electronic design

THE AUTHORITY ON
EMERGING TECHNOLOGIES
FOR DESIGN SOLUTIONS

February 2016 electronicdesign.com

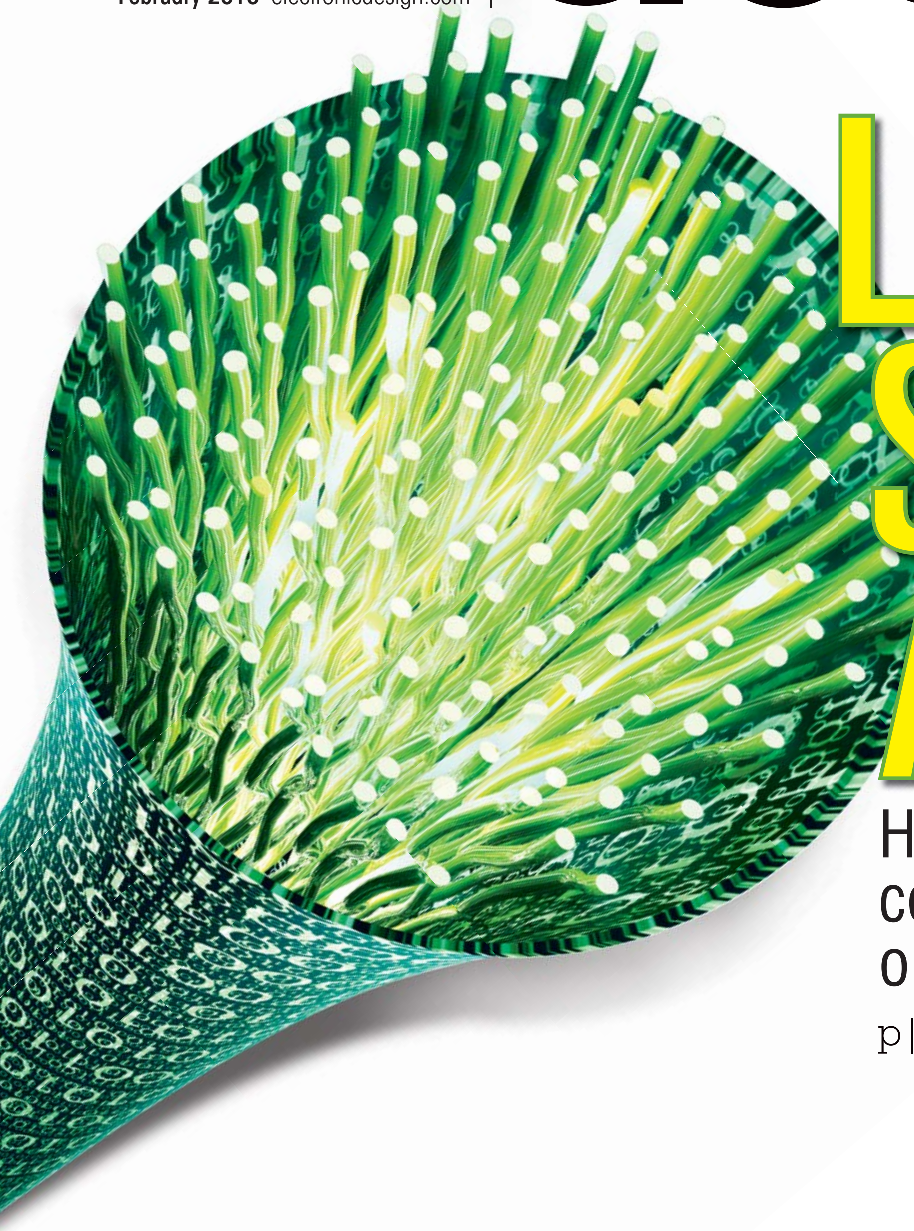

Light Speed Ahead

High-speed connectivity propels optical technology

p| 18

$10.00 Powered by Penton

They. Just. Work.

Discover the versatility of Keysight Data Acquisition Units.

The Keysight 34970A and the 34972A data acquisition (DAQ) units were designed with one goal in mind: simplicity. There's no programming. All you have to do is select the measurements you need and your DAQ will do the rest. Also their built-in 6½ digit multimeter offers ultra-low reading noise. With unprecedented peace of mind and a competitive price, it's easy to see why Keysight is the number one name in DAQs.

Keysight Data Acquisition Units – 34970A and 34972A
Built-in signal conditioning enables 11 measurement functions: temperature, voltage, frequency and more
Measure up to 60 differential channels, up to 300 DC/AC volts
BenchVue software-enabled to control, view and analyze data

Data log and visualize with Keysight's BenchVue software.

Learn more with our test challenge measurement briefs and videos at
www.keysight.com/find/daqpromoUS

Buy from an Authorized Distributor:

microlease Newark element14 TESTEQUITY

KEYSIGHT TECHNOLOGIES

Unlocking Measurement Insights

Agilent's Electronic Measurement Group is now **Keysight Technologies.**

February 2016 VOLUME 64, ISSUE 2

electronic design

InThisIssue

FEATURES

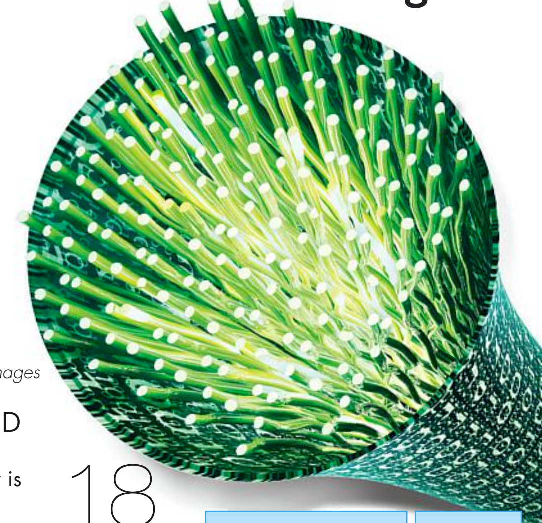

cover art courtesy of Thinkstock Images

18

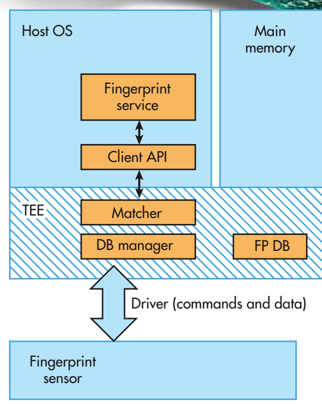

22

NEWS & ANALYSIS

IDEAS FOR DESIGN

COLUMNS & DEPARTMENTS

24

26

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 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 Kansas City, MO, 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. Copyright ©2016 by Penton Media Inc. All rights reserved. The contents of this publication may not be reproduced in whole or in part without the consent of the copyright owner. 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.

on electronicdesign.com

SECURING THE 5G NETWORK

http://electronicdesign.com/communications/securing-5g-network

From an extremely low-latency communication network for the self-driving car to high latency designed to support the long battery life of devices associated with the Internet of Things (IoT), everyone wants a piece of the 5G network. However, each of these applications also brings along a variety of security-related challenges.

WIRELESS SENSOR NETWORKS MONITOR ACTIVE VOLCANOES IN JAPAN

http://electronicdesign.com/iot/wireless-sensor-networks-monitor-active-volcanoes-japan

A new wireless sensor network being installed in Japan could help scientists more accurately predict the behavior of the country's most active volcanoes. The system will gather enormous amounts of data used to forecast volcanic activity, identifying when it might be necessary to issue warnings or evacuations. *(Image courtesy of Kimon Berlin via Flickr)*

blogs

LOUIS FRENZEL
COMMUNICATIONS

- The Internet of Things: Hype, Hope, or Hit?

BILL WONG
EMBEDDED/SYSTEMS/SOFTWARE

- Drones Go Gaming

PATRICK MANNION
TEST AND MEASUREMENT

- DesignCon Mini Tour Highlights USB Type-C, Power, and Memory Test

"UNDERGROUND BATTERY" STORES RENEWABLE ENERGY FOR POWER GRID

http://electronicdesign.com/power/underground-battery-stores-renewable-energy-power-grid

The international accord drafted by 195 countries at the Paris climate talks last year is focused on keeping the average global temperature increase below 1.5°C. The agreement will require a reevaluation of how a renewable power grid will work. Now, a research team led by the Lawrence Livermore National Laboratory (LLNL) in California has proposed a new method for satisfying these requirements. *(Image courtesy of Alstom)*

WIRELESS-CHARGING TECH: TRANSFORMING THE MOBILE WORLD

http://electronicdesign.com/power/wireless-charging-technologies-transforming-mobile-world

The race to create a universal wireless-charging technology is on. New and exciting wireless-charging technological advances and products continue to emerge, with some companies committed to one technology and others working with several. But a lack of interoperability between technologies still remains a major issue, due to the non-existence of a unified wireless-power standard. *(Image courtesy of Powermat)*

join us online

Follow us on Facebook: http://tinyurl.com/odzo6hc
and Twitter: http://tinyurl.com/k5pum39

electronic design
Powered by Penton
FEBRUARY 2016

EDITORIAL

CONTENT DIRECTOR: **NANCY K. FRIEDRICH** nancy.friedrich@penton.com
CONTENT PRODUCTION DIRECTOR: **MICHAEL BROWNE** michael.browne@penton.com
CONTENT PRODUCTION SPECIALIST: **ROGER ENGELKE** roger.engelke@penton.com
PRODUCTION EDITOR: **JEREMY COHEN** jeremy.cohen@penton.com
DISTRIBUTION: **VICTORIA FRAZA KICKHAM** SourceESBeditor@penton.com
EMBEDDED/SYSTEMS/SOFTWARE: **WILLIAM WONG** bill.wong@penton.com
ANALOG/POWER: **MARIA GUERRA** maria.guerra@penton.com
ASSOCIATE CONTENT PRODUCER: **LEAH SCULLY** leah.scully@penton.com
ASSOCIATE CONTENT PRODUCER: **JAMES MORRA** james.morra@penton.com
CONTRIBUTING EDITOR: **LOUIS E. FRENZEL** lou.frenzel@penton.com

ART DEPARTMENT

GROUP DESIGN DIRECTOR: **ANTHONY VITOLO** tony.vitolo@penton.com
SENIOR ARTIST: **JIM MILLER** jim.miller@penton.com
CONTRIBUTING ART DIRECTOR: **RANDALL L. RUBENKING** randall.rubenking@penton.com

PRODUCTION

GROUP PRODUCTION MANAGER: **CAREY SWEETEN** carey.sweeten@penton.com
PRODUCTION MANAGER: **FRAN VAUGHN** fran.vaughn@penton.com

AUDIENCE MARKETING

USER MARKETING DIRECTOR: **BRENDA ROODE** brenda.roode@penton.com
USER MARKETING MANAGER: **DEBBIE BRADY** debbie.brady@penton.com
FREE SUBSCRIPTION/STATUS OF SUBSCRIPTION/ADDRESS CHANGE/MISSING BACK ISSUES
T | 866.505.7173 **F** | 847.763.9673 electronicdesign@halldata.com

SALES & MARKETING

MANAGING DIRECTOR: **TRACY SMITH** **T** | 913.967.1324 **F** | 913.514.6881 tracy.smith@penton.com

REGIONAL SALES REPRESENTATIVES
AZ, NM, TX: **GREGORY MONTGOMERY** **T** | 480.254.5540 gregory.montgomery@penton.com
AK, CA, CO, ID, MT, ND, NV, OR, SD, UT, WA, WI, WY, W/CANADA: **JAMIE ALLEN** **T** | 415.608.1959 **F** | 913.514.3667 jamie.allen@penton.com
AL, AR, IA, IL, IN, KS, KY, LA, MI, MN, MO, MS, NE, OH, OK, TN: **PAUL MILNAMOW** **T** | 312.840.8462 paul.milnamow@penton.com
CT, DE, FL, GA, MA, MD, ME, NC, NH, NJ, NY, RI, PA, SC, VA, VT, WV, EASTERN CANADA:
SHANNON ALO-MENDOSA **T** | 978.501.7303 Shannon.alo-mendosa@penton.com

INTERNATIONAL SALES
GERMANY, AUSTRIA, SWITZERLAND: **CHRISTIAN HOELSCHER** **T** | 011.49.89.95002778 christian.hoelscher@husonmedia.com
BELGIUM, NETHERLANDS, LUXEMBURG UNITED KINGDOM, SCANDINAVIA, FRANCE, SPAIN, PORTUGAL:
JAMES RHOADES-BROWN **T** | +011 44 1932 564999 **M** | +011 44 1932 564998 james.rhoadesbrown@husonmedia.com
PAN-ASIA: **HELEN LAI** **T** | 886 2-2727 7799 helen@twoway-com.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

LIST RENTALS:
SMARTREACH CLIENT SERVICES MANAGER: **DAVID SICKLES** **T** | (212) 204 4379 david.sickles@penton.com

ONLINE

PRODUCT DEVELOPMENT DIRECTOR **RYAN MALEC** ryan.malec@penton.com

DESIGN ENGINEERING & SOURCING GROUP

EXECUTIVE DIRECTOR OF CONTENT AND USER ENGAGEMENT: **NANCY K. FRIEDRICH**
GROUP DIRECTOR OF OPERATIONS: **CHRISTINA CAVANO**
GROUP DIRECTOR OF MARKETING: **JANE COOPER**

PENTON

CHIEF EXECUTIVE OFFICER: **DAVID KIESELSTEIN** david.kieselstein@penton.com
CHIEF FINANCIAL OFFICER: **NICOLA ALLAIS** nicola.allais@penton.com
INDUSTRY GROUP, PRESIDENT: **PAUL MILLER** paul.miller@penton.com
1166 AVENUE OF THE AMERICAS, 10TH FLOOR
NEW YORK, NY 10036 **T** | 212.204.4200

Penton

Electronic Design | Machine Design | Microwaves & RF | Medical Design | Source ESB | Hydraulics & Pneumatics | Global Purchasing | Distribution Resource | Power Electronics | Defense Electronics | Electronic Design Europe | Engineering TV


CONSTANTLY
KEEPING YOU
CONNECTED
MPD DEVICES CAN BE FOUND IN
SERVER EQUIPMENT AROUND THE WORLD
DL2032
DURACELL
MPD
MEMORY PROTECTION DEVICES
MPD IS A GLOBAL MANUFACTURER OF BATTERY HOLDERS AND OTHER ELECTRONIC COMPONENTS. WE BELIEVE THAT OUR COMPONENTS SHOULD FIT EASILY INTO YOUR DESIGNS, WHICH IS WHY WE ARE ALWAYS CREATING INNOVATIVE NEW PRODUCTS.
BATTERYHOLDERS.COM

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

Automotive Ethernet Was the Hidden Trend at CES 2016

There were a lot of trends emerging at the 2016 Consumer Electronics Show, like self-driving and electric vehicles. Chevy's 2017 Bolt (*Fig. 1*) was one of the hot items at the show. It is designed to go over 200 miles between charges. The $33,170 MSRP will be even lower after the federal rebate. The Bolt is a roomy car that does 0 to 60 mph in 7 seconds.

1. Chevy's 2017 Bolt will come in under $30,000 with rebate and is capable of driving 200 miles between charging.

Toyota was showing off artificial intelligence (AI) demos using robotic models that distributed machine learning information among the group of tiny cars. The company has promised to invest $1 billion in Toyota Research Institutes to investigate AI for making cars safer and more intelligent.

The big automotive trend, though, is the use of Ethernet under the hood. This is in the form of 100 Mbit/s, 802.3bw 100Base-T1, also known as BroadR-Reach, and 1 Gbit/s, 802.3bp 1000Base-T1. Both utilize a single twisted pair allowing the use of low-cost connectors and significantly lowering cabling costs. The systems provide full duplex operation and it is possible to provide power over the same twisted pair using the P802.3bu 1-Pair Power over Data Lines (PoDL) standard.

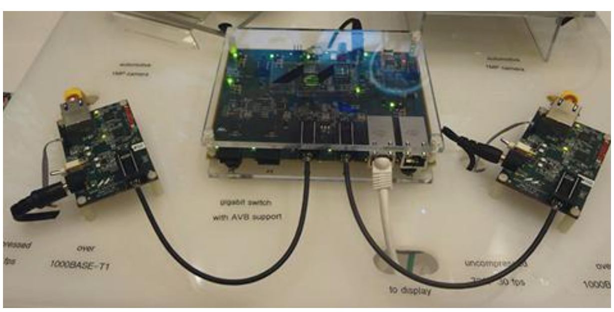

2. Marvell was showing off two-wire, 1000-BaseT1 networking including switches and endpoints.

The OPEN Alliance (One-Pair Ether-Net) Special Interest Group (SIG) is an industry group supporting 802.3bw.

The BroadR-Reach technology was developed by Broadcom and is available to OPEN Alliance members under RAND terms. BMW is already using the technology to connect the four cameras for surround-view in the BMW X5. The faster 1000Base-T1 can handle multiple 4K video streams (*Fig. 2*). The standards utilize the normal Ethernet MAC, but change the PHYs.

The new Ethernet standards are designed for fast startup, but limit the run length to 15 m. This is more than sufficient to handle most transportation-related applications. It is even applicable to larger platforms where switches could be connected using conventional Ethernet cabling that can handle longer connections.

The new standards target automotive applications, but they will be very useful in a range of embedded applications as well. Robotics is one that comes to mind. Ethernet is already in common use and robots would benefit from lower costs and simplified cabling. ed

News

ELECTRIC-VEHICLE CHARGING
Opens Up Semiconductor Market

(Image courtesy of Thinkstock)

The charging infrastructure for electric vehicles is expanding throughout the United States, buffered by large investments from the automotive industry, government initiatives, and international calls to reduce carbon emissions. As a result, the market for semiconductors used in these charging stations is expected to grow rapidly over the next few years, according to a new report from chip industry research firm IHS Technology.

The revenues earned from the semiconductors built into charging stations was around $44 million in 2014, according to IHS Technology. However, the new report predicts that as more electric vehicles hit the road and the demand grows for additional charging stations, the global market for these components will reach approximately $233 million in 2019.

The bill of semiconductors used in charging stations is vast, ranging from the power semiconductors that channel electricity between the charger and the vehicle, to the communication chips linking the station to the smart grid. The report adds that the electric vehicle companies are trending toward systems-on-a-chip (SoCs) that not only provide faster control, but also include the memory chips required for secure communications.

