Electronic Design®

THE AUTHORITY ON
EMERGING TECHNOLOGIES
FOR DESIGN SOLUTIONS

SEPTEMBER 2017 electronicdesign.com

IoT BETS ON THE FARM

High-tech agricultural techniques improve production and create efficiencies p12

$10.00


"I was just working on a new design and wanted to say how much help your website is. I've gotten to the point where I look up part availability (for a reality check), get the datasheets and any notes, then it's designed-in; put the manufacturer's part and Digi-Key number right in the BOM."
Microscopy Automation Provider, Digi-Key Customer since 2013
What Our
Customers
Are Saying...
"Digi-Key continues to be my go-to resource, and much of that is due to the very easy online ordering tools and quick responsive customer service. Keep up the great work!"
Engineer, Digi-Key Customer since 2011
"I've been so pleased with you folks at Digi-Key for nearly 20 years now. I am so happy with everything about your company and your people. You have made my life as an EE so much easier and more productive. You ship promptly so I have parts in the morning, even if I don't realize I need parts until late in the day. I have come to depend on that service. You're the best!"
Engineer, Digi-Key Customer since 1998
SAME-DAY SHIPPING
on orders entered by
8:00 pm CENTRAL
NEXT-DAY DELIVERY
1.800.344.4539
DIGIKEY.COM
Digi-Key® ELECTRONICS
6 MILLION+ PRODUCTS ONLINE | 650+ INDUSTRY-LEADING SUPPLIERS | 100% AUTHORIZED DISTRIBUTOR
Digi-Key is an authorized distributor for all supplier partners. New products added daily. Digi-Key and Digi-Key Electronics are registered trademarks of Digi-Key Electronics in the U.S. and other countries. © 2017 Digi-Key Electronics, 701 Brooks Ave. South, Thief River Falls, MN 56701, USA
ecia MEMBER
ecsn member
CEDA

IN THIS ISSUE

FEATURES

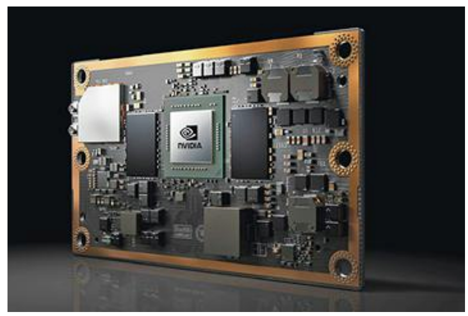

26

COLUMNS & DEPARTMENTS

IDEAS FOR DESIGN

42

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 ©2017 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

Turn Your Car into an Office

The Next Big Things in auto electronics: integrating VPAs like Alexa into Infotainment systems, and combining gesture recognition with Haptic Touch Feedback. By 2022, according to the research firm IHS Markit, nearly 90% of new cars will have some type of speech-recognition capability, and 75% of those cars will also have cloud-based voice control provided by companies such as Microsoft, Amazon, and Google.
For this and more articles on automotive electronics, sign up for Electronic Design's new digital newsletter on the topic (*http://www.electronicdesign.com/newsletters/signup*).

http://www.electronicdesign.com/automotive/turn-your-car-office

Who Are China's Biggest Fabless Chipmakers?

In China, there is growing competition for semiconductor smarts. The number of fabless chipmakers in China rose from 500 to 1,300 over the last five years, and they generated revenues of $24.1 billion in 2016. Here we offer a list of the country's largest fabless suppliers last year along with their 2016 revenues.

http://www.electronicdesign.com/industrial-automation/who-are-chinas-biggest-fabless-chipmakers

11 Myths About Z-Wave Technology

There are bound to be some common misconceptions about the ever-expanding IoT industry and the protocols that support it. To help set the record straight, here are 11 (+1) common myths surrounding Z-Wave technology.

http://www.electronicdesign.com/industrial-automation/11-myths-about-z-wave-technology

5G: The Solution to Broadband Infrastructure

Small towns and rural areas still do not have adequate internet access, if any at all. But, according to Contributing Editor and blogger Lou Frenzel, there are multiple efforts in place to solve that problem.

http://www.electronicdesign.com/communications/5g-solution-broadband-infrastructure

HELPING YOU LIKE THINGS EVERYDAY

MPD DEVICES CAN BE FOUND ON MULTIPLE SOCIAL MEDIA DEVICES , SERVERS, AND DATABASES.

MPD

MEMORY PROTECTION DEVICES

BATTERYHOLDERS.COM

ElectronicDesign.

SEPTEMBER 2017

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
EMBEDDED/SYSTEMS/SOFTWARE: **WILLIAM WONG** bill.wong@penton.com
ANALOG/POWER: **MARIA GUERRA** maria.guerra@penton.com
CONTENT OPTIMIZATION SPECIALIST: **WES SHOCKLEY** wes.shockley@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
CONTENT DESIGN SPECIALIST: **JOCELYN HARTZOG** jocelyn.hartzog@penton.com
CONTENT & DESIGN PRODUCTION MANAGER: **JULIE JANTZER-WARD** julie.jantzer-ward@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
OMEDA **T** | 847.513.6022 **TOLL FREE** | 866.505.7173

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

AL, AR, SOUTHERN CA, CO, FL, GA, HI, IA, ID, IL, IN, KS, KY, LA, MI, MN, MO, MS, MT, NC, ND, NE, NV, OH, OK, SC, SD, TN, UT, VA, WI, WV, WY, CENTRAL CANADA: **JAMIE ALLEN** **T** | 415.608.1959 **F** | 913.514.3667 jamie.allen@penton.com

CT, DE, MA, MD, ME, NH, NJ, NY, PA, RI, VT, EASTERN CANADA:
SHANNON ALO-MENDOSA **T** | 978.501.9116 shannon.alo-mendosa@penton.com

AK, NORTHERN CA, OR, WA, WESTERN CANADA: **STUART BOWEN** **T** | 425.681.4395 stuart.bowen@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

ITALY: **DIEGO CASIRAGHI** diego@casiraghi-adv.com

PAN-ASIA: **HELEN LAI** **T** | 886 2-2727 7799 helen@twoway-com.com

PAN-ASIA: **CHARLES LIU** **T** | 886 2-2727 7799 liu@twoway-com.com

PLEASE SEND INSERTION ORDERS TO: orders@penton.com

PENTON REPRINTS: **WRIGHT'S MEDIA** **T** | 877.652.5295 penton@wrightsmedia.com

LIST RENTALS/
SMARTREACH CLIENT SERVICES MANAGER: **JAMES ADDISON** **T** | 212.204.4318 james.addison@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

1166 AVENUE OF THE AMERICAS, 10TH FLOOR
NEW YORK, NY 10036 **T** | 212.204.4200

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

PENTON IS AN INFORMA BUSINESS

NANCY FRIEDRICH | Director of Editorial Content
nancy.friedrich@penton.com

Can NASA Tame Yellowstone Supervolcano?

Intelligence and data on volcanoes gained via satellites, robots, and artificial intelligence could help prevent catastrophic eruptions.

It sounds like the plot of a disaster movie, only the approach will take years to have any effect. In case you haven't heard yet about NASA's either incredibly brilliant or amazingly outlandish plan to protect us from the supervolcano in Yellowstone National Park (and potentially the other 20 or so supervolcanos on Earth), it is basically the highly complicated equivalent of adding a bit of cold water to a hot bath to cool it down.

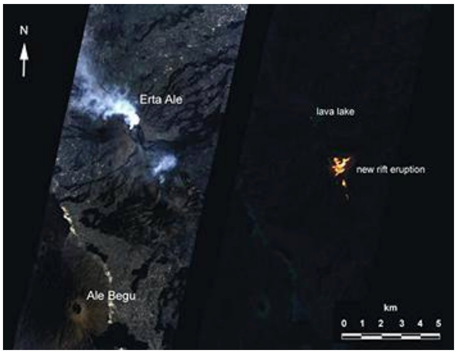

When a fissure opened at the peak of Erta Ale in Ethiopia this past January, an AI system automatically ordered the EO-1 spacecraft to pinpoint the volcano and obtain images. *(Courtesy of NASA)*

According to an article on the BBC website, NASA's idea centers around extracting some of the heat from the volcano to prevent its eruption: "They believe the most viable solution could be to drill up to 10 km down into the supervolcano, and pump down water at high pressure. The circulating water would return at a temperature of around 350°C (662°F), thus slowly day by day extracting heat from the volcano. And while such a project would come at an estimated cost of around $3.46 billion, it comes with an enticing catch which could convince politicians to make the investment." The "catch" is that this drilling approach could be used to create a geothermal plant, which could generate electric power at prices of roughly $0.10/kWh.

Of course, drilling into a high-pressure supervolcano has obvious risks—the biggest being the possibility of accidentally triggering an eruption instead of preventing it. This is why the scientists are thinking that they should drill from the lower sides rather than the top, which would allow the heat to be removed before it reaches the top of the chamber. According to estimates, this approach would result in cooling the Yellowstone volcano only 1 meter per year, which means this project would take thousands of years.

There are many questions to be answered before such a plan would be agreed upon and set in motion. Clearly, the decision-makers should utilize all of the data and information so far obtained by NASA as well as other agencies and organizations (such as the U.S. Geological Survey or USGS). They have been leveraging technologies ranging from radar to robots and artificial intelligence to monitor and investigate volcanoes. NASA's Earth Observing 1 (EO-1) satellite, for example, was equipped with artificial intelligence (AI) software that would prompt it to inform researchers when events of scientific interest took place and direct the spacecraft to take photos during passes.

According to the NASA website, the AI software—dubbed Autonomous Sciencecraft Experiment (ASE) —guided the actions of EO-1 for more than 12 years (concluding this past spring when EO-1's mission ended). Additionally, it managed a "sensor web" comprising both satellites and ground sensors to help decide which events were a priority. Given the ongoing efforts of satellites, ground sensors, radar, and more to obtain data on volcanoes, a supersized version of this AI software may be able to guide the fact-finding needed as plans are discussed around supervolcanos. As farfetched as this plan may sound, it could eradicate a threat and save millions of lives. NASA's scientists are hoping that their blueprints will encourage practical scientific discussion and debate for tackling the threat. ed

News

HOW AUTOMAKERS HANDLE Handoffs to Self-Driving Cars

There is little doubt that fully autonomous cars will reduce the number of deadly collisions caused by human mistakes on the road. But many experts fear that roads could grow more dangerous with semi-autonomous vehicles that need humans to take over in emergencies.

The problem stems from the fact that drivers must be prepared to take over control from a vehicle when it starts raining heavily, sensors fail, or highway lines fade. The semi-autonomous system could lull drivers into a false sense of security, leading to accidents if they fumble with retaking control while playing a smartphone game or have trouble adjusting to the feel of the steering wheel.

It is an extremely tough engineering nut to crack, and many engineers have stopped trying to solve it. That is why Google gutted the steering wheels and pedals in its experimental cars to focus on fully self-driving cars. Ford and Volvo have separately announced plans to aim for higher levels of autonomous driving that allow drivers to read or watch television.

But many companies contend that what Bryant Walker Smith, an assistant professor of law and engineering at the University of South Carolina, terms the “mushy middle” of autonomous driving can improve highway safety and provide insight into new technologies. These manufacturers are tackling “handoffs” in many different ways.

Audi: Last month, the company officially announced the A8, which will be the first production vehicle with Level 3 autonomous driving. This level allows drivers to watch television or read in highway traffic jams. The car sends audio and visual alerts in emergencies, giving drivers eight to ten seconds to retake control. If the alerts are ignored, the car will even tighten the seat beat and pump the brakes, Wired reports.

Continental: Last month, automotive supplier Continental said that its semi-autonomous driving system would be similarly introspective, monitoring drivers’ attention levels with cameras after they push a button to cede control. The system will make visual and audio overtures in emergencies, making noises, buzzing the seat, and flashing warnings on the dashboard.

If drivers ignore the alerts, vehicle using the Cruising Chauffeur system will find a highway shoulder or other safe place to stop. Continental said that the technology, which qualifies as Level 3 autonomous driving, would be ready for production in vehicles by 2020.

Autoliv: The parts supplier has been testing a steering wheel ringed with infrared sensors that detects when a driver's hands are clutching the wheel. Autoliv programmed a vehicle to enter self-driving mode when drivers take their hands off the zForce steering wheel, allowing them to regain control simply by grabbing onto it again. This results in safer handoffs, Autoliv says.

Tesla: After a driver died in a collision while using Tesla's Autopilot system last year, the electric car maker updated its software to make crystal clear to drivers that their hands must always stay on the wheel. Now if drivers ignore reminders to hold the wheel for more than 15 seconds, Tesla's vehicles will slowly grind to a halt.

If the driver ignores too many alerts, the vehicle will disable the so-called Level 2 system – which means that the driver must be prepared to take control at a moment's notice – until the entire car is restarted. In March, Tesla also updated the software to that drivers must hit the turn signal to make automatic lane changes.

General Motors: The SuperCruise system to be installed in the Cadillac CT6 later this year will monitor a driver's gaze with an infrared camera embedded in the steering column. Holding the wheel is optional but the Level 2 system will alert drivers whose eyes have strayed from the road too long. If drivers ignore these alerts, a light bar starts glowing on the wheel and the vehicle resorts to visual and audio warnings.

Mercedes-Benz: The Level 2 autonomous driving mode that Daimler will introduce this year requires more attentiveness. With Drive Pilot, drivers can take their hands off the wheel. But the vehicle will check every 10 seconds that the driver has pressed two capacitive touch buttons on the wheel, otherwise it will send reminders on the dashboard and then make a repetitive bonging noise to grab the driver's attention. ■

IBM PROCESSOR AIMS to Blanket Encryption Over Everything

IBM CLAIMS THAT its new processor can encrypt data on a massive scale, concealing credit-card payments, travel-site bookings, and government payrolls from the prying eyes of hackers.

The company announced that the new silicon chip powers its latest line of mainframes, which can automatically keep entire systems encrypted at all times. The z14 chip devotes around 6 billion transistors—four times more than previous z13—exclusively to encryption, which encodes messages only decipherable with special keys.

IBM's new silicon chip powers its latest line of mainframes, which can automatically keep entire systems encrypted at all times.

The new encryption chip runs at 5.2 gigahertz to process more than 12 billion transactions every day ranging from ATM withdrawals to flight reservations. Manufactured on the 14 nanometer node, it contains 10 computing cores that can encrypt 13 gigahertz of data per second. The z13 could only process around 2.5 billion transactions every day.

