bed for several NAS operating systems.



spinning drives still deliver robust transfer rates, and they consume notably less current. Despite the downsides of using multiple drives, avoid selecting a non-mirrored-drive architecture unless the customer will use the NAS exclusively for connected-computer backup. In the backup-only case, if the NAS drive fails, your customers will likely be able to swap in a replacement drive before any backed-up computer's drive also fails.

Think about it: Your marketing coun-

terparts will be advocating that your customers should use the NAS as a single-point-of-storage contact for all of their precious—often irreplaceable—digital data: music libraries, photographs, videos, financial records, and the like. Unless you use a RAID 1, RAID 5, or other mirrored-drive arrangement, such as Infrant's (now Netgear's) proprietary and flexible X-RAID, an inevitable drive failure will render that information permanently irretrievable. You can surely

convince your customers of the value of redundancy within the NAS, yes? This topic brings up a bigger issue regarding how the NAS market may evolve in the future. Today's NAS suppliers include traditional hard-disk-drive companies, such as Seagate and Western Digital; traditional network-equipment vendors, such as Cisco's Linksys division, D-Link, and Netgear; and start-ups, such as Data Robotics. Hard-disk-drive companies are understandably more loath than companies in the other two categories to admit to the inevitable impermanence of drives. Also, is there a future NAS-supplier role for traditional consumer-electronics companies, such as Samsung or Sony?

All this talk about hard drives inevitably brings up the topic of the solid-state-drive alternative (Reference 5). These drives are increasingly becoming available in cost-effective capacities that make them compelling hard-disk alternatives for client computers. However, the bulk-storage nature of NAS makes it likely that it will continue as a hard-drive candidate at least for the next few years. Near-term pragmatism aside, increased flash-memory density and lower prices are indisputable trends, particularly since the advent of multilevel-cell-storage techniques. As such, solid-state drives' increased reliability and performance, decreased power consumption and heat dissipation, and silent operation will likely in the future encourage their adoption in NAS at hard drives' expense.

OPERATING-SYSTEM BASICS

If you constrain your NAS brainstorming to only networked bulk storage, you might at first glance think that any of a number of operating systems could suffice. Dig a bit deeper, though, and you'll quickly realize that it's more difficult to solve this problem. First, a tangible percentage of your users will likely want to be able to carve up the available capacity into more than one shared-storage resource, with per-share access rights, such as disabled, read-only, or read/write, that customers will define on a per-user and -group basis. They'll access the networked storage from LAN clients running various operating systems and therefore with various

NAS ADAPTERS AND ALTERNATIVES

Although a full-blown NAS (network-attached-storage) system might be the preferable approach for newcomers, some consumers already own USB (Universal Serial Bus)-based external hard-

disk drives that they might want to migrate from a single-computer tether to full network access.

Linksys was among the first to address this market with its NSLU2 network-storage link (Figure A).

Advocates of open-source software have built on the code that Linksys provides to turn the NSLU2 into a full-featured server appliance. More recently, Addonics Technologies released the NASU2 NAS adapter, which is substantially smaller than the NSLU2 and includes built-in Bittorrent, FTP (file-transfer-protocol), and print-server capabilities for both LAN (local-area-network) and WAN (wide-area-network) users.

A NAS adapter isn't the only way to get a USB-inclusive hard-disk drive online. An increasing number of routers, for example, embed USB ports that accomplish a comparable function. Some routers even provide WAN access to the drive either free or with a \$99-per-year MobileMe membership for Apple's routers. Carefully select the file system the networking equipment uses to format the drive. The FAT (file-allocation-table) system, for example, has the advantage of operating-system ubiquity. However, even the newest FAT32 variant exhibits a 4-Gbyte-file-size limit, along with an operating-system- and cluster-size-dependent volume limit. These limitations may be too constraining, depending on how your customers use network storage.




(a)



(b)

Figure A Linksys' NSLU2 (a) jump-started the USB NAS-adapter category, which has now expanded to include products such as Addonics' NASU2 (b), a more fully featured out-of-the-box offering.



supported file-access protocols, such as AFP (Apple-filing protocol), NFS (network-file system), and SMB/CIFS (server-message block/common Internet-file system). They'll also want both configuration and subsequent access to work in a way that doesn't force them to comprehend and grapple with the underlying complexity.

LAN-client backup, another commonly requested NAS feature, is similarly challenging to implement in a simultaneously robust and trouble-free manner. Apple OS 10.5's built-in Time Machine capability, for example, initially supports full backups and subsequently supports incremental backups to AFP-cognizant storage media. However, Apple officially sanctions backups only to its own Time Capsule hardware, which combines a router and a hard-disk drive

(see sidebar "NAS adapters and alternatives"). Windows currently integrates no comparably robust backup features, although both Microsoft and third parties can subsequently augment the operating system with such capabilities. You also might want to include Rsync support for Unix clients. Keep in mind that users often want to back up files that the operating system or an application running on it is currently using. As a case study of the concept, although the Connector client-side software for the Windows Home Server NAS operating system generally runs well, it's not without limitations and quirks. It automatically wakes up PCs once a day, even if they're on standby at the time, but it sometimes fails to put them back to sleep once backup completes. Automatic wakeup also doesn't work if Windows is running

virtualized on another operating system (Reference 6).

Next consider the laundry list of other NAS capabilities that your potential customers might value and, therefore, pay extra for. These features include on-the-fly encryption during storage and subsequent decryption during read-back of information archived on the NAS, along with USB (Universal Serial Bus) ports for printer serving, augmented storage capacity, and networked access to scanners and other USB peripherals. Your customers might also want automatic network discovery through protocols such as UPnP (universal plug and play) and Apple's Bonjour—that is, Zeroconf. Media streaming is also on the list. Protocols such as UPnP AV (audio/video) and DLNA (Digital Living Network Alliance) enable this feature both

x86 ENHANCEMENTS

The relatively archaic 1-GHz Via Technologies processor I used for this project seemed to have sufficient horsepower for entry-level NAS (network-attached-storage) applications. However, as you beef up your software with concurrently running utilities, you might find that your design begins to slow down. If so, consider Via's Artigo A2000, which includes a 1.5-GHz C7 CPU in a NAS-friendly dual-3.5-in. hard-disk-drive case (Figure A). The A2000 also features "a quiet, ball-bearing fan, [which] silently cools

the system with noise levels remaining below 26.8 dB" (Reference A).

Don't forget about Intel's cost-effective, single- and dual-core Atom CPUs and mini-ITX boards that the company based on them. HyperThreading's virtual-multi-core support delivers even more horsepower in some configurations (Reference B). Also, if you're doing PVR (personal-video-recorder) applications, such as video encoding before archiving or transcoding before streaming, you might want to consider using the three-way-super-

scalar, out-of-order architecture that Via includes in its Nano CPU and implements in its VB8002, the first Nano-based and media-server-tailored mini-ITX board.

REFERENCES

- ▲ Artigo A2000 Barebone Storage Server, Via Technologies, www.via.com.tw/en/products/embedded/artigo/a2000.
- Dipert, Brian, "The price of falling prices: evaluating value-oriented x86 CPUs," *EDN*, Jan 8, 2009, pg 30, www.edn.com/article/CA6625435.





Figure A Via's Artigo A2000 NAS-tailored design embeds a 1.5-GHz CPU and a nearly silent system fan (a), whereas Intel's single- and dual-core Atom boards support a more modern CPU architecture (b). For stringent performance demands, consider Via's VB8002 board, which the company based on its most recent Nano CPU, complete with a three-way-superscalar, out-of-order architecture (c).



across the LAN and over a WAN (wide-area-network) connection. Firewall-surmounting technologies, such as UPnP and NAT-PMP (network-address-translation/port-mapping protocol), support the WAN connection.

Customers might also pay for additional file-access and update protocols, such as FTP (file-transfer protocol) and Bittorrent, including built-in servers for them. They might even want approaches such as Telnet and TFTP (trivial FTP). Dynamic DNS (domain-name-service) and NTP (network-time-protocol) support may also be on customers' wish lists, along with a Web server, both for convenient user access to the NAS settings and for enabling the NAS to serve generic Web pages through both HTTP (hypertext transfer protocol) and HTTPS (HTTP-secure) over the LAN and WAN. Also consider iTunes, SqueezeCenter (formerly, SlimServer), and other media servers; iSCSI (Internet-small-computer-system-interface) support for optional SAN (storage-area-network) access; workgroup-tailored servers, such as DNS and e-mail, the latter complete with spam filtering; master-browser capabilities for Windows peer-to-peer workgroups; and direct attachment to OTA (over-the-air), cable, and satellite television tuners for both live-TV viewing and record-and-playback features using network extenders.

Although both open-source and proprietary software exists to implement these capabilities, each incremental concurrent task puts incremental demand on memory, processing, and other system resources. Incremental functions also threaten to exponentially increase the complexity of the perceived customer experience with the end result, along with the probability that functions will negatively interact with each other. With several of the NAS devices that I've tested over the years, multiple applications have insisted on using the same TCP (transmission-control-protocol) and UDP (user-datagram-protocol) ports, and other programs have blocked the NAS from putting its hard-disk drives in spin-down mode, thereby precluding consequent power-consumption reductions and operating-life extensions.

Speaking of operating life, warn your

WARN YOUR CUSTOMERS OF PENDING PROBLEMS WITH THEIR NAS WHILE THE OWNERS CAN STILL RECTIFY THE SITUATION.

customers of pending problems with their NAS while the owners can still rectify the situation. For example, you can send automatic e-mails to inform users of high temperatures, which may indicate clogged or otherwise failing fans and vents, along with SMART (self-monitoring/analysis/reporting technology)-drive-diagnostic results that exceed predetermined thresholds. To get those e-mails to your customers, though, you also need to support spam-blocking safeguards that ISPs (Internet-service providers) now put in place. These potential roadblocks include nonstandard SMTP (simple-mail-transfer-protocol) ports, user-name and password authentication at the SMTP server, and SSL (secure-sockets-layer) capabilities.

Keep in mind, too, that you must support no-brainer updates to the NAS BIOS (basic input/output system) or EFI (extensible-firmware-interface) code, operating system, and applications, both to patch vulnerabilities and bugs and to upgrade features in the field. It would be naive to assume that your customers will remember to regularly search for, download, and install service packs. The built-

in automatic Windows Update capability of Microsoft's Windows Home Server is one notably robust implementation of the concept; Apple's Time Capsules also regularly check for updates and alert users to their availability.

CPUs AND SUCH

Innumerable factors drive your selection of a CPU architecture, the features within that architecture, and a performance option of that feature set, including the software suite's demands, the system's BOM (bill-of-materials)-cost expectations, and the availability of highly integrated and application-optimized IC variants. In addition, consider not only architecture-tailored software from your company but also that of third-party software you might want to license, along with additional utilities that your customers may want to install after the purchase. For example, many enthusiasts have developed freely downloadable add-ons for Windows Home Server on the We Got Served Web site.


Two examples highlight the divergent paths that companies have taken in this regard. First, look at Cisco's Linksys division (**Figure 1**). The company in January 2007 introduced the NAS200, which employed RDC Semiconductor's R3210 CPU, implementing the i486 microprocessor-instruction set. However, the NAS was so performance-strapped that it couldn't support either SMTP-server authentication or SSL cognizance for e-mail alerts; it also could not use its USB port to implement a print server. Similarly, the company initially shipped the NAS200 with support for only the journaling-inclusive XFS (extended file system).

Journaling support is desirable in typical consumer environments, in which a UPS (uninterruptible-power supply) doesn't feed the NAS, which can, therefore, abruptly shut down in the middle of a media write. But journaling and other advanced-file-system capabilities' algorithm processing also steal CPU cycles. In response to user complaints about slow accesses, Linksys added optional support for nonjournaled ext2 (second extended file system) through a firmware upgrade. Marvell's beefier ARM-based 88F5182 Orion SOC (sys-

⊕ See the "Accelerating consumers' NAS adoptions" posts at www.edn.com/briansbrain for supplemental information on this article's topics.

⊕ Go to www.edn.com/090625cs and click on Feedback Loop to post a comment on this article.

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



tem on chip) powers the newer Media Hub NAS line, which has substantially more capabilities than its NAS200 predecessor. Similarly, the latest iterations of Buffalo Technology's LinkStation and TeraStation NAS products have started to use ARM processors; a mix of MIPS and PowerPC CPUs initially fueled these products.

Linksys migrated away from x86, but Netgear seems to be going in the opposite direction. The company in May 2007 acquired Infrant Technologies and its ReadyNAS product line. Infrant began life as a silicon supplier of the Leon SPARC-compatible CPU for NAS. For reasons that likely involved a dearth of stand-alone-IC sales to others, the company later switched gears and decided to become a system supplier selling Leon-based NAS. Netgear no longer manufactures the initial 600, X6, and NV product lines; Duo, NV+, and NVX systems continue to use Leon.

In late 2008, however, Netgear rolled out the ReadyNAS Pro, available in both enterprise- and consumer-target-

ed variants. ReadyNAS Pro leverages a dual-core Intel x86 CPU. It currently occupies the high end of the company's product line, complete with six-drive support. Although Netgear doesn't comment on future product plans, it's not a stretch to imagine the company's embrace of the x86 extending throughout the ReadyNAS line in the future (see sidebar "x86 enhancements"). As a longtime ReadyNAS X6 user, I've been frustrated at the long delays between the debut of new versions of PacketVideo's TwonkyMedia DLNA server on conventional platforms and its availability on Leon-based hardware. Adopting a mainstream-CPU foundation would probably shorten those delays. **EDN**

REFERENCES

- 1 Dipert, Brian, "802.11n: complicated and about to become even messier," *EDN*, May 28, 2009, pg 6, www.edn.com/blog/400000040/post/1970043397.html.
- 2 Dipert, Brian, "Power line: more reliable than wireless? You've got to be kidding ... yes?" *EDN*, Feb 23, 2009, www.edn.com/blog/400000040/post/790041079.html.
- 3 Dipert, Brian, "Power line: Does market success necessitate a function and feature reset?" *EDN*, Feb 23, 2009, www.edn.com/blog/400000040/post/1300041130.html.
- 4 Dipert, Brian, "Hard-disk drive price tracks: competitive impacts," *EDN*, April 7, 2009, www.edn.com/blog/400000040/post/730043073.html.
- 5 Dipert, Brian, "Solid-state drives challenge hard disk," *EDN*, Nov 13, 2008, pg 36, www.edn.com/article/CA6611643.
- 6 Dipert, Brian, "Virtualization: silicon and software salvation or technological tower of Babel?" *EDN*, Oct 2, 2008, pg 34, www.edn.com/article/CA6598366.

