electronic design

THE AUTHORITY ON
EMERGING TECHNOLOGIES
FOR DESIGN SOLUTIONS

06.06.13 electronicdesign.com

CONNECTIVITY
PROVIDES THE LIFELINE OF
INDUSTRIAL CONTROL

$10.00 A PENTON PUBLICATION Periodicals Postage Paid • USPS 100 Approved Poly

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

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

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

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

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

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

Coilcraft®

coilcraftdirect.com
No min order. Next day delivery.

WWW.COILCRAFT.COM

Innovative Power Module Reduces System Size

μIPM™ Power Modules Deliver up to 60% Smaller Footprint

Specifications:

Part Number	Size (mm)	Voltage	IO (DC@ 25°C)	Motor Current**		Motor Power VO=150/75VRMS	Topology
				w/o HS	w/HS		
IRSM836-024MA	12x12	250V	2A	470mA	550mA	60W/72W	3P Open Source
IRSM836-044MA	12x12	250V	4A	750mA	850mA	95W/110W	3P Open Source
IRSM836-025MA	12x12	500V	2A	360mA	440mA	93W/114W	3P Open Source
IRSM836-035MB	12x12	500V	3A	420mA	510mA	108W/135W	3P Common Source
IRSM836-035MA	12x12	500V	3A	420mA	510mA	100W/130W	3P Open Source
IRSM836-045MA	12x12	500V	4A	550mA	750mA	145W/195W	3P Open Source

Features:

- 3-phase motor control IC
- 12x12x0.9mm PQFN package offers up to 60% smaller footprint
- Eliminates the need for heat sink
- DC current ratings from 2A to 4A
- Voltage range of 250V – 500V

μIPM™ Advantages:

- Shortens design time
- Shrinks board space requirements
- Simplicity - Eliminates Heat Sink
- Replaces more than 20 discrete parts to deliver a complete motor drive stage
- Slashes assembly time and cost
- Simplifies procurement and inventory management
- Reference design kits available for quick evaluation on any 3-phase motor

iMOTION™ *

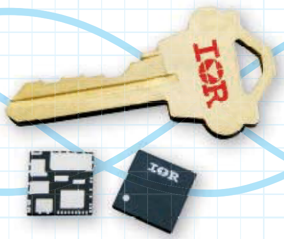

* IR's iMOTION™ (ai mo shan), representing the intelligent motion control, is a trademark of International Rectifier
** RMS, Fc=16kHz, 2-phase PWM, ∆TCA=70°C, TA ≈ 25°C

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

International
IOR Rectifier
THE POWER MANAGEMENT LEADER

06.06.13 VOLUME 61, ISSUE 6

electronic design

InThisIssue

FEATURES

NEWS & ANALYSIS

COLUMNS

14

20

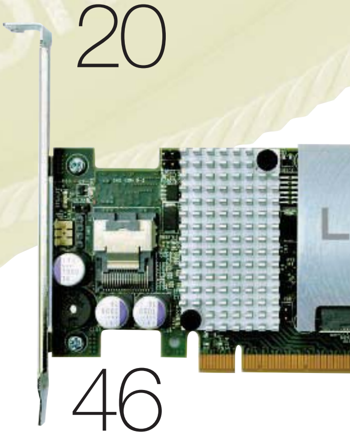

46

80

EDITORIAL MISSION:
To provide the most current, accurate, and in-depth technical coverage of the key emerging technologies that engineers need to design tomorrow's products today.

ELECTRONIC DESIGN (ISSN 0013-4872) is published monthly with an extra issue in June and October by Penton Media Inc., 9800 Metcalf Ave., Overland Park, KS 66212-2216. Paid rates for a one-year subscription are as follows: $120 U.S., $180 Canada, $240 International. Periodicals postage paid at Shawnee Mission, KS, and additional mailing offices. Editorial and advertising addresses: ELECTRONIC DESIGN, 1166 Avenue of the Americas, New York, NY 10036. Telephone (212) 204-4200. Printed in U.S.A. Title registered in U.S. Patent Office. For subscriber services or to order single copies, write to Electronic Design, PO Box 2100, Skokie, IL 60076. POSTMASTER: Send change of address to Electronic Design, PO Box 2100, Skokie, IL 60076. Canadian Post Publications Mail agreement No. 40612608. Canada return address: IMEX Global Solutions, P.O. Box 25542, London, ON N6C 6B2.

InThisIssue

06.06.13
VOLUME 61, ISSUE 6

49

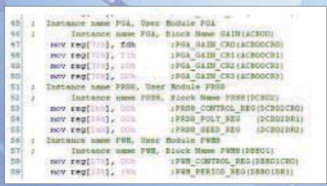

60

76

77

Penton
Design Engineering & Sourcing Group

THE 2008 EDDIE & OZZIE AWARDS
2008 Winner GOLD
EXCELLENCE IN MAGAZINE DESIGN

ASBPE
2008 Winner SILVER
EDITORIAL AWARD OF EXCELLENCE

Putting Technology First

Through the steady advance of civilization, man constantly is gaining new knowledge with which to improve his position. In industry, his facilities for creating and developing new means of performing work are expanding rapidly. Almost daily, new methods and new materials are made available for his use.” These words were used to explain the primary purpose of our group's first publication, *Machine Design*, in September 1929. Almost a century later, they are even more relevant, given the pace at which innovations are now developed and manufactured.

At the heart of any innovation is knowledge—and for knowledge to spread, it must be shared. Since the launch of *Machine Design*, many brands have been conceived to help our technical audience increase their knowledge and improve the way they work. From *Electronic Design* to *Microwaves & RF, Power Electronics Technology, Hydraulics & Pneumatics, Medical Design*, *Global Purchasing,* and *EngineeringTV*, they all serve an audience with specialized expertise. Yet all of them have the same mission: to provide their audiences with reliable information and resources that will help them do their jobs better.

Over the past few years, we have delved more regularly into who our reader is and what he or she needs—in terms of level of technical expertise, age, geographic location, and more. We also have regularly engaged readers to get their feedback about our content. Through audience surveys, our cross-brand “Salary Survey,” or different avenues, we have been able to chart a clear course for the future.

Our goal is to focus on “u”—to be usable, useful, and unified. By being useful, we enable our users to more easily complete tasks and accomplish goals. Our technical information is usable because it is presented in the most concise fashion, allowing it to be quickly digested by today's time-limited design engineer. We also present our content in a unified manner. All of our brands now provide different levels of information in a similar fashion, which enables our audience to easily navigate our sites to find the information they seek.

While we previously provided mostly editor-created technical content, we now join forces with a range of individuals in the various industries we serve. These “Contributed Technical Experts” regularly provide articles for our websites, sharing their technical knowledge and expertise with our audience. Each brand also offers a wider variety of articles, providing hands-on, design-application articles as well as tutorials like our “Essentials” series and “What's the Difference” articles. This multi-pronged approach is key to our technical audience, as many of them are working across different engineering disciplines and require frequent refreshes of technical information.

With social media, online surveys, and other ways of garnering feedback at our disposal, we plan to deepen our engagement with our audience. Through discourses with you, we will continue to adapt the ways that we provide information, ensuring that we help you do your jobs better. We also will be forging a stronger link between our technical brands and SourceESB, which unites various distributors in one place. Our goal is to make Penton's Design Engineering & Sourcing Group a one-stop shop for electronic and mechanical engineers. No matter which brand represents your market, you can be sure of one thing: We will make sure that technology always comes first so that we best meet your needs.

Bill Baumann
VP Market Leader, Design Engineering & Sourcing Group

Nancy Friedrich
Group Director of Editorial Content

electronic design
06.06.13

EDITORIAL

EDITOR-IN-CHIEF: **JOE DESPOSITO** joe.desposito@penton.com
MANAGING EDITOR: **RICHARD GAWEL** richard.gawel@penton.com
CREATIVE DIRECTOR: **DIMITRIOS BASTAS** dimitrios.bastas@penton.com
ANALOG/POWER: **DON TUITE** don.tuite@penton.com
COMMUNICATIONS: **LOUIS E. FRENZEL** lou.frenzel@penton.com
DISTRIBUTION: **VICTORIA FRAZA KICKHAM** SourceESBeditor@penton.com
EMBEDDED/SYSTEMS/SOFTWARE: **WILLIAM WONG** bill.wong@penton.com

ART DEPARTMENT

GROUP DESIGN DIRECTOR: **ANTHONY VITOLO** tony.vitolo@penton.com
SENIOR ARTIST: **JAMES M. MILLER** jim.miller@penton.com
INTERN: **DAYAN S. PELAEZ**

PRODUCTION

GROUP PRODUCTION MANAGER: **JUSTIN MARCINIAK** justin.marciniak@penton.com
PRODUCTION MANAGER: **JULIE GILPIN** julie.gilpin@penton.com

AUDIENCE MARKETING

AUDIENCE MARKETING MANAGER: **BRENDA ROODE** brenda.roode@penton.com
ONLINE MARKETING SPECIALIST: **RYAN MALEC** ryan.malec@penton.com

ELECTRONIC DESIGN EUROPE

EDITOR: **PAUL WHYTOCK** paul.whytock@penton.com
ASSOCIATE EDITOR: **SALLY WARD-FOXTON** sally.ward-foxton@penton.com

LIST RENTALS & CIRCULATION CUSTOMER SERVICE (LIVE)

LIST RENTALS: **MERIT DIRECT**
FREE SUBSCRIPTION/STATUS OF SUBSCRIPTION/ADDRESS CHANGE/MISSING BACK ISSUES
T | 866.505.7173 **F** | 847.763.9673 electronicdesign@halldata.com

SALES & MARKETING

BRAND DIRECTOR, e|DESIGN: **TRACY SMITH** **T** | 913.967.1324 **F** | 913.514.6881 tracy.smith@penton.com

REGIONAL SALES REPRESENTATIVES

NORTHEAST/EASTERN CANADA: **DAVE MADONIA** **T** | 212.204.4331 **F** | 913.514.3966 dave.madonia@penton.com
SOUTH: **BILL YARBOROUGH** **T** | 713.636.3809 **F** | 713.380.5318 bill.yarborough@penton.com
NORTHWEST/NORTHERN CALIFORNIA/WESTERN CANADA: **JAMIE ALLEN** **T** | 415.608.1959 **F** | 913.514.3667 jamie.allen@penton.com
MIDWEST/MID-ATLANTIC: **STEPHANIE CAMPANA** **T** | 312.840.8437 **F** | 913.514.3645 stephanie.campana@penton.com

PLEASE SEND INSERTION ORDERS TO: orders@penton.com
PENTON REPRINTS: **WRIGHT'S MEDIA** **T** | 877.652.5295 penton@wrightsmedia.com
CIRCULATION: **CUSTOMER SERVICE** **T** | 866.505.7173 **F** | 847.763.9673 electronicdesign@halldata.com

INTERNATIONAL SALES

EUROPE: **MARK DURHAM** **T** | +44 (0)7958 564137 mark.durham@penton.com
JAPAN: **HIROKAZU MORITA** **T** | +81 3 3261 4591 **F** | +81 3 3261 6126
TAIWAN: **CHARLES LIU** **T** | 886 2-2727 7799 **F** | 886 2 2728-3686

ONLINE

DIRECTOR OF DIGITAL CONTENT: **PETRA ANDRE** petra.andre@penton.com

DESIGN ENGINEERING & SOURCING GROUP

VICE PRESIDENT & MARKET LEADER: **BILL BAUMANN**
GROUP DIRECTOR OF EDITORIAL CONTENT: **NANCY FRIEDRICH**
GROUP DIRECTOR OF OPERATIONS: **CHRISTINA CAVANO**
GROUP DIRECTOR OF MARKETING: **JANE COOPER**
RESEARCH MANAGER: **JULIE RITCHIE**
MARKETING & EVENTS SPECIALIST: **ADRIAN PIAZZA**
MARKETING COMMUNICATIONS SPECIALIST: **CYNTHIA RODRIGUEZ**

PENTON MEDIA INC.

CHIEF EXECUTIVE OFFICER: **DAVID KIESELSTEIN** david.kieselstein@penton.com
CHIEF FINANCIAL OFFICER: **NICOLA ALLAIS** nicola.allais@penton.com
SENIOR VP, DESIGN ENGINEERING GROUP: **BOB MACARTHUR** bob.macarthur@penton.com

Penton
Design Engineering & Sourcing Group

1166 AVENUE OF THE AMERICAS, 10TH FLOOR
NEW YORK, NY 10036 **T** | 212.204.4200
Electronic Design | Machine Design | Microwaves & RF | Medical Design | Source ESB | Hydraulics & Pneumatics | Global Purchasing | Distribution Resource | Power Electronic | Mobile Dev & Design | Defense Electronics | Auto Electronics | Electronic Design Europe | Engineering TV

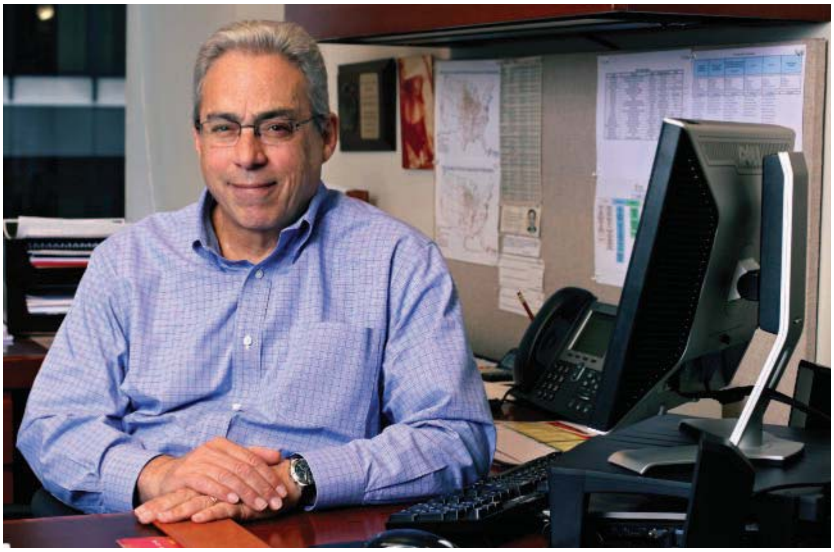

JOE DESPOSITO
Editor-in-Chief
joe.desposito@penton.com

What's The Reach Of Your Wireless Router?

I recently spent some time trying out a Wi-Fi range extender in my home. These devices have been on the market for quite some time. But ever since I purchased a wireless-N router, I've never felt the need for one.

So why did I bother to try the extender? Two of my relatives recently complained to me that wireless-N isn't serving their needs, and they asked me for suggestions. In one case, a smart TV was the straw that broke the wireless-N signal's back. In the other, a recent addition of a family room caused the problem. The new room slipped out of the coverage range of the wireless-N router's signal.

In the case of the smart TV, the router was on one side of the house and the TV was on the other side, with lots of rooms in between. I was asked why Netflix didn't work on the new TV. I knew what the problem was but couldn't solve it by moving the router to a more centralized location. Among other things, it was in the same room as the cable modem and had a wired connection to both the modem and a PC.

With the room addition, the wireless-N router was upstairs and far enough away that it didn't quite cover the new room, so tablets and notebooks could not connect to the Internet from there. The room also seemed to be exhibiting a Faraday cage effect, since wireless worked fine once you stepped out of it.

The Wi-Fi extender I tested is the REC10 from Amped Wireless. Released in April, it's billed as the industry's most powerful, compact Wi-Fi range extender. As with most Wi-Fi extenders, all you have to do is plug it into an outlet and connect to it with your device. I set up the REC10 completely from my tablet via a Web menu and then used the tablet's Wi-Fi signal strength indicator to conduct a test. At a location in my house where I usually get two bars with the wireless-N router, the REC10 boosted the signal strength to three bars from its more centralized location.

The REC10 boasts four amplifiers delivering up to 600 mW power in a 2.75- by 4.0-in. form factor. It is compatible with any brand of router and retails for $79.99. Actually, the REC10 is probably overkill for my home, but it seems perfect for extending Wi-Fi coverage to the smart TV and devices in the new room. ed

PowerOn

BOB ZOLLO | Contributing Technical Expert
bob_zollo@agilent.com

Do You Have Enough AC For Your DC?

When you need 2000 W of dc power for your testing, you may not think too much about where it will come from. The ac line always seems able to provide enough. But in the United States, when operating from the typical and ubiquitous 120-V ac outlet, you might not get enough power off the ac line to generate 2000 W or even 1000 W out of your dc power supply. How can that be?

Power supplies for testing applications have been linear and, thus, limited by size and weight to a few hundred watts. Over time, switching power supplies and other power supply design innovations made dc power supplies much smaller and more efficient. A 2000-W switching power supply is small enough to put on your desktop, but you may be surprised that you can't simply plug it into the nearest 120-V ac outlet because it can't provide enough power into the supply to get 2000 W out.

In North and South America, Japan, Taiwan, and other spots around the globe, ac line voltages are 100 to 120 V ac. In other regions, you have 220 to 240 V ac, meaning you can get twice the power at the same line current. For the rest of this article, I'll stick to talking about the 120 V ac of the United States.

LOOK AT THE SPECS

Let's look at efficiency, which is defined as dc power out over ac power in. Table 1 lists efficiencies for several different types of power supplies. Note that internal functions such as the CPU, cooling fans, I/O, and display consume some of the power for instrumentation power supplies. On a modern power supply, these functions could total 25 to 100 W. This fixed usage will degrade the overall efficiency of the power supply, so it needs more ac power in to get the desired dc power out.

A power supply vendor has to set the required ac line current based on the output power of the supply, its efficiency, and any power consumed internally by its infrastructure. The required ac line current will be the input power divided by 120 V. It could exceed the available line current from a 120-V ac outlet.

Common home or building wiring will have either a 15-A or a 20-A ac receptacle *(see the figure)*. You might think that a power supply will then be limited to either 15 A or 20 A of input current, depending on which receptacle it is plugged into. But the power supply vendor will most likely follow the limitations as expressed by the National Electrical Code (NEC).

Since most people will want to be able to unplug their power supply from the wall and move it to another location, the power supply qualifies as "cord and plug connected equipment not fastened in place." According to the NEC, cord and plug connected equipment cannot draw more than 80% of the circuit rating. That means the 15 A receptacle is valid to deliver only 12 A, or

TABLE 1: POWER SUPPLY EFFICIENCIES

Power supply	Efficiency
Instrumentation dc power supply (linear)	30% to 40%
Fixed output dc power supply (linear)	45% to 50%
Instrumentation dc power supply (switching)	60% to 70%
Fixed output dc power supply (switching)	70% to 95%

TABLE 2: CALCULATING AVAILABLE OUTPUT POWER FOR A DC POWER SUPPLY

Current rating of 120-V ac circuit	15 A
NEC-imposed current derating for cord and plug device	80%
Available current from 15-A, 120-V ac circuit	12 A
Power available from 15-A, 120-V ac circuit at low line (=114 V ac * 12 A) ***This is the ac input power to the dc power supply***	1368 W
Internal (infrastructure) power consumption in the power supply	50 W
Remaining power available to ac-dc conversion	1318 W
Conversion efficiency	70%
Maximum dc output power achievable from a 70% efficient power supply connected on a 15-A, 120-V ac circuit	923 W

1440 W into the power supply. Furthermore, the power supply vendor will have to account for low ac line voltage, which can go as low as 114 V ac. At 114 V ac, the available total ac power is only 1368 W (=114 V ac * 12 A).

But what about your hair dryer that's rated at 1800 W? An electrical device vendor may choose to set its maximum current ratings based on factors other than the NEC. A safety agency, like UL, will permit higher line currents (but never more than 15 A on a 15-A circuit). The power supply vendor may choose to follow the NEC to be more conservative.

Table 2 details the expected dc output power for a sample dc source. To get more dc output power, design the supply to run from a higher voltage, which raises the available input power. While this does provide more power, it also creates an inconvenience for the user, as you need a different, and less common, ac circuit to power the supply.

To achieve full output power in the U.S., 208-V ac input is required, but you need the less common 208-V ac circuit. Agilent's N6700 series dc power supplies are designed so you can also plug it into the common 15-A, 120-V ac outlet.

When the supply's universal input senses that it is plugged into 120 V ac, it automatically limits its output power so its input power will not demand more than 12 A, which is what the NEC says you can draw from the common 15-A, 120-V ac receptacle. So while it can't operate at full power, it can still operate. Users can simply plug in and temporarily use the supply, with some limitations, without a 208-V ac outlet.

BOB ZOLLO is a product planner at Agilent Technologies. He holds an EE degree from Stevens Institute of Technology.

In the United States, standard 120-V ac receptacles include the 15-A receptacle, characterized by two vertical slots, and the 20-A receptacle, characterized by one vertical slot and one horizontal slot.

News & ANALYSIS

ZIGBEE ALLIANCE
Launches Smart Energy Profile 2.0 For HANs

The ZigBee Alliance has finalized its Smart Energy Profile 2.0 (SEP2) standard, which is an application for short-range radios using the ZigBee standard for wireless networking. It implements all the various features needed for home energy management.

Home-area networks (HANs) have been around for years to help consumers monitor and control their energy usage. Multiple applications have been developed using ZigBee and Z-Wave wireless as well as HomePlug power-line communications technology, but no common standard has emerged. This lack of a common standard may be one of the factors slowing the growth of HANs. SEP2 should help solve this problem.

As a spinoff and evolution of ZigBee's original Smart Energy 1.x standard, SEP2 provides IP-based (Internet protocol) information and control for energy management in HANs using both wired and wireless technologies. It supports plug-in and hybrid electric vehicle (PHEV) charging control, deployments in multi-dwelling units like apartment buildings, interfaces for multiple energy services at single premises, and any Internet Engineering Task Force (IETF) IP-compliant protocols such as a previously announced ZigBee IP that offers IPv6 addressability.

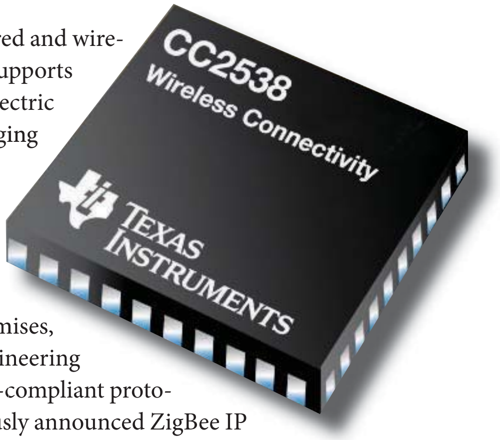

SEP2 is the application layer for the ZigBee stack that attaches to the media access controller (MAC) layer and physical layer (PHY) of the IEEE 802.15.4 radio protocol. It is also designed to work with Wi-Fi radios and devices using the power-line communications standard of the HomePlug Alliance. It offers price and billing information, demand response and load control, and the use of most mainstream protocols like HTTP, XML, and IEC 61968.

The Texas Instruments CC2538 is an example of a single-chip ZigBee solution that supports SEP2 *(see the figure)*. Featuring an ARM Cortex M3 core, it supports smart energy infrastructure applications with scalable memory options and full mesh networks with hundreds of nodes.

The Consortium for Smart Energy Profile (CSEP) was formed to test and certify the interoperability of SEP2 products. The ZigBee Alliance, Wi-Fi Alliance, and HomePlug Alliance founded the organization to ensure that all of the various wired and wireless devices for the Smart Grid ecosystem such as thermostats, appliances, electric meters, gateways, vehicles,

Sandia Device DETECTS BIOTHREATS

Sandia's SpinDx uses centrifugal microfluidics or "lab-on-a-disk" technology to identify anthrax, botulinum, shiga, SEB toxin, and other biothreat agents. Once approved by the Food and Drug Administration and commercialized, it will be used in emergency rooms following bioterrorism incidents. *(photo by Jeff McMillan)* and other devices in the home will talk to one another. Using plug fests and formal testing procedures, CSEP will assist in validating the interoperability of all devices using SEP2. For more information, contact www.csep.org.

The development and deployment of the Smart Grid for nationwide energy monitoring and control has been a slow work in progress. There are more than 120 million homes in the U.S., but only just over 40 million homes now have a smart meter that reports energy usage to one of 11 utilities. Now ZigBee and SEP2 offer a standard way to operate and communicate in a HAN. That should encourage further growth. For more details and a download of the SEP2 standard, go to www.zigbee.org. ■

LOUIS E. FRENZEL

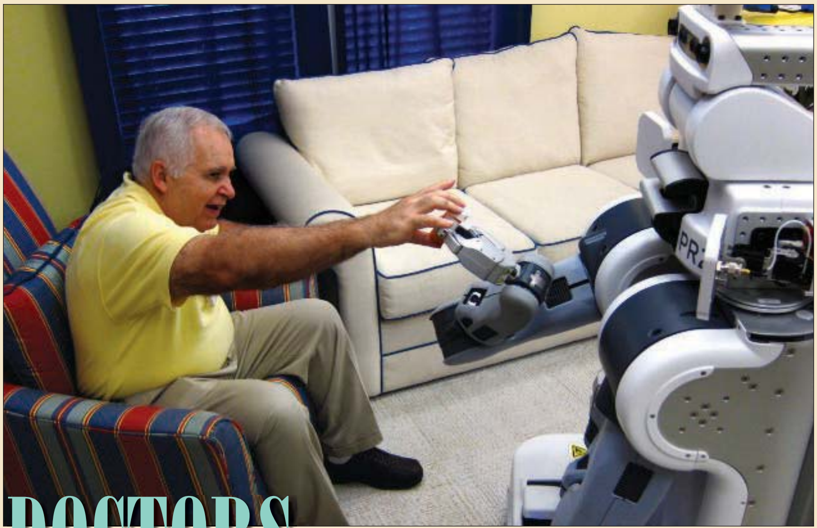

DOCTORS PREFER SOME ROBOT ASSISTANCE

ACCORDING TO A recent Georgia Tech study, more than half of the healthcare providers interviewed said they would prefer a robotic helper over a human one in instrumental activities of daily living, such as housework and administering medication. Yet they thought human assistants were better for direct physical interactions such as bathing, dressing, and feeding patients.