IBM claims that it handles encryption more cheaply and efficiently than rival server systems, which burn through massive amounts of computing power to encrypt and decrypt data. The system's security prowess could be a unique selling point for businesses that typically only encrypt limited lumps of data.

Most corporations have been slow to open their wallets for large-scale encryption. Only around 4% of all the data stolen worldwide since 2013 was encrypted, IBM says. And only around 2% of information in corporate servers is encrypted today, as opposed to almost 80% of mobile data, according to consulting firm Solitaire Interglobal.

"The vast majority of stolen or leaked data today is open and easy to use because encryption has been difficult and expensive," said Ross Mauri, general manager of IBM's Z mainframe business, in a statement. "We created a data protection engine for the cloud era to have an immediate and significant impact on global data security."

To protect encryption keys, IBM created special circuitry that acts like dye packs hidden in bank vaults to foil robberies. When the hardware detects malware or other intruders prying into memory, it can throw out the keys and restore them once the coast is clear again.

Other companies are remaking chips to expedite cryptography in cloud servers. Intel's newest Xeon Scalable processors, for instance, encrypt and decrypt messages without having to save keys in memory, while Advanced Micro Devices added a security subsystem in its Epyc server chips that encrypt data stored in memory.

The mainframe announcement comes at a particularly painful point for IBM. The company, which has been trying to reorient the business toward cloud computing and data analytics, reported revenues of $19.3 billion in second quarter, down from $20.2 billion the same last year. It is IBM's twenty-first consecutive quarter of revenue decline. ■

IN LATEST QUARREL, TOSHIBA CUTS Western Digital Out of Manufacturing Plans

IN THE LATEST FEUD with its American counterpart, Toshiba said that it would continue building its newest fab alone, effectively cutting Western Digital from its supply of 3D NAND memory chips. Western Digital said it won't let that happen.

The companies have been trading barbs for months since Toshiba announced the sale of its memory business to offset devastating losses in its nuclear power unit. Western Digital claims final say on the fate of the Japanese memory chip maker because it partially owns the new Yokkaichi plant.

Western Digital inherited the Fab 6 investment when it bought SanDisk, whose partnership with Toshiba started in 2001. But the Tokyo-based company denies that the joint venture affords Western Digital special rights. And its patience seems to be wearing thin.

"Toshiba is dismayed by Western Digital's pattern of exaggerating SanDisk's right under the relevant agreements," the company said in a statement. "Despite claims to the contrary, Western Digital does not now possess any legal 'rights' to participate in this phase of investment, which is an important investment in the next generation of flash memory."

Western Digital disagreed, saying in a statement shortly after that its "legal rights are clear, and we remain confident that we will receive out share of any capacity from Fab 6." The company, based in San Jose, Calif., also said that it would continue its " constructive dialogue with Toshiba on this and other matters."

These negotiations have gone in and out of court for months. In early August, Western Digital said that a judge had ordered Toshiba to resume sending it wafers and other engineering samples. The ruling also nullified a restraining order that prohibited Western Digital employees from viewing databases shared with Toshiba.

Throughout the quarrel, Toshiba has argued that Western Digital is overstepping its rights. So far, the company has fielded bids from a corporate coalition including Foxconn and Apple. But it prefers to sell its memory chip business to a Japanese consortium led by the government, which placed an $18 billion bid that Western Digital has opposed.

Toshiba's memory unit is seen as the crown jewel of Japan's industrial giant, which sells everything from televisions to electric transformers. Almost four decades ago, Toshiba pioneered the NAND flash chips now vital to smartphones and personal computers and increasingly data centers and factory equipment. It holds a fifth of the global market, second to Samsung.

Toshiba also has the manufacturing might that has become key to success in the memory market. Losing capacity could significantly hurt Western Digital now that the recent shortage of short-term memory and long-term storage has sent prices soaring. Western Digital has submitted its own bids for Toshiba's chip unit.

With everything at stake, Toshiba's talks with Western Digital have been bumpy. Toshiba said that it had negotiated with its American partner over Fab 6 for months. Production is scheduled to start in the middle of next year. It will carefully fabricate advanced 3D NAND chips with 96 layers.

"Toshiba provided an investment proposal to SanDisk earlier this year," Toshiba said in its statement. "Despite numerous meetings and negotiations, including at the CEO to CEO level, Toshiba's proposal was not accepted on the timetable set out in the agreements."

As a result, Toshiba said that it would alone pay 195 billion yen (around $1.76 billion) for clean room equipment that can sandwich memory up to 96 layers. It is Toshiba's attempt to cut off Western Digital from the chips produced on the new technology. Installation could start as early as December. ■

Toshiba's memory unit is seen as the crown jewel of Japan's industrial giant, which sells everything from televisions to electric transformers.

NEW BILL TARGETS COMMON-SENSE Security for Internet of Things

U.S. LAWMAKERS UNVEILED a bill last month that, if passed, would set basic security standards for connected devices from wearables to environmental sensors purchased by federal agencies.

The bill, called the Internet of Things Cybersecurity Improvement Act of 2017, would require devices to have software that can be patched and passwords that can be altered before being sold to the U.S. government. Without such capabilities, experts warn, everything from internet routers to security cameras could be left open to digital threats.

The sponsors of the bill include Republican senators Cory Gardner and Steve Daines and Democratic senators Ron Wyden and Mark Warner, the co-chair of the Senate Cybersecurity Caucus. Last year, Warner also raised concerns to regulators about internet-connected toys recording conversations and collecting data from children.

The legislation comes almost a year after a malicious strain of code called Mirai recruited millions of webcams, routers, and other connected gadgets to attack servers that act like the internet's infrastructure. The so-called Mirai botnet crippled websites in large parts of the United States, making for a spectacular display of the Internet of Things' frailty.

For years, experts have warned that connected devices could be exposed without ways to patch their software or replace hard-coded passwords set at factories. That is particularly vital since sensors and other electronics could be deployed for decades, giving hackers ample time to, for example, steal personal information or take control of traffic lights.

Ray O'Farrell, chief technology officer of cloud computing firm VMWare, said that the bill would provide "reasonable security recommendations" for federal agencies. The bill also requires that devices employ standard protocols and are not sold with known security vulnerabilities.

Drafted with input from experts from the Atlantic Council and Harvard, the bill would create legal protections for "good-faith" researchers that break into devices to uncover previously unknown security flaws. It would also introduce guidelines to report these vulnerabilities.

Under the bill, agencies would also have to keep an inventory of deployed Internet of Things devices. The Office of Management and Budget will also be tasked with laying out guidelines for simpler devices with "limited" software and processing power, which might include wireless sensors or identification tags.

While the legislation will provide companies with a set of guidelines, it does little to directly regulate security, said Jonathan Zittrain, a founder of Harvard University's Berkman Klein Center for Internet and Society. But it could motivate companies eyeing sales to the government, which has a $95 billion technology war chest under President Donald Trump's proposed budget for next year.

"This bill deftly uses the power of the Federal procurement market, rather than direct regulation, to encourage Internet-aware device makers to employ some basic security measures in their products," Zittrain said in a statement. "This will help everyone in the marketplace." ■

MARIA GUERRA | Technical Editor

3 Ways the IoT Revolutionizes Farming

Farmers use high-tech agriculture techniques to improve production output, and exploit sensors and other IoT technologies to create a more efficient operation.

When we think of Internet of Things (IoT) applications, growing crops or raising livestock aren't the first visuals that come to mind. Yet the smart agriculture market is expected to grow from $5.18 billion in 2016 to $11.23 billion by 2022, according to Markets and Markets. That translates into a compound annual growth rate (CAGR) of 13.27% between 2017 and 2022.

Farms are becoming more connected as farmers realize the potential of IoT technologies in helping them minimize operation cost while still achieving better results. Examples include higher crops, lower livestock losses, and less water usage. To help improve farm performance, IoT technology providers continue to develop platforms that can sense, process, and communicate precisely measured environmental data. Behind these IoT platforms is an array of technologies that includes sensing, microcontrollers, transmitters, energy harvesting, LED lights, drones, and more.

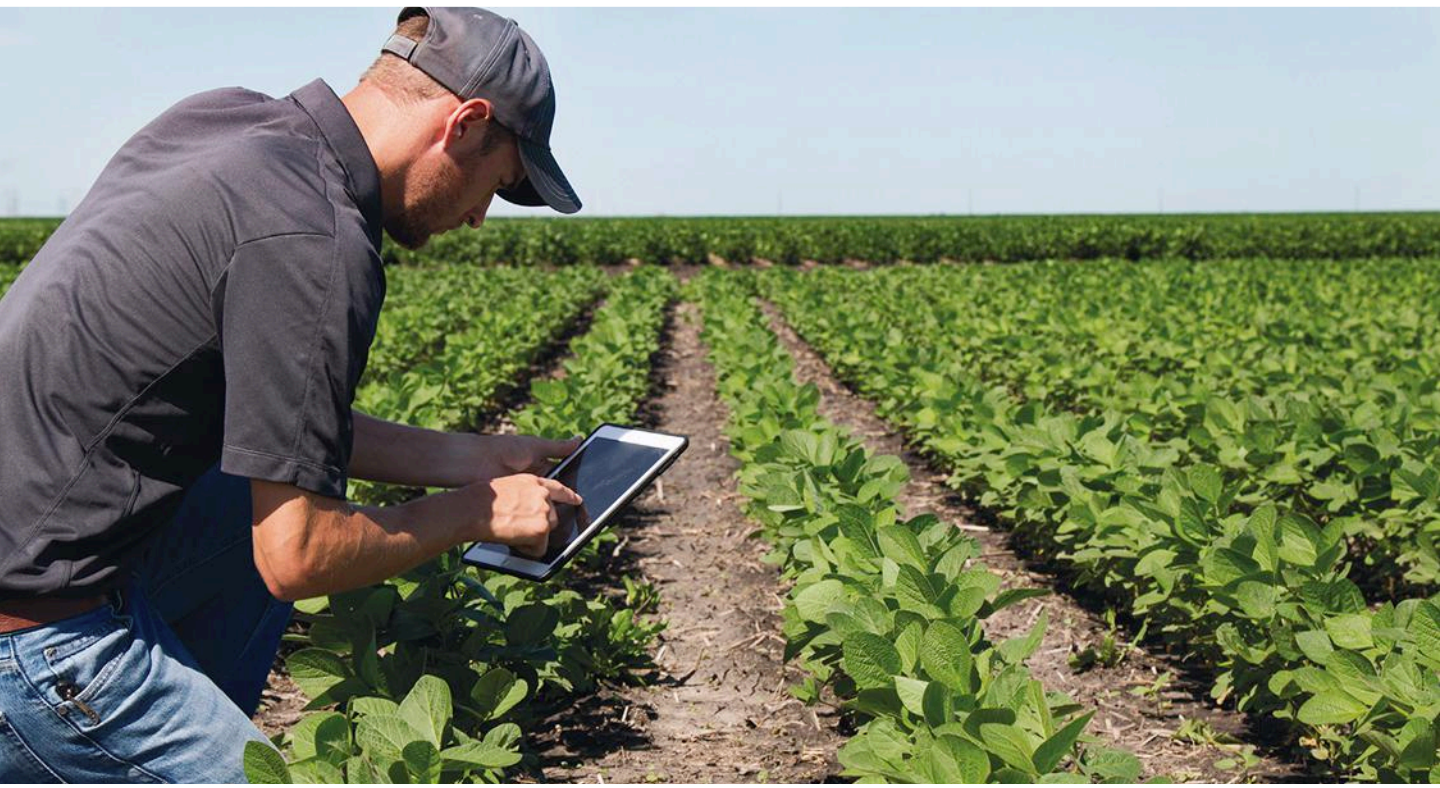

Farms are becoming more connected as farmers realize the potential of IoT technologies in helping them minimize operation cost while still achieving better results.

What follows are three smart farming applications already in play that take advantage of these technology solutions:

1. Livestock monitoring

Here, the IoT allows the entire livestock environment to be constantly monitored. The user is alerted by phone, text, or email if any condition falls outside of a preset parameter. Several systems containing some of the elements of an integrated monitoring system are now available commercially for swine, cattle, broiler, and milk production.

For example, a company called Moocall helps farmers monitor pregnant cows that are preparing to give birth. A battery-powered sensor detects motion associated with oncoming labor and then sends an alert SMS notification (*Fig. 1*). The battery lasts for up to 60 days; the device will send an alert when the battery level drops to 15%. The calving sensor is equipped with an embedded machine-to-machine (M2M) chip that can work over different networks.

1. Moocall works anywhere in the world with GSM signal. *(Courtesy of Moocall)*

Another livestock monitoring application is CattleWatch. This cloud-based hardware/software technology utilizes energy harvesting to power sensor and communication platforms to remotely monitor the health status and location of livestock. As a result, users have access to real-time data from their homes or offices or by smartphone. The CattleWatch system (*Fig. 2*) deploys hub collars that are placed on roughly 2% of the cattle. The remainder of the herd is outfitted with collar units or ear tags powered primarily by lithium batteries.

The hub collars feature miniaturized photovoltaic cells that harvest solar energy, which is then stored in Tadiran TLI Series industrial-grade lithium-ion (Li-ion) rechargeable batteries. TLI Series Li-ion batteries can operate for up to 20 years and 5,000 full recharge cycles, with an extended temperature range of −40 to 85°C and storage of up to 90°C. They're able to deliver high pulses (up to 15 A and 5 A continuous) so that the hub collars can communicate wirelessly with the Iridium satellite network.

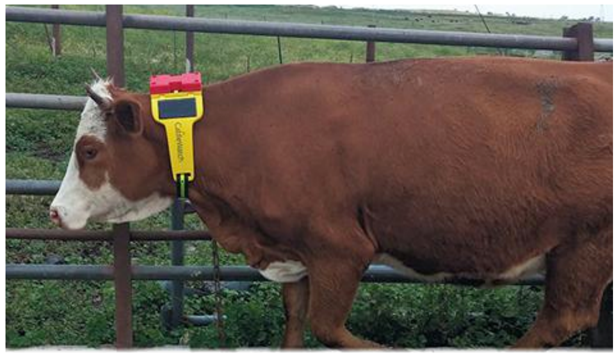

2. The CattleWatch system even broadcasts alerts if predatory animals or poachers are detected. *(Courtesy of CattleWatch)*

All of the regular collars communicate with the solar-powered hub collars to create an in-herd wireless mesh network. It provides valuable, near-real-time insight regarding animal behavior, including herd location, walking time, grazing time, resting time, water consumption, in-heat condition, and other health events. The system even broadcasts alerts upon detection of predatory animals or poachers.