FOR MORE INFORMATION

Addonics Technologies www.addonics.com	MIPS Technologies www.mips.com
ADS Technologies www.adstech.com	Multimedia Over Coax Alliance www.mocalliance.org
Advanced RISC Machines (ARM) www.arm.com	Netgear www.netgear.com www.readynas.com
Apple www.apple.com	PacketVideo www.pv.com www.twonkymedia.com
Buffalo Technology www.buffalotech.com	RDC Semiconductor www.rdc.com.tw
Casetric www.casetric.com	Samsung www.samsung.com
Data Robotics www.drobo.com	Seagate (Maxtor) www.seagate.com www.maxtor.com
D-Link www.dlink.com	Sony www.sony.com
FreeBSD www.freebsd.org	Toshiba www.toshiba.com
FreeNAS www.freenas.org	Tritton Technologies www.trittontechnologies.com
HomePNA Alliance www.homepna.org	Via Technologies www.via.com.tw
Intel www.intel.com	Western Digital www.wdc.com
Linksys www.linksysbycisco.com	Ximeta www.ximeta.com
Marvell Technology Group www.marvell.com	
Microsoft www.micosoft.com	

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



Development Solutions for ARM, 8051 & XE166

Microcontrollers

Microcontroller Development Kits

C and C++ Compilers

Royalty-Free RTX Kernel

µVision Device Database & IDE

µVision Debugger

Complete Device Simulation

Examples and Templates

Keil PK51, PK166,
 and MDK-ARM
 support more than 1,700
 microcontrollers
www.keil.com/dd

RTOS and Middleware Components

RTX Kernel Source Code

TCPnet Networking Suite

Flash File System

USB Device Interface

CAN Interface

Examples and Templates

Keil RL-ARM and ARTX-166
 highly optimised, royalty-free
 middleware suites
www.keil.com/rtos

Download the
 µVision4 Beta Version
<http://www.keil.com/uv4>

www.keil.com
 1-800-348-8051



DESIGN AUTOMATION CONFERENCE

Moscone Center, July 26-31, 2009
San Francisco, California, USA



Connections for Design Success

Now, more than ever, knowing which design tools, IP, test, and manufacturing processes to use makes all the difference for design success. No other event gives you the opportunity to interact with the community of researchers, designers, and tool developers who enable the electronic design process. Invest your time in efficiency and productivity – attend DAC.

Only DAC offers:

- **New DAC User Track:** Over 90 technical paper and poster presentations on design tool use and methodology – by design engineers for design engineers! Learn how your peers have solved challenging problems with novel approaches and solutions.
- Research and innovations from leading EDA and design technologists with over 150 papers, eight panels, three keynotes and nine special sessions on multicore designs, system-level and HW/SW codesign, physical design, power analysis/low-power design, verification and more.
- Conveniently located in San Francisco – close to the East Bay, Silicon Valley and SFO.
- Worldwide attendance from developers, designers, researchers, academics, managers and engineers from leading electronics companies and universities.
- A vibrant exhibition with over 200 companies with solutions covering the electronic design eco-system from concept to silicon. Check out the new IC Design Central Pavilion featuring IP, design and foundry services, PLUS the FREE sessions in the DAC Pavilion and the Exhibitor Forum.

Make DAC your one stop this year for networking, technical insights, and to find all the products and services you need to stay competitive.

DAC – Where electronic design meets.

Register today! www.dac.com



Sponsored by:



Supported by:



In Technical
Cooperation with:



Little-known flash-memory features protect data and IP

FEATURES FROM BLOCK LOCKING TO ENCRYPTED-PASSWORD-ACCESS MECHANISMS CAN PREVENT UNINTENTIONAL DISRUPTION, MALICIOUS DAMAGE, OR COPYING.

You design a system, and somebody messes it up. The damage is sometimes unintentional. For example, a service provider may install its software on your device and corrupt your original code. On the other hand, hackers and IP (intellectual-property) thieves go out of their way to overwrite, copy, or clone data stored in your systems. Whatever the cause, the resulting damage or theft represents no less of a problem. It's not surprising that designers need a way to protect system integrity. What may be surprising is that within its bits and blocks, flash memory holds the key to protecting firmware and even hardware designs.

Flash devices offer a number of data-protection measures, each with its own advantages for read, write, or erase protection. The security options add layers of security to slow down would-be hackers and thieves and provide protection from unintentional modifications. Some flash-security features don't even add cost to the final design, and, although the strongest flash-protection features may cost more than standard flash, they are far more affordable than a nonflash-hardware-encryption engine, hidden operations, authenticated operations, or software-encryption applications.

Manufacturers and even devices from the same manufacturer offer different features. Designers must select the right flash device for the final application after considering a number of factors, such as the built-in security options, performance, density, size, and cost.

FINDING THE RIGHT APPROACH

Evaluating the options starts with identifying the problem that you want to solve. Features perform specific functions, and some come with added cost. First, determine what you must protect amid the bits, data, and code. For example, you might need to protect electronic-system serial numbers, security keys, boot code, or financial information for services access, such as for pay TV. Once you know what you need to protect, determine whether a software or a physical disruption is likely to affect those bits, data, or code. A software attack may come from the Internet or a system application, for example. A physical attack, for instance, could involve the direct removal of

a flash device from a PCB (printed-circuit board).

Finally, identify whether the threat is unintentional or intentional. Unintentional alterations, such as those that bugs in software cause, are typically easier to prevent because the cause of the problem is not elusive or persistent. If the attack comes from a hacker or a thief, quantify how much effort the attacker is willing to make. The amount of time and money a hacker is willing to spend affects how much security the design requires. With these data points, determine which flash-security features provide the right level of protection against the source and intent of the attack. For example, if you must protect the design from data corruption from an Internet attack, block locking provides moderate protection, and OTP (one-time-programmable) blocks provide the best protection (Figure 1). If an IP thief aggressively targets the design by removing the flash device and attempting to read the data using a PROM (programmable-read-only-memory) programmer, protecting the design may warrant paying more for flash-data-encryption features (Figure 2).

DETAILED FEATURE REVIEW

From block locking to advanced encrypted-password access, you can choose the features that address the type and source

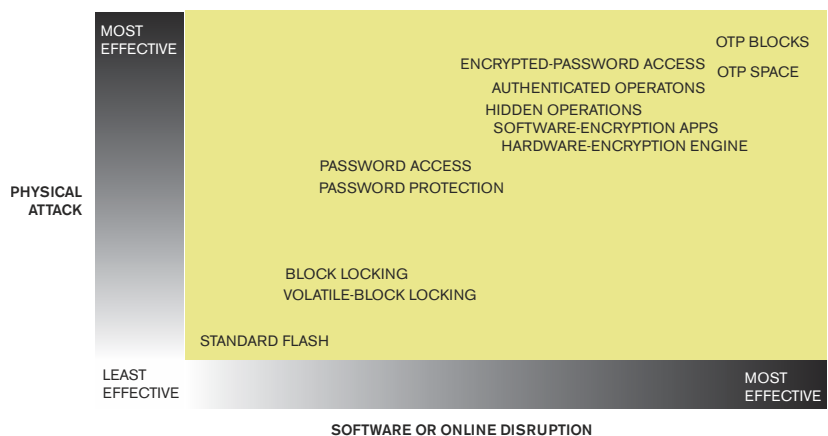


Figure 1 Security alternatives offer differing levels of protection against physical attack and software disruption, whether due to malicious corruption or unintentional data alteration. Note that adding such features as encrypted-password access, authenticated operations, hidden operations, software-encryption applications, and hardware-encryption engines increases system cost.

of an expected attack, with acceptable cost impact on the design. One such feature is password access. Some flash devices offer password-access features that slow down thieves, creating barriers that make the design a less desirable target for copying or cloning. IP thieves must quickly and easily copy system data. Password access adds time, cost, and effort to low-overhead cloning operations. Password access locks either the entire array or selected blocks in the main array from program, erase, or read access, depending on the flash device. You can individually set each block to the desired protection level. Before the system leaves the factory for the end customer, the manufacturer must store a 64-bit password in the password area of the flash device and program a matching password into the system microcontroller or other hidden storage.

When the system receives a command to read, modify, or erase data in the protected blocks, the system processor looks for a match between the number in the microcontroller and the one in the flash device. If the passcode is not valid with both the microcontroller and the flash device, the would-be hacker cannot read or modify the data. If the system detects a matching passcode, a user can read or modify individual blocks. Depending on the flash device, designers can choose from various protection modes, including read, modify, and substitution prevention.

Password protection is both a service-theft deterrent and an IP-copying and -cloning deterrent. Duplicated flash chips can provide premium services to users who don't pay for them, representing lost revenue for the service provider. Password-based read protection is a simple, cost-effective way to thwart attempts to distribute pirated flash chips that enable access to premium services. If the designer uses password protection on the flash device, he leaves the pirate with inoperable chips. When a would-be service thief attempts to read the data stored on the flash device, the device attempts to validate the 64-bit password. Without the password, the device returns only values of zero, rendering the copied chip inoperable. Even if the thief can snoop, discover the password, and copy the data from the chip, the 64-bit password in the pirated chip will not match the password in the microcontroller in the system into which the thief inserted it, again making the chip inoperable.

In the case of cloning, IP thieves must replicate and produce a design before an updated version of the original makes the clone obsolete. Flash-memory-password protection can create a significant enough delay to make the cloner seek an easier target because the flash protects the hardware signature of the system. A flash device with a 64-bit password limits access to legitimate sources. Without the password, an IP cloner who uses a PROM programmer to read the flash chip will read back only zeros.

As with the service-theft deterrent, the cloner can resort to bus-snooping to discover the password, but that step adds time, cost, and effort. The delay gives the legitimate manufac-

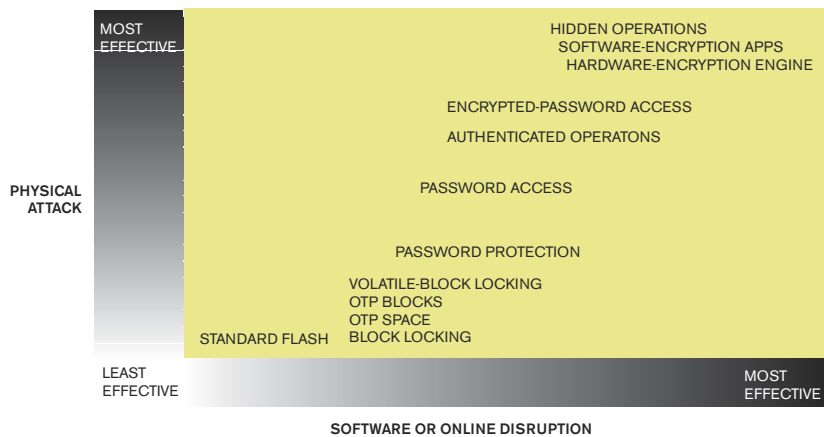


Figure 2 Various flash-security features can protect against cloning and intellectual-property theft. Note that adding such features increases system cost.

turer more time in the market before the clone can compete for revenue, and a deterrent may create enough delay for the original manufacturer to refresh its design before the cloner has an opportunity to produce a viable product. This built-in flash feature offers a cost-effective method for countering the revenue impact of IP loss.

Encrypted-password access for higher-level IP security is another method of protecting IP. A 64-bit password slows down a cloner or a service thief, but an encrypted password adds a significantly higher level of data protection. A few flash devices include an encrypted-password feature through an algorithm that manufacturers implement in silicon. The manufacturer encrypts passwords that pass from the flash and the microcontroller to the processor. The processor deciphers the passwords using the algorithm and confirms a match. A bus snooper, on the other hand, can only read—but not decipher—the encrypted password. Without the unencrypted password, the flash chip is unreadable, and the IP remains protected. Flash devices with an encrypted password typically cost more than those without one because manufacturers implement the algorithm on silicon, which adds to the component cost. However, the cost of the flash chip may be nominal compared with lost service revenue. For products with long refresh cycles, encrypted passwords are essential tools in preventing clones from reaching the market before the manufacturer updates its model.

OTP FLASH

Some flash devices include a system-level OTP area in which you can permanently lock the bits after programming. Once you lock the bits, a hacker cannot program or erase the blocks the bits are mapped to. The OTP has one factory-programmed segment with a unique, unchangeable number. The other segment is blank so that the designer can program it. Flash devices with OTP typically come in varying configurations of segments with as many as 2112 bits.

Several flash devices include an additional OTP feature that allows you to permanently lock blocks in the regular memory array. This implementation of OTP prevents modifications that disrupt system integrity. For instance, a service provider

typically adds its own code to a set-top box after delivery from the manufacturer. A designer cannot know how the service provider's code will affect the system. To protect the erasure or alteration of boot code, the designer can store the boot code in an OTP block and set it to be permanently locked. Then the service provider's code cannot write or erase to the locked block, thereby preventing those inconvenient customer-tech calls.

Most flash devices include some form of hardware-write-protection capability that can prevent programming and erasure of either a block or the entire device. Hardware-write protection works by setting a pin to a certain voltage, either through hard-wiring or by toggling a bit from an I/O pin on a processor. Before executing a modify command, the flash chip checks the pin that corresponds to program/erase protection. If the pin is not at the correct voltage for modification, the chip will not execute the command, and the code or data will not change.

When a valid voltage is present in the program-supply voltage, you can modify the blocks in the main array. If you ground the program-supply voltage, you cannot program or erase the blocks. When you ground the supply voltage, attempts to program or erase will fail, and grounding sets the appropriate status-register fail bit.

Another version of hardware-write protection protects the highest or lowest block against program and erase operations. To protect the highest or lowest block, set $V_{pp}/WP = V_{IL}$, where V_{pp} is the program-supply voltage, WP is write protection, and V_{IL} is the input-low voltage. In this situation, the block is in lock-down mode, and you cannot modify it. To remove a lock-down situation, set $V_{pp}/WP = V_{IH}$, where V_{IH} is the input-high voltage. In this case, you can lock or unlock the block.