DARPA Develops GPS Backup

RESEARCHERS FROM THE Defense Advanced Research Projects Agency (DARPA) at the University of Michigan have developed a timing and inertial measurement unit (TIMU) that can aid navigation when GPS is unavailable. It includes a six-axis IMU and a highly accurate master clock to measure orientation, acceleration, and time.

LOU FRENZEL | Communications Editor
lou.frenzel@penton.com

Most industrial monitoring and control functions require networks and I/O interfaces to function.

CONNE PROVIDES THE LIFELINE OF INDUSTRIAL CONTROL

Industrial control broadly defines the range of electronic equipment used in factories, process control plants, and automated facilities to monitor and control manufacturing and other operations. It involves robots, computers, machine tools, programmable logic controllers (PLCs), sensors, relays, valves, motors, and measuring instruments.

Industrial control is also part of mines, oil and gas production, water and waste water treatment, electrical and gas utilities, and other power generation facilities. Airlines, railways, trucking, and metropolitan transportation systems use industrial electronics as well. Yet one dominant theme throughout all industrial control is communications.

Virtually all equipment and devices rely upon electrical interfaces and networks to function. Over the years, the use of communications equipment has increased and its nature has changed as new technologies have emerged to improve the communications function as well as optimize monitoring and control operations.

TRENDS AND ISSUES

The industrial field generally lags behind other sectors of electronics simply because its technological needs do not follow the consumer or enterprise market trends. But overall, industrial sectors do follow the general trends in communications technology. Key trends and issues include:

- Continued use of fieldbuses: The fieldbuses are the digital LANs of industry. They connect the sensors, controllers, and actuators of most factory automation and process control facilities. Despite the ongoing movement to Ethernet connectivity and wireless, there continues to be growth of several percent per year in the fieldbus market.
- Strong movement to Ethernet: Ethernet has been the local-area network (LAN) of choice for enterprise and even consumer networking for decades, and it dominates. Industry was slow to adopt it but has now embraced it completely. Most new industrial networking efforts use some form of Ethernet. Its proven reliability, low cost, and high availability have made it particularly popular. Special industrial versions of Ethernet have emerged to enhance it for industrial use.
- Significant growth in wireless connectivity: Industry was slow to adopt wireless despite its many benefits. Industrial users assumed it was unreliable and insecure but have learned otherwise since. New and improved wireless standards and equipment have made wireless a key component in most modern industrial settings.
- Fewer proprietary standards and equipment: For decades, industrial communications needs were met with many high-cost proprietary fieldbuses, interfaces, and equipment,

CTIVITY

which are still entrenched in many systems. However, the trend today is to open standards and Ethernet.

- Rapid adoption of the Internet protocol (IP) model: The goal is to give most industrial equipment an IP address so devices and equipment can communicate over Ethernet and the Internet. With the availability of IPv6, that is now possible.
- More video surveillance: security is critical to many plants and facilities, and video is useful. Video also enables improved monitoring that simple sensors cannot provide.
- The lingering hodgepodge of equipment: Most factories, process control plants, and facilities are a real mixed bag of old and new, analog and digital, and proprietary and open standards. A big issue has been the incompatibility and interoperability of different equipment. How can it all work together seamlessly? New standards, equipment, and software are gradually addressing that problem.

FIELDBUSES

In the past, control systems in factories and plants were analog in nature. They used direct connections from controller to actuator or transducer to controller and were based on a 4- to 20-mA control signal. As systems became more complex and as networking technologies evolved, eventually a change from the analog system to a digital system came about.

Direct digital control (DDC) systems using a single computer to control a large part of the system came first. Distributed control systems (DCSs) where each device had some intelligence replaced DDC. Programmable logic controllers (PLCs) are also a major part of the equipment mix. Fieldbuses were developed to tie all these digital components together.

Fieldbuses are digital networks and protocols that are designed to replace the analog systems. These industrial LANs network all of the computers, controllers, sensors, actuators, and other devices so they can interact with one another. A single network cable replaced the dozens or even hundreds of individual analog cables in the older systems. Protocols allowed operators to easily monitor, control, troubleshoot, diagnose, and manage all devices from a central location.

While these fieldbuses reduced the wiring and improved the reliability and flexibility of the system, another issue was created: multiple proprietary systems, with incompatibility and a lack of interoperability between the various components. Devices made to work with one fieldbus and protocol could not work with another.

While some systems could be implemented with a single fieldbus from a single manufacturer, more tended to be "mix and match" systems with components from multiple vendors. The proprietary systems protected the manufacturer's market, but they often created unbearable hardships on the customers and the engineers who had to design, build, and maintain the control systems.

Many fieldbuses have been developed. Some attempts at building a common standard were made but no one system or standard ever emerged. The marketplace weeded out the weaker systems, leaving a few still incompatible standards. ControlNet, DeviceNet, Foundation Fieldbus H1, HART, Modbus, and Profibus PA are the most common fieldbuses.

ControlNet is an industrial network and protocol supported by the Open DeviceNet Vendors Association (ODVA). It is based on the Common Industrial Protocol (CIP), which defines messages and services to be used in manufacturing automation. ControlNet uses RG-6 coax cable with Bayonet Neill-Concelman (BNC) connectors for the physical layer (PHY) and is capable of speeds to 5 Mbits/s using Manchester coding. The topology is a bus with a maximum of 99 drops. Its timing permits a form of determinism in the application.

DeviceNet is another ODVA supported fieldbus. It uses the well-known controller-area network (CAN) technology for the PHY that was originally developed by Bosch for automotive applications. The DeviceNet protocol is similar to that of ControlNet and also uses the Common Industrial Protocol (CIP) at the upper layers. Layers 1 and 2 are CAN bus. The medium is unshielded twisted pair (UTP) using a single-ended non-return-to-zero (NRZ) format with logic levels of 0 V and +5 V. The topology is a bus with up to 64 nodes allowed. Data rate depends on bus length and can be as high as 1 Mbit/s at 25 meters to 125 kbits/s at 500 meters.

Originally developed by the International Society of Automation (ISA) standards group as Foundation Fieldbus, SP50 was one of the earlier digital fieldbuses for replacing 4- to 20-mA loops. The protocol is designated H1, which uses the IEC 61158-2 standard for the PHY and features twisted-pair cabling with a basic data rate of 31.25 kbits/s. The transmission frames are synchronous with start and stop delimiters. Coding is Manchester.

HART, which stands for Highway Addressable Remote Transducer, is a two-way communications path over twisted pair. It retains 4- to 20-mA analog functionality but adds digital signals. The digital signal is a frequency shift keying (FSK) modulated carrier that uses the old Bell 202 modem frequencies of 2200 Hz for a 0 and 1200 Hz for a 1.

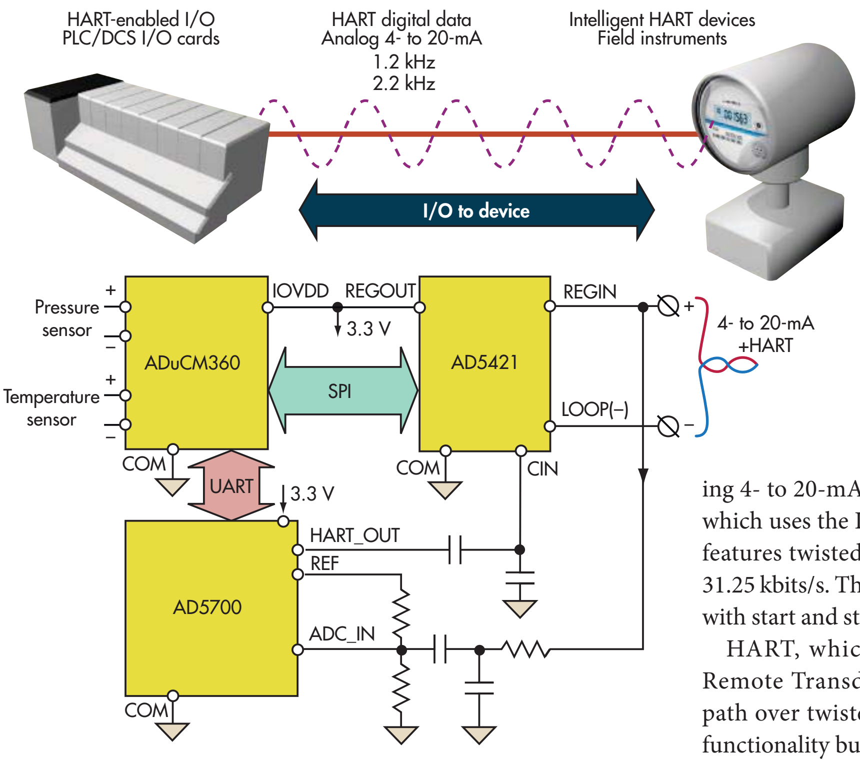

1. A HART field bus device can be implemented with sensors, an embedded controller, a DAC, a dc 4- to 20-mA current source, and a HART modem from Analog Devices.

2. The B&B Electronics APXG-Q5420 is an 802.11b/g wireless router and access point. It features one 10/00 Ethernet port and two DB9M RS-232/422/485 ports for serial connected industrial devices.

The data rate is 1200 bits/s. The FSK signal is phase continuous and doesn't affect the analog signal level because it's ac. Also, the FSK signal is a 1-mA variation around the dc level. The protocol uses OSI layers 1 through 4 and 7. The digital part of the communications is primarily used for commands, provisioning, and diagnostics.

The HART fieldbus is popular as it is compatible with older 4- to 20-mA equipment while adding the digital networking capability. It is still widely used. Typical HART field instruments such as the Analog Devices ADμCM360 consist of an embedded controller and the I/O for the sensors such as a pressure transducer and real-time data (RTD) temperature sensor *(Fig. 1)*.

The on-board 24-bit sigma-delta analog-to-digital converters (ADCs) digitize the sensor information and then send it to the AD5421, a digital-to-analog converter (DAC) and 4- to 20-mA current source for connection to the cable. Digital information is also sent to the AD5700 HART FSK modem.

Modbus is a popular industrial protocol normally used for communications with PLCs. It is simple and the standard is open, meaning anyone can use it. It works with RS-232 interfaces. The basic format comprises asynchronous characters sent and received with a UART. Modbus can be carried over a variety of PHYs and is often encapsulated in Transmission Control Protocol/IP (TCP/IP) and transmitted over Ethernet. It is also compatible with a wireless link.

Profibus was developed in Germany and is popular worldwide. There are versions for decentralized peripherals (DPs) and process automation (PA). The protocol is synchronous and operates in OSI layers 1, 2, 4, and 7. Using RS-485, bit rates can range from 9.6 kbits/s to 12 Mbits/s. With a bus up to 1900 meters long, the data rate is 31.25 kbits/s.

INDUSTRIAL ETHERNET

Ethernet has been the LAN of choice for business and enterprise for decades. It is by far the most successful and widely used networking technology in the world. It is affordable and reliable and is backed up by a strong series of IEEE 802.3 standards that keep it current. Over the past 10 years or so, Ethernet has found its way into the industrial setting for I/O and networking. It is gradually replacing the multiple fieldbuses and proprietary networks or working with them. Some of the benefits of moving to Ethernet are:

- Fewer smaller networks: Most fieldbuses can connect up to 20 to 40 devices. Beyond that, a separate fieldbus network is required for more devices and to connect the two networks, if that's even possible. With Ethernet, you can connect up to 1000 devices on the same network. This improves the efficiency and decreases the complexity of the network.
- Lower cost: With many Ethernet vendors, equipment prices are competitive and the overall cost of building a network is typically lower than building a fieldbus network.
- Higher speeds: Ethernet has much higher speed capability than most fieldbuses. While that speed isn't always needed, it is a benefit and the network grows in size and as faster devices are connected. While 10/100-Mbits Ethernet is the most common, some industrial facilities have already upgraded to 1-Gbit/s Ethernet.
- Connection to the factory or plant IT network: The industrial networks are traditionally kept separate from the business network, but companies are finding that advantages occur when the two networks can be interconnected. Data can be collected and used to optimize the manufacturing process or make improvements or decisions not previously possible.

3. Honeywell's Limitless wireless industrial products use the robust 802.15.4 wireless standard. On the far left is the WGLA mechanical limit switch. The unit second from the left is the Honeywell WDRR, a receiver for wireless limit switches that provides inputs for PLCs. The third from the left is the WPMM panel mount monitor unit, which provides LED visual or audio indication for remote limit switch actions. The far right device is the WLS magnetic proximity limit switch. Honeywell's OneWireless Networking products also comply with the ISA100.11a standard.

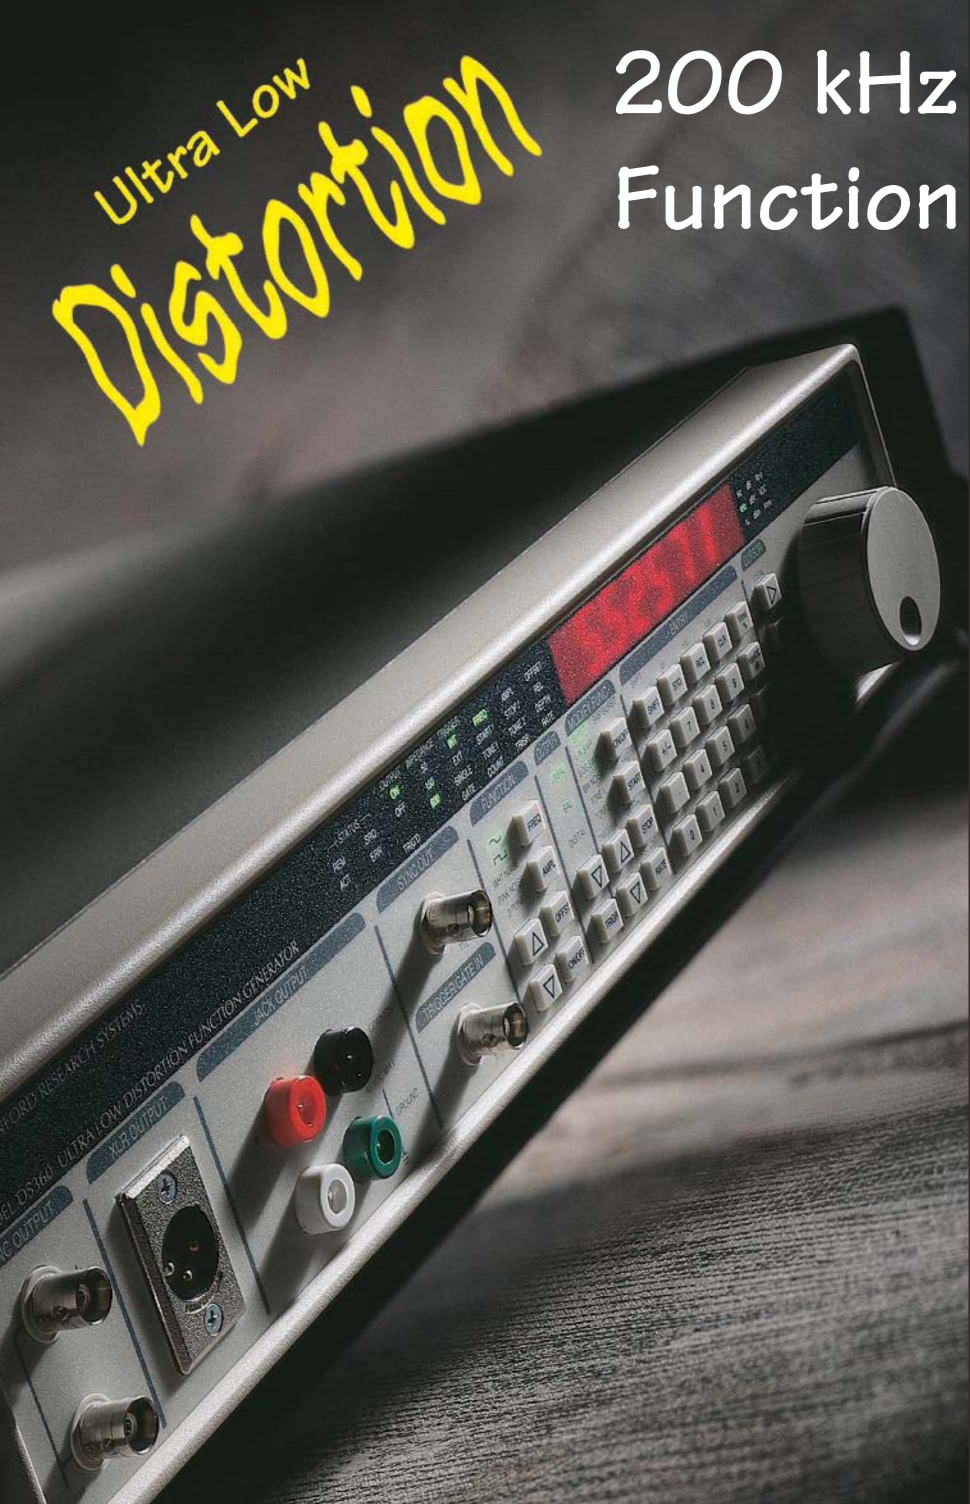

200 kHz DDS Function Generator

0.001% THD
(-100 dB)

DS360...$2795 (U.S. list)

- 25 ppm freq. accuracy
- 1 mHz to 200 kHz
- 4 digit amplitude resolution
- Sine, square, two-tone, burst
- Pink and white noise
- Log & linear freq. sweeps
- 35 µVpp to 40 Vpp
- SPDIF/EIAJ and AES-EBU digital outputs
- GPIB and RS-232 interfaces (standard)

The performance of a low distortion signal source and the features of a digitally synthesized function generator are combined in a unique new instrument, the DS360 from SRS. With 0.001% THD, 25 ppm frequency accuracy and functions that include sine waves, square waves, sweeps, and bursts, the DS360 delivers unparalleled precision and versatility. To make it complete we added two-tone capability (including SMPTE, DIM, CCIF), bandwidth limited white noise, SPDIF/EIAJ and AES-EBU digital outputs, and computer interfaces.

The DS360 - the one source for all your audio frequency needs.

SRS

e-mail: info@thinkSRS.com • www.thinkSRS.com

Stanford Research Systems, 1290-D Reamwood Avenue, Sunnyvale, CA 94089
Phone (408) 744-9040 • Fax (408) 744-9049

- Connection to the Internet: This may not be desirable, but if it is, Ethernet provides a very convenient way to send and receive data over an Internet connection.

There are two main downsides of using Ethernet in the industrial setting. First, the hardware wasn't designed for the demanding environment in factories and process plants. Excessive temperatures, environmental hazards, chemicals, dust, mechanical stresses, and moisture make traditional equipment less reliable. Yet over the years, manufacturers have repackaged Ethernet gear to bear up under such conditions by adding industrial-grade housings and tougher electronics.

Another hazard is excessive noise caused by motors, power switching, and other sources. Ethernet's differential wiring is essentially noise resistant. It can be made noise-free with shielded cable. Most industrial wiring is simply standard, but higher-grade CAT5 or CAT6 cable usually suffices.

On the other hand, the RJ-45 connectors are a source of problems, especially in dirty and high-moisture environments. Special RJ-45 connectors have been developed to solve this problem. These connectors add a dirt- and moisture-sealed cover to an upgraded RJ-45 connector. These connectors meet the rigid IP67 environmental standards for hazardous environments.

A second disadvantage is Ethernet's inherent non-deterministic nature. Many industrial networks rely on timing conditions that must occur within a specific time frame. Many need a real-time connection or something close to it. Determinism means a device or system can respond within a mini-

THE INDUSTRIAL CONTROL MODEL

INDUSTRIAL CONTROL APPLICATIONS involve a monitoring function, a control function, or both. Some physical characteristic like temperature, pressure, light, or position usually is monitored. Some cases involve a measurement, meaning some instrumentation equipment is needed. Otherwise the monitoring results in just a "present" or "not-present" signal.

Data acquisition is an example of monitoring where multiple physical characteristics are detected or measured. This data is collected, stored, processed, and ultimately displayed. Data acquisition may not be involved in control operations, where the measured or monitored parameter becomes a feedback signal to effect them.

In control applications, some physical characteristic is changed. Devices like motors or valves are turned off or on or some other physical variable is modified. The state may be either off or on or may be a proportional variation such as speed. Many applications use a combination or monitoring and control operations, which can be classified as either open loop or closed loop control.

Open loop operations normally involve human monitoring and control. For instance, an operator may visually observe the temperature of the liquid in a tank. If the temperature drops below a certain minimum value, the operator manually turns on a heating element to increase the temperature to the desired value.

Closed loop control automates the temperature condition. A temperature sensor sends a signal to a controller that operates the heating element. Therefore, the temperature of the liquid is self-regulating without human intervention (see the figure). The set point is usually a voltage setting that specifies the desired temperature. It is compared to a signal from a temperature sensor that tells the actual temperature.

If the two signals are the same, no action is required.

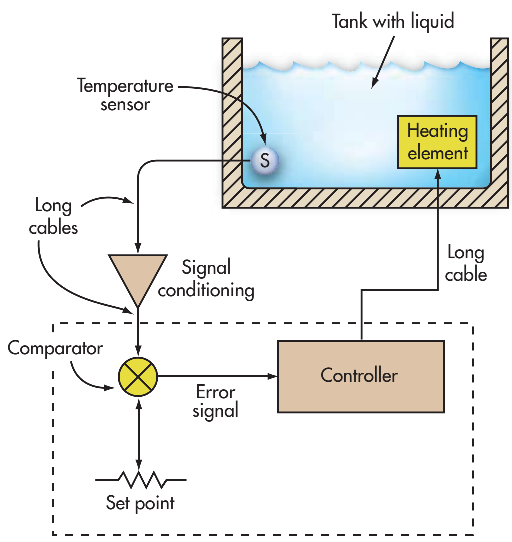

Most industrial control systems use a closed-loop model. The sensors and actuators often are located at a distance, so some interface or network connectivity is required.

If a difference exists, an error signal is generated, initiating the turn-on of the heating element. As soon as the temperature reaches the desired level, the feedback sensor detects it and the controller shuts off the heating element. This is an example of on-off control. Continuous control would provide continuous power to the heating element and vary the current through it to maintain the temperature more closely.

The various parts of this system are often located remotely from one another. Long cables and interfaces are used to tie all of the system's components together. ■

mum time interval. It can respond in less time but no more than a specified time. If a device is deterministic to 10 ms, anything less is okay. The response does not usually have to be repeatable, but that depends on the application.

Ethernet determinism is widely variable. It is a function of the carrier sense multiple access with collision detection (CSMA/CD) access method, cable lengths, number of nodes, and the combinations of hubs, repeaters, bridges, switches, and routers used in the system. To improve the deterministic response, designers of industrial Ethernet systems must keep cables short and minimize the number of nodes, hubs, and bridges. Switches can be added to larger networks to isolate different segments, and that reduces the number of collisions and interactions.

Determinism can also be implemented or improved in some cases by using the IEEE 1588 Precision Time Protocol (PTP). The PTP permits systems with clocks to achieve synchronization among all connected devices, allowing precise timing information transfer within a network. Time stamping and near real-time performance can occur in some applications.

Another feature of Ethernet finding acceptance is Power over Ethernet (PoE). Defined by IEEE standards 802.3af and 802.3at, it allows the transmission of dc power over the Ethernet cables to power remote devices. This is a major benefit in many industrial settings and eliminates the need to install a power source near some remote sensor or other device. The standard defines power levels up to 15.4 or 25.5 W, but higher-power versions up to 51 W are becoming available.

Finally, several enhanced or modified versions of Ethernet have been developed to overcome the timing issues of standard Ethernet or simply make it more compatible with existing equipment and systems. These include EtherCAT, EtherNet/IP, Profinet, Foundation Fieldbus HSE (high-speed Ethernet), and Modbus/TCP. Some use special protocols while others use TCP/IP.

EtherNet/IP is an application layer protocol using CIP, which defines all devices as objects and specifies the messages, services, and transfer methods. CIP is then encapsulated in a TCP or User Datagram Protocol (UDP) packet for transfer over Ethernet.

Profinet is another protocol that uses TCP/IP over Ethernet. It is not Profibus over Ethernet. Instead, it uses two different protocols: one called Profinet CBA for component-based systems and Profinet IO for real-time I/O operations. Profinet CBA can provide determinism in the 100-ms range. It also can deliver determinism to 10 ms. A version of Profinet IO called IRT for isochronous real-time can have a determinism of less than 1 ms.

Foundation Fieldbus HSE uses the H1 protocol over TCP/IP. It also uses a special scheduler that helps to guarantee messages in known times to ensure determinism at some desired level.

EtherCAT gets rid of the CSMA/CD mechanism and replaces it with a new "telegram" message packet that can be updated on the fly. Networked devices are connected in a ring or a daisy chain format that emulates a ring. As data is passed around the ring, message data can be stripped off or inserted by the addressed node while the data is streaming. The one or more EtherCAT telegrams are transported directly by the Ethernet frame or encapsulated into UDP/IP datagrams. Determinism of 30 μs and less can be achieved with up to 1000 nodes.

Modbus/TCP is the popular Modbus fieldbus protocol packaged in a TCP/IP packet. The Modbus checksum is replaced by TCP/IP's 32-bit checksum. Then the TCP/IP packets are carried over standard Ethernet.

Obviously all of these systems aren't interoperable with one another. But they can all coexist on the same Ethernet LAN since they all conform to the Ethernet Layer 1 PHY standard. Those using TCP/IP could be made interoperable with the appropriate software modifications.

INDUSTRIAL WIRELESS

Wireless technology is a relative newcomer to the industrial scene. Conservative engineers preferred the use of hardwired equipment using cables to ensure reliability in a noisy, interference-prone harsh environment. But wiring is expensive and has other disadvantages. In addition, improved wireless techniques and equipment make it a suitable medium for even critical industrial applications.