In the United States, the infrastructure surrounding electric vehicles is still extremely limited relative to the nearly 253 million vehicles that were on the road in 2014, only a tiny fraction of which were electric. According to the Alternate Fuels Data Center, the U.S. Department of Energy's clearinghouse for information on electric vehicles, there are now 11,578 electric charging stations and 29,436 public charging outlets spread across the United States, excluding private charging stations.

Noman Akhtar, an industrial semiconductors analyst for IHS, says that fast charging is a necessary step toward the widespread usage of electric vehicles and building out the infrastructure around them. And higher power ratings are required to support shorter charging times. "Electric vehicle charging stations with higher ratings require more power semiconductors, especially discrete semiconductor components," he says.

Charging equipment for electric vehicles is classified by the rate at which the batteries are charged. AC Level 1 chargers, which are normally compatible with household outlets, add about two to five miles of range to an electric vehicle per hour of charging time, according to the Alternate Fuels Data Center. Although they require special charging equipment, AC Level 2 chargers add approximately 10 to 20 miles of range per hour of charging time.

On the other hand, direct-current (DC) fast chargers represent the core technology within the public charging infrastructure, adding 50 to 70 miles of range in about 20 minutes, according to the Alternate Fuels Data Center. These stations are used by Mitsubishi, Nissan, and Tesla Motors in several different configurations and charging speeds. Tesla Superchargers can apparently provide 170 miles of range in about a half hour, delivering up to 120 kW of DC power directly to the battery.

(continued on p. 12)

DUAL-JUNCTION SOLAR CELL Breaks Efficiency Record

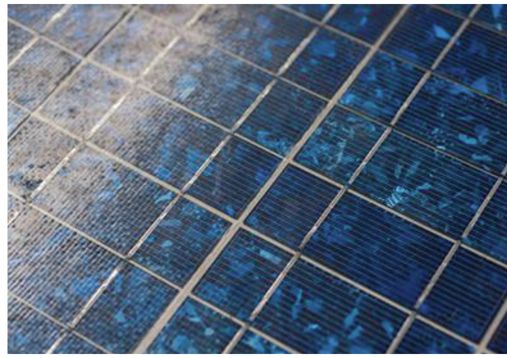

A close-up image of a solar panel. Mass-produced silicon solar cells are typically less than 20% efficient, but a new multi-junction device has demonstrated an efficiency of 29.8%. *(Image courtesy of Thinkstock)*

A RESEARCH TEAM based out of the U.S. National Renewable Energy Laboratory (NREL), has developed a multi-junction solar cell that they claim sets a new efficiency record. The research team partnered with the Swiss Center for Electronics and Microtechnology (CSEM) to create the so-called tandem solar cell, combining two layers of semiconductor material to absorb more of the solar spectrum.

In laboratory tests, the research team demonstrated that the solar cell could convert direct sunlight into electricity at 29.8% efficiency. David Young, a senior researcher at the NREL, points out that these results edged out the theoretical limit of 29.4% for mechanically stacked cells. In addition, the device works without having to concentrate sunlight with reflectors, which can increase efficiency in certain solar cells.

Each research center contributed part of the dual-junction solar cell, which combines III-V and crystalline silicon semiconductors. CSEM scientists developed a silicon sub-cell, on which NREL stacked a layer of gallium-indium phosphide (GaInP). The resulting device has a higher efficiency than either material by itself. The record efficiency for an individual crystalline silicon cell is 25.6%, and 20.8% for single-junction GaInP.

"We believe that silicon heterojunction technology [that combines different crystalline semiconductors] is today the most efficient silicon technology for application in tandem solar cells," says Christophe Ballif, head of photovoltaic research at CSEM.

The researchers gave few additional details in a recent news release, but Young has submitted the team's research paper to the IEEE Journal of Photovoltaics for publication. The experimental results were published in the journal *Progress in Photovoltaics* in an article that reviews solar-cell designs and divulges the highest efficiency in each category.

The review, which includes solar cells up to three times more efficient than conventional solar panels, underlines the fact that efficiency is not everything for commercial solar cells. The highest efficiency ever recorded was 46% for a multi-junction device under highly concentrated sunlight. Soitec, a French company that makes photovoltaic semiconductors, engineered the solar cell in 2014. But the company has since stopped producing this technology.

Despite these higher efficiencies, multi-junction solar cells have been kept out of the commercial market due to their complex structure and high manufacturing costs. These devices have typically been limited to satellites and other spacecraft. On the ground, they have to compete with the gradually falling cost of crystalline silicon, the most prevalent material for conventional solar cells. Mass-produced silicon cells are typically less than 20% efficient, but they are relatively cheap compared to the exotic material normally used in multi-junction cells.

The NREL, which is the primary research laboratory for the U.S. Energy Department, is examining several different semiconductor materials for solar cells. Last November, the laboratory found a way to significantly reduce the amount of energy lost to heat in perovskite-based solar cells. The discovery could one day lead to solar cells that convert up to two-thirds of sunlight to electricity.

The dual-junction research was funded in part by the Energy Department's SunShot Initiative, a program aimed at making solar energy cost-competitive with fossil fuels. Additional funds were provided by the Swiss Confederation and Nano-Tera.ch, a Swiss green technology fund. ■

ELECTRIC VEHICLE CHARGING *(continued from p.10)*

In 2014, according to the IHS report, the average price for semiconductor components in a level-two charging station was about $143. By comparison, semiconductor components used in the latest fast-charging DC chargers now cost more than $1000. DC chargers are almost universally favored in public infrastructure because they provide faster charge times, but it has become increasingly common for hotels and parking garages to invest in lower-cost AC chargers for vehicles parked in the same place for hours.

On another front, semiconductors are also required for the communications modules within the charging stations. In addition to smart-grid compatibility, charging stations will eventually gain features like credit-card readers, billing software, high-resolution displays, automated diagnostics, controlled power flow, and internal metering—all of which will require processors and other semiconductor components.

The IHS Technology report came as the climate change summit in Paris ended earlier this month, punctuated by an international accord reached by 195 countries to lower greenhouse gas emissions and mitigate the effects of climate change. According to a 2013 study by the Environmental Protection Agency, about 27% of the United States' greenhouse gas emissions come from the transportation industry, only slightly less than the emissions from generating electricity. ■

DANIEL BOGDANOFF | Keysight Technologies Inc.

Make Great Oscilloscope

Making good measurements is... well... good. But if you want to make great measurements, know how to best use your oscilloscope to squeeze the most out of sample rates and bits of resolution.

In oscilloscopes today, making a good signal measurement is easy. However, making a great measurement takes some expertise. As edge speeds increase and voltages decrease, your signal's margin of error becomes smaller and smaller. So, making a great oscilloscope measurement could mean the difference between meeting and not meeting your design parameters.

Digital oscilloscopes have turned manual measurements into an automated process, allowing you to make adequate measurements quickly and easily. However, turning that good measurement into a great measurement requires a little extra knowledge and effort. Following some key best practices will give you much higher accuracy and confidence in your measurements. Signal scaling and bits of resolution will impact the quality of your measurement, so it's crucial to understand how to get the most out of your oscilloscope.

START BY MAKING A GOOD MEASUREMENT

First, set up your scope to make a measurement. Simply press the "Measure" button on the front of your oscilloscope. The scope will then automatically display the requested measurement on the screen. You can configure the scope to modify the measurement parameters, track the measurement with cursors, and even display measurement statistics. To improve the quality of this measurement, consider the following best practices.

SIGNAL SCALING

Proper signal scaling is crucial for optimizing measurements. The signal's scaling on screen affects sample rate and bits of resolution, which in turn affects your measurement's accuracy. Both horizontal scaling and vertical scaling influence your measurement in different ways.

1. This signal is improperly scaled; it should be scaled horizontally to only show about one period.

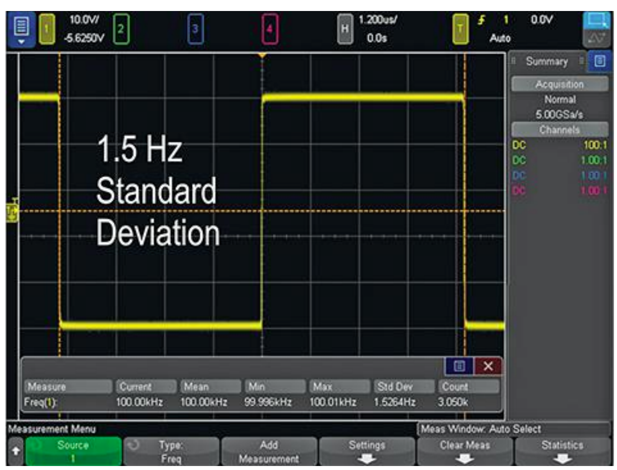

2. By improving the horizontal scaling, you're able to make much better measurements.

Measurements

HORIZONTAL SCALING

Horizontal scaling is important to consider when making time-dependent measurements. When you change the horizontal scaling (time-per-division) of your signal, you're also changing the total signal-acquisition time. Signal-acquisition time in turn affects the sample rate of the scope. This relationship is described by the equation:

Sample Rate = Memory Depth/Acquisition Time

Memory depth is a fixed value, and the acquisition time is fixed by adjusting the time-per-division setting on your oscilloscope. As the acquisition time increases, the sample rate will have to decrease in order to fit the entire acquisition into the scope's memory. Having an appropriate sample rate for time-dependent measurements (frequency, pulse width, rise time, etc.) is important.

For example, take a frequency measurement of a known 100-kHz clock signal *(Fig. 1)*. At 2 ms/division, we can see that the frequency measurement has an average of 99.6 kHz with a standard deviation of 1.48 kHz. Based on the standard deviation value, this gives us a measurement accuracy of about 1.5%. Notice that the sample rate is 100 Msamples/s due to the larger time-per-division setting of the scope.

Looking at the exact same 100-kHz clock signal at 1.2 µs/division instead of 2 ms/division *(Fig. 2)*, we see that the average measured frequency is 100.00 kHz and the standard deviation is 1.53 Hz. This gives us measurement accuracy of around 0.001%. Also note that the sample rate is 5 Gsamples/s, which is the maximum sample rate for the Keysight InfiniiVision 4000 X-Series oscilloscope being used.

By horizontally scaling the signal to only show about one period, we were able to maximize the sample rate and improve the standard deviation of our frequency measurement by a factor of 1000. While 99.6 kHz was a good measurement, our knowledge of horizontal-scaling techniques enabled us to make a measurement three orders of magnitude more accurate.

VERTICAL SCALING

Just as horizontal scaling is important for time-specific measurements, vertical scaling is important for vertically dependent measurements (peak-to-peak, RMS, max, min, etc.).

Take, for example, the same 100-kHz clock signal we discussed for horizontal scaling. This time we will look at a peak-to-peak voltage measurement *(Fig. 3)*. At 770 mV/division,

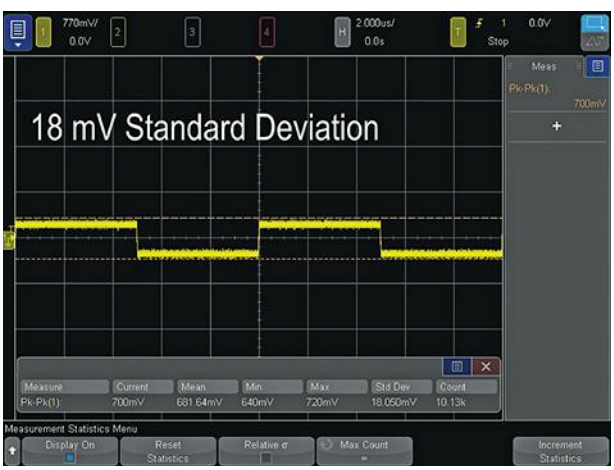

3. This signal is too small on the screen; it should be scaled vertically as much as possible.

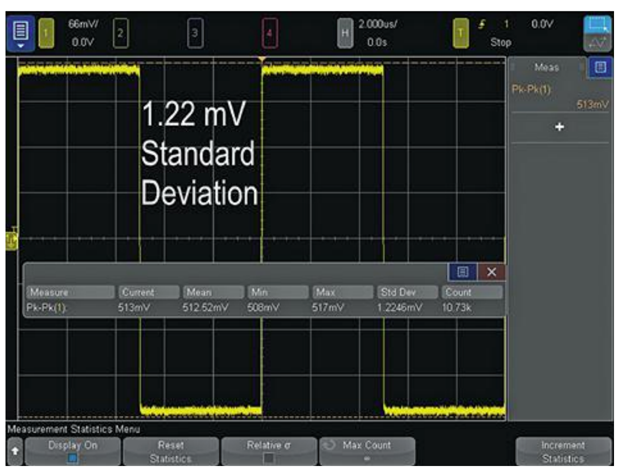

4. By vertically filling up the screen with your signal, your measurements improve significantly.

> “Understanding how to scale your signals properly to take full advantage of sample rate and bits of resolution is a crucial skill for success while testing and debugging.”

we see that the average V_{P-P} value is 681 mV and the standard deviation is 18.1 mV.

By simply increasing the voltage-per-division scaling *(Fig. 4)*, the average value becomes 512.5 mV and the standard deviation becomes 1.2 mV.

Increasing the vertical scaling of the signal enabled a much more accurate V_{P-P} measurement with a standard deviation that's 15 times smaller. Why does the vertical scaling have an effect on the measurements? Just as horizontal (time-dependent) measurements are affected by sample rate, vertical (amplitude-dependent) measurements are influenced by bits of resolution.

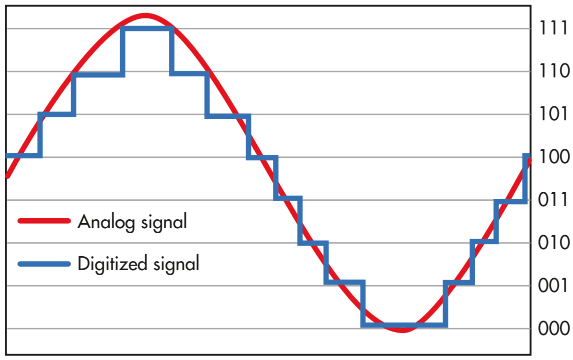

6. A three-bit ADC's analog input and digital output.

BITS OF RESOLUTION

Bits of resolution describes the precision of an analog-to-digital converter (ADC). The higher the bits of resolution, the more vertical quantization levels an ADC can identify when sampling an analog signal.

For example, *Fig. 5* shows a two-bit ADC. The analog-signal input is in red, and the quantized digital-signal output from the ADC is in blue. A two-bit ADC has four vertical quantization levels and doesn't have the ability to be any more precise. Therefore, the analog signal will be quantized down to one of four values.

In a three-bit ADC *(Fig. 6)*, we see that there are eight possible levels for the digital signal. This significantly increases the precision of the ADC, and the analog signal will be more accurately represented.

Most oscilloscopes have eight bits of resolution; the signal viewing area of the oscilloscope's screen represents eight bits. A digital oscilloscope uses the quantized waveform to make measurements. So, the more accurately the ADC portrays the analog signal, the more accurate the measurements. Looking back at *Fig. 3*, we see that by not scaling the signal to the full vertical height of the screen, we aren't taking advantage of all of the oscilloscope's bits of resolution. When we scaled the signal to maximize use of the oscilloscope's resolution *(Fig. 4, again)*, we were able to make a better V_{P-P} measurement.

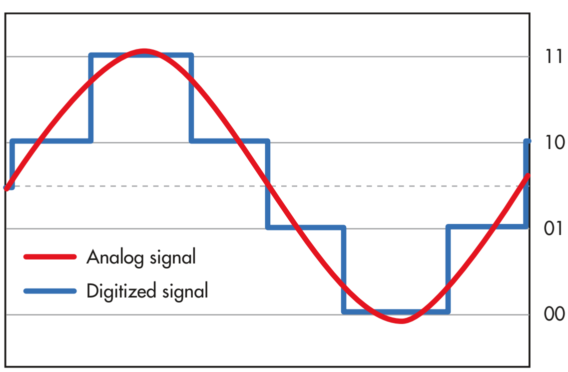

5. A two-bit ADC's analog input and digital output.

This has important ramifications, especially if you're using more than one channel at a time. For example, it's common to make power measurements using two channels (current and voltage). To make a great power measurement, you must scale both signals to the full height of the oscilloscope's screen instead of stacking them above each other on the screen. If you only use half of the screen for voltage and the other half for current, the oscilloscope can't use all bits of resolution on each waveform.

It is important to note, however, that the oscilloscope's bits of resolution are just one part of the system, and the entire signal path must be observed to consider the total effective number of bits. This includes external factors like probing and system noise, as well as internal factors such as your oscilloscope's noise floor. If the system's noise is high enough, an increased number of bits of resolution will only show noise.

Knowing how to best use your oscilloscope is crucial if you want to move from making good measurements to making great measurements. More accurate measurements further boost confidence in verifying and debugging designs. Understanding how to scale your signals properly to take full advantage of sample rate and bits of resolution is a crucial skill for success while testing and debugging. ed

DANIEL BOGDANOFF is product manager for the InfiniiVision series of oscilloscopes at Keysight Technologies. He graduated from Texas A&M with a degree in electrical engineering.

WILLIAM WONG | Technical Editor

High-Speed Interfaces Push OPTICAL CONNECTIVITY

Being somewhat expensive, fiber-optic cables have traditionally been utilized where their advantages—like long distance—were worth the cost. Now higher-speed connectivity is propelling optical technology into more common use.

Courtesy of Thinkstock Images

Being somewhat expensive, fiber-optic cables have traditionally been utilized where their advantages—like long distance—were worth the cost, such as servers and long-haul connections. But now, higher-speed connectivity requirements are pushing optical technology into more common use, even in consumer applications, as well as pushing multiple fiber connections to build out the cloud to meet Internet of Things (IoT) demands.

Optical systems are simple, in theory. There's a light source at one end and a detector at the other. Light goes through optical cable and is not affected by electromagnetic noise, like a wired connection. While there is signal loss, it's much less than a wired connection, making fiber the preferred method for long-distance connections.