HARDWARE PROTECTION

The hardware-based approach provides an inexpensive layer of protection against the malicious code that slithers in through the Internet. Malicious code cannot modify or erase data stored on a flash device that is locked at the hardware level unless it first resets the pin voltage. If a hacker tries to


bypass a router, for example, the malicious code will reach the flash device, and the device will check the pin for the block that stores the boot code. Finding the voltage at a level that doesn't allow modification, the chip does not execute the malicious code, and the router continues to work. Volatile- and nonvolatile-block-locking features use software commands to lock and unlock blocks, protecting data from inadvertent modification. In volatile-block locking, bits in a volatile array are mapped to main-memory-array blocks. You can individually modify, set, and clear these volatile-protection bits. However, they can protect only those blocks that you have not locked with nonvolatile-array bits. When you cycle the system power or reset the hardware, the volatile-protection bits revert back to their original unlocked or locked state.

Nonvolatile-block locking keeps blocks locked or unlocked, as the designer defines, even after a power cycle or reset. A nonvolatile-protection bit is mapped to and can individually lock each block. You can clear nonvolatile-protection bits through a clear-bits command or an erase command. You can use nonvolatile-block locking to ensure that blocks remain locked against inadvertent overwrites even after an unexpected power cycle or reset occurs.

Flash-security features vary by manufacturer and by device. You should consider these features, along with density, performance, technology, lithography, cost, size, and packaging. These factors all work together to enable a design, and you can use them to protect its performance and position in the market. Flash-security features offer an affordable, secure alternative to protecting IP, content, data, or system integrity. **EDN**


AUTHOR'S BIOGRAPHY

Bill Stafford is the director of segment marketing at Numonyx, where he develops application strategies and product requirements for the embedded-flash-memory market and enables the flash-memory ecosystem. A 25-year veteran in the electronics field, he has experience in product engineering, field quality, and marketing for flash memory, PCBs, systems, and aircraft components. Before joining Numonyx, Stafford spent 22 years at Intel Corp after time with the US Department of Defense.



DC-DC Converters

AC-DC Power Supplies



- Expanded Operating Temperatures **-55 to +85C**
- Vibration, *Method 204, Cond. D*
- Shock, *Method 213, Cond. I*
- Altitude, *Method 105, Cond. D*
- Environmental Screening
- Specification Review
- Custom Models Available
- 400 Hz and Now - **800 Hz AC-DC Models**

Thousands of Standard Models 2V to 10,000 VDC Outputs - 0.75 to 2,000 Watts

Call toll free
800-431-1064
for PICO Catalog
Fax 914-738-8225



PICO Electronics, Inc.
143 Sparks Ave, Pelham, NY 10803-1837
E-Mail: info@picoelectronics.com

See full Catalog immediately
www.picoelectronics.com

45V Synchronous Buck

60V Transients

4.5V to 45V_{IN}

LTC3642-3.3

3.3V_{OUT} @ 50mA

I_Q = 12µA

8mm

13mm

Actual Size Demo Circuit

Efficient, Rugged & Easy


Linear Technology introduces the LTC[®]3642, the first in a new family of synchronous high voltage micropower monolithic step-down converters. It addresses applications in 4mA to 20mA industrial control loops, security systems and general-purpose conversion where harsh conditions exist. The LTC3642 provides an extremely compact 45V_{IN} synchronous step-down solution with transient protection to 60V_{IN} and an I_Q of only 12µA. This device provides improved performance in size, efficiency, ease of design, and excellent thermal performance.

▼ Features

- Input Voltage Range: 4.5V to 45V, with 60V Transient Capability
- Low Quiescent Current: 12µA
- 50mA Output Current
- Adjustable Peak Current Limit
- 3.3V, 5V & Adjustable Versions
- Only 3 Externals Required
- 3mm x 3mm DFN or MSOP-8E Packages
- LTC3631: 20mA, 45V_{IN}*
- LTC3632: 100mA, 45V_{IN}*

*Future Products

LTC3642 Efficiency



▼ Info & Free Samples

www.linear.com/3642

1-800-4-LINEAR



Free Telecom,
Datacom and
Industrial DC/DC
Brochure

www.linear.com/48vsolutions

L^T, LTC, LT and LTM are registered trademarks of Linear Technology Corporation. All other trademarks are the property of their respective owners.




designideas

READERS SOLVE DESIGN PROBLEMS

Create a swept-sine function in LabView with just one virtual instrument

Sean McPeak, University of California—San Diego

 Swept sine waves are useful when you want to test a product over a wide frequency range. A large research project included the requirement to determine wave propagation in the open ocean. This application required the generation of a swept sine wave to drive an acoustic transducer. Although many waveform generators have a built-in function for this requirement, you must program it yourself if you want to implement a swept sine with a multifunction data-acquisition card. You can create a swept-sine function in National Instruments' (www.ni.com) LabView with just one VI (virtual instrument). Using this function, you can control start and stop frequencies, sample rate, and the overall duration of the sweep (Figure 1).

The LabView software calculates an array of numbers that represent the swept-sine-wave time series at each sample point as the frequency either increases or decreases, depending on the direction of the sweep. You must handle the frequency change of the output on a point-by-point basis. The basic form of the equation is $Y(I) = V \times \sin((A \times I^2) / 2 + B \times I)$, where $Y(I)$ is the amplitude of the swept sine wave as a function of the sample point, I is the integer that steps through the time series, V is the peak voltage, and A and B are variables. You define A as $2 \times \pi (f_{STOP} - f_{START}) / N$, and you define B as $2 \times \pi f_{START}$ where N is the number of samples, f_{START} is the normalized start frequency, and f_{STOP} is the normalized stop frequency. To normalize the start

DIs Inside

44 Charlieplexing at high duty cycle

48 Serial port tests digital circuits

51 DAC calibrates 4- to 20-mA output current

51 Alarm tells you to close the refrigerator door

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

and stop frequencies, you must change the unit to cycles per sample. You accomplish this task by dividing the f_1 and f_2 frequencies in hertz by the sample rate. You determine the sample rate by deciding how smooth of a transition you want to represent your swept sine wave. A good rule of thumb is to have at least 10 samples/cycle at the high-

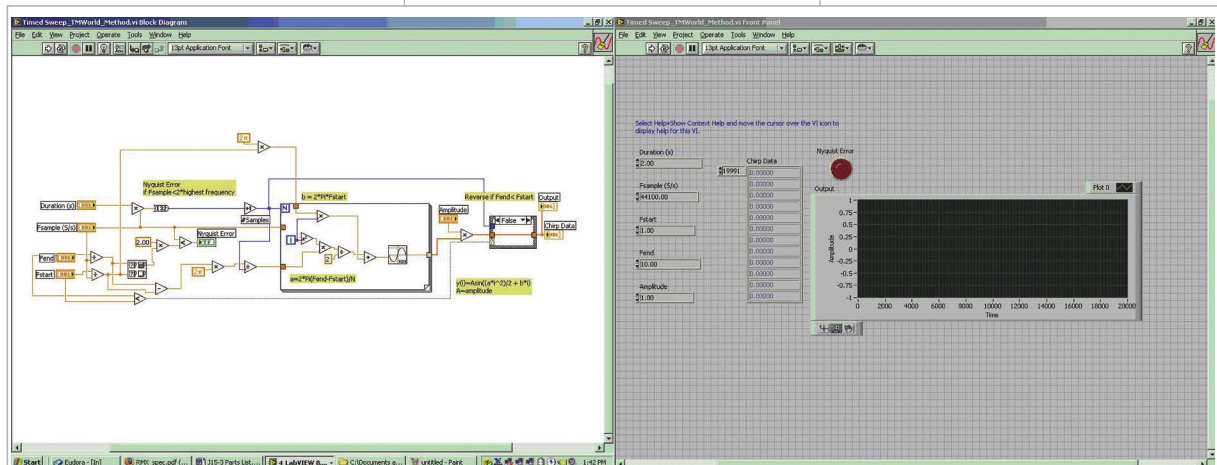


Figure 1 With just one LabView virtual instrument, you can control start and stop frequencies, sample rate, and the overall duration of the sweep.

est frequency. When setting the sample rate, you need to take into account the overall frequency span you are sweeping and the duration of the sweep itself. It is also helpful to compare the results and performance of the LabView data-acquisition-system implementation of the swept sine wave with those of an AWG (arbitrary-waveform generator).

You use two methods of comparison. First, you compare the output of both the data-acquisition and the AWG swept sine wave on a spectrum analyzer. Second, you run them both through an audio-amplifier/speaker system and simply listen to the output. This method is useful in determining sweep rate, duration, and stop and start frequencies. This type of comparison is valid only if the frequencies involved are in the audible range. The LabView VI employs simple array manipulation and uses a “for” loop. The input duration is in seconds, the sample frequency is in samples per second, and the starting and ending frequencies are in hertz. Dividing the sample rate immediately converts the start and end frequencies to cycles per sample. A maximum/minimum block takes the normalized ending and starting frequencies as its inputs and uses the maximum output of the input pair. You use this method to determine whether your design meets the

Nyquist criteria, given the sample rate and highest frequency you require.

This approach drives a simple Boolean variable to alert the user about whether the design meets the Nyquist criteria. You set the “for” loop to run for the total number of samples you want to calculate. You determine this value by multiplying the duration in seconds by the sample rate in samples per second. To guarantee that the loop processes all of the generated samples, you must add one, because the loop stops at $N - 1$.

You implement the output function in the “for” loop with simple algebraic operators and the sine block. The output is an array that reaches the perimeter of the “for” loop. It is important to enable indexing at this node. This action allows the circuit to individually handle each element in the array at the output of the “for” loop. You can also add a simple gain stage to set the peak-to-peak value to whatever point you want. Finally, you use the “rotate-1D-array-block” case structure to flip the array if the ending frequency is lower than the starting frequency. This approach handles cases in which you want a frequency sweep that starts out in a higher frequency and descends to a lower frequency.


You can easily modify and expand

this simple program. One idea would be to use the output array, which is nothing more than the time series representing a prescribed frequency sweep, to feed a loop that would drive a data-acquisition module. The output of the module should accurately represent the frequency sweep, as long as the module’s output sample rate is the same as the sample rate you use for generating the frequency-sweep time series. You should then be able to track the output samples and, when they are complete, reverse the frequency-sweep array. You then again feed this new flipped array to the data-acquisition module. Depending on the maximum and minimum frequencies, sweep duration, and sample rate you use, it may be difficult to flip the array and configure the module quickly enough to not miss a sample. In that case, you can prefill a frequency-sweep array for a set number of passes.

These modifications allow the sweep to continue back and forth for a set period. Another improvement would be to add some real-time FFT (fast-Fourier-transform) capability so that the user can see the sweep in the frequency domain. This approach also adds an increased level of insurance that the circuit properly meets the sweep definition. **EDN**

Charlieplexing at high duty cycle

Luke Sangalli, Digital Designs, Melbourne, Australia

 A few articles have recently appeared describing novel ways to increase the number of LEDs a microprocessor can drive with a limited number of pins (**Reference 1**). The standard multiplexing technique made popular by multidigit seven-segment displays has, in pin-scarce designs, given way to “Charlieplexing.”

Charlie Allen devised this technique while working at Maxim (www.maxim-ic.com), which has since introduced

ICs using the technique (**Reference 2**). Allen used the high-impedance third input state available to most microprocessors for turning off LEDs in a matrix; the LEDs’ respective microprocessor pins’ high or low states individually turn on these LEDs. Using this method, you can drive nine seven-segment LED displays using only nine microprocessor pins rather than the usual 17. For N pins, you can individually address $N \times (N - 1)$ LEDs using Charlieplexing.

One of the gripes people often level at Charlieplexing regards its poor duty cycle. A previous Design Idea com-

TABLE 1 NO. OF LEDs AND DUTY CYCLES

No. of pins	No. of LEDs	Charlieplexing duty cycle (%)	Standard multiplexing duty cycle (%)/no. of pins
Three	Six	33.3	50/five
Four	12	25	33/seven
Five	20	20	25/nine
Six	30	16.6	20/11
Seven	42	14.2	16.6/13
Eight	56	12.5	14.2/15
Nine	72	11.1	12.5/17
10	90	10	11.1/19



HB LED drivers for the widest range of lighting applications

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

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

Architectural and street lighting

MAX16820/MAX16822/MAX16832

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

MAX16826

- 4-string switching driver with I²C

MAX16824/MAX16825

- 3-string linear drivers with SPI™



MR16 lighting

MAX16834

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



Offline lighting

MAX16834

- > 0.90 input power-factor correction
- Compatible with triac dimmers
- Low component count
- Patent pending




Solar-powered lighting

MAX16821

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

MAX16834

- Boost and buck-boost topologies



EV kits available

SPI is a trademark of Motorola, Inc.

www.maxim-ic.com/LED-Light



www.maxim-ic.com/shop



www.avnet.com



For free samples or technical support, visit our website.

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

compares the standard multiplexing method with Charlieplexing (Reference 3). Using Charlieplexing, the maximum duty cycle for a 20-LED display is only 5%. The poor duty-cycle figure is due not to the method, however, but rather to the driving capability of the microprocessor and the parasitic-leakage paths. A single pin cannot usually sink the current a number of LEDs require to effectively light up, so these designs often require one source pin and one sink pin to light only one LED at any time. However, adding a transistor or two resistors allows you to circumvent these issues.

If you rearrange the LEDs in the familiar cross-point array and add a transistor to each column to carry the common current, you'll see the duty cycle of the Charlieplexing method does not differ much from standard multiplexing (Figure 1). For a 20-LED, five-column matrix, each LED remains on for 20% of the time compared with 25% for standard multiplexing, but now using only

BY THE TIME YOU GET TO 90 LEDs, THE PCB REAL ESTATE AND COST OF THE 10 TRANSISTOR/RESISTOR SETS PALE IN COMPARISON TO THE DISPLAY ITSELF.

five pins instead of nine (Table 1).

One of the drawbacks of adding the transistor and resistors to each column is that you need additional components to achieve a reasonable LED brightness when a large number of LEDs is involved. This approach, however, is a better alternative to using a costly IC and no worse than standard multiplexing or "Gugaplexing," which also requires additional transistors and resistors. From a cost and benefits point of

view, consider that, by the time you get to 90 LEDs, the PCB (printed-circuit-board) real estate and cost of the additional 10 transistor/resistor sets pale in comparison to the display itself.