Today, wireless is widely accepted as the technology for connecting sensors, actuators, controllers, computers, and data acquisition systems. In fact, modern factories, process control plants, and similar industrial facilities have been covered with multiple networks, including wireless.

Wireless offers significant benefits to industrial connectivity. First, it is time-critical, offering improved response time in some cases over wired solutions. Wireless also eliminates the endless problems with cables and connectors that are common failure points in industrial settings. It is the connection method of choice for hard-to-wire machines, mobile machines, or machines that rotate.

Wireless is especially favorable in applications with long cable runs or where wiring is prohibitive. Such a run could cost tens or even hundreds of thousands of dollars and take weeks to install by a union electrician. Wireless devices can usually be installed and provisioned in minutes. And, wireless devices can be easily monitored or controlled via the Internet and cloud.

One main disadvantage of wireless devices is their need for battery power, since batteries must be frequently changed. This was an initial major knockout factor for wireless sensors and actuators. Changing batteries is a maintenance problem that more plant engineers would rather not face. Today, wireless devices are very power efficient, extending battery life to years. Solar power also is possible in many applications, and other energy harvesting techniques such as vibration are being adopted.

Initial wireless adoption mostly involved proprietary designs for longer distances. Proprietary technologies are still widely used in industrial, scientific, and medical (ISM) band radios for the 902- to 928-MHz and 2.4-GHz bands. However, standard technologies like Wi-Fi and ZigBee are also deployed. And, special wireless technologies have been developed especially for industrial. They feature a mix of noise-mitigating techniques, security, and low latency for highly deterministic applications.

Standard Wi-Fi can be used in some non-critical industrial applications. However, its disadvantages may outweigh the benefits of low-cost, readily available, and easy to install equipment. A host of other wireless gear crowds the Wi-Fi 2.4-GHz band. Interference may too great for a reliable connection in some applications. Latency is long and also a key problem since precise timing is not its main feature. Nevertheless, Wi-Fi is a good fit for some industrial applications *(Fig. 2)*.

Other popular wireless technologies like Bluetooth and Z-Wave also are generally unsuitable for industrial use. Yet one 2.4-GHz technology has emerged as a prime technology for industrial wireless equipment. The IEEE 802.15.4 standard that uses the 2.4-GHz band is especially robust because its frequency-hopping method ensures an interference-free channel and the direct sequence spread spectrum (DSSS) modulation method survives well in a noisy environment.

The standard's 2-MHz chipping provides a processing gain against interference of 9 dB. Its 250-kbit/s basic data rate is fast enough for most industrial uses. A 128-bit AES encryption engine provides security. Its very low duty cycle makes it a very low-power technology. When enhanced, 802.15.4 becomes ideal for many industrial applications. ZigBee, WirelessHART, ISA100.11a, 6LoWPAN, and several other technologies are based on the 802.15.4 standard *(Fig. 3)*.

4. The B&B Electronics Spectre 3G cellular router is a Wi-Fi 802.11b/g/n hot spot but also provides a high-speed link to CDMA or HSPA+ cellular base-stations.

ZigBee uses the 802.15.4 standard for the physical and link layers and adds a communications stack that implements mesh networking. Mesh networks are commonly used to form wireless sensor networks that are spread out over a large distance. They extend the range of each sensor and make the connectivity more reliable thanks to the multiple redundant paths possible. Even when nodes fail or noise interrupts one path, another path can be used. Some industrial applications use ZigBee, but other enhanced versions of 802.15.4 are more widely incorporated.

WirelessHART is the radio version of the popular HART field bus. It uses the 802.15.4 standard as its base but adds the upper layers of the system. It provides a time-synchronized, self-organizing, self-healing mesh network based on Linear Technologies' Dust Networks technology. It has always-on security with 128-bit AES encryption and full device authentication.

The basic components are a network manager that sets up and maintains the mesh network, a security manager that handles the encryption keys, a gateway that provides the connection to the host network, a repeater that just routes messages to extend range, an adapter that works with existing HART instruments, and a hand-held terminal for routing maintenance and calibration.

The International Society of Automation (ISA) ISA100.11a standard was designed exclusively for the process control industry and general factory automation. The physical and data link layers use the basic 802.15.4 standard. Additional layers of the communications stack implement time division multiple access (TDMA), channel hopping, and mesh-to-mesh routing. This standard also provides a tunneling protocol that permits it to carry existing protocols such as HART, Foundation Fieldbus, Modbus, and Profibus. The frame format complies with the Internet Engineering Task Force (IETF) RFC 4944, an IP-based protocol.

One interesting wireless technology invading the industrial space is 6LoWPAN, which means IPv6 over low-power wireless personal-area networks. This IETF standard is designed to help connect the billions of devices projected for the Internet of Things (IoT) movement. One popular forecast says there will be over 50 billion devices of all sorts connected to the Internet by 2020. This standard will provide some of that connectivity.

IPv6 offers a 128-bit addressing capability, allowing many more IP addresses to be assigned than now possible with the fading IPv4 32-bit capability. 6LoWPAN uses a header compression technique that encapsulates IPv6 packets so they can be carried by 802.15.4 radios. The IETF standard is RFC 4944.

Then there's machine-to-machine (M2M) connectivity *(Fig. 4)*. Known as the Internet of Things (IoT), this movement includes industrial components. M2M applications typically use cellular radios to communicate. These are particularly useful for very remote sensors or other equipment. Low-cost cellular radio modems are widely available for embedding into other units. However, special M2M cellular service with a carrier is required. ed

DON TUITE | Analog/Power Editor
don.tuite@penton.com

THE INTER OF THINGS HITS THE

Making vehicles part of the Internet of Things opens new design possibilities.

Imagine you're getting into your car and you want to continue a hands-free conference call with people at work. Your kids in the backseat are engaged with their Facebook accounts, watching videos, or gaming online. And, you're in a rush to get to where you really need to be.

Before the Internet of Things (IoT) can support these demands, automotive systems need to allocate bandwidth between high priorities like conference calls and low priorities like Facebook. Security is another issue, especially for communications that affect driving safety. As a result, design opportunities abound *(Fig. 1)*.

BREAKING IT DOWN

Reliability is a key factor. Automotive IoT systems must provide non-stop connectivity based on the strength of their current connection and the potential strength of alternatives. If a user is connected to service provider A and the strength of that connection starts to degrade, the hardware must be able to detect the degradation and switch to provider B well before the link fails.

The hardware and software should be able to transparently transition the vehicle's primary connectivity and transfer the link forward to the optimum adjacent cellular provider. In the evolving world of Internet adaptation, though, cellular-to-cellular is only one kind of possible transition. Cellular-to-Wi-Fi is another.

"New mesh Wi-Fi networks that are being created present a huge opportunity to reduce costs to the user," said Mala Devlin, business operations manager of connected industries at Cisco. Best known for Internet routers and switchers, Cisco is staking out vast amounts of territory at the interface between the IoT and transportation.

Dedicated Short-Range Communications (DSRC), defined by IEEE 802.11p, is yet another option. At its roots, DSRC is still Wi-Fi, but some hardware improvements can make it robust. DSRC involves vehicle-to-

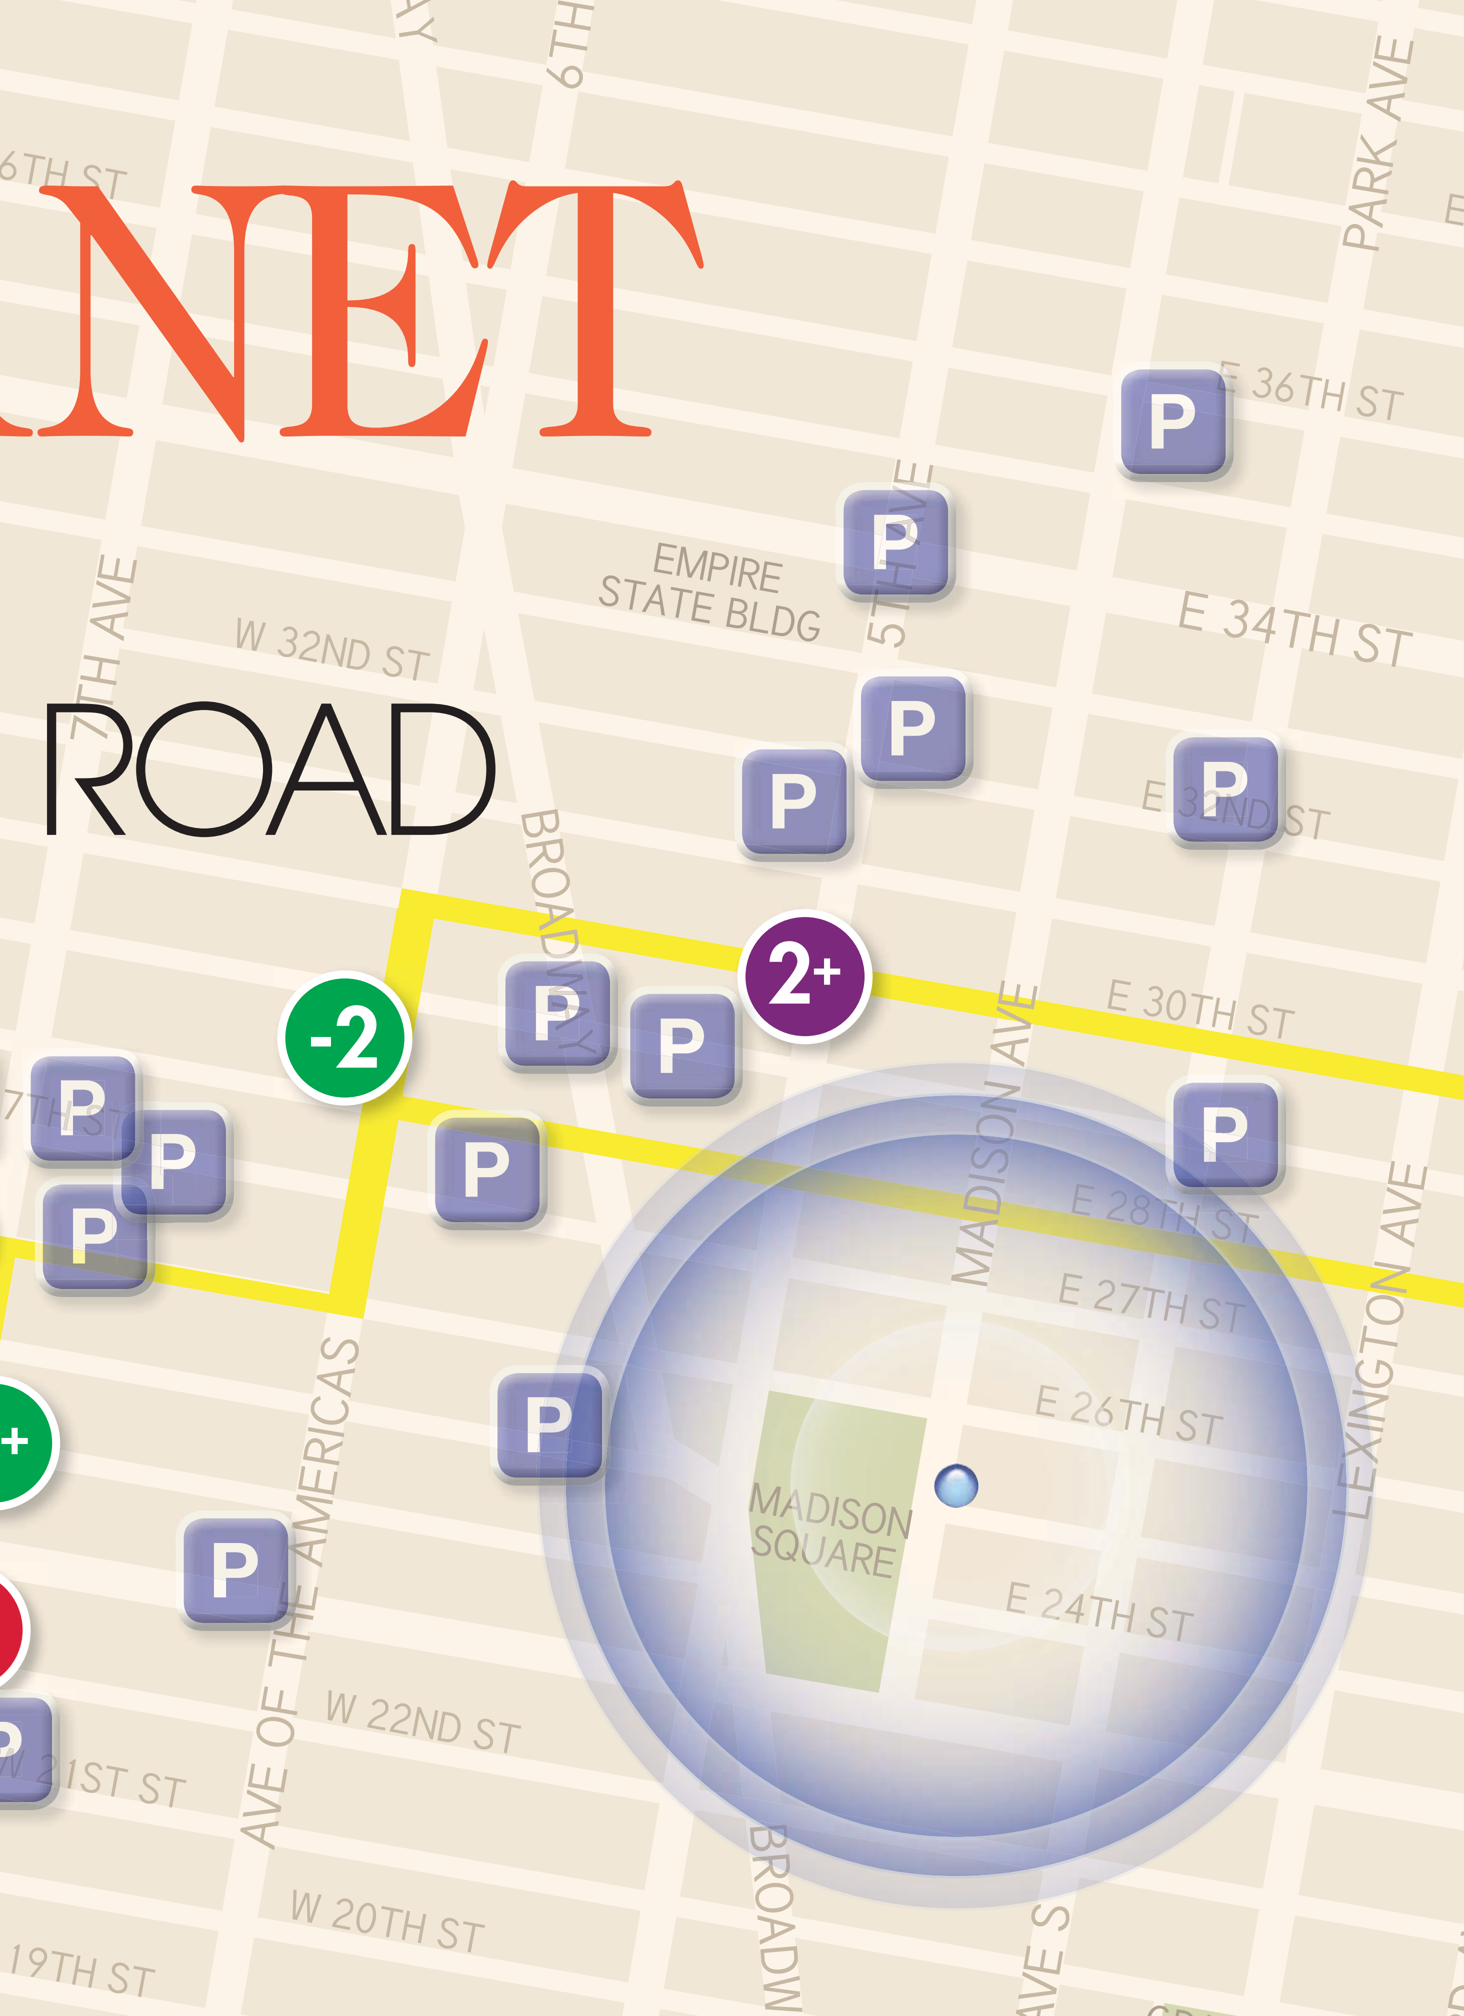
NET
ROAD
PARK AVE
E 36TH ST
EMPIRE
STATE BLDG
5TH AVE
E 34TH ST
W 32ND ST
7TH AVE
E 32ND ST
BROADWAY
2+
-2
E 30TH ST
MADISON AVE
E 28TH ST
E 27TH ST
E 26TH ST
MADISON
SQUARE
E 24TH ST
LEXINGTON AVE
AVE OF THE AMERICAS
W 22ND ST
W 21ST ST
W 20TH ST
19TH ST

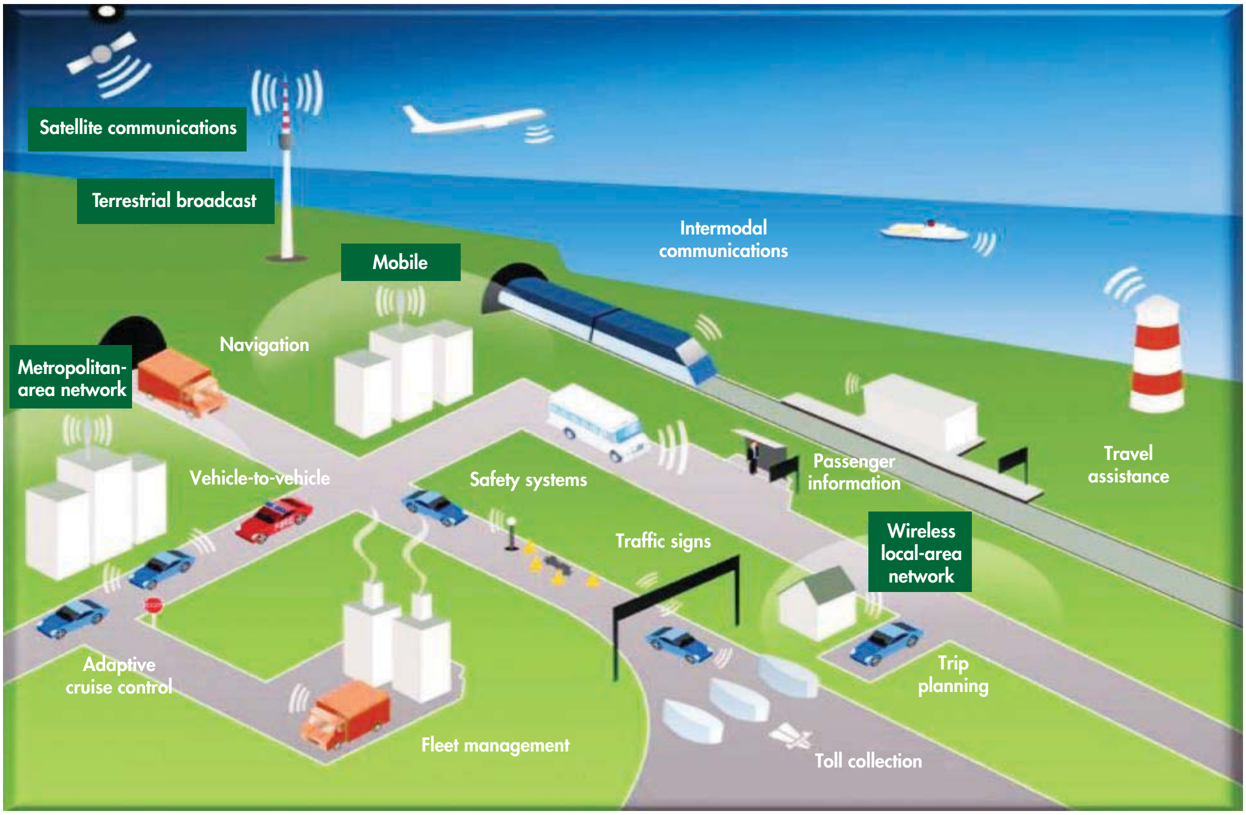

1. For Cisco, a vehicle is just one more thing in the Internet of Things. The key engineering challenges lie in optimizing communications for speed, security, and reliability.

vehicle and vehicle-to-infrastructure communications, helping to protect the safety of the traveling public. It can warn drivers of an impending dangerous condition or event in time to take corrective or evasive actions. The band is also eligible for use by non-public safety entities for commercial or private operations. It operates at 5.9 GHz.

"Moving beyond connectivity issues, security may emerge as the most critical piece of the puzzle," Devlin said. She explained that at Cisco, "Security products and solutions range from VPNs (virtual private networks) to threat defense, to Web filtering, to identity management, and so on." Existing technology could be brought into the vehicular domain in many applications, though it would be necessary to customize it because of differences in scale, footprint, and CPUs, she said.

BUSINESS PATHS

So where are the business opportunities? Cisco is pursuing many possibilities for what could be a common platform with multiple extensions. The key features are high connectivity and speed. Security is important, but it isn't the biggest priority for some people that Cisco works with, Devlin said. Service providers that want to have an aftermarket play represent another business opportunity.

In general, developers for hardware that adds IoT capabilities to vehicles need to look at the situation as a need for a generic platform. "Yes, connectivity is important. But really what we need is an open platform that allows us to securely authenticate and add new applications through which OEMs can monetize various services," Devlin said.

"Looking beyond the feature-level layer, it helps to know what your partners are thinking about in terms of operating systems and silicon. When we work with tier-ones or auto manufacturers… there is a variety of issues that we look at. One is operating systems—what's popular now, what's going to change," she said.

"For example, QNX is fairly widespread today but declining because of its association with RIM," she said. Nevertheless, she added, QNX does have a pronounced market share in the automotive world, so it isn't doomed. However, Linux and Android are increasing in popularity. Similarly, ARM is becoming the predominant CPU architecture because of its low-power requirements.

"When you're building the software stack, you typically start with decisions on which CPU architecture you're going to go with. What's the kernel that you want to have? Then, what are the basic capabilities that are needed?" she said.

P R E C I S I O N • P O W E R • A N A L O G

Speed Monitor.

High Speed, High Bandwidth, and High Current Combination Excels at RF Generation

PA119 POWER AMPLIFIER: 900 V/μs WITH FULL THERMAL PROTECTION

Apex Microtechnology designed the PA119 to drive a variety of loads from DC through the RF frequency range. In addition to its zipped up slew rate, the PA119 boasts 3.5 MHz of power bandwidth and 4 A of output current. For protection, the output stage features a thermal shutoff. An "A" graded model is available for those looking for better input specs.

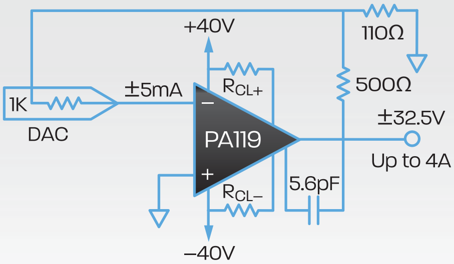

PA119 AS FAST POWER DRIVER
(For Design Discussion Purposes Only)

8-PIN TO-3

Footprint 65.1mm X 42.5mm

Power up at www.apexanalog.com/edpa119

apexanalog.com/edpa119

APEX™ MICROTECHNOLOGY

ISO9001

"Then on top of that, what level of scalability, reliability, and extensibility must the platform provide? As an engineer, that's how I always work. If I look at a spectrum of requirements, I imagine a Venn diagram to see what it takes to build in as much commonality as possible."

For example, accounting for the broad constraints for a product that will allow intelligent, autonomous switching connectivity options such as multiple cellular services and mesh Wi-Fi, a team could design the part of a system that can opt for the best connection solution at any time, based on cost, reliability, and security. Devlin calls that hardware a "policy engine."

Cisco already knows how to do that in the wired Internet. "This involves hardwired decisions. Under these circumstances, you shall switch between these links," she said, along with the ability to learn from what has happened in the past, being able to collect patterns of experiences that point the policy engine toward the advantages of certain kinds of links.

"For example, if you're commuting on the 880 at 8 a.m., and you need to be at work at the Cisco campus at 9, the particular connectivity that you'd use during that time frame might be different than if you were going someplace else at midday," Devlin said.

"What else is needed is the ability to learn, not just from one car's experience but from the experience of a multitude of cars, and to be able to harness that vehicular crowd source by monitoring and accumulating data on what's happening to other vehicles and applying predictive analytics to a spectrum of situations," she said. "That's more than just creating a logic tree based on time of day and location."

In other words, the experience that strangers are having in their cars a mile up the road can affect communications decisions in your car. That's a whole new sense of community. "It's really crowd sourcing at a different level. We've all been frustrated on the road when we suddenly hit a traffic jam and think, 'Gee, I wish I'd known five minutes earlier. I would've passed up the last exit and gone a different way,'" Devlin said.

"This is knowledge-sharing without human intervention. You don't have to tweet or text that 'Hey, I am stopping, and blah, blah, blah.' Instead, every car senses that traffic has slowed down (or sped up) for them, and in effect, tweets the information by means of DSRC, not to you, not to other drivers, but to other cars," she said. "But it will alert you too so you have control. It's through these protocols that I spoke about, such as DSRC vehicle-to-vehicle communication, vehicle-to-infrastructure communication."

This type of information, coupled with in-car mapping intelligence, can do more than simply let drivers know about backups so they can initiate their own alternative routings. It also can automatically access traffic information on all those alternatives and direct cars to different routes to spread the traffic load and improve travel times for every vehicle.

BYOD IN THE CAR

To the horror of corporate IT managers, "bring your own device" (BYOD) is making its way into North American and European businesses. Essentially, technology workers tend to upgrade to newer and more feature-rich gear faster than companies can afford to match. The newest gear is compact and inconspicuous, and inevitably company resources are exposed to the BYODs, which can potentially carry away trade secrets and infect resources. (Think USB drives, Stuxnet, and Iranian uranium-separation centrifuges.)