Laser diodes or LEDs are the usual light source. Lasers include Fabry-Perot (FP), distributed feedback (DFB), and

1. Active fiber-optic cables include transceivers at both ends. Some include copper connections to provide power since it is often sourced at only one end of the cable. *(Courtesy of Corning)*

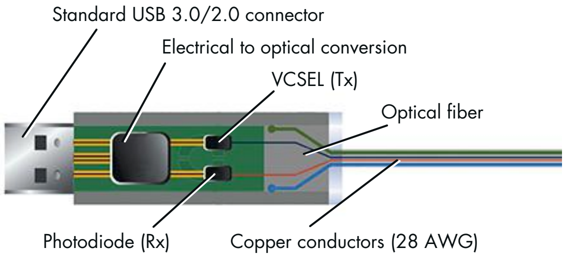

2. Corning's USB 3.0 uses standard USB connectors that incorporate transceivers at both ends. Two embedded fiber cables are used for data transmission, and a pair of copper connections provides power between transceivers.

vertical cavity surface-emitting lasers (VCSEL). Detectors include silicon photodiodes and germanium or InGaAs (indium gallium arsenide) photodetectors. Optical cabling is divided into single- and multi-mode types. The former has a smaller diameter and is normally used with a laser source supporting higher bandwidths and longer distances.

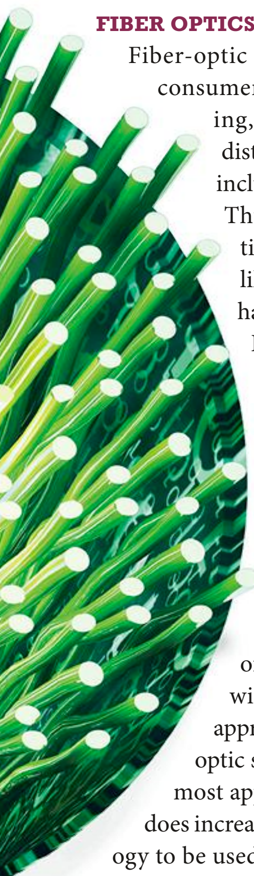

FIBER OPTICS AND CONSUMERS

Fiber-optic cables have tended to be used in specialty consumer applications such as digital video recording, where the connection quality and longer distances offset the cost. The typical connections include high-speed serial interfaces such as USB, Thunderbolt, and PCI Express. Fiber connections are also available for display technologies like DisplayPort and HDMI. In addition, fiber has been used in digital audio applications. S/PDIF (Sony/Philips Digital Interface Format) jacks are common on audio boards and components, including amplifiers and high-definition TVs.

S/PDIF utilizes a passive fiber-optic cable, with the emitters and detectors embedded in the devices. Most other fiber-based systems employ an active cable with electronics at both ends. The typical example is Corning's USB 3.0 cables, which include a pair of fiber-optic links and a pair of copper links with interface logic at both ends *(Fig. 1)*. This approach eliminates the need for specialized fiber optic support at either end of the connection, since most applications assume a copper connection. This does increase the cost of the cable, but allows the technology to be used with existing systems. Typically the copper cables are passive.

Corning is a major supplier of fiber-optic technology; it sells optical USB 3.0 *(Fig. 2)* and Thunderbolt cables in 10, 15-, 30-, and 60-m lengths. Its ClearCurve ZBL (zero bend loss) technology essentially allows the cable to be tied in knots without breaking or signal degradation. Optical cables used to be less robust.

So what exactly is changing? In a word, the Type-C connector *(see "USB 3.1 Type C Connector Is Reversible" on electronicdesign.com)*.

3. Dell's S4810 10-/40-Gb/s Ethernet switch has 48 ports 10 Gb/s SFP+ and four 40-Gb/s QSFP+ uplink ports.

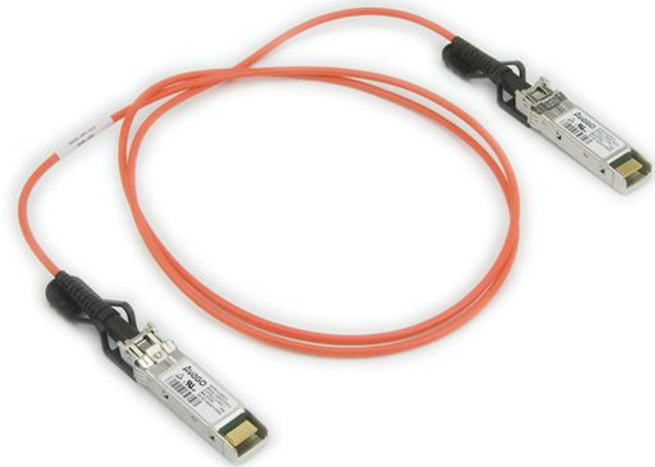

4. Super Microcomputers' 10-Gb/s SFP+ patch cable includes a pair of SFP+ modules with a short fiber-optic cable.

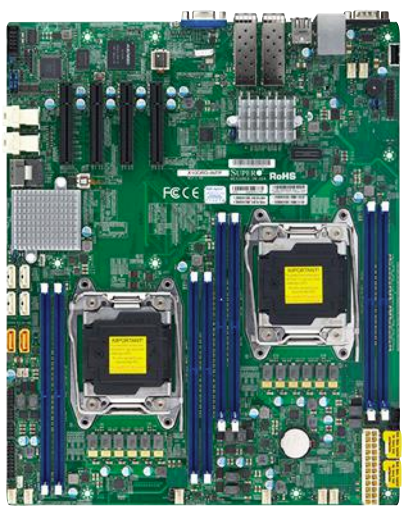

5. SuperMicro's X10DRD-LTP sports a pair of 10-Gb/s SFP+ Ethernet ports.

6. The MXC connector has an array of 64 fibers that can handle up to 1.6 Tbits/s *(Source: Corning)*

7. TE Connectivity's provides MXC cable and connectors with an LC-style latch for easy insertion and removal.

USB 3.1 was the initial driving force behind the reversible Type-C connector. USB 3.1 Gen 2 runs at 10 Gb/s. It is already being tapped for 40 Gb/s for Thunderbolt 3.

The cables are also intelligent, often including redrivers on at least one end of the cable to handle the challenges of using copper connections that are normally under 3 m. This moves copper cable costs closer to optical. Optical cables are still going to cost more, but the mass market should help reduce these costs—especially as higher bandwidths are utilized.

FIBER OPTICS AND THE ENTERPRISE

Networking has been the main use of fiber-optic cable in the past. This is typically Ethernet, but other networking and storage technologies like FibreChannel and InfiniBand support optical connections. Another new architecture is Intel's Omni-Path fabric, which targets the cloud and high-performance computing (HPC). It is designed to scale to tens of thousands of nodes with connection speeds of 100 Gb/s. Speeds continue to move ever higher, and there is even 400-Gb/s Ethernet on the drawing board. Optical connections will be key.

Most of these network adapters and switches embed the transceivers in the hardware or provide a small-form-factor pluggable (SFP) interface. SFP was only the starting point. There is now SFP+ and QSFP+ handling 10-Gb/s and 40-Gb/s interfaces. Dell's S4810 10-/40-Gb/s Ethernet switch *(Fig. 3)* has 48 10-Gb/s SFP+ ports and four, 40-Gb/s QSFP+ uplink ports.

SFP modules include the optical transceivers and optical connections. Modules can support different passive cable connections, including ST, FC, SC, and LC. This allows the cable lengths to be tailored to the installation. Super Microcomputers' (SuperMicro) 10-Gb/s SFP+ patch cable shows a pair of SFP+ modules with a short fiber optic cable *(Fig. 4)*. Unlike Corning's USB 3.0 cables, the SFP+ modules draw power from their host, typically a switch or network adapter board, and they are linked to a pair of passive fiber-optic cables. SFP+ sockets can also be found on motherboards like SuperMicro's X10DRD-LTP *(Fig. 5)*, which sports a pair of 10-Gb/s SFP+ Ethernet ports that handle the networking for a pair of Intel E5-2600 Xeon processors.

Using a pair of optical cables has sufficed for most needs thus far, but higher performance requirements are always on the horizon. This is where the new MXC connector *(Fig. 6)* comes into play. It has 64 fibers, translating into 32 bidirectional links. MXC is designed to use 25-Gb/s VCSEL technology to drive multimode fiber at distances up to 300 m. The connection delivers a total of 1.6 Tb/s of bandwidth.

The connectors are designed to require only seven components instead of more than 25 for existing parallel connections. MXC connectors and cable is available from a number of suppliers, like TE Connectivity *(Fig. 7)*, Corning, Molex, US Conec, and Rosenberger.

8. Pentek's Model 5973 Flexor Virtex-7 Processor and FMC Carrier board supports the optical VITA 66.4 standard.

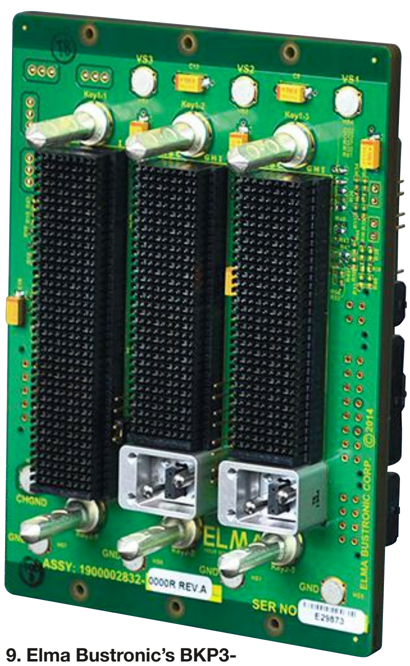

9. Elma Bustronic's BKP3-CEN03 backplane links a pair of VPX VITA 66.4 cards together.

RUGGED FIBER OPTICS

The enterprise is not the only place optical connections are critical. Standards like VITA 66.4 specify rugged board and backplane connections. Boards like Pentek's 3U, Model 5973 Flexor Virtex-7 Processor and FMC Carrier board *(Fig. 8)* supports VITA 66.4. The Xilinx Virtex-7 is linked to the FMC connector and the VPX-P2 VITA 66.4 connector. The latter is done through an optical transceiver.

Pentek's board could be used with Elma Bustronic's BKP3-CEN03 backplane *(Fig. 9)* that links a pair of VPX VITA 66.4 cards together. Elma's motherboard is for development or small-form-factor systems. Usually VPX backplanes are customized for each application and there may be any number of VITA 66.4 connections.

VITA 66.4 may not have the density or throughput of MXC, but it is designed for a much more rugged environment. Still, many of the underlying components are the same or use similar technology, such as the optical transceivers. ed

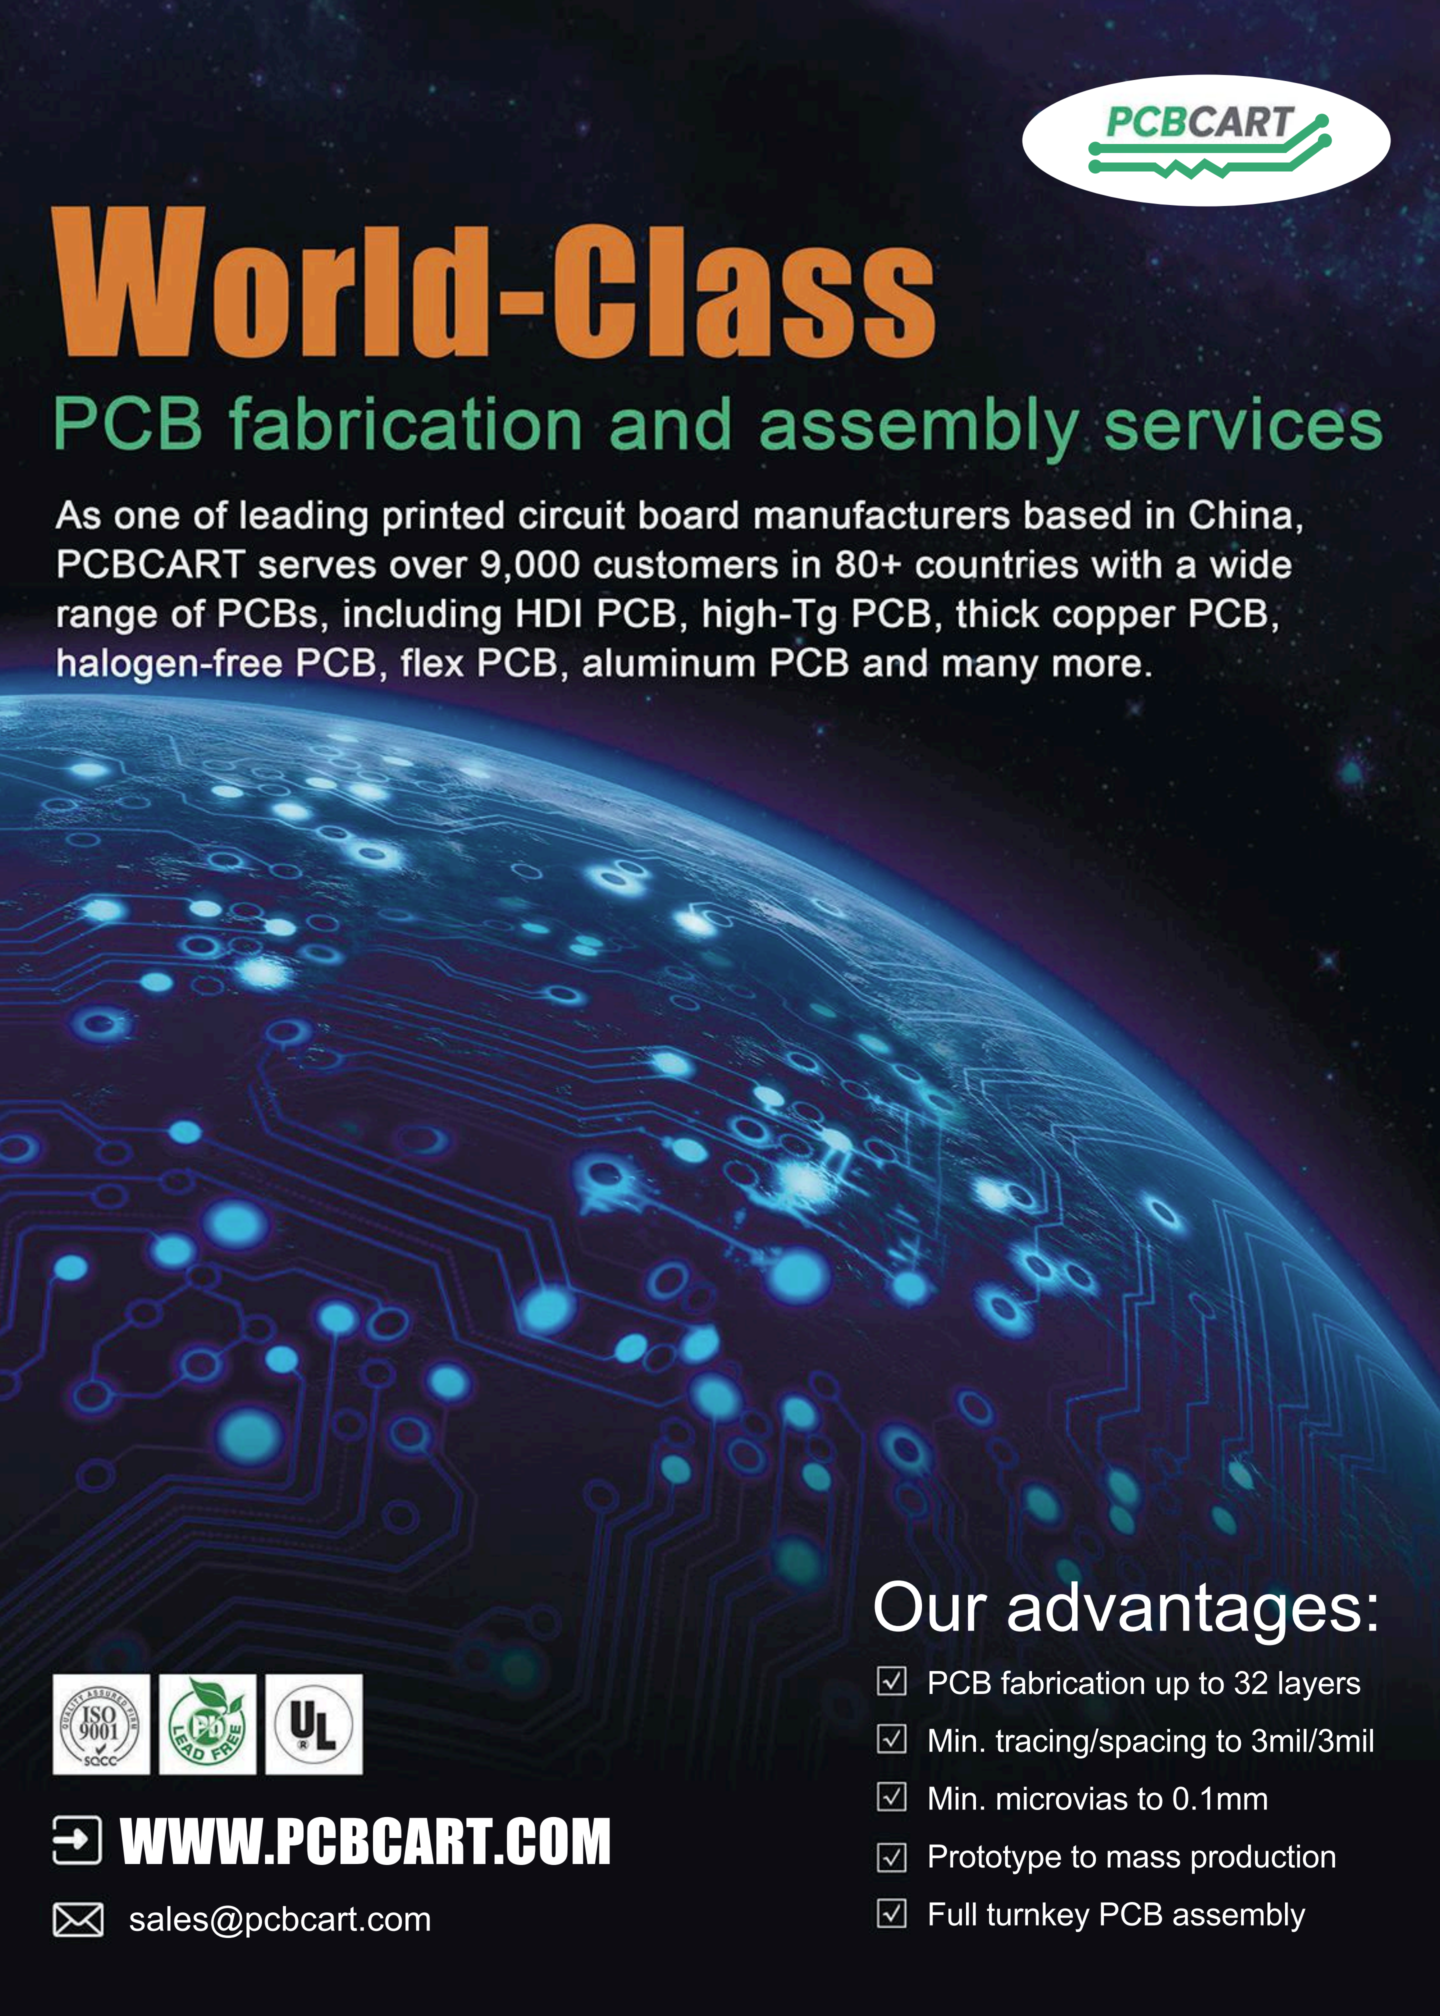
PCBCART
World-Class
PCB fabrication and assembly services
As one of leading printed circuit board manufacturers based in China, PCBCART serves over 9,000 customers in 80+ countries with a wide range of PCBs, including HDI PCB, high-Tg PCB, thick copper PCB, halogen-free PCB, flex PCB, aluminum PCB and many more.
Our advantages:
PCB fabrication up to 32 layers
Min. tracing/spacing to 3mil/3mil
Min. microvias to 0.1mm
Prototype to mass production
Full turnkey PCB assembly
ISO 9001
LEAD FREE
UL
WWW.PCBCART.COM
sales@pcbcart.com

ANTHONY GIOELI | Vice President of Marketing
Synaptics Inc.