Examining the circuit in detail, you'll notice that it has five microprocessor pins, P_1 through P_5 , available, for a total of $N \times (N-1) = 20$ LEDs. When P_3 , for example, is high, the emitter of Q_3 is at approximately 4.4V, and you can turn off D_{13} , D_{23} , D_{43} , or D_{53} if you make P_1 , P_2 , P_3 , or P_5 low. Any pin that you set to input, or high impedance, alternatively turns off the corresponding LED. When P_1 and P_4 are low, P_3 is high, and P_2 and P_5 are in high-impedance states. With P_3 high, transistor Q_3 biases on, all the other transistor bases are either low, which ensures that no current will flow, or high-impedance, which supplies no current into the base to allow the transistor to conduct. All the diodes in the third column can turn on, but only D_{13} and D_{43} have a path directly to ground through P_1 and P_4 , which are low and through the 100 Ω current-limiting resistors.

D_{23} and D_{53} connect to the high-impedance input pins and can conduct only through the 100 Ω resistors attempting to turn on Q_2 and Q_5 . Because of their forward-voltage drop—typically, 2.2V—the emitters of Q_2 and Q_5 will be less than 1.6V, as the following equation shows: $5V_{CC} - 0.6V(Q_3) - 2.2V(D_{23} \text{ or } D_{53}) - 0.6V(Q_3 \text{ or } Q_5) - I_{LED} \times 100\Omega < 1.6V$, where I_{LED} is the current of the LEDs. This scenario does not allow any LED in Column 2 or Column 5 to light up to any level that would have an undesirable effect. **EDN**

REFERENCES

- 1 Lancaster, Don, *Tech Musings*, August 2001, www.tinaja.com/glib/muse152.pdf.
- 2 "Charlieplexing—Reduced Pin-Count LED Display Multiplexing," Application Note 1880, Maxim, Feb 10, 2003, <http://pdfserv.maxim-ic.com/en/an/AN1880.pdf>.
- 3 Gupta, Saurabh, and Dhananjay V Gadre, "Multiplexing technique yields a reduced-pin-count LED display," *EDN*, Oct 16, 2008, pg 68, www.edn.com/article/CA6602447.

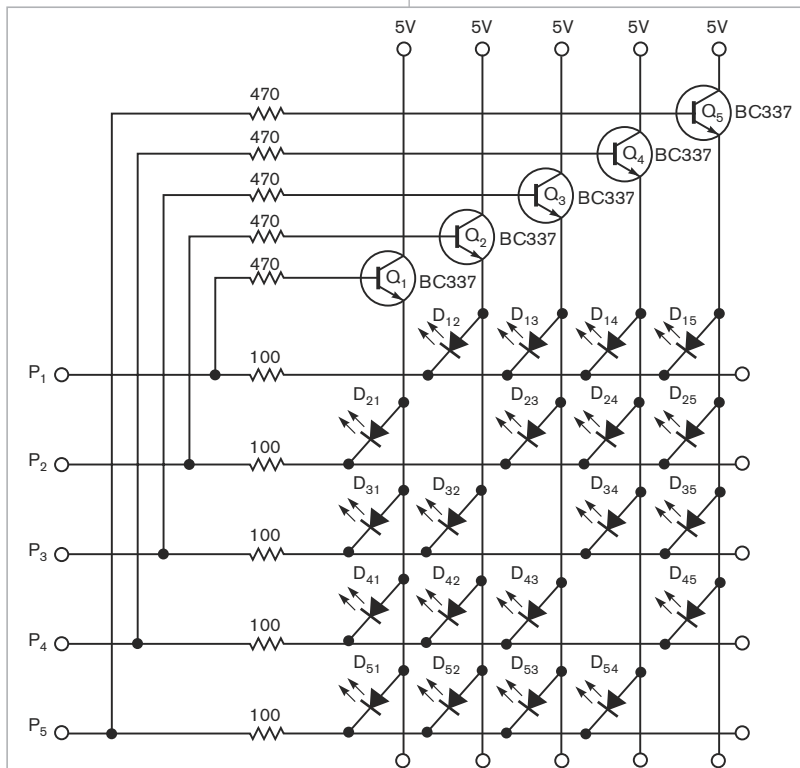


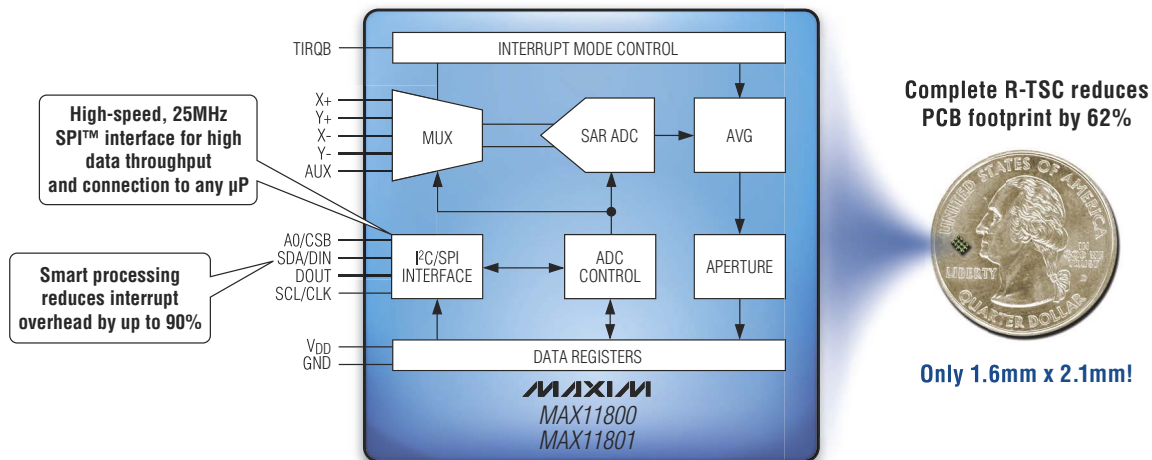
Figure 1 Arranging LEDs in a cross-point array and adding a transistor to each column show that the duty cycle of Charlieplexing is similar to that of standard multiplexing.



Industry's smallest and smartest touch-screen controllers

More features than the competition, yet in a smaller package

The MAX11800/MAX11801 are resistive touch-screen controllers (R-TSCs) available in TQFN and ultra-small WLP packages. These devices have advanced operation modes and other digital functionality that reduce bus loading of the system/applications processor, plus numerous advantages including the reduction of necessary board space and the ability to connect to any host μ P. Ideal applications include mobile phones, MP3 players, personal navigation devices (PNDs), UMPCs, office printers, gaming consoles, and digital photo frames.



- Autonomous mode minimizes application processor involvement, reduces bus loading and host processor resource needs
- Aperture mode* implements spatial filtering, enables recognition of complex touch events
- FIFO increases effective throughput
- Programmable scan rate allows tradeoff between power and accuracy
- Data-tagging identifies type of measurement (X, Y, pressure) and type of touch event
- Programmable averaging reduces noise

Part	Interface	Temp Range (°C)	Package (mm x mm)	Price* (\$)
MAX11800ETC+	SPI	-40 to +85	12-TQFN (4 x 4)	1.55
MAX11800EWC+			12-WLP (1.6 x 2.1)	1.51
MAX11801ETC+	I2C		12-TQFN (4 x 4)	1.50
MAX11801EWC+			12-WLP (1.6 x 2.1)	1.46

SPI is a trademark of Motorola, Inc.

*Patent pending.

†1000-up recommended resale. Prices provided are for design guidance and are FOB USA. International prices will differ due to local duties, taxes, and exchange rates. Not all packages are offered in 1k increments, and some may require minimum order quantities.

www.maxim-ic.com/MAX11800-info



www.maxim-ic.com/shop



www.avnet.com



For free samples or technical support, visit our website.

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

15V_{IN}, 4MHz Monolithic Synchronous Buck Regulator Delivers 5A in 4mm × 4mm QFN – Design Note 467

Tom Gross

Introduction

The LTC[®]3605 is a high efficiency, monolithic synchronous step-down switching regulator that is capable of delivering 5A of continuous output current from input voltages of 4V to 15V. Its compact 4mm × 4mm QFN package has very low thermal impedance from the IC junction to the PCB, such that the regulator can deliver maximum power without the need of a heat sink. A single LTC3605 circuit can power a 1.2V microprocessor directly from a 12V rail—no need for an intermediate voltage rail.

The LTC3605 employs a unique controlled on-time/constant frequency current mode architecture, making it ideal for low duty cycle applications and high frequency operation. There are two phase-lock loops inside the LTC3605: one servos the regulator on-time to track the internal oscillator frequency, which is determined by an external timing resistor, and the other servos the internal oscillator to an external clock signal if the part is synchronized. Due to the controlled on-time design, the LTC3605 can achieve very fast load transient response while minimizing the number and value of external output capacitors.

The LTC3605's switching frequency is programmable from 800kHz to 4MHz, or the regulator can be synchronized to an external clock for noise-sensitive applications.

Furthermore, multiple LTC3605s can be used in parallel to increase the available output current. The LTC3605

produces an out-of-phase clock signal so that parallel devices can be interleaved to reduce input and output current ripple. A multiphase, or PolyPhase[®], design also generates lower high frequency EMI noise than a single-phase design, due to the lower switching currents of each phase. This configuration also helps with the thermal design issues normally associated with a single high output current device.

1.8V_{OUT}, 2.25MHz Buck Regulator

The LTC3605 is specifically designed for high efficiency at low duty cycles such as 12V_{IN}-to-1.8V_{OUT} at 5A, as shown in Figure 1. High efficiency is achieved with a low R_{DS(ON)} bottom synchronous MOSFET switch (35mΩ) and a 70mΩ R_{DS(ON)} top synchronous MOSFET switch.

This circuit runs at 2.25MHz, which reduces the value and size of the output capacitors and inductor. Even with the high switching frequency, the efficiency of this circuit is about 80% at full load.

Figure 2 shows the fast load transient response of the application circuit shown in Figure 1. It takes only 10μs to recover from a 4A load step with less than 100mV of output voltage deviation and only two 47μF ceramic output capacitors. Note that compensation is internal, set up by tying the compensation pin (ITH) to the internal

LT, LTC, LTM, PolyPhase, Linear Technology and the Linear logo are registered trademarks of Linear Technology Corporation. All other trademarks are the property of their respective owners.

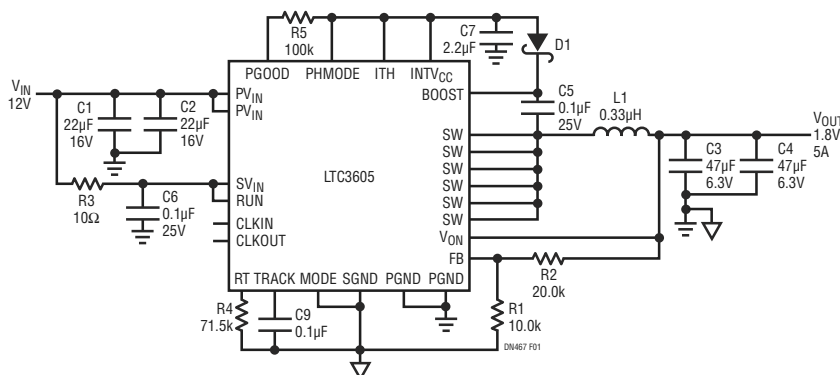


Figure 1. 12V to 1.8V at 5A Buck Converter Operating at 2.25MHz




Figure 2. Load Step Response of the Circuit in Figure 1

3.3V regulator rail (INTV_{CC}). This connects an internal series RC to the compensation point of the loop, while introducing active voltage positioning to the output voltage: 1.5% at no load and -1.5% at full load. The hassle of using external components for compensation is eliminated. If one wants to further optimize the loop, and remove voltage positioning, an external RC filter can be applied to the ITH pin.

1.2V_{OUT}, 10A, Dual-Phase Supply

Several LTC3605 circuits can run in parallel and out of phase to deliver high total output current with a minimal amount of input and output capacitance—useful for distributed power systems.

The 1.2V_{OUT} dual-phase LTC3605 regulator shown in Figure 4 can support 10A of output current. Figure 3 shows the 180° out-of-phase operation of the two LTC3605s. The LTC3605 requires no external clock device to operate up to 12 devices synchronized out of phase—the CLKOUT and CLKIN pins of the devices are simply cascaded,




Figure 3. Multiphase Operation Waveforms of the Circuit in Figure 4. The Switch Voltage and Inductor Ripple Currents Operate 180° Out of Phase with Respect to Each Other

where each slave's CLKIN pin takes the CLKOUT signal of its respective master. To produce the required phase offsets, simply set the voltage level on the PHMODE pin of each device to INTV_{CC}, SGND or INTV_{CC}/2 for 180°, 120° or 90° out-of-phase signals, respectively, at the CLKOUT pin.

Conclusion

The LTC3605 offers a compact, monolithic, regulator solution for high current applications. Due to its PolyPhase capability, up to 12 LTC3605s can run in parallel to produce 60A of output current. PolyPhase operation can also be used in multiple output applications to lower the amount of input ripple current, reducing the necessary input capacitance. This feature, plus its ability to operate at input voltages as high as 15V, make the LTC3605 an ideal part for distributed power systems.

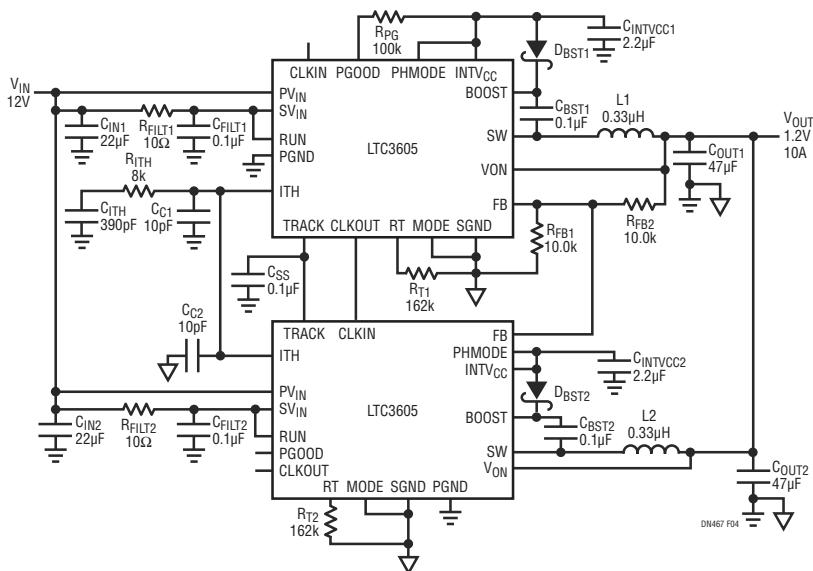


Figure 4. 12V to 1.2V at 10A 2-Phase Buck Converter

Data Sheet Download

www.linear.com

For applications help,
call (408) 432-1900, Ext. 3229

DAC calibrates 4- to 20-mA output current

Ronald Moradkhan and Steven Lau, Maxim Integrated Products, Sunnyvale, CA

Industrial controls make heavy use of 4- to 20-mA current loops to transmit process measurements because current loops retain information in the presence of noise and changes in loop voltage. The loop circuit requires proper calibration to ensure accurate readings. The circuit in **Figure 1** calibrates the loop by generating a current in response to a control voltage:

$$I_{OUT} = \frac{V_{CONTROL}}{R_{SENSE} \times K_{CSA}}$$

where I_{OUT} is the output current, $V_{CONTROL}$ is the control voltage, R_{SENSE} is the sense resistance, and K_{CSA} is the gain of the current-sense amplifier—20 in this case. The circuit comprises IC_2 , a Maxim (www.maxim-integrated.com) MAX5304 DAC; IC_3 , a MAX4376T current-sense amplifier; IC_4 , a MAX420 op amp; and Q_1 , an N-channel IRFL4105 MOSFET. The op amp lets the control voltage set the output current because it forces the voltage on

the negative input equal to that on its positive input. The output current depends on the value of the sense resistor, the gain of the current-sense amplifier, and the control voltage.