2. Precision farming

With rugged and precise IoT sensors, farmers can collect data on weather, soil, air quality, and crop maturity, enabling them to make smarter decisions. For instance, a company called Cropx uses data and sensor devices to help farmers better understand water usage across their fields. The company also informs farmers about the amount of fertilizer and pesticide needed by each patch at specific times, automatically handling daily decisions for farmers. Algorithms and pattern-recognition technology are used to analyze the farmland and determine the various elevations (where they are hilly or flat).

"I think that it's very important to distinguish between data that is coming from above the soil and data coming from below," said Tomer Tzach, Cropx CEO, at the Forbes AgTech Summit. "Data from above the soil is coming from drones, satellite imagery, pictures that are coming from cameras that are put on weather stations, and so forth. And when you take imagery, you're always too little, too late because by the time

As advances push forward in IoT applications—ranging from enhanced battery-storage technology to the production of more complete IoT solutions from semiconductor companies—connected farms are destined to become more efficient and more productive.

Case IH's concept tractor uses onboard video cameras and LiDAR (light imaging, detection, and ranging) sensors to identify obstacles in its path.

you see something, it's already happening and the plant is already suffering. It's a much bigger challenge getting the data out of there, but I think that the data coming out of [the soil] can be much more useful in terms of being predictive."

Analog Devices Inc. (ADI) has been working on a very interesting project called "The Internet of Tomatoes" in its offices in Wilmington, Mass. This precision agriculture experiment leverages technologies like microelectromechanical systems (MEMS) and sensors. For example, inertial-measurement-unit (IMU) sensors, based on multiaxis combinations of precision gyroscopes, accelerometers, magnetometers, and pressure sensors, are used to figure out whether environmental monitoring could improve flavor. ADI integrates hardware solutions with a cloud-based IoT application from ThingWorx to develop a complete solution for farmers, providing them with apps and dashboards built on ThingWorx to better understand and implement improvements.

3. Autonomous tractors

Tractor manufacturers like John Deere and Case IH offer tractors to farmers that drive automatically. Self-driving tractors have been in the market for many years—even longer than semi-autonomous cars on the roads. One advantage of self-driving tractors is their ability to avoid reworking the same crop row by reducing the overlap to less than an inch. As a result, it takes fewer passes to cover each field, therefore saving farmers time and money. In addition, they can make very precise turns without the driver even touching the steering wheel.

Full autonomous tractors navigate using lasers that bounce signals off several mobile transponders located around the field. With supervised tractors, in contrast, a tractor is driven by a person, but followed by autonomous machinery. That machinery copies the steering and speed of the first tractor. This type of tractor offers the advantage of a reduction in human errors when performing tasks like spraying insecticide.

At the moment, no autonomous tractors exist on the market—but all major tractor makers have them in the works. For instance, Case IH (*Fig. 3*) planted soybeans with its concept autonomous tractor, and is taking the machine on a world tour to put it on display at farm shows. This driverless tractor won't be available in the marketplace for at least a few years. However, through experiments and showing it to real-world farmers, manufacturers are getting a better idea of what features and capabilities are desired in an autonomous tractor.

As advances push forward in IoT applications—ranging from enhanced battery-storage technology to the production of more complete IoT solutions from semiconductor companies—connected farms are destined to become more efficient and more productive. Not every farmer agrees with using these technologies, despite the fact that they offer real benefits to farmers. Maybe in the future, as IoT technologies mature, reticent farmers will realize how they can make such solutions adapt to their specific needs.

PROVEN
40
YEAR
OPERATING
LIFE*

Highly remote locations call for Tadiran batteries.

Battery replacement is costly and often dangerous work. Reduce the risk with Tadiran bobbin-type lithium thionyl chloride ($LiSOCl_2$) batteries. With an annual self-discharge rate of just 0.7% per year, Tadiran $LiSOCl_2$ batteries enable low power consuming wireless devices to operate for up to 40 years on a single battery, up to 4 times longer than the competition. Our batteries also feature the highest capacity, highest energy density, and widest temperature range of any lithium cell, plus a glass-to-metal hermetic seal for added ruggedness and reliability in extreme environments.

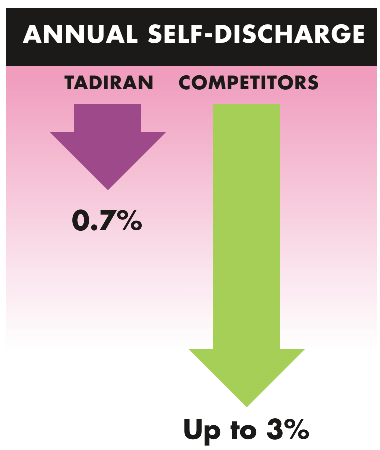

Take no chances. Take Tadiran batteries that last a lifetime.

TADIRAN BATTERIES

Tadiran Batteries
2001 Marcus Ave.
Suite 125E
Lake Success,
NY 11042
1-800-537-1368
516-621-4980

www.tadiranbat.com

* Tadiran $LiSOCL_2$ batteries feature the lowest annual self-discharge rate of any competitive battery, less than 1% per year, enabling these batteries to operate over 40 years depending on device operating usage. However, this is not an expressed or implied warranty, as each application differs in terms of annual energy consumption and/or operating environment.

ALAN GRAU | Icon Labs, alan.grau@iconlabs.com
www.iconlabs.com

What's the Difference Between Device Hardening and Security Appliances?

When it comes to protecting IoT devices from cyber attacks, each approach has supporters, but there are tradeoffs between "device-centric" and "appliance-centric."

Attacks on Internet of Things (IoT) systems continue to make headlines. All devices on publicly accessible networks are being targeted. While the use of IoT devices is increasing at an unprecedented rate, security for these vulnerable devices is painfully and unnecessarily lagging behind. After great financial expense from DDOS attacks and identify and data theft, awareness of the problem is finally growing. Studies utilizing ICS system honeypots have shown internet-connected ICS devices have been attacked within 24 hours of connection to the internet. In our discussions with customers, we commonly hear reports of newly provisioned IoT gateways being probed within 45 minutes.

Industry groups are developing standards, requiring certifications, and pushing legislations. Yet with the excitement to get new devices, software, and services into production, manufacturers continue to deliver products loaded with the security equivalent of a wing and a prayer.

Companies building IoT and other connected devices must ensure their devices are protected from these attacks. But where do they start? What steps can the device developers and manufacturers take to ensure their devices are protected?

Companies deploying IoT solutions and building IoT networks must make certain their networks are protected. Again, where do they start? What steps are needed and what solutions should they deploy? Can they rely on having strong security built into the devices they deploy? Or must they assume all endpoints have limited built-in security, and integrate them into a network relying upon using security appliances for protection?

Each approach has supporters, but there are tradeoffs between the "device-centric" and "appliance-centric" approaches to IoT cyber security.

BUILDING SECURITY INTO THE DEVICE

One approach to IoT security is to build protection directly into the device. This provides a critical security layer—the devices are no longer dependent on the corporate firewall as their sole protection. This is especially critical for mobile devices and IoT endpoints deployed in remote locations.

A security solution for IoT devices must provide protection against a wide range of cyber attacks. It must ensure the device firmware has not been tampered with, be able to secure the data stored by the device, secure in and outbound communications, and it must detect and report attempted cyber attacks. This can only be achieved by including security in the early stages of design.

While there is no one-size-fits-all security solution for embedded devices, solutions are available that provide a framework for OEMs *(Fig. 1)*. A security framework provides OEMs with the core capabilities required to protect their devices and the flexibility needed to customize the solution to the specific requirements of their device, while ensuring that critical security capabilities are included.

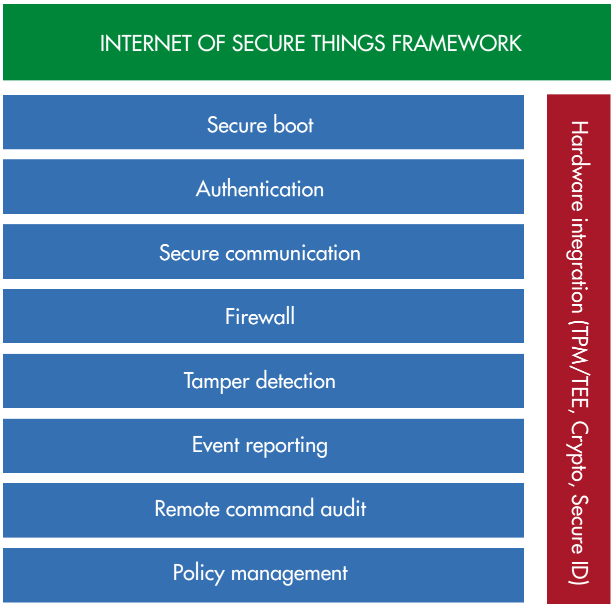

1. Here is a typical security framework for IoT devices.

DEVICE SECURITY REQUIREMENTS

Before selecting an IoT security framework, it is important to step back and look at the requirements at both device and system levels. Security requirements for IoT devices must take into consideration the cost of a security failure (economic, environmental, social, etc.), the likelihood of attack, possible attack vectors, and the cost of implementing a security solution.

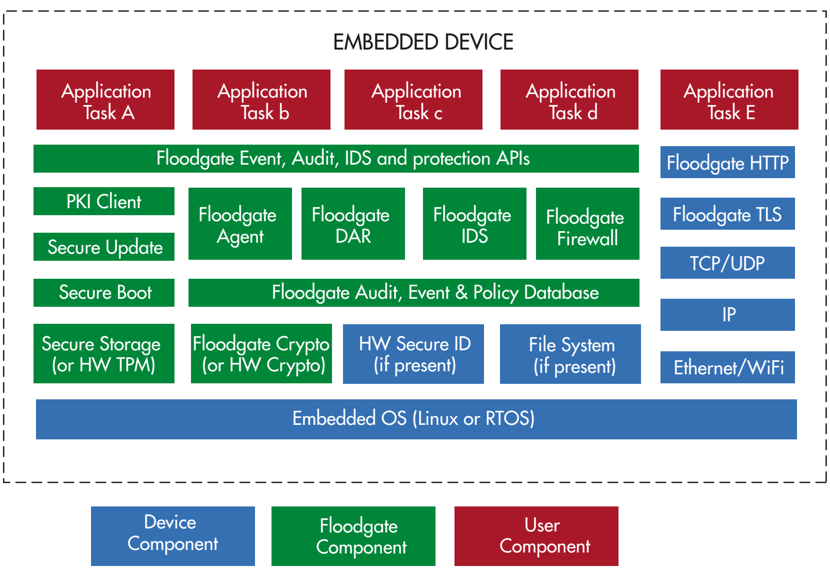

2. The Floodgate Security Framework provides an integrated suite of security building blocks.

Security capabilities needing consideration are:

- Secure boot
- Secure firmware updates
- Secure communication
- Data at-rest protection
- Embedded firewall and intrusion detection
- Key and certificate management
- Authentication
- Integration with security management systems
- Security policy management

SECURITY EVENT REPORTING

A security framework, such as the Floodgate Security Framework *(Fig. 2)*, provides an integrated suite of security building blocks.

SECURE COMMUNICATION

When most engineers think of security, they typically think of secure communication protocols such as SSL/TLS, SSH, and IPSec. In recent years, support for secure communication has been added to many embedded devices. While these protocols provide a first level of defense against protocol-based cyber attacks, they leave other attack vectors unprotected.

Security protocols are designed to protect against packet sniffing, man-in-the-middle attacks, replay attacks, and unauthorized attempts to communicate with the device providing a good starting point for building secure devices.

Small IoT edge devices are adopting wireless protocols such as zigbee, Bluetooth Low Energy (BLE), and other wireless and mesh networking protocols. These protocols have some built-in security. However, it is relatively weak and exploits have been published. Small IoT devices typically run on very low-cost, lower-power processors not supporting TLS or IPSec. For small edge devices, DTLS, which is TLS over UDP, can be used for secure communication.

SECURE BOOT AND SECURE FIRMWARE UPDATES

Secure boot and secure firmware update capabilities ensure an IoT device is running authorized code from the device manufacturer preventing the installation of malware or code modified by hackers *(Fig. 3)*.

Secure boot begins with a first-stage bootloader programmed into a protected or non-writable storage location on the device. This first-stage boot loader validates the authenticity of the second-stage boot loader. The

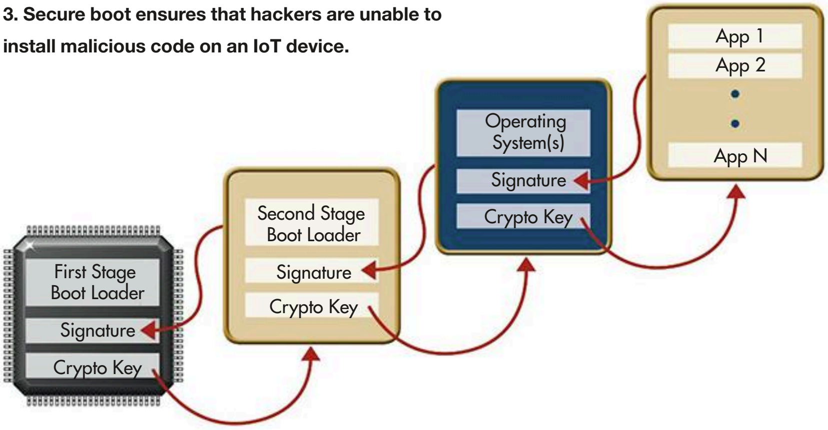

3. Secure boot ensures that hackers are unable to install malicious code on an IoT device.

Most cyber attacks occur in phases, beginning with hackers probing a network looking for, finding, and exploiting a vulnerable device.

second-stage boot loader, which can be more complex and may be stored in reprogrammable flash memory, repeats the process, verifying the operating system and applications are valid.

Secure boot relies on signed code images to enable validation of the image during the secure boot process. The code images are signed by the device OEM using the OEM's private key. The OEM's corresponding public key is used by the device to validate the signature for the firmware image.

Secure firmware update, like secure boot, validates new code images that have been signed by the OEM during the upgrade process. If downloaded images are not valid, they are discarded and the upgrade is not performed. Only valid images are accepted and saved to the device.

DATA-AT-REST (DAR) PROTECTION

IoT devices, unlike enterprise servers, are not locked away deep in a data center. Many are located in the field with the risk of theft or physical attack. Any sensitive data stored on such a device should be encrypted, ensuring it is protected from attempts to read from the device, either by copying the data from the device, or by physically removing the flash drive and reading data directly.