What's the Difference Between Match-on-Host and Match-in-Sensor Fingerprint Authentication?

Think all fingerprint authentication is the same? Think again. Match-in-Sensor is emerging as the go-to security technology.

The biometrics space has recently seen a significant spike in using fingerprint authentication as a simple yet secure method for accessing and protecting data. It's also playing a greater role in safe electronic and mobile transactions.

Using fingerprints for user authentication is immeasurably safer and dramatically easier than requiring users to create, remember, and protect passwords, making it a preferred approach of merchants, banks, users, and third-party clearinghouses. A number of technology advances and aggressive innovation by industry leaders has spawned several different forms of fingerprint recognition, though they're by no means uniform.

One common fingerprint-authentication technique is Match-on-Host technology, whereby the fingerprint sensor reads the fingerprint data and sends it for processing by the host processor or other external processor. While the fingerprint sensor captures the fingerprint data, all of the processing and matching work is performed on the host platform. Then there's Match-in-Sensor technology, a purpose-built, fully encapsulated system-on-a-chip (SoC) architecture that isolates fingerprint enrollment, pattern storage, and biometric matching within the actual fingerprint sensor module.

In the true spirit of "What's the Difference," let's take a look at these two very distinct fingerprint-authentication methods.

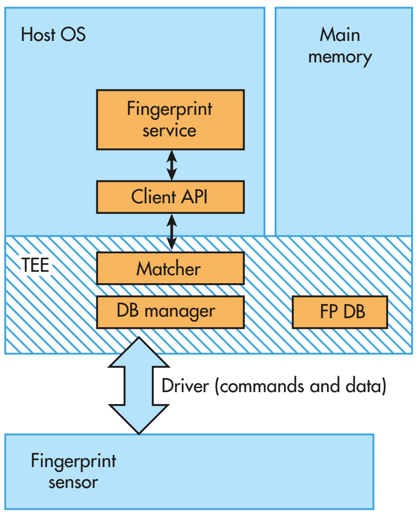

1. Match-on-Host with Trusted Execution Environment (TEE) provides a secure area running within the host environment.

MATCH-ON-HOST: THE CURRENT STANDARD

The fundamental requirement in fingerprint sensing involves positively identifying the user by making a match with a known and secured "template" or record of the user's fingerprint *(Fig. 1)*. The sensor is used initially to capture the data that creates the user's record in an "enrollment" process, and then gets used during every subsequent access attempt to capture fingerprint data to compare with the stored template.

Virtually every implementation of fingerprint sensing today performs the matching process directly on the host system, whether a smartphone, tablet, PC, or dedicated device purpose-built for security. As a result, the Match-on-Host architecture splits the functional requirements between the sensor IC that captures the data and a separate controller IC (often the application processor on a mobile device) used to run the software to make the actual match.

The use of host resources is a natural starting point for any new technology. For this reason, first-generation fingerprint sensors were simple devices limited to a single task: collecting the fingerprint data that would then be used by software running in the host to authenticate the user.

The functions performed in software include identification of fingerprint characteristics, creation of a secure biometric asset (the fingerprint template), storage of the asset, and matching a newly created fingerprint template with the one stored on the device. The host system also provides the security required to protect the integrity and privacy of the fingerprint data. Furthermore, the host system is responsible for detecting biometric forgeries; these so-called anti-spoof techniques are necessary to deter presentation attacks.

Two major selling points of the Match-on-Host architecture have been its low cost and short design-in time, which have enabled fingerprint sensing to be added to devices relatively quickly and cost-efficiently. This momentum, in turn, has facilitated advances made on related fronts, such as the creation of the Universal Authentication Framework (UAF) from the Fast Identity Online (FIDO) Alliance. Nonetheless, for all of its advantages, when it comes to real security, the Match-on-Host approach pales in comparison to the Match-in-Sensor architecture.

MATCH-IN-SENSOR: THE NEXT GENERATION

As the name implies, the Match-in-Sensor architecture integrates the matching and other biometric management functions directly into the sensor IC. The IC contains a high-speed microprocessor, storage for instructions and data, secure communications, and high-performance cryptographic capabilities. To achieve this level of integration, while creating a secure execution environment within the sensor IC, Match-in-Sensor employs a SoC architecture *(Fig. 2)*.

Because integrating multiple functions is the raison d'être of integrated circuits, this advance might not seem worthy of being designated a "next-generation" advance. But the level of additional security afforded by fully integrating the sensing and matching functions is significant enough whereby its expected industry impact should not be understated.

The advanced security with the Match-in-Sensor architecture applies both to the system and to the protection of a user's unique biometric information. System-level security is enhanced with a range of improvements, including:

- Fingerprint data and execution environment of the fingerprint matcher that are physically isolated from the host's operating system, affording immunity from hacks or malware on the host.
- The sensor performing biometric identification autonomously, without reliance on input from the host that might be compromised.
- Input parameters for the matcher that are the live fingerprint information, which is captured, encrypted, processed, and protected on the sensor chip and its enrollment templates.
- Ability to accurately verify authenticity, because the identification result is signed using a sensor-specific private key derived from the hardware.
- Creation, storage, and management of crypto keys that represent the identity credentials being shared—these keys are also used to sign credentials to prevent malware with false information.

Even if the host is completely compromised by a successful attack of any type or origin, it's extremely difficult to force the matcher to generate a false positive result, replay an old result, or in any other way alter or manipulate the match result. This ensures that an identity-authentication subsystem will remain secure even under a worst-case scenario.

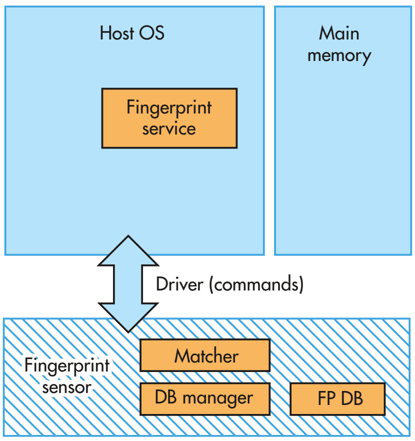

2. Match-in-Sensor solutions isolate fingerprint operations away from the host OS in the sensor itself.

With regard to the user's biometric information, protection is enhanced through a number of features. First, the fingerprint data, including all of the features/characteristics extracted from it and all created templates, is processed only within the sensor's on-chip CPU and storage. None of this information is ever shared or exposed to the host device.

In addition, the enrollment database is located on private flash memory, isolated and physically accessible only by the sensor. Furthermore, the enrollment templates are encrypted and signed by the sensor using proprietary algorithms and strong cryptographic keys before being stored in the private flash memory.

As with the system-level provisions, even if the host becomes completely compromised by a successful attack, the attacker cannot extract any of user's biometric information, avoiding what would arguably be one of the most invasive forms of identity theft imaginable.

CONCLUSION

The support for an ecosystem encompassing fingerprint-authentication technology is steadily expanding, with vendors such as Synaptics developing advanced Match-in-Sensor solutions. These solutions, which are FIDO-certified, include the ability to read fingerprints at various angles, options for visual or haptic feedback, and device- or application-specific optimizations. Moreover, Match-in-Sensor doesn't require transporting or sharing biometric information between the fingerprint module and the host device, thereby eliminating the risk of biometric data being stolen if the system is compromised.

So, if you think all fingerprint-authentication techniques are the same, you may want to take a good, long look at the Match-in-Sensor architecture and its capabilities. ed

ANTHONY GIOELI is vice president of marketing for Synaptics Inc. He has more than 25 years of experience in corporate management, strategy, sales, marketing, business development, and operations in the wireless, semiconductor, IP telephony, biometrics, and software markets. He holds a BSEE from the New York Institute of Technology and an MBA from the University of Southern California Marshall School of Business.

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

PLAY BALL with IoT Indoor Positioning Sensors

Sensors and the Internet of Things are almost synonymous these days. Find out how some are improving basketball.

The Internet of Things (IoT) may be hype to some, but there is a host of products using wireless communication and various sensors to create some interesting IoT-related products. Typically the sensors include things like accelerometers—magnetic and gyroscopic sensors also found in smartphones—but many other types of sensors are available. DecaWave's indoor positioning system is one of these.

DecaWave's ultra-wide band position sensing technology is based on the IEEE802.15.4-2011 standard. The technology is employed in ShotTracker's Shot Tracker (*Fig. 1*). The start kit comes with a a small tag on each player's shoe and the ShotTracker-enabled Spalding basketball connects with portable sensors placed around the court. The software generates a real-time shot chart for each player and automatically tracks other metrics such as possession, passes, assists and turnovers. This platform is designed for individual use. The wrist sensor slips into a wrist band. The position information is relative to the wrist sensor and the net sensor. The latter communicates with the smartphone application. The net sensor also has a three-axis accelerometer allowing it to determine rim bounces and baskets.

1. The Shot Tracker starter kit includes a small tag on each player's shoe and the ShotTracker-enabled Spalding basketball connects with portable sensors placed around the court.

DecaWave's DW1000 is another single-chip indoor-positioning solution (IPS). It can provide position information to within 10 cm.

The team demo at the 2016 Consumer Electronics Show (CES) was a bit more ambitious and used a different set of sensors but the same basic technology. The basketball court was ringed with sensors and a sensor is built into the basketballs used on the court. This allowed the ball to be tracked.

2. The ShotTracker demo at CES 2016 included a basketball court with multiple sensors providing real-time player feedback on a large display screen.

College basketball players typically shoot over 1,000 shots per day in practice. The information provided just by ShotTrack positioning sensors lets players and coaches see how a player is progressing and where improvements can be made.

DecaWave's technology provides accuracy measured in centimeters while the Bluetooth positioning standard is on the order of meters. Bluetooth is more useful for general proximity and Bluetooth beacon technology is being employed in retail advertisements.

Positioning and context information is becoming more important and accessible to designers. There are a number of applications, from factory automation to smart buildings, that can take advantage of IPS and IoT. ed

Simple, Easy Solutions

High Accuracy, High Speed Position Sensors!

MagAlpha Series

Eliminates Optical Encoders!

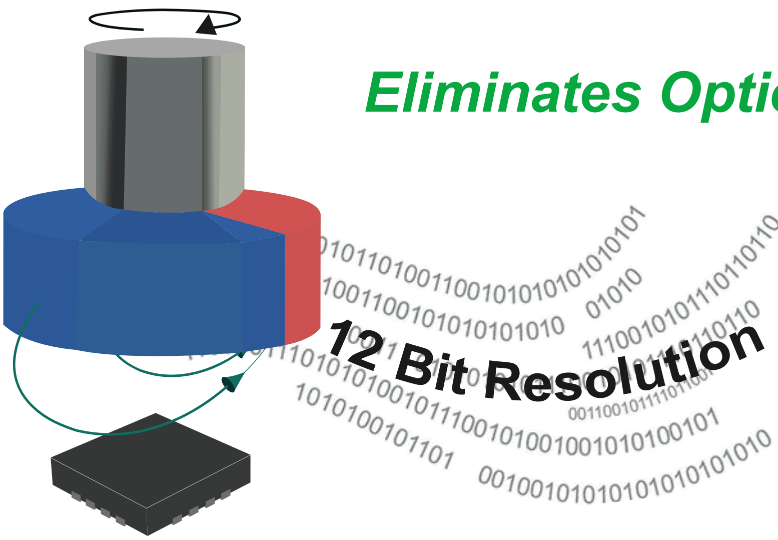

Flexible Mounting

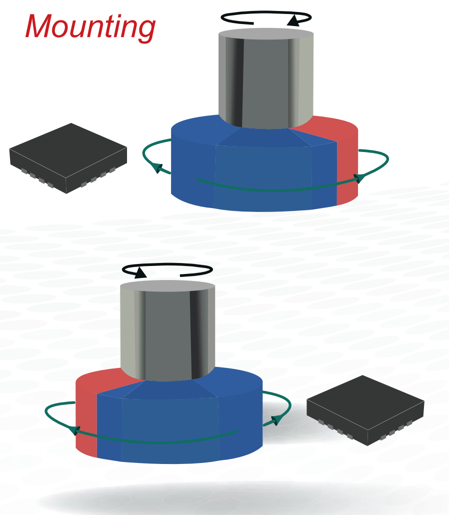

MagAlpha MA300 Features

- UVW Signals for Block Commutation
- 12-Bit Resolution Absolute Angle Encoder
- 500kHz Refresh Rate
- Ultra-Low Latency: 3μs
- Serial Interface for Data Readout and Settings
- 10-Bit Incremental Output (A, B, Z)
- Built-In Linearization for Side-Shaft Mounting
- 7.5mA Supply Current

www.monolithicpower.com

Follow us on:

BILL WONG | Technical Editor

What's the Difference Between M.2 Modules?

What sets different M.2 storage modules apart? It's more than just the capacity.

Using the M.2 interface to connect high-speed modules to a motherboard has become something of a trend. This space is currently dominated by flash-memory M.2 cards of varying capacities, speeds, and interfaces, although it can handle devices such as wireless adapters. The latter is a primary differentiation between modules as the interface that is used. These include SATA 3.0, PCI Express 3.0 (PCIe), and USB 2.0/3.0 using HSIC and SSIC support. Other interfaces can be provided, including audio, UIM, I^2C, and SMBus.

The M.2 modules are 22-mm wide and come in a number of lengths, including 30, 42, 60, 80 and 110 mm. Motherboards can usually handle different-length M.2 modules, but 80 mm is the typical size for a flash-memory module. Wireless adapters are often found on shorter modules. A motherboard does not have to support all lengths.

While a motherboard can incorporate a range of key M.2 socket options, the most popular are B and M *(Fig. 1)*. The B interface provides x2 PCIe, SATA, USB 2.0 and 3.0, audio, UIM, HSIC, SSIC, I^2C, and SMBus support. The M interface provides x4 PCIe, SATA, and SMBus. The sockets are keyed so a module cannot be plugged into a socket that does not support the module's interface. M.2 modules are keyed to fit sockets that will support the interface used by the module.

SATA modules, like SanDisk's X400 *(Fig. 2)*, normally have a B&M edge connector, allowing them to be plugged into a B or M socket. The 80-mm-long X400 holds up to 1 TB of flash memory. It is a single-side M.2 module that is only 1.5 mm high. The 256-GB version is rated at 40 GB/day for five years. It employs SanDisk's sixth-generation X3 Technology.

M.2 modules with a x4 PCIe interface must have a M connector. For example, Samsung's 80-mm-long, 256-GB SM951 M.2 module can be plugged into Super Microcomputer's (Supermicro) C7Z170-SQ gaming motherboard *(Fig. 3)*. Its M-style socket only handles PCIe/NVMe devices *(see "What's The Difference Between SATA And NVMe?" on electronicdesign.com)* like the SM951. Some motherboards, like Gigabytes' GA-170X, have M.2 sockets that can handle an B&M-style M.2 module (like the X400) or a PCIe module (like the SM951).