The DAC provides the control voltage that lets you automate the calibration procedure. By selecting the right value for the sense resistor and by using a suitable resistor divider for R_1 and R_2 at the output of the DAC, you can adjust the circuit's output to 4 mA when the DAC's digital input is zero-scale and 20 mA when the digital input is full-scale. **Figure 1** shows the component values you need to achieve that condition.

With a zero-scale digital input, the DAC output is 0V and the resistor divider produces 0.6V at the op amp's positive input, forcing the output current to 4 mA. With a full-scale digital input, both the DAC output and the midpoint of the resistor divider are at the 3V reference voltage, forcing the output current to 20 mA. A transfer curve relates the output current to the control voltage (**Figure 2**). **EDN**

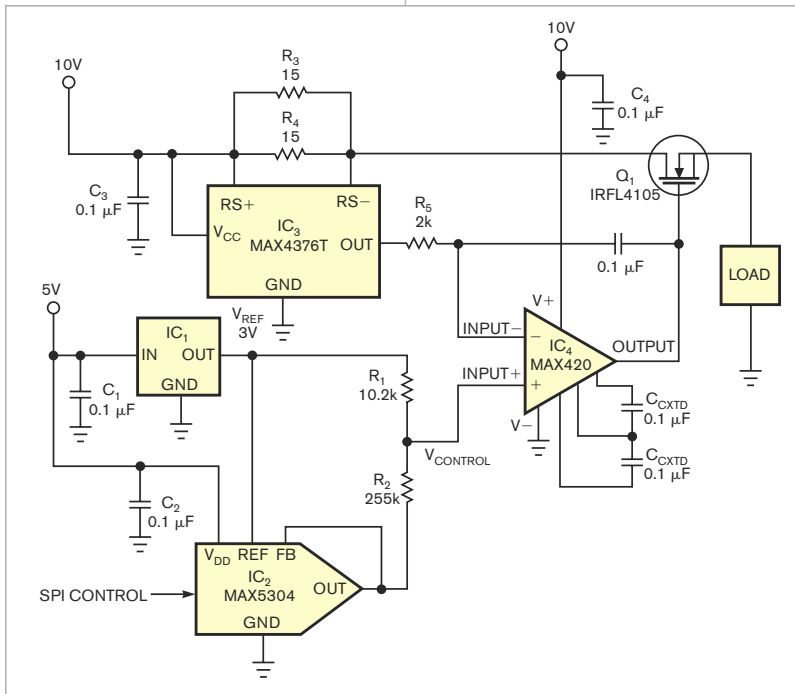


Figure 1 This DAC-controlled 4- to 20-mA transmitter allows digital control of the loop current.




Figure 2 The circuit in Figure 1 produces a linear output current versus digital control voltage.

Alarm tells you to close the refrigerator door

Boris Khaykin, TRW Automotive, Livonia, MI

The circuit in **Figure 1** is a simpler and safer device than a sim-

ilar one I recently read about (**Reference 1**). A few years ago, I built the

circuit that this Design Idea describes, and the gadget still operates with the original 9V battery. The circuit operates by sensing a decrease in resistance of photocell PC_1 that results from light in the refrigerator when its door is open. A counter is in a reset state when PC_1 is in the dark, and its resis-

tance is greater than 30 k Ω . Usually, the dark resistance is greater than 200 k Ω , and current consumption at this state is less than 40 μ A. Oscillator-counter IC₁ starts counting when PC₁'s

resistance is lower than 15 k Ω —that is, when the light bulb in the refrigerator is on. Then, in 20 seconds, it turns on a buzzer for 20 seconds or until someone closes the door. The current at this

state is approximately 2.5 mA.

You can use almost any photocell, such as the Jameco (www.jameco.com) 202403 CDS0018001 with 200-k Ω dark and 3-k Ω light resistance.

This circuit uses a RadioShack (www.radioshack.com) 273-074 buzzer. You can use any similar piezoelectric buzzer with an operating dc voltage of 1.5 to 15V. V₁ can be as low as 3V. The trade-off is that using a voltage this low gives you longer battery life but lower volume of sound. **EDN**

REFERENCE

1 Babu, TA, "Alarm Sounds When Refrigerator Door Remains Open Too Long," *Electronic Design*, March 26, 2009, pg 46, <http://electronicdesign.com/Articles/ArticleID/20806/20806.html>.

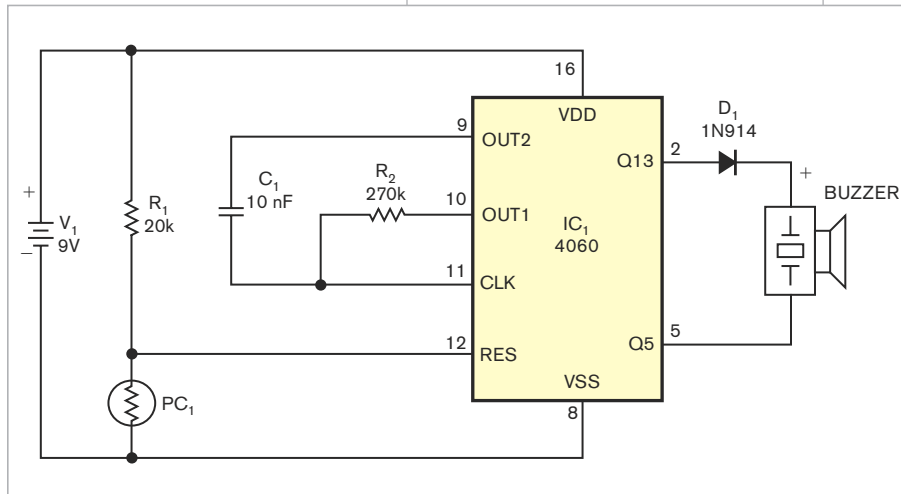


Figure 1 This gadget, placed inside a refrigerator, sounds an alarm when the refrigerator door is open for more than 20 seconds.

Battery Management at its best

AS8510 Battery Sensor Interface

- ▶ Ultimate precision for small signals
- ▶ 2 independent acquisition channels in one device
- ▶ Configurable to fit a wide range of present and future battery management applications



ae austriamicrosystems

Raleigh (919) 676 5292 • Detroit (248) 202 7100
www.austriamicrosystems.com

a leap ahead in analog



No need to duplicate yourself

**BuyerZone saves you time and money
on COPIERS AND MORE.**

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

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

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

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

BuyerZone
Where Smart Businesses Buy and Sell

A division of
 Reed Business Information.

alliedelec.com

Product Info

- 1 **Photos** - Detailed product photographs and expanded views help identify the exact part you need
- 2 **Ratings** - View product reviews and ratings from other users
- 3 **Availability** - Check stock availability and stock on order, so you know exactly when your product will ship
- 4 **Recommended Accessories** - Compatible accessories are prominently displayed next to a product
- 5 **Datasheets** - Download manufacturer datasheets for in-depth product data
- 6 **Overview** - Brief description of the product and its capabilities
- 7 **Specs** - Quick view of the product's technical specifications
- 8 **Similar Products** - Easily find similar products when the search result doesn't quite meet your needs



alliedelec.com



1.800.433.5700

YOU RESEARCH. YOU ORDER. WE DELIVER.

Browse over 1.2 million electrocomponent product solutions and all the information you need to place your next order.



A SPECIAL EDN SECTION

INNOVATORS

EDN



Innovation: genius,
practice, or luck?

Application engineers:
serving the customer

Improving on
PCB design

The Newest Microcontrollers

8051 Cortex™ ARM®
Z8® Flexis™ PSOC® Encore!®
AVR® MSP430 HCS08
Coldfire® PIC® Z80® ST7
HC12 ST6 HC08
HCS12 ST9



New Products

*EDN 2008 Innovation Award
Winners and Finalists:*



MICROCHIP

PIC32 Microcontrollers
[www.mouser.com/
microchip_pic32](http://www.mouser.com/microchip_pic32)



**TEXAS
INSTRUMENTS**
Authorized Distributor

Piccolo™ Microcontrollers:
TMS320F2802x
[www.mouser.com/
tipiccolomcus](http://www.mouser.com/tipiccolomcus)




ATMEL

AVR® XMEGA™
Microcontrollers
[www.mouser.com/
atmel_xmega](http://www.mouser.com/atmel_xmega)

The ONLY New Catalog Every 90 Days

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



a tti company

*The Newest Products
For Your Newest Designs*



www.mouser.com

Over A Million Products Online

(800) 346-6873

BY RICK NELSON, EDITOR-IN-CHIEF

Innovation: genius, practice, or luck?

Three columnists recently touched on this issue, and reached three different conclusions.



David Brooks, writing in *The New York Times*, contends that practice makes perfect (**Reference 1**). He begins by posing “certain paragons of greatness—Dante, Mozart, Einstein—whose talents,” certain romantics would claim, “far exceeded normal comprehension, who had an otherworldly access to transcendent truth, and who are best approached with reverential awe.” Today we know better, Brooks claims, noting, “In the view that is now dominant, even Mozart’s early abilities were not the product of some innate spiritual gift. His early compositions were nothing special.” Mozart, according to Brooks, owed his talent to a father who made him practice.

Is practice, whether with or without innate talent, enough? Not according to Robert H Frank, a Cornell University economics professor, writing in the *The Huffington Post* (**Reference 2**). “There’s no question that hard work and talent make someone more likely to achieve economic success,” he writes. “But for every successful person ... there are hundreds of others who are just as talented and work just as hard, yet earn only modest incomes.” He concludes, “Even talent and the inclination to work hard are themselves heavily dependent on chance.”

In engineering, given sufficient talent and—if Frank is right—luck, what might practice do for us? It might allow us to memorize equations and programming languages, for example. Is rote memorization helpful? Not according to AC Grayling, writing in *The Guardian* (**Reference 3**): “It is a common presumption that if people know a lot, they must be intelligent. Anyone who can reel off capital cities or count to 10 in several languages ... is counted a bright spark.” But, he continues, “There are plenty of very bright people who do not know the world’s capitals and cannot count in other languages, because they have never had a chance to learn them. ... By the same token plenty of people know lots of facts without being creative, thoughtful, quick-witted, humorous, and perceptive—the marks of true intelligence.”

How do these questions relate to innovation? Grayling doesn’t use the term, but for him *innovation* would seem to

be synonymous with intelligence: “Intelligence is a matter of output, not scores in a test. Einstein was unsuccessful at school and no great shakes as a mathematician, but he was creative and insightful. ... A vivid interest in things, and an active desire to understand more about them, is a major characteristic of intelligence.”

Where does this writing leave us? I’m inclined to agree with Frank, but he provides no recipe for success. We can’t revise our genes or command luck to smile on us. All we can do is follow Brooks’ advice and practice. But we can infer one suggestion from Grayling: Organizations wishing to foster innovation must create an environment that fosters in its employees a vivid interest in and active desire to innovate. That’s not an easy task in this day of budget cuts and layoffs—which can distract even the most innovative engineers, wherever their innovative spark originates.

One way to learn to focus is to study how successful innovators operate. The stories of three of them appear in this special section. Jim Williams of Linear Technology warns of the dangers of the rigid setting of goals, which individuals might meet at the expense of companywide innovation. Chuck Grant of Cadence attributes his innovative career in part not to a narrow focus on one area but rather to the chance to gain experience in test, marketing, training, sales, and customer support. And when asked how he became an innovative engineer, Cadence’s Ken Wadland says simply, “I don’t follow rules.”

Read the profiles in this issue, and review our previous “EDN Innovators” profiles at www.edn.com/innovators. Let us know what you think.

REFERENCES

1. Brooks, David, “Genius: The Modern View,” *The New York Times*, May 1, 2009, www.nytimes.com/2009/05/01/opinion/01brooks.html.
2. Frank, Robert H, “Success and Luck,” *The Huffington Post*, May 2, 2009, www.huffingtonpost.com/robert-h-frank/success-and-luck_b_195162.html.
3. Grayling, AC, “Knowledge and genius,” *The Guardian*, May 1, 2009, www.guardian.co.uk/commentisfree/2009/may/01/genius-knowledge-iq-tests.

Breaking Through Performance Barriers

Speedster™ is the World's Fastest FPGA



New Speedster 1.5 GHz FPGAs

Breaking through the speed barriers of traditional FPGAs, fully reprogrammable Speedster devices from Achronix Semiconductor are capable of operating at 1.5 GHz peak performance.

Speedster's unique 40 lanes of 10.3 Gbps embedded SerDes and four independent 1066 Mbps DDR2/DDR3 controllers enable extremely high I/O throughput to match the device's outstanding internal performance. At 3X faster, Speedster opens up new worlds of application design previously unavailable to engineers using traditional FPGAs.

Familiar Silicon. Familiar Tools. Fast Time-to-Market. Speedster uses familiar LUT-based fabric and standard synthesis and simulation tools, so designers can use standard RTL.