Since it's naïve to think that company policies can effectively interdict BYODs, the sensible thing to do has been to build security into the company's internal net. Cisco already offers such security products. Now, extend that protection to airliners, commuter trains, company shuttle buses, and carpool vehicles. That's where security becomes a matter of life and death.

What's at risk is no longer limited to companies' trade secrets and financials. BYOD-introduced malware can turn the tables on DSRC traffic control, increasing driving danger and sabotaging vehicular systems in ways one doesn't want to think about. And at a lower level of paranoia, there are more issues for a protocol engine to contend with.

Examples extend from the stranger one picks up at the park-and-ride to make up a carpool to regular carpool members and to the kids going to the museum on the school trip or heading back from the soccer game. In such cases, the driver wants to give passengers who have their own devices controlled access to the vehicle's connectivity.

That means "identity management" for everyone in the vehicle. Guests in the car wouldn't necessarily have access to the same kind of information or the same level of bandwidth as the driver. Or perhaps if the driver is a real-estate agent showing a client an expensive new home, the system needs to be able to say, "This new device belongs to a red-carpet guest. Provide all-out service."

These decisions shouldn't necessarily be made in the car. This part of system design ties into security and policy management. It's a complex interchange of policy handshaking that ideally needs to happen in the cloud whenever the in-car system detects new devices, allowing a more sophisticated approach to providing access.

Taking that a step further, if the new device wants access to parts of the system that are vulnerable to malware, the in-car gateway could direct it first to a malware-cleaning site. That would never be a total solution, but it would provide a level of security. "That kind of thing is possible once you're able to have that end-to-end security, coupled with the ability to manage even the vehicle's ECUs (electronic control units) from the cloud," Devlin said.

OEMs are very interested in over-the-air updates to their ECUs. It's an economics and liability issue. Automotive OEMs

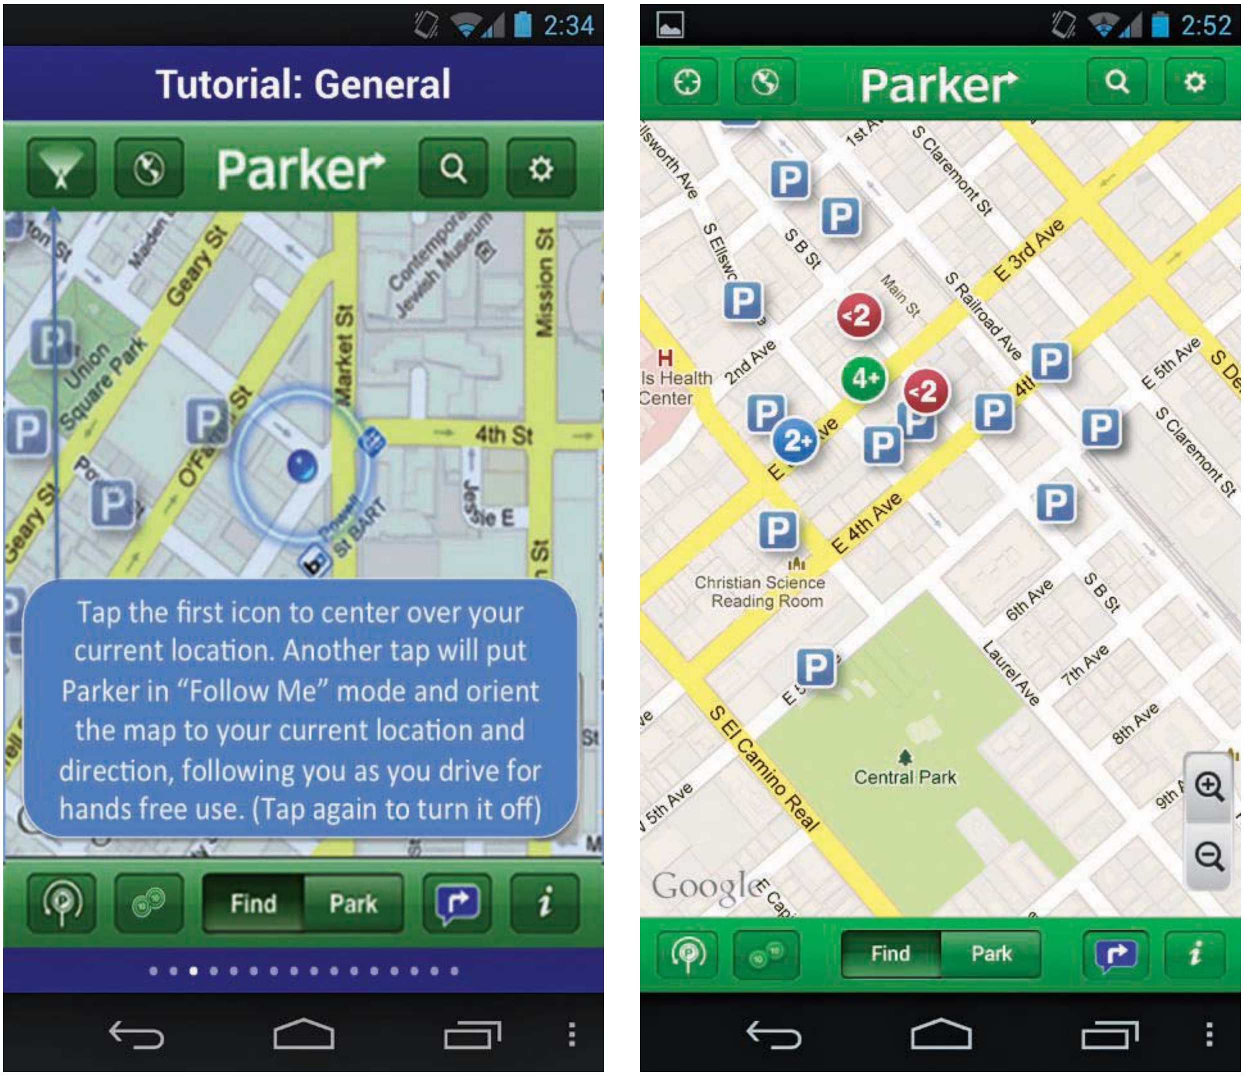

2. Streetline is working with many cities to demonstrate a mobile service that directs automobile drivers to open parking spaces (left). When drivers use the Parker app, a selection of open parking spaces appears on their phone (right).

in particular don't want ECU updates to be a cause for a major vehicle recall, nor do they want to rely on car owners to make appointments to have the work done at a dealer. "It's the ability to 'securely' provide that update to a unit and let it subsequently propagate to sub-tending ECUs that constitutes a huge money saver for OEMs," Devlin said. "But it has to be secure. It has to be reliable. It has to be scalable."

And if the data transfer drops for some reason before the update is complete despite all of its speed, security, and reliability, the system must be able to pick up where it left off once it reconnects. Inevitably, however, there are privacy and control issues that are connected to providing the potential for so much connectivity.

If your ECU can "phone home" for a factory update if its diagnostics indicate a flaw, can a built-in Breathalyzer call the cops (or less intrusively, a designated driver) if you are too drunk to drive? Could you limit the places your kids are allowed to go (known as geo-fencing) or demand a graphic display of speed and location when they come home?

UPGRADING AUTOMOTIVE BUSES

Looking ahead, each vehicle could support its own local network. Would an IP-based (Internet protocol) combination of RF and fiber replace some hardwired buses in the car? Would that be bulletproof enough? Cisco has been involved in a number of discussions on that front.

"There is a huge interest on the part of automotive manufacturers looking at this mesh that I like to think of as an instance of organic networking," Devlin said, though it isn't going to evolve in some predictable way.

"You are looking at a landscape that is being defined as we speak and there are multiple paths in which things can take shape. So we're actually placing bets on multiple fronts, through proof of concepts and experimenting with different ideas with different partners and customers," she said. "We don't have a definitive 'A is going to come before B before C,' but we do see A, B, C, D, E, F. We see the possibilities and we're placing bets among all those possibilities."

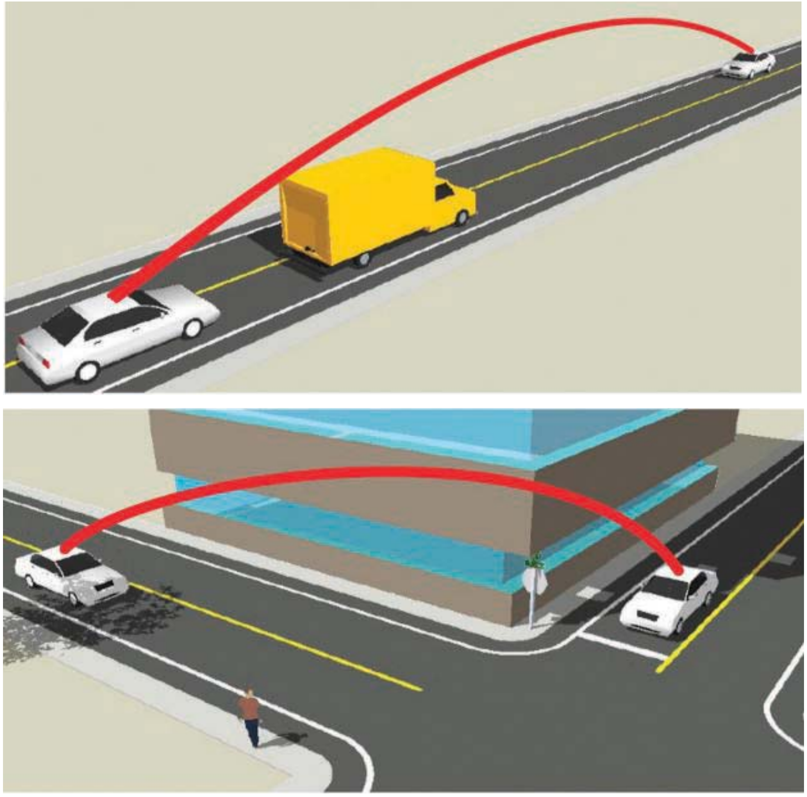

3. The Cohda Wireless DSRC (802.11p) radios employ patented techniques in the receiver to increase data throughput and extend range. This pays off in earlier warnings about road dangers about oncoming traffic just beyond that slow-moving semi you want to pass on a hill (top) or a speeding car just around the corner at a blind intersection (bottom).

SMART PARKING TECHNOLOGY

Your car is already part of the IoT if you park it around the intersection of 3rd Avenue and B Street in the business section of San Mateo, Calif., while shopping or dining. Something similar is true for 64 other cities, and the depth of coverage is growing.

Since last winter, San Mateo has been one of several sites for wireless sensor demonstrations in which drivers use their smart phones to find open parking spots. Another demonstration that allows drivers to find spaces in city lots is taking place in nearby San Carlos, Calif. The demos are a joint effort of Cisco and a company called Streetline. As the experiments feed back data to the cities, Streetline expects to increase penetration and turn the demos into revenue generators.

The technology uses buried sensors at each parking space—essentially, each parking space is an Internet "thing" with its own IP address—and a free consumer app called Parker from Streetline along with the intelligent networking technology platform from Cisco *(Fig. 2)*.

Sensors embedded in the pavement of the selected parking spaces detect when a space is available. Cisco's smart routers communicate with Streetline sensors to aggregate sensor data and, at the same time, communicate with the Streetline cloud center to deliver the availability of the parking spots. The intelligent network platform captures the data and publishes it into Parker, which displays real-time availability of these on-street spaces as well as locations of off-street parking garages and lots along with other information such as pricing and enforcement hours.

The Parker app's hands-free voice navigation system guides drivers to parking facilities using an audible cue when available parking is nearby. It also can toggle between availability and price, including real-time updates as prices are changed or updated. In fact, with the cooperation of the municipality, it's possible to access rates, hours, and time limits before leaving home.

Drivers can even tell the app where they want to go and be guided to the nearest open parking spaces. Friendliest of all, the app can guide drivers back to their cars after dinner or the show. The Parker app on my phone lists coverage in 65 cities. It runs on iOS and Android.

AUGMENTING AWARENESS

Cohda Wireless has investment and technical support from Cisco and NXP. Its hardware supports cooperative intelligent transport systems (CITS) that use vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) communications to enable cooperation between vehicles and road infrastructure to improve safety, mobility, and the environment.

Both the United States and Europe have released 5.9-GHz spectrum for CITS. Major trials are underway around the world. Foundation standards for CITS are the same in both the U.S. and Europe, and they have already been published. Cohda has a unique hardware solution for 802.11p receivers, which it says has advantages over radios from other CITS suppliers that rely on commercial off-the-shelf hardware.

Generally, CITS provides each vehicle with 360° awareness of all other vehicles on the road. If a threat is identified, the driver is immediately warned and can take steps to avoid an accident. Cohda Wireless says its technology extends this 360° awareness beyond buildings and around trucks that block the driver's view, enabling drivers to be aware of all threats, literally around corners.

The company has performed multiple demonstrations at transportation safety events. Typically, the trunks of two cars are fitted with identical DSRC demonstration systems, and antennas are mounted at the rear of the roof.

In the trunks of the demo cars, the Cohda Wireless radio interfaces with an STMicroelectronics STA2062 Cartesio Application processor for portable navigation, in-vehicle navigation, and telematics, which integrates a 32-bit ARM core with a 32-channel GPS subsystem and an array of ports for peripherals, including CAN, USB, and SPI.

The STMicroelectronics platform GPS receiver provides data about current location and time. This information is used to build SAE J2735 basic safety messages, which are then broadcast via the Cohda radio in one of the vehicles at a rate of 10 messages per second. The other Cohda radio receives those messages and forwards them to the processor, which extracts time-stamped positions. Combining knowledge of the local area and remote vehicles, the processor then geometrically assesses potential danger and issues warnings as appropriate.

At several international automotive shows, Cohda has conducted on-the-road comparisons between its products with the patented receiver modifications and competing hardware built using standard Wi-Fi chips. It says that by actual measurement in live conditions, Cohda devices can allow 10 to 15 times more data to be exchanged between vehicles and/or infrastructure and that the useful range is extended from around 20 meters to 200 meters.

Regardless of the type of transmitter in the DSRC radio in the other car, Cohda's system can provide a driver with a "do not pass" warning about oncoming traffic in the opposite lane when the closing time is on the order of 20 seconds, versus 3 seconds, if both radios use standard radios *(Fig. 3)*. When the potential collision involves a vehicle around a blind corner, rather than in a situation passing a big rig, warnings are available when there is still 6 seconds or more to react of time, compared to less than 2 seconds. ed

REFERENCE

1. Dedicated Short Range Communications (DSRC) Service, http://wireless.fcc.gov/services/index.htm?job=service_home&id=dedicated_src

BRIG ASAY | Agilent Technologies
brigham_asay@agilent.com

THE FUNDAMENTALS OF OSCILLOSCOPE TRIGGERING

Hardware and software triggers are essential elements in today's oscilloscopes, adding speed and saving all-too-precious time for data gathering.

A trigger is an act or event that stimulates and initiates or precipitates a reaction or series of reactions. The same definition applies to an oscilloscope trigger, which involves waiting for an event to occur and triggering upon its occurrence. The oscilloscope then captures and displays the electrical signaling (data) that follows the trigger event. For more complex events, an oscilloscope should incorporate advanced triggering capabilities. Overall, triggers often are the deciding factor in oscilloscope purchases.

THE IMPORTANCE OF TRIGGERING

"Banner" specifications such as bandwidth, memory depth, and sample rate typically define oscilloscopes. However, advances have pushed real-time oscilloscope banner specifications far beyond most technology needs. As a result, banner specifications tend to be qualifiers when evaluating an oscilloscope, not the deciding factors.

Oscilloscope users no longer have to settle for oscilloscopes that only trigger on an edge, because they can find scopes with advanced triggering in both hardware and software. Thanks to modern advances in triggering technology, a good triggering system will save hours of time trying to isolate unique problems in a design and, more critically, allow designers to quickly improve their designs. The potential time savings made possible with a good trigger underscore the importance of fully understanding the ins and outs of the triggering system and the available triggers.

THE EDGE TRIGGER

At the heart of every basic, advanced, and software trigger is the edge trigger—the key building block of all triggers. It's conceptually the easiest trigger to understand. The edge trigger comes in three forms: rising, falling, or rising and falling.

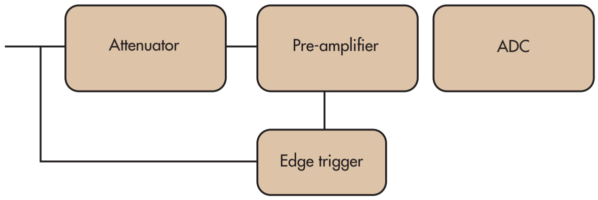

1. The edge trigger, located in the front end of an oscilloscope, is the key component of the entire system.

It resides at the core of the triggering system, providing a window into its workings.

Oscilloscope designs include a front end and back end. The front end typically includes the preamplifier, the edge-trigger chip, attenuators, and the analog-to-digital converter (ADC) *(Fig. 1)*. The edge trigger is the key component of the entire system. The signal path goes through both the edge trigger and the preamplifier. Inside the preamplifier, the data is ignored until the edge trigger confirms that the data is available and should be digitized.

When observing the basic hardware functionality of an edge trigger, it can be categorized as a simple comparator that looks to see if the input signal crosses the entire threshold level *(Fig. 2)*. If the signal crosses the threshold, the preamplifier is told to process the data and the signal is digitized.

EDGE-TRIGGER LIMITATIONS

Of all the specifications, perhaps edge-trigger bandwidth is the least understood. For less expensive oscilloscopes, vendors don't even specify edge-trigger bandwidth.

The edge-trigger chip's comparators will only work to a certain frequency. If the signal frequency exceeds the edge-trigger bandwidth, the edge trigger will not trigger, even though the signal passed both thresholds. That's due to the timing

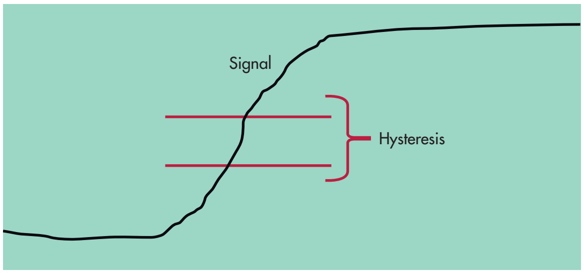

2. An event triggers when the input signal crosses the entire threshold level. The preamplifier then processes the data, and the signal is digitized.

between the comparators—if the signal passes too fast, the edge trigger will know that a signal crossed one threshold, but will not be fast enough to know if the second one was crossed.

Thus, trigger sensitivity is specified to accelerate the edge-trigger bandwidth. Essentially, trigger sensitivity specifies how large the signal needs to be at certain frequencies for a chip to trigger. The larger the signal, the easier it is for the edge trigger to work and, hence, the faster the trigger can function. Say, at 1 GHz, the amplitude must be larger than 1 MHz due to the speed of the signal.

It seems there's always something more to the story for every specification, and the edge trigger is no exception. For instance, how can a 60-GHz oscilloscope trigger on a 60-GHz sine wave when the edge trigger is less than 60 GHz?

A couple options for triggering (auto and triggered) exist inside the scope. The edge trigger is the key qualifier for the "triggered" system to work. For "auto," though, there's no need for the edge trigger. The "auto" triggered mode ignores the edge trigger and lets all signals pass, whether or not they're qualified. It then uses software to align the edges.

For example, in the case of a 60-GHz sine wave, the edge trigger would fail and not trigger. Therefore, the oscilloscope would never enter "triggered" mode. But if the user understood this limitation and changed the oscilloscope to "auto," the oscilloscope would digitize the 60-GHz sine wave through software triggering.

ADVANCED (SMART) TRIGGERS

Basically, the edge trigger's main function is to look for an edge. However, it also keys more advanced triggers. When the edge trigger sees the edge, it "triggers" another block of the oscilloscope—known as the "smart" trigger—to begin its work. The smart triggering system can consist of a number of different triggers, including runt, pulse width, window, and setup and hold.

In many cases, the smart trigger holds the key to simplifying the debugging of a designer's board-level problems. It can be used to find events that occur very infrequently. One misconception is that because the smart trigger is done in hardware, it provides zero dead time. Rather, the smart trigger looks at a data set in hardware, scanning for the anomalous event. If the event occurs in the data that it scans, it triggers. If it doesn't trigger, it waits for the edge trigger to indicate that more data is ready to be scanned.

RUNT TRIGGERING

Runt pulses are pulses that don't reach a valid high or low level. They cause designs to fail and will appear for a variety of reasons. With the handy runt trigger, designers can quickly find runt pulses and evaluate what caused them.

Oscilloscopes can trigger on these pulses by setting two threshold levels and looking for signals that only cross one of them. Once there's a successful trigger on the runt pulse, the scopes can pan and zoom to locate the problem area.

SETUP-AND-HOLD TRIGGERING

Setup and hold is critical for many of today's applications. Setup-and-hold violations can destabilize a system. This instability can cause an IC to oscillate between a high and low state, which introduces inaccuracies into the system.

Setup-and-hold triggering allows users to quickly scan a system for violations. Users simply need to set the instability-inducing specification and then look for the trigger. If the oscilloscope finds a violation of the specified setup or hold time, it will then trigger.

Subsequently, the oscilloscope can be used to find what preceded the failure by looking at typical and atypical failure mechanisms, such as slew rate or asynchronous behavior. If the scope is unable to trigger, then it becomes apparent that the system isn't failing due to setup or hold violations. This debug analysis capability can be invaluable for both analysis and debugging system problems.

3. With software-based triggering, one can, for instance, draw multiple shapes on screen and specify if a signal can go through that shape or zone.

PULSE-WIDTH TRIGGER

The pulse-width trigger, another smart trigger, looks for pulses that are smaller or larger than a certain width. As a result, it's sometimes referred to as the time-qualifying trigger. The pulse-width trigger can work on positive or negative pulses. Typically, when users trigger with the pulse-width trigger, they're looking for pulses that are faster than the reference clock. If this event occurs, it could mean a significant design problem. Pulse-width triggering offers fast and simple debugging in this case.

Still, among all of the smart triggers, the pulse-width trigger seems to lag the furthest behind today's technology. Currently, the fastest pulse-width trigger operates at approximately 200 ps, which is slower than today's fastest data-rate unit intervals (ranging well below 50 ps). Thus, it's important to know the oscilloscope's pulse-width limit. Enhancements through software can help accelerate the pulse-width trigger. In fact, recent advances allow software triggers to trigger on pulse widths as fast as 40 ps.

SOFTWARE TRIGGER

In the last decade, software triggering has advanced to where it's now considered a key part of an oscilloscope trigger. In 2006, Agilent Technologies introduced the InfiniiScan software-based trigger system. Using a software trigger, InfiniiScan can look for narrow pulse widths (half the pulse width of hardware triggers), monotonic edges, pattern trigger, and zone qualify.

Software-based triggers work much like hardware-based smart triggers. The hardware edge trigger finds an edge and the signal is digitized. Then the software scans the acquisition to see if the event has occurred. If so, the software triggers on the event.

The upside of software-based triggering is its extreme flexibility. Software packages make it possible to draw multiple shapes on screen and specify whether the signal can go through the shape or zone *(Fig. 3)*. The downside is that software runs much slower than hardware triggers. If an event occurs infrequently, the software trigger will likely miss it. The hardware trigger is more likely to catch the event, assuming that what can be done in hardware is also possible in software. ed

BRIG ASAY product manager for high-performance oscilloscopes at Agilent Technologies, received a BSEE from the University of Wyoming, Laramie, and an MBA from Northwest Nazarene University, Idaho.

High Voltage ICs and MOSFETs for Industrial Applications

EL Driver ICs | HV Amplifier Array ICs | Ultrasound ICs | PWM Controller ICs

Optical-Switching and Telecom ICs | MOSFETs | Specialty ICs | LED Driver ICs | HV Interface ICs

• Depletion and Enhancement Mode MOSFETs

Low-threshold transistors that utilize an advanced vertical DMOS structure and a proprietary silicon-gate manufacturing process

Applications	Features	Benefits
• Normally-on switches	• Exceptional power handling capability	• Low threshold voltage
• Solid state relays	• High input impedance	• High breakdown voltage
• Proximity switches	• Positive temperature coefficient	• High input impedance
• Linear amplifiers	• No thermal runaway or thermally-induced secondary breakdown	• Low input capacitance
• Constant current sources		• Fast switching speed
• Power supply circuits		
• Input protection circuits		

Proximity Switches

• Fault Protection Switches

AC low side switches with current fold-back protection that protects components from short circuit conditions

Fault Protection Switches

Applications	Features	Benefits
• Solenoid valve control	• Up to 100V input voltage protection	• Protects circuitry from over-current and short circuit conditions
• AC relay control	• Low on resistance	• Minimize voltage drop to save power
• Relay replacement	• Fast switching speed	• Saves circuit board space
• Resettable fuses	• No external supplies needed	
• Downhole drilling		

• Linear and Switching Regulators

Drive single or multiple LEDs at a fixed, constant current to enhance LED performance and extend lifetime

Applications	Features	Benefits
• Industrial lamp indicators	• Wide input voltage range	• Consistent LED brightness
• Signage	• Temperature compensated current	• Simple circuit design
• Constant current source/sink	• Over-temperature protection	• Multiple devices can be combined for higher currents
• LED wall dimmers		
• Start-up power supply		

Industrial LED Lighting

• Piezoelectric Drivers

Designed with advanced HVCMOS technology for applications that require multiple high voltage outputs and current source and sink capability

Industrial Printers

Applications	Features	Benefits
• Inkjet printers	• Output voltages to 220V	• Print various materials
• Electrostatic print heads	• Perform both polarity and blanking functions	• Fast system diagnostic
• Braille drivers	• Forward and reverse shifting options	• Easy board layout

DAVID A. HALL | Contributing Technical Expert
david.hall@ni.com

UNDERSTANDING SIGNAL ANALYZER ARCHITECTURES

The three-stage superheterodye vector signal analyzer uses a reasonably straightforward architecture to cleanly downconvert signals from RF to IF.