The main differences between the x4 PCIe/NVMe interface and the SATA interface are bandwidth and overhead. A single PCIe 3.0 lane can handle a 6-Gb/s SATA 3.0 interface. There is also less overhead with a direct PCIe interface. A higher-speed interface with lower overhead only makes a difference

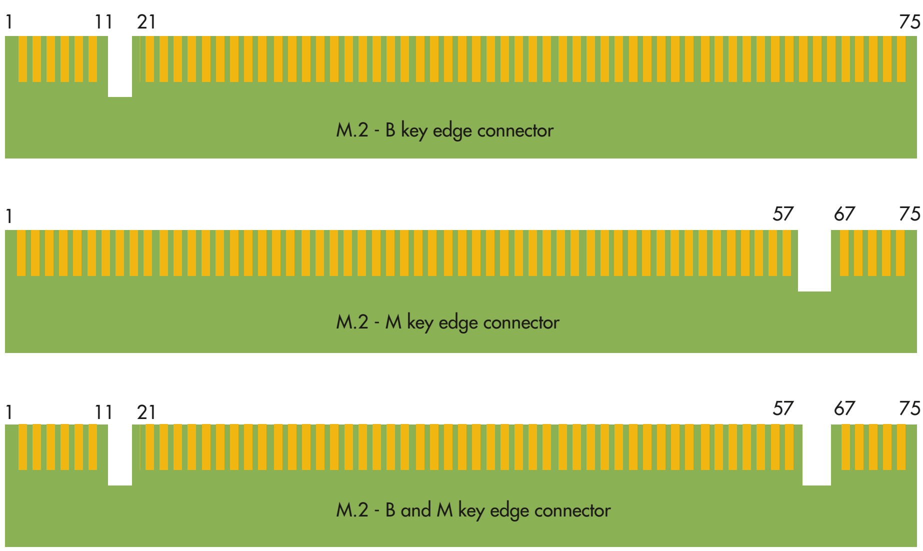

1. The B and M module can plug into a B or M socket, but a B or M module can only plug into a matching B or M socket.

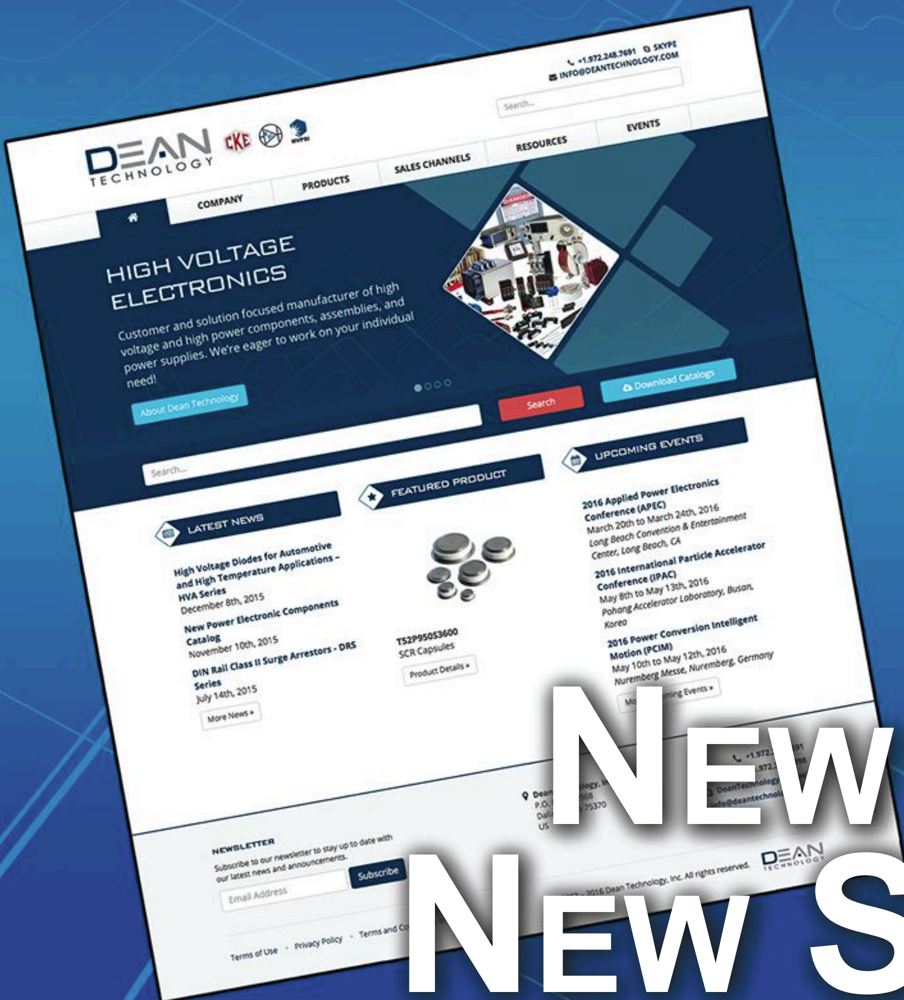

NEW SITE NEW START

INTRODUCING THE NEW WWW.DEANTECHNOLOGY.COM

In a world that demands more information, faster, we're introducing a new site that exceeds those expectations. Built on top of a comprehensive product database, and using the most contemporary technology, our new site offers simple navigation to the most accurate and constant product information in the industry.

And that's just the first step… we're reinventing the delivery of electronics product information while maintaining our commitment to world class engineering and customer support. It is now easier than ever to find exactly what is needed for your high voltage or high power design.

Visit us today and discover how easy it should be.

HVPSI

PRODUCTS BY:
DEAN TECHNOLOGY

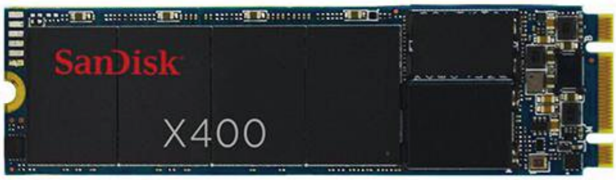

2. SanDisk's X400 can hold up to 1 TB of storage.

if the storage performance is comparable. Of course, flash memory performance varies, but modules like the SM951 can take advantage of a x4 PCIe interface. Its sequential read performance is 2150 MB/s and it can deliver 300K random read IOPS. This performance typically comes at a higher cost than a SATA M.2 module.

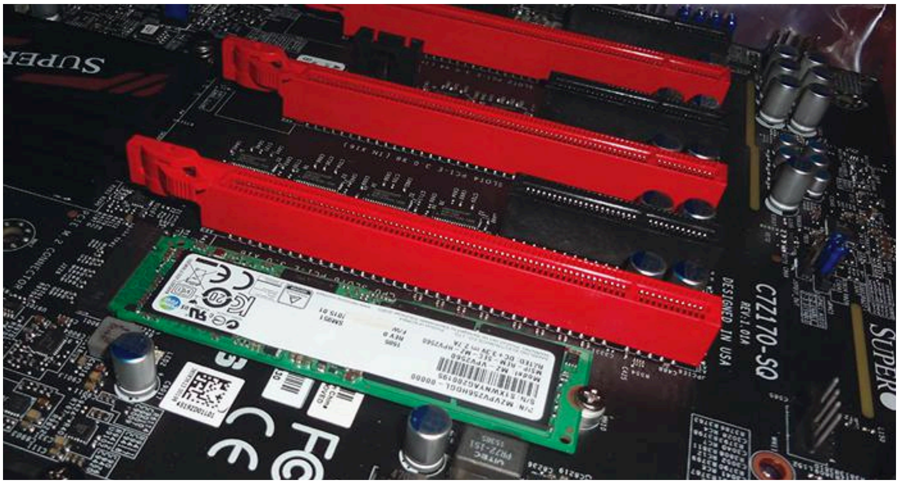

3. Samsung's 256-GB SM951 M.2 module is plugged into Super Microcomputer's C7Z170-SQ gaming motherboard.

Device drivers typically insulate applications from the underlying hardware. This allows designers to choose the appropriate hardware for the application and processor, balancing performance, capacity, and cost.

One or more M.2 sockets can be found on new motherboards. An M.2 module can be the primary storage device, eliminating the need for move conventional hard disk or SSD drives. Intel's compact NUC form factor *(see "What's The Difference Between Mini-ITX And Intel's NUC Platform?" on electronicdesign.com)* is designed to use M.2 internal storage exclusively. ed

BENJAMIN M. BROSGOL, Ph.D. | Senior Technical Staff, AdaCore
www.adacore.com

Comparing Ada and C

Both languages approach the reliability versus efficiency tradeoff from different angles, but each has a place in embedded-systems programming.

This issue is focused on the Internet of Things and the security issues that arise when interconnecting billions of devices, ranging from coffee makers to power grids. This article looks at the subject from a specific and rather basic perspective: Which language(s) should you choose to develop the software, where "software" means both the embedded code that runs the Things and the system programs that manage the networks, etc.? Choice of language is important since it affects the system's reliability, security, and performance, as well as the ease or difficulty in adapting the software as requirements change.

More specifically, this article compares C and Ada, summarizing their strengths and weaknesses and suggesting when to use (or not use) each. These two languages are interesting to look at: C because it's often the default choice for real-time and systems programming, and Ada because it has a successful (but not as well known) record in these same areas.

C and Ada have gone through various updates since their inception. I'll use the most recent version of each—C 11[1] and Ada 2012[2]—as the basis for the comparison. These reflect how the languages are evolving to meet current and future technological trends and challenges, even though at present it's more typical to find earlier versions of the languages in use.

C

In any kind of assessment, it always helps to go back to first principles. What were the main design goals for each language? The introduction to the 1999 version of the C standard[3] distilled the "spirit" of C into a small set of objectives, which have guided and constrained both the original design and each revision:

- Trust the programmer.
- Don't prevent the programmer from doing what needs to be done.
- Keep the language small and simple.
- Provide only one way to do an operation.
- Make it fast, even if it's not guaranteed to be portable.

In keeping with these principles, C offers various data types and data-structuring facilities (arrays, structs, pointers, unions, enums) with straightforward and efficient implementation, conventional algorithmic features (statements, expressions, functions), and modest modularization mechanisms (header files with function prototypes, #include directive, preprocessor).

Standard header files support dynamic memory management (malloc, free), a minimal exception mechanism (setjmp, longjmp), string handling, numerics, input/output, internationalization/locales, operating-system interfacing, and other services. Standard (but optional) and C++ compatible support for concurrent programming, including features that help exploit multicore platforms, have been introduced in C11. It specifies an explicit memory model, and supplies low-level facilities for thread management and communication.

By intent, C has some significant omissions. It doesn't provide generic templates (which can be approximated in part by the preprocessor), programmer-defined operator/function overloading, or object orientation, and its encapsulation support ("information hiding") is rudimentary.

In short, C is very much a WYSISWYG ("What You See Is What You Get") language. When you write a C program, you have a good idea of what the resulting compiled code and data will look like. Thus, C becomes a typical choice for low-level software that needs to interact directly with the hardware. However, a simple WYSIWYG language has two major drawbacks:

- It doesn't easily scale up to very large systems.
- In its focus on efficiency, it can sacrifice checks that are useful or necessary for reliability, safety, or security.

To somewhat address the latter point, "safe" subsets of C have been proposed over the years. Perhaps the best-known is MISRA C,[4] originally intended for automotive software but applicable to other domains as well. Static-analysis tools such as lint and a variety of commercial products have been used to detect potential vulnerabilities, although the language's weak type checking makes this more difficult than for other languages. And various guidelines have been published to facilitate secure coding.[5]

C11 has attempted to address some of the security issues via language features and libraries. For example, the optional Annex K (Bounds-checking interfaces) provides alternative versions of various standard functions, thus helping to prevent certain forms of buffer overflow as well as other vulnerabilities. The optional Annex L (Analyzability) constrains some forms of undefined behavior to be bounded, with the requirement that the implementation not perform an out-of-bounds store.

Will these new features be widely implemented, and will programmers use them? Time will tell. But in my opinion, they look like a patch that may mitigate some vulnerabilities but doesn't alter the original language philosophy. C wasn't designed for programming large-scale high-integrity applications. It's often selected not based on fitness for purpose, but because programmers know it (or it fits smoothly into an organization's software-development infrastructure), or because of perceived inefficiencies in other technologies.

ADA

Ada is very much at the other end of the spectrum. Perhaps a variation of C's principles serves as a first approximation to the "spirit" of Ada:

- Trust the programmer, but verify through appropriate checking since programmers are human and make mistakes.
- Prevent the programmer from doing what shouldn't be done.
- Keep the language kernel small and simple, but provide extension mechanisms in order to increase expressiveness.
- Provide one principal and intuitive

AN ADA PACKAGE

THIS CODELIST ILLUSTRATES a simple Ada package. The package specification, on the top, defines the Vector type as an array of Integer values. Different objects of this type can have different bounds. The Max function returns the maximum value in its parameter V. Its precondition is that V contains at least one element. Its post-condition captures the function's required semantics—the returned value has to be at least as large as every element in V, and it must be an element of V. The Negate procedure performs the unary "-" operation on each element in its parameter V. Its pre-condition (tao avoid overflow) is that no element can be the smallest Integer value. Its post-condition captures the procedure's semantics; V'Old is the value of V at the point of call. The contracts shown are appropriate for use with formal methods, so that they're verified statically, or they could be enabled as run-time checks to support debugging.

The package body contains the implementation of the two subprograms. V'First is the index of the lower bound of V, and V'Last is the index of the upper bound. The "for" loop in Negate illustrates the ability to iterate over a collection (here an array) without explicitly indexing. Note that Ada uses ":=" for assignment, "=" for equality, and "/=" for inequality.

```
package Math_Utilities is
  type Vector is array(Positive range <>) of Integer;

  function Max(V : Vector) return Integer
  with
    Pre  => V'Length>0,  -- V cannot be empty
    Post => (for all  Element of V => Max'Result >= Element) and
            (for some Element of V => Max'Result = Element);

  procedure Negate(V : in out Vector)
  with
    Pre =>  (for all Element of V => Element /= Integer'First),
    Post => (for all I in V'First .. V'Last => V(I) = -V'Old(I));
end Math_Utilities;
```

```
package body Math_Utilities is
  function Max(V : Vector) return Integer is
    Current_Max : Integer := V(V'First);
  begin
    for I in V'First+1 .. V'Last loop
      if V(I) > Current_Max then
         Current_Max := V(I);
      end if;
    end loop;
    return Current_Max;
  end Max;

  procedure Negate(V : in out Vector) is
  begin
    for Element of V loop
       Element := -Element;
    end loop;
  end Negate;
end Math_Utilities;
```

way to do an operation.

- Make it reliable and portable, and depend on the compiler to produce efficient code.

More generally, Ada's main goals were succinctly specified in the introduction to the first version of the language standard:

"Ada was designed with three overriding concerns: program reliability and maintenance, programming as a human activity, and efficiency."

More specifically, Ada was designed from the outset to take advantage of the breakthroughs in software engineering and programming methodology that occurred in the 1970s, with a focus on support for embedded real-time applications. The emphasis was on achieving confidence in program reliability (correctness), through features that include checks either statically or at run time.

Ada is a strongly typed extensible language, with facilities to define new types in various categories: integers, floating point, fixed point, enumeration, arrays, records (structs), and access types (pointers). Unlike C, Ada allows the definition of constrained subranges of scalar values, and checks ensure that objects aren't assigned out-of-range values. Subrange information is very useful to human readers and static-analysis tools.

Ada includes traditional algorithmic features, with a simple set of statements and with code modularization through subprograms (functions). It also has facilities for "programming in the large": encapsulation/data abstraction, separate compilation, packages (somewhat analogous to C header and code files), subprogram and operator overloading, generic templates, and full support for object-oriented programming (OOP). Ada also includes built-in features for exception handling and concurrency, including a structured feature for state-based mutual exclusion that helps avoid race conditions.

The predefined environment of Ada includes packages for character and string handling, I/O, numerics, containers, and operating-system interfaces. Ada also defines an annex with standard support for interfacing with other languages (including C), and optional specialized-needs annexes covering systems programming; real-time, distributed, and information systems; numerics; and safety and security (high-integrity systems).

Ada 2012 introduced contract-based programming features (pre- and post-conditions for subprograms, invariants for encapsulated types). This significant enhancement in effect embeds low-level requirements into the source code, with checks performed either at run time or (with appropriate tool support) statically. The Ada example *(see "An Ada Package")* illustrates the use of pre- and post-conditions; an analog example in C is shown in "C Header and Code File." Ada 2012 also increased the language's multiprocessor/multicore support and added a number of other enhancements.

Ada was intended for embedded systems, and programming at that level may involve getting down-and-dirty with the hardware—writing interrupt service routines, dealing with machine addresses and data representations, handling endianness issues, etc. With Ada, programmers can do all those things—one goal of the Systems Programming Annex is to give the programmer the tools to do anything in Ada that's possible in assembly language.

All of this sounds like a large and complex language. Indeed, the inclusion of generics, OOP, and exceptions makes Ada quite a bit more sizable than C, although subtleties in features

C HEADER AND CODE FILES

THE C HEADER and code files correspond to the Ada package *(see "An Ada Package")*. The pre-condition for max is modeled by an assert statement in the function body. The other Ada contracts are omitted, since C doesn't have quantification expressions.

One of the semantic differences between Ada and C concerns the treatment of array bounds. In Ada, the bounds are accessible through the array object via V'First and V'Last, while in C, the array size needs to be supplied as an explicit parameter to the functions.

```c
// math_utilities.h

typedef int vector[];

int max(vector v, int n);
// n is the number of elements in v

void negate(vector v, int n);
// n is the number of elements in v
```

```c
#include <assert.h>
#include "math_utilities.h"

int max(vector v, int n)
{
    assert(n>0);
    int curmax = v[0];
    for (int i=1; i<n; i++){
        if (curmax < v[i]){
           curmax = v[i];
        }
    }
    return curmax;
}

void negate(vector v, int n)
{
    for (int i=0; i<n; i++){
        v[i] = -v[i];
    }
}
```

LINEAR TECHNOLOGY

DESIGN NOTES

VID Controller with Phase Extender Provides 120A Power Supply for Latest Generation of FPGAs, ASICs and Processors

Design Note 547

Mike Shriver

Introduction

The current drawn by FPGAs, ASICs and processors for high performance servers, network and computing systems continues to rise, and load currents of 100A or greater are becoming common. Meanwhile, the chip's operating voltages are dropping to 0.9V and below with tighter voltage regulation requirements. For many of these applications, the core voltage may need to be adjusted for optimum performance with a VID (voltage identification) interface. A significant challenge is clearly placed on the power supply designer to meet the demands of high efficiency and tight output voltage regulation with a small amount of board space.

One way to meet these demands is to use the LTC3877 and LTC3874 chipset. The LTC3877 is a peak current mode, VID-controlled dual output synchronous step-down controller. The output of phase 1 can be programmed from 0.6V to 1.23V in 10mV increments with a six-bit parallel VID interface. Phase 2 provides an output of 0.6V to 5V, which is set by an external divider. The two phases can be paralleled together or with phases from another LTC3877 or an LTC3874 for higher output current.

The LTC3874 is a peak current mode phase extender chip. It does not have an error amplifier and instead regulates its phase current to the ITH signal from the LTC3877 master. The elegant design of the LTC3874 reduces trace count and board space. The LTC3877 comes in a 44-lead, 7mm × 7mm QFN package; the LTC3874 slave controller comes in a 28-lead, 4mm × 5mm QFN package.