High-Performance Applications

- Networking
- Telecommunications
- Encryption
- High-performance computing
- Digital signal processing
- Test and measurement
- Military and aerospace

See what's possible:
www.achronix.com/speedster

achronix
SEMICONDUCTOR CORPORATION



Jim Williams is a staff scientist at Linear Technology Corp. Longtime EDN readers recognize Williams as a vital contributor of analog-themed articles over the last 30 years. Williams has worked for 27 years at Linear Tech and has held previous roles at National Semiconductor and the Massachusetts Institute of Technology. EDN Technical Editor Paul Rako asked Williams about the changing role of application engineers—from essentially a support function of salespeople to full-blown system designers. Williams explains how a modern view of application engineering can ensure that your company will remain innovative.

Application engineers: serving *the* customer

*Jim Williams
on the
changing role
of application
engineering*



Application engineers have a different job today from in the old days. How have things changed in the years you've been in the business?

I'll direct my comments at analog-application engineers. The big change is that the customer who calls you up or otherwise contacts you is generally not an analog designer. They knew a lot more 20 years ago about what they were doing than you ever would. They had specific questions on specific line items in a data sheet.

Those people are still out there, but they're in the minority now. The customer who's calling you up today, as a rule, is not an analog designer but needs to access analog technology. So rather than a specific question about a specific part or a specific characteristic of a specific part, they'll call you up and they'll tell you, "I have so much space, so much time, so much power, and so much money. What do I do?" They're not asking you detailed technical questions. They're asking you what to do.

That means you're in the service business now; you're providing design services. That also means

that the typical analog customer is not looking at your parts catalog. They're coming to you to solve a problem. They're ignorant of analog technology, but that's no crime. They've got 500 line items on their board to check off. They're not analog designers. They just need to access analog technology, and that's what you're going to do for them. They're not coming to you to ask you questions about components. They're coming to you to solve a problem.

So 20 years ago, what was an application engineer doing?

The classic application-engineering job 20 years ago was [working on] data-sheet support and writing app notes about existing parts. ... [They provided] high-level support for existing parts—how to use existing parts—in the applications they [felt they were] likely to be used in. So it was largely a reactive job.

And now what do you see?

That reactive component is still there, although in this company, the designer writes the data sheet; we insist on that. He may seek assistance with the application engineer, but the designer is responsible for getting the data sheet out the door. Those components that application people were engaged in 20

“The fundamental responsibility of an analog-application engineer is to understand his customer’s problem and provide a solution.”

years ago are still there. But in a well-run analog company, they’re second- or third-tier parts of the job. The fundamental responsibility of an analog-application engineer is to understand his customer’s problem and provide a solution.

For me, that solution may involve this company’s products. I’ve also provided solutions to customers that use this company’s products and competitors’ products because those are the realistic solutions. I’m the customer advocate, and I’m seeking a long-term relationship with a customer. But what I’m doing and what I hope the other application engineers are doing around this company is ... servicing those customers who want to access analog technology but don’t have the expertise to do it and don’t have the time to develop that expertise.

That’s a large part of the job. Part and parcel of it is that you’re garnering a feeling for what needs to be built next. You’re defining new products. And an application note today that’s written by an application engineer may involve support of an existing product, but it’s much more likely to be written around current technical issues.

Do you distinguish between an application note written by an application guy versus the application section of a data sheet written by a design engineer?

There’s a complete difference. An application note, if it’s properly written, stands a good chance of having a 10- or 15-year lifetime because it’s issue-centered; it’s not product-centered. There’s product in it that is illustrating various types of solutions, but, if the application note is really well-written and well-thought-out, it’s applicable even when the parts are long gone from the scene. You’re writing about issues. You’re writing about approaches.

Do you make a distinction between a factory-application engineer and

a field-application engineer? Are those two different jobs, or are those roles converging?

They’re two different jobs, but there’s overlap. The biggest single difference between the two is that the factory-application guy is much more future-oriented, and he’s much more laboratory-based because he needs a laboratory. The field guy, unless he’s got a laboratory at home, hasn’t got access to a laboratory, or, if he does, it probably isn’t as well-facilitated as [that of] the factory guy. So the factory guy can spend a lot more time on research and development for future products. The field guy is busy servicing existing customers with existing problems, which is important, obviously, ... but the factory guy’s not doing that. If he’s focused on servicing existing customers with existing products and getting existing sockets, he’s mortgaging his future. He’s eating his seed corn and developing nothing.

That’s the fundamental difference between the two positions: The factory guy has the luxury of being able to spend more time on futures and on issues. If factory engineers are pushed toward devoting 90% of their effort to getting sockets, you’ve successfully mortgaged the company’s future. Whereas the field guy is with customers all the time. They’ve got today’s problems surrounded by today’s parts that they can get their hands on now. So they’re the heroes of today. But in a well-run application effort, the heroes of tomorrow are in the factory. That doesn’t mean the factory doesn’t provide backup for the field guy when things get sticky. You’ve got to do that, but that plays into tomorrow, because you look at what comes over the wall that the field people can’t handle, and you see trends, which suggest products.

And it doesn’t physically mean that the guy works at a factory. It could be that he works at a local office but is acting as a factory-application guy, right?

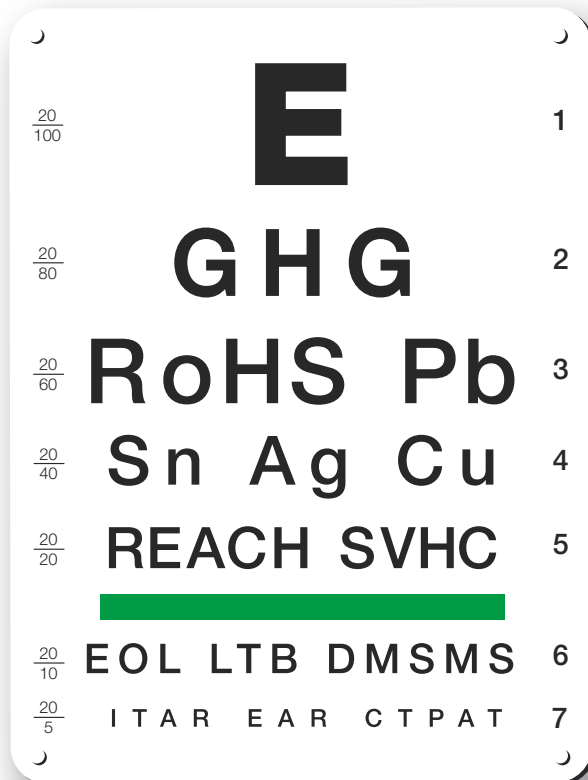
We’re not saying he has to be at the headquarters. ... He needs a lab, and he needs time. The greatest leverage any engineer has in doing his job is time. And, nominally, the field guy is spending the bulk of his time working with customers on existing problems with existing products. The factory guy should be spending the bulk of his time thinking about what needs doing in the general sense.

Do the application groups for field and factory have a responsibility to show customers how to use software tools?

Well, let’s talk about tools. There’s LT Spice. There’s Webench. There are all the various software tools. But there are also screwdrivers, shears, curve tracers, and X-acto knives. The application engineer’s job is to emphasize and show the customer what tool is appropriate for what task. For some tasks, LT Spice is the appropriate tool. For other tasks, a cut-down X-acto knife is the appropriate tool, and that should be reflected in applications. What’s the real down-in-the-dirt way to get from A to Z? Is it Spice? Is it an X-acto knife? Is it a breadboard? Is it cutting copper clad? Is it some fusion of all of those? Tool development, tool use, writing it up, how to measure, how to simulate, when to measure, when to simulate, where to simulate, and when to cut copper clad are all part of the application engineer’s job in educating the customers on how to solve their problems.

Customers like data sheets. Customers like app notes. Customers like publications that educate them. Customers like advice over the phone. Customers like software programs that help them design. But what customers love, what sells a product like nothing else, is that simple little cardboard box arriving in the mail with a breadboard that works when they plug it into their system. No sales pitch, no sales routine, no software program, no phone conversation, no e-mail sells products like a working

Focusing on a green supply chain?



see clearly

IHS knows green

Whether it's compliance with regulations, ensuring sustainability and high reliability, or managing supply chain volatility, IHS understands that green means different things to different people.

Supply chain, sourcing, and engineering professionals rely on IHS information to comply with regulations, demonstrate social responsibility, and maintain uninterrupted product flow.

Design, manufacture and trade anywhere

Providing expertise and information designed to manage product lifecycles and enable green performance, IHS critical insights offer decision-making agility across the supply chain.

- Eco-Friendly and Lead-free Design
- REACH Chemical Performance
- Greenhouse Gas footprint
- High Reliability and Sustainability
- Obsolescence and PCN/EOL Management
- Global Standards, Trade, and Security

Contact IHS and go green
parts.ihs.com/greensupplychain

Worldwide: +1 888 752 0334
US/Canada: +1 303 397 2892



The Source
for Critical Information and Insight™

“There’s no patent on intelligence. Anybody can learn to do anything. It’s just people haven’t got the time.”

breadboard mailed to a customer. I can’t say that loud enough or long enough. Nothing sells like a board that works in the customer’s system. That is the ultimate analog-application support. Nothing beats that.

What about compensation of factory- and field-application engineers? The field people get bonuses for filling sockets. Are you against having factory-app people compensated by filling a socket?

This might get me into trouble in some sectors around here, but I’m against all forms of goal-setting for factory-application engineers because they’ll pursue those goals and let other things slide by the wayside that could be the future success of the company.

If you sit down and you agree with somebody that you’re going to be rewarded if you do this, this, and this, then that’s what a lot of people are going to be tempted to do. Then they’re going to see this, this, and this, which look interesting and, potentially, fruitful, but they know if they do this, this, and this, they’ll get rewarded.

It may be that goal-setting in a field engineer’s job seems unavoidable, because they’re servicing existing customers with existing parts, and that’s the only metric that management can use. But in a factory job, which is nominally R&D-based, if you line up a bunch of goals for somebody and tell them that’s the way to success in this company, you’re destroying the company’s future. You’ll have a bunch of people who are successful according to their goals and a company that stalls in innovation.

I’m curious how you see application engineering helping with reference designs in general.

There are two ways applications can export design expertise—which, ostensibly, I hope they have—into the world. One way is what I call “the noble way.” That’s through application notes that

are topic-centered. The thesis there is that an educated customer is a better customer, and an educated customer will come back for more, I hope, to your company.

The noble way says you educate the customer through publication. The less noble but effective way is the reference design where the customer says, “I haven’t got time to come up to speed. You’re providing me with these publications, but I haven’t got time to come up to speed. I need that little cardboard box with the breadboard.”

That is a reference design. It’s essential because customers are out of time. They haven’t got the time to come up to speed to execute the architectural issues in a circuit themselves. And there’s also a real marketing issue. Through-hole is dead. We can breadboard here at the factory, but, by and large, customers can’t try stuff anymore. The parts are too damned small.

You’re going to have to take the part you’re trying to sell for the problem you’re trying to solve, and you’re going to have to incarnate it on a board-level reference design, because the customers haven’t got the time to develop the expertise they need to execute a design and because they have a hard time breadboarding and playing.

Reference designs and demo boards are important. I would say of the demo boards that get adopted and the reference designs that get adopted, a third get used pretty much the way they are, and two-thirds get used as a place to start. The customer comes back to you and says, “I fired up your demo board and looked at your reference design. It seems to work pretty well, but I needed twists here, there, and the other place.” But the point is, the reference design provides an advanced place to start talking from.

There’s something else that comes to mind with reference designs. In most analog companies 20 years ago, the word from on high was “Don’t get involved in your customer’s design.

Don’t accept responsibility for the design; there are potential legal problems.” Many companies went out of their way to tell their application guys, “Don’t get knee-deep in a design and accept responsibility.” Today, accepting responsibility is a sales tool. That’s turned 180°. There is a range of solutions available from a number of competitors, and, if you want to distance yourself from the competition, you’re going to have to accept responsibility. It’s based on the complexity of the product, and it’s based on its availability in one form or another.

So this way is the new way. That’s how you see application engineering developing?

The analog business is a service business. I can walk into a customer’s facility and hand them data sheets and app notes and parts. That may have been a sale 25 years ago. It isn’t anymore. They haven’t got the time to read the data sheet. They haven’t got the time to develop analog expertise. They may not have the inclination; they’ve got other things to do, and they can’t breadboard with the parts. They say, “So your parts are interesting, your app notes are pretty, your data sheets are pretty, but I need something I can clip into my system that’ll work.” That’s a complete shift from 25 years ago, when management was telling application engineers, “Be very careful about getting knee-deep in your customers’ designs. We don’t want to get sued for a field recall.” That’s 180° out of phase with what’s going on today. You’re ... over your knees, over your head in your customers’ designs.

With many companies, there’s no analog team. There’s no analog engineer.

They’ve got other things to do. There’s nothing magic about the analog field. It’s just people have other things to do. There’s no patent on intelligence. Anybody can learn to do anything. It’s just

Design with Light

Reliable




TURN ON THE
**POWER
OF AVNET**

Sustainable



Durable



Avnet lights the way to reliable, durable, and sustainable LED Solutions.

Choosing LED technology for your design is only the first step on the path. Avnet Electronics Marketing helps light your way to selecting LEDs that meet your reliability, visibility, and availability requirements. At each stage of the design cycle, our team of illumination-focused engineers gives you access to the latest information on LED products, ensuring you find the right solution to fit your specific design needs. When tackling the challenges of thermal management, power driver stage and secondary optics, our experts are your source for leveraging the benefits of LED technology.

As a unit of Avnet Electronics Marketing, LightSpeed brings together the world's foremost LED, high-performance analog and optical/electromechanical manufacturers along with best-in-class technical expertise and supply chain management services – affording you quicker time to market.

Working together, we can help you bring your ideas to light.

For more information and to view the latest issue of Light Matters visit us at: www.em.avnet.com/lightspeed

- > SECURITY
- > MEDICAL
- > INDUSTRIAL
- > AUTOMOTIVE
- > ARCHITECTURE
- > SPECIALITY LIGHTING

LightSpeed



“If you can’t explain how something you did works to a general audience in a general kind of way, you don’t know how it works.”

people haven’t got the time. You walk into a customer’s facility, and you look at their blackboard, if they have one; there’s 50 items to be checked off there, and four of them are analog.

They don’t need to develop in-house analog expertise. They’ve got analog companies who will do it for them. That doesn’t mean there aren’t companies out there that are well-steeped in analog expertise. It just means that most of the products being built today that use analog technology are being built by companies that are not steeped in analog-design techniques and don’t have to be. That’s what they look to us for.

Where do we find this new breed of application engineers?