Many engineers who use spectrum analyzers on a regular basis might be content to know that their instruments will produce a display of "power versus frequency" with little idea of what's going on inside. Despite this, understanding the architecture of these instruments is important to maximizing their performance in practical applications.

An RF signal analyzer must downconvert signals from RF to a frequency range that can be digitized with an analog-to-digital converter (ADC) with anywhere from 12 to 16 bits of resolution. The mixer—and its frequency translation properties—is at the heart of the downconversion process, though I won't focus on mixer theory of operation here *(see "What's Inside Your RF Signal Analyzer?" at electronicdesign.com)*.

Instead, I want to focus on two main vector signal analyzer (VSA) architectures, exploring both the theory of how they downconvert signals to baseband and the tradeoffs between them.

SINGLE-STAGE DOWNCONVERSION

The single-stage downconverter is by far the simplest technique that can be employed to downconvert signals from RF to an intermediate frequency (IF). The premise of this architecture is that a single broadband mixer (hence the term "single stage") will be used to mix signals from a wide range of radio frequencies to a fixed IF filter.

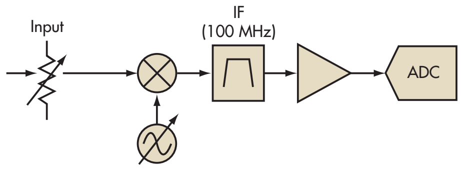

1. The single-stage VSA uses a single mixer to translate the signal of interest to IF.

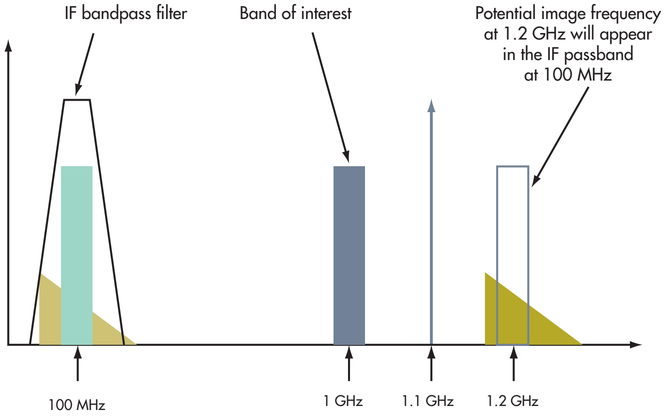

2. Single-stage VSAs don't inherently protect the bandwidth of interest from out-of-band images.

Suppose a single-stage VSA uses a final IF frequency of 100 MHz that will be digitized by a 14-bit, 250-Msample/s ADC *(Fig. 1)*. Given the choice of ADC, our example VSA will be sampling in the second Nyquist zone of our ADC, which would range from 75 MHz to 150 MHz. (The first Nyquist zone would be dc to 75 MHz.) Also, we'll need a tunable local oscillator (LO) to drive the mixer, which is usually a voltage-controlled oscillator (VCO) in single-stage designs.

In practice, the single-stage VSA operates by tuning its LO to a frequency that is offset the signal of interest by exactly the IF. Consequently, to set a center frequency of 1 GHz on our VSA, the instrument must internally tune the LO to either 1 GHz – 100 MHz = 900 MHz, or 1 GHz + 100 MHz = 1.1 GHz. In the former case, setting an LO at 900 MHz is referred to as "low-side" injection, and setting the LO to 1.1 GHz is "high-side injection."

IMAGES IN THE SINGLE-STAGE VSA DESIGN

While our two frequency choices for an LO might seem beneficial, the notion of high-side and injection presents an inherent challenge in the single-stage architecture. Suppose we choose high-side injection in the example above, setting our LO frequency at 1.1 GHz. As it turns out, while we are “high-side injecting” the 1-GHz band, we are simultaneously “low-side injecting” the 1.2-GHz band.

In fact, given that 1.2 GHz – 1.1 GHz = 100 MHz, and that 100 MHz is our IF frequency, any 1.2-GHz signal present at the RF port of the mixer will also pass through our IF filter and obscure the signal of interest. Figure 2 shows how signal content at the image frequency of 1.2 GHz will also mix to exactly 100 MHz in this scenario.

> “The single-stage downconverter is the simplest technique that can be used to downconvert signals from RF to an intermediate frequency. A single broadband mixer will be used to mix signals from a wide range of radio frequencies to a fixed IF filter.”

While this classic imaging scenario is common to all single-stage VSAs, the single-stage design is highly appropriate in a broad range of applications. Single-stage VSAs generally tend to have a smaller footprint and fast tuning times. They also typically cost less and still deliver excellent performance when analyzing modulated signals. As a result, the single-stage design is common where the frequency range of signals present at the RF port is reasonably constrained, such as production test of cellular handsets, wireless connectivity devices, and RFICs.

IMAGE REJECTION

To deal with image rejection, many RF signal analyzers employ multiple mixing stages to adequately remove out-of-band signals from the band of interest. The three-stage design is relatively straightforward, though fully grasping how the frequency plan works requires a bit of mental gymnastics.

From the single-stage example, we learn that the “image frequency” is always exactly “one IF” away from the center frequency of the RF signal analyzer. Thus, if our VSA can filter signal content from the image frequency in the first place, then we can ensure that these images do not obscure the signal we are attempting to analyze. A classic solution is to use a first IF that is high enough in frequency so it will fall outside the frequency range of our instrument in conjunction with a low-pass filter.

Suppose we were building an RF signal analyzer designed to analyze signals from 20 Hz to 3.6 GHz. In this scenario, we could employ a low-pass filter as one of the first components in the VSA design to ensure that signals higher than 3.6 GHz are removed before ever hitting the first mixer. Next, we would choose a first IF that was sufficiently high in frequency to ensure that our 3.6-GHz low-pass filter has “cleaned” the image frequency.

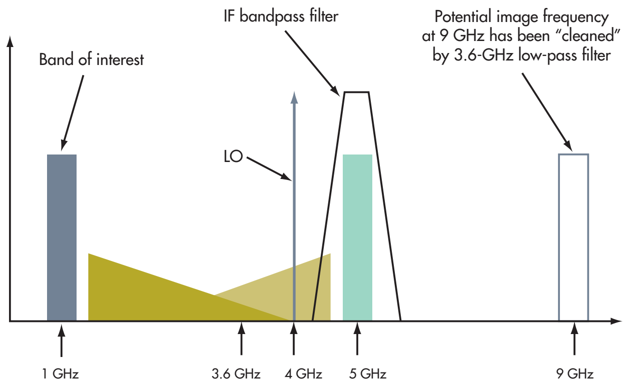

3. Using a high intermediate frequency in conjunction with high-side injection ensures that the frequency content at the image frequency does not mix into our IF.

SUPERHETERODYNE FREQUENCY PLAN SCENARIOS

Given the requirements for a superheterodyne downconverter, consider a 20-Hz to 3.6-GHz VSA design where the first IF stage is at 5 GHz. High-side injecting a 1-GHz signal to produce an IF at 5 GHz would require us to set our LO at 4 GHz (1 GHz + 4 GHz = 5 GHz). Moreover, high-side injecting a 1-GHz band with an LO of 4 GHz means we are simultaneously low-side injecting the 4 GHz + 5 GHz = 9-GHz band. The 3.6-GHz low-pass filter already has “cleaned” this band, ensuring that images at this frequency do not obscure the band of interest *(Fig. 3)*.

After the first IF filter of a superheterodyne signal analyzer, one can be reasonably sure that out-of-band signals have been removed by a combination of both the front-end low-pass filter and the IF bandpass filter. Next, subsequent mixing stages gradually mix the high-IF signal to second and eventually a third IF frequency—using band-pass filters to remove unwanted mixing products *(Fig. 4)*.

In the three-stage design, multiple stages are used to ensure that at each mixing stage, the intermediate frequency at the output of each mixer is far enough from its mixing frequency to prevent imaging from occurring. Here, the rolloff of each of the IF filters primarily determines the likelihood of imaging, and using a second IF at a reasonably high frequency improves

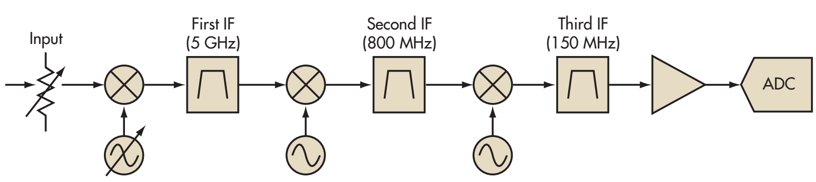

4. The second and third stages of a superheterodyne signal analyzer ensure that a signal is gradually downconverted from the first IF to the final IF.

the instrument's image rejection. In the end, the final IF is low enough in frequency to be digitized by the ADC. After digitization at IF, the VSA will digitally downconvert the IF signal to IQ and perform a complex fast Fourier transform (FFT) to present as spectrum information.

MORE FROM DAVID A. HALL

As a Contributing Technical Expert, David A. Hall writes about test & measurement issues for ***Electronic Design*** and for ***Microwaves & RF***. For more of his work on ***Electronic Design***, go to http://electronicdesign.com/author/david-hall and see:

- VNA's: The Myth, The Mystery, The Electromagnetics
- What's Inside Your RF SIgnal Analyzer?
- Instrumentation Improves Microwave System Simulation

For more of his work on ***MW&RF***, go to http://mwrf.com/author/david-hall and see:

- The Effect Of Moore's Law On RF Instruments
- Find Out Why PXI Is Being Used For RF Instruments

PARTING THOUGHTS

The three-stage superheterodye vector signal analyzer uses a reasonably straightforward architecture to cleanly downconvert signals from RF to IF. Understanding these architectures will help you to understand how to optimize your VSA for making challenging measurements such as intermodulation distortion (IM3)—a topic we'll save for next time! ed

DAVID A. HALL is a senior product marketing manager at National Instruments, where he is responsible for RF and wireless test hardware and software products. He holds a BS degree with honors in computer engineering from Penn State University.

BILL WONG | Embedded/Systems/Software Editor
bill.wong@penton.com

LSI's Kimberly Leyenaar
Discusses Big Data

"Big Data" takes lots of storage, and LSI is at the forefront in providing platforms that can deliver the massive amounts of storage that these applications require. Kimberly Leyenaar, principal Big Data engineer and solution technologist for LSI's Accelerated Storage Division, focuses on discovering innovative ways to solve the challenges surrounding Big Data applications and architectures.

ED: What are the challenges in dealing with Big Data?

LEYENAAR: Most organizations have heard of Big Data, and certainly data mining and business intelligence are no strangers to most businesses. However, the challenges are great. When IT managers begin implementing Big Data solutions, the first question is always, "What data do I keep?" While certainly not a barrier to entry, this has been cited as the number one challenge.

The cost of storing data has gone down significantly, leading many businesses to the "save everything" conclusion for fear that latent information not previously known could be mined from these sources ex post facto. It's well accepted that "more data beats better algorithms." However, this also brings other challenges such as how to handle transfer of data for eventual hardware end-of-life cycles as well as potential legal issues that come with storing large amounts of (potentially private) information. Another challenge is finding the right talent to help implement, manage, and apply the proper analytics in order to ultimately extract the business value within the data sources.

Enter the newest job title, the data scientist, a professional with a special mix of business acumen, statistical and mathematical expertise, software engineering, high-performance computing, patterns, and a little bit of magical pixie dust. Today, dozens of colleges and universities offer data scientist degrees, but the supply is far below the demand as the skill sets needed to effectively implement and fully exploit the value of a Big Data project cannot be found in the general IT population. Because Big Data ecosystems are fairly premature, we are still discovering challenges every day.

ED: What application tools are being used in this arena?

LEYENAAR: While Hadoop is the primary software, the landscape is fertile and I counted over 200 different software applications that, in some way, help IT professionals and data scientists in their quest to derive the value from the Big Data infrastructures. We can divide the application tools into three distinct groups: applications that help build/analyze data queries, applications that help manage server clusters, and finally applications that help manage data sources, structures, and flows.

Hadoop is written in Java, while many business intelligence personnel are accustomed to SQL-based data interaction, and there exists an abundant amount of excellent software meant to alleviate the burden of learning new languages such as Hive, Pig, and Ruby, as well as tools meant to integrate workflows (Oozie) or apply well known query algorithms (Mahout).

Because Hadoop can interface with diverse data sources, there are numerous tools meant to extract data and/or provide connectors to and from other resources (e.g., databases, online data stores) such as Sqoop, as well as software that allows easy data ingest to Hadoop data nodes such as Flume. There is a popular trend to apply organization on top of Hadoop data stores such as noSQL, key/value, or columnar or modular structured data layouts using tools like Cassandra or Hbase.

Finally, because clusters require continuous care and feeding, software such as Zookeeper, Ambari, and Zettaset helps monitor cluster health and maintenance. With so many software choices, tool integration can become daunting. Thus, we will see more Hadoop distributions integrate their own complete toolsets, removing this burden from the end users. We saw this same phenomenon with Linux over the past 20 years, and I anticipate we will continue to see this trend as Hadoop matures.

ED: What types of underlying hardware and software does LSI offer

designers that support these applications?
LEYENAAR: LSI is committed to helping the community solve the problems associated with Big Data. The ability to run a query against a whole dataset and get results quickly is transformative, and it changes the way we all think about data. By improving a Hadoop cluster's efficiency, businesses can run more queries, more batch processes, and that means unlocking the value of their data so they can start innovating their marketing and sales strategies and processes in ways they never knew sooner.

Hadoop clusters are a major investment, and smart IT managers are looking for ways to significantly improve the TCO (true cost to own) and ROI (return on investment) while maintaining cluster availability. The LSI Nytro WarpDrive product line provides the lowest-latency, highest-performing managed PCI Express flash, with capacities up to 3.2 Tbytes, enabling the data to be even closer to the processing and the business closer to their answers. The LSI Nytro MegaRAID provides the best solution with enterprise RAID (redundant array of independent disks) capabilities to add up to 128 devices and integrated on-board flash-based cache with capacities up to 800 Gbytes *(see the figure)*.

Nytro MegaRAID sets the standard for OEM hardened storage, and with eight lanes of 6-Gbyte/s SAS (serial-attached SCSI), it can push up to 4200 Mbytes/s and well over 400k I/Os per second with enterprise features such as RAID, automatic data repair, and SMART (self-analysis, monitoring, and reporting technology) support. The intelligent, integrated cache automatically detects frequently accessed "hot" data and moves it to the high-performing flash.

In our lab, we observed the map-reduce scratchpad operations being

LSI's Nytro MegaRAID supports up to eight lanes of 6-Gbyte/s SAS and transfers up to 4200 Mbytes/s.

moved to the flash (spillover files, intermediate result files, etc.), while HDFS (Hadoop distributed file system) data is read or written sequentially on LSI's high-performing SAS-based architecture. The result is a perfect marriage of high throughput and fast computations that allows most Hadoop jobs to complete in a third less time.

Although Hadoop is designed for inexpensive commodity hardware, Nytro MegaRAID can actually lower the total infrastructure cost while providing highly reliable storage. After all, Big Data is about their data, and since an organization's success will be dictated by its ability to extract value and derive innovation from the data available to it, why would an organization trust poor-quality storage hardware?

ED: How does the OEM partnership with Intel work, and what areas will you be addressing together?
LEYENAAR: LSI has entered into an expanded OEM relationship with Intel, whereby LSI Nytro MegaRAID technology will be available as part of the Intel RAID product family for Intel server boards and systems. Specifically, Intel will offer Nytro MegaRAID technology within its Intel RAID SSD (solid-state disk) cache controllers RCS25ZB040 and RCS25ZB040LX, which include embedded flash of 256 Gbytes and 1 Tbyte, respectively.

The OEM relationship with Intel will broaden LSI's global presence and reach for Nytro MegaRAID technology through Intel's extensive channel network. LSI will also continue to offer Nytro MegaRAID adapters through the LSI worldwide network of distributors, integrators, and VARs (value-added resellers).

LSI Nytro MegaRAID adapters help enterprise customers cut latency and cost-effectively boost performance for applications such as online transaction processing, Web serving, Hadoop/NoSQL, VDI (virtual desktop infrastructure), and other data-intensive workloads. Benchmark testing using Nytro MegaRAID cards have achieved up to a 33% improvement in the time it takes to complete Hadoop jobs and delivered support for up to twice as many VDI sessions compared to a non-caching storage implementation. Nytro MegaRAID cards also enable HDD (hard-disk drive) array rebuilds to complete up to 10 times faster.

ED: How will this partnership improve the handling of Big Data, and what can you do together that would not work as well individually?
LEYENAAR: The partnership will benefit both companies as organizations look to accelerate key business applications while preserving existing investments in HDD and DAS (direct attached storage) infrastructure. The OEM relationship will also benefit both companies as more and more organizations adopt Hadoop implementations and flash-enhanced big data analytics. ed

---

KIMBERLY LEYENAAR is a principal Big Data engineer and solution technologist for LSI's Accelerated Storage Division. She is an electrical engineering graduate of the University of Central Florida and has been a storage performance engineer and architect for more than 14 years.


GO mouser.com
The Newest Products for Your Newest Designs®
The Next Big Thing Is Here.
NEWEST
PRODUCTS
More New Products
More New Technologies
More Added Every Day
Authorized distributor of semiconductors
and electronic components for design engineers.
MOUSER ELECTRONICS
Mouser and Mouser Electronics are registered trademarks of Mouser Electronics, Inc. Other products, logos, and company names mentioned herein, may be trademarks of their respective owners.

U.S. Distributors
Seek Success In Europe

Electronics distributors are expanding in the continent despite ongoing economic weakness across the region.

ALTHOUGH THE EUROZONE ECONOMY continued to suffer this spring, with more weakness expected ahead, U.S.-based distributors remain undeterred in their efforts to grow across the region. Many are capitalizing on ongoing design activity throughout the area while others are growing their production-related business and seeking to build a more global footprint overall.

A handful of those companies spent time at last month's Electronics Distribution Show in Las Vegas emphasizing their global footprint and outlining plans for growth in Europe. Mouser Electronics, which has grown considerably in Europe over the last few years, is one. Company leaders pointed to its strength as a global operation during the yearly conference, which brings together manufacturers, distributors, and manufacturer representatives. Mouser president Glenn Smith pointed to its record growth in

Continued on Page 50

Transportation
Sector Fuels Growth

Distributors respond to new trends in automotive and commercial vehicle markets as demand for creature comforts grows alongside the trend toward lighter, faster, and smaller components.

THE AUTOMOTIVE MARKET HAS been doing well despite sluggish global economic growth. Distributors say the trend is even greater in the broader transportation market, which includes commercial vehicles and agricultural and construction equipment.

Joe Venturella

TTI Inc. formed its transportation business unit about seven years ago, grouping automotive and non-automotive vehicle businesses under one umbrella. The group has outperformed TTI's other segments over the past few years, says vice president Joe Venturella, adding that he expects more of the same in the months ahead. Distribution Resource caught up with Venturella in early May to talk about the transportation sector and where it's heading.

DISTRIBUTION RESOURCE: Many distributors describe the automotive/transportation market as a bright spot in the slow economic recovery. Do you agree? What is your outlook for this market over the next few months?

JOE VENTURELLA: This market segment has done very well for us, and we expect this to continue for the balance of 2013. This segment has outperformed our other segments for the

Continued on Page 54

Continued from Page 49

every region in the first quarter of 2013, emphasizing its 19 locations, technical support, customer service, and marketing personnel worldwide.

"Think of us as a global operation that needs to be aligned with your global marketing teams," Smith told an audience of mostly suppliers during a company update meeting at EDS.

EUROPE'S ROLE

That message rings true in Europe, which remains Mouser's fastest-growing region, according to Mark Burr-Lonnon, who heads the distributor's operations in Europe, the Middle East, and Africa (EMEA). Mouser grew 20% in EMEA last year and is on track to grow another 25% this year, Burr-Lonnon says. He credits the company's focus on new product introductions and customer recruitment as keys to that success.

Mouser has been rounding out its supplier mix, adding 24 new suppliers since January alone while focusing its marketing efforts on attracting new customers. Mouser adds nearly 700 customers per day, with 38% of them coming from Europe and Asia, leaving plenty of opportunity to build the brand in those regions, Burr-Lonnon adds.

"We have carved a niche with design engineers, focusing on new products," Burr-Lonnon says. "We're one of the few guys out there not growing into production [business]. We are where we want to be in Europe and Asia."

Those positive sentiments come despite Europe's ongoing economic crisis. The Eurozone's gross domestic product fell at an annualized rate of 0.9% during the first three months of this year, marking the sixth straight quarter of a recession that began late in 2011. Although economic conditions have improved elsewhere around the world—GDP grew 2.5% in the United States, 2% in the United Kingdom, and 3.5% in Japan during the first quarter—analysts held out little hope that Europe would experience anything resembling a recovery this year. As a result, U.S. distributors are focused on building their brands with a new audience and capitalizing on existing design and production work throughout the region.

Mark Burr-Lonnon, who heads Mouser's distributions in Europe, the Middle East, and Africa, credits the company's focus on new product introductions and customer recruitment as keys to its success.

EXPANDING THEIR REACH

Minnesota-based distributor Digi-Key Corp. is also looking to Europe for growth, although with a different focus. Digi-Key has long served design engineers around the world, but is building business with larger production accounts by placing new resources in Europe. Earlier this year, Digi-Key established an EMEA sales leadership team and added a customer support center in Munich to serve customers in those regions. The distributor also said it will announce more resources covering the Nordics, Baltics, and Eastern and Southern Europe by early summer.

"We're responding to customer demand for a distribution model that supports today's changing global business climate," Chris Beeson, Digi-Key's vice president, global sales and business development, said when the changes were announced in March. "By recruiting experienced sales leaders, we can more quickly build the Digi-Key brand, increase our traditional business, and support larger production business accounts with one-on-one support."

Digi-Key's Europe-based team is divided into four focus areas: EMEA and Asia Pacific; Israel and Russia; the United Kingdom and Ireland; and Central Europe. The team supports a roster of 41,000 EMEA customers and is working to spread the word about Digi-Key to other design engineers and production buyers across each region.

Expansion in Europe is a key part of Digi-Key's strategy to build out its hybrid distribution model, adds Beeson. Production accounts represented about 33% of Digi-Key's $1.4 billion in sales in 2012, he says, adding that he expects the production side to hit $615 million in sales this year. Europe's growing production business represents a small slice of that $615 million, at just about $50 million in business today, leaving plenty of room for growth. Beeson adds that he expects the production business to exceed $1 billion in the next few years.

Digi-Key's focus on developing the supply chain tools necessary to serve those customers is a key reason for that optimism, but Beeson also points to the distributor's dedication to inventory and its commitment to delivering "the right product on time all the time" as a premium value in the marketplace.

"Digi-Key is in the inventory business," Beeson explains. "Our goal is to never let the customer go down."

STRENGTHENING TIES

Texas-based distributor Allied Electronics is taking a different tack when it comes to globalization, but Europe figures prominently in its approach as well. Owned by Britain-based Electrocomponents plc, Allied is a sister company to Britain-based RS Components, which does business across Europe and around the world. The two companies are working to strengthen their ties, aligning their product portfolios globally and taking advantage of being on a single enterprise resource planning (ERP) system.

"We have two very strong brands, and our intention is to be a stronger global force," says Phil Dock, global head of product and supplier management for Electrocomponents plc.

Allied spent much of 2011-2012 upgrading to the SAP platform that RS Components was already using. Now the companies are leveraging the strength that's behind that unified system to create more supply chain visibility and enhance customer-facing services such as design tools, bill of materials management, and quotation and ordering systems.

But before they can take full advantage of those capabilities, the companies must work to better align their product portfolios, which company president Scott McLendon says consist of just a 10% overlap in materials now. Though Allied and RS have many of the same suppliers, the product mix looks different around the world.

"In the next five years, our goal is to have a pretty common [product] range around the world, somewhere around 75% common portfolio and the remaining 25% left up to local [needs]," McLendon said in an interview earlier this year. "This allows you to leverage your supply chain better, leverage your demand, make your global inventory visible to your customer base … and then you build your service proposition around that."

RS Components conducted a similar exercise to unify its product offering across Europe a few years ago, resulting in a £30 million sales increase within 24 months. Global availability can also help on the local level. As OEMs design and build their products in multiple regions, visibility into where the inventory is and how quickly they can access it in a particular location becomes even more important.

"The best global companies look local," says McLendon. "There are certain things you can leverage globally, but at the end of the day all business is transacted locally." ■

Circuit Protection
Gets Smaller, Easier To Use

New distributor offerings feature the latest in high-tech fuses, circuit breakers, and ESD protection.

IN A NOD TO the never-ending need for smaller, lighter, and faster electronic components, some of the industry's top distributors are promoting the newest in high-tech circuit protection devices this spring. These offerings highlight trends taking hold across the industry.

NOW FROM TTI

TTI *(www.ttiinc.com)* is featuring the Littlefuse Nano2 485 series fuse, which packs high voltage and energy protection in a small footprint, among its online offerings. Its fast-acting blow characteristics ensure high-speed protection against surge currents and overload conditions, and its ceramic housing ensures robust operation and long-term reliability even in applications with high ambient operating temperatures.

Also, TTI offers Molex's push-pull style micro-SIM card sockets online. Available with or without a detect switch, they offer maximum printed-circuit board (PCB) and vertical space savings with added anti-shorting and polarized card-insertion features for ultra-slim handheld devices. Measuring 12 by 15 mm, the Molex micro-SIM push-pull card is 52% smaller than its predecessor mini-SIM card (15 by 25 mm).