Data-at-rest (DAR) protection encrypts data stored on the device, providing protection against these attacks. Many IoT devices don't have the computing power to support full disk encryption, but sensitive data such as credit-card numbers or patient information should always be encrypted. Care must be taken to store the encryption key in protected memory on the device or in a secure location such as a USB drive or network server.

The DAR solution should ensure unencrypted data is never stored on the hard drive. Protected data should be encrypted before it is written to a file. Encrypted files should be encrypted in memory and remain in RAM, never written back to the file system without being encrypted ensuring data cannot be leaked due to a power failure.

INTRUSION DETECTION

Many embedded devices lack basic security features, making them easy targets for hackers. As a result, hackers have specifically target embedded devices. Devices such as point-of-sale systems, HVAC systems, and medical devices have been exploited.

Most cyber attacks occur in phases, beginning with hackers probing a network looking for, finding, and exploiting a vulnerable device. Once this initial beachhead is established, hackers use the exploited device to probe deeper into network. The cycle repeats with hackers gradually expanding their reach within the network. Stopping the attacks begins with early detection. Intrusion Detection Systems and Intrusion Detection Software (IDS) are commonplace in enterprise networks and PCs. IDS, as the name implies, detects when a system is under attack or being probed. These solutions can take many forms and detect many different types of attacks, but regardless of form, are in the main, largely absent for embedded devices.

Adding IDS capabilities to embedded devices is critical to providing early warning of a cyber attack. The ability to detect and report potentially malicious activity enables system administrators to take action to block attacks, quarantine compromised systems, and protect their networks. If embedded devices can support basic IDS they will no longer be easy targets for hackers.

PKI AND CERTIFICATE-BASED AUTHENTICATION

A well-known and tested security solution has recently seen a dramatic rebirth in the IoT recently. PKI (Public Key Infrastructure) is a set of technologies and services for managing authentication of computer systems. PKI certificates are very useful in high-security situations. For example, suppose that you needed to securely transmit data between two networked devices. How do you really know you are transmitting the data to the intended device and not to an imposter? One way of ensuring the integrity of the transaction is to use digital

certificates to prove the identities of both machines. Without getting into the details of the public/private key cryptography technology that makes this possible, an IIoT device can verify the certificate holder is the entity specified by the certificate. These services are enabled using public/private key cryptography providing the technical underpinnings of PKI. The result, which is what really matters, is a device can verify, with cryptographic certainty, the holder of the PKI certificate is really who it claims to be and not an imposter.

CRYPTOGRAPHY AND SECURE KEY STORAGE

Secure communication protocols, data-at-rest protection, secure boot, and secure firmware updates all rely on encryption and certificate-based authentication. A security framework must provide support for the cryptographic algorithms used by these features. It must also provide the ability to securely store the encryption keys and certificates used to encrypt data, authenticate firmware, and to support machine-to-machine authentication. Secure key and certificate storage is a critical requirement. If a hacker can discover the encryption keys, they can completely bypass an otherwise robust security solution.

HARDWARE SECURITY MODULE SUPPORT

Many new IoT platforms include a hardware security module providing secure key storage, protected memory regions, and cryptographic acceleration. A security framework must be designed to allow easy integration with hardware-based security features.

Likely candidates for hardware security module support include PUF (Physically Unclonable Functions), security coprocessors such as TPMs (Trusted Platform Modules), and Trusted Execution Environments such as ARM's TrustZone.

PUF uses random patterns in the silicon to differentiate chips from each other and creates a unique random number. The generated random number is used to seed a strong device ID and cryptographic keys creating a hardware root of trust.

Security co-processors are physically separate chips offering true isolation of private keys. A TPM is an industry-standards-based securing chip that offers isolation and facilities for the secure generation of cryptographic keys, and limitation of their use, and true random-number generation. It also includes capabilities such as remote attestation and sealed storage. Its capabilities come at a price, usually moving deployment to higher-end IoT devices.

A hardware security module (HSM) is another physically separate chip and likely at a lower cost than a TPM. Like the TPM, it safeguards and manages digital keys for strong authentication and provides crypto processing. An HSM traditionally comes in the form of a plug-in card or an external device attaching to the protected device, making it somewhat less suited to an IoT device. Depending upon the perceived and likely threat vectors, an HSM may provide an effective solution.

Trust Zone is a single-chip solution segregating execution space into secure and insecure worlds. Insecure apps can't access security-critical assets. Those same security critical assets are isolated from tampering.

IoT SECURITY: THE SECURITY APPLIANCE APPROACH

Security appliances also play a central role in protecting IoT networks from cyber attacks. IoT network architectures are diverse and include a range of devices and computing resources. Not surprisingly, there are equally diverse sets of security appliances for IoT networks. Most of these approaches fall into three main categories, protecting the network and cloud, IoT-specific intrusion detection, and protecting legacy devices.

Network and cloud protection

As with traditional IT networks, security appliances provide a critical layer of defense at the network perimeter and for the data center. The frequency and sophistication of cyber attacks targeting data centers and cloud-based computing resources continues to increase and many new IoT services and connections open up fresh attack vectors for hackers targeting these systems. Network security appliances must not only be deployed to protect these devices, but must also be constantly updated to secure new IoT protocols and services.

Intrusion Detection Systems (IDS) for IoT networks

The deployment of new protocols and services to meet IoT requirements results in new attack vectors hackers can exploit. Companies are developing new network IDS solutions to detect attacks against newer services and protocols.

In some cases, existing network IDS solutions can be enhanced to detect new attacks. These solutions work well for detecting attacks occurring at the network edge or data center, where existing network IDS solutions are deployed.

For mobile or remotely deployed IoT devices, however, these solutions add little value. New types of IDS solutions are required to detect attacks targeting remote IoT endpoints.

There are several challenges to detecting attacks targeting IoT endpoints in the field. The IDS appliance itself must be designed to operate in the same location as the IoT endpoints. In many cases, this requires physical hardening of the device, allowing operation in harsh environments.

The IDS appliance must detect new attacks, many of which are emerging or will emerge in the coming years. They must also support IoT new protocols. Any appliance designed today must be flexible enough to provide protection against new attacks as they emerge.

Finally, economic factors must be taken into considerations. The physical footprint of an IoT network may require deployment of a large number of IDS appliances. Unfortunately, the cost model of many solutions makes them prohibitive for this model.

Protecting legacy devices

Many legacy devices and systems are being connected to the IoT through gateways and proxy services, or using existing network connectivity. Most were manufactured with inadequate security.

Unfortunately, the upgrade process may be difficult, expensive, or impossible. Some devices cannot be upgraded without being returned to the factory. In some cases, the manufacturer may no longer support the device, or may be out of business. Replacing the devices is often simply too expensive to be an option and newer devices may not yet be available with improved security.

For devices and systems that cannot be easily or affordably replaced or upgraded, a "bump-in-the-wire" appliance solution can provide the required security. This type of solution can protect legacy devices that are otherwise vulnerable. The bump-in-the-wire appliance provides security by enforcing communication policies, ensuring only valid communication is allowed with the protected device.

The security appliance must provide the ability to configure communication policies, a set of rules specifying which packets are processed and which are blocked. Smart-grid devices may only need to communicate with a small number of other devices. This can be enforced using communications polices that restrict communication to only what is required. Communication policies define who the device is allowed to talk to, what protocols are allowed, and what ports are open. The policies are then encoded as firewall rules. Rules can be set up to block or allow packets by IP address, port, protocol, or other criteria. Some firewalls support advanced rules allowing additional fine-grained control over the filtering process. The security appliance then filters messages before the device processes the messages, allowing only communication with known, trusted devices.

In a system without a security appliance, a hacker may attempt to remotely access the device using default passwords, dictionary attacks, or stolen passwords. Such attacks are often automated, allowing a huge number of attempts to break the system's password. The same system can be protected by a firewall configured with a whitelist of trusted hosts. The firewall's filters will block attacks from the hacker before a login is even attempted because the IP or MAC address is not in the whitelist, thereby blocking the attack before it even really begins *(Fig. 4)*.

Many attacks are blocked before a connection is even established because each packet received by the devices must pass through the firewall for filtering before being processed. This provides a simple, yet effective layer of protection currently missing from most legacy IoT devices.

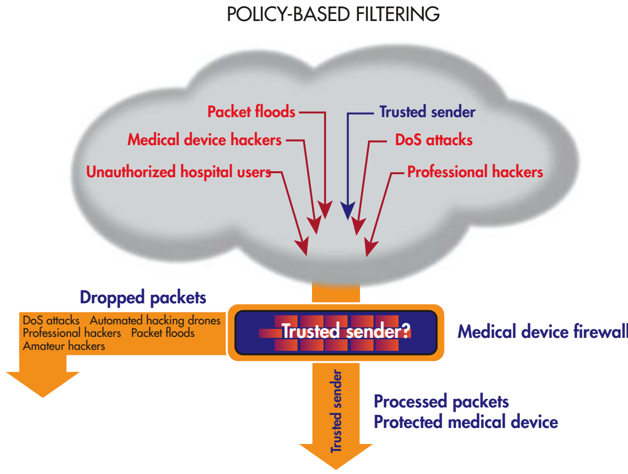

4. Rules-based filtering is used to enforce communication policies, blocking packets from non-trusted senders and isolating devices from attack.

SECURITY APPLIANCE APPROACH VS. DEVICE HARDENING

Two important tradeoffs in considering the hardware versus software approach to IoT security are economic consideration and the protections that can be built into low-cost sensors.

As IoT devices proliferate, the number of required security appliances could explode. The economics of adding security appliances to every IoT device are simply prohibitive.

While this can be addressed with software security built directly into the device itself, this is not without cost of its own. Security software requires additional memory and processing power, and imposes additional overhead on network resources which can dramatically impact battery life for lower power devices. As a result, you are limited in how much security can be added to low end devices such as sensors.

SUMMARY

One of the unique challenges of the IoT is that the network perimeter is often blurry. Network security appliances can protect cloud-based computing resources and any IoT devices that happen to reside within the network perimeter, but do little to protect mobile devices or IoT endpoints located in the field. So while security appliances play a critical role in protecting the IoT, they do not provide the complete solution.

Ultimately, some combination of hardware and software will be required, but building software into IoT devices is a critical missing piece that must be addressed. ed

ALAN GRAU is the president and cofounder of Icon Labs, a leading provider of security solutions for embedded devices. You can reach him at alan.grau@iconlabs.com

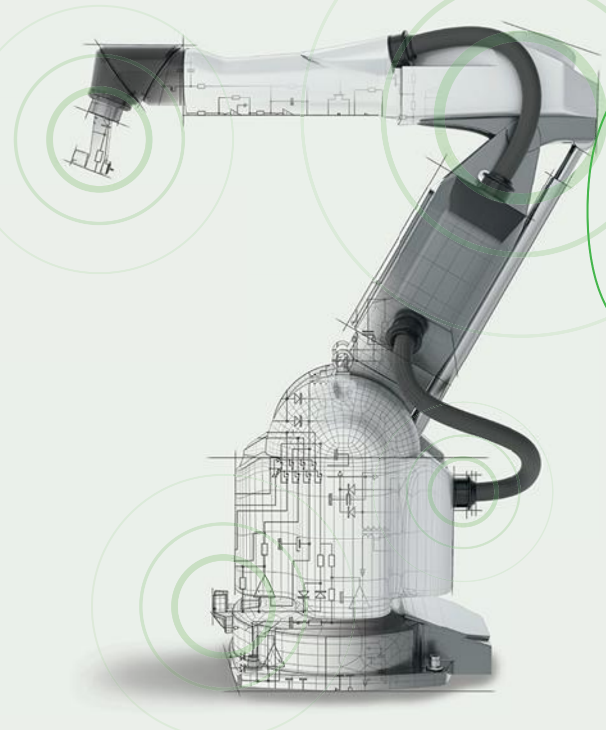

WILLIAM WONG | Technical Editor

NVMe over Fabric Addresses Hyperscale Storage Needs

Hyperscale computing requires hyperscale storage. NVMe over Fabric (NVMe-oF) is one technology that is addressing this need.

Supercomputing or high-performance computing (HPC) has massive storage needs. The hyperscale approach takes a modular approach using high-speed connectivity to link computing, network, and storage resources. This allows resources to be added and adjusted as needed while simplifying management of the system.

These days the fastest processor interface, short of the memory channel, is PCI Express. The Non-Volatile Memory Express (NVMe) interface is based on PCI Express. NVMe is latency-optimized and it is becoming the preferred way to access flash memory because it has less overhead than conventional disk interfaces like SAS and SATA. NVMe performance can also scale because of PCIe's multilane approach. Typically NVMe systems have one (x1) to eight lanes (x8). For example, the M.2 socket supports x4 PCIe/NVMe devices.

NVMe can provide high-speed connectivity to flash memory but other fabrics like Ethernet and InfiniBand provide longer-distance connections than PCI Express. These fabrics have been used to support storage protocols like iSCSI, but these protocols tend to have higher overhead than NVMe so it made sense to develop NMVe-over-Fabric (NVMe-oF). NVMe-oF uses the same NVMe transport layer as a native NVMe implementation *(see figure)*.

NVMe-oF can offer remote direct memory access (RDMA) support that is available with fabrics like Ethernet and InfiniBand. NVMe-oF works with other fabrics like Fibre Channel. Fibre Channel's native transport system can handle NVMe without the need for RDMA support. Each NVMe-oF fabric has its advantages and disadvantages but operating systems and applications will use any of these transparently since they share a common NVMe transport layer.

NVMe-oF promises to provide the scalability needed for hyperscale environments. It provides a significant boost in performance compared to non-NVMe solutions. The choice of which fabric will likely depend upon the existing infrastructure. NVMe is already pushing conventional drive interfaces out because of performance advantages, and NVMe-oF will be doing the same but on a much larger scale.

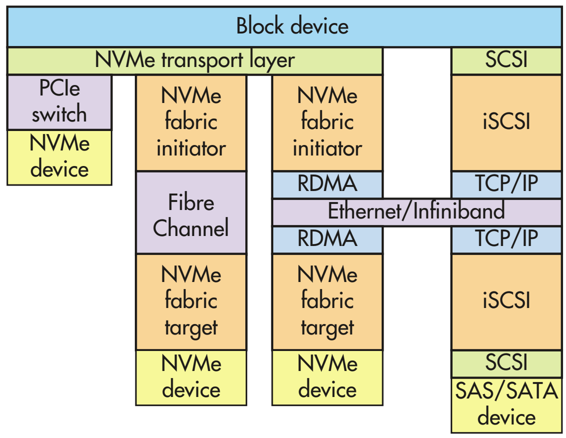

NVMe devices can be accessed through direct connections (left) or using NMVe-over-Fabric with Fibre Channel or RDMA and Ethernet or InfiniBand. NVMe-oF has significantly less overhead than conventional fabric-based connections like iSCSI (right) that employs TCP/IP.