The far side of Alpha Centauri. I don’t know where the hell you find them. You like to think you can breed them, but it takes a long, long time and a lot of burning of fingertips. To some extent, you can breed them. To some extent, they find themselves. You still find kids coming out of college who’ve been playing with electronics since they were in grade school. Those people still exist. They’re aberrant. They’re weirdoes, but they’re wonderful weirdoes.

You still find those people who got addicted early and found a way, despite the surface-mount revolution, to hack electronics in grade school and high school. You do find people who’ve been doing board-level design, sometimes your own customer, who wants to come over—not often, but it happens. Those are the two major sources: lifelong circuit freaks and system guys who’ve been working for board-level-product houses who want to come over and do this kind of work in a semiconductor company. Also, to some extent, inbreeding within the company [and] mentoring [help]. But there is no official, if you will, path toward finding these people. It’s quantum mechanics. If you line up enough people on one side of the fence, experienced system designers who want to work in a semiconductor company

will appear on the other side of the fence. No one knows quite how. It’s a tunneling process.

Do you think application engineers should be writing magazine articles as you do?

Application engineers should definitely write articles, for a couple of reasons. The most obvious reason is that it’s good for the company’s image, but, more important, if you can’t explain how something you did works to a general audience in a general kind of way, you don’t know how it works. It’s a great way to test your own level of understanding of what you just finished doing.

Writing is essential. It’s an essential part of an application engineer’s job and communication skills. It’s harder to find people with really good communications skills than with technical skills. Writing is important because it tests your knowledge of your ability to understand what you really did. It puts your company in a good light, and, career-wise, it puts you in a good light. I can’t imagine working on a difficult problem over a protracted period of time and then not writing it up. Make your contribution. If you’ve got something you think is worth talking about, make your contribution.

I see a difference between an artist and a tradesman. An artist will say, “Here’s the palette, here’s the paint I used, and there’s my canvas. Have at it.” They share all that important stuff, whereas a tradesman would say, “Oh, this is my secret little thing, and I keep it to myself.” You take kind of a higher-plane view of things and act like an artist. Why not keep everything secret?

I think [the “father of the atomic bomb,” Robert] Oppenheimer was right. There are no secrets. The only secret to the atomic bomb is that it works.

Once a skilled technocrat sees that

something can be done, the only question is how. There’s no advantage to secrecy. Whatever advantage you’ll get by protecting some body of knowledge that you have is dwarfed by the goodwill and the good orders and the sense that you’re a problem-solving ally when you disseminate that knowledge.

So customers at system companies appreciate the openness of a modern application engineer?

They’re looking for somebody they can trust technically as a partner. Most technological trade secrets are short-lived at best. You’ll do your company, your professional reputation, your profession, and everybody else a lot more good by saying, “Here’s this problem I had, here’s how I solved it, and here are the results.”

There’s really no advantage to holding back. I don’t recall ever holding back a measurement or design technique from publication in an app note. Remember a number of years ago, when there were soft errors in memories? Intel figured it out. It was [caused by] alpha particles. Intel released it and gave it away, and a lot of people said, “Intel is crazy. Intel shouldn’t do that.” So you’re making friends, you’re generating credibility for yourself and your company, and [disseminating it obviates any] short-term benefit that you can gain by keeping things close to the chest.

There are certain manufacturing processes, trade secrets, and the fab, stuff like that, that you’re going to play close to the vest. But you’re not going to have to play them close to the vest for long because they’re going to be obsolete in six months or a year anyway. But in application engineering, measurement technique, circuit-design technique, whatever, you’ll generate more new customers by printing it all in an app note than you’ll lose because some competitor took your scheme and ran with it.—interview conducted and edited by Paul Rako

V-INFINITY

POWERING INGENUITY

CONTROL

INVENTIVE RESOLUTION

INTERCONNECT

QUALITY CONNECTION

SOUND

DIVERSIFIED SELECTION

THERMAL

COLLABORATIVE COOLING



CUI INC

IN PRODUCTION

SERDES-Capable FPGA Solutions for Wireless and Wireline Applications

LET LOW POWER GUIDE YOUR WAY.

INTRODUCING THE ECP3™. THE LOWEST-POWER
SERDES-CAPABLE FPGA IN THE INDUSTRY.



Built for high-speed communications protocols, the ECP3 helps you keep your power budget down, and applications features high. At half the power and half the price of competitive devices, now you can do more with less.

Learn more about the ECP3 family at:
latticesemi.com/ecp3






In 1995, OrCAD had the most popular PC-based PCB (printed-circuit board)-schematic tool. Another popular tool was available from PADS, which Mentor Graphics now owns. Engineers designed PCB schematics in OrCAD and laid out the board in PADS, a clumsy approach that still persists to this day. Rather than redesign its PCB-layout tool, OrCAD purchased Massteck, a small company offering a good layout tool and, more important, an autorouter that achieved 100% routing on complex boards. PCB designers Al Akermann and Gene Dancause, Massteck's founders, contracted with Ken Wadland, a computer-science guru, to do the autorouting algorithm and database architecture. Just as vital, they hired board designer Chuck Grant, who first did product testing and then eventually wrote software. Cadence in 1999 bought OrCAD. Sadly, OrCAD no longer supports the original Massteck-based OrCAD Layout tool, instead favoring OrCAD Editor, a stripped-down version of Allegro, Cadence's enterprise-class layout software. Wadland and Grant now work for Cadence in a Massachusetts design center. EDN recently had the chance to interview them.

Improving on PCB design

Ken Wadland and Chuck Grant: the innovators behind OrCAD Layout software



Ken Wadland

Where did you get your education?

CHUCK GRANT: I got an economics degree from the University of Rochester. Other than using some computers in those courses, I hadn't been interested in them. But the PCB-design business was growing by leaps and bounds, so they were looking for PCB designers. I just went in and I liked what I saw.

KEN WADLAND: In seventh grade, I started Wadland TV and Radio Repair. In junior high and high school, I started learning assembly-language programming. I went to Worcester Polytechnic Institute and got my degree in mathematics, but I took every computer-science course they offered. I got my master's degree and my doctorate at the University of New Hampshire. I went to Pittsburgh State College, and I created the [computer-science] curriculum there. I was the first chairman of the computer-science department.

What gave you guys the entrepreneurial spirit?

CG: There was a lot of evangelism in that we had a unique way of designing PCBs. The way we divided and conquered was to start from the most difficult part of

the design and completely finish that area. Once you clear out that most difficult area, you just move on to the next one, and you keep stringing them out until you get to the end. When you get to a clear area, you could pop in a via and use the other layer to finish it. We knew that the value added for PCB design is in autorouting. So we said, "Let's take the technique we use ... to divide and conquer as a manual designer and apply it to software."

KW: I couldn't resist the allure of money. [Autorouting] was the only reason I was attracted. At that point, there were only two hard problems in the world: AI [artificial intelligence] and autorouting. Autorouting is a mathematically complicated problem. Beyond being mathematically complicated, there are the aesthetics and the electrical-design rules and everything else that made it just a challenging product.

What made your innovations so effective?

CG: In a start-up company, you wear a lot of different hats. The hats I wore went from testing to marketing to training, sales, and customer support. One of the nice things about a small company is that you really can try all those

things out and see what you like and gain a lot of experience. We had some good PCB designers in that company who were proficient, as well. I decided that I would like to learn how to be a programmer, and, of course, I learned from the best, which is Ken, so that was a great opportunity for me. [Innovation] was driven by the PCB designers because we were used to the time pressures and deadlines and the need to get it right the first time. [We felt] “Hey, we get to invent our own, so let’s do what we want to do, and this is what it’s going to look like.”

KW: Well, we did not have a marketing group separate from our PCB designers. All the people in sales and marketing were PCB designers. So when they gave us the requirements, they understood what they were talking about.

What drove your product innovation?

CG: We really had to be evangelists in interactivity. The marketing machines of all the CAD companies were touting “push the button.” That was their idea, “push the button, and we’ll route your board for you.” Well, real PCB designers know that [approach] never works. You get as much automation as you can, and if you can guide it, all the better. The interactivity we built in was that you could choose where on the board you wanted to start routing. You could start in two different places and have them join up in the middle. So there was interactivity, and then there was the push and shove, which everybody copied. The idea that you could pick a single line, route it between two other lines, click, and [have] everything push out of the way was something that no one else had. We would let you see what it accomplished, and, if it didn’t finish, you might say, “Well, I can fix up that mistake over there, or I’ll take care of it.” But we wouldn’t break any design rules. We would leave unroutes where we couldn’t get them to fit. If you didn’t like it, you could go back to what you had before. We had an undo right from the beginning.

How do you become an innovative engineer?

KW: I don’t follow rules. I don’t like to

start a project unless at least one other person failed at it. Why waste your time on something easy? The downside is: You fail a lot. To do innovation right, you have to ignore the details until you have already established the concept, and then you see if the details work. If you’re too detail-oriented, you reject it too soon. You think of really outlandish things, most of which won’t work. You get a bunch of outlandish things, and then you start pruning them down to one that can actually be usable. If you can’t suspend disbelief for a long enough time, you can’t get these other ideas. One of the things I teach in my classes is: Never write code until you’ve found three ways of doing it. A lot of people [take] the first [approach] they think of. Chances are, it’s the wrong way or a less-than-perfect way. So you want to find three approaches to the problem, analyze all three of them, and then pick the best. People just don’t do that. They just don’t take the extra time to do that extra step.

Why was the program so intuitive and easy to use?

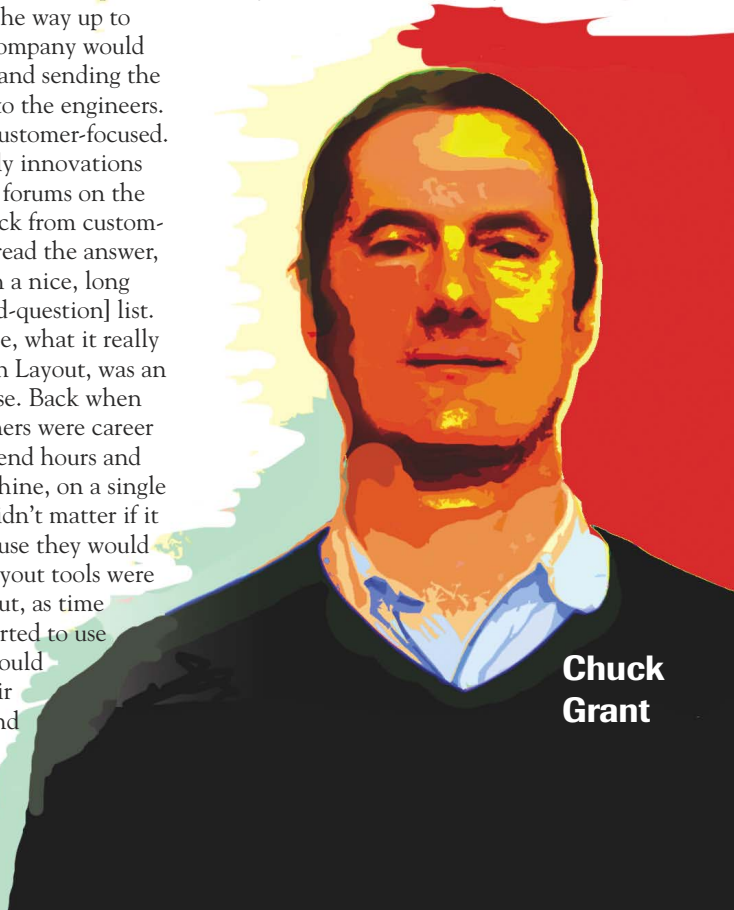
CG: [It was not] that some marketing spec had been translated through an engineer. It was a very direct experience. Those of us all the way up to the president of the company would be using the software and sending the feedback to Ken and to the engineers. We were very much customer-focused. I think one of the early innovations was to use interactive forums on the Internet to get feedback from customers. Everybody could read the answer, and you’d end up with a nice, long FAQ [frequently asked-question] list. The OrCAD influence, what it really brought to the table in Layout, was an emphasis on ease of use. Back when we began, PCB designers were career people, who would spend hours and hours on a single machine, on a single piece of software. It didn’t matter if it was complicated because they would learn it eventually. Layout tools were plenty complicated, but, as time went by, engineers started to use the software. So we would take into account their wishes and requests and tried to make it easier

to get into, and a lot of that happened after the purchase by OrCAD.

Users love the spreadsheet functions in OrCAD Layout. How did you come up with incorporating spreadsheets in the product?

CG: We were unique in that we leveraged existing technology before a lot of other people ever thought of doing that. We knew about Lotus 1-2-3, which we used on a PC for our own accounting, so we leveraged the idea of spreadsheets, not really realizing that they would take over the world later. It was just something that looked like it would work well for us. It allowed us a lot of flexibility. We could lock and unlock nets, establish library components or library symbols, and attach them to the components and all the things we needed to do. It all worked well with that spreadsheet interface, so we went to town with it and used it because it worked so well.

KW: Making a system that’s simple is very hard. One of the early winners was the Lotus 1-2-3 spreadsheet. There was some brainstorming between Gene Dancause and me, but it was my attempt to write a minimum number of lines of software to get maximum capability. We wanted the ability to make



Chuck Grant



"click"

"click"

"click"

Kellysearch clicks. Every time.

Kellysearch finds your next great supplier in three clicks or less. Talk about an elegantly engineered solution. Log on now at kellysearch.com.

Kellysearch.com

Everything you need, nothing you don't

 Reed Business Information

“Every time we got a board to 100%, we jumped up and down. Getting 99.97% was cool, but if you hit 100%? Wow, we were excited.”

changes to more than one object at once. So we came up with this spreadsheet because our database resembled SQL [Structured Query Language]. It was kind of a relational database under the hood, so producing a spreadsheet was a piece of cake. And then editing it through the spreadsheet was just a lot less software to write than writing dialogues for everything you wanted to do. It was basically being lazy.

How did you conceptualize the color-rule interface?

CG: The idea was that the program was WYSIWYG [what you see is what you get], so you could say, “OK, I want to turn on these objects.” This might be the etch and the pads on the top layer, and that’s my Gerber [file], and then you just print it out. So then you have a set of these [files], which are on a spreadsheet, and each one of them establishes a set of rules [of] what’s going to be visible for that artwork layer. If you have the assembly layer, the silk screen, and all of those things that also come with a PCB, ... that’s what will print out on your Gerber.