AVAILABLE FROM DIGI-KEY

Digi-Key *(www.digikey.com)* is touting UL-recognized device circuit breakers from Phoenix Contact. The CB pluggable circuit breaker provides convenience and flexibility for control panel builders and field customization. Using CB pluggable devices for each dc load isolates potential load faults and increases system availability for the end user. Features include:

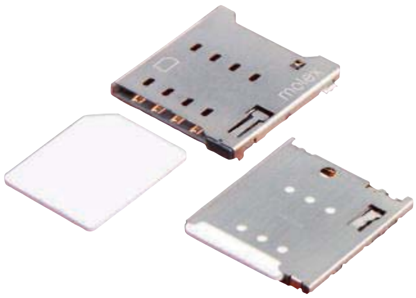

The Bourns Sparctube ST series switching spark gap suits applications where a capacitive discharge is used with a transformer to obtain pulses of up to several thousand volts.

- DIN rail mounting
- Push-in wire termination for reliable fast connections
- Snap-lock plug and base to ensure secure installation
- Codability to prevent plug mismatch
- Thermal-magnetic plugs from 0.5 A to 16 A
- 24-V dc electronic plugs from 1 A to 6 A for sensitive electronic loads
- 12.3-mm width

MOUSER'S LATEST

Offered by Mouser Electronics *(www.mouser.com)*, the Bourns Sparctube ST series switching spark gap suits applications where a capacitive discharge is utilized with a transformer to obtain pulses of very high voltage levels—up to several thousand volts *(see the figure)*. The ST series is a switching gate drive transformer (GDT) that features breakdown voltages between 350 V and 1000 V, including the most utilized versions of 800 V and 840 V. Other features include a fast switching rate, high energy, low loss capability, a wide operating temperature range, and long service life. ■

ALLIED ELECTRONICS®

Voted #1 by Automation Professionals*

Trust Allied to bring you name brand suppliers for all your automation projects

ABB · Hoffman · Phoenix Contact (Inspiring Innovations) · Honeywell · Siemens · Eaton

ebmpapst · Schneider Electric | Square D · Cherry · HellermannTyton · Orion Fans

TDK-Lambda · Panduit · Hammond Manufacturing · Harting (Pushing Performance) · Carling Technologies (Innovative Designs. Powerful Solutions.) · red lion

Belden · Omron · SMC · Crouzet · Altech Corp. · APC by Schneider Electric

crydom · SOLA/HD · Cooper Bussmann · MW Mean Well · power-one (Changing the Shape of Power) · NKK Switches

Turck works · Opto 22 · CDE Cornell Dubilier · Bud Industries, Inc. · SL Power Electronics · Yokogawa

Cosel U.S.A. Inc. · IDEC · Banner (more sensors, more solutions) · Pepperl+Fuchs · Schneider Electric

* For 'Standard Control' products. Source: *Automation World* Magazine – 2011, 2012

ORDER TODAY

Get the products you need from the brands you trust @ alliedelec.com

ecia Authorized Source

alliedelec.com 1.800.433.5700

Continued from Page 49

past couple of years, and we see this continuing in 2013. Our presence in this segment is less than a decade old, and we continue to see growth opportunities for TTI.

Typically, when the OEMs go into production, whether on a combine or a bulldozer or an automobile, their production schedules are set for a 12- to 18-month period. So, we're pretty optimistic based upon meeting with our key customers that the balance of 2013 is going to continue to be pretty positive.

DR: What are your key goals in serving transportation market customers?

VENTURELLA: We have the same goals in place for the transportation customers as we do for the balance of our customer base, and that is to be the most preferred electronics distributor worldwide, delivering the right parts to the customer on time. This is our corporate mission statement. We take that and use the same goals for our transportation customer base.

DR: Are you seeing any particular trends in this customer base?

VENTURELLA: Yes. The transportation segment has embraced outsourcing for probably 30 years now, where the OEMs in the segment have outsourced a lot of their electronics production—at the [printed-circuit board] and wire harness levels. So the customer base we service as a distributor in terms of the actual order transactions are what are referred to in this market as the "tiers."

OEMs have continued to put more and more very stringent demands on their tiers in terms of flexibility of production schedules. For example, one of our key tiers [promises its customers that it] will react to any schedule fluctuation within a two-week period. The challenge that gives to the tier is that somebody in the supply chain has to be carrying inventory to be able to react to a schedule change like that. That's a trend that we've really seen: the OEMs putting a lot more demands on their tiers in terms of delivery flexibility.

One of the strong points for our company is the ability to share forecasts with our customers and to pipeline that inventory to make sure we have it in place so when those schedule fluctuations occur we are ready and prepared to meet those [changes].

DR: What types of products are most in demand among automotive market customers today?

VENTURELLA: The trend there that I think is applicable to all electronics is the smaller, faster, lighter demands of the marketplace today. [Manufacturers] are packing a lot more electronics into the same chassis, engine, etc., so they're looking for smaller components that take up the same space, with the same extreme performance characteristics.

DR: Electronic content in vehicles continues to grow. Does this affect your approach when it comes to service?

VENTURELLA: What we've seen is that our customer base is looking for experts in the subject matter that can bring multiple solutions and multiple application recommendations through a single point of contact. They are looking for a single point of contact that can bring in new products, show them the technical specifications, deliver the samples, and give them the 3D models and other things they need to do their design.

In response, we have added a technical application engineering team as part of my group. They are strategically placed throughout North America with a dual role: they call on key OEMs, working with their design engineering staff and also serving as a technical resource for our field sales organization throughout North America and the world. So when a field sales rep calls on a customer designing a fuel system for a tractor, they have a resource they can go to that's part of the team.

Suppliers have always put significant resources into the automotive space, so most of our key suppliers have full-blown automotive divisions with selling and technical resources. In the last five years or so, with growth in the non-automotive segment our suppliers have been looking at some of their distribution partners to lead the way with similar resources.

DR: Are there any other issues or trends you're seeing in the transportation sector?

VENTURELLA: We think the electronic content in this transportation umbrella that I talk about is going to continue to grow. When you think about it, back in the day when the farmer got into his combine or his tractor to go plow the field, he sat up in the cab and looked at the rows and that's how he kept his vehicle straight. Now most of them have GPS built in. The "comfort" electronics have grown so much, and we think that's going to continue to happen. This will continue to be a big market for us, and we'll continue to invest resources in this business. ■

MORE FROM GLOBAL PURCHASING

FOR MORE ON the always changing distribution market, go to GlobalPurchasing.com and see:

- Taking the Sequester Bull By The Horns
- Reducing Risk Takes Top Priority
- How To Develop A New Or Better Supplier
- Meeting Supply Chain Challenges
- Rating Suppliers Gets Harder
- Top 50 Distributors 2013
- Taking the Sequester Bull By The Horns

UDAY MUDOI | Vitesse Semiconductor
uday@vitesse.com

Use BCs And TCs To Improve LTE And LTE-A Infrastructure Timing

Video and other timing-sensitive applications require a more rigorous solution to packet delivery.

IN ITS MOST recent Mobile Infrastructure and Subscribers Report, Infonetics forecast that the LTE market would almost double in 2013, easily passing the $10 billion mark for the first time. This explosive growth will cause growing pains as cellular carriers struggle to transform the largely circuit-switched network created for voice into an IP-based (Internet protocol), very fast packet-switched network capable of carrying video to millions of subscribers simultaneously.

Not the least of the challenges that wireless carriers face will be posed by a decision that they made years ago to use Global Positioning System (GPS) signals to provide time-of-day (ToD) synchronization for the basestations in their legacy time-division-multiplexed (TDM) backhaul networks.

This choice always involved some complications. In the macro-cell era, areas such as urban canyons that did not have clear line-of-sight (LOS) access to GPS satellites, for example, had to be treated as special cases.

In the coming 4G/LTE era, however, precise timing is the sine qua non for media such as teleconferencing and video. As a result, using GPS for ToD synchronization is already creating a much bigger and widespread challenge that will ultimately require moving to a new synchronization solution. Precise timing is also important because it directly affects radio spectral efficiency and data throughput.

Application	Frequency Transport/air interface	Phase Air interface
	ITU-T G.8261 network target/3GPP max error	
GSM/UMTS/W-CDMA	16/50 ppb	None
UMTS/W-CDMA femtocells	na/250 ppb	
	ITU-T G.8261 network target/3GPP max error	3GPP max error
cdma2000	16/50 ppb	±3 to 10 μs
TD-SCDMA	16/50 ppb	±1.5 μs
	ITU-T G.8261 network target/3GPP max error	3GPP max error
LTE (FDD)	16/50 ppb	None
LTE (TDD)	16/50 ppb	±1.5 μs
LTE-A multi-broadcast single-frequency network	16/50 ppb	±1 to 32 μs
LTE-A hetnet coordination	16/50 ppb	±5 μs
LTE-A coordinated multipoint (network MIMO)	16/50 ppb	±500 ns (within cluster)

1. Previous cellular network generations had relatively lax frequency accuracy requirements and little to no requirements for phase, time of day, or accuracy. In contrast, 4G/LTE and LTE-A networks are much more stringent, making highly precise network timing solutions imperative.

THE NEXT STEP

Most mobile basestations use GPS/GNSS to acquire an all-important signal known as 1 pps, which is used to make the ToD calculation for next-second rollover and to synthesize the fundamental source radio frequencies. Since 4G/LTE will markedly expand the use of picocells and femtocells to ensure coverage and performance, the GPS/GNSS solution becomes even more problematic.

Urban environments—and small cells located indoors—typically don't have an unobstructed LOS to the satellites. To make matters worse, the availability of inexpensive GPS jamming and spoofing technology is already a rising cause of concern among mobile infrastructure planners.

4G/LTE's use of advanced technologies for beam forming such as multiple input multiple output (MIMO) and coordinated multipoint transmission (CoMP) also rely heavily on precise phase alignment between basestation radios. In brief, timing can be everything for 4G/LTE.

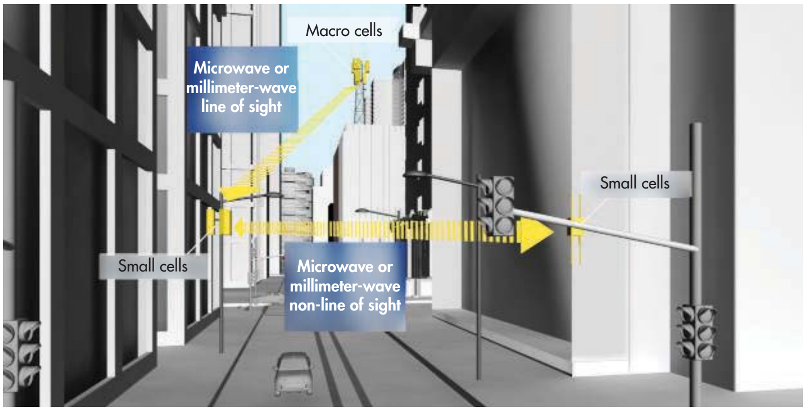

2. Since LTE and LTE-A often are deployed in urban small-cell environments where multiple high-rise buildings prevent line-of-sight (LOS) fixes on a GPS satellite, it is imperative to find alternative, low-cost synchronization solutions for both frequency and time of day. A typical urban small-cell setup uses backhaul over microwave or millimeter-wave links. To compensate for non line-of-sight (NLOS), these small cells on lampposts or traffic lights must have an alternative means of deriving network timing.

Figure 1 shows the relative criticality of timing for phase and frequency in various cellular network generations. The column comparing phase alignment offers the most striking data: from 0 for GSM to 500 ns for LTE-Advanced (LTE-A) coordinated multipoint (network MIMO). The ToD timing requirements are also critically important.

Comparisons with previous cellular generations are just as striking. GSM, for example, has no ToD timing accuracy requirement. UTMS TDD requires ≤2.5-µs accuracy. TD-SCDMA requires ≤3-µs accuracy. On the other hand, LTE-A/MIMO requires ≤0.5 µs, and LTE-A with interband aggregation requires ~0.1-µs accuracy.

The good news is that the transition from SONET/SDH or PDH E1/T1 to IP/Ethernet backhaul networks is already well underway and will continue for several years. The new backhaul network will be a more cost-effective broadband packet network that delivers vastly more bandwidth to basestations at significantly lower cost per bit.

TOWARD A TIMELY TIMING SOLUTION

Today's mobile carriers already deploy Precision Time Protocol (PTP) as described in IEEE 1588 in their backhaul networks. PTP uses time-stamped packets to implement frequency and time/phase distribution and a protocol exchange to derive the actual time from the time-stamped packets.

By necessity, carriers roll out infrastructure upgrades incrementally, which means that some basestation nodes are or will be implementing PTP while others are not. Sending PTP traffic though PTP-unaware network elements creates a packet delay variation (PDV) that's cumulative across the network. An intelligent low-pass filtering algorithm in the PTP ordinary clock slave can slowly filter out these delays.

Unfortunately, filtering allows the slave to recover the frequency but not the equally important ToD and phase alignment parameters. Moreover, the long time constants of the filtering algorithm result in very long frequency acquisition/settling times and also the likelihood of long-term phase wander. Most of the "secret sauce" in intelligent PTP frequency algorithms involves keeping this long-term wander in check.

A good deal of technology backfilling is being applied in attempts to correct the problems of very long frequency acquisition/settling times and the network's susceptibility to long-term phase disorientation. Standards-making bodies are also working toward a more comprehensive solution. But the upshot is that attempts to tweak the basic PTP protocol can provide a viable frequency-only solution but not a synchronization solution for phase and ToD.

As a result, the only way to attain accurate ToD and phase delivery over the network without GPS will be to provide "PTP on-path support." All network elements between the PTP master clock and the PTP ordinary clock slaves must contribute to eliminating or significantly reducing the impairments introduced by static asymmetries in network links as well as the PDV caused by variations in packet traffic through these elements.

PROBLEM ANALYSIS

The first difficulty arises because the timing information is encapsulated in standard IP data packets and passed between network nodes. The same PDVs producing jitter in packet networks also introduce timing inaccuracies in the recovered clock.

Clock recovery is problematic as well. The PTP protocol's timing recovery mechanism works on the assumption that there is fixed and symmetric latency between network nodes—a condition that's rarely true in real-world networks. As previously mentioned, synchronization error caused by PDV and asymmetry is cumulative along the path between the network node generating the master clock and the basestation recovering the clock.

The problem is even more complicated when microwave links are used for backhaul. Unlike wired connections, latency microwave links change dynamically as a function of the modulation format being used or as atmospheric properties vary with changing weather.

A second set of issues stems from the incremental approach carriers must employ for infrastructure upgrades. Without careful planning, delivering the timing precision required to support advanced wireless technologies during these upgrades can quickly become difficult and expensive.

Phase and ToD synchronization are required to eliminate the timing errors created by large asymmetric queuing delays common in switches and routers as well as asymmetric modem delay variations for microwave and wired links. But as discussed earlier, the utilization of GPS for timing synchronization is problematic at best.

IEEE 1588V2 TO THE RESCUE

To address this challenge, later versions of the IEEE 1588 standard defined two additional clock types: the boundary clock (BC) and the transparent clock (TC). Both types provide on-path support to network elements in a packet-based network. A revised standard, IEEE 1588v2, incorporates the BC and TC concepts.

A node that implements a BC regenerates the timing based on the timestamps it receives. In addition to timestamps, a BC requires timestamp correction, a reliable PDV filtering algorithm, and an IEEE 1588v2-aware timing complex that can be synchronized to the network.

This means that implementing a BC node requires a very precise oscillator, a digital phase-locked loop (PLL), and a microprocessor. BCs are expensive and complex to implement. However, they are able to perform the port-level accurate timestamping that's required to meet LTE and LTE-Advanced synchronization.

TC nodes are simpler and less expensive to implement because they simply forward incoming timestamps after correcting errors introduced in the node. A TC can also reuse the physical layer (PHY) or switch that it encounters in the datapath, provided the device is 1588v2-compliant (i.e., it has the added logic required to support accurate timestamping and timestamp correction mechanisms).

A port-based PHY solution is sufficient to implement a highly accurate TC node. Several PHY and switch silicon solutions available today incorporate timestamping correction that compensates for network PDV and asymmetry. At least one of these solutions, Vitesse's VeriTime chipset, includes compensation mechanisms that can reliably maintain a 1588v2 timing domain across microwave links. Clearly, any cost-effective solution to the ToD dilemma maximizes the utilization of TCs and minimizes the number of BCs in the network.

NETWORK DEPLOYMENTS

Care must be taken when creating a 4G/LTE backhaul architecture that can be deployed incrementally and can still maintain timing integrity during an incremental upgrade/build-out. Since 4G/LTE will markedly expand the use of picocells and femtocells to ensure coverage and performance, TCs must play an important role. They are well suited for picocell synchronization in outdoor environments and synchronization down to the femtocell in large indoor multi-floor installations *(Fig. 2)*.

For picocells, TCs can be carried over microwave and millimeter-wave links to satisfy TD-LTE and LTE-A requirements while eliminating GPS/GNSS signals or fiber links. The simpler hardware of TCs fits the model of the small-cell basestation requiring small footprints, low cost, and minimal power consumption.

Low power consumption is particularly important for wireless operators with LTE licenses in higher frequency bands (1900 MHz and the 2.x-GHz bands) with limited indoor coverage. Indoors, the access network itself can generate 1588v2 timing, or in some instances a GPS antenna on the roof of the building can generate time packets for synchronization services inside the building. Furthermore:

- BCs should be deployed throughout the network strategically to reduce the traffic load on the grandmaster clock.
- Timestamping should be performed at the PHY ingress/egress point to eliminate serialization and deserialization delay asymmetries.
- Microwave backhaul equipment should position TCs across each link to minimize time errors from the modem and adaptive modulation schemes that can easily exceed the overall time error budget of the network.
- Accurate timestamping is best done at the PHY I/O interface. With TCs implemented in the backhaul network, timestamping at the PHY/IO interface can lead to simpler servo algorithms that translate to cost-effective PTP clock slave implementations.

By following these guidelines, a CAPEX-constrained operator can move a few key network elements to BC support, paving the way for use of efficient, less costly network elements that use TC clocks during high-volume deployments. This strategic approach to backhaul timing architectures allows carriers to quickly roll out networks of smaller cells that improve their coverage and capacity in a cost-effective manner.

CONCLUSION

The most cost-effective upgrade strategy for LTE-A and future small-cell networks is to deploy distributed TCs liberally and BCs only where necessary to segment timing domains. TCs can increase timing accuracy to the nanosecond range, as shown by Vitesse in a recent submission to the ITU-T standards committee. Silicon advances available today will ensure such solutions will carry only a nominal premium over non-IEEE 1588v2 aware systems designed for 3G networks today. ed

UDAY MUDOI, director of product marketing at Vitesse Semiconductor, holds a BS in electrical engineering from the Indian Institute of Technology, Kharagpur, a master's degree in computer science from North Carolina State University, Raleigh, and an MBA from Columbia University, New York, N.Y.

SACHIN GUPTA | Cypress Semiconductor

Dynamic Microcontroller Reconfiguration Delivers More Than 100% Resource Utilization

Using run-time reconfigurable analog and digital logic can enable greater than 100% utilization of available hardware resources, resulting in a smaller footprint and lower system cost.

WHEN FACING THE challenge of having to design systems that do more, are smaller, and have higher power efficiency, the ability to achieve more than 100% with available resources through dynamic reconfiguration becomes compelling.

Consider a vending machine that needs to accept money and dispense items as its primary function with an occasional need to communicate to a host CPU. Using dynamic reconfiguration, the same set of digital resources that are configured as the timer/counter at the time of vending can be configured as a UART block and pulse-width modulators (PWMs) for baud rate generation while talking to the host.

Dynamic reconfiguration can be realized in nearly every application that performs different tasks at different times. Since a walkie-talkie transmits data only in one direction, for example, dynamic reconfiguration enables the design of a smaller device.

A handheld ticket-dispensing machine could use the logic that operates the thermal printer to serve as the battery-charging circuit. Similarly, an LED-based lantern charges the battery when power is available and uses those resources to control the white LEDs when wall power is not available.

Effectively, battery charging is time-multiplexed with other system functionalities. The result is that a system can fit within a smaller footprint than if separate logic is required for each function.

DYNAMIC RECONFIGURATION

Dynamic reconfiguration is a well-known capability of digital architectures like FPGAs. It is achieved by having a new configuration stored in non-volatile memory and later read written to the appropriate configuration registers. Each of the various configurations is stored in non-volatile memory/flash at the time of device programming.

Programmability is no longer limited solely to the digital domain. Today, analog resources can be dynamically reconfigured and not just to change the specifications of a particular peripheral. Using dynamic reconfiguration, the same analog resource can serve as an analog-to-digital converter (ADC),

1. The analog resources can be dynamically reconfigured into a continuous time programmable analog block.

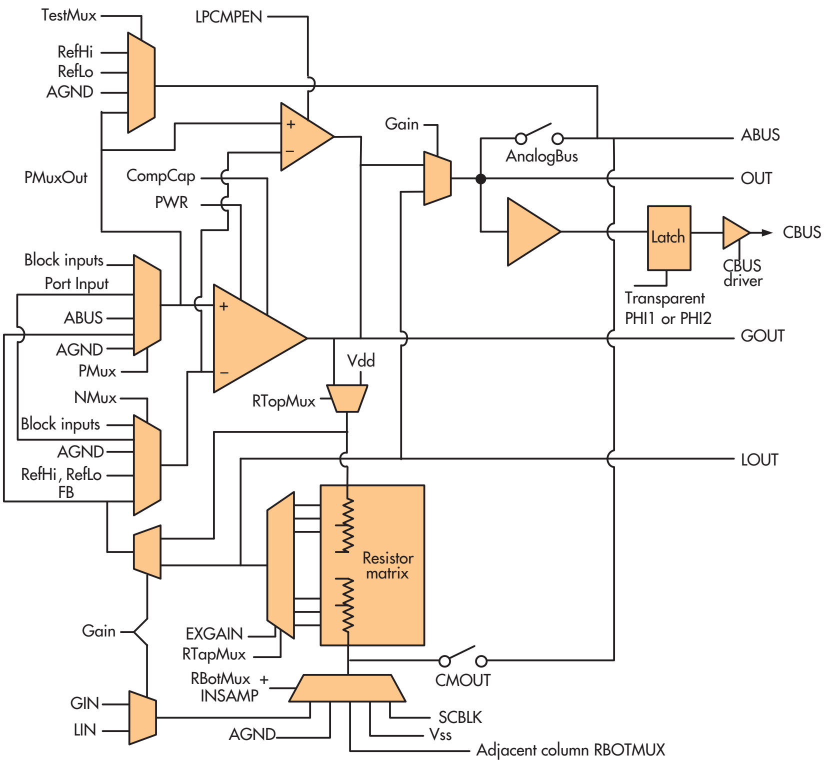

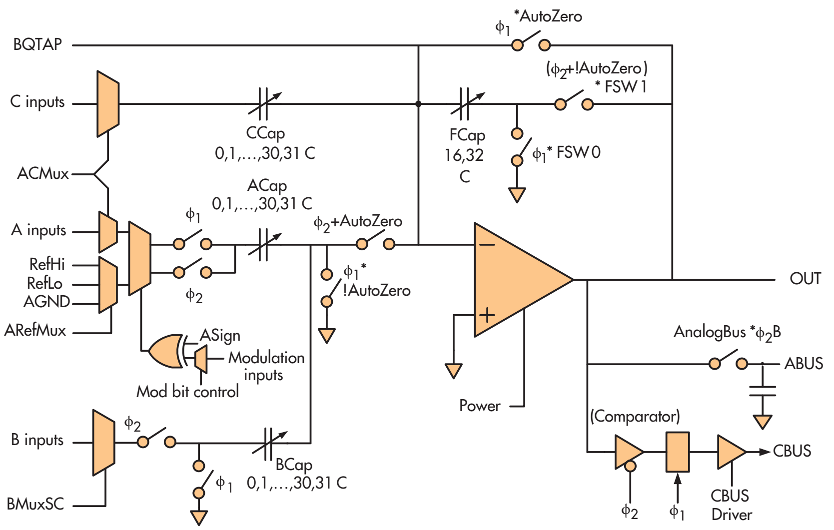

2. Switched capacitor analog blocks are made up of a set of switched capacitors surrounding an op amp.

amplifier, or capacitive touch sensor at runtime based on the application requirements.

To understand how analog resources can be dynamically reconfigured, consider a continuous time programmable analog block *(Fig. 1)*. At first look, this circuit appears quite complicated. In fact, it is a relatively simple circuit that provides the ability to connect various signals at any given input/output to realize different circuits.

For example, this same block can be configured as an inverting or non-inverting amplifier based on the connections. A hysteresis-enabled comparator with the required gain/hysteresis can be achieved by selecting the appropriate resistor value within the resistor matrix. The required reference can be selected using a reference multiplexer. The output of the block can also be routed to some other block or an output pin.

All of these connections and resistor values are configured using configuration registers that can be written at run-time. Thus, a block can act like an attenuator, buffer, inverting or non-inverting amplifier, and even an instrumentation amplifier if combined with additional blocks. System designers can change the functionality of the block using firmware by writing new values to the configuration registers.

A switched capacitor block can emulate a capacitor as a resistor using a switched capacitor technique where the resistance (R) is implemented by controlling the flow of charge:

$$R = 1/f_S C$$

By controlling the switching frequency (f_S), the capacitor values (C), and the way these capacitors are connected around the op amp, a variety of circuits can be implemented without requiring any additional hardware *(Fig. 2)*. The same switched capacitor can also be used to implement a delta-sigma modulator, integrator, filter, or touch sensing block.

Each capacitor has an associated configuration register that can be used to change the capacitor's value. The switching frequency can be controlled using hardware dividers or a PWM that is implemented in a programmable digital block.

Switches are available that can be controlled using registers to keep some capacitors in the circuit all the time or to remove them instead of keeping them in switched capacitor mode. For instance, an integrator needs a capacitor in the feedback path and does not require any switching. The analog blocks can implement a range of devices including an inverting amplifier, an integrator, or a differentiator.

Similar to the analog resources, digital blocks also can support various functions by writing the required configuration value to the configuration registers for the given digital block. Programmable input/output connections enable any pin to be connected to any peripheral with the processor so the same hardware block can be connected to different pins at different times. For example, a digital block being used as a PWM to drive an LED can be used as a UART transmitter at a different time using different output pins.