NVMe-oF is relatively new and replacing existing storage area networks (SANs) will take time but if the rapid adoption of NVMe is any indication, then NVMe-oF will likely follow a similar pattern. Providing a unified block storage environment simplifies system configuration and management while providing better performance.

NVM Express Inc. manages the NVMe-oF specification. This includes the use of RDMA with Ethernet and InfiniBand. The NVMe over Fabrics using Fibre Channel (FC-NVMe) was developed by the T11 committee of the International Committee for Information Technology Standards (INCITS).

NVMe-oF solutions were presented at this year's Flash Memory Summit. Unfortunately, a fire prevented the show floor from opening so most were unable to see the systems in action but they are available. Still, it will be a few years before NVMe-oF becomes commonplace but it will likely become the norm for large systems including HPC applications.

WILLIAM WONG | Technical Editor

Supercomputer Module Doubles Performance for Faster Neural Nets

NVIDIA's Jetson TX2 doubles the performance of its predecessor, or it can provide the same performance at half the power.

NVIDIA's Jetson TX1 was an amazing platform that delivered supercomputer performance in a compact module *(see "Module Delivers Supercomputer Performance" on electronic design.com)*. The 256-core GPU was based on NVIDIA's Maxwell architecture and paired with 64-bit ARM cores.

The latest Jetson TX2 *(see figure, opposite page)* doubles the performance of the Jetson TX1. Alternatively, developers can get the same performance out of the TX2 as the TX1 while cutting the amount of power in half to just 7.5 W. Each approach has its merits. This might double the runtime of a battery powered Jetson TX2 device or it could do twice the work in the same power envelope. For example, it might track twice as many objects or handle two camera input streams instead of one.

The Jetson family supports the CUDA programming environment, as well as deep learning and deep neural nets (DNN), courtesy of the Cuda DNN (cuDNN) runtime targets.

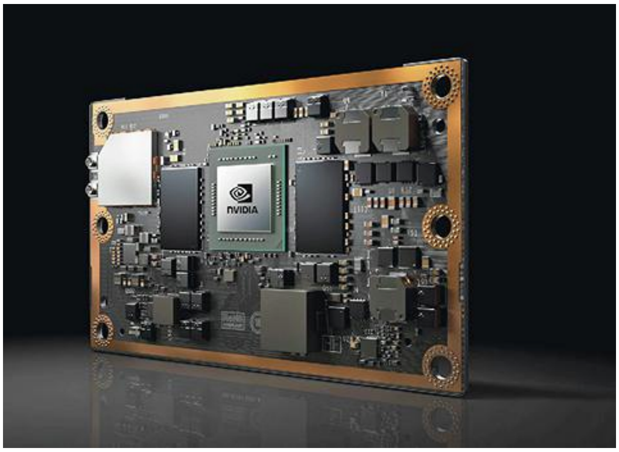

The latest Jetson TX2 doubles the performance of the Jetson TX1.

The cuDNN software can support DNN frameworks like the open-source system from TensorFlow *(see "GPUs and Deep Learning" on electronicdesign.com)*. The modules are small enough to work in midsize and larger drones, providing image recognition and planning support that would not be possible with a less power microprocessor.

The Jetson TX2 is the same 50-mm by 87-mm module size as its predecessor. This allows it to plug into the same carrier boards like ConnectTech's Orbitty that I tested with the Jetson TX1 *(see "Carrying a Jetson TX1" on eletronicdesign.com)*. The Orbitty provides connections for gigabit Ethernet, USB 3.0, USB 2.0 OTG, HDMI, two 3.3V UARTs, I2C, and four GPIOs. It also has a microSD socket. It is designed to operate in temperatures from −40° C to +85° C. The power voltage is between 9 VDC and 14 VDC.

The Jetson TX2 gets it performance from the 256-core NVIDIA Pascal GPU. The architecture was first released in the NVIDIA Tesla P100. The P100 uses High Bandwidth Memory 2 (HBM2) and Chip-on-Wafer-on-Substrate (CoWoS) technology for 5 TFLOPS of double-precision performance and almost 10 TFLOPS of single precision performance. The Jetson TX1 delivers 1 TFLOP of single precision performance while the Jetson TX2 doubles that. Of course, the P100 uses 250 W when running full out.

The Jetson TX2 also has a pair of 64-bit NVIDIA Denver 2 ARM-compatible cores, plus four 64-bit ARM A57 cores. The module doubles the amount of memory to 8 Gbytes of LPDDR4 and a 32 Gbyte eMMC flash module. It retains the 802.11ac WLAN, Bluetooth, and 1-Gbit Ethernet links.

The module is available for $399 in quantities of 1,000. The Jetson TX2 Developer Kit is $599 with an education version available for $299.

WILLIAM WONG | Embedded/Systems/Software Editor

SnapEDA Helps Make Hardware Design Chores a Snap

SnapEDA is a website that provides free schematic models, PCB footprints, and 3D models for a growing list of electronic components.

Have you created a PCB footprint and schematic symbol from scratch for a component that you needed in your hardware design?

Schematic and printed-circuit-board (PCB) design tools come with a wide range of component models and schematic symbols. Most can be customized and utilized with new components, but this takes time and effort. This is especially true when it involves 3D models, simulation information, and so on. This level of functionality may not be needed for a basic PCB layout. However, it becomes invaluable when doing advanced analysis such as thermal simulations.

Managing a large library of components can also be a challenge. Building and maintaining this library becomes much easier if component information is provided by a third party instead of having to be created or recreated. Component vendors often provide this information in one form or another and there are a number of standard exchange formats, but no single format provides all of this information or is used by all tools.

SnapEDA (*https://www.snapeda.com*) is a website that helps simplify the chore of designing a system with new components. It provides free schematic models, PCB footprints, and 3D models for a growing list of electronic components. Designers can simply search for a part on the website and then check out the details *(Fig. 1)*.

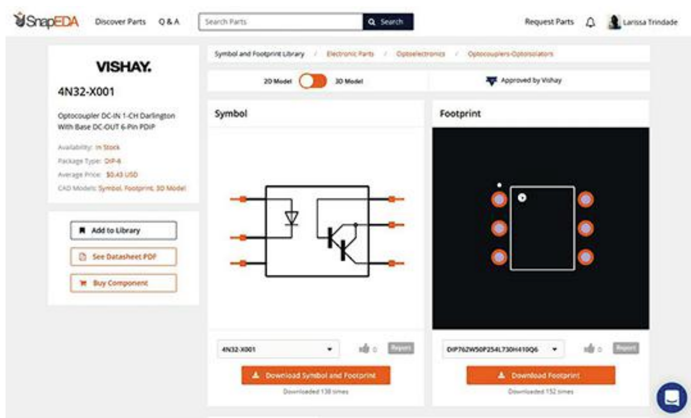

1. SnapEDA's website allows designers to search for components and then examine the details from schematic diagrams to 3D models.

Developers can download information to a range of platforms, including Eagle, Altium, KiCad, OrCAD, Cadence Allegro, Pulsonix, and PADS. Such functionality highlights the value of SnapEDA. Vendors may provide one or more formats, but SnapEDA handles conversion to other formats. It also delivers the proper content to developers.

Users can also view details like the 3D models *(Fig. 2)*. As noted, 3D outline and PCB layout details are just part of the information mix users are able to access. Datasheets and even product availability from a number of sources is available with most parts. It's also possible to purchase parts from a wide selection of partners like Digi-Key, Arrow, and Mouser, or directly from some vendors. The website even provides average pricing and what parts are popular in current designs.

SnapEDA's verification checker is a tool that gets applied to new components. It checks for manufacturing issues such as, for instance, the silkscreen doesn't overlap pads in the PCB information. Not all components have been checked, since new components are always being added. It's possible to request a check on a part if it hasn't already been done. The checker's "report card" for a component is available with the other information about a part, e.g., the datasheet and schematic symbol.

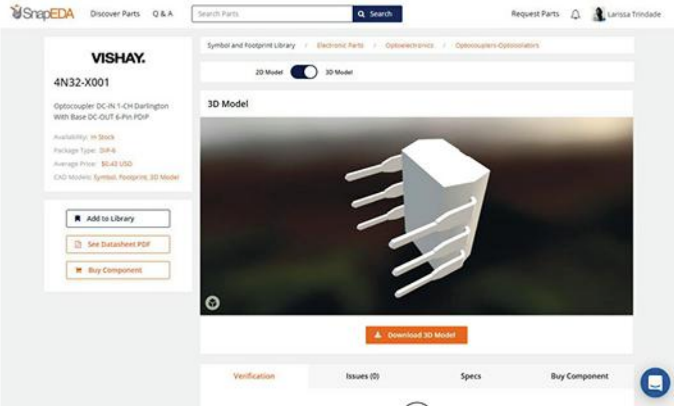

2. In addition to schematic and 2D PCB footprints, SnapEDA offers 3D models.

Of course, SnapEDA's large database makes it easy to find related parts. Downloads are free if the part is already in its library. New parts will be created for only $29, and $79 for a 3D model. Users that sign up can create libraries and download selected component data in bulk.

Obtaining the information provided by SnapEDA may seem like a trivial exercise—until it needs to be done for hundreds of components in a project. Even if SnapEDA only delivers a large fraction of the components you use, it will save a significant amount of time and effort. Best of all, it's free. ed

FOCUSED

& #FIERCELYINDEPENDENT

Good work comes from determination, focus and free thinking. As engineers you know this. We hold ourselves to the same truth – a determination to bring you **bold insights** and information that **helps you do your jobs better** - with a promise to deliver **unbiased, expert content**. Penton Design Engineering and Sourcing is the only independent insight and information source holding the promise of your truth.

Penton Design Engineering and Sourcing

Electronic Design. Machine Design. Power Electronics. Microwaves&RF. SourceESB.

DAN HARMON | Current and Power Measurement Automotive Marketing Manager, Texas Instruments

Minimizing Temperature Drift in Your Current Measurement

Use a current-sense amplifier that integrates a precisely matched, resistive gain network to minimize the temperature-drift effects of the gain error.

As more systems become electrified, thermal management has turned into one of the "hottest" issues facing designers. Using current measurements for thermal management is a leading indicator of system performance and faults, whereas simply monitoring the temperature is potentially a lagging indicator. Accurately monitoring the current consumed, especially over temperature, has become vital as designers pack more functionality into tighter areas.

While room-temperature calibration tends to be relatively straightforward, performing multi-temperature calibration is time-consuming and costly. Identifying ways to minimize the effects of temperature on current measurements can improve system performance and minimize system design margins, as well as potentially lower the total cost of ownership (TCO).

SOURCES OF ERROR IN CURRENT MEASUREMENTS

As I stated in my September 2015 article, "Mitigate Error Sources to Maximize Current-Measurement Accuracy" *(see http://www.electronicdesign.com/test-measurement/mitigate-error-sources-maximize-current-measurement-accuracy),* there are multiple contributing sources of error in current-measurement applications. In the article, I listed these sources of errors:

Amplifier-related errors:
- Input offset voltage (VOS) and VOS drift
- Common-mode rejection ratio (CMRR)
- Power-supply rejection ratio (PSRR)
- Gain error and gain drift

System errors:
- Gain-setting network tolerance, matching, and drift
- Printed-circuit-board (PCB) layout
- Shunt-resistor tolerance and drift

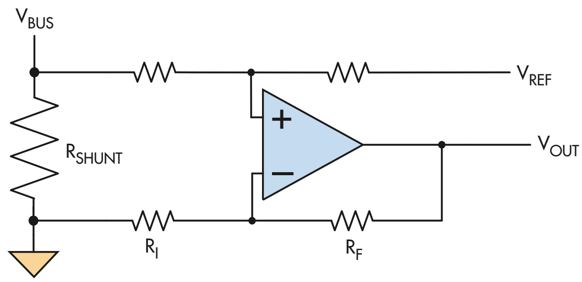

1. This basic differential amplifier configuration is used for current measurement.

You can see that "drift" is part of four of the seven items on both lists, which emphasizes the importance of minimizing the additional errors caused by temperature in a current-measurement implementation.

DISCRETE CURRENT-MEASUREMENT IMPLEMENTATIONS

Many system designers choose a discrete amplifier and external gain network for their low-side current-measurement applications because it's viewed as a low-cost alternative. There are two options when using a discrete circuit for low-side current sensing: a single-ended or differential configuration. *Figure 1* shows the latter.

In either configuration, the gain of the system is defined as $G = R_F/R_I$. The worst-case initial (or room-temperature) gain error is simply the tolerance of the discrete gain resistors. Assuming an application with a gain of 20, where $R_F = 100$ kΩ and $R_I = 5$ kΩ, *Table 1* reveals how that looks for different resistor tolerances.

To understand the effects of temperature, let's assume the same gain implementation with nominal resistor values and then apply various standard temperature coefficients (tempcos). Tempcos are usually specified in terms of parts per million per degree Celsius (ppm/°C). To get from ppm to a per-

TABLE 1: ERROR CALCULATIONS FOR VARIOUS GAIN-RESISTOR TOLERANCES							
Resistor tolerance (%)	Nominal R_F (kΩ)	Nominal R_I (kΩ)	Nominal gain	Minimum gain	Error (%)	Maximum gain	Error (%)
1	100	5	20	20.40	2.0	19.60	–2.0
2	100	5	20	20.82	4.1	19.22	–3.9
5	100	5	20	22.11	10.5	18.10	–9.5
10	100	5	20	22.44	22.2	16.36	–18.2

TABLE 2: ERROR CALCULATIONS FOR VARIOUS GAIN-RESISTOR TEMPCOS							
Resistor tempco	Nominal R_F (kΩ)	Nominal R_I (kΩ)	Nominal gain	Min. gain at 125°C	Error (%)	Max. gain at 125°C	Error (%)
50	100	5	20	20.20	1.01	19.80	–1.00
100	100	5	20	20.40	2.02	19.60	–1.98
250	100	5	20	21.03	5.13	19.02	–4.88
500	100	5	20	22.11	10.53	18.10	–9.52

centage, simply divide by 10,000. *Table 2* shows the calculation results for a variety of resistor tempcos.

The issue here is that the drift could be in opposite directions and different for each resistor, which means minimizing these temperature effects in gain errors requires multipoint calibration (as well as temperature monitoring) to enable temperature compensation in the system management controller. The additional circuitry and resources (calibration means people and time!) can significantly drive up the TCO. Plus, as the tempcos drop, the resistor becomes more expensive, and in turn drives up the TCO.