KW: The problem they posed to me is: “We want to see the colors in different ways at different times.” Some engineers wanted to see the top layer red; some wanted it green; and some wanted to see this [color] differ some of the time. I told them “I can’t do all that. How about if I just give you a color set and you can change the color set when you want?” The other thing that came into it is my campaign to protect people who are colorblind. My son is colorblind. So I refused to have any products that had fixed colors. Everything had to be adjustable. We came up with this idea. We call it “color set.” You set up all the colors the way you like, you save it, and then you pick from those color sets.

Was there any innovation due to the limitations of a 286 PC, and did you program in C?

CG: It was all C. The other innovation that we had was in the way we did software best practices. This [idea] was Ken’s innovation. We ran C, but we

isolated the database from the interactivity by using an object-oriented type of programming so that we prevented database corruption, which was a common problem. In most CAD systems, the programming was done so that you could manipulate an object directly. [We made] a copy of the object, let the user manipulate it, and then let the user commit that [copy] to the database. It would all happen under the hood. What that allowed us to do is to recover from errors without the user’s being aware of it.

KW: We knew from the beginning that a 64-[kbit] database was not going to be big enough. Back then, we had expanded memory and extended memory. One of them came in banks, and the other one worked only in 32-bit mode. We took this “get-put” methodology: ... Under the hood, when you say “get,” we’ll switch in the right bank, or we’ll switch to 32-bit mode and get the data for you. When you say “put,” we’ll do the same game back again. So the application lived in the 16-bit Windows environment, but it could access unlimited memory. We had a complete separation between the UI [user interface], the algorithms, and the database. I had used that concept when I worked with relational databases, so I understood the concept, and I applied that to an object-oriented database.

What was the biggest problem you felt you had to take on?

CG: The biggest problem in marketing was that the PCB designers don’t make buying decisions. They can influence them to some extent, but upper management and engineering make most of the decisions. It was difficult to convince people who’d never done PCB design that interactivity was critical because someone else would come in and say, “We don’t have interactivity; you just push the button.” Well, you can push the button on ours, too, but that’s not really what you want to do. I think we did eventually win that war, where interactivity became important to people.

KW: I’ll tell you the biggest problem that didn’t get solved. We tried to get a good place algorithm. We tried to have

an autoplacement algorithm, and we put enormous resources into it. It’s a simpler problem mathematically in the IC world, where everything’s symmetric. It’s a problem in the PCB world, where every component is a different shape. Things are not interchangeable. If one’s a quad pack, you can’t switch that with a two-pin surface mount. They’re incompatible. So we failed in placing. As far as I know, no one has ever had a successful placement tool in the PCB field.

Was there a personal victory you had?

CG: I will never forget one of those sessions when I had one important customer who absolutely needed to have the ability to move components and have minimization take place as they were moving. [We had] the first product I know of that had interactive component moving and then [PCB]-trace minimization as you were moving it. Initially, it was just a rat’s nest, and then we had a customer who said, “I have to see the real traces.” I definitely remember that one.

Was there a company victory you remember?

CG: In 1994, Jeff Hawes, [a former application engineer at Cadence, OrCAD, and Massteck], and I went to the PCB Design Conference and we entered MAX EDA into the competition. They were assigning points to PCB-design systems, and what they did was pretty comprehensive. You had to design a board, sight unseen. They gave you a net list and the parts. You had to design the board, create the netlist, bring it in, autoroute it, and print out Gerber artwork—all within two days. ... We won that contest in 1994 and 1995, and that was the year OrCAD bought us. When we found out that we had won that competition, we were elated.

KW: I’m not sure there’s any one [victory], but basically every time we got a board to 100%, we jumped up and down. Getting 99.97% was cool, but if you hit 100%? Wow, we were excited.

—interview conducted and edited by Paul Rako



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


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

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

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

REQUEST FREE QUOTES NOW!
Call (866) 623-5525 or visit
BuyerZoneQuotes.com

BuyerZone
Where Smart Businesses Buy and Sell

A division of
 Reed Business Information.

productroundup

MOTION



Motion sensor provides acceleration values in absolute-analog outputs

➔ Providing acceleration values in the form of absolute-analog outputs, the LIS352AX three-axis-accelerometer motion sensor operates with supply voltages of 2.16 to 3.6V and provides stability over temperature and time. Internal regulation voltage makes the measurements insensitive to variations in power-supply voltage that typically occur in battery-operated devices, such as mobile phones and other portable devices. This method enables space saving and reduced cost by connecting the sensor directly to the battery without a separate voltage regulator and by providing compatibility with the power-supply voltage that the application uses. The LIS352AX has high stability over a temperature range for zero-gravity offset and sensitivity, a $\pm 0.3\text{-mg}/^\circ\text{C}$ offset drift, and an accurate output over a full-scale range of $\pm 2g$. Available in a $3\times 5\times 1\text{-mm}$ plastic package, the LIS352AX motion sensor costs \$1.30 (10,000).

STMicroelectronics, www.st.com

MEMS motion sensor measures tilt, shock, and acceleration

➔ Measuring tilt, shock, and acceleration, the ultra-low-power, high-resolution ADXL346 iMEMS accelerometer operates at a 1.8V primary supply voltage. The digital three-axis smart-motion sensor provides selectable measurement ranges and measures dynamic acceleration from motion or

shock and static acceleration, allowing use as a tilt sensor. The device includes an on-chip FIFO-memory block storing as much as 32 samples of X-, Y-, and Z-axis data and offloads the FIFO function from the host processor. This process allows the host processor and other peripherals to enter sleep mode when not in use. Features include a 0.1- to 1600-Hz selectable bandwidth and 150- μA power-consumption ranges at 1600-Hz bandwidth and as low as 25 μA at less

than 10 Hz. The accelerometer measures dynamic acceleration with user-selectable measurement ranges of ± 2 , 4, 8, and 16g and employs built-in orientation sensing using simple register reads. Available in a $3\times 3\times 0.95\text{-mm}$, 16-lead LGA package, the ADXL346 costs \$3.04 (1000).

Analog Devices, www.analog.com

Voice-coil-motor driver uses I²C interface

➔ Suiting autofocus and zoom applications in camera-phone lenses, the A3907ECGTR voice-coil-motor driver has an I²C serial interface and a 2.3 to 5.5V operating range. The driver has a 102-mA maximum output current and operates at -40 to $+85^\circ\text{C}$. Programmed using the I²C interface in 100- μA increments, the I²C inputs set the internal 10-bit DAC output-voltage reference for linear-current control using an output-sink MOSFET and sense-resistor feedback. Additional features include a 10-nA input-supply current in sleep mode and I²C-compatible logic-input levels. Available in a $1.465\times 0.965\text{-mm}$, six-bump WLCSP-CG package, the A3907ECGTR costs 34 cents (1000).



Allegro Microsystems, www.allegromicro.com

INTEGRATED CIRCUITS

Drivers suit white LEDs with flash-mode capability

➔ Claiming an 80% efficiency in flash mode, the TB62730WLG

INTEGRATED CIRCUITS

and TB62750FTG step-up dc/dc converters for white LEDs integrate an NMOS transistor. Requiring one external inductor, the devices drive 800 mA using lithium-ion batteries. Features include overvoltage protection, a 2.8 to 5.5V supply-voltage range, and a temperature-shutdown function. The TB62750FTG also uses an integrated constant-current driver for a red LED. The TB62750 comes in a VQON24-24 package, and the TB62730 comes in a 2x2-mm WLCSP-16 package. Each device costs \$1 (10,000).

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

Gauge driver runs as many as three air-core meters

➔ Aiming at automotive and industrial applications, the triple-air-core-gauge MLX10420 driver has a serial interface and drives two 360° meters and one 90° meter. The device uses sin/cos PWM drivers and serial commands to control two sets of CMOS power bridges. An on-chip oscillator sets 9-bit-per-quadrant-resolution PWM and a 10-bit angle display. Other features include a power-on self-test, which detects open- or short-

circuit outputs for each meter and provides a real-time-angle-tracking scheme to avoid display errors. The device enables communication with the microprocessor using a three-wire serial link and provides an error-status output pin. Available in an SSOP20 package, the MLX10420 costs 98 cents (10,000).

Melexis Microelectronic Integrated Systems, www.melexis.com

Analog Class D amplifier has closed-loop architecture

➔ The 10W CS3511 stereo-analog-input Class D amplifier uses true spread-spectrum-switching-controller technology and a delta-sigma modulator with a patented closed-loop architecture. Features include 0.025% THD+N at 5W and 0.019% at 1W, support for single-ended or fully differential inputs, and operation from 9 or 12V. The amplifier offers built-in protection and error reporting with automatic recovery from overcurrent, thermal-overload, and undervoltage conditions. Measuring 6x6 mm in a thermally enhanced QFN-32 package, the CS3511 costs \$1.39 (10,000).

Cirrus Logic, www.cirrus.com

EDN ADVERTISER INDEX

Company	Page	Company	Page
Achronix Semiconductor	58	LS Research	26
Advanced Linear Devices	20	MathWorks Inc	19
Agilent Technologies	2, 29	Maxim Integrated Products	45, 47
Allied Electronics	54	Mentor Graphics	17
Analog Devices Inc	15	Microsoft Corp	8
austriamicrosystems AG	52	Mill Max Manufacturing Corp	9
Avnet Electronics Marketing	63	Mouser Electronics	4, 56
BuyerZone	53, 71	MP Associates Inc	38
Coilcraft	7	Murata Power Solutions Inc (formerly C&D Technologies Inc)	28
CUI Inc	65	Pico Electronics	C-3
Digi-Key Corp	1		5, 41
Express PCB	26	Trilogy Design	73
IHS International	61	Vicor Corp	27
Ironwood Electronics	73	Xilinx Inc	C-2
Jameco Electronics	C-4		
Keil Software	37		
Lattice Semiconductor	66		
Linear Technology Corp	42		
	49, 50		

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

EDN

product mart

This advertising
is for new and
current products.

Surface Mount Probing Adapters

In-Circuit test Your IC

- BGA, QFN, QFP, PLCC, & SOIC
- SMT adapters solder to target board
- Modular probe boards, controlled Impedance Flex (plug & test)
- High quality ZIF sockets
- Gold plated test points
- 1000's of standard products
- Quick-Turn Custom Solutions



Ironwood
ELECTRONICS 1-800-404-0204
www.ironwoodelectronics.com



How to keep track of it all?

Easily create and manage multi-level parts lists and specs, calculate costs, generate shopping and kit lists, print labels, generate RFQs and POs and much more...

Parts & Vendors Parts List Manager
and Vendor Database

Get the full function DEMO at
www.trilogydesign.com

Trilogy Design / 200 Litton Dr. #330
Grass Valley, CA 95945 / 530-273-1985

Weather or not: All signs point to “no”



Business had been going well at my former place of employment, a start-up company selling micro-processor-controlled weather stations to schools and broadcast entities. The start-up was taking off, orders were picking up, and our booth was popular at trade shows. Always on the lookout for things new and different, our president had become enamored of scrolling LED signs, and he insisted we needed one in the booth at the next show. “Just imagine: Up-to-the-second weather readings scrolling and flashing in bright colors will be right behind the presenter! Who wouldn’t stop to take a look at that?” he said. Although I had reservations about the aesthetic appeal of barometric-pressure readings blinking in rainbow hues, I had to admit that it would be a nice challenge that would show off some of our company’s technical expertise. So I agreed, but I regretted it when I found out that we had limited time to get this show on the road.

The boss had already selected a vendor for the scrolling signs, so I took

a look at the interface specification. I was pleasantly surprised to find that it supported a rich serial protocol for transmitting and updating messages for display. I needed only to convert our serial-data stream from the weather station to a format that the sign understood. The job was starting to seem a lot easier.

We wanted a stand-alone unit that we could later sell, so we went with an 8051-family microcontroller, which would receive weather data on one

UART and transmit display data with the other. To save time and effort, we ordered an OEM board from a vendor we had used before, and we recycled a number of embedded C routines to handle the serial-data streams. The only remaining task was to write some simple C code to transfer weather data from one stream to another.

The initial work went perfectly. I had hard-coded some test messages into the interface, and the sign displayed them perfectly. I set up the serial-input buffer and state machine on the receiving port, and I could see the extracted weather data in the debugging output. I had only to include the weather data in the output stream. I expanded the output buffer to hold the larger data, and I then reset the system.

Panic quickly set in. Instead of my orderly flow of temperatures and precipitation counts, I got only garbled letters and symbols. Something was seriously wrong. Even though we had tested the serial code in a number of other applications, I began to comb through it, looking for any mistake I could have made. I checked and double-checked every inch of the code. After banging my head against the wall for a few hours, I shut off my computer and went home.

The next morning I realized what the problem was. Prying up the main RAM chip on the OEM board from its socket, I found that one of the data pins was bent under the chip, causing all the ASCII characters to get scrambled before I sent them to the sign, which dutifully displayed the gibberish anyway.

After we straightened the pin, it was not long before we had a working sign-interface module, just in time for the next trade show. Even our president was impressed. “Looks great!” he said. “Now, I have this idea for a weather billboard.” **EDN**

JT Kloplic is a technical specialist with United Parcel Service Information Services (Timonium, MD). You can reach him at jtk_1997@yahoo.com.

www.edn.com/tales


PICO




PROVEN Critical Component Integrity

DC-DC Converters 2V to 10,000 VDC Output


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



Military Components
-55° to +85°C Operating Temp
Military Environmental
Screening Available
3.3 to 500 VDC Outputs
M/MV Series
Isolated - Regulated




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



Military Applications
-40° and -55° to +85°C
Operating Temperatures
Wide Input Range
3.3 to 350 VDC outputs
in 1/2 Brick and Full Brick
Military Upgrades Available
LF/LM/FM Series

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





Programmable to 6000 VDC Output
HVP Series

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



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



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

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

PICO

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

www.picoelectronics.com


Complete Listing of Entire Product Line
DC-DC Converters • AC-DC Inductors • Transformers
and Inductors. E Mail: info@picoelectronics.com

Call Toll Free:
800-431-1064

Fax:
914-738-8225



Are you up for a challenge?



What is the missing component?

IC1 – LT1001
or SIMILAR
OP AMP



Industry guru Forrest M. Mims III has created a stumper. Can you figure out what's missing? Go to www.jameco.com/missing to see if you are correct and while you are there, sign up for our free full color catalog. It's packed with components at prices below what you are used to paying.

JAMECO[®]
ELECTRONICS