FEASIBILITY OF DYNAMIC RECONFIGURATION

Depending upon the architecture of the programmable device, dynamic reconfiguration may be possible for many applications. However, the complexity of implementing dynamic reconfiguration may too costly if every register needs to be understood by the developer and its value set manually.

To make dynamic reconfiguration practical, the development tools must provide an efficient way to set these values with minimal effort. Additionally, they must provide a basic reconfiguration infrastructure for run-time implementation. Two features are essential:

- Automatic generation of configuration register values: Development tools must provide an easy way to define and manage multiple configurations. Defining a configuration refers to the selection of the peripherals to use and their internal interconnection. Based on the configuration definition, the tool must generate the required register values for each block used in that configuration *(Fig. 3)*. As can be seen from this code, when the development tools can generate the required register values, this frees the development from manually generating these values. This also enables developers to make frequent changes to configurations without any penalty.

- Ease of switching configurations: If the developer has to manage switching between configurations, this will impose significant overhead as the developer creates code to read values from flash and write to the required registers, especially when frequent changes are made to configurations at the time of development. A single mistake can lead to a system malfunction that could be difficult to detect and locate. Therefore, development tools must abstract the low-level register write details and must provide a high-level application programming interface (API) to perform this task.

```
;  Instance name PGA, User Module PGA
;       Instance name PGA, Block Name GAIN(ACB00)
    mov reg[71h], fdh          ;PGA_GAIN_CR0(ACB00CR0)
    mov reg[72h], 21h          ;PGA_GAIN_CR1(ACB00CR1)
    mov reg[73h], 20h          ;PGA_GAIN_CR2(ACB00CR2)
    mov reg[70h], 00h          ;PGA_GAIN_CR3(ACB00CR3)
;  Instance name PRS8, User Module PRS8
;       Instance name PRS8, Block Name PRS8(DCB02)
    mov reg[2bh], 00h          ;PRS8_CONTROL_REG(DCB02CR0)
    mov reg[29h], 00h          ;PRS8_POLY_REG    (DCB02DR1)
    mov reg[2ah], 00h          ;PRS8_SEED_REG    (DCB02DR2)
;  Instance name PWM, User Module PWM8
;       Instance name PWM, Block Name PWM8(DBB01)
    mov reg[27h], 00h          ;PWM_CONTROL_REG(DBB01CR0)
    mov reg[25h], 19h          ;PWM_PERIOD_REG(DBB01DR1)
```

3. The PSoC Designer IDE generates code for configuration registers based on a user's design.

Further complicating switching configurations is the need to "unload" the values that have been written since the previous configuration was loaded. Unloading may not be required in every situation but will be required in most of them. For example, consider an application with a base configuration and two loadable configurations.

The base configuration includes the resources that are required all the time and do not need to be loaded and unloaded at any point. This configuration is loaded before entering the main application code. In this example application, the base configuration needs to have a UART that talks to the host CPU. The first loadable configuration needs one ADC. The other configuration needs one PWM and a touch sensor.

Assume the UART needs two programmable digital resources and the PWM needs one while the ADC and touch sensor use the same number of resources. The processor handles switching between configurations *(Fig. 4)*. Of course, the development tools must provide a way to load and unload these configurations easily without the need to understand low-level registers and where configuration values are stored in flash memory.

IMPLEMENTING DYNAMIC RECONFIGURATION

To implement dynamic reconfiguration in an application, the various hardware configurations allocating the digital and analog blocks need to be defined. This can be completed using an intuitive graphical user interface (GUI) to simplify the process.

Developers, then, can dynamically switch between these configurations by calling an API within the application. The API calls code that will then set the required register values based on the peripherals and the settings that have been selected for the particular configuration.

To show how dynamic reconfiguration is implemented in applications, let's look at a real-world system. A programmable system-on-chip (PSoC) from Cypress consists of various system resources, CPU core, digital subsystem, analog subsystem, and programmable interconnects.

The digital subsystem comprises various programmable digital blocks that can each be configured independently for different functions. Similarly, the analog subsystem consists of various programmable analog blocks that can be similarly configured. These subsystems are shown graphically as independent blocks in the development tool *(Fig. 5)*.

To configure a programmable hardware block, developers drag a peripheral over and drop it on an empty block. It can then be connected to other peripherals or a pin. The development tool generates the required register values that

are to be stored in flash at programming time. At power up or reset, before the device begins to execute the main application code, the device loads the initial configuration into the registers from flash.

Developers can define a new configuration by right clicking and then selecting a new loadable configuration in the workspace. For each configuration, developers can select different peripherals and connect them as required.

Once all of the configurations are defined, the tool generates the required register values along with assembly code to move these values to their corresponding register as shown in Figure 4. Developers can also load configurations dynamically from the main application using a single line of code: "LoadConfig_[config_name]".

One factor that may require attention from developers is the switching time

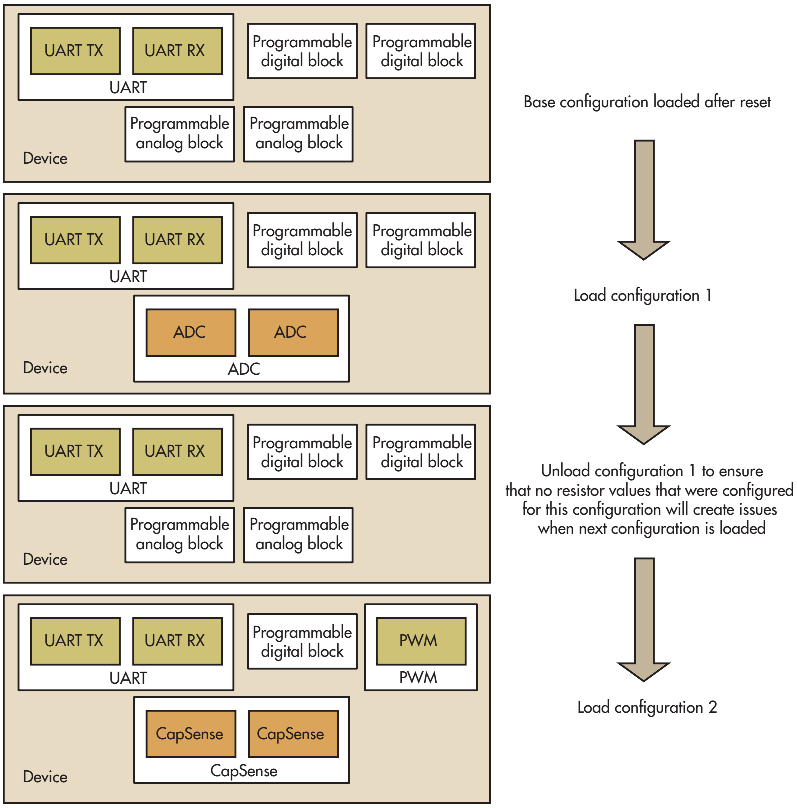

4. It is possible to dynamically switch configuration while the system is running.

LINEAR TECHNOLOGY
DESIGN NOTES

High Input IP3 Mixer Enables Robust VHF Receivers

Design Note 515

Andy Mo

Introduction

An increasing number of applications occupy the 30MHz to 300MHz very high frequency (FHF) band. Television and radio broadcasting, navigation controls and amateur radios are a few examples. Modern RF component development is aimed at much higher frequency bands used for voice and data communications systems. Significant advance in circuit techniques and manufacturing processes are required to meet the demanding performance requirements of the next generation of radios. Applying these techniques to lower frequency designs can significantly improve performance.

The LTC®5567 is a wideband mixer designed and optimized for performance in the 300MHz to 4GHz frequency band. To create very compact circuit implementations, the LTC5567 contains integrated RF and LO transformers. The Input IP3 linearity performance benchmark is an excellent 30dBm for the LTC5567 in its specified frequency range. Going lower in frequency requires the built-in transformers to maintain this linearity as well as conversion gain. With such a high level of linearity to start from, it is worthwhile to modify the mixer circuit design and characterize the performance over lower VHF frequencies. The proof of performance is in the testing.

Impedance Match Design

Figure 1 shows an impedance match design with the LTC5567. Table 1 shows the design values extending input port match below 300MHz, down to 150MHz,

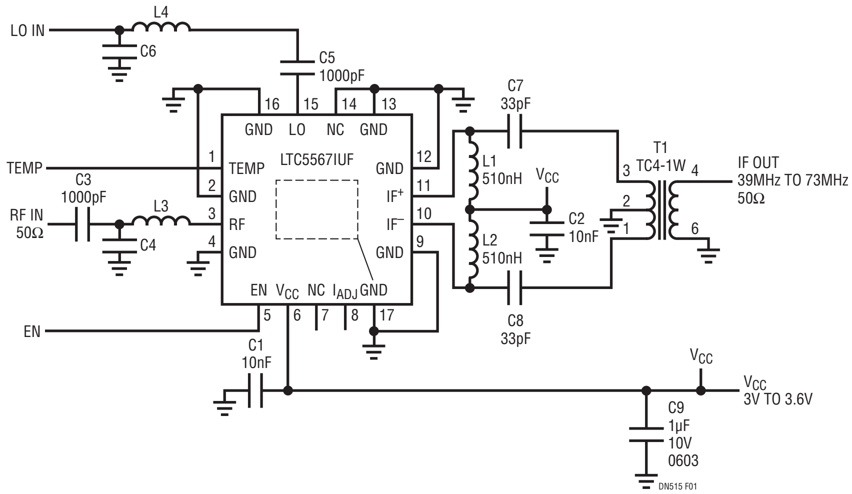

Figure 1. VHF Mixer Design

Table 1. VHF Impedance Match Design Values

MATCH	RF INPUT	L3	C4	LO INPUT	L4	C6
A	150MHz	8.2nH	56pF	200MHz	3.9nH	47pF
B	200MHz	6.8nH	39pF	250MHz	2.7nH	33pF
C	250MHz	3.9nH	27pF	300MHz	1.5nH	27pF

while still achieving outstanding performance. Test results are also provided.

Figure 2 shows the LTC5567 mixer gain and input IP3 versus input frequency. The mixer linearity performance improves as input frequency approaches 150MHz. Input, LO and output port return loss measurements are shown in Figures 3, 4 and 5, respectively. The overall performance is maintained in the VHF range compared to higher input frequencies. As a result, the high IP3 and conversion gain yields maximum dynamic range when used in radio designs. Higher dynamic range minimizes adjacent channel interference, improving selectivity. Operating the LTC5567 below 150MHz input is possible with reduced conversion gain, but not recommended, due to the internal transformer becoming lossy.

Conclusion

The LTC5567 offers very high linearity performance at VHF and UHF input frequencies. High IP3 figures and P1dB in (Table 2) make it an excellent choice for high performance radio design over a wide range of frequencies.

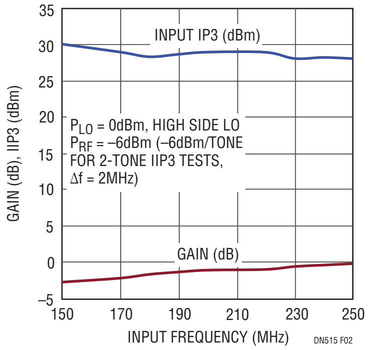

Figure 2. Mixer IIP3 and Gain Performance Results

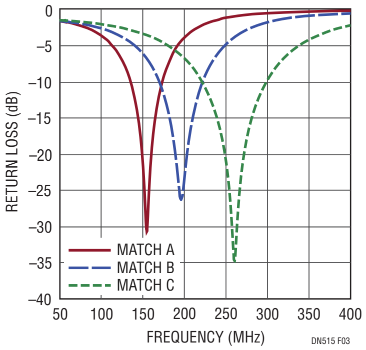

Figure 3. RF Input Return Loss

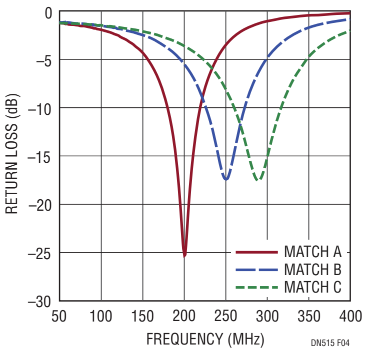

Figure 4. LO Input Return Loss

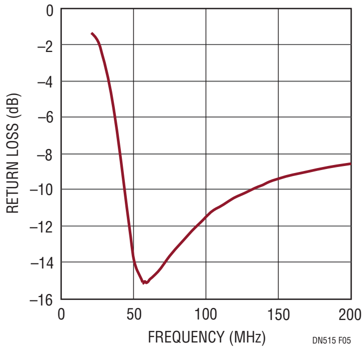

Figure 5. IF Output Return Loss

Table 2. P1dB Compression Point and LO Leakage Over Input Frequency. Output Frequency = 50MHz, HSLO

RF INPUT FREQUENCY (MHz)	P1dB (dBm)	LO TO IF LEAKAGE (dBm)
150	12.29	–35
160	12.9	–42
170	12.9	–42
180	12.75	–42
190	12.70	–41.2
200	11.61	–43
210	12.48	–43
220	12.7	–44
230	11.7	–44
240	11.08	–44
250	12.89	–44

Data Sheet Download

www.linear.com/LTC5567

For applications help,
call (408) 432-1900

dn515f LT/AP 0613 196K • PRINTED IN THE USA

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

International Microwave Symposium

IEEE 2–7 June 2013, Seattle, WA MTT–S

THERE ARE SO MANY WAYS TO EXPERIENCE IMS!

ATTEND A WORKSHOP OR SHORT COURSE!

IMS Workshops and Short Courses offer practical, application-oriented material to advance your career! Whether you're a seasoned professional, recent graduate, or student, IMS will provide a variety of topics that appeal to academia and industry alike. Hear the latest developments in R&D for emerging areas or hone your skills in a specific microwave subject.

MAKE A NEW CONNECTION!

IMS brings together the largest concentration of top engineers and scientists in the RF & Microwave Field. IMS2013 will also be home to over **50 first time exhibitors!** From the exhibition floor, to a variety of social activities, IMS is THE place to network with colleagues and companies both established and new.

VARY YOUR VIEWPOINT!

RF & Microwave technology is always on the cutting edge and that comes with lots of varying opinions. Panel Sessions are a great way to engage with colleagues on the forefront of these hot topics. The open discussion format is perfect for sharing a variety of viewpoints and getting the inside track on future directions!

WHAT WILL YOUR EXPERIENCE BE?

REGISTER TODAY AND **SAVE UP TO 25%!**

SHARE YOUR IMS STORY

COMPLETE CONFERENCE DETAILS ARE AVAILABLE AT:

HTTP://IMS2013.MTT.ORG

IEEE MTT-S

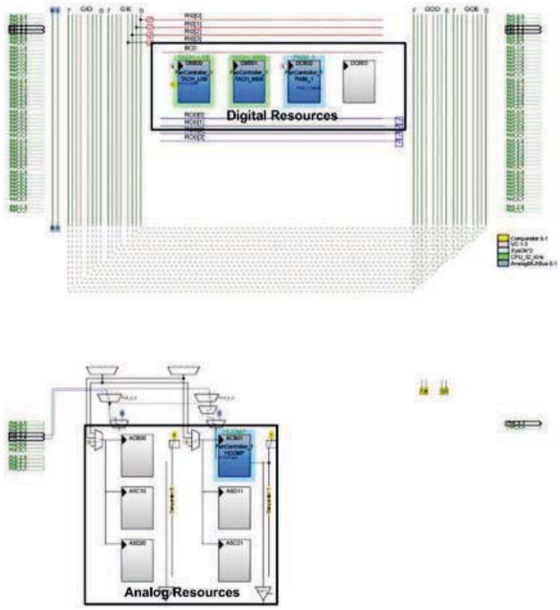

5. The PSoC Designer Chip Editor view shows the connections to the digital and analog resources.

between configurations. Although unloading and loading a configuration is primarily register writes, this process will take time directly proportional to the number of registers that need to be written for a particular configuration. The approximate switching time can be calculated based on the resources used, and the application can adjust for the latency as required.

Dynamic reconfiguration is a powerful feature that can be used in nearly every application where different peripherals can be used in a time-multiplexed manner to reduce system cost and footprint. By changing the functionality of programmable hardware resources at runtime to perform different tasks, system designers can achieve more that just 100% utilization of available hardware resources. ed

SACHIN GUPTA is a senior applications engineer with Cypress Semiconductor. He holds a diploma in electronics and communications from Vaish Technical Institute and a bachelor's degree in electronics and communications from Guru Gobind Singh Indarprastha University, Delhi. He has been working on different SoC products and specializes in capacitive sensing and embedded applications. He is also an expert in board layout reviews for good analog performance.

The NEW

engineering tv

The place where engineering videos gather

Engineering Channels

Recent Videos

Trade Show Coverage

THE SAME GREAT VIDEO CONTENT, IN AN EASY-TO-BROWSE FORMAT

In the past, we've been told by our design engineers that **Engineering TV** is their favorite place to "geek out". Well get your computers ready, because **Engineering TV** is all new and completely redesigned!

Browsing the most compelling videos is easier than ever. Our new engineering channels let you explore the technologies that are of interest to you – everything from Aerospace to Embedded to Robotics. Plus, we film at all the key electronics and design engineering trade shows, so even if you can't attend, you can find out everything that is happening at these important events.

CHECK US OUT TODAY:

www.engineeringtv.com

Single-Cell Regulated Charge Pump Draws Low Quiescent Current

BRUCE D. MOORE | ALERT SOLUTIONS CO., DAVIS, CALIF. bmoorealert@att.net

CAPACITOR-BASED CHARGE PUMPS (OR Q-pumps) generally aren't useful for sourcing large amounts of current, but they work well in niche micropower applications where space is at a premium. They work best in applications where the output voltages are integer multiples of the input voltage. The integer multiples, then, are operating points that result in peak efficiency.

However, Q-pumps can also work well when they are powered from a variable input such as a battery, particularly when quiescent battery drain is more important than heavy-load efficiency. This might be the case when powering a microcontroller that spends most of its time in sleep mode.

Low-voltage microcontrollers such as those in the PIC24 or MSP430 families are generally powered from a regulated supply voltage such as 2.5 V. If clocked slowly, they might draw as little as 25 µA or 50 µA. In standby mode with only the real-time clock running, the current can be vanishingly small, often less than a microamp. This is a good application for the regulated two-stage Q-pump described here, which boosts a single alkaline or nickel-metal-hydride (NiMH) cell to 2.5 V.

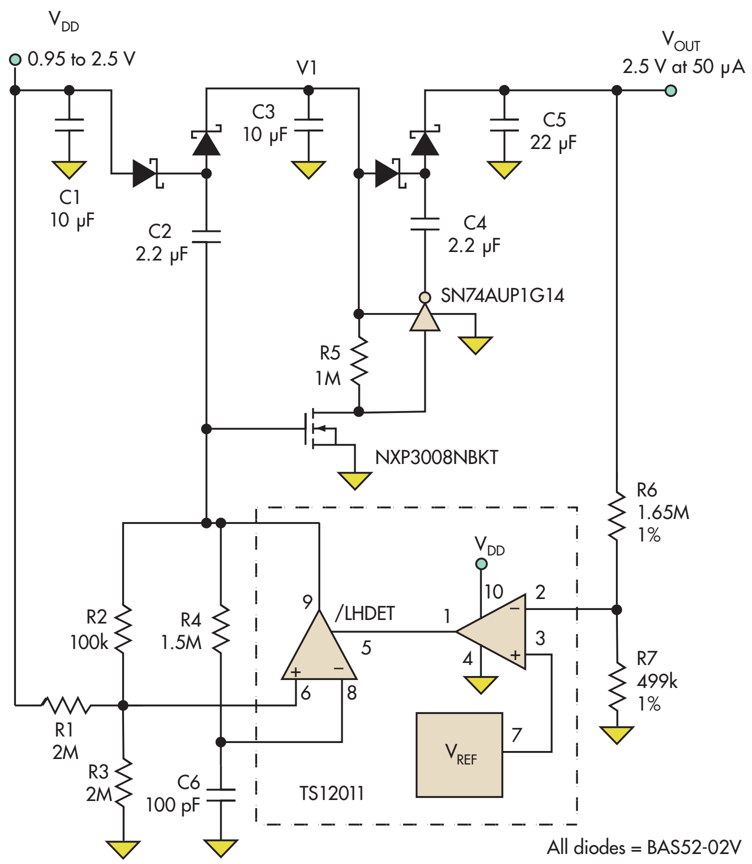

1. The low-voltage regulated charge pump uses the "flying capacitor" topology (C2 and C4) to "stack" and thus increase the output voltage.

The "wings" of a Q-pump, called "flying capacitors," connect first to the input and then to the output. If the capacitor is stacked on the input voltage, it forms a voltage doubler. In the case of a regulated charge pump with a fixed output voltage, the voltage across the flying capacitor may differ significantly from the voltage across the output-filter capacitor.

When you connect two capacitors that are initially charged to different voltages, you get a spark or power dissipation in the switches as the capacitors equalize in voltage. This is why a simple voltage doubler typically exhibits better efficiency than a regulated Q-pump.

This regulated Q-pump has an on-demand oscillator, a feedback regulation loop made from an op amp and reference, and a two-stage pump circuit, plus two flying capacitors, C2 and C4 *(Fig. 1)*. The first pump stage is driven directly by the Touchstone Semiconductor TS12011 comparator, which forms the oscillator, while the second stage is driven by an inverter powered from the output voltage of the first stage. The full-load efficiency varies from 70% to 40% over a 1- to 2.5-V input range, which is comparable to a linear regulator.

The TS12011 analog building block requires very low supply currents (3.2 µA typical) and operates well down at the sub-1-V levels needed for single-cell operation. The comparator output stage has good drive-current capability down below 0.8 V V_{DD}, which is an uncommon feature that allows us to drive the first stage directly from the oscillator.

The SN74AUP family of logic gates similarly uses super-low power. The MOSFETs were also carefully chosen for low-voltage operation, with low gate-threshold voltage specifications and low gate-charge characteristics for low switching losses. Using these high-performance components, the no-load quiescent current is a mere 8 µA.

The RC oscillator is stopped when the output is in regulation via its /LHDET latch input *(Fig. 2)*. When the output is above the regulation threshold set by the reference and feedback divider, the op-amp output drives the latch input low, latching the comparator output in the high state and stopping the oscillator.

A large hysteresis band was chosen for the oscillator, resulting in a large signal swing on the timing capacitor C6. This achieves the most efficient size-versus-current operating point for the oscillator. The maximum frequency is nominally set at 1 kHz with the component values shown, but can be adjusted up to about 3 kHz, where it becomes limited by the propagation delay through the comparator.

The amount of charge transferred with each cycle and the switching frequency determine the output current. Accurate calculations for the output impedance or output current from a regulated two-stage pump are complex and need a large spreadsheet, but you can make some oversimplifications to get into the right range. Assuming the capacitor is completely charged and discharged with each cycle (which obviously isn't true), then:

$$Q = C \times V$$
$$I = Q \times f$$

so:

$$I = C \times V \times F$$

where C is the flying capacitor value, V is the applied voltage during the charging phase, F is the oscillator frequency, and I is the output current you'd get if you could remove all the stored charge during the discharge phase (roughly equivalent to a short-circuit load current).

The modified Dickson Multiplier two-stage pump topology used here is a quadrupler, and both stages must be designed to transfer the needed amount of charge at the worst-case low-battery voltage. The first stage is the most important, because the voltage applied to the capacitor is low. Any drop across the first rectifier subtracts from the applied voltage with a resulting loss in headroom, a problem when trying to multiply up to 2.5-V output levels.

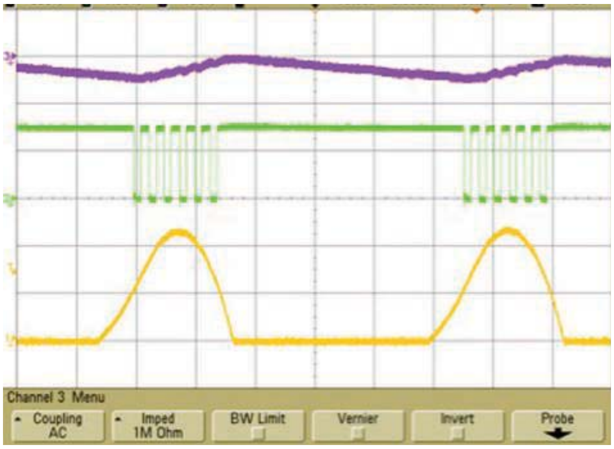

2. The oscillator controls how much charge is pumped and thus the output voltage: V_{OUT} (top waveform), comparator output (middle), and op-amp output (bottom).

The rectifier choice is further complicated in this low-quiescent application by leakage currents, which effectively load down the pump and increase idling current. The BAS52-02V Schottky diode has a good combination of low reverse leakage, low forward drop, small packaging, and wide availability.

Reverse leakage current at high temperature is the weakness of Schottky diodes, but lab measurements show that the typical BAS52 is very good on this parameter, at less than 1 µA at 50°C. For even lower reverse leakage, there is the BAS40-02V, but the tradeoff is about 75-mV higher forward-voltage drop.

The capacitor values and switching frequency are chosen to be somewhat excessive for charge transfer, so they don't get in the way of the voltage-drop/headroom issues. As a result, the frequency is low and the capacitors are relatively large, which also keeps switching losses in check. More importantly, a long on-time allows the voltage across the diode to reach a minimum, with the tail end of the forward voltage/current curve being a little less than 0.2 V for the BAS52.