CALIBRATION

As I mentioned, performing system calibration can minimize (or even eliminate) initial errors. Typical calibration occurs at the total system level and involves two-point calibration. Performing two-point calibration enables you to minimize the error at both low current levels (due to the offset of the system), as well as at higher current levels where gain error is the primary error source.

A video training series from Texas Instruments titled, "Getting Started with Current Sense Amplifiers," *(see https://training.ti.com/getting-started-current-sense-amplifiers)* beginning with session 2.1, discusses this topic in detail. Equation 1 describes the transfer function of the system:

$$V_{OUT} = (I_L \times R_{SHUNT} \times G) + V_{OFFSET} \qquad (1)$$

The intent of calibration is to identify the factors needed to correct for the amplifier's offset voltage (V_{OFFSET}), the gain error of the network (G), and the shunt resistor error (R_{SHUNT}). Applying zero current calculates the necessary offset correction. Using a current close to the nominal load will calculate the total gain correction factor, accounting for both the shunt tolerance and amplifier gain network.

Minimizing the offset error is very straightforward—choose an amplifier with low V_{OS} and low V_{OS} drift. You will have to make a tradeoff between performance and cost. As discussed above, trying to perform two-point calibration at multiple temperatures to calculate the drift component of the offset error is costly and time-consuming.

CURRENT-SENSE AMPS INTEGRATE PRECISION MATCHED GAIN NETWORKS

Current-sense amplifiers are specialized integrated circuits designed specifically to measure current. As shown in *Fig. 2,* these amplifiers have four fundamental benefits compared to a discrete implementation.

The benefit of the precision, matched gain-resistor network really manifests itself in how the gain error changes over temperature. The way the resistors are designed and fabricated on-chip enables near-perfect matching of both the initial error as well as the temperature drift. The absolute accuracy of the resistors isn't a factor in minimizing gain

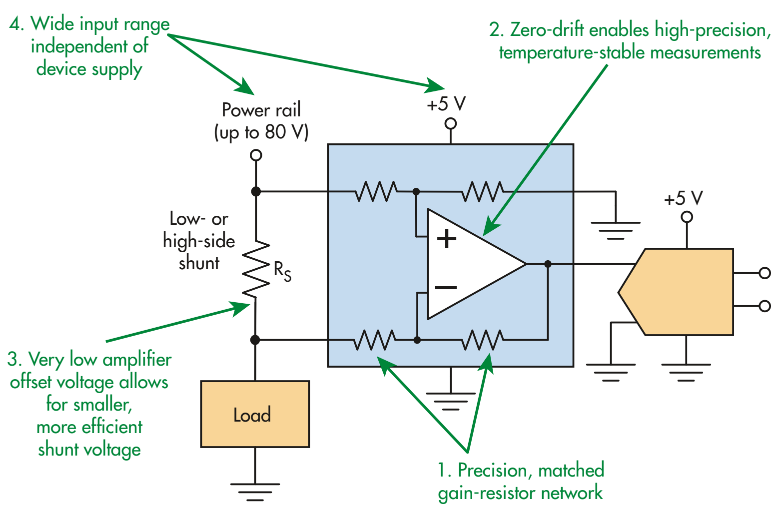

2. A current-sense amplifier offers several key benefits over a discrete current-sense implementation.

DESIGN NOTES

Quad IO-Link Master with Higher Current SIO Channels

Design Note 566

Eric Benedict

Introduction

IO-Link is a communication standard for a point-to-point, 3-wire interface to smart sensors and actuators found in industrial applications. IO-Link extends the traditional interface capabilities of these devices from a simple NC/NO switch interface (standard IO or SIO mode) to a bidirectional intelligent interface capable of sending additional information via coded switching at one of three different speeds (COM1—4.8kb/s, COM2—38.4kb/s or COM3—230.4kb/s). In addition to the data pin (C/Q), the IO-Link Type A interface has a 24VDC power supply pin (L+) and a common return pin (L–).

When an IO-Link master powers up, it interrogates each connected device to determine the proper operational mode for the device: SIO, COM1, COM2 or COM3. This allows for a mixture of legacy and IO-Link enabled devices to operate seamlessly in the same system.

The LTC®2874's rated CQ output current is 110mA. Higher currents up to 440mA may be obtained by paralleling channels. While this exceeds the IO-Link specifications, some non-standard SIO applications may require even larger currents to be sourced and/or need to maintain the functionality of the four independent channels. This article shows how to repurpose the LTC2874's hot swap channels to source larger currents for SIO loads (referred to as SIO+ mode) while maintaining the IO-Link features and capabilities of the LTC2874.

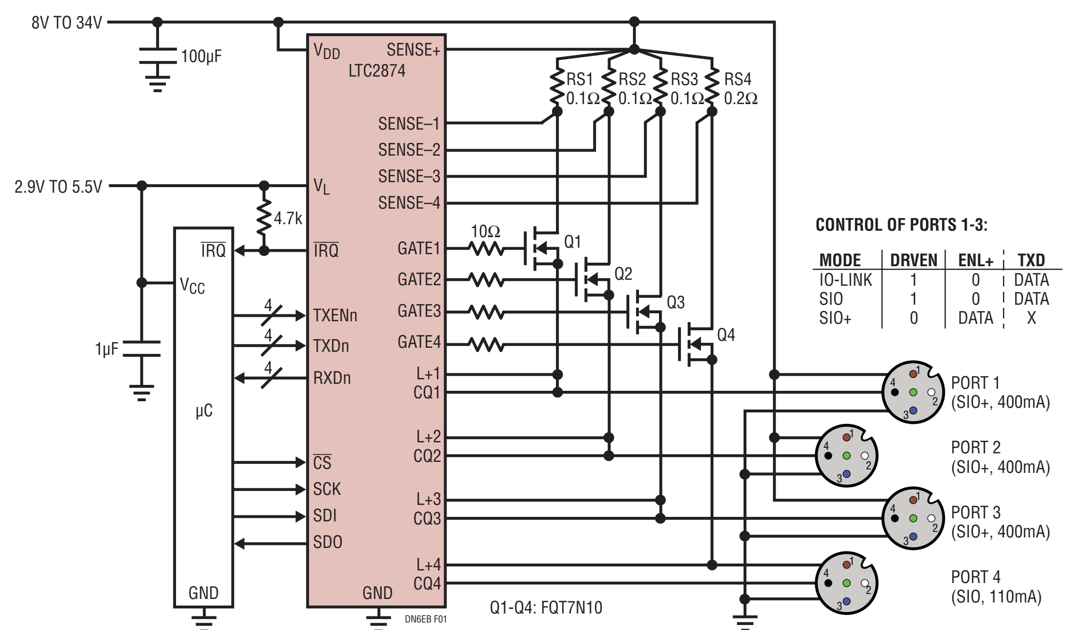

Figure 1. LTC2874, Quad IO-Link Master Configured with Three High Current SIO Ports (SIO+) and One Normal SIO Current Port with L+ Hot Swap

Circuit Description

Arbitrarily large sourcing currents are available in SIO+ mode by connecting the channel's hot swap controller output to its corresponding CQ pin as shown by ports 1–3 in Figure 1. For the high current port(s), the hot swap feature for L+ is not available; however, an external hot swap controller can be added for applications where this is desired. LTC2874 hot swap controllers not used for SIO output are available for normal L+ or other use, as shown by port 4 in Figure 1.

During normal IO-Link or SIO operation, the L+ MOSFET is OFF and the CQ output operates normally via TXEN, TXD and RXD. All IO-Link functionality is maintained, including full speed communications at COM3 speed and wake-up pulse generation.

During SIO+ operation, the L+ MOSFET is controlled via the SPI register interface and CQ is disabled (TXEN is low or under SPI register control). The upper nibble of register 0xE controls the L+ MOSFETs. During SIO+ mode, the switching frequency is limited to approximately COM1 speed.

While the LTC2874 will not be damaged if both the CQ and L+ outputs are active at the same time, this operating mode is not recommended since the output waveform's rise and fall trajectories are non-monotonic. These trajectories arise due to the interaction of the timing differences between the channels and the various current limits and source resistances.

The maximum output current for SIO+ mode is determined by the choice of MOSFET and sense resistor RS. The current limit is set by 50mV/RS. The typical current limit for the circuit in Figure 1 is 500mA. Accounting for tolerance and variation results in an output rating of 400mA for the port. The MOSFET must be selected to handle the voltage, current, and safe operating area (SOA) requirements. See the LTC2874 data sheet for more details.

The output capacitance of the MOSFET contributes approximately 60pF toward the maximum of 1nF that is permitted by the IO-Link standard.

Because this circuit parallels two drivers, the inactive driver acts as a capacitive load on the active driver. When the active driver changes state, it will generate a charging current in the inactive driver. This effect is more noticeable during IO-Link operation due to the larger capacitance of the MOSFET and faster edge rate of the CQ driver. To prevent the charging current pulses from creating ringing when the active driver turns off, minimize the parasitic inductance between the MOSFET source and C/Q driver output.

Figures 2 and 3 show operational waveforms for a single SIO+ capable port driving a resistive load while operating in either SIO+ or normal IO-Link mode. The supply voltage is 24V and the resistive loads are 56Ω and 200Ω, respectively.

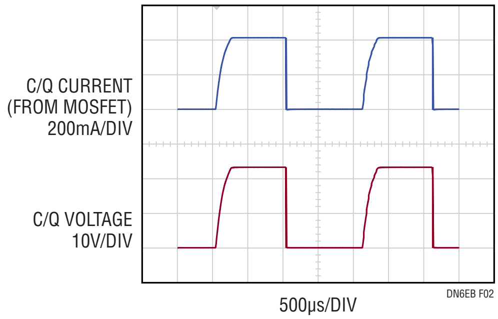

Figure 2. SIO+ Operation

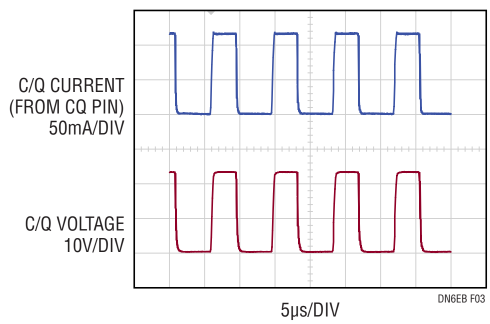

Figure 3. IO-Link Operation of a SIO+ Capable Port with C/Q at COM3 Speed

Conclusion

Arbitrarily large currents for LTC2874 operation in SIO+ mode may be obtained by repurposing the hot swap channels as higher current SIO drivers.

Data Sheet Download

www.linear.com/LTC2874

For applications help,
call (408) 432-1900

DN566 LT/AP 0917 71K • PRINTED IN THE USA

error; what's critical is how well they're matched. If both R_F and R_I are off by 10% in the same direction (either plus or minus), then the ratio remains ideal and the gain error is zero.

When a current-sense amplifier specifies a room temperature gain error of 1%, for instance, in Texas Instruments' INA180 datasheet *(Fig. 3)*, it's actually an indicator of how well-matched the resistors are opposed to the external gain network, where 1% resistors result in ±2% error.

The drift is also typically low on the integrated resistors—20 ppm/°C in the case of the INA180, which minimizes the effect of temperature drift on the total error. *Figure 4* approximates the impact on error for two different amplifier arrangements over temperature. For this comparison, make the following assumptions:

- An ideal 200-mΩ shunt with 0% error and no drift.
- Room-temperature calibration to eliminate the initial offset and gain network errors.
- Offset drift of 5 μV/°C for both configurations.
- Configuration No. 1 uses an external gain network of 20 V/V with a 50 ppm/°C gain network.
- Configuration No. 2 is a current-sense amplifier with matched resistors configured for a gain of 20 V/V and 20 ppm/°C drift.

TEXAS INSTRUMENTS

INA180

www.ti.com SBOS741 – APRIL 2017

7.5 Electrical Characteristics

at T_A = 25°C, V_S = 5 V, V_{IN+} = 12 V, and V_{SENSE} = V_{IN+} – V_{IN-} (unless otherwise noted)

PARAMETER		CONDITIONS	MIN	TYP	MAX	UNIT
E_G	Gain error	V_{OUT} = 0.5 V to V_S – 0.5 V, T_A = –40°C to +125°C		±0.1%	±1%	
	Gain error vs temperature	T_A = –40°C to +125°C		1.5	20	ppm/°C

3. This part of Texas Instruments' INA180 product datasheet electrical specifications table shows offset and offset drift.

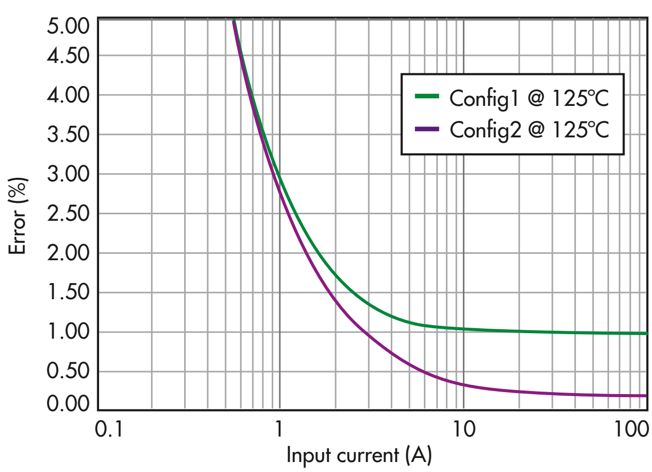

4. The plot compares approximate measurement error over temperature between a current-sense amplifier with 20-ppm/°C drift and a discrete implementation using a gain network with 50-ppm/°C resistors.

SUMMARY

Maximizing accuracy in current measurements over temperature is critical in many applications. While room-temperature calibration is relatively straightforward, performing multi-temperature calibration is time-consuming and expensive. Using a current-sense amplifier that integrates a precisely matched, resistive gain network will minimize the temperature-drift effects of the gain error.

Implementing a low-drift amplifier architecture allows you to minimize the offset error over temperature. On that front, Texas Instruments has developed zero-drift current-sense amplifiers, including the INA180 and INA181, that enable high initial room-temperature accuracy while minimizing the effects of temperature. ed

REFERENCES:

1. Integrating the Current Sensing Signal Path, TI TechNotes (SBOA167), December 2016.
2. Texas Instruments INA180 and INA181 data sheets.

DAN HARMON is the marketing manager for the Current and Power Measurement product line at Texas Instruments. In his 30-year career at TI, he has supported a wide variety of technologies and products including interface products, imaging analog front-ends (AFEs) and charge-coupled-device (CCD) sensors. He also has served as TI's USB-Implementers Forum representative and TI's USB 3.0 Promoter's Group chairman.