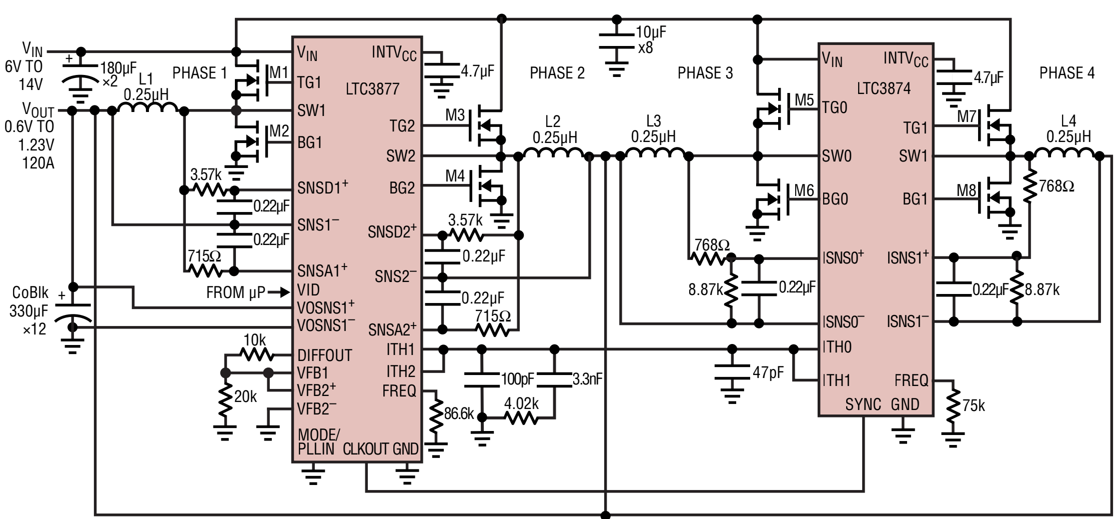

NOTE:
THIS IS A SIMPLIFIED SCHEMATIC. REFER TO THE DATA SHEET FOR A COMPLETE SCHEMATIC.

L1–L4: WURTH 744301025 (0.25µH, 0.32mΩ)
M1, M3, M5, M7: INFINEON BSC050NE2LS
M2, M4, M6, M8: INFINEON BSC010NE2LSI

Figure 1. 4-Phase 120A VID-Controlled Converter Operating at a Switching Frequency of 400kHz

PolyPhase Design with High Accuracy and Efficiency

The 4-phase step-down converter shown in Figure 1 uses the LTC3877 and LTC3874 to provide a VID-controlled output of 0.6V to 1.23V at a maximum load current of 120A at a switching frequency of 400kHz. The LTC3877 yields a total DC regulation accuracy of ±1% for all VID set points over temperature. The differential remote sense amplifier in the LTC3877 senses the output voltage at the regulation point and compensates for voltage drops across PCB trace runs and ground planes. The 4-phase operation results in lower output voltage ripple and faster load step response due to shorter clock delays.

High efficiency is a result of the strong gate drivers and short dead times of the two chips, MOSFET selection and sub-mΩ DCR ferrite inductors. The full load efficiency for a 1.2V output at 120A load is 88.8% as shown in Figure 2.

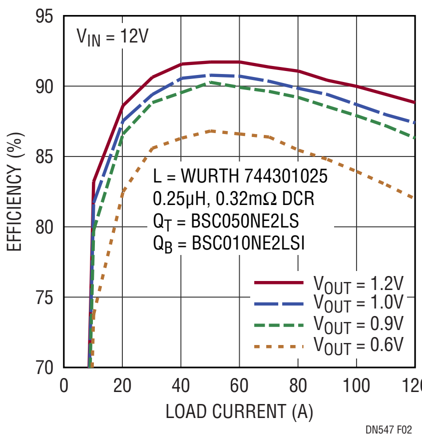

Figure 2. Efficiency of the 4-Phase 120A VID Converter

Sub-mΩ DCR Sensing

Both the LTC3877 and LTC3874 use a proprietary DCR current sensing architecture designed for sub-mΩ DCR sensing, which ensures tight control of the current sharing and current limit. Figure 3 shows the current sharing performance of the 4-phase converter of Figure 1. The inductor used is the Wurth 744301025 (250nH), which has a DCR of 0.32mΩ. The current sharing error is less than 1mV between phases.

More Features

The LTC3877 and LTC3874 both have a phase-lockable frequency range of 250kHz to 1MHz and a FREQ pin

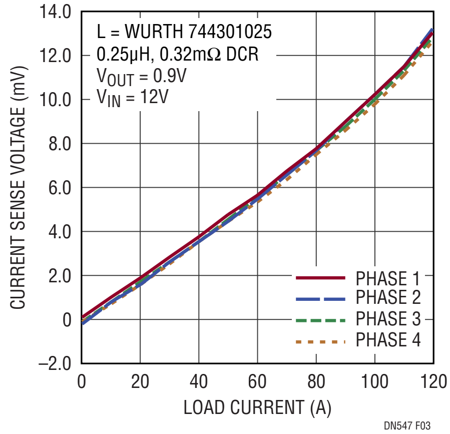

Figure 3. Current Sharing of the 4-Phase 120A VID-Controlled Converter

to set the internal frequency if synchronization is not required. The LTC3877 offers three light load operating modes: Burst Mode® operation, forced continuous mode and pulse-skipping mode. The LTC3874 operates in either forced continuous mode or pulse-skipping mode.

The minimum on-time of the LTC3877 is 40nsec typical, ideal for high step-down converters or small footprint or high bandwidth converters operating at switching frequencies of 500kHz to 1MHz. The minimum on-time of the LTC3874 is 90nsec, typical.

Phase 1 of the LTC3877 provides VID control. If the FPGA, ASIC or processor is not awake or VID programming is not required, then the VID section can be disabled by pulling the VIDEN pin low and setting the output voltage with a divider at the output of the differential amplifier. Both phases of the LTC3877 have differential remote sense amplifiers for precise control of the output voltage. The input voltage range of both chips is 4.5V to 38V.

Other features of the LTC3877 include PGOOD pins, RUN pins and TK/SS pins for each rail. The LTC3874 has its own RUN pins and $\overline{\text{FAULT}}$ pins for quick response to fault conditions.

Conclusion

The LTC3877/LTC3874 chipset provides the power system designer with a highly accurate, efficient and robust PolyPhase® solution for high current rails supplying FPGAs, ASICs and processors.

Data Sheet Download

www.linear.com/LTC3877

For applications help,
call (408) 432-1900, Ext. 3270

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

DN547f LT/AP 0216 111K • PRINTED IN THE USA

such as sequence points don't make C the simple language as is commonly advertised. Skeptics might jest that, while a C program is WYSIWYG, Ada code seems more in the WTF category (acronym intentionally left unexpanded).

QUESTIONS SURROUNDING ADA

Doesn't Ada have some performance challenges? And if Ada is supposed to be used for safety-critical or high-security systems, doesn't the semantic complexity get in the way? How do you certify a system where you need to show traceability from requirements down to object code, or where the implementation's run-time libraries are subject to the same certification requirements as the application software?

These are fair questions. Ada, like any other general-purpose language intended for high-integrity systems, needs to be constrained to a safe subset, only including features with well-defined behavior and a simple (certifiable) implementation. Ada actually anticipated this issue and supplies a compiler directive (pragma Restrictions) that allows programmers to specify features that will not be used. If using such a feature, then the error is detected, generally at compile time but sometimes at run time.

The Ravenscar tasking profile,[6] a set of Ada concurrency features with a small footprint and simple implementation, is part of the Ada standard and is defined through pragma Restrictions. Implementations can supply one or more restricted run-time profiles, corresponding to subsets at different levels of generality (and thus different levels of effort needed for certification).

Another notable example of an Ada subset is the SPARK language.[7] SPARK 2014, an Ada 2012 subset, is designed to facilitate formal proofs of program properties ranging from absence of run-time errors to compliance with a formally specified set of requirements. SPARK eliminates features that are hard to verify, such as pointers, but includes most of Ada's static semantics. Projects like the NSA-sponsored Tokeneer effort[8] demonstrated that ultra-high reliability and security is achievable with formal methods using conventional verification techniques.

CONCLUSIONS

C's emphasis has always been on performance, and its benefits show up most clearly when this requirement is critical (for example, in a software product for a competitive commercial market, where a customer's purchase decision may be strongly influenced by benchmarks). When reliability, safety, and/or security are overriding requirements, C has well-known defects.

Historically, many security holes have been caused by writing past the end of an array, a bug that's detected in Ada. Some can be overcome with external tools (to enforce a "safe" subset or to detect vulnerabilities), or with the help of the new C11 features. However, the language wasn't designed with support for high-assurance systems as a major goal.

Ada's emphasis has always been on the various "ilities" (reliability, readability, maintainability), and its benefits show up most clearly when these requirements are critical (for example in a large, long-lived system where total software lifecycle costs need to be taken into account). Indeed, Ada (and safe subsets such as SPARK) has a long and successful usage history in safety-critical and high-security applications.

So when should Ada not be used? One context is when the need arises for rapid prototyping or scripting. Consider, instead, a dynamically typed language such as Python. Another scenario is when quickness to market is an important goal; then a higher software defect rate may be an acceptable price to pay.

How about when run-time performance (time, space) means the difference between a successful product and an also-ran? It's certainly possible to obtain efficient code from Ada, and indeed technologies such as gcc,[9] which incorporate a common code generator for multiple languages, yields the same performance for Ada and C on language constructs that have the same semantics. You can also improve efficiency by avoiding complex features, or by suppressing run-time checks after verifying through static analysis or sufficient testing that the checks will not fail.

Note that Ada versus C is not an "either/or" decision. They actually get along well together, largely due to Ada's standard interfacing support. An Ada program can import functions or global data from C, lay out data structures to have the same representation as the corresponding C data, and export subprograms or global data for use by an external C function. Therefore, a C program can be extended with functionality provided by Ada, and symmetrically, an Ada program can invoke C services. ed

DR. BENJAMIN BROSGOL, a senior member of the technical staff of AdaCore, has been involved with programming language design and implementation for more than 30 years. He was a Distinguished Reviewer of the original Ada language specification and a member of the design team for the Ada 95 revision.

REFERENCES:
1. ISO/IEC 9899:2011. Information technology -- Programming languages -- C.
2. Ada Reference Manual; ISO/IEC 8652:2012(E); Language and Standard Libraries. Available from www.adaic.org/ada-resources/standards/ada12/.
3. The C Standard Incorporating Technical Corrigendum 1; John Wiley & Sons; 1999.
4. MISRA C:2012 – Guidelines for the Use of the C Language in Critical Systems; www.misra-c.com.
5. Robert C. Seacord, Secure Coding in C and C++, 2nd Edition; Addison Wesley; 2013.
6. ISO/IEC TR 24718:2004. Guide for the use of the Ada Ravenscar profile in high integrity systems (2004).
7. SPARK 2014; www.spark-2014.org/.
8. Tokeneer ID Station Public Release Archive. Available from www.adacore.com/sparkpro/tokeneer.
9. GCC, the GNU Compiler Collection; gcc.gnu.org.

ANIRBAN SENGUPTA | Pricing Manager, Cypress Semiconductor

Knowledge is Power with USB Power Delivery

USB power delivery has changed substantially since USB was first released. The latest standard provides much more power and exchange methodologies.

Since its early days, USB has been positioned as an external bus standard for fast data communications. When observing how the standard has evolved, one can notice a distinct bias toward increasing the speed of data communications *(Fig. 1)*. Moreover, each generation has been branded using terms such as "full-speed," "Hi-speed," "Super-speed," and "super-speed+," which again reemphasized USB as a communications protocol.

However, another area of evolution for USB was sidelined earlier in its development—power delivery. With the focus on speed, power delivery was never the priority for USB. In fact, the first time a specification was made specifically for power delivery occurred only a few years ago, in 2010: The Battery Charging Specification – BC 1.2 increased USB power delivery from 4.5 W to 7.5 W. This specification introduced a mode called CDP (charging downstream port) that allowed for higher charging current (up to 1.5 A) compared to traditional USB.

However, this specification was only a precursor for what was to come. Almost at the same time as the USB 3.1 release, the USB Implementers Forum (USB-IF) released the USB Power Delivery Specification. This new USB power delivery specification allows power transfers of up to 100 W, along with other features and benefits.

A SECONDARY FUNCTION GAINS PROMINENCE

Power management has been part of USB specifications since early generations. Power management in these initial generations was not intended for battery charging, but rather to enable peripherals to power up. In line with this functionality, USB 1.0 described a power source of 5 V @ 100 mA (0.5 W), while USB 2.0 described a power source of 5 V @ 500 mA (2.5 W).

Despite there being no such guidelines in the specifications, designers figured out ways to use USB for battery charging. As a result, USB-based charging products were created and launched in the market. However, due to the lack of a specification, interoperability was a major issue for first-generation USB battery chargers. This interoperability issue, coupled with the rising popularity of battery-charging applications, gave the USB-IF enough reason to consider creating specifications for battery charging. This resulted in the Battery Charging Spec – BC 1.1 (later revised to BC 1.2).

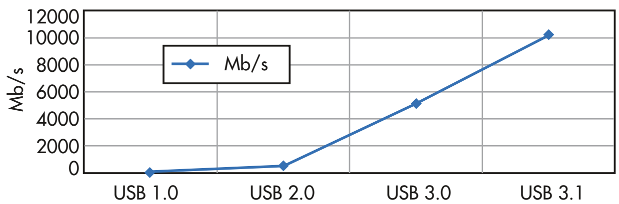

1. USB has always been positioned as an external bus for fast data communications, and has evolved in this way over the years.

It's important to note that BC-1.1 was released as an Engineering Change Notice (ECN) to USB 2.0 and it significantly deviated from the sanctions of USB 2.0. As per USB 2.0, any USB device could be classified as either low power (5 V @ 100 mA) or high power (5 V @ 500 mA). On connection, a USB device was allowed to draw 100-mA current initially while enumerating and negotiating its power budget with the host. Based on the enumeration, the host would either raise the power delivery to 2.5 W or continue at 0.5 W.

The battery-charging spec went on to define more power sources than what was recommended above:

1. Standard downstream port (SDP): Power source compliant with USB 2.0 Spec.
2. CDP: Power source not compliant to USB 2.0. CDP can supply up to 7.5 W (5 V, 1.5 A), and the 1.5-A current can be supplied before enumeration.

While both SDP and CDP were located on typical hosts like desktops and notebooks, the third type of port defined by the BC spec included power sources like wall and auto adapters:

3. Dedicated charging port (DCP): There's no enumeration here, and charging occurs without any digital connection.

DCP supplies up to 1.5 A and 5 V:

While the BC spec did solve quite a few issues, other areas of improvement remained open. Some of the most notable ones were:

- Which port to use? Not all OEMs correctly labeled the CDP and SDP. As a result, the user experience was that some ports charged connected devices quickly while other ports didn't charge at all.
- 7.5 W is not enough: For quite a few devices that could potentially draw power from a USB source, the 7.5-W limit didn't improve things much. Examples include hard drives and external drives.
- Power flow only unidirectional: Power is limited from host to connected device.
- In some cases, data transfer and power delivery failed to work simultaneously.

The USB Power Delivery spec addresses these issues by increasing the maximum power delivery to 100 W (from 7.5 W), along with other benefits:

- Power direction no longer fixed.
- Optimized power management across peripherals.
- Intelligent and flexible system-level management of power.
- Low-power cases can negotiate only for the power required by them.

THE FUTURE OF USB POWER DELIVERY

From 7.5 to 100 W:

Not all peripherals could be powered by a USB cable as long as the maximum power was capped at 7.5 W. Examples include hard-disk drives (HDDs), solid-state drives, printers, and monitors. Consequently, these peripherals required independent power sources at an additional cost. With a power-delivery spec that allows for a theoretical 100-W max power budget, many of these peripherals can now be powered by USB cables.

The power-delivery spec allows for a maximum voltage of 20 V (4× the previous 5-V spec) and a maximum current of 5 A (more than 3× the previous max of 1.5 A). The power-delivery spec also classifies power sources in terms of profiles:

- Profile 1 (Default Startup): 10 W (5 V @ 2 A)
- Profile 2: 18 W (5 V @ 2 A -> 12 V @ 1.5A)
- Profile 3: 36 W (5V @ 2 A -> 12 V @ 3A)
- Profile 4 (Micro B/AB limit): 60 W (5 V @ 2 A -> 20 V @ 3 A)
- Profile 5 (Standard B/AB limit): 100 W (5 V @ 2 A -> 20 V @ 5 A)

When a device is connected to a host, an initial power supply of 10 W occurs to initiate the power negotiation. Based on the final profile selected, the power transfer is 18, 36, 60, or 100 W.

There's also a major change in the process of negotiating power. Earlier generations made use of the data wires for enumeration purposes: USB power delivery specifies negotiation using the V_{BUS} only, without affecting the data bus. This leads to the second benefit:

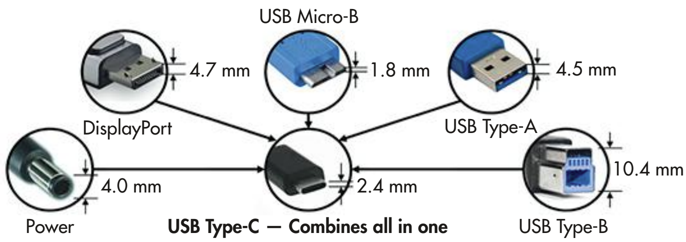

2. The Type-C connector replaces the USB Type-A, mini DisplayPort, and power connectors with a single connector. It leads to sleeker designs by eliminating the need for multiple USB ports.

Data transfer during power delivery:

The power-delivery spec removes the data bus from the process of negotiating power delivery. This makes possible simultaneous data transfer and power flow. This means devices can continue with their primary function while also charging.

Bidirectional power transfer:

Historically, a USB connection allowed for bidirectional data transfer, but power only flows from host to peripheral. The USB Power Delivery spec sanctions bidirectional power delivery. For example, if a laptop is connected to an external monitor that's plugged into the wall, the laptop can now be charged by the monitor.

Power management across peripherals:

The power-delivery spec helps to optimize power flow across multiple peripherals by:

- Allowing devices to draw as much power as they require and draw more power when an application demands it. Say, for example, a phone is charging through a USB connection to a PC. Then we connect a USB RAID array, which needs additional power at the start to get all of the disks spinning, but can subsequently be lowered to a steady state. The system can lower the power delivery to the phone, provide it to the RAID array, and then move it back to the phone when the power is available again.
- Incorporating intelligent/flexible power management across multiple peripherals. For example, battery-backed devices draw more power than devices connected to external power source.
- Allowing low-power devices to negotiate only the power that they require. For example, headsets can draw less power than external HDDs.

Together, these enable any externally powered peripheral to serve as a power source and hub for any host or devices connected to it.