From the simplified charge-transfer equations above, an overly optimistic estimate for first-stage current is:

$$(V_{DD} - V_{FWD} - V_{SAT1}) \times C3 \times F =$$
$$(0.95\ V - 0.2\ V - 0.05\ V) \times 2.2\ \mu F \times 1\ kHz = 1.5\ mA$$

which is more than sufficient to meet a 50-µA load-current requirement. The best-possible peak-output voltage that you can achieve for the intermediate V_1 tap is:

$$2V_{DD} - 2V_{FWD} - V_{SAT1(high)} - V_{SAT1(low)}$$

Similarly, the second-stage peak voltage for V_{OUT} is:

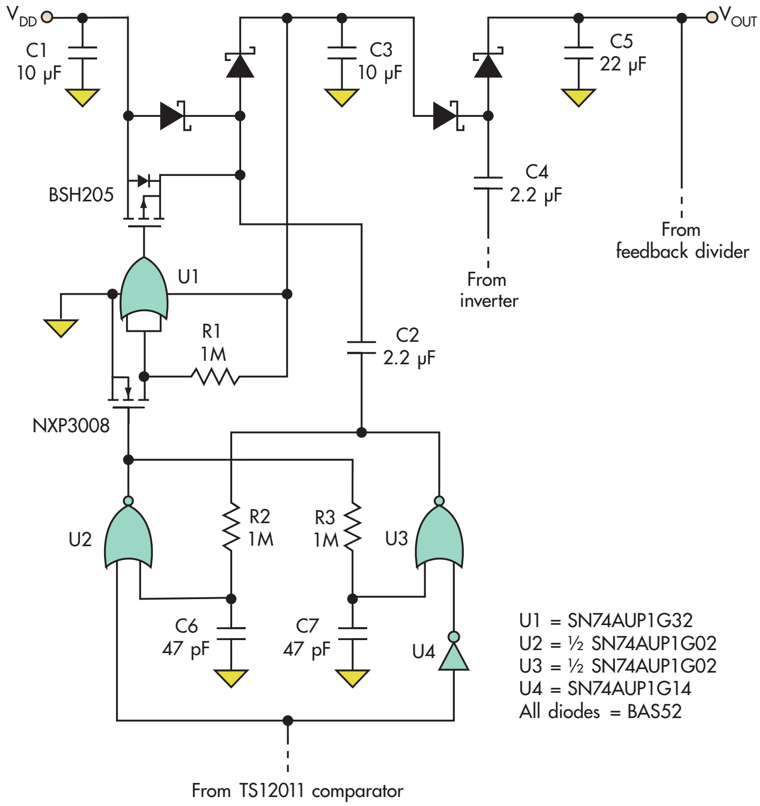

3. To achieve operation below 0.9 V, the first stage is modified by adding a synchronous rectifier across the first diode.

$$2V_1 - 2V_{FWD} - V_{SAT2(high)} - V_{SAT2(low)}$$

Thus, the best-possible peak-load voltage at the minimum input voltage is:

$$V_{OUT} = 4V_{DD} - 4V_{FWD} - 4V_{SAT} =$$
$$(4 \times 0.95\text{ V}) - (4 \times 0.2\text{ V}) - (2 \times 0.05\text{ V}) - (2 \times 0.1\text{ V}) = 2.7\text{ V}$$

where V_{SAT1} and V_{SAT2} are the voltage drops across the comparator and the AUP inverter output stages in the ON state, respectively. Again, this analysis is oversimplified but points out the limits of operation.

This math shows that there is not a lot of room for temperature variation and component tolerances when trying to quadruple the lower end of the voltage range of a single cell. We can improve this, and get to sub-0.9 V operation, by several means: adding another pump stage; changing the output voltage to 2.2 V (but the microcontroller may not be able operate this low); or adding a synchronous rectifier across the first diode.

The third option is chosen here, due to its good efficiency at the operating "sweet spots" *(Fig. 3)*. A low-threshold P-channel MOSFET with low gate charge shunts the first pump diode, reducing the forward drop from 0.2 V to approximately 0.01 V. This adds a few hundred extra millivolts right at the front end, where it's needed most.

A NOR-based logic circuit provides non-overlapping gate-drive signals to the synchronous switch. The goal is to be certain that the MOSFET gate is never driven high when the flying capacitor is high. The amount of dead time is set by the 1 MΩ × 47 pF RC time constants. Here, the dead time is very long since the 1-kHz oscillator is so slow.

Charge pumps offer an alternative to inductor-based boost regulators and can fit nicely with the dual nature of many microcontroller loads, with their low-current RTC mode versus heavier but infrequent run mode. The challenge of designing them for low input voltages can be met by using the right low-voltage components, choosing sensible switching frequencies, and taking full advantage of the V-I characteristics of the rectifiers.

BRUCE D. MOORE is a consulting analog engineer for Alert Solutions Co., doing new product definition and application circuit development for amplifier and power-supply ICs. He graduated with a BSEET from Heald Engineering College, San Francisco.

Hex Buffer, MOSFETs Build A High-Power, Lossless, Virtual Ground

LOUIS VLEMINCQ | BELGACOM, EVERE, BELGIUM louis.vlemincq@belgacom.be

A VIRTUAL GROUND IS useful whenever you need to create a bipolar supply but the dc source is unipolar, as is often the case with battery-operated equipment. Sometimes, the solution is as easy as using a high-ohm resistive divider to provide a mid-supply potential. But if the ground impedance needs to be lower, the simple approach can be enhanced by using a buffer amplifier.

When the ground created this way sees high imbalance currents, though, things become more complicated. The buffer not only needs to be a power buffer, it will also have to dissipate the result of any imbalance. This is wasteful, and it may even require bulky, inconvenient heatsinks.

This circuit addresses these issues simply, cheaply, and effectively *(see the figure)*. It retains the resistive divider plus buffer configuration, but with a big difference. The buffer is now a self-oscillating Class D amplifier, capable of providing many amperes of ground current and of "recycling" any imbalance back to the supply. The circuit is compact and efficient, while providing a stiff, accurate, and powerful ground.

The oscillator is built around a simple CD4069 CMOS hex buffer. Most of its inverters act as gate drivers for the power MOSFETs. Inverter U1 provides the non-inverting input of the amplifier, as the inverting input is implicit and is represented by the CD4069's mid-supply potential. Capacitor C3 provides positive feedback and allows the circuit to self-oscillate.

The entire control circuit has a floating supply referenced to the output. Diodes D3 to D6 and resistor pair R6/R7 generate this supply from the dc input. R2 and R3 create the virtual ground potential. In this example, they are equal, as will generally be the case, but asymmetrical supplies also can be created.

For greater accuracy, a trim potentiometer can be added to the midpoint of R2/R3. With the values shown, the pulse-width modulation (PWM) frequency is approximately 45 kHz, which can be modified by changing C3.

The power MOSFETs and inductor must be sized according to the expected output current. If a very clean ground is required, an additional LC filter can be used after L1 and C6/C7 to clean up the remaining switching residues.

LOUIS VLEMINCQ works as a design engineer in Physcial Layer Engineering at Belgacom, Evere, Belgium.

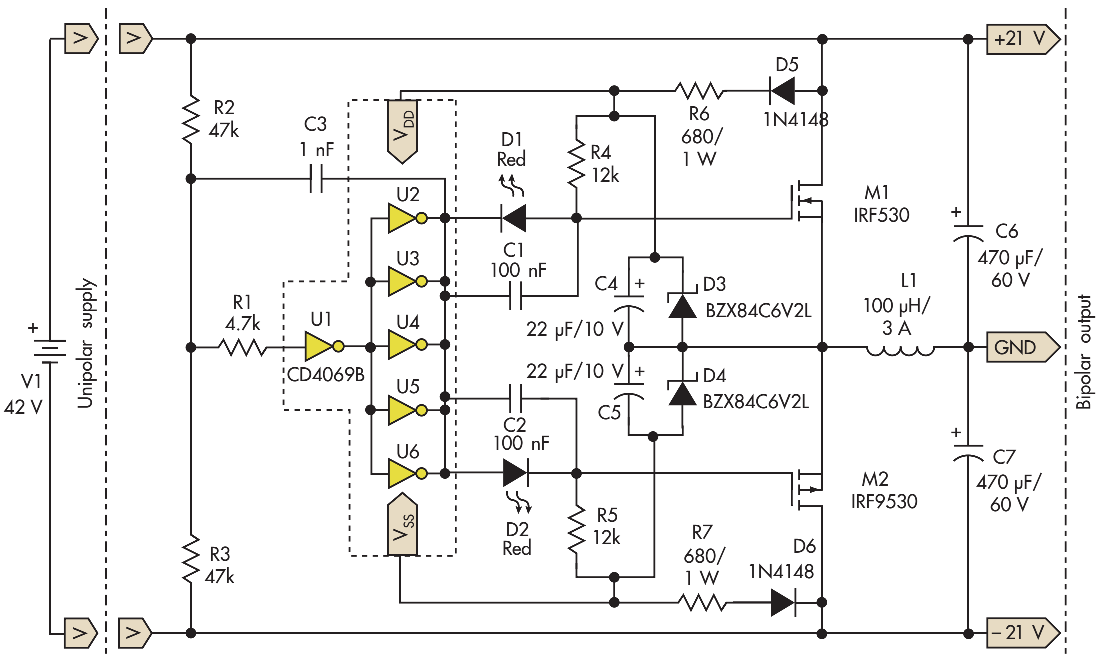

Although it requires more components than the conventional resistor-divider approach, this Class D oscillator/amplifier (based on a CD4069 hex inverter) drives power MOSFETs that make the virtual ground much more efficient.

Micro-Fabricated Magnetic Reed Switch Sets Size And Performance Records

CONVENTIONAL MAGNETIC reed relays provide excellent low-power characteristics combined with a low level of electrostatic discharge (ESD) sensitivity, reasonably good high-power switching capability, long switching lifetimes, high magnetic sensitivity, and low costs. But they're limited in how small they can be made.

The RS-A-2515 RedRock switch from Coto Technology solves the size problem yet retains all of the high-performance properties of classical magnetic reed switches. It is the industry's smallest micro-fabricated magnetic reed switch with a footprint of less than 2.1 mm squared (1.01 by 2.08 mm) and a height of just 0.94 mm *(Fig. 1)*.

The surface-mount technology (SMT) device targets the sweet spot of magnetic switching applications by combining zero power operation and high-current hot switching capability in a small package, compared to other magnetic switching technologies. It features an operating range of 10 to 25 milli-Teslas (mT) and a release range of 5 to 15 mT. It also can switch 300 mW and handle voltages of 100 V dc (70 V ac rms) and currents of 50 mA dc (35 mA ac rms). Contact resistance is just 3 Ω (7 Ω maximum). And, its breakdown voltage rating is 200 V dc.

Manufactured using high-aspect-ratio micro-fabrication (HARM), a MEMS-based process, it offers much higher magnetic sensitivity and high-power switching capability than planar MEMS switches. A structure can be grown at a height orders of magnitude greater than its width. Yet the HARM process allows the RedRock switch to be offered at a relatively low price that's cost-competitive with conventional reed switches when other superior parameters are considered.

Like MEMS IC manufacturing, the HARM process is a micromachined batch process. The RedRock switch

1. Coto Technology's micro-fabricated magnetic reed switch is the smallest in the industry with a footprint of less than 2.1 mm squared (1.01 by 2.08 mm) with a height of just 0.94 mm. It is manufactured using a MEMS-based HARM process.

is hermetically sealed and is made using wafer-level packaging (WLP). Consequently, it can achieve the inherent benefits of MEMS processing, namely a small size, low-cost manufacturability, and item-to-item reproducibility, as well as comparatively good to excellent performance parameters.

The RedRock's contact blades are grown upward from the ceramic base of the switch using a lithographically produced sacrificial mold relative to the switch's substrate. The precise dimensions of this mold and its extremely parallel walls ensure that the thickness of the reed switch blade and the contact gap are controlled to within a fraction of a micron *(Fig. 2)*.

Unlike a typical planar MEMS switch where the blades move parallel to the substrate, the HARM process allows the blades to be made as wide as possible. Also, its footprint isn't increased no matter how wide the blades are made, enabling higher current handling, lower contact resistance, and longer life.

The RedRock's metal cantilever bridges two massively electrically isolated metal blocks that act as magnetic amplifiers, much like the external leads in a conventional reed switch. Magnetic flux from an external magnet builds up in a small gap between the cantilever and one of the blocks and pulls the cantilever into electrical contact with the block. The contacts are coated with ruthenium for maximum contact longevity.

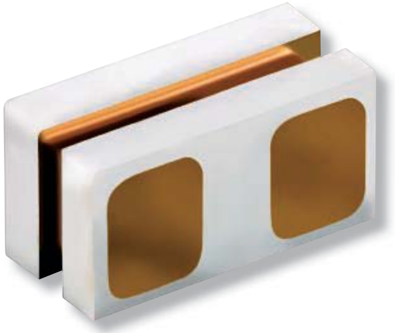

2. The contact blades on the RedRock switch are grown upward from the ceramic base of the switch using a lithographically produced sacrificial mold relative to the switch's substrate. The precise dimensions of this mold and its extremely parallel walls ensure that the thickness of the reed switch blade and the contact gap are controlled to within a fraction of a micron.

Compared to a planar rhodium-coated-contact MEMS switch structure, the Red Rock's blade is 1500 µm long (versus 550 µm for the planar MEMS switch), 200 µm wide (versus 100 µm), and 25 µm thick (versus 6 µm). A contact gap of 4 µm is used in both types of switches.

Although the RedRock switch requires 400 µN of closure force (versus 21 µN) and 45 µN for opening (versus 6 µN), it has vastly improved performance in a contact resistance of just 3 to 5 Ω (versus 500 to 1000 Ω) as well as a breakdown voltage of 200 V (versus 75 V). It also can carry a greater maximum current. The minimum melt current is 250 mA versus 0.7 to 14 mA.

Designed for high-performance applications that require extremely small switch size, the RedRock suits ingestible endoscopic capsules, insulin delivery for implantable pumps, and hearing aid switches. In the automotive sector, applications include level sensing for brake-fluid reservoirs.

The RS-A-2515 RedRock switch is available from stock in evaluation quantities at a unit price of $29.95. Additionally available from stock is an evaluation kit, EVAL, that sells for $49.95 each.

COTO TECHNOLOGY

www.cotorelay.com **ROGER ALLAN**

Low-Power, Single-Chip APU Delivers High Performance

THE SINGLE-CHIP Embedded G-Series accelerated processing unit (APU) family from AMD uses the company's latest 2-GHz Jaguar core and Radeon 8000 graphics processing unit (GPU). Used on AMD's higher-end platforms, the Jaguar supports a 40-bit physical address space and doubles the floating-point bus width to 128 bits. Jaguar also doubles the load/store bandwidth to 16 bytes/cycle, compared to the Bobcat architecture.

AMD's original APU was a multichip solution that integrated the CPU and GPU on a single chip. According to some rumors, Sony's future Playstation 4 will use an APU. Single-chip x86 platforms have been available and popular, but typically the video support was limited. The Radeon GPU is a full-function solution that also can be used for computation chores using OpenCL support. During Design West at 2013 in San Jose, AMD exhibited the 25- by 25-mm package in compact, fanless designs with some impressive graphics.

The chips will be available in a range of configurations with a thermal design power (TDP) from 9 to 25 W. Each chip incorporates a multicore GPU along with new features like multimedia support for low-latency wireless displays. Also, each chip is combined with an Enhanced Universal Video Decode (UVD) Engine 4.2 and Video Compression Engine (VCE) 2.0. And, the Embedded G-Series supports standard display interfaces such as HDMI, DVI, LVDS, and VGA and can manage two displays using the DisplayPort 1.2 interface.

The 28-nm Embedded G-Series APU architecture uses four cores to support four threads, which is comparable to Intel's dual-core Atom with hyperthreading. It has one x4 and four x1 PCI Express links. The 2-Mbyte L2 cache is shared between all cores. The DDR3-1600 interface supports two small-outline dual-inline memory modules (SO-DIMMs). Each chip supports double-data rate (DDR) power states (P-states) to help reduce power requirements. Error correction code (ECC) support will allow the chips to be used in secure or high-reliability applications.

AMD's Embedded G-Series APU chips use only 600 mm^2 of space. Versions with a TDP from 9 to 25 W are available.

The Embedded G-Series chips have a wide selection of peripheral interfaces but use off-chip networking support that typically will be linked via PCI Express. A 6-Gbit/s SATA v3.x interface provides storage support. There also is support for LPC, SPI, USB 2.0/3.0, and an SD reader. The chips boast encryption acceleration and a Security Asset Management Unit (SAMU) 2.1 as well. They support Windows 8, Linux, and Android. And, the use of the Radeon 8000 CPU architecture will simplify video driver support.

AMD has only announced the chips so far, but they will show up in new designs. There likely will be a variant for GizmoSphere, which currently has a multichip APU implementation. It looks like an ideal match for AMD's server group, SeaMicro, which started with the SM10000 and its 512 Atom cores. The SeaMicro fabric chip has x4 PCI Express processor interfaces.

AMD
www.amd.com **BILL WONG**

Stackable Headers Add High-Speed Automotive Interface

COMBINING THE high-speed USCAR-30 HSAutolink interface with the stackable Stac64 connector system, Molex's HS Stac headers (at a 0.80-mm pitch) add flexibility and save space in printed-circuit board trace routing for in-vehicle infotainment and telematics.

The right-angle headers package the interface in a rugged, shroud-stackable formation compatible with the company's Stac64 header family. Such a design creates a multi-bay, high-speed connection system that meets in-vehicle terminal requirements for devices and modules in automotive and commercial transportation applications.

USCAR-30, an automotive industry-standard interface, supports USB 2.0 requirements and other high-speed technologies such as low-voltage differential signaling and Ethernet for infotainment, car stereo and navigation systems, and driver-assist electronic modules, as well as telematics.

MOLEX INC.

www.molex.com *ED* **NEWS STAFF**

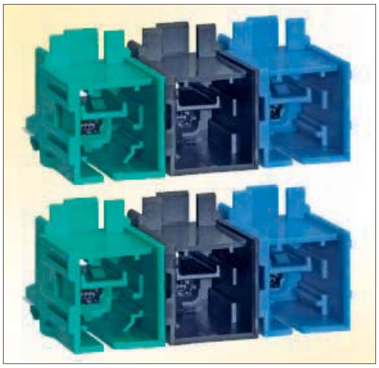

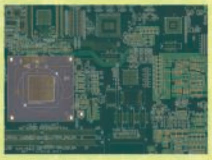

IR's 40V Automotive-qualified COOLiRFET™ Deliver Benchmark On-State Resistance for Improved System Efficiency in Heavy Load Applications

IR introduces a family of automotive-qualified COOLiRFET™ MOSFETs delivering benchmark on-state resistance (Rds(on)) for heavy load applications including electric power steering (EPS), braking systems and other heavy loads on internal combustion engine (ICE) and micro hybrid vehicle platforms.

The family of 22 AEC-Q101 qualified 40V N-channel MOSFETs feature IR's proven Gen12.7 trench technology that delivers ultra-low Rds(on) in D2Pak-7P, D2Pak, DPak, TO-262, IPAK and TO-220 packages. The benchmark D2Pak-7P AUIRFS8409-7P delivers Rds(on) max as low as 0.75mOhm at 10Vgs with a current rating up to 240A. The new devices offer low conduction losses and robust avalanche performance to deliver higher efficiency, power density and reliability. With this performance, many applications using these new COOLiRFET™ devices run significantly cooler than with state-of-the-art MOSFETs.

IR's automotive MOSFETs are subject to dynamic and static part average testing combined with 100 percent automated wafer level visual inspection as part of IR's automotive quality initiative targeting zero defects. The new devices are AEC-Q101 qualified, which requires that there is no more than a 20 percent change in Rds(on) after 1,000 temperature cycles of testing, and feature an environmentally friendly, lead-free and RoHS compliant bill of materials.

More information is available on the **International Rectifier** website at
http://www.irf.com/whats-new/nr130514.html
For more information, contact Sian Cummins,
scummin1@irf.com, 310-252-7148.

For more information on products or services visit our website www.electronicdesign.com, menu item Reader Service. The advertisers index is prepared as an extra service. *Electronic Design* does not assume any liability for omissions or errors.

BILL WONG | Embedded/Systems/Software Editor
bill.wong@penton.com

Supercomputer Robot
Cluster Sails Into The Sunset

When I first interviewed Liquid Robotics CEO Bill Vass, I was very impressed by the company's energy harvesting navigation system and its embedded data acquisition capabilities. The autonomous system completed a record breaking 9000-mile journey across the Pacific using only the waves and solar panels to power itself. It gathered and transmitted data as it sailed and distributed it to the world as part of the PacX Challenge.

It's no speed demon, but the Wave Glider is efficient with an almost unlimited energy supply because there is always wave movement that it can harvest. It comprises a float and fins connected by an umbilical cord *(Fig. 1)*. The float rides on the surface. It houses the computers and sensors along with batteries, solar panels, and communication equipment.

The Wave Glider moves in the desired direction and speed by angling the rudder. The articulated fins are about 6 meters below the float and provide forward propulsion via wave action. Solar panels on the top of the float provide electricity for the batteries that power the onboard computers, sensor payloads, and communications.

The larger Wave Glider SV3 complements the smaller SV2 *(Fig. 2)*. By going large, Liquid Robotics can cram in a waterproof 8U rack that uses salt water cooling. The SV3 also can house Arm-based blades that are currently implemented using dual-issue Texas Instruments OMAP processors. The current system supports up to 24 sockets.

The blades are built on Liquid Robotics' own backplane, which incorporates CAN and Ethernet. CAN provides board and power management. Liquid Robotics found CAN more advantageous than I^{2}C with lower power requirements than Ethernet. CAN alone would probably be sufficient for navigation, most sensors, and system control, but the company has more ambitious plans for the Wave Glider.

Wireless communications can link a swarm of Wave Gliders. The group of computers provides a greater computational resource for the swarm, but Liquid Robotics has taken the idea much further with multiple core clusters in each robot. The clusters run Java virtual machines on top of Linux with a CORBA software backplane, forming a massive multitenancy system or a "cloud on the water."

The company looked at other robotic platforms but did not find one suitable for the dynamics of its environment, so it built its own cloud-based operating system named Regulus. The multitenancy, Java-based cloud approach brings a lot to the table. It allows more flexibility in terms of deployment, so a number of Wave Gliders could be equipped primarily with sensors while others provide number crunching. The computational power can be used for analysis, planning, or data reduction to minimize uploads.

Wave Gliders can handle tasks such as monitoring temperature and chemicals and tracking sharks. They can almost run forever. So if you see some flags floating by, it just might be a swarm of Wave Gliders. ed

1. The Wave Glider consists of a floating section near the surface and an energy harvesting and traction system that is connected by an umbilical cord.

2. The Liquid Robotics Wave Glider SV3 (left) joins its older but smaller sibling, the SV2 (right), for a swim.

Follow Bill Wong on Twitter at #AltEmbedded.

Become a fan at facebook.com/ElectronicDesign.

100’s of Options
100’s of Solutions

SOLUTION 16:

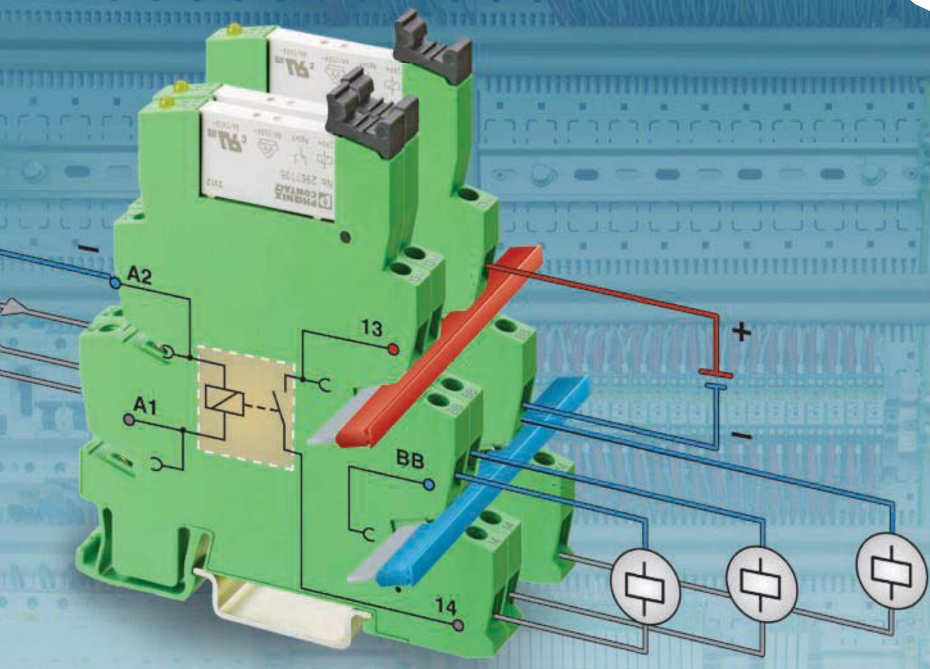

PLC Relay Actuator series from Phoenix Contact

Phoenix Contact has designed hundreds of products to make your job easier. For example, our PLC Relay Actuator Series is specifically for use with common 2-wire output devices.

By eliminating the normally closed (NC) contact — typically unused in most installations — this relay series provides three separate terminals on the output side for connection. As a result, all wires from devices like contactors, valves and small motors connect directly to the relay, without the need of an additional terminal block for power return.

Our PLC Relay Actuator Series:

- **Saves space on the DIN rail**
- **Reduces labor and component costs**
- **Eliminates an extra point of failure**

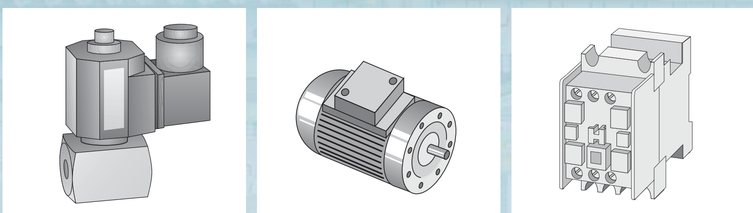

Valve Small Motor Contactor

To learn more, call **1-800-322-3225** or visit **www.phoenixcontact.com/actuatorrelay**

PHOENIX CONTACT
INSPIRING INNOVATIONS