Dan earned a Bachelor's in electrical engineering from the University of Dayton and a Master's in electrical engineering from the University of Texas in Arlington. If you have questions about this article, post a question to the TI E2E™ Community Current Sensing forum.

Surface Mount High Voltage The Right Way

DTI'S SURFACE MOUNT HIGH VOLTAGE DIODES

The cost and quality advantages of surface mount manufacturing technology make it extremely popular in new electronic design. Limited availability of diodes has made the transition to utilize this technology hard for many in the high voltage space - until now!

High voltage surface mount diodes from Dean Technology cover a wide range of electrical performance, and with our new SL series in a longer package we can offer ratings up to 30 kV.

Custom versions are also available for customers that have specific needs not met by our standard product line. We will work with you to get the parts you need for your new designs now.

IN STOCK TODAY! Contact us for a quote or to place your order!

www.deantechnology.com

HVCA CKE HVPSI

Products By: DEAN TECHNOLOGY

Use Current to Drive Solenoid, Relay from Array of Voltages

By JAMES LOCKRIDGE | Applications Engineer, Texas Instruments

RELAYS AND SOLENOIDS are available with many different voltage ratings. Most factory- and process-automation equipment operates from 24-V supplies. However, customers may have control signals for a wide range of ac or dc voltages, such as 12 V, 24 V, 36 V, 48 V, or even 120 V or 220/240 V for some valves and contactors. For each voltage, the coil designs of the variants must be different. Having coils to accommodate all of the possible voltage adds to inventory, bill of materials (BOM), and spare-part headaches.

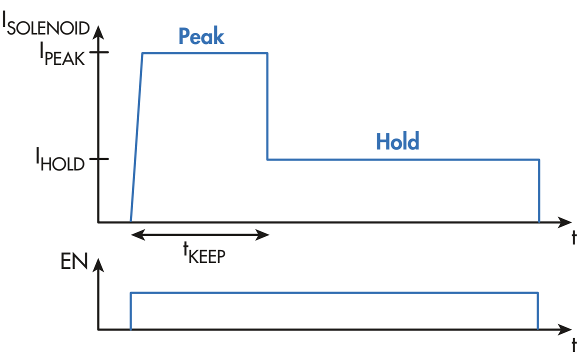

1. Here's an example of how current shaping with PWM can save power in the long term.

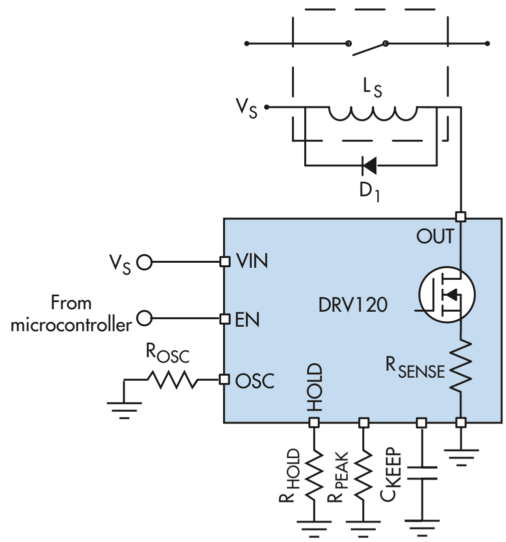

2. Using current control rather than voltage for relay/solenoid driving, here utilizing the DRV120, maintains the desired drive level independent of rail voltage or change in coil resistance with temperature.

One solution could be to design one coil for 12 V, and then use a resistor to limit the current into the coil for each voltage option. However, this wastes energy and dissipates heat from the resistor, especially if the 12-V coil is used for the 220/240-V design.

A more energy-efficient solution is to use pulse-width modulation (PWM) driving and a freewheeling diode to regulate the current in the solenoid. In addition, you could add current-sense feedback with PWM and control the current in the solenoid. Using a PWM driving method can also allow you to shape the current-drive profile to have a higher pull-in current for actuation, then drop to a lower current to hold the solenoid in place and save power in the long term *(Fig. 1)*.

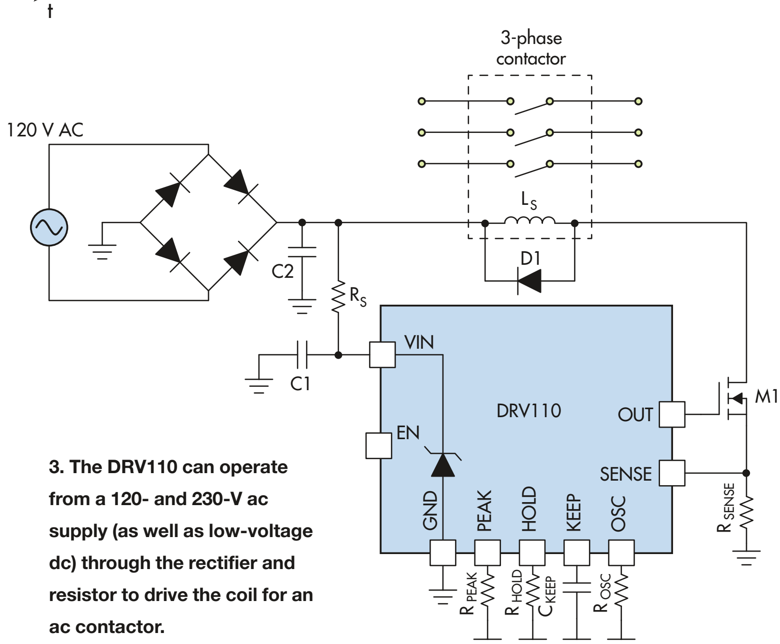

3. The DRV110 can operate from a 120- and 230-V ac supply (as well as low-voltage dc) through the rectifier and resistor to drive the coil for an ac contactor.

While PWM driving without current-sense feedback is easy to implement, variations in coil resistance, temperature, supply voltage, and the other factors can cause unintended solenoid de-actuation. As temperature increases, the resistance of the coil will increase. According to Ohm's law, the increased resistance will cause the current to decrease. Since current is responsible for generating the magnetic force, when the current decreases, so will the magnetic force.

Thus, using PWM driving with current-sense feedback is a more reliable driving method for the coils and relays. With current control, the driver regulates the coil current to the required value independent of resistance, which makes the system more robust over temperature. You could implement current-sense feedback with a MCU and discrete signal-chain components, or use a solenoid/relay driver such as the DRV120 *(Fig. 2)*.

For ac-line power applications with PWM driving, the ac voltage must first be rectified. A driver such as the DRV110 (similar to the DRV120) can then handle the large dc voltage with a current-limiting resistor and a Zener diode. The rectified ac power then drives the solenoid or valve control directly *(Fig. 3)*. Using dc to drive a contactor instead of ac simplifies magnetics design by eliminating the need for a shaded ring in the coil core.

With this current-control technique, you can design only one 12-V coil with your desired force requirements for any ac or dc supply of 12 to 48 V and potentially even up to 120 to 220 V. This allows you to have multiple product offerings while only needing to design for one coil.

Another challenge exists when you want to use a particular relay or solenoid, but the available supply-rail voltage is too high; e.g, when the coil is rated for 12 V but the supply rail is 24 V. While one "obvious" solution is to just switch to a unit with a 24-V coil, doing so may not be easy: Perhaps it also involves adding the new part number as an approved component, which may be a difficult and lengthy process, or there's a need for fast time-to-market. Or maybe the design needs the force specifications of a particular solenoid; however, it's rated for 12 V but you only have a 24-V rail.

Using current control can solve this dilemma, since designers can utilize the 12-V relay or solenoid for both 12- and 24-V systems, and even be prepared for a next-generation product that may have a 36- or 48-V supply rail. Thus, current control will allow continued use of the chosen solenoid or relay, and maybe by using the same solenoid in all three systems, those who procure the parts on the BOM can get a volume discount for purchasing large volumes of that particular part number.

Single-/Dual- Supply Precision Analog Limiter Handles Wide-band Signals

By ALEX RYSIN | Analog Design Engineer

AS THE NAME IMPLIES, a limiter circuit (sometimes called a "clipper") prevents a signal from exceeding desired thresholds, which would otherwise cause saturation and overload of the following stage. The circuit in *Figure 1* precisely limits the wideband input signal V_i at the V_{REF} level and can be used with both single and dual power supplies.

Output signal V_{o1} of the first stage, U1_1 (one-half of an LT1810, a dual 180-MHz, 350-V/µs low-distortion op amp), is precisely limited at the level of V_{REF}, but with inverse polarity compared with the input signal. U1_1 has a gain of -1 for the input signal V_i, and a gain of +2 for the reference signal V_{REF}.

Since the positive input of U1_1 is referenced to V_{REF}, the output signal is referenced to $2 \times V_{REF}$. Thus, while the input signal rises from zero to V_{REF}, the output signal of the first stage decreases from $2 \times V_{REF}$ to V_{REF}. The second stage of limiter U1_2 is an inverting amplifier with the gain of -1. It's referenced to the V_{REF} as well, and restores the original dc level and polarity of the input signal.

For the output of the first stage, $V_{o1} = 2 \times V_{REF} - V_i$ if $V_i \le V_{REF}$ and $V_{o1} = V_{REF}$ if $V_i \ge V_{REF}$, while the output signal of the second stage is $V_{o2} = 2 \times V_{REF} - V_{o1} = V_i$, if $V_i \le V_{REF}$, and $V_{o2} = V_{REF}$ if $V_i \ge V_{REF}$. (Note that the output-dc level, along with the limiting point, may be shifted by an optional network of R6 and R7; otherwise, that network isn't used.)

D1 is a part of the negative feedback of U1_1, making it an "ideal diode"—as soon as V_i reaches the V_{REF} level, U1's output voltage increases until it compensates for voltage drop V_d of the diode. When V_i is less than V_{REF} and D1 is turned off, diode D2 is on and sends the output current of U1 directly to its input. This prevents U1 from negative saturation, which would significantly decrease the switching speed of the first stage.

The voltage level on the output pin of U1_1 is $(2 \times V_{REF}) - V_i + V_d$, and it reaches its maximum level with the minimum level of the input signal. This determines the maximum possible level of V_{REF}, which doesn't cause the distortion of the limiter's output signal at its minimum level of $V_{REF} = [(V+) + V_i - V_d]/2$. Here, V+ is

1. This analog circuit, based on the LT1810 op amp, precisely limits input signal V_i at reference level V_{REF}.

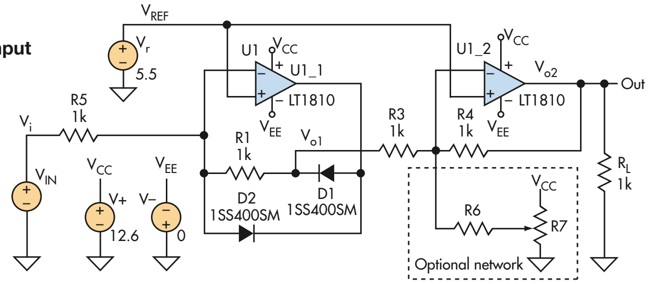

the maximum output voltage of op amp at the chosen V_{CC}, which is almost equivalent to V_{CC} for the rail-to-rail op amp. For the same reason, the minimum level of V_{REF} should not be lower than V_d; however, the input signal may significantly exceed V_{REF}.

The LTspice simulation shows the limiter's responses for a 1-MHz sine-wave signal with 8-V peak-to-peak amplitude *(Fig. 2a)*, and for a 10-V, 2-µs piecewise linear (PWL) signal crossing the limitation level several times *(Fig. 2b)*, both with V_{REF} of 5.5 V. The limiter is quite accurate within the entire output-voltage range of the op amp. The simulation demonstrates that limitation errors don't exceed 8 mV and output setting time is within 30 ns around the limitation point. When compared to the "simple op amp clipper" of Ref. 1, this limiter works at much-higher frequencies and with significantly lower distortion.

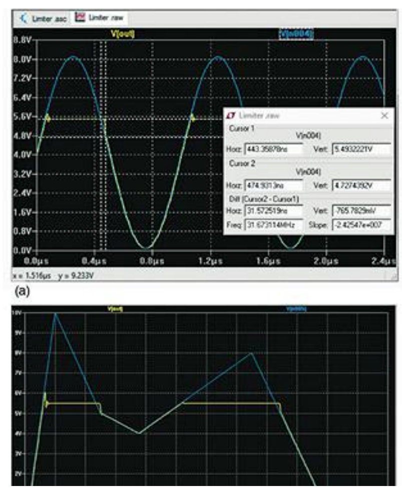

2. The LTspice simulation shows the limiter's response for a 1-MHz sine wave signal (a) and a 2-µs piecewise-linear signal (b).

REFERENCES

1. Thomas Mosteller and Aaron Schultz, "Op Amps Make Precision Clipper, Protect ADC" *Electronic Design,* November 2016
2. Datasheet, LT1810 - Dual 180MHz, 350V/µs Rail-to-Rail Input and Output Low Distortion Op Amps *(http://www.linear.com/product/LT1810)*

From
Concept to Reality

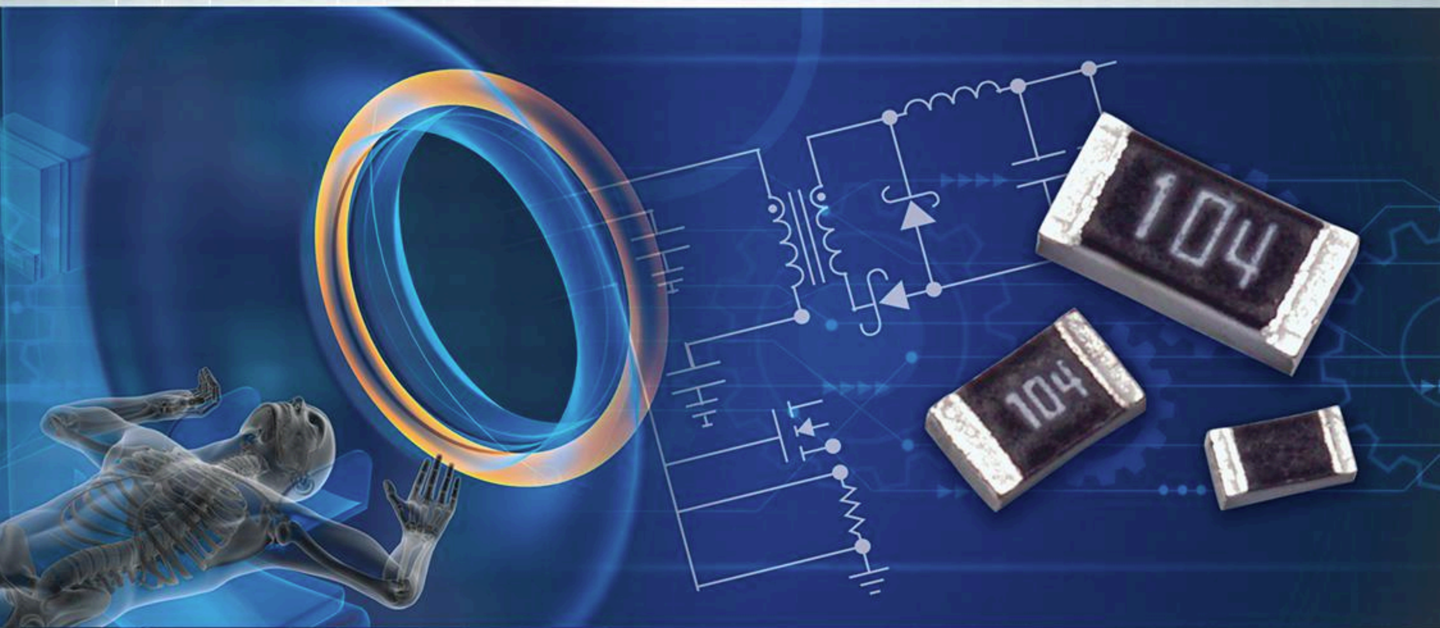

Our new passive component solutions will help you make the leap!