OVERALL SYSTEM-LEVEL POWER MANAGEMENT

USB power delivery can have a substantial impact on overall system-level power-management efficiency. Let's say we need to connect a HDD and display to a laptop. The pre-power delivery implantation would be:

1. Connect a power source to all three system components—HDD, display, and laptop.
2. HDD and display are both connected to the laptop.
3. Data exchange between HDD-laptop and display-laptop
4. Power flows only from laptop to HDD or display.

With the power-delivery spec, the above implementation simplifies to the following extent:

1. Connect the display to a power source.
2. No additional power sources required for laptop or HDD. The user just needs to connect both to the display via a USB connection. Power will flow from display to laptop and HDD.

OTHER CONSIDERATIONS

Taking full advantage of the power-delivery spec involves a few more requirements. The USB-IF covered most of them in the USB Type-C specification to simplify deployment.

So what are these limitations and how does Type-C overcome them? What if a peripheral requires a non-USB signal? For example, display signals are of a non-USB type, and thus external displays are currently connected via DisplayPort and not USB. The power-delivery spec would allow the display to power up a host, but then the spec is not valid for a non-USB port.

Limitations of USB Type-A and Type-B (current USB-IF standard): The USB receptacle on a USB host or USB hub is a Type-A connector. Corresponding type-A plugs are found in cables and small peripherals (e.g., a mouse). However, some larger peripherals (e.g., a printer) use detachable cables. Thus, there's a receptacle required on such peripherals. This receptacle is Type-B. Type-A and Type-B connectors are similar electrically, but different mechanically. As per the USB-IF spec, a cable can deliver power only when one end is connected to Type-A and the other end to Type-B. This was done deliberately to prevent users from connecting host to host (leading to the risk of a short circuit). Power flow occurs only from Type-A to Type-B, and thus with this current setup, bidirectional power flow is not possible.

The new USB-IF Type-C standard solves these problems by:

- Transporting both the USB signal along with a DisplayPort signal over the same connector.
- Replacing Type-A and Type-B connectors with a single connector (Type-C) that supports bidirectional power flow. The plug height is just 2.4 mm—much less than Type-A (4.5 mm) or Type-B (10.4 mm). Hence, it's a useful feature for hosts that are striving for a slim design.

Together, this leads to a simplified implementation of USB with low-cost power delivery up to 100 W. From a product-design perspective, the Type-C connector allows OEMs to achieve sleeker designs by eliminating multiple ports *(Fig. 2)*. For example, the ultra-thin MacBook (the world's first product with a Type-C connector) has replaced three connectors—USB Type-A, mini DisplayPort, and power—with a single Type-C port. ed

Rate Multiplier Eases Measurement of Ultra-Low Frequencies

KAREN HUNSBERGER | INDEPENDENT CONSULTANT

ENGINEERS OCCASIONALLY NEED to measure very low frequencies between 1 and 100 Hz, but getting an accurate reading is difficult. The solutions are to use a suitably long gate time, or get an expensive counter that can compute frequency from period. The inexpensive circuit here effectively upconverts the input frequency to a higher range so it can be read easily and more accurately, and can be implemented as a block within a larger circuit.

At the core of the design are 4089 rate-multiplier ICs (U4 to U6), which output a frequency at N × f(in)/ 16, where N is the 4-bit code applied to the input frequency f(in). By cascading three of these ICs, the equation becomes N x f(in)/4096. Two 7497 6-bit ICs may be used instead of the three 4089 devices to reduce package count.

U1 is a 12-bit counter that measures the period of the input; the count is latched in U2 and U3. This causes the rate multipliers to output a frequency that's directly proportional to the input period. U7 divides this by 4096 to remove jitter and make the period more easily measurable. As with U1, U8 measures this period, which is latched in U9 and U10.

This provides a measure of the period of U7's output, which is inversely proportional to its frequency. Because that frequency was proportional to the input's period, the value latched by U9 and U10 is directly proportional to the input frequency, but at much higher resolution, and can be measured much more quickly. U11 and U12 generate all of the needed timing signals. Different frequencies can be chosen, depending on the range to be measured.

An example helps illustrate the process. For f(in) = 4 Hz and f1 = 4000 Hz, U1 will count to 1000 before being latched by U2 and U3. This will cause rate multipliers U4 to U6 to generate 1000 × f2/4096, or 1 MHz for f2 = 4.096 MHz. U7 divides this down to 244 Hz, and U8 measures the period as f3/244 Hz, or 524 for f3 = 128 kHz.

Therefore, we have multiplied the input frequency by a factor of 131 (524 Hz/4 Hz), so the circuit can resolve to 0.0076 Hz in one second. This would otherwise take two minutes using the original 4-Hz input.

KAREN HUNSBERGER, an independent consultant, holds a B.Sc (Hon) from the University of Waterloo and a B.Ed. from the University of Western Ontario. She enjoys working on alternative energy and farm/garden designs.

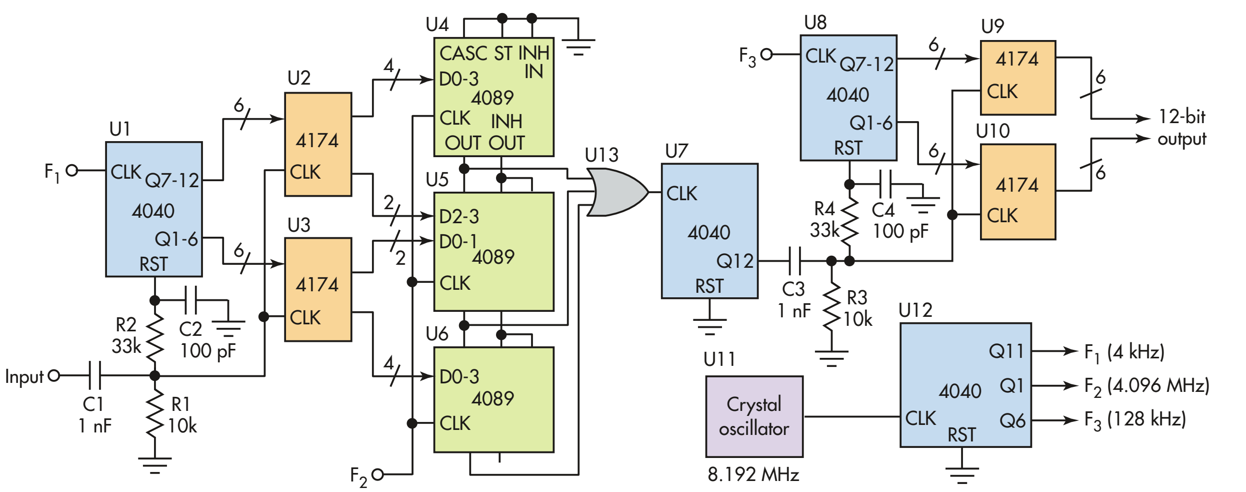

Due to the fact that cascaded rate multipliers upconvert the very low input frequencies, it's possible to measure frequency in a fraction of the normally required time.

Simple Circuit Overcomes MOSFET Gate-Threshold Voltage Challenge

SAJJAD HAIDAR | Electronics Engineering Services, University of British Columbia, Canada, sajjad_haidar@yahoo.com

IN SOME SITUATIONS, it becomes necessary to drive a MOSFET (or IGBT) with a voltage that's lower than its gate-threshold voltage (V_{Th}). This is usually done with a driver or op amp to boost the signal to a level sufficient to drive the device. However, if the power-supply voltage powering the device is below or close to V_{Th}, then even a rail-to-rail op amp will not be able to drive the MOSFET.

A simple technique can be used to drive a MOSFET with a signal that's lower than V_{Th} *(Fig. 1)*. The isolated LED of optocoupler U1 is powered by the load supply (+2.5 V), and potentiometer R1 (100 kΩ) is connected at the output of the optocoupler. Since U1 is a photovoltaic coupler, it acts as a constant-current source up to a certain voltage, with the range determined by the device characteristics.

By changing the resistance of R1, it's possible to bias MOSFET Q1 at various points to set different voltages at R1 *(Fig. 2)*. With R4 set to 68 Ω, the LED forward current is set at approximately 16.5 mA, which provides short-circuit current I_{SC} of the photovoltaic output of about 48 µA. Capacitor C1 ensures low signal-path impedance; its value should be much higher than the gate capacitance of MOSFET Q1. Here, we chose C1 = 0.1 µF, which is higher than the gate capacitance of the MOSFET being used.

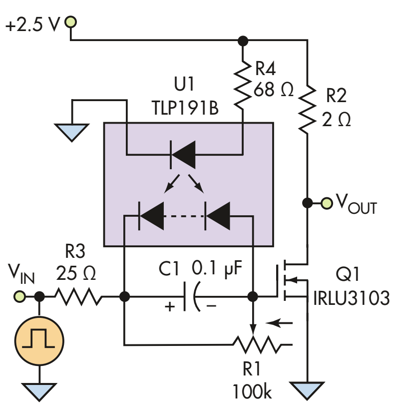

1. The optocoupler as a current source provides drive to a MOSFET gate with a voltage lower than the threshold voltage; R3 is the internal resistance of the test pulse generator.

A resistive load, R2 (2 Ω), is used to test the circuit. R1 is adjusted to set the voltage close to the threshold voltage of Q1, which is measured as approximately 1.9 V. With the TPL191B, the maximum obtainable output voltage is about 7 V (V_{OC} is approximately 8 V), which is above the threshold V_{Th} of most devices. A signal generator is used to deliver pulses of amplitude about 1.2 V at the input (V_{IN}); the signal generator has an internal resistance of 25 Ω (R3). Both input voltage (V_{IN}) and output voltage (V_{OUT}) are measured by an oscilloscope (*Fig. 3*).

The total input voltage to switch the MOSFET to the "on" state is 1.2 V + 1.9 V = 3.1 V, which is higher than the supply voltage (2.5 V). The small delay in off and on times (t_{OFF} and t_{ON}) is primarily due to the pulse generator's resistance (R3). A little overshoot occurs in the output voltage during switch-off, due to the load inductance. The 2-Ω load isn't purely resistive, because it has a little inductance, too. It's possible to use a smaller input signal (V_{IN}) to switch a MOSFET, depending on the drain current and transconductance of the selected MOSFET.

This circuit works from dc to high frequencies, as set by the MOSFET. To see the effect with other MOSFETs, use LTSpice or any other simulation tool with a constant-current source in place of the photovoltaic coupler. In this case, use a current source of 48 μA.

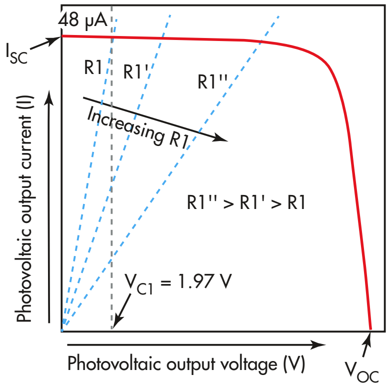

2. The output characteristic of the photovoltaic coupler shows that by adjusting the value of R1, the bias-point of the MOSFET can also be adjusted.

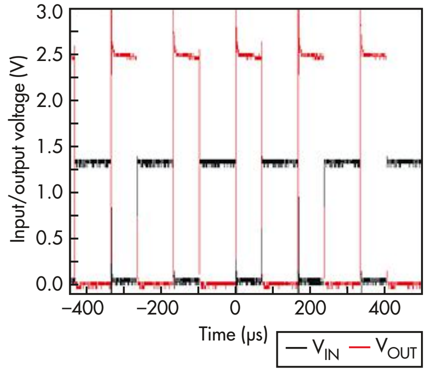

3. The relationship between input voltage and output voltage also shows slight on/off time delays, largely due to the internal resistance of the pulse generator used for test stimulus.

SAJJAD HAIDAR is an Electronics Technologist at the Electronics Engineering Services of the University of British Columbia (UBC). He holds an M.Sc. in applied physics and electronics from the University of Dhaka (Bangladesh). Previously, Mr. Haidar worked in Japan for seven years in the field of tunable solid-state lasers and optoelectronics. He can be contacted at sajjad_haidar@yahoo.com.

REFERENCES:
1. Datasheet, TLP191B optocoupler: http://www.toshiba-components.com/docs/opto/TLP191B_en_ datasheet.pdf
2. Datasheet, IRLU3103 MOSFET: http://www.irf.com/product-info/datasheets/data/ irlr3103pbf.pdf

Wireless NiMH Charger-Receiver Optimized for Hearing Aids

THE LTC4123 from Linear Technology combines a 30 mW wireless receiver with a constant-current/constant-voltage linear charger for NiMH batteries. An external resonant LC tank connected to the device enables the IC to receive power wirelessly from an alternating magnetic field generated by a transmit coil. Wireless charging allows for a completely sealed product and eliminates replacing of primary batteries. Zinc-Air detection allows applications to work interchangeably with both rechargeable NiMH batteries and primary Zn-Air batteries with the same application circuit. The charger-receiver rectifies ac power from the receive coil, and can also accept a 2.2 V to 5 V input to power a full-featured constant-current/constant-voltage battery charger. Other features include programmable charge current up to 25 mA and a single-cell 1.5 V battery charge voltage with ±0.5% accuracy. The device prevents charging when batteries are inserted with reverse polarity and pauses charging if the temperature becomes too hot or too cold.

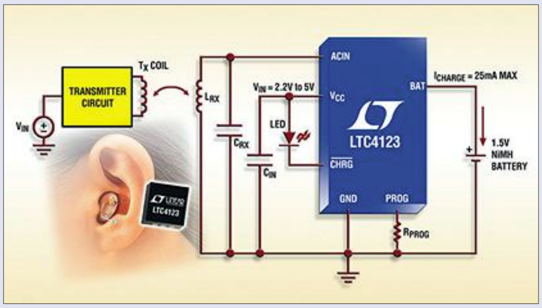

The LTC4123 charger-receiver, housed in a 6-pin 2 × 2 × 75 mm DFN package with backside metal pad, is guaranteed for operation from -20°C to 85°C in E grade and is priced at $2.65 each/1,000.

LINEAR TECHNOLOGY

www.linear.com

2-/4-Port PC/104-Plus CAN I/O Module Targets Rugged Networks

THE JANUS-MM-LP FAMILY of CANbus 2.0 I/O modules from Diamond Systems features independent isolation for each port plus data rates up to 1 Mbps. Z The rugged I/O modules offer 2 or 4 CAN ports and are available in the compact PC/104 and PC/104-Plus form factors. The modules offer two or four opto-isolated CAN 2.0B ports plus 16 digital I/O lines, and models are available with both PC/104-Plus and PC/104 bus configurations. The CAN controllers are implemented as FPGA cores and feature standard and extended device identifiers as well as expanded TX and RX message queues. Each port has its own combination isolator and transceiver chip, and 16 programmable digital I/O lines are organized as two eight-bit ports.

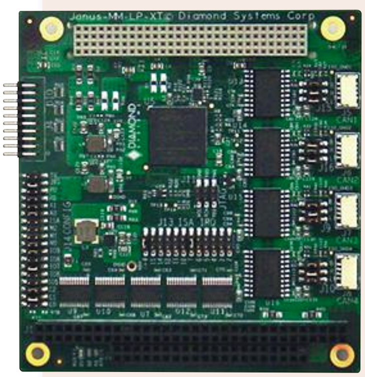

Technical features include data rates up to 1 Mbps; support of standard 11-bit identifier and extended 29-bit identifier frames; 16 8-byte transmit message queues; 31 8-byte receive message queues; 16 receive filters; galvanically isolated transceivers; 500 V port-to-host and port-to-port isolation; jumper-selectable biased split termination; 16 digital I/O lines; latching connectors; PCI and ISA bus interfaces; free basic CAN driver with APIs and monitor program; Linux Ubuntu 12.04LTS and Windows Embedded 7 software support; –40°C to +85°C operating temperature; and MIL-STD-202G shock and vibration compatibility.

DIAMOND SYSTEMS

www.diamondsystems.com

Optically Isolated Error Amps Geared for Power-Supply Feedback

THE LIA135 AND LIA136 low-voltage optically isolated error amplifiers from IXYS Integrated Circuits Division feature an integrated programmable precision shunt reference incorporated into a single eight-pin package. The error amplifier and shunt reference are capable of operating down to 1.6 V, and the devices offer a 1.299-V internal precision voltage reference with a 0.5% tolerance. Integrated, optically isolated error amplifiers, useful for secondary to primary side feedback in ac-to-dc power supplies and dc-to-dc converters, also save board space and replace transformer secondary feedback circuits. The devices' combination of features are optimal for use in isolated power supply design as the bias current for the shunt regulators does not pass through the LED, eliminating bias-related optical current and providing a dynamic range of 1,000:1.

The optocouplers' current transfer ratio is between 500% and 2,000%. The LIA136 boasts enhanced performance characteristics and an extended operational temperature range to 110°C. The LIA135 and LIA136 low voltage optically isolated error amplifiers, available in DIP and SMT packages, are UL E76270 and CSA certified.

IXYS INTEGRATED CIRCUITS DIVISION

www.ixysic.com

Level VI Power Supplies Feature Interchangeable Input Blades

A FAMILY OF LEVEL VI COMPLIANT external ac-dc power supplies from CUI range from 5 W to 36 W and feature interchangeable ac input blades. Designed for global use, the wall plug-in power supplies include ac blade options for North America, Europe, the UK, Australia, and China. The series are all designed to meet the Level VI standards, which aim to significantly lower the amount of power consumed when the end application is not in use or is no longer connected to the system. The 5-, 6-, 18-, 24-, and 36-W multi-blade adapters feature an input voltage range of 90~264 V ac and are available in single output voltages from 5~48 V dc, depending on the series. Protections for over voltage, over current, and short circuit are also included and vary by series.

The SMI5-USB comes with an integrated USB connector and offers UL/cUL, GS and RCM safety approvals. The SMI6 series includes the same, plus CCC and PSE safety approvals. The SMI18 and SMI24 series both offer UL/cUL, GS, and PSE safety approvals, and the SMI36 series holds those, plus the CCC safety approval. The Level VI power supply series is available immediately, with prices starting at $7.93 each/100.

CUI

www.cui.com

12.1-In. TFT LCD Utilizes LED Backlight as Light Source

NLT TECHNOLOGIES' 12.1-IN. diagonal sunlight readable LED-backlit TFT LCD, the NL8060BC31-51C, combines its proprietary T-EVT (Transmissive Enhanced View Technology) for improved viewing in outdoor or other high ambient light environments, and ColorXcell technology for improved color reproduction with lower power. It features 900 cda/m2 brightness, a 1,000:1 high contrast ratio, and a long-life LED backlight that provides 70,000 hours of operation. The T-EVT technology uses the backlight as a light source while minimizing the surface reflection of ambient light, producing high contrast images even in bright sunlight. ColorXcell technology enables the reproduction of images that are comparable in color fidelity to the color content of the original video source. The 12.1-in. diagonal sunlight readable LED TFT display (NL8060BC31-51C) replaces prior-generation CCFLs and is suited for outdoor display systems.

NLT TECHNOLOGIES

www.nlt-technologies.co.jp/en/

Low-Power Bluetooth LE Module Reduces Size by 60%

THE SESUB-PAN-D14580 Bluetooth v4.1 LE micro module from TDK is based on the company's proprietary Semiconductor Embedded in Substrate technology. The single-mode module's ultra-compact footprint and low current consumption suit it for battery-powered wearable devices where small size, light weight, and low power consumption are essential.

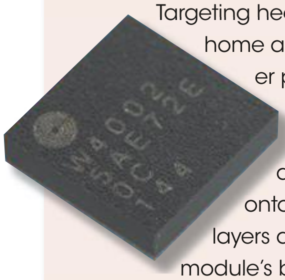

Targeting healthcare and fitness devices, wearables, home and entertainment devices, and computer peripherals, the module integrates a Dialog Semiconductor DA14580 Bluetooth 4.1 chip, 32-bit ARM Cortex-M0 MCU, dc-dc converter and peripheral circuitry onto a thin substrate. I/Os from the substrate layers are routed to a BGA footprint on the module's bottom surface. Interfaces include UART, SPI, and I2C. The module requires a voltage supply of 3.0 V, and consumes 5.0 mA when transmitting, 5.4 mA when receiving and 0.8 µA in standby mode. Output power is rated at 0 dBm (typ.) with a communication range of 10 meters, depending on line of sight and antenna characteristics. Mouser Electronics is now stocking the SESUB-PAN-D14580 Bluetooth v4.1 module.

TDK

www.tdk.com

IoT Messaging Platform Provides Asynchronous Communication

THE FIORANOMQTT enterprise-class messaging platform from Fiorano supports all the features of the lightweight event- and message-oriented MQTT protocol that allows devices/sensors to communicate with remote systems in an asynchronous fashion. Specially designed for resource-constrained devices over unreliable networks with low-bandwidth and high-latency issues, the platform features: the ability for MQTT messages published from devices/sensors to be received by JMS subscribers; client authentication based on LDAP realm; shared high-availability models; client-side persistence for QOS 1 and QOS2 level packets; and the ability to send/receive message packets asynchronously.

The platform supports the MQTT 3.1.1 (OASIS standard) protocol which runs simultaneously with the JMS Server (FioranoMQ). Receiving data from thousands of devices/sensors and delivering them to MQTT and JMS based back-end processing, the FioranoMQTT messaging platform can then pass this data to different enterprise applications such as ERP, PLM, MES, CRM, and/or Big Data processing systems, via Fiorano ESB.

FIORANO

www.fiorano.com

Smart-Touch Computers Utilize SMARC COM Design

ADLINK TECHNOLOGY'S LATEST PANEL PC for industrial applications, the Smart Touch Computer comes in three models: STC-1005/1205/1505. Designed using the Smart Mobility Architecture (SMARC) computer-on-module concept to increase flexibility and scalability when developing and upgrading end solutions, the computers are available with 10.4-., 12.1-, and 15-in. touchscreen displays and are powered by quad-core Intel Atom E3845 processors. A compact aluminum enclosure provides a front bezel with IP65-rated protection to shield from dust and water jets. The STCs are available with either projected capacitive or five-wire resistive multiple touchscreen options.

The computers leverage the SMARC design concept to simplify the process of implementing different solutions to meet specific requirements with one uniform platform. The units are also equipped with standard I/O ports and feature a proprietary expansion slot, allowing customization for specific application requirements. The STC-1005/1205/1505 support VESA and panel mounting and can be installed in a control cabinet or on a support arm. An externally accessible SATA drive and SD card slot are also provided.

ADLINK TECHNOLOGY

www.adlinktech.com

Adapter Simplifies Testing of 0.5 mm Pitch 300-Pad Flex Module

IRONWOOD ELECTRONICS' PB-FLEX300-Z-01 Socket Probe Adapter allows high-speed testing of flex module while accessing the signals using testers via FPC connector. The adapter provides the shortest possible trace length for maximum speed and lowers inductance and capacitance via PCB design technology. The socket probe adapter is designed to interface with 0.5 mm pitch Fine pitch pads of flexible circuit module. The adapter consists of two parts: the Probe board with FPC connectors for multiple groups of signals, and the clamshell socket to accommodate a 300-pad flex module. The spring pin contact flex module socket is mechanically mounted onto the probe board. The spring-pin socket is constructed with a high-performance stamped spring pin with 14-gram actuation force per ball and cycle life of 50,000 insertions. The self-inductance of the contactor is 0.98 nH, insertion loss < 1 dB at 31.7 GHz. The current capacity of each contactor is 1.8 A at 20°C temperature rise. Capacitance to ground is 0.01 pF and operating temperature is –55°C to +180°C.

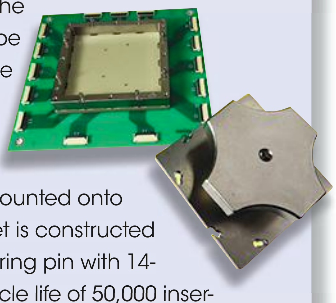

IRONWOOD ELECTRONICS

www.ironwoodelectronics.com

Analysis Functionality Added to Tools for RL78 MCUs

AN UPDATED VERSION OF IAR SYSTEMS' embedded development tools, IAR Embedded Workbench for Renesas RL78, is now available. The new release includes functionality for aiding development and code quality control for applications based on Renesas low-power RL78 MCUs. The update now supports the add-on product C-STAT for static analysis that can detect defects, bugs, and security vulnerabilities as defined by CERT C/C++ and the CWE, as well as for helping keep code compliant to coding standards like MISRA C:2004, MISRA C++:2008 and MISRA C:2012. Version 2.20 also introduces stack usage analysis, which adds listings of the maximum stack depth for each call graph root to the linker map file. The analysis process can be customized to take into account such constructs as calls via function pointers and recursion.

The IAR Embedded Workbench v.2.20 includes the company's C/C++ Compiler, assembler, linker, library tools, and the C-SPY Debugger in the IDE. It is available in several editions, including a Functional Safety version with IEC 61508- and ISO 26262 certifications.

IAR SYSTEMS

www.iar.com

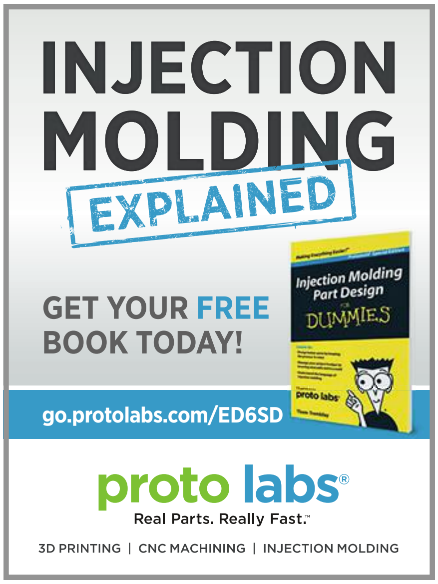

Breakout Test Board for 0.5mm Pitch Flex module

Ironwood Electronics' PB-FLEX300-Z-01 Socket Probe Adapter which allows high-speed testing of flex module while accessing the signals using testers via FPC connector. Ironwood Electronics is an expert in developing high-speed interfaces for processors, peripheral IC's and memory devices and is a development partner with both Agilent and Tektronix. Features of the PB-FLEX300-Z-01 include shortest possible trace length for maximum speed, low inductance, low capacitance via PCB design technology.

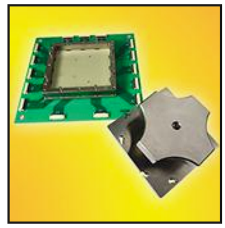

Ironwood Electronics, Inc.

1335 Eagandale Court
Eagan, MN 55121
T: 800-404-0204 • F: 952-229-8201
Email: info@ironwoodelectronics.com
Web Site: www.ironwoodelectronics.com

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

Drones Go Gaming

Skyrocket Toys' Hover Racer may change the way drones are gamed, with its IR emitters and receivers that target and detect other Hover Racers.

As always, the 2016 Consumer Electronics Show (CES) had a host of new products on display, from 4K UHD televisions to flaming-fast hoverboards. Of course, there were lots of drones on hand and the FAA was there as well, letting everyone in the United States know that any drone over 0.55 pounds needed to be registered. Fortunately, many exhibitors were delivering drones that were under this limit.

1. Skyrocket Toys' Hover Racer has IR emitters and receivers spaced around its periphery to target and detect other Hover Racers.

2. The pylons have an IR emitter allowing the drone to track its relative position.

Skyrocket Toys' Hover Racer (*Fig. 1*) is one of those. It features many of the characteristics of other micro-drones, including limited flight time, impact-resistant unibody design, and low cost. It has six-axis digital stabilization and variable modes of flight-control sensitivity.

The Hover Racer can be flown individually, but it is also targeting the gaming aspect of quadcopters that others are pursuing as well. In this case, there will be multiple quadcopters in flight at the same time.

The Hover Racer comes equipped with five sets of IR emitters and receivers that target and detect other Hover Racers. Skyrocket Toys provides IR-emitting pylons that can be detected by the drone's downward-looking sensors (*Fig. 2*). These can be used to market a course and help drones determine their relative position. This allows an app to track the fliers if they are running a race course where each pylon must be passed in the proper order. The system is designed to handle up to four racers.

The other sensors can be used when drones are dogfighting. The emitters run continuously and there is a cone that is illuminated so other drones can detect when they are within range of a cone and which drone is targeting them. The system can be used to simulate a range of weapons and sensors. For example, the system does not have to always alert a player immediately when a drone is detected. A second player can act as an RIO (radar intercept officer) since the Hover Racer can link to two mobile devices at one time. This can help as the flight environment becomes busier. Flying can be hard enough.

Skyrocket Toys' intent is ambitious. They were showing the hardware at CES, which is impressive, especially at the price point they are targeting. The Hover Racer is a neat toy, but it highlights the rapid change in sensors and control systems. Larger and more expensive systems can incorporate more complex sensors, but it is amazing what is packed into this small quadcopter.

Higher-end systems at CES were sporting multiple cameras and other range-finder technology to prevent collisions and improve mapping. Skyrocket Toys' Sky Viper is a more conventional quadcopter that streams and records HD video.

Quadcopters are already being used in battlebot-style games, but Hover Racer is designed to be a bit more sporting and less destructive. Look for even more alternatives to fill this niche in the future.

So, fire up your Hover Racer and let the games begin.

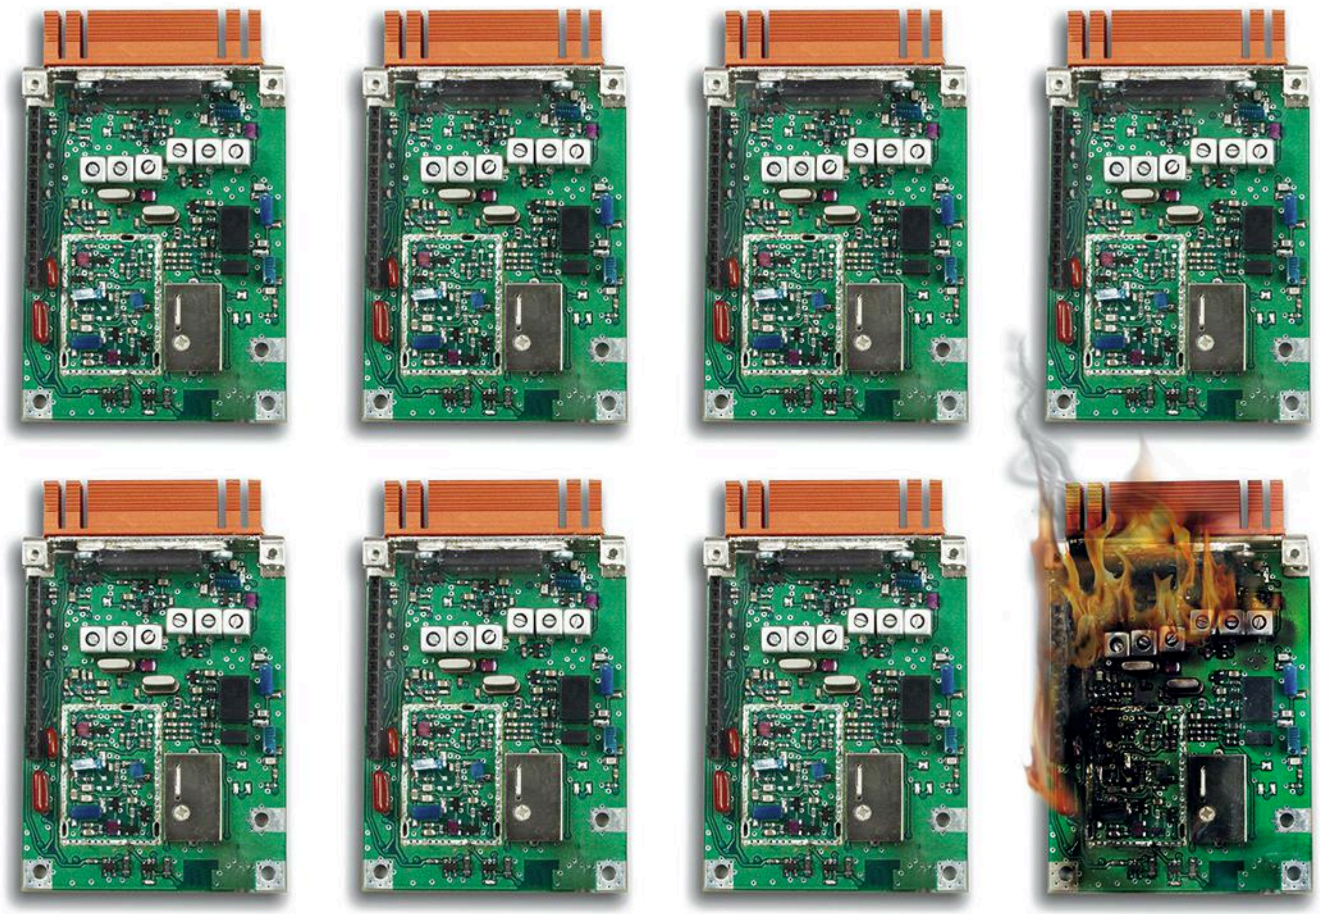

www.bergquistcompany.com 1.800.347.4572

952.835.2322 fax 952.835.0430

18930 West 78th Street • Chanhassen, Minnesota 55317

BERGQUIST

Laminates Division

±270V Precision Difference Amplifier

Wide Common Mode Bidirectional Current Monitor

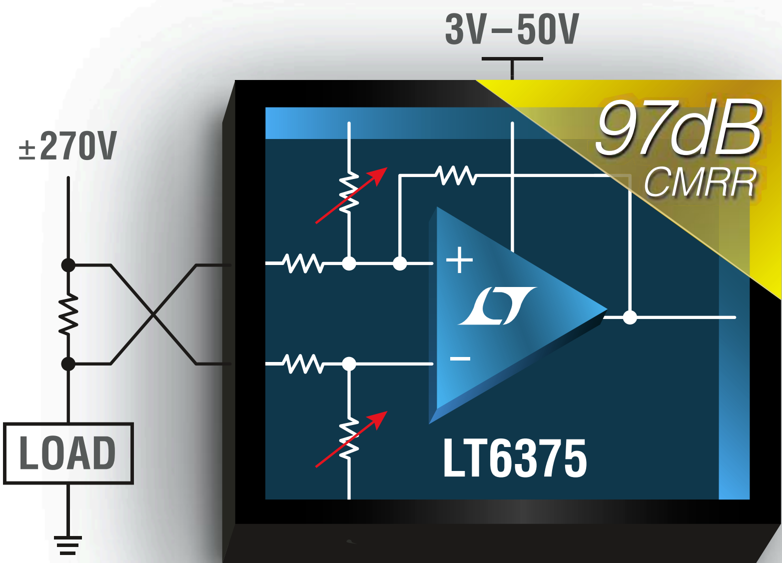

Tailored precision. The LT6375 delivers 97dB min CMRR, ultralow gain error, drift, and nonlinearity for applications where a differential signal must be precisely extracted from a common mode of up to ±270V. Precision matched internal resistors can be configured to match various input ranges, resulting in the best noise, precision and speed for a wide variety of system designs. The Over-The-Top® amplifier at the heart of the LT6375 enables operation from 3.3V to 50V supplies.

Features

- ±270V Common Mode Input Voltage Range
- 97dB Min CMRR
- 0.0035% Max Gain Error
- 1ppm/°C Max Gain Error Drift
- 2ppm Max Gain Nonlinearity
- Selectable Resistor Divide Ratio (7:1 to 25:1 Range)
- 3.3V to 50V Supply Voltage Range

Typical Distribution of Common Mode Error

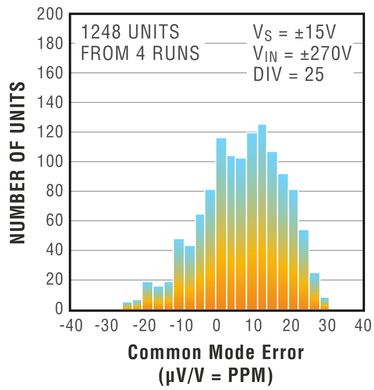

Info & Free Samples

www.linear.com/product/LT6375

1-800-4-LINEAR

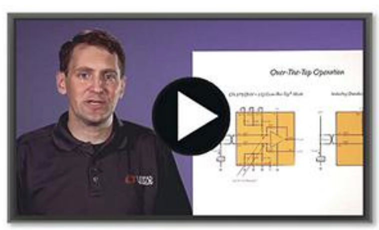

video.linear.com/5972