Whether you're designing the medical device of tomorrow or connecting the Internet of Things, you need cutting-edge product solutions to bring your design to fruition. At KOA Speer, we're the ideal partner to help you do just that. Our constantly expanding line of passive components will give you the solution to make your concept a reality!

More than just Resistors - yes we're the world leader in thick film resistors, but our product line is broad and growing every day, including:

- Shunt/Current Sense Resistors
- Wide Terminal Resistors
- High Voltage Resistors
- Thin Film Chip Resistors
- Surge Resistors

Passive Component Solutions - our engineering group will work with your design team to develop the passive component solution you need. Many of the hundreds of new products we've introduced in recent years are the result of such collaboration.

Award Winning Passive Supplier - we're the industry's most recognized and awarded supplier for achieving the highest product quality, on-time delivery and responsive customer service.

NEW! From KOA Speer

Ultra-Precision Thin Film Resistors (RN73)

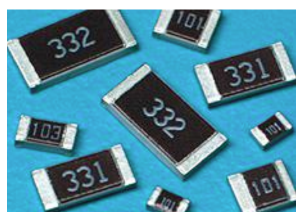

Surge Current Thick Film Resistors (SG73)

High Voltage Chip Resistors (HV73)

KOA®
KOA SPEER ELECTRONICS, INC.

Make the Leap from Concept to Reality, visit KOASpeer.com

New Products

Mini Desktop PC Serves as Media Center

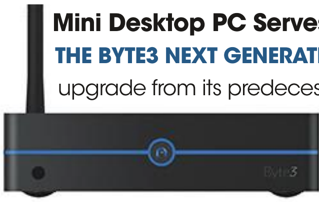

THE BYTE3 NEXT GENERATION fanless mini desktop PC from Azulle is an upgrade from its predecessor, Byte Plus. Lighter, thinner, faster, and quieter, the new PC increases flexibility and is better able to manage multiple tabs and applications.

A quad core Intel Apollo Lake N3450 processor equipped with 4 or 8 GB of RAM and 32 GB of eMMC storage, equips the PC for serving as a media center, living room video conferencing hub and file server. A Micro SD card slot accommodates up to 256 GB and 2280 SATA allows for theoretical speeds as high as 6 GB/s.

The computer can send alternating words and images to LED, LCD and plasma displays or kiosks to target specific market segments with customized messages. I/Os include: SD Card slot; Kensington lock; power input and audio output jacks; and 4x USB, VGA, Ethernet and HDMI ports. Equipped with Wi-Fi, Bluetooth and IR receiver, the Byte3 mini desktop PC will release August 28, 2017, and is available for pre-order now.

AZULLE, *www.azulletech.com*

1U Rackmount Provides Up to 32 GbE LAN Ports

THE PL-81230, a 1U rackmount high-performance networking system from WIN Enterprises, is based on the Intel Xeon E5-2600 V3/V4 processors with up to 22 processing cores, coupled with the Intel C612 series chipset. The CPU features 8-ch DDR4 2,400 MHz memory with 512 GB maximum memory capacity, 55 MB of cache, 80 PCIe lanes, up to 4 NIC modules and multiple Ethernet module bays providing fiber and copper LAN ports with optional bypass function, and up to 8x PCIe X8 slots for up to 32 GbE LAN expansion ports. The platform supports networking functions, such as Unified Threat Management, SPAM filtering, firewall, routing, switching, etc.

I/O elements include two management Ethernet ports, a console port, two USB ports, a Graph LCD module with 5-key keypad, and LEDs for power/ HDD / 2x GPO. The PL-81230 networking system also features internal 2x 2.5" SATA HDD/ SSD swappable, 2U ATX redundant power supply, and onboard Compact Flash m-SATA/M.2 2280 slot for basic network storage applications.

WIN ENTERPRISES, *www.win-ent.com*

Rugged NAS Storage PC Features Core i7

THE MH70S FROM MEN MICRO is a modular turn-key storage PC meeting the requirements of storage intensive applications such as digital video recorders, content servers or NAS. Designed for use in trains, trams or industrial environments, the computer consists of a modular 40 HP CompactPCI Serial system which can be wall or rack-mounted. Independent of the mounting position, cooling can be implemented either by natural convection or with an additional fan tray at the bottom of the system. Complete software support and a multitude of built to order configuration possibilities are available.

The MH70S storage PC features an Intel Core i7 processor supporting Intel AMT, up to five HDD/SSD carriers, and up to eight PoE PSE ports supplying eight devices with a maximum of 100 W in total can be implemented. For robust and safe operation, the HDD carriers can be configured in RAID 0, 1 or 5 while the carriers themselves have their own internal RAID 0, 1 and JBOD hardware configuration possibility.

MEN MICRO, *www.menmicro.com*

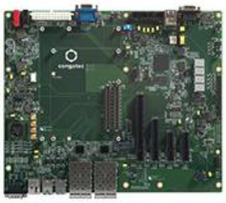

Starter Set Evaluates COM Express Type 7 SoMs

CONGATEC'S NEW CONGA-X7/EVAL COM Express Type 7 quick starter set is designed for the evaluation of server-on-modules in accordance to PICMG's COM Express Type 7 standard. The starter set is suited for the design of Industry 4.0 micro servers communicating in real-time, redundant fog servers, and transcoding cloudlets at the edge of carrier-grade infrastructures. Universally applicable, it targets both commercial and extended temperature ranges for rooftop and public transportation installations.

The evaluation board comes with 4x 10 Gb, 1 GbE, 4x USB 3.0 and an integrated baseboard management controller provides VGA, iKVM and virtual storage. For server extensions, the board executes 32 PCIe lanes via PCIe x16, PCIe x8 and 2x PCIe x4 slots. Further extensions include 4x USB 3.0/2.0 and a COM port, VGA, 2x SATA Gen3 interfaces, PCIe lanes supporting NVMe, and iKVM for input signals over LAN.

A typical COM Express Type 7 quick starter set integrates the conga-X7/EVAL carrier board, its schematics, a choice of COM Express Type 7 server-on-modules, and an accompanying cooling solution.

CONGATEC, *www.congatec.com/us.html*

PC/104-Plus SBC Features Kaby Lake Processor

VERSALOGIC'S NEW LIGER PC/104-Plus SBC, based on Intel's new Kaby Lake processor, combines high performance processing and video with 12 to 14 watts typical power consumption. The dual-core board features hardware-level security using an on-board TPM security chip, and backwards compatibility with systems using PC/104-Plus (ISA or PCI) expansion.

I/O connectivity includes two GbE ports with network boot, 2 mini DisplayPort video outputs, 6 USB ports (2x 3.0, 4x 2.0), 4 serial ports, 8 digital I/Os, I2C, and 3 timers. SPI / SPX interface, Mini PCIe socket with mSATA, and SATA (Rev 3) interface supporting up to 6 GB/sec are also included.

Designed and tested for industrial temperature operation and meeting MIL-STD-202G specifications, the 108 x 96 mm PC/104 footprint Liger (VL-EPM-43) SBC is targeted for compute-intensive high-end applications such as flight navigation, guidance systems and medical imaging. Available in Core i3, i5 and i7 options, the board is compatible with various x86 application development tools.

VERSALOGIC, *www.versalogic.com*

HV MOSFET Platform Enhances Switching, EMI

AOTF190A60L IS THE first product in Alpha and Omega Semiconductor's new aMOS5TM HV MOSFET platform. Engineered for superior switching and EMI performance, the device provides a high-efficiency solution optimized for server power supplies, high-end computers, charging stations and other high-performance applications.

Proprietary aMOS5TM technology reduces the on-resistance per unit area by 30% versus the previous generation, enabling lower conduction losses in the application. Features of the technology include a lower total gate charge with shorter switching plateau time helping to reduce switching losses compared to competitors, improved EMI performance by optimizing parasitic switching parameter behavior, and an enhanced body diode increasing robustness and fast reverse recovery.

Specific applications for the new MOSFET platform include: SMPS with PFC, Flyback and LLC topologies; and micro inverters with dc/ac inverter topology. The product is packaged in a full-molded TO-220F and features a maximum RDS(ON) of 190 mΩ. Priced at $2.24 each/1,000, the AOTF190A60L aMOS5TM HV MOSFET is available now with a lead-time of 12-14 weeks.

ALPHA AND OMEGA SEMICONDUCTOR, *www.aosmd.com*

Octo Core Serverblade Raises Performance/Watt

THE CP6006-SA is Kontron's new generation of 6U server class CPU platforms in the CompactPCI format. Based on Intel Xeon D-1500 family of CPUs with 2 to 16 cores, the product line is designed for multi-CPU server-class applications in demanding environments. The boards will initially be offered with 8-core CPUs with 1.6 GHz at a TDP of 35 W, along with a more powerful variant with 2.0 GHz and a TDP of 45 W. Up to 64 GB of DDR4 ECC memory boosts virtualization support, when virtualized tasks require separated memory in servers and computing nodes.

The product variant CP6006X-SA offers higher data throughput via PCIe 3.0 and 2x 10 GbE on the backplane, targeting it for compute intense applications such as sonar, radar and video stream analytics in areas like commercial avionics, defense, telecommunication, industrial automation and medical image processing. The CP6006-SA 6U CompactPCI server platform is available now as a prototype for testing purposes and regular production is scheduled in the next several months.

KONTRON, *www.kontron.com*

Chassis Provides Scalable Compute/ Acceleration Platform

THE MAXCORE MICRO, a compact compute and acceleration platform from Artesyn Embedded, is designed for applications ranging from small cell baseband processing, scalable video encoding, video surveillance and bump-in-the-wire monitoring to industrial computing. The enterprise-class chassis holds two PCIe Gen 3 FHFL slots (up to 150 W per slot), one for a host server card and the other for any PCIe add-in card. Switchless PCIe connectivity between the slots is up to 100 Gbps. While the compute and acceleration platform can function as a complete stand-alone microserver, its form factor allows three chassis to be fitted side-by-side in a 19-in. rack. The platforhttps://azulletech.com/m can be deployed in a range of environments, from a 19-in. rack to wall-mount. With 150 W of power per slot, it is possible to build a system with as many as 48 Intel Xeon D cores per compact chassis, using two dual-processor 12-core add-in cards such as Artesyn's SharpServer PCIE-7410.

ARTESYN EMBEDDED, *www.artesyn.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

Intel's AI Chip Available in a USB Stick

Intel has made its artificial intelligence (AI) processor available via a USB 3 dongle that works with everything from PCs to Raspberry Pi's.

If you are learning about machine learning (ML) or need to accelerate your latest artificial intelligence (AI) creation, Intel's new Movidius Neural Compute Stick *(Fig. 1)* may be just what you're looking for. The USB 3.0 device contains a 28-nm Movidius Myriad 2 MA2x5x vision processing unit (VPU). Intel picked up Movidius when it was working on getting the dongle out the door. It complements Intel's new Xeon and Xeon Phi. These have also been tuned to handle deep neural networks (DNN). The stick is priced at only $79.

1. Intel's $79 Movidius Neural Compute Stick is a USB 3.0 dongle that contains a Movidius Myriad 2 VPU.

The VPU is optimized for vision applications but can handle all sorts of DNN applications. While the chip has a pair of 32-bit RISC processors to manage resources but the heavy lifting is done by the dozen 128-bit SHAVE vector processors. The SHAVE units support 16- and 32-bit floating point plus 8-, 16- and 32-bit integer operations. The chip's processors share 2 Mbytes of RAM on chip plus a 256 Kbyte L2 cache and the DDR interface provides access to 1 or 4 Gbits of DDR memory in a stacked memory configuration. The Stick has 4 Gbits for DRAM.

The chip runs at 600 MHz at 0.9 V, consuming less than 2.5 W of power. Its image signal processor (ISP) mode is designed to handle video streams directly, although the USB stick requires data to be loaded via the USB interface. The chip has a dozen 1.5 Gbit/s MIPI lanes that can be configured as CSI-2 or DSI interfaces. The chip can be configured directly using I^2C or SPI. There is a 1 Gbit/s Ethernet interface as well. The 4 Gbit LPDDR chip comes in an 8-mm by 9.5-mm BGA package.

The chip can deliver 10 inferences/s using the standard GoogleNet benchmark in continuous inference mode. It does so using only 1 W of power while delivering 100 GFLOPS of performance.

Multiple Myriad 2 chips can be ganged together for more performance. This scales almost linearly since the units work in parallel. A USB hub can be used with the sticks, but deployment of a system will likely utilize the chips directly. Multiple chips can also be used to run different DNNs instead of a single larger one, depending upon application requirements.

The Myriad 2 chip is available to OEMs and has been used in applications like DJI's SPARK drone *(Fig. 2)*. The chip processes captured video to identify objects and avoid them. The drone uses other SoCs for flight control, communication, and so on. The VPU is also handy for chores such as 3D mapping. DJI uses the VPU for facial and gesture recognition, as well as to implement a safe landing mode. The SPARK is priced at $499. ed

2. DJI's SPARK drone employs the Myriad 2 chip to implement visual collision avoidance support.